Development of Multiplexed Bolometer Readout Electronics for mm-wave Space Astronomy

Joshua Montgomery

Master of Science

Physics

McGill University

Montreal,Quebec

2015-02-02

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science in Physics

DEDICATION

In memory of Brian Gawlik and Nate Williams.

ACKNOWLEDGEMENTS

First, I would like to acknowledge my supervisor, Matt Dobbs; not just for the superb guidance he has offered over the last several years, but also for the trust and patience he's shown as I've had the opportunity to take on more responsibilities in the lab. I would also like to thank him in particular for his alacrity over the last month in helping to get this thesis written approximately on time. I must also acknowledge the whole group of people with whom I share the lab now, and who showed me the ropes over the past few years: Amy Bender, for her invaluable introduction to the lab, and for helping to keep me on the right track during the CSA project; Graeme Smecher, for picking up my confused skype calls and setting me straight; Adam Gilbert, for making the lab a positive, kind, environment, and his ability to deflate the anxiety in the most daunting tasks; and Kevin Bandura putting things into perspective with good humor on a daily basis. I wouldn't have gotten this far without all of the fantastic educators I've had along the way. In particular, Nate Williams who showed me the joy in science; and Brian Gawlik, an unparalleled mentor who pushed me without compromise in ways I will always be grateful for. Nor would I have gotten much farther than the door without the unyielding support from my parents, Lauryn Axelrod and Chris Montgomery. Finally, I thank my partner, Kita Mendolia – for her exceptional patience, kindness, and humor, especially over the last month.

ABSTRACT

The polarization anisotropies in the Cosmic Microwave Background contain a rich set of physical signatures. These include known phenomena, such as weak gravitational lensing, which offers a chance to study structure formation at redshifts sparsely explored; and predicted phenomena, such as the propagation of primordial gravity waves through the surface of last scattering ("gravitational B-modes"). The detection and measurement of the latter would be evidence in support of inflationary cosmology theories, the search for which has been at the epicenter of CMB experimental development for nearly a decade. A space-based platform is the best place to conduct such a search: the signature of B-modes is strongest at large angular scales, and full-sky coverage presents a significant advantage over ground-based instruments. Additionally, access to more frequency bands otherwise absorbed by Earth's atmosphere, improves the ability to characterize and subtract galactic foregrounds. To achieve the sensitivities capable of making precision measurements of the CMB polarization, the next generation of instruments will employ focal planes containing $\sim 10,000$ detector elements. This is an order of magnitude increase over the existing instruments. The detectors themselves are arrays of Transition Edge Sensing (TES) Bolometers – photon-noise limited sensors, which have been favored for the last decade of ground-based CMB instruments, but have never flown on a satellite platform. To achieve the requisite focal-plane densities, it is essential to "multiplex" – read out several bolometers simultaneously, over a single pair of wires. In this document, we present an electronics system capable of reading out, and controlling, the next generation of TES focal-planes: The Space-Flight Representative 64x Digital Frequency Multiplexing (DfMUX) readout system. These electronics were commissioned by the Canadian Space Agency, and built in partnership with COM-DEV. The Flight Representative 64x DfMUX system multiplexes together 64 readout channels, a factor of four increase over all current frequency-multiplexed CMB readout systems. We measure a

total readout noise of less than $10\frac{pA}{\sqrt{Hz}}$, which is consistent with prediction, and remains substantially lower than the intrinsic background photon noise of the CMB. We do acknowledge some excess noise at high frequency, and demonstrate that it is due to non-idealities in our testing setup that are well described by our noise model, not the Flight Representative readout electronics. Additionally, the Flight Representative DfMUX system meets all Canadian Space Agency space-flight criteria. these include radiation-hardness; thermal optimization for radiative heat transfer; a per-channel power consumption of 49mW.

ABRÉGÉ

Les anisotropies de polarisation dans le rayonnement de fond cosmologique contiennent un riche ensemble de signatures physiques. Il s'agit notamment des phénomènes connus, tels que les lentilles gravitationnelles faibles, qui offrent la possibilité d'étudier la formation des structures larges à des Z peu explorés et les phénomènes prédits, telles que la propagation des ondes gravitationnelles primordiales à travers la surface de la dernière diffusion ("modes B gravitationnels"). La détection et la mesure de ces dernières serait preuves à l'appui des théories de cosmologie inflationnistes, la recherche de ce qui a été à l'épicentre du développement expérimental CMB pendant presque une décennie. Une plate-forme spatiale est le meilleur endroit pour effectuer une recherche pour modes B gravitationnels: la signature des modes B est la plus forte à de grandes échelles angulaires, et la couverture complète ciel présente un avantage significatif par rapport aux instruments au sol. En outre, l'accès à plusieurs bandes de fréquences normalement absorbés par l'atmosphère terrestre améliore la capacité à caractériser et à soustraire les premiers plans galactiques. Pour atteindre les sensibilités requises pour effectuer des mesures de précision de la polarisation de la CMB, la prochaine génération d'instruments emploiera des plans focaux contenant ∼10,000 éléments de détection. Ceci est un ordre de grandeur de plus que les instruments de la génération actuelle. Les détecteurs eux-mêmes sont des bolomètres à détection de transition (TES) limités par le bruit photonique qui ont été favorisés durant la dernière décennie pour les instruments de CMB au sol, mais n'ont jamais volé sur une plate-forme satellite. Pour atteindre les densités de detecteurs requises au plan focal, il est essentiel de "multiplexer" les signaux, c'est-à-dire lire plusieurs bolomètres simultanément, sur une seule paire de fils. Dans cette thèse, nous présentons un système électronique capable de lire et de contröler les TES modernes utilisés dans la prochaine génération de plans focaux. Le système de mesure en multiplexage en fréquence à 64 canaux (DfMUX, 64x) qualifié pour vol spatial a été commandé par l'Agence spatiale canadienne, et construit en partenariat avec COM DEV. Ce système augmente par un facteur quatre le nombre de canaux offerts par les systèmes actuels CMB. Nous mesurons un bruit total de lecture de moins de 10 $\frac{pA}{\sqrt{Hz}}$, ce qui est conforme aux prévisions, et reste nettement inférieur à celui du bruit de fond de photonique intrinsèque de la CMB. Nous observons un certain excès de bruit à haute fréquence, et démontrons qu'il n'est pas en raison de l'électronique de lecture, mais plutôt à des non-idéalités dans notre configuration de test, qui sont bien prédites par notre modèle de bruit. En outre, le système DfMUX pour vol spatial répond à tous les critères de l'Agence spatiale. Cela comprend la robustesse aux radiations cosmiques, l'optimisation thermique pour le transfert de chaleur par rayonnement, et une consommation de puissance de 49mW par canal.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Science and Technology

1.1 An Introduction to the Science

The very early universe was a hot and dense place. So hot and dense that, even after cooling sufficiently for protons and neutrons to form, matter and radiation remained tightly coupled in a photon-baryonic fluid. Light couldn't travel far in this environment before being absorbed and re-emitted. So it remained for approximately the first 380 thousand years as the universe expanded. Eventually the universe cooled enough for chemistry to occur: protons captured free electrons and formed hydrogen and helium, decoupling photons from baryons. At this point the universe became electromagnetically neutral, and transparent to electromagnetic radiation (a period known as "recombination"). Light streamed outward in all directions from the surface of the cooling plasma (this occurred approximately 400 million years before the first stars ignited). That light, reflected from the "surface of last scattering", is still traveling today; the expansion of the universe since then has elongated its wavelength, shifting its energy it into the microwave, where it is observed in every direction as the Cosmic Microwave Background (CMB). A visual summary this cosmology can be seen in Figure 1–1.

A map of the CMB reveals what looks at first to be random noise. In fact, from the time it was discovered by Penzias & Wilson in 1965 [22] until 1992 when the data from the COBE satellite revealed the first anisotropies [27], the CMB appeared to be uniform. The progression in these measurements can be seen in Figure 1–2. The most recent all-sky map can be seen in Figure 1–3; the largest amplitude variations are fluctuations of just one part in 10^4. This general isotropy is taken as strong evidence for an expansionary event –

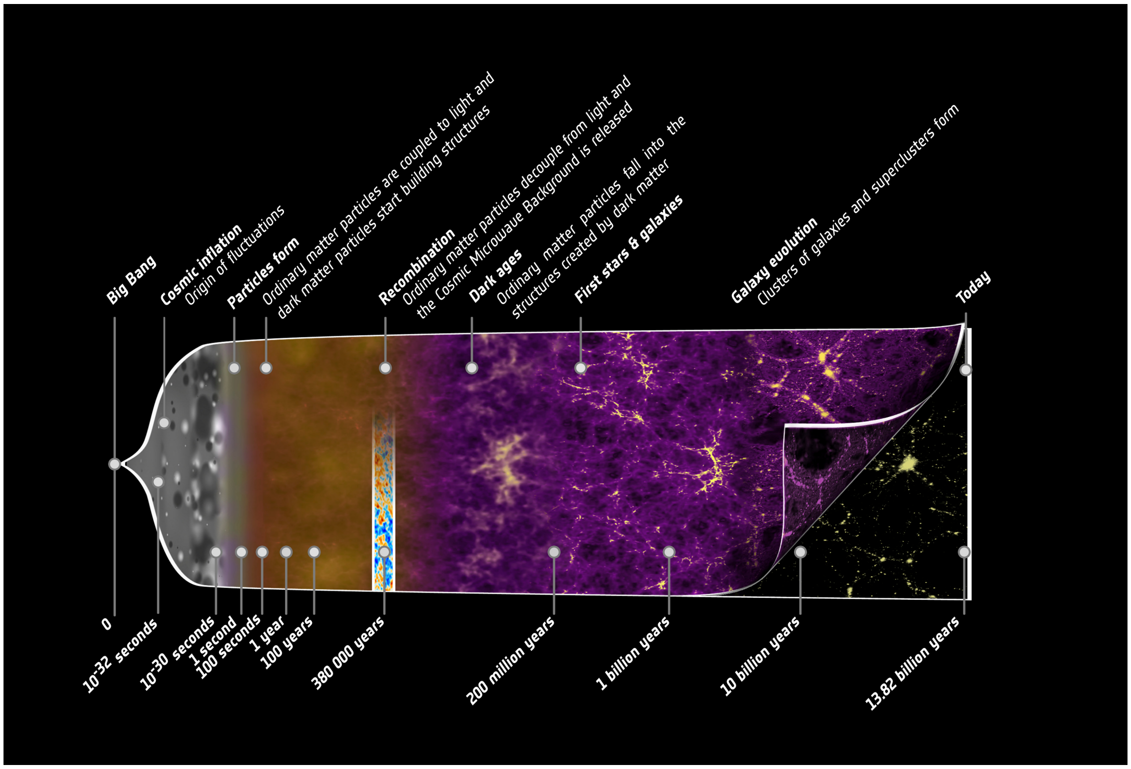

Figure 1–1: A summary of our current understanding of the evolution of the universe. Image Credit: ESA, Planck Collaboration.[1]

a means for parts of the sky that are not now in causal contact to at some point have been close enough to exchange energy.

A leading theory explaining, among other things, the presence and appearance of the small anisotropies in the otherwise homogeneous CMB is known as "inflation". This theory, introduced in 1981 by Alan Guth [9], posits that following the big bang, long before recombination, the universe underwent a period of exponential, superluminal expansion. Quantum fluctuations in the very early universe that would have otherwise rapidly thermalized were taken out of causal contact with their nearby environment during the expansion – thereby

[3] http://www.esa.int/spaceinimages/Images/2013/03/Planck_history_of_Universe_zoom

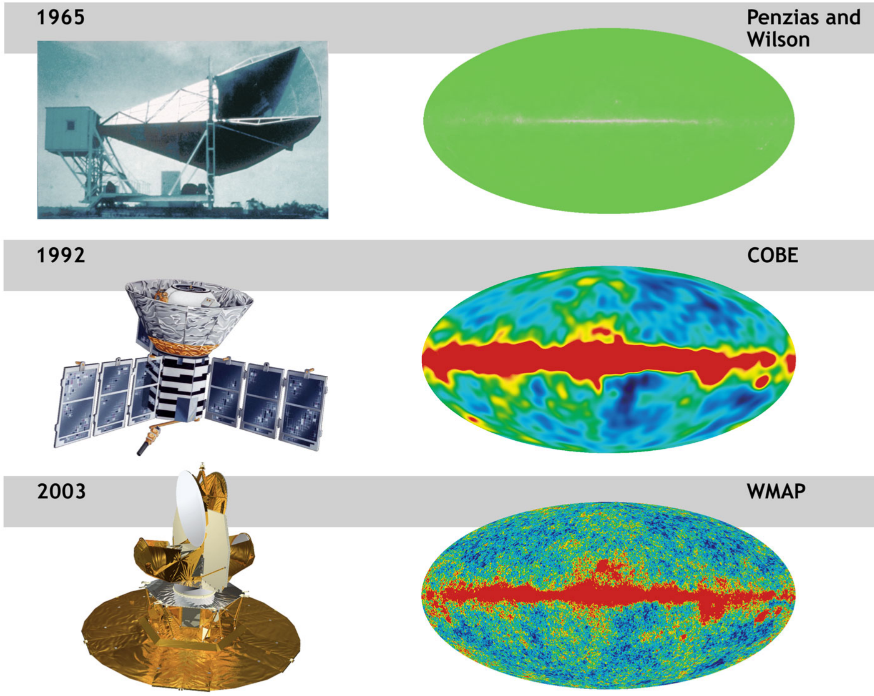

Figure 1–2: A set images showing the progression of CMB temperature anisotropy measurements from its discovery in 1965 to the 2003 WMAP 3-year map. Note that the top image map is a simulation based on the sensitivity of the Penzias and Wilson microwave receiver. Strong contamination of the CMB signal from galactic foregrounds can be seen in the bright band through the galactic plane. Image Credit: NASA, WMAP Science Team.[2]

"freezing" them into the fabric of our universe, while blowing them up to enormous proportions. When the period of inflation ended, those quantum fluctuations had become the dominant structure in the universe: under- and over-densities of matter distributed as a near scale-invariant Gaussian-random field.

At this point, the photons and baryons were still tightly coupled. The over-dense regions began to collapse gravitationally, which was countered outward by the consequential rise in radiation pressure. This process produced acoustic oscillations that propagated through the medium, and the manner in which they did so was influenced by characteristics (such as the

[3] http://wmap.gsfc.nasa.gov/media/081031/index.html

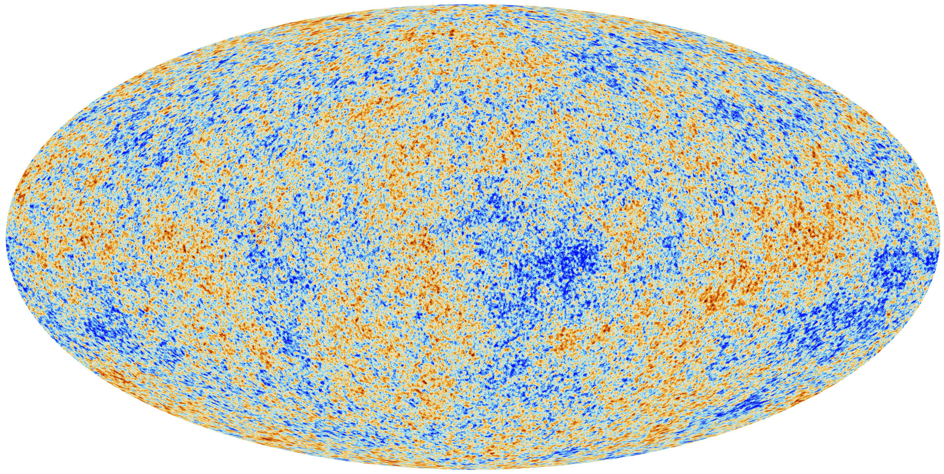

Figure 1–3: A recent all-sky map, with galactic foregrounds removed, produced with the Planck Satellite in 2013. Image Credit: ESA, Planck Collaboration.[3]

composition or geometry) of the universe. During recombination, when radiation decoupled from the baryonic medium, propagation of the acoustic oscillations halted. The remaining topology of over- and under-densities in the universe, and the scales thereof, seeded the structure we see today. The CMB provides a sort of image of the surface of last scattering, of an early state of the primordial universe. By analyzing the scales on which power is distributed in the CMB, we can recover information about the composition of the early universe: its geometry, its energetics, and its dynamics. It is one of our most versatile probes of the universe from before light could travel; a way to measure, with precision, the initial conditions from which everything around us today evolved. One such sought-after measurement is a particular signature in the *polarization* of the CMB that would be left by inflationary gravity waves: a result of the dramatic change in density and volume following the superluminal expansion, imprinted on super-horizon scales throughout the surface of last scattering. These polarization signals offer descriptions of the universe 10^{-35} seconds after the big bang, a universe full of interactions with energy scales on the order of 10^{16} GeV. This signal is often said to be the "smoking-gun of inflation".

The temperature anisotropies in the CMB have been well characterized at scales ranging from the full sky, down to arcminute resolutions. Today, anisotropies in the polarization are the new vanguard of CMB science. There are two types of polarization signals, named in analogy with the electric and magnetic fields: "E-modes", and "B-modes". E-modes are curl-free components of the polarization vector-field on the sky, and can be generated by velocity perturbations in the surface of last scattering. These are an order of magnitude fainter than the temperature perturbations, and were first detected by the DASI experiment in 2002 [18]. B-modes are gradient-free components of the polarization vector-field on the sky, and together with E-modes form a basis for the polarization field. Unlike E-modes, these signals have handedness, and therefore require tensor-field perturbations in the surface of last scattering, or to the CMB photons as they travel through space. B-modes have two potential cosmological origins: weak gravitational lensing of the CMB, and primordial gravity waves. The former, fainter still than E-modes by another order of magnitude, are strongest on small angular scales; these were discovered by the South Pole Telescope (SPT) in 2013 [10], and detected since by the POLARBEAR instrument in 2014 [2]. Most recently, in 2014, the BICEP experiment announced a detection of B-modes on large angular scales, and follow-up is being conducted to determine if the signal is cosmological. The BICEP measurement is the most sensitive large-angular scale polarization measurement to date, and represents a significant step forward in the search for gravitational B-modes. [1]

The CMB polarization field offers both incremental improvements in the science established by studying the CMB temperature field, as well as completely new paradigms

[3] http://www.esa.int/spaceinimages/Images/2013/03/Planck_CMB

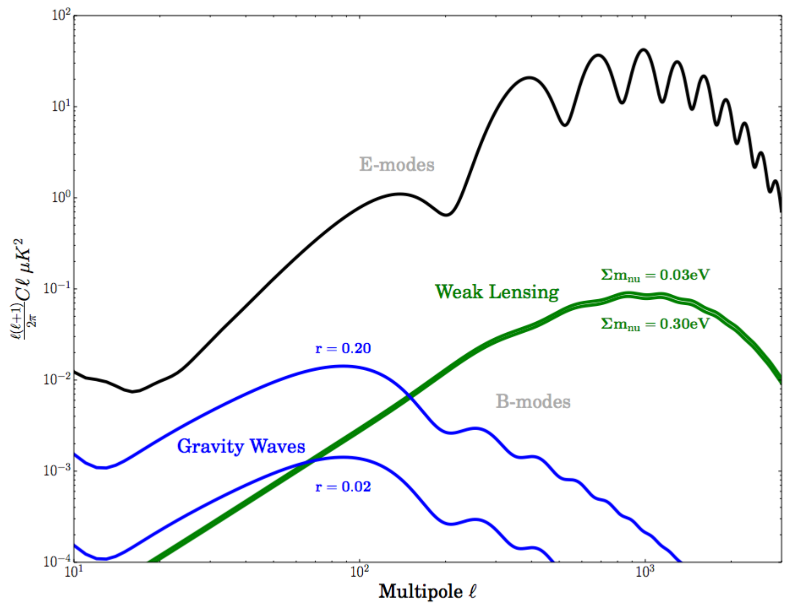

Figure 1–4: A comparison of the relative amplitudes of CMB polarization signals – including the E-mode power spectrum, the weak-lensing B-mode power spectrum, and two possible bounds in the theoretical region of primordial gravity mode signatures. The x-axis is in units of multi-pole moment, such that 1 degree on the sky is approximately $l = 100$, and the y-axis is in μK temperature. [25]

for cosmological inquiry. E-mode measurements are another "view" of the same vector-perturbations responsible for the temperature anisotropies, and study of them will lead to improvements in the cosmological results originally obtained from the temperature field. The weak gravitational lensing field is a fundamentally different object of study: this field results from distortion of the CMB light as it is "lensed" by the gravitational fields of massive objects on its way to us. It is therefore an integrated mass measurement along the line-of-sight from us to the surface of last scattering. The kernel of this field is such that the most significant contributions to that lensing signal will be from objects with redshifts around 2. This seats this measurement between the reaches of the deepest optical surveys (at redshifts less than 1), and the cosmological distance to the CMB (redshift ∼1,000). Right now there are few other probes of large-scale structure at those distances, and a measurement there would be a prized data-point on a very incomplete map of structure evolution in the universe – with the CMB at one end, and optical surveys of nearby objects (such as the Sloan Digital Sky Survey) on the other. Quantifying how structure developed has implications beyond the field of cosmology, such as potentially providing one of the most sensitive measurement of the combined neutrino masses available [12]. Support for inflationary theories notwithstanding, measurements of primordial B-modes may open up a rich source of extremely-early-universe, and high energy, physics. The most rudimentary measurement of them would help pin down the energy scales of the primordial universe.

Detailed measurements of the temperature anisotropies, over a wide range of angular scales, opened the door to precision cosmology from the CMB. Like the temperature anisotropies, detection of the polarization signals herald the arrival of a new scientific resource. To take advantage of this resource, a next generation of instruments is needed. Every sensitivity increase is hard-earned. Detector technology today is such that measurement uncertainty is dominated by the shot-noise in the CMB photon arrival times, rather than the detectors themselves; meaning improvements must be culled from sources other than instrument noise. The number of bolometers in a focal-plane, instrument operating efficiency, and details involved in detector fabrication like optical coupling efficiency, are all quadratic with mapping speed. A factor of 10 increase in sensitivity implies significant technological advancement in a range of technologies.

1.2 An Introduction to the Instruments

Throughout the history of the field of CMB cosmology, both ground-based and space-based observatories have been utilized. Space-based observatories have three important disadvantages, which are not likely to change. The cost of sending a telescope into space will probably always be enormous compared to building an analogous ground-based experiment. The harsh radiation and thermal environment in space, and vibration during launch, place substantial constraints on the components and technology which can be used on satellite platforms, and so design elements, particularly in electronics, face more stringent constraints. This means technology that flies is rarely the current state of the art when it does. Most importantly, payload weight and dimension restrictions make telescopes with appreciable dish sizes virtually impossible to fly, at least until dramatic advancements in payload delivery are made. This last item means that, while larger angular scales are only accessible from space, a space-based platform won't be seeing the small angular scales that the 10-meter South Pole Telescope does from the ground anytime soon. These limitations, difficulties, and financial demands all incentivize ground-based solutions.

That having been said: fundamentally, the earth is a poor place to observe the CMB. Much of the microwave band over which the CMB can be observed is absorbed by water vapor in the atmosphere. This limits ground-based observatories to the highest and driest places on earth, such as the Atacama plateau in Chile, or the geographic South Pole on the antarctic continent.[4] Even there, we are left only windows in the microwave spectrum through which to see the CMB. Shown in Figure 1–5a, the ones most often utilized are at 90GHz, 150GHz, and 220GHz. By contrast, the most recent CMB satellite, Planck, had 9 frequency bands from 30GHz to 857GHz [23]. Using multiple frequency bands provides important degrees of freedom when performing foreground analysis and subtraction. Observing from space opens up the entire frequency range to do detailed foreground measurement and subtraction, which in turn makes even foreground-contaminated sky accessible, and allows a higher sensitivity to the underlying CMB. In this era, when we are no longer instrument-noise limited, foreground subtraction is a crucial tool for improving measurements.

Ground-based telescopes are limited to observing only the part of the sky accessible at their location on earth, and they often choose fields much smaller than that. Choosing to go

[4] Additionally, there have been a number of balloon-borne observatories flown in Antarctica, again striving to get above the atmosphere.

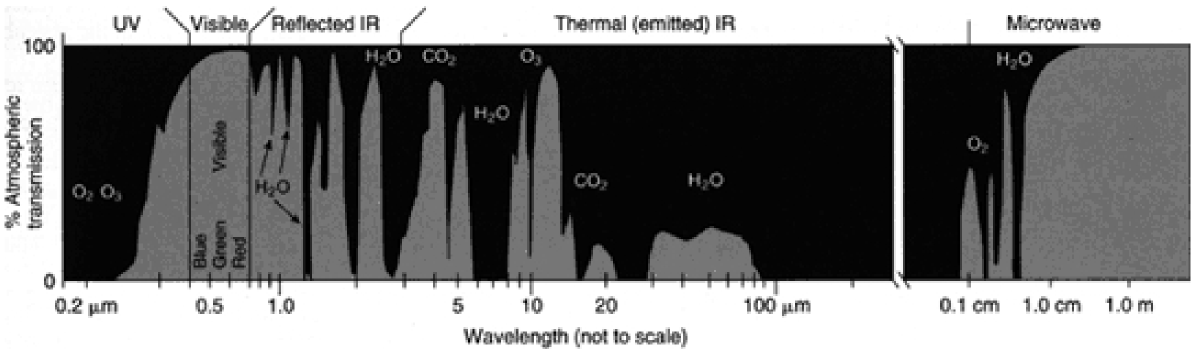

(a) Atmospheric transmission of the electromagnetic spectrum. Notice that for microwaves it is nearly opaque, leaving just three decently sized windows. Image Credit: NASA Earth Observatory, http://earthobservatory.nasa.gov/Features/RemoteSensing/remote_04.php

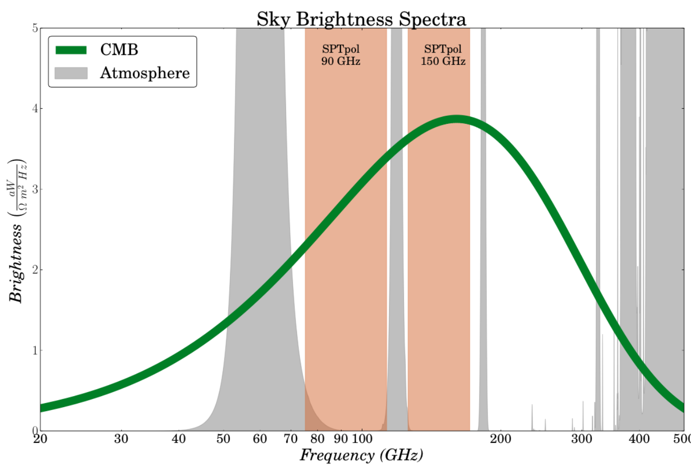

(b) A closer look at the microwave atmospheric transmission windows. In black is the atmospheric brightness at the South Pole, and in color are the pass-bands of the South Pole Telescope (SZ) camera. [25]

"deep" (integrate for a long time) on a small patch of sky plays to the strengths of ground-based telescopes with fine angular resolution. This choice also gives them the freedom to avoid looking through the thick and bright atmosphere at lower elevations, and also regions of the sky with high galactic contamination, where foregrounds are harder to understand.[5] Ultimately, the clearest and most compelling advantage to space-based telescopes is simply that the full sky is observable. This means larger angular scales, which cannot be observed from earth, can be seen from space. These large angular scale signals are thought to contain some of the most exciting science.

The Planck satellite observatory used a focal plane on the HFI instrument of under 100 sensors, of which only a fraction were polarization sensitive [23]. The next space-based CMB observatory will benefit greatly from the technological achievements over recent years, which have allowed ground-based telescopes to deploy focal planes with $\sim$1,000 polarization-sensitive sensors [4], and prepare next generation receivers with $\sim$ 10,000 such sensors. To cope with these advanced focal planes, and substantially larger sensor counts, a similarly new, advanced, readout system is required. Presented in this thesis is a space-flight representative implementation of one such readout system: hardware that was developed in parallel with the hardware that will be deployed to read out 3rd-generation ground-based telescopes with $\sim$15,000 sensors (SPT-3G) [5], and $\sim$ 7,500 sensors (POLARBEAR2) [29].

1.3 Bolometers

Before discussing the particulars of the readout system, let us turn to the sensors they are designed to operate: Transition Edge Sensors (TES). TES detectors are a mature technology that has overtaken the previous forms of bolometers used to observe in the microwave because

[5] One such favorable place viewable from the southern hemisphere is known as the Southern Hole, so called because it is relatively free of polarized dust emission.

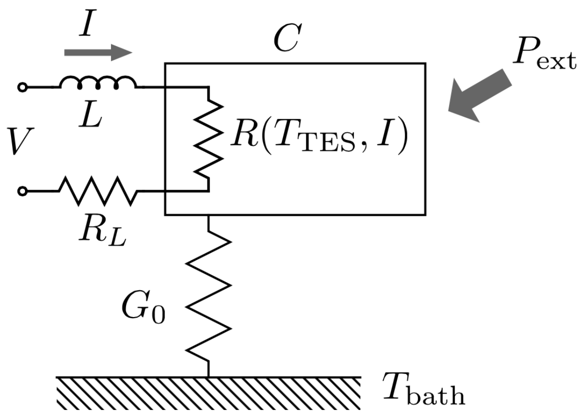

Figure 1–6: A cartoon showing a bolometer weakly coupled to a thermal bath, and receiving both optical power from the sky, and electrical power from a bias voltage. [19]

of their sensitivity, low noise, and scalability. Though they have never flown on a satellite platform, in 2013 the EBEX experiment demonstrated their performance in a space-like environment during a high altitude long-duration balloon flight over Antarctica [3]. In a TES, incident light from the sky is antenna-coupled to an absorber that is weakly linked to a thermal bath (Figure 1–6). A superconducting metal filament (the TES) acts as a thermistor – measuring the temperature of the absorber. This temperature is linearly related to the instantaneous incident power, and allows us to map the power on the sky. The term "bolometer" encompasses the absorbing element, sensor element, and implies a weak link to a thermal bath. A "pixel" can refer to a combination of bolometers that either share the same antenna-coupling or feed-horn, but measure different aspects of the absorbed radiation – be them polarizations, or frequency bands.

The choice of alloys from which the sensors are made vary: previous incarnations have been thin films of aluminum-titanium, the next generation of South Pole Telescope sensors are a gold-molybdenum alloy. These metal films are engineered to enter their superconducting transition at cryogenic temperatures, typically between 400 and 500 mK, with a thermal bath base temperature of 250 mK. The sensors are kept on the edge of their superconducting transition by a combination of optical power from the sky and electrical power provided as a voltage bias.

A TES is an incoherent detector – it measures incident optical power. When incoming radiation is absorbed, the bolometer undergoes a small fluctuation in temperature. The response of a bolometer to the absorption of optical power from the sky can be roughly modeled by Equation 1.1

$$T_{bolo}(t) = T_{bath} + \frac{P_{sky}}{G}\left(1 - e^{\frac{t}{\tau}}\right) \ . \tag{1.1}$$

The time constant τ equals the quotient of the heat capacity of the bolometer C, and the thermal conductivity between the bolometer and the thermal bath, G. From a readout perspective, the time constant defines one of the most important characteristics of these detectors: the bolometers cannot respond to any signal faster than their time-constant, enabling the use of a sinusoidal voltage bias, instead of a strictly DC voltage bias. The sinusoid is filtered out of the response but still deposits electrical power on the bolometer. Typical bolometer time-constants we've operated with are between 1 and 10 ms.

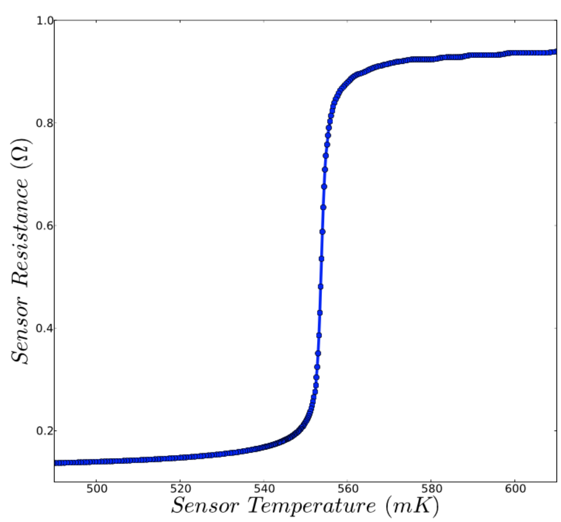

Figure 1–7: A model of the resistance of a TES as a function of temperature. [25]

The utility of these detectors comes from the remarkable steepness of the superconducting transition. The resistance as a function of temperature is shown in Figure 1–7. On the "edge", a small change in temperature results in a very large change in resistance – which is in turn measured by the change in current through the bolometer circuit. This detector technology allows for sensitivities to changes in power of tens of $\frac{\mathrm{aW}}{\sqrt{\mathrm{Hz}}}$.

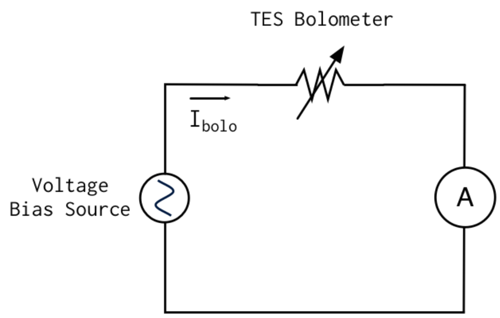

Figure 1–8: A circuit cartoon of a voltage-biased TES bolometer. Because of the fixed voltage bias, any change in the resistance of the bolometer results in a change in the current through the circuit.

The edge of a superconducting transition is not a very stable place. It is only by the application of both optical and electrical power (with a constant voltage) that bolometers can be kept in such a state. We can define the total power deposited on the bolometer as a combination of optical, and electrical,

$$P_{bolo} = P_{optical} + P_{electrical} \ . \tag{1.2}$$

Consider the bolometer circuit as in Figure 1–8 such that the source of electrical power is a fixed *voltage* bias

$$P_{bolo} = P_{optical} + \frac{V_{bias}^2}{R_{bolo}} \ . \tag{1.3}$$

The fixed voltage bias creates negative electro-thermal feedback: when the bolometer absorbs optical power and increases in temperature, its resistance also rises; in response, the electrical power deposited decreases proportionally. When optical power decreases, the inverse happens. The result is a stabilizing effect that keeps the total power constant, and the bolometer in its transition.

If instead we were to provide electrical power by way of a fixed *current* bias, and allow the voltage across the bolometer to vary

$$P_{bolo} = P_{optical} + I_{bias}^2 R_{bolo} \ , \tag{1.4}$$

the bolometer would have strong *positive* electro-thermal feedback. Deposited optical power would drive the bolometer higher into its transition, raising its resistance. In response the current bias would provide even more electrical power, instead of less as in the case of a voltage bias. A bolometer that is current biased is unstable, and will be driven out of the super-conducting transition when optical signal is detected. Bolometers that are strongly current biased are unsuitable for our type of observing.[6]

The total impedance of the circuit in series with the bolometer is made up of the parasitic impedance of everything that is not the bolometer, as well as the series impedance of the bolometer itself. In the absence of parasitic impedances, our circuit can be described accurately by Equation 1.3, and there is clear negative electro-thermal feedback. A more realistic model for us to consider is a fixed voltage bias provided to a circuit that has parasitic impedances in series with the bolometer. For any given bias frequency, the impedance of those elements is fixed, and they create a baseline current, I_{par} through the circuit enclosing the bolometer. I_{par} is a non-trivial function of the parasitic impedances in the circuit, the bolometer resistance, and the voltage bias provided. Without specifying anything about

[6] Current-biased bolometers are not useful for measuring power, but are very sensitive photon-counters. While this behavior is unsavory for us, they have been successfully used this way for high energy spectroscopy [17].

I_{par} aside from the fact that it is non-zero, we can generalize its effect on the circuit in the following way

$$P_{bolo} \sim P_{optical} + \frac{V_{bias}^2}{R_{bolo}} + I_{par}^2 R_{bolo} \ . \tag{1.5}$$

A fixed voltage bias in the presence of parasitic series impedance can therefore be modeled as a fixed voltage bias in addition to a small current bias I_{par}. This is known as a "mixed bias". As the parasitic impedances become large compared to the resistance of the bolometer, the bolometer becomes increasingly unstable. This reveals one of the most onerous aspects of operating bolometers – any device used to amplify their signal in order to read it out must have a very low input impedance, considerably lower than the normal resistance of the bolometers. Considering that this normal resistance is approximately 1Ω, this is an incredibly stringent requirement. Currently, the only devices that meet it, with sufficiently low noise, and that can be operated in cryogenic environments, are Super Conducting Quantum Interference Devices (SQUIDs).

1.4 SQUIDs

Unlike familiar and forgiving analog solid state amplifiers, SQUIDs are rather complicated to operate. The details of SQUID operation, and curious limitations imposed by their use, dictate a significant portion of the readout electronics design. For a more detailed examination of the Josephson Effect and SQUIDs I recommend J. Clarke's SQUID handbook, [14]. For now I note only what will be important when explaining their influence on the design of the readout electronics.

SQUIDs are extremely low-noise transimpedance amplifiers, with a typical input impedance of 150 nH [7]: they sense a change in current through their input coil and generate a proportional change in voltage at their output. The SQUIDs themselves are composed of two elements. The first is an inductive input coil that is in series with the bolometers, which produces a magnetic field as a result of the current through it. The second is a set of superconducting loops with Josephson Junctions, which sense changes in magnetic flux through them and convert these to a voltage. When the input coil is put in series with the bolometer circuit, as in Figure 1–9, perturbations in current through the bolometer circuit can be sensed by the SQUID.

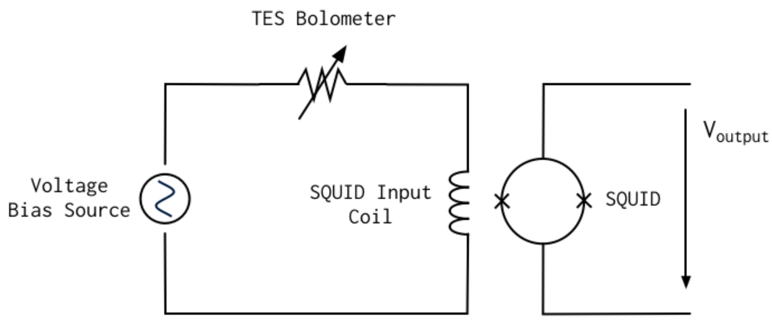

Figure 1–9: A circuit cartoon of a voltage-biased TES bolometer read out by a SQUID.

A Josephson Junction is a set of two superconductors separated by a thin insulating barrier; Cooper-pairs in the superconductor can flow through the insulating junctions via quantum tunneling [15]. Despite the presence of the insulator, and the current flowing through it, resistance of the two superconductors and junction remains zero unless a "critical current" is exceeded. When that critical current is exceeded the superconductors undergo a very rapid state-change and their resistance becomes non-zero: the result is a voltage across the junction [28]. SQUIDs take advantage of this phenomenon, known as the "DC Josephson" effect. External magnetic fields produced by the input coil induce screening currents in the superconductors. Current flows through the junctions until the critical current is exceeded, at which point the superconductors change to a normal state and the current is dumped through a shunt resistor of low fixed resistance. The voltage generated across that resistor is the SQUID output. What makes the SQUID response function odd is the particulars involving the behavior of the screening current in response to the external magnetic fields [14]. Consequently, these are highly non-linear analog amplifiers, with dynamic range and noise

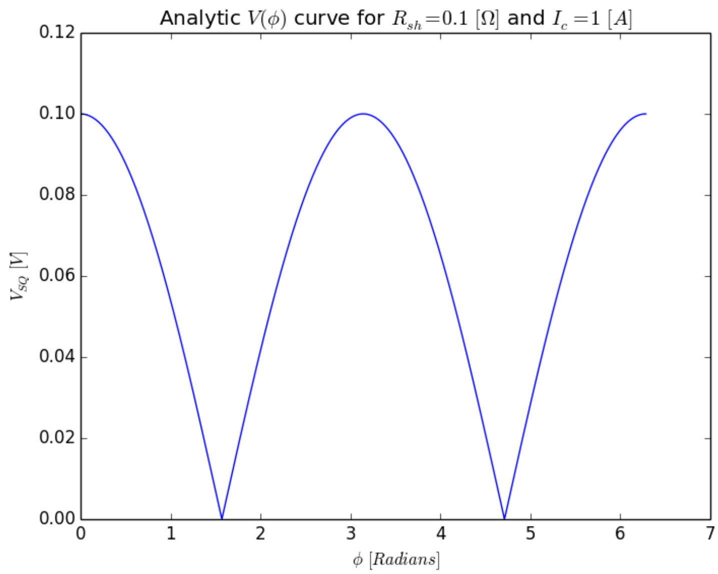

Figure 1–10: An analytic $V(\phi)$ curve, which shows the SQUID response to changes in magnetic flux.

properties that both depend in some way on the input amplitude. The analytic expression of that response function depends on the critical current (I_c) of the SQUID, which is a property of the Josephson Junctions and their particular construction; and the shunt resistance R_{sh}, which converts the current through the superconducting loops into a voltage,

$$V(\phi) = R_{sh} I_c \sqrt{1 - \sin(\phi)^2} \tag{1.6}$$

where ϕ is the magnetic flux through the input coil, a proxy for the current through the bolometer circuit. This function is reproduced as a curve in Figure 1–10.

In order to operate SQUIDs, each time they are cooled two DC biases must be provided and tuned for optimum performance (Figure 1–11). The bias parameters cannot be determined a priori because they depend in part on dynamic factors such as the magnetic flux through the SQUID *when it transitioned into a superconducting state*. This will vary according to the external magnetic environment in ways that are predictable (the orientation of the telescope), and unpredictable (contamination). Because of this, the process of

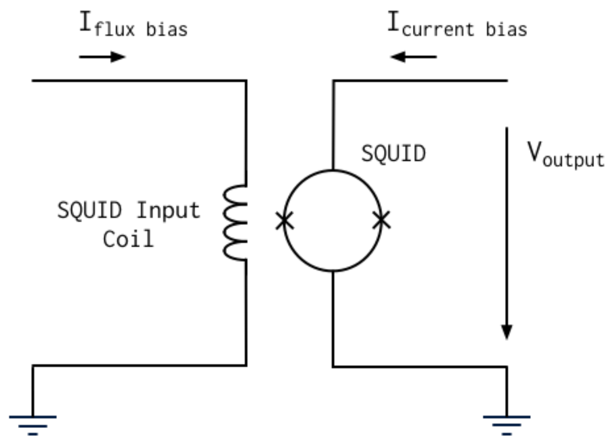

Figure 1–11: A schematic diagram showing how the two SQUID biases are applied. The flux bias is a current through the input coil, the current bias is a current across the SQUID elements themselves.

searching the parameter space by tuning both of these biases is a regularly repeated activity when operating these instruments. A current bias adjusts the baseline current through the Josephson Junctions; this adjusts the achievable peak-to-peak output of the voltage from the SQUID. A flux bias provides current through the input coil, adjusting the static magnetic field through the superconducting loops; this bias is used to maximize the linearity and sensitivity of the SQUID.

Figure 1–12 shows the how adjusting the current and flux biases affect SQUID response. The peak-to-peak amplitude of the response function is a measure of the range of SQUID output voltages over a fixed flux interval. Adjusting the current bias to increase the peak-to-peak amplitude of the curve amounts to increasing the dynamic range of the SQUID, this is demonstrated in Figure 1–12a. The choice of flux bias selects a single point on the $V(\Phi)$ curve, favored for local linearity and good sensitivity (a large slope). An optimal selection of this point is demonstrated in Figure 1–12b.

All told, the combination of bolometer and SQUID allows us to measure a change of power on the order of tens of attowatts on the bolometer and convert that at the input to

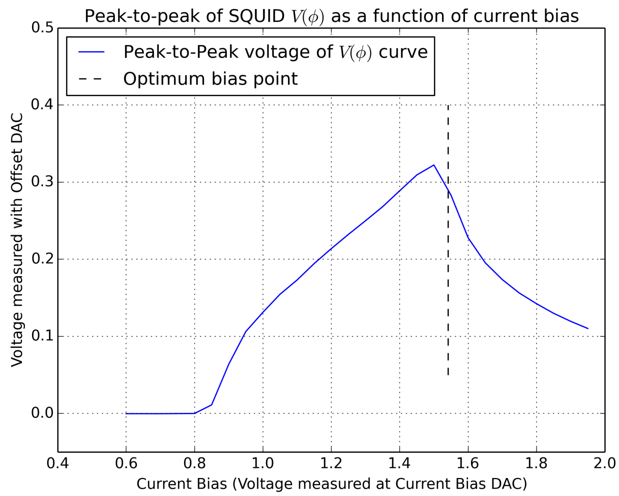

(a) The peak-to-peak amplitude of a SQUID response is a non-trivial function of current-bias, shown here. The black vertical line indicates the optimum current bias point.

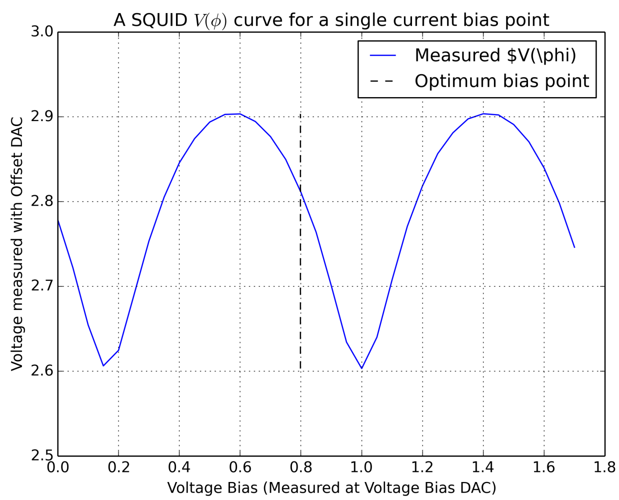

(b) A $V(\phi)$ curve taken at a single current bias value. The black vertical line indicates the optimum flux bias point.

Figure 1–12: The SQUID tuning process involves adjusting the current bias (a) and voltage bias (b) to optimize sensitivity / dynamic range and linearity.

the warm electronics to voltage fluctuations on the order of nanovolts.

1.5 Readout Design

In order to effectively operate and read out a bolometer we have shown that we must be able to deliver a separate voltage bias to each bolometer, and receive the voltage outputs of the amplifying SQUID. In general, the role of readout in all detection electronics is similar – to communicate control commands to, and information from, the detection instruments with as high a fidelity as possible. In the absence of any other considerations, resources would be devoted to precisely controlling and reading out each detection element individually: thus mitigating the introduction of correlations between the individual devices, and tailoring the control parameters to maximize the efficiency of each element. This mandate typically yields to considerations of scalability, cost, and complexity. While each of these will be important, the principal design of our readout system is driven by a constraint peculiar to cryogenically operated instrumentation: the difficulty of maintaining, at cryogenic temperatures, a large thermal load that is constantly conducting heat from the warm surrounding environment.

The cryogenic requirement for these devices is driven not just by the design of the TES (whose transition temperatures can be adjusted), but is actually integral to observing the CMB. Perhaps the easiest way to see this is in terms of thermal emissivity: the CMB is a blackbody of approximately 3 degrees kelvin that we detect in the microwave. For our instrument to see anything but its own thermal emission, the detection elements must be kept at a temperature below 3 degrees kelvin. This is one reason why our bolometers operate at sub-kelvin temperatures.

There are three obstacles to overcome when cooling an instrument to these temperatures, all of which derive from the fact that it does not exist in isolation from the exterior environment – several hundred degrees kelvin hotter. The first, convection, is almost completely neutralized by housing the instrument in an evacuated vessel. Without a buffering medium of gas, the remaining vectors for thermal loading are radiative and conductive. These are addressed by creating, within the cryostat, separate temperature stages bridging the gap between the coldest stage and room temperature. Though to some extent radiative loading is unavoidable – the cryostat looks out at the sky – radiative transfer between elements within the cryostat is minimized by the application of multi-layered insulation (MLI). This insulation is reflective and has a low thermal conductivity between each layer, obscuring each stage from "view" of the other. The most devastating mechanism of thermal loading is conduction of heat across materials within, and making up, the cryostat. These must inevitably construct an unbroken path from even the coldest stage out to room temperature.

The first of these is the mechanical structure supporting the cryostat and its various inner stages. Gradients in temperature across each stage, and in particular the coldest stage, are problematic for the consistency of the instruments. Consequently, the mounting platforms within the cryostat stages are constructed from metals with high thermal conductivity. However, conduction through some elements of the cryostat can be attenuated using standoffs of material with a very low thermal conductivity. Carbon fiber is a popular choice. This technique dramatically cuts down the conductive thermal loading from structural elements; though this source of thermal conduction remains a substantial contribution to the total thermal load.

Sadly, we are obliged to provide an additional source of conductive heat transfer, comprised of many efficient paths directly to our most sensitive components: the electrical wiring to bias and readout the detectors. As seen in Table 1–1, wiring conductance becomes the dominating factor, especially at the colder stages. Thermal conductivity across a wire depends on the size of the wire; its material; and of particular importance, is inversely proportional to the length of that wire. While this seems like a way out for us, unfortunately, long wire lengths also introduce parasitic resistances and inductances, and in some cases can threaten the stability of the readout electronics. We therefore have a strong incentive to keep the wire lengths as short as possible.

Component of Heat Budget	50K Stage	4K Stage
Radiation	71%	3.5%
Wiring Conductance	10 %	64.3%
Structural Conductance	19%	32.2%

Table 1–1: Expected thermal budget for the 50K and 4K stages of the South Pole Telescope (3G) receiver. Numbers taken from internal collaboration technical documents, provided by Brad Benson.

The cryogenic devices used to achieve temperatures between 200mK and 500mK vary in their technologies, methodologies, and achievable base temperatures. For each strategy there is a careful balancing act that rests on the fact that despite the relatively high energy and resource cost, most only have on the order of a a few tens of microwatts of cooling power at the coldest stage, a few hundred microwatts of cooling power at buffering sub-kelvin stages, a single watt of cooling power at the 4K stage, and on the order of a few tens of watts the 50K stage to extract the heat being conducted in. The effective duty-cycle of a cryogenic system is described by the time during which the coldest stage is at base temperature (the "hold time"), and the amount of time required to cycle the system after the energy budget is expended (the "cycle time"). Increasing the cold load proportionally decreases the hold time, and therefore the efficiency of the device. By way of example, the South Pole Telescope has a hold time of approximately 36 hours (depending on optical loading and inclination of the cryostat) after which it must cycle the fridge for four hours.

Returning now to the spirit of the original mandate at the start of this section, we would have each detector element get its own set of wires going into the cryostat and coming back out, as well as its own SQUID. In that scheme the total thermal load associated with a single detector element is dominated by the wiring and the SQUID reading it out. When the total thermal load of the bolometer circuit is small compared to the thermal load of the infrastructure this is relatively insignificant, but it does present a scalability problem.

The first generation of TES bolometer focal planes had sufficiently few detectors that this was acceptable. As larger focal planes became feasible, an increasing proportion of the total thermal budget had to be expended on the thermal load associated with the wiring. For our detectors there is a direct correspondence between observational efficiency and the number of detectors on the sky. A 10% increase in observing (or, approximately, hold) time is equivalent to a roughly 10% increase in the number of detectors on the focal plane. Likewise, a 10% increase in detector number that incurs a 10% penalty in observing efficiency is zero-sum.[7] This fact is what motivates the design of the readout – it must provide a way to reduce the cold load while adding more detectors.

To this end, "multiplexing" the control and readout works by chaining many TES elements together, to form a bundle. This bundle, or "comb", is constructed in such a way that only a single pair of wires per comb traverses the delicate path from 250mK to the room temperature readout electronics, but each sensor may still be read out and have some degree of autonomy from the others. The particulars of how this can be accomplished vary. The implementation presented here is known as the Flight-Representative 64x Digital Frequency Multiplexing (DfMUX) readout system. This is a variation on a new 64x DfMUX system developed for the next generation of ground-based CMB telescopes requiring multiplexing factors of 64 and beyond. A progenitor to the DfMUX system was the Analog Frequency Multiplexing (AfMUX) system, first implemented by the APEX telescope in the Atacoma Plateau in Chile in 2005. With a 280 sensor focal plane they struck this scalability problem, and had to devise a way to decrease the thermal cost per detector element. APEX was able to multiplex together groups of 8 readout channels (7 bolometers plus a calibration channel) on a single pair of wires through to the final cold stage [26].

[7] This is a consequence of the fact that our detectors are incoherent and photon-noise limited.

Since APEX was deployed, advances in detector fabrication, and reductions in the cost to produce them, have resulted in a substantial increase in the number of bolometers that can be made into a focal plane. While planned 3rd generation CMB experiments will have focal planes of >10,000 sensors, a 50 fold increase over 10 years, cryogenic technology has not kept pace. This fact makes multiplexing not merely an argument of efficiency, but a necessity. In order to cool down and operate 3rd generation CMB focal planes such as those being developed for POLARBEAR and SPT-3G, the multiplexing factor must be improved dramatically. The 64x DfMUX system a factor of 8 improvement over the original analog multiplexing employed on the APEX telescope. The planned SPT-3G receiver will employ a multiplexing factor of 64, such that 256 SQUIDs will read out approximately 15,000 TES devices.

1.6 Flight-Representative 64x DfMUX

Developed in parallel with the ground-based hardware, the flight-representative 64x DfMUX readout was commissioned by the Canadian Space Agency to provide the performance of state-of-the-art ground-based readout systems, while also meeting strict space-qualification criteria. These criteria fall into three categories: limits on power consumption, thermal resilience and stability requirements, and the requirement that electronics packages withstand the constant bombardment of cosmic radiation.

Today, scientific satellites that are not destined to travel between planets are powered almost exclusively with solar panels. These must be able to power all of the readout electronics, as well as the telescope control, communications equipment, avionics, and all of the redundancy for these systems. Power constraints manifest as a per-channel power consumption cap, calculated to ensure that the multiplexing capabilities can be matched with a competitive focal plane, while staying within a space-platform power budget. The original AfMUX system consumed 3 watts of power per readout channel; the most recent 16x

DfMUX system (flown on the EBEX long duration balloon flight) had a per-channel consumption of 300 mW. Our constraint was 60mW, and we achieved 49 mW. In some ways we benefit from putting the bolometers in space rather than on the ground (this is covered in more detail in Section 2.6), but approximately 50% of that power is consumed by the digital operations performed in firmware. Making those operations more efficient meant a paradigm shift in the way the sinusoids used to biased bolometers are generated, which is discussed in Chapter 3. The next big improvement to the power consumption came from changing the way bolometers were lowered into their superconducting transitions: a high-power mode was used only to do the initial biasing, and after the bolometers were lowered into the transition we shift into a lower-power mode. Software development for some of these new techniques, including that particular hand-off, was one of the ways I contributed to the project.

Thermal requirements for circuit boards in space are different than those operating in an atmosphere, and their effect on the design and appearance of the boards is immediately apparent.The form factors of the flight-representative circuit boards are approximately double their ground-based counterparts, with considerable spacing between large components. While the ground-based 64x readout system can rely on a very powerful set of fans to cool the motherboards and mezzanines, in space the loss of convective heat transfer makes it difficult to keep the components from overheating. Overcoming this means designing the layout of the circuit board to encourage radiative heat transfer, and making it easy to implement conductive-cooling infrastructure. Large areas of the PCB surface near power-consuming elements are metal, to keep the components mounted in the center from burning themselves up. If lack of atmosphere is the first environmental control problem posed by space operation, extreme temperature variations are the corollary. Although electronics are typically insulated as well as possible from the thermal variations associated with passing in and out of direct sunlight, dramatic swings are unavoidable. To redesign the layout of the PCBs, and conduct environmental testing, we partnered with COM DEV International, a Canadian-based space equipment company. Figure 1–13 shows an environmental chamber at COM DEV where thermal profiling over the range from -20C to 40C was conducted. Some results from the testing are included in Appendix 8.

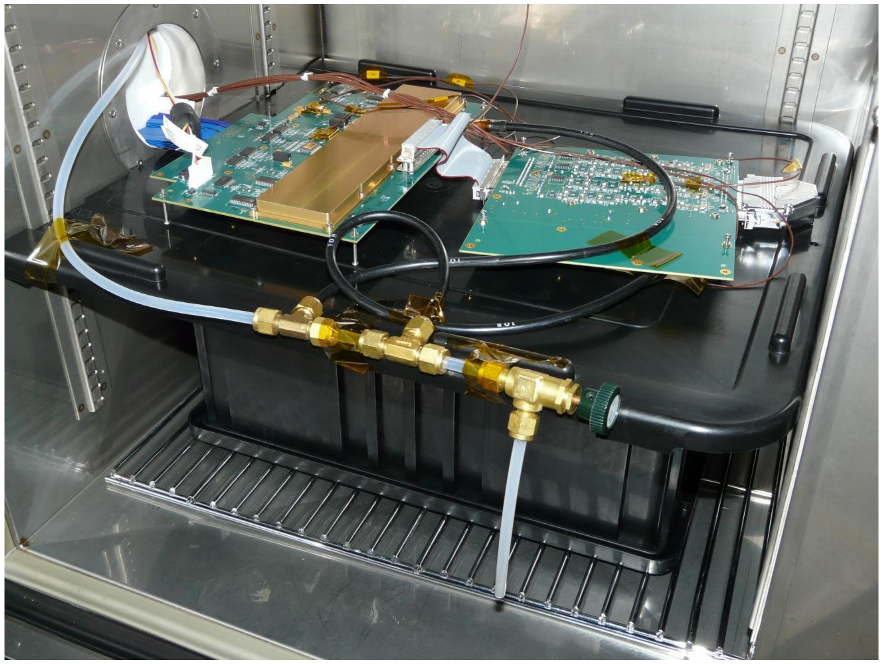

Figure 1–13: A photo of the thermal testing configuration in the COM DEV Tenney Chamber.

The radiation hardness requirement stems from the fact that any electronics onboard a space-craft will be exposed to significant levels of sustained radiation. These can cause "Single-Event Effects", or glitches in electronic equipment, but will also simply degrade performance and shorten the operational lifetime of most components. The predominant source of radiation damage comes from ionizing radiation – where electrons are generated within the components by collisions with high energy photons via the photo-electric effect, or charged particles via elastic scattering. Any semi-conductor is susceptible to damage from ionizing radiation, and digital components in particular are subject to a host of peculiar issues, from "flip-flops" in memory registers and localized "glitching" effects, to "latchups" of components and complete failure. Operating in a high radiation environment results in high electronic noise and unpredictable behavior; eventually damage to the semiconductor lattices will ruin the gate performance of transistors and cause irrecoverable harm to the components.

Electronics systems as a whole can only be insulated from this radiation to a minor degree – it is unfeasible to fly the incredible excess in weight that would be required to fully shield them. The favored option is instead to radiation-harden essential components in order to design a device that is effectively radiation-resistant. Techniques to do this include building in redundancy to digital systems, and in particular memory, as well as using slightly different material. To guard against semiconductor failure, the wafer on to which silicon is typically deposited is swapped with better insulators ("Silicon-on-Insulator", SOI), including sapphire ("Silicon-on-Sapphire", SOS). Finally, shielding of individual chips reduces the problem of weight to a manageable scale, and many chips are sheathed in depleted boron. Because of the relatively high cost of testing and re-engineering, and the relatively small market for such devices, electronics with radiation-hard packages are often outdated and obsolete compared to the available consumer-equivalent analogs. The other strong force working against rapid modernization of space-based electronics packages is risk avoidance. An older chip that has flown on several satellites has proven its reliability, a variable in any engineer's calculus that undermines adoption of new arrivals.

The component of most consequence that was hampered by limited radiation-hard selection was the Analog-to-Digital Converter (ADC) on the mezzanine. The available space-qualified ADC components that met our noise and power criteria are only 12-bits, in contrast to the 16-bit versions used in the ground-based mezzanine. The dynamic-range loss and resultant increase in digitization noise by going to 12-bits had to be calculated to guarantee it was a safe choice. This calculation is described in Chapter 5. Another element of the design dictated by the dearth of radiation-hard components was that the most suitable Digital-to-Analog Converter (DAC), used to generate the voltage biases, was too power-hungry when running in its recommended configuration. We had to measure for ourselves that it could perform adequately when constrained to only 1/5th its typical output range, without adversely changing its noise characteristics. Its intrinsic noise also dictated the transfer function of the synthesizer chain. This is covered in Section 2.6.

DfMUX readout systems are split into three separate electronics boards, of which two were made as flight-representative versions: a motherboard (not included in the CSA flight-representative project), a mezzanine, and a SQUID controller board. The differences in these boards, why they are separate, and how they relate to each other and the flight-representative project is covered in Section 2.4. Aside from the challenges and additional testing described above, the process of commissioning this readout hardware is similar that of commissioning previous systems. Each device was incrementally evaluated on the bench-top, and then with existing hardware, before being tested in conjunction together, and finally end-to-end with cold hardware. In the case of the SQUID controller board, an adapter board was fabricated and firmware was modified to allow it to interface with older motherboards and mezzanines. In all there were three revisions to the non-flight-representative SQUID controller before fabricating a flight-representative version. The new mezzanine architecture is incompatible with older motherboards, and so were controlled using Kintex-7 FPGA Evaluation boards, provided by Xilinx Inc. Using the evaluation-boards, we could analyze mezzanine performance and conduct firmware development in an environment with already established and supported communication and control resources. In all, there were two revisions of a mezzanine before building a flight-representative version. These boards became the ground-based hardware for 64x DfMUX, to be used in the SPT-3G and POLARBEAR2 telescopes.

Evaluation criteria, beyond power-consumption and thermal stability, included the ability to successfully operate 64 bolometer channels (with as many bolometers as existing cold-hardware could facilitate), stably at depth-in-transition of 70%; and to ensure a readout noise contribution of less than $10\frac{\mathrm{pA}}{\sqrt{\mathrm{Hz}}}$. A detailed description of the system noise is given in Chapter 5. In February of 2014, having met our goals, we submitted our evaluation and hardware to the Canadian Space Agency. These were formally accepted, and the contract concluded. This document is a characterization of that hardware, with particular focus on the evaluation criteria related to bolometer operation and noise performance. My role in the project included some limited design prototyping of a circuit in the SQUID controller board, and simulation work, but was primarily in commissioning the boards. This entailed writing new, and modifying existing, software algorithms to operate bolometers and tune SQUIDs, as well as characterizing the performance of the hardware both on the benchtop and end-to-end with bolometers.

CHAPTER 2
Operational Principals and Hardware Review

So far we have motivated multiplexing itself, and introduced the basic cold components that are instrumental to detecting the CMB. In this chapter we will take a closer look at how we accomplish that multiplexing, the specific tasks required to perform it, and the hardware that we've designed to do so.

We start with a general overview of the technique and the implementation methodology employed here. The latter half of the chapter will show how these functions map onto physical electronics. This includes a general tour of the individual electronics boards, their purpose, and the motivations behind their design – highlighting specific elements that represent a departure from previous incarnations of the DfMUX system.

2.1 Techniques in Multiplexing

There are two different strategies being employed in the field of CMB cosmology to achieve the sort of multiplexing introduced in Chapter 1. Both involve bundling a number of bolometers through a single set of wires to cut down on thermal loading. The first method still only reads out one bolometer at a time, but rapidly switches between them. Provided the switching can occur fast enough, the results yield time-streams of each bolometer. For multiplexing factors of approximately 40, the resulting "frame-rate" (the frequency at which each element is revisited) is $\sim$20 kHz [13]. This technique is known as Time Domain Multiplexing (TDM), and has been used in a number of CMB experiments – including, recently, the balloon based "SPIDER" [21] and South Pole based "BICEP2" [20].

Our design allows instead for the continuous readout and biasing of every bolometer simultaneously. Rather than subdividing the elements in time, we do so in frequency. Every bolometer sits behind an LC resonant circuit in parallel with the voltage bias source, as in Figure 2–1.The voltage bias source generates sinusoids at the resonance frequencies of each LC filter, but each bolometer sees only its own "carrier" sinusoid, as the rest are filtered out. To bias each bolometer on a comb, the waveform produced by the voltage source is a linear superposition the individual carriers, and only requires a single differential pair of wires to carry it from the warm electronics into the cold-stage. Once at the cold-stage, the waveform is separated into the component carriers using the LCR filters. Similarly, the currents flowing through each bolometer, modulated by sky signal, are summed at the input of the SQUID they share, such that the output waveform from the SQUID is a linear superposition of all bolometer outputs. This signal is digitized and then digitally demodulated at each carrier frequency before being analyzed. When the frequency spacing between filters is sufficiently large, each component bias sinusoid (and output signal) may be thought of as independent elements. The carrier frequencies are chosen in the range between 200KHz and 5MHz.[1] As covered in Section 1.3, since the periods of these carriers are all well above the thermal time constant for typical bolometers, each bolometer effectively sees its own separate "DC" bias, whose amplitude is controlled independently of other neighboring bolometer channels. This technique in general is known as Frequency Domain Multiplexing (FDM). In our case the signals are generated and demodulated digitally, hence the terminology Digital Frequency Multiplexing (DfMUX).

2.2 Implementing FDM at High Multiplexing Factors

In Section 1.4 it was stressed that usable SQUID dynamic range is a valuable commodity, largely owing to non-linearities and physical limitations in their output range; it is therefore

[1] Higher frequencies pose separate challenges when designing the cold-hardware. The warm-electronics can support bias frequencies up to 10MHz, but to date bolometers have only been biased using frequencies as high as 5MHz.

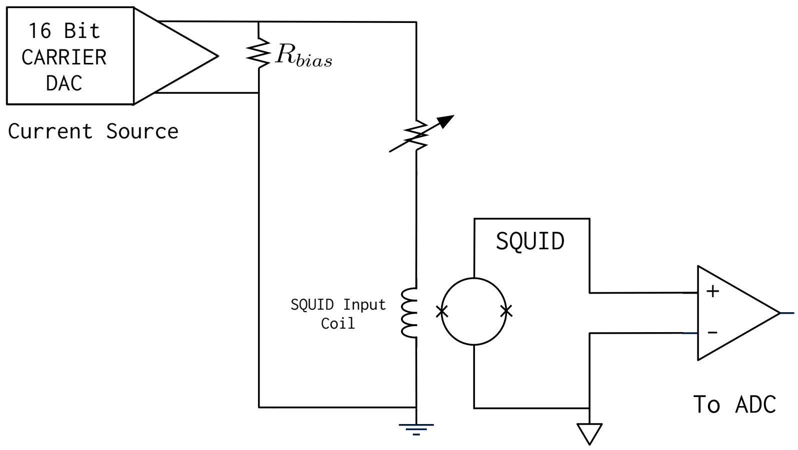

Figure 2–1: A cartoon circuit diagram where the bias sinusoid, and bolometer modulations to that sinusoid, are sensed directly by the SQUID.

a priority not to squander it. The biasing scheme introduced in Section 2.1, with a circuit design as laid out in Figure 2–1, has a flaw in this regard. The majority of that limited range is devoted not to sensing the perturbations we are interested in, but rather to the swings of the carrier sinusoids, which are themselves considerably larger than the sky signals. If actually implemented in this way, the SQUIDs would saturate at multiplexing factors of just a few.

Our solution is to insert a "nuller" signal at the input of the SQUID, to remove the large carrier sinusoids before they can be sensed. This means a second set of wires must travel between the warm electronics and the cryostat, as seen in Figure 2–3; but because the SQUID is mounted at only the 4 kelvin stage, the nuller traces do not directly load the coldest stage. In the original DfMUX system, up to multiplexing factors of 16, the nulling waveform was a fixed duplicate of the multiplexed carrier waveform, 180 degrees out of phase. This removed the carriers so the residual signals that remained were our desired signal perturbations. Nulling in this way was an efficient use of the SQUID dynamic range, reserving it for perturbations to the biases, which should be dominated by sky-signal.

Unfortunately, this method alone is inadequate for higher multiplexing factors, such as those required by 64x DfMUX. Each LC filter has a finite bandwidth, and to prevent crosstalk between channels these must be spaced sufficiently far apart from one another in frequency. Upper limits on multiplexing factor are in this way linked to the available bandwidth within which LC resonances may be placed. Recall that the input impedance of the SQUID comes from the inductance of the input coil, and is therefore sensitive to frequency. Even when biasing bolometers using relatively low-frequency sinusoids ($\sim$ 1 MHz), the series impedance of the SQUID is enough to cause destabilizing mixed-bias [6]. In order to suppress the effective input impedance and extend the usable bandwidth of the SQUID, a broadband analog-feedback circuit called a Flux-Locked Loop (FLL) was used. A FLL is a negative-feedback loop that feeds the output of the SQUID through an amplifier and across a resistor, before being shunted back into the SQUID input coil (Figure 2–2).

The "loop-gain" is a measure of the effectiveness of this sort of negative-feedback, and the Flux-Locked Loop had a selectable loopgain (with 3 options), between 5-10, which sufficiently suppresses the SQUID input impedance to operate bolometers. In order to advance to higher multiplexing factors, the FLL had to be abandoned: due to phase-shifts introduced by the wiring lengths, it drives the loop unstable at frequencies above $\sim$1.3 MHz, limiting the number of LC resonances that could fit in-band to approximately 16.

In order to open up higher bandwidths, and the door to higher multiplexing factors, in 2012 McGill introduced a digital feedback mechanism to replace the Flux-Locked Loop and shift some of the dynamic range demands from the SQUIDs to the warm electronics: this is called "Digital Active Nulling" (DAN) [6]. DAN has already been used at frequencies below 1.3 MHz on the EBEX experiment, and currently on South Pole Telescope Polarimeter. The 64x DfMUX hardware will be the first system without a Flux-Locked Loop circuit,

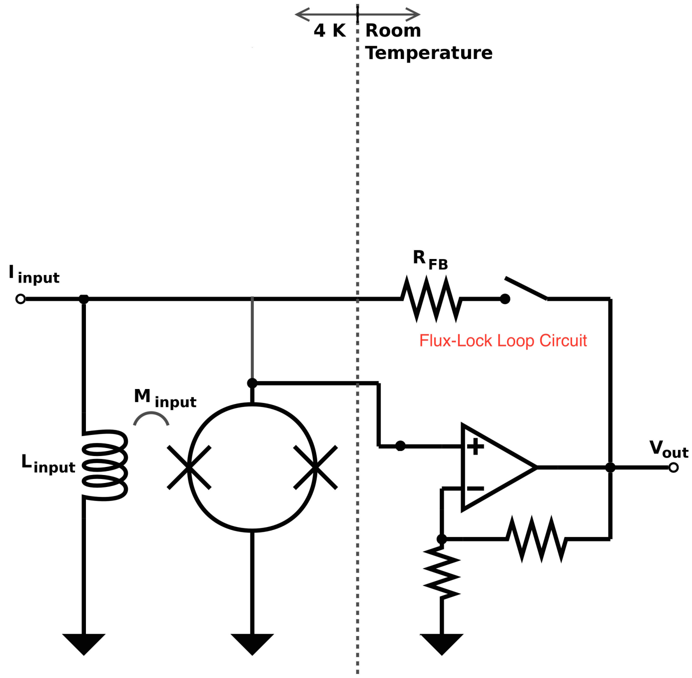

Figure 2–2: A circuit diagram showing the Flux Lock Loop circuit, where R_{fb} specifies that loopgain of the FLL. [7]

using DAN exclusively to read out bolometers at frequencies up to 5MHz. I will describe what DAN does only to the extent that is necessary here, for a comprehensive review, see [6].

2.3 "Digital Active Nulling" in a Nutshell

When operating with Digital Active Nulling, the nuller is more than a static mirror of the carrier waveform, it is actively modified according to measurements of the SQUID output, in order to create negative feedback that suppresses the SQUID input impedance at select frequencies (those being used to bias bolometers). Unlike the carrier, which is always a well defined superposition of sinusoids that have a fixed frequency, phase and amplitude, the component sinusoids of the nuller are being continuously adjusted to zero the output of the SQUID in select bands. As a result, this feedback *is not* applied with a constant loop-gain across the entire band, but rather in discrete channels with a fixed width in frequency, centered on the carrier bias frequencies. This digital form of feedback does not suffer from

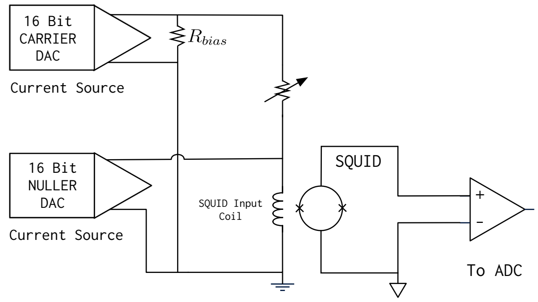

Figure 2–3: A cartoon circuit diagram where the bias sinusoid is removed using a "nulling" signal before reaching the SQUID, such that only bolometer modulations are sensed directly by the SQUID.

the same stability limitations that the Flux-Locked Loop does, and can effectively suppress the input impedance of the SQUID at discrete frequencies out to 10MHz.

To see why DAN suppresses the SQUID input impedance, refer to Figure 2–3: in the limit of perfect feedback there becomes a virtual ground at the input of the SQUID coil. In practice there must always be some residual current through the SQUID input coil (and therefore some effective impedance), else there would be no signal to feed back on. The degree to which the input impedance is suppressed is measured by the loop-gain of the feedback. DAN has an adjustable narrow-band loop-gain that is a strong function of frequency, and is applied within discrete frequency intervals. The shape of the DAN response can be seen in Figure 2–4. It is tailored to have high loop-gain across the entire bolometer response region, before rolling off. Where the bolometer response is strongest, and where we are most interested in it, is at frequencies out to approximately 100Hz, meaning DAN must have sufficient loop-gain at bandwidths of about 200Hz around the central frequencies.

While DAN is enabled, the record of the phase and amplitude of the individual nuller channels create time-streams analogous to the demodulated SQUID output time-streams. These "DAN-Channel" time-streams become the new data-product, as they are adjusting to

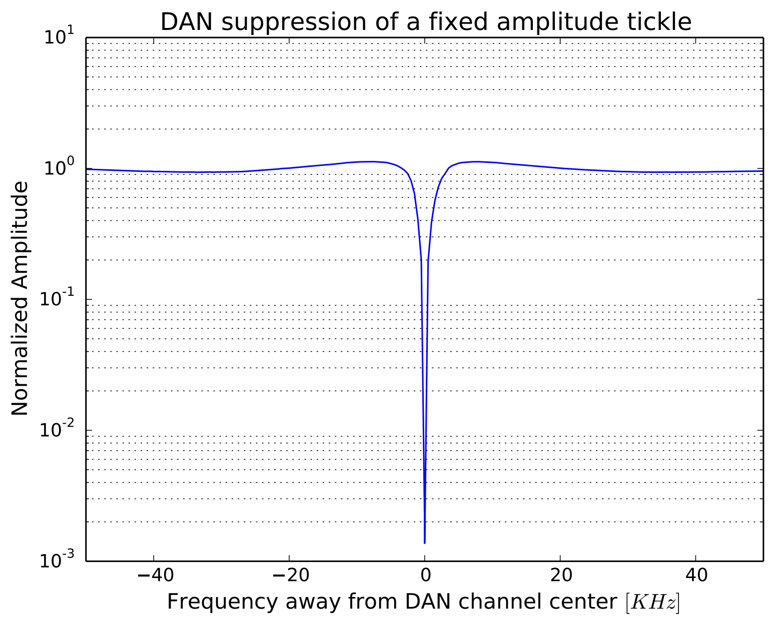

Figure 2–4: The DAN loopgain as a function of frequency can be seen here in the suppression of a fixed-amplitude tickle sinusoid as it sweeps across the DAN bandwidth.

null both the carrier sinusoids *and* the modulations to them from the sky power.

2.4 Hardware Overview

The 64x DfMUX readout provides the following general functions: hosting a user interface and streaming data; performing the digital computations required to generate the carrier sinusoids, demodulate the sky-signal modulated carriers, and calculate the DAN feedback; synthesizing the analog sinusoids; tuning SQUID bias parameters; amplifying the analog SQUID outputs; and digitizing those amplified signals. These tasks are broken down into three distinct electronics components: digital signals processing motherboard, synthesizer/digitizer mezzanine, and SQUID pre-amplifier and controller board (refer to Figure 2–1 for specifics). Of these three, only the mezzanine and the SQUID controller board are part of the flight-representative project.

The motherboard contains digital signal processing and networking hardware, which are responsible for streaming the data, hosting the API and user interface, and running the firmware that does digital calculations. It is not included in the flight qualified project in part because FPGA (Field Programmable Gate Array, the signal-processing workhorse) technology changes rapidly, and the exact configuration of the motherboard communications and data handling on a satellite mission will likely be tailored to the project in important ways. In contrast, the SQUID controller board and mezzanine are relatively stand-alone devices that will not fall rapidly into obsolescence. The motherboard used for this project was a Xilinx Kintex-7 FPGA Evaluation Board.

Function	Electronics Board	Flight Representative
Hosting Control Interface	Motherboard	No
Performing Digital Computations	Motherboard	No
Bias and Nuller Synthesis	Mezzanine	Yes
Bias and Nuller Amplification	Mezzanine	Yes
SQUID Output Digitization	Mezzanine	Yes
SQUID Tuning	SQUID Controller Board	Yes
SQUID Output Amplification	SQUID Controller Board	Yes

Table 2–1: Division of labor between the warm electronics components.

The decision to physically split the functions described in Table 2–1 into separate electronics components is strategically motivated. On one hand, signals from the SQUID exiting the cryostat are still extremely weak, and should be amplified immediately to avoid attenuation or being drowned in RF contamination. This amplification happens within an RF-tight box attached directly to the cryostat. One the other hand, the FPGA, DACs, and ADCs together consume the bulk of the electrical power required by the system. The devices which regulate the power to each of these components are strong potential sources of electronic interference, and should be separated from that RF-tight environment. The FPGA also requires a high-speed connection to the synthesizing and digitizing components: a fast digital clock must be distributed from the motherboard to the DACs and ADCs, and the data transfered to and from them aligned with that clock. Minimizing the length these digital signals must travel is important. The consequence of all of this can be seen in Figure 2–5: the mezzanine is mounted directly to the motherboard (justifying its name) with a high data-rate connection, and isolated from the RF-tight environment, where signals are most sensitive to contamination. The SQUID controller board, residing within that RF-tight box, is connected to the mezzanines via a cable that transmits only analog signals during data acquisition.[2]

Each mezzanine has four separate modules, containing identical synthesis and digitization signal paths. There is then a one-to-one mapping between mezzanine and SQUID controller boards, which themselves have four identical signal paths and modules. Each module provides the signals to bias a full comb of bolometers, and the tuning, amplification, and digitization components to read out a corresponding SQUID. At the multiplexing factor of 64 that this system was designed for, a single mezzanine/SQUID controller pair can readout 256 bolometers using 4 SQUIDs. This hierarchy is shown in Figure 2–5.

2.5 The SQUID Controller Board

The SQUID controller board (Figure 2–6) provides some conditioning of the synthesis signals [3] , biases the SQUIDs, and amplifies the SQUID output enough to be transmitted to the mezzanine. Mentioned in Section 1.4, SQUIDs require two types of DC biases to be adjusted in order to operate in a linear, stable, low-noise region of high sensitivity: a current bias, and a flux bias. These are provided by two DACs for each SQUID controller board module. The current bias DAC applies a voltage, converted to a current with a series resistor, across the superconducting loops of the SQUID. The flux bias DAC does the same through the input coil. Both are 14 bit DACs, providing finely configurable biases.

[2] Low-power digital communications are also transmitted, but only during SQUID tuning.

[3] This is for the purpose of meeting noise criteria and matching the relative amplitudes of the carrier and nuller at the SQUID input. The former will be covered in more detail in the next section.

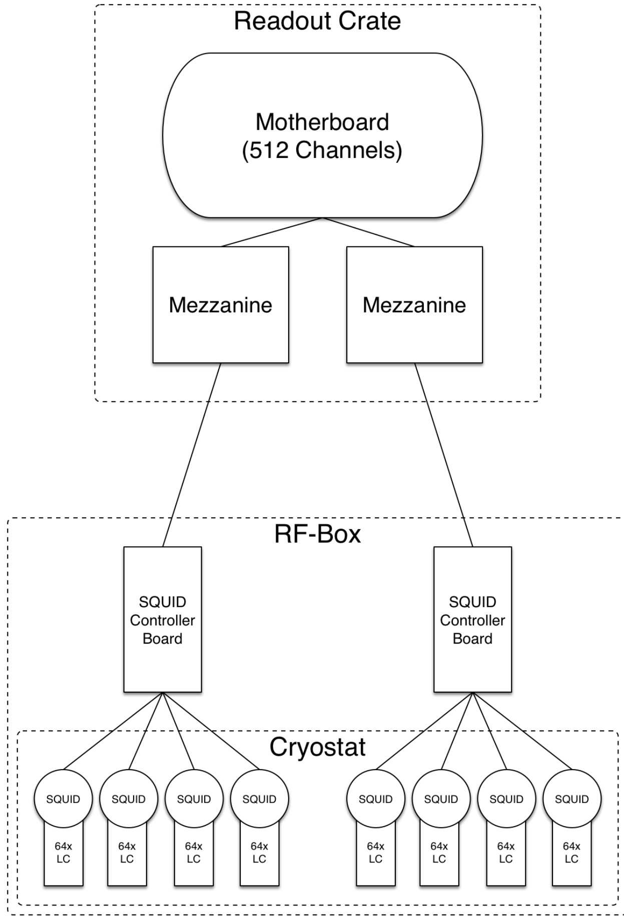

Figure 2–5: A block diagram showing the motherboard, mezzanine, and SQUID controller board connections into the cryostat. The topmost numerology of two-mezzanines-per-motherboard is correct for the existing ground-based motherboard, but is not a requirement. However, this configuration was designed to maximally take advantage of the resources in current FPGAs; it is likely the same calculus will be applicable to a space-flight configuration.

Because the total range of voltages that a SQUID can produce in general is much larger than the peak-to-peak voltage it can produce *at a single current bias* (refer to Figure 1–12a), there is an intrinsic mismatch between the range of all possible SQUID voltage outputs, and the range of expected *signal* voltage outputs. Naturally, we would prefer to match the rails of the downstream amplification stages to the expected signal amplitudes when operating at a fixed SQUID bias point. This necessitates an "offset" DAC, which does not affect the SQUID biasing parameters, but does translate the differential voltage across the first-stage amplifier (the first amplifier that the SQUID output sees on the SQUID controller board)

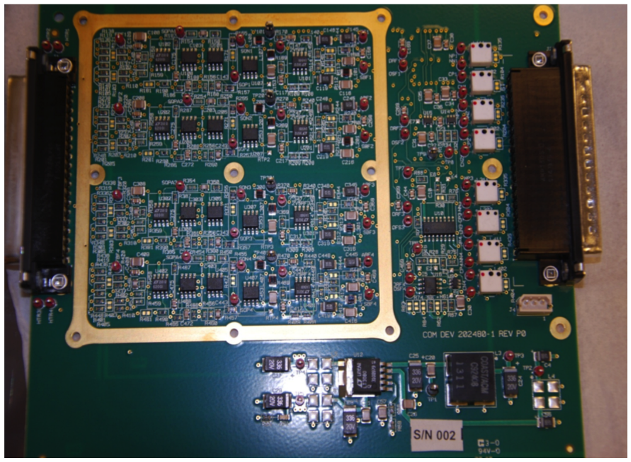

Figure 2–6: Image of the flight representative SQUID controller board.

up or down. Zeroing the first stage amplifier following a successful SQUID bias frees up the full dynamic range of the amplifier chain for signal readout, similar to taring a scale. These elements are all shown together in Figure 2–7.

In addition to biasing the SQUIDs, it is imperative to keep them fixed to that bias value. To guard against slow drifts in the SQUID flux bias point, the output of the first stage amplifier is used as an input to a low frequency integrating circuit that feeds back into the SQUID input (Figure 2–8). This acts as an exclusively low-frequency Flux-Locked Loop that prevents low-frequency drifts from shifting the flux-bias point over time, degrading the SQUID tuning. Our biggest concern in this regard comes from changes in external magnetic fields. Changing magnetic fields create a dynamic flux offset which can introduce a time-dependent transfer function, and confusing correlations with quantities such as pointing position. In ground-based operations this can happen when the telescope scans azimuthally through the earth's magnetic field lines.

The integrator circuit is enabled immediately following the SQUID tuning procedure, before bolometers are enabled, after zeroing DC level of the amplifier circuit with the offset DAC. Once enabled, the integrator will attempt to maintain the flux through the SQUID coils. Despite the fact that the integrator applies corrections to the SQUID input using the

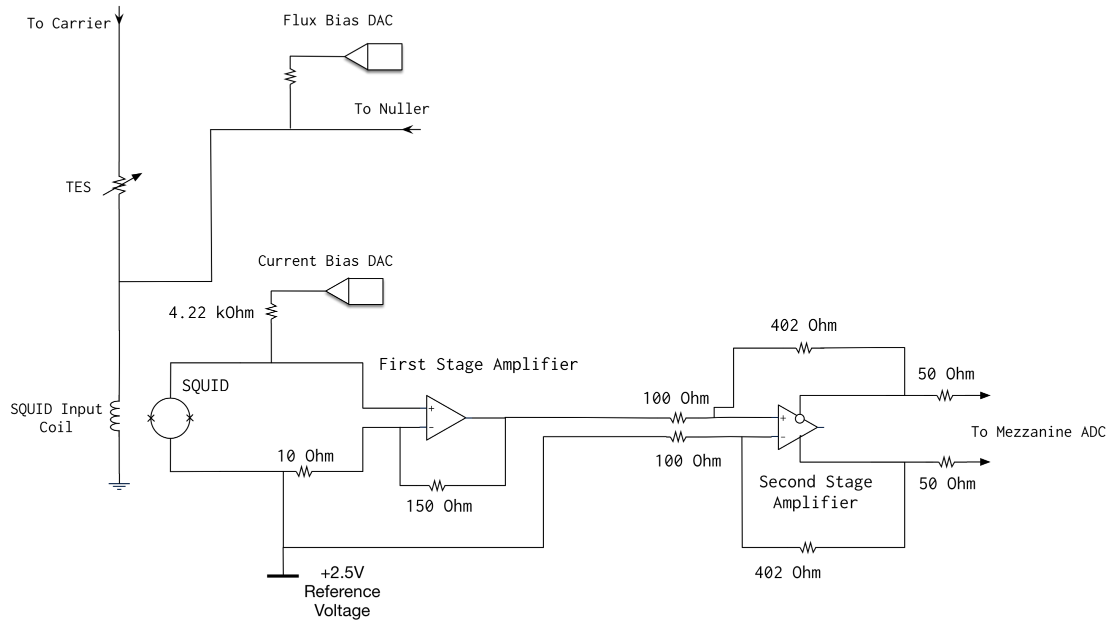

Figure 2–7: A cartoon circuit diagram of the SQUID-amplifying path that shows the biasing and offset DACs, and ends at the first-stage amplifier.

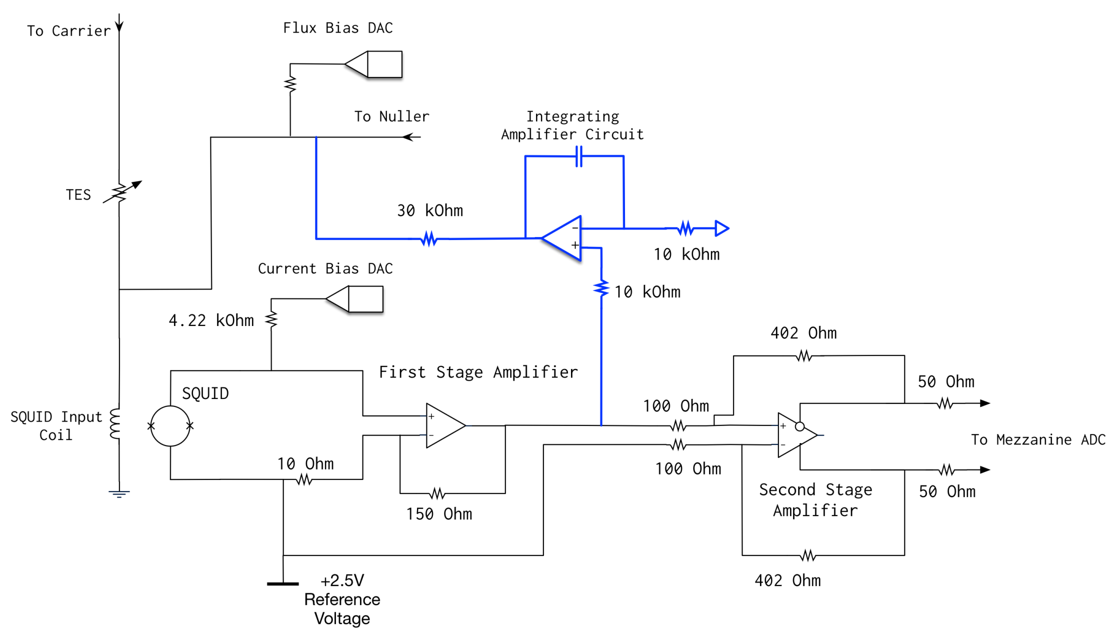

Figure 2–8: A cartoon circuit diagram that includes the low-pass integrator circuit, highlighted in blue.

nuller lines, these are analog corrections, applied downstream of the DAC. Since they are not reflected in the digital waveform generated by the FPGA, they cannot be mistaken for signal. Moreover the 3dB point for the filtering on the integrator is 10 kHz, far below the lowest LC resonance of the bolometer comb.

To monitor the operating point, and quantify the performance, of a SQUID on the fly, we can measure the SQUID transimpedance. This is the first derivative of the $V(\Phi)$ curve, and it defines the change in voltage at the SQUID output for a given change in current at the SQUID input. We can find this empirically by injecting a small sinusoid of known amplitude along the nuller lines and measuring the demodulated amplitude from the SQUID output. To verify the operation of the integrating circuit, we simulate a change to the external magnetic field by sweeping the flux bias DAC voltage after enabling the low-frequency feedback, and recording the amplitude of the transimpedance monitoring signal (Figure 2–9). The output of the SQUID, where in the absence of the integrator we would expect a $V(\Phi)$ curve, is flat.

2.6 The Mezzanine

The mezzanine performs the synthesis of the nuller and carrier waveforms, and digitization of the SQUID controller board output signals. Like the SQUID controller board, it has 4 independent modules that each contain the signal paths required generate carrier and nuller synthesizers and to digitize the SQUID output. Within each module, the signal paths for carrier and nuller synthesis are identical – using 16-bit DACs. These DACs are operated with an adjustable reference current supplied to them by separate 12 bit DACs. By adjusting the reference current they may be transitioned into and out of a lower-power mode. The digitization of each SQUID output channel is performed by a 12-bit ADC operating at 20MSPS.

As stated in Chapter 1, the signal synthesis on the mezzanine is the largest source of electrical power consumption after the FPGA. It is here where the balancing act between

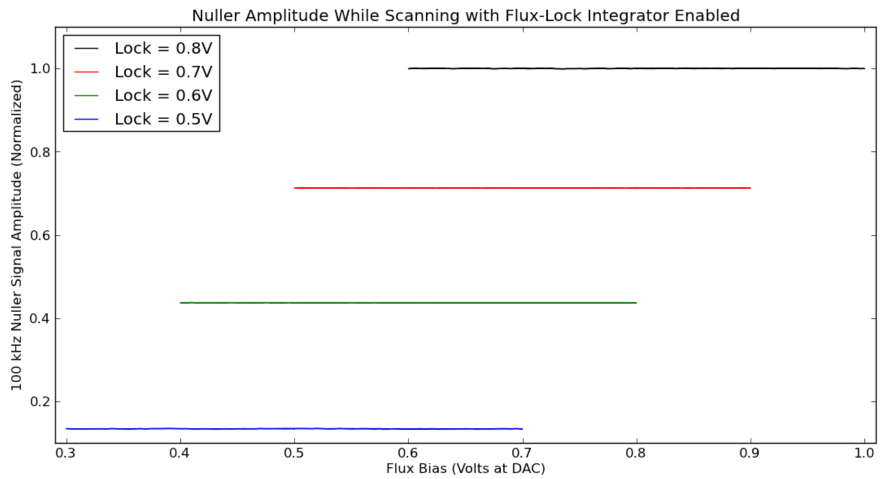

(a) The apparent amplitude of a fixed nuller signal as the SQUID flux bias is swept. Each colored curve is taken after disabling the integrator, shifting flux bias voltages, and then enabling the integrator, locking to those points.

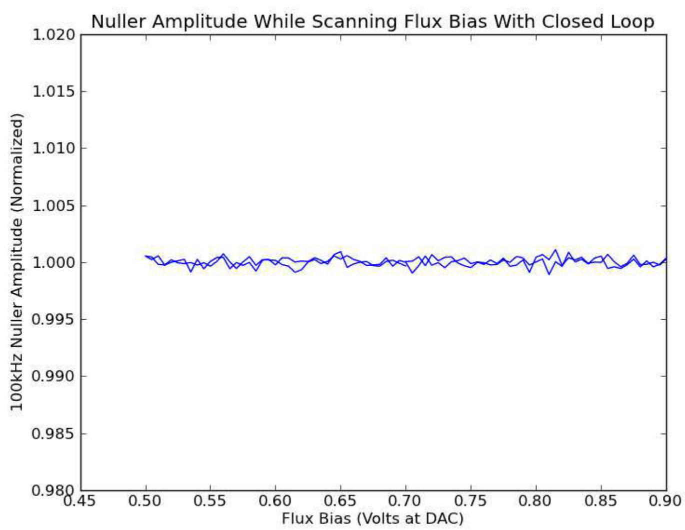

(b) Stability of the locked bias point when integrator is engaged. Note that deviations never exceed 0.25%.

Figure 2–9: A demonstration that the integrating circuit succeeds in locking the total flux through the SQUID input coil at low frequencies, and is stable. Image Credit: Amy Bender

Figure 2–10: Image of the flight representative mezzanine.

power constraints and the needs of the system play out. The readout system must provide enough power to bias bolometers, not exceed the power limitations imposed by space-based operations, and have a readout noise contribution that is no larger than that of current DfMUX systems. Since the current instruments are photon-noise limited, there are few gains from substantially improving readout noise – so the difficulty in the last constraint is in maintaining this noise performance while scaling up the multiplexing factor.

Aside from the SQUID, the carrier and nuller DACs are the largest sources of noise in the readout system; as such, the design of the mezzanine and downstream electronics is heavily influenced by the need to limit their contribution to the overall readout noise. Our noise constraint for the synthesis chain is derived from the intrinsic SQUID noise – as uncorrelated noises sources add in quadrature, there are quickly diminishing returns for improving any source of noise significantly below the one dominant source. For very high-performing SQUIDs this sets our allowable maximum DAC noise contribution, measured at the bolometer, at 3 $\frac{\mathrm{pA}}{\sqrt{\mathrm{Hz}}}$. Simultaneously, our stated power restrictions are to be below 60 mW/channel total – for 64x multiplexing, based on power-consumption breakdowns, this leaves something on the order of 500mW per DAC.

The 16-bit AD768 DAC chosen for this project has a maximum output swing of 20 mA Peak-to-Peak with an output noise of 60 $\frac{\text{pA}}{\sqrt{\text{Hz}}}$. To achieve the desired noise performance, the DAC output is divided down by a factor of 20 in the downstream electronics, leaving 1 $\text{mA}_{\text{P-P}}$ per comb. To conform to our per-channel power budget for satellite operations, the DAC must be further limited to 1/5th its maximum output swing – bringing its total power dissipation to 440 mW (note that there is no noise benefit from operating the DAC in this lower-power mode). After these constraints, the total available output current from the carrier chain at the bolometers is 200 $\mu\text{A}_{\text{P-P}}$, or just 70.7 μA_{RMS}. The question then becomes – is this sufficient power to operate a comb of 64 bolometers?

A single bolometer for a typical ground based telescope has a saturation power of 15 pW, and will require a voltage bias amplitude of about 3 μV_{RMS} for $R_{bolo} = 1\Omega$. This corresponds to a necessary minimum supplied current of 3 μA_{RMS} per bolometer. Fortunately, bolometers optimized for space-based applications avoid optical loading from the atmosphere, and so can be designed with lower saturation powers of typically 2 pW. As such, they require a bias of approximately 1 μV_{RMS} for $R_{bolo} = 1\Omega$; this sets a necessary minimum supplied current of 1 μA_{RMS} per bolometer.

A rule of thumb, allowing for component variation, and flexibility in the bolometer fabrication design, is to be able to provide approximately 3 times the minimum required bias power for an ideal comb. In the conceptually simplest, and practically most pessimistic, mode of operation every synthesized carrier sinusoid would, at some point in time, be phased-up. Operating this way would limit our multiplexing factor to far below the target goal of 64x, as doing so would be equivalent to operating with no safety factor at all. Fortunately, by randomizing the relative phases of each carrier frequency, the available bandwidth can be used far more effectively. This is a technique designed to accomplish what is known as Crest Factor Minimization.

2.6.1 Crest Factor Minimization

The crest factor (sometimes called the "peak-to-average ratio") is the ratio of the maximum excursion of a waveform amplitude to the root-mean-square (RMS) waveform amplitude.

$$\text{Crest Factor} = \frac{|A_{Peak}|}{A_{\text{RMS}}}. \tag{2.1}$$

Thus the crest factor for a waveform consisting of a single carrier sinusoid is $\frac{A_{Peak}}{A_{\text{RMS}}} = \sqrt{2}$. Crest factor minimization refers to the practice of reducing the crest factor through a judicious choice of the waveform parameters. In our case, the free parameters are carrier phases. For N carriers of equal amplitude A, the RMS grows proportional to $\sqrt{N}$. In the worst-case scenario, such that the all carriers are in phase, the waveform peak amplitude would be $N \cdot A$, and the crest factor would be $N \cdot A / \left(\frac{\sqrt{N} \cdot A}{\sqrt{2}}\right) = \sqrt{2N}$.

Crest factor minimization for purposes similar to these has been explored before in [8]; their algorithm adjusts the phase of many sinusoids to search for a set of parameters that provides the smallest possible crest factor. Our case is slightly different from the one best suited to the methods presented in [8]. Temperature-dependent complex stray impedances in the cold-electronics impose an element of randomness to the phase-shifts that carrier signals are subject to, which cannot be determined a priori. Also, in the case of the nulling comb, DAN is constantly making adjustments to the phases of each nuller. We are therefore uninterested in narrow global minimization, instead we look for the smoothness of the parameter-space to ensure the existence of broad minima. No simple analytic formula to determine a minimum crest factor exists in the literature that we are aware of. Instead, the problem is usually addressed with numerical techniques, wherein a set of phases is chosen randomly, and the crest factor is estimated by simulating the waveform and monitoring the RMS value and maximum excursion.

Our current practice of randomizing the phases of each carrier frequency at runtime has worked for lower multiplexing factors, and is far more efficient than having to use any precomputed set of phases. To ensure that this same practice can be used with the new multiplexing factors and tighter dynamic range constraints, I designed a set of simulations to sample the parameter space. To minimize inter-modulation distortion (IMD) product surprises, and generally control where inter-modulation distortion shows up, all carrier frequencies of the DfMUX system are required to be multiples of a base frequency – typically 117 Hz. This is approximately the extremum of our science band – which ensures no IMD appears in our science-band around each channel. It also makes it easy to calculate the crest factor of a waveform that is a linear combination of those base frequencies, as the summed waveform will have a repetition period of no more than 8.6 ms. For our 20 MHz sampling rate, this restricts the number of samples required to simulate a full period to 20 MHz/117 Hz = 170,941 – an easy computation. The simulations were designed to probe two qualities of the parameter space: first, that given random sets of phases, the variance in crest factor, and the typical value, were acceptably low; second, that the local parameter space around each point remains smooth with a similarly low variance.

Using these techniques I find a typical crest factor of 4.25 with a spread of $\sim 4\%$. This means the maximum waveform value for a 64 carrier waveform will be about $\frac{4.25\sqrt{64}\cdot A}{\sqrt{2}}$, which is a factor approximately $3\sqrt{64}$ larger than the maximum value for a single sinusoid; or equivalently $\frac{3}{\sqrt{64}}$ the pessimistic "in-phase" case.

Based on these results, using Equation 2.2, for a multiplexing factor of 64 we require a minimum DAC dynamic range of $68\mu A_{P-P}$. When the DAC is operated using 1/5 of its full range output swing, our total supplied dynamic range is 3 times the required minimum – consistent with the recommended safety factor of 3. When the DAC is programmed to provide its maximum output swing, that factor increases to over 14 times the required minimum.

$$\frac{A_{dr}\ \mu A_{P-P}}{2} = 3\sqrt{N}\cdot A_{Peak},\ \text{with}\ \ A_{Peak} = \sqrt{2}\cdot 1\mu A_{\mathrm{RMS}} \tag{2.2}$$

An increase in the achievable multiplexing factor in the cold-electronics would relax the power constraint per DAC (as it scales as power-per-channel), allowing the current mezzanine and SQUID controller boards to accommodate multiplexing factors beyond those which are being used in this project.

The design of the mezzanine thus addresses the three primary concerns with regard to noise constraints, power consumption constraints, and bolometer bias demands. Concerns over the DAC noise are addressed by dividing down the output noise with resistive networks. The degree to which this must be divided is independent of the dynamic range required of the DAC, and determined based on the SQUID noise – one of the dominant noise sources that cannot be improved by means of the warm-electronics.

Juggling the power consumption constraints with power demands of biasing a large comb of bolometers is accomplished by operating the DAC in a low-power consumption mode. We find this section that this low-power mode is both sufficient to meet our power consumption criteria and operate a comb of 64 space-rated bolometers with a safety favor of 3. The DAC power consumption, and thus dynamic range, is digitally adjustable, and in the event of a change in desired multiplexing factor, the DAC can provide up to 5 times more output without any modifications to the board, software, or firmware. This would not necessarily require concessions in the power consumption constraint – as a higher multiplexing would maintain a similar power-per-channel.

CHAPTER 3
Signal Production and Data Capture

Fundamentally, anything we are interested in controlling, or recording, in the DfMUX system is an analog signal. However, we interact with this analog system digitally: all of the carrier and nuller sinusoids are digitally computed in firmware on-board the FPGA, before being converted into analog signals on the mezzanine. Likewise, in order for the firmware to interact with the analog signals output by the SQUID, they must first be digitized, which also takes place on the mezzanine. Moreover, the analog waveforms that are produced and digitized are complex waveforms consisting of the superposition of many sinusoids – but the data-streams we wish to record are the baseband responses of each bolometer channel, and the manner in which we set nuller and carriers is channelized into the component sinusoids. Digital signal processing that takes place in firmware is responsible not just for generating and digitizing these waveforms, but modulating and demodulating them. The exact manner in which we do so represents one of the most significant innovations of this technology – and is largely the reason why we are able to operate a 64x multiplexing readout system on just a fraction of the total power consumption of any previous DfMUX hardware.

This chapter begins with an introduction to the digital data-streams that we have access to, before describing the signal processing techniques involved in generating synthesizer combs, and demodulating the resulting wave-functions down to base-band.

3.1 Data-streams

We have four windows through which we see and interact with the system – the first are down-sampled data-streams of the demodulated frequency channels (either from DAN, or directly from the ADC), which show us a narrow bandwidth of $\sim$ 100Hz around a central frequency. This is ideal for observing an individual bolometer, but does not contain much information about the system as whole. In addition to those, we are able to directly capture up to approximately 1.3 million contiguous samples from the carrier and nuller DACs, and the ADCs, at the full 20MSPS. It's difficult to overstate the utility of these snapshots. Looking at the waveforms going to the DACs and coming from the ADCs, at the full sampling rate, gives us perspective on the whole 10MHz bandwidth of the readout system. It also allows us to separate source signal products from contaminate signal products. As such it is our primary means of searching for noise sources within our band, including crosstalk, distortion, and RF pickup.

3.2 Modulation

The Legacy 16x DfMUX readout system employed Direct Digital Synthesis (DDS) within the FPGA as a means to produce the waveforms generated by the carrier and nuller. For DDS synthesis of a single sinusoid, discrete values of a normalized sine wave from 0 to $\frac{\pi}{2}$ are stored in memory on the FPGA, called Look-Up-Tables (LUTs), and then used to recreate a full sinusoid on-the-fly. The amplitude is provided by a programmable multiplicative factor, and the frequency of the sinusoid is determined by the rate at which these values are accessed and accumulated, a process known as Phase-to-Amplitude conversion. The density with which the interval from 0 to $\frac{\pi}{2}$ is sampled and stored in the LUT, and precision thereof, are such that this method achieves 4.7mHz frequency resolution.[1] Each frequency channel of any sort requires its own DDS "instance" (though they may share LUTs to some extent) – an allocation of memory and signal-processing resources within the FPGA to perform those operations. In the Legacy 16x DfMUX readout there were 16 frequency channels per DAC,

[1] This value comes from the fact that the Phase-to-Amplitude converters use a 32 bit index and provide values at 20MHz.

2 DACs per module, 2 modules per mezzanine, and 2 mezzanines per motherboard. This amounts to 128 individual synthesizer DDS implementations in the firmware per FPGA.

In such an implementation, every individual DDS instance must be calculating subsequent samples at the same rate as being output by the DAC. All 128 DDS instances had to therefore run look-ups and perform multiplication at 20 MSPS throughput. These sort of full-speed DDS operations are expensive to implement on an FPGA, requiring a large number of static memory blocks, and high-speed, power-hungry, Digital Signal Processing blocks (DSPs). DSPs are scarce resources within an FPGA. They are dedicated circuits, containing memory and processing functions, that have been optimized for high-precision multiplication-and-addition and memory access, all at rates many times faster than ordinary logic on the FPGA fabric. Since resource requirements of DDS implementations scale linearly with multiplexing factor, to use DDS for the carrier and nuller modulation in the 64x readout system would have meant 4 times as many individual synthesizer DDS elements in firmware, per module. This encroaches the limits of FPGA technology, and would have demanded enough power to jeopardize our power budget.[2]

For this reason a new synthesis chain was developed, employing Polyphase Filter Banks (PFBs). A complete treatment of PFB synthesis is not possible here – for a detailed overview, see [11] – but it is different from DDS in that it does not build the final waveform out of individual sinusoids that have been generated separately. Instead, it computes the summed waveform directly, using an inverse Fast Fourier Transform (FFT) whose inputs determine

[2] The ground-based motherboard still maintains a 2-to-1 mezzanine-to-motherboard ratio, meaning that the firmware and FPGA it runs on must produce 8 times the number of readout channels.

the amplitudes, phases, and frequencies of a multiplexed array of carriers or nullers.[3] See Figure 3–1 for the data-flow diagram: a single DDS block computes the value for 128 readout channels (in blocks of 32), based on the phase, amplitude, and frequency information set by the user. This value is then accumulated in frequency bins stored in RAM; such that the bin each output is directed to is determined by the user-inputted frequency.[4] The values in those frequency bins describe the frequency-domain spectrum of our desired time-domain waveform. We construct that waveform using an inverse FFT, such that the imaginary output becomes the nuller waveform, and the real output becomes the carrier waveform. These are then filtered to remove artifacts of the finite FFT-length transformation, before being sent to the DAC to be digitized.

The nature of an FFT offers insight into why this is a solution to the scaling problem. In a continuous-input FFT, the rate at which each individual frequency bin is updated is a fraction of the total input data-rate, where the specific fraction depends on the FFT length. Similarly, for an inverse-FFT, the sampling rate of the inputs are a fraction of the total data rate that is output in the time-domain. In our case, each channel is accumulated in the appropriate PFB bin at 156 KHz, despite the 20 MSPS digital throughput to the nuller and carrier DACs. The end result is that the computational requirements to generate 512 synthesizer channels (per mezzanine) are enormously lower than would have been demanded by 512 separate DDS instances, all updating at 20 MSPS. The overall efficiency gains in moving to the PFB synthesis, which was considerably less computationally demanding, are

[3] The exact manner in which the inputs to the PFB control the parameters of the 64 component sinusoids is non-trivial and will not be explored in this document. This inverse-FFT analogy glosses over the details in an effort to highlight operational features.

[4] This means, for instance, that all 128 channels could be accumulating in only a single frequency bin. Think of the bins as building a histogram for a spectrum that will be produced by the inverse-FFT.

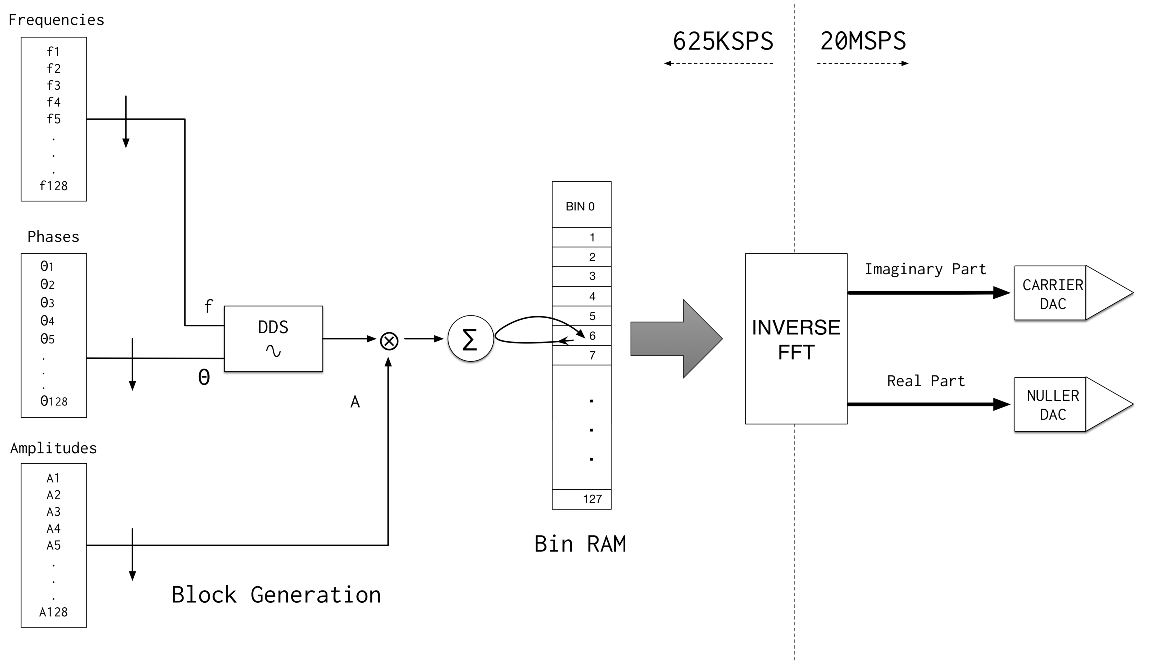

Figure 3–1: A PFB Data-Flow diagram, showing one DDS instance building up the spectrum in an inverse-FFTs frequency bins, which is then used to construct a time-domain waveform that corresponds to the desired spectrum. The black vertical arrows denote that the channels are being stepped through iteratively, in sync. The summation symbol indicates that each value from the DDS is added to the value in the frequency bin it corresponds to, which is accumulated until the next cycle of the FFT.

largely responsible for what became a factor of approximately 5 improvement in power consumption on a per-channel basis.

The gains in moving to a PFB are not a free lunch; one consequence of this increased efficiency is that we lose some degrees of freedom. The carrier and nuller signal path now share resources, and a few previously independent features are interwoven. In particular, the phase-generation of each nuller and carrier channel are combined into a single firmware module, so nuller and carrier frequencies of the same readout channel must always be set together. This is not a substantial loss – the ability to independently set phase and frequency parameters for the carrier and nuller of a single readout channel was a quirk of having independent DDS modules, and has never been assumed or taken advantage of in standard operation. Another repercussion of intertwining the nuller and carrier synthesis is that overloads in one can, under the right conditions, overflow and present as overloads in the other (of the same module). This can make debugging more difficult, but ultimately we don't envision any mode of operation such that one of the synthesizer channels can saturate without being catastrophic for the data-quality of all bolometers on the comb. This makes it irrelevant whether or not a rail in one synthesizer can degrade another, as there is no additional failure mode introduced by this counter-intuitive behavior.

The PFB increases the complexity of the system, and there are a few idiosyncrasies introduced in the details: the 10MHz bandwidth is divided into 64 frequency bins, each approximately 156 kHz wide. Actual synthesis frequency resolution is much finer than that – about 18.6mHz, on par with the resolution of the legacy DDS system. In DDS synthesis, a frequency is set by adjusting how many steps a Phase-to-Amplitude converter takes each time it accesses stored values in a LUT; with PFB synthesis, signals of different frequencies are produced by driving a slow sinusoid of an appropriate low frequency into the correct PFB bin, which mixes it up. The frequency bins themselves are fundamentally filters, and therefore have a frequency-dependent transfer function: as synthesizer frequencies traverse each PFB bin, a gain in the overall amplitude of the consequent sinusoid is applied. The frequency response is shown in Figure 3–2, and is symmetric around a minimum at the center, with a gain variation between the center and edges of each bin on the order of 2%. This gain variation can be derived analytically, and is known exactly. For sky-data it is irrelevant, as the final count-to-kelvin calibration will always be performed with a source on the sky, and this variation is contained within that measurement. For transfer function measurements, calibration, and other laboratory data it is a small correction made off-line. Finally, as a last note on modulation, it is disingenuous to say that a PFB does not use DDS – the sinusoids used to feed the inputs to the PFB bins are the product of a single DDS implementation, which calculates and accumulates the next output value of each channel into the appropriate

PFB bin before the inverse-FFT transforms those bins into a timestream. Instead of using the DDS to make both coarse adjustments to the desired frequency, as well as have the dynamic range for fine adjustments, the coarse adjustments to frequency come from *which* PFB bins is being used as an input, and the DDS sinusoid transmitted to that bin performs the fine-tuning. Fundamentally, what is being taken advantage of here is the difference in power consumption between running hundreds of DDS blocks at 20 MSPS, and a single one per module.

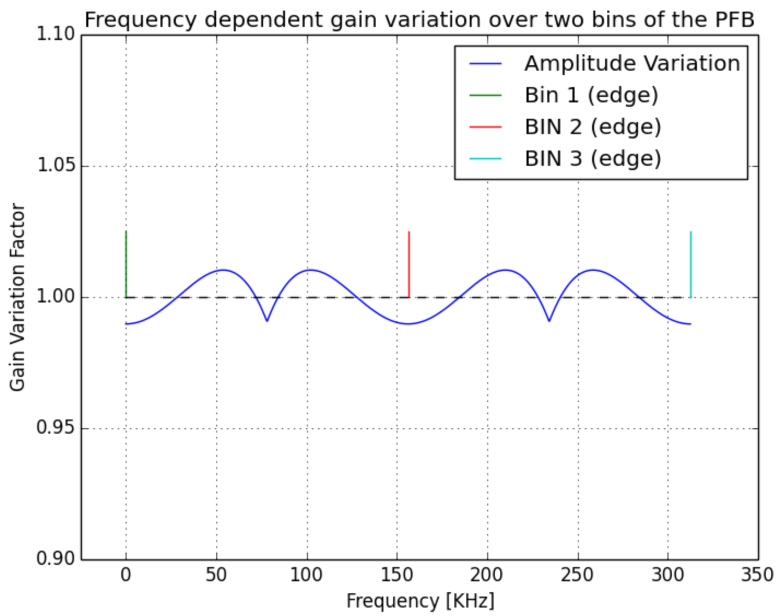

Figure 3–2: The gain variation across individual bins of the Polyphase Filter Bank.

3.3 Demodulation

As described in Section 1.3, voltage-biased bolometers in each comb modulate the resistance of the readout circuit, and therefore the current through it – which we record. The current through the bolometer circuit is itself being modulated at the resonant frequency of the LCR-filter the bolometer participates in. Our signal of interest is therefore amplitude perturbations to a large carrier sinusoid. This is an *amplitude modulation* encoding, analogous to how AM radio is broadcast. At present, the demodulation chain in place is a scaled-up version of one that has been successfully used in the legacy 16x DfMUX systems. The process involved is, fundamentally, a digital implementation of the familiar path of a super-heterodyne receiver: the SQUID output waveform is digitized on the mezzanine ADC at the full 20MSPS data rate, and then mixed down to baseband frequencies using DDS blocks to generate the mixing sinusoids. The DDS logic that was introduced in the previous section for synthesis are applied again here, such that each demodulation channel has its own programmable mixing sinusoid to demodulate at any frequency across the 10MHz bandwidth. Unlike analog mixers, the digital implementation allows for complex demodulation, shifting the entire signal down to baseband without producing the image frequencies that are a nuisance in analog demodulation. Doing this we recover both sidebands of the original carrier frequencies, and DAN uses the amplitude and phase information of these baseband frequencies to control feedback through the nuller. The data saved to disk is first down-sampled to 192Hz. The decimation used to do this has the potential to introduce aliasing from frequencies outside the 192Hz bandwidth for each demodulator channel: there are several layers of digital filtering between the demodulation and the decimation to avoid this.

This method of demodulating the ADC signals suffers from the same scalability problem as the synthesis chain. Significant gains in power consumption can be made by converting to a PFB demodulator, especially when going to higher multiplexing factors. For the 64x implementation, owing to the efficiency gains from the PFB synthesis, it was not necessary to do so. In fact, the coexistence in firmware of both types of logic allows for an interesting comparison. Figure 3–3 shows the FPGA logic occupancy for a complete version of the firmware running on a Kintex-7 FPGA. In purple are all 1024 (PFB) synthesis channels (carrier and nuller), in yellow is just 512 (DDS) demodulator channels. A future upgrade to a PFB demodulator would further reduce the power consumption, and is a strictly firmware upgrade that can be implemented with no changes to the operation or hardware of the device.

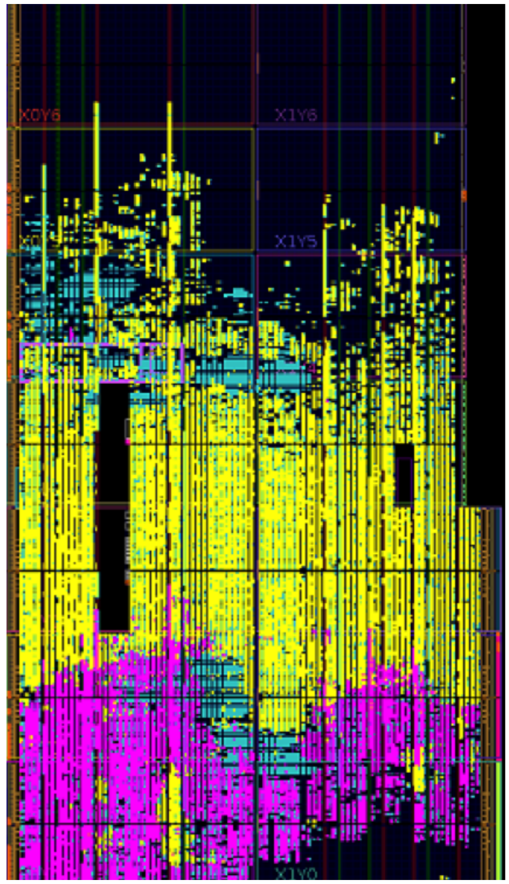

Figure 3–3: The occupancy of the logic on the fabric of the FPGA. This is a 2D representation of the topology of the FPGA Kintex-7 chip. Lit areas represent logic and memory that will be utilized as a result of the firmware programmed. Resources in purple are used for PFB synthesis, in yellow are the logic of the DDS demodulator. Image Credit: Graeme Smecher.

3.4 Data products

In order to construct maps and extract science data, the relevant product from each bolometer is just the time-stream of the amplitude and phase of the modulations. Both of these are represented as a time-stream of vectors in an "I&Q" plane, time-indexed by Inter-Range Instrumentation Group (IRIG-B) standard timestamps. The plane is defined for each bolometer with respect to the phase of the carrier signal, and the demodulator is aligned with the carrier such that the bias sinusoid appears maximally in I and minimally in

Q. In the absence of any signal, we would expect to see the carrier signal as a large DC offset in I, and for Q to record only quiescent noise, as in Figure 3–4. Thus a periodic amplitude modulation in phase with the carrier would appear as a periodic amplitude modulation in I, and would not register in Q. This seeks to ensure that output signals originating from the bolometer as perturbations to the carrier bias are projected into I, and only uncorrelated noise in Q. The utility in this is that rather than taking the magnitude of the signal from a bolometer, we get to exclude some of the Johnson noise of the bolometer which is suppressed in I, but appears primarily in Q.[5] Whether streaming data from the nuller or demodulator, both I and Q are recorded, but only the component signals in-phase with the carrier are used for science.

As a telescope scans across the sky, the underlying power fluctuations in the CMB illuminate the bolometer array, with a beam width defined by the optics of the telescope. As the incident power from the sky increases and decreases, the electrical power provided by the bias voltage responds accordingly in order to keep the total power deposited on the bolometer fixed. If the sky power drops, the resistance of the circuit drops, and more current is shunted through that leg of the circuit, and therefore into the SQUID input coil. In DAN we see this as an increase in the current (the amplitude) delivered by the appropriate DAN channel, to minimize the larger current at the input of the SQUID. As a bolometer moves from a cold spot on the sky to a hot one, the DC level of the I-component of the demodulated time-stream will do the opposite: decrease. By relating the time-index to pointing position, these amplitude variations can be used to paint a picture of hot and cold spots on the sky. Using this to get an absolute measurement of the power from the sky is difficult, especially for an entire focal-plane of bolometers. Every bolometer will have an offset related to its thermal conductivity, depth in the transition, and the shape of its particular superconducting

[5] This will be covered in more detail in Chapter 5.

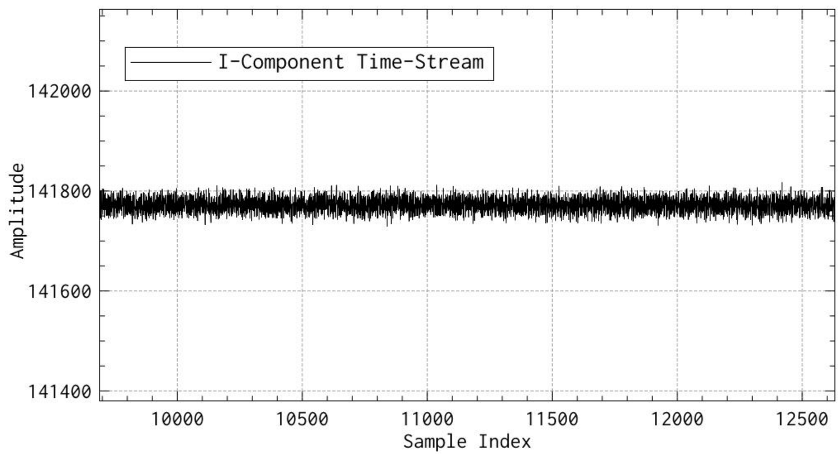

(a)

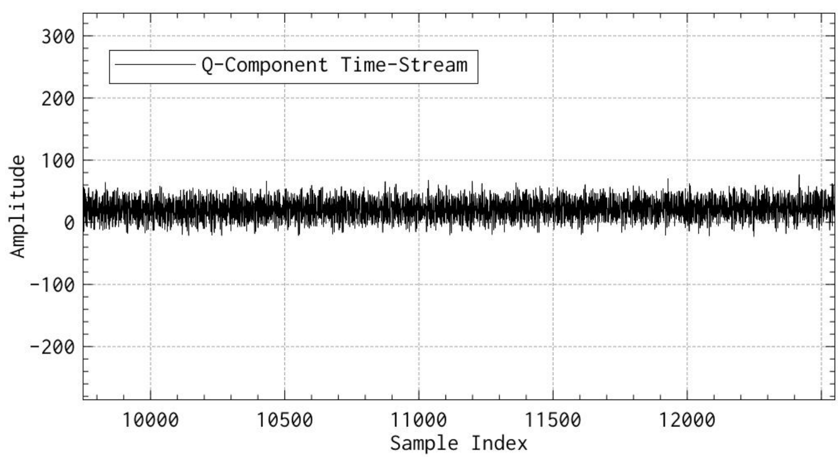

(b)

Figure 3–4: Example I (a) and Q (b) channels for an off-resonance time-stream.

transition. Maps are first constructed by building up relative measurements. The relative offsets of each bolometer are measured by pointing the array at a calibration source on the sky frequently throughout observations. Once all of the bolometers are normalized, their timestreams can be combined to take advantage of the sensitivity of the full comb.

Thermal time constants and drift within the cryostat must be slow compared to the rate that the target sky signal power moves through the beam when the telescope scans, otherwise they strongly contaminate the signal. The 1/f noise of the system as a whole, including the noise from thermal drift and changing atmospheric loading, determines the lower-limit of the rate at which the telescope must scan across the sky. The size of the beams, and the speed at which the telescope sweeps them across the sky, in turn determines how quickly the time-streams must be sampled to preserve all of the sky-signal. This is how the down-sampling rate is calculated. There is no reason to sample much faster than the rate at which fluctuations on the sky can be resolved. The down-sampling used for this project is designed to accommodate the beams and scan speed of the South Pole Telescope. This provides a decent margin of safety, since SPT has a 10-meter dish – at ∼1', the SPT beam is far narrower than one any satellite that flies this technology will have. It also scans rather rapidly, at approximately 1.6 degrees per second. For comparison, the Planck Satellite scanned at 6 degrees per second, with a main beam of approximately 1 degree [24]. We therefore oversample compared to what a satellite with similar characteristics as Planck would need by a factor of ~ 15.

CHAPTER 4
Transfer Functions

We have already noted that the measurement of current through the bolometer circuits serves as a proxy for getting at the underlying changes in sky power, but in fact there is an additional layer of abstraction – the amplitudes we actually record are in units of a variety of different "counts", which discretize the total dynamic range of the (12 bit) ADC, (16 bit) DAC, and (24 bit) DAN accumulators. Referring these units back to fixed physical quantities at the SQUID input coil and bolometers is essential to making sense of the data; assessing the performance of the system; and making algorithmic decisions in real-time.[1] The conversion factors required come from a combination of analytic and empirical methods.

This chapter goes through exactly how these conversion factors are derived, and the validation of them. It starts with a frequency-independent analysis of the synthesis and demodulation electronics, and finishes with an account of the corrections necessitated by higher-frequency effects. The focus will be on the transfer function of warm electronics components. Though the behavior of the cold electronics certainly complicates our understanding of the system as a whole, these considerations will be reserved for Chapter 5, as their effects on the noise of the system offer a way to understand them.

Ultimately the relevant absolute calibration of any instrument will be done on the sky, fully integrated, using sources with (relatively well) known surface brightness – but the development and validation of the readout relies on comparisons between bench-top measurements of noise and amplitudes, and those predicted by analytic calculations. It

[1] Specifically, this plays a huge role in the algorithms that manipulate the bolometer bias points.

gives us a start-to-finish confidence in our understanding of the system as a whole, and a quantitative expectation for noise performance.

On the mezzanine, the signal paths of the nuller and carrier are schematically identical, it is only once they reach the SQUID controller board that the nuller is divided down disproportionately to the carrier. From there, the nuller goes directly to the SQUID input on the cryogenic "SQUID card" (at the 4K stage), while the carrier is converted to a voltage-bias on that same card, and then applied to the LCR filters at the milli-kelvin stage. A complete cartoon circuit diagram of the signal paths that determine the transfer function of the warm electronics, including the cold electronics, is given in Figure 4–1. In the sections that follow each path shown there will be isolated to and discussed with additional figures.

4.1 Transfer Function Design

Before describing the arithmetic of the circuits, let us consider how decisions regarding the gain stages of the synthesizer and demodulator are made. For the synthesizer we've already mentioned power demands and constraints, and also a desire to divide down the intrinsic DAC noise before the signal reaches the cold electronics. In addition to reducing our absolute synthesizer noise contribution, dividing down the DAC output is further incentivized by a desire to minimize digitization noise – accomplished by exercising as many bits as possible without railing the DAC. Digitization noise quickly becomes sub-dominant to the DAC noise, provided we are using a reasonable amount of the dynamic range.[2] We also benefit from the fact that a large multiplexed array of sinusoids provides ample dithering for individually low-amplitude signals.

Finally, the differences in the carrier and nuller transfer functions are chosen in an effort to roughly match their amplitudes *at the SQUID input*, such that the programmable gains

[2] See Chapter 5.

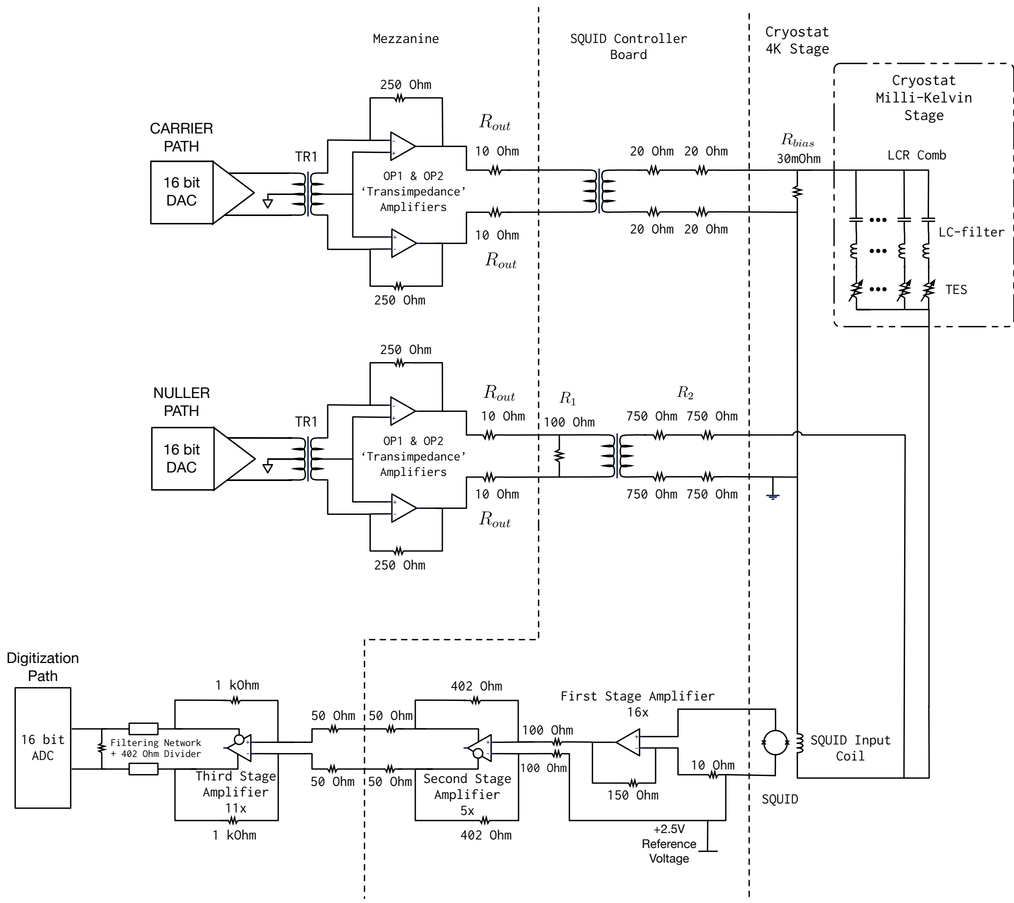

Figure 4–1: A cartoon circuit diagram showing the signal paths for both synthesizer and digitzation chains. All elements used to derive the DC transfer function are included in this image.

at the mezzanine DACs can remain symmetric, and neither has wasted dynamic range at the SQUID input. It also helps construct some user intuition when, for a bolometer that is normal (above its transition), carrier and nuller sinusoids of the same amplitude at the DAC will approximately cancel at the SQUID input coil.

For the digitization & amplification chain, the amount of gain used is determined by the sensitivity and upper range of the ADC. Our 12-bit ADC has a 2 volt peak-to-peak range. The gain chosen seeks to project the largest expected *signal* amplitudes into this full range; ensure that no amplifiers upstream of the ADC saturate before it will, to avoid accidentally locking ourselves out of the full dynamic range available; and leave some safety margin for innovations in SQUID transimpedance. A typical SQUID output voltage can be between 0 and 10mV, but the expected signal variations after tuning are less than half that.

4.2 Frequency-Independent Transfer Functions

There are a number of electronic filtering networks in the signal paths considered here, designed to attenuate signals outside of our band and prevent contamination, distortion, and additional loading on the cold electronics, typically in the high-frequency regions. Several transformers, also in the signal path, have roll-ons in the very low-end of our bandwidth, far lower than will ever be relevant for bolometers. The calculations in this section do not explicitly address these elements, they are done for signals at frequencies within our band, and assume no other frequency-dependence. As we will see, this is a pretty good approximation for the lower half of our bandwidth, but corrections to it must be applied when considering the full 10MHz range – these are detailed in Section 4.5.

When operating within flight-power constraints, each synthesizer DAC has a differential output of $I_{DAC} = 8mA$ peak-to-peak.[3] A "slow" 12-bit DAC can adjust the ADC peak-to-peak amplitude on the fly, increasing the reference current to the DACs such that $I_{DAC}^{P-P} = 8 \times I_{ref}$, up to a reference current of 5mA. Low power mode is defined as a reference current of 1mA. The derivations in this section assume a full-scale DAC output swing when confined to low power-mode, such that $I_{DAC}^{P-P} = 8 \times 10^{-3}\ A$.

At the output of the synthesizers, the differential signal goes through a center-tapped one-to-one transformer, labeled $TR1$ on Figure 4–2, with an effective current-gain of one half. The two operational amplifiers that follow, $OP1$ & $OP2$, form a differential transimpedance

[3] The DACs have a 20mA current flow through each of two legs, making a maximum differential current of 40mA. We limit this to 1/5th the total output swing in low-power mode, leaving 8mA peak-to-peak.

amplifier (converting an input current to an output voltage with some gain). The gain for the transimpedance amplifier is determined by a pair of 249Ω feedback resistors on each amplifier. The amplifiers each have an output impedance of 10Ω (R_{out}). The R_{out} resistors will affect the nuller and carrier signal paths on the SQUID controller board in different ways; to avoid double-counting them, we will exclude them in this section. Therefore, the mezzanine output voltages (defined in Equation 4.1 as V_{Mezz}^{P-P}) refer to the voltage at the output of the *transimpedance amplifiers*, not the mezzanine itself. Equation 4.1 is valid for both carrier and nuller signal paths

$$V_{Mezz}^{P-P} = (2\times 249)\ [\Omega]\times\frac{I_{DAC}}{2}\ [A] = 249\times\ I_{DAC}\ [V]\ . \tag{4.1}$$

4.2.1 The Nuller

The output impedance of the nuller transimpedance amplifiers are in parallel with a 100Ω resistor (labeled R_1 in Figure 4–3) at the input of the SQUID controller board. R_1 is transformer-coupled to two sets of two 750Ω resistors in series on each line, which form an effective series impedance of 3kΩ (labeled R_2). The effective impedance of R_1 in parallel with the R_2 is 96.77Ω. This forms a voltage divider with the two 10Ω output resistors on the mezzanine (term 2 of Equation 4.2). R_2 converts that voltage back to a current before it leaves the SQUID controller board for the cold electronics (term 3 of Equation 4.2).

The total nuller transfer function of the warm electronics becomes

$$In_{SQUID}^{P-P} = V_{Mezz}^{P-P}\ \ [V]\times\frac{96.77\ \ [\Omega]}{20+96.77\ \ [\Omega]}\times\frac{1}{4\times 750\ \ [\Omega]} \tag{4.2}$$

$$In_{SQUID}^{P-P}\ \ [A] = 0.0687\ \times I_{DAC}^{P-P}\ \ [A]\ . \tag{4.3}$$

Source	Gain Factor or Output	Location
AD768 DAC, output source	$I_{out}^{P-P} = 8mA \times I_{ref}$	Mezzanine
Current Dividing Transformer	0.5x	Mezzanine
Transimpedance Amplifier	498x ($\times[\Omega]$)	Mezzanine
Voltage Divider	0.8287x	Mezzanine & SQCB
Final Output Impedance	$0.00\overline{33}$x ($\times[\Omega^{-1}]$)	SQCB
Total Gain Factor	0.0687x	Warm Electronics

Table 4–1: A table of the nuller transfer function, going from the DAC output current to the current at the input of the SQUID coil. In low-power mode $I_{ref} = 1mA$, but can go as high as 5mA.

In_{SQUID}^{P-P} is the peak-to-peak current through the input coil of the SQUID, referred back to the current at the output of the nuller DAC. A summary of each of the nuller gain stages derived here can be found in Table 4–1.

It is common to refer to DAC output in "Normalized" units, which are expressed as a peak-amplitude between 0-and-1. A nuller signal in Normalized units, referred back to the current at the SQUID, is

$$In_{SQUID}^{P-P}\ [A] = (5.498\times10^{-4})\times A^{Peak}\ [Normalized]\ . \tag{4.4}$$

The second commonly used unit is raw DAC counts – a signed 16 bit number. Herein we refer to these simply as "Readout Counts" ("ROCS"). Converting Equation 4.4 in terms of ROCS

$$1\ [A_{norm}] = 2^{15}\ [ROCS] \tag{4.5}$$

$$In_{SQUID}^{P-P}\ [A] = (1.678\times10^{-8})\times A^{Peak}\ [ROCS]\ . \tag{4.6}$$

There is a third unit used when discussing specifically nuller amplitudes, and in many ways is the most important of all: "DAN Readout Counts" ("DROCS"). These are signed 24 bit numbers, and also by convention are peak-amplitude values.

$$1\ [A_{norm}] = 2^{23}\ [DROCS] \tag{4.7}$$

$$In_{SQUID}^{P-P}\ [A] = (6.554{\times}10^{-11}){\times}A^{Peak}\ [DROCs]\ . \tag{4.8}$$

To form a series of universal conversions, we remove the implicit conversions between peak-to-peak amplitude and peak-amplitude present in the above equations. While in low-power mode, our final conversion factors are as follows

$$C_1[\frac{uA}{A_{Normalized}}] = 274.9 \tag{4.9}$$

$$C_2[\frac{nA}{ROC}] = 8.389 \tag{4.10}$$

$$C_3[\frac{pA}{DROC}] = 32.77 \tag{4.11}$$

These are ubiquitous unit conversions within the readout control code and analysis; C_1 is often used to predict the exact current driven through the SQUID input coil by a small sinusoid programmed with normalized units – one half of the calculation required to construct a probe that measures SQUID transimpedance. C_3 is the crucial conversion factor for determining the current through the bolometer circuit from amplitude variations in the DAN time-streams.

4.2.2 The Carrier

The signal path of the carrier across the SQUID controller board is considerably simpler than that of the nuller. Each of the two traces of from the mezzanine transimpedance amplifiers exit through their 10Ω output resistors, as in Section 4.2.1, and are transformer-coupled on the SQUID controller board to two 20Ω resistors, shown in Figure 4–4.

This means that the only modification to Equation 4.1 necessary is an equivalent series resistance of 100Ω. At the output of the carrier path on the SQUID controller board (and the input to the cryogenics), the carrier signal is

$$Ic_{SQCB}^{P-P} = \frac{V_{Mezz}^{P-P}\ [V]}{100\ [\Omega]} = 2.49\ \ I_{DAC}^{P-P}\ \ [A]\ . \tag{4.12}$$

On the SQUID card, Ic_{SQCB}^{P-P} is driven through a $30m\Omega$ resistor (R_{bias} on Figure 4–4), creating the voltage bias that travels through a superconducting strip-line to the comb of LC-filters and bolometers at the mK stage,

$$V_{Bias}^{P-P} = 0.03[\Omega]\times Ic_{SQCB}^{P-P}[A] = 0.0747\times I_{DAC}^{P-P}\ [V]\ . \tag{4.13}$$

A summary of the full carrier transfer function can be found in Table 4–2.

Source	Gain Factor or Output	Location
AD768 DAC, output source	$I_{out}^{P-P} = 8mA\times I_{ref}$	Mezzanine
Current Dividing Transformer	0.5x	Mezzanine
Transimpedance Amplifier	498x ($\times[\Omega]$)	Mezzanine
Series Impedance	0.01x ($\times[\Omega^{-1}]$)	Mezzanine & SQCB[4]
Bias Resistor	0.03x ($\times[\Omega]$)	SQUID card (4K)
Total Gain Factor	0.0747x ($\times[\Omega]$)	Warm Electronics & SQUID card

Table 4–2: The carrier transfer function, from a current at the DAC output to a voltage across the bolometer comb. In low-power mode $I_{ref} = 1mA$, but can go as high as 5mA.

Since the carrier is never part of the DAN feedback path, the only relevant units are Normalized and ROCS. Often, carrier amplitudes are known by the values set by users or user-generated code, making the most useful conversion between Normalized units and volts across the bolometer. In order to find the general conversion factors for the carrier chain we convert from peak-to-peak amplitudes, to peak-amplitudes

$$1.0\ \ [Normalized] = \frac{V_{Bias}^{P-P}}{2} = 0.0747\ \times\ \frac{I_{DAC}^{P-P}}{2}\ [V]\ , \tag{4.14}$$

Source	Gain Factor	Source Location
First Stage Amplifier	16x	SQUID controller board
Second Stage Amplifier	5x	SQUID controller board
Third Stage Amplifier	11x	Mezzanine
Filtering Network	0.5x	Mezzanine
Total Gain Factor	440x	Warm Electronics

Table 4–3: Gain stages of the demodulator path.

and solve for our next unit-conversions

$$C_4[\frac{uV}{Normalized}] = 298 \tag{4.15}$$

$$C_5[\frac{nV}{ROC}] = 9.12 \tag{4.16}$$

C_4 is used extensively in the software, together with C_3, to calculate bolometer resistance and electrical power across the comb in real time. This is particularly important in algorithms that drop bolometers to specific depths in their superconducting transition. Real-time measurements of bolometer resistance is also used to determine the number that successfully tuned, and identify any that have gone fully superconducting ("latched").

4.2.3 The Demodulator

Between the output of the SQUID coils on the 4K SQUID card, and the 12 bit ADC on the mezzanine, there are three active gain stages, and a voltage-dividing filtering network. These are highlighted in Figure 4–5, and summarized in Table 4–3.

To calculate conversion factors akin to those in the above two sections, which will enable us to recover volts at the output of the SQUID coil from the raw ADC counts, we work backwards. The unit used to describe the raw ADC readout counts is called an "AROC"

---

[4] Equation 4.12.

("ADC Readout Counts"), and is also a peak-amplitude quantity. Although the ADC is 12-bits, we cast the numbers in firmware into 14 bit values (multiply by 4) – which allowed us to use an existing demodulation framework in firmware, since the legacy 16x DfMUX systems employed 14-bit ADCs. The ADC has a peak-to-peak range of 2 volts, so a saturating signal will be

$$1\ [V_{ADC}^{peak}] = 2^{13}\ [AROCS] \tag{4.17}$$

Next, to calculate how many volts at the output of the SQUID would rail the ADC, we divide by all of the gain stages in Table 4–3

$$2^{13}\ [AROCS] = \frac{1\ [V_{ADC}]}{440} = 0.002\overline{27}\ [V_{SQUID}]\ . \tag{4.18}$$

This quantifies what's been said earlier regarding the relative amplitudes of a SQUID voltage output, versus a signal voltage variation at a single SQUID tuning. The offset DAC allows us to tune SQUIDs which can produce between 0-10mV, while still matching the ADC dynamic range to a ~2mV signal. Reducing Equations 4.17 & 4.18 yields our desired conversion factors

$$C_6[\frac{uV_{ADC}}{AROCS}] = 122.07 \tag{4.19}$$

$$C_7[\frac{uV_{SQUID\ output}}{AROCS}] = 0.2774 \tag{4.20}$$

4.3 Measuring SQUID Transimpedance

Without DAN, recovering the current through the bolometer circuit would require a measurement of the SQUID transimpedance. DAN relieves us of this inconvenience, but the transimpedance of the SQUID is still an important quantity to know. It allows us to assess the performance of the SQUID, and quantify the success of a SQUID tuning operation. In Section 4.2.1 we suggested that C_1 is one half of the necessary ingredients to measure this. C_7 gets us the rest of the way there. The transimpedance of the SQUID (Z_{SQUID}) is what determines how a current at the SQUID input coil is converted to a voltage at the SQUID output

$$V_{SQUID\,output} = Z_{SQUID}[\Omega] \times I_{SQUID\,input} \tag{4.21}$$

$$Z_{SQUID}[\Omega] = \frac{V_{SQUID\,output}}{I_{SQUID\,input}} \; . \tag{4.22}$$

Consider a "tickle" signal delivered to the SQUID input coil by a 5×10^{-4} Normalized unit nuller sinusoid, which when demodulated at the ADC is measured have an amplitude of 62 ADC Counts. The current through the SQUID input, and voltage at the SQUID output, can then be calculated and entered into Equation 4.22

$$I_{SQUID\,input} = 5 \times 10^{-4}\;[A_{Norm}] \times C_1\;[\frac{uA}{A_{Norm}}] = 1.37 \times 10^{-7}\;[A] \tag{4.23}$$

$$V_{SQUID\,output} = 62\;[AROCS] \times C_7\;[\frac{uV}{AROCS}] = 6.88 \times 10^{-5}\;[V] \tag{4.24}$$

$$Z_{SQUID} = \frac{6.88 \times 10^{-5}\;[V]}{1.37 \times 10^{-7}\;[A]} = 504[\Omega] \; . \tag{4.25}$$

A resulting transimpedance of about 500Ω at the frequency of the tickle signal is nominal for highest performing SQUIDs used in ground-based DfMUX systems.

4.4 Experimental Verification

Each of the conversion factors derived in Section 4.2 were experimentally verified through several methods. A self-contained check of internal consistency is made simpler by the separate synthesizer and digitization chains. We exploit this by wiring a resistor of known value in a loop-back configuration, from the synthesizer outputs on the SQUID controller board, back to the input where voltage from the SQUID is sensed. This configuration is illustrated in Figure 4–6. Our synthesizer conversion factors (C_1 & C_4) allow us to calculate the expected current through the resistor resulting from a programmed amplitude signal at the DACs, and therefore make a prediction about the voltage across that resistor. We test that prediction by measuring the amplitude at the ADC in AROCS, and recovering the voltage drop across the resistor using C_7. This internal consistency check is then verified using an external multimeter probe across the resistor. The results of these tests agree, at low frequencies where other corrections are not needed, to within the acceptable tolerance from component variation (approximately 3%, owing to scatter in the actual resistance of the components used [5]).

A separate method to validate the synthesis chain comes from driving the outputs of the SQUID controller board directly into a spectrum analyzer. A frequency produced by the warm electronics is identified as a peak in the spectrum, and the amplitude of that peak is compared to the expected value.

A final method to validate the demodulation chain is known as "signal injection", the setup for which can be seen in Figure 4–7. It also serves as a bench-top-validation of the DAN feedback path. A function generator injects sinusoidal voltage, at a programmed amplitude and frequency, into the input of the amplification path of the SQUID controller board. A DAN channel is then enabled at that frequency, which nulls out the signal. The residual amplitude at the ADC is recorded, and checked to be consistent with the expected DAN loop-gain; and the amplitude of the DAN channel is referred back to the voltage being produced by the function generator using C_3. This was a critical validation tool during the DAN implementation commissioning, hardware commissioning, and transfer function validation – areas of this project that I was closely involved with.

4.5 Frequency-Dependent Corrections

Thus far we have neglected any frequency dependent corrections to the transfer function. In the legacy systems, which operated below 1 MHz, it was reasonable to approximate the

[5] Resistive components are generally 1% tolerance components

response as frequency independent, but with the new 10MHz bandwidth these effects can no longer be ignored. The primary mechanisms by which the response of the synthesizer or demodulator chains can be sensitive to frequency are loss in components, such as the in-line transformers on the SQUID controller board; the intrinsically spectral-shaping effects of filtering networks, which do not completely avoid our band; and capacitive coupling on the traces of the PCB or within individual ICs, although this is much harder to confirm. In particular, the filtering that takes place to low-pass filter signals in the synthesis chain will, even when optimally designed, have a significant ripple across the pass-band. This is especially true with the 5% capacitor component tolerances.

Deriving the exact corrections is difficult, but measuring them is straightforward. In general the synthesizer chain will be most affected, while we expect the digitization chain to be very flat. In the latter there are no transformers, and the filtering is a simple anti-alias design. In the synthesizer path we worry about loading the SQUID with any superfluous power at high frequencies, which leads us to use strong analog filtering. The digitization chain frequency response is verified using a function generator to inject signal at the SQUID controller board. For the synthesis chain the bandwidth is populated by nuller and carrier tones, which are piped into a spectrum analyzer. Both of these are variations on the techniques introduced in Section 4.4.

Those measurements demonstrate a flat response as a function of frequency for the digitization chain, as expected; and a synthesizer response that rolls off strongly at about 4.5 MHz (Figure 4–8). The roll on below 100 kHz is the transformer on the SQUID controller board. The roll-off at 4.5 MHz is consistent with transformer cut-off frequencies and our understanding of the filtering networks. In part due to the expected variations in the analog signal path from passive component tolerances, we don't need to know this frequency dependent transfer function to better than 5-10%. If we did need to, it would be important to calibrate the spectrum analyzer in detail, which we have not done.

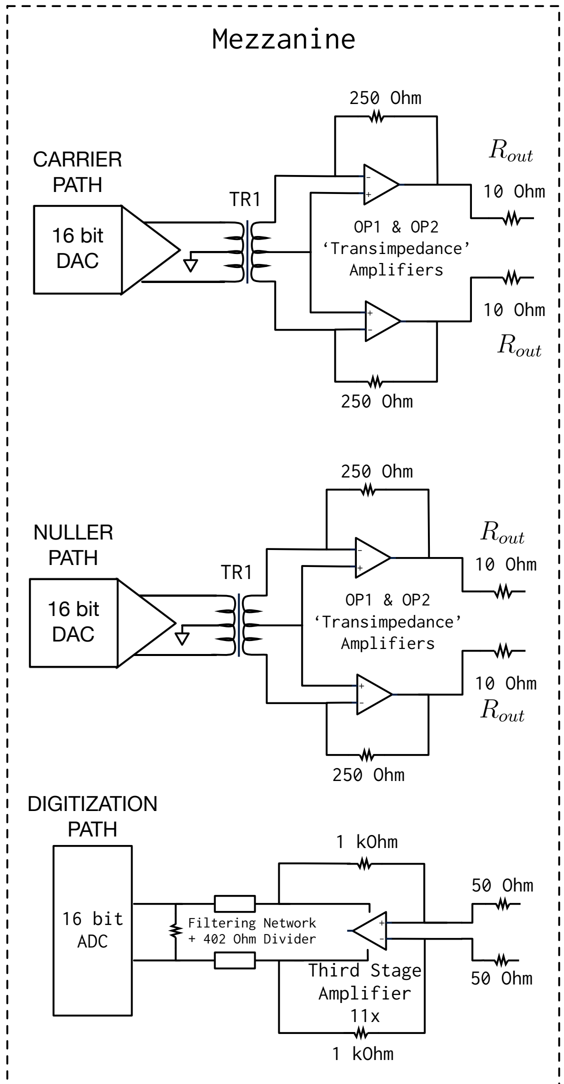

Figure 4–2: A cartoon circuit diagram of the signal paths on the mezzanine.

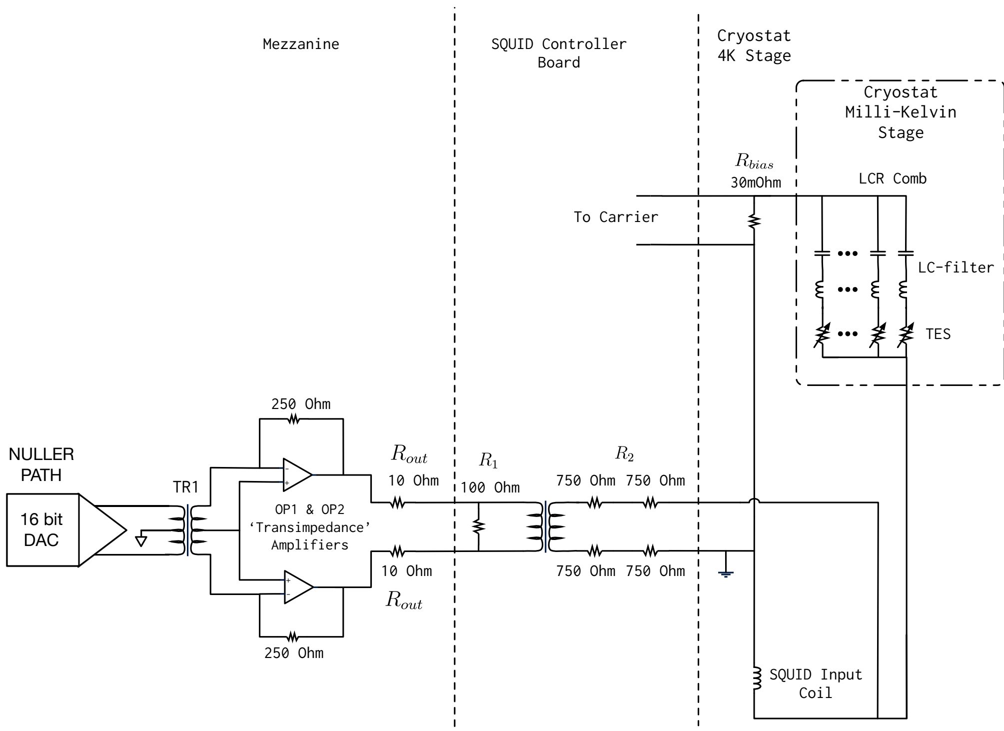

Figure 4–3: A cartoon circuit diagram of the signal paths of the nuller chain the mezzanine.

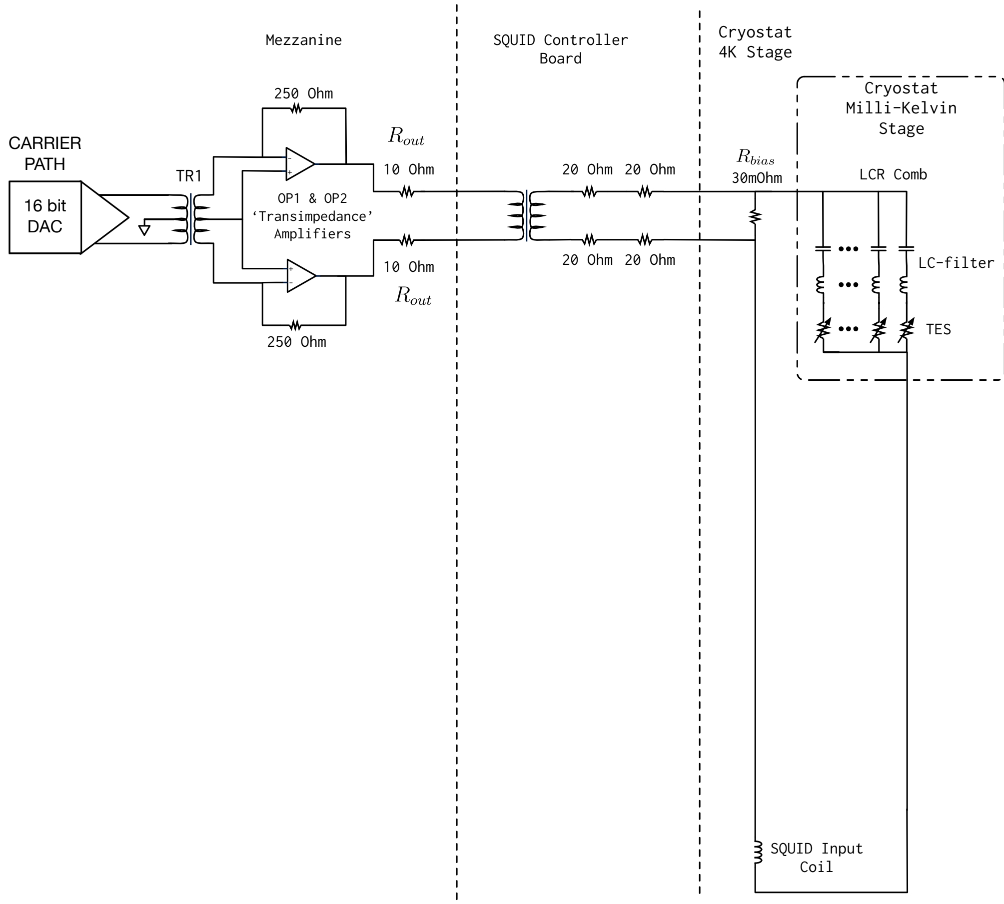

Figure 4–4: A cartoon diagram of the signal path of the carrier chain.

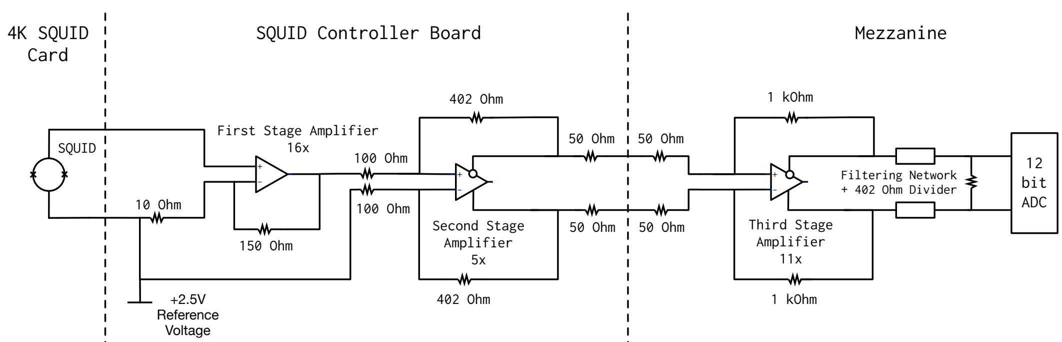

Figure 4–5: A cartoon circuit diagram of amplification and digitization chain. Note: This has been flipped with respect to the other images to facilitate reading.

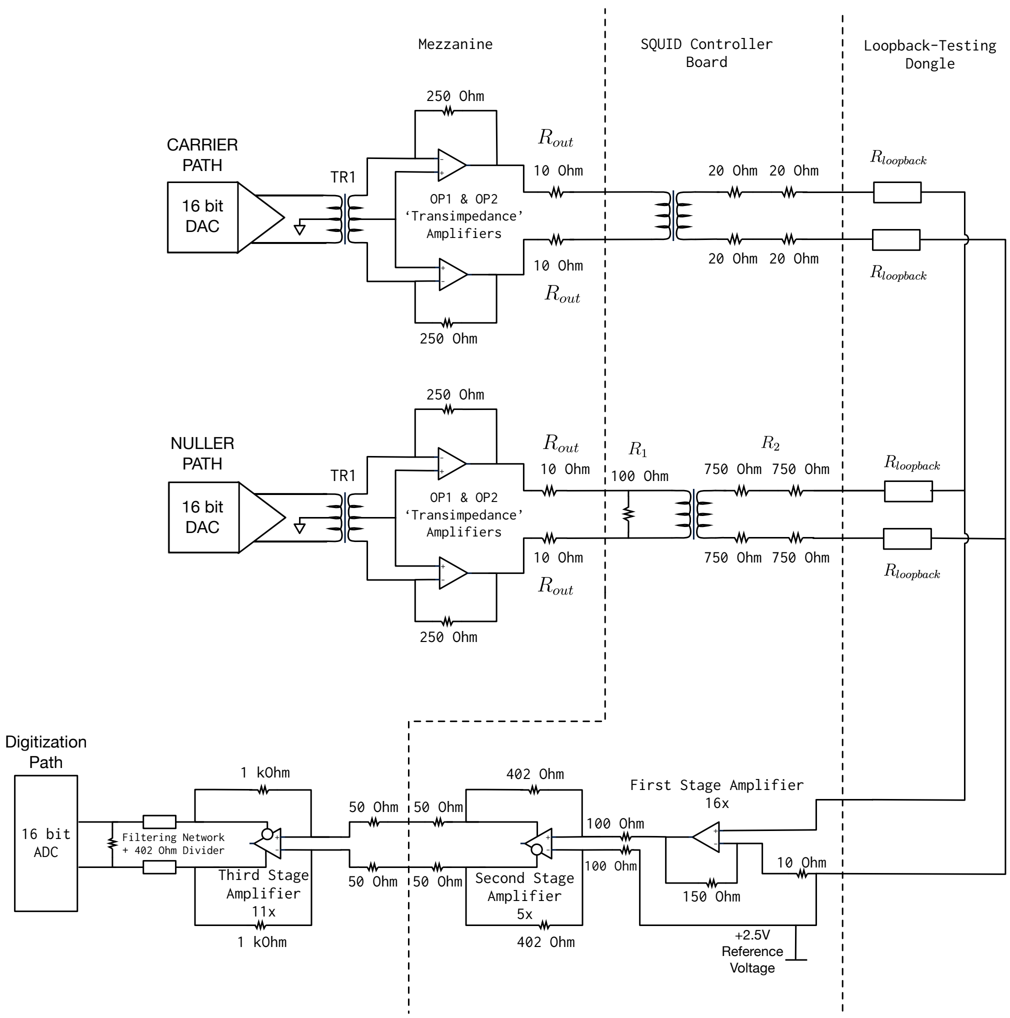

Figure 4–6: A diagram of the loop-back testing circuit.

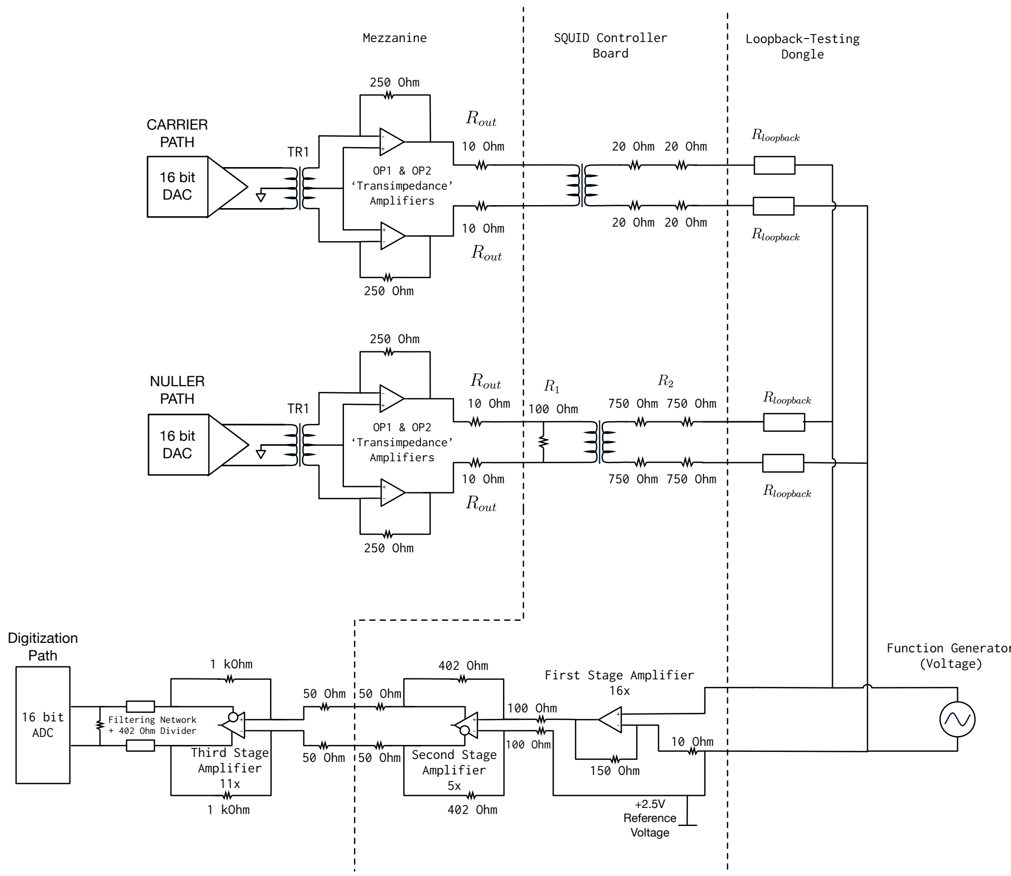

Figure 4–7: A diagram of the signal-injection bench-top testing setup.

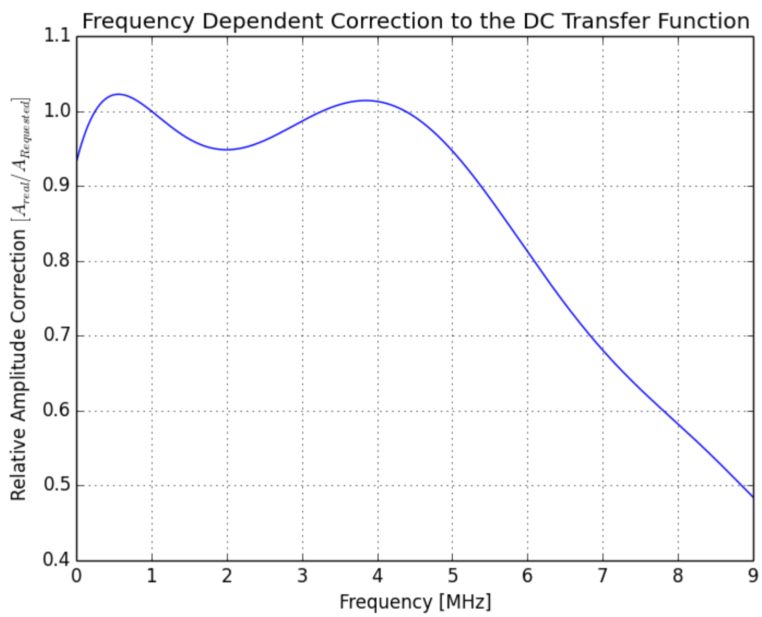

Figure 4–8: The frequency response of the synthesis chain as measured with a spectrum analyzer.

CHAPTER 5
Noise Models and Predictions

Assessing the power consumption, thermal properties, or the ability to bias a comb of bolometers, are all relatively straightforward, if technically challenging, operations. Quantifying the noise of the system can be more complicated. In large part this comes from the fact that the readout contribution to noise is buried amongst a host of additional sources, all of which combine in ways that are sensitive to the dynamic state of the system. On a fundamental level, there is no way to entirely disentangle the readout noise from the noise of the system as a whole; it must be assessed in an end-to-end environment.

Major sources of noise in the system can be broken into the those originating in the carrier and nuller synthesizer paths, the digitization path, SQUID noise, and bolometer noise. The sections that follow first describe how we see and measure noise in the system. We then derive and quantify the known noise sources for the synthesizer and demodulator paths, SQUID, and bolometer. Finally those derivations will be paired with some assumptions about the properties of cold components that would be representative of a deployed scientific-instrument in order to make a total noise prediction.

5.1 Noise Measurement Strategy

Measuring the noise in our output data – a DAN channel – is simple. A time-stream from the channelized and down-sampled 192Hz streamer is recorded for an interval of 3 minutes. This data consists of the magnitudes in I and Q as a function of time; where I and Q are the projected amplitude of the magnitude-phase vectors onto a plane defined by a component that is aligned with the phase of the carrier sinusoid (I) and an orthogonal out-of-phase component (Q). In Section 3.4 we asserted that the I-component of this data was the most valuable, as the bolometer response to all power deposition (which describes all of our signal) will be in-phase with the voltage bias sinusoid, and only uncorrelated noise will be present in the out-of-phase Q-component.

This is correct but for the fact that the I-dimension of the I&Q plane is aligned with the carrier phase *at the synthesizer*. This is not quite the same as being aligned with the bolometer responsivity – which is the phase of the voltage bias *as seen by the bolometer*. Stray complex impedances in the cold-electronics result in a phase-shift of the carrier sinusoid between the warm electronics and the bolometer circuit. This phase-shift is typically on the order of a few degrees, which means that the I&Q plane as it is recorded is mis-aligned with the responsivity of the bolometer. The extent of that mis-alignment determines the degree to which we are needlessly degrading the signal-to-noise of our data by adding an additional source of incoherent noise.

We can correct for this misalignment of the I&Q plane by performing a Principal Component Analysis on the I&Q data-streams (an Eigen-Decomposition of the covariance matrix whose rows are the I and Q data), and transforming to a basis that maximizes the variance in I and minimizes it in Q. This transformation aligns I with the carrier phase as seen by the bolometer, because the largest sources of variance come from power deposition on the bolometer, which will be along the responsivity axis. Note that this is true even in the case of a "dark detector" which is not absorbing substantial optical power, in part because – as we shall see in the following sections – the largest sources of noise in the bolometer circuit are power-noise from the bolometer itself, and therefore will be aligned with bolometer responsivity.

To evaluate noise, we take the Amplitude Spectral Density (ASD) of the resulting in-phase time-stream and then calculate the median value from 0.5-40Hz, which spans the typical science band with additional allowances for future faster bolometers. The "white-noise" of a readout channel is then quoted as a single number. An example of such an ASD is shown in Figure 5–1.

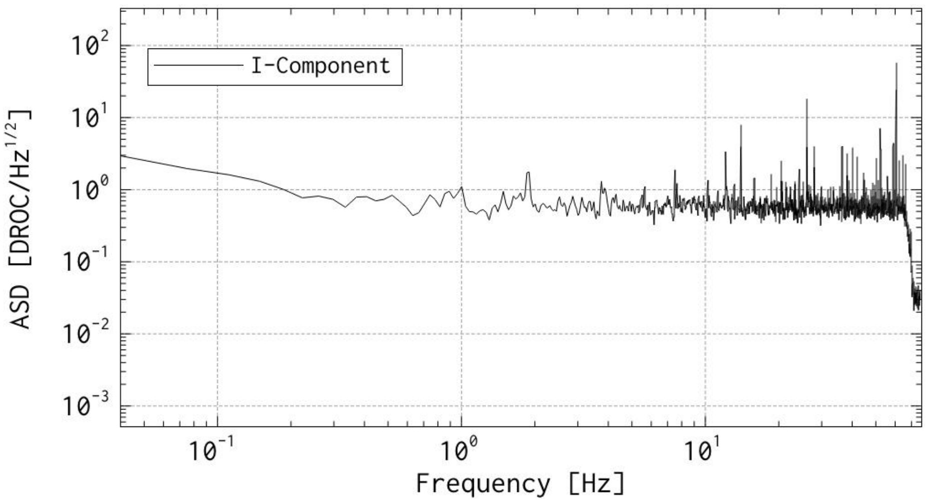

Figure 5–1: An example Amplitude Spectral Density plot of an overbiased bolometer. The band in which the noise is evaluated is 0.5-40Hz. The discrete lines are discussed in 6.2.2

There are three different measurements of the white noise we are interested in. The first is the noise of an overbiased ("normal") bolometer, wherein we have delivered so much electrical power that the TES is pushed into its normal state – and is effectively a resistor. The typical normal resistance of a bolometer with the desired properties for CMB observation is $R_{bolo}^{norm} = 1\Omega$. The second is the noise of a bolometer operated in its superconducting transition – typically at 70% of its normal resistance. At this depth bolometers typically exhibit a loop-gain of approximately 10. For the calculations in this chapter we will assume that $R_{bolo}^{trans} = 0.7\Omega$ and $\mathcal{L}_{bolo} = 10$. The last state we are interested in is the noise present in a DAN feedback channel with a carrier bias amplitude representative of a bolometer bias, but at a frequency that does not correspond to an LC-resonance. This is known as an "off-resonance" channel, and it is our most useful proxy for measuring the pure "readout" noise, independent of bolometer noise sources (the validation of which are not the object of this project).

In each of these three states, the relative strengths of the components of the total readout noise differ: the readout channel of a bolometer operated in-transition will suffer from all components of the system noise; an overbiased bolometer is relieved of one major component of bolometer noise, only present when in its superconducting transition; and an off-resonance channel is independent of all bolometer noise sources, and additionally sees all carrier-signal-path noise strongly attenuated. There is no direct way to measure the quantity we are most interested in, which is the representative readout noise independent of all noise sources of the cold electronics, but an off-resonance channel is a sufficiently close probe that it was defined as our proxy measurement criterion for noise performance during the evaluation with the CSA. This document endeavors to do slightly better, and model each of the noise sources directly.

A final note is that in modeling the noise components of the readout system we make a choice to evaluate them as a current noise at the input of the SQUID coil. This is simply because it is convenient to do so – that is what our measurements are probing, and doing so eliminates the need to apply transformations to both predictions and measurements simultaneously. A consequence of this is that, in some cases, assumptions regarding the state and properties of the system will need to be applied. For this section those assumptions are made to be consistent with hardware properties of recent CMB instruments, which is after-all the target application of this readout system. Corrections to the predictions to take into account some idiosyncrasies of the specific hardware used at McGill will be made in Chapter 6.

5.2 Types of Noise Sources in the Warm Electronics

There are four types of noise that are non-negligible in the warm electronics: intrinsic DAC noise, intrinsic amplifier noise, quantization noise, and Johnson noise. All noise sources are quoted as Amplitude Spectral Densities (ASDs): an RMS amplitude per square root of bandwidth. This may be peculiar to some; often the natural units to use for calculating noise are in *power*. However, our instrument senses perturbations in power, and converts them to modulations in amplitude. The relevant noise schema for us is therefore variations in amplitude that can be misconstrued as signal power. We will demonstrate in this section how to calculate the noise of these elements at their source, and then in the following three sections we will use the tools laid out here to referred them to a current noise through the SQUID.

5.2.1 DAC Noise

Our DACs are digital current-sources whose internal networks of amplifiers, transistors, and current-ladders have an associated electronic noise. The actual sources of this noise vary (including transistor noise, current-ladder noise, thermal noise, 1/f) – and over the whole bandwidth that the IC can produce it has some spectral character. However, over our 10MHz bandwidth it is uniformly distributed, and manufacture quoted as 60 $\frac{pA}{\sqrt{Hz}}$ at the DAC output. This means that the noise source can be modeled as a current source, at the DAC output, producing an RMS noise of 60 $\frac{pA}{\sqrt{Hz}}$ in addition to the amplitude programmed at the DAC. Note that this noise is uncorrelated with the amplitude and frequency of any signals the DAC may be generating. The AD768 DAC is used in both the carrier and nuller synthesis, but the different transfer functions (detailed in Chapter 4) mean that the DAC noise contribution at the SQUID coil will be different for the carrier and nuller signal paths.

5.2.2 Amplifier Noise

Operational amplifiers are active components that have their own intrinsic noise. Like the DACs, their intrinsic noise is given in a manufacturer's data-sheet, and subsequently confirmed by bench-top measurements. These are typically given as an input voltage noise, which means that the effective-noise contribution at the input of the amplifier is independent of the gain of that amplifier (and therefore particular implementation of the device) – but the total amplifier noise at the output will be multiplied by it.

The amplifiers on the synthesis signal path at the output of the mezzanine are being used in a transimpedance configuration (See Figure 4–2), the actual voltage-gain is much less than 1 (in this respect, the 249 Ω gain resistor is forming a voltage gain with the effective output impedance of the DAC, which is several $k\Omega$). The result is that their effective input noise is divided down substantially, and does not play a role in the noise of the synthesis chain. In contrast, all of the amplifiers used in the digitization chain are configured as traditional voltage-amplifying stages, whose purpose is to apply a gain greater than 1. These *do* contribute substantially to the digitization chain noise. The most important amplifier in this respect is the first stage amplifier on the SQUID controller board, which has a 0.95 $\frac{nV}{\sqrt{Hz}}$ input noise and a gain of 16. This means that at the output of this amplifier, one would measure 15.2 $\frac{nV}{\sqrt{Hz}}$. Although the subsequent amplifying stages also have gains greater than one, they are all sub-dominant to the first stage amplifier, whose noise is multiplied up by each following stage.

5.2.3 Quantization Noise

Quantization noise (also known as "digitization noise" when referring to quantization that results from the digitization of an analog signal) comes from the fact that a continuous signal is being quantized (either in signal generation or recording) to the precision of the bits available at the device (the DAC or ADC, respectively). This noise can be understood as an error on the signal that is introduced by digital truncation. Provided more than just a few (~ 6) bits are being used, this error is "white", in the sense that it is distributed uniformly over the bandwidth. The transfer-function design discussed in Chapter 4.1 ensures that during normal operation we are using as much of the full 16-bit dynamic range of the DAC as possible, so we are always in the regime where quantization noise is white. A measure of the quantization noise (when producing sinusoidal waveforms) is given as a Signal-To-Noise-Quantization-Ratio

$$SNQR \sim \frac{1}{2}(1.761 + 6.02Q)\ [dB_{amplitude}]\ . \tag{5.1}$$

We use this ratio to relate the absolute noise amplitude at the digital device in RMS, to a full-scale signal in peak-to-peak amplitude (the natural unit when we define the ranges of our DAC and ADC)

$$A^{RMS}_{Quant} = \frac{1}{2\sqrt{2}} \times \frac{A^{P-P}_{full-scale}}{10^{\frac{SNQR}{10}}}\ . \tag{5.2}$$

Finally, to convert to an ASD, we divide by the square root of our bandwidth

$$ASD_{Quant} = \frac{A^{RMS}_{Quant}}{\sqrt{10^7}} \quad \left[\frac{A^{RMS}}{\sqrt{Hz}}\right]\ . \tag{5.3}$$

The DACs are 16 bit devices, whose full-scale output (in low-power mode) is 8mA peak-to-peak. From Equation 5.3, the effective quantization noise at the DAC output is then 11.12 $\frac{pA}{\sqrt{Hz}}$. Similarly, our ADCs have 12 bits of precision, which are all used when the input is 2V peak-to-peak. This works out to a voltage noise at the ADC input of 44.61 $\frac{nV}{\sqrt{Hz}}$.

5.2.4 Johnson Noise

Johnson-Nyquist noise (also known as "Thermal noise"), is most often talked about in regards to resistors, but is a property of all ohmic devices. It is related to the thermal excitation of charge-carriers within the material, exchanging energy with their surroundings at equilibrium, and will therefore be present in a resistor independent of any external applied voltage or current. It has a white spectrum, and can be modeled as an amplitude spectral density of either a current source in parallel, or a voltage source in series, with a resistance

$$v_{johnson} = \sqrt{4k_BTR}\ \left[\frac{V}{\sqrt{Hz}}\right] \tag{5.4}$$

$$i_{johnson} = \sqrt{\frac{4k_BT}{R}} \left[\frac{A}{\sqrt{Hz}}\right] . \tag{5.5}$$

Where K_b is the Boltzmann Constant in units of $[JK^{-1}]$, T is the temperature of the conductor in kelvin, and R is the effective resistance of the element.

The Johnson noise of a resistive network is equivalent to the Johnson noise of the equivalent resistor of that network, and so the Johnson noise of the carrier and nuller networks can be reduced to a signal element. This is not necessarily true if there are active gain stages between elements, as in the case of the digitization chain, which requires a slightly different analysis in Section 5.5.

5.3 Sources of Noise in the Nuller Signal Path

The nuller signal path is the easiest to calculate and apply; it simply drives a current directly to the SQUID input coil, and so the natural reference frame in which to evaluate the noise also happens to be at the location we are most interested in. To get there only "conventional" warm electronics are in involved, and we need not consider the more pernicious elements of the cold hardware. To first order we simply assume that the current driven by the nuller goes exclusively through the SQUID input coil, such that the transfer function applied to the sources of noise in the nuller circuit is independent of frequency and dynamic system state.

5.3.1 Nuller-chain DAC and Quantization Noise

The nuller DAC noise and quantization noise are both calculated as noise sources at the DAC output – and so the transfer function relating the amplitude of a current at the output of the nuller DAC to a current at the SQUID input coil is just directly applied. This is a multiplicative factor of 0.0687, derived in Section 4.2.1, and presented in Table 4–1. When applied to the 60 $[\frac{pA}{\sqrt{Hz}}]$ fixed DAC noise described in Section 5.2, this gives a current noise at the SQUID input coil of 4.122 $[\frac{pA}{\sqrt{Hz}}]$.

From Section 5.2 we know the quantization noise is 11.12 $[\frac{pA}{\sqrt{Hz}}]$ at the DAC output, and so at the SQUID coil it multiplies down to 0.77 $[\frac{pA}{\sqrt{Hz}}]$.

5.3.2 Nuller-chain Johnson Noise

The Johnson noise of the nuller-chain at the SQUID input coil is entirely dominated by the resistors on the SQUID controller board. One way to see this is that the amplifiers on the mezzanine form a virtual ground at their output. Therefore, the Johnson noise current resulting from the two 249Ω gain resistors are sunk directly into the "virtual short" between the outputs of both amplifiers. This is equivalent to saying that the amplifiers suppress their own current noise, and the degree to which they do is related to the gain.

Another way to see this is that those resistors, and the two 10Ω resistors at the mezzanine output, do produce Johnson noise. However, most of this is confined to the loop made by the 100Ω parallel resistor on the SQUID controller board (R_1 in Figure 2–3), and never makes it to the SQUID input coil. Therefore, the Johnson noise seen by the SQUID input coil is completely dominated by the loop formed by the 100Ω and four 750Ω resistors on the SQUID controller board. These form an effective impedance of 3100Ω in series with the SQUID input coil. Using Equation 5.5, this provokes a Johnson noise current at the SQUID coil of 2.31 $\frac{pA}{\sqrt{Hz}}$.

5.3.3 Nuller-chain Noise Totals

A summary of the noise sources from the nuller chain, referred to the input of the SQUID coil, is given in Table 5–1.

5.4 Sources of Noise in the Carrier Signal Path

Unlike the nuller, the carrier noise can only be derived numerically from the circuits as a voltage-noise across the comb. To refer it to a current at the SQUID input coil requires

Source	Location	Noise [$\frac{pA}{\sqrt{Hz}}$]
DAC Noise	DAC	60
DAC Digitization Noise	DAC	11.12
DAC Noise	SQUID Input	4.12
DAC Digitization Noise	SQUID Input	0.77
Johnson Noise	SQUID Input	2.31
Total	SQUID Input	4.79

Table 5–1: The nuller-signal-path noise contributions at their sources and as a current noise at the SQUID input coil.

making an assumption about the effective real impedance of the comb. That impedance varies according to resistance of the bolometers, and the impedances in the LC comb. The latter includes the discrete impedances from the LC resonances, but also any stray series impedance z_s. z_s comes largely from the Equivalent Series Resistance (ESR) in the capacitors of the LC-filters, which for our system varies considerably across the band – at the upper frequencies it can be as much as 1.5Ω. However, with the more advanced cold-hardware components currently being fabricated, a typical series impedance is closer to 0.25Ω, which is roughly consistent with the series impedances present in the Legacy 16x DfMUX systems. A large series impedance has serious implications for bolometer noise, but the carrier noise at the SQUID input coil doesn't distinguish between the sources of resistance in the LCR comb, and in this case a series impedance actually attenuates the contribution of the carrier-chain noise to the total readout noise at the SQUID input coil.

Because the effective resistance of the comb is strongly frequency dependent, and the bolometer resistance is dynamic, the relative contribution of carrier-path noise as a current at the SQUID input coil is sensitive to the state of the comb at the frequency of interest. In this section the carrier-path noise sources will be quantified as a voltage noise across the comb (an account of the noise that is independent of the cold hardware); as general equations to calculate the current at the SQUID input coil for a given resistance across the comb; and numerically for a few common comb configurations. Our five comb configurations consist of an overbiased bolometer ($R_{bolo} = 1\Omega$), and a bolometer in the transition ($R_{bolo} = 0.7\Omega$), both in the presence of no series impedance, and with a $z_s = 0.25\Omega$; as well as an off-resonance channel ($R_{eff} = 10\Omega$). Results are presented in Table 5–2.

5.4.1 Carrier-chain DAC and Quantization Noise

Using the transfer function factor of 0.0747 given in Table 4–2, and the intrinsic DAC output noise of $60\frac{pA}{\sqrt{Hz}}$, we recover a voltage noise across the comb of $4.482\frac{pV}{\sqrt{Hz}}$. To refer this to a current at the SQUID input coil for a given resistance across the entire comb

$$i^{carrier}_{DAC\ Noise} = \frac{4.482}{R_{comb}} \left[\frac{pA}{\sqrt{Hz}}\right] . \tag{5.6}$$

As before, the digitization noise at the DAC is $11.12\ \frac{pA}{\sqrt{Hz}}$. Applying the same 0.0747 transfer function, this becomes $0.833\frac{pV}{\sqrt{Hz}}$ across the comb. Referred to the SQUID input coil, the equation is

$$i^{carrier}_{Quant\ Noise} = \frac{0.833}{R_{comb}} \left[\frac{pA}{\sqrt{Hz}}\right] . \tag{5.7}$$

5.4.2 Carrier-chain Johnson Noise

The Johnson noise from the carrier path is separated into the warm (300K components) and cold (4K 30mΩ bias resistor). For the same reasons described in Section 5.3, in calculating the Johnson noise of the warm carrier path we do not consider the 249Ω resistors on the mezzanine amplifiers. The relevant equivalent resistance is then just 100Ω. Using Equation 5.5, this works out to a current noise at the bias resistor of 12.87 $[\frac{pA}{\sqrt{Hz}}]$, or a voltage noise across the comb of $0.386\ \frac{pV}{\sqrt{Hz}}$. The conversion to a current at the SQUID input coil requires a division by the effective comb impedance,

$$i^{carrier}_{Johnson\ (300K)} = \frac{0.386}{R_{comb}} \left[\frac{pA}{\sqrt{Hz}}\right] . \tag{5.8}$$

For the 30mΩ resistor at 4K, the Johnson noise from the bias resistor itself across the comb is 2.574 $\frac{pV}{\sqrt{Hz}}$, or referred to the SQUID input coil

$$i^{carrier}_{Johnson\ (R_{bias})} = \frac{2.574}{R_{comb}} \left[\frac{pA}{\sqrt{Hz}}\right] . \tag{5.9}$$

5.4.3 Carrier-chain Noise Totals

All of the values computed in 5.4, including examples of the comb configurations introduced at the beginning of this section, can be found in Table 5–2.

Source	Location	Noise Calculated Independent of R_{comb}		
DAC Noise	DAC	60 $[\frac{pA}{\sqrt{Hz}}]$		
Digitization Noise	DAC	11.12 $[\frac{pA}{\sqrt{Hz}}]$		
DAC Noise	Bolometer	4.48 $[\frac{pV}{\sqrt{Hz}}]$		
Digitization Noise	Bolometer	0.83 $[\frac{pV}{\sqrt{Hz}}]$		
Johnson (300K)	Bolometer	0.386 $[\frac{pV}{\sqrt{Hz}}]$		
Johnson (R_{Bias})	Bolometer	2.57 $[\frac{pV}{\sqrt{Hz}}]$		
Total	SQUID	$\left(\frac{1}{R_{comb}}\right)$ 5.245$[\frac{pA}{\sqrt{Hz}}]$		
		Noise for 5 comb configurations $[\frac{pA}{\sqrt{Hz}}]$		
Source	Location	$R^{norm}_{bolo} = 1\Omega$ ($z_s = 0.25\Omega$)	$R^{trans}_{bolo} = 0.7\Omega$ ($z_s = 0.25\Omega$)	Off-Resonance
DAC Noise	SQUID	4.48 (3.58)	6.40 (4.72)	0.448
Digitization Noise	SQUID	0.83 (0.66)	1.12 (0.87)	0.083
Johnson (300K)	SQUID	0.386 (0.309)	0.552 (0.406)	0.0386
Johnson (R_{Bias})	SQUID	2.57 (2.06)	3.67 (2.71)	0.257
Total	SQUID	5.25 (4.19)	7.49 (5.52)	0.525

Table 5–2: The carrier-signal-path noise contributions when referred to locations independent of the comb, and in the lower section presented again with several variations of bolometer resistance and series resistance.

5.5 Sources of Noise in the Demodulation Signal Path

Like the carrier-path, the demodulation-path also cannot be referred back to a current at the input of the SQUID coil without making an assumption about the state of a dynamic element of the cold-electronics: the SQUID itself. The natural reference for the noise contributions to the demodulation path is as a voltage at the SQUID output; in order to refer that to a current at the SQUID input we must divide by the SQUID transimpedance. A high-performing SQUID will have a transimpedance of approximately 500Ω. A summary of the noise in the demodulation path can be found in Table 5–4.

5.5.1 Demodulation-chain Amplifier Noise

Of the amplifiers, only the first stage amplifier plays a significant role. This is because the input noise of each subsequent amplifier, when referred back to the SQUID, is divided by all of the gain stages between it and the SQUID. The first stage amplifier has a gain of 16, which already makes the second stage amplifier (at twice the noise and a gain of 5) sub-dominant by a factor of 3. A breakdown of the first, second, and third stage amplifier noises is presented in Table 5–3.

Source	Intrinsic Input Noise $\left[\frac{nV}{\sqrt{Hz}}\right]$	Referred to SQUID input coil ($Z_{SQUID} = 500\Omega$) $\left[\frac{pA}{\sqrt{Hz}}\right]$
First Stage Amplifier (LT6200)	0.95	1.92
Second Stage Amplifier (AD8138)	5	0.63
Third Stage Amplifier (AD8138)	5	0.13
Total Amplifier Noise	1.01	2.02

Table 5–3: The demodulator-signal-path amplifier-noise contributions as a current noise at the SQUID input coil. Note that for the total amplifier noise sum they add in quadrature.

5.5.2 Demodulation-chain Johnson Noise

For the same reason that the second and third stage amplifier noises are insignificant relative to the noise of the first stage amplifier, Johnson noise from all resistors behind the first stage amplifier are also insignificant compared to those in front of it. The only resistors which really contribute to the Johnson noise are the two that determine the gain of the first stage amplifier: a 10Ω resistor and 150Ω feedback resistor (refer to Figure 4–5). Together, these source an equivalent voltage noise of 1.63 $\frac{nV}{\sqrt{Hz}}$, which when referred back to a current at the input of the SQUID coil (assuming $Z_{SQUID} = 500\Omega$), is 3.26 $\frac{pA}{\sqrt{Hz}}$.

5.5.3 Demodulation-chain Quantization Noise

From Section 5.2.3, the ADC digitization noise at the ADC itself is 44.61 $\frac{nV}{\sqrt{Hz}}$. To get this as a voltage at the SQUID output we divide by the total gain of the demodulation path, 440, from Table 4–3. This yields $0.101\frac{nV}{\sqrt{Hz}}$, or just 0.203 $\frac{pA}{\sqrt{Hz}}$ for $Z_{SQUID} = 500\Omega$.

5.5.4 Demodulation-chain Noise Totals

A summary of the demodulation-chain noise contributions is given in Table 5–4.

5.6 SQUID and Bolometer Noise

There are four final sources of noise to consider; the first is an intrinsic SQUID noise, for which we use as a model the SQUIDs that have flown on the balloon-based EBEX experiment in 2014. These had a measured SQUID noise of ~ 3.5 $\frac{pA}{\sqrt{Hz}}$ [3]. The SQUID noise can vary significantly from SQUID-to-SQUID, and it is likely that the SQUIDs used for testing this hardware have worse noise performance. This will be considered more in Chapter 6, but as before these calculations assume a high-performing SQUID.

The other three sources originate with the bolometers. Noise characteristics of bolometers change depending on whether or not they are normal (in which case they are simply 1Ω

Source	Location	Noise
Digitization Noise	ADC	44.62 $[\frac{nV}{\sqrt{Hz}}]$
Digitization Noise	SQUID Output	0.101 $[\frac{nV}{\sqrt{Hz}}]$
Amplifier Noise	SQUID Output	1.01 $[\frac{nV}{\sqrt{Hz}}]$
Johnson Noise	SQUID Output	1.63 $[\frac{nV}{\sqrt{Hz}}]$
Digitization Noise	SQUID Input ($Z_{SQUID} = 500\Omega$)	0.203 $[\frac{pA}{\sqrt{Hz}}]$
Amplifier Noise	SQUID Input ($Z_{SQUID} = 500\Omega$)	2.02 $[\frac{pA}{\sqrt{Hz}}]$
Johnson Noise	SQUID Input ($Z_{SQUID} = 500\Omega$)	3.26 $[\frac{pA}{\sqrt{Hz}}]$
Total Digitization Chain Noise	SQUID Input	$\frac{1920}{Z_{SQUID}}$ $[\frac{pA}{\sqrt{Hz}}]$
Total Digitization Chain Noise	SQUID Input ($Z_{SQUID} = 500\Omega$)	3.84 $[\frac{pA}{\sqrt{Hz}}]$

Table 5–4: The total demodulator-signal-path noise contributions as a both voltage noise at the SQUID output, and current noise at the SQUID input coil.

resistors and suffer only from Johnson noise), or if they are in their superconducting transition state. While superconducting, the bolometers begin “responding” to photons, such that Poisson (“shot”) noise becomes relevant. The characterization of this readout system does not require an optical instrument; as such the bolometers are operated “dark”, looking only at the cold blackbody of the cryostat. We therefore do not consider photon shot-noise in our noise predictions; however, since it is a fundamental noise of the measurement, it is noteworthy as a benchmark below which the noise of the instrument itself must be sub-dominant. This is what is meant by the terminology “photon noise” dominated.

The final source is one that is typically the dominant noise source of a dark system: “phonon noise”. This noise is present in any superconductor, and is exacerbated by the weak thermal link between the bolometer and its temperature bath. Phonon noise will be be described in Section 5.6.2. A complete summary of the SQUID and bolometer noise can be found in Table 5–5.

5.6.1 Bolometer Johnson Noise

When held normal, the Johnson noise can be calculated using Equation 5.5. With a resistance of 1Ω, at a temperature of 800mK, this is 6.65 $\frac{pA}{\sqrt{Hz}}$. When a bolometer is lowered into the transition, and is experiencing strong electro-thermal feedback, Equation 5.5 no longer accurately describes the system. In such a scenario, consider the Johnson noise as a voltage noise source is series with the bolometer as in Figure 5–2.

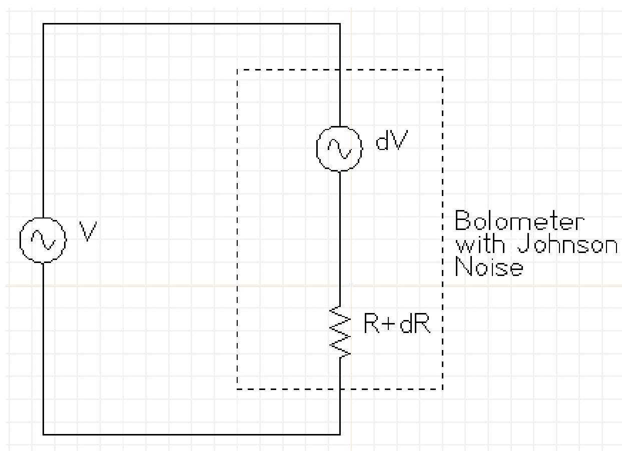

Figure 5–2: An cartoon diagram modeling bolometer Johnson noise as a voltage noise source in series with the bolometer. Image Credit: Tijmen De Haan.

As the output of this voltage noise source varies, it deposits power on the bolometer which changes the bolometer resistance. The negative electrothermal-feedback of the bolometer counters this variation in the resistance, resulting in an overall suppression of the Johnson noise by a factor of $\frac{1}{1+\mathcal{L}}$ [16], where the loop-gain, $\mathcal{L}$, for a typical bolometer in its transition is 10.

The precise derivation of this for an AC-biased bolometer is sufficiently complicated to be beyond the scope of this document, but a consequence is that when the bolometer signal is demodulated using a sinusoid in phase with the bolometer responsivity, the Johnson noise is suppressed as above by $\frac{1}{1+\mathcal{L}}$. However the out-of-phase component of the time-stream does not benefit from this suppression. A modified Equation 5.5 for the Johnson noise as a current noise in the presence of electro-thermal feedback is given in Equation 5.10. Note that when the bolometer is over-biased, it has a loop gain of 0.

$$i_{bolo\ Johnson} = \frac{1}{1+\mathcal{L}}\sqrt{\frac{4k_BT}{R}}\ \left[\frac{A}{\sqrt{Hz}}\right] \tag{5.10}$$

For a bolometer in its transition, with a resistance of 0.7Ω, a temperature of 485mK, and a loopgain of 10, $i_{bolo\ Johnson} = 0.562\ [\frac{pA}{\sqrt{Hz}}]$ – a substantial suppression.

5.6.2 Bolometer Phonon Noise

Phonon noise is a type of thermal noise, similar to but distinct from, Johnson noise. In the case of phonon noise, the agents of energy exchange between the bolometer and thermal environment are phonons – quantized vibrational modes – not charge carriers. This is defined as a Noise Equivalent Power at the bolometer,

$$NEP_{phonon} = \sqrt{4k_B(\gamma T_{bolo}^2)G}\ . \tag{5.11}$$

Where G is the thermal conductance between the bolometer and the 285mK thermal bath (typically $\sim 100\ \frac{pW}{K}$), T_{bolo} is 485mK, and γ is a correction factor that takes into account the thermal gradient between the bolometer and the 258mK bath. The factor of $\gamma^{\frac{1}{2}}T_{bolo}$ in Equation 5.11 is the "effective temperature" of the bolometer. A typical γ for our test-setup is ~ 0.5 [3], which results in an NEP of 25.5 $\frac{aW}{\sqrt{Hz}}$.

In deriving this we've only had to make assumptions about quantities that are insensitive to the actual operation of the bolometer – the thermal conductivity is fixed by the detector fabrication, as is the transition temperature T_{bolo}; and γ is independent of bolometer bias. If we wish to convert the noise power above into a current at the SQUID input, we must know the *responsivity* of the bolometer: the change in current through the bolometer for a given change in power. In particular, Equation 5.11 becomes

$$i_{phonon} = S\sqrt{4k_B(\gamma T_{bolo}^2)G}\ , \tag{5.12}$$

where the responsivity $S = \frac{\partial I}{\partial P}\left[\frac{A}{W}\right]$. A typical responsivity factor for a 0.7Ω bolometer is $\sim 6.5{\times}10^5 \frac{A}{W}$, which predicts a phonon noise at the SQUID input coil of 16.57 $\frac{pA}{\sqrt{Hz}}$.

5.6.3 Increased Responsivity In the Presence of Series Impedance

We've mentioned several times in this document how a stray series impedance in the bolometer circuit can degrade performance by providing a mixed voltage-current bias. In Section 1.3 we stated that a current bias, because it spoils the negative electro-thermal feedback, has a deleterious effect on bolometer stability. This idea is tightly coupled with the the fact that it can dramatically increase bolometer responsivity; and in so doing it exacerbates bolometer phonon noise (Equation 5.12). A complete derivation of the increased responsivity due to a series impedance with the bolometer is not possible here, but can be found in [7]. The equation for bolometer responsivity, which has been derived with a stray series impedance z_s, is as follows

$$S = -\frac{1}{V_{bias}}\frac{\mathcal{L}}{\mathcal{L}+1}\left(1+\frac{2z_s}{R_{bolo}^{(0)}}\frac{\mathcal{L}}{\mathcal{L}+1}\right)\ , \tag{5.13}$$

such that $R_{bolo}^{(0)}$ is the first order term of a series expansion of R. In the limit of high-loop-gain (as would be experienced by a bolometer in its transition) this becomes

$$\lim_{\mathcal{L}\to+\infty} S = -\frac{1}{V_{bias}}\left(1+\frac{2z_s}{R_{bolo}^{(0)}}\right)\ . \tag{5.14}$$

For a 0.7 Ω bolometer with a series impedance $z_s = 0.25\Omega$, the responsivity goes up by a factor of 1.65. This implies that in the presence of a 0.25 Ω stray impedance, the phonon noise at the SQUID input coil goes up to 27.34 $\frac{pA}{\sqrt{Hz}}$. Note that a typical voltage bias amplitude for a bolometer in the transition is about 2.1uV (RMS).

5.6.4 Bolometer Photon Noise

Bolometer Photon noise is described as the Poisson-statistical variability in the arrival of absorbed photons from the CMB, and is given as a power-noise at the bolometer. Like the bolometer Johnson noise, the bolometer responsivity is required to refer that power noise back to a current at the SQUID input coil. This is given in Equation 5.15

$$i_{photon} = S\sqrt{2h\nu P_{rad}}\,. \tag{5.15}$$

Where h is the Planck constant, ν is the frequency of observed photons, and P_{rad} is the incident optical power.

This relationship holds assuming that the photons arriving are uncorrelated, but must be modified in the case that they are (as in polarized detection). This modification is treated as a separate component, known as Photon Correlation noise, and given as

$$i_{photon} = S\sqrt{\zeta_{corr}\frac{P_{rad}^2}{\Delta\nu}}\,. \tag{5.16}$$

Where $\Delta\nu$ is the bandwidth of the detection instrument, and ζ_{corr} is a correlation factor between 0 (completely uncorrelated) and 1 (completely correlated) to account for different degrees of correlation. For detectors only sensitive to one polarization, such as in current CMB bolometer instruments, ζ_{corr} is fixed at 1. In earlier spider-web absorber bolometers, which did not distinguish between the polarizations of incoming light, this was typically set to 0.3 [7].

When observing the CMB, using example parameters from the South Pole Telescope instrument [1] , the bolometer photon noise is 20.66 $\frac{pA}{\sqrt{Hz}}$, and the photon correlation noise $26.33\frac{pA}{\sqrt{Hz}}$. This value is provided as reference, and will not be included in the noise prediction tables. As we shall see, the total photon (fundamental) noise contribution of approximately

[1] $\zeta = 1$, $\nu = 150GHz$, $\Delta\nu = 39GHz$, and $P_{rad} = 6.3pW$.

Source	Location	Noise
Bolometer Phonon Noise	Bolometer	25.5 $[\frac{aW}{\sqrt{Hz}}]$
SQUID Noise	SQUID Input	3.5 $[\frac{pA}{\sqrt{Hz}}]$
Bolometer Johnson Noise ($R_{bolo}^{norm} = 1\Omega$, $\mathcal{L} = 0$)	SQUID Input	6.65 $[\frac{pA}{\sqrt{Hz}}]$
Bolometer Johnson Noise ($R_{bolo}^{trans} = 0.7\Omega$, $\mathcal{L} = 10$)	SQUID Input	0.562 $[\frac{pA}{\sqrt{Hz}}]$
Bolometer Phonon Noise ($z_s = 0\Omega$)	SQUID input	16.57 $[\frac{pA}{\sqrt{Hz}}]$
Bolometer Phonon Noise ($z_s = 0.25\Omega$)	SQUID input	27.34 $[\frac{pA}{\sqrt{Hz}}]$

Table 5–5: Bolometer noise sources and typical values. Assumed quantities are: $S = 6.5{\times}10^5 \frac{A}{W}$, $G = 100\frac{pW}{K}$, $T_{bolo} = 485\ mK$, $\gamma = 0.5$.

33 $\frac{pA}{\sqrt{Hz}}$ is above, or equivalent to, the total noise when operating a bolometer in-transition (with $z_s \sim 0.25[\Omega]$), as it was in Legacy 16x DfMUX systems. [3]

5.6.5 SQUID and Bolometer Noise Totals

A summary of SQUID and bolometer noise sources, and their typical values assuming some bolometer properties based on EBEX and SPT bolometers, can be found in Table 5–5.

5.7 Prediction

At the SQUID input coil all of the noise sources detailed in the preceding sections are uncorrelated, and thus sum in quadrature. A large table that summarizes the results of all derived noise quantities is presented in Table 5–6. Total noise predictions that combine all of these sources for our five of example comb configurations (normal and in-transition bolometer, both with and without a small stray series impedance, and for frequencies that are off-resonance) are presented in Table 5–7. Notably, the prediction for noise performance off-resonance, given high-performance cold components that would be typical on a deployed instrument, is 7.08 $\frac{pA}{\sqrt{Hz}}$.

Source	Type	Noise $[\frac{pA}{\sqrt{Hz}}]$
SQUID	Intrinsic	3.5
ADC-chain	Total	$\frac{1920}{Z_{SQUID}}$
Nuller-chain	Total	4.79
Carrier-chain	Total	$\left(\frac{1}{R_{comb}}\right)$ 5.245

		Carrier-chain and Bolometer Noise for 5 comb configurations $[\frac{pA}{\sqrt{Hz}}]$		
Source	Type	$R_{bolo}^{norm} = 1\Omega$ ($z_s = 0.25\Omega$)	$R_{bolo}^{trans} = 0.7\Omega$ ($z_s = 0.25\Omega$)	Off-Resonance
Carrier-chain	Total	5.25 (4.19)	7.49 (5.52)	0.525
Bolometer Noise	Total	6.65 (6.65)	16.58 (27.35)	0

Table 5–6: Upper: A summary of the bolometer noise sources and typical values expressed generally where possible. Lower: specific comb configurations, with both overbiased and in-transition bolometers. In parenthesis are values in the presence of a 0.25Ω series impedance. Assumed quantities are: $S = 6.5{\times}10^5$, $G = 100[\frac{pW}{K}]$, $T_{bolo} = 485\ [mK]$, $\gamma = 0.5$, $\mathcal{L}_{bolo}^{trans} = 10$.

Prediction	Noise ($z_s = 0.25$) $[\frac{pA}{\sqrt{Hz}}]$
Off-Resonance	7.08 (7.08)
Overbiased	11.49 (10.57)
In Transition	19.51 (29.22)

Table 5–7: A summary table of the expected noise with different configurations. Assumed quantities are: $Z_{SQUID} = 500\Omega$, $R_{bolo}^{OB} = 1.0\Omega$, $R_{bolo}^{trans} = 0.7\Omega$.

5.8 An Additional Note: $\sqrt{2}$ Demodulation and Sideband Factors

Although we have focused on the noise as a current at the input of the SQUID coil, which is effectively what we are measuring with the channelized DAN data, there are two additional steps: demodulation, and decomposition. In all but the power-noise sources (Phonon, Photon), the noise is uncorrelated across both sidebands of the bolometer bias frequency. When we demodulate the current waveform, those sidebands add incoherently, and thus so does the noise in each band. The *demodulated* noise contributions from each of the above sources (again, aside from Phonon and Photon) is therefore multiplied up by a factor of $\sqrt{2}$.

This is, however, undone in the act of decomposing that timestream into I and Q. The broadband noise that is uncorrelated with the carrier voltage power (so, non-power noise sources) in the demodulated timestreams is split evenly between I and Q.

Phonon and Photon noise are power noise. Power noise sources are split into correlated sidebands, and therefore not "enhanced" by demodulation when the sidebands are combined. Because they are aligned with the carrier voltage sinusoid, they only appear in the in-phase projection of the demodulated timestream, and thus are also not divided down by a factor of $\sqrt{2}$ by taking I-only. [7]

The result is that noise seen in an in-phase demodulated timestream is the same as the noise referred back to a current noise at the input of the SQUID coil, though this is serendipitous, which often leads to confusion regarding these factors.

CHAPTER 6
Test Setup and Additional Pathologies

In Chapter 5, we suggested that there are some peculiarities in the McGill testing hardware and measurement setup, which result in parameters that deviate from the assumptions used to calculate noise predictions in a science-deployment environment. The corrections fall into two categories: updates to a single value (such as SQUID transimpedance), or a modification of a quantity that was used in calculations calculations as a fixed parameter, but actually varies across our bandwidth (such as stray series impedance).

In some cases these corrections are because the components on hand were simply not deployment-grade. In other cases we were faced with the fact that the some of the equipment used was not optimized for low-noise performance, and that warm hardware development for high-frequency and high-multiplexing operation outpaced the development of cold development.

This chapter explores the practical considerations involved in our system setup and noise measurement. We present more accurate parameters to be used in predicting the performance; and in the case where the parameter is known to vary, we present simulations showing its effect on the overall noise performance, and experimentally verify these predictions to the extent that we can. We start with a description of the cold hardware, and finish with a discussion of narrow-band noise sources, and their effects on system performance.

6.1 64x from 16x: Cold Electronics For Flight Representative Testing

In December of 2013, as the hardware commissioning of the the flight-representative readout electronics was underway, no cold-hardware specifically designed to be capable of $>$ 16x multiplexing, or frequencies above 1MHz, existed. In order for us to evaluate a 64x readout system, it was necessary to first devise hardware it could operate. It is worth noting that the design, fabrication, and testing of cold-hardware suitable for 64x multiplexing is an active area of research and development. For the SPT-3G and POLARBEAR2 experiments, this is taking place at universities and national laboratories across Canada and the USA. It is only in December of 2014 that this development has progressed to the point of testing dedicated 64x LC-boards, with older-generation bolometers. A full integration of next-generation bolometer wafers and LC-boards designed specifically for higher multiplexing is scheduled to take place in March, 2015.

With that in mind, we sought to construct a set of cold hardware that was sufficiently operable to demonstrate the warm electronics readout hardware, but were not expecting it to be representative of final deployable cold-hardware. To do this we combined several existing 16x LC-boards and a non-science grade bolometer wafer, fabricated for the SPT Polarimeter.

6.1.1 Bolometer Wafer: SPTpol5_B2

The SPTpol5_B2 wafer used here was rejected for use on the telescope due to a combination of low yield, high scatter in the normal bolometer resistance, and a time-constant that was slightly faster than desirable. The median stray series impedance measured for bolometers across the whole wafer, when operated in a legacy 16x configuration, was 0.29Ω. The median bolometer normal resistance was 1.54Ω. This means that when operating the bolometers at a typical depth in transition of $0.7R_{bolo}^{norm}$, they have a typical resistance of 1.08Ω.

The B2 wafer supported enough bolometers for a high multiplexing comb, but these traces had to be ganged together and then attached to a set of Legacy 16x LC boards (Figure 6–1).

This was accomplished using a small PCB "combiner-board", installed between the SQUID striplines and the LC boards at 285mK. The additional soldering, and unavoidably

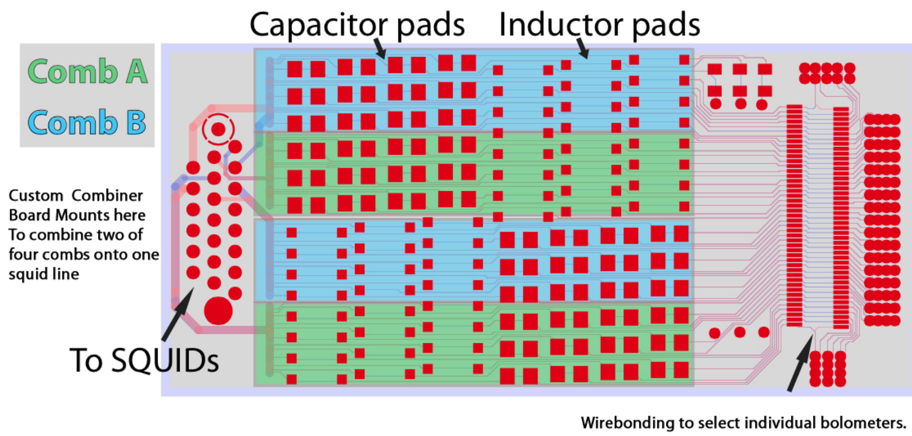

Figure 6–1: A schematic image of an LC-board with two combs. The custom combiner board takes the SQUID input and combines several modules onto a single SQUID channel – modifications to the wirebonding on the focal-plane side are how we select individual bolometers from the wafer. [25]

non-ideal layout of superconducting traces, wirebonds, and solder, next to one another on a small PCB has the potential to introduce strays at the cold stage in ways that are difficult to quantify. The process of plucking and re-wire-bonding the wafer traces was a delicate one not without losses; 44 bolometer / LC-channel pairs were able to be recovered in this way onto a single comb.

6.1.2 64x With Commercial Capacitors

The LCs themselves posed a separate challenge: each of the LC-boards that we tied together to get to a 64x multiplexing factor were initially designed to share the same 1MHz of bandwidth. Frequencies are selected with 24uH inductors in series with commercial ceramic capacitors, such that the precise capacitance is varied to tune the LC resonance. Unfortunately, these ceramic capacitors suffer from high Equivalent Series Resistance, a problem that is compounded at higher frequencies and lower capacitances. ESR is an effective real resistance originating from losses in the capacitor dielectric, and acts as a series impedance in the bias circuit. Next-generation LC boards being fabricated for SPT-3G and POLARBEAR2 are specifically engineered, "inter-digitated", capacitors whose dielectric is just the vacuum space between finger-like traces, etched onto a substrate. These are specifically designed as such in order to exhibit extremely low ESR.

For our setup, we used the same ceramic capacitors employed in the 1MHz LCs, with capacitances that range from 40pF to 10nF. The result is that at the frequencies below 1MHz we recover the same series impedance as measured during the characterization of the Sptpol15_B2 wafer, $\sim 0.29\Omega$. However, as the frequency climbs, so does ESR, with a relationship that is linearly proportional. The highest series impedance in any LCR resonance (which will be dominated by the capacitor ESR) is measured at $\sim 1.5\Omega$.

Another modification to the modern LC-combs designed for this higher bandwidth is the use of larger, 60uH, inductors. This is partially motivated by the difficulty in fabricating inter-digitated capacitors with large capacitance, but also has the important benefit of making the resonance peaks narrower. The bandwidth of an LCR filter is given by $\Delta\omega = \frac{R}{L}$. Increasing the inductance, and thereby narrowing the bandwidth, minimizes potential crosstalk between neighboring peaks, and allows for a denser comb of bolometers over the same bandwidth. The 24uH inductors used to commission this system resulted in wide peaks with overlapping bandwidths that caused higher crosstalk between readout channels, and made it difficult to find the bandwidth in which to fit 64 channels.

Overlapping LC-resonances are problematic primarily when considering the efficacy of the bolometers as detectors, causing "signal" to leak from one to another. However, leakage current crosstalk in the comb can influence noise properties of both bolometers, and off-resonance channels. Consider the LCR network at a single frequency as just a network of resistors, such that the admittance of each LC-filter at that frequency determines the value of the resistors in the legs. In an ideal comb, the network would look like a set of resistors with high resistances ($\sim 10\Omega$), and a single resistance of approximately 1Ω, such that the voltage bias provided to the comb would provoke a large current through the "resonant" leg, and negligible current through the other legs. Overlapping filters, as in Figure 6–2 mean that instead, for every voltage bias frequency, a non-negligible amount of current flows through legs with nearby resonances in frequency space. The consequence of this is that as the resistance of a bolometer in one leg changes, the current through that leg changes substantially (as we hope it does), but the current through a neighboring leg also changes appreciably. By this mechanism, each bolometer can, by just absorbing power, alter the bias being applied to its neighbors.

To see how this affects off-resonance frequencies, consider that, as the large-frequency-spacing approximation fails, the effective impedance of the comb at an off-resonance frequency becomes sensitive to its proximity to neighboring LC-filters. In this scenario, we expect scatter in the off-resonance channel noise that correlates with the comb impedance at that frequency, and comes from the variance in the carrier noise that results from a change in comb impedance.

Although the "signal"-crosstalk that results from leakage current does not directly affect our noise measurements – it can, in the most pathological cases, prevent us from biasing the bolometers, or dropping them into the transition due to destabilizing transient perturbations.

6.1.3 Mapping out the Comb

The LCR resonances are mapped by taking a carrier network analysis – sweeping the bandwidth using a fixed amplitude carrier sinusoid, and measuring the demodulated ADC output. The results can be seen in Figure 6–2. Of the 44x bolometers and resonances, four bolometers were unusable due to close proximity to other resonances. An additional seven suffered from varying degrees of instability or exceptionally high noise, due to either issues

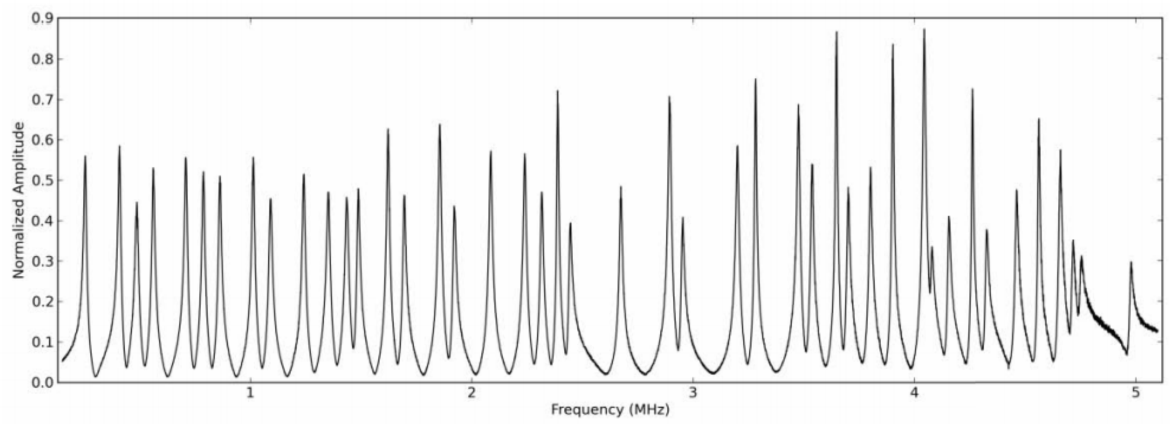

Figure 6–2: A carrier network analysis of the 44x comb. The amplitude on the Y-axis is normalized arbitrarily.

with the bolometers themselves[1] , or exceptionally high scatter in normal resistance or ESR. The best of these could occasionally be biased, but often latched during the drop into the transition. Figure 6–2 makes clear that there is not much room to spare in our bandwidth. Once we populate the off-resonance channels to fill the 64 readout channels, it becomes clear that our model of frequency-independent off-resonance noise is grossly naïve. An example population of 28 off-resonance channels is shown in Figure 6–3. These were selected to maximize the distance between themselves and any neighboring LC-resonance peak.[2]

6.2 Narrow-Band Interference

In Chapter 5 we discussed broadband noise sources, but there is another potential source of noise that comes from narrow-band Electro-Magnetic Interference (EMI) and pickup. Our

[1] Examples of which are large thermal conductivity due to issues with the etching, or malformed transition edge due to unusual geometry or chemical makeup

[2] Note that such a selection is not actually optimum for noise – as the asymmetries in the LC-peaks mean that maximum comb impedance is not found at simply the midpoint between two peaks. As such, this methodology maximizes the number of off-resonance channels that are populated, but does not minimize the scatter in the carrier noise each will experience.

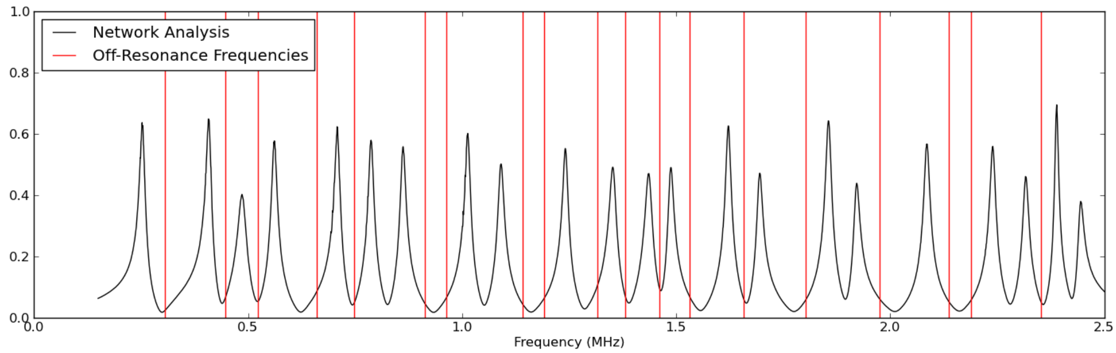

(a)

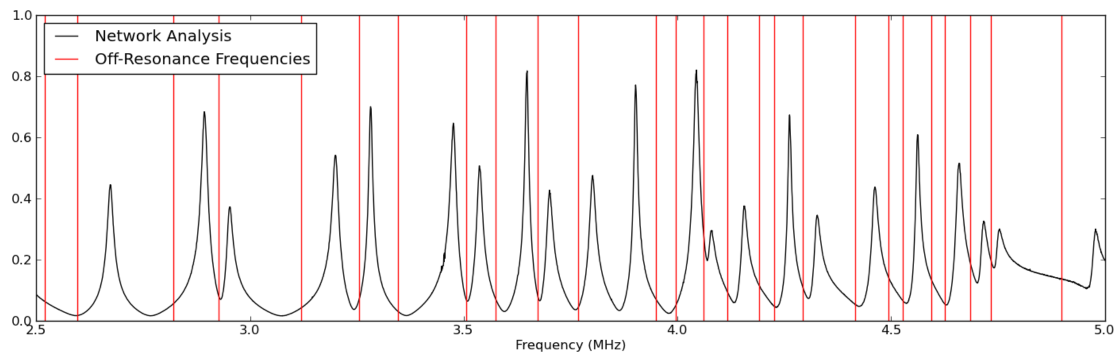

(b)

Figure 6–3: The frequencies of 28 off-resonance channels are selected to maximize the frequency separation between each other, and any LC-resonance peak. Shown here is the network analysis of our comb with the off-resonance channel frequencies superimposed. The images divide the badnwidth into 0-2.5MHz (a), and 2.5-5MHz (b).

bandwidth from 0.1-10MHz is unfortunately situated in a frequency regime commonly used by power switching and communications on electronic devices. Even interference that falls outside of our band can deposit power on the bolometers and the SQUID, and contribute to the total measured noise. By way of example, even in the "dark" configuration of the McGill test-setup, the bolometer-SQUID circuit is sufficiently sensitive that one can clearly see interference from a cell-phone receiving a text-message nearby, or a bluetooth device announcing itself; both of these transmission types are at frequencies greater than 1GHz, but are picked up by the cabling and transmitted into the cryostat where they deposit enough power on the detectors and/or current through the SQUID to be seen as signal in the time-streams. Evaluating the noise in an RF-environment filled with these sources would be extremely difficult, and so the entire testing setup is operated within an RF-tight room (Figure 6–4a). There are no transmission lines that traverse to or from the room, aside from a well filtered power supply. Communication with the readout electronics from outside the RF-room takes place over fiber-optic cable that is fed along a thin, narrow, pipe, which does not admit RF frequencies below a cutoff that is many orders of magnitude above our band. Mounted on the cryostat itself is an additional RF-tight box that contains the SQUID controller board electronics (for reason described in Section 2.4), and can be seen in Figure 6–4b.

6.2.1 Sources of Electro-Magnetic Interference

The RF-room isolates the equipment from EMI in the environment, but does nothing to guard against sources of interference originating with electronics in the RF-room itself. Of these, the spectrum is dominated by electronic noise from the "CryoBoard", a custom thermometry readout and cryogenic control board designed at McGill; and the Kintex-7 FPGA Evaluation motherboard itself (hereafter referred to as the "K7 Board").

Electronic noise enters the system via two means: radiatively, such that emission is coupled into the system via antenna-pickup; and directly, such that EMI generated from

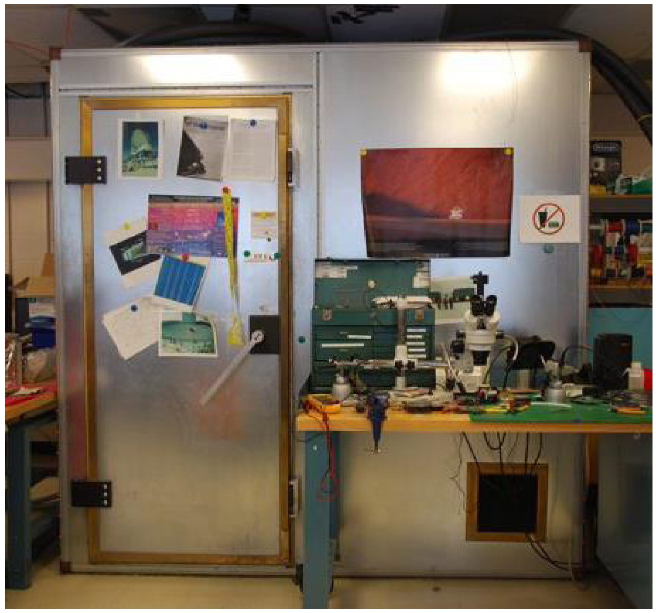

(a) The RF-sealed room where the cryostat and electronics are housed to minimize pickup from the noisy EM environment.

(b) The McGill Dark Cryostat, in the testing configuration with another RF-tight box at the electronics output. This box, seen in the lower right of the image, contains the SQUID controller board. The cabling you see connected is for the thermometry, Helium-3 Fridge control, and the readout.

Figure 6–4: The testing setup of the McGill cryostat in an RF-tight room.

the operation of the electronics itself is not fully isolated from the analog signal paths. A spectrum of data taken by the ADC when operating with tuned SQUIDs, but no synthesizer channels enabled, gives us a "baseline" measure of the EMI environment (Figure 6–5).

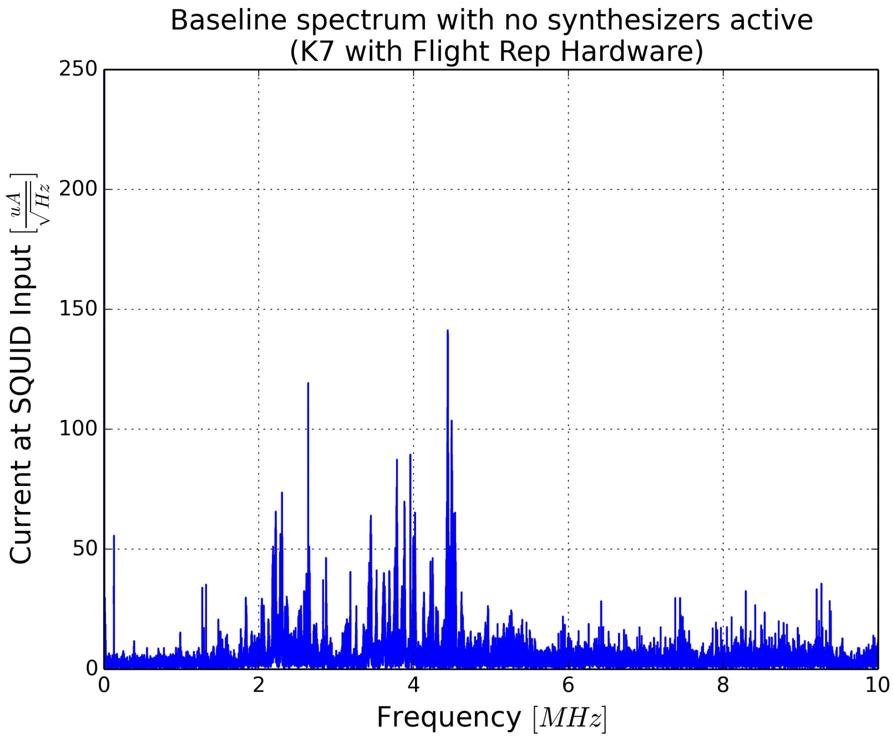

Figure 6–5: A spectrum of the ADC output obtained when looking at tuned SQUIDs, but no carrier or nuller sinusoids. This is a measure of the RFI environment of the electronics. Notice the strong peaks at 2.5, 4, and 4.5 MHz.

In the case of the CryoBoard, the bulk of the EMI is produced in a wide forest near 4MHz. Previously, signals from the CryoBoard at this frequency could simply be filtered out by strong Capacitor-Input (Pi-) Filters in the analog signal path, at the SQUID controller board. The 10MHz bandwidth of the 64x system necessitates much more lenient filtering to avoid attenuation within the band. Noise contamination from the CryoBoard can also enter the signal path radiatively on the 300K side from the electronics themselves and cabling; but also in the cryostat through wires that control heater voltages applied to gas-switches, which thermally couple the fridge and mK-stage. By powering down the CryoBoard during data-taking we can neutralize this EMI, but it makes us vulnerable to thermal drifts of the mK-stage as the gas switches cool. This is addressed to some extent by the use of a temporary battery-pack, which allows us to continue applying a voltage to the switch heaters. By doing this we lose the ability to read-out thermometry, and so it is a measure that may only be used for short intervals, but is sufficient to garner the time-streams used in performing the noise analysis.

We don't have the option to power down the motherboard while taking data, and so EMI from it is unavoidable; sadly, we also suspect that a considerable amount of the total EMI contamination originates with K7 Board. The density and compact nature of this device, and its many features, make it an ideal general developer test-bed, but these design features are antithetical to noise minimization best-practices. In particular there are many different resources on the board that we do not use (such as the GTXs and PCIe) but exist, and are supplied power and clocking. The plethora of supporting power circuits include separate buck converters – ubiquitous devices for performing voltage and current stepping, and which contain the noxious combination of large amplitude signals and high-frequency switching; and clocks of a variety of frequencies and qualities. An annotated image of the various components on a K7 Board can be seen in Figure 6–6.

Some "loud" sources of noise on a motherboard are unavoidable; however, there are measures that can buffer the sensitive signal paths from these contaminates. For instance, on the custom motherboard built for ground-based operation, we synchronize the buck converter switching frequencies, so they appear only as a two single strong lines, at 500KHz and 1 MHz (and their harmonics). This was not possible on the K7 Board – there were simply too many independent bucks, with natural switching frequencies scattered between several hundred KHz and several MHz.

For the most part, EMI to and from digital signal paths is harmless, it is the contamination of analog signal paths that we wish to avoid. The motherboard is host exclusively

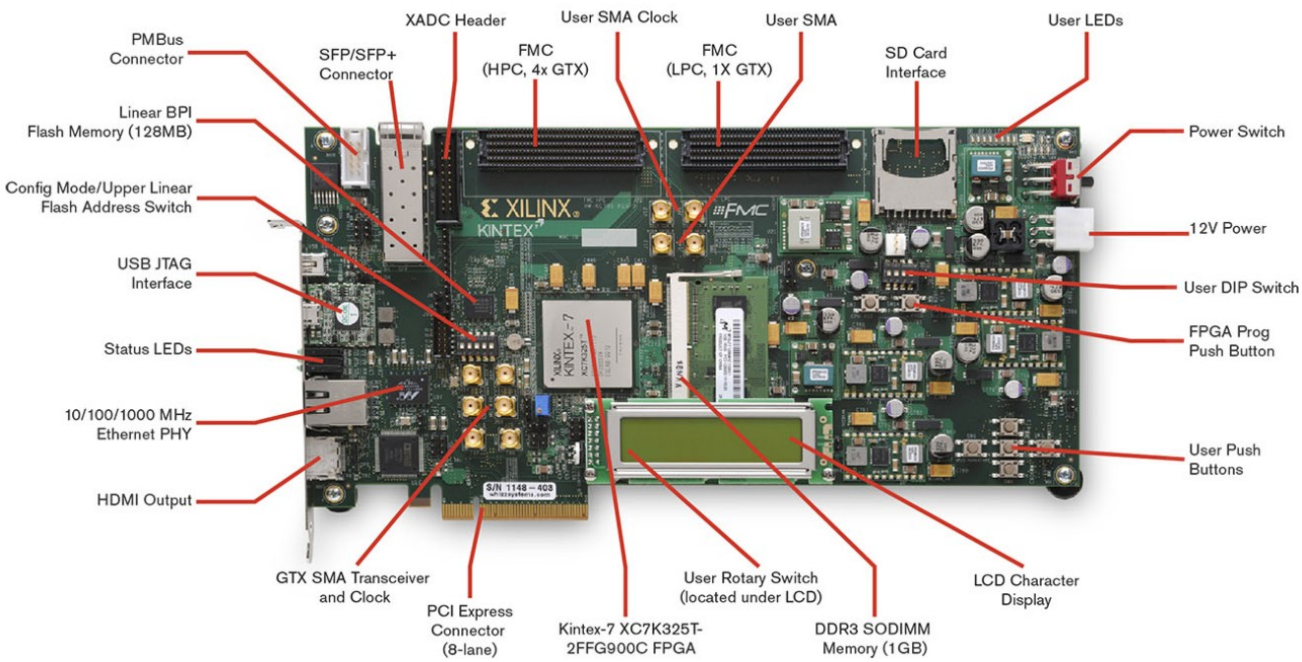

Figure 6–6: An annotated image of the Kintex-7 FPGA Evaluation Board. Note the density and large number of resources. Functions that were not used for DfMUX operation include the GTX and SMA transceivers and ports, PMBus, XADC Header, one of the two FMC slots and communications, and PCIe Express connection. Image Credit: Xilinx Inc.

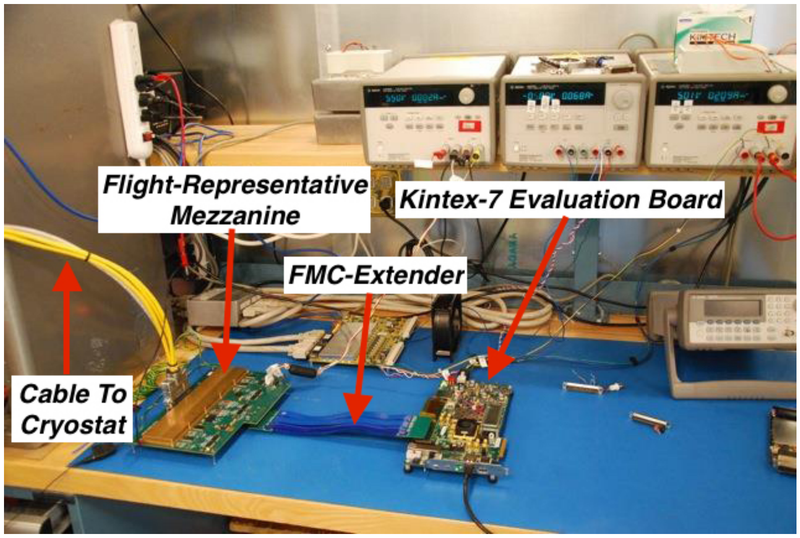

Figure 6–7: An image of the Kintex-7 FPGA motherboard connected to our mezzanine via an FMC extension.

to digital processes, but the mezzanine handles the digitization and synthesis of all analog signals, and mounts directly to this motherboard. One manner in which we were able to improve the EMI at the ADC was to move the mezzanine a few inches away from the motherboard using an FMC extension cable (Figure 6–7).

Some our suspicions regarding the K7 Evaluation board could be confirmed directly by the use of an antenna-coupled spectrum analyzer to probe the local RFI environment. We were able to localize large signals correlating with spikes in the ADC spectra at a number of components on the K7 Board, and trace those signals along the length of the FMC-extension and onto the mezzanine. However, the strongest indication that it was the source of much of our narrow-band contamination comes from a comparison to the custom 64x ground-based motherboard. The same baseline spectrum, using that hardware (Figure 6–8), suggests that much of the noise forest present in the spectra during the flight representative hardware commissioning was due to the K7 motherboard. Although the ground-based system uses non-flight-representative versions of the SQUID controller board and mezzanine, it is unlikely that either of these are significant contributors. The SQUID controller board and mezzanine do not have components capable of generating narrow-band interference, aside

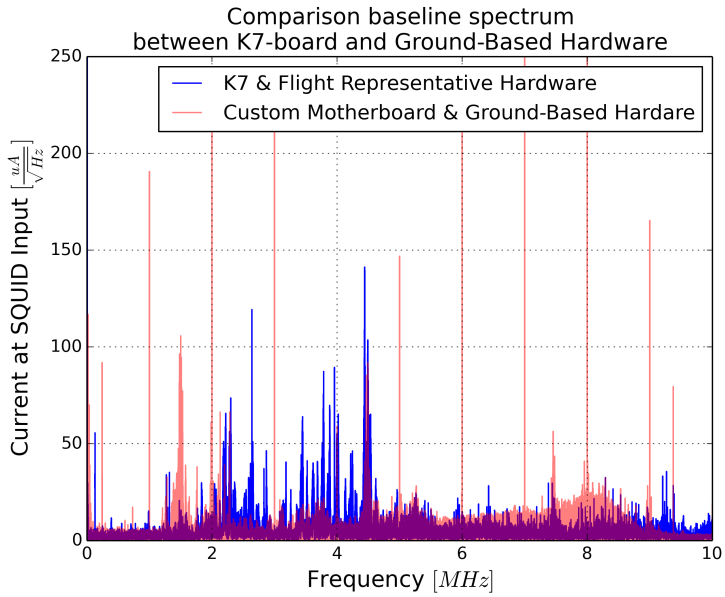

Figure 6–8: A comparison of the ADC output when looking at tuned SQUIDs when using the K7 motherboard and Flight Representative Hardware (blue) and custom motherboard with ground-based readout electronics (red).

from a single set of buck-regulators on the mezzanine which are carefully synchronized.

6.2.2 Direct Narrow-Band EMI Contamination

A readout channel is an exceptionally narrow-band thing compared to the bandwidth as a whole, with a relevant bandwidth of only a few hundred Hz. This is part of the reason that the lines of the buck regulators, when properly synchronized, are “ok” for the performance of the system: we simply avoid placing DAN channels nearby.

Even in the face of the forest of contamination shown in Figure 6–5, where most readout channels are contaminated by at least some of these discrete lines, the degree to which they effect noise directly is minimal. For narrow-band signals, a single line in the bolometer bandwidth doesn’t consist of an appreciable amount of power. A larger issue is that sufficiently strong lines can drive a bolometer into instability, in part by providing what is effectively an additional bias. We have not sought to try and quantify narrow-band direct EMI noise contributions on either a case-by-case basis, or as a general contribution; but have dropped the channels with the highest noise due to EMI from our measurement, instead replacing the readout channels with off-resonance channels to maintain a 64x channel count.

6.2.3 SQUID Loading

Direct contamination isn't the only way that EMI engenders poor instrument performance. Consider the SQUID response curve, $V(\Phi)$, introduced in Figure 1–10. Changing the amplitude of a DC current at the input coil changes the flux through the SQUID, moving it along that curve (we take advantage of this to adjust the flux bias). An oscillating current through the input coil likewise oscillates the SQUID over a local region of that curve, near the point selected with the flux bias. This local region increases with the amplitude of the input current signal. Because we use the SQUIDs as amplifiers, we value a large transimpedance – or equivalently we wish them to output a large change in voltage for a small change in input current. To this end we bias them at an operating point where the slope of the $V(\Phi)$ curve is large. Notice that because the transimpedance is related to the peak-to-peak amplitude of the $V(\Phi)$ curve (as explained in Section 1.4), a SQUID with high transimpedance will have a larger region of local linearity than SQUID with low transimpedance.

A high transimpedance is still not sufficient for SQUIDs to be good amplifiers, it's also important that they remain *linear*. That requirement is intrinsically difficult, because the response function is anything but. One of the two principal purposes of DAN is to suppress the SQUID non-linearity by ensuring that input signal amplitudes to the SQUID remain low, such that the local region of the curve being moved through is small enough to be approximately linear, and the transimpedance remains constant. For any typical DAN channel, the amplitude of either the carrier, or the nuller, would be more than sufficient to saturate the SQUID, let alone drive it non-linear. The DAN loop gain is high enough that a full comb of 64 such pairs can be operated in this way without doing so.

EMI at any frequency that does not fall within the very narrow DAN bandwidths is not suppressed – and the SQUID bears the full brunt it. As the total RMS of the input current waveform as seen by the SQUID coil increases, it enlarges the local region of the $V(\Phi)$ curve being exercised. As the curve in that region begins to look non-linear, the transimpedance acquires an amplitude-dependent variability, and the SQUID exhibits amplitude-dependent gain. We call this "SQUID loading". The way in which we measure SQUID transimpedance involves an integration over several samples of a demodulated (downsampled) signal, and is therefore really a measure of the weighted-average of the transimpedance in that interval of time.

6.2.4 EMI Begets EMI: Inter-Modulation Distortion Products

The presence of a non-linear gain element in the signal path gives rise to Inter-Modulation Distortion (IMD): tones attendant on the basis EMI frequencies, formed at harmonics, and at the sum and differences of all combinations of the principal frequencies. The amplitudes of these frequencies depend on the exact function of the non-linearity, and in this case increase as the SQUID is driven farther into non-linearity. Due to the relationship between the size of the linear local region and the peak-to-peak amplitude of the response curve, this effect worsens as the transimpedance falls: the same excursions in Φ can cause worse IMD at lower transimpedance than at high transimpedance. Thus, the ability of the SQUID to tolerate larger RMS amplitude inputs is closely related to the *initial* loading through it due to EMI.

Once this effect is strong enough for IMD products to appear, they do so as a forest. We are intentionally careful in choosing bias frequencies with a common multiple to exercise some control over the IMD products (ensuring they remain outside of our signal band), but a single large narrow-band contaminate can contribute several lines from its own harmonics alone. Through this process, even a handful of large EMI contaminates quickly generates a dense quasi-broadband distribution of signals, all of which are loading the SQUID. The additional mixing tones that arise from a single readout channel, let alone 64, can generate substantial IMD products; and as the amplitude of any tone increases, so will the total amplitude of the resulting IMD products, and therefore the total loading on the SQUID.

Bolometers are another non-linear element of our system – one whose linearity we can degrade by lowering them into their superconducting transitions. Figure 6–9 shows a spectrum of the SQUID output as seen at the ADC, first with overbiased bolometers and then after putting them into the transition. The spectrum is clearly degraded by the presence of the in-transition bolometers compared to when those same bolometers are kept overbiased. This is despite the fact that the actual amplitudes of the carrier and nullers have decreased.[3]

6.2.5 SQUID Loading and Digitization-Path-Noise

By lowering the transimpedance of the SQUID, narrow-band EMI – relatively inconsequential as a direct contaminate – can have significant ramifications for overall system noise. In Section 5.5 we calculated the ADC noise as a current at the input of the SQUID coil. This noise has an inverse proportionality to SQUID transimpedance (See Table 5–4). From the perspective of the SQUID as an amplifier, this makes sense: decreasing the gain of the SQUID preserves the signal-to-noise at the output of the SQUID coil, but does not change the effective noise from the digitization chain. Therefore, a lower SQUID transimpedance corresponds to a lower overall signal-to-noise. A simulation of this effect using the off-resonance noise parameters derived in Chapter 5, can be found in Figure 6–10.

We can demonstrate this by intentionally degrading SQUID transimpedance while measuring the white noise of several off-resonance channels. The SQUIDs used at McGill have a maximum transimpedance when operated with the Flight Representative Hardware, and

[3] Although it is possible to lower bolometers far enough into their transition that the total nuller amplitude increases compared to that of the initial overbiased state, to first order they decrease or remain zero-sum.

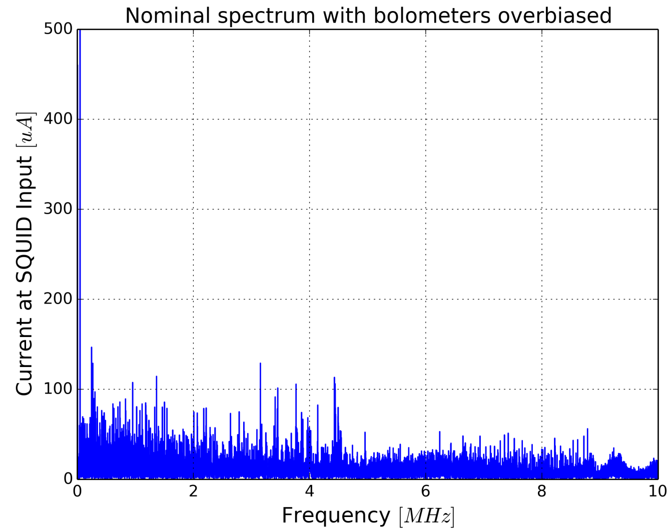

(a) Spectrum of the SQUID output as seen at the ADC when a comb of 40 bolometers are overbiased.

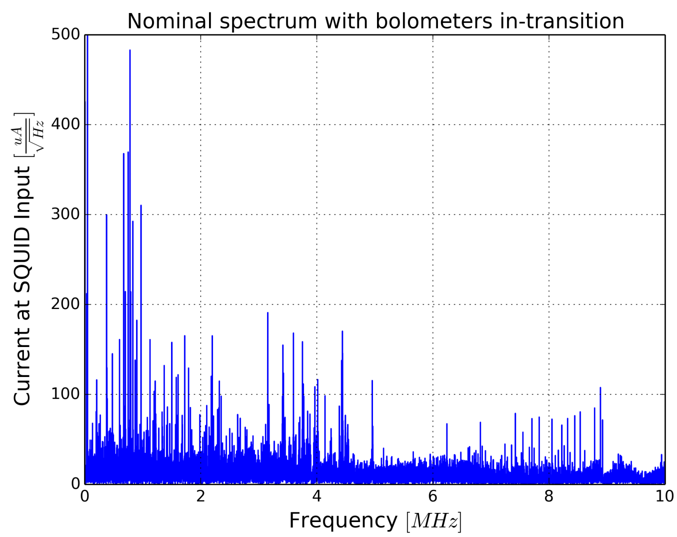

(b) Spectrum is the SQUID output as seen at the ADC when a comb of 40 bolometers is being operated in the transition. Note that this is from one of the final data-sets. We would expect this spectrum to be more contaminated as the bolometers are lowered into their transition the nuller amplitudes signals increase.

Figure 6–9: Spectra of the ADC input when a comb of bolometers is being operated both overbiased, and in the transition.

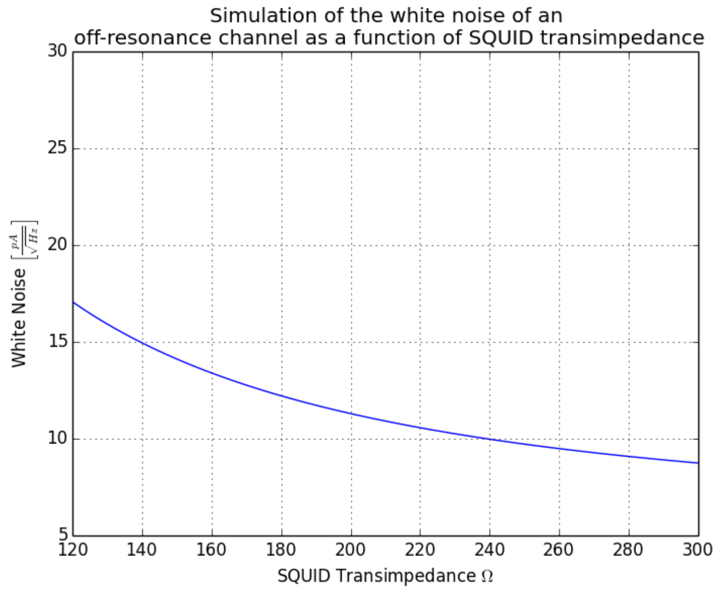

Figure 6–10: A simulation of the effect of varying the SQUID transimpedance on the overall noise measured off-resonance.

optimally tuned, of approximately 280Ω.[4] To modulate the transimpedance after tuning the SQUID, we sweep the flux bias away from its optimum value. At each step, the transimpedance is measured using the methods described in Section 4.2.3, along with the white-noise at several off-resonance frequencies. The result is shown in Figure 6–11.

A slight modification of this test also allows us to verify the assertion that a "DC" change transimpedance as above (by altering the flux bias point) is equivalent to an effective change in transimpedance due to the RMS of the input waveform, and also that IMD products do indeed load the SQUID. Instead of changing the SQUID transimpedance directly by modifying the SQUID bias, we do so by modifying loading from IMD products. An off-resonance DAN channel at 6MHz is enabled, and the carrier amplitude varied to modulate the resultant IMD. The amplitude of the 6MHz DAN channel is adjusted through the range

[4] Note that the optimal SQUID bias point is not necessarily the point of highest transimpedance, since linearity is also a factor. In fact, it is typically just off the transimpedance maximum.

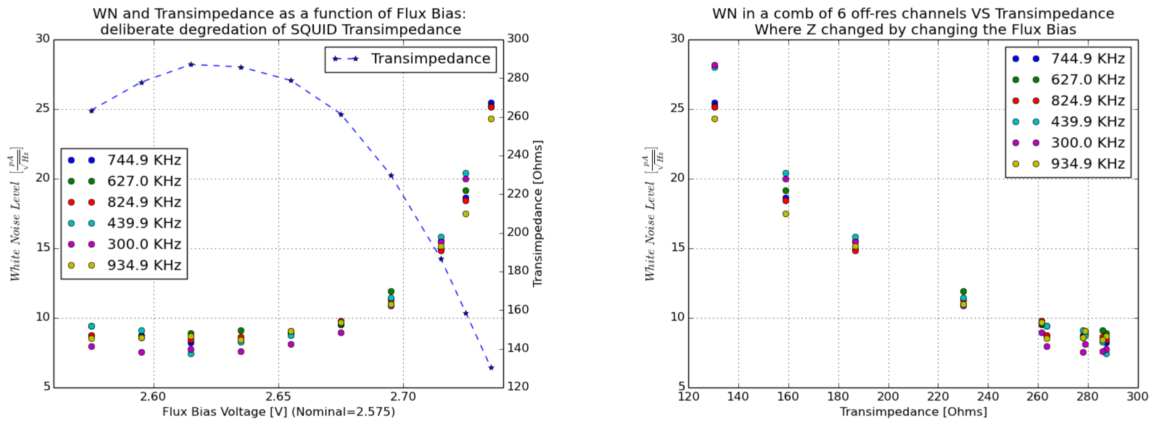

(a) Measured off-resonance white noise and SQUID transimpedance as a function of SQUID flux-bias.

(b) Measured white-noise of several low-frequency off-resonance channels as a function of SQUID transimpedance.

Figure 6–11: The relationship of the transimpedance of the SQUID on the digitization-chain white-noise is demonstrated by "de-tuning" the SQUID while measuring the transimpedance and noise.

from a typical bias carrier amplitude for a bolometer in-transition (0.025 Normalized units) to 0.13 Normalized units.

The results of this test can be seen in Figure 6–12. Increasing the total RMS amplitude of the carrier waveform does degrade the weighted-average SQUID transimpedance, and therefore noise performance – in a way which is consistent with "de-tuning" the SQUID (Figure 6–12b).

To first order this agrees with the analytic model presented in Figure 6–10. This is especially true in the higher-transimpedance regimes. However, ability of the model to predict the noise performance at lower transimpedance deviates significantly, such that at the lowest transimpedance measured the noise is underestimated at approximately the 30% level. This suggests that there is another mechanism by which the total noise is a function of SQUID transimpedance. One possible second-order effect is that the region of local linearity is itself sensitive to SQUID transimpedance. This has not been fully explored, largely due to the fact that our operating transimpedance is sufficiently far away from that regime that factoring

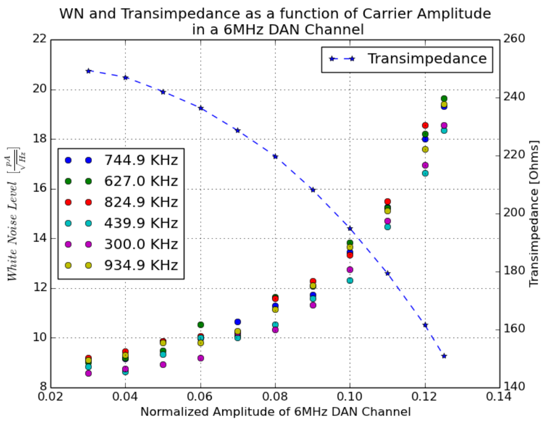

(a) The total measured off-resonance white and SQUID transimpedance as a function of SQUID flux-bias.

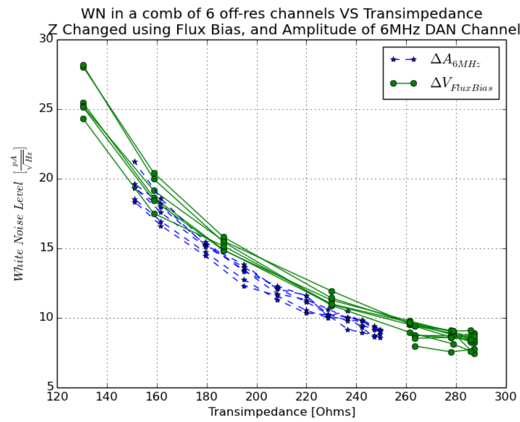

(b) The measured white-noise of several low-frequency off-resonance channels as a function of SQUID transimpedance.

Figure 6–12: These plots indicate first that increased loading on the SQUID degrades the transimpedance (and therefore the noise performance of the system as a whole) in a similar manner as “de-tuning” the SQUID; and also that our model of loading that is proportional to the total RMS amplitude of the synthesizer waveforms (through either intermodulation distortion products, or some form of cross-talk) is well motivated.

the second order effect into our noise predictions is not required.

6.2.6 Notes on SQUID Loading Across DfMUX Platforms

The SQUIDs used in this project have been previously characterized using the 16x Legacy DfMUX system, and since with the 64x ground-based warm electronics. In both cases, where the motherboards were specifically designed for low noise performance, we’ve achieved transimpedances of $\sim 380\Omega$, even with the CryoBoard enabled. For the flight-representative system, using the K7 Board, loading on the SQUID is responsible for a transimpedance degradation of nearly 200 Ω, down to $Z_{SQUID} = 190\Omega$. We made significant gains by disabling the CryoBoard, recovering 270 Ω, but the remaining 100 Ω of transimpedance degradation is thought to be entirely the result of contamination from the K7 Board.

Note that the non-linear, self-reinforcing, nature of IMD and SQUID loading offers an explanation for why the CryoBoard appears to be a more significant source of SQUID loading in this configuration than with the ground-based 64x system. The relative effect of the CryoBoard EMI on a SQUID with a 270Ω baseline transimpedance is unsurprisingly larger than for a SQUID with a baseline transimpedance of 380Ω (which would have a larger local region of linearity).

6.3 Warm-Electronics Crosstalk

There are two ways in which crosstalk between the analog signals can result in SQUID loading. The first is relatively conventional – any large-amplitude carrier or nuller signals which couple to the return lines in the cryostat can induce additional current at the SQUID input coil. From a noise perspective, this form of crosstalk doesn't matter much – it would occur at the frequencies of existing DAN channels, and thus be compensated for and removed by DAN instead of loading the SQUID.[5]

A far more problematic form of crosstalk would be from the carrier or nuller into the demodulation chain *in the warm* electronics. Such a signal would be seen as a large residual to a DAN channel, and thus be mistaken by DAN as a current through the SQUID, to be nulled. In order to zero the actual current through the SQUID, plus the current from crosstalk, DAN will shunt a corresponding additional current through the SQUID input coil. This is the same as effectively moving the virtual ground from the input of the SQUID coil, to just behind it.

[5] This is not necessarily true if we are operating more than one SQUID at a time. The type of crosstalk described above, taking place between the carrier or nuller on one module into the return-lines of another one, would not appear at frequencies of DAN channels on that other module, and therefore would load the SQUID directly.

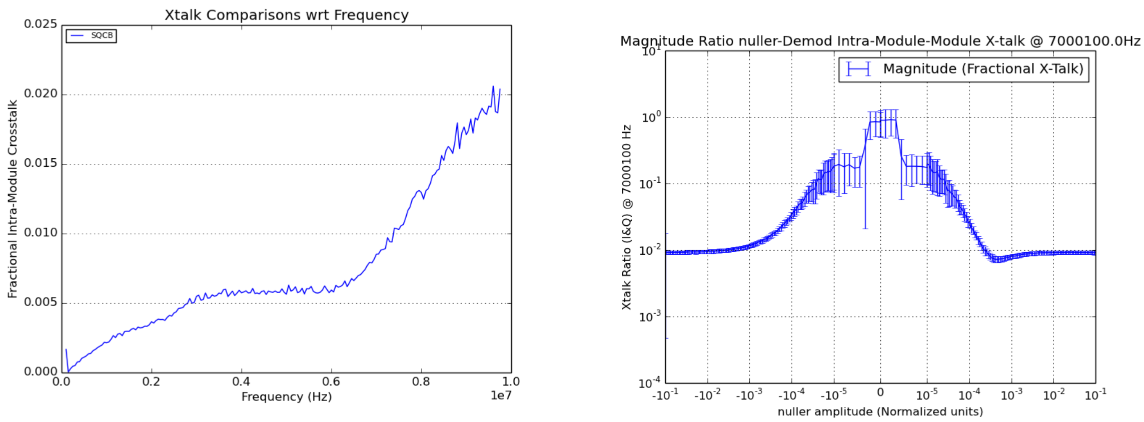

(a) Fractional nuller-to-demodulator path crosstalk in the warm-electronics as a function of frequency for a fixed amplitude. The amplitude chosen was representative of a typical synthesizer output during normal operation.

(b) Fractional nuller-to-demodulator path crosstalk in the warm-electronics as a function of amplitude for a single frequency. Note that typical amplitudes used for even single-channels are $> 10^{-2}$ Normalized units.

Figure 6–13: Nuller-to-demodulator path crosstalk in the warm-electronics.

We have measured the crosstalk in the warm electronics extensively, the results of which can be seen in Figure 6–13. Even at the highest frequencies, this is below 2%. At the frequencies that current cold-components can be operated, the nuller-to-demodulator crosstalk is less than 1%.

This level of crosstalk is consistent with existing Legacy 16x DfMUX systems. The magnitude and linear nature of the crosstalk as a function of amplitude (Figure 6–13b) is consistent with capacitive coupling between neighboring traces on the PCB and in the DB37 connectors. Note that at the lowest amplitudes on Figure 6–13b, the probe signal is dropping below the dynamic range of the ADC, and so the convergence to 1 is expected.

Although there are a number of additional pathologies that differentiate our testing setup from that of the "representative ideal", for which noise estimates were calculated in Chapter 5, corrections for them fall within the scope of the theoretical models established throughout this document. In particular, degraded SQUID performance due to increased loading can be quantified, and accounted for with the existing framework that defines the relationship between the SQUID transimpedance and noise contribution from demodulation-chain sources. We've also identified that the greatest source of uncertainty in how to update these predictions comes from the variation in comb impedance (both as scatter in the bolometer normal resistance, as well as the large range in capacitor ESR). Finally, our LC-width and channel spacing dictate that we are no longer in the large LC-spacing regime, and so we expect to see some variation in their white noise that has a frequency dependence. The underlying revelation was that we may not be able to consider each leg of the comb as independent, as we could under the large-LC-spacing approximation.

CHAPTER 7
Results and Analysis

We evaluate the noise of the system when operating a full comb of bolometers (plus off-resonance channels) at the three configurations for which our noise model makes predictions – off-resonance, for bolometers when overbiased, and for bolometers in their superconducting transitions. These measurements are compared to the noise-model presented in Chapter 5, but updated based on the non-idealities of our testing setup.

7.1 Corrected Noise Predictions

The previous sections have outlined a number of ways the expected noise of our test-setup should differ from the expectations laid out in Chapter 5. We've shown that the effective impedance as seen from an off-resonance channel varies due to overlap with LCR resonance frequencies, rather than being a fixed 10Ω. This variation can be expected to increase the scatter in the off-resonance noise. However, the effective impedance while off-resonance should never appreciably approach that of an on-resonance frequency, as even a factor of 2 (rather than the ideal factor of 10) attenuation already makes the carrier noise contribution sub-dominant. This fact is even more pronounced when we take into account the corrected bolometer resistances and ESR, which are significantly higher than those used in the calculations in Chapter 5.

Although the increased comb resistance should decrease the overall carrier noise contribution, we can expect that benefit will overshadowed by the resulting increase in bolometer responsivity (Equation 5.14), which multiplies the bolometer phonon noise (Equation 5.11), our largest single source of noise in the system.

Finally, the digitization-chain noise will go up in our predictions relative to those in Chapter 5, due to the substantially lower transimpedance of our SQUIDs relative to those of deployment-quality devices. This transimpedance difference, from 500Ω to 270Ω, is strongly a function of the non-nulled loading on the SQUID input coil.

The increased loading stems predominantly from the commercial Xilinx Kintex-7 FPGA Evaluation board we used as a motherboard to the system. This hardware was not designed for low-noise applications: it lacks synchronization, isolation, and strong filtering of the buck-converting regulators; physical layout choices that would minimize digital crosstalk into the analog signal paths of the mezzanine; and the spartan approach to resources undertaken to minimize errant signals on low-noise-optimized platforms. This explanation is suggested by comparing the spectrum at the ADCs when no synthesizers are present with the spectra with carrier and nuller synthesizer tones, which exacerbate IMD; and also by comparisons to the ground-based 64x readout system, which records transimpedances of $\sim 100\Omega$ better with the same SQUIDs.

Although we are unable to quantitatively evaluate the narrow-band EMI itself, we can measure the SQUID transimpedance with good precision, and showed in Section 6.2.5 how the relationship between SQUID loading and noise follow the behavior predicted by the analytic model; though we do note that the model for this relationship begins to fail at extremely low transimpedances $Z_{SQUID} <\sim 150$. Table 7–1 shows quantities and predictions that were used in Chapter 5 alongside updated values and to reflect our specific test-setup.

7.2 Measurement Results

Results of the white noise measurement are shown in Figure 7–1a, where the off-resonance prediction is given as a dotted line. Figures 7–1b and 7–2 present the three component measurements separately. Note that the off-resonance noise shown was taken when the in-transition noise was recorded, and the overbiased data corresponds to the same

Quantity	"Ideal" model	Corrected model
SQUID Transimpedance (Z_{SQUID})	500 Ω	270 Ω
Stray Series Impedance	0.25 Ω	[0.29-1.5] Ω
Overbiased Bolometer Resistance	1Ω	1.54 Ω
In Transition Bolometer Resistance	0.7 Ω	1.08 Ω
Bolometer Responsivity	$S_{z_s=0.25} = 10.7\times10^5 \left[\frac{A}{W}\right]$	$[0.9S_{ideal} - 2.14S_{ideal}]$
Noise	"Ideal" model $\left[\frac{pA}{\sqrt{Hz}}\right]$	Corrected model $\left[\frac{pA}{\sqrt{Hz}}\right]$
SQUID Noise	3.5	3.5
ACD-chain	3.84	7.11
Nuller-chain	4.79	4.79
Carrier-chain (Off-resonance)	0.525	0.525
Carrier-chain (Overbiased)	4.12	[2.86 – 1.73]
Carrier-chain (In-Transition)	5.52	[3.83 – 2.03]
Bolometer Noise (Overbiased)	6.65	5.36
Bolometer Noise (In-Transition)	27.35	[24.62 – 57.75]
Measurement Type	"Ideal" model	Corrected model
Off-Resonance	7.08	9.28
Overbiased	10.57	[11.08 – 10.84]
In-Transition	29.22	[26.58 – 58.52]

Table 7–1: A comparison of the assumed values and noise predictions from Chapter 5, with measured values and corrected predictions tailored for the McGill testing setup.

run, but just before cooling the stage below the superconducting transition temperatures of the bolometers, and lowering them into the transition.

The values are largely consistent with expectations, especially at frequencies below ~ 3.5 MHz, where the conditions of the cold-components are not creating an environment dominated by less well understood, high-frequency, effects. The theory line for off-resonance channels is, more accurately, a predicted lower-bound. Variance in the comb impedance means we should expect some frequency sensitivity. This can be observed directly from the plot by noting the strong correlation between points. Those correlations stem from the fact that off-resonance channels are often packed between neighboring LC peaks, such that they are either bounding a minimum in comb impedance or a maximum.

7.3 Better ESR modeling

The up-turn in the noise of the in-transition detectors is consistent with phonon noise, which is sensitive to stray series impedance due to the increased responsivity of the bolometer. Although so far we have modeled the ESR using two points (to provide a bound), from which we derived the upper and lower limits in Figure 7–2b, we can actually do a little bit better.

The ESR of a capacitor is related to frequency by the equation $ESR(f) = \frac{1}{2\pi fCQ}$ where Q is the "quality factor" of our capacitors, which is to be calculated at ~ 520 and C is their capacitance. The quality factor of a capacitor is related to the loss tangent, and is a product of the capacitor construction and material choice. Due to the manner in which we tune the LC resonances, the capacitance is also a function of frequency: we keep the inductance of our LC-chips fixed and adjust capacitance to vary the resonant frequency. We can therefore re-arrange this to give Equation 7.1.

$$ESR(f) = \frac{2\pi fL}{Q} \, , \tag{7.1}$$

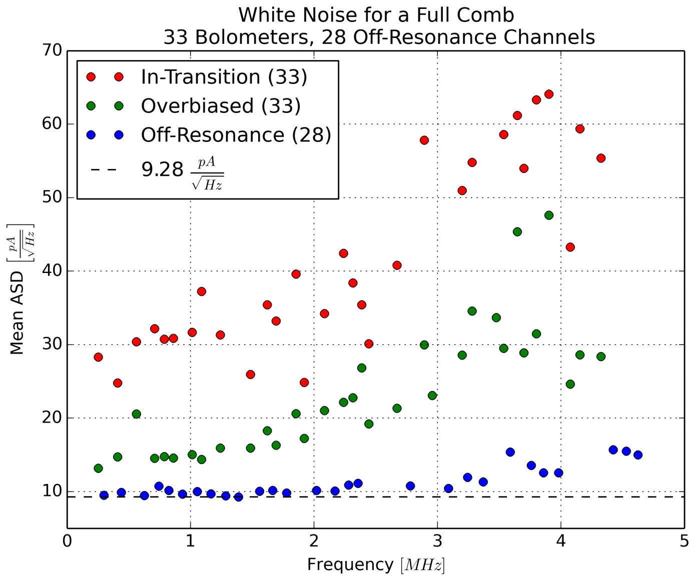

(a) Measured white-noise noise for the full comb. The in-transition and off-resonance points were measured simultaneously; and the overbiased points were measured before dropping them into the transition. Note: One exceptionally high noise bolometer channel is kept off-scale. See Figure 7–2b for complete set.

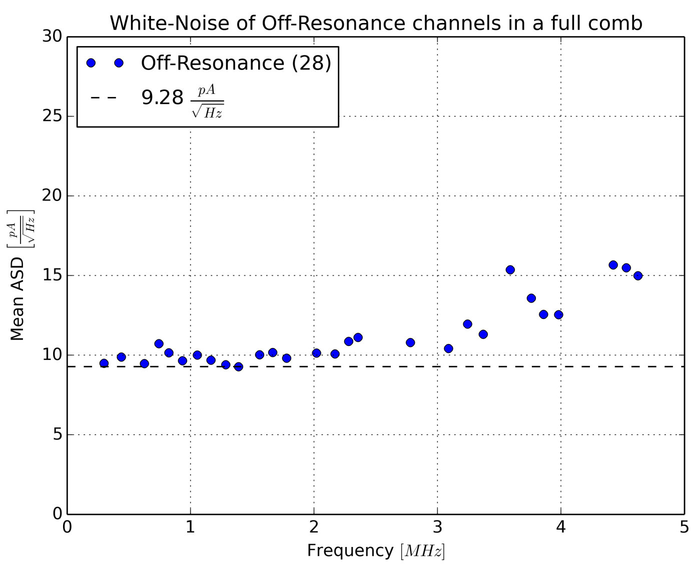

(b) A closer look at just the off-resonance channels from Figure (a). The dotted line is the prediction based on noise modeling. The divergence from the prediction at high frequency is unexpected, but may be the result of leakage current-noise from bolometers at nearby frequencies due to the wide LC-resonances.

Figure 7–1: Noise measurements from a full comb. Note that one bolometer channel is off-scale, one channel is being used as a transimpedance monitor, and two others were biasing bolometers at high frequencies, which became unstable upon entering their transitions and had to be disabled. Error bars for these points are omitted as the uncertainties in the measurements themselves are too small to be seen on this scale.

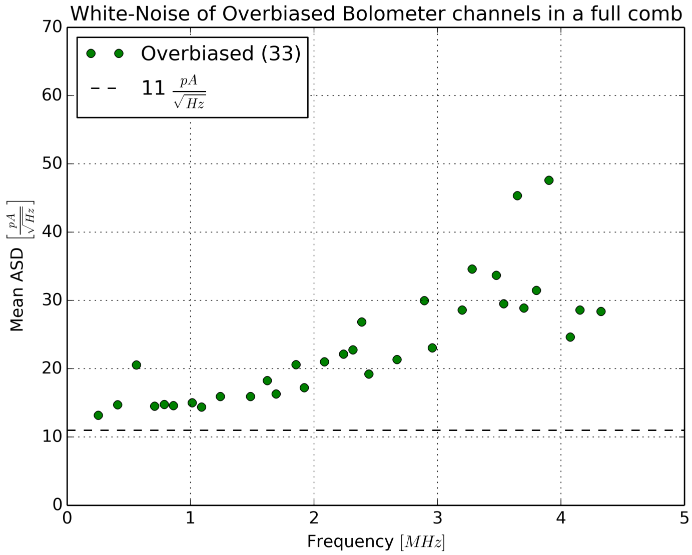

(a) Measured white noise of the bolometers before they were dropped into the transition. The dotted line is a prediction based on noise modeling. The excess noise, as well as the strong correlation with frequency, suggests that the bolometers are not completely normal, and thus are susceptible to phonon noise. This is explored in Section 7.4.

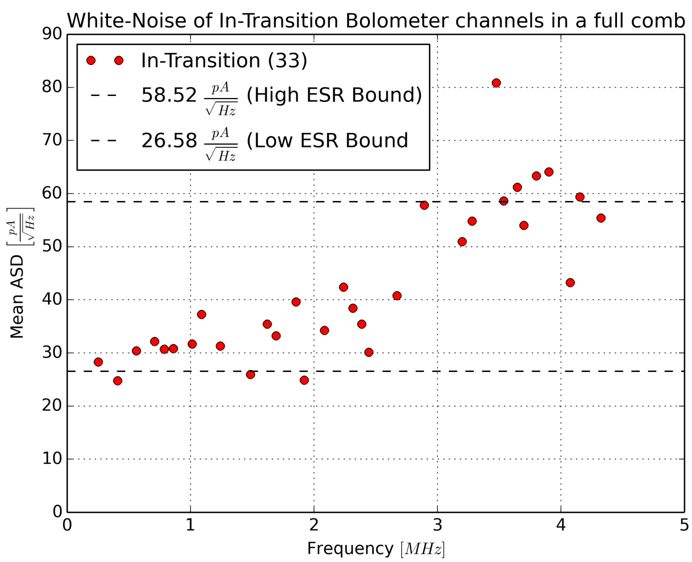

(b) Measured white noise from the bolometer channels when operating in-transition channels. The lines show upper and lower expectation values based on ESR.

Figure 7–2: Plots showing the bolometer noise when overbiased and in-transition in closer detail.Error bars for these points are omitted as the uncertainties in the measurements themselves are too small to be seen on this scale.

where L is 24 μH and Q is 520. In the case of fixed Q, we expect a linear relationship between the ESR and the frequency of the bolometer channel. In Figure 7–3 we show simulations of the expected noise as a function of first ESR and frequency using the above parameters. Notice that although the phonon noise is only one component in a quadrature sum, the line we see is very nearly linear. This is because even at relatively low ESR ($\sim 0.25\Omega$) the noise components proportional to ESR (primarily phonon) are already larger than all other components together.

Also worth noting is that because we fix a Q for this model, the ESR continues falling even in the low frequency regime. We know this is not entirely accurate – although the ESR for any single capacitor (and therefore a single Q) gets larger at higher frequencies, and falls with lower frequencies, the Q factor of the capacitor is actually approximately inversely proportional to its capacitance. Since the capacitance used in the LCR filters is, in turn, proportional to one over the *square* of the desired operating frequency, an increasing capacitance will eventually work to counter the falling ESR at lower frequencies. This only remains true provided the capacitors used across the bandwidth remain similar enough in construction, geometry, and materials.

Additionally, scatter in the normal bolometer resistance, and the fact that the ESR of each channel is not independently measured and modeled, creates uncertainty in the noise contributions from the carrier-chain noise and bolometer Johnson noise. More significantly, it means there is some scatter in the actual depth of transition the bolometers reach – it is impossible to know precisely how far into the transition they have been lowered, without separating the ESR and bolometer resistance. The bolometer phonon noise and Johnson noise components are functions of the bolometer loop-gain, which will vary across the comb according to transition depth, though it has been fixed in our predictions at 10.

The updated ESR model appears to describe the frequency correlation very well, and is shown with the in-transition data in Figure 7–4. Including support in the model for

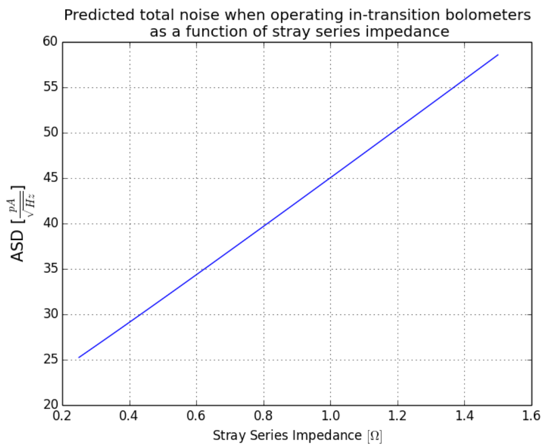

(a) A simulation of the effect of varying the stray impedance in series with the bolometer on the total noise when bolometers are in their transition.

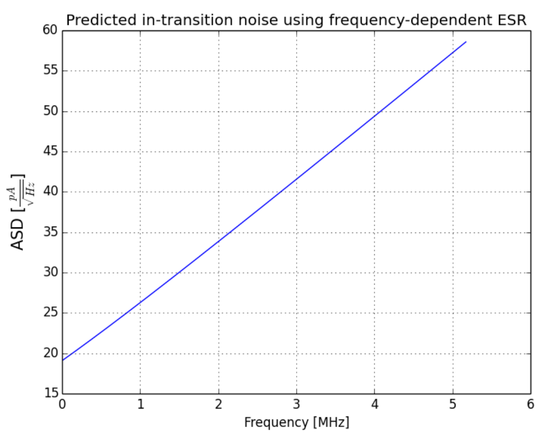

(b) A simulation of the effect of the noise as a function of frequency, when a bolometer is in-transition, using a Q-factor of 520.

Figure 7–3: Updated noise simulations for an in-transition bolometer channel as a function of ESR and frequency.

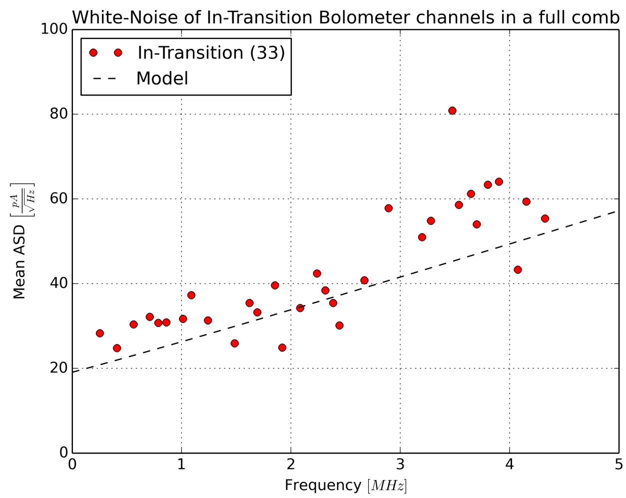

Figure 7–4: In-transition bolometer noise shown with the updated, frequency-dependent model.

frequency-dependent Q would undoubtedly result in a higher accuracy prediction, but uncertainties in the fundamental properties of the bolometers, their dynamic state during data-taking, and scatter in other quantities such as the individual characteristics of the capacitors, limit the returns on such an increase in model precision. Ultimately, bolometer noise characterization is not the object of this study; and this is sufficient to demonstrate that the noise in the readout system is well described using conventional models.

7.4 Evidence of Soft Bolometer Transitions

The most discrepant measurement is that of the overbiased bolometers. The large excess noise is striking, but even more so is the strong frequency dependence; which follows a similar shape as seen once they are lowered into the transition. The only noise source we know of that has such a frequency dependence is bolometer phonon noise – a property that becomes relevant in a bolometer as it goes superconducting. One of the fundamental assumptions in leaving out the phonon noise when calculating an overbiased bolometer noise expectation, is that it is completely normal. Yet some measurements of the transition, such as those shown in Figure 7–5 seem to demonstrate that this transition is soft, and that the saturated region

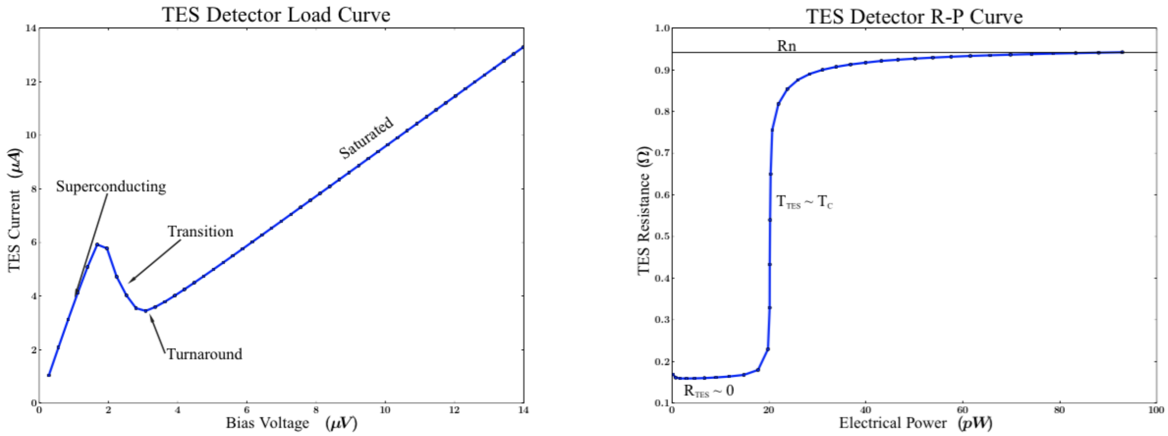

Figure 7–5: Two load-curves showing bolometer response in IV, and in RP. [25]

above the turning point of the transition (which can be seen in the IV curves in Figure 7–5) is still not sufficiently normal to disregard the bolometer phonon noise component.

This measurement suggests that the model of a "soft-transition" bolometer is correct, and that even at 800mK an overbiased bolometer will have some non-zero loop-gain.

7.5 Frequency Dependence in Off-Resonance Channels

The frequency dependence in the off-resonance channels is notable because, given our assumptions, there should be no frequency-dependent noise source contributing to the off-resonance noise. We explore two possible explanations, neither of which would directly contribute noise, but do modulate the *transfer function*, thereby modulating existing noise.

7.5.1 Stray Series Inductance

A small stray inductance in series with the bias resistor of the carrier would have a negligible impedance at low frequencies, but as frequencies increase would contribute to the effective impedance of the bias resistor. Instead of multiplying the carrier waveform by a constant $R_{bias} = 30m\Omega$, it would be better to multiply it by up by $R_{eff} = R_{bias} + j\omega L$.

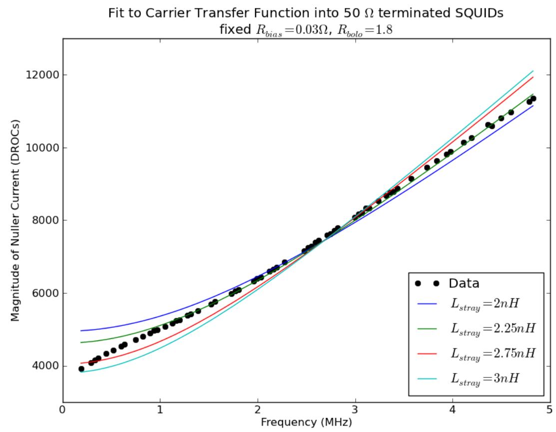

Figure 7–6: A fit for a stray inductance in series with the bias resistor – using SQUIDs terminated with 50Ω. Image Credit: Amy Bender.

Assuming this stray exists, we can measure it using specially configured SQUIDs – which, rather than connecting to a comb of LCR-filters, are terminated with 50Ω. A DAN channel is used to suppress the other frequency-dependent stray in the system – the SQUID input coil itself. By sweeping a carrier of fixed amplitude through the bandwidth, and recording the nulling amplitude required to zero it, we identify any frequency-dependent elements in the carrier signal path (that are not also shared by the nuller signal path). We then fit the result for a stray inductance at the bias resistor. In the absence of our stray inductance we expect this to be flat. The result, shown in Figure 7–6, suggests a stray inductance of approximately 2.25 nH. This value is consistent with what might be expected from soldering on surface-mount passives. It only becomes relevant in series with a very low resistance element such as the bias resistor.

At 4 MHz $Z_{straybias} = 63m\Omega$, a stray of this magnitude will be applying an additional gain of approximately 3 to the carrier signal path. This seems considerable, except that off-resonance channels see the carrier noise strongly attenuated. A factor of 3 increase in the carrier contribution represents only a 1.3% increase in the expected off-resonance amplitude.

Even in the channels for which the carrier noise is more substantial, its overall contribution to the total noise is still sufficiently small as to trivialize this change in bias (from a noise standpoint). In the presence of a 2.25nH stray series inductance, the noise of an in-transition bolometer at 4MHz will have increased by less than a percent. The change for an overbiased bolometer is slightly more dramatic – owing to the reduction in competing phonon noise, but the largest increase in white noise we could expect is still only on the order of 9%, not enough to explain the discrepancies.

7.5.2 Leakage current

Another possible explanation relies on the fact that we know the large-LC-spacing approximation is incorrect. Recall that the noise sources so far discussed are *broadband* sources of noise, and so do populate the bandwidth of our system. However, they do so from within an LC-resonance, and are attenuated as seen by an off-resonance channel. This attenuation is finite, and we expect it to be worse than our initial presumptions. If it is the case that broadband noise from bolometers in LC-resonances near to off-resonance channels is strong enough to be non-trivially seen at the off-resonance frequencies, then the frequency dependence shouldn't be surprising to us, since we have just shown the underlying bolometer noise to have a strong frequency dependence.

This is an avenue worth pursuing, and detailed simulations of the comb may shed some light on whether this is a feasible mechanism to explain the high-frequency behavior of our off-resonance channels. It has not been undertaken, in part because we expect much of the higher-frequency behavior to improve with the use of modern cold-components, which are being commissioned currently.

These measurements, and the phenomenological crosschecks, such as those in Section 6.2.5, demonstrate a sufficient understanding of the noise, and that the end-to-end performance of the system as a whole meets specifications and agrees with our theoretical models.

The most discrepant elements in the measurement derive from difficulty in quantitatively predicting high-frequency effects, due largely to uncertainties in the bolometer normal resistances and of the properties of the commercial ceramic capacitors that contribute large stray series resistances. We've been able to show that these frequency-sensitive effects can be explained using elements already contained within the noise model, and most importantly, are exacerbated far above and beyond what would be present in a science-grade instrument by the compromises necessitated in using Legacy 1MHz cold hardware used for these tests. The largest cause of frequency-correlated noise – ESR – has already been addressed in the latest generation of "inter-digitated" capacitors, which do not suffer from the fundamental limitations of ceramic capacitors.

CHAPTER 8
Conclusion

The polarization anisotropies of the Cosmic Microwave Background contain both newly uncovered, and as yet undiscovered, science. These signatures within the CMB are fainter than the temperature anisotropies by between one, and several, orders of magnitude. In order to reach the required sensitivities, we must leverage the advancements already made in the fields of detector technology, with an ability to massively multiplex large arrays of these devices.

The most advantaged means of measuring the polarization of the CMB on large scales, where the most exciting physics is thought to be found, is with a space-based observatory. Such an instrument would not only have the full sky accessible, but also the ability to exploit a frequency spectrum far more diverse than is accessible from within Earth's atmosphere. Multi-band observations are crucial tools for the removal of galactic foregrounds, which are ever more significant at these larger angular scales and fainter signal amplitudes. With focal planes of below 50 sensors, no space-based CMB observatory to date has employed multiplexing. This will not be possible for the next generation of instruments; current ground-based projects are now developing focal-planes with 10,000 TES bolometers.

In this document, we've presented a space-flight representative digital frequency multiplexing readout system, which satisfies the Canadian Space Agency criteria for space-based instruments. The CSA criteria include constraints on radiation hardness, thermal stability, and power consumption. This instrument represents the current state-of-the-art, with a demonstrated multiplexing factor of 64 – a factor of 4 beyond any currently deployed ground-based frequency multiplexing readout system.

Today, readout noise in such systems is sub-dominant to both intrinsic bolometer noise, and the fundamental photon-noise of the CMB. With the 64x flight-representative readout electronics, we've increased the multiplexing factor without a degradation in noise performance, and have demonstrated that the measured noise qualities are well understood – conforming to existing models of noise sources from the warm and cold electrical components. Discrepant features in the noise measurements are the result of pathologies within the cold-components, not the readout electronics, and current development at other institutions has addressed the most alarming of these.

A challenge with any space-based instrument is obsolescence between development and flight; the electronics presented here are robust to this in several ways. Existing cold hardware can only make use of ∼5MHz of the total supported 10MHz of bandwidth. As cold-hardware development progresses, more of that bandwidth will open up to support increasing numbers of bolometer channels. When that happens, there is a clear path forward for expanding the supported multiplexing factor from 64 to 128, which does not face any unknowns, and could be run on the existing hardware with only firmware modification.

Appendix A
Thermal Testing of the Flight Representative Electronics

Thermal tests were conducted over a range of temperatures between -20 and 40C, and the electronics demonstrated suitable robustness to the temperature fluctuations. The top plots in the matrix of images in Figure 8–1 show that the transfer function of the warm electronics, and the noise in a single readout channel, are acceptably insensitive to thermal variation (2% and 10% respectively). Note that a 10% change in the measured readout noise level does not effect instrument performance. The bottom plots demonstrate the SQUID bias voltages change by less than 0.5% over the full thermal range.

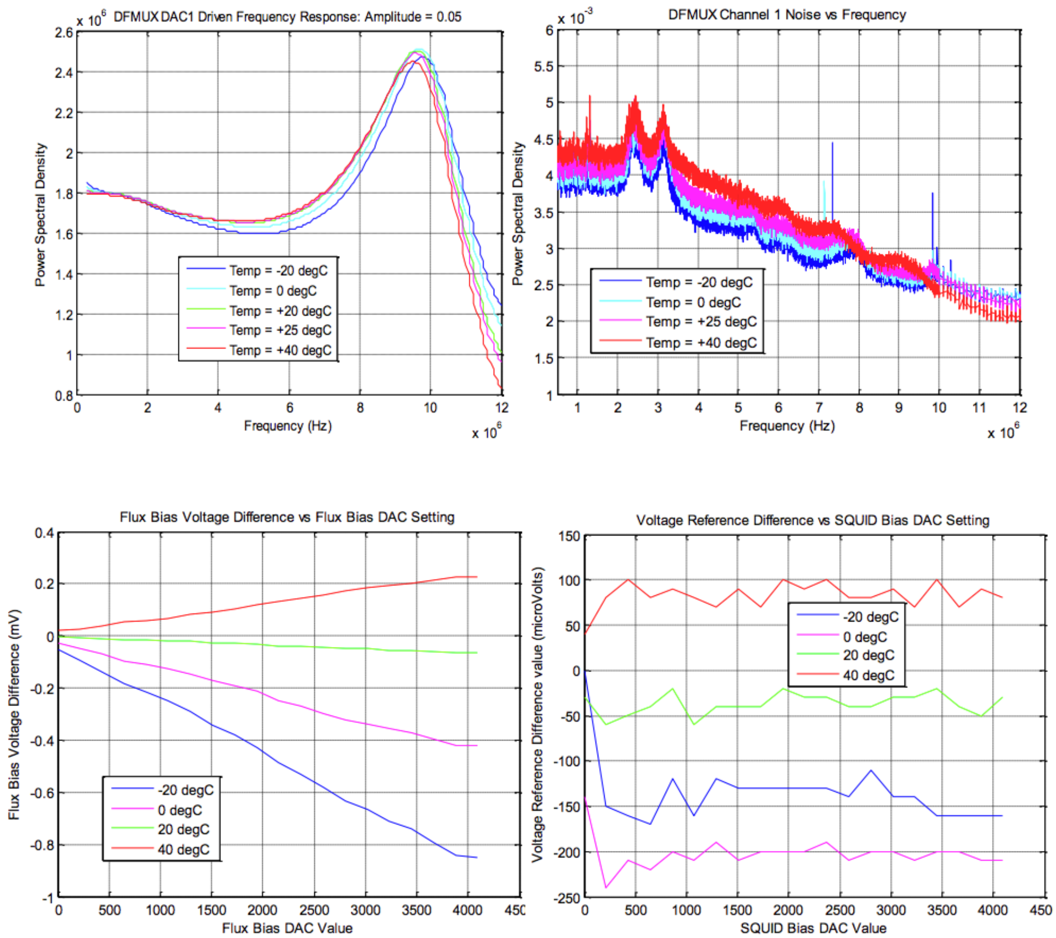

Figure 8–1: A set of plots summarizing the results of the thermal testing. Upper left: the frequency-dependent transfer function is stable as a function of temperature – and the degree to which it varies is consistent across the bandwidth. Upper right: with the instrument configured to use a loop-back dongle, as detailed in 4.4, white noise in a readout channel is measured at frequencies across the bandwidth. Because of peculiarities in the COM-DEV testing setup, the spectral shape of this isn't too significant, the important detail is that it is consistent across temperatures. The lower two plots show the variation in the voltage output of the SQUID controller board DACs as a function of temperature. Note that the y-scale is in mV, which is a trivially small variation.

References

[1] P. A. R. Ade, R. W. Aikin, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Brevik, I. Buder, E. Bullock, C. D. Dowell, L. Duband, J. P. Filippini, S. Fliescher, S. R. Golwala, M. Halpern, M. Hasselfield, S. R. Hildebrandt, G. C. Hilton, V. V. Hristov, K. D. Irwin, K. S. Karkare, J. P. Kaufman, B. G. Keating, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, E. M. Leitch, M. Lueker, P. Mason, C. B. Netterfield, H. T. Nguyen, R. O'Brient, R. W. Ogburn, A. Orlando, C. Pryke, C. D. Reintsema, S. Richter, R. Schwarz, C. D. Sheehy, Z. K. Staniszewski, R. V. Sudiwala, G. P. Teply, J. E. Tolan, A. D. Turner, A. G. Vieregg, C. L. Wong, K. W. Yoon, and Bicep2 Collaboration. Detection of B-Mode Polarization at Degree Angular Scales by BICEP2. *Physical Review Letters*, 112(24):241101, June 2014. `arXiv:1403.3985`, `doi:10.1103/PhysRevLett.112.241101`.

[2] The Polarbear Collaboration: P. A. R. Ade, Y. Akiba, A. E. Anthony, K. Arnold, M. Atlas, D. Barron, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, M. Dobbs, T. Elleflot, J. Errard, G. Fabbian, C. Feng, D. Flanigan, A. Gilbert, W. Grainger, N. W. Halverson, M. Hasegawa, K. Hattori, M. Hazumi, W. L. Holzapfel, Y. Hori, J. Howard, P. Hyland, Y. Inoue, G. C. Jaehnig, A. H. Jaffe, B. Keating, Z. Kermish, R. Keskitalo, T. Kisner, M. Le Jeune, A. T. Lee, E. M. Leitch, E. Linder, M. Lungu, F. Matsuda, T. Matsumura, X. Meng, N. J. Miller, H. Morii, S. Moyerman, M. J. Myers, M. Navaroli, H. Nishino, A. Orlando, H. Paar, J. Peloton, D. Poletti, E. Quealy, G. Rebeiz, C. L. Reichardt, P. L. Richards, C. Ross, I. Schanning, D. E. Schenck, B. D. Sherwin, A. Shimizu, C. Shimmin, M. Shimon, P. Siritanasak, G. Smecher, H. Spieler, N. Stebor, B. Steinbach, R. Stompor, A. Suzuki, S. Takakura, T. Tomaru, B. Wilson, A. Yadav, and O. Zahn. A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR. *The Astrophysical Journal*, 794(2):171, 2014.

[3] Francois Aubin. *Detector readout electronics for EBEX: a balloon-borne cosmic microwave background polarimeter*. PhD thesis, McGill University, 2012.

[4] J. E. Austermann, K. A. Aird, J. A. Beall, D. Becker, A. Bender, B. A. Benson, L. E. Bleem, J. Britton, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, T. M. Crawford, A. T. Crites, A. Datesman, T. de Haan, M. A. Dobbs, E. M. George, N. W. Halverson, N. Harrington, J. W. Henning, G. C. Hilton, G. P. Holder, W. L. Holzapfel, S. Hoover, N. Huang, J. Hubmayr, K. D. Irwin, R. Keisler, J. Kennedy, L. Knox, A. T. Lee, E. Leitch, D. Li, M. Lueker, D. P. Marrone, J. J. McMahon, J. Mehl, S. S. Meyer, T. E. Montroy, T. Natoli, J. P. Nibarger, M. D. Niemack, V. Novosad, S. Padin, C. Pryke, C. L. Reichardt, J. E. Ruhl, B. R. Saliwanchik, J. T. Sayre, K. K.

Schaffer, E. Shirokoff, A. A. Stark, K. Story, K. Vanderlinde, J. D. Vieira, G. Wang, R. Williamson, V. Yefremenko, K. W. Yoon, and O. Zahn. SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope, 2012. URL: `http://dx.doi.org/10.1117/12.927286`, `doi:10.1117/12.927286`.

[5] B A Benson, P A R Ade, Z Ahmed, S W Allen, K Arnold, J E Austermann, A N Bender, L E Bleem, J E Carlstrom, C L Chang, H M Cho, J F Cliche, T M Crawford, A Cukierman, T de Haan, M A Dobbs, D Dutcher, W Everett, A Gilbert, N W Halverson, D Hanson, N L Harrington, K Hattori, J W Henning, G C Hilton, G P Holder, W L Holzapfel, K D Irwin, R Keisler, L Knox, D Kubik, C L Kuo, A T Lee, E M Leitch, D Li, M McDonald, S S Meyer, J Montgomery, M Myers, T Natoli, H Nguyen, V Novosad, S Padin, Z Pan, J Pearson, C Reichardt, J E Ruhl, B R Saliwanchik, G Simard, G Smecher, J T Sayre, E Shirokoff, A A Stark, K Story, A Suzuki, K L Thompson, C Tucker, K Vanderlinde, J D Vieira, A Vikhlinin, G Wang, V Yefremenko, and K W Yoon. SPT-3G: a next-generation cosmic microwave background polarization experiment on the South Pole telescope. In *Proceedings of the SPIE*, page 91531P. Fermi National Accelerator Lab. (United States), July 2014.

[6] Tijmen de Haan, Graeme Smecher, and Matt Dobbs. Improved performance of TES bolometers using digital feedback, 2012. URL: `http://dx.doi.org/10.1117/12.925658`, `doi:10.1117/12.925658`.

[7] M A Dobbs, M Lueker, K A Aird, A N Bender, B A Benson, L E Bleem, J E Carlstrom, C L Chang, H M Cho, J Clarke, T M Crawford, A T Crites, D I Flanigan, T de Haan, E M George, N W Halverson, W L Holzapfel, J D Hrubes, B R Johnson, J Joseph, R Keisler, J Kennedy, Z Kermish, T M Lanting, A T Lee, E M Leitch, D Luong-Van, J J McMahon, J Mehl, S S Meyer, T E Montroy, S Padin, T Plagge, C Pryke, P L Richards, J E Ruhl, K K Schaffer, D Schwan, E Shirokoff, H G Spieler, Z Staniszewski, A A Stark, K Vanderlinde, J D Vieira, C Vu, B Westbrook, and R Williamson. Frequency multiplexed superconducting quantum interference device readout of large bolometer arrays for cosmic microwave background measurements. *Review of Scientific Instruments*, 83(7):3113, July 2012.

[8] M.A. Lindeman *et al.* Carrier phase optimization for frequency division multiplexing of low temperature detectors. *Journal of Low Temperature Physics*, 167:701.

[9] Alan H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, Jan 1981. URL: `http://link.aps.org/doi/10.1103/PhysRevD.23.347`, `doi:10.1103/PhysRevD.23.347`.

[10] D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. J. Bock, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H-M. Cho, A. Conley, T. M. Crawford, T. de Haan, M. A. Dobbs, W. Everett, J. Gallicchio, J. Gao, E. M. George, N. W. Halverson, N. Harrington, J. W. Henning, G. C. Hilton, G. P. Holder, W. L. Holzapfel, J. D. Hrubes, N. Huang, J. Hubmayr,

K. D. Irwin, R. Keisler, L. Knox, A. T. Lee, E. Leitch, D. Li, C. Liang, D. Luong-Van, G. Marsden, J. J. McMahon, J. Mehl, S. S. Meyer, L. Mocanu, T. E. Montroy, T. Natoli, J. P. Nibarger, V. Novosad, S. Padin, C. Pryke, C. L. Reichardt, J. E. Ruhl, B. R. Saliwanchik, J. T. Sayre, K. K. Schaffer, B. Schulz, G. Smecher, A. A. Stark, K. T. Story, C. Tucker, K. Vanderlinde, J. D. Vieira, M. P. Viero, G. Wang, V. Yefremenko, O. Zahn, and M. Zemcov. Detection of B-Mode Polarization in the Cosmic Microwave Background with Data from the South Pole Telescope. *Phys. Rev. Lett.*, 111:141301, Sep 2013. URL: `http://link.aps.org/doi/10.1103/PhysRevLett.111.141301`, `doi:10.1103/PhysRevLett.111.141301`.

[11] F.J. Harris, C. Dick, and M. Rice. Digital receivers and transmitters using polyphase filter banks for wireless communications. *Microwave Theory and Techniques, IEEE Transactions on*, 51(4):1395–1412, Apr 2003. `doi:10.1109/TMTT.2003.809176`.

[12] D. Huterer, D. Kirkby, R. Bean, A. Connolly, K. Dawson, S. Dodelson, A. Evrard, B. Jain, M. Jarvis, E. Linder, R. Mandelbaum, M. May, A. Raccanelli, B. Reid, E. Rozo, F. Schmidt, N. Sehgal, A. Slosar, A. van Engelen, H.-Y. Wu, and G. Zhao. Growth of Cosmic Structure: Probing Dark Energy Beyond Expansion. *ArXiv e-prints*, September 2013. `arXiv:1309.5385`.

[13] K. D. Irwin, L. R. Vale, N. E. Bergren, S. Deiker, E. N. Grossman, G. C. Hilton, S. W. Nam, C. D. Reintsema, D. A. Rudman, and M. E. Huber. Time-division SQUID multiplexers. *AIP Conference Proceedings*, 605(1), 2002.

[14] A. Braginski J. Clarke. *The SQUID Handbook: Fundamentals and Technology of SQUIDs and SQUID Systems*, volume 1. Wiley-VCH Verlag GmbH & Co. KGa, 2004.

[15] B. D. Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1:251–253, July 1962. `doi:10.1016/0031-9163(62)91369-0`.

[16] G.C Hilton K.D Irwin. *Cryogenic particle detection*, volume 99. Springer, 2005.

[17] Caroline A. Kilbourne, Simon R. Bandler, Ari D. Brown, James A. Chervenak, Enectali Figueroa-Feliciano, Fred M. Finkbeiner, Naoko Iyomoto, Richard L. Kelley, F. Scott Porter, and Stephen J. Smith. Uniform high spectral resolution demonstrated in arrays of TES x-ray microcalorimeters, 2007. URL: `http://dx.doi.org/10.1117/12.734830`, `doi:10.1117/12.734830`.

[18] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel. Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420(6917):772–787, 12 2002.

[19] Martin Lueker. *Measurements of Secondary Cosmic Microwave Background Anisotropies with the South Pole Telescope*. PhD thesis, University of California, Berkeley, 2010.

[20] R. W. Ogburn, IV, P. A. R. Ade, R. W. Aikin, M. Amiri, S. J. Benton, J. J. Bock, J. A. Bonetti, J. A. Brevik, B. Burger, C. D. Dowell, L. Duband, J. P. Filippini, S. R. Golwala, M. Halpern, M. Hasselfield, G. Hilton, V. V. Hristov, K. Irwin, J. P. Kaufman, B. G. Keating, J. M. Kovac, C. L. Kuo, A. E. Lange, E. M. Leitch, C. B. Netterfield, H. T. Nguyen, A. Orlando, C. L. Pryke, C. Reintsema, S. Richter, J. E. Ruhl, M. C. Runyan, C. D. Sheehy, Z. K. Staniszewski, S. A. Stokes, R. V. Sudiwala, G. P. Teply, J. E. Tolan, A. D. Turner, P. Wilson, and C. L. Wong. The BICEP2 CMB polarization experiment. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1, July 2010. `doi:10.1117/12.857864`.

[21] A. Orlando, R. W. Aikin, M. Amiri, J. J. Bock, J. A. Bonetti, J. A. Brevik, B. Burger, G. Chattopadthyay, P. K. Day, J. P. Filippini, S. R. Golwala, M. Halpern, M. Hasselfield, G. C. Hilton, K. D. Irwin, M. Kenyon, J. M. Kovac, C. L. Kuo, A. E. Lange, H. G. LeDuc, N. Llombart, H. T. Nguyen, R. W. Ogburn, C. D. Reintsema, M. C. Runyan, Z. Staniszewski, R. Sudiwala, G. Teply, A. R. Trangsrud, A. D. Turner, and P. Wilson. Antenna-coupled TES bolometer arrays for BICEP2/Keck and SPIDER, 2010. URL: `http://dx.doi.org/10.1117/12.857914`, `doi:10.1117/12.857914`.

[22] A A Penzias and R W Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419–421, July 1965.

[23] Planck Collaboration, Ade, P. A. R., Aghanim, N., Alves, M. I. R., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Aussel, H., Baccigalupi, C., Banday, A. J., Barreiro, R. B., Barrena, R., Bartelmann, M., Bartlett, J. G., Bartolo, N., Basak, S., Battaner, E., Battye, R., Benabed, K., Benoît, A., Benoit-Lévy, A., Bernard, J.-P., Bersanelli, M., Bertincourt, B., Bethermin, M., Bielewicz, P., Bikmaev, I., Blanchard, A., Bobin, J., Bock, J. J., Böhringer, H., Bonaldi, A., Bonavera, L., Bond, J. R., Borrill, J., Bouchet, F. R., Boulanger, F., Bourdin, H., Bowyer, J. W., Bridges, M., Brown, M. L., Bucher, M., Burenin, R., Burigana, C., Butler, R. C., Calabrese, E., Cappellini, B., Cardoso, J.-F., Carr, R., Carvalho, P., Casale, M., Castex, G., Catalano, A., Challinor, A., Chamballu, A., Chary, R.-R., Chen, X., Chiang, H. C., Chiang, L.-Y, Chon, G., Christensen, P. R., Churazov, E., Church, S., Clemens, M., Clements, D. L., Colombi, S., Colombo, L. P. L., Combet, C., Comis, B., Couchot, F., Coulais, A., Crill, B. P., Cruz, M., Curto, A., Cuttaia, F., Da Silva, A., Dahle, H., Danese, L., Davies, R. D., Davis, R. J., de Bernardis, P., de Rosa, A., de Zotti, G., Déchelette, T., Delabrouille, J., Delouis, J.-M., Démoclès, J., Désert, F.-X., Dick, J., Dickinson, C., Diego, J. M., Dolag, K., Dole, H., Donzelli, S., Doré, O., Douspis, M., Ducout, A., Dunkley, J., Dupac, X., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Fabre, O., Falgarone, E., Falvella, M. C., Fantaye, Y., Fergusson, J., Filliard, C., Finelli, F., Flores-Cacho, I., Foley, S., Forni, O., Fosalba, P., Frailis, M., Fraisse, A. A., Franceschi, E., Freschi, M., Fromenteau, S., Frommert, M., Gaier, T. C., Galeotta, S., Gallegos, J., Galli, S., Gandolfo, B., Ganga, K., Gauthier, C., Génova-Santos, R. T., Ghosh, T., Giard, M.,

Giardino, G., Gilfanov, M., Girard, D., Giraud-Héraud, Y., Gjerløw, E., González-Nuevo, J., Górski, K. M., Gratton, S., Gregorio, A., Gruppuso, A., Gudmundsson, J. E., Haissinski, J., Hamann, J., Hansen, F. K., Hansen, M., Hanson, D., Harrison, D. L., Heavens, A., Helou, G., Hempel, A., Henrot-Versillé, S., Hernández-Monteagudo, C., Herranz, D., Hildebrandt, S. R., Hivon, E., Ho, S., Hobson, M., Holmes, W. A., Hornstrup, A., Hou, Z., Hovest, W., Huey, G., Huffenberger, K. M., Hurier, G., Ilić, S., Jaffe, A. H., Jaffe, T. R., Jasche, J., Jewell, J., Jones, W. C., Juvela, M., Kalberla, P., Kangaslahti, P., Keihänen, E., Kerp, J., Keskitalo, R., Khamitov, I., Kiiveri, K., Kim, J., Kisner, T. S., Kneissl, R., Knoche, J., Knox, L., Kunz, M., Kurki-Suonio, H., Lacasa, F., Lagache, G., Lähteenmäki, A., Lamarre, J.-M., Langer, M., Lasenby, A., Lattanzi, M., Laureijs, R. J., Lavabre, A., Lawrence, C. R., Le Jeune, M., Leach, S., Leahy, J. P., Leonardi, R., León-Tavares, J., Leroy, C., Lesgourgues, J., Lewis, A., Li, C., Liddle, A., Liguori, M., Lilje, P. B., Linden-Vørnle, M., Lindholm, V., López-Caniego, M., Lowe, S., Lubin, P. M., Macías-Pérez, J. F., MacTavish, C. J., Maffei, B., Maggio, G., Maino, D., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Marinucci, D., Maris, M., Marleau, F., Marshall, D. J., Martin, P. G., Martínez-González, E., Masi, S., Massardi, M., Matarrese, S., Matsumura, T., Matthai, F., Maurin, L., Mazzotta, P., McDonald, A., McEwen, J. D., McGehee, P., Mei, S., Meinhold, P. R., Melchiorri, A., Melin, J.-B., Mendes, L., Menegoni, E., Mennella, A., Migliaccio, M., Mikkelsen, K., Millea, M., Miniscalco, R., Mitra, S., Miville-Deschênes, M.-A., Molinari, D., Moneti, A., Montier, L., Morgante, G., Morisset, N., Mortlock, D., Moss, A., Munshi, D., Murphy, J. A., Naselsky, P., Nati, F., Natoli, P., Negrello, M., Nesvadba, N. P. H., Netterfield, C. B., Nørgaard-Nielsen, H. U., North, C., Noviello, F., Novikov, D., Novikov, I., O'Dwyer, I. J., Orieux, F., Osborne, S., O'Sullivan, C., Oxborrow, C. A., Paci, F., Pagano, L., Pajot, F., Paladini, R., Pandolfi, S., Paoletti, D., Partridge, B., Pasian, F., Patanchon, G., Paykari, P., Pearson, D., Pearson, T. J., Peel, M., Peiris, H. V., Perdereau, O., Perotto, L., Perrotta, F., Pettorino, V., Piacentini, F., Piat, M., Pierpaoli, E., Pietrobon, D., Plaszczynski, S., Platania, P., Pogosyan, D., Pointecouteau, E., Polenta, G., Ponthieu, N., Popa, L., Poutanen, T., Pratt, G. W., Prézeau, G., Prunet, S., Puget, J.-L., Pullen, A. R., Rachen, J. P., Racine, B., Rahlin, A., Räth, C., Reach, W. T., Rebolo, R., Reinecke, M., Remazeilles, M., Renault, C., Renzi, A., Riazuelo, A., Ricciardi, S., Riller, T., Ringeval, C., Ristorcelli, I., Robbers, G., Rocha, G., Roman, M., Rosset, C., Rossetti, M., Roudier, G., Rowan-Robinson, M., Rubiño-Martín, J. A., Ruiz-Granados, B., Rusholme, B., Salerno, E., Sandri, M., Sanselme, L., Santos, D., Savelainen, M., Savini, G., Schaefer, B. M., Schiavon, F., Scott, D., Seiffert, M. D., Serra, P., Shellard, E. P. S., Smith, K., Smoot, G. F., Souradeep, T., Spencer, L. D., Starck, J.-L., Stolyarov, V., Stompor, R., Sudiwala, R., Sunyaev, R., Sureau, F., Sutter, P., Sutton, D., Suur-Uski, A.-S., Sygnet, J.-F., Tauber, J. A., Tavagnacco, D., Taylor, D., Terenzi, L., Texier, D., Toffolatti, L., Tomasi, M., Torre, J.-P., Tristram, M., Tucci, M., Tuovinen, J., Türler, M., Tuttlebee, M., Umana, G., Valenziano, L., Valiviita, J., Van Tent, B., Varis, J., Vibert, L., Viel, M., Vielva, P., Villa, F., Vittorio, N., Wade, L. A., Wandelt, B. D., Watson, C., Watson, R., Wehus, I. K., Welikala, N., Weller, J., White, M., White, S. D. M., Wilkinson, A., Winkel, B., Xia, J.-Q., Yvon, D., Zacchei, A., Zibin,

J. P., and Zonca, A. Planck 2013 results. I. Overview of products and scientific results. *A&A*, 571:A1, 2014. URL: http://dx.doi.org/10.1051/0004-6361/201321529, doi:10.1051/0004-6361/201321529.

[24] Planck Collaboration, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Baccigalupi, C., Banday, A. J., Barreiro, R. B., Battaner, E., Benabed, K., Benot, A., Benoit-Lvy, A., Bernard, J.-P., Bersanelli, M., Bielewicz, P., Bobin, J., Bock, J. J., Bond, J. R., Borrill, J., Bouchet, F. R., Bowyer, J. W., Bridges, M., Bucher, M., Burigana, C., Cardoso, J.-F., Catalano, A., Challinor, A., Chamballu, A., Chary, R.-R., Chiang, H. C., Chiang, L.-Y, Christensen, P. R., Church, S., Clements, D. L., Colombi, S., Colombo, L. P. L., Couchot, F., Coulais, A., Crill, B. P., Curto, A., Cuttaia, F., Danese, L., Davies, R. D., de Bernardis, P., de Rosa, A., de Zotti, G., Delabrouille, J., Delouis, J.-M., Dsert, F.-X., Diego, J. M., Dole, H., Donzelli, S., Dor, O., Douspis, M., Dunkley, J., Dupac, X., Efstathiou, G., Enlin, T. A., Eriksen, H. K., Finelli, F., Forni, O., Frailis, M., Fraisse, A. A., Franceschi, E., Galeotta, S., Ganga, K., Giard, M., Giraud-Hraud, Y., Gonzlez-Nuevo, J., Grski, K. M., Gratton, S., Gregorio, A., Gruppuso, A., Gudmundsson, J. E., Haissinski, J., Hansen, F. K., Hanson, D., Harrison, D., Henrot-Versill, S., Hernndez-Monteagudo, C., Herranz, D., Hildebrandt, S. R., Hivon, E., Hobson, M., Holmes, W. A., Hornstrup, A., Hou, Z., Hovest, W., Huffenberger, K. M., Jaffe, A. H., Jaffe, T. R., Jones, W. C., Juvela, M., Keihnen, E., Keskitalo, R., Kisner, T. S., Kneissl, R., Knoche, J., Knox, L., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J.-M., Lasenby, A., Laureijs, R. J., Lawrence, C. R., Leonardi, R., Leroy, C., Lesgourgues, J., Liguori, M., Lilje, P. B., Linden-Vrnle, M., Lpez-Caniego, M., Lubin, P. M., Macas-Prez, J. F., MacTavish, C. J., Maffei, B., Mandolesi, N., Maris, M., Marshall, D. J., Martin, P. G., Martnez-Gonzlez, E., Masi, S., Massardi, M., Matarrese, S., Matsumura, T., Matthai, F., Mazzotta, P., McGehee, P., Melchiorri, A., Mendes, L., Mennella, A., Migliaccio, M., Mitra, S., Miville-Deschnes, M.-A., Moneti, A., Montier, L., Morgante, G., Mortlock, D., Munshi, D., Murphy, J. A., Naselsky, P., Nati, F., Natoli, P., Netterfield, C. B., Nrgaard-Nielsen, H. U., Noviello, F., Novikov, D., Novikov, I., Osborne, S., Oxborrow, C. A., Paci, F., Pagano, L., Pajot, F., Paoletti, D., Pasian, F., Patanchon, G., Perdereau, O., Perotto, L., Perrotta, F., Piacentini, F., Piat, M., Pierpaoli, E., Pietrobon, D., Plaszczynski, S., Pointecouteau, E., Polegre, A. M., Polenta, G., Ponthieu, N., Popa, L., Poutanen, T., Pratt, G. W., Przeau, G., Prunet, S., Puget, J.-L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renault, C., Ricciardi, S., Riller, T., Ristorcelli, I., Rocha, G., Rosset, C., Roudier, G., Rowan-Robinson, M., Rusholme, B., Sandri, M., Santos, D., Sauv, A., Savini, G., Scott, D., Shellard, E. P. S., Spencer, L. D., Starck, J.-L., Stolyarov, V., Stompor, R., Sudiwala, R., Sureau, F., Sutton, D., Suur-Uski, A.-S., Sygnet, J.-F., Tauber, J. A., Tavagnacco, D., Terenzi, L., Tomasi, M., Tristram, M., Tucci, M., Umana, G., Valenziano, L., Valiviita, J., Van Tent, B., Vielva, P., Villa, F., Vittorio, N., Wade, L. A., Wandelt, B. D., Yvon, D., Zacchei, A., and Zonca, A. Planck 2013 results. VII. HFI time response and beams. *A&A*, 571:A7, 2014. URL: http://dx.doi.org/10.1051/0004-6361/201321535, doi:10.1051/0004-6361/201321535.

[25] James Tomlinson Sayre. *Measuring Polarization of the Cosmic Microwave Background with the South Pole Telescope Polarization Experiment*. PhD thesis, Case Western Reserve University, 2014.

[26] D. Schwan, P. A. R. Ade, K. Basu, A. N. Bender, F. Bertoldi, H.-M. Cho, G. Chon, John Clarke, M. Dobbs, D. Ferrusca, R. Gsten, N. W. Halverson, W. L. Holzapfel, C. Horellou, D. Johansson, B. R. Johnson, J. Kennedy, Z. Kermish, R. Kneissl, T. Lanting, A. T. Lee, M. Lueker, J. Mehl, K. M. Menten, D. Muders, F. Pacaud, T. Plagge, C. L. Reichardt, P. L. Richards, R. Schaaf, P. Schilke, M. W. Sommer, H. Spieler, C. Tucker, A. Weiss, B. Westbrook, and O. Zahn. Invited Article: Millimeter-wave bolometer array receiver for the Atacama pathfinder experiment Sunyaev-Zeldovich (APEX-SZ) instrument. *Review of Scientific Instruments*, 82(9):–, 2011. URL: http://scitation.aip.org/content/aip/journal/rsi/82/9/10.1063/1.3637460, doi:http://dx.doi.org/10.1063/1.3637460.

[27] G F Smoot, C L Bennett, A Kogut, E L Wright, J Aymon, N W Boggess, E S Cheng, G de Amici, S Gulkis, M G Hauser, G Hinshaw, P D Jackson, M Janssen, E Kaita, T Kelsall, P Keegstra, C Lineweaver, K Loewenstein, P Lubin, J Mather, S S Meyer, S H Moseley, T Murdock, L Rokke, R F Silverberg, L Tenorio, R Weiss, and D T Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *The Astrophysical Journal*, 396:L1–L5, September 1992.

[28] W. C. Stewart. Currentvoltage characteristics of Josephson junctions. *Applied Physics Letters*, 12(8):277–280, 1968. URL: http://scitation.aip.org/content/aip/journal/apl/12/8/10.1063/1.1651991, doi:http://dx.doi.org/10.1063/1.1651991.

[29] Takayuki Tomaru, Masashi Hazumi, Adrian T. Lee, Peter Ade, Kam Arnold, Darcy Barron, Julian Borrill, Scott Chapman, Yuji Chinone, Matt Dobbs, Josquin Errard, Giullo Fabbian, Adnan Ghribi, William Grainger, Nils Halverson, Masaya Hasegawa, Kaori Hattori, William L. Holzapfel, Yuki Inoue, Sou Ishii, Yuta Kaneko, Brian Keating, Zigmund Kermish, Nobuhiro Kimura, Ted Kisner, William Kranz, Frederick Matsuda, Tomotake Matsumura, Hideki Morii, Michael J. Myers, Haruki Nishino, Takahiro Okamura, Erin Quealy, Christian L. Reichardt, Paul L. Richards, Darin Rosen, Colin Ross, Akie Shimizu, Michael Sholl, Praween Siritanasak, Peter Smith, Nathan Stebor, Radek Stompor, Aritoki Suzuki, Jun-ichi Suzuki, Suguru Takada, Ken-ichi Tanaka, and Oliver Zahn. The POLARBEAR-2 experiment, 2012. URL: http://dx.doi.org/10.1117/12.926158, doi:10.1117/12.926158.