Design and testing of tuning algorithms for the E and B EXperiment (EBEX)

Kevin MacDermid

M. Sc.

Department of Physics

McGill University

Montreal,Quebec

August 2008

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

Your file Votre référence
ISBN: 978-0-494-67037-8
Our file Notre référence
ISBN: 978-0-494-67037-8

ACKNOWLEDGEMENTS

First I would like to acknowledge the tireless support of my supervisor, Matt Dobbs who took the time to help me understand all parts of the experiment throughout my Masters and has patiently edited more drafts of this document than any single person should have to bear.

Next, I would like to thank everyone who helped develop this project, including Johannes Hubmayr and François Aubin, Graeme Smecher and Eric Bissonnette. A special thanks to Trevor Lanting as well, for helping me get a grasp on SQUID operation and good cryogenic practices that were so vital in commissioning the test dewar.

Finally, I would like to acknowledge my friends and office mates in 230, who have made Montreal a great place to live. In particular, thanks to James Kennedy for hours of amusement over the last two years and more links to Onion articles than many people are fortunate enough to receive in an entire lifetime, and Sarah Dyck, for taking care of me throughout the writing process and long before.

ABSTRACT

Detection of B-mode polarization from the cosmic microwave background would provide compelling evidence for the inflationary paradigm and has thus become a principal goal for experimental cosmologists during the last 5-10 years. A number of high sensitivity experiments have been developed and many are under construction, including the E and B EXperiment (EBEX), a balloon-borne experiment scheduled to take data in 2010. The design of EBEX is presented here, including the scientific motivation for the experiment, and an overview of all its components, with particular attention paid to the read out electronics for its array of over 1300 bolometric detectors. In the read out of such a large array of bolometers, which must be kept at $\sim$ 250mK, running many signals down the same wire (i.e. multiplexing) is vital. The digital frequency multiplexing (dfmux) electronics designed and tested at McGill addresses this as well as providing the control required to tune the bolometers and super conducting quantum interference devices (SQUIDs) used in readout. How these electronics accomplish both of these goals is described with an in-depth description of the recently designed tuning algorithms required to take the detectors from initialization to fully-operational. Finally, some tests of the readout system on cold bolometers in a test cryostat are presented. These measurements are encouraging with 98% success rate of automated tuning on a test sample of 47 bolometers. Improvements to achieve a 100% success rate have been proposed and will be implemented for the upcoming integration of the bolometer camera with the telescope in

November 2008. A test-flight for the experiment is planned for spring 2009, and the science flight for 2010 from Antarctica.

ABRÉGÉ

La détection des modes B de polarisation du rayonnement cosmologique fossile (CMB) serait une évidence forte du paradigme inflationniste. Cette détection est par conséquent devenue un des buts des cosmologistes expérimentateurs au cours de la dernière décennie. De nombreuses expériences de haute sensibilité ont été développées alors que plusieurs autres sont présentement en construction, incluant "E and B EXperiment" (EBEX), un télescope situé sur un ballon. Le plan de cette expérience est présenté dans ce document en incluant ses motivations scientifiques ainsi qu'une vue d'ensemble de ses composantes en insistant sur le système d'acquisition de données des 1300 bolomètres. Afin d'acquisitionner autant de données des bolomètres, qui doivent être conservés à une température de 250mK, le multiplexage est primordial. L'électronique de multiplexage développée et testée à l'université McGill est utilisée à cette fin en plus de fournir le contrôle requis afin de régler les bolomètres et les SQUIDs utilisés comme instruments de mesure. L'accomplissement de ces objectifs est décrit en profondeur via la description des algorithmes de réglage récemment développés qui permettent de rendre les détecteurs opérationnels. Finalement, quelques tests du système d'acquisition de données sur des bolomètres sont présentés. Ces mesures sont encourageantes avec un taux d'efficacité de 98% pour le réglage automatique d'un échantillon de 47 bolomètres. Les améliorations nécessaires afin d'atteindre un taux d'efficacité de 100% ont été proposées et seront incorporées

lors de l'intégration prochaine du plan focal de bolomètres avec le télescope en novembre 2008. Un vol d'essai est prévu pour EBEX au printemps 2009 et le vol scientifique est prévu pour 2010 en Antarctique.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Introduction and Science Motivation

The "Inflationary Big Bang" model is currently the leading theory of universe formation. It states that the universe underwent a momentary exponential expansion shortly after the big bang. This theory is consistent with all available astrophysical measurements [5] but is far from the only possible scenario [6]. While it is currently the leading theory, and more or less has been since it was proposed in 1981 by Alan Guth, it lacks strong confirmation. That is to say, while inflation solves a number of problems in cosmology, which will be further discussed in section 1.1, it leaves some serious questions of its own, including just why the universe underwent exponential expansion at all.

There is one prediction of the inflationary paradigm that has yet to be measured, often called the "smoking gun" of inflation, the signal of the "inflationary gravity wave background" (IGB). If the universe underwent exponential expansion as the inflationary paradigm suggests then relativity predicts that gravity waves should have been produced and that inflationary expansion would have stretched them to super-horizon scales. The only known way to detect these gravity waves is to look for the imprint they left on the Cosmic Microwave Background radiation (CMB) in the form of a specific polarization pattern, called B-mode polarization.

Detection of B-mode polarization is one of the main science goals of the E and B experiment (EBEX). EBEX is a balloon borne experiment that will measure the polarization of the CMB to provide a glimpse of the universe at its very earliest stages. It consists of a 1.5m aperture telescope and an array of 1406 bolometric transition edge sensors behind a half wave plate and polarizing grid. It it scheduled to fly its test flight in the Spring of 2009, and its 21 day long duration flight over Antarctica the following year. EBEX will either detect B-mode polarization or set an upper limit 10 times tighter than current bounds.

Of course, B mode polarization is not the only goal of EBEX (or it would just be called BEX). EBEX also intends to make a cosmic-variance limited measurement of the E-mode polarization pattern from $16 < l < 1500$, where l is the spherical multipole moment and thus inversely proportional to the angular scale on the sky. Additionally, EBEX intends to provide some of the first data about the polarization of Galactic foregrounds, which will be interesting in its own right and provide vital information for the design of future polarization experiments.

This document will attempt to work from the general toward the specific. This chapter is dedicated to providing a scientific background and motivating the need for the E and B experiment (EBEX). The following chapter outlines the experimental approach that will be taken by EBEX, finishing with a description of the detectors. Next, the detectors themselves and the read out electronics are described, followed by an in-depth description of the tuning algorithms used to prepare the system for taking data. Finally, a description of the test setup at McGill and test measurements of the complete readout system in operation will be given.

1.1 Inflation and the Cosmic Microwave Background Radiation

The current model of universe formation, 'The Big Bang Model', hypothesizes that the universe began as a single point in space-time and that it began to expand in an explosion of matter and energy. At first the universe was so dense and hot that it was entirely a plasma as any atoms that might temporarily form were immediately split again by the high-energy photons that permeated space.

As it expanded, however, it cooled, and eventually, nearly four hundred thousand years after the big bang, it cooled enough to let atoms of hydrogen form. At this moment the universe changed from opaque to transparent as the photons previously scattering from the free electrons were suddenly free to travel across the universe largely unimpeded. Now, some 13 billion years later, these first freed photons are still available for us to observe and have been named the Cosmic Microwave Background radiation, or CMB.

What is perhaps most surprising about the CMB signal is that it has a particularly uniform power from every direction. It all but perfectly fits a blackbody curve at 2.7K at every point, with variations only on the order of one in10^{-5} [7]. The reason this is surprising is that if the CMB was created less than 400 000 years after the big bang then there should not have been time for a region on one side of the sky to have communicated with another, due to information only being able to travel at the speed of light and the "horizon" at last scattering being greater than 400 000 light years across. This presents a problem, commonly referred to as the "horizon problem," just how did one side of the CMB become homogeneous with the other if they had never been causally connected?

Additionally, the tiny variations themselves present another mystery for cosmologists. If the universe began at a singularity, how did one location in the universe become so much different than another in terms of matter/energy density?

To answer these questions the theory of inflation was formed, which provided a solution to both these problems in the form of a universe that expanded exponentially during the first 10^{-35} seconds of its existence. If this were the case, then it is possible for quantum fluctuations present in the initial universe to be "stretched" to cosmic proportions and act as seeds for the anisotropies visible across the observable CMB and, indeed, the visible universe. Furthermore, inflation suggests that, at some point in the distant past, all of the universe was in causal contact, and hence it could become homogenous as required to solve the horizon problem.

While inflation neatly wraps up the horizon problem and provides a means for anisotropies to be created, it is difficult to prove that it actually occurred and leaves questions as to what caused it, and why. Fortunately, there remains another effect of inflation that is still visible today, and that is an imprint left on the *polarization* of the CMB. It is this so-called "smoking gun" of inflation that the E and B Experiment (EBEX) is attempting to measure. Further details of the polarization of the CMB follow in the next section, the details of how EBEX will attempt to measure the polarization pattern follows in the third chapter.

1.2 Polarization of the CMB

The polarization of the CMB is a powerful tool for testing our theories of universe formation and early evolution as it provides a complementary set of information to the well-measured intensity map provided by WMAP and others [8]. That polarization should be produced from the CMB anisotropies is a fundamental prediction of the gravitational instability paradigm [1]. In fact, any theory of universe formation including fluctuations on the surface of last scattering also suggests polarization anisotropy due to polarization dependence of the Thomson scattering cross-section,

$$\frac{d\sigma_T}{d\Omega} \propto |\hat{\epsilon} \cdot \hat{\epsilon}'|^2 \tag{1.1}$$

where $\hat{\epsilon}$ is the polarization of the incoming vector and $\hat{\epsilon}'$ is the polarization vector of the scattered radiation. This can be thought of as the incident light setting up oscillations of the electron in the direction of its electric field vector, in other words along the axis of its polarization. Our oscillating electron then transmits a spectrum with an intensity peak in the direction normal to, with polarization parallel to, the incident radiation.

If the incident radiation field were isotropic the random distribution of incident polarization vectors would lead to an unpolarized outgoing radiation field. On the other hand, if the incident radiation field had a quadrupolar variation in intensity or temperature, with peaks at $\pi/2$ separations, the result is some degree of linear polarization, as shown in figure 1–1. This is the basis for polarization of the CMB, quadrupolar variations in the surface of last scattering will create a linear polarization for every point on the CMB.

If the temperature anisotropies we observe are truly caused by fluctuations on the surface of last scattering, then the quadrupolar variations they produced must

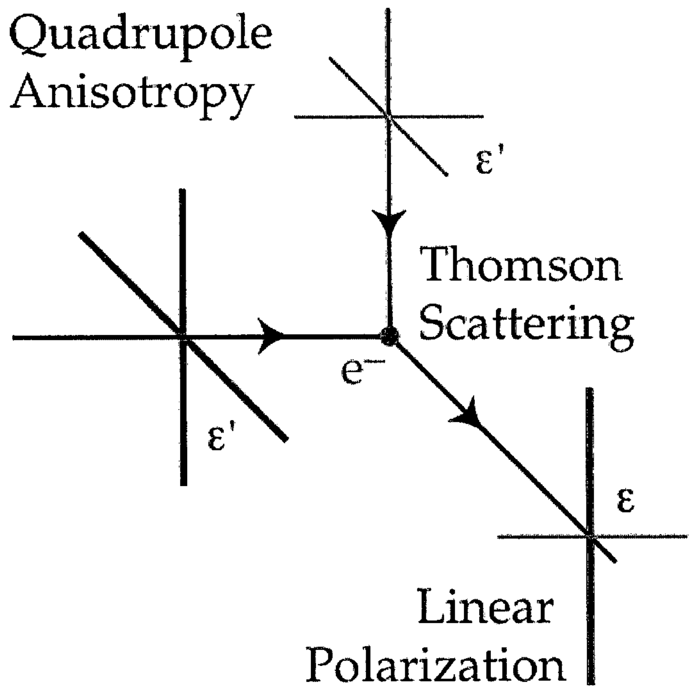

Figure 1–1: Thomson scattering of radiation. Blue (thick) lines represent hotter radiation, and red (thin) lines represent colder radiation. Image credit Hu, White. [1]

have polarized the CMB at some level. What's more, the different sources of temperature anisotropies, scalar, vector, and tensor, produce different patterns in both the polarization itself and its correlation with the temperature anisotropies. So, by studying the polarization of the CMB, we can separate the components of the temperature power spectrum and constrain cosmological models.

Before moving on it is worth mentioning the degree of polarization expected and, to some degree, measured. From the discussion earlier in this section it is apparent that the level of polarization is directly related to the quadrupolar anisotropy at last scattering. What may not have been so apparent is that the fraction of polarized light is strongly dependent on the duration of last scattering. This follows from the fact that photons of different energies must be able to reach the same electron for that electron to "see" a quadrupole. This was unlikely to happen in the early plasma, as the mean free path of a photon was too short for it to travel to a region of significantly different temperature without being scattered. Thus, only during recombination, while the density of free electrons is diminishing, can photons travel across a temperature gradient and provide the quadrupolar variation required to produce linear polarization [9].

So the question of how much light will be polarized depends on the duration of last scattering. For the standard thermal history it is thought to be around 10% [1] on the scale of tens of arcminutes. Thus, while the temperature anisotropies are at the 10^{-5} level fractionally, the polarization appears at the 10^{-6}, or several μK or less. The small size of this signal is why it has remained largely unmeasured until now. Experiments are only now reaching the sensitivity to make high signal to noise measurements of these tiny temperature signals.

Generally, when discussing CMB polarization, two "modes" are discussed, so-called "E-modes" and "B-modes." Here the "E" and "B" are meant to create an analogy between these modes and the electric and magnetic fields, respectively. In particular, the E-mode polarization pattern (E-modes) is curl-free while the B-mode pattern (B-modes) is not. Each mode can only be generated by certain types of perturbation on the surface of last scattering which are describe in increased detail in Appendix A.

E-mode polarization of the CMB has fairly recently been detected by the Degree Angular Scale Interferometer (DASI) [10] in 2002, and has even had its power spectrum measured, to some degree, by subsequent experiments such as CBI [11], BOOMERANG [12] and more recently QUaD [2]. These results are interesting, as they lend credence to the existence of density variations on the surface of last scattering; however, they cannot lend significant support to the model of inflation on their own as they can be created by scalar perturbations that occur in any model with variations on the surface of last scattering.

Much stronger evidence in support of the inflationary paradigm could be present in the currently undetected B-mode polarization signal, as the only known means of creating super-horizon B-mode CMB polarization that would be visible to us is to have gravity waves pass through the surface of last scattering [13]. Unfortunately, current predictions of the magnitude of the B-mode signal put it at over 10 times smaller than the E-mode polarization signal, just hundreds of nano Kelvin. This presents a significant challenge for experimental cosmologists.

1.3 Foregrounds

While CMB polarization data can provide fascinating clues as to the evolution of the early universe, it is no simple matter to measure it precisely enough. Beyond the technological challenges of measuring such a faint signal is the fact that not a great deal is known about the polarized foregrounds that will have to be removed during analysis.

The main sources of polarized foregrounds are synchrotron emission at low frequencies and polarized dust emission at high frequencies. While it is appreciated that these foregrounds will significantly impact the search for the signal of the IGB, little is known about either between 100 and 450GHz [14].

One common method is to extrapolate from the WMAP data at 23GHz and 90GHz to make estimates of the synchrotron emission spectrum [15]. Such extrapolations have the synchrotron emission drop to below $10^{-2}\mu$K by 70 GHz, which should be significantly below the B-mode CMB polarization peak of $\sim 10^{-1}$ at 150 GHz. Thus synchrotron is less of a concern for measurements of the IGB signal.

Similarly, for dust it is assumed that the synchrotron emission seen in the WMAP 23 GHz signal traces the galactic magnetic field and that the dust aligns very well with this magnetic field. Assuming some small fraction of this aligned dust provides a polarized signal a spectrum of dust emission can be produced. This spectrum crosses synchrotron below 100 GHz and is on the order of $10^{-1}\mu K$ at 150 GHz, thus providing a significant foreground to IGB detection.

There are several methods employed to avoid foreground confusion of the signal. For dust the first and simplest is to avoid dusty regions such as the galaxy. The second is to make multiple frequency measurements and remove the foregrounds during analysis by comparing their spectra to that of the well known CMB spectrum.

1.4 Summary

Detection of B-mode polarization of the CMB would provide some exciting details of how the universe formed and lend significant weight to the inflationary paradigm. It also provides a distinct experimental challenge due to its low amplitude and the presence of, largely unmeasured, foregrounds. The E and B EXperiment plans to meet this challenge with the design presented in the following chapter.

CHAPTER 2
The E and B Experiment (EBEX)

The previous chapter has hopefully made clear how important the primordial B-mode polarization signal is. The rest of this document will describe the approach we will take to measure this signal. This chapter provides an overview of the EBEX experiment as a whole, while later chapters describe the detectors and readout in increased detail.

2.1 Experiment Overview

As discussed in section 1, EBEX is a long-duration balloon borne experiment that will measure the polarization of the CMB. During its long duration flight it will fly above much of the atmosphere at 120 000 feet, so that the background loading contribution from the atmosphere is negligible. This reduces both spurious signals from atmospheric fluctuations and the photon shot noise contribution from this loading, which would otherwise be one of the dominant terms in the instrument's noise budget.

EBEX consists of a 1.5m telescope that provides a resolution of less than 8 arcminutes over a 4° field of view. It will be equipped with over 1300 bolometeric transition edge sensors maintained at sub-Kelvin temperatures. These detectors will be split up into three frequency bands, 150, 250 and 410 GHz to measure polarized foregrounds, allowing us to gauge their importance to future CMB polarization experiments as well as subtract them out of our signal. As our bolometers are not inherently polarization sensitive, it will employ rotating half wave plate and a polarizing grid to modulate and select a single linear polarization, respectively. In fact, EBEX has two focal planes and the polarizing grid reflects one polarization to one and the orthogonal polarization to the other. The choice to use non-polarization sensitive bolometers was a conservative one in EBEX, it was deemed more reliable to re-use tested technology rather than trust that polarization sensitive bolometers could be produced and tested in time for the launch.

EBEX is currently under construction as it is scheduled to fly its test flight in the spring of 2009, and its two week long duration flight in the 2010/2011 Austral summer. The plan is to test integration of all parts of the experiment in November 2008 before taking it apart and shipping all the pieces to Palestine, Texas for re-integration and the day-long test flight in 2009.

2.2 Optics

Light from the sky is captured by the 1.5m aluminum primary mirror before being reflected to a 1m diameter ellipsoidal secondary, forming a "Gregorian Mizuguchi-Dragone" telescope [14]. A diagram of this design is presented in figure 2–1. One feature of this off-axis design is that it leaves the primary mirror unblocked, avoiding polarization created by having pieces of metal, such as supports, in the field of view. The telescope forms an image of the sky at the window of cryostat and there the cold optics begin.

The cold optics begin with a window on the top of the cryostat, as shown in figure 2–2. From there it passes through a series of low-pass filters (to block IR), followed by the first cold lens. After the first cold lens the light passes through

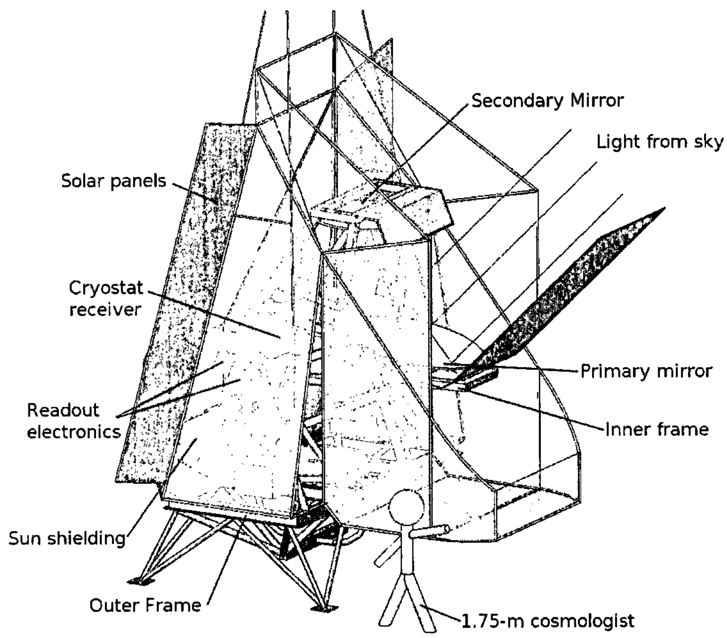

Figure 2–1: A model of the EBEX experiment and balloon gondola. Note that light from the sky comes unimpeded to the primary mirror, before being reflected to the secondary and into the cryostat. This minimizes polarization signals induced by supports, etc, in front of the primary.

the rotating half-wave plate that will be discussed in the next section before being focused by a final cold lens and sent to each of the two focal planes by the polarizing grid.

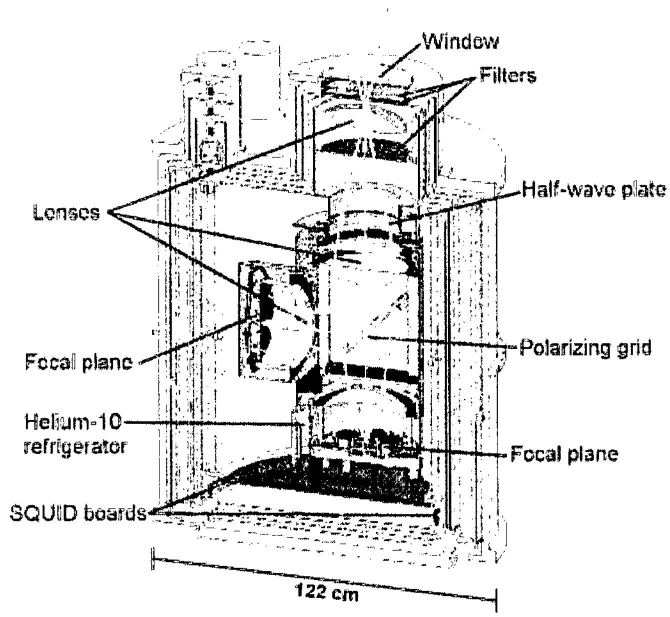

Figure 2–2: A model of the EBEX cryostat and cold optics. Incoming light is put through first a cold lens, then the rotating achromatic half-wave plate before hitting the polarizing grid and being sent to one of the two focal planes.

2.3 Polarimetry

The bolometers that make up the focal planes are not polarization sensitive, thus all separation into polarizations must be performed before the light falls on the focal planes. The separation of the incoming light from the sky into its constituent polarizations is performed by a wire-grid polarizer, which reflects light with linear polarization parallel to the wires while allowing light with its polarization perpendicular to the wires to pass through largely unattenuated. Note that, since a wire-grid thousand degrees of sky. EBEX will focus on the peak at $l = 100$ instead, by surveying 350 square degrees of sky [14]. The field, shown in figure 2–3, was selected to minimize foreground dust emission and provide overlapping coverage with QUAD, a ground-based E mode polarization experiment near the South Pole [2].

EBEX is designed with a resolution of about 8 arcminutes, adequate to probe angular scales up to $l = 1000$, where the as yet unmeasured signal of B-mode lensing is expected to be seen. Conical feed horns atop the focal plane adjust the area illuminated on the primary mirror so that the resolution does not change with frequency.

During its long-duration flight EBEX will integrate to an average level of 0.6 μK in the Q and U Stokes parameters and 0.5 μK in T per 8 arcminute beam. The EBEX field will provide 19 688 such "pixels" which will be sufficient to make a high signal-to-noise detection of B-mode signal down to a tensor to scalar ratio (τ) of 0.1. If the B-mode signal is not detected EBEX will put a 2σ upper limit on $\tau < 0.02$, improving current limits by a factor of ~ 5. [14].

2.5 Cryogenics

All of the cold optics, including the HWP and its superconducting bearing, as well as some cold electronics, most notably the super-conducting quantum interference devices (SQUIDs) must be maintained at temperatures below 5K. This will be achieved in EBEX using a liquid helium dewar, which will cool the main plate of the dewar to 4.4K or below. At the moment this is the only means to achieve these temperatures without requiring a great deal of power, which translates to an impossible weight for a balloon-borne experiment such as EBEX.

polarizer is used EBEX circular polarization will be split equally between the two focal planes, and thus will not be measured. This is acceptable as there is expected to be no significant circular polarization present on the sky at the scales EBEX is interested in.

To achieve the low noise measurement required for B mode detection EBEX will employ a rotating half-wave plate (HWP) to modulate the polarization seen by each focal plane. The HWP will spin at 6 Hz, and due to the half-wave nature of the plate each physical rotation causes the polarization vector to rotate 4 times, leading to a signal at 24 Hz. Modulation significantly reduces the effect of systematic errors as well as removing the otherwise significant low-frequency noise.

In the case of EBEX the half wave plate must be achromatic, as the frequency range of the experiment is broad. This is achieved by stacking 5 chromatic sapphire HWPs, each 1.62mm thick, at angles of $(0,25,75,25,0)^\circ$ with respect to the crystal axis of the first. Furthermore, this whole sandwich is maintained at a temperature of ~ 5 K in EBEX to reduce loading on the detectors. To accomplish this it floats atop a superconducting magnetic bearing within the cryostat with the rest of the cold optics [16]. The superconducting magnetic bearing ensures low power consumption and very low noise.

2.4 Sky Coverage, Resolution and Sensitivity

Since the IGB is created by inflation, the scale of the gravitational waves is very large, this leads to peaks in the B mode polarization signal at $l \sim 5$ and ~ 100, or several thousand square degrees and $\sim$350 square degrees respectively. The larger scale peak ($l \sim 5$) is only available to an instrument that surveys more than several

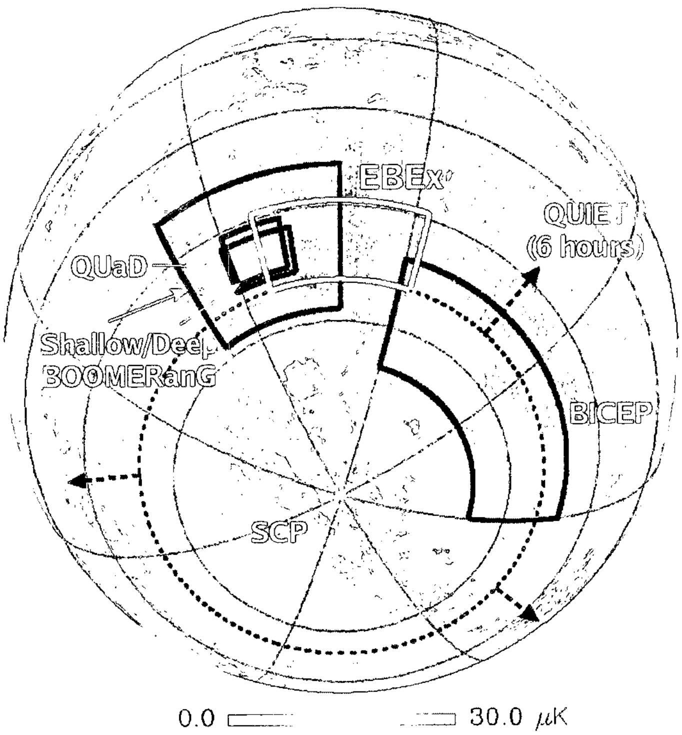

Figure 2–3: The EBEX scan area in relation to those of several other experiments. The overlapping coverage with QUAD is intended to strengthen foreground characterization by providing an additional frequency band at 100GHz [2]. The EBEX scan area has low dust emission at $\sim 3\mu K$ mean brightness, which should reduce the significance of this polarized foreground.

The cold optics will be cooled to 1 K to reduce the radiation load on the detectors. This will be achieved using a ^{4}He closed-cycle adsorption refrigerator made by Chase Research Cryogenics in the UK. A second, closed-cycle, three-stage adsorption refrigerator will keep the detectors cooler than 0.27K. Both of these refrigerators work by pumping on a small enclosed helium bath to reduce the liquid's boiling temperature. The pumped helium adsorbs onto a charcoal surface where it remains until the following fridge cycle, so there is no loss of helium from within these fridges. Of course, nothing is free, the heat they remove from the optics box and detectors is ultimately dumped on the liquid helium bath of the cryostat, which limits the duration of the experiment.

2.6 Focal Plane and Detectors

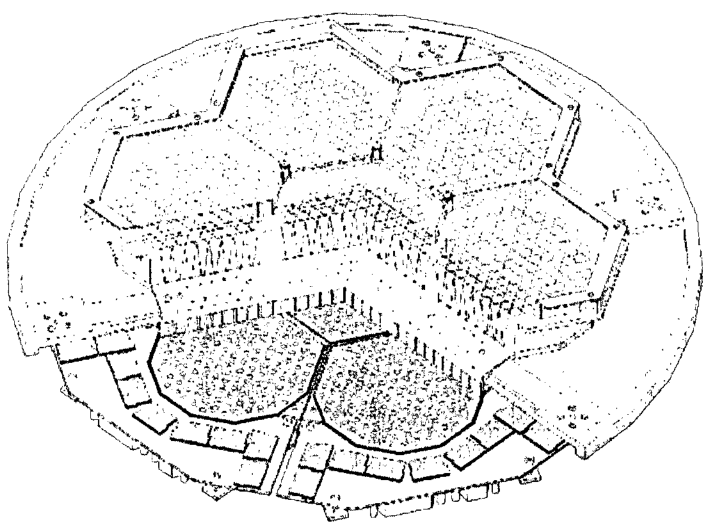

Figure 2–4: A diagram of an EBEX focal plane with part of the feed horn and wave guide array cut away to show the mechanical structure. Atop the feed horns sit the band-defining filters, four of which select 150 GHz, two select 240 GHz, and one 450 GHz. At the bottom the transition edge sensor bolometers are shown, these will be discussed in section 3.1.

EBEX has two focal planes, one for each linear polarization. Each focal plane is made up of seven wafers in a hexagonal close-packed pattern, see figure 2–4, with each wafer having 139 transition edge sensor (TES) bolometers, shown as small white circles. TES bolometers are extremely sensitive incoherent detectors that will be discussed in more detail in section 3.1.

The conical feed horns shown just below the band-defining filters in figure 2–4 couple the light from the sky to the individual detector while the wave guides sitting just below define the high frequency cut-off for each of the frequency bands. The filters sitting atop the feed horns define the low frequency cut-off.

With seven wafers of 139 bolometers, each focal plane can contain 973 detectors for a total of 1946 detectors if every pixel were used. In reality some of these will be left dark, and the lithographic techniques to create bolometer wedges are not perfect so a total of just over 1300 detectors are expected to look at the sky. If two wires were run from 4K to each detector the wiring would impose a significant heat load on the focal plane and it would be extremely difficult to maintain the temperatures sub-Kelvin required during detection. As such, considerable effort has gone into creating readout technologies using multiplexed bolometer readout. This readout is the major contribution from McGill and is the subject of much of the rest of this document.

CHAPTER 3
Detectors & Readout

The previous chapter outlined EBEX as a whole, this one will focus more specifically on the focal planes and the bolometer detectors that form them. Recall that the EBEX bolometers are not inherently polarization sensitive, so by the time the light has reached the focal plane it has already passed through the half-wave plate and polarizing grid, where the polarization selection has been made. This chapter describes how this light is converted into an electrical signal and how this signal is digitized and stored.

3.1 Bolometers and Transition Edge Sensors

Bolometers are incoherent detectors ideal for detecting the CMB as they exhibit no quantum noise. They consist of an absorbing element, with heat capacity C, which is weakly thermally linked to a heat sink at temperature T_s. When the bolometer receives a pulse of radiation power P from the sky its temperature increases to

$$T_{bolometer} = T_s + \frac{P}{G}\left(1 - e^{-t/\tau}\right), \tag{3.1}$$

where G is the thermal conductance of the link and $\tau = C/G$ is the thermal time constant of the bolometer.

EBEX uses what are called "transition edge sensor (TES) spider-web bolometers." Each part of this considerable mouthful describes a particular element of their construction and operation. To begin with the last part of the name, they are called "spider-web" bolometers due to their distinctive shape, which can be seen in the right panel of figure 3–1. This shape is created by applying a $1\mu m$ layer of low-stress silicon nitride to a wafer then etching away all but the traces shown in the figure [17]. What remains is a web of silicon nitride effectively hanging over a silicon substrate. The weak thermal link is provided by the legs at the edge to the rest of the silicon bath. The characteristic, spider-web, shape is designed to lower their cross section to cosmic rays while maintaining their high millimeter-wavelength photon absorption cross section.

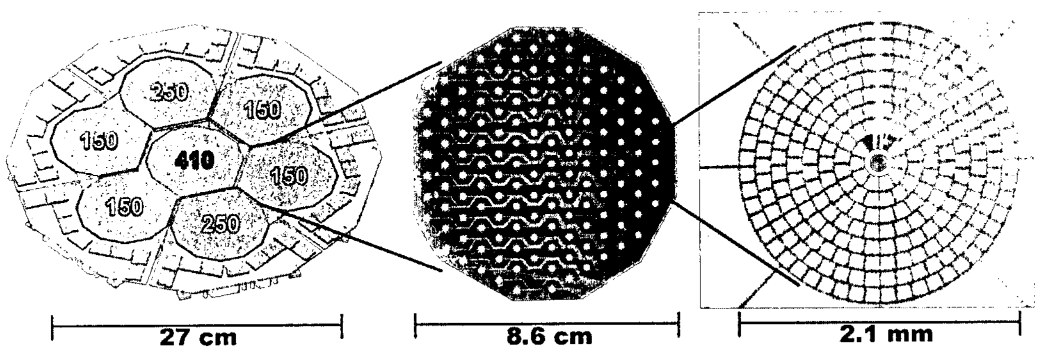

Figure 3–1: **Left Panel:** A diagram of the EBEX focal plane with the frequency band of each bolometer wedge printed on it. The preference toward 150 GHz is because the signal is expected to be largest, and the noise lowest, in this band. **Center Panel:** Photograph of an EBEX bolometer wafer. The small light dots are individual bolometers, the possibly-visible gray traces are the superconducting aluminum leads. **Right Panel:** A close up on a single bolometer. Its distinctive "spider-web" design is intended to provide a large cross section to millimeter-wavelength photons while providing little surface area for a cosmic ray to strike.

The amplitude of the polarization signal on the CMB is just hundreds of nano kelvin, so we must be capable of measuring *very* small changes in temperature. The first part of the name of the EBEX bolometers, "transition edge sensor," (TES) refers to the thermistor capable of measuring these variations. A TES is a superconductor biased into its superconducting transition so that, as its name implies, it sits on the "edge" of being a superconductor. Within this transition the resistance of the TES is strongly sensitive to changes in temperature, perfect for the detection of CMB polarization.

The challenge with TES bolometers is keeping them in the superconducting transition. It would be all but impossible to keep the focal plane temperature precisely within the transition region, especially since it may vary slightly from one bolometer to the next. Furthermore, any variation in the incoming power from the sky (such as those we are attempting to measure) would be sufficient to push the TES out of its transition.

The trick to keeping the TES stable in this balancing act is voltage-biasing. The total power present at the bolometer, which effectively sets its temperature since the thermal link is constant, can be separated into two components,

$$P_{total} = P_{sky} + P_{electrical} \tag{3.2}$$

The power from the sky, P_{sky}, will change as the telescope scans the sky, while the electrical power, $P_{electrical} = V_{bias}^2/R_{TES}$, is supplied by the readout electronics. Keep in mind that any increase in the total power at the bolometer will serve to push it up the transition curve, increasing its resistance, while any decrease will lower its resistance. With this is mind the negative feedback of this system is evident, whenever the sky power increases it drives the resistance of the TES up, which lowers the electrical power due to the constant voltage bias and the presence of R_{TES} in the denominator. Conversely, any decrease in sky power is offset by a corresponding increase in electrical power. The system is said to exhibit *strong electro-thermal feedback*, which results in stability and linearity.

In EBEX, each bolometer has a TES lithographically produced at the center, as may be just barely visible in the right panel of figure 3–1. These TES make use of the proximity effect, wherein different superconductors lain one atop the other can tune the transition temperature of the stack. EBEX uses aluminum and titanium to set the superconducting transition temperature, T_c, between 450 and 550 mK, just above their operating temperature of 250mK. They have a normal resistance of about 1Ω and are typically electrically biased to about 0.8Ω during operation, as this has empirically been found to be stable.

As mentioned in the opening to this section TES Bolometers are the most sensitive detectors of microwave radiation above around 100 GHz available, due to their lack of quantum noise [18]. As EBEX intends to observe near where the CMB peaks, at 150 GHz, and above, bolometers are the most appropriate technology to use. In fact, since individual bolometers approach the fundamental limit of photon (statistical) noise at the detectors, the only way to obtain the high sensitivity required by EBEX is to create an array of many pixels. Throughout the rest of this document, bolometer will always refer to spider-web transition edge sensor bolometer.

3.2 SQUIDs

As the bolometers are voltage biased, the sky signal appears as fluctuations in the current through the bolometer. While the bolometers are very sensitive to incoming power the signals they produce are still tiny, with current fluctuations on the order of pA. In EBEX, these signals are converted to voltage and amplified by

Super Conducting QUantum Interference Devices (SQUIDs), before being digitized by an analog-to-digital converter (ADC) on the digital frequency multiplexing board that will be discussed in section 3.5

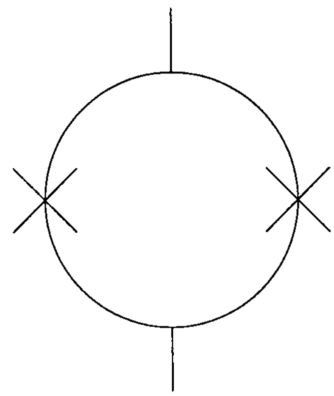

Figure 3–2: The circuit symbol used to indicate a SQUID. The two "X"'s indicate small insulating breaks in the ring of superconductor called Josephson junctions.

A SQUID consists of a loop of superconductor with two small breaks, called "Josephson junctions," interrupting it. The symbol used to indicate a SQUID is shown in figure 3–2. They make use of two fundamental properties of superconductors. The first is that the magnetic field through a loop of superconductor is quantized. Any attempt to change the magnetic field results in the generation of a current through the coil that cancels out the changes to the magnetic field. The second is the "Josephson effect," a full description of which is beyond the scope of this document (see [18]). For our purposes it is sufficient to know that the Josephson effect allows a voltage to be present across the SQUID whenever a current is circulating, as the two small insulators interrupt the flow of the super-current. The combination of these two elements makes a SQUID an extremely sensitive detector of changes in magnetic flux as any variation in magnetic flux is transduced to a voltage on the SQUID.

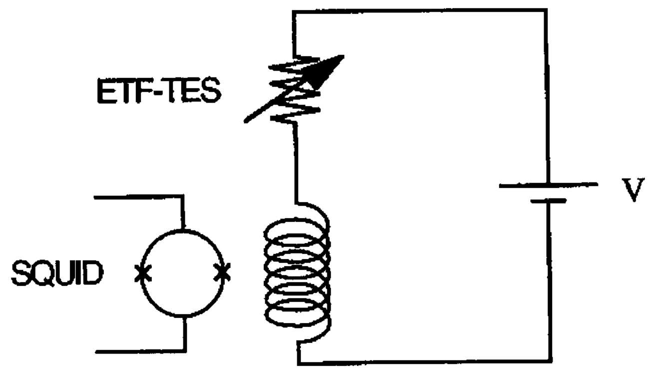

Figure 3–3: The basic implementation of a SQUID for bolometer read out. A steady bias voltage is supplied to both the bolometer (ETF-TES) and the inductive coil. Any change in resistance of the bolometer will be converted to a varying magnetic field by the coil and show up as a voltage on the SQUID. Image credit [3].

Thus, for a SQUID to detect the output of the bolometer, it will have to be converted to magnetic flux. This can easily be achieved by an inductor, as shown in the single bolometer readout circuit in figure 3–3. In this circuit a steady voltage bias is applied across the bolometer, which is essentially a variable-resistor, and the inductor. Changes in the resistance of the bolometer alter the current that flows through the inductor, which is coupled to the SQUID coil. The result is that variations in bolometer resistance, from the polarization signal, appear as considerably larger variations in voltage at the SQUID.

SQUIDs have a number of properties that make them ideal for bolometer read out. The first is their low input impedance. The bolometer has a normal resistance of 1Ω, so the input impedance must be $\ll 1\Omega$. This is accomplished with the above circuit design so long as the impedance of the inductor is low, which will be true for the low frequency signals from the bolometers.

The second property is noise lower than that of the bolometer, so as not to degrade the signal. A typical bolometer might have a current noise of $\sim 10pA/\sqrt{Hz}$, while a series-array SQUID readout, using many SQUIDs connected back to back, can achieve a current noise of $\sim 2pA/\sqrt{Hz}$. Since this is an incoherent noise source it will add in quadrature with the bolometer noise and not significantly degrade the polarization signal.

Finally, the SQUID needs to have a high forward gain, so that the signal will be increased above the noise in subsequent stages of the electronics, particularly those that are at room temperature, instead of the cryogenic temperatures near the detectors. So, while it is typical to use current-based signals, due to their lower susceptibility to noise, this is not an issue in the readout as the bolometer noise, after being amplified by the SQUID trans-impedance of 500Ω, is the dominant noise source.

Although SQUIDs are extraordinary flux to voltage transducers for bolometer signals they are not without their complications. First, they are extremely sensitive to long-wavelength transmissions, such as radio and cell phones, as well as magnetic fields, such as that of the Earth. The radio and cell phone transmissions are not a significant problem as the SQUIDs must sit within the cryostat to superconduct and the cryostat is effectively a Faraday cage. The magnetic fields, on the other hand, are dealt with by a mu-metal shield that surrounds each SQUID board inside the cryostat. Mu-metal is designed to have very high magnetic permeability, making it very effective at screening static or low-frequency magnetic fields.

Secondly, the response of SQUIDs is not linear but is rather somewhat sinusoidal, as is shown in section 4.1.3. This is dealt with in a similar fashion as the bolometers, using feedback although, in the case of the SQUIDs the feedback is entirely electrical. The SQUID signal passes through an operational amplifier, through a feedback resistor and back to the inductor coil. So long as the SQUID response curve has a negative slope, as it does on the downward section of the sinusoid, negative-feedback occurs and the SQUID is "locked" to one small region on its response curve. Much like the situation with the bolometers, this feedback provides both stability and linearity.

Much of the description so far describes a single SQUID coil; however, the SQUIDs used in EBEX are actually series-arrays of individual SQUID coils, with around 100 individual coils in each. They function as outlined so far, except that the signal-to-noise ratio is improved since the signal is coherent, so that it adds directly from one coil to next, while the SQUID noise is incoherent and adds in quadrature. It is these series-arrays that can achieve the low noise requirement mentioned earlier. In the rest of this document "SQUID" will always refer to a series-array of SQUID coils.

3.3 Frequency Domain Multiplexing

It would be extremely difficult, if not impossible, to run a separate wire for each individual pixel in EBEX, as the heat load on the cold stage would be enormous.

Instead, EBEX uses frequency domain multiplexing to read out eight bolometers with one set of wires, coupled to a single SQUID.

The frequency domain multiplexing done by the EBEX readout system is similar to the operation of an AM radio. Each bolometer is preceded by its own inductive-capacitive filter, which selects the small frequency band (a radio station in our analogy) that will pass through to the bolometer. The bandpass of the RLC circuit formed depends on the resistance of the bolometer, which is not fixed; however, for a bolometer similar to those used in EBEX, with a resistance of approximately 1Ω, the bandwidth is about 25kHz. These frequencies range from around 300 kHz to around 1 MHz, all of which are much faster than the bolometer can respond as our bolometer electrical time constants are 2-3ms. Thus, to the bolometer, these carrier frequencies look like a steady, DC, voltage bias and the signal from each bolometer amplitude modulates this carrier signal. In frequency space, this modulation shows up as side bands on the specific carrier signal received at the bolometer, effectively separating the bolometer signal in frequency space.

The schematic of this setup is shown in figure 3–4 for a single comb. As is visible in the diagram, the input "carrier comb" waveform (consisting of the 8 different frequencies summed together) is provided by the digital multiple frequency synthesizer (DMFS) on the digital frequency multiplexing (dfmux) board. This carrier comb passes through the bolometers, with each individual carrier frequency component passing through a single LC filter and on to a single bolometer. The signal is then "amplified" by the SQUID, low-pass filtered, and sent to the digital multiple frequency demodulator. Here it is mixed with a square-wave at the same frequency to return it to base band and digitally filtered by a series of finite-input-response (FIR) filters which reduce the data rate to 100-400 Hz (depending on the settings). The filtered data is finally streamed over ethernet to the data hard disks, where it will remain until the balloon completes its flight and the data analysis can begin.

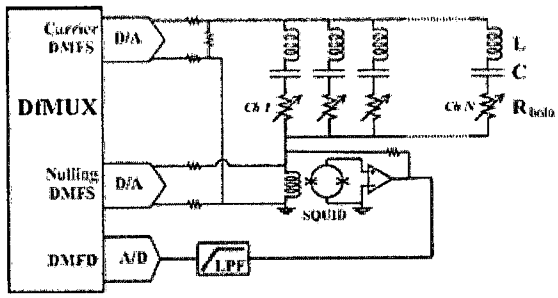

Figure 3–4: Schematic of the digital frequency multiplexing (dfmux). Here the dfmux board can be seen on the left providing the carrier and nuller signals from digital multiple frequency synthesizers (DMFS), as well as demodulating the bolometer output at the digital multiple frequency demodulator (DMFD). Note the Low Pass Filter (LPF) that precedes the DMFD is simply intended to remove high frequency noise sources. Image credit [4]

There is one additional complication which arises due to this scheme. As the carrier is considerably larger than the signal being measured much of the dynamic range of the SQUID and ADC would be committed to its measurement, were it present at the inductor coil. Since the carrier provides no useful information this would be a considerable waste of resources which would severely limit the number of detectors that could be multiplexed on a single wire. To rectify this the dfmux board generates a nulling wave tuned to be exactly the same as the carrier, as seen at the inductor coil, except that it is 180 degrees out of phase. Ideally, this "nuller" perfectly cancels the carrier allowing all of the dynamic range of the SQUID and

ADC to be used to measure the polarization signal. In practice, a small fraction, about 1%, of the carrier appears at the SQUID, which is small enough that it has no significant impact on the data.

At this point all of the readout elements have been discussed, and a schematic of how they fit together electrically has been presented in figure 3–4. The discussion that follows will center on the specific EBEX implementation of this system, including discussions of the circuit boards involved and the connections between them.

3.4 Experiment Schematic

This section describes the physical implementation of the readout components mentioned in the previous section, starting at the focal plane. Each comb of eight multiplexed bolometers is connected to a single SQUID. This SQUID sits on a small board, with 7 others, within the cyrostat to maintain the $\sim$ 4K temperatures required for its operation. All 8 of these SQUIDs are connected through a cable feed-through to a single "SQUID control board" that is just outside the cryostat, at 300K. Each SQUID control board is connected to 8 SQUIDs or 64 bolometers, due to multiplexing.

A SQUID control board has all of the analog electronics to provide the voltages required to tune a SQUID but must receive its power and commands from a dfmux board through a 25 pin control cable connected directly from one to the other. Each dfmux board can read-out 4 SQUIDs, so the output of each SQUID control board must be connected to *two* separate dfmux boards. Additionally each dfmux board has two "mezzanine" boards that sit atop it , as shown in figure 3–5. These mezzanine boards provide the ADCs for the input signal. So, the output of the SQUID board actually goes to 4 mezzanines, on two dfmux boards. Thankfully, this unfortunate complexity does not hinder the performance of the system, just those that have to explain it. For clarity figure 3–5 shows diagrammatically the connections between these elements.

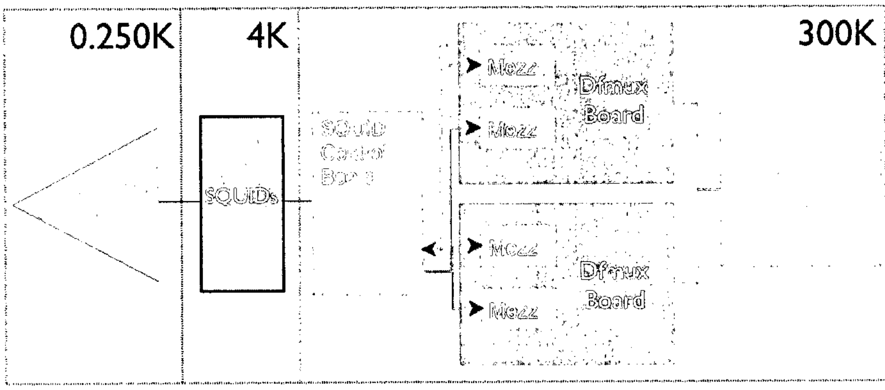

Figure 3–5: A schematic diagram of the readout system. Notice that the SQUID control board gets its input from one dfmux board but sends its output to *two*. All connections without arrows go both directions.

Once the signals have arrived at the dfmux boards they are further processed and sent, over ethernet, to be stored on the data hard drives that will sit within a pressure vessel on EBEX. In fact, all outside communication by the dfmux boards is done over ethernet. So, in our example setup in figure 3–5, each dfmux board has an IP address associated with it, 192.168.0.100 and 192.168.0.101, where they can receive commands to check their status and change their settings.

3.5 The Digital Frequency Multiplexing (dfmux) Board

The dfmux board is the heart of the readout electronics, providing the carrier and nuller combs, demodulating the return signal, and even controlling the SQUID bias levels. It is capable of all of this due to the powerful Xilinx field programmable gate array (FPGA) that sits at the center of the board and is highlighted in figure 3–6. An FPGA is similar to a typical computer processor except the connection of the gates is configurable. With the proper connections the FPGA is capable of performing many computations in parallel with low power consumption, perfect for a balloon experiment such as EBEX.

The dfmux board FPGA is largely dedicated to the generation of the carrier and nuller, in the DMFS, and the demodulation of the return signal, in the DMFD. It is capable of generating up to 32 separate frequency channels on each of its four "modules" or wires at sampling rates up to 25 MHz. Each of these wires connects to a separate comb of, currently, 8 bolometers, which are read out by a single SQUID, as shown in figure 3–4. The data measured for each of these channels is streamed out over ethernet in UDP format by a section of code called the "data streamer." During the EBEX flight the data will be sent using IP multicasting, wherein multiple listeners can receive the same transmission without interfering with one another. This will allow the data to be simultaneously stored to the data hard drives and evaluated by the flight control computer for problems.

While much of the FPGA is dedicated to the synthesis of sine waves for the carrier, nuller and demodulator, it also has a section configured to act as an embedded processor, making the board similar to a small computer. Like a computer, the board has an operating system that runs on its 64 MB of memory, uClinux. uClinux is a reduced version of linux which allows the board to run typical linux programs such as telnet, and even includes a small web server running on-board for setting and recording parameters during testing. It also allows the board to run tuning and testing algorithms stored in its 64 MB of on-board memory, allowing the full parallelization of tuning that will be described in section 4.5.

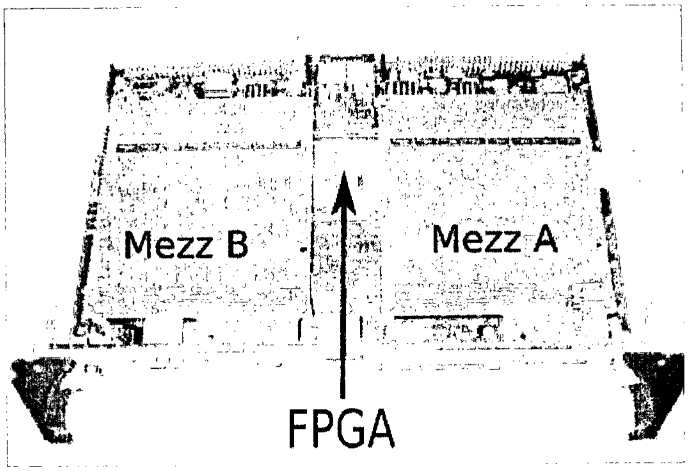

Figure 3–6: A labelled photo of the dfmux board. The Xilinx FPGA that generates the combs, demodulates and like a computer processor can be seen in the middle of the motherboard. The lighter boards to the left and right (called "Mezz A" and "Mezz B") hold the analog electronics and connect to the top of the motherboard below.

In terms of operating the dfmux board, EBEX presents a challenge unseen in ground experiments. While the dfmux board was designed to be low-power, and is nearly an order of magnitude lower power than its analog predecessor, each board (with mezzanines) still dissipates about 16W. At ground level this does not present a problem, as the heat can be convected away by air. However, at balloon altitudes this becomes a much greater concern, as the thinner atmosphere can no longer remove enough heat to keep the boards from overheating. To address this issue a student at Minnesota,Vanessa Cheesbrough, designed the hardware required to conduct this heat away from the dfmux board and onto the gondola, where it can be radiatively transmitted away from the experiment. The testing of this design was undertaken by the author first at Minnesota and later at a NASA long-duration balloon facility in Palestine, Texas. The results of these tests were encouraging, with the board staying within acceptable operating temperatures so long as the test heat sink temperature remained fixed. The details of the tests and results are presented in appendix B for those interested.

With the inclusion of the dfmux board the electronics of the readout are complete. The dfmux board generates the carrier and nulling combs, which bias the bolometers. Each bolometer's resistance varies with the polarization signal from the sky, which produces a current at the input coil of a SQUID. The SQUID amplifies the signal and converts it to voltage which is digitized by an ADC on a mezzanine board. The signal is processed digitally by the FPGA on the dfmux motherboard before being streamed to the data hard disks for analysis post-flight. Of course, the set up of this system is not as simple as the summary makes it sound. Many parameters, including the frequencies to be used in the generation of the carrier and nulling combs and the proper voltages to be supplied to the SQUIDs must be determined before any useful data can be taken. The next section describes how these parameters are determined and used to achieve a full-tuned detector array.

CHAPTER 4
Readout Software

This chapter and the next are the author's main contribution to the project, a suite of algorithms, written in python, that can be used to test, setup and tune the system from initialization to full operation. These algorithms have been written and tested over the last year by the author with much appreciated help from Aubin at McGill and Hubmayr at the University of Minnesota.

The next several sections describe the purpose and structure of the individual algorithms, while the final section describes how these algorithms can be called remotely to run on the dfmux board's processor. The intricate details of how one runs these algorithms, such as the specific arguments they require, will not be included here but are available in a separate user manual for those interested in using the system.

In some algorithm descriptions a plot of the data returned by the algorithm will be shown for illustration. It should be noted that these plots are not necessarily the raw data returned by the algorithm as they may have had conversions applied to convert the raw measurements, in ADC counts, to physical units elsewhere in the signal chain. These plots do not indicate error bars as they are system diagnostics so the error bars would be too small to see.

For reference, a flow-chart visually outlining how the algorithms work together to achieve a fully tuned readout system is presented in figure 4–1 on page 36.

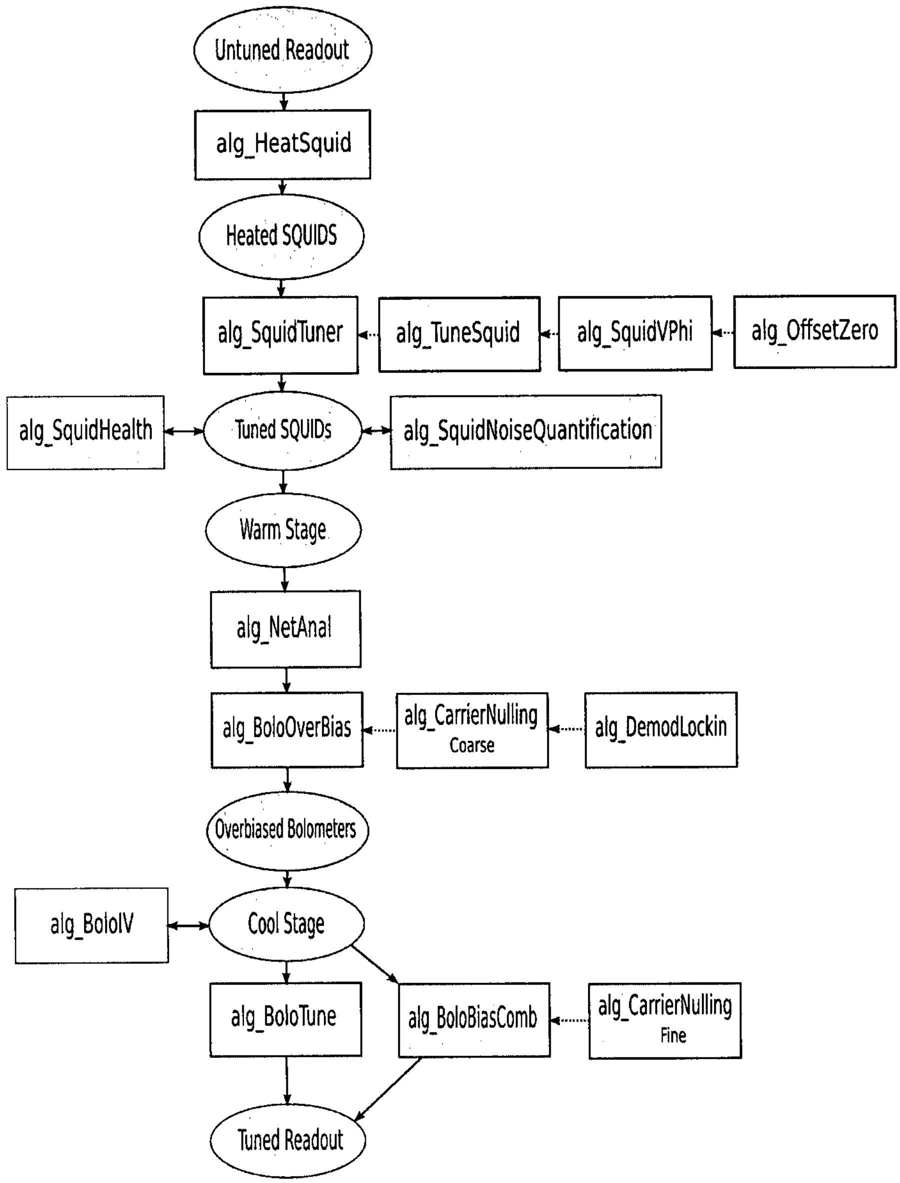

Figure 4–1: A flowchart of algorithm operation. Algorithms are presented in rectangles and system states in ellipses. Striped rectangles indicate algorithms mainly used for quality control. Finally, dashed lines indicate that the output of a given algorithm feeds into another.

4.1 Squid Algorithms

4.1.1 alg_HeatSquids

One of the simplest algorithms, it sets a voltage on the 100Ω resistors near each of the SQUIDs. The heat dissipated in the resistor warms the SQUID above its critical temperature, forcing the SQUID to go normal. After heating the SQUIDs the algorithm turns the voltage off again and waits the requested amount of time, typically 1200 seconds, for them to cool again.

To understand why this is done it is necessary to recall that each "SQUID" is actually an array of 100 individual SQUID coils in series, as discussed in section 3.2. Also, keep in mind that an individual SQUID coil quantizes the magnetic field through its loop, essentially trapping whatever flux was present when it became superconducting [18]. When the cryostat is first cooled different areas cool at different times, depending on the exact nature of their thermal link with the helium bath. This means that different SQUID coils within the same series array can become superconducting at different times, and trap different flux levels. If this is the case each element of the array responds differently to the bolometer signals and the signal-to-noise ratio of the array drops significantly. Thus, it is vital for series-arrays of SQUID coils to cool as simultaneously as possible, so that they will behave identically during operation. We find that heating them just above their critical temperature and allowing them to cool together accomplishes this.

4.1.2 alg_OffsetZero

On the SQUID control board the SQUID output is connected to the positive terminal of an operational amplifier (op amp) which amplifies the signal before passing it on to the rest of the signal chain. This op amp is typically referred to as the 'stage 1' amplifier, as it is the first stage of amplification applied to the SQUID signal. When attempting to tune or characterize SQUIDs the voltage present at this terminal is of interest, as it is a direct probe of the SQUID output voltage. The way this value is measured is somewhat tricky. Rather than probing the SQUID output directly we allow a voltage to be set on the other terminal of the stage 1 amplifier, called V_{offset}. Additionally, there is an ADC on the mezzanine boards which is connected directly to the output of the stage 1 amplifier, the voltage it reads is often called V_{demod} as it is the voltage at the input to the demodulator. By stepping V_{offset} from -8V (the lower rail) to 8V (the upper rail) the curve shown in figure 4–2 can be generated. When V_{offset} is exactly equal to the DC output of the SQUID, the stage 1 amplifier will have an output, V_{demod}, of 0. So, determining the DC output of the SQUID, which we require to tune and characterize SQUIDs, involves finding the V_{offset} required to achieve a zero the stage 1 amplifier.

Tuning a SQUID can require up to a hundred measurements of the SQUID output. If each measurement involved stepping V_{offset} from -8V to 8V this would take a very long time to complete. Instead, alg_OffsetZero attempts to find the zero V_{demod} point in the least number of V_{offset} steps, and thus the least amount of time. It accomplishes this using two methods, a coarse bisection method to find the linear region followed by linear fit routine to get to the required accuracy.

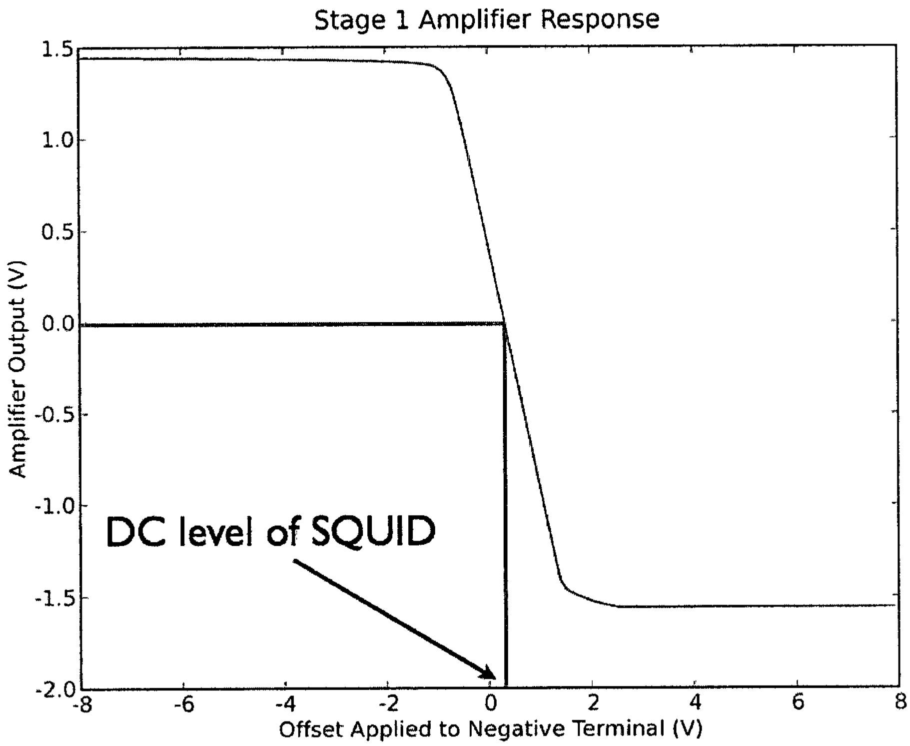

Figure 4–2: The blue curve is the output of the stage 1 amplifier with a SQUID attached to one terminal and the variable voltage V_{offset} applied to the other. The red lines highlight the point on the curve where the output of the first stage amplifier, often called 'V_{demod}' is equal to zero. This implies that both terminals are matched, so the DC level of the SQUID is equal to V_{offset}.

The bisection routine starts by ensuring that there is an offset voltage that gives a V_{demod} reading of zero and will issue an error if no such zero exists between the upper and lower V_{offset} limits. It then iteratively reads V_{demod} in the middle of the range and throws away the half without the zero by checking in which section the V_{demod} of the endpoints changes from positive to negative. By so doing it reduces the V_{offset} range in which the zero could appear by half per measurement. Once the endpoints are not on the rail of the amplifier ($\pm \approx$1.4V) anymore the algorithm moves on to the linear fit routine.

The linear fit routine takes the last point determined by the bisection routine and one other (offset by 0.25V), fits a line and extracts the zero. It will repeat this with smaller offsets until the required tolerance is achieved, at which point V_{offset} and V_{demod} are returned.

Alg_OffsetZero is capable of getting within 0.005 V (the default tolerance) of the true value in 7 iterations, requiring approximately 9 reads from the dfmux and taking only a fraction of a second. This tolerance was chosen to be the equivalent of several least significant bits on the ADC and can not be significantly improved with the current hardware. This algorithm acts as a base for tuning SQUIDs and taking V-Phi measurements, which are described in the next two sections.

4.1.3 alg_SquidVPhi

Alg_SquidVPhi takes one or more response curves for a particular SQUID channel. It relies heavily on alg_OffsetZero in the following way. The algorithm consists of setting a particular current across the SQUID, the SQUID bias, then stepping the flux bias, the current through the inductor coupled to the SQUID, within a user-defined range. At each point alg_OffsetZero is called to determine the DC output of the SQUID. Concisely, alg_SquidVPhi measures the DC output of the SQUID at each flux bias point for each SQUID bias requested by the user. Please see figure 4–3 for a plot of the output of this algorithm, which should help to clarify the description. The curves plotted are called "V-Phi curves," as they plot the output response of the SQUID (V) to changes in the flux through the loop (which is quantized into flux quanta, Phi, for superconductors [18]). These curves are generally used to characterize SQUIDs as well as to ensure the SQUID tuning algorithm has chosen the proper bias parameters.

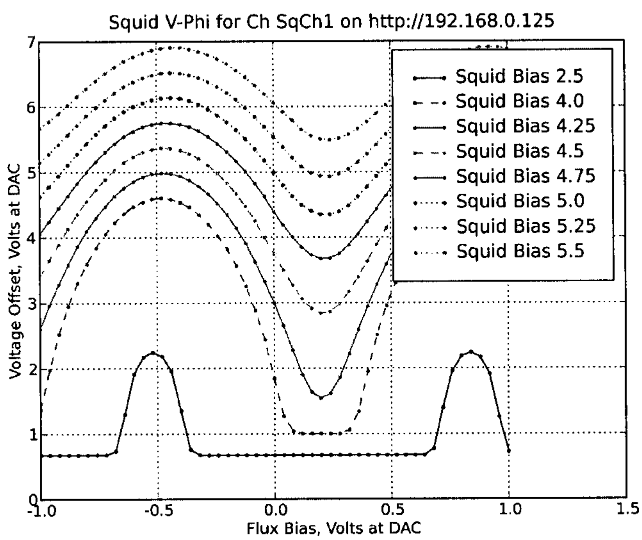

Figure 4–3: A plot of the output of alg_SquidVPhi, also called a V-Phi curves. Each line represents a particular SQUID bias setting in Volts at the SQUID control board. Note that the curve becomes more sinusoidal as the SQUID bias is increased but that the peak-to-peak amplitude (effectively the dynamic range of the SQUID) decreases with higher SQUID bias.

4.1.4 alg_TuneSquid

The SQUID tuning algorithm, it can be used for quality control in a similar fashion to alg_SquidVPhi but is more commonly used to determine the SQUID bias and flux bias combination required for proper SQUID operation. Recall from section 3.2 and figure 4–3 that the SQUID response is sinusoidal, which is not ideal. Feedback is employed to linearize them but the flux bias and current bias at which to connect the feedback resistor and close the loop must be carefully chosen.

To understand the factors involved in the choice of bias point notice that there is a trade-off in figure 4–3 between higher dynamic range at lower SQUID bias currents to a more sinusoidal, and thus more linear (with feedback), response at higher SQUID bias currents. Empirically, the curve with ninety percent of the maximum peak-to-peak amplitude has been shown to be stable and provide strong response to input signals. Alg_TuneSquid finds this curve by running a first SQUID V-Phi at a user-selected SQUID bias reference and using this curve to determine the flux bias of the maximum and minimum SQUID response. As is visible in figure 4–3 each curve has its maxima and minima at more or less the same flux bias, so alg_TuneSquid is able to skip many points on the V-Phi curves for subsequent SQUID biases. Rather than plot the entire curve the algorithm simply takes three points near the locations of the extrema of the first curve and fits these with a second order polynomial. Each set of points provides a single equation and, by solving the system of equations produced, it is relatively straight-forward to calculate the location of the extrema. After running through all the user-requested SQUID biases it uses a linear fit to determine the SQUID bias that has a peak to peak amplitude of ninety percent of the maximum at a current bias higher than the max (over-biased) and sets the SQUID bias to this value. Figure 4–4 is a plot of this process.

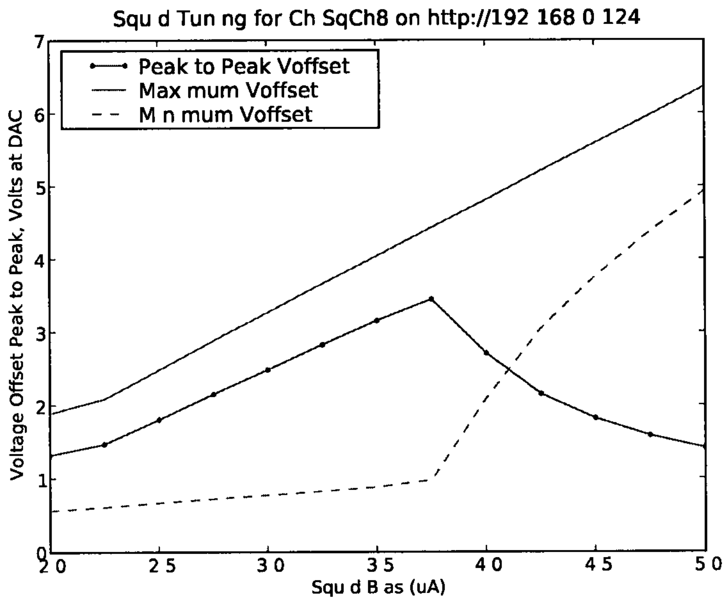

Figure 4–4: A plot of some of the output of alg_TuneSquid. Along the x-axis are the SQUID bias currents chosen for the test. At each of these SQUID bias currents the maximum and minimum of a v-phi are shown as the solid line and dotted line respectively. The line with circles is the difference between these two, the peak-to-peak response of the SQUID at that SQUID bias.

Next, the flux bias of the SQUID needs to be set so that the SQUID will have a strong linear response to input signals and exhibit the negative feedback required to keep it stable. The strong response is achieved by setting a flux bias on the steepest part of the sine curve while the negative feedback condition requires that this be on the downslope of the V-Phi curve. The point that has been empirically determined to work well is the average of, first, the flux bias half way between the maximum and minimum and, second, the flux bias corresponding to the midpoint of the offset voltage, as shown in figure 4–5.

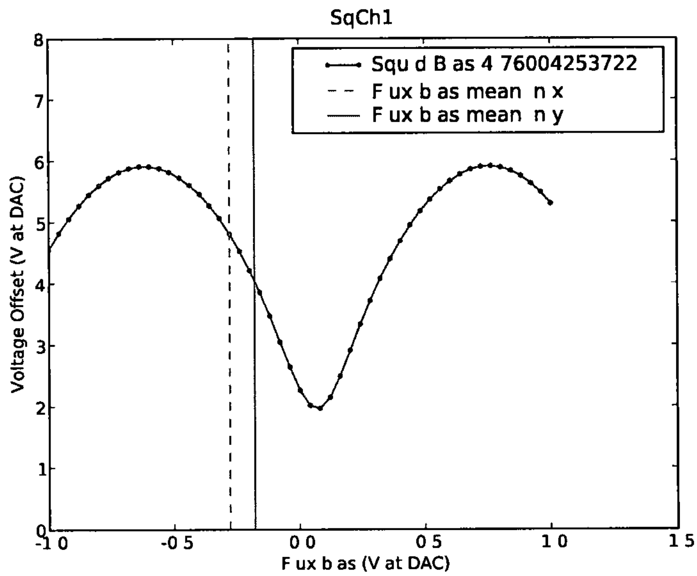

Figure 4–5: A plot of some of the output of alg_TuneSquid. This is the v-phi curve done at the optimal SQUID bias of 90% of the maximum peak-to-peak to select the flux bias. The vertical line marks the flux bias corresponding to the mean in voltage offset, while the vertical dashed line marks the midpoint between flux bias of the maximum and minimum. Alg_TuneSquid averages these to select the operating flux bias that has empirically been shown to be stable.

Once the proper point is determined the feedback resistor is set to the user-selected value of either $10k\Omega$, $5k\Omega$ or $3.3k\Omega$. The default is $5k\Omega$, as this has been found to be stable during our tests. The algorithm also returns data to the user, including the SQUID biases tested, the values of the maxima and minima, and the peak-to-peak voltage for each SQUID bias.

4.1.5 alg_SquidTuner

Alg_SquidTuner is an extension of alg_TuneSquid which will attempt to do a "smart" tuning of the SQUID requiring the least amount of time and user supervision. It will be particularly useful during the EBEX flight, when no "user" is present, and is used extensively by the automated tuning script described in section 4.5.

Alg_SquidTuner has several tuning paths that it can take, depending on what the user supplies and the success of each attempt. If it is provided with full SQUID tuning parameters (i.e. a SQUID bias, a flux bias and the expected DC level of the squid) it will implement these settings and check if $V_{demod} \approx 0$ to see if it has successfully selected a point on the V-Phi curve. If so, it assumes that the tuning has succeeded, it "locks" the SQUID by setting the feedback resistor, and returns a small dictionary with the parameters set and the V_{demod} level detected. This is somewhat dangerous as it is possible that these settings are incorrect but that the point happened to fall on the V-Phi curve by chance. This is unlikely, however, as the response of the stage 1 amplifier is quite sharp (see figure 4–2) so the chances of accidentally hitting a value within the required range are quite small.

In the case that this fails, due to $|V_{demod}|$ being greater than 0.1, or if some of the required parameters are not specified, alg_SquidTuner will resort to running alg_TuneSquid with a single SQUID bias. The peak-to-peak amplitude of this single V-Phi will be checked against a user-defined minimum, which defaults to 2V. If the curve passes this check alg_TuneSquid determines the correct flux bias, the SQUID is tuned and locked, and the data from the tuning is returned.

In the worst case, both of these might fail. In this situation alg_SquidTuner, undeterred, will run a full alg_TuneSquid over the default range of parameters. This will take the most time but ensures a properly tuned SQUID in most situations.

4.1.6 alg_SquidHealth

Alg_SquidHealth checks each of the SQUID channels specified by the user to see if any have had enough flux through the coil to push them out of their linear response region. This is typically called "flux jumping," and ruins the output of the SQUID, which often needs to be retuned before it can be recovered. alg_SquidHealth checks for SQUIDs that have jumped by measuring the output of the stage 1 amplifier (V_{demod}), and comparing it to the jump level for the chosen feedback resistor calculated using the following formula

$$V_{flux\_jump} = 1\phi_o \cdot 26\mu A/\phi_o \cdot R_{feedback} \cdot G_{stage1} * 0.80 \tag{4.1}$$

where ϕ_o is a single quantum of flux through the SQUID coil, $R_{feedback}$ the value of the feedback resistor in ohms, and G_{stage1}, is the gain of the first stage amplifier, which is 5 for the latest revision mezzanine boards. The final term, 0.80, is a safety margin to ensure that all flux jumps are detected, even if the DC output is slightly lower than expected.

If a SQUID channel is found to have jumped, the algorithm will try to fix it by opening and closing the feedback loop a user specified number of times or until the V_{demod} level drops below the jump level once again. Occasionally, opening and closing the loop causes the SQUID to "release" its trapped flux, effectively recovering it without requiring retuning.

The return data states each channel and whether or not it has jumped. If a SQUID channel has jumped it lists whether or not it was fixed. This is intended to be used during operation to find and potentially fix problematic SQUIDs between scans, for instance.

4.1.7 alg_SquidNoiseQuantification

This algorithm does what its name implies, it measures SQUID noise to ensure proper function for quality control purposes. To do this it starts by taking a user-defined number of samples of data with the SQUID turned off (all of its voltages set to 0 and no feedback connected).

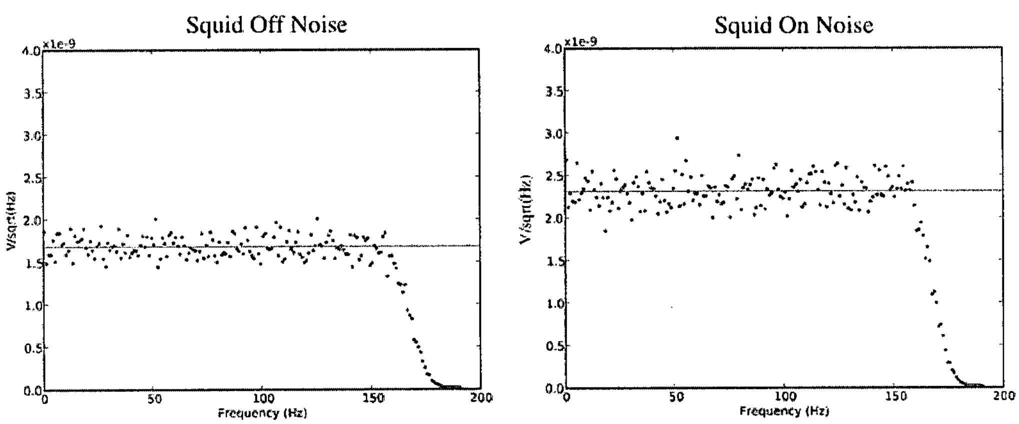

Figure 4–6: Two plots together showing a Fourier transform of the output of alg_SquidNoiseQuantification. The left plot is data taken with the SQUID off, the right with the SQUID tuned using alg_TuneSquid. The horizontal line is the average noise level in the regime before the low pass filter turns on at ~ 150 Hz.

After this baseline is established the algorithm tunes the SQUID, using alg_TuneSquid but does not lock it, so as not to involve feedback in the noise measurement. Alg_SquidNoiseQuantifi then takes another set of data samples, and returns both sets of data for the user to analyze. Figure 4–6 shows plots of Fourier transforms of this data, these plots allow diagnosis of unexpected noise sources as well as possible SQUID problems.

4.2 Multiplexing Algorithms

4.2.1 alg_DemodLockin

This algorithm synchronizes the phase of the demodulator with that of the carrier at the input to the dfmux board. These components must be synchronized so that the mixing of the demodulator square wave with the signal (modulated by the carrier) does not cancel part of that signal. Naively, it appears to be easier to synchronize the phases of the demodulator and carrier when they are produced on the dfmux board but this does not take into account the stray inductances and capacitances present in the cryostat, many of which depend on temperature, that will change the phase of the carrier.

Fortunately, synchronizing after the cryostat is made easy by an underlying dfmux board function, implemented by Graeme Smecher, which allows us to set an extra demodulator channel 90 degrees out of phase with the channel we are trying to measure. This function effectively provides both the real and imaginary (or I and Q) components of the carrier signal from the cryostat.

Alg_DemodLocking determines the phase angle of the carrier relative to the demodulator by taking 32 samples of I and Q using the function described above, averaging them and calculating the angle of the resulting vector. It then changes the demodulator phase by this value and takes one final point that is returned to the user as a measure of how well it worked. Generally this works so well that the remaining imaginary component is no more than several ADC counts in size, less than 1% of the total signal.

4.2.2 alg_CarrierNulling

Recall from section 3.3 that it is necessary for the nulling comb to be exactly the carrier comb except 180^o out of phase. Also note that the nuller does not pass through the bolometers as, if it did, it would negate the bias provided by the carrier. Hence, setting the nuller involves similar phase difficulties to those encountered in the previous section along with the additional challenge of selecting the proper amplitude, both of these issues are addressed by alg_CarrierNulling.

Before describing how alg_CarrierNulling tunes the nuller first note that it contains two possible methods of nulling, "fine" and "coarse." Each of these methods invokes a slightly different algorithm and while the names fine and coarse seem to imply that one is more accurate than the other this is not the case. The difference resides in *how* the nulling is performed, not the quality of the output. Coarse nulling will be described first.

In coarse nulling, the algorithm starts by turning off any nuller that may have been set and getting a measurement of the carrier alone. It averages 32 points, including both I and Q, and stores this as the first point. Next, it sets the amplitude of the nuller to a small value, 10% of the carrier setting, and takes a second average of 32 points, once again including I and Q in the measurement. This measure of the sum of the raw carrier with a small nuller is stored as the second point. By subtracting the real and imaginary parts of the first point from the second the angle and magnitude of this small nuller are determined. Once these values are known it is possible to calculate the adjustments required to properly remove the carrier.

First, the required adjustment to the nuller phase can be determined as the phase of the carrier alone and of the nuller are known. The nuller phase is adjusted from its original phase setting, to one that is 180^o out of phase with the carrier according to the following equation,

$$\Delta\theta_{nuller} = \theta_{pt1} - \theta_{pt2-pt1} + 180 \tag{4.2}$$

where $\Delta\theta_{nuller}$ is the change needed to the phase of the nuller, θ_{pt1} is the phase of the carrier alone and $\theta_{pt2-pt1}$ is the phase of the 10% nuller that connects the first point to the second. The 180 is required to put the nuller out of phase with the carrier. If the value calculated exceeds 360 degrees it is not a problem as the driver used to set the phase reduces the phase to the proper range before setting it.

After the phase adjustment has been performed the needed magnitude adjustment is calculated by finding the ratio of the carrier to the nuller amplitudes and multiplying the nuller amplitude setting by this factor. If this attempts to set the nuller amplitude to less that one percent of the output or greater than is possible at the current gain setting the algorithm checks to see if the user has allowed it to set the wire gain. If so, it will step up or down the coarse gain of the nuller until it finds a solution or eventually raises an exception at the highest or lowest gain. Otherwise it throws an exception back to the user indicating that the nuller amplitude is out of range. A diagram of the process of coarse nulling is presented in figure 4–7.

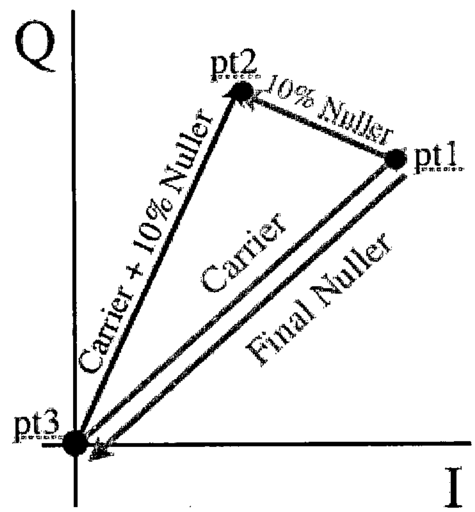

Figure 4–7: A diagram of the operation of the coarse nulling part of alg_CarrierNulling. The first point, pt1, is measured with only the carrier present. Next, a small amplitude nuller is turned on and pt2 is measured. By subtracting these points the length and angle of this small nuller (10% Nuller) can be found and used to calculate the final nuller required to balance the carrier and return to the origin, pt3. All of this is necessary as the system does not know the absolute phase of the carrier or nuller.

Coarse nulling is great for nulling individual bolometer channels, generally while the bolometers are warm, however it is not ideal for cold bolometers as turning off the nuller completely can result in a SQUID becoming saturated and flux jumping. If this happens while the bolometers are cold then it is unlikely that the comb can be recovered without warming the entire stage.

That is why "fine" nulling exists. Fine nulling is unlikely to cause flux jumps as it never turns off the nuller during its operation. Instead, it assumes that the nulling is almost correct and measures the first point without changing any settings. After taking the first point, it changes the nuller amplitude by only 5 percent of its amplitude and takes a second point. With both points taken it does similar calculations to those performed by coarse nulling to achieve a comparable result so long as the nulling was close during the first measurement.

Depending on which mode of nulling is used the return data can vary somewhat. In both cases, however, all of the components in each point are returned, including a point taken before and after the nulling, used to ensure that it worked as expected. The coarse version of the nulling is typically used when the bolometers are warm, as there is no possibility of latching the bolometers. The fine version is more conservative in its changes, so it is often run after the stage is cooled to deal with the slight changes in carrier amplitude and phase that result from the bolometers cooling from $\sim$ 800mK to 250mK.

4.3 Bolometer Algorithms

4.3.1 alg_NetAnal

Alg_NetAnal performs a network analysis on the requested comb, specified by its motherboard, wire and channel. In this context a network analysis entails determining the locations, in frequency space, of the resonant peaks of the RLC circuits for each bolometer. To get reasonable results the bolometers must be normal but their leads must remain superconducting so that the resistance of the leads does not skew the measurement. This is typically achieved by warming the cold stage up to around 800mK.

To measure the resonant peaks this algorithm steps a single component of the carrier comb in frequency in user-defined steps. At each point it sets the demodulator 9 Hz away from the carrier and looks at the average magnitude of 101 data points, including both the real and imaginary components in the calculation of the magnitude.

This first scan is called the "rough scan." After its completion the algorithm goes back to the points where the maxima were measured and does a more precise scan by stepping in, typically smaller, user-defined sub-increments. It takes points exactly the same way as the rough scan, and adds these to the return data.

The data returned by this algorithm contain a list of the frequencies tested and another of the amplitudes measured at each. The peaks of the data can be fitted to determine the exact frequency required to run each bolometer, as is shown in figure 4–8

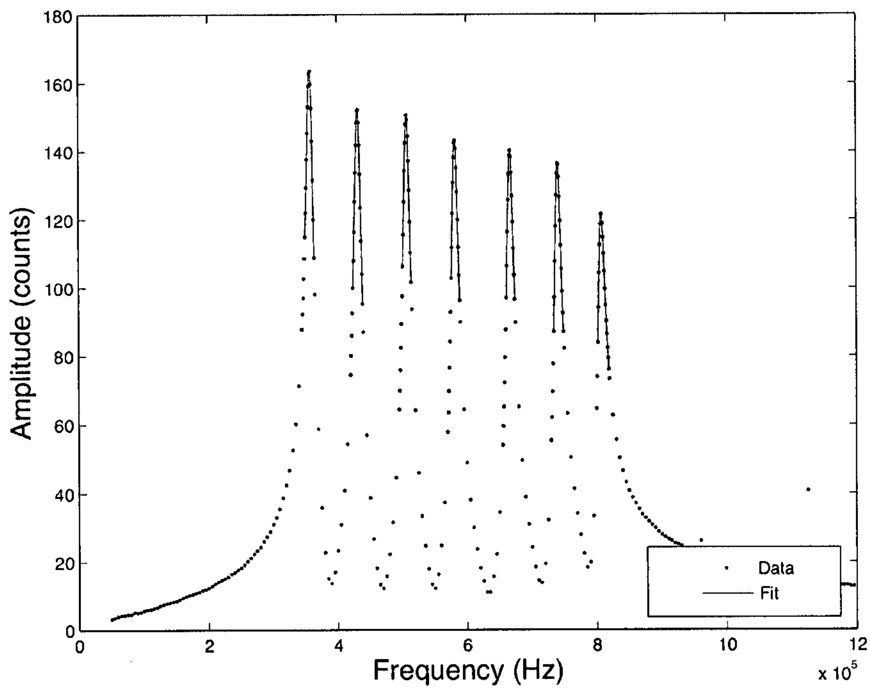

Figure 4–8: A plot of the output of alg_NetAnal overlaid with the fits for the peaks. Each peak is the resonance of a bolometer RLC circuit, the fits, which use the resonance of a perfect LC circuit but also take into account stray inductances and capacitances elsewhere in the readout chain, accurately provide these frequencies for multiplexing.

4.3.2 alg_BoloOverBias

This algorithm is meant to be run while the bolometer is still normal (stage is 800mK) and serves to provide the power required to keep the bolometer normal when the stage is cooled. This is often referred to as "over-biasing" the bolometer. If it were not done, the bolometer would latch into its superconducting state as soon as it were cooled and, since it would have zero resistance, it would be impossible for our system to supply the power required to put it back into its transition.

Over-biasing is accomplished by setting the desired frequency (found using alg_NetAnal), gain and amplitude of the carrier on the requested channel. It then sets the nuller gain to the user-selected value and runs alg_CarrierNulling in coarse mode to null the bolometer channel. Nulling is key as the combined carrier signal from all the bolometers on a given comb is sufficient to flux jump a SQUID.

Its return data includes the parameters that were set and the average of 14 points after nulling, which should be below 50 ADC counts if the nulling was successful.

4.3.3 alg_BoloOverBiasComb

This algorithm is an extension of alg_BoloOverBias. It allows the user to specify a list of frequencies and amplitudes so that they can tune a whole comb at once. It will exit with an error if the list is longer than the number of channels per mux module available on the dfmux board.

For each frequency given it calls alg_BoloOverBias to do the actual setting. If an error is raised by any particular instance of alg_BoloOverBias it does not stop the execution of subsequent channels but the error is returned in the return data from alg_BoloOverBiasComb. The final return data is a combination of the data from each execution of alg_BoloOverBias, with the possibility of some of them listing an error.

This algorithm is intended to be used during an automated tune up just before cooling the bolometers from ~ 800mK to their operating temperature of 250mK. It should be used after the system has been debugged so that individual channel errors are unexpected, as these errors will not be reported until it completes all the channels on that comb.

4.3.4 alg_BoloIV

Performs a bolometer current voltage (I-V) curve for a bolometer. It does this by stepping down the carrier amplitude in user-defined steps and reading the average magnitude of 32 points at each amplitude from the dfmux board. The start and stop points must be given in fraction of the maximum carrier amplitude.

The return data contains a list of the carrier amplitudes tested, the magnitudes measured and the phase of the measured points, which can be used to look for possible nulling problems. The curves provided by alg_BoloIV can be used to ensure the bolometer is functioning properly and determine the carrier amplitude required for drop them into their transition. Figure 4–9 is a plot of the output of alg_BoloIV for a working bolometer.

4.3.5 alg_BoloTune

alg_BoloTune is basically the same as alg_BoloIV except for the units used. Instead of being designed for taking test measurements alg_BoloTune is meant to allow the user to tune the bolometer to whatever fraction of the normal resistance they

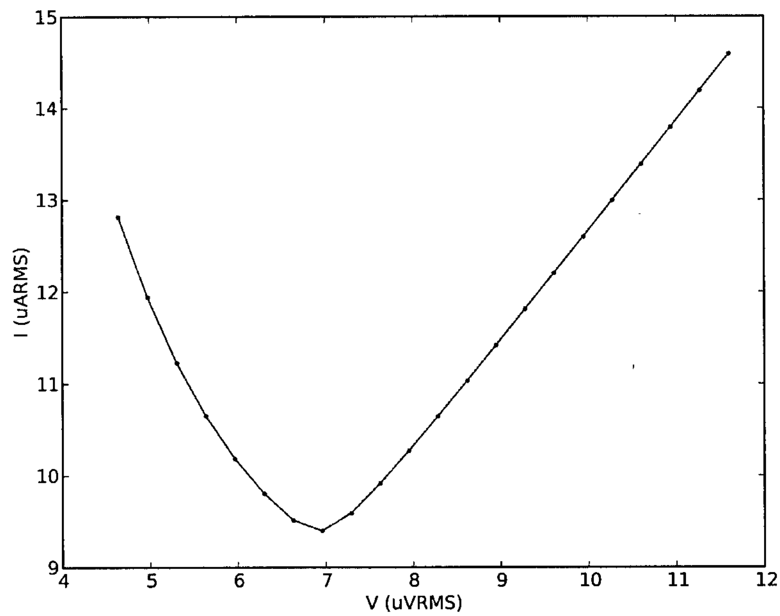

Figure 4–9: A plot of the output of alg_BoloIV for a working bolometer channel. The curve is generated from right to left, with the right part being simply a resistor I-V curve as the bolometer has not yet entered the superconducting transition. As the power on the bolometer is decreased it enters the superconducting transition and the curve flattens out, then turns around as the resistance of the bolometer drops toward 0.

want. It goes about this the same way alg_BoloIV does, by stepping the carrier amplitude progressively down in user-defined steps and reading the average magnitude of 32 points of data. The difference is that alg_BoloTune calculates the fractional resistance at each of these points by taking the ratio of the voltage to current and dividing by the same ratio done while the bolometer was normal (at the start of the algorithm), as shown in the following equation,

$$R_{frac} = \frac{V_{point}/I_{point}}{V_{over-biased}/I_{over-biased}} \tag{4.3}$$

where the subscript "point" refers to the particular point on the I-V curve and "over-biased" refers to the value at the beginning of the algorithm, when the supplied power was still high enough that the bolometer was normal.

The return data is much the same as alg_BoloIV's, though it also includes a list of the fractional resistance at each point. This data can be used to ensure that the bolometer was tuned correctly, after which all the tuning is complete and the bolometer is prepared to take data. A plot of the output data of alg_BoloTune is shown in figure 4–10 on page 59.

4.3.6 alg_BoloBiasComb

alg_BoloBiasComb is a short algorithm to bias all the bolometers in a comb into their transitions. It assumes that the frequency for each bolometer channel has already been set when they were being over-biased (see section 4.3.2). It also requires that alg_BoloIV or alg_BoloTune have already be run on the desired wire and the carrier amplitudes required to drop the bolometers into their transition have been stored. For each channel in the user-supplied wire it sets the carrier amplitude to the

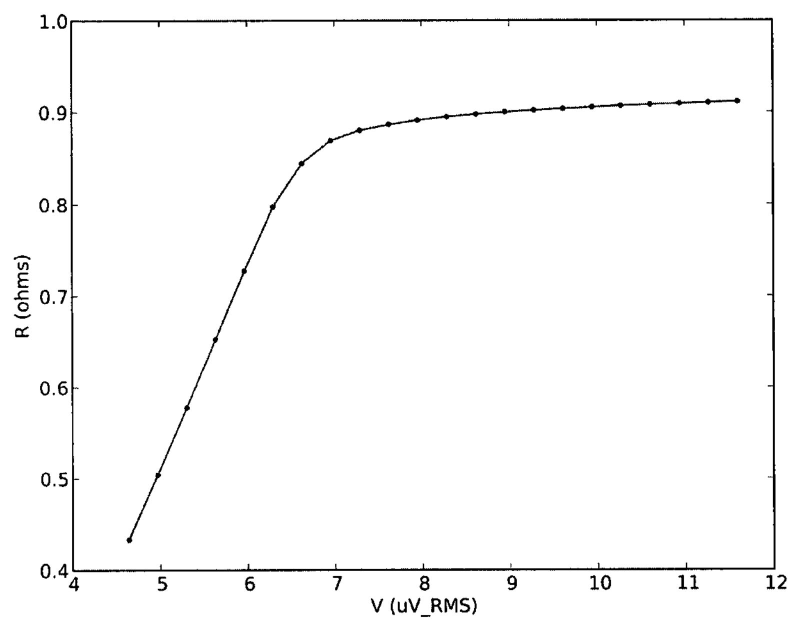

Figure 4–10: A plot of the output of alg_BoloTune. As the bias voltage is lowered the bolometer enters the superconducting transition and its resistance drops.

corresponding user-supplied amplitude and then runs a fine nulling on that channel. If the amplitude in the list is 0 that channel is skipped.

Like alg_BoloOverBiasComb, if any particular channel throws an exception it does not stop the algorithm but rather the error is added to the return dictionary and the following channel is attempted. In the end the algorithm returns the nulling data for any of the channels that worked, 'skipped' for any that were skipped (by setting the amplitude to 0), and the error text for any that fail.

This is intended to be the last algorithm run during a tune up, biasing the bolometers into their transitions, ready to begin measuring the polarization on the sky.

4.4 Algorithm Management

Since the dfmux board has its own embedded processor the algorithms described in the last section can be run directly on the board, allowing considerable parallelization and reducing the total tune up time of the experiment to the tune up time of a single dfmux board. This necessitates a manager to accept user-requests for algorithms to be run and to keep track of those waiting to be run and already finished. This is accomplished by two separate pieces of code, dae_AlgManager.py (AlgManager) and dae_AlgProcess.py (AlgProcess). While a detailed explanation of the operation of these two python scripts is not necessary here, an outline will be provided.

At any given time the dfmux board is running one AlgManager process and up to four AlgProcess processes. The AlgManager acts first as a web server, as all user communication with the dfmux board is done over ethernet. It listens on port

2727 (by default) and awaits algorithm requests. Those requests are expected in Javascript object notation (JSON), an expandable, well documented format that is supported in most popular programming languages. After decoding the algorithm request it determines whether an AlgProcess is idle. If so, AlgManager stores the requested algorithm and arguments in a file and signals AlgProcess to start running it. If not, the request is added to a queue, stored in memory, until an AlgProcess becomes available.

Once signaled, and AlgProcess retrieves the arguments and runs the algorithm. Once it completes the current request it saves the data to file and signals AlgManager that it has completed this algorithm before looping back to the beginning and awaiting further requests. In this way the AlgProcesses do all of the computation while the AlgManager is always prepared for new user requests.

The return of the algorithm's data is dependent upon user-request. As AlgProcess runs an individual algorithm it stores the data in a file on the dfmux board. When a user makes a specific type of request, a "check" request for a given algorithm, AlgManager reads the contents of this file and sends it over ethernet. The file on the dfmux board is deleted during this process to conserve the limited on-board memory, so the user must be careful to store the data requested. If the AlgProcess has not completed the requested algorithm a small JSON string indicating this fact is returned to the user instead.

While the outline of these two pieces of code is fairly simple, the synchronization and interplay between them is actually quite complicated. In addition, both contain considerable error handling as a great deal of data would be lost of they were to crash during the long duration flight. How these processes will be used during flight is outlined in the following section, and some recent tests involving their use are presented in section 5.4.

4.5 Automated Tuning

During the actual EBEX flight it will not be possible to send more than a trickle of data to the system, so the flight control computer will need to act as the "user" mentioned above, making algorithm requests and interpreting the results. Since all of the detectors will have previously been tested on the ground it pays to make use of what was learned in those tests. To this end a "parameter file" has been designed to store all of the necessary parameters of a particular detector tune up. This data is, like the algorithm requests, stored in JSON format, in this case on the hard disk of the flight control computer. It will include such things as the SQUID bias for every channel, the carrier amplitudes for over-biasing, and the SQUID to dfmux motherboard connection mapping.

Implementation of a program to read these parameters and send commands to the on board algorithm manager to tune the system for the actual flight is ongoing at the University of Minnesota. Due to the nature of the flight control program (fcp), it will have to be written in C directly into the main flow of the system. Here at McGill, a similar program is being developed to act as an example and permit earlier testing of the routine. Ours is written in python, like the algorithms, and is called simply. "tune_readout.py." Here I will describe the steps that it performs to illustrate how a tuned system can be achieved using the parameter file. Note that each step in tune_readout.py is sent through the algorithm manager so adding additional dfmux boards to the setup does not significantly increase the tune up time.

The scripts starts by tuning the SQUIDs as there can be no useful communication with the bolometers until the SQUIDs are properly tuned and locked, using feedback. It uses alg_SquidTuner exclusively to do this. For each SQUID board it loops over requests to the AlgManager for each SQUID channel and then waits for each to finish. The data returned by each is printed to the screen for testing purposes and saved to a log file on the computer hard drive for later analysis.

Upon completion of the SQUID tuning it begins to set up the bolometers. It starts by waiting for the user to warm the bolometers to $\sim$ 800mK. In EBEX the flight control computer will be able to trigger and control this process but in our test setup it requires user intervention. Once the stage is warm the automated tuning algorithm over-biases the bolometers using alg_BoloOverBiasComb. The parameter file contains all of the frequencies and over-bias amplitudes so this is a fairly quick process, requiring only a little time for alg_CarrierNulling to complete on each bolometer channel. The tuning script again waits for each algorithm to complete and both prints and stores the data returned in each case.

Finally, the stage is cooled to 250mK with the bolometers remaining normal due to the over-bias power being deposited on them. Again, this will be performed by the flight control computer on EBEX, but is currently done by a user in our test setup. With the stage cooled, the script is prepared to complete the tuning by biasing the bolometers into their transition using alg_BoloBiasComb. Once again, the correct transition amplitudes must be stored in the parameter file from a previous tune up so it does not take long to drop each bolometer into its transition and re-null using the fine nulling method of alg_CarrierNulling. The data from each algorithm is again printed to the screen and stored, as it was during the previous two steps. As always, once the bolometers are dropped into their transitions they are ready to take data.

Of course, tuning the readout is not quite so simple as it sounds here. There are many problems that can occur during each step that will need to be accounted for. The following chapter outlines some test measurements performed without and with the automated tuning script and should provide an idea of some of these challenges and what we plan to do to surmount them.

CHAPTER 5
Measurements

Designing algorithms for tuning bolometers is all well and good but often they function better on paper (or computer screen) than on real hardware. Like any software project, testing is key. Until April 2008 testing was performed mostly by Johannes Hubmayr at the University of Minnesota as he had a cryostat and a set of test bolometers on which to run the code. This method was hardly ideal as it took a considerable amount of time for new code to get tested and corrected.

In April we commissioned our own test cryostat with bolometers intended for use in POLARBEAR [19], a ground based CMB polarization experiment that also intends to take data in 2010. Having our own dedicated test bed improved debugging of the code significantly as code could be written and tested as soon as it was completed by the person who had written it. This chapter starts by describing our test cryostat before moving on to the results obtained using the algorithms described in the previous section.

5.1 Test Setup

Over the last year the commissioning of a test dewar has proceeded here at McGill to permit the testing of all stages of the readout, as well to determine the properties of bolometers destined for the experiments we collaborate on. The test setup consists of a custom three stage (300K, 50K, and 4K) cryostat, fabricated by Precision Cyrogenics (Indianapolis, IN). This cryostat is cooled by a mechanical pulse tube cooler (PTC) fabricated by CyroMech (Syracuse, NY). The PTC provides over a Watt of cooling power, sufficient to cool the main plate of the dewar to about 3.5K without expending any helium, just electrical power. The main plate temperature achieved by the PTC is more than adequate for the SQUIDs, which become superconducting around 9K, but is not cold enough for the bolometers to operate, as they become superconducting around 550mK.

The additional cooling required by the bolometers is provided by a three stage Chase Research Cryogenics (Sheffield, UK) adsorption fridge. This closed system fridge contains one helium-4 bath and two helium-3 baths each thermally sunk to the last. At each stage a carbon-coated surface pumps on the bath to lower its boiling point resulting in the "ultra" stage, to which the bolometers are thermally connected, being cooled to 250mK. The helium gas at each stage adsorbs onto the large surface area presented by the carbon coated surfaces and, since each vessel is entirely enclosed, no helium is lost.

All of the bolometer wedges tested in our test setup were fabricated by Berkeley [20] and most are candidate detectors for the South Pole Telescope (SPT) [21], with only the first sample being a test wafer for POLARBEAR. This is necessary as the process to fabricate EBEX bolometer wedges is still being perfected so EBEX wedges are sent straight to the University of Minnesota to be tested in the EBEX cryostat after the shortest possible delay. The EBEX bolometers are also fabricated at Berkeley and are very similar to those used in SPT, so substituting one for the other makes little difference to tests of the readout system function. One minor difference to keep in mind during the following sections is that SPT bolometer wedges have 7 bolometers per comb while EBEX will have 8 for the North American test flight and 12 for the long-duration flight.

5.2 Tuning Outline

During testing the PTC is run continually so that the main plate of the dewar is always just below 4K. Our adsorption fridge is cycled once a day, at midnight, so that the bolometers are superconducting in the morning, sitting at 250mK. If we want to perform bolometer or SQUID tests we start by dumping a small amount of power onto the adsorption fridge's cold stage, warming the bolometers to their non-superconducting (normal) state. This is usually accomplished by heating the stage to around 800mK. With the bolometers normal it is possible to tune the SQUIDs, as super-currents through the bolometers that otherwise would have been present at the inductor coil are no longer a concern. It is also possible to over-bias the bolometers, as while they are normal they have non-zero resistance. This is achieved using alg_BoloOverBias (section 4.3.2), which ensures that sufficient power is present at the bolometer that it will stay normal even when the stage is cooled back down.

If we want to perform bolometer measurements we then cool the stage by turning off the heat load we set at the beginning. It typically takes around 25 minutes to cool the stage and bolometers back to 250mK. Recall that, even with the stage at 250mK the bolometers are still normal as there is enough power provided by the over-bias voltage to keep them from entering their transition. The final step in tuning the bolometers involves stepping down each bolometer's bias until the bolometer sits within the superconducting transition. This is often performed by alg_BoloTune (section 4.3.5), though it can also be done by alg_BoloBiasComb (section 4.3.6) if the required bias settings are known from a previous tuning. Once the bolometers have been lowered into their transitions and nulled, tuning is complete. What has been described so far is the general method used to tune bolometers, similar to the one described in section 4.5 with some details skipped for brevity. It pays to keep this scheme in mind for the sections that follow.

In addition to the outline of typical measurements it should also be noted that there are two ways to run any of the tuning algorithms mentioned above and described in detail in chapter 4. The first is to run the algorithms directly from a computer connected to the same network as the dfmux boards. In this method the computer performs all of the calculations involved in the algorithm and the board deals only with specific commands such as setting a SQUID bias voltage or the frequency of the carrier signal for a particular comb. In this approach, which will henceforth be called "remote testing," each algorithm is run sequentially. This is ideal for testing individual components as the failure of any given element can immediately be addressed and the data from each step of the tune up is immediately available to the user. Unfortunately, it is also the slower method as it does not use the processing power of each individual dfmux board.

To tune the array quickly, the second method, running the algorithms on the dfmux board itself, is preferable. In this method the algorithm manager described in section 4.4 is used to split up the algorithm processing onto the dfmux boards themselves. Since each board deals with its four combs simultaneously with every other board, this method involves running the algorithms in parallel, drastically reducing the time required to tune a large array of bolometers. This method is more challenging as there are many algorithms being run simultaneously and the user can only detect problems with an algorithm after it has completed.

The next section describes tests done using the first, "remote testing," method. These tests were mostly performed by Aubin at McGill, using the algorithms described in chapter 4. Meanwhile, the author assisted by correcting bugs as they turned up and simultaneously developed and tested the automated tuning method required for the EBEX flight that is described in section 4.5 with test results presented in section 5.4.

5.3 Remote Testing

5.3.1 SQUID Testing

Four SQUID boards have been tested in the McGill cryostat since it was completed, each containing 8 SQUIDs. For each SQUID, V-Phi curves have been performed for the SQUID bias currents in the typical range using alg_SquidVPhi. The output of this algorithm on one SQUID has already been shown in figure 4–3. NIST, who provided the SQUIDs, guarantees their peak-to-peak response to be greater than 3V peak-to-peak when properly cooled and biased. So far, this agrees with what has been seen for all of the 32 SQUIDs tested.

After getting acceptable V-Phi curves, noise measurements were taken for all of these channels using alg_SquidNoiseQuantification. Once again, the output of the algorithm has been seen in figure 4–6 and will not be repeated here. To make a meaningful comparison of the raw data, taken in ADC counts at the demodulator, with specifications and other experiments it must first be converted to voltage at the output of the SQUID. This conversion is performed using the following equation,

$$V_{squid} = C_{demod} \cdot \left(\frac{2}{16384}\right) \cdot T \cdot \left(\frac{100+Y}{10000}\right) \cdot \left(\frac{1}{G_{digitalfilter}}\right) \qquad (5.1)$$

where the first term, C_{demod}, is the measured ADC counts at the demodulator and the second term is the conversion between ADC and volts for this 14-bit ADC. The next term, T, is the transfer function for the analog electronics on the mezzanine board and squid controller board, which is 7.18×10^{-4} for our system. The fourth term, $\left(\frac{100+Y}{10000}\right)$ represents the effect of the user-selectable demodulator gain. The user can select the demodulator gain of $(0, 1, 2, 3)$ corresponding to a value of $Y = (10000, 2000, 200, 0)$. For SQUID noise measurements the demodulator gain is set to 3, so that this term becomes 1×10^{-2}. The final term involves the gain of the digital filters, $G_{digital_filter}$, applied to the data received at the demodulator. The gains of each of the finite input response (FIR) filters is set in firmware, and will not be discussed here. Alg_SquidNoiseQuantification uses the output of FIR stage 5, meaning the input passes through, first, a square wave mixer, next a cascaded integrator-comb filter, and then five separate FIR filters. Cumulatively these produce a $G_{digitalfilter} = 1.6678$. Putting these typical values into equation 5.1 reduces it to the following,

$$V_{squid} = 5.26 \times 10^{-10} (\text{V/ADC Counts}) \cdot C_{demod} \qquad (5.2)$$

This conversion is applied to the output of alg_SquidNoiseQuantization to convert it to volts at the output of the SQUID. This data is then used to generate tables such as table 5–1, which act as a reference for each of our SQUIDs. So far the noise for all of our SQUIDs, referred to the output of the SQUID, has been found to be about $1.5 \pm 0.3 \frac{nV}{\sqrt{Hz}}$, which agrees with specifications and other experiments.

Channel	Noise SQUID off $\frac{nV}{\sqrt{Hz}}$	Noise SQUID tuned $\frac{nV}{\sqrt{Hz}}$	SQUID noise $\frac{nV}{\sqrt{Hz}}$
1	1.69 ± 0.12	2.31 ± 0.18	1.6 ± 0.3
2	1.73 ± 0.14	2.27 ± 0.18	1.5 ± 0.3
3	1.71 ± 0.14	2.44 ± 0.19	1.7 ± 0.3
4	1.76 ± 0.14	2.30 ± 0.19	1.5 ± 0.3
5	1.72 ± 0.14	2.25 ± 0.19	1.5 ± 0.3
6	1.70 ± 0.13	2.3 ± 0.7	1.5 ± 1.0
7	1.68 ± 0.13	2.2 ± 0.16	1.4 ± 0.3
8	1.73 ± 0.137	2.3 ± 0.4	1.5 ± 0.7

Table 5–1: Example of noise table constructed from the output of alg_SquidNoiseQuantification using equation 5.2

5.3.2 Bolometer Testing

Following successful SQUID tests, measurements of bolometers in this cryostat began in April 2008 with a sample of POLARBEAR bolometers. Tests on SPT wedges with many more bolometers (around 140 per wedge) followed starting in May 2008. At the moment the testing of the third SPT bolometer wedge is under way.

Each bolometer test begins with a net analysis, using alg_NetAnal, to determine the resonant frequencies of the LC circuits preceding each bolometer, as it is impossible to over-bias or bias the bolometers without knowing the proper frequencies. The final output of the net analysis is a plot like figure 4–8 which can be used to generate a list of frequencies for a given bolometer comb which looks something like table 5–2. Each frequency listed is determined by fitting a peak in the data from alg_NetAnal. The resistance is calculated by looking at the width of the peak, since the bandwidth of an RLC circuit is directly related to the resistance.

Peak ID	R(Ω)	Amplitude (counts)	frequency (Hz)
1	0.94	209	410735
2	0.95	198	490605
3	0.91	175	565943
4	0.87	188	794835
5	0.89	150	868043

Table 5–2: Table of bolometer parameters produced from the output alg_NetAnal. The second row is the resistance of the bolometer, estimated by looking at the width of the peak, the third is the amplitude of that peak and the fourth is the frequency.

Finally, I-V curves are produced, using alg_BoloIV or alg_BoloTune, for each bolometer. Recall from chapter 4 that these curves involve slowly decreasing the power on the bolometer and allowing it to enter into its superconducting transition, as shown in figure 4–9. These curves provide a wealth of information about the bolometer, including its thermal conductance. The thermal conductance of a bolometer must be tuned carefully as a higher thermal conductance will increase the bandwidth of the experiment, while a lower thermal conductance lowers the overall noise. Unfortunately the techniques used to produce bolometer arrays cannot always obtain the desired thermal conductance so each wedge must be tested to ensure the bolometers are suitable for the experiment.

I-V curves also provide vital information for the automated tuning tests described in the next section. In particular, the bias point for each bolometer can be determined from their I-V curve and used by the automated tuning to quickly set up that bolometer.

5.4 Automated Tuning Tests

As described in section 5.2, the second way to run the tuning algorithms involves running them directly on the dfmux motherboards themselves. This method involves more separation between the user and the system, as the output from a given algorithm is only returned upon algorithm completion, but it can make use of parallelization to achieve a fully tuned system in the least amount of time. In addition, during the EBEX flight it will be difficult to transmit data to the experiment so an automatic tuning routine is necessary to take the place of the "user" in the remote testing procedure. The development and testing of this automated tuning routine has been the author's focus for the last several months.

The individual steps involved in automated tuning have been outlined in section 4.5 and described briefly at the beginning of this chapter. As such, they will not be repeated here. This section will, instead, describe recent tests of the automated system on seven combs of an SPT wedge in the McGill test cryostat.

These combs had earlier been tested remotely, so it was possible to construct a parameter file describing them, much as it will be for the EBEX bolometers before the flight. With the parameter files constructed the automated tuning script, tune_readout.py, was run. It made use of two dfmux boards, one with four channels, the other with three, to perform the algorithms and determined the time required to perform each step. These times are listed in table 5–3.

Tuning Step	Software Time (s)	Thermal Time (s)
Heating SQUIDs	30	1200
Tuning SQUIDs	24	0
Overbias and Null Bolometers	142	0
Cool Bolometers to 250mK	0	1500
Bias and Null Bolometers	76	0
Totals	272	2700

Table 5–3: Times taken for each step of automated tuning, the second column is the time taken for the software to operate, while the third is the time required for components to heat or cool. Note that the first and third step involve waiting for thermal time constants, these cannot be improved in software but may be different from these test cryostat measurements in EBEX. Also, the SQUID tuning time is for the ideal case of SQUID parameters that do not change from day to day. If the SQUID tuning parameters do change during flight this time will increase to about 141 seconds.

There are several notes about the times taken. First, note that time was saved heating of the bolometers from 250mK, after an adsorption fridge cycle, to 800mK, required for over-biasing, as this was done while the heating of the SQUIDs (step 1 in table 5–3) was performed.

Second, note that the amount of time these principally thermal steps, such as warming and cooling the bolometers, require depends on the precise hardware used. In our test cryostat, bolometer heating takes about 15 minutes while bolometer cooling takes about 25, while in EBEX these times have yet to be tested.

Third, the SQUID tuning time listed in the table is an ideal case in which all the properties of the SQUIDs do not change with time. In this case all that is required to tune the SQUID is to set the SQUID bias, flux bias and V_{offset} to their previous values. This may not be the case during the EBEX long-duration flight since, as the experiment circumnavigates the Earth the magnetic field passing through the SQUID will change, likely causing a single V-Phi to be required at the previously determined SQUID bias level. If this is the case, this tuning step will increase to around 141 seconds and will automatically be detected and dealt with by alg_SquidTuner.

Finally, remember that, due to parallelization, this essentially represents the *entire* time required to tune the readout system, no matter how many additional dfmux boards are added. That is a slight exaggeration, as it might take a few more seconds to send commands out to a much greater number of boards, but in any reasonable configuration it should take no more than 30 additional seconds to send the extra requests over ethernet.

At this point a great deal has been said about the time taken to tune the 7 tested combs but nothing has been said about the yield. In fact, the yield during these tests was good but not quite perfect. Recall that these 7 combs are from SPT, and thus could have a maximum of 49 individual detectors, with 7 per comb. In fact, previous tests with the remote system indicated that 47 of these actually functioned. Out of these 47 every channel tuned successfully during the automated tune up except one.

The one channel that failed showed a problem during nulling. This particular bolometer needed a nuller amplitude very near the limit of the nuller gain setting it had previously set. The system does not currently have the intelligence to step up the gain of this wire and correct this problem so it reported an error after running. Methods to correct this error for future tests have been devised, the simplest involving storing the required nuller gain for a wire in the parameter file as it has not been seen to change much from one tune up to the next; however, these changes may not be completed in time to include them in this document.

Even without this correction, the automated tuning had a success rate of nearly 98% on this small sample. The next hurdle for the automatic tuning system will be running on a full EBEX bolometer wedge (140 bolometers) during the upcoming integration of the bolometer camera with the telescope in November 2008. We are confident that, with fairly minor tweaks and corrections, the system will be able to tune this wedge with a high success rate.

CHAPTER 6
Summary & Conclusions

Detecting B Mode polarization has been called the "critical next step in extending our knowledge of both the early Universe and fundamental physics at the highest energies" [22] and was highlighted by the National Research Council in their 2003 report 'Connecting Quarks with the Cosmos' [23].

The search for this signal is ongoing with many projects hailing from many collaborations. At the moment the main competitor to the dfmux system is the time domain multiplexing system (TDS) developed by NIST and the University of British Columbia [24]. The TDS uses SQUIDs as switches to selectively turn on each bolometer pixel, thus separating them in the time domain akin to how the dfmux system separates them in frequency. At the moment the TDS has been shown to achieve multiplexing factors of 32:1, while the dfmux has shown multiplexing factors of 8:1. There is no fundamental reason that the multiplexing factor of the dfmux cannot be improved drastically with few modifications to the existing system. In fact, one of the future development paths involves testing new firmware for the dfmux FPGA that increases the multiplexing factor from 8:1 to 16:1.

The other main development path for the dfmux concerns noise. In particular, tests of bolometer noise using the dfmux board need to be performed and the output analyzed to ensure that readout noise will not cause a problem during operation.

Some preliminary noise tests have been performed to date but more extensive analysis is required to draw conclusions from the available data.

As for EBEX, its construction is nearly complete, with the first complete integration of the experiment scheduled for November 2008. The integration will bring together parts from the many universities collaborating in its production, including the readout system that has been designed and tested at McGill. Fortunately for us, all of the required tuning algorithms have been written and tested on our test system and it appears that only minor issues remain. We look forward to completing the debugging, integrating our equipment into EBEX and putting our system through rigorous testing during the 2009 test flight and through the ultimate test of the 2010 long-duration flight over Antarctica.

APPENDIX A
Polarization Modes

During the introductory chapter of this thesis the three distinct sources of CMB polarization are mentioned briefly without delving into detail about what those modes are and how each produces a different polarization pattern. This appendix addresses those omissions to a certain degree, though a full treatment is beyond the scope of this document but can be found in [1], for example.

A.1 Scalar Modes

Of the three sources of CMB polarization probably the most familiar are the scalar modes. These modes represent fluctuations in the energy density of the primordial plasma at last scattering that are thought to have arisen from quantum fluctuations in the very early universe.

To understand what kind of polarization pattern is generated by scalar waves it helps to consider the simplest case, a plane wave with an electron at the trough. The higher effective temperature (where the effective temperature takes into account the gravitational potential) leads to flows from the peaks toward the trough. Thus, the electron sees a quadrupolar pattern with cold regions perpendicular to the propagation direction of the wave and hot ones along it. What's more, the azimuthal symmetry requires that the matter flows are parallel to the disturbance, hence the flow is irrotational. The quadrupole pattern seen by our electron can be described in terms of its spherical multipole as having $l = 2$ (as for all quadrupoles) and $m = 0$.

This pattern will be referred to as Y_2^0, and can be shown to generate strictly a curl-free, E-mode polarization pattern. Here "E" is making an analogy to the curl-free nature of the electric field.

E-Mode polarization has been detected but not precisely characterized. The first experiment to make a detection was the Degree Angular Scale Interferometer (DASI) [10]. It made a low signal-to-noise measurement of the presence of polarization in 2002 and has been upgraded to measure its angular spectrum more recently.

A.2 Vector Modes

The next most obvious source of CMB polarization would likely be the vector modes. Vector perturbations represent vortical motions of the matter where its divergence vanishes but it's curl is non zero.

To again consider the simplified plane wave perturbation, the velocity field is now perpendicular to the wave vector, with its direction reversing in peaks and troughs. The radiation field at either side exhibits a dipole pattern due to the Doppler effect of the bulk motion. In this case the quadrupolar variations disappear at the peaks and troughs but are maximized between. This is fairly easy to imagine if we place our electron in the middle of a wave and "look" toward the crest a dipole will be evident, while "looking" the opposite direction shows a dipole in the opposite orientation. Thus, from the point in the middle there is a quadrupolar variation in temperature with $m = \pm 1$. As for the overall polarization signal, vector modes generate strictly B-mode (non-zero curl) polarization but the original progenitor of these perturbations has no associated density perturbation so this mode is expected to become very small during the expansion of the universe [5]. So far measurements have shown this to be the case, as only E mode polarization has been detected, due to its higher intensity [10].

A.3 Tensor Modes

The final source of CMB polarization are tensor fluctuations on the surface of last scattering. Tensor fluctuations are really perturbations to the metric, essentially gravitational waves. A plane gravitational wave perturbation will "stretch" space in the plane of the perturbation along one direction at the peak and a perpendicular direction at the trough. This stretching of space essentially lengthens the wavelengths of the photons leading to a quadrupolar variation with an $m = \pm 2$ pattern.

As in the other two cases, Thomson scattering produces a polarization pattern from this quadrupolar variation in temperature. In this case Q and U are generated in almost equal proportions and it can be shown that gravitational waves can generate both E and it's complementary B fields [1]. Unfortunately, due to cosmic variance it is impossible to determine precisely how much of the E mode spectrum that has already been measured is due to tensor perturbations [13]. Hence, the direct detection of the B-mode component is required to confirm the presence of large scale gravity waves at last scattering.

APPENDIX B
Dfmux Thermal Tests

Each dfmux board uses about 16W of power, which is dissipated as heat. As it is raised to high altitudes during the EBEX flight it will have much less atmosphere to convect that heat away. As such, a student at the University of Minnesota, Vanessa Cheesbrough, has designed hardware to conduct the heat away from the board and into the gondola, where it can radiatively transmitted down to the Earth or out to space.

She has accomplished this using two copper "nanospreaders" thermally epoxied to the board. One runs laterally across under the two mezzanines and across the top of the aluminum cage which surrounds the digital electronics on the motherboard. The other runs from atop the FPGA straight toward the backplane, between the two VME power connectors on the back of the dfmux board. A photo makes this much more evident, please see figure B–1. These nanospreaders are proprietary hardware so their exact composition is not known; however they appear to be a form of heat pipe, which is a thermal conductor in which a liquid is able to convect the heat way rather than relying entirely on conduction. The fact that the heat transfers by convection as opposed to conduction serves to increase the conductance by an order of magnitude as compared to copper, which is a common heat sink material.

In this design all of the heat from the mezzanine boards must first pass to the crossing point of the nanospreaders before it can be transferred out to the backplane

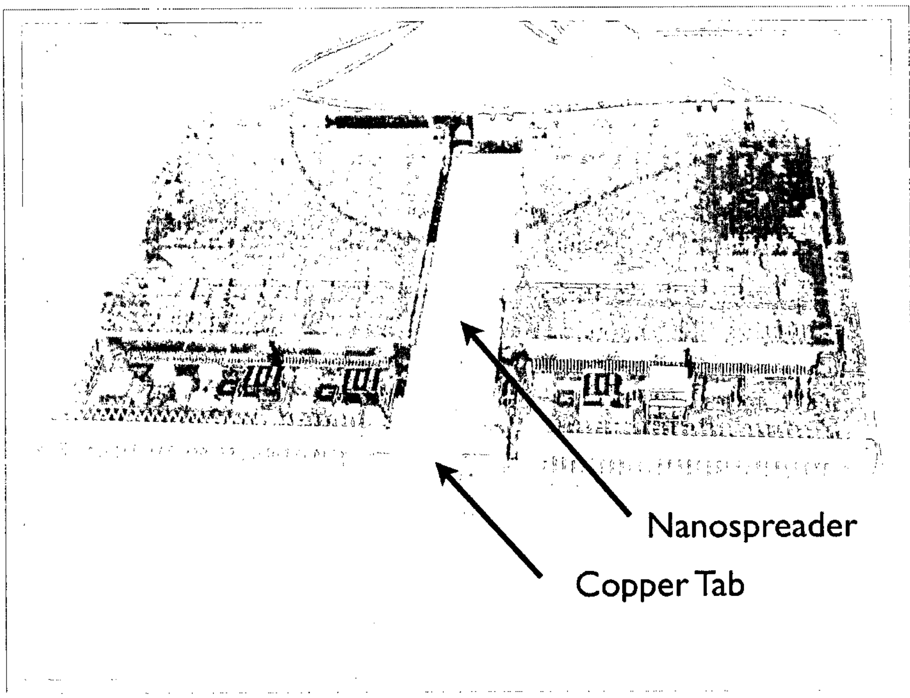

Figure B–1: A photo of the dfmux board with heatsinks attached. Note that there is one, barely visible, running across under the mezzanine boards. The copper tab in the fore fits into a slot in the back of the readout crate.

of the board. Here it arrives at a large copper tab which slots into the aluminum of the gondola behind it. A silver-infused thermal grease is used to ensure sufficient heat flow between this tab and the readout crate backplane. The back of the readout crate connects directly to the gondola, so heat that reaches the backplane will be spread throughout the gondola and removed radiatively from the experiment.

During the Fall of 2007, tests of the thermal design were performed at the University of Minnesota in their thermal vacuum chamber. Due to a shortage of dfmux boards and wires in the feedthrough of the chamber only one board was tested during this run. The measurements shown in figure B–2 are using diodes built onto the motherboard itself.

The results were encouraging. The mezzanine boards got hotter than the dfmux motherboard, as expected since they are heat sunk to the motherboard, but even they only went to around 55^{o}C, within the 0 to 70^{o}C range permitted by the electronics. Several additional tests were performed by changing the main plate temperatures to account for the conditions that will be present as EBEX rises to its flight altitude. It was found that the board temperature is very strongly dependent on the pressure, with rises in pressure of just 1 Torr causing board temperatures to make a noticeable, albeit small, drop. Otherwise the board and its mezzanines maintained the temperature difference visible in figure B–2, which was acceptable.

Similar tests were performed in the Spring of the following year using NASA's much larger thermal vacuum chamber in Palestine, Texas. In this much larger chamber a readout crate with 6 dfmux boards and a full complement of mezzanines were tested.

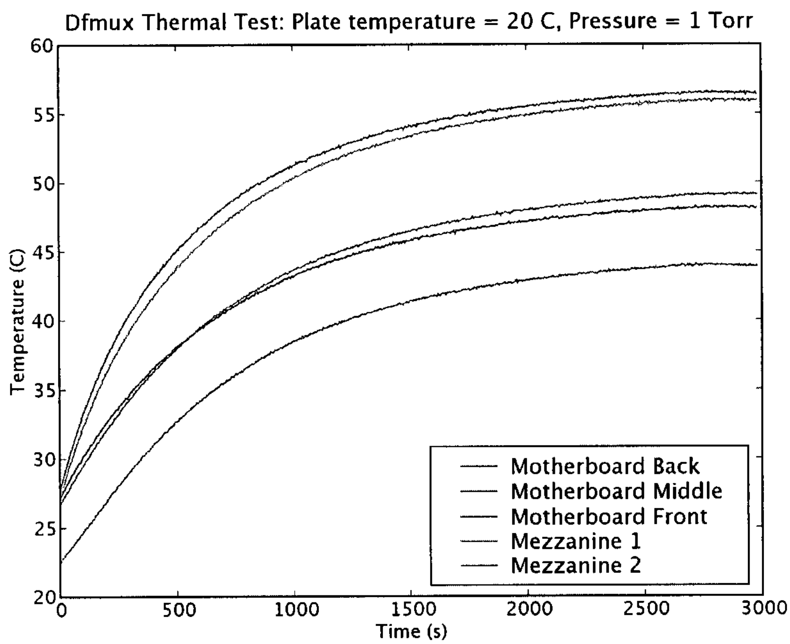

Figure B–2: A plot of the temperatures on the motherboard during a vacuum test run in Minnesota. The hottest components are the mezzanine board, as they are heat sunk to the motherboard. The coolest region is the front of the motherboard, that is the point the furthest from where the heat sinking is done. The reason it remains so cool is that little heat is generated in this part of the board and almost none flows this direction.

Initially these tests were concerning, as several mezzanine boards heated above 65 ^oC as soon as the pressure was lowered to 5-6 Torr, slightly higher than the pressures expected during flight. At this point the boards were turned off to avoid possible damage as we attempted to track down the problem. As it turned out, the overheating of the boards was largely caused by a weak connection to the "gondola," the final heat sink in the setup. In the actual EBEX flight the readout crate, into which the thermal tabs slot, will be connected directly to the gondola of the experiment, which is a very large heat sink. Unfortunately, for these tests no such heat sink was available so the read out crate was connected to an aluminum beam, which touched the temperature controlled wall of the chamber on one end. The aluminum had lower thermal conductivity than expected leading to a gradient of over 15^{o}C from one end to the other, sufficient to drive temperatures of the dfmux board into dangerous regions.

While no replacement heat sink was readily available at the test facility thermistors were placed along the aluminum beam and the chamber was cooled further to account for the gradient. Additionally, the outside of the readout create was painted white to increase radiative cooling. The results from these second sets of tests were much more encouraging, with the mezzanines reaching from 50^{o}C to 55^{o}C, much like the one during the University of Minnesota test. A slight temperature increase of middle boards versus edge boards due radiative heating, shown in Figure B–3 for the motherboards and Figure B–4 for the mezzanines, was found but did not present a major concern.

One anomaly is also visible in B–4. The mezzanines on a single motherboard were 4-5°C hotter than all of the others. This is suspected to have been due to a subtle issue with the heat sinking of that particular board although the connection was examined carefully in the lab afterward but the cause of this problem has not been identified. Fortunately, even if this problem appeared during operation it still would not over-heat the board, just put it closer to the upper limit of 70°C.

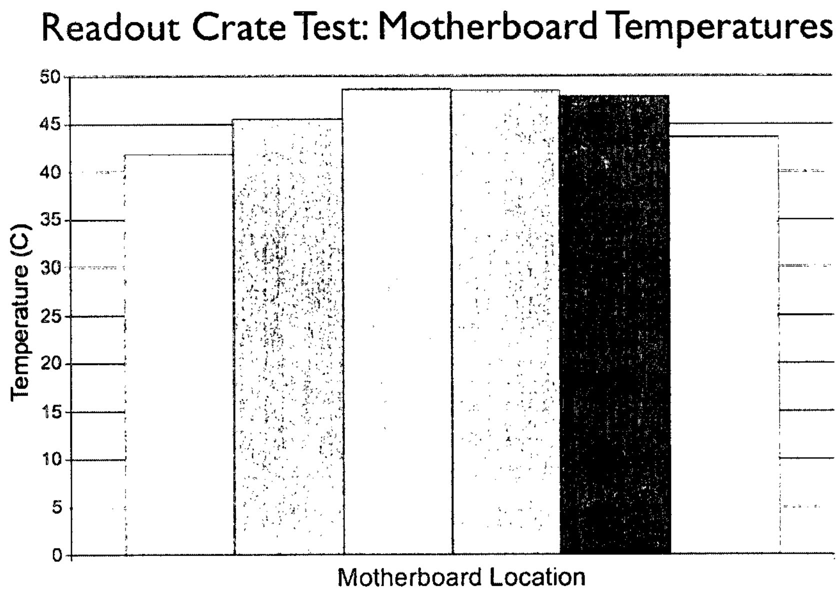

Figure B–3: A histogram of the temperatures of each of the motherboards when they were heatsunk to was near 15°C. Each block represents a the motherboard in that position in the readout crate. Only the hottest sensor, in the middle of the motherboard, is shown here.

Overall the heat sinking design appears functional and no major changes have been made within the readout crate. The heat sinking of the readout creates to the rest of the EBEX gondola will be tested at the integration of the experiment in

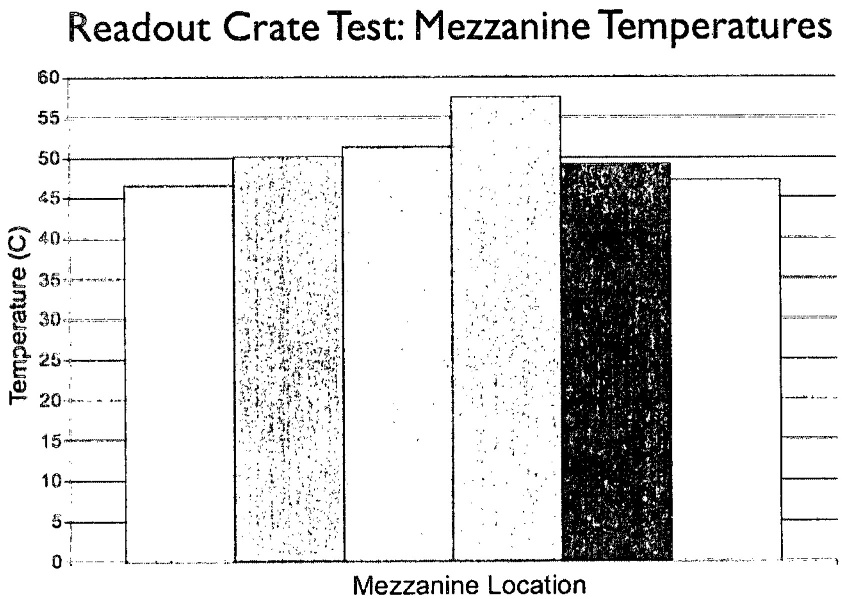

Figure B–4: A histogram of the temperatures of the hotter mezzanine of each of the motherboards when the "gondola" heatsink was near 15^oC. Each block represents the mezzanine in that position. Notice that one mezzanine is considerably warmer than the others, it is suspected that the heatsinking of this board was subtly different than the others.

November 2008. It is hoped that the thermal connection between these two pieces will be strong enough to deal with this heat load. If not, thermal straps of copper or potentially gold may be required to ensure the readout crate, as a whole, does not heat up significantly.

References

[1] W. Hu and M. White. A CMB polarization primer. *New Astronomy*, 2:323–344, October 1997. doi:10.1016/S1384-1076(97)00022-5.

[2] P. Ade, J. Bock, M. Bowden, M. L. Brown, G. Cahill, J. E. Carlstrom, P. G. Castro, S. Church, T. Culverhouse, R. Friedman, K. Ganga, W. K. Gear, J. Hinderks, J. Kovac, A. E. Lange, E. Leitch, S. J. Melhuish, J. A. Murphy, A. Orlando, R. Schwarz, C. O'Sullivan, L. Piccirillo, C. Pryke, N. Rajguru, B. Rusholme, A. N. Taylor, K. L. Thompson, E. Y. S. Wu, and M. Zemcov. First Season QUaD CMB Temperature and Polarization Power Spectra. *ApJ*, 674:22–28, February 2008. doi:10.1086/524922.

[3] K. D. Irwin. An application of electrothermal feedback for high resolution cryogenic particle detection. *Appl. Phys. Let.*, 66(15):1998–2000, 1995.

[4] M. Dobbs, E. Bissonnette, and H. Spieler. Digital Frequency Domain Multiplexer for Millimeter-Wavelength Telescopes. *IEEE Transactions on Nuclear Science*, 55:21–26, 2008. doi:10.1109/TNS.2007.911601.

[5] B. A. Bassett, S. Tsujikawa, and D. Wands. Inflation dynamics and reheating. *Reviews of Modern Physics*, 78:537–589, April 2006. doi: 10.1103/RevModPhys.78.537.

[6] J. Urrestilla, N. Bevis, M. Hindmarsh, M. Kunz, and A. R. Liddle. Cosmic microwave anisotropies from BPS semilocal strings. *Journal of Cosmology and Astro-Particle Physics*, 7:10–+, July 2008. doi:10.1088/1475-7516/2008/07/010.

[7] G. Smoot. Summary of results from cobe - preprint. *arXiv:astro-ph/9902027v1*, 1999.

[8] G. Hinshaw, J. L. Weiland, R. S. Hill, N. Odegard, D. Larson, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, N. Jarosik, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, M. Halpern, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Five-Year Wilkinson Microwave

Anisotropy Probe (WMAP) Observations: Data Processing, Sky Maps, and Basic Results. *ArXiv e-prints*, March 2008.

[9] A. de Oliveira-Costa. The Cosmic Microwave Background and its Polarization. In A. Adamson, C. Aspin, C. Davis, and T. Fujiyoshi, editors, *Astronomical Polarimetry: Current Status and Future Directions*, volume 343 of *Astronomical Society of the Pacific Conference Series*, pages 485–+. December 2005.

[10] E. M. Leitch, J. M. Kovac, N. W. Halverson, J. E. Carlstrom, C. Pryke, and M. W. E. Smith. Degree Angular Scale Interferometer 3 Year Cosmic Microwave Background Polarization Results. *ApJ*, 624:10–20, May 2005. doi: 10.1086/428825.

[11] J. L. Sievers, C. Achermann, J. R. Bond, L. Bronfman, R. Bustos, C. R. Contaldi, C. Dickinson, P. G. Ferreira, M. E. Jones, A. M. Lewis, B. S. Mason, J. May, S. T. Myers, N. Oyarce, S. Padin, T. J. Pearson, M. Pospieszalski, A. C. S. Readhead, R. Reeves, A. C. Taylor, and S. Torres. Implications of the Cosmic Background Imager Polarization Data. *ApJ*, 660:976–987, May 2007. doi:10.1086/510504.

[12] T. E. Montroy, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, P. Cabella, C. R. Contaldi, B. P. Crill, P. de Bernardis, G. De Gasperis, A. de Oliveira-Costa, G. De Troia, G. di Stefano, E. Hivon, A. H. Jaffe, T. S. Kisner, W. C. Jones, A. E. Lange, S. Masi, P. D. Mauskopf, C. J. MacTavish, A. Melchiorri, P. Natoli, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, G. Polenta, S. Prunet, S. Ricciardi, G. Romeo, J. E. Ruhl, P. Santini, M. Tegmark, M. Veneziani, and N. Vittorio. A Measurement of the CMB ¡EE¿ Spectrum from the 2003 Flight of BOOMERANG. *ApJ*, 647:813–822, August 2006. doi:10.1086/505560.

[13] U. Seljak and M. Zaldarriaga. Signature of Gravity Waves in the Polarization of the Microwave Background. *Physical Review Letters*, 78:2054–2057, March 1997.

[14] P. Oxley, P. A. Ade, C. Baccigalupi, P. deBernardis, H.-M. Cho, M. J. Devlin, S. Hanany, B. R. Johnson, T. Jones, A. T. Lee, T. Matsumura, A. D. Miller, M. Milligan, T. Renbarger, H. G. Spieler, R. Stompor, G. S. Tucker, and M. Zaldarriaga. The EBEX experiment. In M. Strojnik, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume

5543 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 320–331. November 2004. doi:10.1117/12.563447.

[15] A. Amblard, A. Cooray, and M. Kaplinghat. Search for gravitational waves in the CMB after WMAP3: Foreground confusion and the optimal frequency coverage for foreground minimization. *Physical Review D*, 75(8):083508–+, April 2007. doi:10.1103/PhysRevD.75.083508.

[16] S. Hanany, T. Matsumura, B. Johnson, T. Jones, J. R. Hull, and K. B. Ma. A Cosmic Microwave Background Radiation Polarimeter Using Superconducting Bearings. *ArXiv Astrophysics e-prints*, April 2003.

[17] P. L. Richards J. M. Gildemeister, Adrian T. Lee. Monolithic arrays of absorber-coupled voltage-biased superconducting bolometers. *Appl. Phys. Let.*, 77(24):4040–4042, 2000.

[18] T. M. Lanting. *Multiplexed readout of superconducting bolometers for cosmological observations.* Ph.D. thesis, University of California, Berkeley, 2006.

[19] A. T. Lee, H. Tran, P. Ade, K. Arnold, J. Borrill, M. A. Dobbs, J. Errard, N. Halverson, W. L. Holzapfel, J. Howard, A. Jaffe, B. Keating, Z. Kermish, E. Linder, N. Miller, M. Myers, A. Niarchou, H. Paar, C. Reichardt, H. Spieler, B. Steinbach, R. Stompor, C. Tucker, E. Quealy, P. L. Richards, and O. Zahn. POLARBEAR: Ultra-high Energy Physics with Measurements of CMB Polarization. In *American Institute of Physics Conference Series*, volume 1040 of *American Institute of Physics Conference Series*, pages 66–77. August 2008. doi:10.1063/1.2981555.

[20] J. M. Gildemeister, A. T. Lee, and P. L. Richards. A fully lithographed voltage-biased superconducting spiderweb bolometer. *Applied Physics Letters*, 74:868–+, February 1999. doi:10.1063/1.123393.

[21] J. Ruhl, P. A. R. Ade, J. E. Carlstrom, H.-M. Cho, T. Crawford, M. Dobbs, C. H. Greer, N. w. Halverson, W. L. Holzapfel, T. M. Lanting, A. T. Lee, E. M. Leitch, J. Leong, W. Lu, M. Lueker, J. Mehl, S. S. Meyer, J. J. Mohr, S. Padin, T. Plagge, C. Pryke, M. C. Runyan, D. Schwan, M. K. Sharp, H. Spieler, Z. Staniszewski, and A. A. Stark. The South Pole Telescope. In C. M. Bradford, P. A. R. Ade, J. E. Aguirre, J. J. Bock, M. Dragovan, L. Duband, L. Earle, J. Glenn, H. Matsuhara, B. J. Naylor, H. T. Nguyen, M. Yun, and J. Zmuidzinas,

editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 5498 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 11–29. October 2004. doi:10.1117/12.552473.

[22] J. Bock, S. Church, M. Devlin, G. Hinshaw, A. Lange, A. Lee, L. Page, B. Partridge, J. Ruhl, M. Tegmark, P. Timbie, R. Weiss, B. Winstein, and M. Zaldarriaga. Task Force on Cosmic Microwave Background Research. *ArXiv Astrophysics e-prints*, April 2006.

[23] B. O. P. Committee On The Physics Of The Universe, D. O. E. Astronomy, and N. R. C. O. T. N. A. Physical Sciences. *Connecting quarks with the cosmos : eleven science questions for the new century.* Connecting quarks with the cosmos : eleven science questions for the new century / Committee on the Physics of the Universe, Board on Physics and Astronomy, Division on Engineering and Physical Sciences, National Research Council of the National Academies. Washington, D.C. : National Academies Press, c2003., 2003.

[24] E. S. Battistelli, M. Amiri, B. Burger, M. Halpern, S. Knotek, M. Ellis, X. Gao, D. Kelly, M. Macintosh, K. Irwin, and C. Reintsema. Functional Description of Read-out Electronics for Time-Domain Multiplexed Bolometers for Millimeter and Sub-millimeter Astronomy. *Journal of Low Temperature Physics*, 151:908–914, May 2008. doi:10.1007/s10909-008-9772-z.