

A 145-GHz Interferometer for Measuring
the Anisotropy of the Cosmic Microwave
Background

William Bertrand Doriese

A DISSERTATION PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF PHYSICS

NOVEMBER 2002

© Copyright 2002 by William Bertrand Doriese.
All rights reserved.

Abstract

This thesis presents the design, construction, testing, and preliminary data analysis of MINT, the Millimeter INTerferometer. MINT is a 145-GHz, four-element interferometer designed to measure the anisotropy of the Cosmic Microwave Background (CMB) at spherical harmonics of $\ell = 800$ to 1900. In this region of ℓ -space, the CMB angular power spectrum should exhibit an exponential damping due to two effects related to the finite thickness of the last-scattering surface: photon diffusion and line-of-sight projection. Measurements in this region have already been made at 31 GHz by the Cosmic Background Imager (CBI). MINT's goal is to complement CBI by extending these results to a higher frequency that is much less prone to extragalactic point-source contamination. MINT's mission is also complementary to that of the Microwave Anisotropy Probe (MAP) satellite.

MINT observed the CMB in November and December, 2001, from an altitude of 17,000 feet in the Atacama Desert of northern Chile. We describe the performance of the instrument during the observing campaign. Based on radiometric hot/cold-load tests, the SIS-mixer-based receivers are found to have an average receiver noise temperature (double sideband) of 39 K in a 2-GHz IF bandwidth. The typical atmosphere contribution is 5 K. We assess the phase stability, gain stability, pointing accuracy, and overall sensitivity of the interferometer via observations of Mars and Jupiter, and find that the instrument is sufficiently stable to allow an ultimate experimental sensitivity at the few- μ K level needed for detection of the CMB anisotropy.

Acknowledgments

Friday, March 27, 1998 must rank alongside the all-time best days that anyone ever had. That day, the fledgling MINT experiment's working group received a visit from Nobel-laureate Bob Wilson, who generously gave us a day of his time to teach our group about interferometry and talk about the issues we would face in designing our experiment. Although we'd been tossing around ideas for a few months, the design of the instrument really began to take shape after that meeting. That night, I had my first date with my future wife, Rachel, whom I had met the night before and who quickly became the love of my life.

I grew up immersed in science. In a household of two parents who not only were professional scientists, but who kept readily available in our home several microscopes, at least one early Hewlett Packard programmable calculator, and an endless supply of answers and questions, how could I not? Although my particular choice of field, physics, was thrust upon me by circumstance as much as anything else (freshman organic chemistry met at 8:30 A.M. at Haverford College, and anyone who knows me knows that that was not going to happen), it was clear to me from a very early age that I would become a scientist. I was also strongly influenced by wonderful teachers all through my years in the Lower Merion School District and at Haverford College, and by a worldview-changing visit to the University of Pennsylvania Observatory sponsored by family friend Dr. William Blitzstein.

Lyman Page has been an ideal thesis advisor for me here at Princeton. His "dive right in" attitude was a perfect experience and confidence builder for me especially, as I had little hands-on lab experience when I arrived as a new grad student. An ill-placed static shock that put an \$8000 D-band frequency tripler out of commission was deemed a learning experience, not a stamped ticket to another research group. Lyman has always been available to talk and to listen, and during the MINT design somehow managed to be able to keep straight in his head what I was working on in

order to advise me at night, even though he was up to his ears in MAP satellite work most of the day. And, although I tease him often about it, I do appreciate one of the most important lessons Lyman has taught me: that a quick-and-dirty hack job very often is a preferable alternative to a day in the machine shop making something that looks beautiful.

Joe Fowler and I seemed to click as friends and colleagues right from the very beginning of his time in Princeton. Joe is in large part responsible for getting the many disparate parts of the MINT telescope smoothly integrated together. He designed much of the software, and really just filled in wherever needed. It would have been understandable for Joe, as a new postdoc, to join the experiment and demand a large, independent part of the experiment to work on by himself. Instead, Joe has truly been a leader; as such, he had a hand in every part of the project. He was also the one who delved into the hard math most often, including figuring out the sensitivity of the instrument and designing the observing strategy. I think the most remarkable thing about Joe is his ability to, with a few well-placed questions, get quickly up to speed on whatever I've been working on and to quickly recognize the interesting points. Joe demonstrated this quality time and again, most notably as I was designing the phase-switching system, when he helped me debug the circuit designs before board fabrication, but also during the final push to ship the instrument in the late summer of 2001, when he quickly learned about and then took over the tuning of one of the receivers while I went to get married. I'm sure this ability will help Joe in his eventual roles as teacher and research advisor.

Suzanne Staggs has been our other intrepid leader on this project, although she seemed to have the hardest of all possible jobs. Suzanne was kept so busy with her other ongoing polarization projects that for most of the development of MINT, her main role in the project was offering helpful suggestions while listening to the rest of us drone on in weekly meetings. Suzanne provided great leadership in the field in Chile, where her experience from running numerous other CMB experiments really shone through. She also was the second reader for this thesis.

Huan Tran and I worked together on this project from the very beginning. Our theses should best be read as two complementary sections of a larger work. Huan is

responsible for large parts of the experiment, notably the correlator and the channelizer. Huan and I seemed to compete in our early graduate school years to see who could come up with the most bizarre sleeping schedule. By all accounts, Huan won.

MINT was truly a team effort. Asad Aboobaker, Toby Marriage, and Yeong Loh, made large contributions to the project as graduate students. We also had strong undergraduate senior thesis work from Toby (I try to get out, but they keep pulling me back in) Marriage, as well as Daniel Wesley, Charles Dumont, and Jamie Hinderks. Other notable undergraduate contributors were Long Tran, Mark Tygert, Paul Oreto, Mark Morales, and Ziggy Kermish. Special thanks to Mark Halpern for building the RF loads and to Gaelen Marsden and Juan Burwell for serving time in Chile.

I cut my teeth on CMB experiments as part of the MAT/TOCO team. I thank, in addition to folks I've already mentioned, Mark Devlin, Tom Herbig, Eric Torbet, Jason Puchalla, Amber Miller, Angel Otarola, and everyone else involved in the project for the wonderful learning experience I was provided.

As a Princeton Physics graduate student, you live and die by your prelims and generals study groups. I thank Walter Briskin, Kentaro Nagamine, Michael Nolta, and Huan Tran for dragging me kicking and screaming through the exams.

While working on MINT, I received a lot of outside assistance. Bob Wilson's SMA group gave us their phase-locked loop design, and I had the great fortune to work closely with Bob as well as Robert (Kim) Kimberk, Todd Hunter, and Ray Blundell, all at Harvard. Dick Plambeck was generous with his time and ideas while helping us squeeze the maximum performance out of MINT's cryogenics, which are of his design.

I also received tons of help from within the Physics Department. I am grateful for enlightening discussions I had with Jim Peebles, Paul Steinhardt, and Norm Jarosik. There's nothing like learning from the masters! Stan Chidzik walked me through the design and also the soldering of the PC boards for the phase-switching system. The amazing machinists of the Jadwin machine shop, Ted Griffith, Laszlo Varga, Bill Dix, Glenn Atkinson, Mike Peloso, and John Mellodge, built much of the telescope, sometimes with nothing more than very sketchy drawings. I also thank all the folks that shepherded me through daily life in the Physics Department so seamlessly: John Washington, Kathy Warren, Helga Murray, Angela Qualls, Laurel Lerner, Joe Horvath, Claude Champagne, Richard Fernholz, Eva Zeisky, Harley

Pretty, and Kathy Patterson. Also thanks to Denis Barkats and Juan Burwell for being great office-mates. I'm going to miss your smiling faces, guys!

Finally, I thank my defense committee, Lyman, Bob Austin and Igor Klebanov, for being willing to help me through my final step at Princeton.

Contents

Abstract	iii
Acknowledgments	iv
1 Introduction: The Cosmic Microwave Background	1
1.1 Introduction	1
1.2 Characterizing the Anisotropy: The Angular Power Spectrum	2
1.3 Primordial Fluctuations	6
1.4 Foregrounds, Recent Experiments, and MINT	8
2 Overview: Interferometry and the MINT Telescope	11
2.1 Introduction	11
2.2 How Does the Interferometer Work?	15
2.3 The Digital Correlator	24
2.4 The Channelizer	28
2.5 Telescope Pointing	29
2.6 Calibration Noise Source	30
3 Design of the Receivers	32
3.1 Introduction	32
3.2 Optics	33
3.3 The Receiver Front End	37
3.4 The Receiver Back End	40
3.5 Cryogenics	42
4 SIS Tuning and Receiver Sensitivity Measurements	49
4.1 Introduction: Definitions and Conventions	49

4.2	Tuning the SIS's	54
4.3	Looking Outside the Dewar	60
4.4	The Digitizer Histograms	62
4.5	Determining T_{sys} During CMB Observations	64
5	Local Oscillators and Phase-Switching	73
5.1	Introduction	73
5.2	Background on Phase-Locked Loops	74
5.3	The MINT Phase-Locked Loop	75
5.4	Phase-Switching	79
5.5	Transient Response: Recapturing Lock	84
6	Observations & Analysis	87
6.1	Introduction	87
6.2	Converting Raw Correlator Data to Visibility	88
6.3	Jupiter Observation	91
6.4	Calibration Noise Source	97
6.5	Mars Calibrations	100
7	Conclusion	108
A	Calibration Odds & Ends	111
A.1	Diode Linearization	111
A.2	Digitizer Histogram Linearization	112
A.3	Digitizer Attenuator Calibration	113
A.4	LO Attenuator Calibration	114
B	Data Cuts	116

Introduction: The Cosmic Microwave Background

1.1 Introduction

A blizzard of observational cosmology results over the last several years has yielded new insight into the structure and composition of our universe. Anisotropy measurements of the cosmic microwave background radiation (CMB) showed that the spatial geometry of the universe is best described as flat (Miller et al. 1999; Mauskopf et al. 2000; de Bernardis et al. 2000; Lee et al. 2001). This conclusion was strengthened by a joint analysis of the CMB and the large-scale structure of galaxy clusters (Ostriker and Steinhardt 1995; Turner 1997; Efstathiou et al. 2002; Percival et al. 2002). At the same time, various analyses of mass distribution in galaxies and clusters confirmed that there is not nearly enough energy density in matter alone to close the universe (Bahcall et al. 1995; Carlberg et al. 1996; Bahcall and Fan 1998). The final piece of the puzzle was laid in place with the discovery, by two independent groups (Riess et al. 1998; Perlmutter et al. 1999), using type-Ia supernovae as standard candles, that the expansion of the universe appears to be accelerating. These interlocking measurements paint a compelling picture: that we live in a universe that has the critical energy density necessary to make it spatially flat, but that most of this energy density consists not of ordinary (baryonic) matter or even dark matter, but of a

repulsive, dark energy field that has the general properties of Einstein's cosmological constant, Λ (e.g. Bahcall et al. 1999).¹

While the idea that Λ plays an important role in the universe has required a change in thinking and will likely lead to interesting new physics (e.g. Caldwell et al. 1998), the underlying standard model of cosmology has held up remarkably well to the onslaught of recent observational tests. A particular success story has been the CMB, the oldest light in the universe. Even before the precise measurement of its spectrum by the COBE FIRAS experiment (Mather et al. 1994) showed it unequivocally to be thermal in origin, the CMB was widely believed to be the relic radiation from the hot, dense, early universe. The discovery of the CMB anisotropy on large angular scales by the COBE DMR (Smoot et al. 1992) showed that the universe's primordial, large-scale density fluctuations were consistent with inflationary cosmology (Guth 1981). Over the past decade, the CMB has been developed into a precise tool for measuring cosmological parameters, with experimental measurements being compared to detailed theoretical predictions about the anisotropy and as-yet-undiscovered CMB polarization for a variety of models. This thesis concerns a measurement of the temperature anisotropy of the CMB using an instrument called the Millimeter Interferometer (MINT).

1.2 Characterizing the Anisotropy: The Angular Power Spectrum

The early universe was filled with a smooth plasma of free protons, electrons, and photons, with the protons and photons tightly coupled by electrons through Coulomb and Thomson scattering (Peebles and Yu 1970).² By a redshift $z \sim 1100$, or $\sim 10^5$ years after the big bang, the expansion of the universe had caused the temperature of the photon-baryon fluid to drop enough that electrons and protons combined into hydrogen atoms. The photons, no longer coupled to the suddenly neutral matter, are what we see today as the CMB. Because the so-called recombination of electrons and

¹A recent estimate (Wang et al. 2002) of the proportions of these various components gives $\Omega_b = 0.05$, $\Omega_{DM} = 0.29$, and $\Omega_\Lambda = 0.66$.

²For a readable review, see Tegmark (1995) or Hu and Dodelson (2002). Many of the arguments advanced in this chapter are based on these two papers.

protons was rapid, the CMB is approximated to have been emitted at the *surface of last-scattering*, and its photons provide us with a near snapshot of the universe at recombination. CMB anisotropy experiments measure temperature fluctuations as a function of angle on the sky, and have found the last-scattering surface to have been remarkably smooth, with density fluctuations of only one part in $\sim 10^5$. The CMB is considered an ideal cosmological observable for two reasons: First, except for the interaction with a few foregrounds, which we discuss in Section 1.4, the photons have simply free-streamed to us since the very early universe. Second, since the fluctuations are tiny, the theory that describes their formation can operate in the realm of linear perturbation theory, making the effects relatively easy to calculate.

A few basic physical effects translate the conditions at the surface of last scattering into the temperature of the CMB radiation we receive. First, a slightly over-dense region of plasma is intrinsically hotter, so its photons are emitted with an increased temperature. Second, photons originating in a potential well are gravitationally red-shifted to a lower temperature as they climb out of the well. In the case of adiabatic initial fluctuations (density and temperature fluctuations are tied together) that are consistent with inflation, over-dense regions occur in potential wells. The density effect thus partially cancels the potential effect with the net result that over-dense, hotter regions of the plasma are seen as slightly colder than average in the CMB (Sachs and Wolfe 1967). A third effect is peculiar velocity, with photons last scattered by matter moving toward us receiving a Doppler blue-shift to higher temperature. In most situations, the Doppler contribution plays a minor role, as it is smoothed by a projection effect. We describe the causes of these primordial density, potential, and velocity perturbations in Section 1.3.

The common practice (e.g. Peebles 1973) in discussing the temperature anisotropy of the CMB is to represent the measured fluctuations, $\Delta T(\theta, \phi) = T(\theta, \phi) - T_{avg}$, in terms of spherical harmonics, $Y_{\ell, m}$:

$$Y_{\ell, m}(\theta, \phi) \propto P_{\ell}^m(\cos\theta)e^{im\phi}, \quad (1.1)$$

where the P_ℓ^m 's are the associated Legendre polynomials. The temperature fluctuations are described by³

$$\Delta T(\theta, \phi) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell,m} Y_{\ell,m}(\theta, \phi). \quad (1.2)$$

The positive integer ℓ is an inverse angular wavenumber, corresponding roughly to an angular period $\theta \sim \frac{2\pi}{\ell}$. The coefficients, $a_{\ell,m}$, are complex numbers, as are the $Y_{\ell,m}$'s. The $\Delta T(\theta, \phi)$ are real. The $Y_{\ell,m}$'s and $a_{\ell,m}$'s have the property, $Y_{\ell,m} = Y_{\ell,-m}^*$.

We calculate the variance of the temperature fluctuations, averaged over the sky, as this is what is measured by experiment:⁴

$$\langle (\Delta T)^2 \rangle_{sky} = \left\langle \left(\sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell,m} Y_{\ell,m}(\theta, \phi) \right)^2 \right\rangle_{sky} \quad (1.3)$$

$$= \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{\ell'=1}^{\infty} \sum_{m'=-\ell'}^{\ell'} a_{\ell,m} a_{\ell',m'} \langle Y_{\ell,m}(\theta, \phi) Y_{\ell',m'}(\theta, \phi) \rangle_{sky} \quad (1.4)$$

$$= \frac{1}{4\pi} \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{\ell'=1}^{\infty} \sum_{m'=-\ell'}^{\ell'} a_{\ell,m} a_{\ell',m'} \int Y_{\ell,m}(\theta, \phi) Y_{\ell',m'}(\theta, \phi) d\Omega. \quad (1.5)$$

Using the orthonormality relation of the spherical harmonics,

$$\int Y_{\ell,m}(\theta, \phi) Y_{\ell',m'}^*(\theta, \phi) d\Omega = \delta_{\ell,\ell'} \delta_{m,m'}, \quad (1.6)$$

we find that

$$\langle (\Delta T)^2 \rangle_{sky} = \frac{1}{4\pi} \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell,m} a_{\ell,-m} \quad (1.7)$$

$$= \frac{1}{4\pi} \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell,m} a_{\ell,m}^*. \quad (1.8)$$

Taking an average over an ensemble of universes,⁵ each of which has no preferred direction, allows us to sum over all m for each ℓ :

$$\langle (\Delta T)^2 \rangle = \sum_{\ell=1}^{\infty} \frac{2\ell+1}{4\pi} \langle a_{\ell,m} a_{\ell,m}^* \rangle. \quad (1.9)$$

³In a general spherical harmonic decomposition, ℓ runs from 0 to ∞ . However, since we have removed the average temperature, the $\ell = 0$ term vanishes, so the sum starts with $\ell = 1$. In practice, the dipole ($\ell = 2$) term cannot be measured, due to the proper motion of the earth.

⁴The measured variance also depends on the experiment's *window function*, or the weighted ℓ -space coverage. Accordingly, we leave explicit the ℓ -dependence of $\langle (\Delta T)^2 \rangle$.

⁵For small angular scales (large ℓ), an ensemble average is the same as an angular average, which is what an experiment really measures. For small ℓ , this is not the case, resulting in the increase in error-bars called *cosmic variance*.




Figure 1.1: A CMB angular power spectrum, or ℓ -space, plot showing the data from several previous experiments in which Princeton was involved. The COBE/DMR points establish scale-invariant initial fluctuations. The TOC098 data set from the MAT experiment localized the first acoustic peak. The upper limit at $\ell = 400$ is of 95% confidence. The two model curves are (solid) a Λ -model and (dashed) standard cold-dark-matter. This figure originally appeared in Miller et al. (1999), and is reproduced with permission.

We define the angular power spectrum, $C_\ell = \langle a_{\ell,m} a_{\ell,m}^* \rangle$. If the temperature fluctuations are normally distributed, as they are in most models, they are fully specified by the C_ℓ 's. Approximating the sum over ℓ in Equation 1.9 as an integral gives

$$\langle (\Delta T)^2 \rangle \approx \int_{\ell=1}^{\infty} \frac{2\ell+1}{4\pi} C_\ell d\ell. \quad (1.10)$$

Finally, we convert to a logarithmic interval, and approximate $\frac{2\ell+1}{4} \approx \frac{\ell+1}{2}$, so that:⁶

$$\langle (\Delta T)^2 \rangle \approx \int_{d \ln \ell = 0}^{\infty} \left(\frac{\ell(\ell+1)}{2\pi} C_\ell \right) d \ln \ell, \quad (1.11)$$

so that $\frac{\ell(\ell+1)}{2\pi} C_\ell$ is the contribution to the variance, $\langle \Delta T^2 \rangle$, per logarithmic interval in ℓ . This is the quantity that theorists and experimentalists use to communicate with each other when comparing CMB results. An example of a C_ℓ plot with some theoretical models and experimental data is shown in Figure 1.1.

1.3 Primordial Fluctuations

According to a simplified version of standard inflationary theory,⁷ the nascent universe underwent a period of superluminal, exponential expansion driven by the potential energy of a scalar field. Quantum fluctuations in this scalar field grew (in extent, not in amplitude) with the inflationary expansion, seeding perturbations in the metric that were scale-invariant (equal fluctuation power per logarithmic interval of scale). After the inflationary epoch, the photon-baryon fluid was left among gravitational potential wells at all length scales, and the wells began to attract the fluid. Radiation pressure of the compressing photons resisted infall, setting up acoustic oscillation modes that vibrated at the sound speed ($\sim 1/\sqrt{3}$ the speed of light). At recombination, the modes that were captured at maximal compression or rarefaction formed a series of peaks in the CMB's angular power spectrum. The lowest- ℓ peak corresponds to a mode that had just enough time to reach its first maximum compression, and as such is said to represent the size of the sound horizon at recombination. Localizing this first acoustic peak at $\ell \sim 200$ (see Figure 1.1) was a major triumph of CMB anisotropy observation in the late 1990's (Miller et al. 1999; Mautskopf et al. 2000; de Bernardis et al. 2000; Lee et al. 2001). Because the overall geometry of the universe (open, closed, or flat) strongly impacts the location in ℓ -space of the acoustic peaks,

⁶These conventions, both $\ln \ell$ and $\frac{2\ell+1}{4} \approx \frac{\ell+1}{2}$ are purely historical, and are used because the integrand of Equation 1.11 has a flat power spectrum at low ℓ for scale-invariant, adiabatic models (e.g. Peebles 1982). The difference between $\frac{2\ell+1}{4\pi}$ and $\frac{\ell+1}{2\pi}$ is significant only at low ℓ .

⁷Inflation is a widely accepted paradigm, and its predictions fit the data well, but of course it isn't the only possible answer. For an interesting alternative based on extra dimensions and branes, see Khoury et al. (2001).

localization of the first peak has constrained the universe to be flat, within the context of currently accepted models (e.g. Wang et al. 2002). Recent experiments have discovered the second acoustic peak and hint at a third (see Figure 1.2), providing further confirmation of the standard cosmological model.

Angular scales with ℓ significantly lower than $\ell = 200$ represent modes that were too big to allow the photon-baryon fluid to self-interact before recombination. These modes are said to be outside the sound horizon. Since no acoustic oscillations could occur, the CMB has retained the imprint of the initial fluctuations seeded by inflation. This prediction has been verified by the large-angular-scale measurements of the COBE DMR (Wright et al. 1996, see Figure 1.1), which found a featureless power spectrum for $\ell < 20$.

At the small angular scales ($\ell > 800$) where MINT is sensitive, the angular spectrum is predicted to exhibit an exponential decrease in power. There are two effects that create this high- ℓ *damping tail*, and both stem from the fact that recombination, while rapid, is not instantaneous. What we have thus far approximated as the surface of last scattering is perhaps better viewed as a “slab” of last scattering, with the optical depth for Thomson scattering, τ_T , decaying over time as the free-electron density dwindles. The first cause of power damping at small angular scales is line-of-sight projection. The received CMB photons from any given direction come from a variety of depths in the last-scattering slab; Jones and Wyse (1985) calculate that the photons are last scattered over an effective redshift range of $\Delta z \sim 80$. An observation of the CMB integrates along the line of sight, washing out fluctuations on scales smaller than the slab thickness.⁸ The second cause of the damping tail is photon diffusion, or Silk damping (Silk 1968). As τ_T falls during recombination, photons can travel further without encountering electrons. Photons in acoustic modes with small spatial extent random-walk out of the potential wells, suppressing the oscillation of these modes. Detailed numerical calculations (see, e.g. Hu and Dodelson 2002) show that Silk damping should set in beyond the third acoustic peak ($\ell > 800$). A first observation of the damping tail was made by the CBI experiment (Padin et al. 2001, see Figure 1.2).

⁸As Tegmark (1995) notes, the damping due to projection is much like a coarse telescope beam smoothing out fine-scale fluctuations in an astronomical source, except the projection smoothing is radial rather than transverse.

There are a few effects important to MINT's ℓ -range that occur between last scattering and the present. We classify these as foregrounds, and discuss them in the next section.

1.4 Foregrounds, Recent Experiments, and MINT

CMB photons free-stream to us from the surface of last scattering only if they encounter no ionized gas along their path. However, the universe is widely thought to have become reionized at some point in the recent cosmological past, due to the ultraviolet emission from quasars and/or the first generation of stars. Moreover, we know the universe to be locally ionized. The effect of this global reionization on the CMB angular power spectrum depends on the optical depth, τ_{ri} , of the free electrons between the observer and the last-scattering surface. The earlier that reionization occurs, the greater is the optical depth, and the more the CMB power spectrum is washed out on fine angular scales. Current CMB anisotropy data place a limit on the redshift of global reionization of $z_{ri} < 20$, with a corresponding limit on the optical depth of $\tau_{ri} < 0.3$ (Wang et al. 2002).

Reionization can also occur locally in patches on the sky, in the hot gases within large galaxy clusters. The observational consequence for CMB photons, called the Sunyaev-Zeldovich (SZ) effect (Sunyaev and Zel'dovich 1972), is extra power at high ℓ , and is due to inverse-Thompson scattering off of the hot electrons. The SZ effect shifts the CMB photon population from low to high frequency, and would result in a temperature *decrement* in the direction of the cluster at MINT's 145-GHz central signal frequency. Bond et al. (2002) report anomalous power beyond $\ell = 2100$ which may be due to the SZ effect, but further work is needed to verify this result. In any case, the SZ effect is not expected to interfere with MINT's observation of the damping tail in the range $800 < \ell < 1900$.

There are many secondary foregrounds of a more local origin that have the potential to mask CMB anisotropies. An exhaustive discussion of many possible foregrounds (Tegmark and Efstathiou 1996), including dust emission, synchrotron radiation, free-free emission, and extragalactic point sources, finds that for most of the radio frequencies and angular scales of interest these foregrounds are small compared




Figure 1.2: A collection of recent experimental data. The two theoretical curves are produced by CMBFAST (Seljak and Zaldarriaga 1996) and correspond to a standard cold dark matter model and to the best-fit parameters by Wang et al. (2002) to all CMB anisotropy results published as of late 2001. Experimental data are from: BOOMERanG, Netterfield et al. (2002); CBI, Pearson et al. (2002); Mason et al. (2002); DASI, Halverson et al. (2002); MAXIMA, Lee et al. (2001); VSA, Scott et al. (2002). This plot was produced by Michael Nolta.

to the underlying CMB anisotropy signal. The one major exception is the contribution of extra-galactic point sources: radio galaxies that are randomly distributed and too small to be resolved by a CMB experiment. Using a two-component model (IR and radio) to determine the contribution of galaxies, Toffolatti et al. (1998) find that the point source contribution would overwhelm a CMB anisotropy measurement in the $\ell \sim 1000$ range below ~ 90 GHz. Insensitivity to this point-source foreground is the fundamental reason for designing MINT to operate at 145 GHz. Measurements of the damping tail have been made from $800 < \ell < 1600$ by the Cosmic Background Imager, a 31-GHz CMB interferometer (Pearson et al. 2002; Mason et al. 2002, see Figure 1.2). However, the point source contamination in the CBI data is severe enough that it must be observed by another radio telescope and subtracted (Padin et al. 2002). Subtraction of data can bias a result. According to the calculations by

Toffolatti et al., MINT should more be than 500 times less sensitive to point sources than is CBI, so we should not need to correct for or remove point sources when analyzing MINT data. A measurement of the damping tail in agreement with the CBI results at MINT's high frequency would provide yet another compelling confirmation of the cosmological standard model.

Overview: Interferometry and the MINT Telescope

2.1 Introduction

MINT, the Millimeter INTerferometer, is a compact interferometer designed to measure cosmic microwave background (CMB) fluctuations at small angular scales from a high-altitude site in northern Chile. Similar in many respects to the CBI and DASI instruments (Padin et al. 2002; Leitch et al. 2002), MINT is designed to take advantage of two important characteristics of interferometers, namely the strong rejection of unwanted signals (e.g. atmospheric fluctuations, man-made radio interference, ground pickup, cross-talk) that are not *correlated* in more than one receiver, and that the output of an interferometer is the Fourier transform (FT) of the sky brightness distribution, which means that it makes a direct measurement of the angular power spectrum discussed in Chapter 1. An advantage of direct Fourier-space measurement is that an interferometer is insensitive to the $1/f$ amplifier gain fluctuations that necessitate Dicke-switching in a differencing radiometer (such as the MAT/TOCO experiment), as sky-differencing occurs instantaneously.

The MINT instrument consists of four D-band (~ 145 GHz) heterodyne receivers, a 1-GHz-sample-rate digital correlator, and support electronics on a two-axis (azimuth, elevation) telescope base. A steel tripod supports the four-foot-diameter azimuth bearing, upon which rests the elevation fork. The four receivers and the electronics all move as a single, rigid unit in the aluminum elevation cage, which is supported in the fork by a 2-inch, stainless-steel axle. External to the main telescope




Figure 2.1: An overhead-view engineering drawing of the MINT elevation cage. The baselines (bold lines) are of lengths are 13.750, 21.149 (x4), and 40.000 inches, corresponding to $\ell = 1062, 1634,$ and 3090. MINT’s (u, v) -coverage is given in Figure 2.5 and Table 2.1.

base is more support equipment, including the “cooling box,” a water pump with a thermally regulated reservoir to cool the electronics; the “pointing rack,” which contains the support electronics, including a PC, for the azimuth and elevation motors; and the “compressor rack,” which contains the compressors and motor drivers for the receiver cryogenics. When the telescope is deployed on Cerro Toco in Chile, data and commanding radio links allow the telescope to be run remotely from the nearby town of San Pedro de Atacama (~ 30 km away). A 50-kW Katolight diesel generator supplies electrical power for the experiment.

Figure 2.1 shows an overhead view of the elevation cage. Four of the baselines¹ form a rhombus with sides of 21.149”. The other two baselines are the diagonals of the rhombus, with lengths 13.750” and 40.000”. The rest of the electronics on the

¹A baseline is any pair of receivers in an interferometer. The baseline length is measured between the optical centers of the receivers. In a general interferometer with n receivers, there are $\frac{n(n-1)}{2}$ possible baselines.

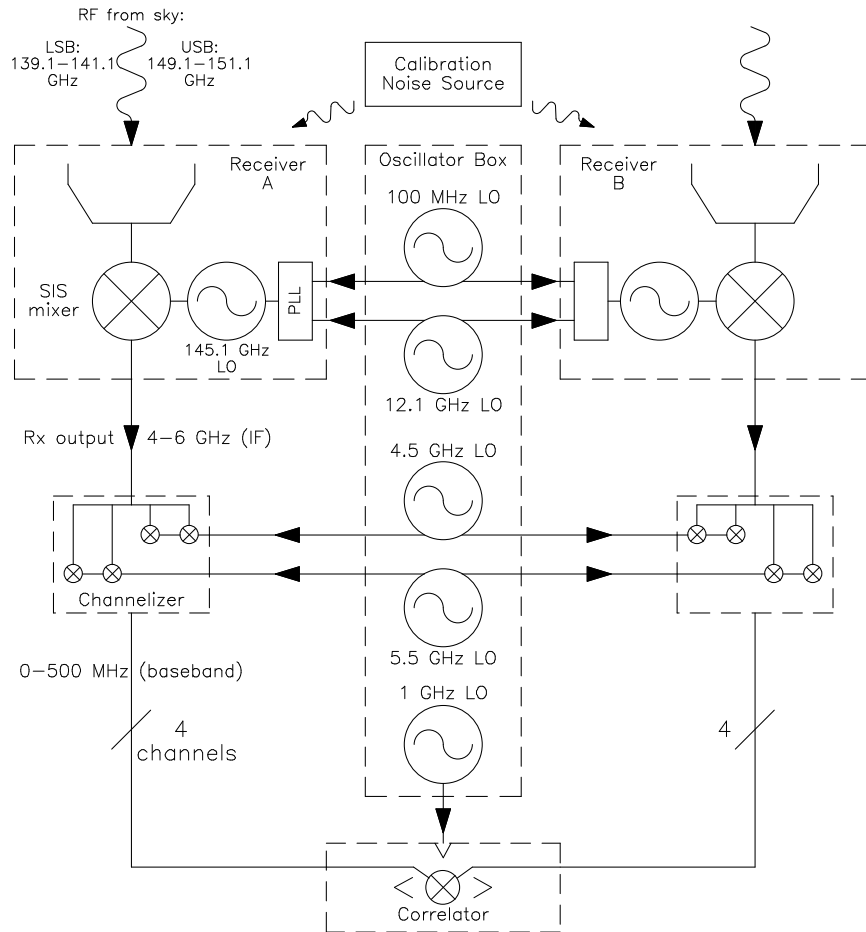


Figure 2.2: The signal path and frequency conversions for a single MINT baseline. The receivers are double-sideband and the correlator is a digital-lag correlator. All oscillators are phase-locked to the 100-MHz source.

cage are mounted in a pair of standard 19-inch instrument racks (labeled “Rack 1” and “Rack 2” in the figure).

A diagram of the signal path through the MINT telescope is shown in Figure 2.2. The receivers gather incoming radiation (radio frequency, or RF) of roughly 145 GHz and mix it down to an intermediate (IF) band centered at 5 GHz. The channelizer breaks each receiver’s IF band into four baseband channels of 0–500 MHz, thus preparing them for digital sampling at a 1 GHz rate. The correlator digitizes all 16 signal bands and performs a large digital processing job.

MINT's compact, rigid receiver cage offers several key advantages.² First, because the receivers, channelizer, and correlator never move with respect to one another, the instrument is stable to phase fluctuations that could otherwise be caused by flexing cables and other moving parts. Without phase stability, time integration of the signal is impossible. Second, the baseline projection angle never changes, which means the (u,v) coverage (see Section 2.2) does not change as the telescope observes the sky at different angles. This would be a poor feature in an instrument designed to make maps, as it would yield incomplete (u,v) -coverage and thus a distorted map, but it is highly advantageous for CMB observations, where we want to concentrate our integration time on a small portion of (u,v) space to improve signal-to-noise. Third, a ground shield can be placed around all the receivers, limiting the potentially devastating systematic error of correlated ground pickup. For example, a 300 K ground signal, if observed in sidelobes at a level of -50 dBi over 2π sr, even if it were only 10% correlated in two receivers, would produce a signal of $150 \mu\text{K}$, swamping the expected CMB signal. As its ground shield, MINT employs a hexcell, four-sided, inverted frustum of height $48''$ and flare angle 10° , designed so the receivers can't see the ground by any number of non-diffracted bounces. The ground shield moves with the elevation cage. Enclosure of a substantially larger instrument would be hampered by the weight of a large shield and the stress generated by the strong desert wind. Finally, a small instrument is easy to work with. MINT is designed to fit, with minimal removal of externally jutting parts, inside the Jadwin Hall freight elevator. With a set of four heavy-duty casters that mount to the azimuth base, the instrument can be pushed around by hand by only a few people. The entire instrument spent more than a month on the roof of Jadwin Hall in the summer of 2001 in pre-deployment tests.

This chapter is organized as follows. We examine the theory of interferometry in Section 2.2, with special attention to MINT's digital lag correlator and double-sideband receivers. Descriptions of several instrument systems follow, including the

²Large radio interferometers are designed to achieve fine angular resolution with long baselines (for example, the VLA has ~ 0.05 arc sec maximum resolution at 43 GHz in its longest baseline of 30 km). CMB interferometers, on the other hand, are designed with their receivers as close together as possible for high sensitivity to the angular scales of current interest, 0.1 to 1° .

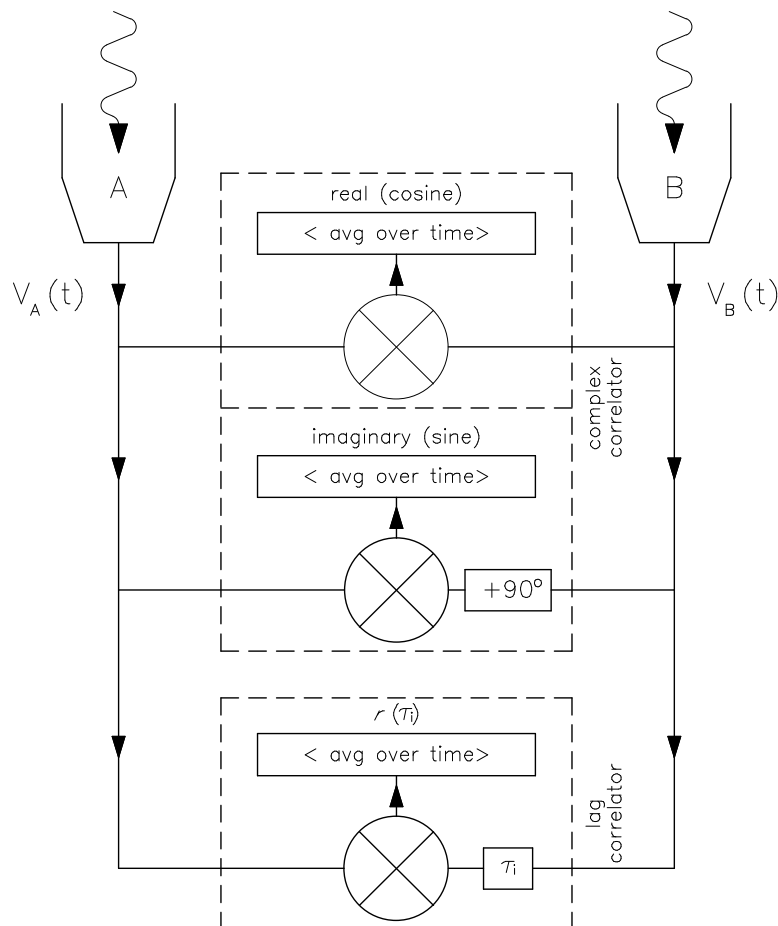


Figure 2.3: A baseline of an interferometer with several different correlators. The action of a correlator is to multiply the two incoming voltage signals together and then average over time. The top correlator produces an output proportional to the cosine of the phase difference of the two incoming signals and the middle correlator, with 90° of added phase, produces an output proportional to the sine; these two together make up a “complex correlator.” The bottom correlator, which produces an output that is a function of the inserted instrumental delay, or lag, τ , is called a “lag correlator.” MINT’s correlator is of the lag type.

correlator (Section 2.3), the channelizer (Section 2.4), the telescope pointing (Section 2.5), and the calibration noise source (Section 2.6). Other systems, such as the receivers and oscillators, are given their own chapters.

2.2 How Does the Interferometer Work?

We review some key ideas about interferometry important for describing the MINT telescope design and the point-source analysis of Chapter 6. Wherever possible, we

strip away non-critical details to work with the simplest possible example, and then add components back to the model as necessary. We develop the equivalence between lag correlators and complex correlators, and then use the two descriptions interchangeably. The ideas in this section are based on two sources, the invaluable textbook, *Interferometry and Synthesis in Radio Astronomy* (Thompson et al. 1994, hereafter, TMS), and a very useful MINT internal group memo, *MINT Phasing Requirements and Sideband Separation* (Fowler 2001).

Figure 2.3 shows the important components of an interferometer: two receivers that each produce a voltage output, and a correlator that multiplies the voltages together and averages over time. We start our discussion by investigating the correlator output with two receivers that are sensitive only to a single frequency, ν_R , and with a phase difference ϕ between the two correlator inputs, so that the multiplier output (labeled “real” in Figure 2.3) has the form³

$$V_A(t) \times V_B(t) = [V_A \cos(2\pi\nu_R t)] \times [V_B \cos(2\pi\nu_R t + \phi)] \quad (2.1)$$

$$= \frac{1}{2} V_A V_B [\cos \phi + \cos(4\pi\nu_R t + \phi)]. \quad (2.2)$$

Time-averaging removes the high-frequency term, leaving the fundamental result that the output of the correlator, r , is proportional to the product of the amplitude of its inputs and the cosine of the phase difference between them.

We introduce a second multiply-and-average correlator stage, but with a 90° phase shift in one of the arms inside the correlator. As before, the phase difference between the correlator inputs is ϕ , and we call the total phase difference of the signals that are multiplied, Φ . The output of the correlator is now $r \propto \cos \Phi = \cos(\phi - \frac{\pi}{2}) = \sin \phi$, or the sine of the phase difference of the receiver outputs. In general, the correlation of the two receiver outputs can be viewed as a complex number, $r \propto V_A V_B e^{i\phi}$, and a pair of correlators that measures the real ($\cos \phi$) and imaginary ($\sin \phi$) parts simultaneously is known as a “complex correlator.” The real and imaginary parts of the complex correlation each carry half the information necessary for full image reconstruction. In our CMB observations, where we are solely interested in measuring the variance of the sky, the real and imaginary values in the complex correlation yield

³At this point, we make no assumptions about the origin of the receiver input radiation; it could be from the sky, or from some near field source such as injected noise.

two independent measurements. Measurement of the full complex correlation is the ideal toward which correlators are designed.

MINT's correlator is a different type, known as a *lag correlator* (bottom of Figure 2.3). It measures the correlation function, $r(\tau) = \langle V_A(t)V_B(t - \tau) \rangle$, using a variable instrumental time delay, τ . A lag correlator and a complex correlator make equivalent measurements. The total phase difference in a lag correlator between the multiplied signals is $\Phi = \phi + 2\pi\nu_R\tau$, and so the measured correlation function is (still with receivers sensitive to a single frequency):

$$r(\tau) = \langle V_A(t)V_B(t - \tau) \rangle = \frac{1}{2}V_AV_B \cos(\phi + 2\pi\nu_R\tau). \quad (2.3)$$

The Fourier transform of $r(\tau)$ yields the correlation of the receiver outputs as a function of frequency:

$$\tilde{r}(\nu) = \int_{-\infty}^{\infty} e^{-i2\pi\nu\tau} r(\tau) d\tau \quad (2.4)$$

$$= \frac{1}{4}V_AV_B [e^{i\phi}\delta(\nu - \nu_R) + e^{-i\phi}\delta(\nu + \nu_R)]. \quad (2.5)$$

Since $r(\tau)$ is real, $\tilde{r}(\nu) = \tilde{r}^*(-\nu)$, and so the negative frequency term ($\delta(\nu + \nu_R)$) carries no new information. We drop it and find that $\tilde{r}(\nu)$ is indeed proportional to $r \propto V_AV_B e^{i\phi}$, as advertised (Fowler 2001). In the case of a broad-band receiver, one would integrate $\tilde{r}(\nu)$ over frequency to show this equivalence.⁴ For the remainder of this thesis, we use descriptions of the complex correlator and the lag correlator interchangeably.

When the interferometer looks at the sky, the phase difference between the receiver outputs is the result of a geometrical delay, or path difference (see Figure 2.4). We assume a point source of astronomical distance, so the rays incident on the two receivers are well approximated as parallel. If the source is at an angle, θ , from the normal to the baseline, the path difference is $P.D.(\theta) = D \sin \theta$, where D is the baseline length. The phase difference between the receiver inputs is then $\phi = 2\pi\nu_R D \sin \theta / c$ (c is the speed of light), meaning the output of the interferometer is

$$r(\theta) \propto e^{i2\pi\nu_R D \sin(\theta)/c}. \quad (2.6)$$

⁴As Fowler (2001) notes, multiple frequencies might be expected to generate cross terms in the correlator output, for example in Equations 2.2, 2.3, and 2.5. The multiplication stage of the correlation does indeed mix frequencies, but averaging over a time T_{avg} averages away any cross terms that differ in frequency by more than $\Delta\nu > 1/T_{avg}$. Since in MINT, T_{avg} is 0.5 sec and each channel has a bandwidth of 500 MHz, different frequencies are very nearly independent.




Figure 2.4: An interferometer looking at a point source on the sky.

Finally, we introduce a brightness distribution on the sky, $B(\theta)$ (power per unit angle, area, and bandwidth), and a receiver power-reception beam, expressed as an effective area, $A(\theta)$. Since $A(\theta)$ is typically highly directional, we adopt the small angle approximation $\sin \theta \rightarrow \theta$. To get the total interferometer response, we integrate over the source.⁵

$$\mathcal{V}(D, \nu_R) \propto \int A(\theta)B(\theta)e^{i2\pi\nu_R D\theta/c} d\theta, \quad (2.7)$$

which yields the result that an interferometer measures the FT of the receiver response to the brightness distribution on the sky. In interferometry parlance, $\mathcal{V}(D, \nu_R) = |\mathcal{V}|e^{i\phi_v}$ is called the *complex visibility*. The sinusoidal corrugations on the sky in either the real or imaginary part of the visibility of a point source are known as *fringes*.

Until now, we have limited our discussion to a δ -function in frequency space, and have not mentioned the frequency conversion that occurs in two mixing steps

⁵We assume that the source is made up of incoherent point sources. Also, we make the simplification here that the sky is one-dimensional to show the Fourier relationship. We generalize to a 2-D beam, source, and FT later in this section.

in MINT. In general, as we suggested in Equation 2.5, the time-based correlation function output of a lag correlator is the FT of the cross-power spectrum of the receiver responses.⁶ Averaging the visibility over frequency reduces the interferometer beam on the sky by an envelope function. In the particular case of MINT, with its small fractional RF bandwidth ($\frac{12 \text{ GHz}}{145.1 \text{ GHz}}$) and close-packed antennae, the effect is small. When mixers operate in double-sideband (DSB) mode, as do the MINT SIS's, the correlator output departs from the intuition developed in Equation 2.7. A DSB interferometer cannot observe the real and imaginary parts of the complex correlation simultaneously. To show this, we introduce a local oscillator with frequency ν_{LO} and a signal frequency ν_{RF} into a mixer, whose function is to multiply the signals together, thus producing the difference frequency of its inputs (intermediate frequency, or ν_{IF}). As usual, we restrict discussion of the receiver to a monochromatic output, ν_{IF} , meaning that the receiver is sensitive to two incoming RF signal frequencies: $\nu_{RF} = \nu_{LO} \pm \nu_{IF}$. We allow the LO's in the two receivers to differ in phase by ϕ_{LO} and observe a point source at angle θ as in Figure 2.4. The phase difference of the signals multiplied in the correlator is then

$$\Phi_{L,R}^U = 2\pi(\nu_{IF} \pm \nu_{LO})\frac{D\theta}{c} \pm \phi_{LO} \quad (2.8)$$

for the real part of the upper(+)/lower(-) sideband, and

$$\Phi_{L,I}^U = 2\pi(\nu_{IF} \pm \nu_{LO})\frac{D\theta}{c} \pm \phi_{LO} - \frac{\pi}{2} \quad (2.9)$$

for the imaginary part. The complex correlator output due to just the upper sideband is then

$$r_U \propto e^{i2\pi(\nu_{IF} + \nu_{LO})D\theta/c + \phi_{LO}} \quad (2.10)$$

which, other than the multiple frequency terms, is similar to Equation 2.6. The lower sideband response is similar, with just some sign flips. However, in the double sideband case, we have (assuming equal signal power in the sidebands)

$$r_{DSB,R} \propto \cos \Phi_{U,R} + \cos \Phi_{L,R} \quad (2.11)$$

$$= 2 \cos \left(\frac{\Phi_{U,R} + \Phi_{L,R}}{2} \right) \cos \left(\frac{\Phi_U^R - \Phi_L^R}{2} \right) \quad (2.12)$$

$$= 2 \cos \left(2\pi\nu_{IF}\frac{D\theta}{c} \right) \cos \left(2\pi\nu_{LO}\frac{D\theta}{c} + \phi_{LO} \right) \quad (2.13)$$

⁶This is a statement of the Wiener-Khinchin relation.

for the real part and, similarly for the imaginary part,

$$r_{DSB,I} \propto \cos \Phi_{U,I} + \cos \Phi_{L,I} \quad (2.14)$$

$$= 2 \sin \left(2\pi\nu_{IF} \frac{D\theta}{c} \right) \cos \left(2\pi\nu_{LO} \frac{D\theta}{c} + \phi_{LO} \right). \quad (2.15)$$

The complex correlation is then

$$r_{DSB} \propto e^{i2\pi\nu_{IF}D\theta/c} \cos \left(2\pi\nu_{LO} \frac{D\theta}{c} + \phi_{LO} \right) \quad (2.16)$$

which is qualitatively different from the above expressions, such as Equations 2.6 or 2.10, as the complex correlation is now enveloped by a cosine, so it covers on average only half the sky. For instance, if the LO's are phase-matched ($\phi_{LO} = 0$), the interferometer is sensitive to only the real part of the sky's FT. In essence, because its output is zero, the imaginary part of the complex correlator is wasted. We note also that the envelope function is independent of ν_{IF} , so if the imaginary part of the complex correlation disappears for one choice of IF frequency, it will disappear for all. Losing sensitivity to half the correlation is non-ideal for many reasons, but has the potential to hurt us the most when observing the calibration noise source, which, due to unknown delay in the waveguide that brings the signal to each receiver, could end up being invisible with the wrong choice of LO phases.

The solution to the double-sideband problem is to spend half the observing time with an extra $\frac{\pi}{2}$ of LO phase difference in each baseline, and to combine the results (post-observation) in a way that separates the sidebands. We return to the lag correlator to develop the sideband separation algorithm, since it will be applied to MINT correlator data. We want to measure the phase difference and amplitude of the upper and lower sidebands of the correlator inputs, ϕ_U , ϕ_L , U , and L . With LO's phase-matched, the correlation function is

$$r_0(\tau) = U \cos(\phi_U + 2\pi\nu_{IF}\tau) + L \cos(\phi_L + 2\pi\nu_{IF}\tau), \quad (2.17)$$

and when LO's are 90° out of phase, it is

$$r_{90}(\tau) = U \cos\left(\phi_U - \frac{\pi}{2} + 2\pi\nu_{IF}\tau\right) + L \cos\left(\phi_L + \frac{\pi}{2} + 2\pi\nu_{IF}\tau\right). \quad (2.18)$$

Taking the FT of these functions, and dropping the negative frequency terms which, as in Equation 2.5, carry no additional information, we find:

$$\tilde{r}_0(\nu) = \left[\frac{U}{2} e^{i\phi_U} + \frac{L}{2} e^{i\phi_L} \right] \delta(\nu - \nu_{IF}) \quad (2.19)$$

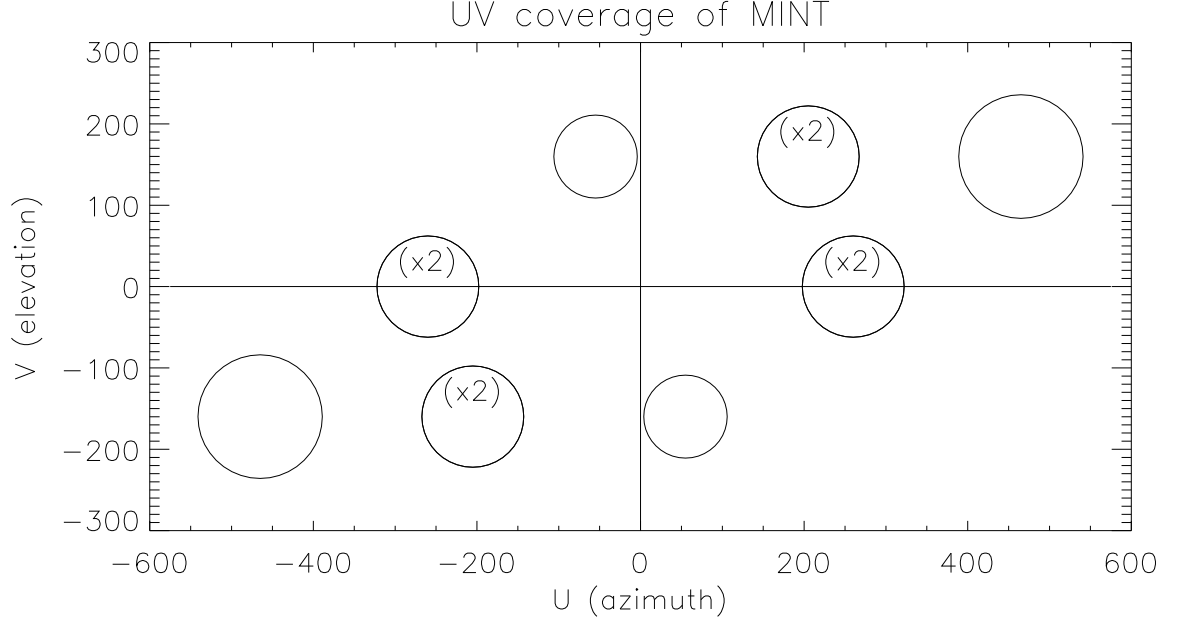


Figure 2.5: MINT's (u, v) coverage, with FWHM contours shown for the (non-Gaussian, DADRA-simulated, see Chapter 3) power-reception window-functions of each baseline. These window functions are formed from the cross correlation of the aperture-voltage-reception pattern of the receiver primary mirrors in each baseline (e.g. White et al. 1999). As there are two pairs of parallel baselines in the rhombus, two of the (u, v) circles are double-covered, as noted by (x2). The points (u, v) and $(-u, -v)$ contain identical information, because since $T(\vec{x})A(\vec{x})$ is real, its FT, the complex visibility, obeys the relation $\mathcal{V}(u, v) = \mathcal{V}^*(-u, -v)$.

$$\tilde{r}_{90}(\nu) = \left[-i\frac{U}{2}e^{i\phi_U} + i\frac{L}{2}e^{i\phi_L} \right] \delta(\nu - \nu_{IF}). \quad (2.20)$$

Finally,

$$\text{USB} = \delta(\nu - \nu_{IF})Ue^{i\phi_U} = \tilde{r}_0(\nu) + i\tilde{r}_{90}(\nu) \quad (2.21)$$

and

$$\text{LSB} = \delta(\nu - \nu_{IF})Le^{i\phi_L} = \tilde{r}_0(\nu) - i\tilde{r}_{90}(\nu). \quad (2.22)$$

The specific form of this data processing algorithm used in MINT, with attention to the discrete FT and the complications caused by the secondary frequency conversion (channelizer), is explained in further detail in Chapter 6.

Our discussion of baselines and sky signal distribution has thus far been limited to a single dimension (see Equation 2.7). Converting to two dimensions, with a flat-sky (small-angle) approximation, the angle θ is replaced by \vec{x} . The source brightness becomes $B(\vec{x})$, and the effective area of the beam is written $A(\vec{x})$. The baseline, D , is really a 2-D vector, \vec{D} . The important physical quantity involving the baseline is

Baseline	x (in.)	y (in.)	u (x/λ)	v (y/λ)	ℓ ($2\pi u $)
D–A	−4.470	13.003	−21.64	62.94	1062
C–B	37.828	13.003	183.10	62.94	3090
D–C	−21.149	0.000	−102.37	0.000	1634
A–B	21.149	0.000	102.4	0.000	1634
D–B	16.679	13.003	80.73	62.94	1634
A–C	−16.679	−13.003	−80.73	−62.94	1634

Table 2.1: MINT baselines, (u, v) -coverage, and ℓ -coverage. The wavelength at the center of MINT’s RF band, 145.1 GHz, is 0.2066 cm.

not \vec{D} itself (see Equation 2.6), but rather $\vec{D}\nu/c = \vec{D}/\lambda$, which we call \vec{u} , and its components (u, v) .⁷ The visibility is then

$$\mathcal{V}(\vec{u}) \propto \int A(\vec{x})B(\vec{x})e^{i2\pi\vec{u}\cdot\vec{x}}d\vec{x}, \quad (2.23)$$

and \vec{u} and \vec{x} are a Fourier pair. In the small-angle approximation, we have a direct connection back to the l -space of CMB anisotropies, via $l = 2\pi|\vec{u}|$. Figure 2.5 shows the (u, v) coverage of MINT. Thus the MINT ℓ -coverage is centered on windows of 1062, 1634, and 3090 (see Table 2.1).

The real part of the MINT visibility response to a point-source ($B(\vec{x}) \propto$ a δ -function at the center of the receiver beams) is pictured in Figure 2.6. The Fourier relationship between $A(\vec{x})B(\vec{x})$ and $\mathcal{V}(\vec{u})$ is easy to see in light of this diagram, as the fringes pick out a small range of Fourier components on the sky. If $A(\vec{x})$ were uniform on the sky, then there would be no envelope superimposed on the fringes, and the interferometer would be sensitive to exactly one Fourier frequency. This intuition holds in l -space as well, where as larger receiver dishes are used and $A(\vec{x})$ becomes more tightly constrained, the l -coverage becomes broader.

Finally, in this section, the overall normalizations of calculated correlations have been designated “proportional” rather than “equal.” In Chapter 6 we account for

⁷The standard interferometry direction convention, intended for baselines that have fixed positions with respect to the ground during observations, is to use u as the east-west component of \vec{u} , and v for north-south. In MINT, since the baselines move with the (Az, El) motions of the telescope, we assign u to the azimuth direction (u increases with azimuth, from north toward east) and v to elevation, (v increases with elevation).

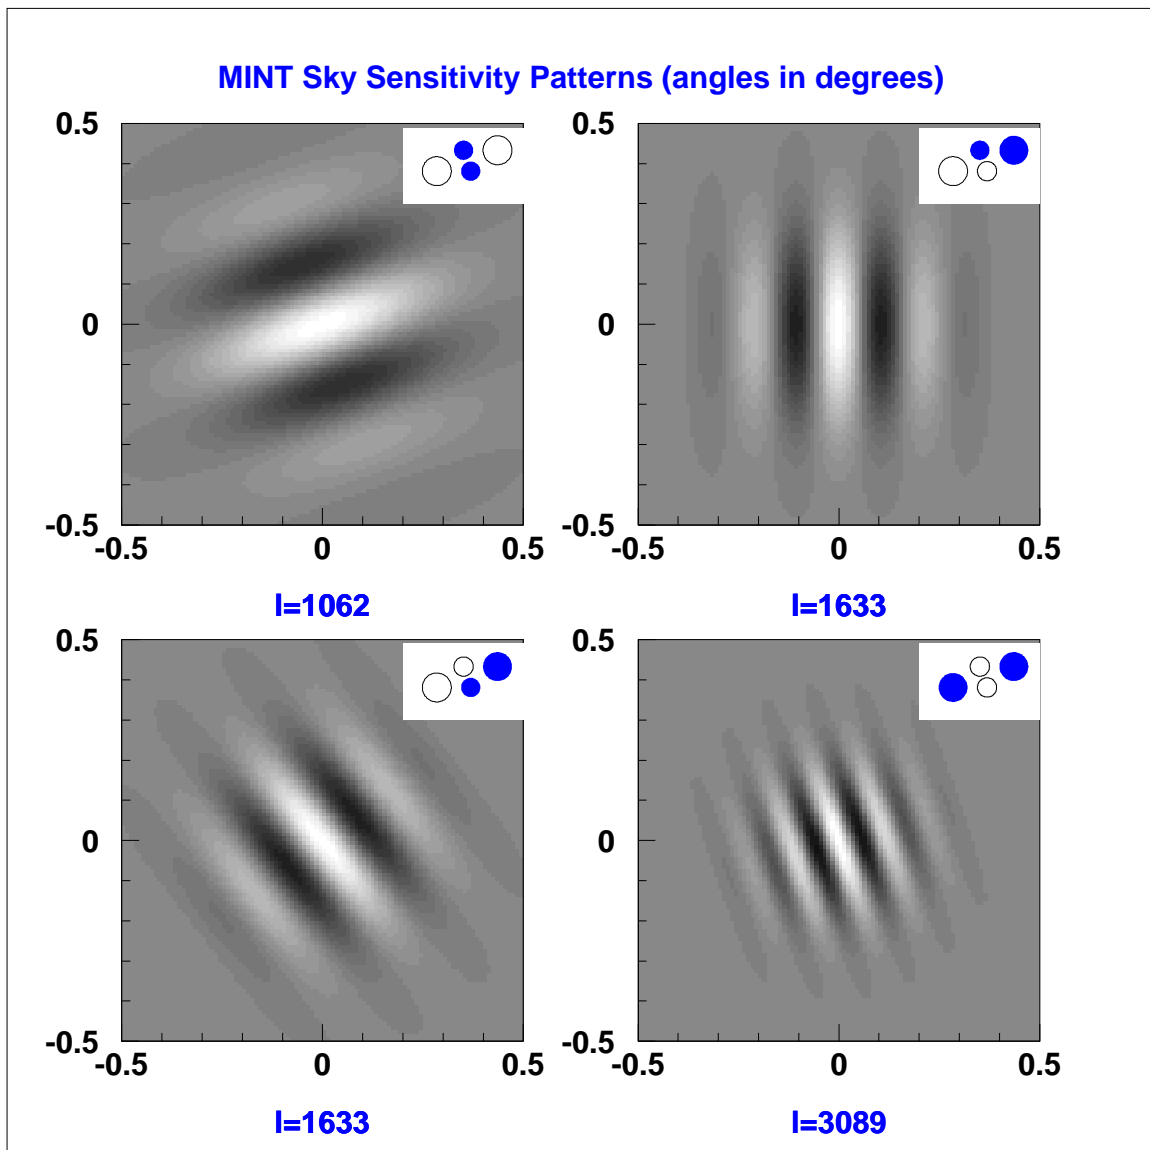


Figure 2.6: The real part of the MINT visibility response to a point source as a function of the point source position. Inset drawing of the baselines (as projected on the sky) has the relevant receivers darkened. Baseline pairs $[(A,B),(D,C)]$ and $[(A,C),(B,D)]$ are parallel, so their point source responses are identical and are not pictured twice. The larger optics (receivers B and C) have primary mirrors 1.5 times as large as the smaller, and so have beams about 1.5 times smaller. Figure created by Joseph Fowler.

the appropriate numerical factors, and compare the telescope measurements with the calculations based on point source observations.

2.3 The Digital Correlator

The MINT correlator is of the lag variety as described in Equation 2.3, designed to make the measurement $r(\tau) = \langle V_A(t)V_B(t - \tau) \rangle$. However, since MINT's correlator is digital, there are two complications. First, the receiver voltages are sampled discretely in time, at a 1-GHz sampling rate. The discrete lags, τ_i , are spaced at 1 nsec intervals from -7 to $+8$ nsec. Second, the voltage amplitudes are converted from a continuous analog value to a quantized digital value. The rest of this section is about the implementation of and analysis of these two effects. The description of the correlator presented here is brief; the reader interested in more details of the correlator design and testing is referred to the Ph.D. thesis of Tran (2002).⁸

In the previous section, receiver outputs were assumed to be noiseless, but in reality, the correlated CMB signal is a paltry few μK buried in a system temperature of 40 K or more. The measurement the correlator actually makes is

$$\rho(V_A(t), V_B(t), \tau) = \frac{\langle V_A(t)V_B(t - \tau) \rangle}{\sqrt{\langle V_A^2(t) \rangle \langle V_B^2(t) \rangle}}, \quad (2.24)$$

which is unitless and ranges from -1 (anti-correlated) to $+1$ (correlated) and except during calibration procedures, is very nearly zero. If V_A and V_B have a bivariate Gaussian probability distribution, the usual assumption, then ρ is the cross-correlation coefficient of that distribution. This is the function that is Fourier transformed and averaged over frequency to yield the complex correlation, and then interpreted as the visibility on the sky.

The correlator design is based on a compromise between fidelity in measuring $\rho(\tau)$ and electronic simplicity. The signals V_A and V_B are Nyquist sampled, with sampling rate of 1 GHz and signal band-limited to 0–500 MHz. Therefore, if the correlator could

⁸The MINT correlator was designed by Huan Tran and Toby Marriage, with CAD help from Stan Chidzik, and is based on a design by David Hawkins of O.V.R.O. Jamie Hinderks (1999) built a 1-bit prototype correlator as his senior thesis project. A full prototype of the digitizer was the senior thesis of Toby Marriage (2000). Charles Dumont (2001), in his senior thesis, provided the first tests of the correlator with two receivers.




Figure 2.7: The digitizers in the MINT correlator are 4-level digitizers, with thresholds of $-V_0$, 0 , and V_0 . We label the bins 0–3, from most negative to most positive. Each channel of each receiver is sampled every nanosecond, producing a raw data stream of 32 Gbps that the correlator must process in real time.

convert the signal voltages with infinitely fine quantization (IFQ), it would measure $\rho(\tau)$ with no loss in sensitivity (this is shown explicitly in TMS, §8.2). Degradation in fidelity due to a finite number of bits per quantization occurs for two reasons (see TMS, §8.3). First, quantization is an inherently non-linear operation, so the digital multiplication-and-average is not strictly proportional to $\rho(\tau)$. Second, since the “sharp edges” of the digitized signal contain new frequency components outside the initial 0-500 MHz band, the digitized signal no longer meets the Nyquist condition. In both cases, more bits are better. MINT uses a four-level (2-bit) digitization scheme (see Figure 2.7) with threshold levels of $-V_0$, 0 , and V_0 , where V_0 is 0.25 V.

Once the signals are digitized, the task of the correlator is to multiply them, with lags digitally inserted, and then to sum the results over time. An efficient multiplication table is developed for the task. The four digitizer levels are assigned the values $(-3, -1, 1, 3)$.⁹ A full multiplication table based on this choice is shown in Table 2.2(a), and can achieve a signal-to-noise ratio (SNR) in determining $\rho(\tau)$ 88 % that of IFQ. This table would take 4 bits plus a sign bit to represent all outcomes. Parts (b) and (c) of Table 2.2 show the progression toward (d), the actual multiplication table that is used in the correlator. Through elimination of the least-significant terms in the multiplication table, this new algorithm (called “deleted inner-product”) gains efficiency, requiring only 3 unsigned bits per multiplication, but still maintains 87 % S.N.R. of the I.F.Q. scheme (only a 1 % drop).

⁹The most general bit-value assignment for a four-level digitizer that reflects the signal distribution’s symmetry about 0 is $(-n, -1, 1, n)$. The value n is typically chosen to be an integer, as the digital math is easier, and then is optimized for best signal-to-noise.

<table border="1" style="border-collapse: collapse; margin: auto;"> <thead> <tr> <th colspan="2" rowspan="2"></th> <th colspan="4">Digitizer level, Rec A.</th> </tr> <tr> <th>0</th> <th>1</th> <th>2</th> <th>3</th> </tr> </thead> <tbody> <tr> <th rowspan="4" style="writing-mode: vertical-rl; transform: rotate(180deg);">Digitizer level, Rec B.</th> <th>0</th> <td>9</td> <td>3</td> <td>-3</td> <td>-9</td> </tr> <tr> <th>1</th> <td>3</td> <td>1</td> <td>-1</td> <td>-3</td> </tr> <tr> <th>2</th> <td>-3</td> <td>-1</td> <td>1</td> <td>3</td> </tr> <tr> <th>3</th> <td>-9</td> <td>-3</td> <td>3</td> <td>9</td> </tr> </tbody> </table> <p style="text-align: center;">(a) Standard Mult. Table</p>			Digitizer level, Rec A.				0	1	2	3	Digitizer level, Rec B.	0	9	3	-3	-9	1	3	1	-1	-3	2	-3	-1	1	3	3	-9	-3	3	9	<table border="1" style="border-collapse: collapse; margin: auto;"> <thead> <tr> <th colspan="2" rowspan="2"></th> <th colspan="4">Digitizer level, Rec A.</th> </tr> <tr> <th>0</th> <th>1</th> <th>2</th> <th>3</th> </tr> </thead> <tbody> <tr> <th rowspan="4" style="writing-mode: vertical-rl; transform: rotate(180deg);">Digitizer level, Rec B.</th> <th>0</th> <td>9</td> <td>3</td> <td>-3</td> <td>-9</td> </tr> <tr> <th>1</th> <td>3</td> <td>0</td> <td>0</td> <td>-3</td> </tr> <tr> <th>2</th> <td>-3</td> <td>0</td> <td>0</td> <td>3</td> </tr> <tr> <th>3</th> <td>-9</td> <td>-3</td> <td>3</td> <td>9</td> </tr> </tbody> </table> <p style="text-align: center;">(b) Deleted Inner Product</p>			Digitizer level, Rec A.				0	1	2	3	Digitizer level, Rec B.	0	9	3	-3	-9	1	3	0	0	-3	2	-3	0	0	3	3	-9	-3	3	9
			Digitizer level, Rec A.																																																												
		0	1	2	3																																																										
Digitizer level, Rec B.	0	9	3	-3	-9																																																										
	1	3	1	-1	-3																																																										
	2	-3	-1	1	3																																																										
	3	-9	-3	3	9																																																										
		Digitizer level, Rec A.																																																													
		0	1	2	3																																																										
Digitizer level, Rec B.	0	9	3	-3	-9																																																										
	1	3	0	0	-3																																																										
	2	-3	0	0	3																																																										
	3	-9	-3	3	9																																																										
<table border="1" style="border-collapse: collapse; margin: auto;"> <thead> <tr> <th colspan="2" rowspan="2"></th> <th colspan="4">Digitizer level, Rec A.</th> </tr> <tr> <th>0</th> <th>1</th> <th>2</th> <th>3</th> </tr> </thead> <tbody> <tr> <th rowspan="4" style="writing-mode: vertical-rl; transform: rotate(180deg);">Digitizer level, Rec B.</th> <th>0</th> <td>3</td> <td>1</td> <td>-1</td> <td>-3</td> </tr> <tr> <th>1</th> <td>1</td> <td>0</td> <td>0</td> <td>-1</td> </tr> <tr> <th>2</th> <td>-1</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <th>3</th> <td>-3</td> <td>-1</td> <td>1</td> <td>3</td> </tr> </tbody> </table> <p style="text-align: center;">(c) Rescaled</p>			Digitizer level, Rec A.				0	1	2	3	Digitizer level, Rec B.	0	3	1	-1	-3	1	1	0	0	-1	2	-1	0	0	1	3	-3	-1	1	3	<table border="1" style="border-collapse: collapse; margin: auto;"> <thead> <tr> <th colspan="2" rowspan="2"></th> <th colspan="4">Digitizer level, Rec A.</th> </tr> <tr> <th>0</th> <th>1</th> <th>2</th> <th>3</th> </tr> </thead> <tbody> <tr> <th rowspan="4" style="writing-mode: vertical-rl; transform: rotate(180deg);">Digitizer level, Rec B.</th> <th>0</th> <td>6</td> <td>4</td> <td>2</td> <td>0</td> </tr> <tr> <th>1</th> <td>4</td> <td>3</td> <td>3</td> <td>2</td> </tr> <tr> <th>2</th> <td>2</td> <td>3</td> <td>3</td> <td>4</td> </tr> <tr> <th>3</th> <td>0</td> <td>2</td> <td>4</td> <td>6</td> </tr> </tbody> </table> <p style="text-align: center;">(d) Biased</p>			Digitizer level, Rec A.				0	1	2	3	Digitizer level, Rec B.	0	6	4	2	0	1	4	3	3	2	2	2	3	3	4	3	0	2	4	6
			Digitizer level, Rec A.																																																												
		0	1	2	3																																																										
Digitizer level, Rec B.	0	3	1	-1	-3																																																										
	1	1	0	0	-1																																																										
	2	-1	0	0	1																																																										
	3	-3	-1	1	3																																																										
		Digitizer level, Rec A.																																																													
		0	1	2	3																																																										
Digitizer level, Rec B.	0	6	4	2	0																																																										
	1	4	3	3	2																																																										
	2	2	3	3	4																																																										
	3	0	2	4	6																																																										

Table 2.2: Development of the multiplication table used in the correlator. (a) A standard multiplication table, with outer digitizer levels given the values ± 3 and inner levels assigned ± 1 . (b) Inner products have been set to 0. (c) The common factor of 3 has been removed. (d) Finally, a bias of 3 has been added, so each multiplication can be represented in three unsigned bits. Using unsigned multiplications is important, because the time-integration step of the correlation operation is an accumulator, which uses a counter that doesn't count backward.

Figure 2.8 shows a sketch of the correlator electronics. The pre-digitizer signal-conditioning circuitry in part (a) allows the signal amplitudes to be kept in a narrow range, as the sensitivity of the measurement depends the rms of each voltage being close to the optimal value of $1.104 V_0$ (see TMS, §8.3). The digitally variable attenuator can assume any value from 0 to 31.5 dB in increments of 0.5 dB. In the field, all digitizer power levels are automatically checked hourly, and the attenuators reset if necessary, to keep the power within 0.5 dB of optimal (resulting in less than 0.3 % loss in sensitivity). The post-digitizer calculations are implemented via a Xilinx

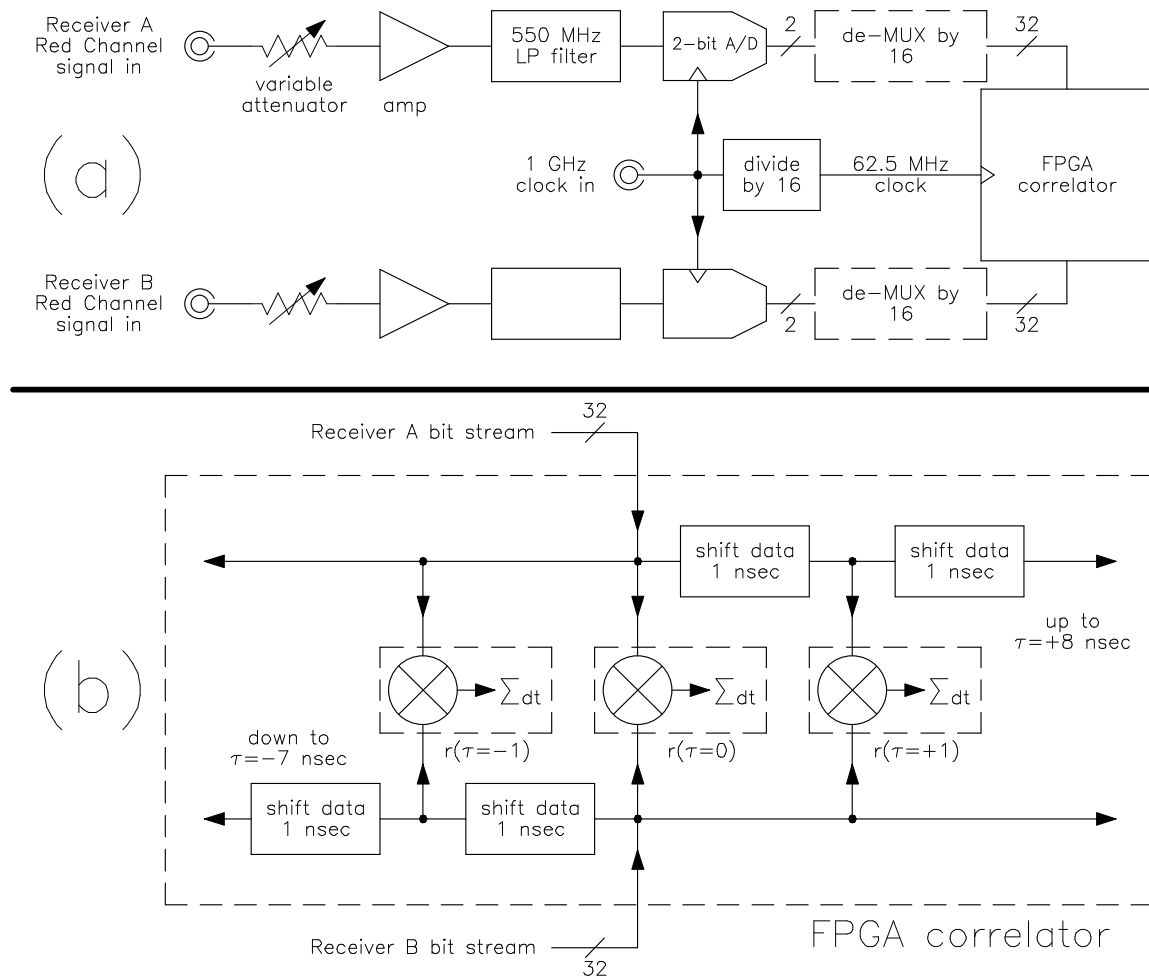


Figure 2.8: A highly simplified block diagram of the MINT 1-GHz digital correlator. (a) The signal-conditioning circuitry and digitizer leading into the FPGA. (b) A sketch of the correlation function calculations in the FPGA. Each multiplication (\otimes) is calculated via the multiplication table in Table 2.2(d). To be fully representative of MINT, the diagram of the correlator would include two more receiver inputs and their interconnecting five additional baselines. In total, the full correlator also processes four independent frequency channels.

Field-Programmable Gate Array (FPGA)¹⁰ which is clocked at 62.5 MHz, so the digitizer must be highly de-multiplexed (parallelized) as it is fed into the FPGA. There are four separate 13" x 13" correlator boards, one for each frequency channel. All four boards are contained in a water-cooled aluminum block housing.

¹⁰The MINT correlator uses a Virtex series, part number XCV-1000, which at more than 10^6 system logic gates, was the largest commercially-available FPGA when the correlator was designed.

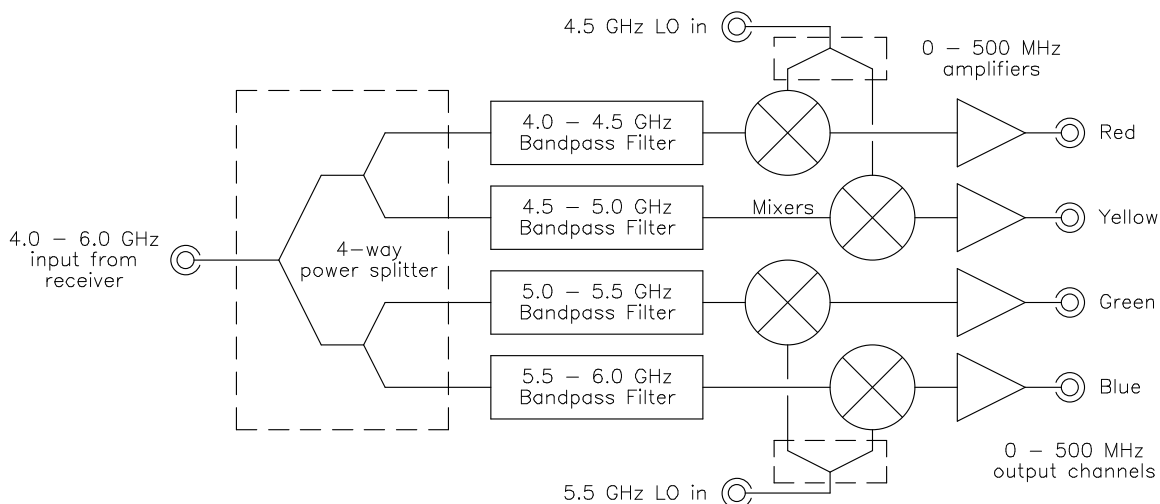


Figure 2.9: Block diagram of the channelizer. The 4-6 GHz IF band is broken down into four 0-500 MHz channels to prepare for digitization. In the full channelizer, there are four blocks like that pictured: one for each receiver.

2.4 The Channelizer

The conversion between the 4-6 GHz receiver outputs and the 0-500 MHz baseband correlator input channels occurs in the channelizer. This custom integrated microwave circuit, designed and built by Huan Tran, is an elegant replacement for what would otherwise be a tangle of discrete, commercial components (filters, amplifiers, mixers, power splitters) joined by potentially unreliable coaxial cables. It consists of a pair of 13.1" x 13.1" microstrip boards with drop-in mixers, amplifiers, and chip resistors. Each board contains two copies of the circuitry shown in Figure 2.9, for a total of four receiver inputs and 16 channel outputs.

Following the signal path through the channelizer in the figure, the output of a receiver (4-6 GHz) is first broken into four equal parts by a cascade of power splitters. Next, each part is passed through a different bandpass filter, defining the four channels of 4.0-4.5 GHz, 4.5-5.0 GHz, 5.0-5.5 GHz, and 5.5-6.0 GHz.¹¹ After filtering, each channel is mixed with an LO of either 4.5 GHz or 5.5 GHz (half the mixers operate in USB mode and half in LSB), down-converting it to the 0-500 MHz baseband. Finally, the signals are amplified, and then passed to the correlator.

¹¹We label these channels Red, Yellow, Green, and Blue, respectively.

2.5 Telescope Pointing

The MINT azimuth and elevation drives are each monitored by a 17-bit, Heidenhain encoder, and are controlled by a LabVIEW servo loop on the pointing PC. Azimuth motion is driven by a motor geared to the azimuth crane bearing.¹² Elevation is driven by a linear actuator, a rotary motor and jack-screw combination that has a pivot point on the azimuth bearing and a pivot point on the elevation cage. During CMB observations, the instrument drift scans (stares at one position and lets the sky rotate through) for 7.5 sidereal minutes at a time, and then makes a small move to another sky location. The motors are turned off between moves to minimize electrical noise during observations; the azimuth bearing contains a brake independent of the motor, and the elevation drive has enough friction to lock the elevation cage in place, even with no current flowing. Azimuth is measured to the east of north. The telescope has a range of motion in azimuth from 60° to 300° , and in elevation from 59° to 100° .¹³ The GPS location of the telescope is:

$$\text{latitude} = 22.9581^\circ \text{ S}, \quad (2.25)$$

$$\text{longitude} = 67.7858^\circ \text{ W}, \quad (2.26)$$

$$\text{altitude} = 17080 \text{ feet above sea level.} \quad (2.27)$$

After leveling the base and zeroing the encoders by mechanical means at the beginning of the campaign, there are still minor deviations from ideal. During the day in Chile, since the instrument is unable to observe the CMB, the dead time is often used to make observations of the Sun that allow the pointing solution to be calculated. Pieces of a corrugated cardboard box with white, thick paper taped to the top are placed over the receiver cross-talk shields (see Figure 3.2) to block the strong optical and IR solar radiation but partially transmit the signal band. The Sun, whose absolute position on the sky is known from ephemerides, is then beam-mapped with the total power channel (see “DC level” in Section 3.4) of each receiver at various positions in the sky. The receivers are found to be aligned relative to one another at the $0^\circ.05$ level, and the absolute azimuth and elevation pointing are found to require

¹²The telescope base and cage were designed by Ted Griffith and built by Laszlo Vargas, Glenn Atkinson, Bill Dix, Ted Lewis, and John Mellodge in the Princeton Physics machine shop.

¹³Elevation is typically only defined up to 90° (zenith). The MINT telescope can look up to 10° backward past zenith.

a few minor corrections. By fitting to the data from solar scans taken through the campaign, Toby Marriage (2002) models a small encoder offset in each axis, as well a slight tilt of all of the receivers in common with respect to the elevation axis. The encoder corrections are:

$$\text{Az encoder offset correction} = -0.31^\circ \quad (2.28)$$

$$\text{El encoder offset correction} = +0.08^\circ. \quad (2.29)$$

The receivers are found to be tilted 0.12° to the East when the azimuth gear is pointed south, resulting in an additional correction that is a function of elevation. No evidence was found of measurable drifts in the overall pointing solution, nor of relative motion of the receivers during the observing campaign to a level of $0^\circ.01$.

A final, smaller pointing correction, discovered during analysis of Mars point-source calibrations, is that the azimuth bearing is tilted so that its West side is slightly higher than its East side. The resulting pointing correction to the elevation is:

$$\text{El correction for Az bearing tilt} = +0.034^\circ \times \sin(\text{Az}). \quad (2.30)$$

Details of the Mars analysis appear in Section 6.5.

2.6 Calibration Noise Source

In order to monitor phase and amplitude stability of the interferometer, a correlated noise injector is integrated into the telescope. A broadband noise source is power-split and distributed to the receivers in coin-silver, D-band waveguide. The waveguide enters the receiver optics chamber through a slot in the bottom of the primary mirror (see Figure 2.10(b)). The signal exits the waveguide, reflects off a small diagonal mirror to the secondary mirror, and then reflects back into the receiver. The signal coupled into each receiver is on average about 5 K.

Since a commercial D-band noise source is not available, we built one.¹⁴ The noise signal begins at Q-band (33-50 GHz) (see Figure 2.10(a)) with a commercial noise source of ENR (excess noise ratio) 20 dB.¹⁵ After amplification, the noise is

¹⁴The noise source was built by Yeong Loh and Asad Aboobaker. Post-campaign tests of the noise source are described in Aboobaker (2002).

¹⁵ENR is defined as the ratio of the noise temperature to the physical temperature, or 290 K.




Figure 2.10: (a) A block diagram of the broadband calibration noise source. The Q-band amplifiers have gain only below 47.3 GHz, meaning the noise source is lower-sideband only. The mechanical rotary switch is needed to ensure that the noise source is completely turned off. (b) Signal path of the noise source into the receiver. The waveguide runs in a notch cut into the bottom of the mirror.

tripled in frequency. A rotary waveguide switch transmits the signal to the power splitters, which feed each receiver. The Q-band amplifier pair exhibits a sharp gain roll-off above 47.3 GHz, resulting in a frequency-tripled noise signal that is visible only in the lower sideband of the receivers (see Figure 2.2).

Design of the Receivers

3.1 Introduction

The use of superconductor-insulator-superconductor-based (SIS) receivers in CMB experiments at Princeton is a long tradition. Timbie and Wilkinson (1984) used SIS-based receivers in a two-element, 46-GHz interferometer. More recently, the MAT/TOCO experiment (Miller et al. 2002), a chopping radiometer, had an SIS-based channel centered at 144 GHz. Both of these experiments used mixers designed by Tony Kerr of the NRAO (Feldman et al. 1983; Kerr et al. 1993).

Continuing in this tradition, the four MINT receivers, which are described in this chapter, are built around SIS mixers. They operate in double-sideband (DSB) mode with a primary local oscillator (LO) frequency of 145.1 GHz and an intermediate frequency (IF) band of 4.0–6.0 GHz. The optics for the four receivers are described in Section 3.2. Section 3.3 details the cryogenically-cooled front end of the receiver, including the SIS and the HEMT IF amplifier. The receiver back end (RBE), containing warm amplification and support electronics, is presented in Section 3.4. Cryogenics, including efforts to stabilize the temperature of the SIS block, are described in Section 3.5. The MINT receiver is of essentially the same design as that used in the MAT/TOCO experiment, using the same SIS's and HEMT amplifiers from NRAO; on several occasions we refer the reader to the Ph.D. thesis of Miller (2000), where details of many of the components have been described in exhaustive detail.¹ The MAT/TOCO D-band receiver worked well in the 1998 campaign, but did not achieve

¹The MAT/TOCO used a slightly different LO frequency, 144 GHz, than MINT's 145.1 GHz. Otherwise, the MINT and MAT/TOCO receivers use the same parts.




Figure 3.1: Sketch of a typical Cassegrain antenna. The primary is a paraboloid, and the secondary a hyperboloid. The two diamonds mark focal points. The dashed line is a geometric ray trace.

the expected sensitivity. Here, we emphasize improvements in the design of the MINT receivers based on the experience garnered from the MAT/TOCO experiment.

More specialized topics relating to the receivers have been given their own chapters. The SIS tuning process and measurements of sensitivity, both in the lab and in the field, are found in Chapter 4. The LO system is described in Chapter 5.

3.2 Optics

Each MINT receiver employs an on-axis, Cassegrain optics system, with a concave-paraboloid primary mirror, a small convex-hyperboloid secondary mirror, and a cold, corrugated feed horn. A focused Cassegrain system has the secondary mirror positioned so that one of its foci coincides with the focal point of the primary. The other focus of the secondary then becomes the focal point of the combined system (see Figure 3.1), which in MINT is designed to coincide with the aperture of the feedhorn.

The important characteristics of the MINT optics² are listed in Table 3.1. MINT is a *heterogeneous array*, meaning that some receivers have different optics than others. There are two distinct optical designs. Receivers A and D have 30 cm-diameter primary mirrors, the biggest they could be without having their flared radiation shields touch. Figure 3.2 shows a cross-sectional view of optics A and D. Each secondary is held in place by a spider of three G-10 support legs spaced 120° apart, with each

²The primary mirror, secondary mirror, conical radiation shield, and G-10 support spider were designed by Yeong Loh and built in the Princeton Physics Department's research machine shop. Many optics calculations and design details are found in Loh (2000a,b).

Parameter	Sm. Optics (A & D)	Lg. Optics (B & C)
Primary Diameter	30 cm	45 cm
Primary Focal Length	12 cm	18 cm
Secondary Diameter	6.503 cm	9.754 cm
Secondary Inter-focal Length	13 cm	19.5 cm
Secondary Vertex-Center Distance	4.126 cm	6.189 cm
Top Feed Horn to Primary Vertex	1 cm	1.5 cm
Feed Horn Aperture Diameter	0.89 cm	0.89 cm
Feed Horn beam FWHM	21°	21°
Hole in Primary, Diameter	4.425 cm	4.425 cm
Maximum Gain (simulated)	51.3 dBi	54.8 dBi
FWHM (simulated)	0.44°	0.30°

Table 3.1: Parameters of the Cassegrain optics. The big optics are 1.5 times the linear size of the small optics. The feed horn is produced by Custom Microwave Inc., Longmont, CO, and its quoted beam is from the optics simulation program CCHORN, written by YRS Associates, Sherman Oaks, CA. Details of the feed horn are found in Miller (2000).

leg presenting only a 1/16"-wide cross-section to the primary. The radiation shields, intended to prevent cross-talk between receivers, are rolled at the top with a radius of about 1 cm (5λ) to limit diffraction. The primary and secondary mirrors are machined from solid aluminum using a computer-numerically-controlled (CNC) lathe. Their surfaces are accurate to 0.001" and are smooth to about 0.0001". The optics for receivers B and C are scaled copies of A and D, expanded by a factor of 1.5 about the top of the feed horn (the focal point of the Cassegrain optics) in every linear dimension. Only the feed horn itself and the hole in the primary mirror remain the same size. The surfaces of the mirrors are defined by the following equations in cylindrical coordinates ($+z$ is up, the top of the feedhorn is $z = 0$, and coordinates are in cm). For the small primary and secondary mirrors:

$$z_p = 1 + \frac{r^2}{48}, \text{ and} \quad (3.1)$$

$$z_s = 6.5 + 4.126\sqrt{1 + \frac{r^2}{25.226}}, \quad (3.2)$$

and for the bigger ones:

$$z_p = 1.5 + \frac{r^2}{72}, \text{ and} \quad (3.3)$$




Figure 3.2: To-scale cross section of the A and D (smaller) optics.

$$z_s = 9.75 + 6.189\sqrt{1 + \frac{r^2}{56.759}}. \quad (3.4)$$

Simulated beam patterns³ for the two sets of optics are shown in Figure 3.3. The large optics have a maximum gain of 54.8 dBi (dB above isotropic) and a FWHM of 0.30°, while the small optics have a maximum gain of 51.3 dBi and a FWHM of 0.44°. At the time of this writing, detailed beam-maps do not exist for either the small or large optics in their current configurations. However, as a test of the reliability of the simulations, an earlier version of the smaller optics was beam-mapped. The data and their comparison to the simulated beam map are shown in Figure 3.4. The measurements compare well to the simulations, especially in the important main lobe

³Optics simulations are from the program DADRA, *Diffraction Analysis of a Dual Reflector Antenna*, written by YRS Associates, Sherman Oaks, CA.




Figure 3.3: Simulated beam patterns for the large (solid) and small (dashed) optics.




Figure 3.4: Beam map (solid) and simulated beam pattern (dashed) for an older version of the small optics. Beam map results have been scaled to match the simulation at the peak, as the experimental apparatus measures only relative gain.




Figure 3.5: Path of signals in the receiver front end.

and first sidelobe, giving us confidence to trust the unmapped optics for use in the 2001 campaign.

3.3 The Receiver Front End

The SIS mixer and HEMT amplifier, contained in a cryogenic vacuum enclosure, or dewar, make up the receiver front end. A diagram of the front end is shown in Figure 3.5. Radiation enters the dewar through the 0.022”-thick polypropylene vacuum window. The vacuum window is kept free of ice via a forced-dry-air loop, sealed with shrink-fit window insulation⁴ over the radiation shield (see Figure 3.2). The quartz window⁵ is used as an infrared blocker to reduce radiative loading on the SIS junction. The thicknesses of both materials are chosen for maximum transmission at 2.1 mm. With a spacing of 0.082” between the two windows, transmission of the pair is optimized at $95 \pm 2\%$. A further improvement over the TOCO receiver is that the quartz window is thermally isolated from the feedhorn, so that IR and optical loading can change the quartz temperature without directly affecting the SIS stage. Next, the signal enters the corrugated feed horn, is passed through a circular-to-linear (C-L) polarization converter,⁶ and then enters WR-5.8 waveguide. Since the rectangular waveguide carries only a single linear-polarization mode, the C-L

⁴The material used is Dennis Shrink-To-Fit Window Insulation Kit, 2506, W. J. Dennis & Co., Elgin, Illinois.

⁵The z-cut quartz, 0.0205” thick, 0.700” diameter, is produced by V. A. Optical Labs, Inc., San Anselmo, CA.

⁶Built by Custom Microwave Inc., Longmont, CO.




Figure 3.6: The frequency mapping between RF and IF signals. The SIS operates in DSB mode.

converter gives the receiver sensitivity to circularly polarized signals at its input.⁷ A circularly polarized receiver is preferable to a linear one because the signal doesn't change as the telescope rotates with respect to the sky, and also because many sources of spurious signals, such as the G-10 spider legs or the ground, might be expected to be somewhat linearly polarized. Next, the signal enters the 20 dB branch-line coupler,⁸ where the LO signal also enters. A 20-dB coupler is necessary to attenuate thermal noise (~ 50 K) that accompanies the LO, and also preserves 99% of the sky signal. The signal and LO enter the SIS, where they are mixed. Now at intermediate frequency (IF), the signal passes through the bias-T, which couples the bias voltage into the SIS, then an isolator, and finally is amplified by the HEMT amplifier. The amplified IF signal is carried out of the dewar on a stainless-steel, coaxial cable to the receiver back end.

The relationship between the receiver RF and IF signal bands is shown in Figure 3.6. The SIS is a double-sideband device, and has roughly balanced power in the two sidebands. We define 4-6 GHz as the MINT IF band, because it is this portion of the mixed-down signal that is eventually measured by the correlator. However, the inherent IF band of the receiver, as measured by Miller (2000), is substantially wider, with measurable gain in a range of 3.5-7.0 GHz. The shape of the inherent IF bandwidth is primarily determined by the IF components, namely the HEMT and the bias-T.

⁷The receivers are sensitive to only one handed-ness, either RCP or LCP. We actually don't know which, but it doesn't matter, since all four receivers use identical converters.

⁸The branch-line coupler is designed by NRAO and built by Zax Millimeter Wave Corporation, San Dimas, CA.

Receiver	HEMT Serial #	stage	V_D (V)	I_D (V)	V_G (V)
A	C-6	1	2.000	0.401	-0.089
		2	2.001	0.702	-0.315
		3	2.000	0.700	-0.242
B	C-23	1	1.749	0.501	-0.505
		2	2.004	0.903	-0.330
		3	2.005	0.904	-0.328
C	C-63	1	1.124	0.639	-0.320
		2	2.001	1.061	+0.120
		3	2.003	1.058	+0.074
D	C-27	1	2.000	0.400	-0.007
		2	2.000	0.901	-0.362
		3	1.999	0.899	-0.348

Table 3.2: HEMT bias parameters for the 2001 observing campaign. For ease of readout, I_D is converted to a voltage with a $300\ \Omega$ resistor; the true current in amps is recovered by dividing the readout voltage by 300. V_D and I_D are held fixed, and V_G is monitored as a health check of the amplifier stage. For comparison to Miller (2000), HEMT C-6 was used in MAT/TOCO channel D1, and C-27 in D2.

The NRAO SIS is a 6-junction device, with layers of superconducting Nb and insulating Al-Al₂O₃. The design includes three tuning parameters: an adjustable mechanical backshort, the SIS junction bias voltage (V_j), and the LO power level. The sliding backshort is coupled to a micrometer and is controlled by a rotary feedthrough mounted on the outside of the dewar. V_j is controlled by a servo card in the RBE, and can be supplied externally (for tuning in the lab), or from an internal reference (for observations in the field). LO power is adjusted via a D-band variable attenuator in the RBE. A tuning algorithm based on these three parameters is described in Chapter 4.

The NRAO HEMT is a three-stage amplifier, with gain of ~ 33 dB and a noise temperature of about 5–8 K. The HEMT, like the SIS, has tuning parameters, namely the drain voltage (V_D) and drain current (I_D) of each stage, which are controlled by a servo card in the RBE. NRAO tested each device before shipment and gave suggested tuning values, which we used in the field (Table 3.2).

3.4 The Receiver Back End

Each of the four RBE's⁹ contains the warm support electronics for its receiver. The subsystems within each RBE (see Figure 3.7) are the local oscillator, cryogenic thermometry and temperature servo, warm IF amplification, SIS bias control, and HEMT bias control.

Great care is taken to control sources of systematic error in the RBE. The inner components are contained in two nested, metal cages, with a layer of Styrofoam insulation in between. Temperature is controlled via a combination of circulating antifreeze (whose temperature is itself regulated) and a heater system with a proportional-integral servo loop. During the campaign, the temperature of each RBE is set to 303 K and remains constant to ± 0.1 K during observations. The ground of the high-current RBE heater is decoupled from case ground to avoid large ground currents mingling with sensitive signals. In addition, the box is well shielded from RF pickup. Every component in the RBE except the RBE heater card receives its DC power from a voltage regulator contained in the RBE itself.¹⁰ All signal wires to and from the outside are passed through pi-filter feedthroughs on the wall of the inner cage to prevent high-frequency pickup. Both cages are sealed with tape that has conductive glue.

The IF signal enters the RBE from the dewar through a hermetic co-axial feedthrough and is amplified by a warm amplifier. Due to budgetary constraints, the RBE's do not all use the same model amplifiers. After power splitting, the signal power level is detected and the remaining signal exits the RBE, in some cases after a second round of amplification. The power-detection signal, called the "DC level,"¹¹ is not part of the interferometer signal, but is used as a diagnostic for monitoring the atmosphere and the health of the receiver. The IF signal that exits the RBE goes to the channelizer to be broken into bands and mixed down to the baseband of 0–500 MHz before digitization (see Chapter 2).

⁹Two undergraduate senior theses were based on assembling and testing RBE's. Daniel Wesley (2000) built RBE D, and Charles Dumont (2001) built RBE's B and C.

¹⁰The PLL box and the LO contain their own voltage regulators. All other components receive power from the regulators in the strip at the bottom of Figure 3.7.

¹¹"DC level" is a legacy terminology from the MAT/TOCO experiment, which measured total power in two different ways: a modulated power signal due to the chopping radiometer, and the absolute, or DC power level, measured after a low-pass filter.

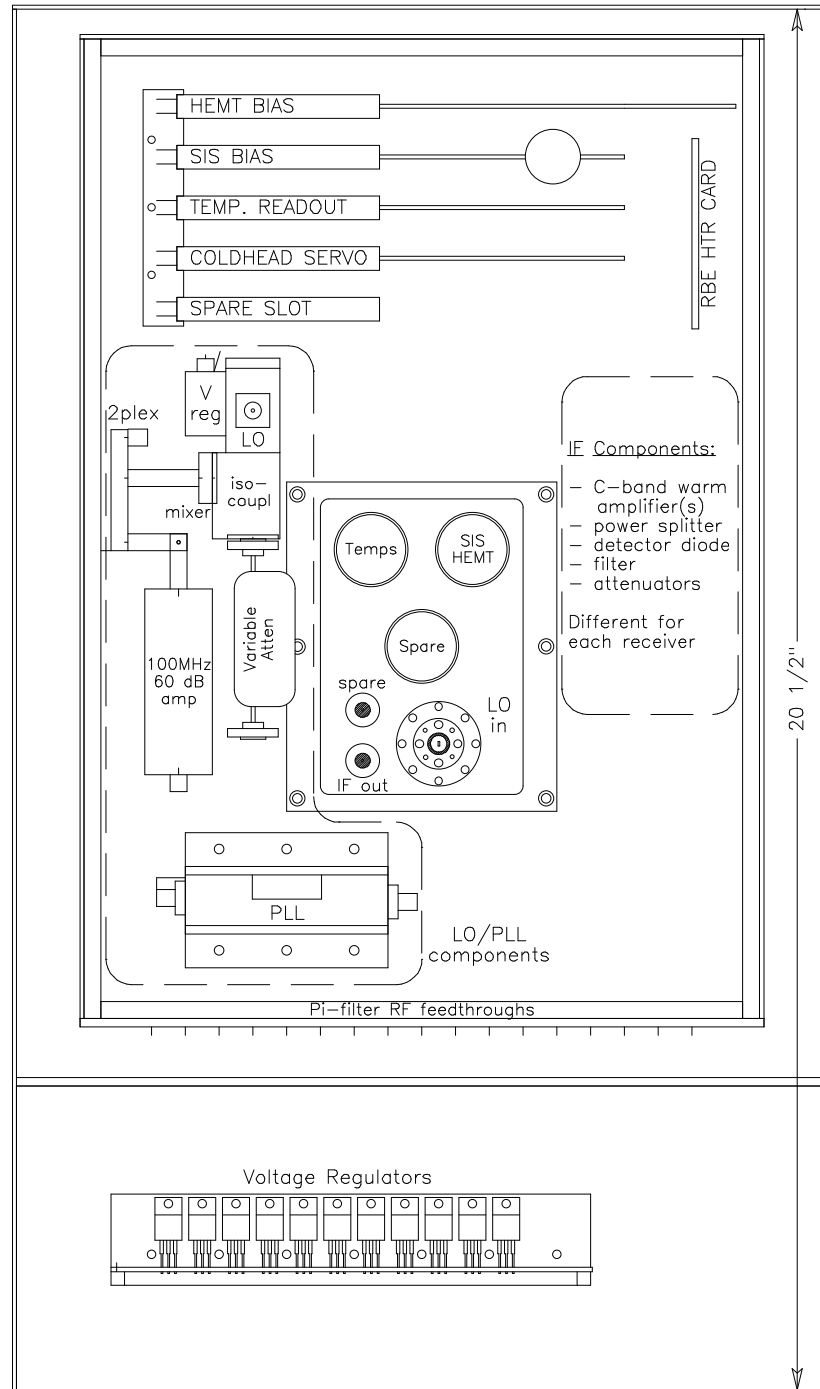


Figure 3.7: Scale drawing of the components in a MINT RBE. Each receiver uses a different model of warm IF amplifier, so they are not pictured (see Figure 3.8). Also not pictured are the external connections to the RBE: on the panel above the voltage regulators are about 30 BNC jacks and some switches. On the right are SMA cables for IF signal output, two LO-PLL reference oscillator inputs, and LO IF-monitor output, as well as cooling water input and output tubes.

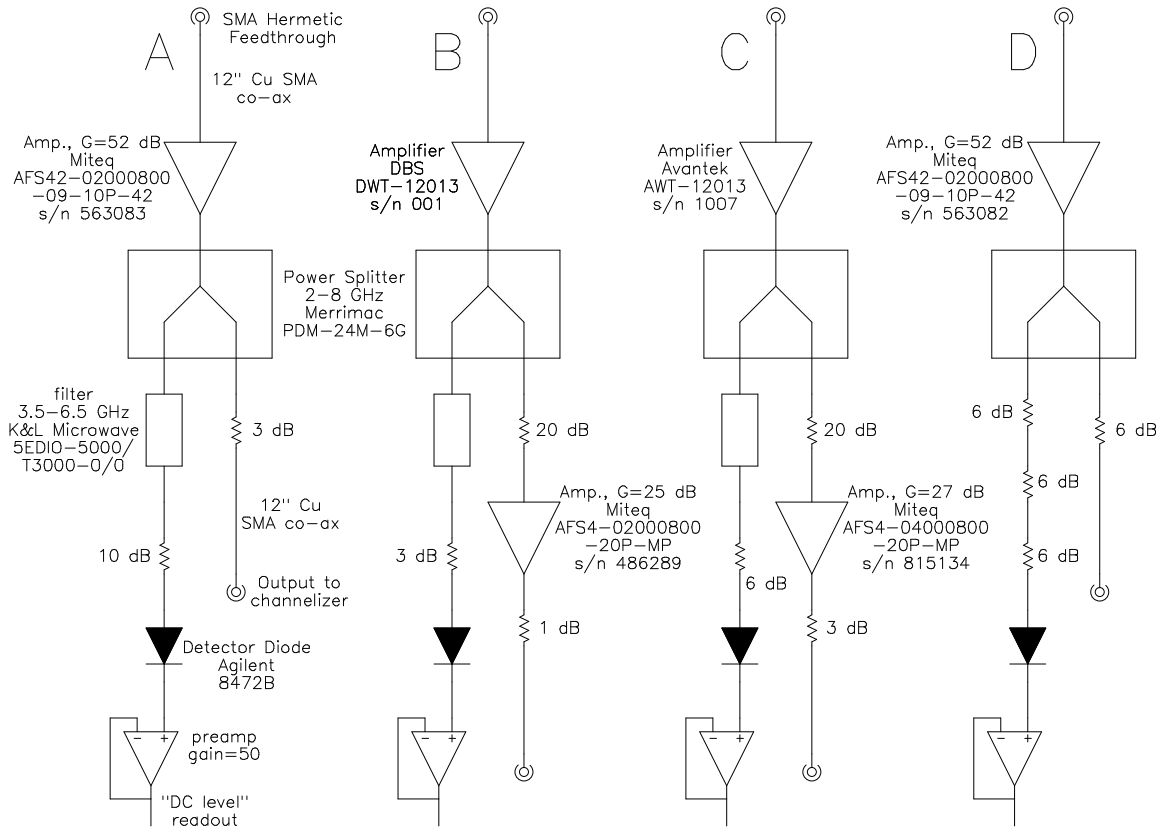


Figure 3.8: The four RBE IF receiver chains. Although all receivers are different, an effort is made to keep the pairs (A,D) and (B,C), which have identical optics, phase-matched by using the same model amplifiers and the same model and number of attenuators.

3.5 Cryogenics

Each MINT receiver is cooled by a three-stage Gifford-McMahon (GM) refrigerator designed by Richard Plambeck of the University of California, Berkeley (Plambeck et al. 1993).¹² The Plambeck design, in use for several years on the Berkeley-Illinois-Maryland Array (BIMA), involves the addition of a third stage to a standard two-stage CTI 1020 CP coldhead.¹³ Each coldhead has its own small Helium compressor, a 3-phase, air-cooled CTI 8200 which is designed to run the much smaller CTI 350 coldhead. In order to accommodate the telescope's cable wrap, the coldheads are linked to their compressors via 40' flex-lines. We report no loss of cooling capacity

¹²The MINT coldheads are built commercially by Cryostar Associates, Yuma, AZ.

¹³CTI is a division of Helix Technology Corporation, Mansfield, MA.

Receiver	SIS	HEMT	1st Stage
	T _{cryo2}	T _{cryo3}	T _{cryo5}
A	3.7	12.8	50.9
B	4.1	13.6	54.7
C	3.6	—	65.4
D	3.6	13.6	50.0

Table 3.3: Averages of several cryogenic temperatures for the MINT receivers for the 2001 observing campaign. The SIS temperature is actively controlled with a servo heater. Receiver C’s HEMT thermometer was not read out correctly through the data stream. The Lakeshore diodes (standard curve 10) used as thermometers are accurate to ± 0.2 K.

despite these long flex-lines.¹⁴ Table 3.3 shows the average temperature of the SIS (servoed, 3rd stage), HEMT (2nd stage) and 1st stage of each receiver for the 2001 observing campaign.

One of the key features of the Plambeck design is that the coldhead motor cycles at less than half the speed of a standard GM refrigerator. CTI coldheads and the various available clones are driven at 72 RPM by a synchronous motor that expects a 60-Hz 2-phase, quadrature feed.¹⁵ Changing the motor speed requires special equipment. One solution, a switching, AC-frequency converter,¹⁶ is undesirable for two reasons: first, the output amplitude is not adjustable; and second, the converter broadcasts so much interference at its switching frequency (~ 14 kHz) that it prevents nearby computer monitors from being read! A better solution is a master-slave pair of analog-oscillator-driven linear amplifiers¹⁷ with both variable speed and amplitude control and a faithful, sinusoidal output. With this pair of drivers, which powers the motors in all four coldheads, we report quiet operation at 25 Hz line frequency with a drive

¹⁴There was some thought of using a “ballast tank” to aid the undersized compressor. The ballast tank was not found to improve the performance of a coldhead in the lab, and so was not used during the campaign.

¹⁵There are 50 steps per cycle of the synchronous motor, so the motor’s cycle is $50\times$ slower than that of the AC line that feeds it.

¹⁶We tested the Bronco BAC16108 from Warner Electric, Bristol, CT.

¹⁷We settled on a pair of Powertron 500S motor drivers, custom-configured as a master and a quadrature slave, and produced by Industrial Test Equipment, Port Washington, NY.




Figure 3.9: Comparison between an electronic low-pass filter and a filter for thermal variations. Electrical resistance is analogous to resistance to heat flow, and electrical capacitance is analogous to heat capacity.

amplitude of 45 V (rms) in each leg. A significantly higher drive voltage makes the motor buzz.¹⁸

An undesirable side effect of running the refrigerator slowly is a large, cyclic temperature swing on the third stage. We measure 390 mK peak-to-peak variations on an unloaded cold finger in an empty dewar that, if coupled directly to the SIS, would cause an intolerable variation in the response of the receiver. To combat this, the third stage employs a two-tiered system, using a specially-designed, high-heat-capacity block and an active electronic servo loop. Using a combination of heat capacity and thermal resistance, it is possible to create a low-pass filter for thermal variations. Figure 3.9 shows the idea. With a heat capacity, C (J/K), in series with a thermal resistivity, R (K/W), we have a direct analogy to an electronic, passive, one-pole, low-pass filter with $f_{3dB} = 1/(2\pi RC)$. Since the refrigerator cycle and its associated thermal variations have frequency 0.5 Hz, the condition for this “circuit” to attenuate the variations is $RC \gg 1$ sec. This is a difficult combination of R and C to attain. Any net heat flow through R in Figure 3.9 causes an undesirable average temperature difference between the left and right blocks, given by $T_2 - T_1 = \Delta T = RP$, where P is the power flowing through R . In designing the cold stage, we decide we can tolerate up to a 0.2 K drop across this joint with an estimated 20 mW, largely due to the servo heater, being drawn away by the cold finger, giving $R \leq 10$ K/W.

With this R , to attenuate the fluctuations from 390 mK p-p down to 10 mK p-p requires $C \geq 1.2$ J/K. A heat capacity of this size corresponds to more than 0.5 kg

¹⁸We give special thanks to Paul Amundsen at CTI for lending us the coldhead motor that made these tests possible.




Figure 3.10: Utilizing the heat capacity of stainless steel to damp thermal variations. (a) A thin slab of area, A , and thickness, L , in contact with a thermally fluctuating heat bath. (b) The MINT heat capacity block. The various copper and stainless-steel sections are joined by nickel brazing. Heat flows vertically in the diagram, across the faces of the stainless steel bars.

of stainless steel, the metal among the common machining materials that has by far the highest specific heat ($C_p = 2.2 \text{ J/kg-K}$) at 4 K.¹⁹

The poor thermal conductivity of stainless steel causes further complications to building a usable thermal capacitor. A block of stainless with large R and C acts as a low-pass filter to transmission from one side to the other, meaning that thermal variations of high frequency don't penetrate more than a certain depth into the block. We refer to Figure 3.10(a) for the following argument. A slab of thickness, L , and cross-sectional area, A , is connected to a heat bath that fluctuates thermally with frequency, f . The heat capacity of the slab is given by $C = C_p \rho A L$, and the thermal resistance from one face of the slab to the other is $R = L/(A\lambda)$, where λ (W/cm-K) is the thermal conductivity of the material. If we make the approximation that the R and C of the slab are in series, as in Figure 3.9, then the thermal time constant for a fluctuation to be communicated to the entire slab is given by

$$t_{RC} = RC = \frac{C_p \rho L^2}{\lambda}. \quad (3.5)$$

L is somewhat analogous to the “skin depth” of electromagnetic wave propagation in an imperfect conductor (Jackson 1975), and in fact has the same frequency dependence, $L \propto 1/\sqrt{f}$. Using $C_p = 2.2 \text{ J/kg-K}$, $\rho = 7.9 \text{ g/cm}^3$, and $\lambda = 2.4 \text{ mW/cm-K}$ for stainless steel 304 at 4 K, we find $t_{RC} = (7.2 \text{ sec/cm}^2)L^2$. The MINT refrigerators

¹⁹This and subsequent cryogenic data concerning C_p , λ , and ρ are taken from Touloukian and Buyco (1970).




Figure 3.11: Block diagram of the full third-stage cryogenic system. Peak-to-peak thermal variations of a refrigerator cycle, given at several stages along the chain, are from measurements on the fully-assembled receiver A.

run with a period of 2 sec, so in order for a piece of stainless to fully communicate a fluctuation to its entire thickness, and thus utilize its entire C , it must have t_{RC} much less than this, say 0.2 sec. The penetration depth of fluctuations of period 0.2 sec is $L = 0.16$ cm, or about $1/16''$.

Figure 3.10(b) shows the design of the MINT heat capacity block. Copper is used as a conductor because of its high thermal conductivity of $\lambda = 3$ W/cm-K at 4 K. The cold finger of the refrigerator is attached to the copper base plate. Since each piece of stainless is connected to the copper conductors (heat bath) on both faces, we increase the slab thickness to double the penetration depth calculated above, or $1/8''$. The whole block is welded together through a process called *nickel brazing*, in which the piece is heated in a hydrogen atmosphere to avoid oxidizing the stainless steel and copper surfaces. In total, the block contains roughly 32" of $1/8'' \times 1''$ stainless steel 304 bar stock, giving a heat capacity of $C = 1.1$ J/K.

A diagram of the full MINT third stage cryogenics is found in Figure 3.11. The cold finger of the refrigerator is coupled to the copper and stainless heat-capacity block via a thin stainless wafer that serves as a thermal resistor ($R_1 = 5$ K/W). Another stainless wafer, R_2 , of the same resistance, connects the SIS stage to the heat-capacity block. Due to contact resistance in joints, the resistance of each of these two wafers is likely enhanced by about 2 K/W. The SIS mixer block and other associated RF and IF receiver parts (see Section 3.3) are attached directly to the SIS block. The discrepancy between the 390 mK variations with a naked cold finger and the 160 mK measured on the cold finger in the fully assembled dewar is reasonable; the cold finger is an imperfect heat bath and is weakly coupled to the heat capacity

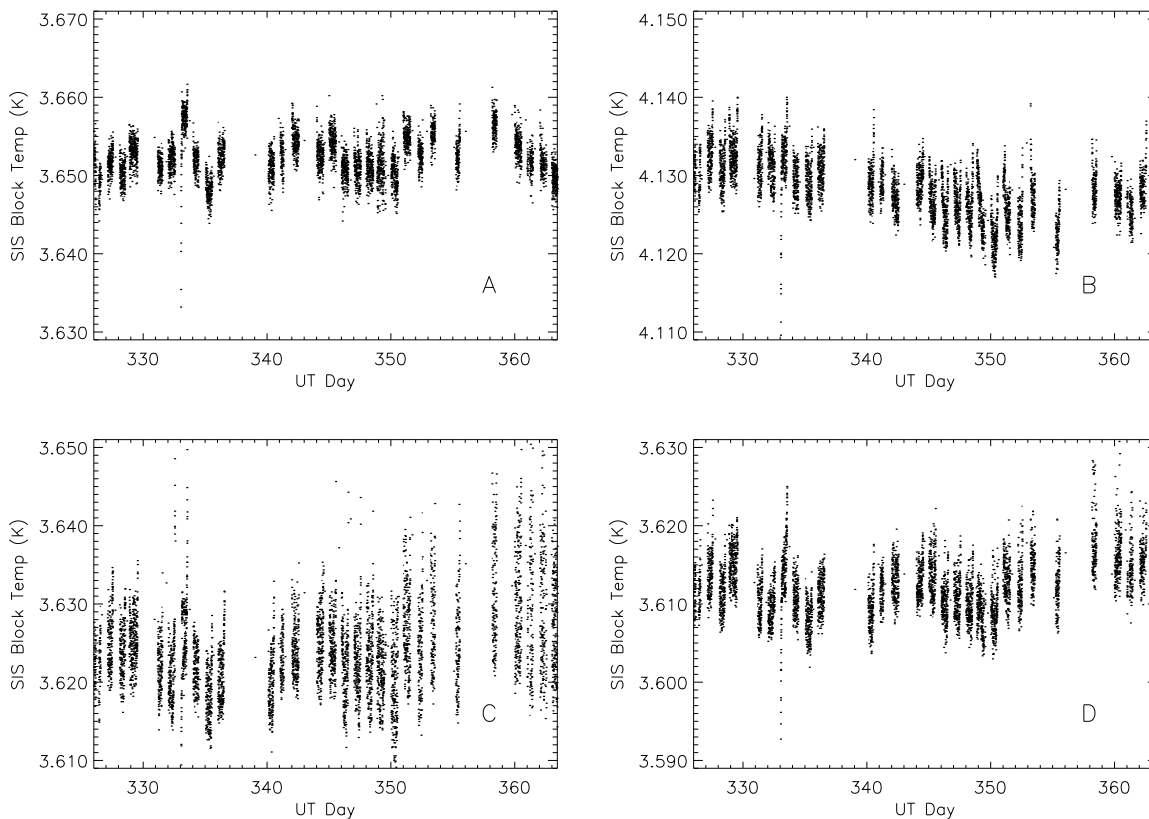


Figure 3.12: SIS-block temperature for each of the four receivers through the whole 2001 observing campaign. All graphs have the same scale.

block. The low-pass filter of R_1 and C does its job, reducing the fluctuations by another order of magnitude. A Silicon-diode thermometer and a heater resistor are attached to the copper SIS stage, and the temperature of the block is served with a PID loop designed by undergraduate Mark Tygert. The thermal resistor R_2 acts to decouple the heater resistor from the heat capacity block, making the job of the heater much easier, as the heat capacity of the SIS stage is small.

The set-point temperature of each receiver is chosen so that the power delivered to a cold finger by the servo heater, averaged over a cycle, is about 18-20 mW. With the servos running, the four SIS's see p-p temperature fluctuations of only 4.5, 8, 10, and 8 mK (A, B, C, and D, respectively) within a refrigerator cycle. Figure 3.12 shows each receiver's SIS temperature for the whole campaign. The servo loop works well over the long term, keeping three of the receivers in narrow 15 mK ranges and the fourth within a 40 mK range over the entire 40-day observing campaign despite

large variations in the daily operating conditions. This outstanding temperature stability is responsible for the single biggest improvement in receiver sensitivity over the MAT/TOCO receiver, whose SIS's wandered in physical temperature by up to 1 K during the 1998 campaign (Miller 2000). Furthermore, as we will see in Chapter 4, the ability to keep the SIS at the same temperature looking at the sky as when looking at a cold load in a closed dewar is a key to accurate tuning.

SIS Tuning and Receiver Sensitivity Measurements

4.1 Introduction: Definitions and Conventions

MINT, like all cosmic microwave background anisotropy experiments, measures signals so small ($\sim 30 \mu\text{K}$) that they are detected only by finding excess noise through statistical means after time-integration of a large data set; therefore we need a good understanding of the noise in the measurements. The uncertainty in a radiometric temperature measurement, ΔT , is given by the Dicke radiometer equation:

$$\Delta T = \frac{T_{sys}}{\sqrt{t\Delta\nu}} . \quad (4.1)$$

In this equation, $\Delta\nu$ is the bandwidth of the receiver, t is the observing time, and T_{sys} is the system temperature, the equivalent input noise power added by the receiver, expressed as a temperature. This chapter discusses methods of measuring T_{sys} and $\Delta\nu$, and tuning the receivers to minimize ΔT . In addition, we present a method to calculate T_{sys} in all of MINT's receiver channels during the observing campaign. T_{sys} is the calibration factor for converting the unitless correlator output to a sky visibility that has temperature units; this process is detailed in Chapter 6.

Receiver noise is typically measured via the “Y-factor” method, which is based on the ratio of the receiver outputs (usually the voltage from a square-law, power-detecting diode) when looking at hot and cold blackbody loads:¹

$$Y = \frac{V_{hot}}{V_{cold}}. \quad (4.2)$$

The receiver output is proportional to the sum of the power due to the load and the intrinsic noise power of the receiver, i.e.:

$$V_{hot} = G(P_{hot} + P_{rec}), \quad (4.3)$$

$$V_{cold} = G(P_{cold} + P_{rec}). \quad (4.4)$$

Solving Equations 4.2–4.4 for the receiver noise power gives

$$P_{rec} = \frac{P_{hot} - Y P_{cold}}{Y - 1}. \quad (4.5)$$

The noise power in a single waveguide mode due to a thermal load is expressed in terms of its physical temperature, T_{ph} , by the Callen & Welton (C-W) dissipation-fluctuation theorem (Callen and Welton 1951; Kerr et al. 1997):²

$$P_{C-W}(T_{ph}) = h\nu\Delta\nu \left(\frac{1}{e^{\frac{h\nu}{k_B T_{ph}}} - 1} + \frac{1}{2} \right). \quad (4.6)$$

We can represent *any* noise power, regardless of assumptions about its spectrum, by defining an equivalent noise temperature,

$$T_{noise} = \frac{P_{noise}}{k_B \Delta\nu}. \quad (4.7)$$

Combining Equations 4.6 and 4.7, we get an expression for the C-W noise temperature of an emitter as a function of its physical temperature:

$$T_{C-W} = \frac{h\nu}{k_B} \left(\frac{1}{e^{\frac{h\nu}{k_B T_{ph}}} - 1} + \frac{1}{2} \right). \quad (4.8)$$

¹Often, the load is a microwave absorber, viewed at room temperature (295 K) and then at the Nitrogen boiling temperature (77 K).

²The Callen & Welton formula is the Planck blackbody radiation law plus the $\frac{1}{2}$ -photon per unit bandwidth due to the zero-point fluctuation noise. Since all C-W-corrected load temperatures come out $\frac{1}{2}$ -photon higher than when using the Planck correction, C-W receiver temperatures are $\frac{1}{2}$ -photon lower (3.5 K at MINT’s 145 GHz). However, the C-W corrected temperatures for the sky and the CMB are $\frac{1}{2}$ -photon higher than Planck-corrected values, so in the end the C-W and Planck conversions yield the same value for T_{sys} . See Kerr et al. (1997) for a good discussion.




Figure 4.1: (a) A graphical interpretation of the Y -factor measurement. The y -intercept is the receiver output at zero input power. Tracing back to the negative x -intercept converts to temperature units. (b) The multi-point method used for tuning the MINT receivers. Several points are fit to a line to minimize the systematic effects of gain drifts during long tuning runs.

Finally, the receiver noise power, as measured from the Y -factor (Equation 4.5), can be re-expressed as the receiver noise temperature:

$$T_{rec} = \frac{T_{hot} - Y T_{cold}}{Y - 1}, \quad (4.9)$$

where T_{hot} and T_{cold} are the C-W noise temperatures of the hot and cold loads. A simple way to view the progression from Equations 4.5 to 4.9 is that P_{rec} , P_{hot} , and P_{cold} can be expressed in temperature or any other units, as long as they are *proportional* to power, and that the C-W temperature converts physical temperature to a noise temperature that is proportional to power.

For the T_{rec} measurements discussed in this chapter, we use a small modification to the Y -factor measurement (see Figure 4.1). Instead of just a pair of points, we measure several and fit a line to them. The negative x -intercept of the best-fit line is the receiver temperature. The slope of this line, the *responsivity*, R , is equivalent to a gain and gives the conversion between receiver output units and Kelvin. A laboratory receiver-tuning system built around this measurement is described in Section 4.2, We also use this method in Section 4.5 to find the receiver temperature of each channel for the observing campaign.

We describe two different kinds of power measurements of the receiver output: detector diodes and digitizer histograms. The first is the common square-law detector diode which gives a voltage output that is a function of power. The lab tuning uses four diodes, one for each channel of the channelizer output. In the field, each receiver also contains a diode total-power monitor used to monitor gain stability and

the atmosphere (see the “DC level” in Chapter 3). The less conventional digitizer histogram power measurement, whose mathematical underpinnings are advanced in Section 4.4, allows us to assess the sensitivity of each channelizer channel during the campaign. Both the diodes and the digitizer histograms exhibit non-linearity at high power, but calibration of and correction for these small effects is straightforward (see Sections 4.4, 4.5, and Appendix A).³

For all receiver temperature measurements quoted in this chapter, we mean the following:

- 1) The receiver sees well-matched load inputs of several different temperatures.
- 2) The physical load temperatures are converted to T_{C-W} according to Equation 4.8.
- 3) The receiver output (diode or histogram) is corrected for offset and non-linearity to make its output proportional to power.
- 4) T_{rec} is calculated via the graphical Y -factor measurement of Figure 4.1.

The system temperature, T_{sys} , of Equation 4.1 has more components than just the T_{rec} that we’ve discussed so far. In general, ground-based CMB anisotropy measurements have noise contributions from the atmosphere and also from the CMB itself, which has a “physical” temperature of 2.73 K, and thus a C-W temperature of 4.07 K:

$$T_{sys} = T_{rec} + T_{sky} \quad (4.10)$$

$$= T_{rec} + \epsilon T_{air} + (1 - \epsilon) T_{CMB} \quad (4.11)$$

$$= T_{rec} + (1 - e^{-\tau}) T_{air} + (e^{-\tau}) T_{CMB} \quad (4.12)$$

The atmospheric optical depth, τ , ranges from 0 to ∞ , and the related quantity, the atmospheric emissivity, $\epsilon = 1 - e^{-\tau}$, can take values from 0 to 1. In D-band, the primary atmospheric culprits are water vapor, oxygen, and ozone. The emissivity at the Cerro Toco site is such that the D-band atmosphere contribution, $T_{atm} = \epsilon T_{air}$,⁴ is typically 5 K during clear weather. In Section 4.5, we present measurements of T_{sys} , including all of these contributions, throughout the 2001 campaign.

³The diodes have a small DC offset and a gain roll-off toward high power. The digitizer measurements exhibit a well characterized but poorly understood gain *increase* toward higher power.

⁴A note on terminology: we use T_{air} for the physical temperature of the air (about 250 K above Cerro Toco), T_{atm} for the contribution of the air to the system temperature, and T_{sky} for the sum of the atmosphere and CMB contributions.




Figure 4.2: A non-flat passband (solid) with an equivalent, flat passband (dashed).

Beside the system temperature, the other important contributor to the sensitivity of a receiver is the bandwidth, $\Delta\nu$. In an ideal receiver, the passband has flat gain inside sharply defined bounds, and zero gain elsewhere. A general receiver passband, $G(\nu)$, shown in Figure 4.2, can be described in terms of its *equivalent* noise bandwidth:

$$\Delta\nu_{equiv} = \frac{(\int_{-\infty}^{\infty} G(\nu)d\nu)^2}{\int_{-\infty}^{\infty} G^2(\nu)d\nu}. \quad (4.13)$$

To get the $\Delta\nu$ used in Equation 4.1, these integrals are taken over the IF band. As we will see in Section 4.2, the MINT passbands are not flat. Measurements of the equivalent passbands of the IF channels are presented in that section.

Bandwidth is important in another respect: the HEMT IF amplifier produces noise in a much wider bandwidth than that in which it has useful gain.⁵ Filtering the IF passband reduces T_{rec} by restricting the measurement to regions of high gain-to-noise. Due to the budgetary necessity of using spare MAT/TOCO parts, the filter employed before the DC level diode (see Chapter 3) is not ideally suited for MINT (see Section 3.4). The filter has a passband of 3.5–6.5 GHz, much wider than the 4–6 GHz region of high gain that is spanned by MINT’s channelizer channels. Due to the out-of-band noise, the DC levels have a higher receiver temperature than the more tightly filtered channelizer channels.

The measurements of T_{rec} quoted in this chapter are double sideband (DSB). The receivers are inherently DSB, so this is the most natural measurement to report. The

⁵Laboratory measurements discovered significant noise output power from 2–12 GHz, whereas the HEMT has useful gain over the much narrower range of \sim 3–7 GHz. The MINT channels are designed to cover the high-sensitivity range of 4–6 GHz.

receiver temperature can be broken down into contributions from the mixer and the IF chain (e.g. Blundell et al. 1992):

$$T_{rec} = T_M + LT_{IF}, \quad (4.14)$$

where L is the conversion loss in the mixer. In the course of lab testing, we measured T_{IF} , which includes the cold HEMT, isolator, and several sections of lossy, stainless, co-axial cable (the receivers are diagrammed in Section 3.3), to be about 7 K for receiver A's 4.0-4.5 GHz band. In an SIS mixer, L is typically a few dB. T_M could theoretically be as low as 0 K in a DSB mixer. We make no further attempt to isolate the contributions of T_M , T_{IF} , and L .

4.2 Tuning the SIS's

The NRAO mixer employed in the MINT receivers is a 6-junction superconductor-insulator-superconductor (SIS) device, with layers of superconductor Nb and insulator Al-Al₂O₃ (Kerr et al. 1993). The physical phenomenon that gives an SIS junction its utility as a quantum mixer is photon-assisted tunneling of (single-electron) quasiparticles from superconductor to superconductor across the insulating barrier. The effect was first observed by Dayem and Martin (1962), explained by Tien and Gordon (1963), and developed into a full quantum mixer theory by Tucker (1979). The (I, V) curve of the SIS is shown in Figure 4.3. The onset of quasiparticle tunnel current occurs at a voltage bias of $\frac{2\Delta}{e}$, where Δ is the superconductor energy gap and is a strong function of junction temperature. The SIS is voltage-biased just below this point, in order that a single RF photon adds enough energy to allow an electron quasiparticle to tunnel. The LO power level mainly affects the mixer gain (see Tucker and Feldman (1985)), and is optimized at about 300 nW at the mixer (Kerr et al. 1993). The backshort position controls coupling of the RF signal and the LO to the stripline antenna. This coupling is cyclic with a period of $\frac{1}{2}$ the waveguide wavelength:⁶

$$\lambda_{WG} = \frac{\lambda}{\sqrt{1 - \left(\frac{f_c}{f}\right)^2}}. \quad (4.15)$$

⁶The period is *half* the waveguide wavelength because there are two nodes and two anti-nodes per wavelength.




Figure 4.3: LO-pumped (I, V) curve of the 6-junction SIS. The sharp rise at $6(2\Delta/e)$ is due to the onset of quasiparticle tunneling across the insulating barrier. The optimum bias point is $\frac{1}{2}$ -photon-per-junction below this point.

The cutoff frequency, f_c , for WR-5.8 waveguide is 102 GHz, giving $\frac{\lambda_{WG}}{2} = 1.45$ mm for the center of the RF band, 145.1 GHz. Since the micrometer that controls the backshort travels 0.5 mm/turn, and the external rotary feedthrough is geared by a factor of five, the backshort coupling has a period of $\frac{\lambda_{WG}}{2} = 14.5$ turns of the feedthrough.

The automated SIS tuning rig (see Figure 4.4(a)) enables the search of a vast 3-dimensional parameter space and the rapid minimization of receiver temperature across the four IF channels. In about 24 hours of minimally attended operation, the rig searches through over 2500 combinations of LO power, junction bias, and backshort position. The RF load⁷ fits over the feedhorn and is contained within the dewar. In a tuning run, the computer heats the load and allows it to thermalize

⁷Designed and built by Mark Halpern & Gaelen Marsden (University of British Columbia).

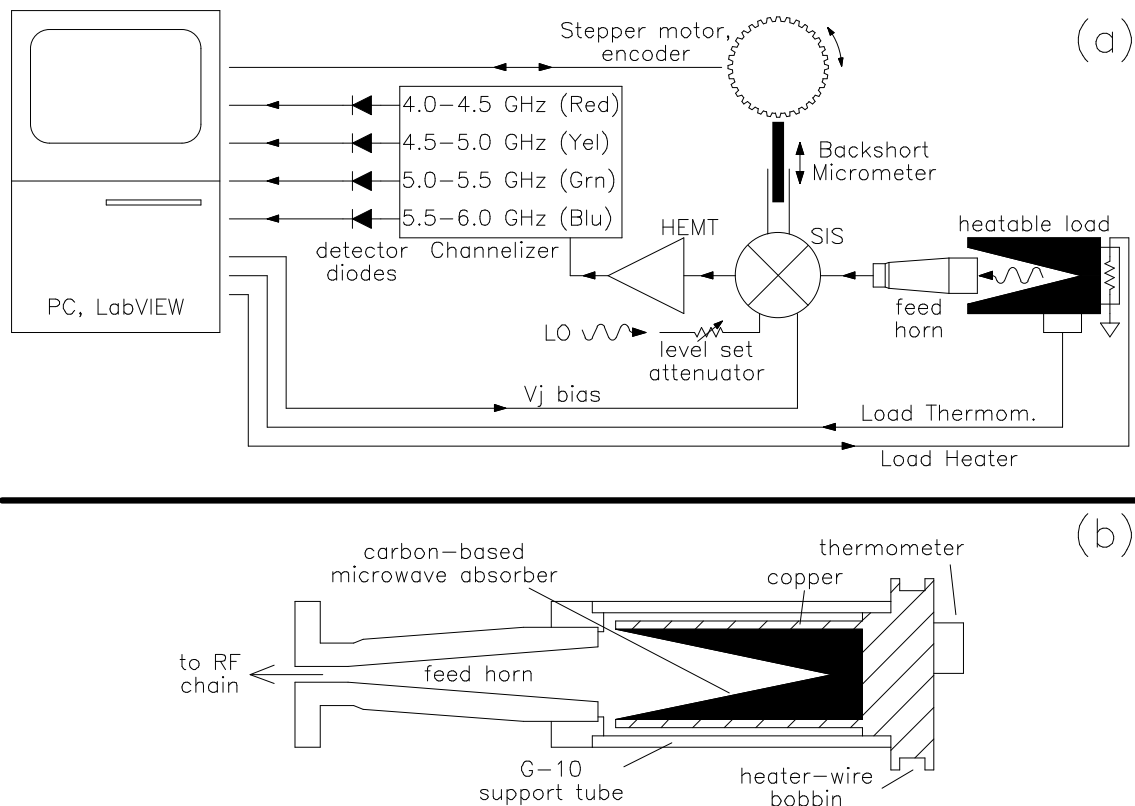


Figure 4.4: (a) Sketch of the SIS tuning rig. The PC controls the backshort position, the junction bias voltage, and load heater voltage and reads out the load temperature and the detector diodes from the four channelizer channels. The LO attenuator is adjusted by hand before each tuning run. (b) The Halpern D-band RF load. The G-10 support tube thermally isolates the load from the feedhorn and SIS. Total length of the load is about 2.5 in. Drawing is not to scale.

at some temperature, and then sweeps through a variety of bias voltages (V_j , or *junction voltage*) and backshort positions, recording the diode voltages from the four channelizer channels⁸ at each setting. The process is repeated for several (typically 3–4 total) different load temperatures, ranging from 4 K to about 20 K. The LO attenuator is not motorized, so one LO power level is probed per run. At the end of the run, the computer calculates the system temperature for each channel at a given V_j and backshort via the graphical Y -factor method. No diode linearization corrections are necessary during lab tuning due to the low output power levels encountered. Data from a typical run are shown in Figure 4.5.

The main design philosophy of the tuning system is to minimize systematic errors in the measurement of T_{rec} , and the key is the Halpern load (see Figure 4.4(b)). Many

⁸See Section 2.4 for a description of how the channelizer breaks up the 4–6 GHz IF band.




Figure 4.5: Selected T_{rec} curves for a typical tuning run for the four channels of receiver B. The LO attenuator's micrometer setting is 40 for this run, and represents the optimal LO power for Receiver B. In each channel's graph, the 5 curves (ascending) correspond to backshort rotary feedthrough settings of 1.0, 2.0, 3.0, 4.0 and 5.0 turns from all-the-way-in. The optimal setting for receiver B was chosen to be 0.0 turns (not pictured).

important systematic errors in a receiver temperature measurement stem from a lack of control or understanding of the load. For instance, using a load external to the dewar that is cooled with liquid nitrogen can introduce absorptive water vapor or frost. Often, precise determination of the temperature of the load is difficult. Alignment repeatability is important if there are moving parts or beam switches. Another alternative, a waveguide-termination load, must be thermally isolated from the mixer block with lossy waveguide sections, introducing systematic errors via the calculation of the loss. By contrast, the Halpern load doesn't move, can be easily varied over a factor of five or more in temperature, and includes whatever loss occurs in the feed-horn. Another important systematic error in obtaining reliable T_{rec} measurements is $1/f$ gain drift. In a typical tuning run, we find the receiver gain to be stable to about 1% over several hours, allowing the measurement of around 30 bias voltages and 20 backshort positions per tuning run. Expanding the Y -factor measurement to include several points rather than the conventional two aids in eliminating gain-drift errors, as any data set not well fit by a line is rejected.

If we assume the effective bandwidths of the four channelizer channels are the same and the channels are uncorrelated, the equivalent receiver noise temperature over the full (2-GHz) IF band as a function of the receiver temperatures of the individual channels is given by

$$T_{rec} = \left(\frac{1}{4} \left(\frac{1}{T_{rec1}^2} + \frac{1}{T_{rec2}^2} + \frac{1}{T_{rec3}^2} + \frac{1}{T_{rec4}^2} \right) \right)^{-\frac{1}{2}}. \quad (4.16)$$

It is this equivalent T_{rec} that we aim to minimize. Table 4.1 contains the optimal tuning parameters and T_{rec} measurements for each receiver. We note the general trend that the receiver temperature is higher for higher IF frequencies. This makes sense in the context of Figure 4.6, where we see that the receiver gain rolls off toward the higher end of the IF band.

As a final check of our understanding of the receiver performance, we measure the equivalent bandwidth, $\Delta\nu$, of the four channelizer channels for receivers A and D. The diode's power-spectral density (PSD , $V/\sqrt{\text{Hz}}$) is measured by spectrum analyzer at $\nu = 1$ kHz (above the $1/f$ -knee, below the rolloff from the low-pass filter in the diode

Receiver	A	B	C	D	units
V_j	1.410	1.478	1.380	1.435	V
LO attn., lab	45	42.5	56	50	micrometer
LO attn., sky	41.7	42.5	56	51	micrometer
Backshort	2.75	0.0	3.75	0.0	turns
CH servo	1.643	1.630	1.646	1.643	V
T_{rec} (equiv.)	27.7	21.8	40.4	29.4	K
T_{rec} Ch. Blu	39.5	33.1	67.9	41.6	–
T_{rec} Ch. Grn	27.5	21.9	55.0	31.4	–
T_{rec} Ch. Yel	23.9	19.2	32.9	27.2	–
T_{rec} Ch. Red	25.7	19.3	31.9	24.4	–
$\Delta\nu$ Ch. Blu	490	–	–	480	MHz
$\Delta\nu$ Ch. Grn	430	–	–	420	–
$\Delta\nu$ Ch. Yel	490	–	–	500	–
$\Delta\nu$ Ch. Red	430	–	–	450	–

Table 4.1: Lab-derived SIS tuning parameters and sensitivities for the 2001 observing campaign. The V_j quoted is the readout from the SIS bias card, which, due to the small voltages involved, multiplies the junction bias by 100. The “LO attenuator” is the micrometer reading (see Appendix A for calibration of the attenuator), with “lab” referring to the optimal value with the Halpern load and “sky” the setting chosen for use with the receiver looking outside. The backshort position is measured in CW revolutions of the rotary feedthrough, with 0 as far in as the backshort will go. The coldhead (CH) servo voltage, which refers to the setpoint of the SIS temperature servo as measured by a Lakeshore curve-10 diode (see Chapter 3 for details of the servo system), is included for future users of the receiver, as the best V_j bias point is a strong function of temperature. Channel colors are defined in Section 2.4. Measurements of $\Delta\nu$ were not made for receivers B and C. Tran (2002) gives the equivalent bandwidths of the channelizer channels alone, and finds the typical channel to be substantially wider than its nominal 500-MHz bandwidth.

preamp), and is combined with measurements of the responsivity, R , from the tuning data to calculate $\Delta\nu$ via:⁹

$$PSD = \frac{T_{sys} R \sqrt{2}}{\sqrt{\Delta\nu}}. \quad (4.17)$$

The results are found in Table 4.1. The shape of each channel’s passband, and thus the equivalent bandwidth, is determined by the product of the gains of the many RF and IF receiver components (see Chapter 3) as well as the channelizer passband. Measurements of the channelizer bands alone are given in Tran (2002).

⁹The $\sqrt{2}$ is from a conversion between $\sqrt{\text{sec}}$ and $1/\sqrt{\text{Hz}}$.




Figure 4.6: Receiver RF passband. This measurement was made by Miller (2000) using the MAT/TOCO receiver, channel D1. The data are shifted in center frequency from 144 GHz to 145.1 GHz to account for MINT’s different LO frequency. The upper and lower sidebands corresponding to MINT’s four channels are represented by the boxes.

4.3 Looking Outside the Dewar

For all its speed and accuracy, the drawback of receiver tuning using an internal load is that the optimal tuning can change when the receiver is opened to the outside. There are two main changes that could occur. The first is that the SIS block changes temperature due to a different parasitic loading. Since the superconductor gap voltage is a strong function of temperature, a change in temperature causes the optimal V_j setting to shift. We measure the effect of changing the temperature of the SIS in the lab; the results are plotted in Figure 4.7. This effect is minimized in the MINT receivers by servoing the SIS block temperature, so that the average temperature in the closed dewar is within 10 mK of that with the dewar open. Further, IR loading down the feedhorn that might heat the SIS junction even while the block is servoed (due to a finite thermal resistance between the stripline circuit and its housing) is blocked by the quartz window. We make no effort to re-bias the SIS for the open-dewar observing mode.

The second potential problem in opening the dewar to look outside is that LO power levels change because of reflections. After entering the SIS input waveguide




Figure 4.7: Receiver temperature vs. V_j for two different SIS block temperatures, 4.0 K (solid) and 4.5 K (dashed). As expected, the gap voltage, and thus the optimal bias point, decrease with increasing temperature. The 0.5 K temperature increase shown, if not compensated for via re-biasing, would cause T_{rec} to increase from 51 K to 58 K. Even if re-biased optimally, the SIS is still noisier when it is warmer. These data were taken with receiver A, through the DC level diode with no band-limiting filter, accounting for the high T_{rec} value.




Figure 4.8: Diagram of the optical path of the LO signal as it couples to the SIS. The quartz window, the polypropylene window, and the secondary mirror all provide opportunities for secondary reflections. Neither spacing nor size of the components is to scale.

through the branch-line coupler, the LO signal reflects off the backshort to form a standing wave with peak intensity (when tuned correctly) at the SIS stripline antenna (see Figure 4.8). A large fraction of the LO power is reflected back out the feedhorn toward the sky. Secondary reflections could occur at any or all of: the quartz IR-blocking window, the polypropylene vacuum window (although an effort is made to minimize the reflection from this pair, see Section 3.3) or the secondary mirror, causing changes in the intensity of the standing wave at the SIS antenna. Fortunately, the SIS current (I_j) is a sensitive monitor of LO power (Tucker and Feldman 1985). We measure the (I, V) curve of the SIS with optimal tuning and the cold load in place in the closed dewar. After the receiver is in its final configuration, the LO power is adjusted so that the new (I, V) curve matches the old. In two of the four receivers, the LO attenuator setting is changed by this process (see “LO attn., sky” in Table 4.1). Receiver A’s optimal LO micrometer setting moves from 45 to 41.7 (~ -2.5 dB) while receiver D’s changes from 50 to 51 ($\sim +0.5$ dB). Without these adjustments, the final receiver temperatures would be higher by several K. The optimal LO power settings for receiver B and C do not change from the cold load to the open dewar.

As we will see in Section 4.5, we were not entirely successful in the 2001 campaign at reproducing the low receiver temperatures measured in the lab. Receivers A & B behaved significantly worse in the field than in the lab. It was perhaps overly optimistic not to attempt to retune the receivers in the field. For any future MINT observations, the receivers will be fine-tuned, using the information already gained from lab tuning to reduce the section of parameter space to be covered.

4.4 The Digitizer Histograms

The digitizers in the correlator provide an alternative to diodes for making power measurements. The analog input signal to the correlator is assumed to be Gaussian random noise, with probability density of voltage, $\Psi(V)$, given by

$$\Psi(V)dV = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{V^2}{2\sigma^2}} dV, \quad (4.18)$$

where σ is the root-mean-square (rms) voltage of the signal. The rms power (Watts) of the signal, which is carried on 50- Ω coaxial cables and stripline, is $P(W) = \frac{V^2}{R} = \frac{\sigma^2}{50\Omega}$.




Figure 4.9: Formation of the digitizer histogram. The distribution shifts toward bins 0 and 3 with increasing power levels as the distribution describing the voltage widens.

The digitizer samples the signal in four levels (see Figure 4.9) using thresholds of 0, $V_0 = 0.25$ V, and $-V_0$, at a sampling rate of 1 GHz. One of the many diagnostic signals the correlator reports after each 0.5-sec integration period (roughly 500 million samples) is the number of digitizer samples that fall into each of the four levels. The fraction of readings in the inner two bins, ϕ , is a function of σ , given by

$$\phi(\sigma) = \frac{\text{bin1} + \text{bin2}}{\text{bin0} + \text{bin1} + \text{bin2} + \text{bin3}} \quad (4.19)$$

$$= \int_{-V_0}^{V_0} \Psi(V) dV \quad (4.20)$$

$$= \frac{1}{\sigma\sqrt{2\pi}} \int_{-V_0}^{V_0} e^{-\frac{V^2}{2\sigma^2}} dV. \quad (4.21)$$

Using the error function,

$$\text{erf}(x) = \frac{2}{\pi} \int_0^x e^{-t^2} dt, \quad (4.22)$$

we rewrite ϕ as

$$\phi = \text{erf}\left(\frac{V_0}{\sigma\sqrt{2}}\right). \quad (4.23)$$

Therefore, the signal power, in mW, is given by

$$P(\text{mW}) = \frac{1000\sigma^2}{50} = 10 \left(\frac{V_0}{\text{erf}^{-1}(\phi)}\right)^2 = 0.625 \left(\frac{1}{\text{erf}^{-1}(\phi)}\right)^2. \quad (4.24)$$

Programmable attenuators in the correlator board are checked hourly during observations to keep the power level at the digitizers in a narrow range, as the correlator has maximum sensitivity when $\sigma = 1.104 \cdot V_0$ (see Section 2.3). The actual attenuation levels of the attenuators may deviate by up to 0.5 dB from the nominal settings;

we correct for this using lab measurement data. In addition, lab measurements allow us to correct the histogram-derived power for non-linearity at high power. Both sets of corrections are described and tabulated in Appendix A. The steps in calculating true correlator input power in a given channel are to apply Equation 4.24 to the raw ϕ measurement, correct for nonlinearity, and then multiply by the true attenuator value. We use this method in the next section to calculate the system temperature of each channel through the observing campaign.

4.5 Determining T_{sys} During CMB Observations

We develop a method for determining the system temperature of each of the 16 receiver channels at all times during the campaign. The primary calibration uses the DC level diodes (recall that these measure power across a 3.5–6.5-GHz bandwidth) to observe an ambient-temperature microwave absorber and then the sky a few times during the campaign, thus determining T_{sys} and R for each receiver. With known R and T_{sys} (assumed to be constant through the campaign), the DC levels are then used as a sky monitor, yielding a sky temperature at all times during the campaign. The DC levels are chosen for this measurement because they are stable; all electronics involved are contained within the dewar/RBE system under tight temperature control and with no connectors exposed to the elements.¹⁰ Next we measure a campaign-average for T_{rec} in each channel as follows. The sky is used as a calibrated, variable load (varying by 50 K through the campaign) analogous to the cold load described in Section 4.2, allowing a graphical Y -factor measurement of T_{rec} in all channels via the digitizer histograms. Finally, the varying sky temperature is added back to the constant T_{rec} to determine T_{sys} for each channel through the campaign.

¹⁰Several of the individual channelizer channels show gain instability during the observing campaign, with the most likely culprit being loose or faulty connectors in the SMA cables connecting the receivers to the channelizer. This gain instability is the reason that the channelizer channels are not used directly as a primary calibrator. However, due to the peculiarity of having a unitless correlation function as the correlator output, a variable gain does not compromise MINT’s measurement, as long as the T_{sys} of the channel can be accurately reconstructed.

Three times during the campaign (early A.M. hours of Dec. 18, Dec. 21, and Jan. 6) we recorded the DC level output of each receiver while observing an ambient-temperature, eccosorb load and then the sky at two different elevation angles. The DC level voltage when the receiver is looking at the absorber is given by

$$V_{\text{out}}(\text{mV}) = R(\text{mV/K}) (T_{\text{load}} + T_{\text{rec}}). \quad (4.25)$$

The load temperature, T_{load} , was measured to be 270, 267, and 269 K, respectively, on the three nights. When looking at the sky, the receiver output is (see Equation 4.12)¹¹

$$V_{\text{out}} = R \left(T_{\text{rec}} + e^{-\tau_z \csc \theta_{el}} T_{\text{CMB}} + \left(1 - e^{-\tau_z \csc \theta_{el}} \right) T_{\text{air}} \right), \quad (4.26)$$

where τ_z is the optical depth of the atmosphere at zenith ($\theta_{el} = 90^\circ$). The two different sky measurements are made at $\theta_{el} = 60^\circ$ and 90° . A first-order approximation in τ ,¹² $e^{-\tau} = 1 - \tau$, simplifies Equation 4.26, allowing us to write a series of linear equations and solve for T_{rec} , τ , and R . With three measured DC levels,¹³ V_a , V_{90} , and V_{60} , for absorber, sky at $\theta_{el} = 90^\circ$ and sky at $\theta_{el} = 60^\circ$, respectively, we find

$$T_{\text{rec}} = \frac{\csc 60^\circ (V_{90} T_{\text{load}} - V_a T_{\text{CMB}}) - (V_{60} T_{\text{load}} - V_a T_{\text{CMB}})}{\csc 60^\circ (V_a - V_{90}) - (V_a - V_{60})}, \quad (4.27)$$

$$T_{\text{atm}} = \tau_z T_{\text{air}} = T_{\text{air}} \left(\frac{T_{\text{load}} - T_{\text{CMB}}}{T_{\text{air}} - T_{\text{CMB}}} \right) \left(\frac{V_{90} - V_{60}}{\csc 60^\circ (V_{90} - V_a) - (V_{60} - V_a)} \right) \quad (4.28)$$

and

$$R = \frac{V_a}{T_{\text{rec}} + T_{\text{load}}}. \quad (4.29)$$

We assume that T_{air} , the physical temperature of the air, is 250 K.¹⁴ The calibration results are compiled in Table 4.2.

Next, we calculate the sky temperature for each data frame during the campaign. The DC level output of each receiver is converted to T_{sys} by dividing by the average R measured during the eccosorb/sky-dip tests (see Figure 4.10). Subtracting the average T_{rec} , also as calculated from these tests, leaves the sky temperature (atmosphere

¹¹The relation $\tau \propto \csc(\theta_{el})$ assumes a planar blanket of air over the telescope.

¹²This approximation is appropriate as we measure $\tau = 0.11, 0.07$, and 0.03 on the three nights.

¹³The DC level diode corrections from Appendix Section A.1 are used for these and all subsequent DC level measurements.

¹⁴An error in assumption about T_{air} would change T_{atm} by less than 1% and would leave T_{sky} unchanged.




Figure 4.10: T_{sys} vs. time for the 2001 observing campaign. The points are the DC level in each receiver (recorded here once per 200 sec) divided by the average responsivity as measured in the calibration tests. The dashed lines represent the receiver temperature (assumed to be constant through the campaign), also as measured by the average of the calibration tests. The difference between the measured T_{sys} values and the receiver temperature is due to the sky (atmosphere plus ~ 4 K for the C-W corrected and slightly attenuated CMB). All graphs have the same vertical scaling.

–	Dec. 18	Dec. 21	Jan. 6	avg.
A T_{rec} (K)	54.6	57.2	55.4	55.7
B T_{rec} (K)	37.8	–	40.0	38.9
C T_{rec} (K)	40.6	42.6	–	41.6
D T_{rec} (K)	46.1	49.1	44.7	46.6
T_{atm} (K)	27	17	7	–
A R (mV/K)	2.91	2.69	2.78	2.79
B R (mV/K)	4.06	–	3.92	3.99
C R (mV/K)	2.81	2.76	–	2.78
D R (mV/K)	2.79	2.71	2.74	2.75

Table 4.2: Results from the three in-field eccosorb/sky-dip measurements. The atmosphere temperature measurements are an average of all working receivers for the given night.

plus CMB) as measured by each receiver (see Equation 4.11). We average together the sky temperature results from the four receivers. We have made the important assumption here that R and T_{rec} of the DC levels are constant through the observing campaign. We gain confidence in this assumption in noting that the calculated values of T_{sky} vary by only about ± 1 K from receiver to receiver, even though they are rendered from four independent sets of measurements. The receivers have no important electronics in common, and are housed in separate physical enclosures with independent temperature servos. Figure 4.11 shows the distribution of measured atmosphere temperature.

If the sky temperature derived from DC level measurements is to be useful as a calibrator in the four MINT frequency channels, the emitted atmospheric power must be flat across the bands. Figure 4.12 shows the simulated atmospheric contribution during good weather as a function of frequency across the MINT RF bands. The atmosphere is flat to better than 1 K within the bands.

Now calibrated over the whole campaign, the sky is used as a variable-temperature load. In analogy to Figure 4.1, we plot the digitizer-histogram-derived power against sky temperature to determine a receiver temperature¹⁵ for each channel (see Figures 4.13 and 4.14). If the data are close to a line, then T_{rec} and R are nearly constant during the observations. In Figure 4.10 and subsequent analysis, we find no

¹⁵Recall that T_{rec} is the $-x$ -intercept; see Figure 4.1.




Figure 4.11: (a) Histogram of the measured atmosphere (T_{atm}) contribution to T_{sys} during the 2001 campaign. This graph only includes data that survive instrument-function cuts (see Chapter 6). (b) Fraction of observing time with T_{atm} less than the given temperature (e.g. the atmosphere contribution was below 10 K more than 50 % of the time). This graph uses an average of 9.5 hours per night between Nov. 18 and Dec. 29, 2001. The five nights when the telescope was tarped due to snow have been assigned $T_{atm} = 250$ K (opaque).

–	Rec. A	Rec. B	Rec. C	Rec. D
Chan. Red	37	30	33	28
Chan. Yel	39	–	33	35
Chan. Grn	55	40	37	42
Chan. Blu	83	51	39	55
Avg.	46	38	35	37

Table 4.3: Receiver temperatures (K) of each channel during the campaign, as derived from the data in Figures 4.13 & 4.14. The band-averaged T_{rec} for each receiver is calculated via Equation 4.16.




Figure 4.12: Simulated atmosphere temperature spectrum at the Cerro Toco site with the MINT RF bands super-imposed. The simulations, described in Torbet (1997), assume 0.5 mm precipitable H_2O and Lorentz line profiles. The large slope on the left is due to an oxygen line at 118 GHz and on the right to a water line at 180 GHz. The smaller lines are from ozone. The large variability in the atmosphere during the campaign comes from the water line.

evidence to suggest variable T_{rec} in any of the receivers, so we make the assumption that only the responsivities of the various channels change during the campaign.¹⁶ The data fit a line with a spread of a few Kelvin for all the A channels, the three functioning B channels, and C Green and Blue. In C Red and Yellow, and all of the D channels, a single fit to the data is poor, but the striation in the plots suggest regions of constant responsivity to which lines can be fit. In each of these channels, the fits converge to a common receiver temperature, strengthening our confidence in the assessment that R is variable and T_{rec} is constant.¹⁷ The derived values of T_{rec}

¹⁶Changes in T_{rec} would require state changes in the cryogenic receiving elements. Responsivity, or gain, changes, are much easier to explain, as any of the warm electronics could produce them with changes in temperature, loose connectors, etc. In addition, T_{rec} should vary in a channelizer channel only if it also varies in the DC level of a given receiver.

¹⁷The variability in the responsivity of the D channels is explained by a loose co-axial cable between the receiver and the channelizer that was discovered at season's end. The reason for the variable gain in the two C channels is less clear, although a further analysis shows that the gains of the two channels are highly correlated. Perhaps there is a shared bad connector or component in the channelizer between the LO input and the mixers. In addition, the power measured from the

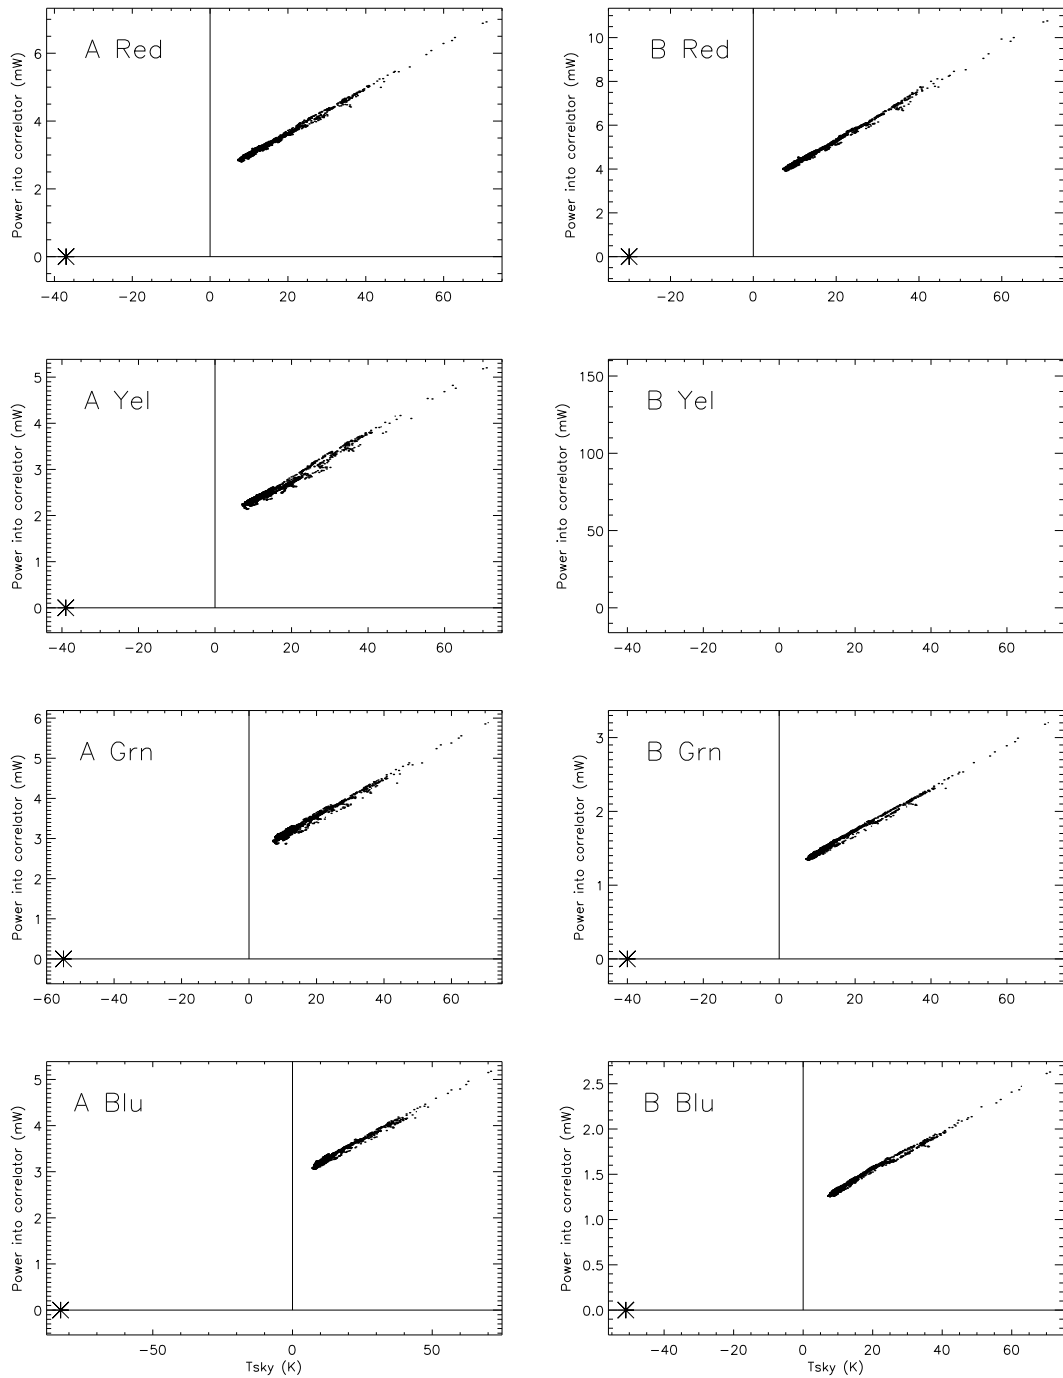


Figure 4.13: Receivers A and B, all channels, power input to correlator vs. sky temperature. In a method analogous to the graphical Y -factor method pictured in Figure 4.1, we calculate T_{rec} (denoted by $*$) for each channel using the variable sky as the load. The B-Yellow channel was broken during the campaign.

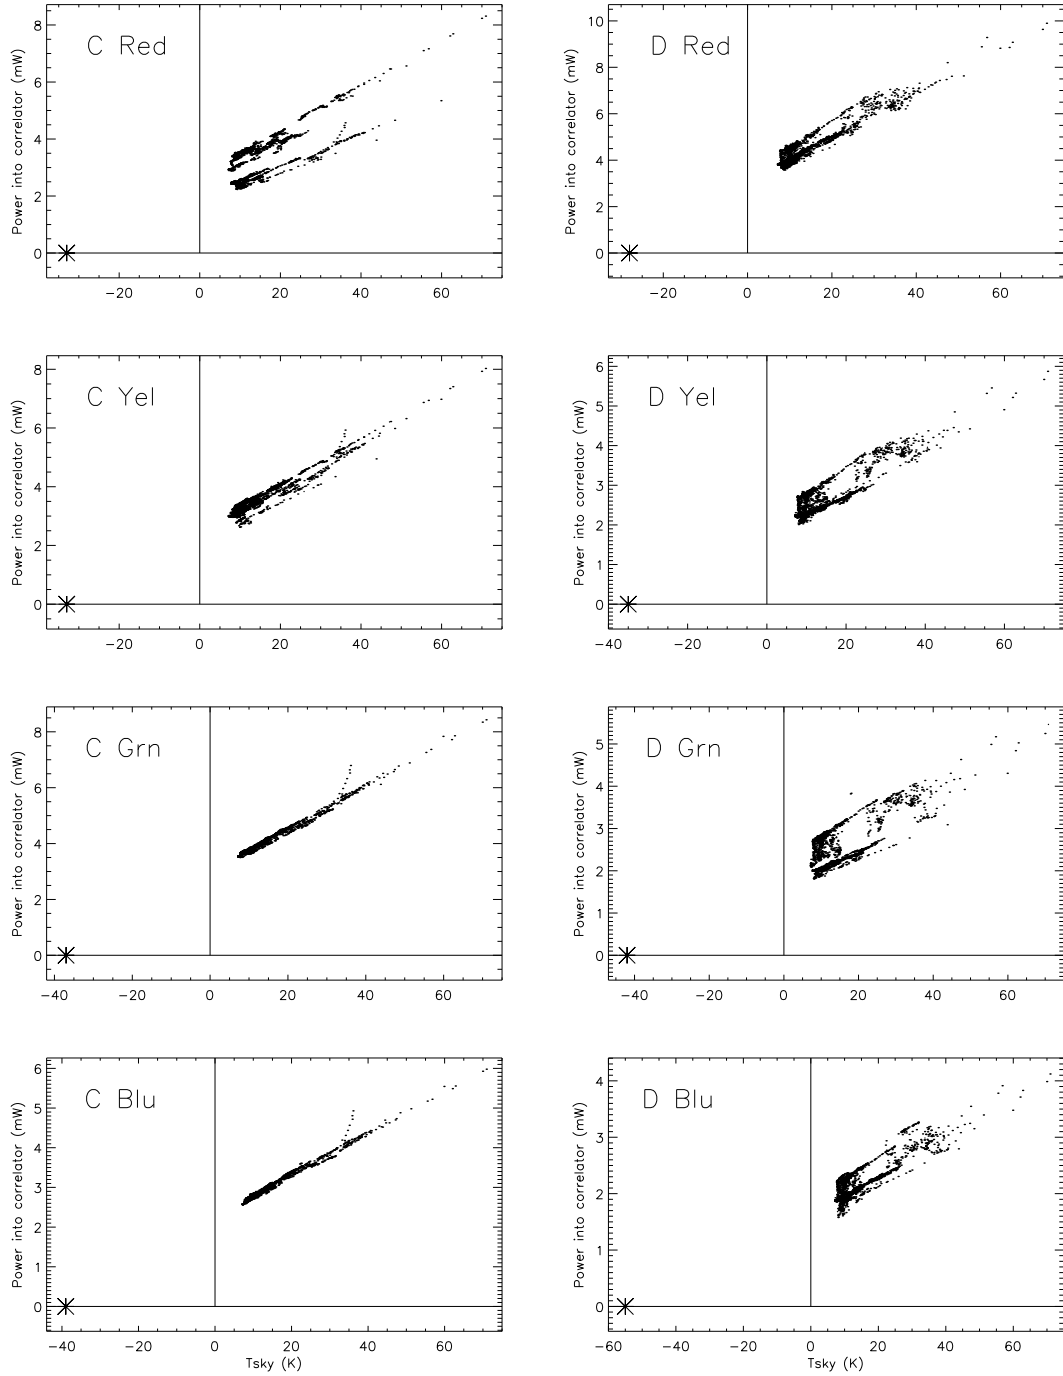


Figure 4.14: The continuation of Figure 4.13 for Receivers C and D. The Red and Yellow channels of Receiver C had non-constant gains through the season (cause unknown). Receiver D suffered from a loose SMA cable in the RBE output (discovered at the end of the campaign) that affected all four D channels. Even though the gain varies through the campaign, it is constant over short periods, enabling us to fit a series of lines with the same x -intercept ($-T_{\text{rec}}$).

for each channel are listed in Table 4.3. In any data frame (indexed by i), the system temperature of a channel is given by

$$T_{\text{sys},i} = T_{\text{rec}} + T_{\text{sky},i}. \quad (4.30)$$

The $T_{\text{sys},i}$ measurements are used in Chapter 6 to calibrate the correlator data.

calibrated noise spikes, constant through the campaign and injected every 8 minutes, tracks the gain of all channels well.

Local Oscillators and Phase-Switching

5.1 Introduction

The RF local oscillator system in MINT serves two purposes. The first is to provide a phase-stable tone to all four receivers in the telescope. Each of the MINT receivers has its own LO, and all are phase-locked to a common reference signal. Section 5.2 gives background on the *phase-locked loop* (PLL), the common circuit element that locks a local oscillator to an external reference. Section 5.3 describes the design and behavior of the more complicated MINT PLL.

The second main function of this system is phase switching, of which there are two levels. Fast phase switching, which allows the MINT correlator to read a modulated signal rather than a DC signal, undergoes a full cycle every 5 msec. The fast phase switching uses 180° shifts to reject IF correlated noise and crosstalk. The second, slower level of phase switching inserts an extra 90° of phase periodically into each baseline, in essence interchanging the real and imaginary parts of the correlation function. This slow switching, which allows post-separation of the sidebands, goes through a full cycle every 2 seconds. The design of the phase-switching system is discussed in Section 5.4. Finally, Section 5.5 explores the behavior of the PLL as it relocks after switching phase.

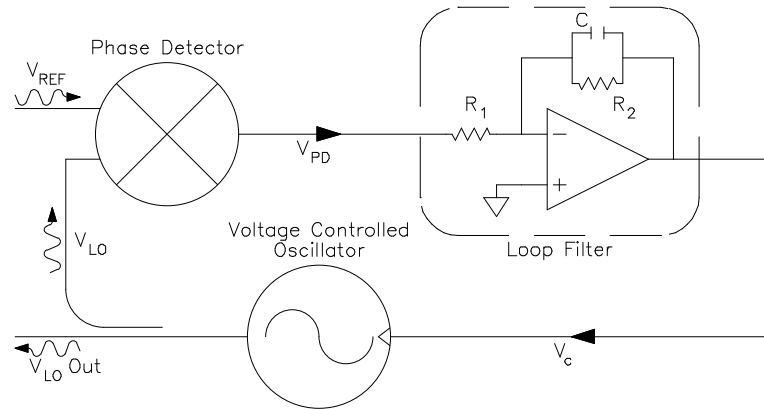


Figure 5.1: The canonical phase-locked loop.

5.2 Background on Phase-Locked Loops

The phase-locked loop (PLL) is a common circuit element; every FM radio has at least one, and a color television has several. The electrical engineering literature concerning PLL's is vast. The reader interested in exploring this topic further is referred to two books, those of Gardner (1979) and Best (1999) and the references contained therein. This section is intended to provide the reader with the background to understand the MINT oscillator system, and borrows material from Chapter 2 of Gardner.

The canonical PLL is pictured in Figure 5.1. There are three basic parts: the voltage-controlled oscillator (VCO), an oscillator whose output frequency depends on its bias voltage; a phase detector, or mixer, which produces a signal dependent on the phase difference between its inputs; and a loop filter, which conditions the signal that is fed back to the VCO. We model the VCO as having a linear change in frequency with control voltage:

$$\Delta\omega_{LO} = K_{VCO}V_c, \quad (5.1)$$

K_{VCO} is called the *VCO gain factor* and has units of (rad/sec-V). The phase detector works by multiplying its two inputs together. Assuming the reference and LO signals are sine waves of the same frequency, the output of the phase detector is a sinusoidal function of the phase difference plus high-frequency terms:

$$V_{PD} = G\{V_{REF} \sin(\omega t + \theta_{REF})\}\{V_{LO} \cos(\omega t + \theta_{LO})\} \quad (5.2)$$

$$= \frac{GV_{REF}V_{LO}}{2} \sin(\theta_{REF} - \theta_{LO}) + \text{high-freq. terms} \quad (5.3)$$

where G is the gain of the multiplier. If the frequencies of the reference and LO signals are similar but not the same, $\theta_{REF} - \theta_{LO}$ is a slowly varying function of time. In either case, to first order in phase difference and without the high-frequency terms, which are typically filtered away, the phase detector output can then be written

$$V_{PD} = \frac{GV_{REF}V_{LO}}{2} (\theta_{REF} - \theta_{LO}) \quad (5.4)$$

$$= K_{PD}(\theta_{REF} - \theta_{LO}). \quad (5.5)$$

K_{PD} is known as the *phase-detector gain factor* and has units of (V/rad). The loop filter could conceivably take any form; however, the most widely used is the one shown in Figure 5.1, which is an integral-plus-proportional circuit. With this filter, which is used in MINT, the loop is known as a type-II, second-order PLL.

The expressions for the above loop parameters can be rearranged to yield two key descriptions of the loop performance. The *damping factor* is given by

$$\zeta = \frac{R_2}{2} \sqrt{\frac{K_{VCO}K_{PD}C}{R_1}} \quad (\text{unitless}) \quad (5.6)$$

and the *natural frequency* is written as

$$\omega_n = \sqrt{\frac{K_{VCO}K_{PD}}{R_1C}} \quad (\text{rad/sec}). \quad (5.7)$$

By adjusting the loop parameters, we vary the damping factor in order to strike a balance between the competing performance factors of noise rejection and lock speed. Noise rejection is best at $\zeta = 0.5$ and the time-to-lock falls with increasing ζ . For these and a variety of other reasons, the damping factor is typically set to $1/\sqrt{2}$ (Gardner 1979). The natural frequency may be roughly described as the frequency beyond which the loop fails to be able to track phase changes. Again, by adjusting the loop parameters, it can be selected to match the noise-jitter characteristics of the oscillator to be locked.

5.3 The MINT Phase-Locked Loop

The MINT phase-locked loop design came about after extensive discussions with Bob Wilson, Robert Kimberk, and Todd Hunter at the Harvard-Smithsonian Center for

Astrophysics. The PLL board, containing the phase detector and loop filter, was designed at Harvard for the Sub-Millimeter Array telescope project (Hunter et al. 2001), and is used in MINT with minor modifications.

The MINT PLL is similar to the canonical PLL discussed in the previous section, the main difference being that the LO is down-converted in frequency before being compared to the reference signal. We describe the PLL function with reference to Figure 5.2, working around the loop and discussing each component, starting with the phase detector. The phase detection in this loop is digital, with one-bit digitization (compare to zero) before the multiplication step. The principal advantage of a digital phase detector is that unlike the analog case (see Equation 5.5), the phase detector gain, K_{PD} , does not depend on the amplitude of the input signals. Instead, K_{PD} is a settable parameter of the phase detector chip. The loop filter is the proportional-integral circuit discussed in the previous section, with R_2 left as a potentiometer so we can vary the damping factor (see Equation 5.6). The LO, a D-band Gunn diode,¹ runs at a center frequency of 145.1 GHz, and is specified for 0 dBm minimum output power. Tuning is achieved via a varactor, a voltage-tuned capacitor that changes the resonant frequency of the Gunn’s waveguide cavity. The unit contains an integrated 10-dB isolator and 20-dB power splitter on its output. The varactor tuning is well-approximated by the linear Equation 5.1 (see Figure 5.4). Most of the power goes to the dewar, but the -20 -dB port of the power splitter feeds a sub-harmonic mixer.²

The mixer is where the MINT system starts to diverge from the canonical PLL. Two reference LO’s are provided for this loop, one at 100 MHz and one at 12.1 GHz. The 12.1-GHz signal is fed into the mixer, which is tuned to enhance even harmonics. The 12th harmonic of the 12.1-GHz reference LO (145.2 GHz) mixes with the 145.1-GHz LO to produce a 0.1-GHz, or 100-MHz version of the LO signal. Since the mixer operates on high harmonics, its efficiency is very small (-60 to -70 dB), so a high-gain amplifier is needed to boost the 100-MHz LO signal. An attenuated version of the amplified, mixed-down LO is produced by the PLL box as the “ -20 dB IF Monitor.” Looking at this signal with a spectrum analyzer allows diagnosis of whether

¹The LO’s, model ZVTISC 6/1/145.1/0.15, are built by Zax Millimeter Wave Corporation, San Dimas, CA.

²This mixer, model DMA, contains an anti-parallel pair of diodes to enhance even-harmonic mixing. It is made by Pacific Millimeter Products, Golden, CO.

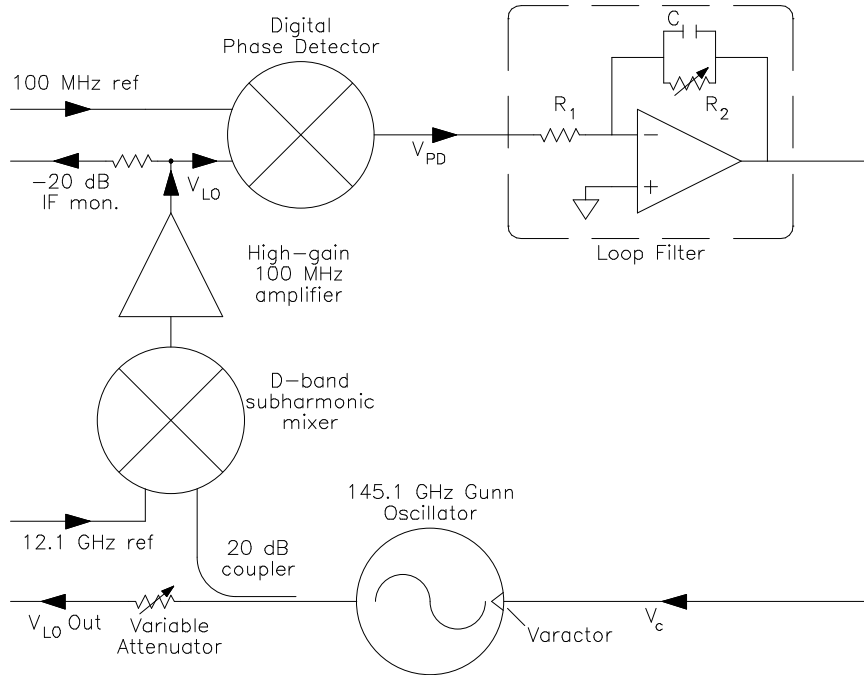


Figure 5.2: The MINT phase-locked loop. The phase detector is digital, and consists of two Analog Devices chips: the AD99687, Dual Ultra-fast Comparator; and the AD9901, Digital Phase/Frequency Discriminator. The LO is mixed down to 100 MHz.

or not the loop is locked. Finally, the 100-MHz mixed LO and the 100-MHz reference signal are compared by the phase detector. This completes the loop.

Locking the local oscillators of each of the four receivers is now reduced to a problem of managing the various free and fixed loop parameters in order to set ω_n and ζ optimally. The procedure followed is to get close to the desired settings for each LO by calculation, and then fine-tune R_2 and K_{PD} by hand until the LO is locked stably with minimal jitter about the desired 145.1 GHz. The natural frequency of the loop is set to roughly the width of the noise peak of the unlocked oscillator's power spectrum (see Figure 5.3); we choose $\omega_n = 1.3 \times 10^7$ rad/sec. The damping factor, as described in the previous section, is set to $\zeta = 1/\sqrt{2}$. The phase-detector gain can be varied via an external control voltage continuously from $K_{PD} = 0$ to $K_{PD} = 0.287$ V/rad. The middle of the range, $K_{PD} = 0.144$ V/rad, is used in these calculations to allow maximum adjustability. Figure 5.4 shows the tuning curves of the four oscillators. The slopes of these curves at 145.1 GHz give the VCO gain factor




Figure 5.3: The unlocked power spectrum of receiver D's D-band local oscillator, mixed down to 100 MHz. The vertical scale is 10 dB/div, from -20 to -120 dB, and the horizontal is 1 MHz/div, centered at 100 MHz. As the peak is about 2 MHz wide, we choose $\omega_n = 2$ MHz as a preliminary guess. The four unlocked oscillators have similar power spectra.




Figure 5.4: Frequency response of the 145.1-GHz LO to varactor voltage. Tuning is a strong function of temperature, which is servoed in the field. For these measurements, the temperature setpoint of the RBE's is 303 K.




Figure 5.5: The locked power spectrum of receiver D's local oscillator. The vertical and horizontal scales are the same as in Figure 5.3. The noise floor of this measurement is -95 dB, 2.5 blocks from the bottom of the plot.

for the four oscillators, which range between $K_{VCO} = 1.6 \times 10^8$ and 2.5×10^8 rad/sec-V. With $R_1 = 1$ k Ω , Equation 5.7 can be inverted to give $C = 140$ to 210 pF for the range of K_{VCO} values. Restricting the design to standard, commercially available parts, we choose $C = 200$ pF for all four oscillators. Finally, inverting Equation 5.6 yields a range of R_2 from 530 to 660 Ω for the four loops. We choose a 2 k Ω trimpot as R_2 for each loop. Leaving both K_{PD} and R_2 adjustable allows independent control over ω_n and ζ . Optimal lock is defined to be when the area under the power spectrum, excluding the central spike, is minimized. The resulting locked power spectrum, obtained from the -20 dB Monitor port of the PLL of receiver D, is shown in Figure 5.5. Table 5.1 contains the values of the various loop parameters used during the 2001 observing campaign.

5.4 Phase-Switching

By implementing a switching system in the local oscillators that rapidly adds and removes 180° of phase in each baseline, we reject unwanted correlated signals in the IF chain. Since $\cos(\theta + 180^\circ) = -\cos\theta$ and $\sin(\theta + 180^\circ) = -\sin\theta$, rotating the phase of a receiver's LO by 180° multiplies the IF output of that receiver by -1 . Every time one of the two receivers in a baseline is switched by 180° , the correlator

Parameter	Units	PLL A	PLL B	PLL C	PLL D
R_1	k Ω	1	1	1	1
R_2 (pot)	Ω	887	1250	670	624
C	pF	200	200	200	200
K_{VCO}	rad/sec-V	2.5×10^8	1.6×10^8	2.3×10^8	1.9×10^8
LGC	V	3.350	3.111	3.502	3.139
K_{PD}	V/rad	.072	.094	.056	.092
ω_n	rad/sec	9.5×10^6	8.7×10^6	8.1×10^6	9.3×10^6
ζ	none	0.84	1.07	0.55	0.58

Table 5.1: PLL tuning parameters for the 2001 MINT observing campaign. LGC is the average loop-gain-control voltage applied to the digital phase detector to set K_{PD} . R_2 corresponds to the trimpot R_{19} on the PLL circuit board.

also multiplies that baseline by -1 before adding it into the accumulating correlation function. The result is that RF signals that pass through the SIS mixer are multiplied by -1 twice, and thus are unchanged and add normally. Any signals entering the receiver after the mixer, such as potentially devastating correlated noise that could mimic a CMB signal (for example, an FM radio station near 100 MHz), are multiplied by -1 only once, and so quickly average away. Cross-talk between analog signal lines that might occur within the channelizer or correlator is similarly rejected.

Phase switching of the D-band LO's is carried out via the 100-MHz reference that is piped to each MINT receiver. Each PLL, by definition, works to keep the two signals entering its phase detector (refer to Figure 5.2), the 100 MHz reference, and the mixed-down LO, in a constant phase relation. Thus a phase change θ_{ref} in the reference signal is tracked by an equal phase change θ_{mix} in the 100-MHz LO mixing product. But what does this mean for the phase of the D-band LO, θ_{LO} ? From simple trigonometric identities, one can show that *phase changes are preserved across mixing*, i.e. that $\theta_{ref} = \theta_{mix} = \theta_{LO}$.³ In other words, to create a phase change on the D-band LO, we shift the 100 MHz reference by that same phase.

³Because the D-band LO is mixing in the lower sideband of the subharmonic mixer, the relation is actually $\theta_{ref} = -\theta_{LO}$. This is not important, since the sign of a phase change is a matter of convention, and all four receivers have the same negative sign.

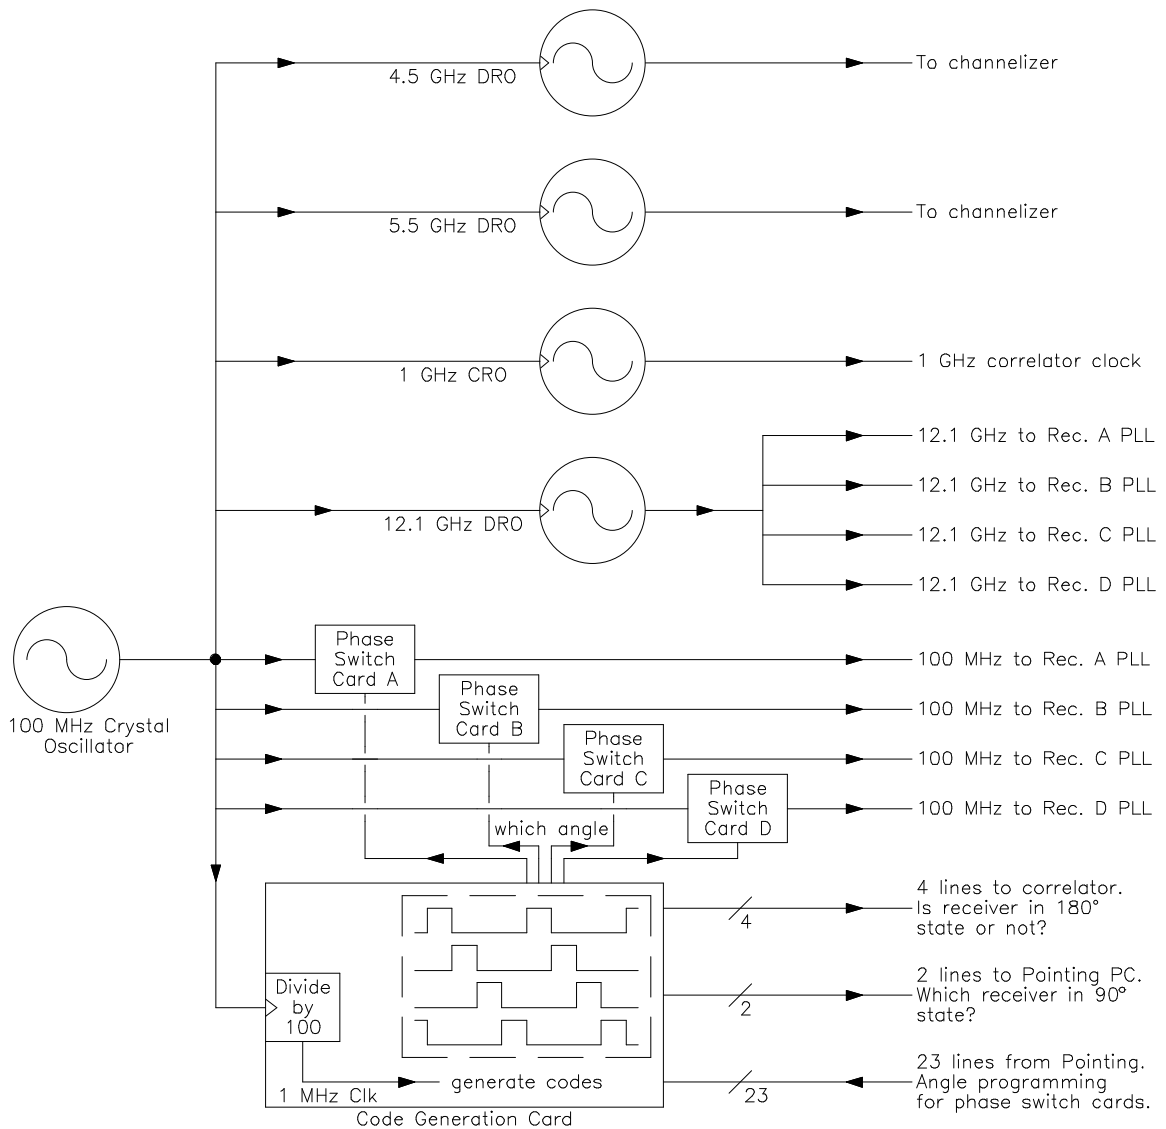


Figure 5.6: Outline of the interconnections in the oscillator box. All boards and oscillators are locked to the 100 MHz crystal.

An outline of the oscillator system for the entire telescope is shown in Figure 5.6.⁴ A digital board, the Code Generation Card (CGC), decides which receivers should be at which phase angles, and sends a two-bit code to each of the four Phase Switch Cards (PSC's). The PSC's can rotate the phase of the 100 MHz reference to any desired angle. The key chip in the PSC is a vector modulator,⁵ which splits the

⁴The assembly and testing of the oscillator box are described in Aboobaker (2001).

⁵Agilent HPMX-2005, Vector Modulator.

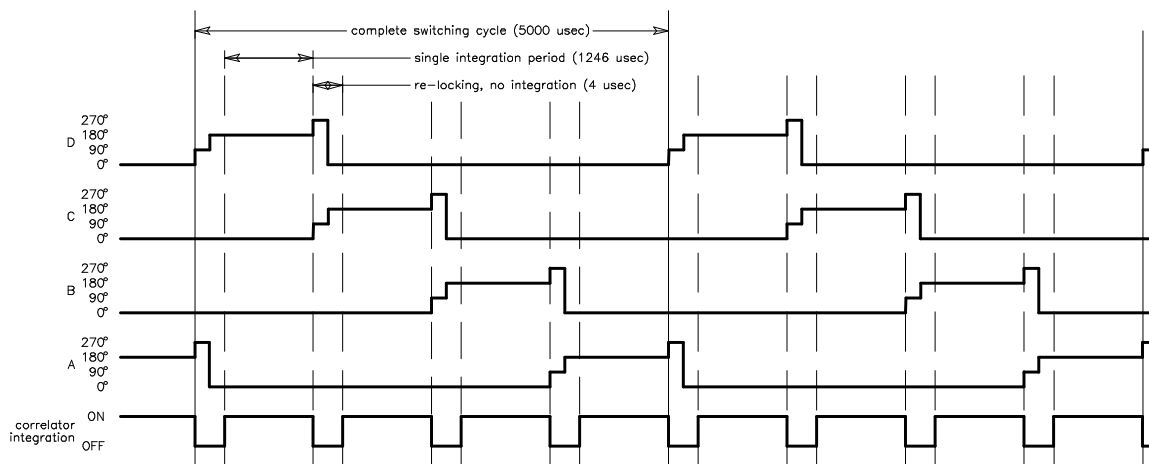


Figure 5.7: The timing of the 180° phase switching. The system undergoes a phase state change every $1250 \mu\text{sec}$ with $4 \mu\text{sec}$ (two steps of $2 \mu\text{sec}$) of dead time. The smallest complete switching cycle is $5000 \mu\text{sec}$. The correlator integrates for 100 complete switching cycles, or $500000 \mu\text{sec}$, before each readout.

incoming 100 MHz signal into in-phase and quadrature (shifted by 90°) signals, and then adds them back together in varying strengths to create the phase shift. Digital signals are passed from the CGC to the correlator, describing whether or not each receiver is in the 180° state, so the correlator knows when to multiply by -1 .

Figure 5.7 is a timing diagram for the 180° phase switching. Each of the four receivers spends $1/4$ of the integration period rotated by 180° from nominal. This is an unusual phase-switching algorithm. A typical interferometer uses *Walsh codes*, a series of square waves and square-wave-like functions that run at different frequencies in different receivers (Beauchamp 1975). If the 0° and 180° states of a receiver are represented by $+1$ and -1 , the requirement for valid phase switching codes (whether Walsh or otherwise) is that they be orthogonal, in the sense that the integral over a complete switching cycle of the product of any two codes is 0. An alternative formulation of this rule is that the product of any pair of codes has a 50% duty cycle. If this condition is not met, correlated IF noise is not averaged away. For an interferometer with exactly two or four receivers, one can use the much simpler prescription that we have used for MINT, with the advantage that all the receivers have identical codes. There are no special receivers in MINT.

Because our particular PLL's may not recover well from a direct 180° phase glitch (see Section 5.5), a 180° switch is carried out in two rapid 90° increments, after each




Figure 5.8: The difference between the 90° state and the normal state for receiver D.

of which the PLL's are given 2 μsec to re-lock. The correlator then integrates for 1246 μsec , and the whole process repeats. After four integrations, or 5000 μsec , the switching has gone through a complete cycle, as every receiver has been in the 180° state once. At the end of 100 complete switching cycles, or 500000 μsec , the CGC sets all four of its correlator lines high for 1 μsec , informing the correlator to write data to disk and start a new integration cycle. The three timing parameters, the length of the PLL re-lock period, the length of a single integration, and the length of a full integration cycle, are settable with switches on the CGC.

The phase switching system has a further task to accomplish: it must periodically introduce a 90° switch into each receiver. Adding 90° to a baseline in effect interchanges the real and imaginary parts of the correlation function of that baseline, a necessary step in separating sidebands of a double-sideband interferometer (see Chapter 2). The 90° phase switching in MINT happens on a much longer timescale than the 180° switching, but the pattern is the same. In each full integration cycle (500000 μsec) a different receiver is put into the 90° state, so that after 2 seconds, each receiver has had its turn once. To switch a receiver into the 90° state, the CGC interchanges $0^\circ \leftrightarrow 90^\circ$ and $180^\circ \leftrightarrow 270^\circ$ in the LO of that receiver (see Figure 5.8).⁶ During the integration periods, a receiver that would have seen $(\dots, 0^\circ, 180^\circ, 0^\circ, 0^\circ, \dots)$ instead sees $(\dots, 90^\circ, 270^\circ, 90^\circ, 90^\circ, \dots)$. During the 2-second full integration cycle, a *baseline* spends 0.5 sec integrating in each of the states $(0^\circ, 0^\circ, 90^\circ, -90^\circ)$, in a different order for each baseline.

⁶This particular algorithm is chosen, rather than the more straightforward $0^\circ \rightarrow 90^\circ$, $90^\circ \rightarrow 180^\circ$, $180^\circ \rightarrow 270^\circ$, and $270^\circ \rightarrow 0^\circ$, because it avoids forcing 180° switches at the beginning and end of the cycle (all switches are $\pm 90^\circ$). Note that the two algorithms are different only during the re-lock periods.

A pathology in the phase switching system, discovered only after the observing season, caused small deviations from the intended switches of up to 2.2° . The 12.1 GHz reference oscillator, as a result of itself being internally phase locked to the 100-MHz reference, also carries faint harmonics of 100 MHz around its peak ($\dots, 11.9, 12.0, 12.2, 12.3, \dots$ GHz). Upon entering the subharmonic mixer, the various frequency components of the 12.1-GHz signal mix with each other, producing a small but measurable 100-MHz interference tone that adds as a vector to the 100-MHz signal tone from the D-band LO, causing a small phase shift. Since the 12.1 GHz reference stays constant when the LO is phase-shifted, the interference tone results in a *different* phase error in each phase-shift state (the effect is about 1° if the interference tone is 35 dB below the signal). In lab measurements after the observing campaign, we measured phase errors of up to 0.2° , 2.2° , 2.1° , and 2.2° in PLL's A, B, C, and D, respectively. These phase errors have two effects on the performance of the telescope. The first effect, a slight decrease in the sensitivity of the instrument (far less than 1%), occurs because the IF signals are not multiplied by exactly $+1$ and -1 . This is unrecoverable, as the correlator averaging is done in hardware before data readout. The second effect is a slight mixing of the real and imaginary parts of the sky visibility (at roughly the 3% level in this case). This effect could be included in the data-covariance matrix as a real-imaginary cross term during analysis.

5.5 Transient Response: Recapturing Lock

When the 100-MHz reference signal to one of the PLL's undergoes a phase step, the PLL temporarily loses lock. The time for the phase error in a type-II, second-order loop to settle to 1% is (Hunter et al. 2001; Franklin et al. 1994)

$$t_{1\%} = \frac{4.6}{\zeta\omega}, \quad (5.8)$$

meaning that a MINT PLL should re-acquire lock within about $1 \mu\text{sec}$. However, this analysis is only valid over a certain range of phase steps. Figure 5.9 shows the difficulty. Our expression for ω_n (Equation 5.7) contains a linear approximation in the derivation of K_{PD} , the phase-detector gain factor. Clearly, nothing can have a linear relation vs. phase for all angles, because phase is cyclic. In the AD9901 phase

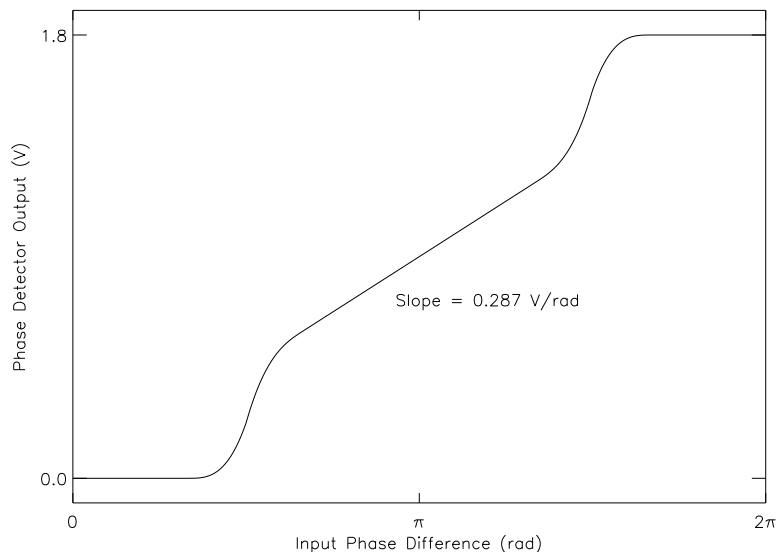


Figure 5.9: Response of the AD9901 phase detector as a function of input phase difference. The linear approximation (Equation 5.5) clearly breaks down at phase differences far from π radians. Landing in these “dead” regions after a 180° phase switch could lead to unpredictable loop behavior. Adapted from Analog Devices, datasheet for AD9901, Ultrahigh Speed Phase/Frequency Discriminator.

detector, K_{PD} rolls off to zero 180° away from the equilibrium phase offset. A sudden 180° phase step thus has the potential to behave very differently from a 90° step, and it is for this reason that the phase switching algorithms are designed to forbid 180° phase steps. As K_{PD} goes to zero, ω_n and ζ (Equation 5.6) go to zero, so the re-locking time (Equation 5.8) can diverge as the loop sits at an unstable equilibrium.

Tests of the re-lock ability of a MINT PLL are shown in Figure 5.10, where a direct 180° phase step is compared to a 90° step. As expected, the loop re-locks about $1 \mu\text{sec}$ after a 90° step. Interestingly, the loop appears to have no trouble re-locking after a 180° step, so the conservative design of two steps per phase switch appears to have been unnecessary.




Figure 5.10: Transient response of the D-band LO varactor bias to phase steps. Phase steps of 180° and 90° are introduced at $t = 0$ and $t = 3 \mu\text{sec}$, respectively. The 180° is similar to the 90° response, much better than expected.

Observations & Analysis

6.1 Introduction

MINT observed the CMB from November 22 to December 29, 2001 from Cerro Toco in northern Chile. In this chapter we present analysis intended to assess the stability and sensitivity of the instrument. The conversion of raw correlation functions to visibilities, including sideband separation, is covered in Section 6.2. In Section 6.3 we analyze an observation of Jupiter that occurred at the end of the observing campaign. Finally, we present a pair of interlocking assessments of the phase stability of the instrument. The calibrated noise source (Section 6.4), observed throughout the campaign during CMB mode, provides a good measure of phase *stability* by measuring phase changes. The noise source is bright (~ 5 K antenna temperature) so it provides high signal-to-noise after only one 2-second data frame, and can be observed as often as desired during CMB data collection. However, the noise source contains phase structure inherent to its delivery via waveguide to the centers of the primary mirrors that is not reflective of the true phase response of the telescope. Observations of Mars (Section 6.5) provide a true measure of the telescope's phase response. Unfortunately, during MINT's season, Mars was dim (a few mK antenna temperature) and was only in observing range for a few hours per night around sunset.

At time of this writing, the analysis of the MINT CMB data is not complete. The mode of observation was a modified drift-scan near zenith (for details see Tran 2002, § 2.2.5). Formalism concerning analysis and the response of the interferometer to the CMB, as well as signal-to-noise calculations are presented in Fowler (2002) and Tran (2002, § 2.2). Based on these analyses, the sensitivity derived from the Mars

observations, and observing time in the 2001 campaign (see Appendix B), it appears that MINT will be able only to place upper limits on the CMB anisotropy over the ℓ -range of 800–1900. A preliminary guess is a 200–300- μK , (95%-confidence level) upper limit in each of two window functions, $\ell = [800, 1300]$ and $\ell = [1300, 1900]$.

6.2 Converting Raw Correlator Data to Visibility

The raw 16-lag correlation function undergoes a number of data-processing steps to become the visibility, a function of frequency with temperature units. The first steps in the analysis are to remove the bias due to the digital multiplication algorithm (Table 2.2) and to normalize the result by the number of multiplications in the accumulation period. Because the correlator ignores data for 4096 nsec each time the LO phases switch (every 1250 μsec , or 400 times per accumulation), the actual number of multiplications per 0.5-second accumulation is $500,000,000 - 4096 \times 400 = 498,361,600$.¹ The raw correlator data undergo 14 bits of truncation before readout, and have a bias of 3 from multiplication. Thus, the correlation function, $r_{4DI}(\tau)$, in units of a single multiplication using Table 2.2(c)'s 4-level deleted-inner-product scheme, is given by

$$r_{4DI}(\tau) = \frac{r_{raw}(\tau) + \frac{1}{2}}{\left(\frac{498,361,600}{2^{14}}\right)} - 3, \quad (6.1)$$

where $r_{raw}(\tau)$ is the raw correlator output as a function of lag, τ , and the $\frac{1}{2}$ is the average loss due to truncation in a binary number. Next, we correct for the non-linearity that results from the quantization-and-multiplication algorithm (Tran 2002; Thompson et al. 1994, § 8.3) to find the unitless correlation function, $\rho(\tau)$:

$$\rho(\tau) = \left(\frac{\pi}{2}\right) \frac{r_{4DI}(\tau)}{E_A E_B + E_A + E_B}. \quad (6.2)$$

The value E is calculated for each of the two receivers in the baseline (here designated A and B) via

$$E = \exp\left(-\left(\text{erf}^{-1}\phi\right)^2\right), \quad (6.3)$$

¹The correlator is clocked at 62.5 MHz, or 16 times slower than the 1 GHz-clocked digitizers. As a diagnostic of its stability, the correlator returns the number of non-blanked accumulation cycles. This number is consistently $\frac{498,361,600}{16} = 31,147,600$.

where ϕ , the inner-fractional population of the digitizer histograms, and erf^{-1} , the inverse error function, are as defined in Section 4.4. The correlation function now corresponds to the unitless cross-correlation coefficient of Equation 2.24:

$$\rho(V_A(t), V_B(t), \tau) = \frac{\langle V_A(t)V_B(t - \tau) \rangle}{\sqrt{\langle V_A^2(t) \rangle \langle V_B^2(t) \rangle}}. \quad (6.4)$$

Finally, converting the measured cross-correlation coefficient to temperature units, $r_T(\tau)$, relies on the system temperatures calculated in Chapter 4. Since ρ is unitless, we multiply by the geometric mean of the system temperatures:

$$r_T(\tau) = \sqrt{T_{\text{sys}A} T_{\text{sys}B}} \rho(\tau). \quad (6.5)$$

All correlation functions after this analysis step have units of temperature.

Next, we find the visibility by taking the discrete Fourier transform (DFT) of the correlation function,

$$\mathcal{V}(\nu) = \sum_{\tau=-7 \text{ ns}}^{+8 \text{ ns}} r(\tau) e^{-i2\pi\nu\tau}. \quad (6.6)$$

The 16 frequency bins represented in $\mathcal{V}(\nu)$ are (± 500 MHz, -437.5 MHz, -375 MHz, \dots , $+312.5$ MHz, $+375$ MHz, $+437.5$ MHz). Since $r(\tau)$ is real, $\mathcal{V}(\nu) = \mathcal{V}^*(-\nu)$, so the negative frequency bins carry no additional information. We drop them, leaving the nine bins: $\nu = (\pm 500$ MHz, -437.5 MHz, -375 MHz, \dots , -125 MHz, -62.5 MHz, 0).² The DSB cancellation problem in Equation 2.16 is corrected by separating the sidebands. Each baseline spends $\frac{1}{4}$ of the 2-sec phase-switching cycle in four LO states, with relative LO phases of ($0^\circ, 0^\circ, +90^\circ, -90^\circ$). These phase-states occur in a different order in each baseline.³ We have derived (Equations 2.21 and 2.22) the sideband separation algorithm based on 0° and 90° LO states:

$$\text{USB} = \tilde{r}_0(\nu) + i\tilde{r}_{90}(\nu) \quad (6.7)$$

$$\text{LSB} = \tilde{r}_0(\nu) - i\tilde{r}_{90}(\nu). \quad (6.8)$$

²The ± 500 -MHz bin represents an average of the visibility at $+500$ MHz and -500 -MHz; this frequency wrap-around is a general property of the DFT. There are 16 independent numbers in the (real) time-based data, $r(\tau)$, and 16 independent numbers in the (complex) transform, as the 0 bin and the ± 500 -MHz bin have only real parts.

³See Section 5.4 for a description of the phase switching of the receiver LO's.




Figure 6.1: A demonstration of the sideband separation algorithm using a single data frame taken during Jupiter observations. Data are from the Red channel (4.0–4.5 GHz) of baseline D–B. (a) The time-based correlation function in each of the four 90° LO states. (b) The magnitude and phase of the complex visibility in both sidebands. The data are plotted in terms of their IF (4–6 GHz) frequency.

We modify this method, following the prescription of Fowler (2001), to account for the additional sign changes due to the upper- and lower-sideband down-conversions in the channelizer, by defining

$$p = \begin{cases} +1: \text{USB in 2nd LO} \\ -1: \text{LSB in 2nd LO} \end{cases} \quad (6.9)$$

Calculating \mathcal{V}_0 as the average of the two 0° LO states, and \mathcal{V}_{90} as the average of the $+90^\circ$ state and the negative of the -90° state, we find:

$$\mathcal{V}_{USB} = \mathcal{V}_0 + i(p)\mathcal{V}_{90} \quad (6.10)$$

$$\mathcal{V}_{LSB} = \mathcal{V}_0 - i(p)\mathcal{V}_{90}. \quad (6.11)$$

The upper sideband has frequency components spaced by 62.5 MHz from 0 to 500 MHz, and the lower sideband goes from -500 MHz to 0. A mathematical curiosity that results from this process is that the zero-frequency term in each sideband is actually an average of the zero-frequency terms of both sidebands. The same is true of the ± 500 -MHz terms. An example of the sideband-separation algorithm, based on Jupiter data, is pictured in Figure 6.1.

There is one more step toward interpreting the visibility if the incoming signal is astronomical in origin: we must account for the attenuation in the atmosphere. We follow the algorithm in Section 4.5 for finding T_{sky} , and then calculate ϵ according to

$$\epsilon = \frac{T_{sky} - T_{CMB}}{T_{air} - T_{CMB}}. \quad (6.12)$$

The constants T_{air} and T_{CMB} are defined in Section 4.1 (see Equation 4.12). The atmospheric attenuation is corrected by multiplying the visibility by $\frac{1}{1-\epsilon}$.

6.3 Jupiter Observation

Jupiter is a bright point source, and would be an ideal candidate for a MINT calibrator. Unfortunately, during the 2001 observing campaign, Jupiter never rose high enough in the sky ($\theta_{el} > 60^\circ$) to be observed directly by MINT. By placing a plane mirror over the telescope (see Figure 6.2) to aim the receiver beams much closer to the




Figure 6.2: A sketch of the mirror used to observe Jupiter.

horizon, we were able to observe Jupiter on the morning of January 4, 2002, beginning at about 7:00 UT.⁴ Receiver C had already been shut down for the season when Jupiter was observed, limiting the observation to the D–A, D–B, and A–B baselines. Because of the missing receiver and also because of unknown losses, alignment and curvature of the mirror, the Jupiter observation does not provide a satisfactory phase or amplitude calibration of the data. Rather, the analysis here is intended to provide insight into what the interferometer sees and an idea of what issues we will encounter during data analysis.

The Jupiter observation consists of a series of drift scans, each with a different declination offset from Jupiter. The total scan covers a square in (RA, Dec) about $0^\circ.6$ on a side, and represents ~ 2000 seconds of data. We grid the scan into pixels $0^\circ.025$ on a side, averaging the complex visibilities in each pixel.⁵ Figure 6.3 shows the real and imaginary parts of the visibility (the fringes) for the D–B baseline. The observed Jupiter fringes agree well with the predicted shape of the point-source visibility (Equation 2.23), exhibiting a sinusoidal fringe of period, $|\vec{x}| = 1/|\vec{u}|$ (\vec{x} is

⁴Since the large ground screen had to be removed to accommodate the mirror, in a process that bumped the telescope enough that it could not have occurred safely during the campaign, we were only able to observe Jupiter well once, after all CMB observations were completed.

⁵There are typically 3–4 of the 2-second data frames per pixel.




Figure 6.3: The real and imaginary part of Jupiter’s complex visibility, as measured in the D–B baseline.

telescope pointing angle, with Jupiter’s position taken as the origin), enveloped by the cross-beam power of the two receivers.

To produce the plots in Figure 6.3, the visibility in each pixel is averaged across frequency with uniform weighting to increase the signal-to-noise. Frequency averaging works well as long as the phase-vs.-frequency profile is approximately constant. The D–B baseline is well-behaved, with a rms phase error of $\sim 10^\circ$. In general, however, the phase profile of a given baseline is not flat, so averaging over frequency reduces the calculated signal and thus the signal-to-noise. The average of visibilities over frequency bin i , when all have the same amplitude $|\mathcal{V}|$ and each has a different phase, θ_i , is given by⁶

$$\mathcal{V}_{avg} = \frac{1}{N} \sum_{i=1}^N |\mathcal{V}| e^{i\theta_i}. \quad (6.13)$$

The amplitude of the average visibility is given by

$$|\mathcal{V}_{avg}| = \sqrt{\left(\frac{|\mathcal{V}|}{N} \sum_{i=1}^N \cos \theta_i\right)^2 + \left(\frac{|\mathcal{V}|}{N} \sum_{i=1}^N \sin \theta_i\right)^2}. \quad (6.14)$$

If we assume phase fluctuations are Gaussian-distributed across the band, with rms angle θ_{rms} , the amplitude of the average visibility is⁷

$$|\mathcal{V}_{avg}| = \frac{|\mathcal{V}|}{\theta_{rms} \sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{\theta^2}{2\theta_{rms}^2}\right) \cos \theta \, d\theta. \quad (6.15)$$

⁶A version of this derivation appears in Tran (2002, §7.2.2).

⁷The sine terms integrate to zero because they are odd and the Gaussian function is even.




Figure 6.4: Average magnitude and phase of the visibility as measured by the three baselines that observed Jupiter. All raw Jupiter visibilities have this phase profile removed before frequencies are combined. The average magnitude in each baseline is lower than the peak value reported in Table 6.1 because the plotted data represent an average over the central $1\text{-}\sigma$ of the cross-beam. Note: the Yellow channel of receiver B did not function during the campaign.

For $\theta_{rms} < 30^\circ$, the fractional signal loss due to averaging, $f_{loss} = 1 - \frac{|V_{avg}|}{|V|}$, is well approximated by $f_{loss} \sim 1 - \cos \theta_{rms}$.

To avoid the loss in sensitivity due to phase fluctuations, we correct the visibility phase in each frequency bin so that all frequencies have the same phase.⁸ In order to flatten the phase of the Jupiter observation, we first *spatially* average the visibility data to produce a high-signal-to-noise phase template. The fringe phases, given by $\theta_{phase}(\vec{x}) = 2\pi\vec{u} \cdot \vec{x}$, are removed before we average over the central $1\text{-}\sigma$ of the cross-beam. The resulting phase profiles are shown in Figure 6.4. The D-B phases are well-behaved, but the D-A and A-B baseline phases vary over the full 360° . Returning to the un-binned data, we remove the phase profiles from the measured Jupiter visibility in each data frame. Finally, we produce a beam map by taking the magnitude of the visibility in each bin. The results are shown as surface and contour plots in Figure 6.5.

What signal size should we expect from Jupiter in the three baselines? During the observation, Jupiter had an angular diameter of $\theta_D = 46.9$ arc sec.⁹ We

⁸By “correcting by” or “removing” a phase, θ_c , we mean multiplying the complex visibility, $\mathcal{V} = |\mathcal{V}|e^{i\theta_{\mathcal{V}}}$, by $e^{-i\theta_c}$, thus rotating it in the complex plane to $\mathcal{V}_c = |\mathcal{V}|e^{i(\theta_{\mathcal{V}} - \theta_c)}$.

⁹This and all other astronomical data on the Jupiter and Mars (sky position, etc.) come from the ephemeris program xephem by Elwood C. Downey.

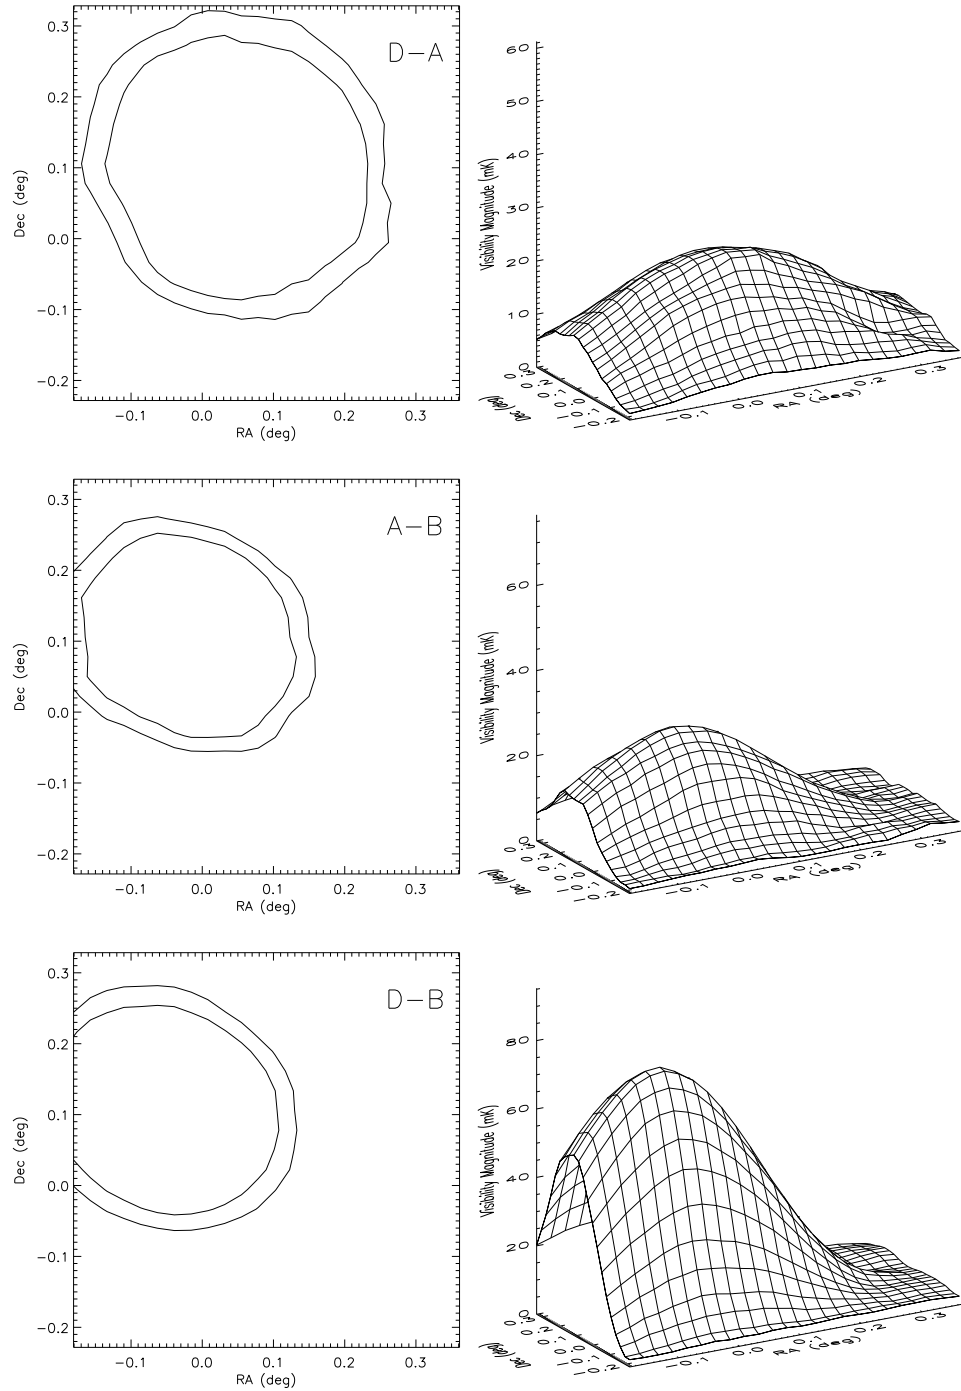


Figure 6.5: The magnitude of the Jupiter visibility as a function of the telescope pointing away from Jupiter. The peak magnitudes in the three baselines are 21.1, 26.8, and 72.3 mK, in the baselines D–A, A–B, and D–B, respectively. The two levels in each contour plot are $1/\sqrt{e}$ and $1/2$ of the measured peak value. The maximum of each surface plot’s z-scale is the predicted peak signal for the baseline. The cause of the non-circularity seen in the contour plots is unknown, but could be explained by curvature in the plane mirror.

	measured $ \mathcal{V}_{peak} $	predicted $ \mathcal{V}_{peak} $ perfect alignment	predicted $ \mathcal{V}_{peak} $ measured alignment	meas./pred. measured alignment	measured FWHM (avg) (degrees)	predicted FWHM (degrees)
D–A	22 mK	66 mK	61 mK	35%	0.42	0.44
A–B	27 mK	98 mK	77 mK	36%	0.35	0.36
D–B	74 mK	98 mK	95 mK	77%	0.36	0.36

Table 6.1: Jupiter peak visibility magnitudes.

found in Section 3.2 that the optics have forward gains of $G_f = 51.32$ dBi (small) and 54.80 dBi (large). We assume Jupiter to have a brightness temperature of $T_B = 150$ K at 145 GHz (using the 90-GHz value of Hedman 2002). The expected visibility magnitude at the peak of the beam map is given by:

$$|\mathcal{V}_{peak}| = T_B \sqrt{G_{f1} G_{f2}} \frac{\pi (\theta_D/2)^2}{4\pi}. \quad (6.16)$$

These values are listed in Table 6.1, along with the measured values. During the course of the Jupiter analysis, we determined that each baseline’s beam pattern had a different center, meaning that the receiver beams were misaligned from one another by up to $0^\circ.22$. Having found no evidence for this in previous alignment checks using the sun (Marriage 2002), we deduce that the beam misalignment is due to curvature in the plane mirror placed over the telescope. Since the peaks of the beam patterns do not coincide when the beams are misaligned, there is a reduction in the expected peak signal. This reduced value is listed in Table 6.1 as “predicted $|\mathcal{V}_{peak}|$, measured alignment.” The measured signal, especially in the baselines involving receiver A, is too low. The reason for this is still unknown, and is being investigated. A final check of our understanding of the telescope behavior comes from the measured shape of the beam map. The measured FWHM of the beam in all three baselines agrees well with values calculated based on the optics simulations (see Table 6.1).

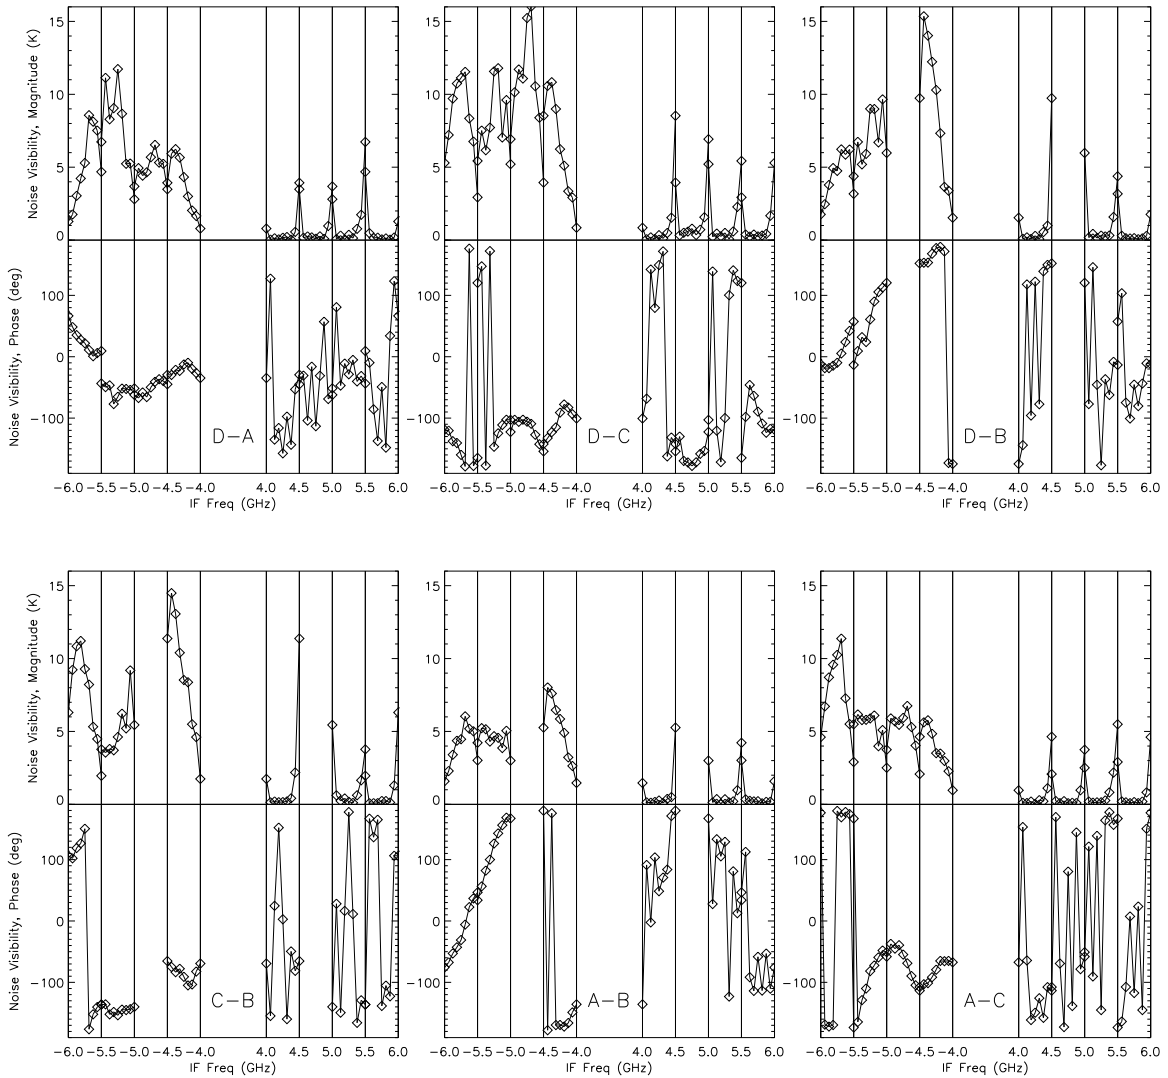


Figure 6.6: Amplitude and phase of the visibility for the fiducial noise spike. This phase profile is removed from all other noise spikes. The resulting visibilities are analyzed for non-zero slopes and offsets.

6.4 Calibration Noise Source

The broadband calibration noise source, described in Section 2.6, is observed for a few seconds every ~ 7.5 minutes throughout the observing campaign. The resulting spike in correlated power provides a test of the phase and amplitude stability of the measured visibility over time. In order to compare noise-spike phases through the campaign, we define a “fiducial” noise spike.¹⁰ Noise data are processed in the

¹⁰Data file 2415, records 147–148 (December 14, 5AM UT) contain the fiducial noise spike.

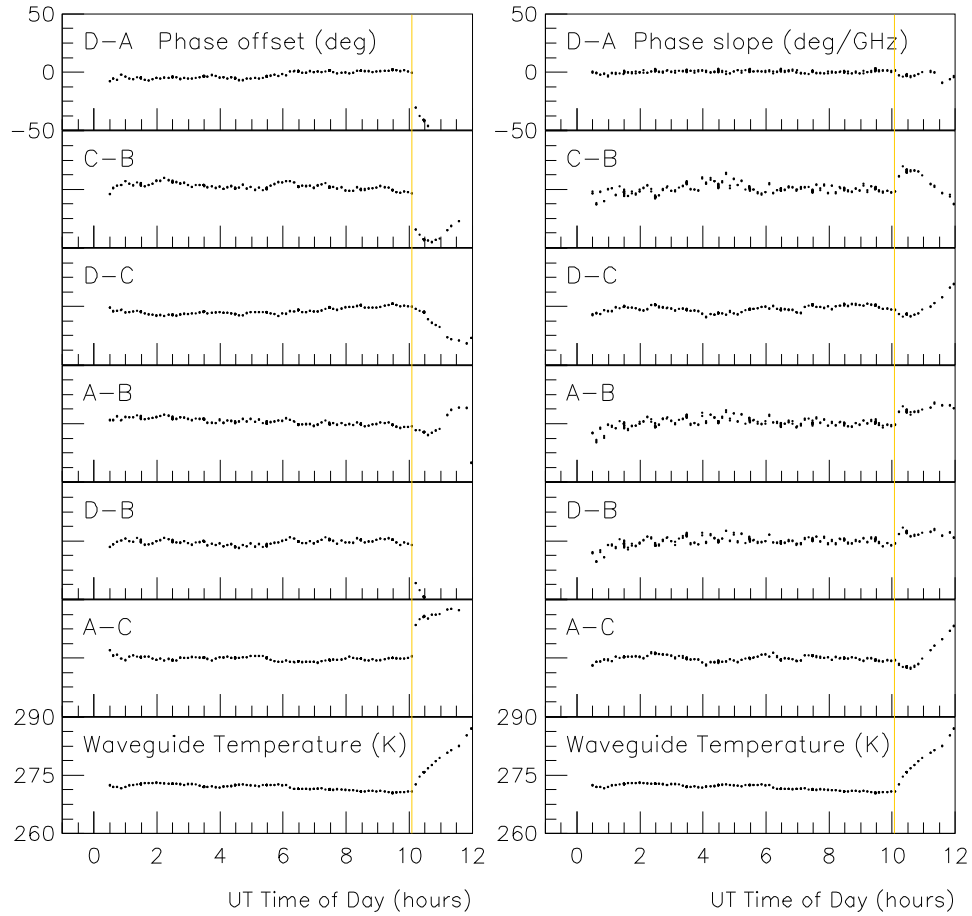


Figure 6.7: Phase fluctuations away from the fiducial phase of the noise source: offset and slope across the band. Results are plotted for all baselines in the Blue channel (other channels track to within a few degrees) for night 333. All phase plots in a column have the same vertical scale. The bottom plot shows the temperature of one of the noise source waveguides, which increases rapidly after sunrise (grey line, $\sim 10:00$ UT, 7 A.M. local). The phase of the noise source degrades at sunrise as well.

normal way, with conversion to temperature units and separation of sidebands. The visibility amplitude and phase of the fiducial noise spike for all six baselines appear in Figure 6.6. The noise spike looks substantially different from the Jupiter data of Figure 6.4: it is visible only in the lower sideband;¹¹ and its phase has much more structure, most likely due to various path length differences and reflections owing to the waveguide that delivers the signal to each receiver. We remove the phase of the fiducial noise spike from all other noise spikes during the campaign, and then fit a line to the remaining phase of each channel, baseline, and spike, recording the offset and

¹¹The LSB-only profile is an inherent property of the MINT noise source (see Section 2.6).

Baseline	f_{loss}	θ_{eff}
D–A	6.0%	20°
C–B	5.7%	20°
D–C	2.5%	13°
A–B	2.3%	12°
D–B	7.4%	22°
A–C	5.0%	18°

Table 6.2: The sensitivity loss due to time averaging visibilities with varying phases. Phases are as measured by the noise source through the campaign. The value $\theta_{eff} = \cos^{-1}(1 - f_{loss})$. Table provided by Joseph Fowler.

slope of the line. A phase offset would be due primarily to a delay in the RF band, most likely in the noise source waveguide. A delay at baseband would translate to a phase slope. An IF delay would produce both a phase slope and an offset. Plots of a single night’s offset and slope data for the Blue channel (5.5–6.0 GHz) are found in Figure 6.7. We note a few interesting trends. First, all baselines are phase-stable to a few degrees rms over the night (before sunrise). Second, at sunrise, the phases become unstable. The instability tracks the temperature of the waveguide, indicating that thermal expansion of the waveguide and waveguide joints during sunrise may be the culprit. Third, at sunrise, the baselines A–B and C–D are more stable than the other four, which has a plausible explanation in light of Figure 2.10. A slightly loose joint between the first power splitter following the noise source and, say, the A–B power splitter would cause the observed pattern. The C–D baseline would be unaffected, and the A and B phase changes would be in common, leaving A–B unchanged. It is also possible that this phase change is not due to the noise source, but to the telescope itself (the A and B optics could thermally flex in common, for instance). Since the noise source provides the only real-time assessment of phase stability, we have no way of knowing if the problem is in the noise source or the instrument itself, and decide to cut all CMB data after sunrise (see Appendix B).

Once the sunrise cut has been applied, the remaining phase data appear to be random noise, showing no real trends through the campaign. As derived in Section 6.3, a fluctuation of θ_{rms} across the band in each baseline causes a fraction signal loss of $f_{loss} \sim 1 - \cos \theta_{rms}$ in that baseline. We average f_{loss} over the campaign, and find

that even if all the detected phase fluctuations were inherent to the telescope, (i.e. none had to do with the noise source or its waveguide) the maximum fractional loss in sensitivity would be only 7.4% (see Table 6.2). The table also shows the effective rms phase fluctuation, $\theta_{eff} = \cos^{-1}(1 - f_{loss})$. This quantity is different from the one that could be more simply derived by taking the rms of the distribution of phases, because in general, $\langle \cos \theta \rangle \neq \cos \langle \theta \rangle$.

The noise source also provides a measure of amplitude stability for the measured visibility. Fowler (2002) finds that the amplitude of the noise source visibilities vary by 5% rms through the campaign. As with the phase, we have no way to separate inherent telescope gain changes from amplitude changes of the noise source, so 5% is an upper limit on the rms gain fluctuations of the telescope.¹²

6.5 Mars Calibrations

During the 2001 observing campaign, Mars was visible to MINT ($\theta_{el} > 60^\circ$) during the late afternoon and early evening. It was observed rising (East); overhead (South in azimuth, and backward ($\theta_{el} > 90^\circ$) in elevation, net pointing to the North); and setting (West), whenever instrument maintenance and weather allowed. Mars, unlike Jupiter, was observed *directly*, without the plane mirror, and so provides phase and amplitude calibrations that will be directly applied to the CMB data. Unfortunately, Mars was dim during the campaign, so each individual observation had poor signal-to-noise. In fact, integration over all 34 usable Mars calibrations does not yield nearly the SNR of the single Jupiter observation. Table 6.3 and 6.4 give some information about all 48 attempted observations. The scan patterns of the calibrations that yielded data are pictured in Figure 6.8.

The analysis of the Mars observations parallels the Jupiter analysis. First, we remove the fringe phases, and average all the calibrations together over the inner $1\text{-}\sigma$ of the cross-beam to produce a phase template for each baseline (see Figure 6.9).

¹²We have ignored, in our discussions of phase and amplitude fluctuations, the pathological case of the noise source and the instrument having anti-correlated fluctuations, where a change in the telescope is compensated by an opposing change in the noise source. Anti-correlated fluctuations would derail the argument that noise-source-based measurements represent the worst case scenario of possible fluctuations.




Figure 6.8: Scan patterns for the 34 used Mars calibrations. Each point represents 8 seconds of data. The axes are (Az,El), Mars position subtracted from telescope pointing, with a range of ± 0.3 in each. The circles in each plot represent the FWHM of the large and small beams.

Cal. #	Date	UT Beg. (hrs)	UT End (hrs)	# Recs.	R/O/S	Diam. (arcsec)	Why not used
0	Nov. 21	20.400	21.120	–	–	–	Moon contam.
1	Nov. 21	21.120	22.272	–	–	–	Moon contam.
2	Nov. 21	22.272	23.966	–	–	–	Moon contam.
3	Nov. 22	20.088	21.024	–	–	–	Moon contam.
4	Nov. 22	21.864	22.248	–	–	–	Moon contam.
5	Nov. 22	22.656	23.999	–	–	–	Moon contam.
6	Nov. 23	20.232	20.976	526	R	7.7	–
7	Nov. 23	21.288	22.224	846	O	7.7	–
8	Nov. 24	0.7200	1.4400	–	–	–	Not on Mars
9	Nov. 24	19.704	20.160	485	R	7.6	–
10	Nov. 25	19.992	20.904	873	R	7.6	–
11	Nov. 25	22.512	23.424	913	S	7.6	–
12	Nov. 27	20.400	20.832	–	–	–	Not on Mars
13	Nov. 27	22.272	23.400	–	–	–	Not on Mars
14	Nov. 29	19.776	20.688	1199	R	7.4	–
15	Nov. 29	22.224	23.400	1806	S	7.4	–
16	Nov. 30	19.968	20.808	–	–	–	Not on Mars
17	Nov. 30	22.200	23.712	476	S	7.4	–
18	Dec. 1	19.680	20.544	885	R	7.3	–
19	Dec. 1	21.168	21.840	488	O	7.3	–
20	Dec. 1	22.200	23.750	508	S	7.3	–
21	Dec. 2	19.752	20.640	832	R	7.3	–
22	Dec. 4	22.248	23.592	1460	S	7.2	–
23	Dec. 5	22.200	22.680	244	S	7.2	–

Table 6.3: A list of all Mars calibrations listed in the Observing Log. “# Recs.” is the number of data records (2 seconds each) used from the calibration. R/O/S refers to whether the calibration was when Mars was rising, overhead, or setting. “Moon contam.” refers to observations during which the Moon was close enough to Mars to contribute correlated power in the sidelobes. “Not on Mars” refers to scans reported in the Observing Log as being Mars calibrations, but that were either not Mars calibrations at all, or were attempted calibrations whose scans were too far from the planet to be useful. Continued in Table 6.4.

Cal. #	Date	UT Beg. (hrs)	UT End (hrs)	# Recs.	R/O/S	Diam. (arcsec)	Why not used
24	Dec. 6	22.176	23.232	726	S	7.1	–
25	Dec. 7	19.488	20.160	704	R	7.1	–
26	Dec. 7	22.176	23.448	979	S	7.1	–
27	Dec. 8	19.632	20.280	636	R	7.1	–
28	Dec. 8	22.152	23.400	1104	S	7.1	–
29	Dec. 9	19.536	20.160	720	R	7.0	–
30	Dec. 9	22.176	23.400	1253	S	7.0	–
31	Dec. 10	19.464	20.136	767	R	7.0	–
32	Dec. 10	22.200	23.400	1798	S	7.0	–
33	Dec. 11	19.704	20.160	525	R	6.9	–
34	Dec. 11	22.200	23.400	1794	S	6.9	–
35	Dec. 12	19.512	20.112	672	R	6.9	–
36	Dec. 12	20.400	20.976	–	–	–	Not on Mars
37	Dec. 12	21.552	21.720	–	–	–	Not on Mars
38	Dec. 12	22.224	23.352	1723	S	6.9	–
39	Dec. 13	19.488	19.944	512	R	6.9	–
40	Dec. 13	22.200	23.352	1770	S	6.9	–
41	Dec. 14	19.440	19.992	694	R	6.8	–
42	Dec. 15	19.440	19.920	718	R	6.8	–
43	Dec. 15	22.128	23.280	1506	S	6.8	–
44	Dec. 16	22.272	23.280	–	–	–	Bad data
45	Dec. 25	22.128	22.920	1129	S	6.5	–
46	Dec. 26	22.152	22.848	–	–	–	Bad data
47	Dec. 28	22.128	22.848	894	S	6.4	–

Table 6.4: The continuation of Table 6.3 “Bad data” refers to occasions when the LO phase-switching angles were set incorrectly by the Pointing PC, causing the sideband separation algorithm to yield nonsensical results.

This phase template is not only removed from all Mars calibration data before averaging over frequency; we will also remove it from all CMB data. Next we take a step backwards, and take an average over frequency and the inner $1\text{-}\sigma$ of each defringed calibration individually. The result, one complex visibility per baseline per calibration, is used to assess the phase stability of the instrument over the campaign. Some results are reported in Table 6.5.

	Mars	Jupiter	Mars	Mars	Noise
	$ \mathcal{V}_{peak} $	$ \mathcal{V}_{peak} $	f_{loss}	f_{loss}	f_{loss}
Baseline	meas./pred.	meas./pred.	All cals	Rising Only	
D–A	33%	35%	4.1%	2.2%	6.0%
C–B	75%	–	6.3%	1.8%	5.7%
D–C	66%	–	2.6%	2.0%	2.5%
A–B	32%	36%	3.6%	1.8%	2.3%
D–B	74%	77%	4.8%	3.2%	7.4%
A–C	33%	–	3.3%	1.0%	5.0%

Table 6.5: Signal loss due to phase fluctuations as measured during Mars calibrations. The fractional peak signals in Jupiter and Mars are strongly correlated. The average loss due to phase fluctuations in the setting Mars calibrations is only 2%.

In the course of evaluating phase stability from the Mars data, we discover a small telescope pointing error that, by shifting the fringes on the sky, has a signature in the visibility phases of the various calibrations. We plot the visibility phase of each calibration and baseline, θ_i , vs. Mars’ azimuth (see left panels of Figure 6.10) and find an interesting pattern. The phases of observations made in the East ($Az \sim 70^\circ$) have a phase offset from those made in the West ($Az \sim 290^\circ$) of $\sim 60^\circ$ in four of the six baselines. The affected baselines have a non-zero v -component of $\vec{u} = (u, v)$, and the sign of the phase offset matches the sign of v . This pattern is explainable by an East-to-West tilt in the azimuth bearing, such that when the telescope looks East its elevation angle is a little higher than recorded by the encoder, and when it looks West it is pointing a little low. A fit to the data finds a tilt of $0^\circ.034$. When the gear tilt is corrected (Equation 2.30), the phase offset disappears (right panels of Figure 6.10).

We search for other possible systematic phase effects akin to the azimuth gear tilt, looking for correlations among phases: one baseline vs. another, to discover phase effects related to a single receiver; vs. Mars elevation, time of day, day of year, to probe for other offsets or shifts over time in the pointing; vs. outside temperature to look for thermal effects, and find no significant correlations. There is one correlation in the *distribution*, but not the average, of the visibility phases that is apparent in Figure 6.10: that phases from observations made in the East ($Az \sim 70^\circ$) have a much noisier distribution than those in the West ($Az \sim 290^\circ$). The Eastern observations are




Figure 6.9: Average magnitude and phase of the visibility for Mars. Magnitude is plotted as a fraction of the expected peak signal; since Mars' angular diameter, and thus its antenna temperature, changes through the campaign, we must normalize before averaging. As in the equivalent Jupiter plot (Figure 6.4), data are averaged over the central $1\text{-}\sigma$ of the cross-beams, so the average magnitude shown in this plot is lower than that quoted in Table 6.5.

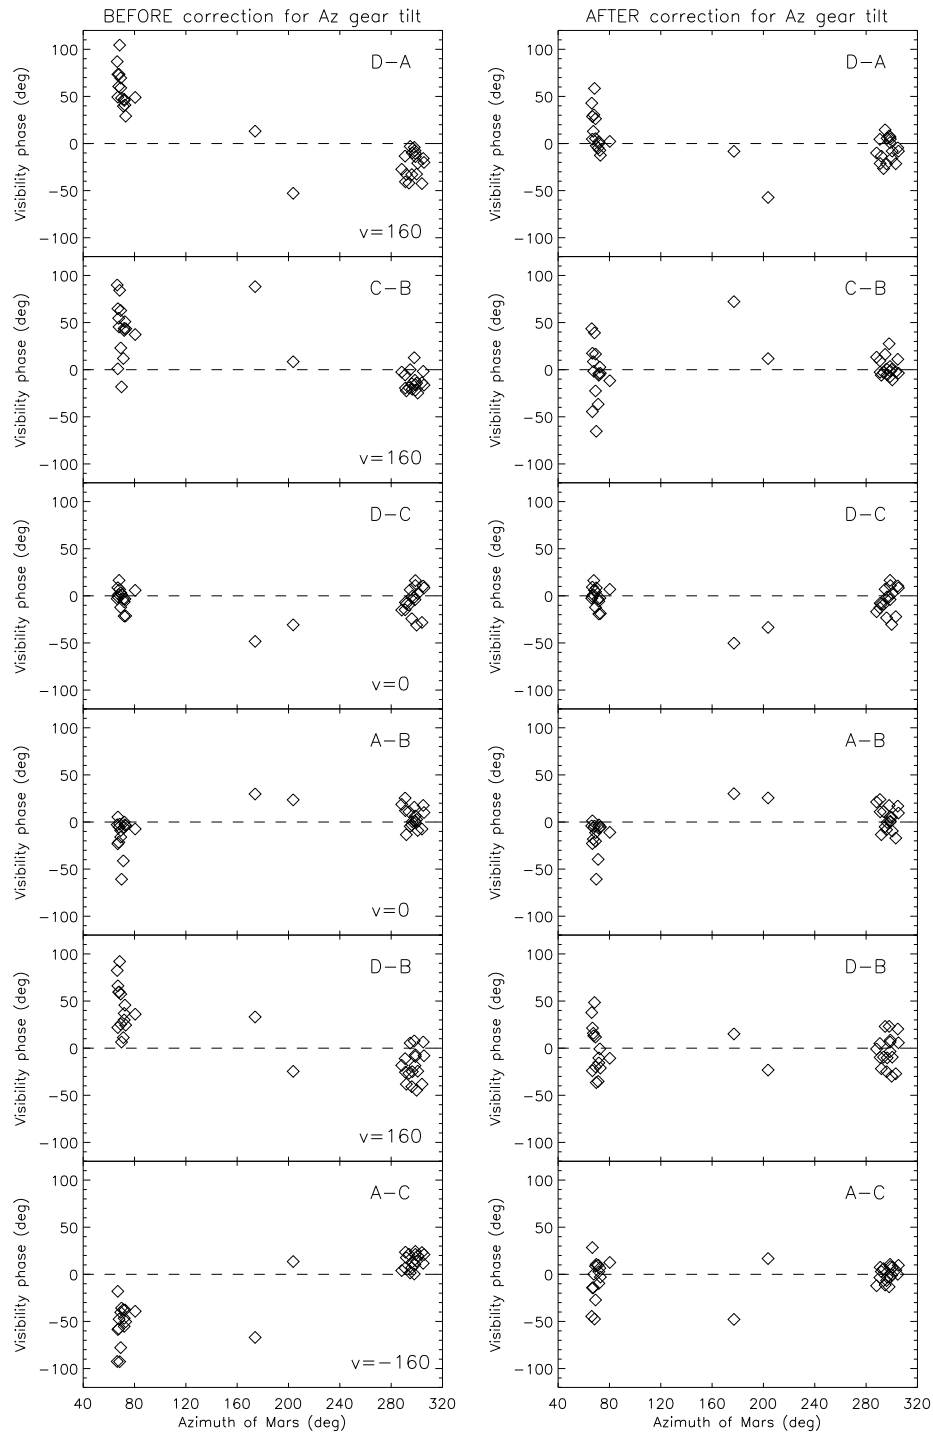


Figure 6.10: Average phase of the Mars visibility vs. azimuth of the calibration. Each of the 34 calibrations provides one data point. On the left are the raw data. The East data have a $\sim 60^\circ$ average phase offset from the West data in baselines that have a non-zero v component of \vec{u} . On the right are the data with the correction for the tilt in the azimuth bearing of $0^\circ.034$.

of Mars rising, and occur at ~ 5 P.M. local when the sun is still shining on the telescope. Those toward the West are of Mars setting, and occur after sunset at ~ 8 P.M. Because of thermal expansion effects due to sunlight on the telescope, and also possible mechanical stresses from wind (late afternoon is generally the windiest time of day at the site) we have strong reason to believe that the rising Mars calibration data are not representative of the stability of the instrument during CMB observations.

We believe the setting-Mars ($Az > 250^\circ$), azimuth-tilt-corrected data in the right-hand panels of Figure 6.10 represent the fundamental limit on phase stability of the instrument over time (see Table 6.5). Akin to the analysis of signal loss due to noise source phase fluctuations in Section 6.4, we calculate the signal loss over the campaign due to Mars phase fluctuations:

$$f_{loss} = 1 - \langle \cos \theta_i \rangle. \quad (6.17)$$

We show this value for each baseline in Table 6.5, comparing the value using all Mars calibrations to that averaged over only the setting calibrations. The telescope is certainly stable enough over the 38-day observing campaign for CMB work, with average signal loss due to phase fluctuations of 2%.

Finally, we calculate the fraction of Mars' expected peak visibility magnitude that we measure (Table 6.5). We follow the data processing steps for Jupiter, removing the phase-frequency template, averaging the data over pixels, and then taking the magnitude of each pixel. The expected magnitude is calculated via Equation 6.16, using $T_B = 198$ K (this is the 171 GHz value of Goldin et al. 1997).¹³ We find (in good agreement with the Jupiter results, where the comparison is available) that the measured peak visibility amplitude is substantially less than expected, especially in baselines involving receiver A. We currently have no good explanation for what is causing this decrease in sensitivity; perhaps observations of Jupiter planned for October–November, 2002 from the roof of Jadwin Hall will shed light on the matter. If this effect were real, and we have no reason to doubt that it will be confirmed by further testing, it would leave MINT in a position only to place a weak upper limit (200–300 μ K, 95%-confidence level) on the CMB anisotropy at $\ell \sim 1000$ and $\ell \sim 1600$ (Fowler 2002).

¹³The expected magnitude is calculated for each calibration separately, since Mars' angular diameter changes significantly through the campaign (see Tables 6.3 and 6.4).

Conclusion

Overall, the MINT telescope worked well during the 2001 Chilean observing campaign. Here is a partial list of the design and observational accomplishments:

- Design, construction, testing, and field-use of four broadband, SIS-based, 145-GHz receivers. We have developed reliable methods for tuning the receivers in the lab, and adapting them for the field. The receivers performed well in the field under difficult environmental conditions.

- Design, construction, testing, and field-use of a frequency down-converter and band splitter (the channelizer), designed as a microwave integrated circuit.

- Design, construction, testing, and field-use of a 1-GHz-sample-rate, 2-bit digital correlator. The correlator pushes the envelope on high-speed digital design, and was stable in the field.

- Integration of these components into a stable and systematic-error-free interferometer. The instrument has gain stability of 5%, and phase stability such that only $\sim 2\%$ of sensitivity is lost when integrating over nearly 40 days of data. Tran (2002, §7.2.3) has shown that the data show no excess correlated noise or spectral lines at the $15 \mu\text{K}$ level in an overnight integration. We believe we have successfully avoided the twin scourges of correlated atmosphere and ground pickup.

- Several months operating experience in Chile. Along with the time spent operating the MAT/TOCO experiment in 1997 and 1998, this experience in remote observing and local infrastructure development will be invaluable to the next big project, the Atacama Cosmology Telescope (ACT).

- Verification of Cerro Toco specifically and the Atacama Desert region of northern Chile in general as a viable site for CMB anisotropy experiments. A 10.5 hour

period of the night, beginning just before midnight UT, had a typical atmosphere temperature of 5 K.

While the CMB data have not been fully analyzed at the time of this writing, it has become apparent through the analysis of calibration data in Chapter 6 that MINT will not have achieved the sensitivity needed to detect CMB anisotropy during its first campaign. We give some reasons for the failure and items that could be improved:

- The receivers, especially receiver A, need to be tweaked in tuning to match the specific conditions in the field in Chile. The largest source of error is likely the LO power transmitted to the SIS, because of reflections, temperature shifts of the RBE from indoor to outdoor operation, etc. The LO could be adjusted in the field through the use of the simple sky dips and hot-load/sky tests described in Chapter 4.

- The large unexplained deficit in the visibility amplitude when observing planets needs to be investigated. An investigation will occur in the Fall of 2002 on the roof of Jadwin Hall.

- A series of unusual mechanical failures during the campaign ate away at MINT's observing time. Based on experience from MINT, MAT/TOCO, and the findings of the CBI and ALMA groups, any future observing campaigns should be targeted to begin in late July or early August and run through the first week in January.

- MINT's sensitivity is hurt by its small IF bandwidth of only 2 GHz. Even a factor of two improvement would be a big step forward. We understand that NRAO has a wider-bandwidth SIS/HEMT system in development.

MINT's primary contribution to the field is the demonstration that a stable 150-GHz interferometer can be operated in the Atacama Desert. The high frequency has the advantage of immunity to extra-galactic point sources, as compared to other interferometers which operate at ~ 30 GHz. With the addition of 220- and 260-GHz channels and expansion to slightly longer baselines, MINT could observe the SZ effect in decrement, null, and increment. Other CMB interferometers, such as CBI and DASI, have 13 times the number of baselines and five times the IF bandwidth. MINT could be expanded to this scale, but at a cost of many millions of dollars.

There are no immediate plans to take MINT back to Chile. Extensive post-campaign tests will occur on the roof of Jadwin Hall during fall, 2002. Among other things, we will look for cross-talk between receivers and the cause of the mysterious point-source signal deficit.

Calibration Odds & Ends

A.1 Diode Linearization

For square-law power detection in MINT, we use the Hewlett Packard (now Agilent Technologies) 8472, a low-barrier Schottky diode, to which we add a video resistor of about $1\text{ k}\Omega$ and a pre-amp of gain $G \sim -50$. The voltage output response of the diode rolls off at high power (see Fig. A.1). We measure this effect using a calibrated




Figure A.1: (left) Output diode voltage vs. input power for a typical diode. The asterisks are measurements. The line is the fit to the low-power data, and is forced to include the origin. (right) The output voltage correction necessary to linearize this diode.

Diode	a_0	a_1	a_2	a_3	a_4
A DC level	0.0064	1.0194	-0.1186	0.0461	-0.0049
B DC level	0.0093	0.9821	0.0833	-0.0742	0.0377
C DC level	0.0055	0.9855	0.0555	—	—
D DC level	-0.0005	0.9739	0.0886	0.0091	0.0009
Chan. 1 (Blu)	-0.0020	0.9816	0.0546	0.0024	-0.0001
Chan. 2 (Grn)	-0.0005	0.9827	0.0396	0.0065	-0.0006
Chan. 3 (Yel)	0.0004	0.9686	0.0689	-0.0012	0.0002
Chan. 4 (Red)	-0.0008	0.9763	0.0516	0.0030	0.0000

Table A.1: Fourth-order correction coefficients for detector diode voltages. The calibration measurements for Diode C were not made at a sufficiently high input power level to allow fitting of the third- and fourth-order coefficients.

input power source for each of the eight diodes used in MINT.¹ The resulting 4th order correction curve is described by

$$V_c(\text{V}) = a_0 + a_1 V_o + a_2 V_o^2 + a_3 V_o^3 + a_4 V_o^4. \quad (\text{A.1})$$

The coefficients a_n are listed in Table A.1 for the eight diodes. During the campaign, the power into the DC level diodes is attenuated to a level that they remain almost exclusively in the linear regime. Only during hot-load (270 K) observations do the corrections approach even 20%. The high-power ranges are included more for legacy users of these diodes than for any use during MINT analysis.

A.2 Digitizer Histogram Linearization

The digitizer histogram power measurement, described in Section 4.4, also has well-characterized non-linearity at high input power, except that unlike the diodes, the measured power *increases* with increasing input. We made measurements of this phenomenon in the lab on two channels (Correlator 2, Digitizer 2, and Corr. 1, Dig. 3),² and found that they behaved the same, so we assume the non-linearity is

¹The four receivers (A-D) each have a diode for the DC level readout. We also used a diode for each channelizer channel during the lab tuning and for sanity-checks in the field.

²We use the numbering convention of 0-3 for both the digitizers and correlators.

Dig.	Corr.	0.5 dB	1 dB	2 dB	4 dB	8 dB	16 dB
0	0	0.54	1.03	2.02	3.93	8.14	15.95
1	0	0.55	1.06	2.07	4.02	8.24	16.01
2	0	—	—	—	—	—	—
3	0	0.57	1.06	2.08	4.02	8.25	15.89
0	1	0.55	1.05	2.07	4.01	8.22	16.00
1	1	0.54	1.04	2.05	3.98	8.19	15.98
2	1	0.52	1.05	2.07	4.02	8.24	15.94
3	1	0.54	1.03	2.05	3.97	8.18	16.03
0	2	—	—	—	—	—	—
1	2	0.56	1.06	2.07	4.01	8.27	15.98
2	2	0.55	1.05	2.07	4.02	8.22	16.05
3	2	0.55	1.04	2.06	4.01	8.20	16.04
0	3	0.56	1.07	2.10	4.04	8.26	16.06
1	3	0.56	1.07	2.08	4.04	8.24	16.01
2	3	0.56	1.06	2.09	4.04	8.26	16.05
3	3	0.57	1.07	2.10	4.04	8.25	16.06

Table A.2: Measured attenuation of the digitizer attenuators. Dig. 0, Corr. 2 is the broken channel, and Dig. 2, Corr. 0 had a malfunctioning readout during these measurements. All measurements have error ± 0.02 dB.

a common property of the digitizers. The correction applied to the measured power (in mW) at the digitizer is well approximated by

$$P_c(\text{mW}) = P_o - 0.0063P_o^2. \quad (\text{A.2})$$

A.3 Digitizer Attenuator Calibration

The digitizer boards contain settable attenuators with nominal values of 0.5, 1, 2, 4, 8, and 16 dB that can be added in series to create any attenuation between 0 and 31.5 dB in steps of 0.5 dB. In practice, the attenuators deviate slightly from their nominal values. We present lab measurements of the actual attenuation in Table A.2.

Microm.	Trans. (linear)	Atten. (dB)
70.0	0.98	0.1
64.0	0.92	0.4
61.0	0.86	0.7
59.0	0.81	0.9
57.3	0.75	1.3
55.8	0.69	1.6
54.4	0.63	2.0
53.2	0.58	2.4
52.0	0.52	2.8
50.9	0.46	3.4
49.6	0.40	4.0
48.4	0.35	4.6
47.1	0.29	5.4
45.6	0.23	6.4
44.0	0.17	7.7
41.8	0.12	9.2
38.6	0.058	12.4
35.6	0.029	15.4
34.7	0.023	16.4
33.7	0.017	17.7
32.2	0.012	19.2
29.4	0.0058	22.4
27.5	0.0029	25.4

Table A.3: Measured transmission of Millitech LSA-06 as a function of micrometer setting.

A.4 LO Attenuator Calibration

The level-set attenuator (Millitech, LSA-06) used to vary the LO power in the SIS tuning rig works by protruding an absorbing vane into the waveguide. In Chapter 4 we quote the position of the micrometer knob as a measure of the LO power. A calibration of the LSA attenuation vs. its micrometer reading, made by Eric Torbet during preparation for the MAT/TOCO experiment, is found in Table A.3 and Figure A.2.

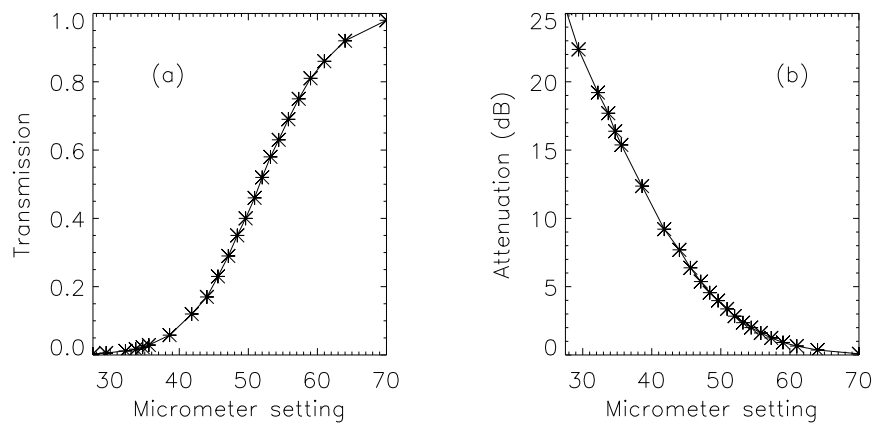


Figure A.2: (a) Transmission (linear) and (b) attenuation (dB) of the Millitech LSA-06.

Data Cuts

Description Of Cut	Hrs. Remaining After Cut	% of Total Data Cut
All Data	912	–
Daylight (10:05 UT–23:35 UT)	399	56.3%
Telescope tarped (usually due to snow)	322	8.4%
Calibrations (Mars or Hot Load/Sky)	303	2.1%
Data PC drive full, no data written	294	1.0%
Pointing PC down or El encoder failed	258	3.9%
Unlocked PLL or Coldhead Temp. Servo	252	0.7%
Thermalizing after tarp removed, Correlator readout unsynchronized, Bad angle loaded into PSC	236	1.8%
Noise source on, telescope moving, Other minor cuts (bad digital reads, etc)	225	1.2%

Table B.1: The preliminary MINT data cuts. There are 38 days (912 hours) of raw data. We end up with a yield of 24.7% usable data. Not included here are any cuts based on sky temperature.

Bibliography

- Aboobaker, Asad M., 2001. *The MINT Oscillator System*. Experimental project, Physics Department, Princeton University. Available on the MINT website at: <http://imogen.princeton.edu/mintweb/publications/>.
- Aboobaker, Asad M., 2002. *A 139-141 GHz Noise Source for Calibrating the MINT Telescope*. Generals project, Physics Department, Princeton University. Available on the MINT website at: <http://imogen.princeton.edu/mintweb/publications/>.
- Bahcall, N. A. and X. Fan, 1998. The most massive distant clusters: Determining Omega and delta 8. *Astrophys J*, 504:1–6.
- Bahcall, N. A., L. M. Lubin, and V. Dorman, 1995. Where is the dark matter? *Astrophys J Lett*, 447:L81–L85.
- Bahcall, N. A., J. P. Ostriker, S. Perlmutter, and P. J. Steinhardt, 1999. The cosmic triangle: Revealing the state of the universe. *Science*, 284:1481–1488.
- Beauchamp, Kenneth G., 1975. *Walsh Functions and their Applications*. Academic Press, London.
- Best, Roland E., 1999. *Phase-Locked Loops: Design, Simulation, and Applications*. McGraw-Hill Professional Publishing, New York.
- Blundell, R., R. E. Miller, and K. H. Gundlach, 1992. Understanding noise in SIS receivers. *Int J IR Millimeter Waves*, 13:3–14.
- Bond, J. R. et al., 2002. The Sunyaev-Zeldovich effect in CMB-calibrated theories applied to the Cosmic Background Imager anisotropy power at $\ell > 200$. *Astrophys J*. Submitted, see astro-ph/0205386.

- Caldwell, R. R., R. Dave, and P. J. Steinhardt, 1998. Quintessential cosmology: Novel models of cosmological structure formation. *Astrophys Space Sci*, 261:303–310.
- Callen, Hebert B. and Theodore A. Welton, 1951. Irreversability and generalized noise. *Phys Rev*, 83:34–40.
- Carlberg, R. G., H. K. C. Yee, E. Ellingson, R. Abraham, P. Gravel, S. Morris, and C. J. Pritchett, 1996. Galaxy cluster virial masses and Omega. *Astrophys J*, 462:32.
- Dayem, A. H. and R. J. Martin, 1962. Quantum interaction of microwave radiation with tunneling between superconductors. *Phys Rev Lett*, 8:246–248.
- de Bernardis, P., P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, B. P. Crill, G. De Gasperis, P. C. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, A. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, S. Prunet, S. Rao, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforna, and N. Vittorio, 2000. A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404:955–959.
- Dumont, Charles, 2001. *Digital Correlation between Two Cryogenic SIS Mixers*. Senior thesis, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Efstathiou, G., S. Moody, J. A. Peacock, W. J. Percival, C. Baugh, J. Bland-Hawthorn, T. Bridges, R. Cannon, S. Cole, M. Colless, C. Collins, W. Couch, G. Dalton, R. de Propris, S. P. Driver, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackcon, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, P. Norberg, B. A. Peterson, W. Sutherland, and K. Taylor, 2002. Evidence for a non-zero Lambda and a low matter density from a combined analysis of the 2dF Galaxy Redshift Survey and cosmic microwave background anisotropies. *Mon Not R Astron Soc*, 330:L29–L35.

- Feldman, M. J., S. K. Pan, A. R. Kerr, and A. Davidson, 1983. SIS mixer analysis using a scale model. *IEEE Trans Magnetics*, 19:494.
- Fowler, Joseph, 2001. *MINT Phasing Requirements and Sideband Separation*. Internal group memo, MINT group, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Fowler, Joseph, 2002. *MINT Signal-to-Noise, Second Edition*. Internal group memo, MINT group, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Franklin, G. F., J. D. Powell, and A. Emami-Naeini, 1994. *Feedback Control of Dynamic Systems*. Addison-Wesley.
- Gardner, Floyd Martin, 1979. *Phaselock Techniques*. John Wiley and Sons, New York.
- Goldin, A. B., M. S. Kowitt, E. S. Cheng, D. A. Cottingham, D. J. Fixsen, C. A. Inman, S. S. Meyer, J. L. Puchalla, J. E. Ruhl, and R. F. Silverberg, 1997. Whole-disk observations of Jupiter, Saturn, and Mars in millimeter/submillimeter bands. *Astrophys J Lett*, 488:L161.
- Guth, A. H., 1981. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys Rev D*, 23:347–356.
- Halverson, N. W., E. M. Leitch, C. Pryke, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd, 2002. Degree angular scale interferometer first results: A measurement of the cosmic microwave background angular power spectrum. *Astrophys J*, 568:38–45.
- Hedman, Matthew, 2002. *The Princeton IQU Experiment and Constraints on the Polarization of the Cosmic Microwave Background at 90 GHz*. Ph.D. thesis, Princeton University, Physics Department, Princeton, New Jersey.

- Hinderks, James, 1999. *The Theory, Design, and Implementation of a High-Speed Digital Correlator*. Senior thesis, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Hu, Wayne and Scott Dodelson, 2002. Cosmic microwave background anisotropies. *araa*. Submitted, see astro-ph/0110414.
- Hunter, Todd R., Robert Kimberk, and Robert W. Wilson, 2001. *Design and Performance of a Digital PLL for Gunn Oscillators*. SMA Technical Memo #143, Smithsonian Astrophysical Observatory, Harvard University. Available on the SMA website at:
http://sma-www.harvard.edu/private/memos/tech_no.html.
- Jackson, John David, 1975. *Classical Electrodynamics*. John Wiley & Sons, New York.
- Jones, B. J. T. and R. F. G. Wyse, 1985. The ionisation of the primeval plasma at the time of recombination. *Astron Astrophys*, 149:144–150.
- Kerr, A. R., M. J. Feldman, and S. K. Pan, 1997. Receiver noise temperature, the quantum noise limit, and the role of the zero-point fluctuations. In *Proc. 8th Intl. Symp. Space THz Techn.*
- Kerr, A. R., S. K. Pan, A. Lichtenberger, and N. Horner, 1993. A new mixer for the 2-mm band. In *Proc. 4th Intl. Symp. Space THz Techn.*
- Khoury, J., B. A. Ovrut, P. J. Steinhardt, and N. Turok, 2001. Ekpyrotic universe: Colliding branes and the origin of the hot big bang. *Phys Rev D*, 64:123522.
- Lee, A. T., P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, S. Hanany, V. V. Hristov, A. H. Jaffe, P. D. Mauskopf, C. B. Netterfield, E. Pascale, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu, 2001. A high spatial resolution analysis of the maxima-1 cosmic microwave background anisotropy data. *Astrophys J Lett*, 561:L1–L5.

- Leitch, E. M., C. Pryke, N. W. Halverson, J. Kovac, G. Davidson, S. LaRoque, E. Schartman, J. Yamasaki, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd, 2002. Experiment design and first season observations with the Degree Angular Scale Interferometer. *Astrophys J*, 568:28–37.
- Loh, Yeong Shang, 2000a. *MINT Optical Layout*. Experimental project, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Loh, Yeong Shang, 2000b. *MINT Optics*. Generals project, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Marriage, Tobias A., 2000. *The MINT Digitizer*. Senior thesis, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Marriage, Tobias A., 2002. *MINT 2001 Pointing Calibration with Sun as Source*. Experimental project, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- Mason, B. S. et al., 2002. The anisotropy of the microwave background to $\ell = 3500$: Deep field observations with the Cosmic Background Imager. *Astrophys J*. Submitted, see astro-ph/0205384.
- Mather, J. C., E. S. Cheng, D. A. Cottingham, R. E. Eplee, D. J. Fixsen, T. Hewagama, R. B. Isaacman, K. A. Jensen, S. S. Meyer, P. D. Noerdlinger, S. M. Read, L. P. Rosen, R. A. Shafer, E. L. Wright, C. L. Bennett, N. W. Boggess, M. G. Hauser, T. Kelsall, S. H. Moseley, R. F. Silverberg, G. F. Smoot, R. Weiss, and D. T. Wilkinson, 1994. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *Astrophys J*, 420:439–444.

- Mauskopf, P. D., P. A. R. Ade, P. de Bernardis, J. J. Bock, J. Borrill, A. Boscaleri, B. P. Crill, G. DeGasperis, G. De Troia, P. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, S. Hanany, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, A. T. Lee, S. Masi, A. Melchiorri, F. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, P. L. Richards, G. Romeo, J. E. Ruhl, E. Scannapieco, F. Scaramuzzi, R. Stompor, and N. Vittorio, 2000. Measurement of a peak in the cosmic microwave background power spectrum from the North American test flight of BOOMERaNG. *Astrophys J Lett*, 536:L59–L62.
- Miller, A. D., J. Beach, S. Bradley, R. Caldwell, H. Chapman, M. J. Devlin, W. B. Dorwart, T. Herbig, D. Jones, G. Monnelly, C. B. Netterfield, M. Nolta, L. A. Page, J. Puchalla, T. Robertson, E. Torbet, H. T. Tran, and W. E. Vinje, 2002. The QMAP and MAT/TOCO experiments for measuring anisotropy in the cosmic microwave background. *Astrophys J Suppl Ser*, 140:115–141.
- Miller, A. D., R. Caldwell, M. J. Devlin, W. B. Dorwart, T. Herbig, M. R. Nolta, L. A. Page, J. Puchalla, E. Torbet, and H. T. Tran, 1999. A measurement of the angular power spectrum of the cosmic microwave background from $\ell = 100$ to 400. *Astrophys J Lett*, 524:L1–L4.
- Miller, Amber. D., 2000. *A Measurement of the Cosmic Microwave Background from the High Chilean Andes*. Ph.D. thesis, Princeton University, Physics Department, Princeton, New Jersey.
- Netterfield, C. B., P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, C. R. Contaldi, B. P. Crill, P. de Bernardis, P. Farese, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, W. C. Jones, A. E. Lange, L. Martinis, S. Masi, P. Mason, P. D. Mauskopf, A. Melchiorri, T. Montroy, E. Pascale, F. Piacentini, D. Pogosyan, F. Pongetti, S. Prunet, G. Romeo, J. E. Ruhl, and F. Scaramuzzi, 2002. A measurement by boomerang of multiple peaks in the angular power spectrum of the cosmic microwave background. *Astrophys J*, 571:604–614.
- Ostriker, J. P. and P. J. Steinhardt, 1995. The observational case for a low-density universe with a non-zero cosmological constant. *Nature*, 377:600.

- Padin, S., J. K. Cartwright, B. S. Mason, T. J. Pearson, A. C. S. Readhead, M. C. Shepherd, J. Sievers, P. S. Udomprasert, W. L. Holzzapfel, S. T. Myers, J. E. Carlstrom, E. M. Leitch, M. Joy, L. Bronfman, and J. May, 2001. First intrinsic anisotropy observations with the Cosmic Background Imager. *Astrophys J Lett*, 549:L1–L5.
- Padin, S., M. C. Shepherd, J. K. Cartwright, R. G. Keeney, B. S. Mason, T. J. Pearson, A. C. S. Readhead, W. A. Schaal, J. Sievers, P. S. Udomprasert, J. K. Yamasaki, W. L. Holzzapfel, J. E. Carlstrom, M. Joy, S. T. Myers, and A. Otarola, 2002. The Cosmic Background Imager. *Publ Astron Soc Pac*, 114:83–97.
- Pearson, T. J. et al., 2002. The anisotropy of the microwave background to $\ell = 3500$: Mosaic observations with the Cosmic Background Imager. *Astrophys J*. Submitted, see astro-ph/0205388.
- Peebles, P. J. E., 1973. Statistical analysis of catalogs of extragalactic objects. I. Theory. *Astrophys J*, 185:413–440.
- Peebles, P. J. E., 1982. Large-scale background temperature and mass fluctuations due to scale-invariant primeval perturbations. *Astrophys J Lett*, 263:L1–L5.
- Peebles, P. J. E. and J. T. Yu, 1970. Primeval adiabatic perturbation in an expanding universe. *Astrophys J*, 162:815–836.
- Percival, W. J., W. Sutherland, J. A. Peacock, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. Cannon, S. Cole, M. Colless, C. Collins, W. Couch, G. Dalton, R. DePropris, S. P. Driver, G. Efstathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, S. Moody, P. Norberg, B. A. Peterson, and K. Taylor, 2002. Parameter constraints for flat cosmologies from CMB and 2dFGRS power spectra. *Mon Not R Astron Soc*. Submitted, see astro-ph/0206256.
- Perlmutter, S., G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer,

- B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and The Supernova Cosmology Project, 1999. Measurements of Omega and Lambda from 42 high-redshift supernovae. *Astrophys J*, 517:565–586.
- Plambeck, R. L., N. A. Thatte, and P. B. Sykes, 1993. A 4-K Gifford-McMahon refrigerator for astronomy. *Bull Am Astron Soc*, 25:886.
- Riess, A. G., A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry, 1998. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron J*, 116:1009–1038.
- Sachs, R. K. and A. M. Wolfe, 1967. Perturbations of a cosmological model and angular variations of the microwave background. *Astrophys J*, 147:73–90.
- Scott, P. F. et al., 2002. First results from the VSA. III. The CMB power spectrum. *Mon Not R Astron Soc*. Submitted, see astro-ph/0205380.
- Seljak, U. and M. Zaldarriaga, 1996. A line-of-sight integration approach to cosmic microwave background anisotropies. *Astrophys J*, 469:437–444.
- Silk, J., 1968. Cosmic black-body radiation and galaxy formation. *Astrophys J*, 151:459–471.
- Smoot, G. F., C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson, 1992. Structure in the COBE differential microwave radiometer first-year maps. *Astrophys J Lett*, 396:L1–L5.
- Sunyaev, R. A. and Ya. B. Zel'dovich, 1972. The observation of relic radiation as a test of the nature of X-ray radiation from the clusters of galaxies. *Comm Astrophys Sp Sci*, 4:173–178.

- Tegmark, M. and G. Efstathiou, 1996. A method for subtracting foregrounds from multifrequency CMB sky maps. *Mon Not R Astron Soc*, 281:1297–1314.
- Tegmark, Max, 1995. Doppler peaks and all that: CMB anisotropies and what they can tell us. In *Proc. Enrico Fermi, Course CXXXII, Varenna*.
- Thompson, A. R., J. M. Moran, and G. W. Swenson, 1994. *Interferometry and Synthesis in Radio Astronomy*. Krieger Publishing Company, Malabar, Florida. Referred to in text as TMS.
- Tien, P. K. and J. P. Gordon, 1963. Multiphoton process observed in the interaction of microwave fields with the tunneling between superconductor films. *Phys Rev*, 129:647–651.
- Timbie, P. T. and D. T. Wilkinson, 1984. A novel interferometer using SIS mixers. *Bull Am Astron Soc*, 16:517.
- Toffolatti, L., F. Argueso Gomez, G. de Zotti, P. Mazzei, A. Franceschini, L. Danese, and C. Burigana, 1998. Extragalactic source counts and contributions to the anisotropies of the cosmic microwave background: predictions for the Planck Surveyor mission. *Mon Not R Astron Soc*, 297:117–127.
- Torbet, Eric, 1997. *Atmospheric Emission in Chile*. Internal group memo, MAT/TOCO group, Physics Department, Princeton University.
- Touloukian, Y. S. and E. H. Buyco, 1970. *Thermophysical Properties of Matter*. Plenum, New York.
- Tran, Huan T., 2002. *A mm-Wave Interferometer for Studies of the Cosmic Microwave Background Anisotropy*. Ph.D. thesis, Princeton University, Physics Department, Princeton, New Jersey.
- Tucker, J. R., 1979. Quantum limited detection in tunnel junction mixers. *IEEE J Quantum Electron*, 15:1234–1257.
- Tucker, J. R. and D. T. Feldman, 1985. Quantum detection at millimeter wavelengths. *Rev Mod Phys*, 57:1055.

- Turner, M. S., 1997. The case for Λ CDM. In Turok, Neil, ed., *Critical Dialogues in Cosmology*, pp. 555–568. World Scientific, Singapore.
- Wang, X., M. Tegmark, and M. Zaldarriaga, 2002. Is cosmology consistent? *Phys Rev D*, 65:123001–123014.
- Wesley, Daniel, 2000. *Construction and Characterization of a Receiver for CMB Interferometry*. Senior thesis, Physics Department, Princeton University. Available on the MINT website at:
<http://imogen.princeton.edu/mintweb/publications/>.
- White, Martin, John E. Carlstrom, Mark Dragovan, and William L. Holzapfel, 1999. Interferometric observation of cosmic microwave background anisotropies. *Astrophys J*, 514:12–24.
- Wright, E. L., C. L. Bennett, K. Gorski, G. Hinshaw, and G. F. Smoot, 1996. Angular power spectrum of the cosmic microwave background anisotropy seen by the COBE DMR. *Astrophys J Lett*, 464:L21–L24.