

**A MM-WAVE
INTERFEROMETER FOR
STUDIES OF THE COSMIC
MICROWAVE BACKGROUND
ANISOTROPY**

Huan T. Tran

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

June 2002

© Copyright by Huan T. Tran, 2002. All rights reserved.

Abstract

The small angular scale, or high ℓ , component of the Cosmic Microwave Background anisotropy spectrum contains cosmological information about the early universe. In particular, the angular power spectrum above $\ell \sim 600$ should be damped due to photon diffusion between hot and cold regions and due to averaging along the line of sight. Fluctuations and damping at these scales have already been discovered at 30 GHz. At these frequencies, however, contamination from point sources can interfere with interpretation of these results. This work describes the design, construction, and performance of a 150 GHz, four-element interferometer built specifically to probe $\ell > 1000$ at frequencies where minimum point source contamination is expected. A new microwave channelizer and fully digital correlator are discussed in detail, and the performance during the first season of observation at a high Chilean site is examined.

Advisor: Professor Lyman Page Jr.

Acknowledgments

There is nobody more responsible for the completion of my thesis than Lyman Page. As my advisor, he has taught me more than I am aware of. I started at Princeton unenthusiastic about my future as an academic but through Lyman's example or pedagogic osmosis my life is now a single-minded pursuit of experimental beauty.

The Gravity group is a unique construction. Given the lack of funding, it is a surprise that it provided such a vibrant community, full of some incredibly smart people. All of the experimental groups, MINT, MAP, PIQUE, and Pulsar mingled freely with the theorists and all profoundly influence every person who wanders through the second floor of Jadwin.

Within the Gravity Group, I have to start with thanking the people who contributed to MINT. The "core" members were Joseph Fowler and Randy Dorwart. The three of us sacrificed our careers, lives, and unmeasurable amounts of sleep to a project that very few thought we could pull off. Tobias Marriage is not included in the core only because he was an undergrad for most of the time, but his contributions were so numerous and vital that they should not be acknowledged separately. In addition to Toby, the MINT project has suffered an embarrassing wealth of excellent undergrads and young grad students. Mark Tygert was there at the beginning, and James Hinderks, Charles Dumont, and Daniel Wesley wrote excellent senior theses with MINT. Yeong Loh provided the early optics and Asad Aboobakar came out to Chile and provided comic relief. A particular summer undergrad sticks in recent memory, my little brother, Long Tran. I've never seen anybody work so hard and impress so many people in such a little time.

From the larger experimental Gravity Group, many other people deserve thanks. Dave Wilkinson played the role of grandfather superbly. I do not know how after seeing so much science he still has abundant enthusiasm for inviting undergraduates to share in his immense experimental knowledge. Suzanne Staggs worked closely with us on MINT, helped me get a job, and served as a reader. Josh Gundersen, Matt Hedman, and Dennis Barkats were always around to set me straight. I am also grateful that the juggernaut called MAP was around to monetarily revive the Gravity

Group, distract Lyman, and bring so many people to Princeton and really make it an ideal classroom.

Nothing would be possible in the physics department without the extraordinary staff. All of the staff members, from the master machinists, including Ted Griffiths, Laszlo Varga, and William Dix, to the electronics technician, Stan Chidzik, and even to the purchasing staff, Kathleen Warren, Helga Murray, and Claude Champagne, contributed their creativity and hard effort toward accomplishing our scientific goals.

MINT could not have been designed and built in such a short time without the help of outside interferometry experts. David Hawkins of OVRO served as a model of scientific openness. He freely shared his ideas and designs on digital correlators. The MINT correlator was inspired by his careful designs, and he more than once helped us out of tough problems.

Walter Brisken and Mike Nolta really made Jadwin feel like home. Both were optimized as friends and roommates (imaginary or not) through a process of dragging to my schedule. Along with these two, Ken Nagamine and Randy also helped me get through Prelims. Eric Splaver and I took the very unwise option of attempting Generals in May, but somehow we still passed.

My family played a pivotal role in helping me through this process. Although they barely understand what I do, our story of escape, immigration, and success in our new home serves as a constant reminder of what is possible. I don't think that anything I do from now on will compare to what my family has already accomplished.

Of course, life at Princeton was not all about the lab. Romulus Johnson was my partner in misery early on and rescued me from myself more than I want to admit. Antonio Garcia and Stephanie Harves made me feel that scientists can get along with other people. Kimberly Bohman is that rare friend who has been with me for the full process, providing perspective, understanding, diversion, and life lessons, all for only the price of coffee. And finally, I wish I could but I cannot thank Janet Klein enough. She has not only read the entire thesis, but has provided heart and home, sometimes against her better judgment.

Contents

Abstract	iii
Acknowledgments	iv
1 Introduction	1
1.1 The Angular Spectrum	2
1.1.1 Low ℓ : The Sachs-Wolfe Plateau	4
1.1.2 Mid ℓ : Acoustic Oscillations	5
1.1.3 High ℓ : The Damping Tail	6
1.2 Previous Measurements	10
1.3 Foregrounds	11
1.3.1 Point Sources	11
1.3.2 Sunyaev-Zel'dovich Effect	12
2 Interferometry and Observing Strategy	13
2.1 Interferometer Response	14
2.1.1 Flat Sky Approximation	14
2.1.2 Primary Beam	14
2.1.3 Visibility	16
2.2 Sensitivity	19
2.2.1 Signal	19
2.2.2 Noise	26
2.2.3 SNR	27
2.2.4 Measuring δT_ℓ	28
2.2.5 Observation strategy	30

3	Instrument Overview	36
3.1	Advantages of Interferometry	36
3.1.1	Stability	37
3.1.2	Antenna Layout and Instantaneous Beam Switching	38
3.2	Signal Path	39
3.2.1	Antennas	41
3.2.2	Receiver	43
3.2.3	Phase Lock	47
3.2.4	Channelizer and Correlator	49
4	Channelizer	50
4.1	General Layout	50
4.2	Integrated Microwave Circuits	51
4.3	General Board Properties	53
4.4	Circuit Design	54
4.5	Transmission Lines	55
4.6	Filter Design	55
4.6.1	Filter Synthesis	56
4.6.2	Linear Simulation	57
4.6.3	Momentum Simulation	59
4.7	Power Splitter Design	60
4.7.1	Narrow Band Power Splitter	61
4.7.2	Broad Band Power Splitter	63
4.8	Combined Simulation	66
4.9	Assembly	67
4.9.1	Prototyping	67
4.9.2	Enclosure	67
4.10	Performance	69
4.10.1	Bandpass	69
4.10.2	Phase	71
4.11	Field Performance	76

5	Correlator Theory	77
5.1	Statistics	78
5.2	Sampling	79
5.3	Quantization	80
5.4	Multiplication Table	82
5.5	Corrected Correlation Coefficient	84
5.6	Expected Correlator Output	85
5.7	Summary	87
6	Correlator Implementation	88
6.1	Digitizer	89
6.1.1	Data Path	90
6.2	Correlator	91
6.2.1	Technology	91
6.2.2	Algorithm	92
6.3	Board Level Control	96
6.3.1	Threshold Servo Loop	96
6.3.2	Synchronization	96
6.4	Correlator Output	97
6.4.1	Mean	97
6.4.2	<i>rms</i>	99
6.4.3	Correlator Data Stream	99
6.4.4	Calibration Spikes	100
6.5	Summary	105
7	Interferometer Performance	107
7.1	Correlator Output Reduction	107
7.1.1	Correlation in Temperature Units	107
7.1.2	Sideband Separation	109
7.1.3	Correlator Synchronization	111
7.2	Stability	111
7.2.1	Amplitude Stability	112
7.2.2	Phase Stability	116

7.2.3	Integrating Down	126
7.3	Conclusion	127
A	Channelizer Details	131
A.1	Physical Parameters	131
A.2	Bandpass filter Design Equations	132
A.3	Jumper	135
B	Correlator Details	138
B.1	Physical Parameters	138
B.2	Enclosure	140
C	Tables	141
	List of Figures	148
	List of Tables	153

Introduction

The Cosmic Microwave Background (CMB) is the relic radiation left over from a time when the Universe was dense and hot enough to ionize hydrogen. As the universe expanded and cooled, the radiation field remained in thermal equilibrium with the local baryonic matter until it eventually ceased interaction with the hydrogen during a period of decoupling and has since not electromagnetically interacted with matter until reaching the observer. As such, the CMB is a fossil from the time of decoupling.

There are three interesting aspects of this electromagnetic radiation. The first aspect studied in detail was the spectrum, now known to follow a blackbody with $T_{CMB}=2.728$ K to high accuracy (Fixsen et al. 1996). Another aspect is the polarization. The current body of theory suggests that the CMB should be polarized to a few parts in 10^6 , just below the detection threshold of current experiments (Hedman et al. 2001; Keating et al. 2001; Subrahmanyam et al. 2000), but a detection is considered imminent. Finally, the aspect that is the subject of this thesis is the spatial variations of the intensity, or the anisotropy of the CMB.

The spatial anisotropy spectrum of the CMB has proven to be a powerful tool in understanding the universe. For example, Wang et al. (2001) have shown that current CMB anisotropy data can constrain many cosmological parameters. The trend in the ground-based CMB anisotropy field has been toward increased sensitivity at smaller angular scales. The work presented in the following chapters outlines an instrument that follows this trajectory, one that pushes the resolution and sensitivity of CMB instruments further.

1.1 The Angular Spectrum

The CMB radiation is spatially uniform to a few parts in 10^5 . The spatial distribution, therefore, is usually represented in terms of fluctuations from uniformity or deviations from isotropy. The anisotropy spectrum is commonly expanded as a spherical-harmonic decomposition of the temperature fluctuations. The multipole moments of the temperature distribution are,

$$a_{\ell m} = \int Y_{\ell m}^*(\hat{n}) \frac{\Delta T}{T} d\hat{n} \quad (1.1)$$

In the case of a gaussian random field, the multipole moments are fully characterized by the power spectrum:

$$C_\ell \equiv \langle |a_{\ell m}|^2 \rangle, \quad (1.2)$$

where the brackets denote ensemble averaging and it has been assumed that there is no preferred direction in the universe, implying that C_ℓ should have no azimuthal dependence. These C_ℓ 's represent the variances of the $a_{\ell m}$'s and, hence, are related to the variance of the fractional temperature fluctuations at an angular scale indicated by multipole moment ℓ . The particulars of a cosmological model will determine the C_ℓ 's. A typical plot looks like the one in Figure 1.1. Plotted along the y axis is the quantity:

$$\delta T_\ell = \sqrt{\frac{\ell(\ell+1)C_\ell}{2\pi}} T_{CMB} \quad (1.3)$$

For reasons that will be clarified in Section 1.1.1, this quantity (or its square) is the commonly accepted value to plot. It has the feature of equal power per unit logarithmic interval, which is why a logarithmic ℓ -axis is preferred.

The plot can be divided into three regions: low, middle, and high ℓ . Each region is dominated by different physical processes in the early universe, and each region has required slightly different observational approaches.

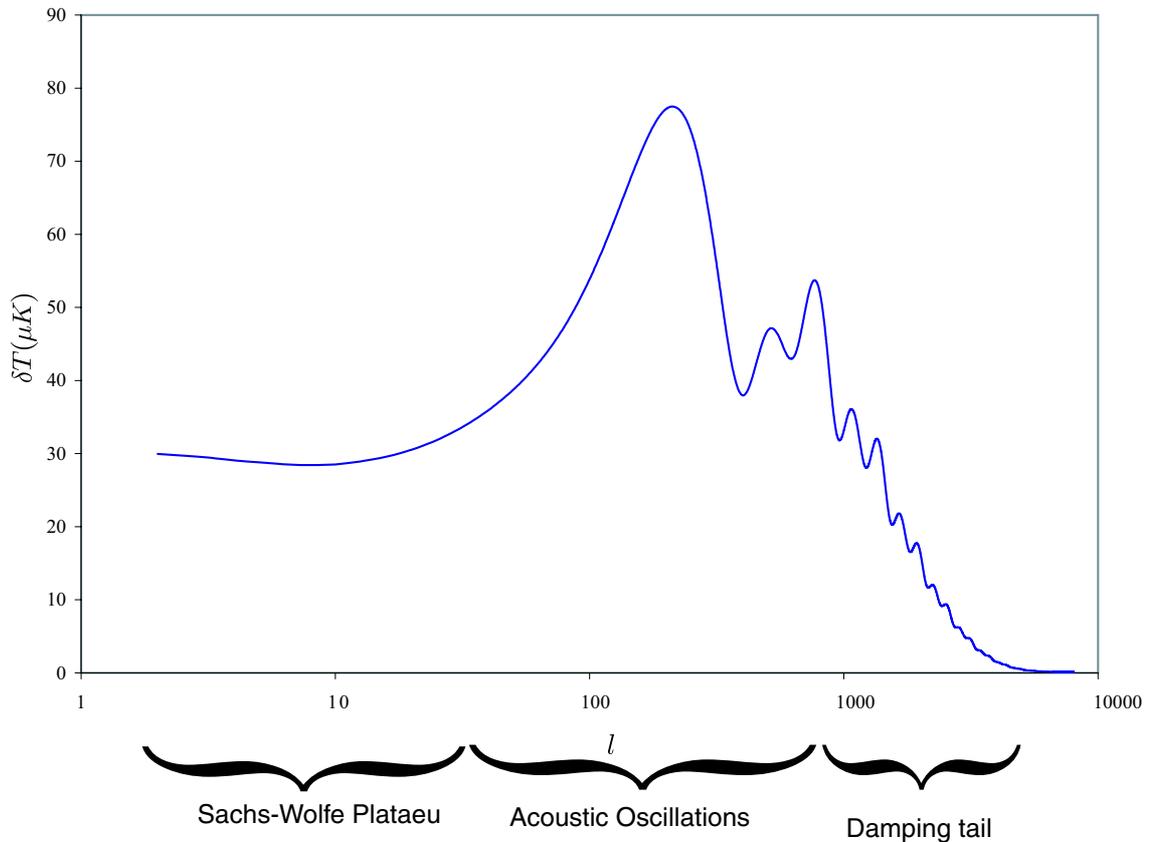


Figure 1.1: Typical theoretical plot of δT vs ℓ .

Analytical vs. Numerical approach

The early universe was a mix of many species: baryons and electrons, photons, neutrinos, and Cold Dark Matter (CDM), each with their own properties and sometimes complicated interactions with other species. The equation governing the dynamics of each species is the Boltzmann equation. The fluctuations about homogeneity are small, and therefore may be analyzed with linear perturbation theory.

There have historically been two different approaches to the understanding of CMB anisotropy. The first approach to analyzing this equation was a massive numerical integration, usually taking many hours (Bond and Efstathiou 1984). The technique was refined (Seljak and Zaldarriaga 1996) into a code called CMBFAST, which runs in a few minutes and has the same accuracy as earlier numerical methods. CMBFAST may be used to predict how cosmological parameters impact the C_ℓ

spectrum, and indeed it is used to show how measured spectra constrain cosmological parameters.

Although the CMBFAST method is accurate to 1-2% and convenient, numerical methods sometimes obscure physical intuition about how the spectrum scales with cosmological parameters. The numerical methods may derive the behavior of the spectrum under different parameters, but an understanding of the physical processes requires analytic equations derived from physical reasoning that faithfully reproduce the numerical results. At the same time that CMBFAST was developed, fully analytic approximations were developed (Hu and Sugiyama 1995) as well as partially analytic approximations (Seljak 1994). The story of the anisotropies outlined in this chapter is distilled from the many resources: Hu and Dodelson (2002); Hu (1995); Hu et al. (1997); Tegmark (1995), all of which draw upon the analytical approach.

The analytical approach relies on some simplifying features of the early universe. Before decoupling and the formation of hydrogen, photons and baryons were tightly coupled through electromagnetism. It is therefore possible to treat them as a single photon-baryon fluid. Neutrinos and CDM interacted only gravitationally with the other species. Furthermore, before decoupling but after matter-radiation equality (see below) CDM is the dominant component of the energy density. It is useful then to analyze the photon-baryon fluid as interacting in a gravitational potential formed by the CDM. After decoupling, the photons traveled unimpeded to the observer, and were subjected only to the large-scale geometry of the universe and to evolving potentials.

1.1.1 Low ℓ : The Sachs-Wolfe Plateau

The low ℓ or large angular scale part of the anisotropy spectrum is sometimes referred to as the Sachs-Wolfe Plateau (Sachs and Wolfe 1967). It includes scales large enough that during the time of decoupling the universe was not old enough to allow regions separated by that physical scale to be in causal contact. Fluctuations on this scale are said to be “outside the horizon” at the time of decoupling. Matter and radiation fluctuations at this scale did not evolve before decoupling and hence, fluctuations measured here must have been formed by some primordial process.

The theory that fits the current data set is known as inflation. According to

this model of the early universe, quantum vacuum fluctuations grew rapidly during a period of inflation. These were the initial fluctuations that evolved into the CMB anisotropy. The spectrum of fluctuations produced by such an event is nearly scale-invariant and can be characterized by a random gaussian field. Furthermore, current measurements indicate that the fluctuations were adiabatic, meaning that different species that contribute to the energy-density remained in constant proportion throughout space.

In general, there are three aspects of the early plasma that imprint themselves in the observed anisotropy: temperature, velocity, and potential. Matter that is outside of the horizon scale does not evolve and hence, at the Sachs-Wolfe plateau there is no bulk velocity and therefore no doppler effect. Furthermore, adiabatic inflation results in gravitational potential fluctuations that are anti-correlated with density (or intrinsic temperature) fluctuations. The result is that photons that come from intrinsically hotter regions must climb out of potential wells, becoming redder in the process. In fact, the potential contributes more to the apparent temperature than the intrinsic temperature, so intrinsically hot regions appear to the observer as apparently colder. The opposite happens for intrinsically colder regions. This process is called the Sachs-Wolfe effect, and leads to a flat anisotropy spectrum, if the C_ℓ 's are plotted according to Equation 1.3, explaining the wide use of this convention¹.

The first detection of anisotropy at any scale came from the COBE satellite which had just enough sensitivity to detect fluctuations statistically on these large scales. The measurements cover the C_ℓ spectrum from $2 < \ell < 20$ and agree well with a flat spectrum. The COBE data presently used to normalize the spectrum; that is, they set the magnitude of the primordial fluctuations, a parameter that is not predicted by inflationary theory (Tegmark and Rees 1998).

1.1.2 Mid ℓ : Acoustic Oscillations

The mid- ℓ range covers scales that are inside the horizon at or before decoupling, and have enough time to evolve. The upper end of this range is limited by how tightly the

¹The two alternatives are $\frac{\ell(\ell+1)C_\ell}{2\pi}$ and $\frac{\ell(2\ell+1)C_\ell}{4\pi}$. While the second version is closer to the power-per-logarithmic interval, the first version is a natural result of integration of the Sachs-Wolfe effect. See Dodelson (2002) for details.

photons are coupled to the baryons in the plasma. This is the domain of the acoustic peaks. The dominant feature in the spectrum is a characteristic peak at $\ell \sim 200$.

In the most naive picture, the general relativistic equations can be simplified to acoustic harmonic oscillator equations for each Fourier mode k . The restoring force is formed by gravitational potential during the rarefaction phase and by photon pressure during the compression phase. As the universe gets older, the horizon gets larger, and progressively larger scales begin to evolve. The largest peak in the spectrum corresponds to acoustic scales that have had just enough time to fully compress. The second peak is located at scales that have had enough time to compress and then fully rarefy. Both peaks are positive definite on the C_ℓ spectrum because only the power is plotted. This alternating series of full compression and full rarefaction scales leads to the familiar set of acoustic peaks on the C_ℓ plot.

1.1.3 High ℓ : The Damping Tail

The work described here focuses on measuring the high- ℓ part of the anisotropy spectrum. The high ℓ range is characterized by an exponential damping envelope of the acoustic peaks. There are two competing effects leading to this characteristic shape. Radiation driving tends to amplify the peaks but Silk damping overcomes the amplification, resulting in an overall damping of fluctuations. Geometric effects also play a role in the overall damping envelope.

Radiation Driving

The energy content of the very early universe is dominated by radiation. Because radiation is redshifted as the universe expands, the energy density of radiation falls as a^{-4} , where a is the scale factor in the Friedman-Robertson-Walker metric that sets physical lengths as the universe expands (Misner et al. 1973). Non-relativistic matter density, on the other hand, falls as a^{-3} . At some time known as t_{eq} , matter becomes the dominant component of the universe. In most models, t_{eq} happens before recombination, so there are scales that cross the horizon before t_{eq} and scales that cross after. Those scales that cross before are physically smaller and are influenced by radiation driving.

In the radiation era, the gravitational potentials decay with time. The decay

happens because radiation dominates CDM before t_{eq} . The photon-baryon fluid, therefore, creates the potentials under self-gravity, rather than the CDM. Radiation pressure acts to keep overdense regions from forming. As the universe evolves, when a certain scale crosses the horizon, material begins to compress due to self-gravity. As material builds in a region, radiation pressure acts to halt the growth of potential, so as the universe expands, the potential decays. The critical timing of the potential decay acts as a time-dependent driving force, which amplifies the rarefactions. In other words, during the rarefaction phase, the compressed fluid does not have to overcome as much gravity as was present during the compression phase, so there is less restoring force holding back the fluid. Furthermore, the decay is timed to act as a near resonant driving force. The effect of radiation driving is roughly doubled by a complementary dilation effect due to the evolving geometry of the decaying potential wells. Potential wells that initially stretched space will evolve to a more compressed space, blueshifting the photons inside.

Silk Damping

Small scale fluctuations in the photon-baryon fluid were suppressed by radiative diffusion damping (Silk 1968). The early plasma was not a perfectly coupled photon-baryon fluid. In actuality, photons do not couple directly to baryons, but rather to electrons through Thompson scattering. The electrons in turn couple to the baryons through electromagnetic fields. The imperfect coupling is only important on scales that are approximately the size of the Thompson mean-free path²,

$$\lambda_T = \frac{1}{an_e\sigma_T}, \quad (1.4)$$

where n_e is the electron density and σ_T is Thompson cross section,

²The development in this section ignores many important features. Namely, the factor of a in the denominator is due to the definition of λ 's as comoving distances. In general, the analytic approach to understanding anisotropies requires decomposing all perturbations into Fourier modes characterized by the comoving wavenumber, $k = 2\pi/\lambda$. To translate k into the observed anisotropies in ℓ -space, a projection is required. This mixes many k 's into a single ℓ , but for high ℓ , the relation $k * r_\theta \sim \ell$ is sufficient, where r_θ is the angular diameter distance, and for flat geometries it is just the distance to the surface of last scattering. There is a factor of a in the denominator of Equation 1.4 to convert the physical mean free path to a comoving one.

$$\sigma_T = \frac{8\pi}{3} \left(\frac{e^2}{m_e c^2} \right)^2. \quad (1.5)$$

The diffusion of photons from hot regions into colder regions and vice versa leads the observer to believe that hot photons that originated in hot regions actually are coming from cold regions. The mixing of hot and cold photons damps the inhomogeneities by an exponential factor of order³ $\exp(-\ell^2/\ell_d^2)$. The characteristic damping scale is ℓ_d , set by a number of factors, but given loosely as the geometric mean of the mean free path and the horizon scale. To heuristically justify this statement, consider that the distance a random walking photon can travel is $\lambda_D \sim \sqrt{N}\lambda_T$. The number of collisions, N, is given by the age of the universe at decoupling, η_* divided by the average time per interaction, $1/\lambda_T$ in units where the speed of light is set to $c=1$. This gives: $\lambda_D \sim \sqrt{\eta_*\lambda_T}$. In these units, η_* is also the horizon scale. The damping scale expressed as a wavevector is:

$$k_D \sim \frac{2\pi}{\sqrt{\eta_*\lambda_T}}. \quad (1.6)$$

It is useful to compare ℓ_D to ℓ_{peak} , the position of the first acoustic peak, by constructing the dimensionless quantity

$$\frac{\ell_D}{\ell_{peak}} \sim k_D \frac{\eta_*}{\sqrt{3}} \sim \frac{2\pi}{\sqrt{3}} \sqrt{\frac{\eta_*}{\lambda_T}}, \quad (1.7)$$

where $\eta_*/\sqrt{3}$ is the *sound* horizon scale at decoupling, which differs from the particle horizon only in that the speed of sound is $\sqrt{3}$ slower than the speed of light. It signifies the distance that sound could travel before decoupling and hence indicates the scale of the first acoustic peak.

³The derivation of the damping envelope, $\exp(-\ell^2/\ell_d^2)$, requires a second order correction to the tight coupling approximation used to generate the potentials. This is beyond the scope of this introduction, but for an example derivation see Peebles (1980). A simple understanding of the form of the damping envelope comes from considering the effect of diffusion as a convolution of the last scattering surface with a gaussian of width $\sim \lambda_D$. In the power spectrum domain, this convolution appears as multiplication by a gaussian of width $\sim 1/\lambda_D \sim k_D$. Projection onto the last scattering surface turns k_D into ℓ_D . It is also important to note that this form of the damping envelope only applies at scales roughly equal to and larger than the damping scale. There is no damping, for instance, on scales that are outside of the horizon at decoupling.

Equation 1.7 should not be too taken literally because it does not take into account the changing values of λ_T through the process of decoupling. A detailed analysis (Hu and Dodelson 2002), arrives at $k_D \frac{\eta_*}{\sqrt{3}} \sim 10$. Since the acoustic peak is at roughly $\ell_{peak} \sim 100$, this implies that the damping scale is $\ell_D \sim 1000$, and that acoustic peaks beyond the third are heavily damped.

Thickness of last scattering.

The last effect to consider arises because recombination did not occur instantaneously. As the photons decouple, the relative ionization of the electron-baryon fluid quickly diminishes, decreasing n_e and leaving fewer electrons to scatter. This is important because photons travel freely to the observer only if there are very few free electrons. Therefore, for an observer, the “surface” of last scattering is actually a 3-D volume.

An observer looking at a particular point on the 2-D sky is actually seeing the last scattering surface integrated along the line of sight, through a thickness corresponding to the time it took for the universe to go from opaque to transparent. This thickness is referred to as the thickness of last scattering and has the canonical value of $\Delta z \sim 100$, measured in redshift (Jones and Wyse 1985). Looking through the last scattering surface averages hot and cold regions that are smaller than the thickness. The effect is similar to convolving a beam over a temperature map. A larger Δz leads to more suppression of the peaks (White 2001).

Angular Diameter Distance

Finally, in the case of a non-flat universe, the apparent damping scale in ℓ -space can be modified by the angular-diameter distance relation. Curved geometries will magnify or shrink the apparent size of features in the distant universe. This effect shifts the damping scale to higher ℓ in the case of an open universe and to lower ℓ in the case of closed. Measuring ℓ_D is another test for the total energy density, Ω_{tot} . There is currently, however, strong evidence from the acoustic peaks that the universe is flat and Ω_{tot} equals the critical density (Wang et al. 2001).

Cosmological parameters

Aside from the angular diameter distance relation, the damping scale is also dependent on the horizon scale at decoupling, η_* , and the Thompson mean free path, λ_T . If the radiation density at decoupling is fixed by the present value, then the only term that influences η_* is the matter content, Ω_m , as curvature and the cosmological constant do not contribute much at early times. The dependence on the matter content is roughly $\eta_* \propto (\Omega_m h^2)^{-1/2}$ (Kamionkowski et al. 1994). With more matter in the universe, the expansion is slower and the horizon at decoupling is smaller.

The dependence on the Thompson mean free path implies that there is also a simple relationship between ℓ_D and the electron density n_e . The early universe was neutral, so every electron was matched to a proton, $n_e = n_b$, and the damping scale is proportional to $(\Omega_b h^2)^{1/2}$. More baryons means more scattering sites and photons are not allowed to wander as far away from their origin. The process is not quite that simple, and competing effects during recombination tend to reduce the dependence on the baryon density to $(\Omega_b h^2)^{1/4}$ (Hu and Sugiyama 1996).

A full numerical solution to the Boltzmann equation is required for an accurate understanding of how the damping scale depends on $\Omega_m h^2$ and $\Omega_b h^2$. The effects of damping are charted by dividing the numerically generated power spectra by spectra from a fictitious universe with no damping. A fitting formula that describes the numerical behavior of the damping scale accurately to 1% can be constructed from a two-power-law formula in both $\Omega_m h^2$ and $\Omega_b h^2$ (Hu and White 1997).

1.2 Previous Measurements

Experimenters in the last decade have made great progress in filling out the C_ℓ spectrum. A cosmological parameter extraction from roughly two dozen experiments is presented in Wang et al. (2001). The experiments are shown to be consistent with one another and with standard cosmological predictions with few surprises, except that fundamental theories conceived over 30 years ago seem to be correct.

The current body of data shows the main peak at very high confidence. Additionally, data through the mid- ℓ range are consistent with second and third doppler peaks, but do not detect them with high confidence. The definitive measurements at

mid- ℓ are currently underway with the MAP satellite.

The Cosmic Background Imager team (Padin et al. 2001a) has made an initial measurement of the damping tail at 30 GHz. Two bandpower measurements at $\ell = 603$ and 1190 have values of $59_{-6.3}^{+7.7} \mu\text{K}$ and $29.7_{-4.2}^{+4.8} \mu\text{K}$ respectively. An analysis of these points presented in White (2001) has shown that the measurements are consistent with standard cosmology. This study also set the damping scale at $\ell_D \sim 1000$.

The work presented in this thesis describes an instrument, the Microwave Interferometer (MINT), designed to produce a measurement complementary to the CBI results with similar precision. Both instruments are close-packed interferometers located at nearly the same site. The major difference between the two instruments, outside of cost, personnel, complexity and number of elements, is the higher operational frequency of MINT, about 150 GHz, which places different systematic limits on the measurements. The high frequency allows the construction of a smaller and more manageable instrument to probe the same angular scales. It also leads to more immunity from sources of astronomical contamination.

1.3 Foregrounds

The subject of foregrounds is nearly as varied as the study of the primary anisotropies themselves. Outlined here are only two interesting foregrounds. Other foregrounds, such as galactic dust or free-free emission are thought to be small enough to neglect at 150 GHz and $\ell > 1000$. Tegmark et al. (2000) is an exhaustive survey of foreground issues and the information in this section is drawn from this paper.

1.3.1 Point Sources

Point sources are generally extragalactic in origin and come from two populations, radio sources including blazars and far-IR sources including early galaxies. Both types have generically power law spectra. The contribution to the C_ℓ 's at 150 GHz and $\ell = 1000$ is roughly ten times smaller than at 30 GHz, where the contamination is of order $10 \mu\text{K}$, with the detectable sources removed. Thus, analysis of any instrument operating at low frequencies requires subtraction of known sources and

statistical estimation and subtraction of undetected sources based on models. The measurements quoted for CBI include a 1.6% and 8.6% correction in the lower and upper ℓ measurements respectively.

1.3.2 Sunyaev-Zel'dovich Effect

The thermal Sunyaev-Zel'dovich Effect (SZE) is caused by hot electrons in galaxy clusters inverse-Compton scattering CMB photons. The process shifts the photons to higher energy causing a decrement in signal at frequencies below the spectral peak of the CMB (~ 220 GHz) and an increase in signal at frequencies above. The SZE is an interesting cosmological probe in itself. For a summary of the current results, see Carlstrom et al. (2001). As a foreground, the SZE becomes important at the characteristic scales of clusters ($1'$), and the temperature fluctuations due to the SZE become equal to the primary anisotropies well into the damping tail at $\ell \sim 3000$. In the canonical view, the SZE will not be a major contaminant at the damping scale, and the effect can be suppressed by observing closer to the SZE “null” at 220 GHz.

The contribution to the anisotropies can be modified, however, if the diffuse SZE is included. The diffuse SZE arises from intra-cluster gas or gas that is filamentary. These effects are poorly constrained and may amount to large contributions.

Even though SZE dominates the power in the anisotropy spectrum at the highest ℓ , the contributions may be removed in a map by subtracting known clusters. Additionally, a multi-frequency map will be able to distinguish between the SZE and the primary anisotropies.

The motivation for MINT was to be complementary to both MAP and CBI. MINT has smaller angular resolution than MAP and slightly higher frequency coverage. CBI also probes higher angular scales than MAP. Although CBI is much larger and has many more elements than MINT, they both cover similar ℓ -range. Given the vastly different observing frequency, however, the response to foregrounds will be very different. MINT is not expected to be at all sensitive to point source contamination.

Interferometry and Observing Strategy

This chapter develops the fundamental measurable quantity for the interferometer, the visibility. The goal is to estimate the sensitivity of MINT. While the theoretical groundwork for interferometry has been in place for decades, quantifying how interferometers perform for diffuse sources in close-packed, detector noise limited configurations has matured only recently.

With the recent construction of a number of interferometers designed specifically to study CMB came many papers aimed at analyzing the sensitivity, particularly from the CBI, DASI and VSA groups. These papers differ from earlier treatments, (Partridge et al. 1987; Fomalont et al. 1984; Subrahmanyan et al. 1993), in that the emphasis is on analyzing the visibilities directly produced by interferometers instead of a map derived from the visibilities. A Bayesian approach to analyzing the visibility is formalized in Hobson et al. (1995). An explicit connection to the commonly accepted ℓ -space representation is presented in White et al. (1999) (hereafter referred to as WCDH), as well as a summary of visibility-based analysis of interferometer data.

Central to the development here is work done by Joseph Fowler on heterogenous arrays and the signal to noise ratio (SNR) of the MINT configuration. The chapter is in some ways a reworking of memoranda that examine the sensitivity issues for the specific case of MINT.

2.1 Interferometer Response

2.1.1 Flat Sky Approximation

For this development, the flat-sky approximation is adopted, in which the angular sky is projected onto a 2-D plane. Vectors in angular space, $\vec{s} = (\theta, \phi)$ measured in radians, are taken to be linear vectors \vec{x} on the flat sky also measured in radians. This approximation is commonly adopted in interferometry because the main telescope beam covers a small fraction of the sky.

This approximation has many simplifying effects. Instead of decomposing the sky temperature, $T_s(\vec{x})$, into spherical harmonic amplitudes $a_{\ell m}$, the power spectrum of the sky, $S(\vec{u})$, is calculated via the Fourier Transform. The variables \vec{x} and \vec{u} are fourier conjugates. Furthermore, assuming that the sky has no preferred directions means that S is only a function of $|u|$. The approximate relationship between the $C_\ell = \langle |a_{\ell m}|^2 \rangle$ and the power spectrum of the sky is (from WCDH):

$$u^2 S(u) \simeq \frac{\ell(\ell+1)}{4\pi^2} C_\ell \Big|_{\ell=2\pi u} \quad \text{for } u > 10 \quad (2.1)$$

This relation can be understood intuitively to within factors of 2π . The C_ℓ 's go as $a_{\ell m}^2$, and hence are proportional to the total power in each m mode. There are $2\ell+1$ m modes for each ℓ , so the total power in a given ℓ mode is $(2\ell+1)C_\ell$. The extra ℓ on the right hand side of Equation 2.1 is inserted to make C_ℓ plots measure power per unit logarithmic interval¹. One factor of u on the left hand side is used to offset this factor of ℓ . The other factor of u is from integrating over one dimension, resulting in the one-dimensional C_ℓ plot.

2.1.2 Primary Beam

Interferometers share many components with traditional total power receivers. In particular, each element of an interferometer is composed of the front end of a coherent

¹Some feel that it is slightly unfortunate that C_ℓ plots have the factor $2(\ell+1)$ instead of the more natural $(2\ell+1)$. In some ways, this is just a matter of historical legacy and at low ℓ makes little difference. The specific form was adopted because under a scale invariant spectrum, it produces a flat C_ℓ plot at low ℓ . Refer to Section 1.1.1 for more details.

receiver, including the antenna and low noise amplifier or mixer. The power response of the antenna, $A(\vec{x})$, is known as the primary beam function. In general, the function can be of arbitrary shape, but for the sake of specificity, we will approximate the beam function with a cylindrically symmetric gaussian:

$$A(\vec{x}) = e^{-\frac{|\vec{x}|^2}{2\sigma_p^2}}. \quad (2.2)$$

MINT has two different sized antennas. The gaussian width, σ_p , may be determined by matching the solid angle, $\Omega = \int_{beam} d\vec{x} A(\vec{x})$, of the measured beam with that from the gaussian. We find $\sigma_p = .00355$ for the small antenna and $\sigma_p = .00247$ for the large antenna (in radians). See Section 3.2.1 for a more detailed discussion of the beams.

For a source of infinitesimal size dx , temperature $T_s(\vec{x})$, and located at \vec{x} , the antenna temperature, dT_{ant} , is given by:

$$dT_{ant} = \eta_{R-J} T_s(\vec{x}) A(\vec{x}) \frac{dx}{\Omega}, \quad (2.3)$$

where

$$\eta_{R-J} = \frac{x^2 e^x}{(e^x - 1)^2} \quad \text{and} \quad x = \frac{h\nu}{kT}, \quad (2.4)$$

and

$$\Omega = 2\pi\sigma_p^2 = \frac{\lambda^2}{A_e}. \quad (2.5)$$

The wavelength is λ and A_e is the effective area of the antenna. The equation $\Omega = 2\pi\sigma_p^2$ is only correct for a gaussian beam (Kraus 1966). The factor η_{R-J} is the Raleigh-Jeans correction factor. It converts from blackbody temperatures on the sky, known as thermodynamic temperature to temperatures measured in the receiver, or antenna temperature (Partridge 1995). The correction is important for shorter wavelengths. For MINT, $x = 2.55$ and $\eta_{R-J} = 0.60$. Note that the temperature of the source must be known in order to make this correction. For CMB fluctuations, $T=2.728$ K to good accuracy (Fixsen et al. 1996).

Finally, the voltage response of each receiver of an interferometer is proportional to the electric field, not the intensity of the radiation, so the appropriate expression is the square root of dT_{thm} :

$$\sqrt{dT_{ant}} = \sqrt{\frac{\eta_{R-J} T_s(\vec{x}) A_e}{\lambda^2}} g(\vec{x}) dx, \quad (2.6)$$

with

$$g(\vec{x}) = \sqrt{A(\vec{x})} = e^{-\frac{|\vec{x}|^2}{4\sigma_p^2}}, \quad (2.7)$$

where $g(\vec{x})$ is the point response of a receiver to the electric field (also a gaussian), normalized to $g(0) = 1$.

2.1.3 Visibility

The fundamental unit of an interferometer is an antenna pair, referred to as a baseline. MINT has four antennas and thus, six baselines. The two antennas are separated by a dimensionless vector \vec{u} , the magnitude of which is measured in wavelengths. In general, the antennas are allowed to point in any direction on the sky relative to \vec{u} , but for MINT, the pointing direction, \vec{s}_0 , is fixed to be perpendicular to \vec{u} .

For the current analysis, there will be some simplifying assumptions made. First, it is assumed that the source on the sky is monochromatic. This is in general not a good supposition, but since, in the case of MINT, the ratio of the bandwidth to the center frequency is small, the assumption is reasonable. The outputs of the receivers are fed to the correlator, which for the moment, can be thought of as simply a multiplier followed by an accumulator. It is also assumed that the accumulation time is long compared to the period of the RF radiation. Additionally, the main result will be derived in thermodynamic temperature units instead of the more conventional flux density units.

Figure 2.1 is a diagram of a fundamental antenna pair. Both antennas are pointed perpendicular to \vec{u} . Consider a contribution to the signal from a direction \vec{x} , a vector projected onto the flat sky. The angle θ is the component of \vec{x} in the plane containing \vec{u} and \vec{s}_0 , and ϕ is the component out of that plane. The geometry indicates that a signal emanating from \vec{x} must travel a distance $|u| \sin \theta$ longer to get to the left antenna. The extra time is known as the geometric delay,

$$\tau_g = \frac{|u| \sin \theta}{\nu}, \quad (2.8)$$

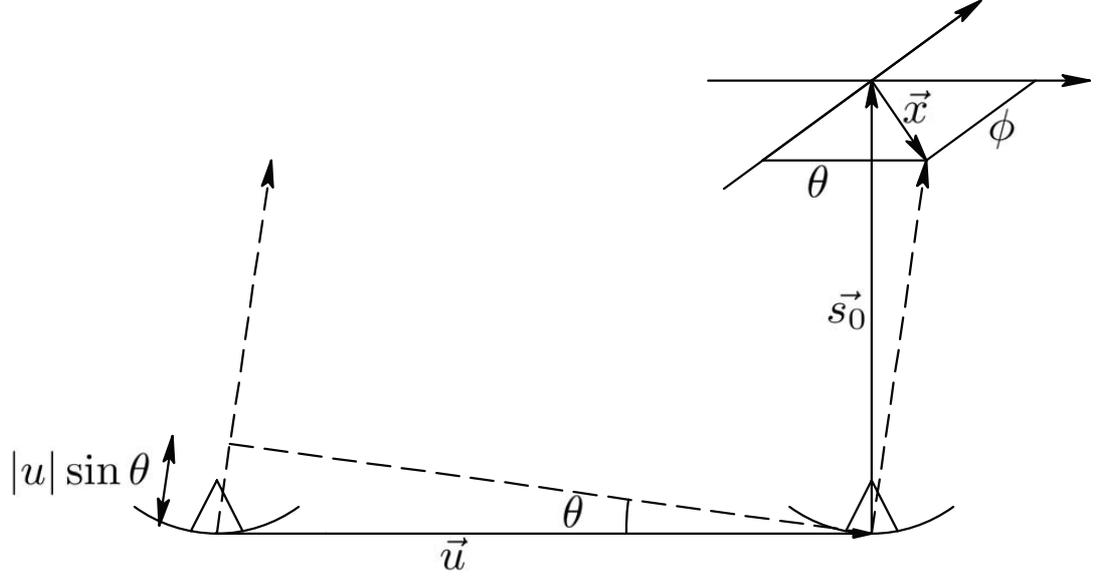


Figure 2.1: Simplified Baseline. The dashed arrows point in the same direction. The vector \vec{x} is the projection of the dashed arrows onto the 2-D sky plane. The center of each antenna beam points along \vec{s}_0

where ν is the frequency. We divide by ν instead of c because \vec{u} is measured in wavelengths. Also, under the small angle approximation,

$$\tau_g = \frac{\vec{u} \cdot \vec{x}}{\nu}. \quad (2.9)$$

Given a temperature distribution on the sky, $T_s(\vec{x})$, the time-varying infinitesimal electric field entering an antenna from a direction \vec{x} varies as:

$$dE(\vec{x}, t) \propto \sqrt{\frac{\eta_{R-J} T_s(\vec{x}) A_e}{\lambda^2}} g(\vec{x}) \sin(2\pi\nu t) dx. \quad (2.10)$$

The signal into the other antenna is delayed by τ_g . In the correlator, the signals from two antennas are multiplied. The output of the multiplier is:

$$dF(\vec{x}, t) \propto \frac{\eta_{R-J} T_s(\vec{x})}{\lambda^2} \sqrt{A_{e_1} A_{e_2}} g_1(\vec{x}) g_2(\vec{x}) \sin(2\pi\nu t) \sin(2\pi\nu(t - \tau_g)) dx \quad (2.11)$$

$$\propto \frac{\eta_{R-J} T_s(\vec{x})}{\lambda^2} \sqrt{A_{e_1} A_{e_2}} g_1(\vec{x}) g_2(\vec{x}) [\cos(2\pi\nu\tau_g) - \cos(4\pi\nu t) \cos(2\pi\nu\tau_g) - \sin(4\pi\nu t) \sin(2\pi\nu\tau_g)] dx. \quad (2.12)$$

The last two terms vary fast compared to the accumulation time and are centered on zero. The accumulator in the correlator will average them down so they may be ignored. The output of the accumulator is then:

$$dF(\vec{x}) = \frac{\eta_{R-J} T_s(\vec{x})}{\lambda^2} \sqrt{A_{e_1} A_{e_2}} g_1(\vec{x}) g_2(\vec{x}) \cos(2\pi\nu\tau_g) dx. \quad (2.13)$$

In a complex correlator, an additional multiplier is employed with one of the inputs shifted by $\pi/2$. This essentially shifts the argument of Equation 2.13 by $\pi/2$. The full output of a complex correlator is:

$$dF_{\cos}(\vec{x}) = \frac{\eta_{R-J} T_s(\vec{x})}{\lambda^2} \sqrt{A_{e_1} A_{e_2}} g_1(\vec{x}) g_2(\vec{x}) \cos(2\pi\nu\tau_g) dx \quad (2.14)$$

$$dF_{\sin}(\vec{x}) = \frac{\eta_{R-J} T_s(\vec{x})}{\lambda^2} \sqrt{A_{e_1} A_{e_2}} g_1(\vec{x}) g_2(\vec{x}) \sin(2\pi\nu\tau_g) dx. \quad (2.15)$$

To obtain the response from the entire sky, both equations are integrated over dx . The two equations together can then be recognized as the complex Fourier transform of $T_s(\vec{x})g_1(\vec{x})g_2(\vec{x})$. With Equation 2.9 the full output of a complex correlator can be written as:

$$V(\vec{u}) = \frac{\eta_{R-J} \sqrt{A_{e_1} A_{e_2}}}{\lambda^2} \int dx T_s(\vec{x}) g_1(\vec{x}) g_2(\vec{x}) e^{2\pi i \vec{u} \cdot \vec{x}}. \quad (2.16)$$

$V(\vec{u})$ is called the visibility function. The Fourier transform relationship leads naturally to the use of \vec{u} and \vec{x} as conjugate variables. The action of a baseline is to measure the visibility function around the point in the plane that contains \vec{u} , known as the u - v plane². The exact weighting of how the interferometer responds in this u - v plane is specified through the window function, described in the next section.

²The u and v of u - v plane refers to the traditional labels of the orthogonal axes

2.2 Sensitivity

2.2.1 Signal

Given a model of the temperature distribution of the sky, it is possible to estimate the signal in a given baseline. The goal is to find out how the particular choices for the physical sizes and locations of the antennas affect our ability to measure the power spectrum, $S(u)$.

The Fourier transform of a product is the convolution of the Fourier transform of the factors:

$$\mathcal{F}[T_s(\vec{x})g_1(\vec{x})g_2(\vec{x})] = \mathcal{F}[g_1(\vec{x})g_2(\vec{x})] \star \mathcal{F}[T_s(\vec{x})]. \quad (2.17)$$

Given that a single baseline only samples the visibility at one baseline \vec{u}_0 , Equations 2.16 and 2.17 must be modified to look like:

$$V(\vec{u}_0) \propto \widetilde{g_1 g_2}(\vec{u}_0 - \vec{u}) \cdot \widetilde{T_s}(\vec{u}), \quad (2.18)$$

where the tilde is shorthand notation for the Fourier transform. The dot product here is the evaluation of the convolution for a single baseline \vec{u}_0 . To evaluate $\widetilde{g_1 g_2}(\vec{u}_0 - \vec{u})$, $g(\vec{x})$ is approximated as a gaussian as described in Equation 2.2. The product of two gaussians is another gaussian and the Fourier transform of a gaussian is also gaussian:

$$g_1(\vec{x})g_2(\vec{x}) = \exp\left(\frac{-|x|^2}{4\sigma_1^2}\right) \exp\left(\frac{-|x|^2}{4\sigma_2^2}\right) = \exp\left(\frac{-|x|^2}{2\sigma_3^2}\right) \quad (2.19)$$

where $\sigma_3^2 = 2 \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}$

$$\widetilde{g_1 g_2}(\vec{u}) = \mathcal{F}\left[e^{\frac{-|x|^2}{2\sigma_3^2}}\right] = 2\pi\sigma_3^2 e^{-2\pi^2\sigma_3^2|u|^2}. \quad (2.20)$$

Evaluating Equation 2.20 at $\vec{u} - \vec{u}_0$ and using Equations eq:visibilty and 2.18 yields the following expression for the visibility as a function of baseline \vec{u}_0 :

$$V(\vec{u}_0) = \frac{\eta_{R-J} \sqrt{A_{e_1} A_{e_2}}}{\lambda^2} 2\pi \sigma_3^2 \int d\vec{u} e^{-2\pi^2 \sigma_3^2 |\vec{u} - \vec{u}_0|^2} \widetilde{T}_s(\vec{u}). \quad (2.21)$$

For CMB experiments, the fundamental theoretical construct is the correlation matrix. Consider visibilities collected while the primary beams are pointed in a direction fixed on the celestial sphere. The associated correlation matrix that contains the visibilities measured at different baselines \vec{u}_i and \vec{u}_j is $C_{ij}^V = \langle V^*(\vec{u}_i) V(\vec{u}_j) \rangle$, where the brackets refer to ensemble averaging. Given that the theoretical power spectrum, $S(|u|) = |\widetilde{T}_s(|u|)|^2 / T_{CMB}^2$, is diagonal in \vec{u} , The double integral over \vec{u} collapses to a single integral over u :

$$C_{ij}^V = \eta_{R-J}^2 \frac{\sqrt{A_{e_{1i}} A_{e_{2i}} A_{e_{1j}} A_{e_{2j}}}}{\lambda^4} 4\pi^2 \sigma_{3i}^2 \sigma_{3j}^2 \times \int d\vec{u} e^{-2\pi^2 \sigma_{3i}^2 |\vec{u} - \vec{u}_{0i}|^2} e^{-2\pi^2 \sigma_{3j}^2 |\vec{u} - \vec{u}_{0j}|^2} S(u) T_{CMB}^2. \quad (2.22)$$

Here, $A_{e_{1j}}$ refers to the effective area of the first antenna of the pair comprising the baseline \vec{u}_i , while $A_{e_{2j}}$ refers to the second and σ_{3i}^2 is defined in Equation 2.19.

The MINT baselines are somewhat correlated, as will be shown graphically in Figures 2.2 and 2.3. If, however, this small correlation is ignored and the baselines are analyzed separately, this complication can be ignored. To estimate the sensitivity, only the diagonal part of C_{ij}^V is important. For baseline \vec{u}_0 , the correlation matrix reduces to:

$$C^V = \eta_{R-J}^2 \frac{A_{e_1} A_{e_2}}{\lambda^4} 4\pi^2 \sigma_3^4 \int d\vec{u} e^{-4\pi^2 \sigma_3^2 |\vec{u} - \vec{u}_0|^2} S(u) T_{CMB}^2. \quad (2.23)$$

Since $T_s(\vec{x})$ is a real function, $V(\vec{u}_0) = V^*(-\vec{u}_0)$. Therefore, every measurement of the visibility can be conjugate-reflected through the origin of the u - v plane. A 2-D window function, W_{2d} may be defined for each baseline:

$$W_{2d}(\vec{u}) = \eta_{R-J}^2 \frac{A_{e_1} A_{e_2}}{\lambda^4} 4\pi^2 \sigma_3^4 e^{-4\pi^2 \sigma_3^2 |\vec{u} - \vec{u}_0|^2} \quad (2.24)$$

When multiplied by $S(u) T_{CMB}^2$ and integrated over $\vec{u} = u d u d \theta$, W_{2d} gives C^V . There is a separate window function for each of the six MINT baselines, but not all of the

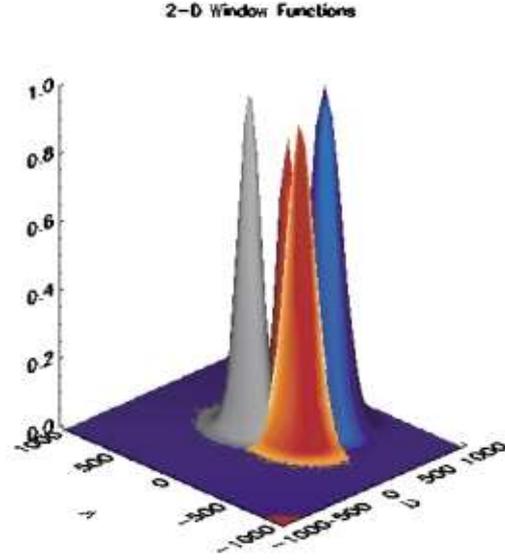


Figure 2.2: 2D MINT window functions. The conjugates of the baselines are not shown. The factor $\eta_{R-J}^2 = 0.36$ is omitted for clarity.

window functions are unique for MINT. Instead of σ_1 and σ_2 , the width of g_1 and g_2 , the window function can be written in terms of σ_{1p} and σ_{2p} , the measured widths of the two primary gaussian beams, $A_1 = g_1^2$:

$$W_{2d}(\vec{u}) = 4\eta_{R-J}^2 \left[\frac{\sigma_{1p}\sigma_{2p}}{\sigma_{1p}^2 + \sigma_{2p}^2} \right]^2 e^{-\frac{|\vec{u}-\vec{u}_0|^2}{2\sigma_{wf}^2}}$$

$$\text{where } \sigma_{wf}^2 = \frac{\sigma_{1p}^2 + \sigma_{2p}^2}{16\pi^2\sigma_{1p}^2\sigma_{2p}^2}. \quad (2.25)$$

Figure 2.2 is a plot of the four unique MINT window functions and Figure 2.3 is a $1-\sigma$ contour plot of the same. As can be seen from the contour plot, the MINT window functions are slightly correlated in that they sample overlapping areas of the $u-v$ plane, which leads to small off-diagonal elements in C_{ij}^V . Also note that four receivers produce $n(n-1)/2 = 6$ baselines, but MINT was intentionally designed with redundancies. There are four equal-length baselines containing both a small and a large antenna. Furthermore, these four baselines are composed of two pairs of baselines that are oriented in the same direction and therefore identical in $u-v$ space. For more discussion on these correlations and redundancies, see Section 2.2.5.

Since $\widetilde{T}_s(\vec{u})$ is only dependent on $|u|$, it is sometimes easier to work in polar

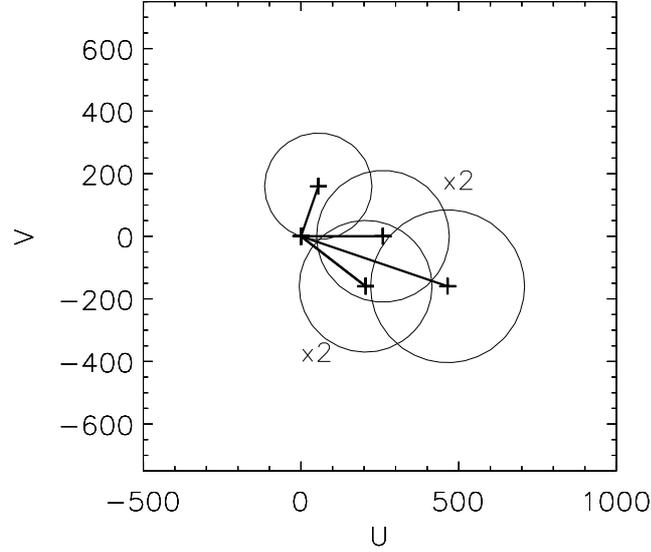


Figure 2.3: 1- σ contours of the 2D MINT window functions. The conjugates of the baselines are still omitted. The window functions labelled “x2” are doubly redundant. The “+” signs mark $[u, v] = [0, 0]$ and the center of the window functions. The heavy lines represent the \vec{u} , or baseline vectors, as in Figure 2.1.

coordinates. The one-dimensional window function³ is defined as:

$$\begin{aligned}
 W(u) &= \int d\theta W(\vec{u}) = 4\eta_{R-J}^2 \frac{A_{e_1} A_{e_2}}{\lambda^4} \int_0^{2\pi} d\theta \widetilde{g_1 g_2}(\vec{u}_0 - \vec{u}) \\
 &= 4\eta_{R-J}^2 \left[\frac{\sigma_{1p} \sigma_{2p}}{\sigma_{1p}^2 + \sigma_{2p}^2} \right]^2 \int_0^{2\pi} d\theta e^{-\frac{|u|^2}{2\sigma_{wf}^2}}. \quad (2.26)
 \end{aligned}$$

A plot of the 1D window function appears in fig 2.4. Equation 2.23 may be written in terms of the 1D window function:

$$C^V = \int_0^\infty u \, du \, W(u) S(u) T_{CMB}^2. \quad (2.27)$$

³For gaussian beams, it is possible to evaluate this expression analytically in terms of the Bessel function, but for the current purposes, a numerical integration is sufficient. When inserting measured or simulated beams into $\widetilde{g_1 g_2}$, it is necessary to evaluate this expression numerically.

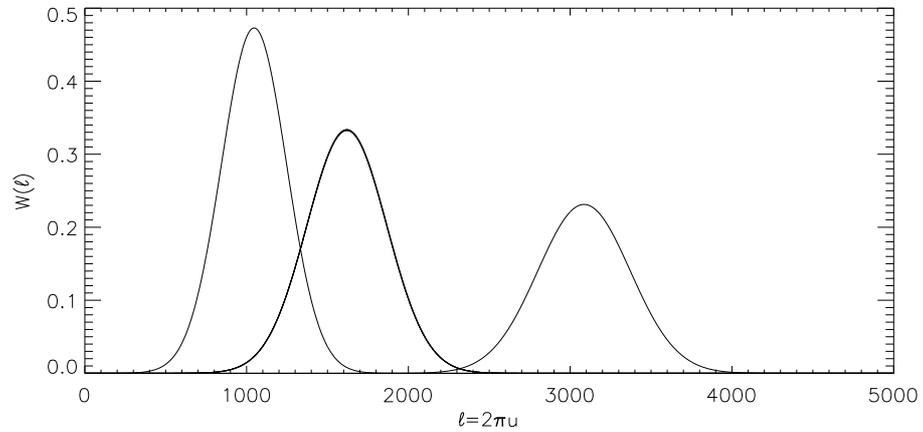


Figure 2.4: 1D MINT window functions assuming gaussian beams. The middle window function is 4-fold redundant. The window functions are unitless. As in Figure 2.2, the factor $\eta_{R-J}^2 = 0.36$ is omitted.

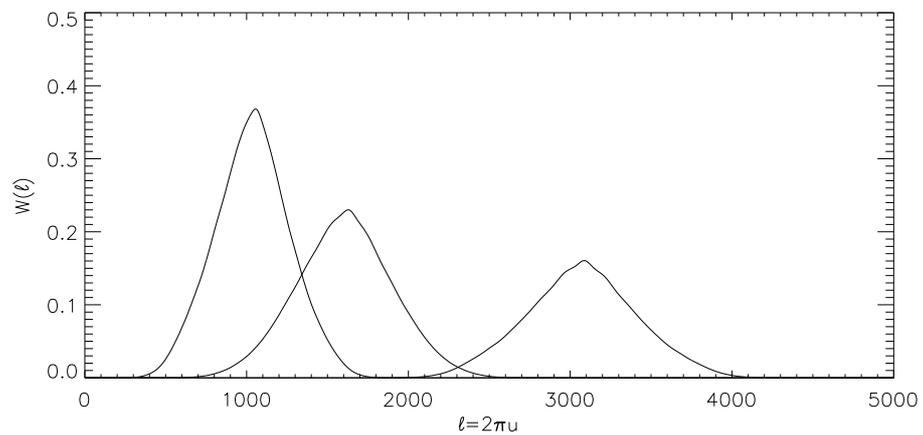


Figure 2.5: 1D MINT window functions given simulated beams. See Figure 3.7 for the beam shapes. As before, the factor $\eta_{R-J}^2 = 0.36$ is omitted.

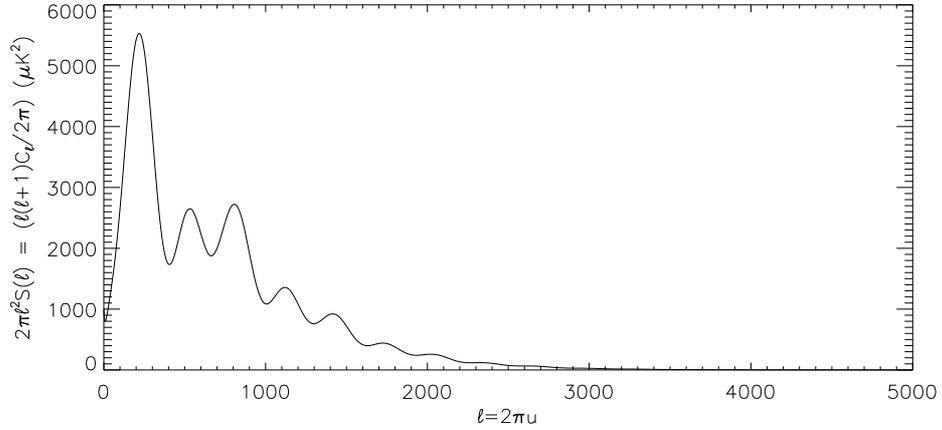


Figure 2.6: Best-fit power spectrum from the BOOMERanG collaboration. The parameters used to generate the plot are $[\Omega_{tot}, n_s, \Omega_b h^2, \Omega_{cdm} h^2, \Omega_\Lambda, \Omega_m, \Omega_b, \tau_c, h] = [1.00, 1.03, 0.023, 0.13, 0.65, 0.33, 0.05, 0.15, 0.66]$ (Netterfield et al. 2001). The parameters are as quoted for the flat “strong H_0 ” model with the exception of Ω_Λ which was corrected from 0.65 to force $\Omega_{tot} = 1.00$. Note that the ℓ -axis is linear for clarity.

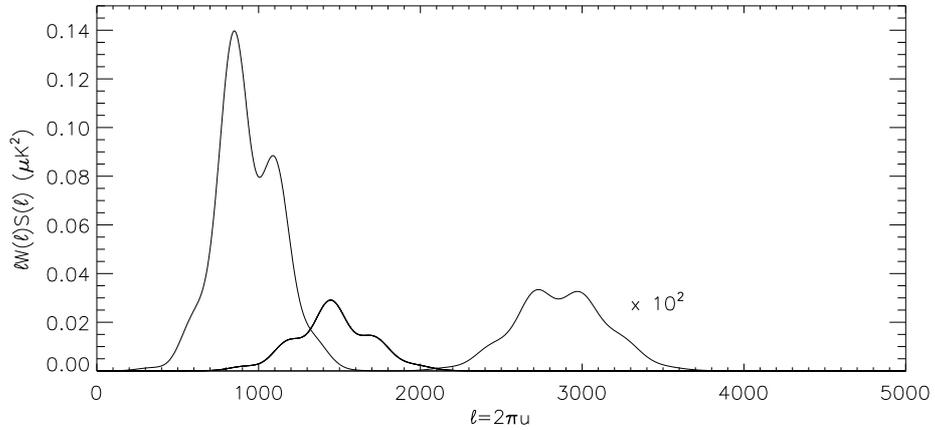


Figure 2.7: Signal before integration over ℓ , again assuming gaussian beams. The plot is shown to display the contribution of different features in the power spectrum to the MINT signal. The signal from the largest baseline is multiplied by 100 for clarity. Note that signals are heavily weighted toward lower ℓ 's because there is inherently more power there. The biggest contribution is from the shortest baseline whereas the largest baseline measures virtually no signal. These signals must be multiplied by the factor $\eta_{R-J}^2 = 0.36$ to convert to antenna temperature.

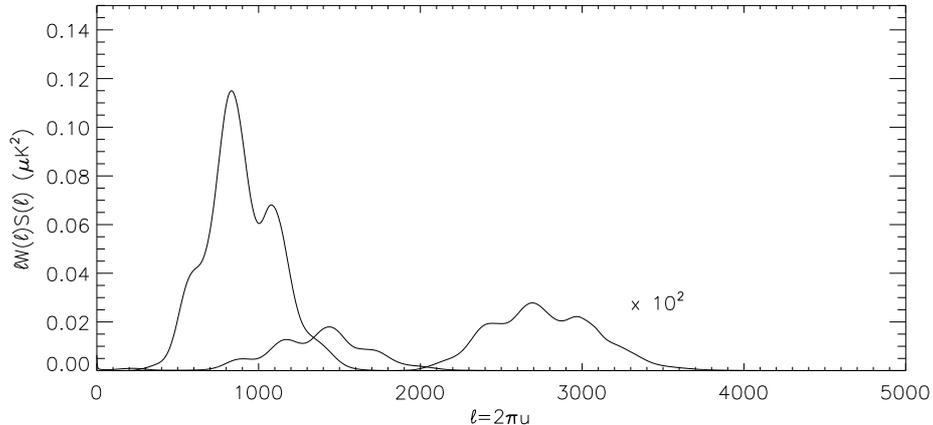


Figure 2.8: Expected contributions given simulated beams. See the caption of Figure 2.7 for explanation.

Beam type	ℓ	$\delta T = \sqrt{C^V}$
Gaussian	1062	$4.6\mu K$
	1634	$2.1\mu K$
	3090	$0.29\mu K$
Simulated	1062	$4.4\mu K$
	1634	$2.0\mu K$
	3090	$0.28\mu K$

Table 2.1: Expected signals in each of the MINT baselines. The center of the window functions are in the “ ℓ ” column. The factor $\eta_{R-J} = 0.60$ excluded from the earlier plots is now included. This is the expected *rms* signal from a temperature calibrated-interferometer.

The window functions may also be evaluated numerically using the simulated beam shapes (see Section 3.2.1). The results of the numerical analysis appear in Figure 2.5. Of course, the signal level $S(u)$ differs for different cosmologies. The best fit cosmology as measured by the BOOMERanG experiment (Netterfield et al. 2001) (see Figure 2.6) is used as a model. Figures 2.7 and 2.8 show the relative contributions of the different peaks in the spectrum to the signal for each baseline. A summary of the expected signal sizes appears in Table 2.1.

2.2.2 Noise

Given a system temperature T_{sys} , the signal detected at SNR=1 in a total power receiver with negligible gain fluctuations is calculated using the Dicke equation:

$$\delta T = \frac{T_{sys}}{\sqrt{\Delta\nu t}}, \quad (2.28)$$

where $\Delta\nu$ is the effective bandwidth, t is the accumulation time, and so $\sqrt{\Delta\nu t}$ estimates the independent number of samples, given Nyquist sampling (Kraus 1966). Since an interferometer produces the correlation of two receivers, one might expect a modified Dicke equation. The equivalent equation for the real or imaginary correlation of two noise-dominated signals is (Wrobel and Walker 1999):

$$\delta T = \sqrt{\frac{T_{sys1}T_{sys2}}{2\Delta\nu t}}. \quad (2.29)$$

This formula is accurate for a single-sideband interferometer. The error on the magnitude of the visibility comes from the sum of errors from both the real and imaginary correlations. Since the terms are uncorrelated, the errors add in quadrature and the minimum measurable magnitude of the visibility is

$$\delta T^{vis} = \sqrt{\frac{T_{sys1}T_{sys2}}{\Delta\nu t}} \quad (2.30)$$

Assuming that the noise correlation matrix is diagonal,

$$C^N = \frac{T_{sys1}T_{sys2}}{\Delta\nu t}. \quad (2.31)$$

As an initial estimate, it may be assumed that all of the receivers have the same $T_{sys} = 35$ K and that the bandwidth is $\Delta\nu = 2$ GHz. C^N does not depend on baseline length or orientation, so all the diagonal elements of the C^N matrix are approximately the same.

$$C^N \approx \left[\frac{0.78 \text{ mK}\sqrt{\text{sec}}}{\sqrt{t}} \right]^2 = \frac{0.61(\text{mK})^2\text{sec}}{t} \quad (2.32)$$

There is, however, a complication. MINT uses double-sideband mixers to produce the 2 GHz IF output. The effect is explained in Thompson et al. (1986), and the specifics in Dorwart (2002). Essentially, downconversion in the upper-sideband has the opposite frequency sign from the lower-sideband, leading to cancellation of the imaginary component. In section 3.2.3, a process for separating the sidebands and recovering the imaginary part is discussed. The result is that MINT only sees either the sine or cosine channel at any given instant. The time, therefore, in Equation 2.31, is cut in half, which adds an extra factor of 2 to C^N .

$$C^N = \frac{2T_{sys1}T_{sys2}}{\Delta\nu t} \quad (2.33)$$

2.2.3 SNR

The MINT signal to noise may be estimated by dividing C^V by C^N :

$$\frac{C^V}{C^N} = 4\pi^2 \eta_{R-J}^2 \frac{A_{e1}A_{e2}}{\lambda^4} \sigma_3^4 \frac{\Delta\nu t}{2T_{sys1}T_{sys2}} \int d\vec{u} e^{-4\pi\sigma_3^2|\vec{u}-\vec{u}_0|^2} S(u) T_{CMB}^2. \quad (2.34)$$

To make Equation 2.34 more useful, some simplifying assumptions are made. Figure 2.2 is a plot of the 2-D window functions given in Equation 2.24. The height of the window function shows that the factor $4\pi^2 \frac{A_{e1}A_{e2}}{\lambda^4} \sigma_3^4 \approx 1.0$. It is also assumed that $T_{sys1} \approx T_{sys2}$. If it is assumed that $S(u)$ is constant over the window function, the integral may be approximated as $\Delta u^2 S(u_0)$ where Δu is the approximate width of the 2D window function. Equation 2.34 may then be approximated as:

$$\frac{C^V}{C^N} \approx \frac{\eta_{R-J}^2}{2} \Delta\nu t \frac{T_{CMB}^2}{T_{sys}^2} \left[\frac{\Delta u}{u_0} \right]^2 u_0^2 S(u_0), \quad (2.35)$$

where the factor $\frac{u_0^2}{u_0^2}$ has been inserted to produce $u_0^2 S(u_0)$, the commonly plotted power spectrum. The *rms* signal-to-noise is the square root of this equation. This equation may be directly compared with equation 17 in WCDH. The factor $\left[\frac{\Delta u}{u_0} \right]^2$ may be called the ‘‘filling factor’’ and refers to the fraction of the u - v plane sampled by the window function relative to the total part of the u - v plane that contributes to $u_0^2 S(u_0)$. The magnitude of the baseline vector, u_0 , is proportional to the antenna

separation and Δu is proportional to the diameter of the antenna, so the filling factor may be made larger by filling in as much of the area of the interferometer as possible, or, in other words, compacting the array.

Equation 2.35 differs from Equation 17 in WCDH by the factor of 2 in the denominator, which is due to the double sideband downconversion in MINT.

2.2.4 Measuring δT_ℓ

It is generally assumed that $\Delta T(\vec{x}) = T_s(\vec{x}) - T_{CMB}$ is a random normally distributed field. The aim is to measure the standard deviation of this distribution, δT_ℓ , within a certain range of angular scales as defined by the window function. The first task is to construct an estimator for δT_ℓ from what is actually measured, the visibility, V .

Before constructing the estimator, the observation strategy must be examined. Equation 2.34 gives the square of the signal-to-noise for a single pointing on the sky if a time t is spent observing that point. Knowing the visibility at that single point, however, is of limited value because the anisotropy is the change in visibility over the sky. To measure δT_ℓ , the parent distribution must be adequately sampled, which means that many points must be measured on the sky. Each measured visibility is labelled V_i with $i = 1, 2, \dots, n$ where n is the number of independent fields on the sky. In the following section, the relative error in determining δT_ℓ will be estimated for a single baseline. The estimate for a single baseline may be combined to estimate the overall relative error. The visibility, V_i , is formed from both a signal and a noise component,

$$V_i = D_i + N_i, \quad (2.36)$$

where D_i is the signal and N_i is the noise. Note that the above equation is complex. All terms have both a real and imaginary component and the equation may be thought of as two separate equations for either the real or imaginary parts. The first attempt at an estimator⁴ might be $\langle V^* V \rangle_i$ with the property that

⁴This construction is usually considered the “quick and dirty” way to calculate the answer, but is adequate for determining the error bars. A serious analysis would involve likelihood analysis.

$$\langle V^*V \rangle_i = \langle D^*D \rangle_i + \langle N^*N \rangle_i, \quad (2.37)$$

where the brackets denote ensemble averaging and it has been assumed that the real and imaginary parts of the signal and noise are gaussian and uncorrelated and furthermore, that the noise is independent of the signal. There is a cross term formed by the very small D_i and the uncorrelated N_i , which averages down to zero faster than any other terms of interest. The quantity $\langle D^*D \rangle$ can be identified as the power spectrum, S , as in Section 2.1.1. The above equation is effectively a statement of variances,

$$2\sigma_V^2 = 2\sigma_D^2 + 2\sigma_N^2, \quad (2.38)$$

with $2\sigma_D^2$ being the quantity to be estimated. The σ^2 's alone are the variances of either the real or imaginary components. The extra factor of 2 comes from combining the real and imaginary parts of the visibility. A useful statistic to define is,

$$K^2 = \sum_{i=1}^n \frac{V_i^*V_i}{\sigma_V^2}, \quad (2.39)$$

where K^2 follows the χ^2 distribution with mean $2n$ and variance $4n$. The real and imaginary components are treated separately, which effectively means that the degrees of freedom are double the number of observed fields. The final statistic to be formed is:

$$\Delta^2 = \frac{2\sigma_V^2 K^2}{2n} - 2\sigma_N^2. \quad (2.40)$$

The expectation value of Δ^2 is the quantity of interest, $2\sigma_D^2$. The *rms* of Δ^2 is just $2\sigma_V^2/\sqrt{n}$. The $2\sigma_N^2$ noise term does not contribute to the *rms* because it is constant for all points on the sky. It is now possible to estimate the fractional error on Δ^2 :

$$\frac{\delta\Delta^2}{\Delta^2} = \frac{rms(\Delta^2)}{\langle\Delta^2\rangle} = \sqrt{\frac{1}{n}} \left(1 + \frac{\sigma_N^2}{\sigma_D^2} \right). \quad (2.41)$$

ℓ	$\delta T = \sqrt{C^V}$	t_{un}	$n_{optimum}$
1062	$4.4\mu K$	17.6 hr	20
1634	$2.0\mu K$	77.1 hr	5
3090	$0.28\mu K$	4046 hr	.09

Table 2.2: The equivalence time as calculated from Equation 2.45. t_{un} is defined as the time that required to reach a S/N of 1 in a given pointing. $n_{optimum}$ is the optimum number of observing spots given a total observing time of 360 hours.

To estimate the error for δT , which is called⁵ $\delta(\delta T)$, it must be recognized that $\delta T \propto \sqrt{\Delta^2}$, so that,

$$\frac{\delta(\delta T)}{\delta T} = \frac{1}{2} \frac{\delta \Delta^2}{\Delta^2} = \frac{1}{2\sqrt{n}} \left(1 + \frac{\sigma_N^2}{\sigma_D^2} \right). \quad (2.42)$$

2.2.5 Observation strategy

The actual error bars are strongly dependent on how the observing time is divided. To make this clear, an important quantity, t_{un} , is defined as the time it takes to achieve a signal-to-noise ratio of unity. Setting the noise from Equation 2.33 equal to the visibility:

$$C^N = \frac{2T_{sysi}T_{sysj}}{\Delta\nu t_{un}} = C^V, \text{ so that} \\ t_{un} = \frac{2T_{sysi}T_{sysj}}{\Delta\nu C^V}. \quad (2.43)$$

For the favored model in Netterfield et al. (2001), $\sqrt{C^V}$ are given in Table 2.1. If the approximations $T_{sys} \approx 35$ K and $\Delta\nu = 2$ GHz are taken, the values of t_{un} may be tabulated (Table 2.2).

Equation 2.42 may be rewritten as:

$$\frac{\delta(\delta T)}{\delta T} = \frac{1}{2\sqrt{n}} \left(1 + \frac{t_{un}}{t} \right). \quad (2.44)$$

⁵This rather awkward double use of the symbol δ results from the common use of δ to refer to temperature fluctuations. It should be understood that the first δ refers to “error on” and the second delta refers to temperature fluctuations. The fractional error on δT is therefore $\frac{\delta(\delta T)}{\delta T}$

In the limit of long observations per spot, $t \gg t_{un}$, the sensitivity reduces to $\delta(\delta T)/\delta T = 1/(2\sqrt{n})$ which is known as the sample variance. This is the limiting sensitivity of the observing strategy; observing for too long on one spot has diminishing returns. On the other hand, in the limit of little observing time, $t_{un}/t \gg 1$, the error on the bandpower, $\delta(\delta T)/\delta T$ is dominated by the second term and the error bars are SNR-limited.

There is an optimum time per spot given an estimate of the total observing time, t_{tot} . This optimum time, \hat{t} , may be found by substituting $n = t_{tot}/t$ and minimizing Equation 2.44 with respect to t with the result:

$$\hat{t} = t_{un} \quad \text{and} \quad \left. \frac{\delta(\delta T)}{\delta T} \right|_{min} = 1/\sqrt{n}. \quad (2.45)$$

If MINT had the ability to track the sky, the number of fields would be chosen to be $n = t_{tot}/t_{un}$ widely separated on the sky. This number is different for each baseline so some compromise would be made. The last column of Table 2.2 lists the optimum number of spots with a total observing time of 30 days and 12 hours per day.

MINT, however, was not designed to track. The first scan strategy considered was a drift-scan, wherein the telescope points at a fixed azimuth and elevation, say at zenith, and the sky rotates though the field. This strategy has numerous instrumental advantages, the best of which is that it keeps the telescope stationary. An estimate of the number of independent fields is the number of beam widths on the sky as defined by the FWHM of each beam. The total number of degrees that pass though the MINT beams is given by $\frac{1}{2}360^\circ \cos(\text{latitude})$, assuming that the telescope is pointed at zenith. The latitude at the Chilean site is $\sim -22.9^\circ$ and is given in Table C.1, a summary of the GPS measurements. The factor of $\frac{1}{2}$ is from observing only at night. With the drift-scan strategy, there are 166 $1/2^\circ$ beam widths, well above the optimum value.

There were two different things done to decrease the sampling (or under-integration) of the sky. The first concerned the array design. The layout of the antennas was chosen to maximize redundancy. The four antennas are arranged in a rhombus with the large antennas placed at the far corners. In terms of baselines, this arrangement produces one short baseline formed from the correlation of the signal from the two small antennas, 4 medium baselines with the same length from a large and a small antenna,

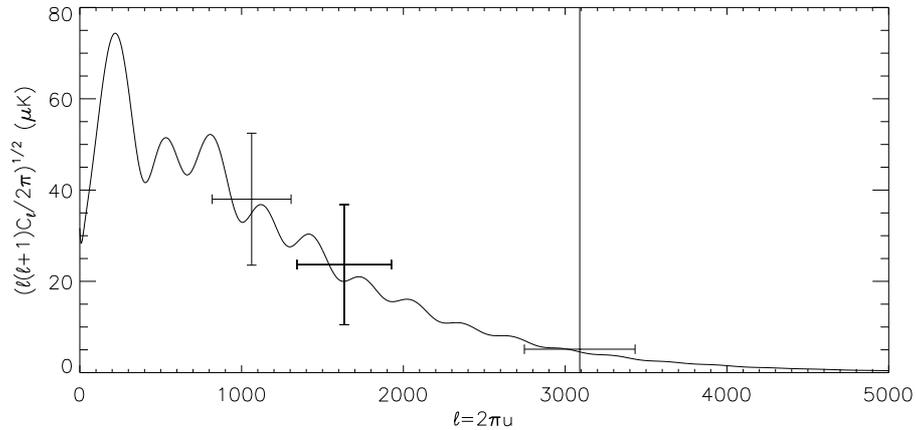


Figure 2.9: Expected error bars from observing 166 uncorrelated fields for 30 12-hour days. The horizontal error bars are the FWHM of the window functions. This is the estimate of the error bars if MINT employed a zenith-drift-scanning technique. The estimate for the scan strategy that was actually used appears in Figure 2.11.

and one long baseline from the two large antennas. See Figure 3.1 for the antenna layout. The middle baseline lengths are identical and each is made from a small and large antenna, producing identical 1-D window functions sampling the same ℓ -range. Furthermore, the rhombus layout leads to 2 pairs of identically oriented baselines. Since each member of a pair samples exactly the same u - v space, it is equivalent to doubling the integration time per sample. As seen from Figure 2.3, the pairs of middle baselines are only slightly correlated. If they were completely uncorrelated, this would be equivalent to observing more fields and the signals may be averaged to reduce the error bar on $\delta T/T$ by $\sqrt{2}$. Figure 2.9 is an estimate of the error bars in this over-sampling regime.

The second technique that reduces the sampling is a modification of the scan strategy. The actual scan pattern is a hybrid between drift-scanning and tracking. Figure 2.10 is a diagram of the scan strategy. In the modified drift-scan, a single patch of sky is allowed to drift through the beam eight times per night. To accomplish this, the telescope azimuth is pointed east, and does not change during the scan. At the beginning of the campaign, 24 evenly spaced “marks” on the sky were chosen centered evenly (for convenience) on integer Right Ascension (RA) hours. The algorithm works by first defining 8 elevation bins, each defined by the time when the mark passes the

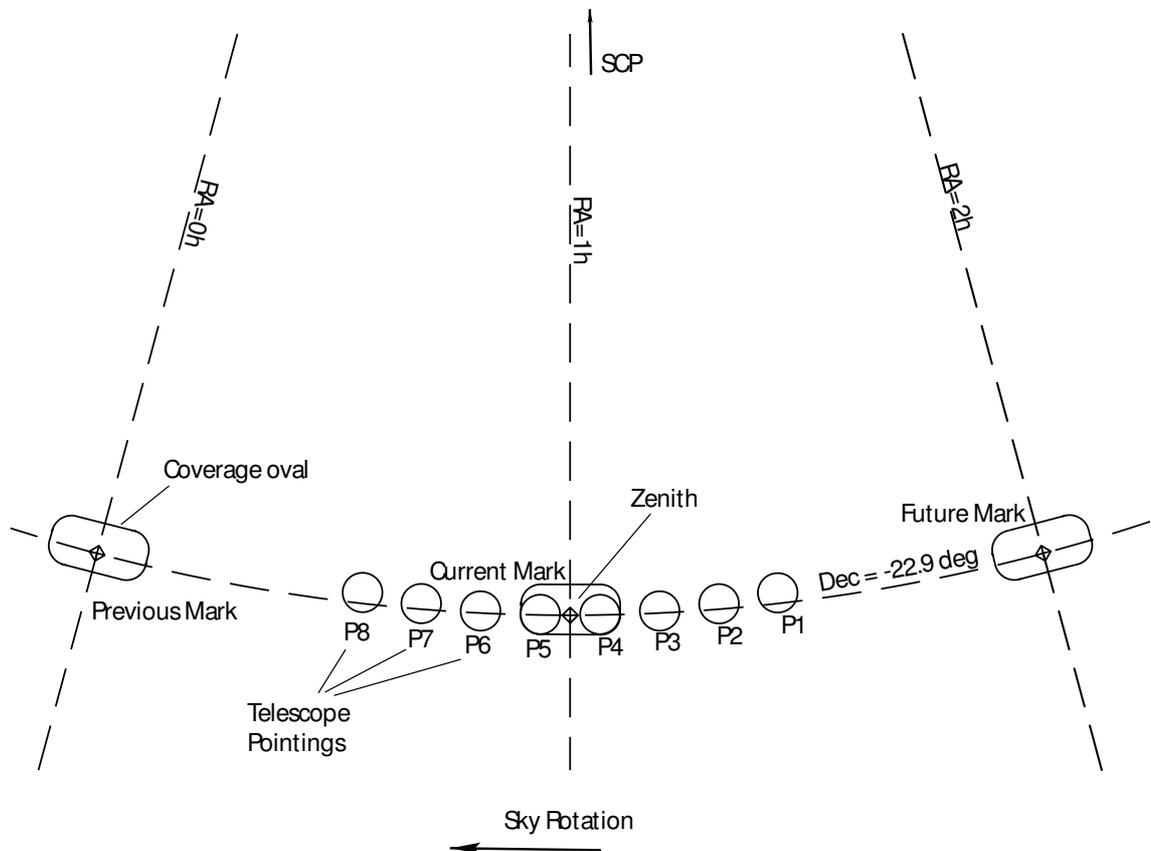


Figure 2.10: Scan strategy and sky coverage. In this RA- δ projection, where δ =declination. The South Celestial Pole (SCP) is off the top of the figure and the zenith is in the lower middle. The “Marks” are represented as diamonds, and are evenly centered on integer RA hours. As a mark approaches zenith, the telescope slews in elevation and points to P1. It stays there for 7.5 minutes, allowing the mark to pass through, before slewing to P2 and so forth. After it gets to P8, the telescope slews again to P1 and starts scanning the next mark. The times that the telescope spends at the given pointings is given in Table 2.3. This snapshot of the sky is taken just as the current mark is passing through zenith and the telescope is moving from P4 to P5. The coverage ovals show the area of the sky that passes through the main beam. They are thicker in the δ direction than one beam width because the pointings (P1-P8) do not all lie at exactly the same declination. P4 and P5 are the closest to $\delta = -22.9^\circ$, the latitude of the site. This is because the telescope does not change azimuth when it slews in elevation. This technique is an approximation to tracking the mark, and is done to reduce sky coverage by a factor of 8 over drift-scanning alone.

label	Time range from transit sidereal minutes	Center of elevation bin actuator degrees
P1	-30.0 – -22.5	83.957
P2	-22.5 – -15.0	85.684
P3	-15.0 – -7.5	87.410
P4	-7.5 – 0.0	89.137
P5	0.0 – 7.5	90.863
P6	7.5 – 15.0	92.590
P7	15.0 – 22.5	94.316
P8	22.5 – 30.0	96.043

Table 2.3: The elevation bins defined for the modified scan strategy. The actuator readings differ from the actual elevation because the actuator is allowed to point the telescope backward, and hence, to readings higher than 90° . The labels refer to Figure 2.10.

though. Defining 8 bins for each mark means that the mark passes through each bin for $1 \text{ hour}/8 = 7.5$ sidereal minutes. The first bin consists of a range of elevation that contains points on the sky that are between -30 and -22.5 sidereal minutes from crossing the zenith. The next bin is defined from -22.5 to -15 minutes and so forth until 8 bins are defined. Each bin corresponds to a fixed elevation range (see Table 2.3). Each elevation bin is ~ 3.5 1/2-degree beam widths wide. The algorithm then finds which bin contains any of the marks and then points the telescope to the center elevation of the bin. The mark, along with some length of leading and trailing RA, is allowed to drift through the beam for 7.5 sidereal minutes at which point it moves to the next bin. The same patch of sky surrounding the mark is allowed to drift through the beam for another 7.5 sidereal minutes. This is repeated for a total of eight times on any given mark until another mark enters the first bin. This strategy roughly reduces the number of points on the sky from 166 to 21, a much better match to Table 2.2. Figure 2.11 shows the expected error bars given this scan strategy.

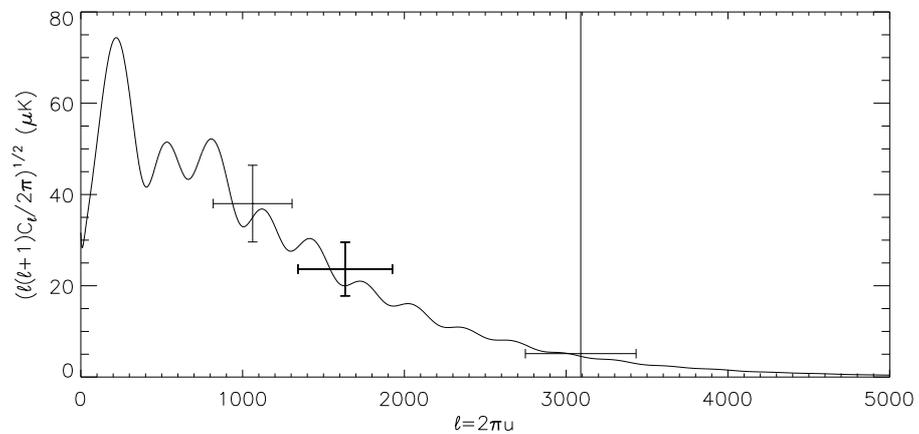


Figure 2.11: Expected error bars from observing 21 uncorrelated fields for 30 12-hour days. This is an estimate of the error bars given the modified scan strategy that was actually used. Compare these estimates to Figure 2.9. The reduction in the error bars is a result of reducing the number of observed fields.

Instrument Overview

3.1 Advantages of Interferometry

The use of interferometers to measure the CMB had been considered early in the search for degree-scale CMB anisotropy. The first instrument designed for CMB studies was built as a Ph.D. thesis at Princeton (Timbie 1985). The next significant attempt at using interferometry to detect the anisotropy was with the VLA (Partridge et al. 1987). The first to actually detect anisotropy was the CAT instrument located in Cambridge (Jones 1997). Another significant detection came from a two-element interferometer at Tenerife (Dicker et al. 1999). These instruments began to show the promise of interferometry in measuring the CMB. Following the success of the CAT instrument, the Cambridge group began construction of the larger-scale VSA (Very Small Array), designed to measure the anisotropy in the mid ℓ range. Observations with the VSA are currently underway.

Recently, there have been a number of interferometers that have made breathtaking measurements of the CMB. Of particular interest are the results from DASI (Halverson et al. 2001), located at the South Pole, and CBI (Padin et al. 2001a), located in the Atacama desert in Chile. Together these two instruments were designed to cover the anisotropy spectrum over most of the acoustic peak range. The DASI instrument has already produced a high quality spectrum over the mid ℓ range and the CBI results are eagerly anticipated. Both instruments share many subsystems and inspired some of the design principles for MINT.

The way that MINT fits into the bigger picture of current CMB experiments is discussed in Section 1.2. What is important to point out here is the uniqueness of the

MINT design relative to other interferometers. MINT will measure the C_ℓ spectrum in a very similar range to that of CBI, but since the frequency of operation of MINT is roughly five times higher, the size is five times smaller. More importantly, the high frequency makes MINT almost immune to point source contamination, one of the primary limitations of the measurements from CBI.

There are also some major instrumental differences. MINT is a heterogeneous array, which better fills the available aperture but makes the analysis slightly more complicated. Both CBI and DASI use HEMT amplifiers as the critical component of the receiver. MINT uses SIS mixers. SIS mixers, however, present the best sensitivity at 150 GHz. Additionally, the SIS mixers operate in double-sideband mode, adding complications to the operational design and data reduction procedure. Almost all interferometers require the breaking up of bandwidth into smaller portions and downconversion to baseband. The MINT instrument accomplishes this in one monolithic device (the channelizer) as opposed to an assembly from components. Finally, MINT employs a fully digital correlator for frequency resolution, stability, and systematic control, bringing with it a host of experimental challenges. All other CMB interferometers use analog correlators.

3.1.1 Stability

The resolution of an interferometer is determined by the separation of antennas, not by the size of a single dish. The much touted ability of interferometers to resolve fine details is not actually the principal reason to use them for CMB studies. The angular scales of interest in studying the damping tail are large enough that single-dish, small-scale (~ 1 m) instruments are still possible, particularly at the high frequency of MINT. In fact, the added complexity that comes with more than one receiver and the correlator outweighs the advantages of using smaller antennas. As discussed in Chapter 2, interferometers also suffer from a decrease in signal relative to a filled-dish instrument due to the filling factor. For a small-scale instrument, single dish filled antenna instruments are both easier to make and are more sensitive.

The real limitation of CMB instruments, however, is systematic error. The signal is 5-6 orders of magnitude weaker than the background noise, making stability paramount. Instrumental offsets can overwhelm the signal.

The two major systems that interferometers use to tackle this problem are correlators and the 180° phase switches. Neither of these systems is possible in total power instruments. All uncorrelated signals average to zero as \sqrt{t} in the correlator. This includes the base temperature of the atmosphere, CMB, and the receiver. The 180° phase switch is used to reduce internal instrumental offsets due to unwanted correlations. Together these systems produce a very stable platform to measure the CMB.

3.1.2 Antenna Layout and Instantaneous Beam Switching

The first set of instruments to detect the degree-scale anisotropy were beam-switching experiments, whereby a telescope beam was switched, or swept, across the sky, and differences were taken in the time stream data to produce δT measurements. Interferometers, on the other hand, perform this measurement instantaneously. The response pattern of an interferometer may be thought of as sinusoidal corrugations attenuated by an overall beam envelope. With a complex correlator, the interferometer also produces a cosinusoidal beam. Together these beams form the Fourier transform of the sky, the desired measurement.

The orientation and spacing of the corrugations is determined by antenna layout. The MINT layout was designed to maximize sensitivity subject to the constraint of desired ℓ -coverage. The layout appears in Figure 3.1. The vectors that connect the centers of the antennas are shown at the bottom, known as the baselines or \vec{u}_i , as in Figure 2.1. The length of \vec{u}_i is inversely proportional to the wavelength of corrugations on the sky, and the corrugations are perpendicular to the vector, as shown in Figure 3.2. The 2-D corrugations are found from the Fourier transform of the 2-D window functions in Figure 2.2. The longer the baseline, the denser the corrugations and the higher the ℓ coverage.

The response pattern that is sensitive to the lowest ℓ 's is from the smallest baseline, A-D. The ℓ coverage is centered on $\ell=1062$. Estimates from Chapter 2 suggest that this single pair alone has the sensitivity to detect fluctuations. As the baselines move higher in ℓ , however, the power is dramatically damped. To maximize the sensitivity in the next ℓ -range, the other two antennas are placed to construct a rhombus. The four sides of the rhombus lead to 1-D window functions centered on $\ell=1634$. The

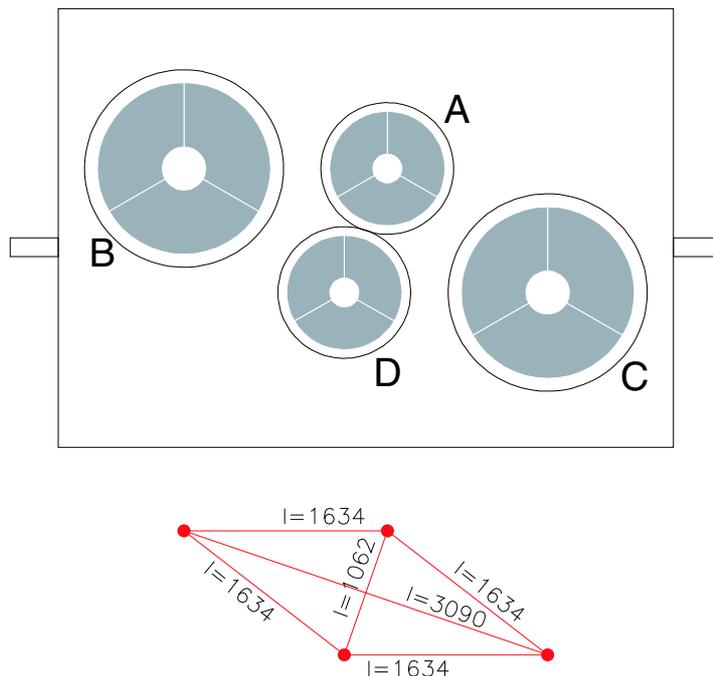


Figure 3.1: Overhead view of antenna layout. The top drawing shows the placement of both the large and small optics. The physical distance between the centers of antennas B and C is 1 m. The bottom drawing shows the baselines, \vec{u}_i , resulting from the antenna locations. The numbers indicated along the baselines are $\ell = 2\pi u$. Figure is courtesy of Joseph Fowler.

two pairs of parallel baselines measure exactly the same response patterns on the sky, effectively doubling the observing time. To further maximize sensitivity, the outlying antennas were enlarged to better fill in the available area. The largest baseline is formed from the B-C pair. Although the results from Chapter 2 conclude that this baseline will have no chance of detecting the anisotropy at such high ℓ , this baseline remains the most sensitive to point sources, including calibrators, due to the large collecting area.

3.2 Signal Path

Modern interferometers employ complex receiving systems that process the signal in many steps. The general problem is to convert the high frequency signals into a form that the correlator can process. The cosmic signals start out centered on \sim

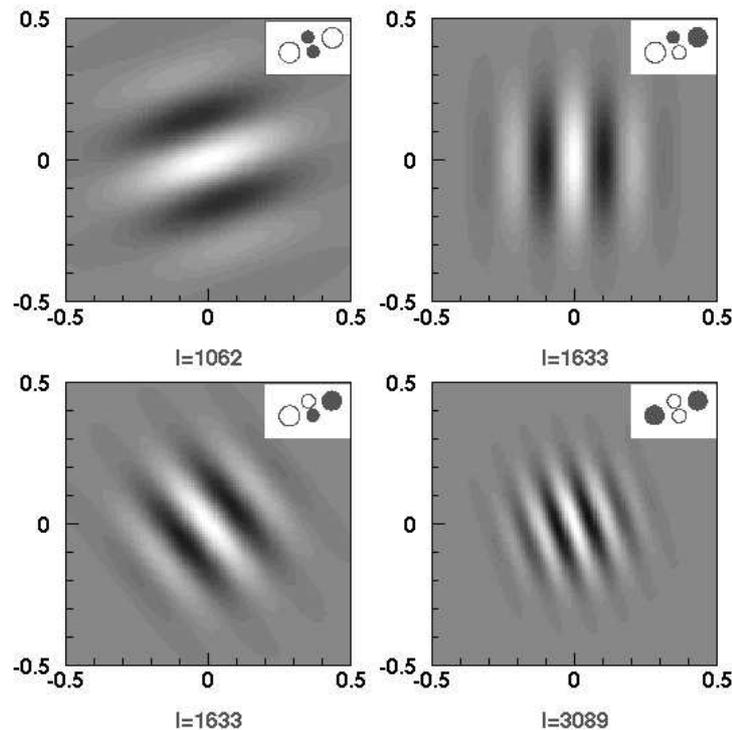


Figure 3.2: Response patterns for MINT baselines. Axes are in degrees. The cartoon in the upper-right corner of each pattern show the pair of antennas that lead to the pattern. Although there are six MINT baselines, only four are unique. Figure is courtesy of Joseph Fowler.

150 GHz. The correlators can process 500 MHz wide baseband signals (signals from 0-500 MHz). MINT accomplishes the task with two separate stages of downconversion, and a filterbank.

Figure 3.3 is a diagram of the signal path. It represents one elemental baseline. The signal is collected by the antennas and fed to the SIS mixers. The mixers down-convert the cosmic signal to a 4-6 GHz IF band. The relative phase between receivers of the downconversion is maintained by the phase-locked loop. This IF signal is then processed by the channelizer, a monolithic filterbank and downconverter that outputs four 500 MHz-wide baseband signals for each receiver. The signals are fed to the correlator, which calculates 24 16-lag cross correlations and outputs them to the data computer for storage and transmission to the ground station.

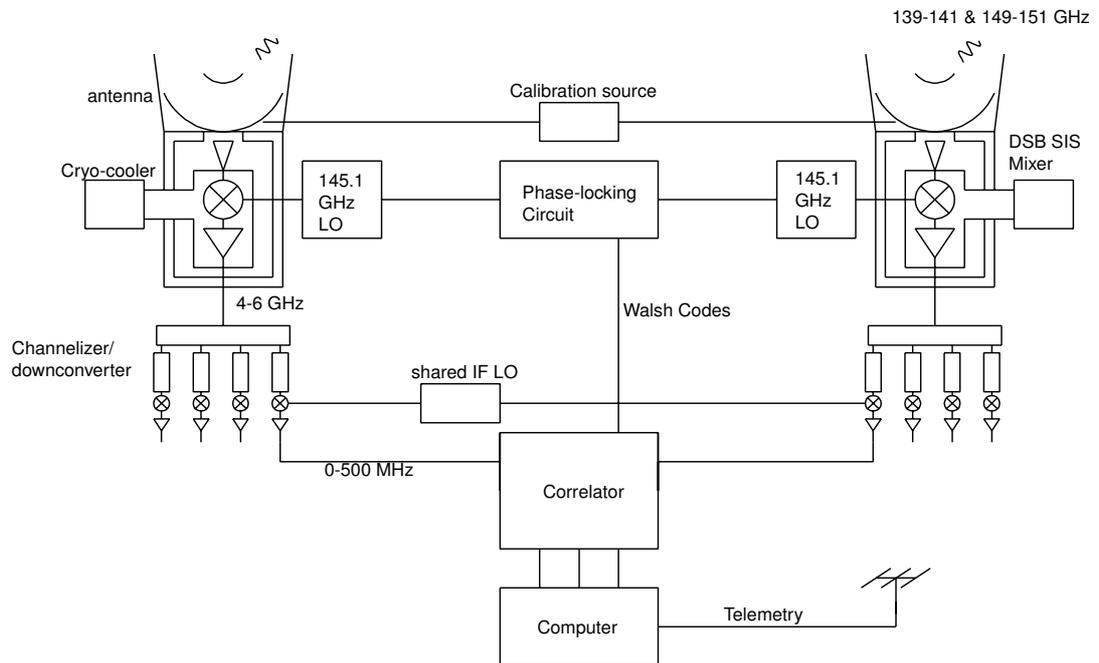


Figure 3.3: The MINT signal path for one baseline

3.2.1 Antennas

Cassegrain optics focus the cosmic signal down to corrugated feed horns. They consist of a parabolic primary mirror and a hyperbolic secondary mirror. There is a hole in the primary that allows the rays to pass through to prime focus, where the feed horn is located. The secondary is held in place with G10 support legs that are epoxied in place. The entire structure is enclosed by a shield, which is meant to decrease antenna-to-antenna crosstalk. The design is similar to the antennas used on CBI (Padin et al. 2001b) and other microwave experiments, where the measured crosstalk was ~ -110 dB, a level that will produce sufficiently small false correlations. The MINT optics should have comparable, if not slightly better, isolation because the shield on MINT is about twice the relative height of CBI's and the MINT primary is illuminated with about 10 dB more edge taper, leading to less diffraction. There are two antenna sizes; the larger is a linearly scaled version of the smaller. A diagram of the smaller antenna appears in Figure 3.4. The antennas were primarily designed by

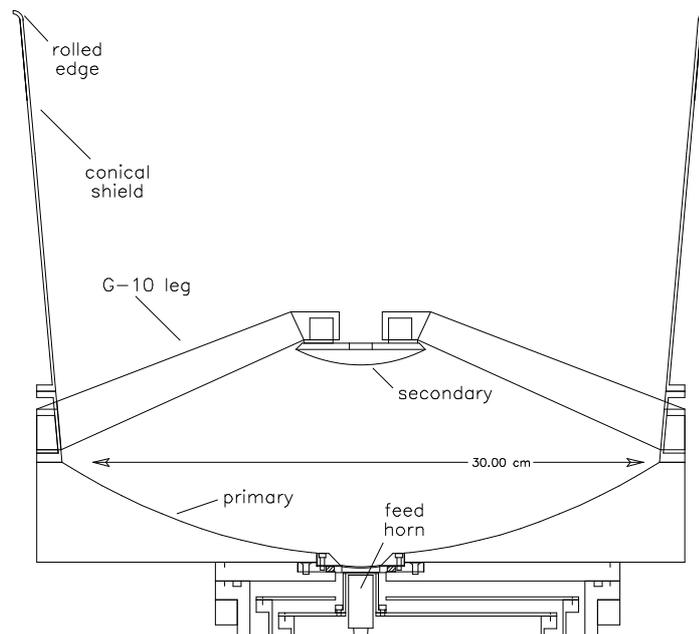


Figure 3.4: Cutaway diagram of the smaller optics. The larger ones are linearly scaled in every dimension by a factor of 1.5. The 30 cm paraboloidal primary rests on top of the dewar. The hyperbolic secondary is held in place by a three-legged G-10 support spider. The entire antenna is encased inside a conical shield, rolled at the top to minimize diffraction. Figure is courtesy of William B. Dorwart and also appears in Dorwart (2002).

Yeong Loh and Joseph Fowler. A more detailed account can be seen in Loh (2000a) and Loh (2000b).

Briefly, the antennas were designed to minimize sidelobes and concentrate as much power as possible into the main beam. The design process consists of setting geometrical parameters using analytical formulae and modeling performance using *DADRA* (*Diffraction Analysis of a Dual Reflector Antenna* (Rahmat-Samii et al. 1996)), an analysis package which solves for the electromagnetic response of metal surfaces. The parameters are then tweaked to obtain the desired result.

To measure the beam profile, the optics were mounted on the MINT base and pointed at a chopped noise source. The signal was synchronously detected using a lock-in amplifier. After the first mirror was fabricated and tested, it became clear that there was a problem with the simulations (see Figure 3.5). The original optics

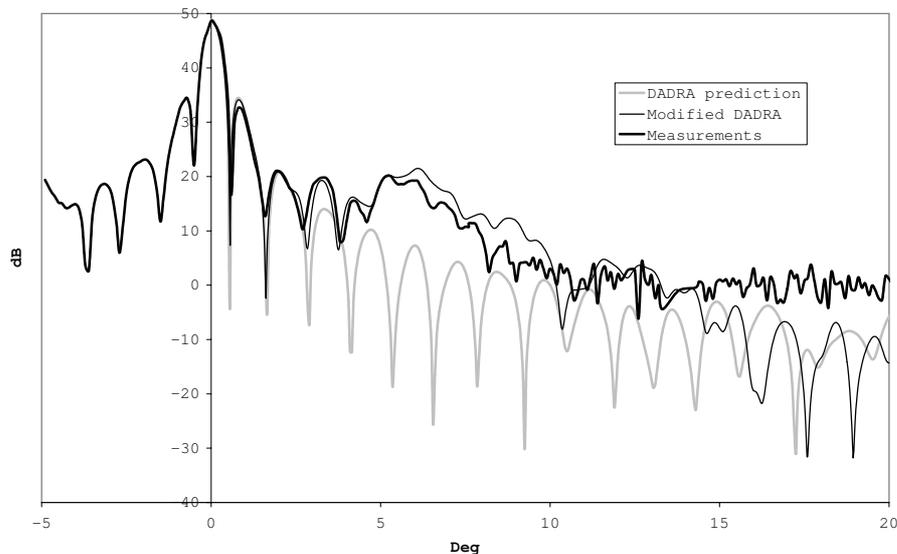


Figure 3.5: Beam map from the first run of optics testing. The extra power in the measured beams near 5° is from rays that have made a “double bounce” off the primary which couples in stray radiation from farther off-axis. After it became clear that the DADRA simulations were incorrect, the code was modified to generate the corrected prediction. The y -axis is in dB gain over a hypothetical isotropic antenna (dBi). The measured data were normalized to the predictions at the peak. These optics are not used in the final interferometer.

had an overly small opening in the primary for the feed horn. This extra surface area of the primary formed a cavity with the outer edge of the secondary. The original *DADRA* code was not designed to account for more than one reflection from a given surface. The new prediction seen in Figure 3.5, now much closer to the measurement, was made using code modified by Joseph Fowler to handle extra bounces.

Based on the new simulations and measurements, it was decided that too much power was outside of the main lobe. A new set of secondaries was designed and fabricated. The predicted beams appear in Figures 3.6 and 3.7.

3.2.2 Receiver

The two main types of receivers used for CMB interferometry are based on HEMT (High Electron Mobility Transistor) amplifiers and SIS (Superconductor Insulator Superconductor) mixers, both developments from the NRAO (Pospieszalski et al.

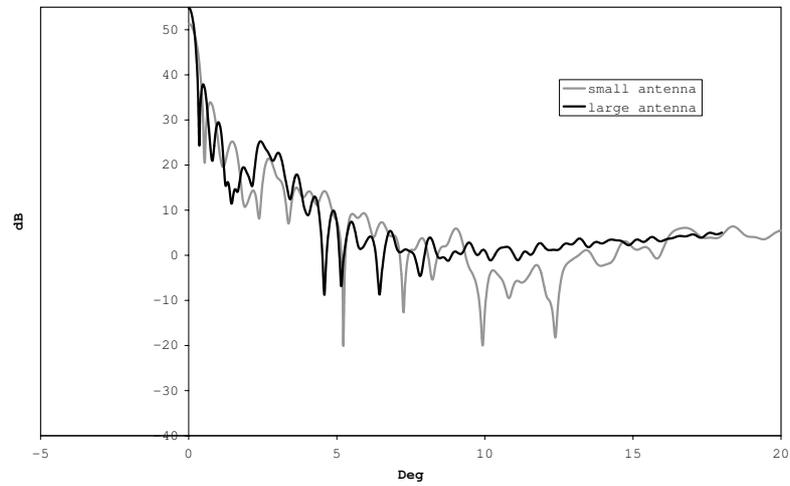


Figure 3.6: Predicted beam map from Modified DADRA Code. MINT employs 2 large antennas and 2 small antennas.

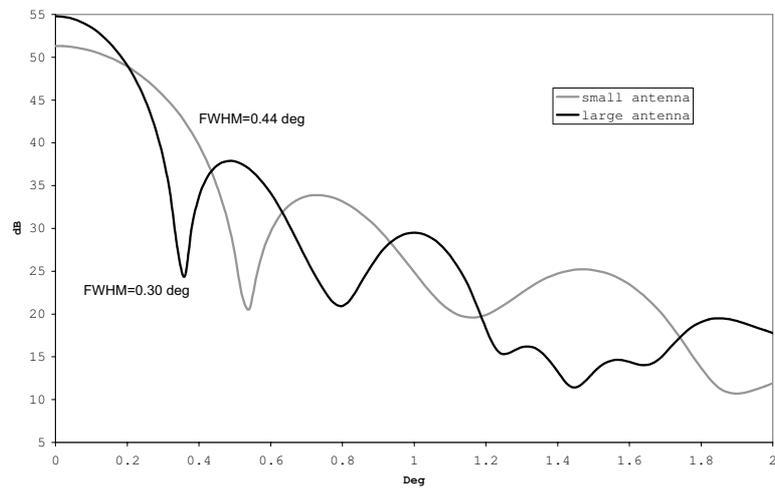


Figure 3.7: An enlarged view of the main lobes of both antennas

1993; Kerr et al. 1993). The three major dedicated CMB interferometers, DASI, CBI, and VSA, employ HEMT amplifiers. Although great advances have been made in pushing the upper-frequency limit, the usable frequency of these devices is so far limited to less than 110 GHz. MINT is the first dedicated CMB interferometer to use SIS mixers. SIS's are already used on BIMA (Berkeley Illinois Maryland Association) and OVRO (Owens Valley Radio Observatory), two mid-scale interferometers, and they are featured in the planned large-scale interferometers, the SMA (Sub-Millimeter Array), ALMA (Atacama Large Millimeter Array) and the SZA (Sunyaev-Zel'dovich Array).

The receivers were designed, built and tested by Randy Dorwart. In-depth details of the receiver may be found in Dorwart (2002), Wesley (2000), and Dumont (2001). The cryogenic components of the receiver, including the feed horn, are housed in a vacuum dewar, which is cooled via a 3-stage Gifford-McMahon mechanical cryocooler. An attached RBE (Receiver Back End)¹ or “backpack,” contains the warm components and support electronics for the entire receiver.

SIS mixers are made of a sandwich of thin layers of superconductors and insulators, forming Josephson junctions. A mixer possesses a non-linear I-V curve, which produces the multiplication of two input signals. The multiplier can be used as a downconverter, which converts radiation from frequency ν_{RF} to frequency ν_{IF} where

$$\begin{aligned} \text{Upper Sideband: } \nu_{IF} &= \nu_{RF} - \nu_{LO}, \\ \text{Lower Sideband: } \nu_{IF} &= \nu_{LO} - \nu_{RF}. \end{aligned} \tag{3.1}$$

For a single ν_{IF} there is an ambiguity as to what the original ν_{RF} was. It could have been from either $\nu_{RF} = \nu_{LO} \pm \nu_{IF}$, named the upper and lower sidebands. SIS devices are more sensitive when used in double sideband mode, with an important caveat when used in interferometers. The sensitivity of a double sideband interferometer is $\sqrt{2}$ lower than an equivalent single sideband interferometer with the same bandwidth and system temperature². This reduction is evident in Equation 2.33 as a reduction in time by a factor of 2 and is discussed in more detail in Dorwart (2002).

¹This is not really the backend of the interferometer. The actual backend consists of the channelizer and correlator. RBE is a legacy term referring to a time when the science detectors were

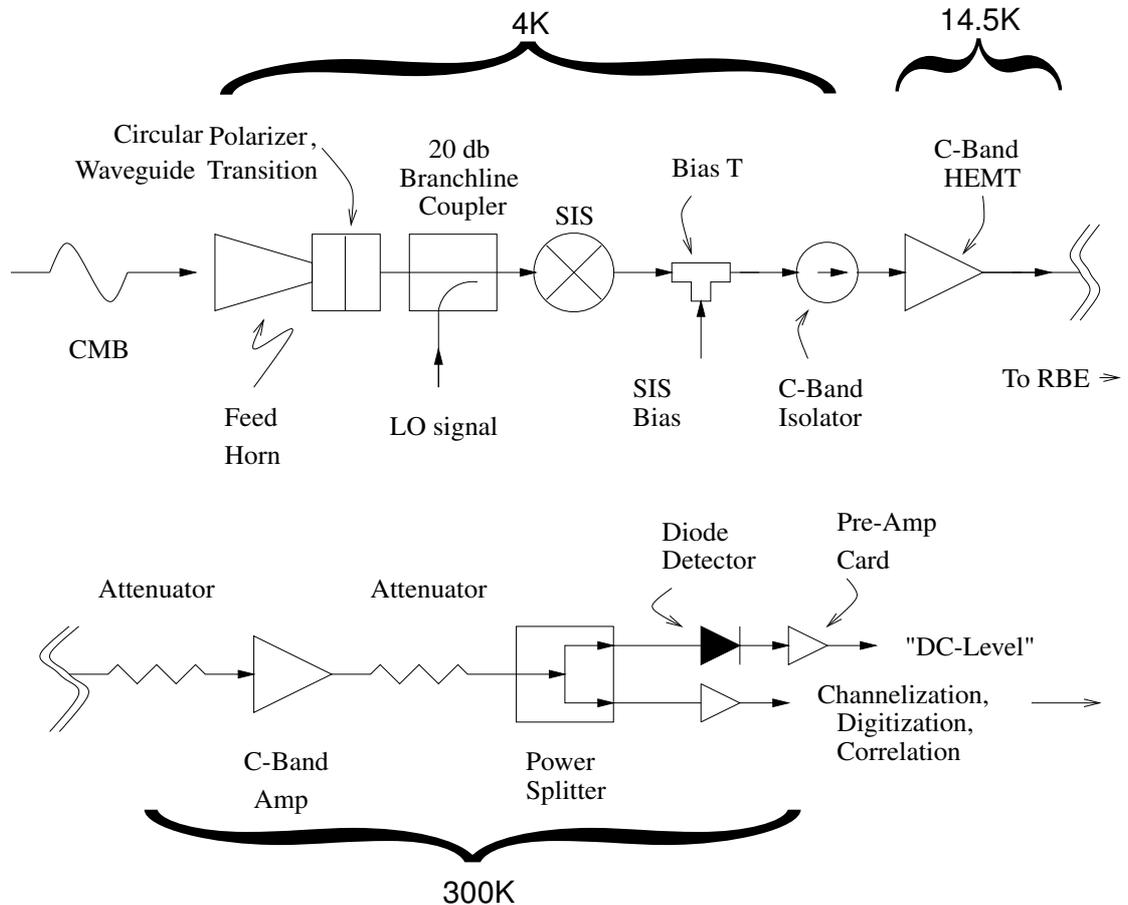


Figure 3.8: Microwave components of the MINT receiver. Adapted from Figure 3.1 in Wesley (2000). In two of the dewars, there is a bandpass filter between the power splitter and detector diode. More details may be found in Dorwart (2002) and Dumont (2001).

Figure 3.8 is an outline drawing of the relevant microwave components in the receiver, including the SIS. The unpolarized microwave radiation enters a corrugated feed horn from the left. The radiation is polarized and transmitted into rectangular waveguide with the combination circular polarizer/waveguide transition. Next is the

actually in the backpack. There still are detectors in the backpack used for diagnostics.

²The bandwidth here refers to the IF bandwidth of the mixer. The sensitivity of a double sideband interferometer is $\sqrt{2}$ better than the same interferometer with a RF filter that blocks one of the sidebands. The reason is because the system temperature of the filtered interferometer is twice that of double sideband one. See Thompson and D'Addario (2000) for a complete discussion.

branchline coupler, which combines the LO power with the CMB radiation. After the coupler, the radiation enters the SIS cavity, whose size is tunable via a backshort. The SIS downconverts the radiation into the IF band, which is coupled out of the cavity by the bias T. The bias T couples in the DC bias voltage necessary for operation of the SIS, but allows the IF to pass through. The next device is an isolator followed by a 3-7 GHz C-band cryogenic HEMT, located on a higher temperature stage of the 3-stage cryo-cooler. The IF signal is then fed out of the receiver dewar and into the backpack, which contains attenuators, more C-band amps, and a power splitter. The attenuators are used to roughly set appropriate power levels for the amps, both in the backpack and later on in the channelizer. The outputs of the power splitter are fed to a detector, whose DC output is monitored as a diagnostic, and to the IF processor, which eventually performs the correlation.

3.2.3 Phase Lock

It is imperative to maintain the relative phase of the LO signals on each SIS on very long (\sim hour) timescales in order to measure signal on the sky. The conceptually simplest scheme to do this is to use a high-power LO and then power split into four separate long-run waveguides to feed each mixer.

There are at least two problems with the simplest scheme. LO power is fed to the SIS's via rectangular waveguide. The attenuation of waveguide at 145 GHz is large, ~ 10 dB/m. The LO must therefore be physically near the SIS. In terms of flexibility in placement of the receivers, the best place for the LO is inside the backpacks.

The other problem is temperature stability. The wavelength of D-band radiation is approximately 2 mm (smaller in waveguide). The approximate length of the waveguide run to the dewar would be 1 m. Given that the thermal expansion coefficient of silver³ is $18.9 \cdot 10^{-6} \text{ K}^{-1}$, the expansion of one 1-m length of waveguide will differ by $\frac{1}{20}\lambda$ from another 1-m section of waveguide if the change in temperatures of the two waveguides differ by 5 K.

Both of these problems are solved with the Phased Locked Loop (PLL). Again, the details of this circuit are available in Dorwart (2002), Wesley (2000), and Dumont (2001). In the most basic terms, the PLL locks a high frequency slave oscillator

³the waveguide is actually made of *coin* silver.

to a lower frequency master oscillator. It does so by comparing a reference phase from the master to the derived phase from the slave. In the MINT scheme, the master is a 100 MHz oscillator and the slaves are 145.1 GHz. The MINT scheme also requires an intermediary 12.1 GHz oscillator to derive the comparison phase, which is also phase-locked to the 100 MHz oscillator. This design allows the placement of the 150 GHz LO's in the backpacks, with 100 MHz and 12.1 GHz signal cables run to each receiver. This solves the waveguide attenuation problem and decreases the susceptibility to phase variations due to differential length changes by roughly a factor of 10.

Phase Switch

The PLL makes possible the critical technique used to suppress false correlations from instrumental offsets. A phase switch implemented at the LO periodically changes the relative phase of the reference signal by 180° . Signals downconverted at the SIS mixer will retain this 180° phase shift, whereas signals added after the mixer will not. Demodulating the phase-shifted signal will retain the signals before the SIS, but suppress the signals that come in after the SIS.

For a two-element interferometer, the phase switch can be a simple half duty square wave. With more elements, a more complicated switching cycle must be used, which ensures that all baselines have oscillating relative phase. The use of the phase switch effectively suppresses the crosstalk in all of the IF processing systems, particularly in the channelizer and correlator. The relative phase of a real sky signal will modulate after the phase switch. Crosstalk induced after the phase switch appears as a common-mode signal between two receivers, which does not have a modulating phase. After demodulation the real sky signals have steady phase and the crosstalk signals have modulated phase, which when accumulated, will average to zero. The phase switch is also effective in removing $1/f$ noise in the IF amplifiers, as all signals that vary on timescales slower than the phase switch period are also suppressed.

The phase switch is also critical in sideband separating. A slow 90° phase modulation cycle shifts the relative phase of the LO on 0.5 sec timescales. The switch is done in between accumulation cycles. The 90° switch exchanges the real part of the

visibility with the imaginary. As mentioned in Section 2.2.2, a double sideband interferometer is not sensitive to the imaginary part of the sky, so this switch allows the imaginary part to be measured. At the same time, this 90° phase switch makes possible a sideband-separation procedure, discussed further in Section 7.1.2 and described in detail in Dorwart (2002).

3.2.4 Channelizer and Correlator

After leaving the RBE, the 4-6 GHz (C-band) signals enter the channelizer before being passed to the correlator. These devices are the main instrumental focus of this thesis.

The channelizer is a monolithic integrated microwave device. Its primary function is to condition the signals for the correlator. It has the dual capabilities of a filterbank and downconverter. The receivers output a wide signal, ranging from approximately 3 to 7 GHz. The correlator, however, can only accept a 0-500 MHz signal. The filterbank parses the incoming signal into 4 bands or channels, each 500 MHz in width with the first starting at 4 GHz. All signal below 4 GHz and above 6 GHz is rejected. This is where the band edges of the entire interferometer are defined, resulting in an overall IF bandwidth of 2 GHz. The next stage of the channelizer is the downconverter, which mixes the channels down to baseband (0-500 MHz). The last element is an amplification stage, which brings the power level up to levels needed by the correlator. A complete description of the channelizer appears in Chapter 4.

The correlator is a digital signal processing device which replaces the detector in a conventional microwave telescope and analog correlators in other interferometers. Its task is to calculate the correlation function (see Equation 5.1) between the signals in identical frequency bands of two separate receivers. The theory behind digital correlators is presented in Chapter 5 and the implementation in hardware and software is in Chapter 6.

Channelizer

The field of microwave engineering has evolved to incorporate Computer Aided Design (CAD) in a circuit analysis-oriented approach. Sophisticated solutions to Maxwell's equations have given way to analytical approximations and numerical simulations. While a field theory approach is still required at frequencies of >100 GHz, the majority of modern microwave engineering is concerned with designing planar devices and integrating them into monolithic circuits.

The channelizer is a monolithic circuit with drop-in components. It allows the use of an exploding market of surface-mountable microwave devices in the <10 GHz regime. Among the benefits of monolithic circuits are the elimination of the inter-device connectors, an increase in phase control and stability, and a reduced cost per device for large production runs.

The design and testing of the channelizer involved many people. Special thanks goes to Andrew Harris who introduced us to microwave circuit techniques through the design of the WASP correlator (Harris et al. 1998). Joseph Fowler and Zigmund Kermish helped out with the measurement of the performance. None of this would have been possible without the generous academic donation of the design software from Agilent.

4.1 General Layout

The channelizer conditions signals from the receivers so that the digitizers can read them. The job of the channelizer is to split the IF signal into four 500 MHz bands,

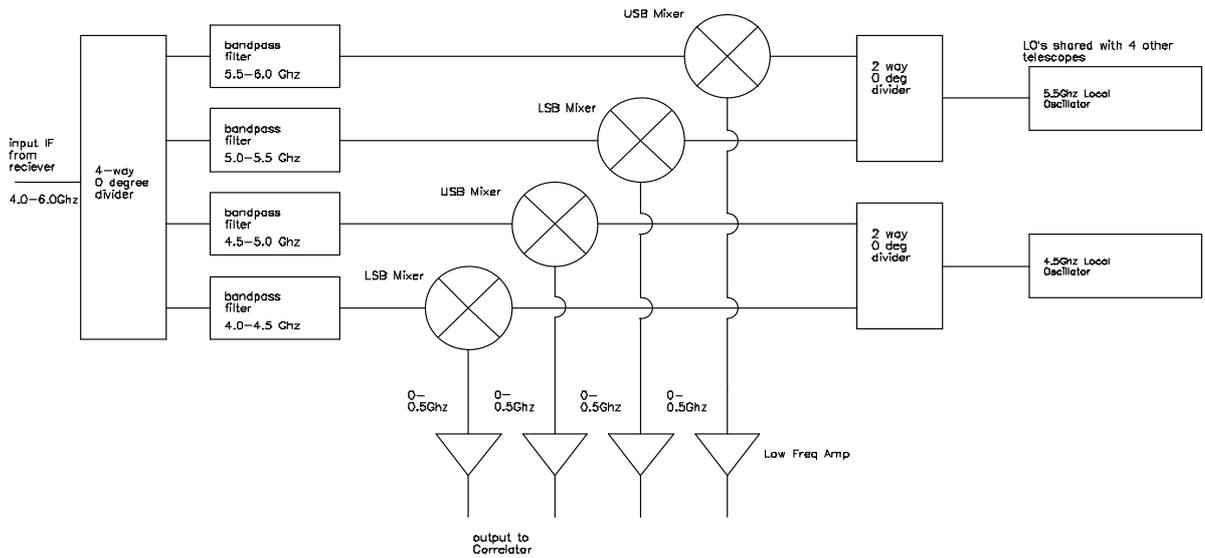


Figure 4.1: Channelizer Layout.

then downconvert each band to baseband. A schematic of the design appears in Figure 4.1.

The frequency splitting is accomplished using a 4-way power splitter followed by bandpass filters tuned to select the separate bands. The bandpass filters are followed by mixers, whose LO ports are connected to the output of a power splitter, which is fed by one of two LO's. The LO's can be shared by adjacent bands if one mixer is used in Upper Side Band (USB) and the other is used in Lower Side Band (LSB). The mixers are followed by amplifiers to bring the signal level expected by the digitizers.

4.2 Integrated Microwave Circuits

It is possible to assemble a channelizer from discrete components, as was often done in the past. Each component is connected to the next using SMA barrels. This not a cost-effective approach and it also involves many connectors, each with a finite failure rate. The alternative is to construct an integrated microwave circuit.

Microwave circuit boards have many of the same advantages of integrated electronic circuit boards such as smaller size and high reproducibility. The higher frequency, however, creates new challenges. All physical tolerances scale with frequency,

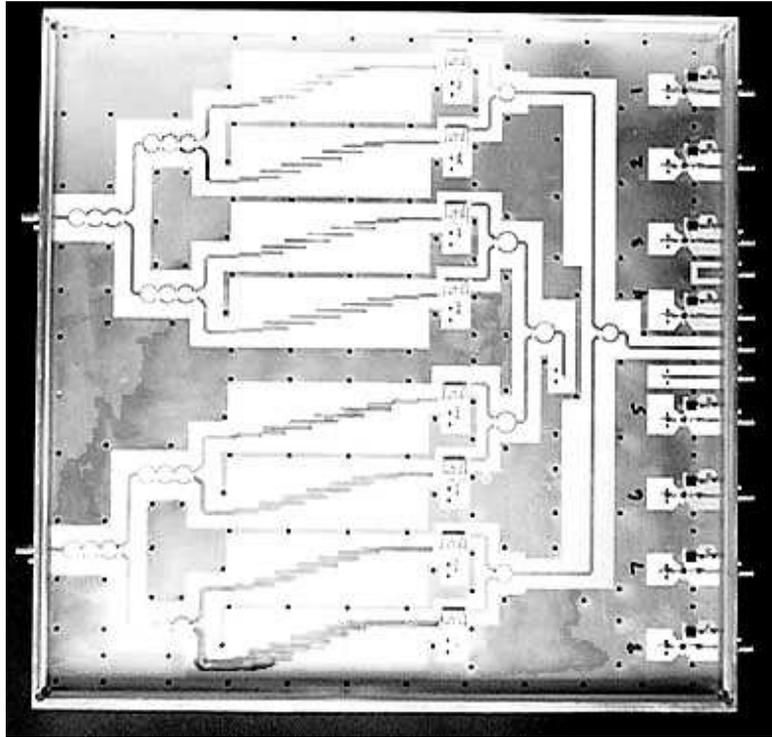


Figure 4.2: The MINT channelizer. The board size is 13.1"×13.1". Two full channelizers fit on one board. Signals enter from the left side, and leave on the right. LO signal is sent in on the right, and split between the mixers. There are only 13 connectors on this board and only two such boards are needed for MINT.

and, hence, more sophisticated and costly board fabrication techniques must be used. The tolerances on the board substrate are also tightened, most notably in terms of thickness and dielectric uniformity. Another issue is that the size of normal electric components (lumped elements) are comparable to the microwave wavelengths and therefore do not have uniform properties over broad frequency ranges. In many cases, this forces the microwave engineer to use distributed elements to reproduce capacitors and inductors. This last restriction may in many ways be seen as a benefit because many components can be constructed into the circuit board itself. The channelizer incorporates both the power splitters and the bandpass filters into the layout of the circuit board. In addition, the mixers and the amplifiers come in small surface-mount packages, which allows for a truly connector-less design.

A picture of the MINT channelizer board appears in Figure 4.2. The compactness of microwave circuits allows the placement of two channelizers on one board. The

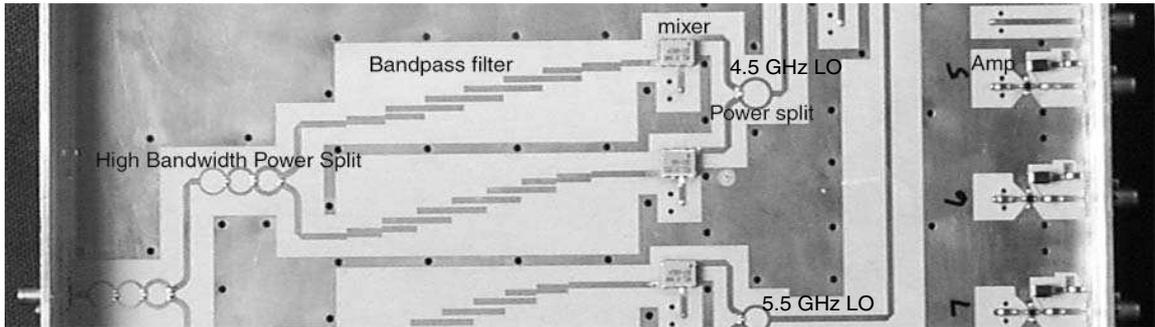


Figure 4.3: A picture focusing on one downconversion chain of the channelizer. In this blown up view, the components can be seen in greater detail. After leaving the IF port on the mixer, the signals traverse a short distance on the top layer before sinking through the board and under the LO lines, and then reemerge at the far right hand side before entering the amplifiers.

MINT interferometer requires two boards to process the IF signals from four receivers.

Figure 4.3 focuses on a single downconversion chain. The only part of the signal path that is not evident from the photograph is the transmission line which connects the IF¹ output of the mixer and the input of the amplifier. After leaving the mixer, the signal makes a transition underneath the board through a “via.” It then reemerges on the far right of the board immediately before the input of the amplifier².

The main planar circuit elements that make up the channelizer are the power splitters, filters, transmission lines, right angles, and vias. The surface mount elements include the mixer and amplifier. The design process for the filter and power splitter is outlined in the sections below.

4.3 General Board Properties

The circuit board is fabricated from a Teflon-based dielectric substrate sandwiched between two thin copper layers. The channelizer material is known by the industry trademark name Duroid 6002, which has a dielectric constant $\epsilon = 2.94$, considered

¹The IF (Intermediate Frequency) of this mixer should not be confused with the IF of the SIS mixer. There are two stages of downconversion on MINT, leading to two different IF's. The first IF is from 4-6 GHz, which becomes the RF input to the next level of downconversion in the channelizer. The channelizer IF is 0-500MHz or baseband.

²This solution was the best available at time of the design. In the months following the fabrication of the channelizer a better surface-mount solution presented itself. See Appendix A.3.

medium-low. The thickness is 0.076 cm (30 mils). For an even smaller design, a higher dielectric like $\epsilon \approx 10$ could be used. Our design, however, was constrained by the minimum possible gap, defined by the board fabricator, Filtran Microcircuits³, to be 76.2 μm (3 mil)⁴.

Multiple sheets of Duroid may be laminated together to form a multi-layer board. The MINT design, however, uses the simplest possible structure type known as microstrip. Microstrip structures are directly etched from one of the copper layers. This layer is the component or top layer of the board. The bottom copper layer is left mostly intact and serves as the current return or “ground” plane. The ground layer is breached for *vias* which bring the ground to the component layer. A via is a hole drilled through the substrate and then plated through with copper. A special *jumper* is also needed, discussed in Appendix A.3, which is used to cross signal lines.

Of the possible structure types, microstrip structures are the easiest to fabricate. The artwork is usually either photo- or laser-etched. There is no additional lamination required. Microstrip has the added benefit that surface mount components can be directly connected to the microstrip.

4.4 Circuit Design

The MINT channelizer design starts with analytical models that produce physical parameters that can be verified and tweaked in simulation. Microwave design is very computer-intensive, with certain simulations running for three days or more on a late 1990’s desktop-class microcomputer . We used the industry standard ADS (Advanced Design System) software donated for educational purposes by Agilent. The package is a full design suite, which incorporates design entry, simulation, and board layout into a single environment.

³Filtran Microcircuits is based in Ottawa, Canada see <http://www.filtranmicro.com>

⁴Note: 3 mils is merely the limitation for standard fabrication techniques. Gaps of less than 1 mil are possible at extra cost.

4.5 Transmission Lines

The most common structure on any microwave board is the transmission line. The most important number for a transmission line is its characteristic impedance. For a lossless line, the impedance is completely real, meaning the current and the voltage at every point on the line are in phase with each other. The impedance is then given simply by $\sqrt{L/C}$, where L and C are the inductance and capacitance per unit length. There are many different types of transmission lines, and each has different relationships among physical parameters and the characteristic impedance. The channelizer uses predominantly microstrip lines with the exception of the jumper, which is made from stripline (see Appendix A.3). As most devices have 50 Ω ports, most transmission lines have a characteristic impedance of $Z_0 = 50 \Omega$.

In the simplest approximation, there are only 2 parameters involved in the equations that describe microstrip transmission lines (see Pozar (1998) sec 3.8.), the dielectric constant and the width of the line divided by the depth of the dielectric, W/d . In general, the impedance of microstrip goes up for decreasing width at fixed depth and also goes up for decreasing ϵ . Therefore, for a fixed Z_0 a smaller trace can be used with either a thinner board or a higher dielectric constant. For $\epsilon = 2.94$ and $d = 0.076$ cm, a $W = 0.194$ cm (76.7 mil) line will have a characteristic impedance of $Z_0 = 50 \Omega$.

There are other parameters, such as conductor thickness, resistivity, and loss tangent of the dielectric. These all serve to modify the impedance equations and introduce loss. These effects are best analyzed by the simulation, and for the specific parameters of our substrate, all are small. For instance, the loss tangent of Duroid 6002 is $\tan \delta = .0012$ at 10 GHz, which leads to only -0.2 dB of loss over 10 inches of transmission line. Adding in the finite conductivity of copper, $\sigma = 5.9 \cdot 10^7 (\Omega\text{m})^{-1}$, the loss increases by an additional -0.2 dB. The net loss should be 0.4 dB across the board.

4.6 Filter Design

The job of the channelizer filters is to define each 500 MHz wide sub-band. The specific filter type is a 6-pole Tchebyscheff with .01 dB ripple within the passband.

At radio and lower frequencies, filters can be built from discrete, lumped-element capacitors and inductors. In the microwave regime, capacitors and inductors are not available in wide values or frequency ranges. It is possible to design distributed element capacitors from circuit gaps and inductors from transmission line spirals, but neither act purely as one or the other at high frequency. Instead of incorporating either capacitors or inductors into a design, certain circuit elements are used where the reactive impedance can be calculated. $\lambda/8$ transmission line stubs, for instance, can be used either in series or shunted. A low-pass filter can be constructed completely from these elements. To make a bandpass filter, a more complicated circuit element is needed. The coupled transmission line consists of two $\lambda/4$ transmission line sections of equal width placed closely together in a parallel configuration. These sections are cascaded to form the filter.

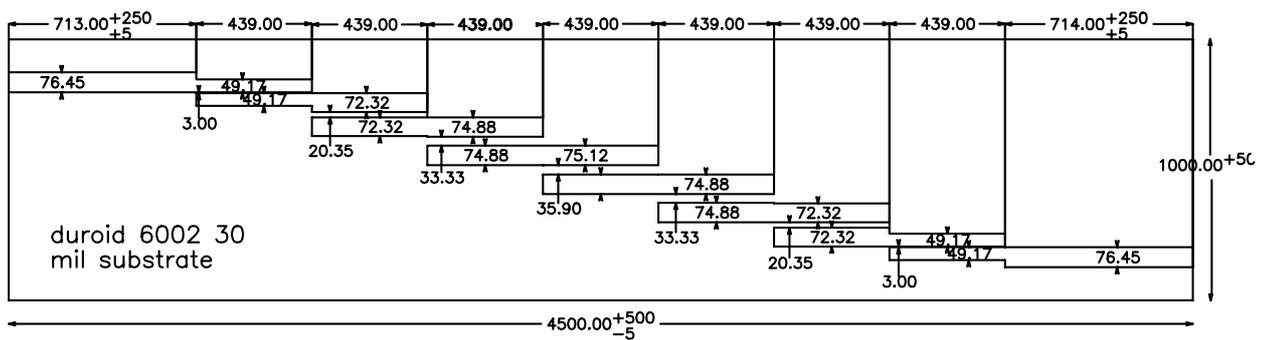
The filter design process starts with designing a filter based on lumped-elements. Then a series of conversions are used to derive the needed impedances, which are then used to specify the physical parameters of the circuit elements.

4.6.1 Filter Synthesis

With the band edges and ripple chosen as inputs, it is possible to derive the specific physical dimensions of the coupled transmission line sections. The general procedure follows very closely the design method laid out in Matthaei et al. (1980) Section 8.09. The only difference is that we used equations derived in Gupta et al. (1996) Section 8.5.1 instead of graphical solutions presented in Matthaei et al. (1980) to derive the impedance of microstrip lines. The procedure is outlined in Appendix A.2.

The output of the equations are the length of each coupled microstrip pair, the separation between lines and the width of the lines. Figure 4.4 is an outline drawing of the artwork for a 4.0–4.5 GHz filter. The device is symmetric about the middle section, so for a 7-section device, there are only 4 unique sections.

The coupled microstrip sections at the extreme ends have a gap of $76.2 \mu\text{m}$, the minimum set by the fabricator. The wider the relative bandwidth, the smaller this gap. This dimension constrained the entire board design, forcing the use of low ϵ and relatively thick material.



Attribute	Value
Customer	Princeton University
Contact	Huan Tran huantran@princeton.edu
Phone	609 258 5077
Fax	609 258 6853
Part Name	7 section Bandpass Filter 4.0–4.5 Ghz
Quote #	05440
Substrate	Duroid 6002 30 mil dielectric constant=2.94
Metalization	1/4 ounce copper front and back side ground plane
Coating	gold plating front and back 5–15 microinches
Units	inches
Tolerance	Standard easy (.25 mil)

Note:
Dimensions are slightly different from initial Quote

We will be doing the final machining here at Princeton
The overall size of the board therefore, should not be less than 4.5 X 1.0
inches, but can exceed that by any convenient amount.

Figure 4.4: The filter dimensions sent to the board fabricator. These dimensions were slightly adjusted relative to the parameters derived from Gupta et al. (1996). The design procedure is discussed in Section 4.6.3.

4.6.2 Linear Simulation

In principle, the physical parameters produced by the design equations may be used directly in the fabrication of the bandpass filter. Unfortunately, the design equations are necessarily approximations that do not faithfully reproduce the true behavior of the coupled microstrip elements, particularly at higher frequencies. Simulation is therefore necessary to verify and modify the design.

There are two stages of simulation. The first is called linear simulation, wherein the analytical design equations derived in Gupta et al. (1996) are used to calculate the properties of each section. The sections are then strung together to calculate the properties of the *network*⁵. The parameters are entered into the simulation via a schematic window as in Figure 4.5. Circuit elements are selected from a library and the parameters are entered into property fields.

In the case of coupled microstrip lines, the simulator uses the same equations

⁵“Strung together” means matrix multiplication of the transmission or *ABCD* matrices. The *ABCD* matrix relates the voltage and current at the output port to the input port of a two-port device. Calculating the *ABCD* for a *network* made of two elements is a matter of multiplying the matrices of the two elements. See Pozar (1998) Section 4.4 for a full description of transmission matrices.

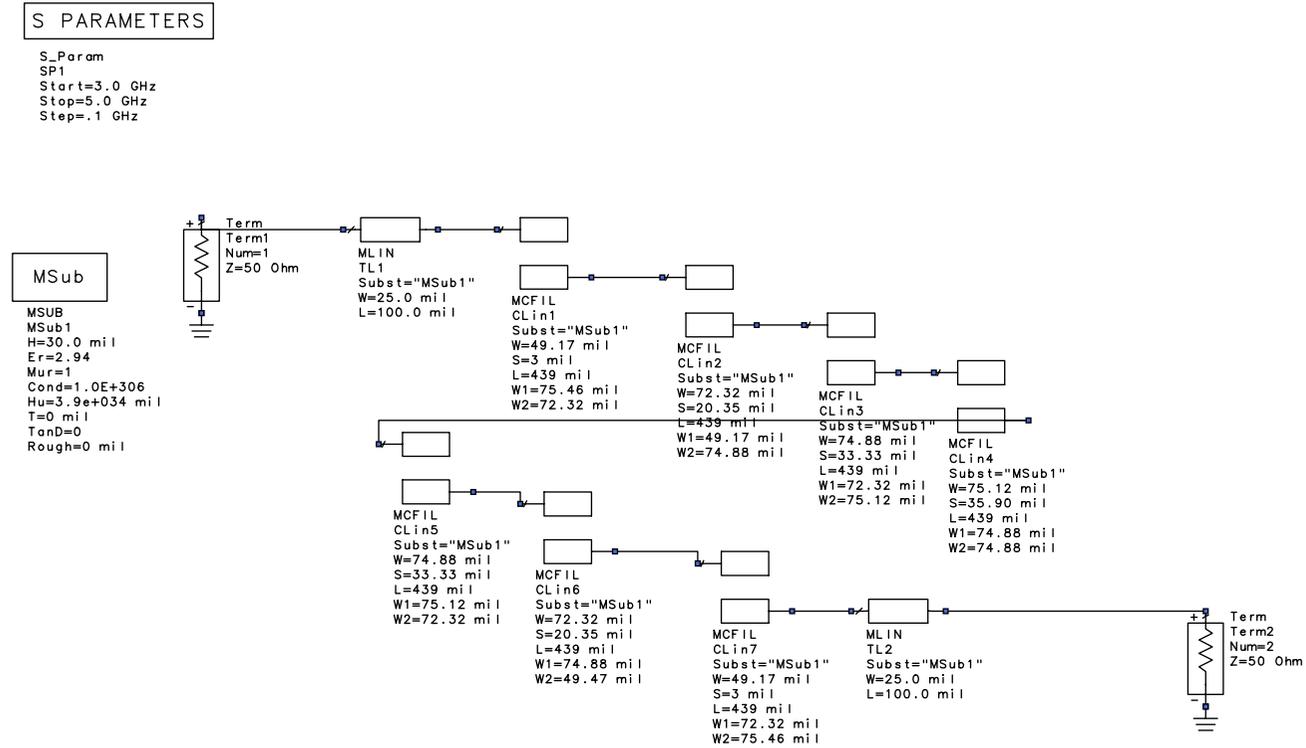


Figure 4.5: A schematic entry window in from Advanced Design System. The circuit consists of a transmission line section followed by 7 coupled-line elements, then another transmission line. There are what look like 2 resistors on either end, which are the test “ports” used for simulation. The “Msub” box has the entries for the microstrip substrate properties, and the “S parameters” box evokes the simulator, which has the frequency ranges for the simulation.

from Gupta et al. (1996) that we used to derive the physical parameters. The linear simulation serves as a verification step. It ensures that all elements work together as expected. One mode of linear simulation produces the S-scattering matrix⁶. The linear results tend to be accurate to within 1-5% for a certain range of parameters. The primary benefits of a linear simulation are ease and speed. It is not, however, wise to fabricate based on linear simulations alone.

4.6.3 Momentum Simulation

The Momentum Simulator uses the actual physical layout of the device to perform the simulation. ADS conveniently translates the parameters entered into the circuit diagram into real shapes in the layout side of the program. The user then pieces together the parts to make the device.

The simulator first defines a mesh for the simulation. Here, the user has some input as to how fine the mesh is laid out. The default resolution is 30 cells per wavelength. This, along with the maximum frequency of the simulation determines the mesh. *Momentum* will draw rectangles and triangles that approximate the physical layout, which gives an idea of the minimum feature that affects the results.

Momentum uses the “method of moments” (See Agilent (2001)) to solve for currents and then the fields are inferred from the currents. The distributed element is broken into discrete cells. Each cell is replaced with a capacitor to ground and an inductor to each neighboring cell. The major task of the simulator is to solve for the currents by inverting large matrices. This is a very computer intensive process, with simulations of the prototype filter taking 8 hours or more on a ~ 60 Mflop desktop computer. The momentum simulation is more accurate than linear simulation because the accuracy of the simulation depends on the mesh size and not on analytical approximations.

The results of the momentum simulation were generally different from the results of the linear simulation. The chosen criteria for a good filter performance was based on the location of the 3 dB band edges. If the simulated location of the 3 dB points

⁶The S_{ij} scattering matrix relates the voltage amplitude incident on port j to the outgoing voltage on port i . It is defined as: $S_{ij} = \frac{V_i^-}{V_j^-}$. S_{11} , for instance, measures the reflected amplitude from port 1 while S_{12} measures the transmitted amplitude from port 1 to 2. The parameters of the S matrix are measured by a Network analyzer. See Pozar (1998) section 4.3 for more details.

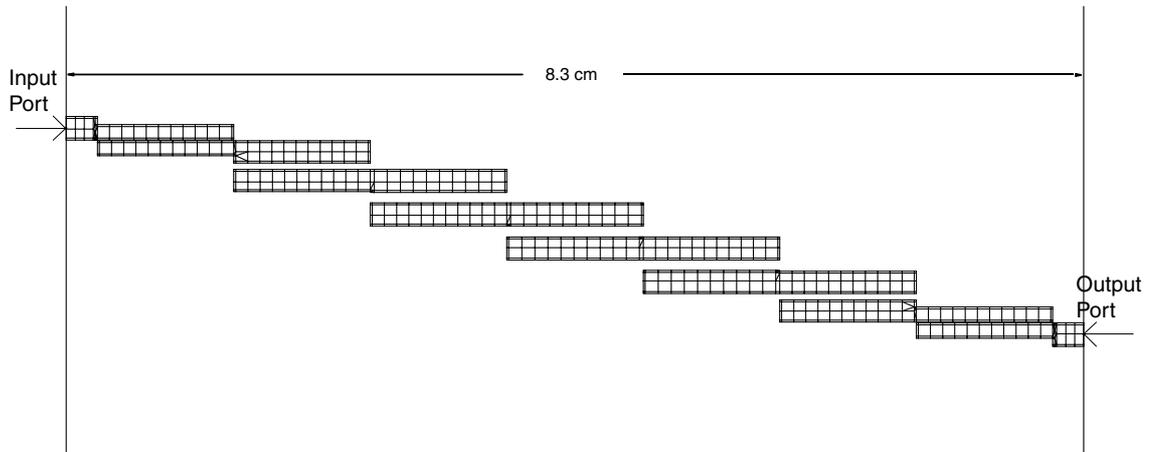


Figure 4.6: Meshing of the filter by momentum simulator. The arrows indicate where the test ports are placed.

differed from the desired location, the initial parameters that went into the analytical part of filter synthesis were modified to bring the momentum simulation closer. Eventually, the parameters and simulations converged to the desired performance. Figure 4.7 shows the simulated S parameters for a prototype filter. Momentum simulations also differ from linear simulations in some non-trivial ways, such as depth of the peaks in the return loss.

4.7 Power Splitter Design

The other major distributed element device is the power splitter. The MINT power splitters are based on the Wilkinson power splitter design, a lossy 3-port device. There are two types of power splitters used, a narrowband and a broadband version. The narrowband versions split the monochromatic LO signals that feed the mixers. Since there are two different LO frequencies, 4.5 and 5.5 GHz, there are two different narrowband power splitters. The broadband splitter splits the incoming 2–6 GHz signal. As it will be seen, the narrowband versions have just enough bandwidth to be used as the broadband splitter. It is however, more prudent to over-design the bandwidth to account for simulation inaccuracy.

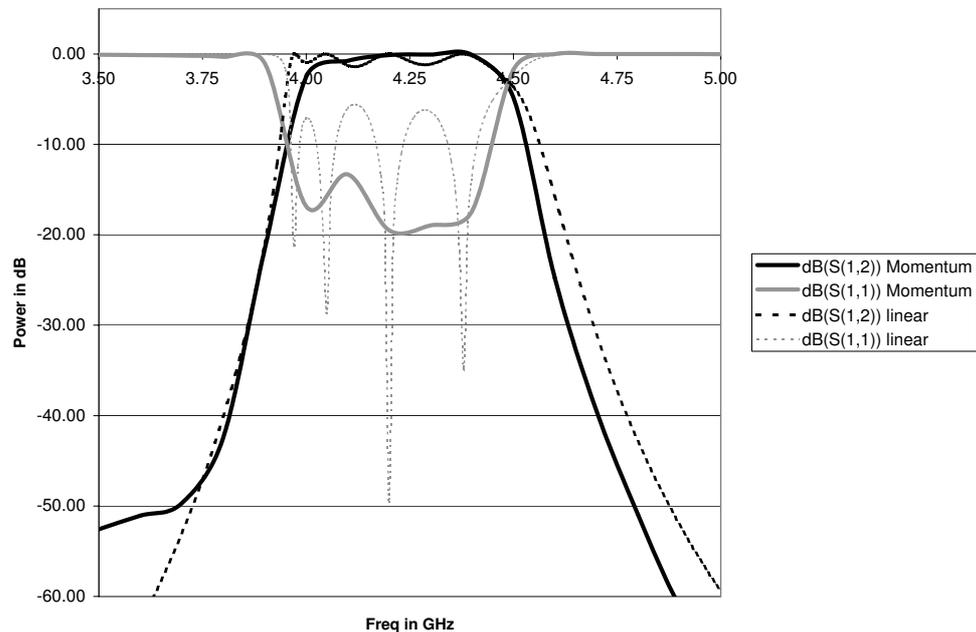


Figure 4.7: Results of linear and momentum simulation, using the modified parameters in Figure 4.5.

4.7.1 Narrow Band Power Splitter

The theory of operation for the narrow-band version is described in many textbooks, for example, Pozar (1998), sec 7.3. The narrow-band version consists of a T-junction power-divider feeding two quarter-wave transformers which are connected via a resistor at the output port end. To ensure that all of the ports are matched to $50\ \Omega$, the quarter wave sections are $\sqrt{2} \cdot 50\ \Omega$ and the resistor is $100\ \Omega$. The important feature of this type of power divider is that the output ports are isolated; that is, power input into one of the output ports is not seen on the other output port. The disadvantage is that since there is a resistor, the device is lossy. Half of the power would be lost into the resistor were it driven at the output port. When used as a power combiner, only half of the power makes it through. The channelizer uses the device as a splitter, where all of the power is split evenly between the output ports. The performance of the resistor also tends to limit the maximum frequencies of possible designs.

Power splitters can have many different geometries. Circular sections were chosen because they are the most gentle way to bend the signals through the 1/4-wave

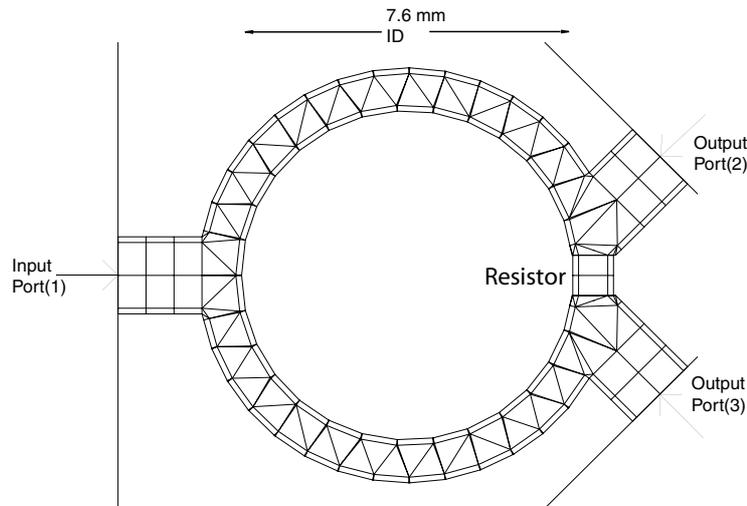


Figure 4.8: Meshing of 4.5 GHz power splitter. The ports are labelled “input” and “output” in the case of a device that acts as a splitter. When the device is used as a combiner, the ports are reversed. The resistor is also meshed and labelled above. For a description of the simulation process, see Section 4.6.3.

transformers. Furthermore, the output ports must optimally diverge. Too slow a separation causes excess coupling between the output ports and too fast a separation causes poor transmission.

The transition section that accommodates the resistor pad and output port is of comparable physical size to the $1/4$ -wave transformers. Any linear simulation that does not take this into account will differ greatly from a momentum simulation. The meshing for the power splitter appears in Figure 4.8.

The narrow band power splitters are used to divide the channelizer LO power between four mixers on a single board. Since there are two different LO frequencies, 4.5 and 5.5 GHz, there should be two different sized designs, each optimized for the appropriate frequency. The 5.5 GHz filter is slightly smaller because the $1/4$ wave circular sections are shorter.

The simulated S-parameters of the 4.5 GHz splitter are shown in Figure 4.9. The design was considered satisfactory when the dip in the isolation between output ports occurred at the design frequency. From the simulation, we see that the useful bandwidth of this device is about 2 GHz. The results of the simulation for the 5.5GHz device are also shown in Figure 4.10.

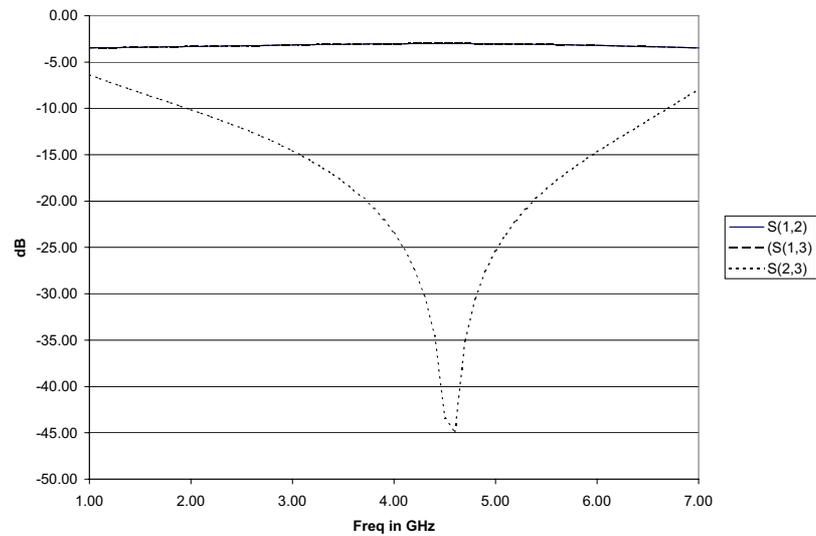


Figure 4.9: Momentum simulation results for the 4.5 GHz power splitter. The throughput is $S(1,2)$ and $S(1,3)$, which lie on top of each other, verifying that the momentum simulation is symmetric. Since the T-Junction is a broadband device, $S(1,2)$ and $S(1,3)$ are nearly flat at -3 dB. A lower $S(2,3)$ means that in the event of the splitter acting in reverse as a combiner, there is less crosstalk between the two inputs. The characteristic dip in $S(2,3)$ comes from the departure of the $1/4$ -wave transformers away from the design frequency. The maximum dip was fine-tuned to be near 4.5 GHz. The edges of the band for a power splitter are defined to be where $S(2,3) \sim 20$ dB. For this splitter, the bandpass is about 3.5-5.5 GHz.

The momentum simulation includes an attempt at simulating the thin-film resistor. It appears in Figure 4.8 as the small bridge that connects the output parts. Thin film chip resistors come with resistive material on top of a ceramic substrate with conductor tabs on the ends of the chip to connect to the substrate. The tabs lead to some stray inductance which is somewhat difficult to model. A real resistor can be made to approximate the 2-D simulation by “flip” mounting. The resistive material is thus brought closer to the circuit by mounting it upside-down. The resistors in the MINT channelizer are flip mounted.

4.7.2 Broad Band Power Splitter

With many microwave devices, constructing a broadband version of a narrow-band version is just a matter of cascading sections. Such is the case for Wilkinson power

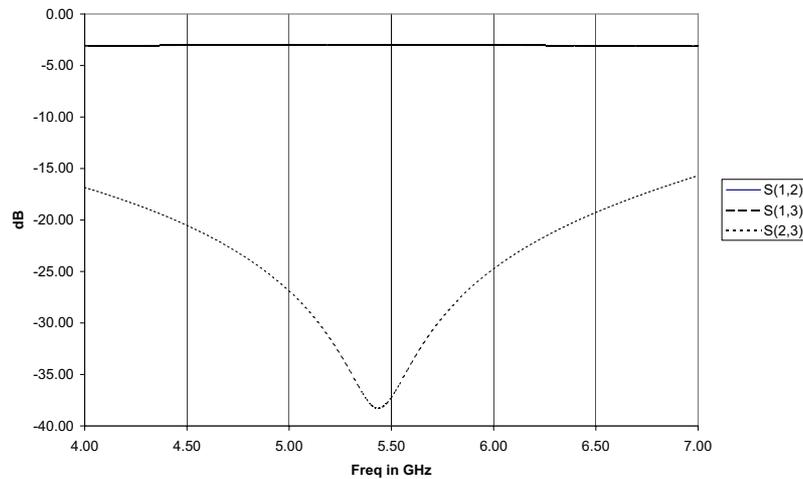


Figure 4.10: The S parameters for the 5.5 GHz LO power splitter. See the caption for Figure 4.9 for details.

Section	Transmission line Impedance Ω	Spec'd Resistor Value Ω	Actual Value $\pm 1\%$ Ω
1	86.98	107.18	105
2	70.71	211.46	210
3	57.485	400	402

Table 4.1: The parameters from Li et al. (1994) used to design the broadband power splitter. The actual resistors were defined by the manufacturer. All values are in Ohms.

splitters. The only difference is that each section must have 1/4-wave transformers of different characteristic impedance, and each resistor must be a different value. In practice, these values are arrived at by simulation. Li et al. (1994) have performed this arduous task and tabulated the optimum parameters. They were able to build a very broadband splitter, from 2-18 GHz with 7 sections. Since we were not interested in such broad bandwidth, we could settle for the much smaller 3-section splitter. We used the values as listed in Li et al. (1994) and verified them with simulation.

Each section of the power splitter is cut from a perfect ring, similar to the narrow band splitters and in keeping with the design in Li et al. (1994). Most of the splitter was hand drawn, using layout-only functions instead of transforming circuit elements.

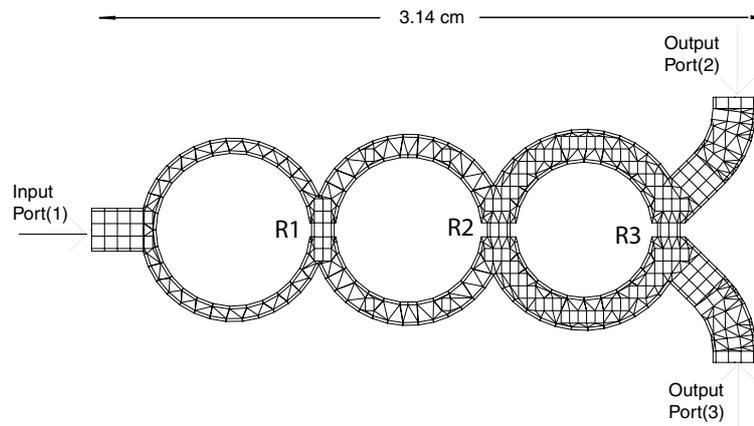


Figure 4.11: Meshing of broadband power splitter. The three resistors are labelled.

There is a gap of 25 mils cut in each section which serves as the mounting pad for the chip resistor. The simulated performance of this device appears in Figure 4.12.

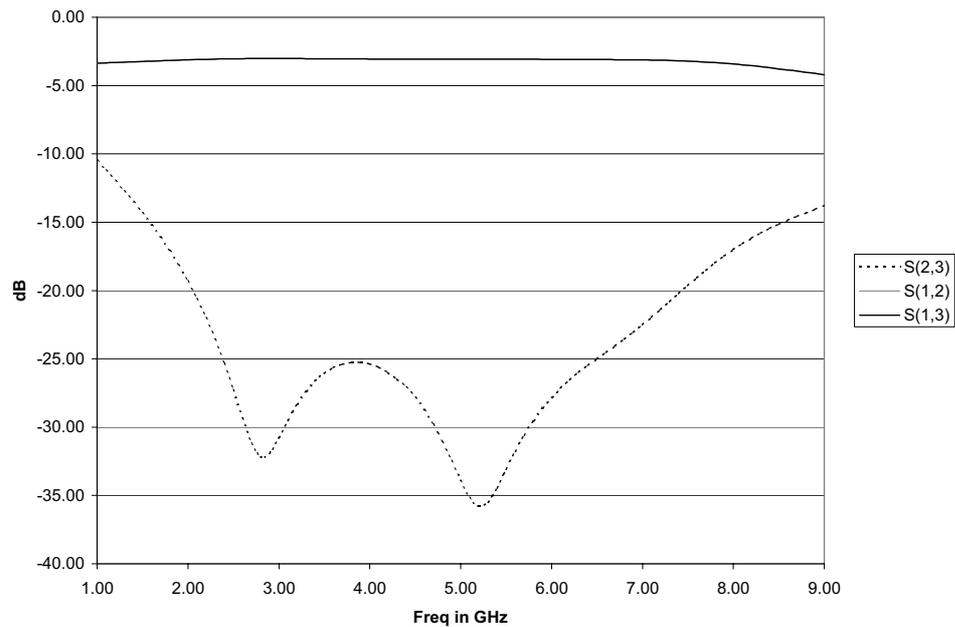


Figure 4.12: The S parameters for the broadband power splitter. A 3-section device should have 3 characteristic dips. The last dip is most likely suppressed due to poor performance of the chip resistor at high frequency. This device has $S_{23} < 20$ dB between 2GHz - 7.5 GHz.

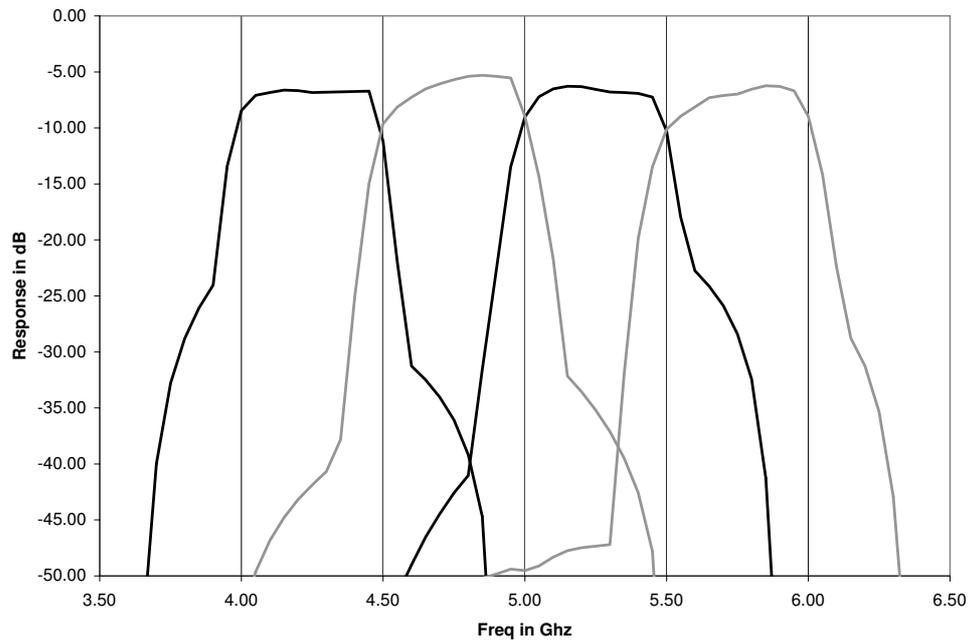


Figure 4.13: Bandpasses for the filterbank resulting from combining the momentum simulations for three cascaded broadband power splitters and four bandpass filters. The maximum level is -6dB, exactly what would be expected after splitting the incoming signal 4 ways.

4.8 Combined Simulation

ADS also allows the use of tabulated S-parameters of a circuit element in simulation. The tabulation may come from measurement on a network analyzer or in this case, from the momentum simulation. As a final test of system-level performance, the simulation results for all the filters and wide-band power splitters were combined in a large scale linear simulation. This gave the performance of the filterbank section of the channelizer. The results appear in Figure 4.13. The comparison of the simulations to actual performance appears in Section 4.10.1.

4.9 Assembly

4.9.1 Prototyping

There were two stages of prototyping before the final channelizer was built. The first stage involved fabrication of only one of the bandpass filters as a sanity check. After we were satisfied with the performance, the entire channelizer was fabricated with a provision to incorporate or omit the downconversion and amplification section. This was done so that we could test the filterbank section on the network analyzer⁷, the results of which appear in Section 4.10.1. The second prototype showed that the mixer chosen did not have optimal performance and was difficult to mount because it required wire bonding. The final channelizer design, appearing in Figure 4.2, incorporates an inexpensive surface mount mixer.

4.9.2 Enclosure

Design of the enclosure is an important aspect of microwave design. For the channelizer, it serves many purposes: protection, grounding, cooling, and microwave isolation. Since the Duroid substrate is flexible, the bottom part of the enclosure is used to flatten the board. It is very important to keep the board flat against this surface, as it provides the primary ground and any gaps will produce unwanted resonances. This can be accomplished by silver-bearing epoxy, but we used a lid to sandwich the substrate against the bottom, which was milled to 0.254 cm (0.1”) to allow it to conform to the lid.

The enclosure shields the rest of the telescope from the channelizer by acting like a Faraday cage. Each channel is also isolated from other channels with walls, which match ground traces on the board.

The critical dimensions on the enclosure are the separation of the circuit elements from the walls and from the top. Both tend to lower the impedance of transmission lines. Gupta et al. (1996) report that a height of top vs. substrate height of 5 gives negligible effects. For our design, this ratio is about 25.

The cavities also lead to unwanted resonances. The resonant frequency can be

⁷Having mixers in the signal path would make the network analyzer useless.

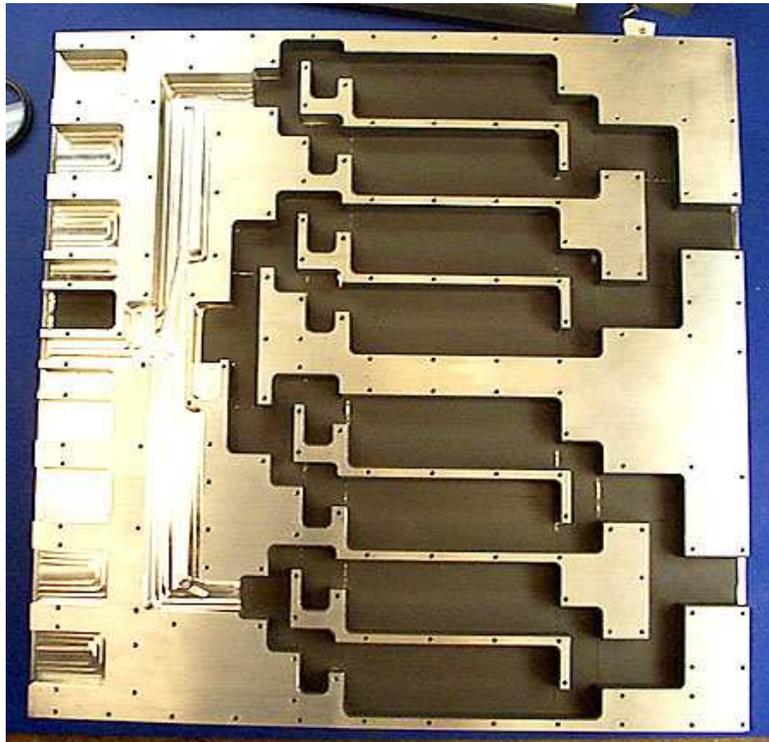


Figure 4.14: View of the enclosure lid for the channelizer. Both the lid and the box are made from solid pieces of aluminum. They are hollowed out on a CNC mill.

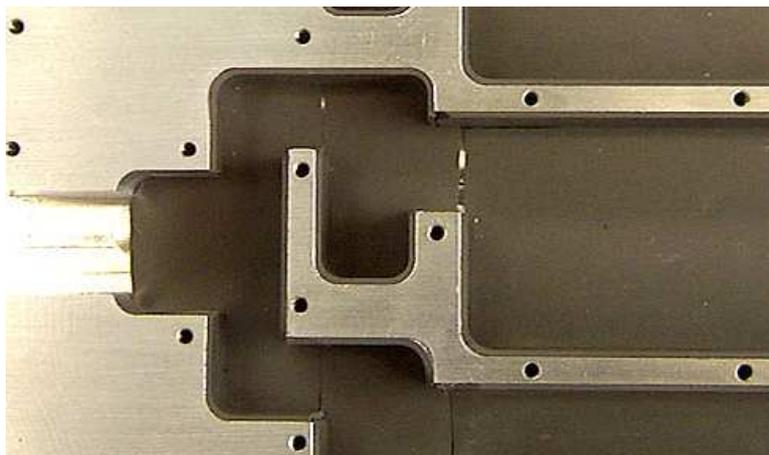


Figure 4.15: Detail of the channelizer lid to show the eccosorb lining. The lining was cut to shape and has an adhesive backing.

modeled, but it is usually difficult to do this reliably. To suppress the cavity resonances, all the walls are covered with Eccosorb, a microwave absorbing material. The Eccosorb comes in many different forms; the type used here consists of 0.8 mm rubber mats with adhesive on one side.

Fully assembled, one channelizer board with enclosure measures 2.54 cm (1.00") thick. Two channelizer boards inside enclosures are stacked to form the channelizer unit with IF outputs, LO inputs, and power coming into the front and RF coming in through the back. Two external power splitters with outputs 2.54 cm apart are directly connected to the LO inputs. The entire assembly is enclosed in foam and attached to a temperature-servoed plate.

4.10 Performance

The physical parameters of the channelizer appear in appendix A.1, including power requirements.

The aspects of the channelizer performance that are pertinent to interferometry are the overall bandpass, both absolute and in comparison to other channelizers, and the relative phase between outputs.

4.10.1 Bandpass

The bandpasses of the filterbank alone were measured on a network analyzer. They were done on a prototype board, but the results should apply to the final channelizer version because the power splitters and filters are identical. The results are shown in Figure 4.16 and are plotted on top of the simulations.

To test the bandpasses of the entire channelizer with the mixers and amplifiers in the signal chain, a slightly more complicated setup was used. A programmable sweeper fed a monochromatic signal (with frequency between 3.5 and 6.5 GHz) to the input of the channelizer. A 550 MHz low-pass filter followed by a power meter was placed at the desired output. Both the sweeper and power meter were controlled and read out via GPIB and LabView software. The sweeper output was varied in small steps, and a few measurements of the power at the output of the channelizer were taken and averaged.

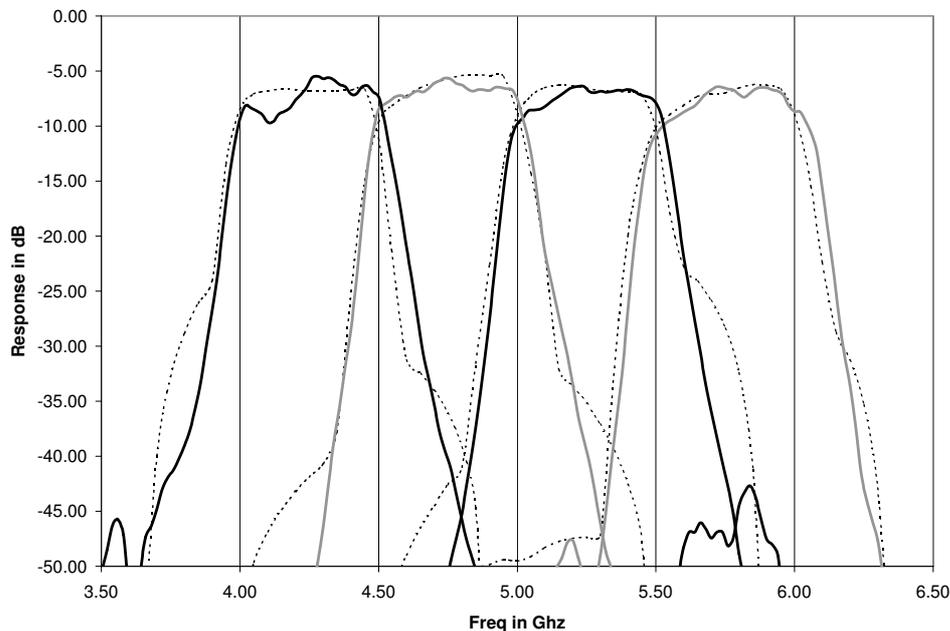


Figure 4.16: VNA measurements of prototype filterbank (solid lines) plotted on top of the simulation results shown in Figure 4.13 (dotted lines). The quantity shown is S_{12} , with all the other ports terminated. The VNA measurements were scaled up by 3dB to match the simulations. This extra loss is presumably due to losses in the microstrip which were not simulated or to losses in the connectors. As can be seen, there is generally good agreement between simulations and measurements.

Figure 4.17 shows the results of the swept measurements. The measurements are normalized to the input and they imply about 12 dB of overall gain in the channelizer. The performance is degraded relative to the filterbank alone, possibly due to impedance mismatch at the input to the mixer. If we were more concerned with this performance, it would be necessary to carry out another round of prototyping to design an impedance matching network for the mixer. Alternatively, the gain of the amplifier could be increased and a matched attenuator could be placed at the input to the mixer.

From these measurements we may compute a quantity useful when estimating the sensitivity of a system known as the noise bandwidth(Kraus 1966):

$$\Delta\nu_{eff} = \frac{(\int f(\nu)d\nu)^2}{\int f^2(\nu)d\nu}, \quad (4.1)$$

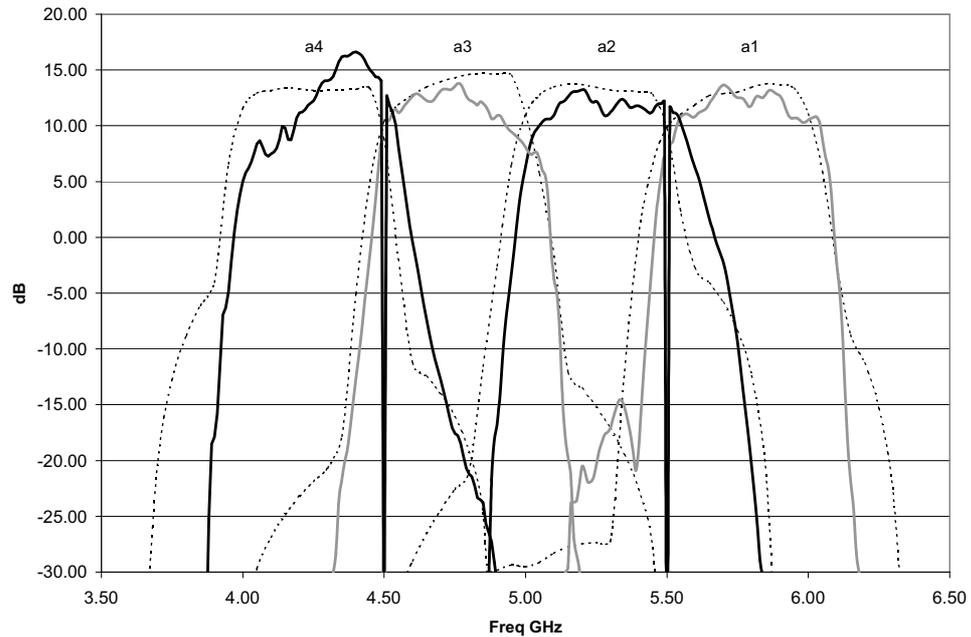


Figure 4.17: Bandpasses through the entire channelizer. These results are typical for all channelizer halves, but in this case represent the response for channelizer A. The channels are numbered from high frequency to low. The solid lines are the power measured at the output of the channelizer normalized to the input. The positive values in band show that the channelizer has an overall gain of about 12 dB. The dotted lines are the filterbank simulations. There are large dips at the LO frequencies because the high pass bypass capacitor on the output of the channelizer filters out DC.

where $f(\nu)$ is the power response as shown in Figure 4.17. Table 4.2 lists the computed effective bandwidths.

Another important figure of merit is the degree to which identical channels differ in bandpass. Figure 4.18 plots the differences in all channels relative to channelizer A. The differences are approximately 1-3 dB.

4.10.2 Phase

A more important figure of merit is the relative output phase of the four channelizer. This quantity has a direct effect on how accurately the visibility can be measured and in general, if the differences are stable in time, must be calibrated out.

A noise injector and fast digital oscilloscope were used to measure the phase. A

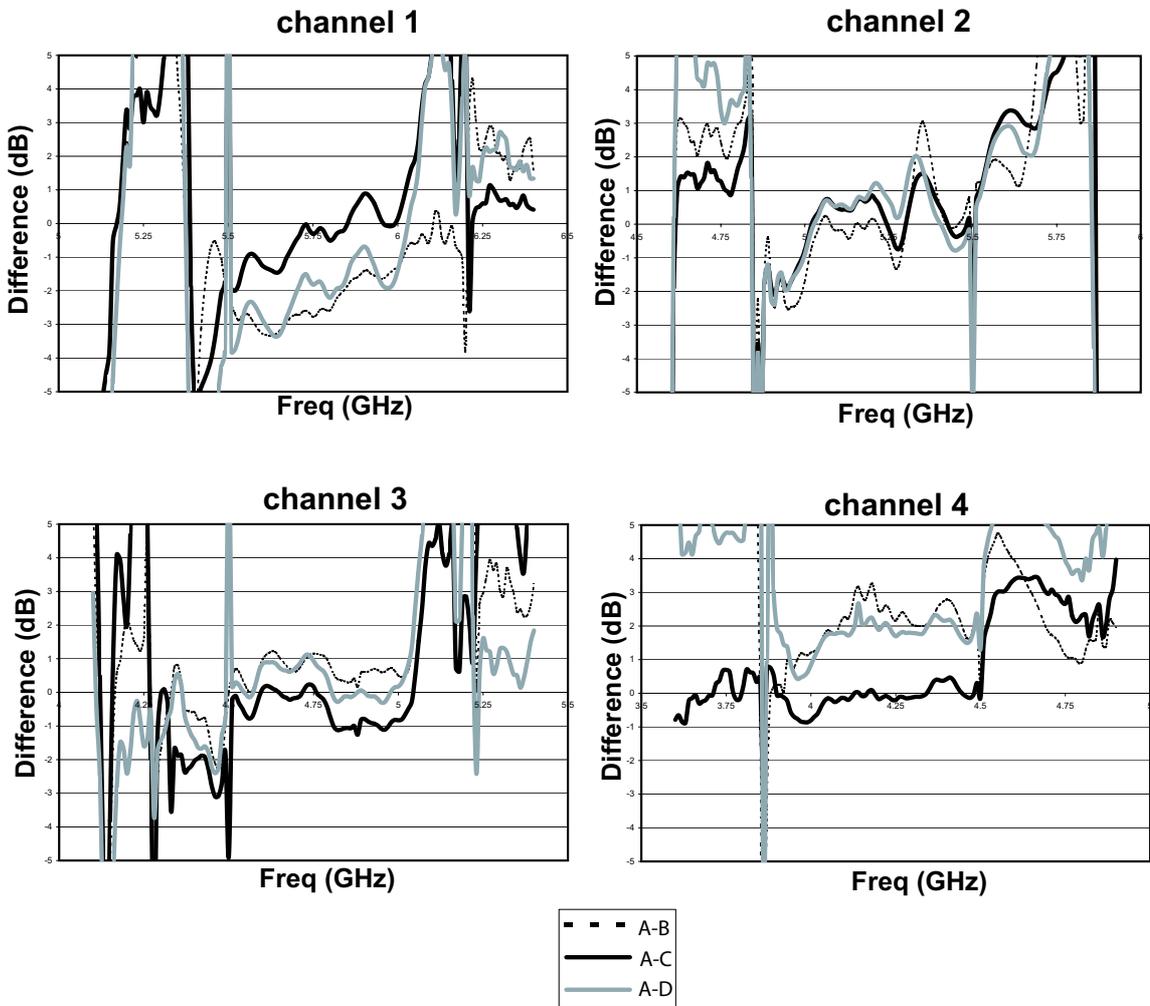


Figure 4.18: Bandpass difference between identical channels of the 4 separate channelizers. All differences are taken relative to channelizer A. If the phase roughness in channelizer A is larger than the rest of the channelizers, it will show up as a common component in these differences.

Band	$\Delta\nu_{eff}$
1	542 MHz
2	584 MHz
3	518 MHz
4	387 MHz

Table 4.2: The noise bandwidth as computed from Equation 4.1. The design target was 500 MHz. The large values of channel 1 and 2 mean that the band is wider than target whereas the small value for channel 4 is from both a narrower than target combined with a very non-flat bandpass. The values should be true for all channelizers, as the bandpasses between channelizers are nearly identical (See Figure 4.18). These bandpasses are only true for the channelizer itself. The total bandpass of the system includes the bandpass of the SIS mixers, C-band amplifiers, and digitizer roll-off.

broadband noise source was amplified then split with a broadband power splitter and fed into the inputs of the two channelizers under measurement. On the LO port, a signal that was 0.25 GHz away from the nominal LO frequency was injected to shift the normal channelizer IF output from 0-500 MHz to 250-750 MHz. This was done to prevent any aliasing at the low end of the IF. The two output signals were then simultaneously sampled at 2 GHz and 8-bit resolution with the digital oscilloscope.

The stream of time-sampled data were then transferred to computer for processing. In software, the 256-lag correlation function was computed (see Equation 5.1). The FFT of the correlation function yields the relative phase of the signals. Figure 4.19 shows the results of this procedure.

A slope in the phase plots indicates path length difference. Overall phase offsets are slightly harder to explain, but could be due to a path length difference⁸ in the channelizer LO path, either in the external power splitter or in a slight misalignment in the placement of the mixer. It is also possible that the LO inputs to the mixers have slightly different impedance, which could also lead to a phase offset.

An offset is easier to calibrate out than a slope, amounting to a simple rotation in the imaginary plane for the correlator output. The slope is impossible to correct with

⁸The offsets in the measured phase are at most on order $\pi/4$ or $\lambda/8$ in wavelength. At 5 GHz, this amounts to 7.5 mm in free space. This is probably an unreasonable amount. The actual length in microstrip is smaller, but not by much. The exact amount of the offset, however, is not important; only the stability matters.

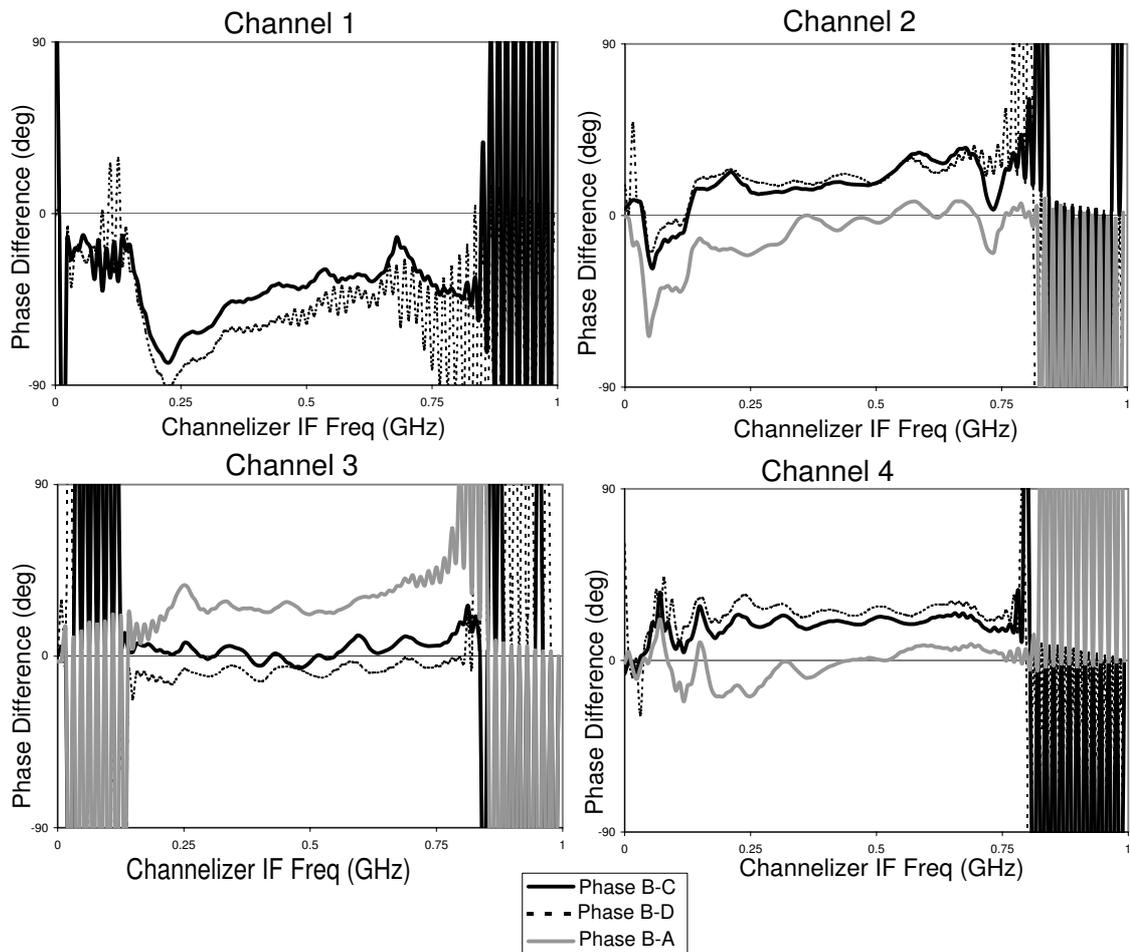


Figure 4.19: Relative phases at the IF output of the channelizer as measured relative to channelizer B. The frequency range of interest is 0.25-0.75 GHz. The phase at the edges diverges because of lack of signal and quantization noise. Note that B-A is missing from channel 1 due to bad data.

a broadband correlator without the use of compensating delays. Luckily, MINT's digital correlator measures the correlation over 8 bins across the 500 MHz IF of the channelizer. This gives us some ability to correct for slopes in the relative phase. To get a better idea of the relevant phase error, we can subtract a linear fit to the phase differences. Figure 4.20 has the phase measurements with a linear fit subtracted, and Table 4.3 lists the *rms* of the phase error. The decrease in signal due to phase error scales roughly as the cosine of the error (see Section 7.2.2). All values in the table are $\leq 10^\circ$, which is almost negligible.

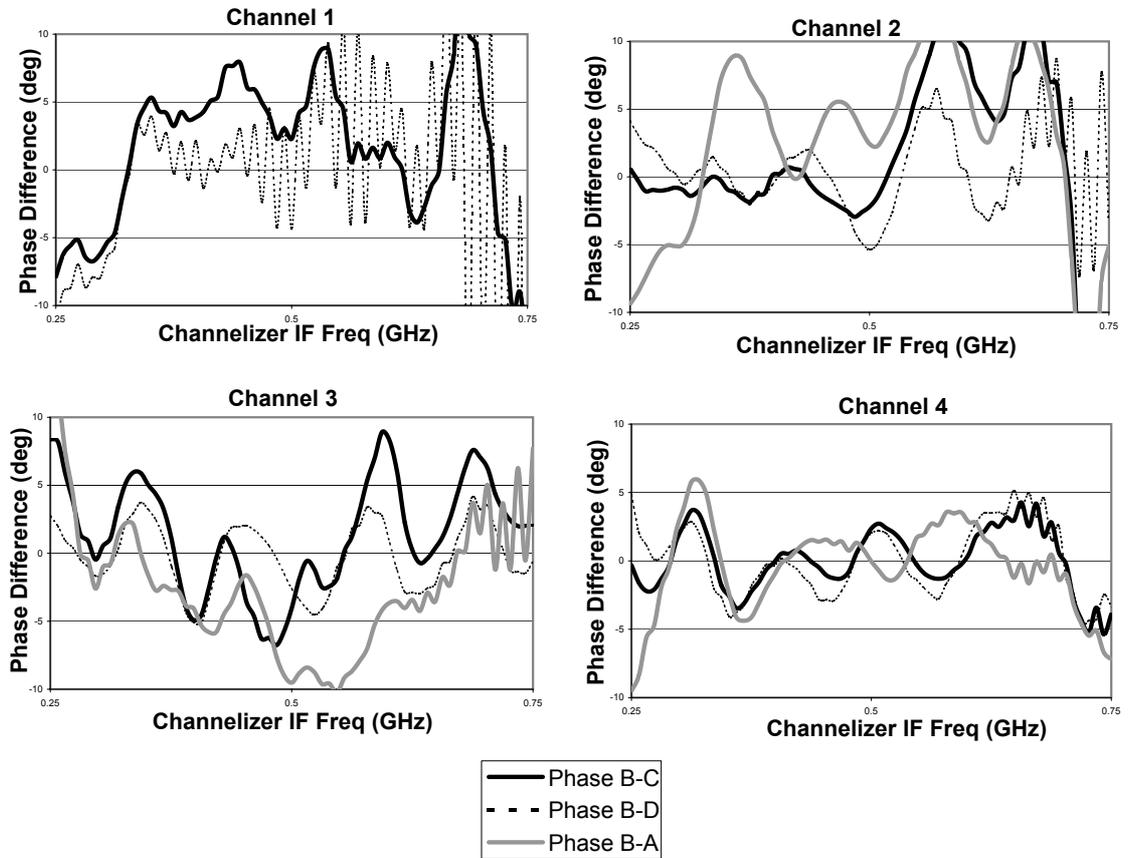


Figure 4.20: Relative phases at the IF output of the channelizer with a linear fit removed. On this enlarged view, channel 1, pair B-D obviously has a processing artifact, possibly due to poor quantization.

channelizer pair: channel:	b-c	b-d	b-a
1	5.7	10	-
2	6.9	3.4	7.3
3	4.3	2.5	5.5
4	2.3	2.6	3.3

Table 4.3: The *rms* of the phase errors after a linear fit has been subtracted. All are in degrees. Again, the pair B-A for Channel 1 is missing due to bad data.

4.11 Field Performance

The channelizer performed successfully over the 3-month campaign. The monolithic, low connector count design proved to be robust and reliable. Only one band in one channelizer failed during the campaign, and the failure was coincident with a system temperature test which required a few connection cycles on the output. The loss was likely due to a broken output connector.

Correlator Theory

Digital correlators are complex electronic instruments. The mathematical theory behind the operation of these devices was refined in the 1960's in anticipation of the VLA, which makes use of a 50 MHz digital correlator. There are many aspects to the theory of digital correlation, including SNR reduction relative to an analog correlator and correction factors based on quantization. These issues will be explored in this chapter for the specific case of the MINT correlator.

What the digital correlator lacks in simplicity and sensitivity compared to analog correlators it makes up for with stability and functionality. Digital electronics are inherently stable to thermal variations and crosstalk. In addition to returning the same information as an analog correlator (that is, both the real and imaginary correlation across the band), they also return the correlation function. As will be seen, the correlator can further subdivide the visibility into frequency bins. This can allow for observing emission lines, or, in the case of a thermal source, allow for phase profile corrections.

The equations outlined here are drawn from many sources. The textbook of interferometry, Thompson et al. (1986), summarizes the theory developed for the VLA (Very Large Array) correlator, much of which translates directly to the MINT correlator. Very special thanks goes to David Hawkins of OVRO, who is designing a wide-band correlator for OVRO. The MINT correlator design was inspired by his design, and his memo on correlator theory (Hawkins 1998), was very useful in understanding the MINT correlator.

5.1 Statistics

The task of the correlator is to calculate the cross-correlation function of two functions $x(t)$ and $y(t)$:

$$r(\tau) = \int_0^T dt x(t)y(t - \tau), \quad (5.1)$$

where T is the integration time and τ is a time delay known as the *lag*. By the Wiener-Khinchin relation, the Fourier Transform (F.T.) of the auto-correlation function is the Power-Spectral-Density (PSD), a real function. Likewise, the F.T. of the cross-correlation function is the Cross-Power Spectral Density (CPSD), a complex function:

$$CPSD(\nu) = \tilde{x}^*(\nu)\tilde{y}(\nu) = \mathcal{F}(r(\tau)). \quad (5.2)$$

The discrete formulation of this theorem is known as the *discrete correlation theorem* (see Equation 6.58 in Brigham (1988)), and is important when the signal is discretely sampled. The real part of the *CPSD* is exactly the real part of the visibility, as in Equation 2.14, except that it is given as a function of frequency. The imaginary component of the visibility, Equation 2.15, is the imaginary part of the *CPSD*. To recover the output of an analog correlator, a sum over ν is performed.

Assuming that the input signals are gaussian and normalized¹ to have the same variance σ^2 , the probability distribution $p(x, y)$, should follow the bivariate gaussian distribution:

$$p(x, y) = \frac{1}{2\pi\sigma^2\sqrt{1-\rho^2}} \exp\left[\frac{-(x^2 + y^2 - 2\rho xy)}{2\sigma^2(1-\rho^2)}\right] \quad (5.3)$$

$$\text{and } \rho = \frac{\langle x(t)y(t) \rangle}{\sqrt{\langle x^2(t) \rangle \langle y^2(t) \rangle}}.$$

This gives the probability of $x(t)$ being within the range $[x, x + dx]$ and $y(t) \in [y, y + dy]$ at the same time, t . The parameter ρ is known as the correlation coefficient. This formula represents zero lag, $\tau = 0$. Nonzero lags have $(t - \tau)$ inserted:

¹The actual signals are normalized via variable attenuators on the correlator board.

$$\rho(\tau) = \frac{\langle x(t)y(t-\tau) \rangle}{\sqrt{\langle x^2(t) \rangle \langle y^2(t-\tau) \rangle}}. \quad (5.4)$$

The input signals at the correlator are dominated by amplifier, mixer, and sky noise, all of which are uncorrelated between inputs. The correlation coefficient is due to signal from the CMB fluctuations predicted in Section 2.2.1, which is very small and therefore $\rho \ll 1$. In this limit, $(1 - \rho^2) \approx 1$ and Equation 5.3 may be approximated to first order in ρ as:

$$p(x, y) = \left[\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \right] \left[\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right) \right] \left(1 + \frac{\rho xy}{\sigma^2}\right), \quad (5.5)$$

where it can be clearly seen that for $\rho = 0$, the bivariate gaussian distribution is just the product of two gaussian distributions.

Given that the variances of $x(t)$ and $y(t)$, $\langle x^2(t) \rangle$ and $\langle y^2(t) \rangle$ are both σ^2 , and from the definition of ρ (Equation 5.3),

$$\langle x(t)y(t) \rangle = \rho\sigma^2. \quad (5.6)$$

5.2 Sampling

The functions $x(t)$ and $y(t)$ are continuous. Digital signals differ from these signals in two ways. The actual input to the correlator is both time-sampled and quantized. The case of sampling is considered first.

The signals are bandlimited at frequency $\Delta\nu$, and the Nyquist sampling rate is $2\Delta\nu$. For MINT, $\Delta\nu \sim 500$ MHz and the sampling rate is 1 GHz.

The result, $r_\infty = x_i y_i$, of a single multiplication at lag $\tau = 0$ is used to examine the statistics of this correlator. The subscript ∞ refers to no quantization (or infinitely fine quantization), and x_i and y_i refer to the samples.

The signal to noise in the output of a sampled but unquantized correlator is:

$$SNR_\infty = \frac{\langle r_\infty \rangle}{\sigma_\infty} = \frac{\langle r_\infty \rangle}{\sqrt{\langle r_\infty^2 \rangle - \langle r_\infty \rangle^2}}. \quad (5.7)$$

The statistics of the Nyquist sampled x_i and y_i follow that of the unsampled case,

$$\langle r_\infty \rangle = \langle x(t)y(t) \rangle = \rho\sigma^2. \quad (5.8)$$

The expectation of r_∞^2 may be found from:

$$\langle r_\infty^2 \rangle = \langle x^2 y^2 \rangle = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^2 y^2 p(x, y) dx dy = \sigma^4(1 + 2\rho^2), \quad (5.9)$$

where $p(x, y)$ is the bivariate gaussian distribution given in Equation 5.5. The *rms* per multiplication is the square root of the difference between Equation 5.9 and the square of Equation 5.8,

$$\sigma_\infty = \sigma^2 \sqrt{(1 + \rho^2)}. \quad (5.10)$$

Therefore the signal to noise for sampled but unquantized signals is (assuming again that $\rho \ll 1$):

$$SNR_\infty \approx \rho. \quad (5.11)$$

5.3 Quantization

Ideally, the signals would be quantized with arbitrarily fine resolution. In practice, however, there is a tradeoff between resolution and electronic complexity. The MINT digitizer quantizes the incoming signals into 2-bit, 4-level numbers. The binary code is assigned according to Table 5.1. The goal of this section is to calculate the SNR relative to the ideal case of Equation 5.11.

The quantities that may be controlled are the threshold value, V_0 , and the relative weights of the levels. With 4-level quantization, there is only one degree of freedom in the relative weights, the value n in Table 5.1. The expected output of a single multiplication may be constructed given these levels. There are six possible products given this quantization table: $-n^2, -n, -1, 1, n,$ and n^2 . There are two ways to achieve $-n^2$, from $-n \cdot n$ and $n \cdot -n$. There are four ways to get $-n$, four ways

binary code	weight	voltage range
00	-n	$-\infty < V < -V_0$
01	-1	$-V_0 < V < 0$
01	1	$0 < V < V_0$
11	n	$V_0 < V < \infty$

Table 5.1: Quantization table. V_0 is the threshold voltage.

to get n and so forth. With these considerations, the expected output of a single multiplication is:

$$\langle r_4 \rangle = 2n^2 p_{nn} - 2n^2 p_{n\bar{n}} + 4np_{1n} - 4np_{1\bar{n}} + 2p_{11} - 2p_{1\bar{1}}, \quad (5.12)$$

where p_{nn} refers to the probability that both samples have value n , which is the same as the probability that both have value $-n$. The bars over the subscript refer to the negative values. $p_{1\bar{1}}$ is the probability that one sample has value 1 and the other has value -1 and so forth. In other words, the expected value for one multiplication is found by multiplication of each possible outcome by its probability of occurrence.

To calculate the probabilities, the bivariate gaussian distribution is integrated over the proper ranges. For instance, to calculate $p_{1\bar{n}}$,

$$p_{1\bar{n}} = \int_0^{V_0} \int_{-\infty}^{-V_0} p(x, y) dx dy, \quad (5.13)$$

where again $p(x, y)$ is the bivariate gaussian distribution given in Equation 5.3. The variance is found by $\sigma_4^2 = \langle r_4^2 \rangle - \langle r_4 \rangle^2$. The first term of the variance is:

$$\langle r_4^2 \rangle = 2n^4 p_{nn} + 2n^4 p_{n\bar{n}} + 4n^2 p_{1n} + 4n^2 p_{1\bar{n}} + 2p_{11} - 2p_{1\bar{1}}, \quad (5.14)$$

where the square of the possible outcomes of multiplications are weighted by the probability of occurrence. The SNR calculated from $\langle r_4 \rangle / \sigma_4$ is dependent on the threshold, V_0 , and the weight n . A numerical integration of the above equations as a function of the threshold and weight is required to determine the optimum values. These are determined in Thompson et al. (1986) (see eq 8.48). The optimum SNR is:

	00	01	10	11
00	n^2	n	$-n$	$-n^2$
01	n	1	-1	$-n$
01	$-n$	-1	1	n
11	$-n^2$	$-n$	n	n^2

Table 5.2: Multiplication table for standard scheme.

	00	01	10	11
00	9	3	-3	-9
01	3	1	-1	-3
01	-3	-1	1	3
11	-9	-3	3	9

Table 5.3: Multiplication table for standard scheme with the optimal integer value $n=3$.

$$\frac{SNR_4}{SNR_\infty} = 0.881, \quad (5.15)$$

with the optimum integer weight and threshold set to $n=3$ and $V_0 = \sigma$:

5.4 Multiplication Table

To further simplify the electronics, MINT makes use of a modified multiplication scheme. Instead of using the weights implied by Equation 5.12 and Table 5.2, a deleted inner product multiplication scheme in Table 5.4 is used. Additionally, the deleted inner product table is scaled and biased to further simplify the correlation algorithm. This modification allows the accumulator to be a pure 3-bit adder, the details of which are explained in the next chapter. The actual multiplication scheme appears in Table 5.5.

The values of the expected correlator output, $\langle r_{4d} \rangle$, and the variance, σ_{4d}^2 , may be calculated in a similar manner as the full 4-level multiplier (see Equation 5.12). This is calculated in Hawkins (1998). The results, given $\rho \ll 1$ are:

	00	01	10	11
00	9	3	-3	-9
01	3	0	0	-3
01	-3	0	0	3
11	-9	-3	3	9

Table 5.4: Multiplication table for modified deleted inner product scheme.

	00	01	10	11
00	6	4	2	0
01	4	3	3	2
01	2	3	3	4
11	0	2	4	6

Table 5.5: Deleted inner product scheme, scaled and biased. The values from Table 5.4 are divided by 3 and then added to 3. The MINT correlator uses this scheme.

$$\langle r_{4d} \rangle = \frac{2\rho}{\pi} \cdot [E^2(l - 2k + m) + 2E(k - m) + m], \text{ and} \quad (5.16)$$

$$\sigma_{4d} = \sqrt{\phi^2(l^2 - 2k^2 + m^2) + 2\phi(k^2 - l^2) + l^2}, \quad (5.17)$$

where

$$\begin{aligned} \phi(V_0/\sigma) &= \operatorname{erf}\left(\frac{V_0/\sigma}{\sqrt{2}}\right) = \frac{2}{\sqrt{\pi}} \int_0^{\frac{V_0/\sigma}{\sqrt{2}}} e^{-t^2} dt \\ \text{and } E(V_0/\sigma) &= \exp\left(-\frac{1}{2}[V_0/\sigma]^2\right), \end{aligned} \quad (5.18)$$

where m is the weight for the inner product, k is the weight for middle product and l is the weight for the outer product. For the deleted inner product table, $(m, k, l) = (0, 3, 9)$. With this scheme, the optimum SNR is:

$$\frac{SNR_{4d}}{SNR_{\infty}} = 0.872, \quad (5.19)$$

achieved when

$$V_0 = .906\sigma. \quad (5.20)$$

The SNR for the deleted inner product scheme is about 1% less than the full multiplier. This is a small price to pay for the reduced logic needed to implement the multiplier. The output of the normal multiplication scheme ranges from -9 to $+9$, which requires 5 bits to represent. The outputs of the deleted inner product scheme ranges from 0 to 6, which only requires 3 bits to represent.

5.5 Corrected Correlation Coefficient

The correlation result from Equation 5.16 is dependent on the value of the threshold voltage V_0 . In practice, this value is not fixed and must be monitored. The next chapter details the monitoring scheme. The purpose of this section is to find the relationship between actual correlation coefficient ρ , and the measured correlator output as a function of V_0 . As can be seen from Tables 5.4 and 5.5, the proper values to use for the weights are² $(m, k, l) = (0, 1, 3)$. Substituting these into Equation 5.16 and solving for ρ gives:

$$\rho = r_{4d} \cdot \frac{\pi}{2} \frac{1}{E^2 + 2E}, \quad (5.21)$$

where E is the exponential given in Equation 5.18. A plot of the factor $\frac{\pi}{2} \frac{1}{E^2 + 2E}$ appears in Figure 5.1. The correlator is able to servo the threshold to within 0.5 dB of the optimum value, which means that the range of possible thresholds is $V_0 \in [0.855\sigma, 0.959\sigma]$. The correction factor varies by about 6% over this limited range. In the field, the threshold was reset every hour to be within this range.

Equation 5.21 is relevant for an auto-correlator at lag $\tau \neq 0$. It is in general not relevant for a cross correlator because there are two factors of E , one for each

²The multiplication table for the deleted inner product scheme (Table 5.4) implies that the correct weights should be $(m, k, l) = (0, 3, 9)$. The rescaling and biasing, however, confuses the issue. Equation 5.17 cannot account for a bias, so it is appropriate to incorporate only the rescaling into the weights, meaning we should use $(m, k, l) = (0, 1, 3)$.

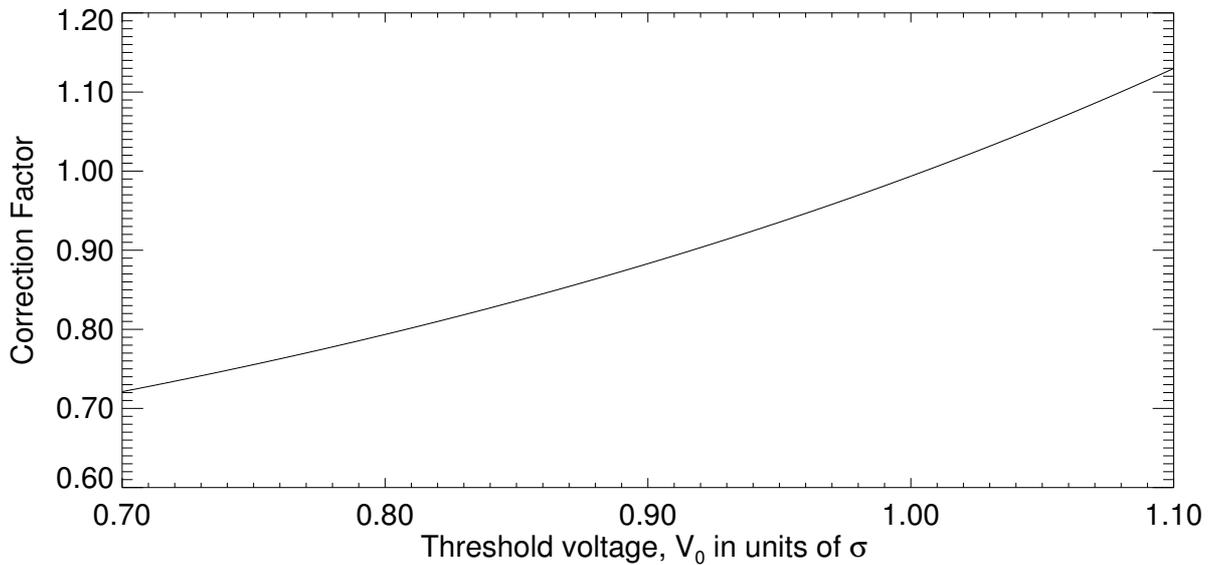


Figure 5.1: Correction factor from Equation 5.21. The correction factor is used to convert the output of the digital correlator to the correlation coefficient, a number that is linearly proportional to the output of an analog correlator.

digitizer because each has a different threshold voltage. Equation 5.21 can, however, be generalized for a cross correlator:

$$\rho = r_{4d} \cdot \frac{\pi}{2} \frac{1}{E_1 E_2 + E_1 + E_2}, \quad (5.22)$$

where the subscripts label the separate digitizers. Figure 5.1 applies when the threshold changes identically for both digitizers, a likely scenario when the receivers have the same sensitivity and the power level change is dominated by a change in the temperature of the atmosphere.

5.6 Expected Correlator Output

In this section the exact statistics are calculated for the output of the correlator. Equation 5.16 is the expected value of a single multiplication which is a linear estimator for ρ for a given V_0 . The square root in Equation 5.17 leads to the *rms*, σ_{4d} , of a single multiplication. Of course, correlation also requires accumulation over time,

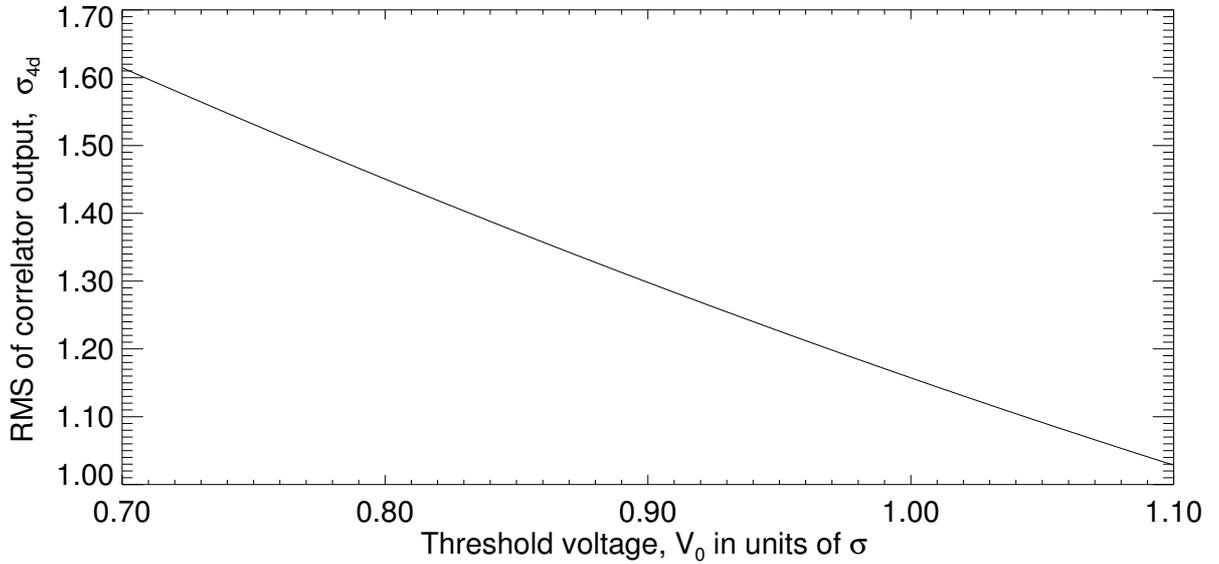


Figure 5.2: The *rms* of the correlator output for a single multiplication. Here, σ is the *rms* of the incoming signals to the correlator and σ_{4d} is the *rms* of the output.

so the expected output of an accumulated correlator is :

$$\sum_i^{N_T} r_i = N_T r_{4d} \pm \sqrt{N_T} \sigma_{4d}. \quad (5.23)$$

In other words, the output *integrates down* as $1/\sqrt{N_T}$ where N_T is the number of multiplications that go into the accumulation. A plot of the *rms*, σ_{4d} , appears in Figure 5.2. The value at the optimal threshold is $\sigma_{4d}(V_0) = 1.289$.

As with the expected correlation coefficient, Equation 5.17 for σ_{4d} is only true if the thresholds for two digitizers change in exactly the same way. A generalized form of Equation 5.17 is:

$$\sigma_{4d} = \sqrt{7\phi_1\phi_2 - 8(\phi_1 + \phi_2) + 9}, \quad (5.24)$$

where the subscripts again refer to digitizer and the values for $(m, l, k) = (0, 1, 3)$ have been substituted.

A comparison of the actual performance with these predicted values of $\sum_i^{N_T} r_i$ and σ_{4d} are given in Section 6.4.

In the raw form, the correlator output is uncalibrated. It is simply a correlation coefficient, defined in terms of percentage correlation of the incoming signals. To calibrate in temperature, the temperature of the incoming signals must be known. The calibration procedure is discussed in Section 7.1.1.

5.7 Summary

The input to the correlator is two signal streams, each sampled at 1 GHz and 4 levels of resolution. The correlator then calculates the correlation function, using the multiplication scheme given in Table 5.5, designed to simplify the algorithm and optimize the SNR. The multiplication is performed for 16 lags, or time delays between incoming signals. The result of the multiplication is accumulated for N_T samples, with the relative *rms* of the output decreasing as $1/\sqrt{N_T}$.

The following chapter examines the implementation of the correlator in digital logic. Also shown is a comparison between the predicted statistics and the measured performance.

Correlator Implementation

The correlator is actually two major subsystems integrated into one unit. Correlators on large interferometers are usually just that—the electronics associated with calculating the actual correlation function. The digitization is done in a completely separate unit, with high speed data cables connecting the two. The MINT correlator incorporates the digitizers into the same card containing the correlator logic. All correlators are housed in a single rack-mounted, water-cooled unit. See Appendix B.2 for details of the enclosure.

The MINT design exploits many new technologies and benefits from the semiconductor industry’s push to higher logic densities and faster clock rates. The correlator is implemented as a Field-Programmable Gate Array (FPGA). FPGA’s are logic devices whose functionalities are defined in software, then programmed into the chip after it has been integrated into a circuit board. While programmable gate arrays have been a standard building block in circuit design for many years, the technology has only recently been able to compete with more traditional and costlier methods implementing custom logic.

The development of the correlator involved many people. The first attempt at high speed digital electronics coupled to FPGA’s was made by James Hinderks for his senior thesis (Hinderks 1999). Hinderks was successful in producing a 1-bit 2-input 500 MHz digitizer, which fed a single-lag correlator. Tobias Marriage carried on the work in his senior thesis by designing a digitizer board with two 1 GHz inputs and 2-bit quantization (Marriage 2000). Coupled to an external correlator chip, this board produced the first correlation functions. As a following summer project, Marriage went on to design and lay out the full digitizer/correlator, incorporating the earlier

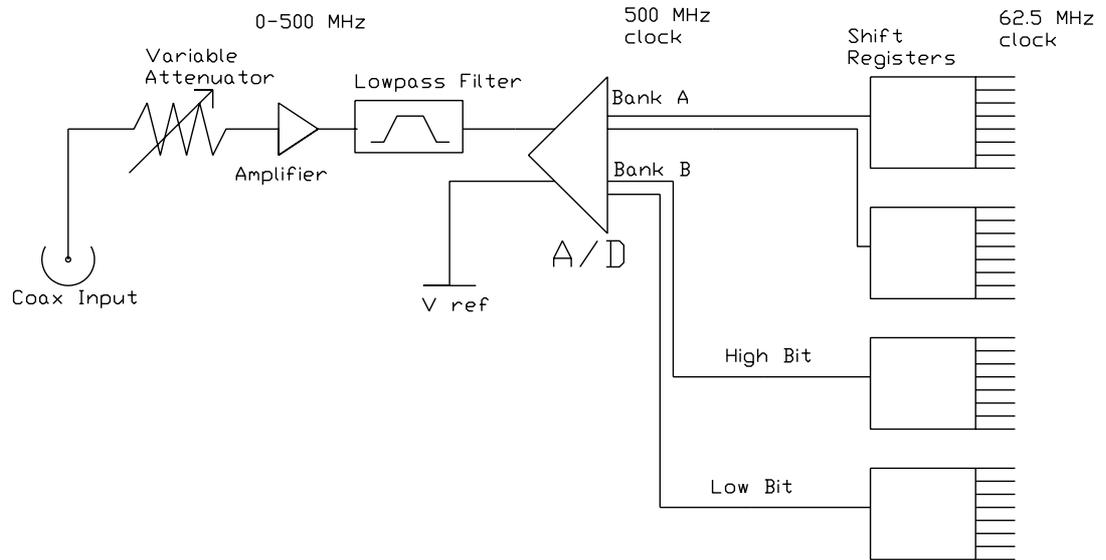


Figure 6.1: Schematic of a single digitizer section. There are 4 digitizers per board. After conversion to ECL levels, the outputs of each of the four digitizers is fed to the correlator.

digitizer design with an on-board FPGA correlator. Finally, Mark Tygert developed the OVRO correlator code into the first working version for the MINT correlator.

6.1 Digitizer

The digitizer is responsible for the analog processing of the 0-500 MHz signal from the channelizer, quantizing and sampling the signal into 1 ns 2-bit samples, then slowing the data rate to 62.5 MHz, the input clock rate of the correlator. These steps are accomplished in three major sections: the RF section, the A/D converter, and the deserializer. Figure 6.1 is a schematic of all three sections.

All components are commercially available off-the-shelf parts, although the critical components have only recently become available. Of particular importance is the A/D converter itself, as the rest of the digitizer is designed around it. There are many companies racing to develop faster and more accurate digitizers. The new GHz-digitization rate A/D converters enable digital correlators to compete with the bandwidth of analog correlators.

6.1.1 Data Path

RF section

The RF section is composed of a variable attenuator, an amplifier and a lowpass filter. The objective is to bring the incoming signal to a level that the A/D converter requires for optimal sampling (*rms* of 0.25V), and to filter out any residual signal above the Nyquist frequency of 500 MHz. A variable attenuator is required to compensate for fluctuating power levels out of the receivers, due mostly to changing sky temperature. The attenuator is a 6-bit digital surface-mount attenuator with a 0.5 dB step size and a maximum of 31.5 dB of attenuation. The levels are externally controlled via a control computer. Section 6.3.1 is a discussion of the servo loop used to control the attenuator. Section 7.1.1 details how the attenuator setting is involved in temperature calibration.

A/D Converter

The A/D converter converts the incoming 500 MHz analog data into two separate 6-bit 500 MHz digital data streams. Of the 6 bits, only the highest 2 are retained. The two separate output streams are interleaved so that two sequential samples appear at any given time on the two separate streams, that is to say that the output data are already deserialized or demultiplexed (demuxed) by a factor 2. This achieves a lower data clock rate at the expense of more signal lines.

Deserializer

The correlator FPGA cannot operate at the output clock rate from the digitizer. While there is some flexibility with the clock rate, a reliable clock rate for the correlator is 62.5 MHz, or 8 times slower than the initial demux. To slow the clock rate even further, each of the bits from the digitizer are fed into a 8 level shift register. The shift register has an output clock for every 8 input clocks, and it is this clock that drives the correlator. This additional level of demux results in an overall demux factor of 16, with 32 output lines per digitizer due to the fact that each sample is two bits wide.

6.2 Correlator

After deserialization, all data lines from four separate digitizers are sent to a single FPGA correlator chip. Placing four digitizers and a correlator on one board eliminates the need for high-speed interconnects that are often problematic in other designs.

The correlator has the computationally intensive task of calculating the correlation function for each of the six possible pairings of the four inputs from the digitizers. It does so in real time, and accumulates the correlation function for an externally defined time period (0.5 sec) set by the oscillator box control circuitry.

6.2.1 Technology

The particular design scheme for the correlator has only recently been made possible by the advent of commercially available high-density, high-speed FPGA's. Much like microprocessor technology, programmable logic has increased in speed and power on very short time scales. The design of the MINT correlator used the largest commercially available FPGA one year before deployment. This chip was just large enough for the placement of the entire 16-lag correlator logic into one chip.

The implementation of FPGA logic into a digital processing design differs radically from the traditional ASIC (Application-Specific Integrated Circuit) approach. Using ASIC's requires specification of a chip at the gate level, then fabricating many thousands of chips to be tested with very little margin for error. This approach is favorable under the limit of large production runs, particularly because ASIC's have a speed advantage over FPGA's. Until recently, all large scale digital correlators, including the VLA, have used proprietary ASIC's.

The advantages of using FPGA's in the prototyping and small run regime are hard to overstate. The ability to design the logic in software, with multiple levels of abstraction, and then immediately implement and re-implement the logic into an in-system device makes for a short and affordable development cycle. Not only is the internal logic flexible, but so, too, are the external connections because the I/O pins are reassignable, allowing for multiple levels of redundancy in case of failure or poor circuitboard design. In fact, the physical circuit may be designed before the logic for the FPGA is complete.

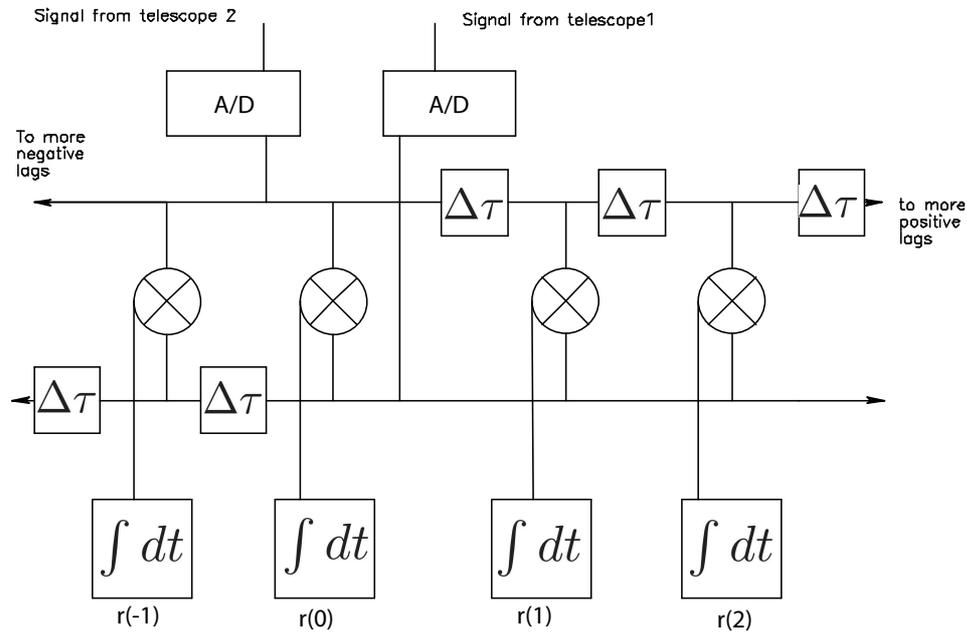


Figure 6.2: General correlation algorithm for digital lag-correlators. The input signals are actually 2 bits wide. Each digitized signal is fed directly to a digital multiplier with the result read into an accumulator to calculate $r(0)$, or zero lag. To calculate other lags, one or the other signals are delayed by $\Delta\tau$, which is just one clock cycle. In the demultiplexed scheme, the diagram is slightly more complicated in that each lag actually requires 16 multiplications. Lags are then implemented as shifts in the 16-wide buses. See Figure 6.4 for a detail of a single lag.

6.2.2 Algorithm

The correlator is used to calculate the function:

$$r(\tau) = \int_0^T dt x(t)y(t - \tau). \quad (6.1)$$

Figure 6.2 is a diagram of how this is implemented in a digital correlator. The correlator takes as input two digitized input streams from two telescopes. To calculate $r(\tau = 0)$, the two inputs are simply multiplied together and accumulated. To calculate the other lags, a time delay, $\Delta\tau$, is inserted into one or the other input stream before multiplying and accumulating. The time delay is an integral multiple

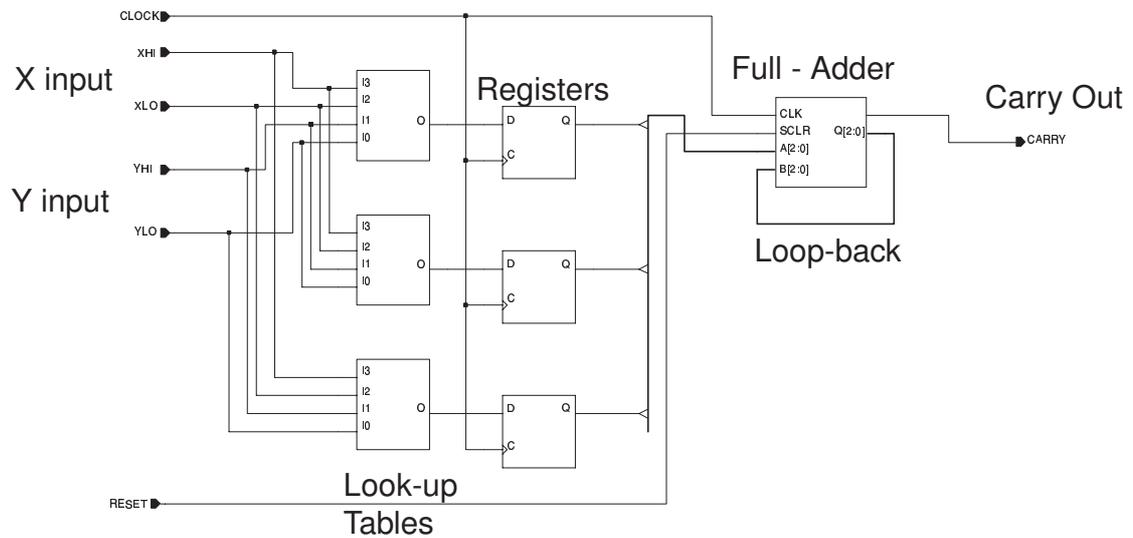


Figure 6.3: Schematic of the multiplier-adder kernel. The inputs, X and Y come from two different receiver via two digitizers and a distribution network inside of the correlator. XHI and XLO refer to the high and low bit from the X input. The clock line comes from an internal clock-distribution network. The reset is asserted after every readout of the correlator and clears the accumulator.

of the Nyquist sampling period of 1 ns, and is easily implemented as a memory element. In this algorithm, the entire correlation function is updated in real time, with new data presented at the accumulator on every clock cycle. All lags are calculated simultaneously in parallel.

Figure 6.2 is a slight simplification because it does not include the demux by factor of 16. In actuality, there are 16 input samples from each telescope and the correlator is 16 times slower than the Nyquist sampling rate of the A/D converter. The above schematic is easily adapted to accommodate the slower clock. Instead of one sample multiplication contributing to the accumulation for a single lag, there are 16 multiplications contributing, one for each level of demux.

Multiplier-Adder Kernel

At the heart of the correlator algorithm is the multiplier-adder kernel. It is the most repeated element, and, as such, every effort has been made to streamline the amount of logic it uses. A schematic of the kernel appears in Figure 6.3. Lookup tables

implement the biased and deleted-inner-product multiplication scheme in Table 5.4. Each lookup table takes 4 bits of input, one pair from each receiver (the distribution network from each receiver is not shown in Figure 6.3). There are three separate lookup tables, each one being responsible for one of the bits in the 3-bit output. The deleted inner product scheme allows the use of two fewer lookup tables than the full multiplier, resulting in a 40% reduction in logic for the multiplier and a $\sim 40\%$ reduction in the adder. The output of the lookup table is fed into a register, which holds the value for the next clock cycle. In this case, registers are used to allow the circuit to operate at a higher frequency at the expense of latency. The outputs of the registers are wired to a bus and then sent to one of the inputs of a 3-bit full adder. The other input is fed by the lower 3 bits of the previous adding operation. The highest bit is the carry output, which goes high whenever the result of adding is higher than 111_2 . An adder wired this way acts as a 2-bit accumulator that bleeds out the highest bit whenever the lower bits are filled. Dropping the low order bits reduces the inherent resolution of the multiplier by 3 bits.

Lag and carry cascade

To calculate a single lag, 16 such kernels are required because of the demux by 16 factor. Figure 6.4 is schematic of a single lag, in this case, lag 2. For simplicity, only 8 of the 16 multipliers are shown. The current samples presented at the input of the correlator are $x(0) - x(15)$ and $y(0) - y(15)$, with sample 15 being the latest out of the A/D converter. Samples -1 through -16 are stored from the previous correlator clock cycle. To calculate positive lags, the x samples are shifted by two samples relative to y . To calculate negative lags, the y 's are shifted relative to x 's. Each pairing of samples is then fed into the multiplier adder.

The carry outputs of each multiplier adder are added together through a carry-cascade chain. Each element of the carry cascade is just a 1-bit full adder with the output wired to the carry input and the carry bit fed to the next level of cascading. It operates much like the adder described before, except that there is a carry input wired to the carry output. It takes 2 bits of input and calculates the sum while outputting the highest bit and “remembering” the result of the previous operation and also including it in the sum.

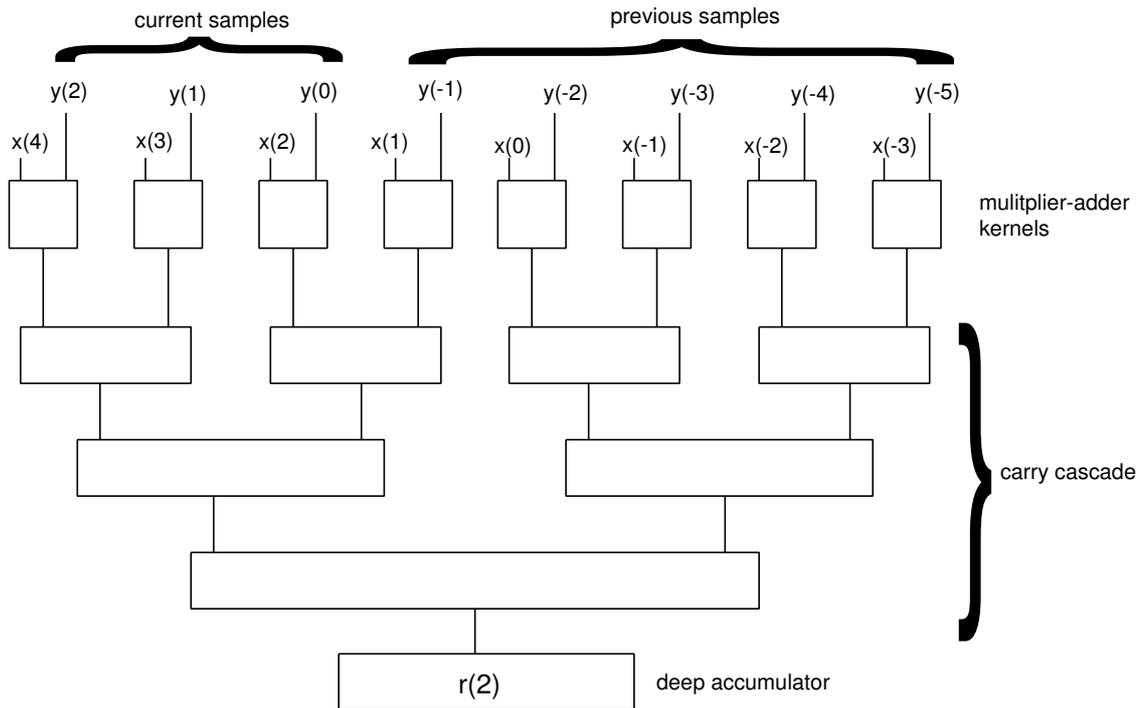


Figure 6.4: Schematic of a single lag for $\tau=2$. Only 8 of the required 16 multiplier-adder kernels are shown. To add another 8 kernels, an additional level of carry cascade would be required.

Four levels of cascading are required to reduce 16 carry outputs to 1. With each level of cascading the inherent resolution is reduced by another bit. This reduction combined with the 3-bit loss from the adder results in 7 bits of lost resolution, a figure that is important in calculating the correlator statistics (see Equation 6.2). The loss of resolution is counterbalanced with deep accumulation. The single-bit output of the last carry cascade element is fed into a deep accumulator, whose minimum depth is determined by the length of the accumulation cycle. For a 1/2 sec accumulation, there are $0.5 \text{ sec}/16 \text{ ns} = 3.125 \cdot 10^7$ clock cycles. If the carry bit is high for each clock, then the accumulator must be at least $\text{round_up}[\log_2(3.125 \cdot 10^7)] = 25$ bits deep.

Determination of Readout Size

The 7 bits of resolution lost in the adder/carry cascade together with the 25 bits from the accumulator amount to 32 bits of inherent resolution. In systems with gaussian statistics, the lowest one half of the bits on a large accumulation are mostly noise. This argument implies that only about 16 bits must be read out. The MINT design reads out 18 bits for slightly improved SNR. The discarding of bits allows for a simpler and faster readout scheme, and a proportionally smaller data set.

6.3 Board Level Control

There are many functions that must be controlled externally on the board. The control is accomplished through digital commands issued at a computer and read by an onboard controller chip, very similar to the correlator itself, but with less logic capacity and a slower clock rate.

6.3.1 Threshold Servo Loop

The controller chip sets the attenuation level of the variable attenuators on the RF section of the digitizer. A threshold servo-loop determines the appropriate threshold from a histogram of the data from each digitizer. The histogram of each digitizer is computed in the correlator chip and appears in the data stream that is read out to a control computer. In the computer, the statistic ϕ is calculated (see Equation 6.4) and compared to the optimal value (Equation 6.5). The difference is used to determine the appropriate attenuation level. Once the level is determined, the computer sends a command to the controller to set the level. Every hour after the initial servo the levels are checked. If the power levels have strayed, the threshold is set again.

6.3.2 Synchronization

Demux synch

The digitizer section requires multiple levels of control to present synchronized data at the correlator. The first level is the demux reset. There are 2 separate shift register chains on each bank of the 500 Hz output of each A/D converter. Each shift register

is composed of 2 actual ECL-logic chips. A single reset signal is wired to all such chips on the entire correlator. The signal originates in the controller chip and happens whenever a reset command is issued from the control computer. After conversion from TTL logic the signal is passed through 3 registers to quench metastability (Dover and Pearson 2000). Within a single digitizer, the reset synchronizes each bank so that the 32 bits of output are ordered correctly.

Clock Swallow

On a single correlator board, the digitizers also need to be synchronized relative to one another. The factor of 2 demux inside the A/D chip leads to separate A and B banks, and there is an ambiguity as to whether all A banks on all A/D's are synchronized. Although it was not used in the first season of MINT observations, there is a “clock swallow” circuit that will erase a single clock cycle going to any digitizer to synchronize it with the others. The signature of poor digitizer-to-digitizer synchronization is a misaligned correlation function. If the peak of the correlation function does not appear at lag zero, then the two digitizers are not synchronized. Since there is only a two-state ambiguity in the synchronization state (either the A's are aligned or the B's are aligned), the digitizers can be at most 1 ns apart in synchronization. There is, though, a software fix for this problem implementable in the data reduction procedure. See Section 7.1.3 for details.

6.4 Correlator Output

6.4.1 Mean

The statistics of the correlator are given in Section 5.6. There are two quantities to monitor: the expected value of the output given zero correlation ($\rho = 0$) and the *rms* of the output. Given zero correlation, $\rho = 0$, the expected correlator output $\langle r \rangle$ should also be zero. The complication is that the multiplication table, Table 5.5, is biased.

The accumulated average value of each multiplication from Table 5.5 is 3. The expected value of the correlation is just 3 times the number of multiplications. The

number of multiplications is given by the length of the accumulation time in clock cycles. Along with the accumulators at the end of each lag, there is a single accumulator for each correlator that counts the total number of clock cycles. The input to this accumulator is always wired high. The output of this accumulator is very stable, and it returns a value of $N_T = 31,147,600$.

The total accumulation cycle is 1/2 second long, set by the oscillator control circuitry. To find the expected number of correlator cycles, the accumulation time is divided by the clock cycle length, 16 ns. The result is $N_T = 31,250,000$. The difference between the expected and the measured value from the clock-cycle counter is 102400 cycles, which is accounted for by noting that each switch of the LO phase shifter induces a blanking time of $2^9 = 512$ cycles. “Blanking” refers to ignoring data. This time is needed for the PLL to lock and settle. In each 1/2 second accumulation cycle, the phase is switched 200 times, which accounts for the difference.

The expected value for the correlator output, r , for a single lag is $3 \cdot 16 \cdot N_T / 2^{n_{drop}}$, where 3 is the mean value of each multiplication, n_{drop} is the number of bits dropped from the accumulation, and the factor of 16 is due to the fact that there are 16 multiplier kernels for each lag. The first place where bits are dropped is in the adder of the correlator kernel. The adder is a 4-bit adder, where we strip off the top bit, which is equivalent to dropping 3 bits. The next place is in the carry-cascade chain, where a bit is dropped for each level in the cascade. To go from 16 carry bits to 1 bit, 4 bits of resolution are dropped. The last place that bits are dropped is on the output of the accumulator itself, where only 18 bits are read and 7 are dropped. This leads to a total of 14 dropped bits. So finally, the expected value for the correlator output for a single lag given $\rho = 0$ is:

$$\langle r_{\rho=0} \rangle = \frac{3 \cdot 16 \cdot N_T}{2^{n_{drop}}} = \frac{3 \cdot 16 \cdot 3.11476 \cdot 10^7}{2^{14}} = 91252.734, \quad (6.2)$$

where the value substituted for N_T is the measured value from the clock cycle counter, accounting for the blanked bits.

6.4.2 *rms*

The *rms* for the accumulator output is given by Equation 5.23. The predicted *rms* of the correlator output is¹ :

$$\sigma_R = \sqrt{\langle r^2 \rangle - \langle r \rangle^2} = \frac{\sqrt{N_T} \sigma_{4d}}{2^{n_{drop}}} = 1.756, \quad (6.3)$$

where we have again divided by the number of dropped bits. The value of σ_{4d} was taken at the optimum threshold, i.e. $\sigma_{4d}=1.289$.

6.4.3 Correlator Data Stream

Figure 6.5 shows the typical output for a correlation between the outputs for receivers A and B (see Figure 3.1 for locations). The expected value for ρ is small (10^{-6}) for a CMB anisotropy signal. It may therefore be assumed that $\rho = 0$ for this short sample. The top panel shows the DC levels of the diodes for two receivers, which are proportional to the power and measured before the signal enters the channelizer (see Figure 3.8 for location of the diodes). The next two panels show the occupation of the lower digitization level, calculated from the digitizer histograms from the following formula:

$$\phi = \frac{N_{10} + N_{01}}{N_{00} + N_{01} + N_{10} + N_{11}}, \quad (6.4)$$

where N_{00} is the number of samples that have the binary value “00”, N_{10} is the number of “10” samples and so forth. The more samples that correspond to the “inner” samples, “01” and “10”, which happens when the incoming signal power is smaller, the higher the number is. The optimal value, from a SNR viewpoint, for ϕ is:

$$\phi_{opt} = \text{erf} \left(\frac{.906}{\sqrt{2}} \right) = .6356 \quad (6.5)$$

¹Instead of using N_T here it is more proper to use $2/\Delta\nu_{eff}$, where the effective bandwidth is defined both by the channelizer (Table 4.2) and from the bandwidth of the receivers. The latter quantity has not been measured.

This optimal value is also plotted on Figure 6.5. The measured values of ϕ can be used to calculate the expected *rms* using Equation 5.24. This modifies the value for σ_r in Equation 6.3, but for this particular data sample the difference is negligible.

The bottom two panels show the correlator output for both sideband separation phase cases. All traces show calibration spikes. If the temperatures of the receivers are the same, then the heights of the spikes in the DC levels imply that the coupling of the calibrator is stronger in B than in A, which is also evident in ϕ . The value of ϕ goes down when the calibrator is on because in the absence of attenuator servoing, more power on the digitizer places more samples above the threshold, which lowers ϕ . The two lowest panels show the accumulator output for lag $\tau = 0$. The two phase states of the LO are shown. The lower calibrator values for r during the $\Delta\phi = \pi/2$ implies a negative correlation coefficient.

To find $\langle r \rangle$ and σ_r from the data set, the calibration spikes must be excised. Figure 6.6 shows the accumulator output once the calibration spikes are removed. The mean of the top panel is 91252.210 and the mean of the bottom panel is 91252.219. In both cases, the measured value is less than the expected value from Equation 6.2 (91252.734) by about 0.5. This is a result of truncation when the bits are dropped at the accumulator readout, which rounds down the output by 0.5.

To examine the *rms*, a histogram of the output is inspected (Figure 6.7). A Gaussian fit shows that the correlator output closely follows a normal distribution. The calculated *rms* of the output is ≈ 1.893 in both cases – slightly higher than the value 1.756 from Equation 6.3. This is due to slightly smaller effective bandwidth (see footnote 1) resulting from non-flatness in the channelizer and receiver.

6.4.4 Calibration Spikes

The response of the correlator to actual signal can be tested by examining the calibration spikes. The amount of correlated signal could, in principle, be measured by the change in signal detected by the total power diode when the noise source is on. There is a diode on each receiver that monitors the power level entering the channelizer. The expected correlation coefficient, ρ , would just be the geometric mean of the fractional change in power of the two diode levels. The problem with this technique is the unequal bandpasses between channels and the possibly non-flat output of the

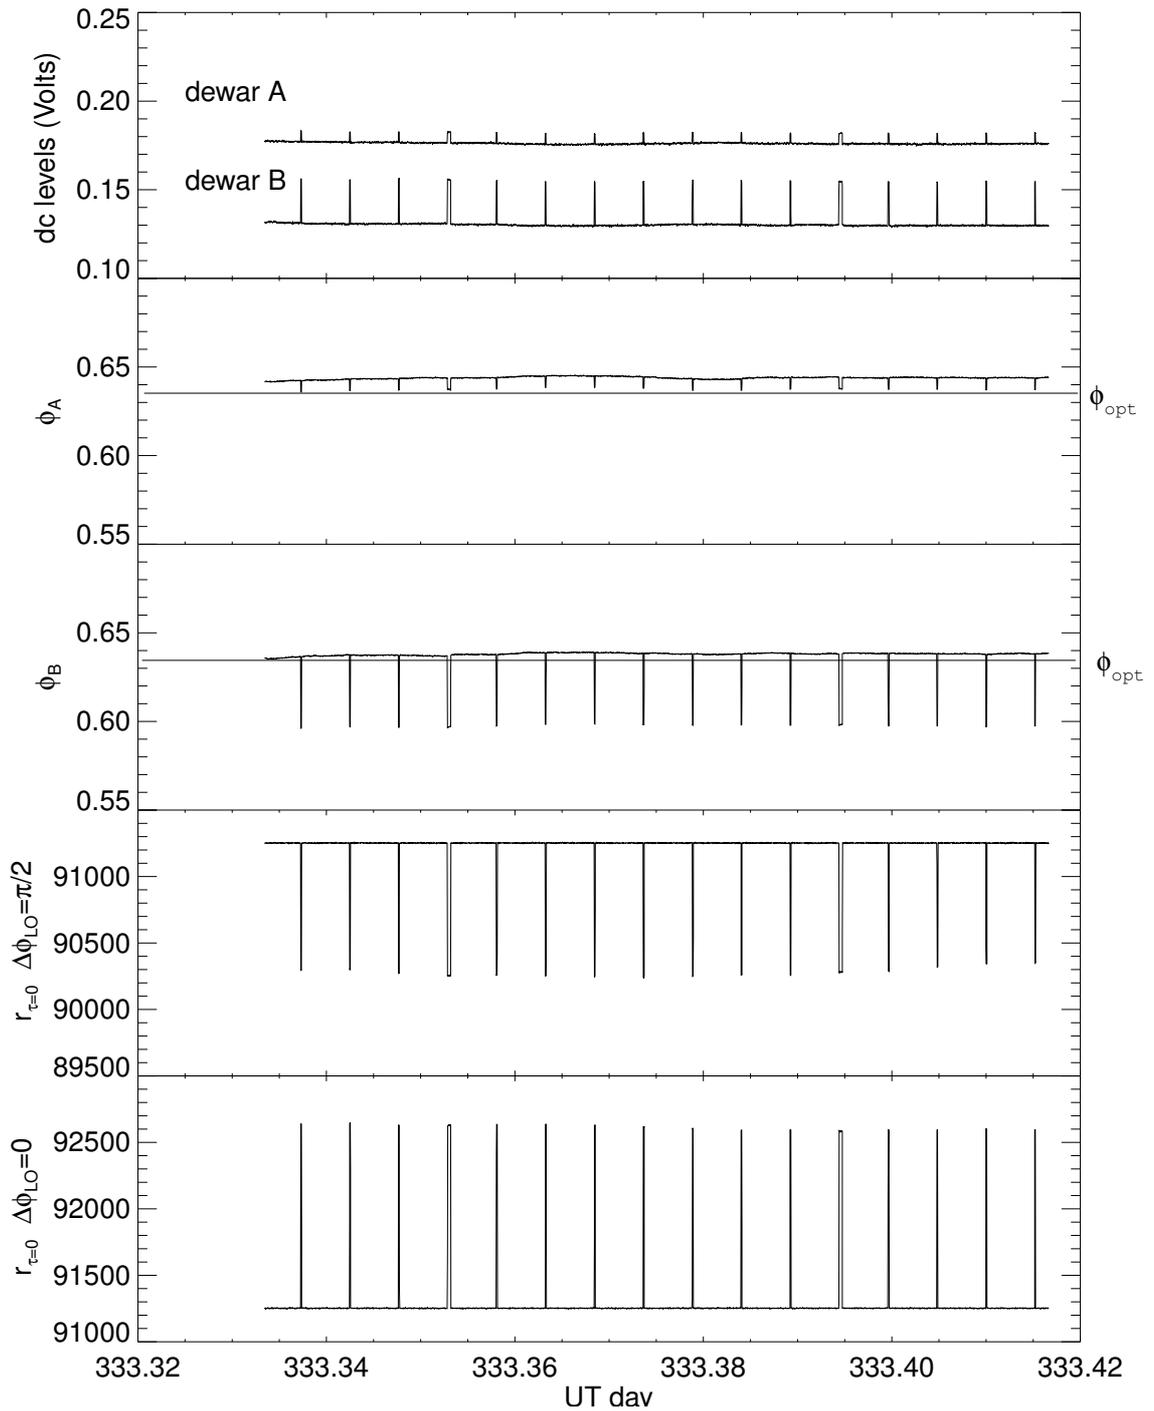


Figure 6.5: Two hours of typical correlator 1 output, taken at the coldest part of the night. The UT (Universal Time) fractional day is the time axis. The top panel shows the DC levels of the diodes, proportional to the total power from the receivers. For display purposes the DC level for B was lowered by 0.05V. The second and third panels are the occupation level monitors, ϕ , calculated from the histograms. The straight line in each panel is the optimal level, ϕ_{opt} . The bottom panels are the raw correlation output for the zeroth lag ($\tau = 0$) and for the two possible phase states of the LO's. The spikes in all the panels correspond to the noise source turning on.

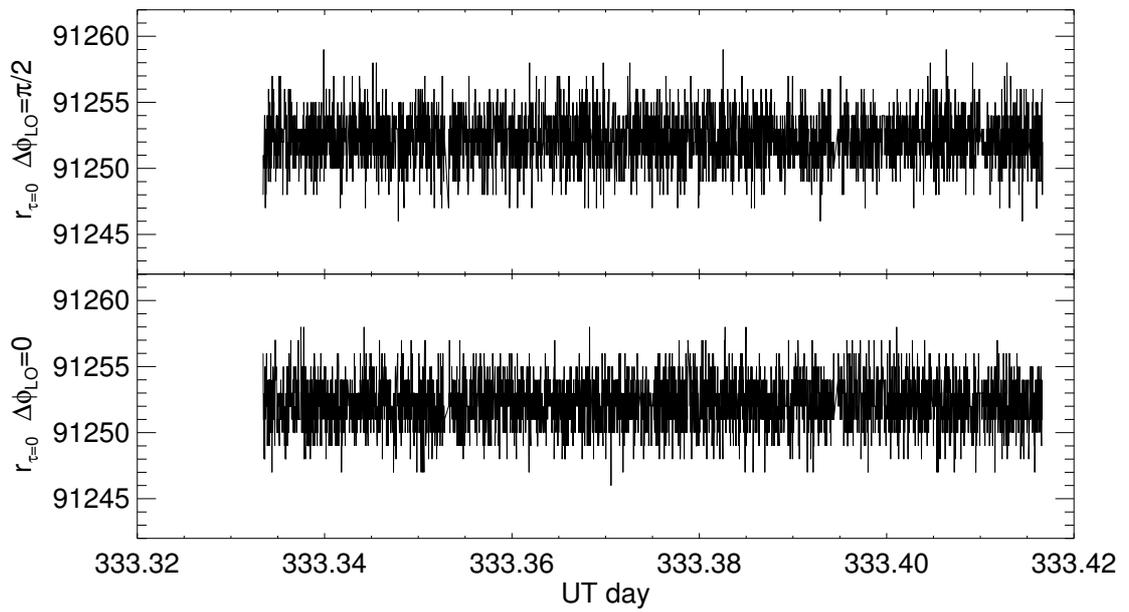
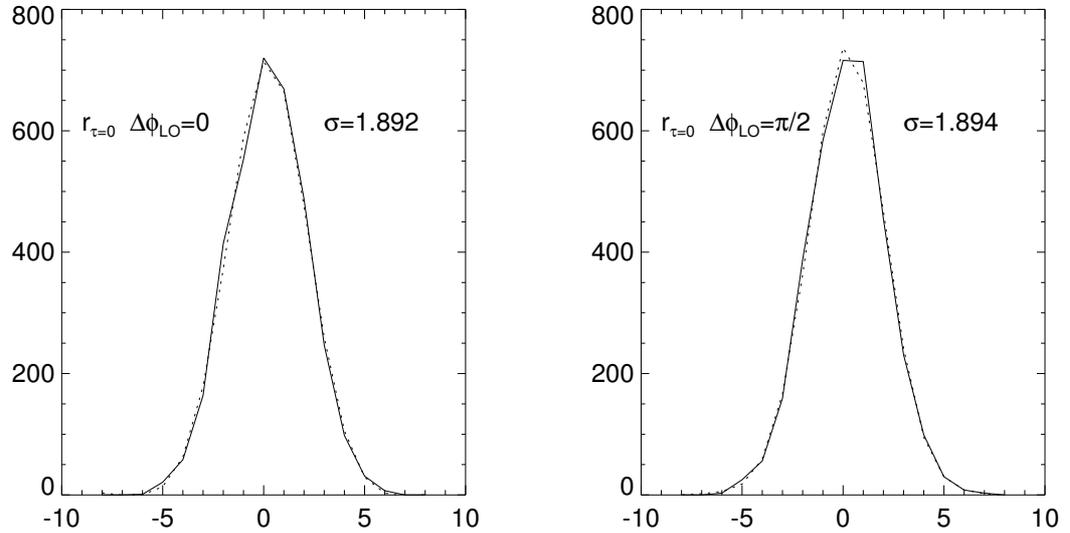


Figure 6.6: Correlator output with calibration spikes removed.

Figure 6.7: Histogram of the correlator output for two hours. The solid line is the actual histogram and the dashed line is a Gaussian fit. The x scale is centered on the value 91252. The σ that appears on the plot is the actual *rms* of the data, not from the fit.

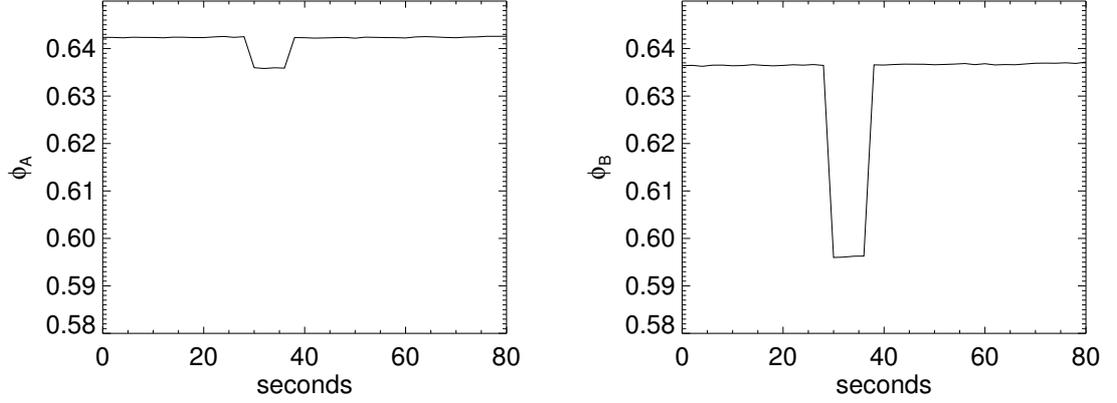


Figure 6.8: The fractional occupation of lower levels, ϕ , for a short time around a calibration spike, for two digitizers on receivers A and B.

noise source. A more accurate determination of the calibrator noise in each digitized signal comes from the threshold monitor.

Figure 6.8 shows how ϕ changes around the first calibration spike from Figure 6.5. The *rms* of the incoming signals change with the calibrator as:

$$\sigma = \frac{V_0}{\sqrt{2}} \frac{1}{\text{erf}^{-1}(\phi)}, \quad (6.6)$$

where erf^{-1} is the inverse error function, V_0 is the threshold voltage, which in this case is a constant 0.25 V set by the digitizer. The fraction of the power due to the calibrator is:

$$P = \frac{\sigma_{on}^2 - \sigma_{off}^2}{\sigma_{on}^2}, \quad (6.7)$$

where σ_{on} is calculated using Equation 6.6 when the noise source is on and σ_{off} is calculated when the noise source is off. To predict the expected correlation the geometric mean of the two powers is calculated to be:

$$\rho_{exp} = \sqrt{P_A \cdot P_B}, \quad (6.8)$$

where P_A and P_B are the application of Equation 6.7 separately for receivers A and B. For this particular spike in Figure 6.8, the expected correlation coefficient is $\rho_{exp} =$

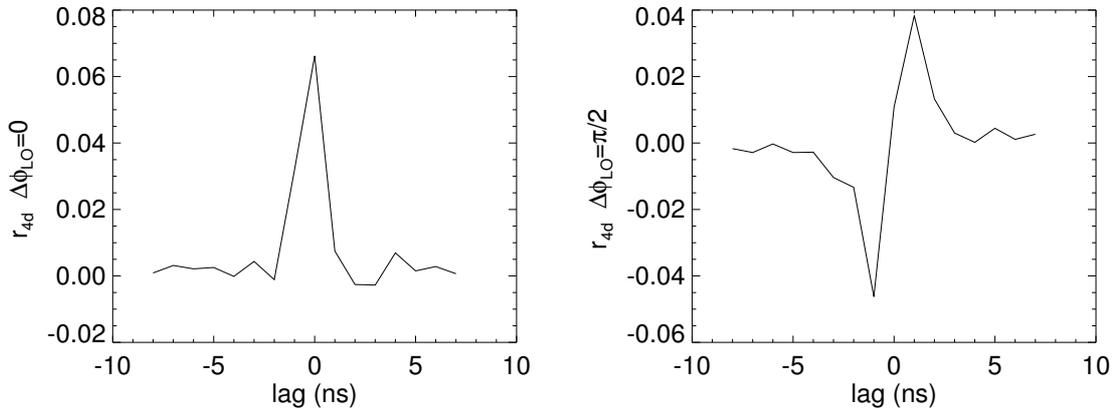


Figure 6.9: Correlator output for both phase states given as a function of lag. Data were taken with the noise source on.

0.063. To compare this with the actual correlator output, the raw correlator output is converted to the mean of the output for one multiplication, r_{4d} . This is derived from the correlator output r by first subtracting the bias 91252.2, then dividing by the number of accumulation cycles, found by taking the mean correlator output divided by the expected value for each multiplication, or $91252.2/3$.

Figure 6.9 is a plot of $r_{4d}(\tau)$ for both cases of the LO phase shift. The $\Delta\phi_{LO} = 0$ case may be thought of as the real part of a complex correlator and the $\Delta\phi_{LO} = \pi/2$ case as the imaginary. As can be seen, there is power in both sine and cosine channels at lag $\tau = 0$. The total correlation, $|\rho|$ comes from finding the magnitude using the two plots, and then multiplying by the correction factor, Equation 5.22. Figure 6.10 is a plot of the result of this operation. As can be seen from the solid line, lag $\tau = 0$ does not correspond to zero delay between the two incoming signals, where the power is presumably maximum. This implies that there is some uncompensated instrumental delay between the two receivers. To account for this, a reasonable guess for the shape of the correlation function, a Gaussian, is fit to the data. The peak of the Gaussian is then ρ . Comparing this with the expected value $\rho_{exp} = 6.3\%$, there is a deficit of 5%. The levels in all the baselines and correlators may be checked this way. Table 6.1 summarizes the results.

The results show that the measured correlated power is always less than the expected results by no more than about 15%. There are two possible reasons for

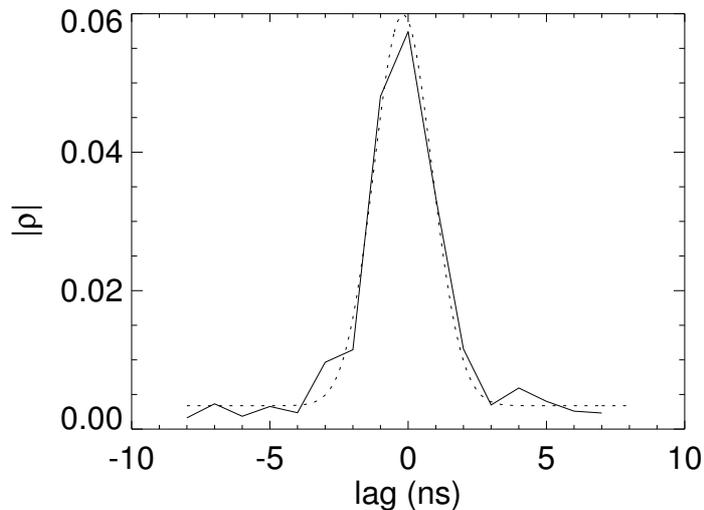


Figure 6.10: Maximum correlation interpolation. The solid line is the magnitude of the correlation function calculated from the data in Figure 6.9. The dashed line is a Gaussian fit, whose peak indicates the maximum correlation and the instrumental delay between calibration paths.

this. Non-flatness in the phase across the 500 MHz band is responsible for some loss in efficiency. Correcting for this with the inherent frequency resolution of a digital correlator brings the results about 1% closer to the expected values. The larger source of uncorrelated power is due to differences in the bandpasses between receivers.

6.5 Summary

The mean of the correlator output can be predicted to the 7th digit. The (rms) of the output differs from the prediction by $\sim 7\%$, with the difference due likely to the effective noise bandwidth being less than 500 MHz from the receiver through the channelizer.

The calibrator on/off tests show that the relative population of the inner samples, ϕ , can be used to monitor the power level entering the correlator. Specifically, the correlated output of two channels agrees well with the values predicted from the inputs.

correlator	baseline	ρ_{exp}	ρ_{meas}	% difference
1	1-2	0.046	0.044	4.0
	0-3	0.112	0.106	5.2
	1-3	0.082	0.078	4.7
	0-2	0.063	0.060	5.4
	0-1	0.113	0.109	3.6
	2-3	0.046	0.044	3.3
2	1-2	0.036	0.033	3.0
	0-3	0.077	0.069	10
	1-3	0.090	0.075	16
	0-2	0.031	0.028	9.6
	0-1	0.039	0.035	12
	2-3	0.070	0.059	15
3	1-2	0.055	0.052	4.7
	0-3	0.091	0.0074	92*
	1-3	0.128	0.116	9.1
	0-2	0.039	0.0029	93*
	0-1	0.082	0.0054	93*
	2-3	0.061	0.053	14
4	1-2	0.074	0.071	4.3
	0-3	0.059	0.054	8.9
	1-3	0.088	0.083	5.3
	0-2	0.050	0.043	13
	0-1	0.080	0.076	5.3
	2-3	0.054	0.049	10

Table 6.1: The expected and measured correlation coefficients. The baseline heading refers to the digitizer pair. (*)Note that digitizer 0 on correlator 3 was considered broken and malfunctioning. Therefore, any baseline involving this digitizer will not return the proper correlation coefficient

Deep thanks goes to David Hawkins, whose design was the inspiration for the MINT correlator. His help and suggestions were invaluable, both at the board level and at the FPGA code level. Special thanks also goes to Stan Chidzik, who provided circuit layout expertise and master soldering.

Interferometer Performance

7.1 Correlator Output Reduction

The data reduction procedure for a digital lag correlator is somewhat more complicated than the procedure for an analog correlator. The analog correlator, for one, has an output that is directly proportional to the correlation coefficient. In general, the output of a digital correlator must be corrected with some nonlinear factor to produce the correlation coefficient. Another issue is the frequency resolution of digital correlators. The MINT correlator is an X-F type, referring to the fact that the correlator does the “X” or multiplication in real time but requires an off-line Fourier transform, “F”, to be performed on the correlation function to produce the visibility. The MINT procedure has the added complication of double-sideband splitting. These data reduction issues are discussed in this chapter.

7.1.1 Correlation in Temperature Units

The fundamental data unit is referred to as a *record*. Each record contains 2 seconds of data. Within each record are four sequential readouts of all correlators with each readout containing an accumulation of the correlation function for 0.5 seconds. Each accumulation function also has a different relative LO phase state, $\Delta\phi_{LO}$, determined by the Walsh cycle. The relative phase between any pair of receivers passes through the four states: $\Delta\phi_{LO} = 0, \pi/2, 0, -\pi/2$, the exact order of which depends on the baseline. The first operation is to subtract the bias level and divide by the number of accumulations, as done in Section 6.4.4, resulting in the average correlator output for

one multiplication, r_{4d} . The first level of data reduction comes from averaging the two $\Delta\phi_{LO} = 0$ states and subtracting the $\Delta\phi_{LO} = -\pi/2$ case from the $\Delta\phi_{LO} = \pi/2$ case. This reduces the data set by a factor of two and produces two correlation functions per baseline for each data record: $r_{4d}(\tau)_{\Delta\phi_{LO}=0}$ and $r_{4d}(\tau)_{\Delta\phi_{LO}=\pi/2}$.

The next step is temperature calibration, which first involves correction for the threshold level. The threshold correction factor is given in Equation 5.22. The result of applying the correction factor is the correlation coefficient, $\rho(\tau)$. Multiplying the correlation coefficient by the geometric mean of the two system temperatures gives a temperature calibration. The system temperature changes with time, due mostly to the changing temperature of the atmosphere. This time-dependent system temperature may be constructed using the threshold-level monitor combined with the attenuator setting and the assumption that the overall system gain is constant. The *rms* of the signal coming into the A/D converter is Equation 6.6:

$$\sigma = \frac{V_0}{\sqrt{2}} \frac{1}{\text{erf}^{-1}(\phi)}.$$

In front of the A/D converter is the variable attenuator, which attenuates the signal by α dB. The power before the attenuator then goes as:

$$P \propto 10^{\alpha/10} \sigma^2. \quad (7.1)$$

This quantity changes as the power level into the attenuator changes. Under normal conditions, the power level changes either because the system temperature changes or the gain changes, assuming that the cosmic signal is small. Assuming that the gain remains constant, the current T_{sys} may be found by comparing this measured σ to a fiducial *rms* taken at time t_0 , when the system temperature is measured by other means:

$$T_{sys}(t) = T_{sys}(t_0) \frac{\sigma^2(t) 10^{\alpha(t)/10}}{\sigma^2(t_0) 10^{\alpha(t_0)/10}}. \quad (7.2)$$

The system temperature was measured with hot/cold load tests in the field. The details of this test are in Dorwart (2002), and a summary of the results appears in

Appendix C. The system temperatures were taken with the noise calibrator off during a cold night. Although they were not taken at this time, $t_0=334.1$ UT day is used as the fiducial time. The exact time chosen affects the overall temperature calibration but not the relative calibration. Since the system temperature is stable from night to night, the effect is small. A temperature-calibrated correlation function is given by:

$$r(\tau) = \sqrt{T_{sys1}T_{sys2}} \cdot \rho(\tau), \quad (7.3)$$

with T_{sys1} and T_{sys2} referring to the system temperatures of the two receivers involved. The quantity $\rho(\tau)$ is found by applying the correction factor in Equation 5.22:

$$\rho(\tau) = \frac{\pi}{2} \frac{1}{E_1 E_2 + E_1 + E_2} \cdot \frac{r_{raw}(\tau) - 91252.2}{91252.7/3}. \quad (7.4)$$

The quantity $E_j = \exp[-\text{erf}^{-1}(\phi_j)]$ must be calculated for each digitizer involved in forming the correlation function, and $r_{raw}(\tau)$ are the 18-bit correlation functions in the raw data stream.

7.1.2 Sideband Separation

With the temperature calibration applied to the two LO phase state cases, it is possible to perform the sideband separation procedure. Figure 7.1 outlines this process. The first step is to take the FFT¹ of both correlation functions, producing real and imaginary parts. Following the convention of Thompson et al. (1986), the following quantities are labelled:

$$\begin{aligned} f_1 &= \Re[\text{FFT}(r(\tau)_{\Delta\phi_{LO}=0})] \\ f_2 &= \Im[\text{FFT}(r(\tau)_{\Delta\phi_{LO}=0})] \\ f_3 &= \Re[\text{FFT}(r(\tau)_{\Delta\phi_{LO}=\pi/2})] \\ f_4 &= \Im[\text{FFT}(r(\tau)_{\Delta\phi_{LO}=\pi/2})], \end{aligned} \quad (7.5)$$

¹or, regular discrete Fourier transform, as the case may be. The number of lags was chosen to be a power of 2, allowing the use of the FFT.

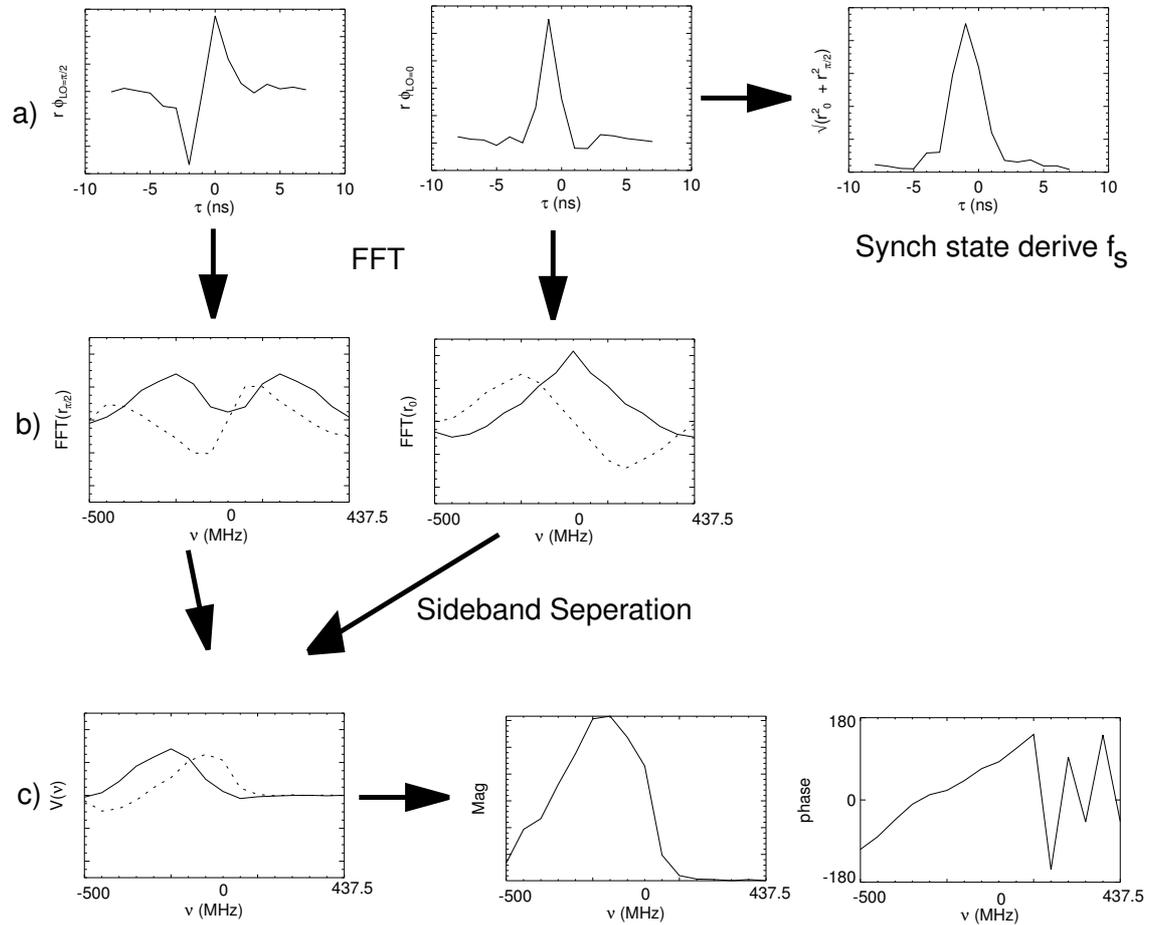


Figure 7.1: Data reduction procedure. The data are shown for a single data readout during a calibration spike. a) The data comes out of the correlator as a correlation function. There are two states, r_0 and $r_{\pi/2}$ which come out sequentially. The modulus of the two functions is used to derive the synchronization state. The location of the peak of $\sqrt{r_0^2 + r_{\pi/2}^2}$ is the time delay, f_s , which is applied to correct the phase. b) The FFT of the two phase states. The solid line is the real part and the dashed is the imaginary. c) The FFT's of the correlation function are used to generate the Visibility function, both real (solid line) and imaginary (dashed line) parts. The real and imaginary parts are used to generate the magnitude and phase. The phase here has not yet been corrected for the slope. From the magnitude plot it is clear that there is only power in one of the sidebands. The lack of calibrator power in the upper-sideband has been verified in the lab (Aboobaker 2002).

with the script letters referring to the real or imaginary part. The upper and lower sideband visibilities may then be constructed as:

$$\begin{aligned} V_{upper} &= f_1 - f_4 + i(f_2 + f_3) \\ V_{lower} &= f_1 + f_4 + i(f_2 - f_3). \end{aligned} \tag{7.6}$$

7.1.3 Correlator Synchronization

The flow chart in Figure 7.1 adds one extra step, which involves the synchronization state of the correlator. An ambiguity in the resetting of the digitizers leads to an arbitrary lag of -1,0, or +1 ns between digitizers (see Section 6.3.2). To trace the delay, the sum of the squares of the 0 and 90 degree states (before transforming and sideband separating) are produced. In the synchronized case, the peak of this function should be at 0 lag. The actual integer location of this peak is taken to be the relative delay, f_s . A shift in time appears in the visibilities as a slope in the phase, clearly seen in Figure 7.1 row c. Subtracting $2\pi f_s$ from the phase corrects for this slope. See Section 7.2.2 for details.

7.2 Stability

The two most important aspects of system performance are stability and sensitivity. The system sensitivity is established primarily by the overall system temperature and secondarily by other effects such as telescope alignment. The best way to measure the sensitivity is by astronomical calibration, such as observations of planets. The sensitivity issues are addressed in Dorwart (2002).

The stability of the system is measured in two ways: first, by monitoring the noise source calibration spike, and second, by performing “integrating down” tests. There are two aspects of monitoring the calibration spike, both phase and amplitude. The amplitude of the calibration spike traces three coupled quantities: the temperature of the calibrator, the system gain, and the system temperature. There are, however, complications with regard to separating these quantities. This is discussed in the next section.

signal path is described in Section 3.2.

Things that affect the amplitude stability include coupling between components, reflectivity of the optical surfaces, SIS conversion efficiency, and amplifier gain fluctuations. Of these, the amplifiers are the largest contributor to gain changes. There are five sets of amplifiers that process the noise source signals. Four of the amplifiers are also involved in measuring signals from the sky and come after the SIS downconversion. All amplifiers are temperature controlled, some more rigorously than others. The least well controlled amplifiers are in the correlator. Although the temperature of the amplifier was not monitored directly, the temperature of the nearby digitizer was. Over one night, the temperature varied by $\sim 5^\circ$.

The correlator is not in itself sensitive to changes in the gain in any of the four IF amplifiers. The reason is that the correlator only returns a correlation coefficient. The correlation coefficient is a measure of the percentage of the incoming signal that is correlated. If there is a gain change in any of the IF amplifiers, both the calibrator signal and the noise from the receiver will increase in the same proportion, leaving the correlation coefficient unchanged. In the temperature calibration procedure presented in Section 7.1.1, a change in the IF gain would be interpreted as a change in the system temperature, even though the system temperature might be constant. Through Equation 7.3, the calibrator would then also be assigned a changing temperature, even if its strength is the same. This effect is difficult to separate from changes in the output temperature of the calibrator, making external calibration via a stable astronomical source necessary to track the gain.

Figure 7.3 is a plot of the calibrator temperature reduced using the prescription outlined in Section 7.1. The top panel is a plot of the data from all correlators in the D-A baseline for one night. The middle panel shows the noise switch encoder position and the bottom panel is a plot of the outside temperatures. The sunrise time can be clearly deduced from the outside temperatures. There is a hint of sunset at the beginning of the plot. All subsequent plots will show data from only the stable part of the evening, from roughly UT midnight to 0.425 UT day.

From the encoder plot it is also clear that the switch did not come to the same position on every turn-on, with results reflected in the data. By calibrating for the encoder position it is possible to remove the effect. To correct for the encoder position,

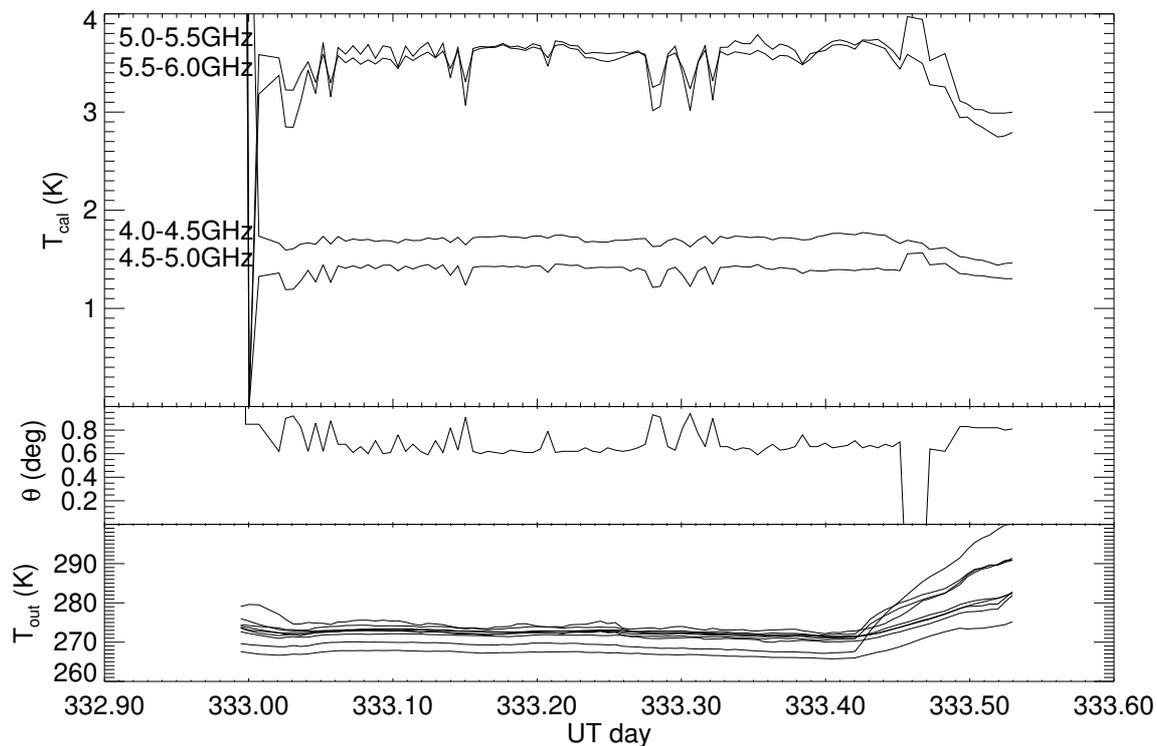


Figure 7.3: Temperature calibrated calspike power from the correlated baseline D-A, from correlator 1. The power is summed over frequency. The middle panel shows the encoder position of the noise switch and the bottom panel displays the readings of temperature sensors that are mounted on the outside of the telescope.

it is assumed that there is some multiplicative function of θ , the encoder position, which modifies the power relative to some fiducial encoder position θ_0 ($\theta=0$ has no significance). The form of the correction is:

$$T_{cal} = \frac{T(\theta_0)}{Q(\theta)} * T_{meas} \quad (7.7)$$

where $Q(\theta)$ is the multiplicative function. To find this function, refer to Figure 7.4, a plot of calibrator temperature vs. noise encoder position (for positive positions only). The solid lines are quadratic fits to the distributions, which are taken to be the multiplicative correction. Figure 7.5 is a plot of the same data as in Figure 7.3, with the correction in Equation 7.7 applied along with a cut for the stable part of the night.

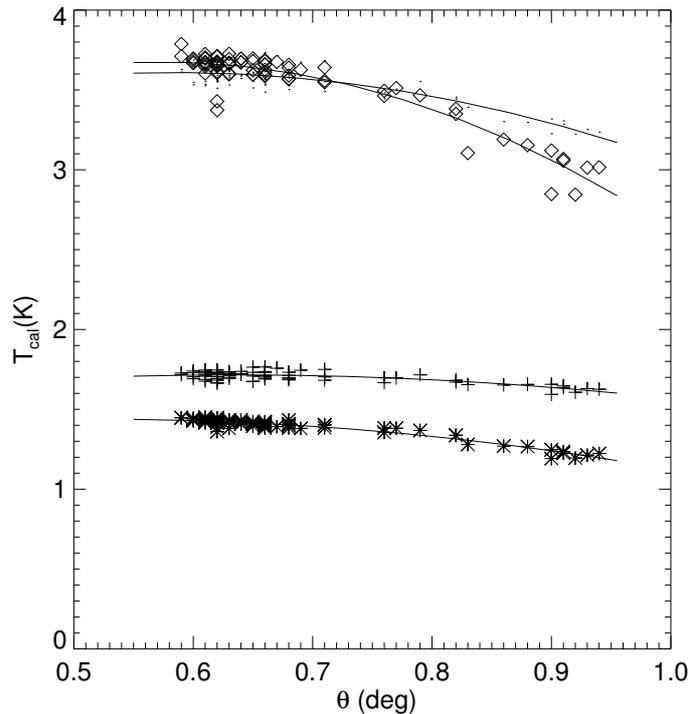


Figure 7.4: Calibrator signal vs. switch encoder position. The data are a sub-sample of the data plotted in Figure 7.3. The data were cut for the stable part of the night and for obvious outliers. Each symbol represents a different correlator in the D-A baseline. The solid lines show quadratic fits.

The quantity of interest is the relative temperature of the calibrator. Figure 7.6 is a plot of the calibrator temperature relative to a fiducial value at UT day = 333.17. It can be seen that within this one night, the calibrator is steady to within a few percent. The figure shows the relative calibrator strength for all baselines and correlators. From the plots, it is clear that there is some common component to all traces. A common fluctuating gain is likely to come from a localized source, such as in the noise source itself or from gain fluctuations in the channelizer or correlator. Gain fluctuations in amplifiers usually trace temperature variations. The temperatures of the channelizer and correlator, however, do not follow the gain variations, so it is likely that the source of the common mode fluctuations is in the noise source.

Figure 7.7 is the same as Figure 7.6, but extended to show the entire season. Again, there appears to be a strong common component among all baselines and

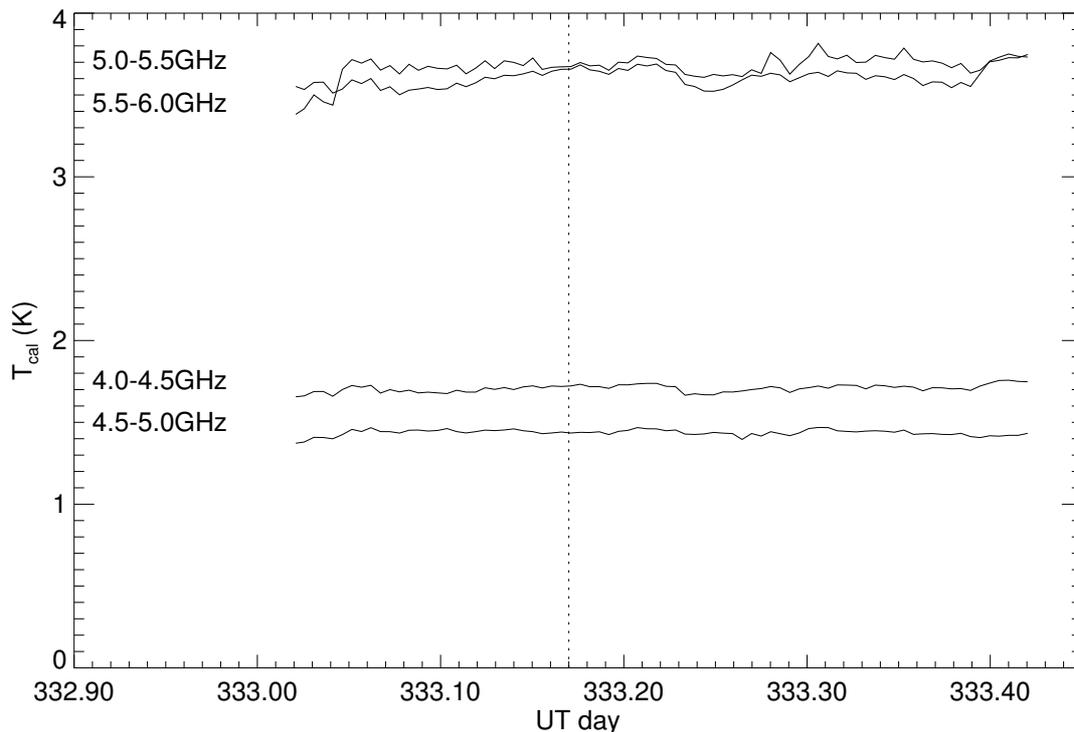


Figure 7.5: Calibrator signal corrected for noise switch encoder position. The values are corrected to the signal at $\theta = 0.55$. The dashed line represents the fiducial time, t_0 . Figure 7.6 is generated by comparing $T_{cal}(t)$ to this time.

correlators up to about UT day 340, when the site experienced heavy snows and the telescope was frequently tarped. The somewhat abrupt changes in common components between baselines and correlators may be attributed to possible disturbances to the waveguide that couples the calibrator to the optics. If the dominant variation in relative calibrator strength is due to the calibrator itself, the inherent variation in the amplitude stability of the interferometer is smaller than the variation in Figure 7.7, and the calibrator should not be used to correct for the gain.

7.2.2 Phase Stability

An equally important aspect of interferometer stability is the phase stability. Phase stability refers to how well the relative phase between two receivers is maintained. A night-to-night fluctuating phase leads directly to a decrease in sensitivity. For sources

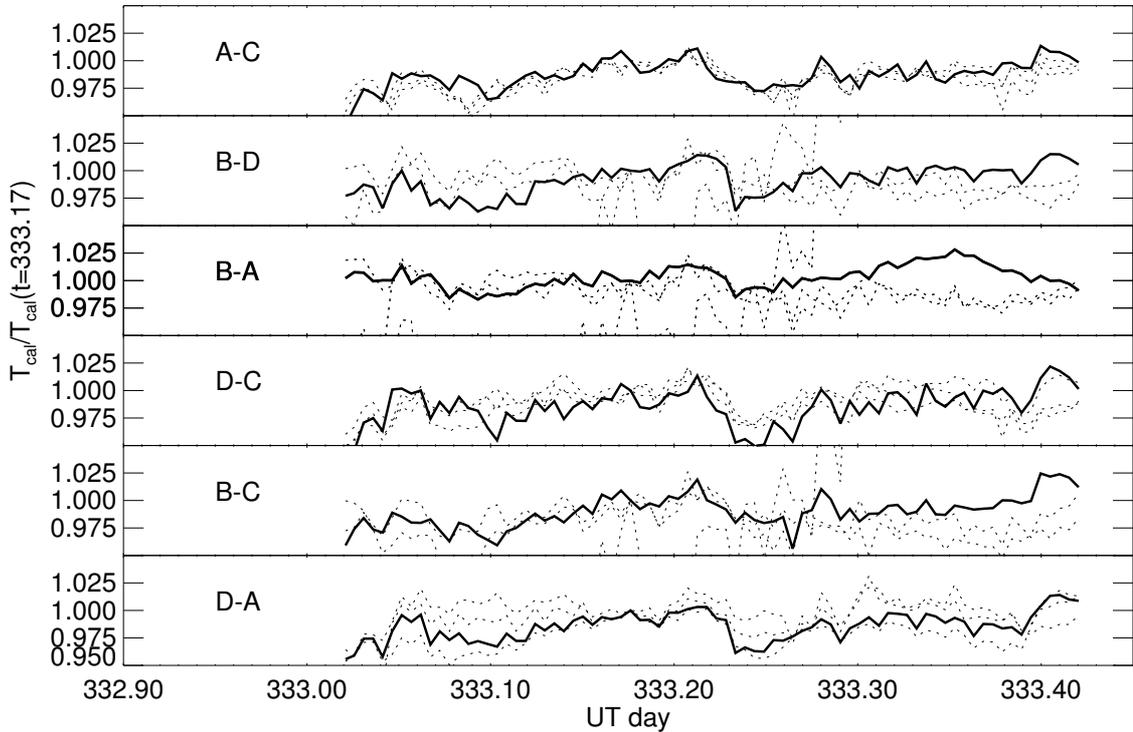


Figure 7.6: Relative calibration signal. Each panel shows the relative calibrator strength in a specific baseline. All panels are limited to $\pm 5\%$. The strengths are shown relative to the strength at $t=333.17$ UT days. The heavy lines are for correlator 1 (4.0-4.5 GHz), and the dashed lines are for the rest. Panels that show a baseline involving dewar B include one defective digitizer as indicated by the “out of family” dashed line.

on the sky, changing the phase confuses the real and imaginary parts of the visibility.

There are many physical sources of phase variations in an interferometer. The primary source is from differential temperature changes, which lead to differing path lengths from the source to the correlator. The path length includes both a high frequency (RF) and low frequency (IF) component. The RF component includes everything before the SIS mixer, including the optical components. The IF components include the RBE components, the long (~ 1 m) coax cables that lead to the channelizer, the channelizer itself, and the transmission lines that lead to the correlator. A length change in the RF (150 GHz) has 30 times the effect in phase as an equal change in length at the IF (5 GHz) due to the much higher frequency.

Other sources of phase instability include the LO phase that feeds the SIS. The

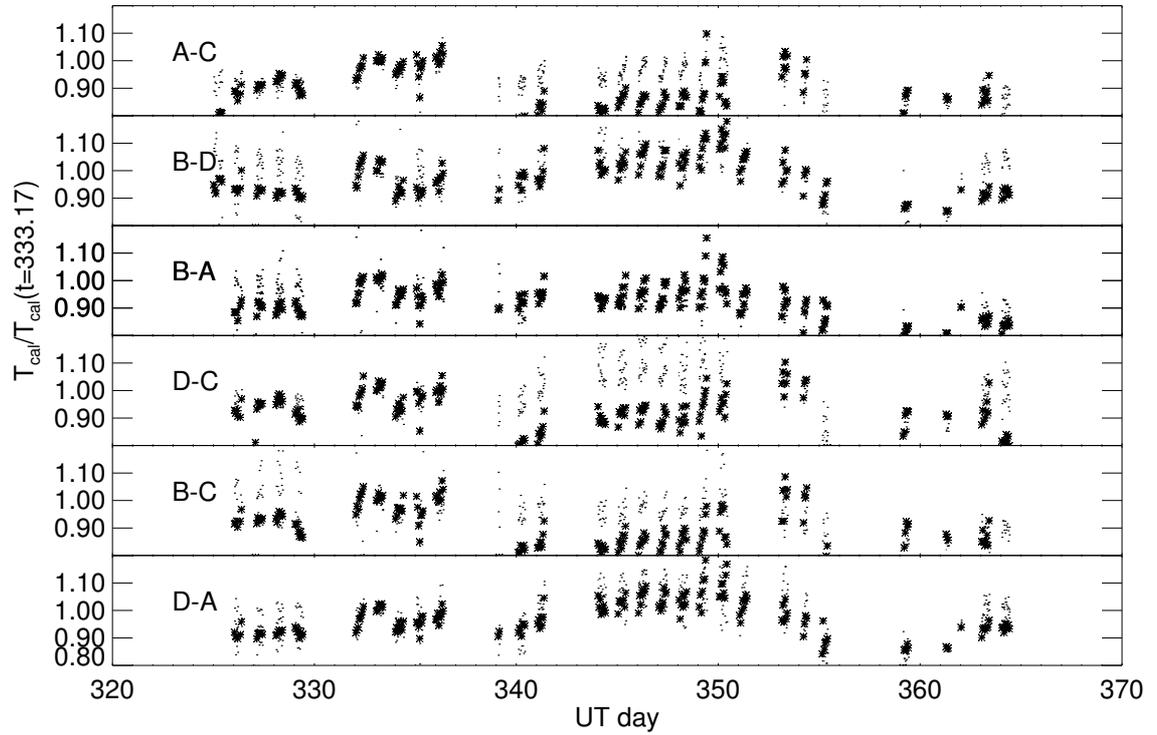


Figure 7.7: Relative Calibration signal for the entire season. Each panel shows the relative calibrator strength in a specific baseline. All panels are limited to $\pm 20\%$. The strengths are shown relative to the strength at $t=333.17$ UT days. The asterisks represent correlator 1 (4.0-4.5 GHz) and dots represent the rest. Only every 10th calibration spike is shown for clarity.

LO phase is controlled by the PLL. The largest source of instability is from diurnal thermal cycling of the 12 GHz 1 m coax lines that feed each receiver. Again, a change at the RF has a tenfold effect as a change in the PLL lines².

The relative phase on the sky is also monitored with the noise source. Using the calibrator introduces a large source of instability that is not present in data from the sky. Relatively long (~ 1 m) waveguide is used to route the broadband 150 GHz noise signal to all four antennas. The radiation transmitted through the waveguide is then coupled to the receiver with a small pickoff mirror located near the horn. The pickoff

²In analog correlators, there is an additional internal stability issue due to changing transmission line lengths. With a digital correlator, the phase stability depends on the stability of the reactive coupling at the input the digitizer. This is less dependent on temperature than transmission line lengths. Analog correlators will also suffer similar phase fluctuations due to reactive coupling at the inputs of any electronic component

mirror reflects radiation from the waveguide to the secondary, and eventually to the feed horn. Since they carry high frequency signals, differential thermal cycling of the 1-meter waveguide or changes in any optical dimension lead to the dominant term in the phase stability of the calibrator signal.

The calibrator is not sensitive to phase changes that affect the locations of the antennas while leaving the relative lengths of the waveguide fixed. Flexure in the telescope frame due to temperature changes or gravity may not be reflected in the calibrator signal, but would be seen in astronomical sources. Gravity effects are limited by observing only in a small range of altitude directions.

Phase Derivation

The relative phase of the calibrator between the two receivers is a product of the reduction procedure outlined in Figure 7.1. There are two important caveats. First, the synchronization state of the correlators must be removed to compare phase between nights. This step is discussed in Section 7.1.3. Second, the phase is not flat across the bands, so there is not a single number that describes the phase state.

The procedure to extract the phase is outlined in Figure 7.8. The sideband separation procedure returns the magnitude and phase for 16 frequencies across both upper and lower sideband. The solid line in the right panel represents the raw phase output with the calibrator on. The phase shows an obvious slope due to a time shift in the lag domain. The first step is to apply the correction described in Section 7.1.3, which brings the phase to the flatter dotted line. This line shows the residual non-flatness in the phase across the band³. The next step is to subtract a phase profile taken at some fiducial time. In this case, the fiducial time is the same as that for the amplitude, UT day=333.17, which for Figure 7.8 is ~ 1 hour away. The resultant phase profile is generally very flat, and can be well fit to a line.

It is clear from Figure 7.8 that the noise source is much stronger in one sideband than the other⁴, resulting in a poor determination of the phase in the low-power upper

³The non-flatness is not surprising given the complicated geometry of the radiation paths from the waveguides to the feedhorns. If the calibrator were a point source in the far field, one would measure a flat phase or the instrumental phase. Not only are the mirrors different sizes, but also subtle differences in the alignment of the pickoff mirror can lead to large differences in phase.

⁴The cause of the power imbalance was traced to the Q-band amplifier in the noise source. The gain of the amplifier falls off at 47 GHz. Upon upconversion in the tripler, the power falls

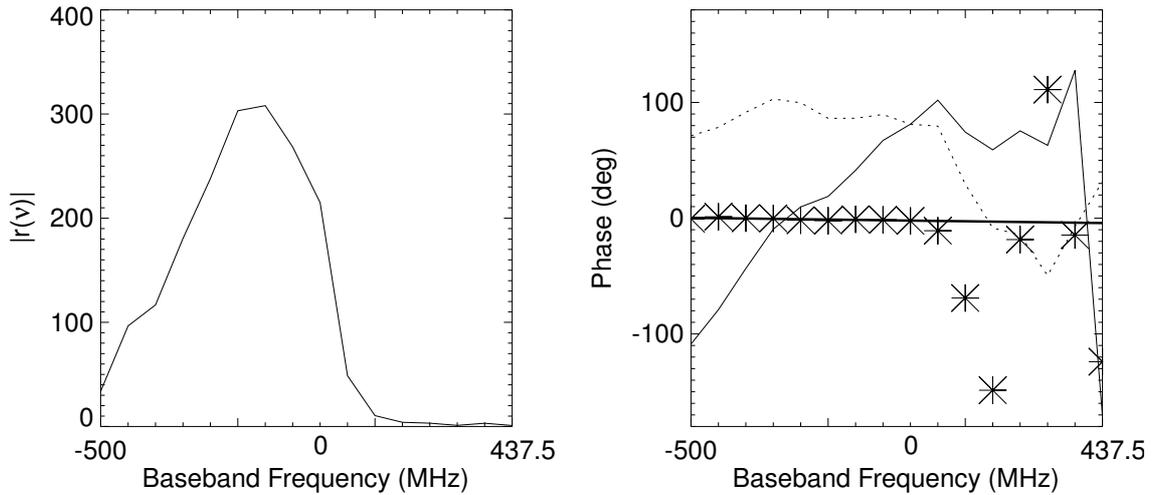


Figure 7.8: Phase correction technique. The left panel shows the visibility magnitude of a single calibration spike, separated into upper and lower sideband. The fact that the signal appears mostly in one sideband leads to poor determination of the phase in the other sideband as evident from the solid line in the right panel. The first level of correction comes from subtracting an integer phase slope, f_s . The dashed line represents the result. The next step is to subtract a fiducial phase, which leads to the asterisks. Finally, a line is fit to one of the sidebands. The phase offset indicates a change in high-frequency path length, whereas a slope implies a change in either the IF or RF path lengths.

sideband. For the linear fit only the sideband with signal is considered. One sideband is sufficient for determining the phase stability.

Phase Stability Characteristics

To assess the phase stability of MINT the linear fit to the difference between the current and fiducial phase profile is monitored. Because the phase profile difference as a function of frequency is so close to a line, it reduces the number of parameters to monitor from 16 (a separate phase for each frequency) to 2, both slope and offset. Changes in slope and offset indicate two different effects. Generally speaking, differential changes in transmission line lengths lead to changes in the time lag between

off at 141 GHz, leading to diminished power in the upper sideband (Aboobaker 2002). This also implies that the amplitude stability tests from the previous section only apply to the lower sideband, although there is no reason to believe that it will be much different for the upper sideband.

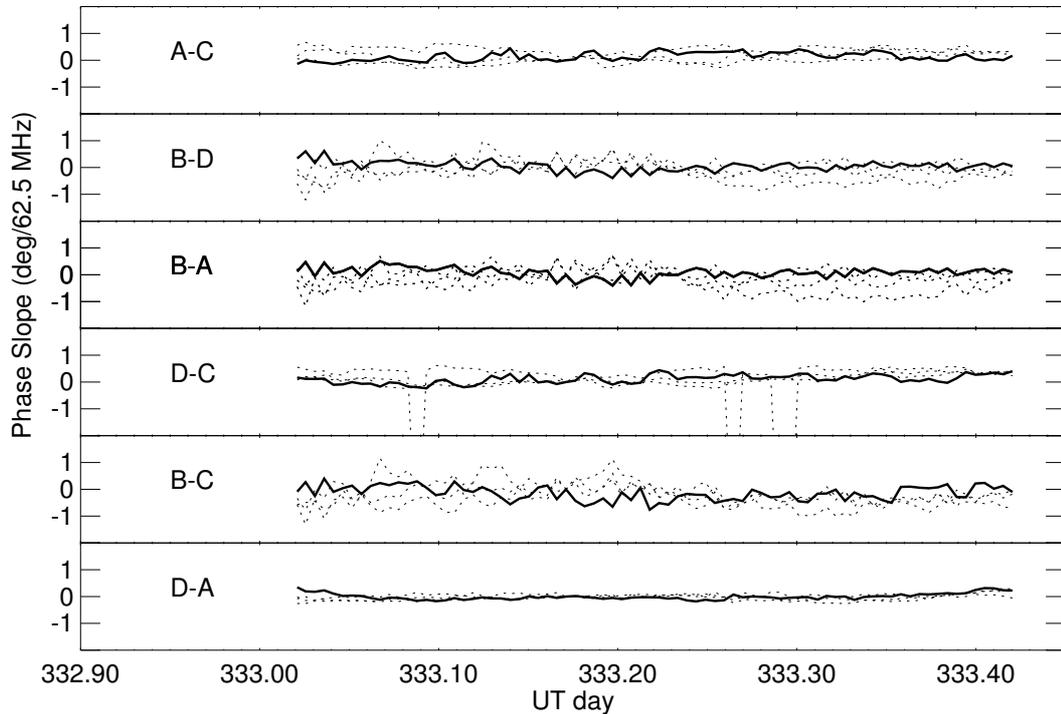


Figure 7.9: Phase slope for one night. The solid line is Correlator 1 and the dashed lines are the rest. The high degree of flatness implies that there are no length changes at the baseband frequency.

receivers, which is manifested as a phase slope. For a given change Δl , the corresponding slope in phase is $2\pi\Delta l/c$, independent of frequency. It is exceedingly difficult to cause large changes in phase slope with differences in line length. A $1^\circ/62.5$ MHz phase slope corresponds to roughly 1 cm of optical path length difference, a change that is very unlikely to be caused by thermal length changes alone. These facts are borne out in Figure 7.9, where the phase slope for one night is plotted. The slope is within $1^\circ/62.5$ MHz for the night. Figure 7.12 shows the slope for the entire season fluctuating at the $5^\circ/62.5$ MHz level. Changes this large cannot be accounted for by length changes but are possibly due to slight changes in the alignment of the noise source waveguide with the rest of the optics.

Unlike the phase slope, the phase offset is sensitive to changes in length in the 5 GHz IF and very sensitive at the 150 GHz RF. A 1° phase offset requires a 0.3 mm optical path difference at 5 GHz and 30 times less or $10 \mu\text{m}$ at 150 GHz. Phase

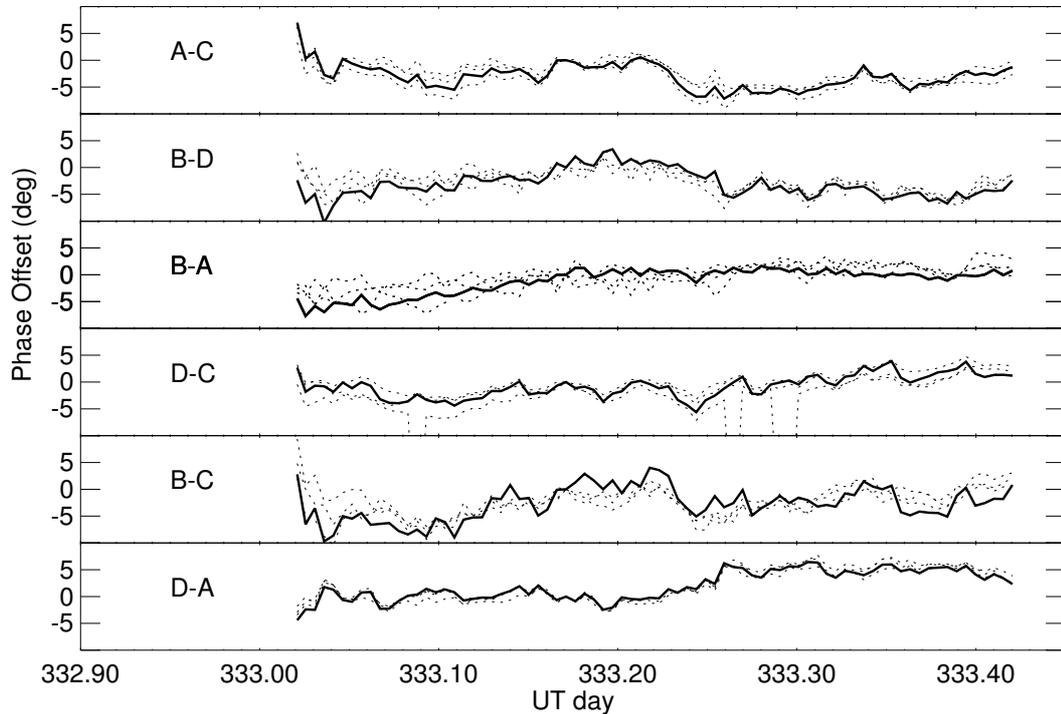


Figure 7.10: Phase offset plotted for a single night. The solid line is Correlator 1 (4.0- 4.5 GHz) and the dashed lines the rest. The fact that all correlators for a given baseline follow each other implies that the phase offset drift originates at either the SIS IF (4-6 GHz) frequency, or, more likely at RF (~ 150 GHz). A cut at 333.01 and 333.425 UT days has been applied to remove twilight and morning, when the telescope is undergoing rapid temperature changes.

offset fluctuations may also be caused by errors in the PLL circuits that control the 150 GHz LO's. There is yet another PLL that controls the phase of the channelizer LO's at both 4.5 and 5.5 GHz. The PLL's are in a regulated thermal environment and operation of the PLL's in the lab, however, indicate that all of them are more stable than the observed fluctuations.

Figure 7.10 is a plot of the offset for one night. Barring any coincidental cancellations between changes in the noise source waveguide and the rest of the radiation path, the interferometer appears to be very phase stable within one night of observing.

The picture is different, however, when the entire season is plotted, as shown in Figure 7.11. Along the right hand side of the plot is the *rms* variation of the phase offset for the entire season. The values are given for correlator 1 but are approximately

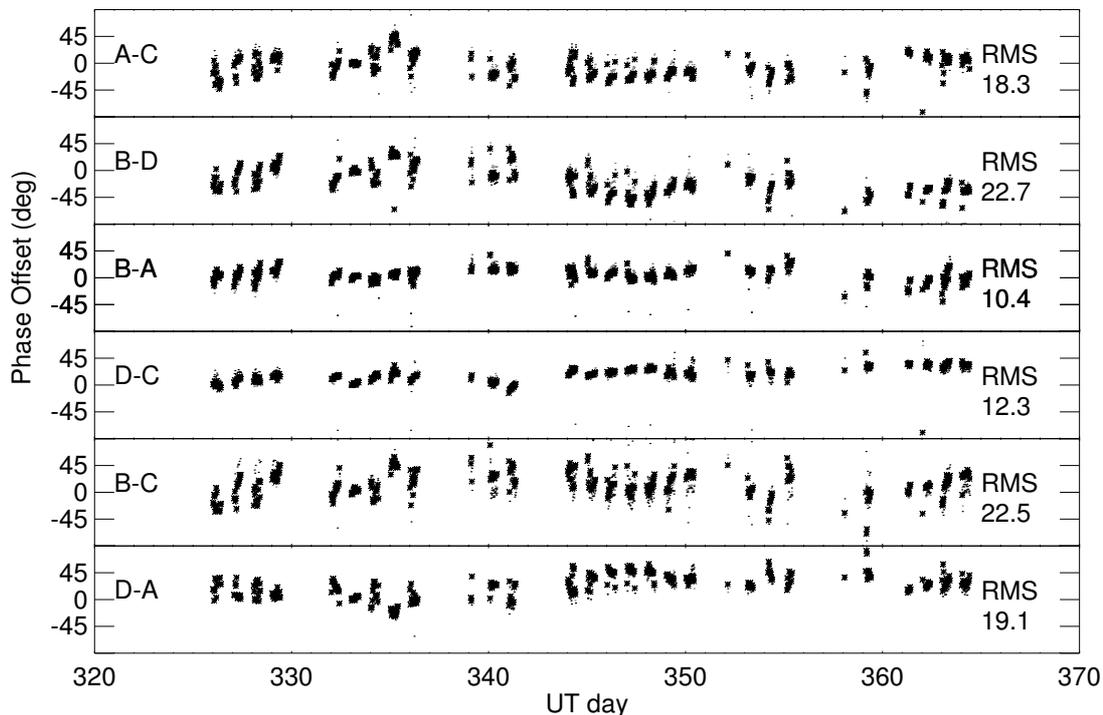


Figure 7.11: Phase offset for the entire season based on the noise source. The small asterisks are correlator 1 and the dots are the rest. Only every 5th point is plotted.

true for all correlators. There is a strong correlation between correlators for a single baseline. This can only happen only if the source of the fluctuation occurs before the signals are split at the channelizer. Since length changes at RF are 30 times as effective as changes at IF, it is likely that the length variations happen in RF.

It is now left to determine if the length variations are dominated by components involved in measuring the sky, that is, in the antenna to SIS RF chain, or if the variations are dominated by the noise source and waveguide distribution alone. As mentioned earlier in Section 3.2.3, a 5 K differential temperature change in the 1 m waveguide will cause a $\frac{\lambda}{20}$ phase difference. The optics add extra optical path, which should increase the phase changes. After leaving the waveguide, the radiation is reflected from the secondary, which is about 10 cm away for the small optics, and then back into the horn, for a total round trip of 20 cm. For the larger optics, this distance is 1.5 times larger, or about 30 cm. To compensate for this difference in

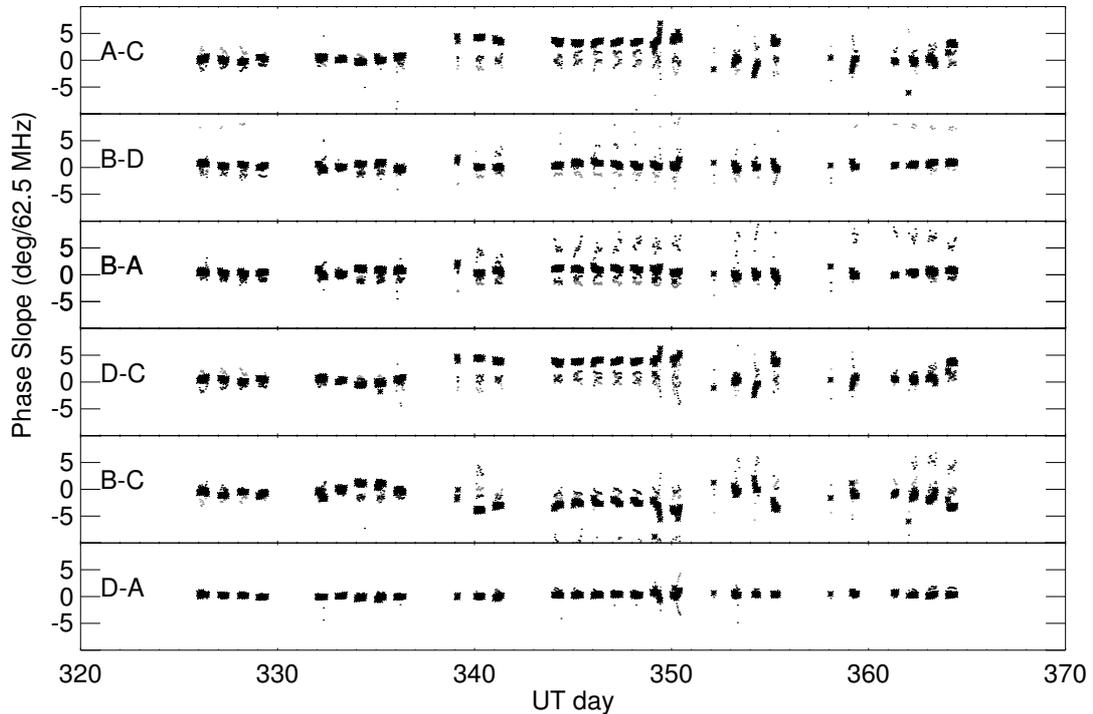


Figure 7.12: Phase slope for the entire season based on the noise source. The small asterisks are correlator 1 and the dots are the rest. Only every 5th point is plotted.

path for the differently sized antennas, an extra ~ 10 cm section⁵ was added to the waveguide leading to the smaller mirrors. The optics are composed of aluminum, and the struts that hold the secondary are composed of G10 fiberglass-epoxy. The thermal coefficient of expansion of both aluminum and G10 are very similar to that of silver. The optics, therefore, have about 30% of the effect of the noise source waveguide in producing phase variations due to differential thermal expansion.

Even if it is assumed that the phase fluctuations are due to temperature variations, it is difficult to distinguish between changes in the optics and changes merely in the waveguide because both components contribute roughly equally to the phase fluctuation. They may, however, be separated by astronomical calibration, which involve only the optics. Tracking the phase with daily Mars observations is presented

⁵Corrected for the difference in optical path in waveguide vs. air. Waves travel slower in waveguide than in air. To match the phase between calibrator signals at the inputs to the receivers, the longer path in air must be matched by a longer path in waveguide. The ratio of air distance to waveguide distance is 6.355 cm:8.932 cm (Aboobaker 2002).

in Dorwart (2002).

Baseline-Baseline Correlation

There also seems to be a strong correlation between the offsets measured in baselines A-C, B-D, and B-C, which are anti-correlated with D-A. These phase offsets are small and do not seem to be correlated with any other readouts on the telescope, including waveguide temperatures. Baselines B-A and D-C do not appear to be correlated with any of the other baselines. It is possible that this is due to the optical path in receivers B and A changing in common mode while D and C are stable or vice-versa. This would explain the correlation in the other baselines while minimizing the phase offsets in the baselines B-A and D-C.

Effect of Phase Fluctuations on Sensitivity

A phase fluctuating over the season will affect the overall sensitivity. To estimate the effect, consider a point source at the phase center of a baseline, so that the relative phase from two receivers is zero. The visibility in this case is all real. Inserting a relative phase θ into the baseline mixes the real and imaginary parts. The amplitude of the visibility is now:

$$|V|^2 = \cos^2 \theta + \sin^2 \theta \quad (7.8)$$

If the magnitude of the visibility were calculated before θ was allowed to change, there would be no loss in sensitivity. The problem, however, is that instead of averaging the magnitude, the real and imaginary components are averaged separately. To evaluate this case, a more sophisticated version of the equation above must be used:

$$|V|^2 = \left[\int_{-\infty}^{\infty} \cos \theta P(\theta) d\theta \right]^2 + \left[\int_{-\infty}^{\infty} \sin \theta P(\theta) d\theta \right]^2, \quad (7.9)$$

where $P(\theta)$ is the probability density of θ . If θ is taken to be normally distributed,

$$P(\theta) = \frac{1}{\sqrt{2\pi} \theta_{rms}} \exp\left(-\frac{\theta^2}{2\theta_{rms}^2}\right). \quad (7.10)$$

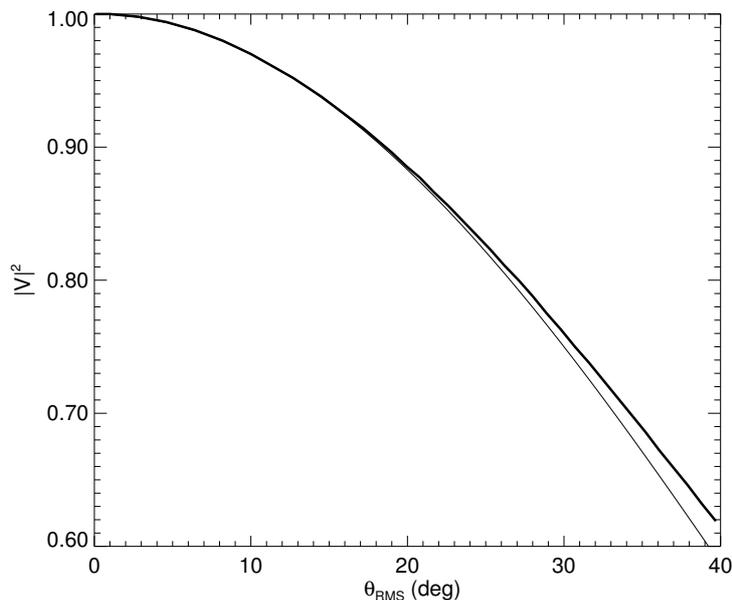


Figure 7.13: The decrease in sensitivity as a function of the *rms* fluctuations in phase offset. The thick line is a plot of Equation 7.9, and the thin line is a plot of $\cos^2(\theta_{rms})$, showing that a cosine rule of thumb applies for this range of angles.

Equation 7.9 may be evaluated numerically to obtain Figure 7.13. As can be seen from the figure, a θ_{rms} of 20° leads to a reduction in sensitivity of about 10%. Even if the entire fluctuation seen in phase offset over the season (Figure 7.11) is due to thermal changes in the optics, not correcting for phase will introduce only a small error.

7.2.3 Integrating Down

The final test of interferometer stability is the integrating down test. It is a crucial test of how stable the mean of the data is. It is performed by subdividing the data set, taking the mean of the pieces, then finding the *rms* of the means. Figure 7.14 shows the results of this test for one night. The data were first temperature-calibrated and sideband-separated using the process described above. The correlation functions were then synchronized according to the prescription outlined in Section 7.1.3 using the synchronization state from the nearest calibration spike. The form of the *rms* used is:

$$rms = \sqrt{\frac{1}{N-1} \sum (x_i - \bar{x})^2}. \quad (7.11)$$

The length in time of each subdivision is the abscissa. The first point is for 1 second, and the rest are powers of 2 times 1 second. The solid thin and dotted lines show the respective results for the real and imaginary parts of the visibility for correlator 1 and baseline A-D. The dots represent the real and imaginary parts for the rest of the data. The scatter in the data increases as the integration time becomes longer. This is a result of the error in determination of the *rms*, which occurs as a result of fewer samples.

The Dicke equation for either the real or imaginary part of the correlation (Section 2.2.2) is:

$$\delta T = \sqrt{\frac{T_{sys1} T_{sys2}}{2\Delta\nu t}}. \quad (7.12)$$

On a log-log plot, this appears as a straight line with a slope of $-1/2$ and an offset determined by the sensitivity. Inserting 40 K into T_{sys1} and T_{sys2} yields the thick solid line in Figure 7.14. The fact that the data follow the slope of this line indicates that the correlation functions are indeed stable for one night. If they were not stable, there would be excess fluctuation power, which would show up in the plot as a rise in the data at the timescales of the fluctuation. It should be noted however, that this test cannot be used as a system temperature calibration for the data, as it uses an input system temperature from the hot/cold tests.

7.3 Conclusion

Given the system temperature measured via the hot and cold tests, the estimates from Chapter 2 imply that MINT should have the sensitivity to measure the fluctuations near the damping scale. The only potential obstacle to detection is stability; in terms of amplitude and phase of the noise source, the MINT system is shown to be sufficiently stable. Also, given the integrating down tests, the mean of the interferometer is stable on one-night timescales. These observations bode well for the final analysis of the MINT data set.

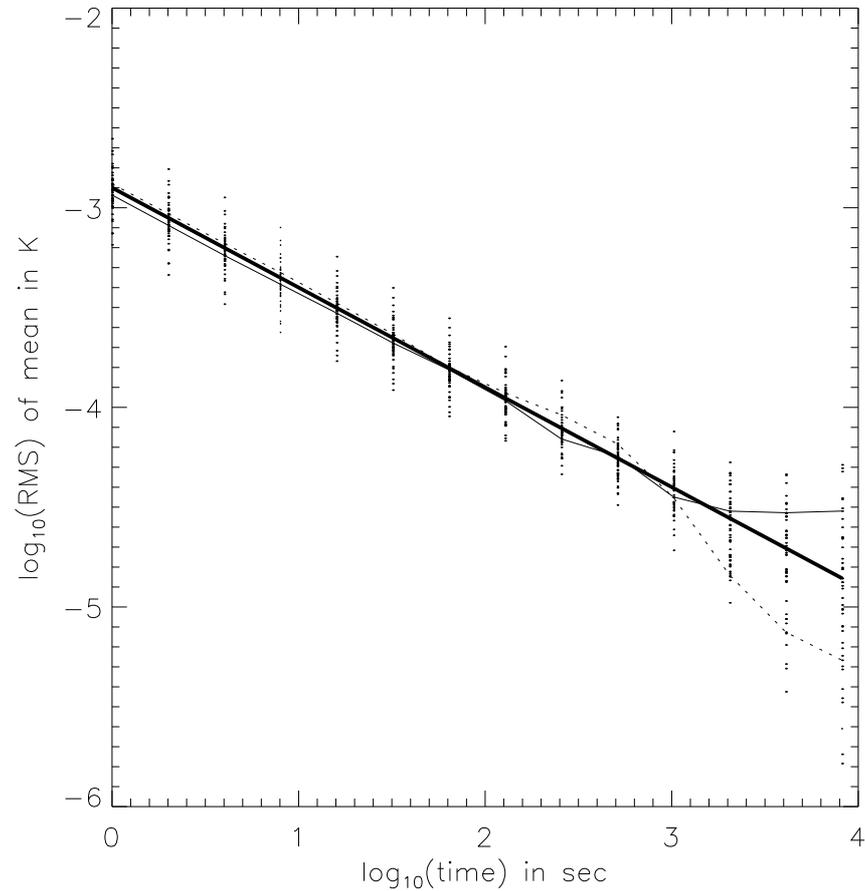


Figure 7.14: Integration down plot. The data are taken from one night, UT day 333. The thin solid line is for the real part of the visibility for baseline A-D on correlator 1 and the dashed line is the imaginary part. The dots represent the rest of the baselines and correlators, both real and imaginary. The thick straight line is the function $\delta T = \sqrt{\frac{T_{sys1}T_{sys2}}{2\Delta\nu t}}$, where both T_{sys} are 40K and $\Delta\nu$ is 500 MHz. This radiometer equation is appropriate for either the real or imaginary part of a correlation receiver (see Section 2.2.2).

The next step in the analysis is to integrate the data into sky bins and look for excess variance that is not explainable by receiver noise. Any excess is presumably due to cosmic signal. A full analysis, however, must include correlations in the data samples, which are processed through likelihood analysis.

The short-term future for MINT may include a second season, pending the results of the initial analysis. A second season will decrease error bars on the C_ℓ plot by more than a factor of $\sqrt{2}$ if the second observation season is longer. This is a good assumption because the first season was delayed in the setup phase.

In the longer term, the existing MINT system may be improved by increasing the bandwidth. The inherent bandwidth of the receivers is about a factor of 2 larger than what is used by the correlators. This would entail incorporating a totally new channelizer to process the signals outside of the current MINT band. The current correlator design can still be used, but must be doubled in quantity. The upgraded platform may be used to “go deeper” on the CMB to resolve any differences between Sunyaev- Zel’dovich effect foregrounds, point sources, and the CMB.

Alternatively, a subtle change in the telescope pointing code could allow tracking astronomical sources, which would allow very deep exposures on fields and open up the possibility of map-making. The CBI has already shown the promise of such a platform in the mapping of low-redshift, large angular scale SZE clusters (Udomprasert et al. 2000).

In the next ten years, both large-scale surveys and deep images will be a powerful probe of both galactic astronomy and cosmology. Large-scale surveys require raw sensitivity, and are ideally suited to large focal plane bolometer arrays. Arrays of ~ 1000 bolometer pixels on single detector are now possible and plans are underway to install these detectors onto single ~ 6 m dishes in both Chile and Antarctica. The raw sensitivity of focal plane arrays is hard to match, but is not a stringent requirement for deep exposures on small numbers of clusters. The insensitivity to sources of systematic error make interferometers the natural choice for this type of measurement.

To examine the larger population of medium redshift clusters, finer angular resolution than the current MINT capability is required. This can only happen with the MINT design if it is made larger by about a factor of 3-5. Already the Taiwanese-led

AMiBA (Lo et al. 2001) project is planning a 90 GHz 19-element interferometer on a 6 m platform. The design draws upon the early conceptual design of MINT, and the lessons learned from the MINT experiment will no doubt prove useful in the design of future interferometers.

Channelizer Details

A.1 Physical Parameters

Parameter	Value
Power supply voltage	9.09 V
Input Current	1.3 Amp
Channelizer enclosure size	11.8 x 11.8 x 2 inches, all channelizers
Rack mounted	19 inch rack mount 4 inches high
Nominal Gain	+12 dB (see Figure 4.17)
Nominal 4.5 GHz and 5.5 GHz LO input	+13 dBm

Table A.1: Physical Parameters

A.2 Bandpass filter Design Equations

The field of microwave filter design is lively and complicated. Many companies stake their livelihood on their ability to meld craft with engineering to produce designs that never appear in the literature. Our purpose here is not to develop the equations for filter design. A good development appears in Pozar (1998) or Matthaei et al. (1980). The aim is merely to show the exact procedure used in designing the MINT bandpass filters.

We start by noting that bandpass filters are based on low-pass filter prototypes with a transformation in frequency:

$$\omega' \leftarrow \frac{2}{\Delta} \left(\frac{\omega - \omega_0}{\omega_0} \right),$$

$$\text{where } \Delta = \frac{\omega_2 - \omega_1}{\omega_0} \quad \text{and} \quad \omega_0 = \frac{\omega_1 + \omega_2}{2}. \quad (\text{A.1})$$

This maps bandpass response to low-pass response. The band-edges are ω_1 and ω_2 and Δ is the fractional bandwidth, while ω_0 is the center frequency. The only inputs here are ω_1 and ω_2 .

Next we must choose the general class of filter. For simplicity's sake, we followed an example in Matthaei et al. (1980), Section 8.09. Here, a Tchebyscheff .01 dB ripple design is used for its low VSWR in the passband. The next step is to decide on the number of sections. The more sections there are, the faster the roll-off will be¹. In general, one uses a graph to decide if a certain number of sections gives appropriate attenuation at a frequency determined by the transformation equations above. These graphs appear in Pozar (1998), Figure 8.27(a) or Matthaei et al. (1980), Figure 4.03-4. Again, we follow the example in Matthaei. Here an $n=6$ design is outlined.

Once the number of sections is known, we can consult a table for the values of the filter elements. The appropriate table is in Matthaei et al. (1980), Table 4.05-2(a). The values for $n=6$ are copied in Table A.2. These values refer to the inductance or capacitance of the filter elements, depending on the exact design of the filter. See Figure 4.04-1 in Matthaei et al. (1980) for a definition of these parameters.

The next step is to calculate the admittance inverter parameters. The coupled

¹Faster roll-off also leads to faster phase-wrapping at the band-edges.

g_0	1
g_1	.7813
g_2	1.3600
g_3	1.6896
g_4	1.5350
g_5	1.4970
g_6	.7098
g_7	1.1007

Table A.2: Element values from table 4.05-2(a) in Matthaei et al. (1980).

microstrip sections can be modelled as admittance inverters. See Pozar (1998) section 8.5 for a discussion of admittance inverters.

The J-admittance inverter parameters are found by:

$$\begin{aligned}
\frac{J_{01}}{Y_0} &= \sqrt{\frac{\pi\Delta}{2g_0g_1}} \text{ and} \\
\frac{J_{j,j+1}}{Y_0} &= \frac{\pi\Delta}{2} \frac{1}{\sqrt{g_jg_{j+1}}} \text{ for } j = 1 \text{ to } n-1, \\
\frac{J_{n,n+1}}{Y_0} &= \sqrt{\frac{\pi\Delta}{2g_ng_{n+1}}} \\
\text{and } Y_0 &= \frac{1}{Z_0} = \frac{1}{50\Omega}.
\end{aligned} \tag{A.2}$$

Given the J-admittance parameters, we can calculate the even and odd mode impedance² parameters:

$$\begin{aligned}
(Z_{0e})_{j,j+1} &= \frac{1}{Y_0} \left[1 + \frac{J_{j,j+1}}{Y_0} + \left(\frac{J_{j,j+1}}{Y_0} \right)^2 \right] \text{ and} \\
(Z_{0o})_{j,j+1} &= \frac{1}{Y_0} \left[1 - \frac{J_{j,j+1}}{Y_0} + \left(\frac{J_{j,j+1}}{Y_0} \right)^2 \right].
\end{aligned} \tag{A.3}$$

The even and odd mode impedances uniquely define the physical parameters for a coupled microstrip section. The length is given by $\lambda/4$ at frequency ω_0 in the medium

²The even and odd mode impedances come from a consideration of the impedance in a two branch device when the currents are travelling in the same (even mode) direction or in opposite (odd mode) direction. The impedance is then given as a superposition of the two modes: $\sqrt{Z_e Z_o}$

of propagation. It is not clear what exactly is meant by this, so as an approximation, we started with the length given by a $\lambda/4$ length of 50Ω Microstrip. This was only used as an initial guess for optimization. The other two parameters are width W and separation s .

The approximations for the impedances are given in Gupta et al. (1996), beginning with equation 8.96, which are restated here:

$$\begin{aligned} Z_{0e} &= Z_0 \frac{\sqrt{\epsilon_{re}(0)/\epsilon_{re}^e(0)}}{1 - Q_4 \sqrt{\epsilon_{re}(0)} \cdot Z_0/377}, \\ Z_{0o} &= Z_0 \frac{\sqrt{\epsilon_{re}(0)/\epsilon_{re}^o(0)}}{1 - Q_{10} \sqrt{\epsilon_{re}(0)} \cdot Z_0/377}. \end{aligned} \quad (\text{A.4})$$

Z_0 is the characteristic impedance of a microstrip line of the same width. The equations for this may be found in Pozar (1998), sec 3.8.

The constants are given by:

$$\begin{aligned} Q_1 &= 0.8695u^{0.194}, \\ Q_2 &= 1 + 0.7519g + 0.189g^{2.31}, \\ Q_3 &= 0.1975 + \left[16.6 + \left(\frac{18.4}{g} \right)^6 \right]^{-0.387} + \frac{1}{241} \ln \left[\frac{g^{10}}{1 + (g/3.4)^{10}} \right], \\ Q_4 &= \frac{2Q_1}{Q_2} \frac{1}{u^{Q_3} e^{-g} + (2 - e^{-g})u^{-Q_3}}, \\ Q_5 &= 1.794 + 1.14 \ln \left[1 + \frac{0.638}{g + 0.517g^{2.43}} \right], \\ Q_6 &= 0.2305 + \frac{1}{281.3} \ln \left[\frac{g^{10}}{1 + (g/5.8)^{10}} \right] + \frac{1}{5.1} \ln(1 + 0.598g^{1.154}), \\ Q_7 &= \frac{10 + 190g^2}{1 + 82.3g^3}, \\ Q_8 &= e^{-[6.5 + 0.95 \ln(g) + (g/0.15)^5]}, \\ Q_9 &= \ln(Q_7)(Q_8 + 1/16.5), \\ Q_{10} &= Q_4 - \frac{Q_5}{Q_9} \exp \left[\frac{Q_6 \ln(u)}{u^{Q_9}} \right], \end{aligned} \quad (\text{A.5})$$

with $u = \frac{W}{d}$ and $g = \frac{s}{d}$, where d is the thickness of the dielectric substrate. For the MINT channelizer, this is 30 mils. The effective dielectric constants are given by:

$$\begin{aligned}\epsilon_{re}^e(0) &= \frac{\epsilon_r + 1}{2} + \frac{\epsilon_r - 1}{2} \left(1 + \frac{10}{v}\right)^{-a_e b_e}, \\ \epsilon_{re}^o(0) &= \epsilon_{re}(0) + [0.5(\epsilon_r + 1) - \epsilon_{re}(0) + a_o] e^{-c_o g^{d_o}},\end{aligned}\quad (\text{A.6})$$

with

$$\begin{aligned}v &= \frac{u(20 + g^2)}{10 + g^2} + g e^{-g}, \\ a_e &= 1 + \frac{1}{49} \ln \left[\frac{v^4 + (v/52)^2}{v^4 + 0.432} \right] + \frac{1}{18.7} \ln \left[1 + \left(\frac{v}{18.1} \right)^3 \right], \\ b_e &= 0.564 \left(\frac{\epsilon_r - 0.9}{\epsilon_r + 3} \right)^{0.053}, \\ a_o &= 0.7287 [\epsilon_{re}(0) - 0.5(\epsilon_r + 1)] (1 - e^{-0.179u}), \\ b_o &= \frac{0.747\epsilon_r}{0.15 + \epsilon_r}, \\ c_o &= b_o - (b_o - 0.207) e^{-0.424u}, \\ d_o &= 0.593 + 0.694 e^{-0.562u}.\end{aligned}\quad (\text{A.7})$$

For Duroid 6002, $\epsilon_r = 2.94$ and $\epsilon_{re}(0)$ is the effective dielectric constant for a microstrip line of the same width:

$$\epsilon_{re}(0) = \frac{\epsilon_r + 1}{2} + \frac{\epsilon_r - 1}{2} \frac{1}{\sqrt{1 + 12d/W}}.\quad (\text{A.8})$$

These equations are all that is needed to solve for all W 's and s 's for each section. Obviously, it is quite complicated. These equations are also non-invertible so we must use a graphical or numerical technique to solve for W and s . Table A.3 is a sample calculation for a 4.0-4.5 GHz filter. These parameters are used as a starting point for optimization.

A.3 Jumper

Placing LO power splitters on the output side of the channelizer places a barrier that the output signals must cross. The solution with the highest isolation between crossed

i, j	J_{ij}	$(Z_{0e})_{ij}$	$(Z_{0o})_{ij}$	$W(\text{mils})$	$s(\text{mils})$
0,1	.486	86.144	37.509	47.211	2.540
1,2	.179	60.571	42.643	71.214	17.362
2,3	.122	56.839	44.648	74.360	29.416
3,4	.115	56.396	44.921	74.669	31.610
4,5	.122	56.839	44.648	74.360	29.416
5,6	.179	60.571	42.643	71.214	17.362
6,7	.486	86.144	37.509	47.211	2.540

Table A.3: Example calculated values for a 4.0-4.5 GHz bandpass filter.

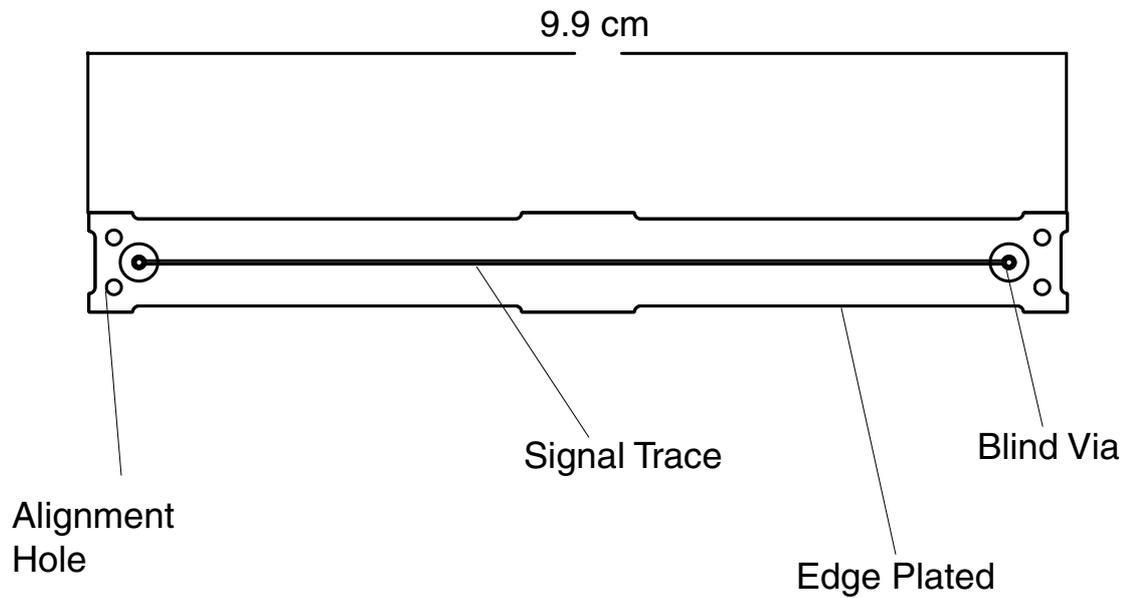


Figure A.1: Outline of the long jumper used to cross the output signals with the LO lines. There is also a smaller jumper used to cross the 4.5 GHz LO line with the 5.5 GHz LO line.

lines and the least impedance mismatch problems was a multi-layer board with vias, much like circuit boards at lower frequencies. The cost of going to a multi-layer board is much higher and it removes the primary ground from the case.

The way to solve this was to make a small “jumper” which would span the short distance that it would take to cross lines. Figure A.1 is an outline drawing of the jumper. The jumper is a 3-layer circuit. Because it has a ground plane on either side, it is a stripline instead of a microstrip circuit. This, of course, means that the design equations are different from microstrip circuits. Stripline circuits are, in general, thinner than microstrip circuits. See Pozar (1998) for more details. The actual parameters were reached using simulation.

At either end of the jumper, there are “blind” vias to the transmission line as well as through-pads on the ground layer. Blind vias do not extend through the entire board and are generally more expensive than through vias. The jumpers are laminated with silver epoxy (Epo-tex H20E) on the underside of the channelizer. There are holes at the end of the jumper for alignment pins to facilitate the lamination process.

The jumper is also “edge plated,” which ensures that the ground connection is solid around the perimeter of the device. There are also oval shapes cut from the edge to improve the ground area.

Soon after the fabrication of the jumper, the microwave industry presented a better alternative to the jumper. A crossover is now available from Anaren in their Xinger line of product. It is a surface-mount device with high isolation. The company offers both a RF-RF and RF-DC devices, with 25 dB of isolation.

Correlator Details

B.1 Physical Parameters

Parameter	Value
Correlator PCB Size	15.0 x 15.0 x .0625 inches
Rack mounted	19 inch rack mount 8 inches high
Nominal input RF power	-8 dBm in 0-500 MHz band

Table B.1: Physical Parameters

Name	Voltage	Current	Description
V_{ee}	-5.2V	8 A	Main supply for digital side ECL
AV_{ee}	-5.2 V	4 A	Main supply for analog side ECL
V_{tt}	-2.0 V	3 A	Termination supply for digital ECL signals
AV_{tt}	-2.0 V	1 A	Termination supply for analog ECL signals
V_{cc}	5.0 V	1 A	Main analog and TTL supply
3.3	3.3 V	3 A	LVTTL reference
2.5	2.5 V	4 A	FPGA supply
$V_{amp}/linear$	12 V	1.2 A	Supply for amplifiers and power supply fans

Table B.2: Voltage and current levels for a single correlator board. The power comes from the PBX Corr, which is composed of high-power switching power supplies with LC filters on the output. There is also one linear supply for the fans and amplifiers.

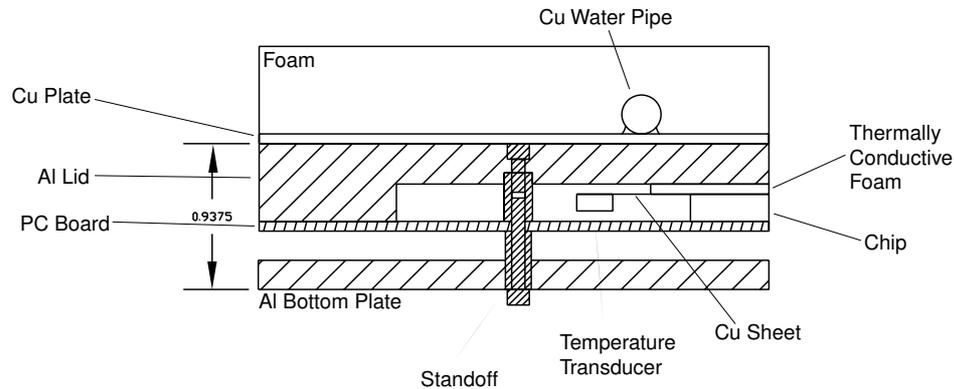


Figure B.1: Detail cutaway view of correlator enclosure sandwich. There are chips on both sides of the PCB.

B.2 Enclosure

The correlator is housed in an enclosure very similar to that of the channelizer. The desire to have stringent RF shielding combined with the cooling requirements in low atmospheric pressure led to the idea of using a single large heatsink per board. Figure B.1 is a cutaway diagram of the enclosure sandwich. The PCB is sandwiched between two aluminum plates. The bottom plate is flat with standoffs and the lid is milled from a 1/2" plate. The thickness of the lid follows the chip height across the board. To account for slight chip-to-chip variations in height, a .100" thick layer of conductive foam is inserted between the chip and the lid. A temperature transducer is epoxied to a copper sheet, which is inserted between the chip and the foam for the digitizers and correlators.

A copper tube for cooling fluid is soldered to a copper plate, then mounted directly to the lid. The fluid is a 50/50 mixture of ethylene glycol and water. The water is cooled or heated via a cooling box, which selectively cools via fans on a heat exchanger or heats via a servo-controlled water heater.

Tables

latitude	-22.9581° (south)
longitude	-67.7858° (west)
altitude	5200 m

Table C.1: The GPS coordinates of the MINT site as reported by the onboard GPS sensor.

Band:	4.0-4.5	4.5-5.0	5.0-5.5	5.5-6.0
Receiver:				
A	38.3	41.1	52.7	77.4
B	30.6	33.6*	37.3	45.5
C	34.5*	29.9*	35.6*	37.3*
D	32.7	36.3	41.8	49.5

Table C.2: T_{sys} in Kelvin derived from a hot/cold load test. A T_{atm} of 4-8K is included. The details are in Dorwart (2002). At the time of the measurement, not all of the channels were functional. The asterisks refer to data taken at an earlier time.

Bibliography

- Aboobaker, Asad, 2002. A 139-141 GHz Noise Source for Calibrating the MINT Telescope. Generals project, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Agilent, 2001. Momentum. This is a software manual presently available at: <http://contact.tm.agilent.com/tmo/hpeesof/docs/adsdoc2001/pdf/mom.pdf>.
- Bond, J. R. and G. Efstathiou, 1984. Cosmic background radiation anisotropies in universes dominated by nonbaryonic dark matter. *Astrophys J Lett*, 285:L45–L48.
- Brigham, E. O., 1988. *The Fast Fourier Transform*. Upper Saddle River, N.J. 07458: Prentice-Hall, 1988.
- Carlstrom, J. et al., 2001. The Sunyaev-Zel'dovich Effect: Results and Future Prospects. *IAP conference July 2000*. Astro-ph/0103480.
- Dicker, S. R., S. J. Melhuish, R. D. Davies, C. M. Gutiérrez, R. Rebolo, D. L. Harrison, R. J. Davis, A. Wilkinson, R. J. Hoyland, and R. A. Watson, 1999. Cosmic microwave background observations with the Jodrell Bank-IAC interferometer at 33GHz. *Mon Not R Astron Soc*, 309:750–760.
- Dodelson, Scott, 2002. *Modern Cosmology*. In Press.
- Dorwart, William B., 2002. *A 150-GHz Interferometer for the Cosmic Microwave Background Anisotropy*. Ph.D. thesis, Princeton University, Dept of Physics, Princeton New Jersey. This document is not yet published, but will be available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Dover, Rennie Wm. and Todd Pearson, 2000. Metastability and the ECLinPS

- Family. Application note AN1504D from ON semiconductor. Available at <http://www.onsemi.com/pub/Collateral/AN1504-D.PDF>.
- Dumont, Charles, 2001. *Digital Correlation between Two Cryogenic SIS Mixers*. Senior thesis, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Fixsen, D. J., E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, 1996. The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set. *Astrophys J*, 473:576.
- Fomalont, E. B., K. I. Kellermann, and J. V. Wall, 1984. Limits to the small-scale fluctuations in the cosmic background radiation. *Astrophys J Lett*, 277:L23–L26.
- Gupta, K. C., Ramesh Garg, Inder Bahl, and Prakash Bharta, 1996. *Microstrip Lines and Slotlines 2nd Edition*. Artech House.
- Halverson, N. W. et al., 2001. DASI First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum. Astro-ph/0104489.
- Harris, A. I., K. Isaak, and J. Zmuidzinas, 1998. WASP: wideband spectrometer for heterodyne spectroscopy. In *Advanced Technology MMW, Radio, and Terahertz Telescopes*, pp. 383–394.
- Hawkins, David W., 1998. Correlator theory. Available at <http://www.ovro.caltech.edu/dwh/>.
- Hedman, M. M., D. Barkats, J. O. Gundersen, S. T. Staggs, and B. Winstein, 2001. A Limit on the Polarized Anisotropy of the Cosmic Microwave Background at Subdegree Angular Scales. *Astrophys J Lett*, 548:L111–L114.
- Hinderks, James, 1999. *The Theory, Design, and Implementation of a High-Speed Digital Correlator*. Senior thesis, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Hobson, M. P., A. N. Lasenby, and M. Jones, 1995. A Bayesian Method for Analysing Interferometer Observations of Cosmic Microwave Background Fluctuations. *Mon Not R Astron Soc*, 275:863.

- Hu, W. and S. Dodelson, 2002. Cosmic Microwave Background Anisotropies . *Astronomy and Astrophysics Ann Rev*. Submitted, see astro-ph/0110414.
- Hu, W. and N. Sugiyama, 1995. Anisotropies in the cosmic microwave background: an analytic approach. *Astrophys J*, 444:489–506.
- Hu, W. and N. Sugiyama, 1996. Small-Scale Cosmological Perturbations: an Analytic Approach. *Astrophys J*, 471:542.
- Hu, W., N. Sugiyama, and J. Silk, 1997. The Physics of Microwave Background Anisotropies. *Nature*, 386:37–43.
- Hu, W. and M. White, 1997. The Damping Tail of Cosmic Microwave Background Anisotropies. *Astrophys J*, 479:568.
- Hu, W. T., 1995. *Wandering in the Background: a Cosmic Microwave Background Explorer*. Ph.D. thesis, UC Berkeley, University of California, Berkeley Dept of Physics, Berkeley, CA 94720.
- Jones, B. J. T. and R. F. G. Wyse, 1985. The ionisation of the primeval plasma at the time of recombination. *Astronomy and Astrophysics*, 149:144–150.
- Jones, M. E., 1997. Results from Cat and Prospects for the Vsa. In *Microwave Background Anisotropies*, p. 161.
- Kamionkowski, M., D. N. Spergel, and N. Sugiyama, 1994. Small-scale cosmic microwave background anisotropies as probe of the geometry of the universe. *Astrophys J Lett*, 426:L57–+.
- Keating, B. G., C. W. O’Dell, A. de Oliveira-Costa, S. Klawikowski, N. Stebor, L. Piccirillo, M. Tegmark, and P. T. Timbie, 2001. A Limit on the Large Angular Scale Polarization of the Cosmic Microwave Background. *Astrophys J Lett*, 560:L1–L4.
- Kerr, A. R., S. K. Pan, A. W. Lichtenberger, and F. L. Lloyd, 1993. *Proceedings of the Fourth Intl Symp on Space Terahertz Texhnology*.
- Kraus, J. D., 1966. *Radio astronomy*. New York: McGraw-Hill, 1966.

- Li, C. Q., S. H. Li, and R. G. Bosisio, 1994. CAD/CAE Design of an Improved Wideband Wilkinson Power Divider. *Microwave Journal*.
- Lo, K. Y., T. H. Chiueh, R. N. Martin, K.-W. Ng, H. Liang, U.-L. Pen, C.-P. Ma, M. Kesteven, R. Sault, R. Subrahmanyam, W. Wilson, and J. Peterson, 2001. AMiBA: Array for Microwave Background Anisotropy. In *20th Texas Symposium on relativistic astrophysics*, p. 172.
- Loh, Yeong Shang, 2000a. Mint optical layout. Prelims project, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Loh, Yeong Shang, 2000b. Mint optics. Generals project, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Marriage, Tobias A., 2000. *The MINT Digitizer*. Senior thesis, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- Matthaei, George L., Leo Young, and E. M.T. Jones, 1980. *Microwave Filters, Impedance-Matching Networks, and Coupling Structures*. Artech House Books, Dedham, MA.
- Misner, C. W., K. S. Thorne, and J. A. Wheeler, 1973. *Gravitation*. San Francisco: W.H. Freeman and Co., 1973.
- Netterfield, C.B. et al., 2001. A measurement by boomerang of multiple peaks in the angular power spectrum of the cosmic microwave background. Astro-ph/0104460.
- 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. Holzapfel, S. T. Myers, J. E. Carlstrom, E. M. Leitch, M. Joy, L. Bronfman, and J. May, 2001a. First Intrinsic Anisotropy Observations with the Cosmic Background Imager. *Astrophys J Lett*, 549:L1–L5.
- Padin, S., J. K. Cartwright, M. C. Shepherd, J. K. Yamasaki, and W. L. Holzapfel, 2001b. A wideband analog correlator for microwave background observations. *IEEE Trans Instrum Meas*, 50:1234–1240. Available at:<http://www.astro.caltech.edu/tjp/CBI/Papers/>.

- Partridge, R. B., 1995. *3K: The Cosmic Microwave Background Radiation*. ISBN 0521352541, Cambridge University Press, 1995.
- Partridge, R. B., E. A. Richards, E. B. Fomalont, K. I. Kellerman, and R. A. Windhorst, 1987. Small-Scale Cosmic Microwave Background Observations at 8.4 GHz. *Astrophys J*, 483:38–50.
- Peebles, P. J. E., 1980. *The large-scale structure of the universe*. Research supported by the National Science Foundation. Princeton, N.J., Princeton University Press, 1980. 435 p.
- Pospieszalski, M. W., W. J. Lakatos, R. Lai, K. L. Tan, D. C. Streit, P. H. Liu, R. M. Dia, and J. Velebir, 1993. *Microwave Symposium Digest IEEE MTT S International*, 2:515–518.
- Pozar, David M., 1998. *Microwave Engineering 2nd ed*. Wiley and Sons.
- Rahmat-Samii, Y., W. Imbriale, and V. Galindo, 1996. *Diffraction Analysis of a Dual Reflector Antenna (DADRA)*. Software User's Manual.
- Sachs, R. K. and A. M. Wolfe, 1967. Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. *Astrophys J*, 147:73.
- Seljak, U., 1994. A two-fluid approximation for calculating the cosmic microwave background anisotropies. *Astrophys J Lett*, 435:L87–L90.
- Seljak, U. and M. Zaldarriaga, 1996. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *Astrophys J*, 469:437.
- Silk, J., 1968. Cosmic Black-Body Radiation and Galaxy Formation. *Astrophys J*, 151:459.
- Subrahmanyan, R., R. D. Ekers, M. Sinclair, and J. Silk, 1993. A Search for Arcminutescale Anisotropy in the Cosmic Microwave Background. *Mon Not R Astron Soc*, 263:416.
- Subrahmanyan, R., M. J. Kesteven, R. D. Ekers, M. Sinclair, and J. Silk, 2000. An Australia Telescope survey for CMB anisotropies. *Mon Not R Astron Soc*, 315:808–822.

- Tegmark, M., 1995. Doppler peaks and all that: CMB anisotropies and what they can tell us. Available at: <http://www.hep.upenn.edu/max/varena.html>.
- Tegmark, M., D. J. Eisenstein, W. Hu, and A. de Oliveira-Costa, 2000. Foregrounds and Forecasts for the Cosmic Microwave Background. *Astrophys J*, 530:133–165.
- Tegmark, M. and M. J. Rees, 1998. Why is the Cosmic Microwave Background Fluctuation Level 10^{-5} ? *Astrophys J*, 499:526.
- Thompson, A. R. and L. R. D’Addario, 2000. Relative Sensitivity of Double- and Single-Sideband Systems for both Total Power and Interferometry. ALMA memo 304, available at <http://www.eso.org:8082/index.html>.
- Thompson, A. R., J. M. Moran, and G. W. Swenson, 1986. *Interferometry and synthesis in radio astronomy*. New York, Wiley-Interscience, 1986, 554 p.
- Timbie, P. T., 1985. *A novel interferometer to search for medium scale anisotropy in the 2.7K background radiation*. Ph.D. thesis, Princeton University, Dept of Physics, Princeton New Jersey 08544.
- Udomprasert, P. S., B. S. Mason, and A. C.S. Readhead, 2000. The Sunyaev-Zel’dovich effect with the Cosmic Background Imager. *Constructing the Universe with Clusters of Galaxies IAP 2000 meeting Paris*.
- Wang, X., M. Tegmark, and B. Jain, 2001. Is Cosmology Consistent? *American Astronomical Society Meeting*. See also astro-ph/0105091.
- Wesley, Daniel, 2000. *Construction and Characterization of a Receiver for CMB Interferometry*. Senior thesis, Princeton University. Available at: <http://imogen.princeton.edu/mintweb/publications/>.
- White, M., 2001. Constraints from the Damping Tail. *Astrophys J*, 555:88–91.
- White, M., J. E. Carlstrom, M. Dragovan, and W. L. Holzapfel, 1999. Interferometric Observation of Cosmic Microwave Background Anisotropies. *Astrophys J*, 514:12–24. Referred to as WCDH in text.
- Wrobel, J.M. and R.C. Walker, 1999. Sensitivity. *Synthesis Imaging in Radio Astronomy II Astronomical Society of the Pacific Conference Series vol 180*.

List of Figures

1.1	Typical theoretical plot of δT vs ℓ	3
2.1	Simplified Baseline	17
2.2	2D MINT window functions	21
2.3	1- σ contours of the 2D MINT window functions	22
2.4	1D MINT window functions assuming gaussian beams	23
2.5	1D MINT window functions given simulated beams. See Figure 3.7 for the beam shapes. As before, the factor $\eta_{R-J}^2 = 0.36$ is omitted.	23
2.6	Best-fit power spectrum from the BOOMERanG collaboration	24
2.7	Signal before integration over ℓ , again assuming gaussian beams	24
2.8	Expected contributions given simulated beams	25
2.9	Expected error bars from observing 166 uncorrelated fields for 30 12-hour days	32
2.10	Scan strategy and sky coverage	33
2.11	Expected error bars from observing 21 uncorrelated fields for 30 12-hour days	35
3.1	Overhead view of antenna layout	39
3.2	Response patterns for MINT baselines	40
3.3	The MINT signal path for one baseline	41
3.4	Cutaway diagram of the smaller optics	42
3.5	Beam map from the first run of optics testing	43
3.6	Predicted beam map from modified DADRA Cod	44
3.7	An enlarged view of the main lobes of both antennas	44
3.8	Microwave components of the MINT receiver	46
4.1	Channelizer Layout.	51

4.2	The MINT channelizer	52
4.3	A picture focusing on one downconversion chain of the channelizer . .	53
4.4	The filter dimensions sent to the board fabricator	57
4.5	A schematic entry window in from Advanced Design System	58
4.6	Meshing of the filter by momentum simulator	60
4.7	Results of linear and momentum simulation, using the modified parameters in Figure 4.5.	61
4.8	Meshing of 4.5 GHz power splitter	62
4.9	Momentum simulation results for the 4.5 GHz power splitter	63
4.10	The S parameters for the 5.5 GHz LO power splitter	64
4.11	Meshing of broadband power splitter	65
4.12	The S parameters for the broadband power splitter	65
4.13	Bandpasses for the filterbank resulting from combining the momentum simulations for three cascaded broadband power splitters and four bandpass filters	66
4.14	View of the enclosure lid for the channelizer	68
4.15	Detail of the channelizer lid to show the eccosorb lining	68
4.16	VNA measurements of prototype filterbank (solid lines) plotted on top of the simulation results shown in Figure 4.13 (dotted lines)	70
4.17	Bandpasses through the entire channelizer. These results are typical for all channelizer halves, but in this case represent the response for channelizer A	71
4.18	Bandpass difference between identical channels of the 4 separate channelizers	72
4.19	Relative phases at the IF output of the channelizer as measured relative to channelizer B	74
4.20	Relative phases at the IF output of the channelizer with a linear fit removed	75
5.1	Correction factor from Equation 5.21	85
5.2	The <i>rms</i> of the correlator output for a single multiplication	86
6.1	Schematic of a single digitizer section	89

6.2	General correlation algorithm for digital lag-correlators	92
6.3	Schematic of the multiplier-adder kernel	93
6.4	Schematic of a single lag for $\tau=2$	95
6.5	Two hours of typical correlator 1 output, taken at the coldest part of the night	101
6.6	Correlator output with calibration spikes removed.	102
6.7	Histogram of the correlator output for two hours	102
6.8	The fractional occupation of lower levels, ϕ , for a short time around a calibration spike, for two digitizers on receivers A and B.	103
6.9	Correlator output for both phase states given as a function of lag . .	104
6.10	Maximum correlation interpolation	105
7.1	Data reduction procedure	110
7.2	Diagram of some relevant components involved with interferometer stability	112
7.3	Temperature calibrated calspike power from the correlated baseline D- A, from correlator 1	114
7.4	Calibrator signal vs. switch encoder position	115
7.5	Calibrator signal corrected for noise switch encoder position	116
7.6	Relative calibration signal	117
7.7	Relative calibration signal for the entire season	118
7.8	Phase correction technique	120
7.9	Phase slope for one night	121
7.10	Phase offset plotted for a single night	122
7.11	Phase offset for the entire season based on the noise source	123
7.12	Phase slope for the entire season based on the noise source	124
7.13	The decrease in sensitivity as a function of the <i>rms</i> fluctuations in phase offset	126
7.14	Integration down plot	128
A.1	Outline of the long jumper used to cross the output signals with the LO lines	136
B.1	Detail cutaway view of correlator enclosure sandwich	140

List of Tables

2.1	Expected signals in each of the MINT baselines	25
2.2	The equivalence time as calculated from Equation 2.45	30
2.3	The elevation bins defined for the modified scan strategy	34
4.1	The parameters from Li et al. (1994) used to design the broadband power splitter	64
4.2	The noise bandwidth as computed from Equation 4.1	73
4.3	The <i>rms</i> of the phase errors after a linear fit has been subtracted. All are in degrees. Again, the pair B-A for Channel 1 is missing due to bad data.	75
5.1	Quantization table. V_0 is the threshold voltage.	81
5.2	Multiplication table for standard scheme.	82
5.3	Multiplication table for standard scheme with the optimal integer value $n=3$	82
5.4	Multiplication table for modified deleted inner product scheme.	83
5.5	Deleted inner product scheme, scaled and biased	83
6.1	The expected and measured correlation coefficients	106
A.1	Physical Parameters	131
A.2	Element values from table 4.05-2(a) in Matthaei et al. (1980).	133
A.3	Example calculated values for a 4.0-4.5 GHz bandpass filter.	136
B.1	Physical Parameters	138
B.2	Voltage and current levels for a single correlator board	139

C.1	The GPS coordinates of the MINT site as reported by the onboard GPS sensor.	141
C.2	T_{sys} in Kelvin derived from a hot/cold load test	142