The Princeton IQU Experiment and Constraints on the Polarization of the Cosmic Microwave Background at 90 GHz

Matthew Hedman

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 Matthew Hedman, 2002. All rights reserved.

Abstract

The polarization anisotropies in the Cosmic Microwave Background provide important information about the structure and dynamics of the early universe. This cosmological polarization has thus far eluded detection, which is not surprising because it should be an extremely small signal (an order of magnitude smaller than the temperature anisotropies, or about 5 μK *rms* variations in 3K thermal radiation). The Princeton IQU Experiment, or PIQUE, is an effort to measure this tiny polarized signal using two correlation polarimeters operating at W-band (84-100 GHz) and Q-band (35-46 GHz). This work describes the construction, characterization, observations and preliminary results from the W-band system, which observed the sky around the north celestial pole over two succesive winters (2000 and 2001), and acquired 500 hours of useful data. These data yield one of the tightest constraints on the cosmological polarization to date.

Acknowledgments

Many thanks go to the members of the PIQUE team. Suzanne Staggs, of course, was the one who got me involved in this project, and she deserves credit for designing most of the experimental system (including the clever chopping-plate calibration system). As my advisor, she has been extremely patient with my many questions and has often provided useful advice and information, not to mention innumerable suggestions on previous drafts of this thesis. Josh Gundersen was a great help with many of the quirks of the telescope, and also provided the first analysis of the planetary observations. Bruce Winstein provided many stimulating conservations and has also been responsible for a good portion of the final data analysis. Denis Barkats and Jeff McMahon have done a lot of work on this project, including much painful wiring. All these people also deserve high praise for being willing to watch over the telescope at all hours of the night and to rush over at a moment's notice to stow the telescope whenever rain and snow threatened.

Thanks must also be given to Ted Griffiths, who was responsible for both the outer ground screen and the LO waveguide, and Al Dietrich, Laszlo Varga, William Dix, Glenn Atkinson, Ted Lewis and John Mellodge of the physics department metal shop, for making the inner ground screen and many other parts for the telescope.

Other people who helped out with PIQUE include: Chris Herzog, who calculated the polarized signal from the chopper plate and designed and built the mounting hardware needed to move this plate; Jamie Hinderks, who wrote the data acquisition software (and Marc Schreiber, who modified it); Chris Barnes, who provided useful tools for performing near-field maps of the optics; Matias Zaldarriaga, who devloped (and explained) the formalism for analyzing CMB polarization data, and especially Norm Jarosik, who not only was a careful steward of the irreplaceable HEMTs and the indispensible VNA, but also on many occasions provided sage advice on experimental issues.

The support of the larger Gravity Group has also been terrific. Special thanks must go to the MAP team for providing the HEMTs that made this project possible. Both Lyman

Page and David Wilkinson have often provided useful information on likelihoods and planets and so on. David Wilkinson has also been kind enough to act as reader.

I must also thank my fellow graduate students. I could never have passed generals without the help of Eric Splaver, Laura Cadonati, Nick Darnton, Olgica Bakajin, Ken Nagamine and Diego Casa. More recently, Randy Dorwart, Mike Nolta and Huan Tran have been kind enough to share their expertise in the world of CMB experiments.

Finally, thanks to my mom, my dad, and my brother, for believing in me and offering words of encouragement when I needed them most.

Contents

Chapter 1

Introduction

The Cosmic Microwave Background (CMB), as the relic radiation from the hot, dense phase of the early universe, is an invaluable cosmological probe. In particular, the small variations in the CMB from point to point across the sky encode an immense amount of information regarding the structure and composition of the early universe. Indeed, variations in the intensity, or temperature, of the CMB have now been measured with sufficient accuracy and precision that they place important constraints on cosmological models (Wang et al., 2001). However, these temperature anisotropies are not the only structures observable in the CMB. The CMB should also have an anisotropic polarized component, which represents a source of additional cosmological information.

Even though there have been several searches for this cosmological polarization in the three decades since Penzias and Wilson (1965) first reported the detection of the "unpolarized" microwave background, the polarized component of the CMB has so far eluded detection. This is not that surprising, since the current body of theory strongly predicts that the polarization of the CMB should be an order of magnitude smaller than the temperature anistropies (roughly one part in 10^6 or about 5 μK rms at most), which is below the detection threshold of all measurements to date. However, with the increasing sensitivity of microwave receivers, the limits on the polarized component of the CMB have steadily improved (see Figure 1.1, Staggs et. al. 1999), and the tightest limits are now around 10μK (Subrahmanyan et. al., 2000; Hedman et.al., 2001; Keating et. al., 2001).

Current efforts to detect and measure the cosmological polarization can be divided into

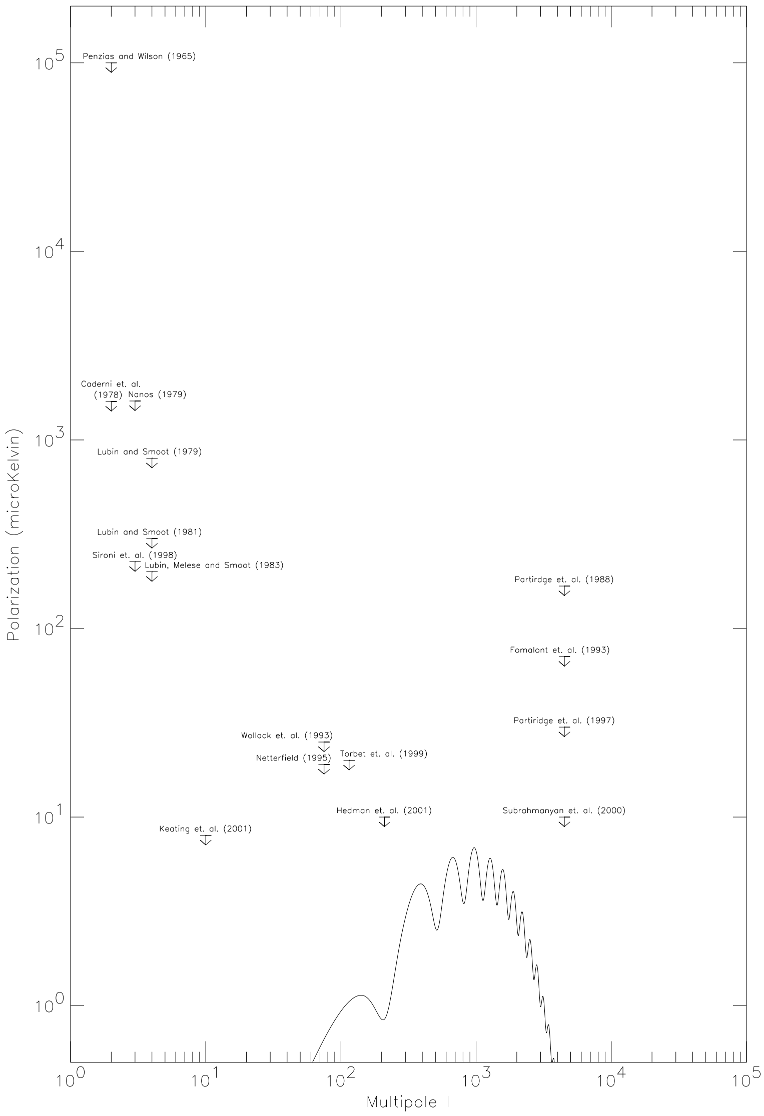

Figure 1.1: Status of measurements of the CMB polarization as of the fall of 2001. The curve is the predicted polarized signal from a flat, ΛCDM cosmological model as a function of angular scale (parametrized by the multipole moment l). The arrows indicate the level and angular scales of the published limits. Since most of these limits were made before the full formalism for presenting CMB polarization results was worked out (Zaldarriaga, 1998), no effort is made to show the exact range of angular scales covered by each experiment, or the precise value of the limits.

two groups. On the one hand, there are experiments which were designed primarily to measure the temperature anistropies, but which have some polarization sensitivity as well. The recent Saskatoon, TOCO and ATCA experiments (Wollack et al., 1993; Netterfield, 1995; Torbet et al., 1999; Subrahmanyan et al., 2000) are in this tradition, as are the ongoing CBI, DASI, BOOMERanG and MAXIPOL projects (Cartwright et al., 2001; Kovac, 2001; Ruhl, 2001; Hanany, 2001). On the other hand, there are experiments which were designed exclusively to look for CMB polarization. POLAR, which has produced the strongest limit on CMB polarization yet (Keating et al., 2001a), falls within this group, along with the ongoing COMPASS and Polatron experiments (Staggs et al., 1999; Pilhour et al., 2001; Timbie, 2001). The project described here is another of these dedicated polarimeters: the Princeton IQU Experiment or PIQUE (the IQU refers to the Stokes parameters that describe the polarization state of electromagnetic radiation; see Section 2.1).

The following chapters detail the construction, characterization, observations and results of the PIQUE experiment. In order to motivate this project, as well as justify certain aspects of the experimental design, this chapter provides a brief overview of the theoretical predictions regarding the origin and size of the polarized signal from the CMB, as well as a summary account of other sources of polarized microwave radiation that can potentially interfere with measurements of the cosmological polarization.

1.1 The Origin of Cosmological Polarization

The polarized component of the CMB arose naturally from the small inhomogeneities in the primordial plasma that filled the early universe. Such an inhomogeneous plasma emits polarized radiation due to Thomson scattering between the photons and the free electrons within the plasma. This process is the classical interaction between charged particles and electromagnetic waves, so the origin of the cosmological polarization does not rely on any particularly strange or exotic physics.

Heuristically, Thomson scattering occurs when an incident electromagnetic wave causes a charged particle to oscillate back and forth, and the accelerated charge re-emits radiation

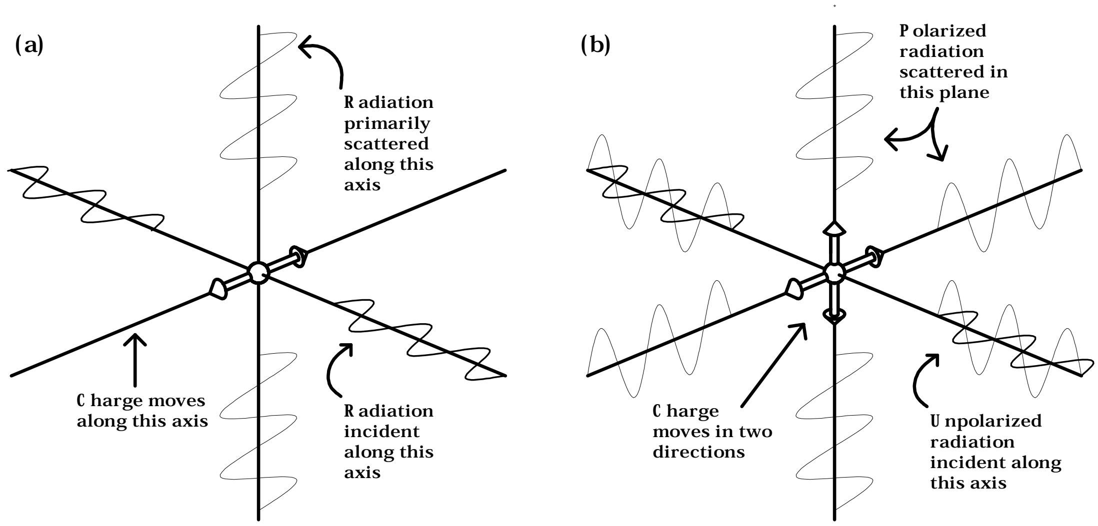

Figure 1.2: Simple examples of Thomson scattering with the charge being illuminated from a single direction. The two cases correspond to the incident light being either (a) linearly polarized or (b) unpolarized. For clarity, only the electric field of the radiation propagating along the coordinate axes is illustrated.

over a wide range of angles. The intensity and polarization of the scattered light depends on both the direction and polarization of the incident radiation, as given by the cross section (Jackson, 1975):

$$\frac{d\sigma_T}{d\Omega} = (\frac{e^2}{mc^2})^2 |\hat{\epsilon} \cdot \hat{\epsilon}'|^2,$$

where $\hat{\epsilon}$ and $\hat{\epsilon}'$ are the incident and scattered polarization directions, respectively.

For a simple example of this process, consider the case in which a charge is illuminated from a single direction with linearly polarized light, as illustrated in Figure 1.2a. In this situation, the charge moves along a single axis, and the light emitted by the moving charge is simple dipole radiation. If, on the other hand, the incident light is unpolarized, as shown in Figure 1.2b., then the motion of the charge may be decomposed into oscillations along two orthogonal axes. With the charge confined to move in a single plane, radiation scattered into this plane must be polarized parallel to it, thus demonstrating that Thomson scattering can indeed generate polarization from initially unpolarized light.

Of course, the electrons in the early universe were not illuminated from a single direction, but were instead bathed in a nearly isotropic sea of photons. If this incident radiation were

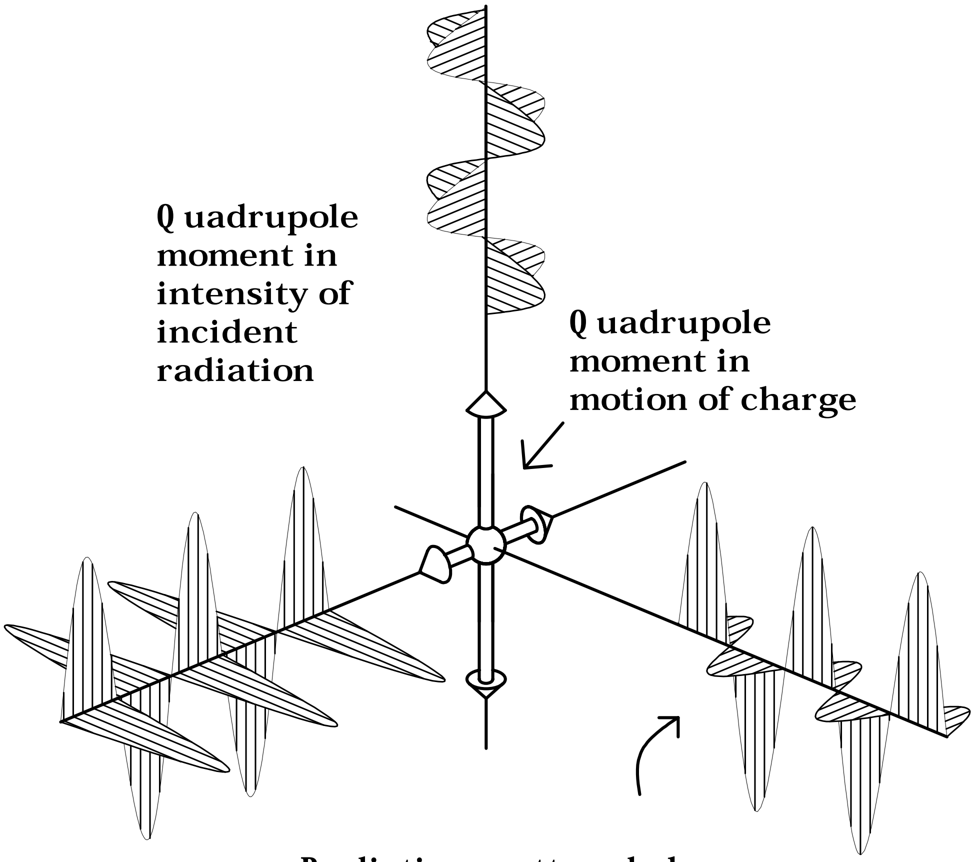

Figure 1.3: Thomson scattering with an anisotropic incident radiation field. The incident (unpolarized) radiation is more intense along one axis than the other, which causes the charge to oscillate more strongly along one axis than the other. The radiation emitted by the accelerating charge along the third axis consequently has a finite polarization. Only a quadrupole moment in the intensity of the unpolarized radiation field can generate polarization in this way, since other moments do not define unique axes along which the motion of the charge is a maximum or a minimum.

perfectly isotropic, then there would be no favored direction, and the scattered light could not be polarized. However, since the primordial plasma was not completely homogeneous, the incident radiation on the charged particles was not, in general, perfectly isotropic. In particular, there could have been quadrupole moments in the intensity of the radiation incident upon these charges. Such a quadrupole moment, as shown in Figure 1.3, produces scattered light with a finite polarized component.

At first, it might appear that these polarization-generating quadrupole moments were the result of variations in the temperature of the primordial plasma. In fact, such temperature gradients were unable to generate detectable amounts of cosmological polarization. Imagine that the two sources of the incident radiation illustrated in Figure 1.3 are from thermal sources at different temperatures. In this case, the induced motion of the charge along the third axis (not shown) scatters radiation between the two sources, bringing them into thermal equilibrium and washing out the quadrupole moment before much polarized radiation is generated. The expected polarized signals from the CMB were therefore produced by more dynamic aspects of the early universe.

One of the more exotic sources of cosmological polarization was primordial gravity waves. Gravity waves passing through the primordial plasma caused the wavelengths of photons propagating in orthogonal directions to be alternately stretched and compressed, producing a quadrupole moment in the redshifts of the radiation (see Figure 1.4). These quadrupole moments in the *apparent* temperature of the plasma in turn generated a polarized signal with a distinct signature (see Chapter 9), that, in principle, allows the gravity wave content of the early universe to be quantified (which has relevance to the dynamics of inflation, etc.; see Kamionkowski and Kosowky (1999)). Unfortunately, the polarized signal due to gravity waves is expected to be very small (less than 10% of the still undetected polarization signal), and therefore this particular component of the cosmological polarization may remain elusive for some time to come.

The primary source of cosmological polarization is a far more mundane process involving velocity gradients in the primordial plasma. Consider a flow field with a gradient, such as that shown in Figure 1.5, which might occur when a plasma flows into a potential well.

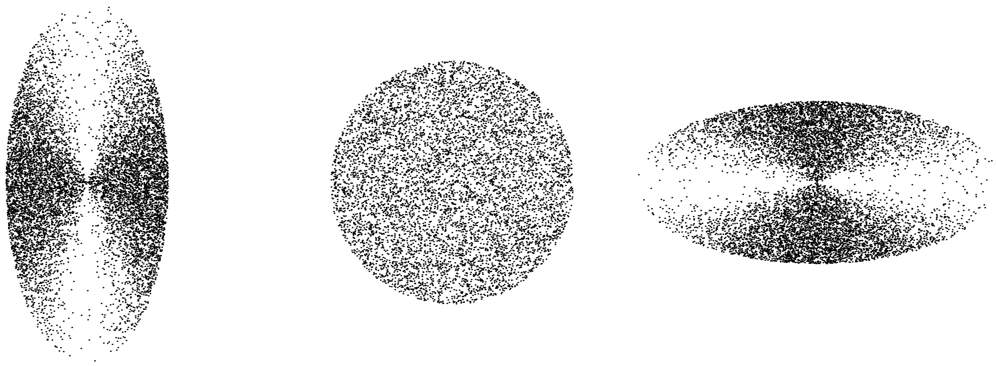

Figure 1.4: Generating polarization with gravity waves. The central figure illustrates a circular patch of homogeneous plasma in the absence of gravity waves. The right and left hand figures show how this patch is distorted by the passage of a (very intense) gravity wave. Just as the shape of the region is stretched along one axis and squashed along the other, the wavelengths of the photons propagating along these axes are distorted. This yields quadrupole moments in the apparent temperature of the radiation, as indicated by the shading. These quadrupole moments can in turn generate polarized signals as shown in Figure 1.3.

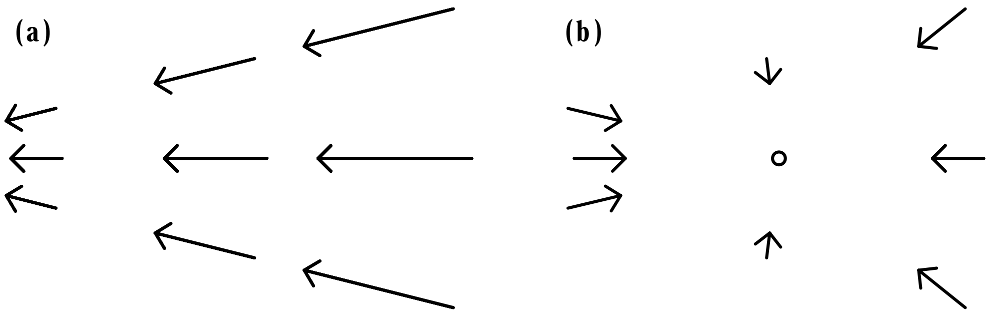

Figure 1.5: Velocity fields and quadrupole moments. (a) The velocity field for a small region of plasma. Material is converging to a point to the left under the influence of gravity, but is slowed down due to photon pressure. (b) The same region in the reference frame of one of the charged particles in the plasma. In this frame the fluid is moving in towards this particle from all directions, and the relative velocities have a distinct quadrupole moment. This results in a quadrupole moment in the redshifts of the radiation arriving at the charge, providing the quadrupole moment in the intensity of the radiation field necessary for the production of polarized radiation (see Figure 1.3)

In the reference frame of one of the charged particles in this plasma, the fluid appears to be converging on this charge. Due to the gradient in the flow, the relative velocities are different along different directions. The redshifts of the radiation emitted from the surrounding plasma towards the particle consequently have a quadrupole moment. This anisotropy in the redshifts creates a quadrupole moment in the apparent temperature of the incident radiation, which, in turn, yields polarization. The polarization produced in this manner comprises most ($> 90\%$) of the polarized component of the CMB.

1.2 The Structure of Cosmological Polarization

Since the polarization of the CMB arose from gradients in the velocity field of the primordial plasma, the structure of this cosmological polarization depends upon the dynamics of the early universe. The detailed measurements of these polarization anisotropies will therefore provide important cosmological information. Furthermore, fairly robust predictions regarding the size of the polarized signal can be made using the currently available data on the temperature anisotropies of the CMB. These temperature fluctuations reflect variations in the density of the plasma, not velocity fields (but see Hu and Sugiyama, 1995), so the temperature anisotropies do not directly determine the size of the cosmological polarization. Instead, the correlation between the temperature and polarization anistropies is due to the dynamics of the primordial plasma itself, which coupled the density variations to the bulk flows. The accurate modeling of these dynamics typically requires computer programs like CMBFAST (Seljak and Zaldarriaga, 1996) to solve the appropriate general relativistic Boltzmann equations. However, the fundamental features of these CMB anisotropies can be understood intuitively, as shown below.

1.2.1 Power spectra and the Structure of CMB Anisotropies

The structure of the CMB anisotropies is quantified using power spectra which indicate the *rms* amplitude of the variations as a function of angular scale (see Chapter 9 for a more

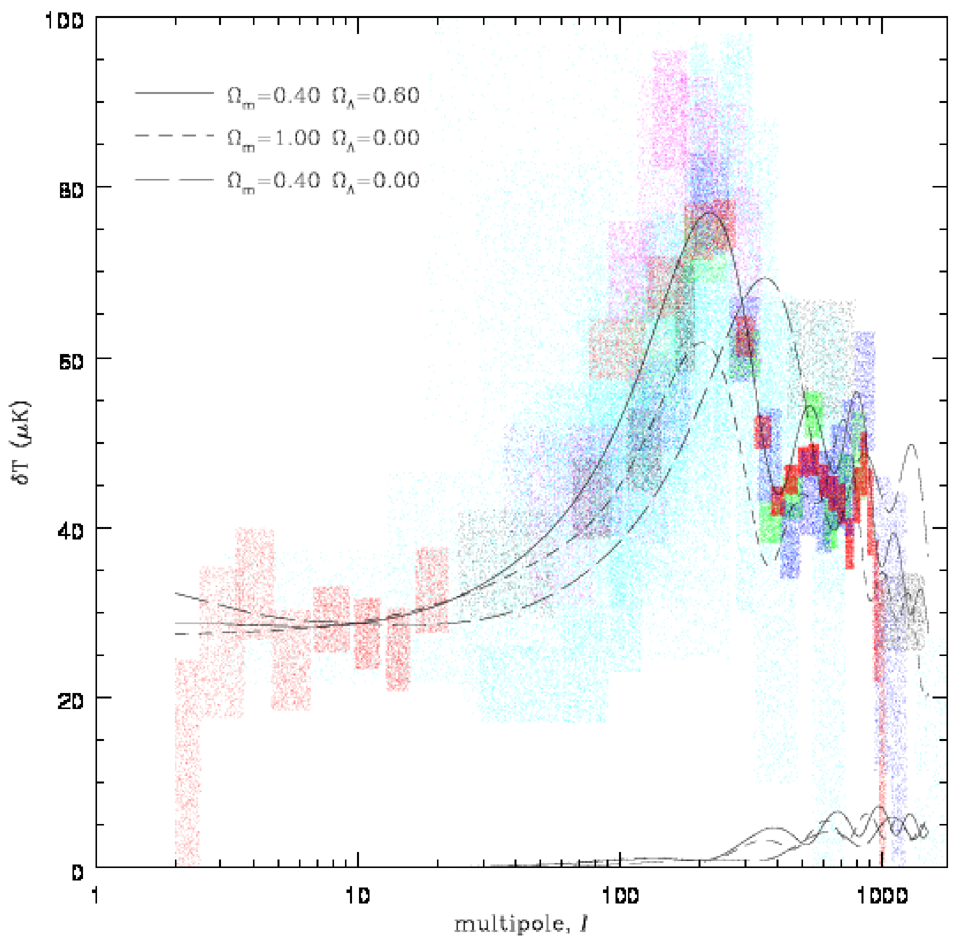

Figure 1.6: CMB power spectra, showing the size of the temperature and polarization anisotropies as a function of angular scale. Conventionally, the angular scale is parametrized by the multipole moment l from the spherical harmonic decomposition Y_m^l. (For an angular size θ, $l \simeq 120°/\theta$.) The ordinate axis gives the *rms* amplitude of the variations per logarithmic interval in l (which technically makes this a plot of the spectral density, or the square root of the power spectrum). The curves are the power spectra predicted by various cosmological models obtained from CMBFAST (Seljak and Zaldarriaga, 1996), for both the temperature anisotropies and the polarization anisotropies. The polarization power spectra are the three curves peaking in the lower right corner of the plot. Note that alterations in the assumed cosmological model produce concordant changes in the temperature and polarization power spectra. The shaded boxes represent the results of measurements of the actual temperature anisotropies (Wang et al., 2001). A fixed number of dots per box has been used to give more visual weight to experiments with smaller error bars.

detailed discussion of how these power spectra are obtained by decomposing the fluctuations on the sky into spherical harmonics). The power spectra of both the temperature and polarization anisotropies (as well as the non-trivial power-spectra of the cross-correlation between the temperature and polarization) for different cosmological models can be calculated using Boltzmann codes like CMBFAST. The exact shapes of these curves depend upon the parameters assumed by the model (see Figure 1.6), which allows cosmological information to be extracted from measurements of the power spectra. The complex relationships between the detailed shape of the predicted power spectra and the assumed cosmological parameters cannot be dealt with here. However, an introduction to the relatively simple physics underlying the basic pattern of bumps and wiggles common to all the illustrated power spectra will be sufficient to demonstrate the close relationship between the temperature and polarization anisotropies and the plausibility of current expectations regarding the detectability of the cosmological polarized signal (for more detail, see Hu et al. (1997); Hu and Sugiyama (1995); Kodama and Sasaki (1984)).

1.2.2 Initial Conditions and Inflation

The one factor affecting the shape of the power spectra that involves exotic physics is the assumed model of structure formation, which determines how variations in the primordial plasma were created in the first place. Rather than detailing the zoo of proposed structure formation models, only the adiabatic inflationary model will be considered here (this choice of model is motivated by the fact that it fits the observed temperature anisotropies better than the others, see Stompor et al. (2001)). Inflationary models generate density perturbations through the rapid growth of quantum fluctuations during the inflationary epoch. The perturbations generated in this way have a very simple structure: random gaussian noise with a nearly scale-invariant spectrum (Liddle and Lyth, 2000). Adiabatic models further stipulate that while the overall density of the universe varied, the fraction of the material in any particular form (neutrinos, photons, baryons, cold dark matter, etc.) was initially a constant throughout space (Liddle and Lyth, 2000). Since the variations in the densities of different components did not cancel each other out, this model also yields fluctuations in the curvature of space-time (i.e. the gravitational potential). Therefore, in the adiabatic inflationary scenario, a well defined pattern of initial perturbations was created in the first fraction of a second after the big bang.

1.2.3 The Evolution of Perturbations and Acoustic Oscillations

The evolution of the initial perturbations laid down by inflation was determined entirely by the equations of general relativity. These equations are imposing nonlinear formulae. However, since the perturbations were small (CMB anisotropies are only about one part in 10^5), the relevant equations can be well approximated with linear expressions. This linearity not only makes the equations much easier to understand, but also allows the perturbations to be decomposed into normal modes. Each normal mode consists of a sinusoidal (plane wave) variation in the density of the universe and the curvature of space- time, and since each mode evolved independently, the evolution of the structure at each length scale (which corresponds roughly to a single angular scale in the CMB) can be studied separately.

The dynamics of a particular mode depended strongly on the size of its wavelength relative to the distance a photon could have traveled since the big bang. This distance, known as the horizon scale, increased with cosmic time faster than the wavelength, so at sufficiently early times this distance was much smaller than the wavelength of any given mode. As long as the perturbation wavelength greatly exceeded the horizon scale, regions of different densities were not in causal contact, so there was no tendency for material to flow between different regions. The amplitude of the density perturbation therefore remained essentially constant and there was little, if any, bulk flow. Modes with sufficiently long wavelengths remained in this frozen state all the way until decoupling, when the universe was no longer warm enough to maintain the primordial plasma, and the photons that form the CMB were released. These frozen (acausal) modes, which correspond to the largest angular scales on the sky, produced variations in the temperature of the CMB which reflect the initial perturbations laid down during inflation. However, since there was little bulk flow on these large scales, the polarized signal should be extremely small (but see Keating et. al., 1997).

Modes with short enough wavelengths reached a point before decoupling when their wavelength became comparable to the horizon scale. Regions of different densities then came into causal contact and local forces acted to redistribute material. Two forces were particularly important: gravity and photon pressure. Gravity tended to move material from underdense regions to overdense regions, while photon pressure acted in the opposite direction. The initial perturbations laid down by inflation were such that these two forces were out of balance, so material flowed back and forth between underdense and overdense regions. These acoustic oscillations continued until decoupling, when the resulting density contrasts and bulk flows left their mark on the small-scale ($< 1^\circ$) anisotropies of the CMB.

1.2.4 Acoustic Oscillations and Apparent Temperature Variations

The ebb and flow of the plasma during the acoustic oscillations can be better understood if gravity and photon pressure are combined into a single restoring force. This can be done by considering how gravitational redshifts affect the apparent temperature of a source. Imagine an object at some intrinsic temperature (the temperature an observer located at the source would report) is sitting at the bottom of a potential well. As the thermal radiation from the source climbs out of the potential well, the photons are redshifted such that the source appears to be cooler than its intrinsic temperature. Conversely, a source located at a potential maxima will appear warmer than its intrinsic temperature when the photons blueshift as they fall down the slope in the potential. Note that while the apparent temperature of a source depends on the location of the observer, the difference in the apparent temperatures of two sources is the same for all observers.

In the primordial plasma, overdense regions were both warmer than underdense regions (by the ideal gas law) and at lower gravitational potentials. In this situation, the gravitational redshifts cause the apparent temperature variations to be smaller than the intrinsic temperature variations (see Figure 1.7). In principle, the gravitational redshifts could balance the intrinsic temperature variations exactly, so that there would be no variations in the apparent temperature. In this perfectly balanced case, the plasma would appear homogeneous, and there could be no net force acting to redistribute the plasma. If

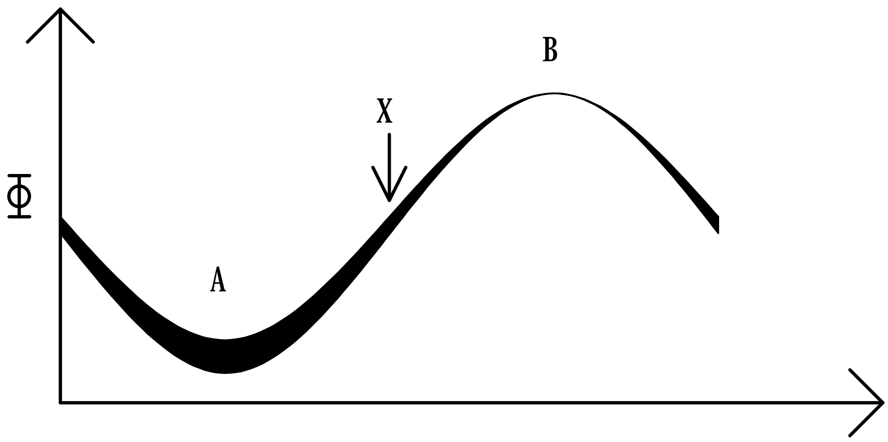

Figure 1.7: Apparent temperature fluctuations. A section of the plasma is shown with a sinusoidal variation in the (intrinsic) temperature of the plasma (illustrated by the varying thickness of the line) and the gravitational potential (indicated by the varying height of the line). For the observer at point X the radiation emitted from point A is redshifted, causing the plasma to appear colder, while the radiation emitted from point B is blueshifted, thus appearing warmer. Therefore the observer at X sees an apparent temperature difference between points A and B that is smaller than the intrinsic temperature difference.

the variations in the gravitational potential were slightly smaller (relative to the intrinsic temperature variations), not only would overdense regions appear warmer than underdense regions, but also the net gravitational force would be slightly weaker than the force due to photon pressure (material would tend to move out of overdense regions into underdense regions). Conversely, slightly larger variations in the gravitational potential would result in the overdense regions appearing cooler than underdense regions and the gravitational force overcoming the pressure force (material would tend to move out of underdense regions into overdense regions). In both of these cases material moves from apparently warmer regions to apparently cooler ones. Therefore the plasma dynamics can be described in terms of a single restoring force acting to minimize the apparent temperature variations.

The initial conditions laid down by inflation were such that each mode originally had a nonzero apparent temperature variation (Kodama and Sasaki, 1984). Consequently, once the horizon scale became larger than the wavelength of a mode, the size of the apparent temperature variation in that mode oscillated about zero like a harmonic oscillator.

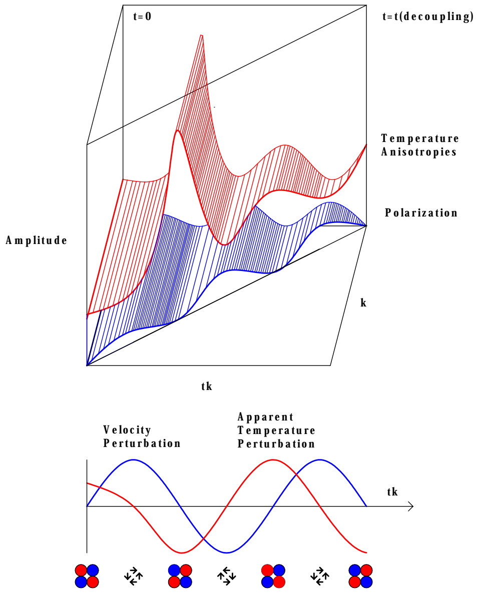

Figure 1.8: The evolution of structure in the CMB. At the bottom of this figure is a graph that shows the evolution of the amplitudes of the apparent temperature and velocity perturbations of a single mode as a function of tk, where k is the mode wavenumber and t is cosmic time (The ratio of mode wavelength to horizon scale, which determines the evolution of the oscillation, is directly proportional to tk). The size of the velocity and temperature perturbations oscillate about zero as material ebbs and flows between apparently hot and cold spots, as shown by the icons at the bottom of the plot. The upper part of the figure shows a graph with the axes, k, tk and the rms of the variations. The left side of this box corresponds to $t = 0$, so on this wall of the box the initial conditions are displayed. The polarization is zero here because the velocity perturbations were initially zero, while the temperature fluctuations have a flat power spectrum due to the assumption of scale invariant initial conditions. Each of the modes evolved as a function of tk as shown in the lower part of the figure, indicated by the rippled surfaces (which are positive definite because the rms of the variations are shown). Since the lower axes are k and tk, a single moment of time corresponds to a vertical plane intersecting the z-axis, and increasing time corresponds to rotating the plane. The plane which forms the edge of the surfaces corresponds to the epoch of decoupling. On this plane, modes with larger k (smaller wavelength) correspond to higher values of tk (a later point in the oscillations), so that the power spectrum at decoupling recapitulates the evolution of a single mode. These power spectra are in terms of length scales: the angular power spectra of the CMB have similar shapes because $l \sim k$.

1.2.5 Acoustic Oscillations and CMB Anisotropies

The oscillations in the apparent temperature of the primordial plasma were responsible for most of the observed structure in the CMB power spectra shown in Figure 1.6. Earthbound telescopes observe from a single point in space, so the measured temperature anisotropies of the CMB are approximately the variations in the apparent temperature of the plasma at decoupling. At progressively smaller angular scales on the sky (higher l), the observed modes have smaller wavelengths, so they began oscillating earlier prior to decoupling, and at decoupling they had progressed through more oscillations, yielding the familiar pattern of acoustic peaks in the power spectrum of the temperature variations (see Figure 1.8).

The oscillations in the plasma were responsible for the peaks in the polarization power spectrum as well. As the material oscillated back and forth between the overdense and underdense regions, the magnitude of the velocities in the plasma rose and fell 90° out of phase with the oscillations in the apparent temperature variation (The material hardly moved when the apparent temperature variations were large and the restoring force was reversing the direction of flow, and conversely the velocities were large when the apparent temperature variation was nearly zero). The polarization power spectrum therefore has peaks where the temperature power spectrum has troughs, and vice versa. The peaks in the polarization power spectrum are significantly lower than those of the temperature anisotropies because the thermalizing effect of Thomson scattering discussed in Section 1.1 greatly reduced the efficiency of converting velocity perturbations into cosmological polarization (Kosowsky, 1999).

If the plasma were an ideal fluid, there would be fluctuations in the temperature and polarization of the CMB on infinitely small scales. However, the fluid was not ideal, and the photons propagated finite distances between scattering events. On distances shorter than this mean free path (which became large at decoupling), variations in the temperature and polarization of the photon gas were suppressed, resulting in the damping tails at high l apparent in both power spectra.

This short, qualitative summary of the physics behind the CMB anisotropies is enough to demonstrate that the temperature and polarization anisotropies are connected through sensible physical processes (see Figure 1.9), and that it is reasonable to use the existing measurements of the temperature anistropies as a guide for estimating the size of the cosmological polarization. The TOCO, BOOMERanG, MAXIMA and DASI temperature anisotropy experiments all show a strong peak in the power spectrum with a peak amplitude of about 80μK at angular scales of about 0.5° (Miller et al., 1999; Netterfield et al., 2001; Lee et al., 2001; Halverson et al., 2001), which implies a flat universe with a significant cosmological constant. The location of this first acoustic peak also sets the angular scale of the expected peaks in the polarization power spectrum, and the height of the measured peak indicates the strength of the polarized component. The cosmological polarization therefore should be largest on sub-degree angular scales, where it will be roughly 5μK rms at most (Figure 1.6).

1.3 Foregrounds

Unfortunately, the CMB is not the only source of polarized microwaves. While the atmosphere should be completely unpolarized, Galactic synchrotron radiation, interstellar dust emission and extragalactic point sources produce polarized signals that can contaminate measurements of the cosmological polarization. There is little experimental information on these polarized foregrounds, and no data at all on the angular scales and frequencies relevant for CMB experiments. Rough estimates of the foregrounds must therefore be extrapolated from other frequencies using their (still somewhat uncertain) spectra (see Chapter 9 for details). These estimates, illustrated in Figure 1.10, suggest that synchrotron radiation is the strongest source of polarized microwaves at frequencies below 50 GHz, while for frequencies above 200 GHz dust emission is dominant. Between these two frequencies is a window where the cosmological polarization can (hopefully) be measured with relatively little foreground contamination (The location of this window is relatively insensitive to uncertainties in the normalization of the foreground spectra, see Tegmark et. al. 1999).

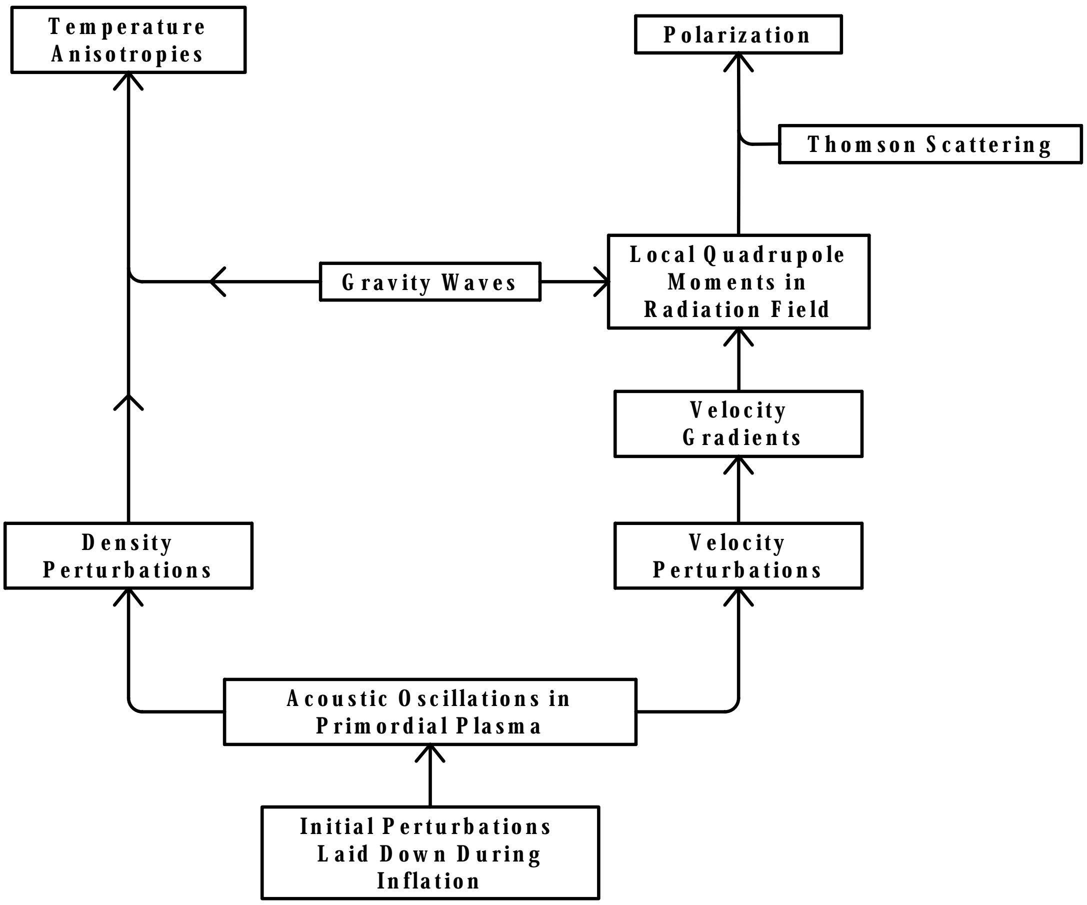

Figure 1.9: Summary of the physics behind the structure of the CMB

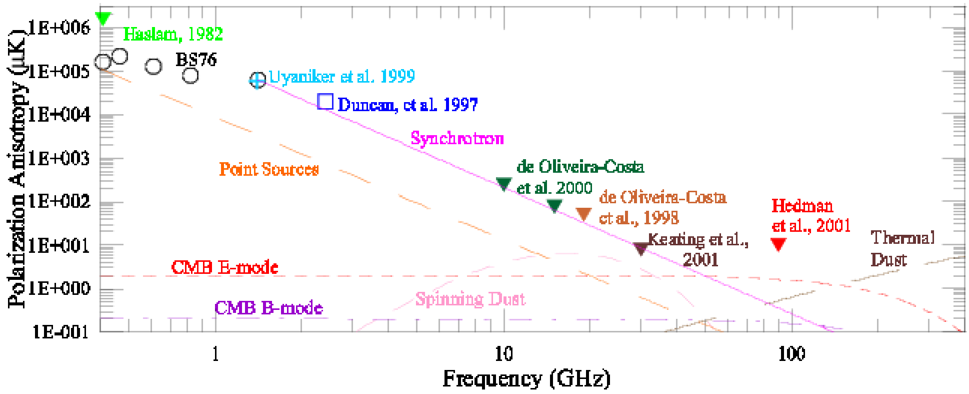

Figure 1.10: Polarized foregrounds on sub-degree angular scales

Chapter 2

Overview

PIQUE was built specifically to search for the polarized signal from the CMB, which should be (1) very faint (only a few μK at most), (2) strongest on sub-degree angular scales, and (3) least contaminated by foregrounds around 100 GHz (see Chapter 1). Both these predictions regarding the signal and a host of more practical considerations strongly influenced the experimental design, which is summarized in this chapter. More detailed information on various aspects of the instrument can be found in later chapters. The polarimeter itself is discussed at length in Chapter 3; the optics system is described in Chapter 4; and an overview of the observations is given in Chapter 5. The tests used to calibrate the operational telescope are reviewed in Chapter 6 and the atmosphere at the site is studied in Chapter 7. The last two chapters deal with the data obtained from the two observing seasons. Chapter 8 describes how corrupted data were removed from the time stream to produce the final data set, which was subjected to the likelihood analysis described in Chapter 9 to produce constraints on the cosmological polarization.

2.1 The Stokes Parameters

The IQU in PIQUE refers to the Stokes parameters, which provide the necessary formalism for describing the polarization states of electromagnetic radiation (Kraus, 1986). The intensity of the radiation is given by the parameter I, while the parameters Q and U are measures of the linear polarization. These two parameters allow both the magnitude and

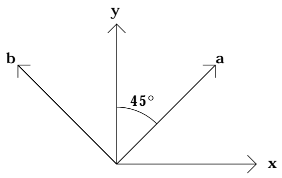

Parameter	a-b System	x-y System
I	$< E_a^2 > + < E_b^2 >$	$< E_x^2 > + < E_y^2 >$
Q	$< E_a^2 > - < E_b^2 >$	$2 < E_x E_y >$
U	$2 < E_a E_b >$	$< E_y^2 > - < E_x^2 >$

Figure 2.1: The definitions of the Stokes parameters

orientation of the linearly polarized component to be determined.[1] A fourth parameter, denoted V, measures the degree of circular polarization. The CMB is expected to have no detectable circular polarization, and PIQUE is insensitive to V, so this parameter will not be considered further.

The definitions of these parameters and the relationships between them can best be understood by decomposing the electromagnetic field into its orthogonal components in two different coordinate systems rotated by 45° with respect to each other , as shown in Figure 2.1 (A little vector algebra demonstrates the definitions of I, Q and U in these two coordinate systems are equivalent). In this formalism, Q is the difference in power carried by the vertical and horizontal components of the electric field, while U is the difference between the power carried in the two diagonal components. If Q and U are known, then it is straightforward to determine both the magnitude ($\sqrt{Q^2 + U^2}$) and the orientation ($\frac{1}{2}\tan^{-1}(Q/U)$) of the linearly polarized component. However, the observations and the analysis presented here are more naturally described in terms of Q and U themselves.

The definitions of Q and U depend on the orientation of a set of coordinate axes, so a coordinate system must be specified to render these definitions unambiguous. This work will follow the convention established by the International Astronomical Union (IAU), which states that a polarized signal oriented radially with respect to the celestial poles has positive Q and zero U (IAU, 1973). This convention requires that an azimuthally polarized signal will have negative Q and zero U, a polarized signal rotated 45° counterclockwise from radial with have zero Q and positive U, and a polarized signal rotated 45° clockwise from radial will have zero Q and negative U (Figure 2.2).

[1]For CMB measurements I derives mainly from the unpolarized radiation, which is roughly a million times stronger than the polarized component.

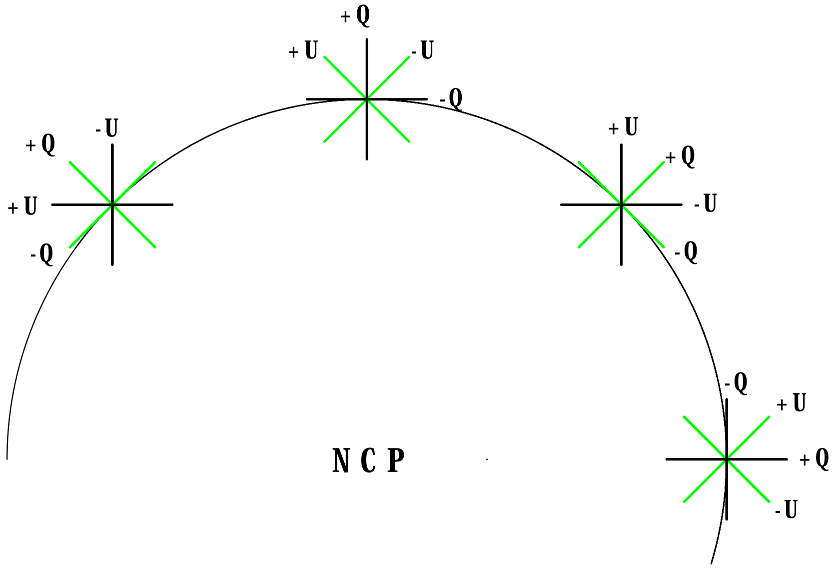

Figure 2.2: The International Astronomical Union (IAU) convention for the Stokes parameters.

While the Stokes parameters, as defined in Figure 2.1, have units of power, the quantities measured by CMB experiments conventionally have units of temperature. For thermal radiation (in the Rayleigh-Jeans limit, see Section 3.3.2), the intensity I is directly proportional to the physical temperature T of the source, so $T = I/K$, where K is a constant unit conversion factor. By analogy, one can construct parameters equivalent to Q and U with units of temperature using the same conversion factor: $Q_T = Q/K$ and $U_T = U/K$.

The Stokes parameters with units of temperature must be handled with care. Non-bolometric microwave receivers are sensitive to only one polarized component of the radiation field, and the power carried in this single component has often been used to estimate I and T. This is reasonable for essentially unpolarized sources such as the CMB. However, when the polarization is the parameter of interest, equating the response of such a receiver with the effective temperature of the incident radiation can cause confusion. Say two receivers measure two orthogonal components of the electric field, yielding two different estimates of the effective temperature, T_x and T_y. It might be tempting to say that the stokes parameter Q_T is the difference between these two estimates, $Q_T = T_x - T_y$, but this is incorrect. Such a calculation is exactly analogous to equating T with the sum of the two estimates, $T = T_x + T_y$. Of course, the correct expression for T is the average of the two estimates, $T = (T_x + T_y)/2$, and similarly the correct expression for Q_T is $Q_T = (T_x - T_y)/2$ (for examples of situations in which this factor of two is important, see Sections 2.2.1, 3.3.2).

2.2 The Instrument

The PIQUE instrument consists of a single correlation polarimeter viewing a 1.2 m off-axis parabolic dish through a corrugated feed horn. Two separate polarimeters were used in this telescope at different times: a Q-band (35-46 GHz) system, which will be discussed in detail elsewhere (Barkats, 2002), and a W-band (84-100 GHz) system, which is described at length here. Note that the W-band polarimeter operates in the frequency range where astrophysical foreground contamination is expected to be minimal (see Section 1.3).

2.2.1 The Correlation Polarimeters

Correlation polarimeters can extract small polarized components from largely unpolarized radiation by taking advantage of the phase information in the incident electromagnetic fields (see Figure 2.3). Say the polarimeter is set up to measure the difference in the power carried by the two (real) orthogonal components of the electromagnetic field E_a and E_b in a certain coordinate system (referred to as the "detection" axes). The instrument splits this radiation into two orthogonal components E_1 and E_2 in a coordinate system rotated $45°$ with respect to the detection axes (referred to as the "line" axes); thus $E_1 = E_a + E_b$ and $E_2 = E_a - E_b$. These two components are amplified, filtered and down converted before they finally reach the two inputs of a multiplier. This device produces an output voltage proportional to the product of the two input fields (see Appendix A): $E_1 E_2 = (E_a - E_b)(E_a + E_b) = E_a^2 - E_b^2$. The DC voltage out of the multiplier is therefore, ideally, proportional to one of the Stokes parameters (Q or U) of the input radiation. In reality, the multipliers also produce undesirable nonzero DC output voltages proportional to E_1^2 and E_2^2, the total power in each of the two input lines (see Appendix A). The real polarized signals

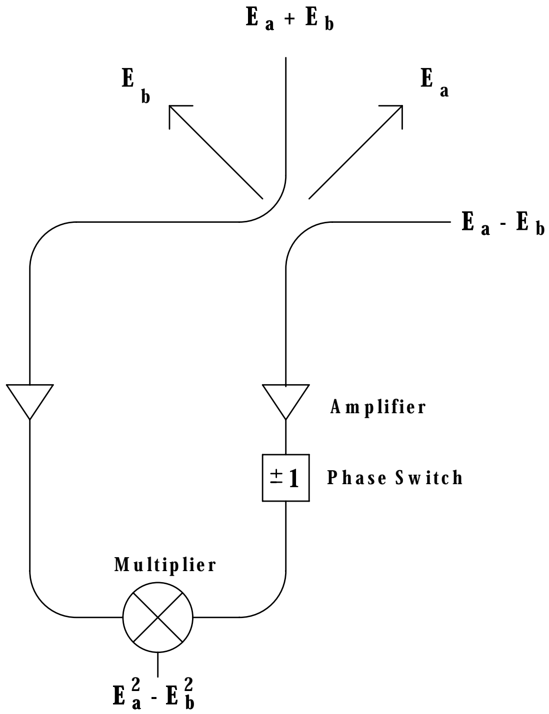

Figure 2.3: The essentials of a correlation polarimeter. A correlation polarimeter measures one of the Stokes parameters ($E_a^2 - E_b^2$) by splitting the incident electromagnetic field into two components ($E_a - E_b$ and $E_a + E_b$) and then multiplying these two components together. The phase switch in one line serves to select true polarized signals from spurious offsets by modulating only the polarized signal.

are extracted from these offsets with the aid of a phase switch (see Figure 2.3). This device multiplies the electric field in one of the lines by a ± 1 square wave, which modulates the sign of the desired polarized signal. Appropriate filtering and processing of the multiplier output then isolates this modulated signal from the unmodulated offset terms.

Modulating the desired signal, or Dicke-switching, also allows the extremely small cosmological polarization to be extracted from the largely unpolarized microwave radiation from the CMB, atmosphere and the radiometer itself. This unpolarized radiation (which is about 10^8 times larger than the expected polarized signal) generates variations (i.e. noise) in the output voltage of the polarimeter, which must be averaged out to find any polarized signal. The smallest polarized signal measurable with a correlation polarimeter is (Kraus, 1986):[2]

$$\Delta Q_{min} = T_{sys}\sqrt{\frac{1}{\Delta\nu * t_{int}} + \left(\frac{\Delta G}{G}\right)^2}.$$

The system temperature T_{sys} is the effective temperature of the radiation entering the polarimeter, which includes significant contributions from the atmosphere and the radiometer itself (the contribution from the CMB is an order of magnitude smaller), which are referred to as the atmospheric noise temperature and the receiver noise temperature, respectively (see Chapters 3 and 7). The two terms under the square root correspond to two different types of noise. In the first term, $\Delta\nu$ is the bandwidth of the polarimeter and t_{int} is the integration time. This term represents the random gaussian fluctuations (thermal noise) due to the finite temperature of the sources of the unpolarized radiation. With sufficient integration time, this term can be made arbitrarily small. In the second term $\Delta G/G$ is the fractional variations in the gain G of the amplifiers. These variations have a characteristic frequency spectrum $\Delta G/G \propto f^{-\alpha}$ with $\alpha \simeq 1$. This so-called $1/f$ noise is characterized by long time scale drifts. Increasing the integration time includes ever larger drifts, and therefore this type of noise cannot be reduced through repeated measurements.

Dicke switching controls the $1/f$ component of the noise because the amplitude of the modulated signal is only affected by drifts on time scales shorter than the modulation period (Kraus, 1986). Phase switches can modulate the signal at very high frequencies ($> 10^3$Hz, as opposed to the few Hz possible with mechanical chopping strategies) and therefore strongly suppress the $1/f$ noise (but see Section 3.3.3). The smallest measurable modulated signal is therefore given by

$$\Delta Q_{min} = \frac{T_{sys}}{\sqrt{\Delta\nu}}\frac{1}{\sqrt{t_{int}}} = \frac{\mathcal{S}}{\sqrt{t_{int}}},$$

where $\mathcal{S}$ is the sensitivity of the polarimeter. In order to minimize the integration time needed to detect the cosmological polarization, the sensitivity must be as small as possible,

[2]This expression differs by a factor of $\sqrt{2}$ from the expression for the smallest measurable temperature difference ΔT_{min} for a correlation receiver given by Kraus (1986) in Table 7.3. Kraus assumes temperature is estimated from a measurement of the power carried in one polarized component of the radiation. As shown in Section 2.1, the polarization $Q_T = (T_x - T_y)/2$, where T_x and T_y are estimates of the source temperature based on measurements of one polarized component. With standard error propagation the smallest measurable polarization can be shown to be $\Delta Q_{min} = \frac{1}{2}\sqrt{\Delta T_x^2 + \Delta T_y^2} = \Delta T/\sqrt{2}$

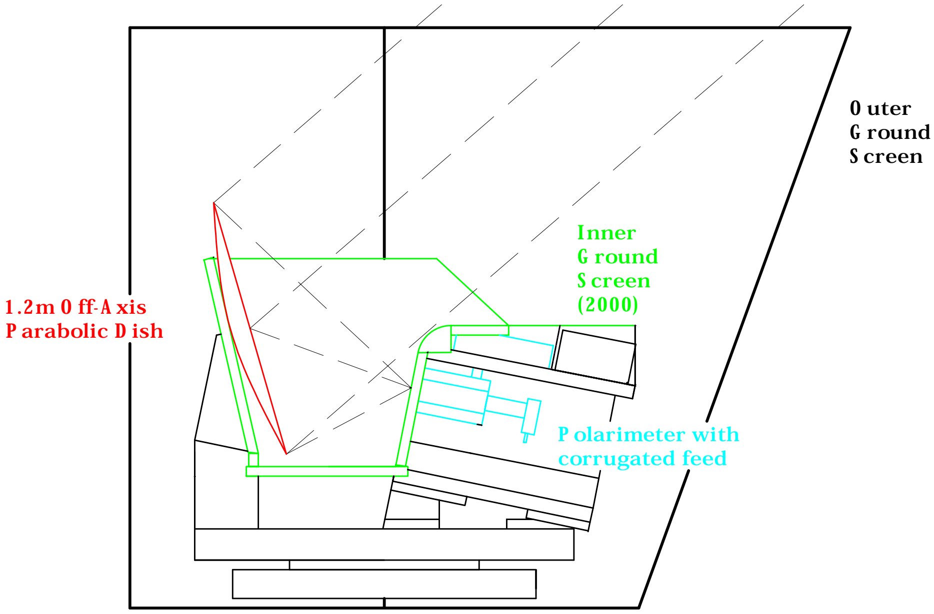

Figure 2.4: PIQUE optics. The corrugated feed horn of the polarimeter views the sky around the NCP through the parabolic mirror. A crane bearing in the telescope base allows the telescope to scan in azimuth at a fixed elevation

requiring both large bandwidths and low noise temperatures (for the W-band system, $\Delta\nu >$ 10 GHz and $T_{sys} <$ 150K, see Chapter 3).

In addition to the rapid Dicke switching allowed by the phase switches, correlation polarimeters have the advantage that they take the product of two fields, not the difference of two powers. Accurate measurements of power differences are difficult because drifts in the relative gains of the two lines can generate spurious nonzero output signals even if the input radiation is completely unpolarized. This is not a problem in a correlation polarimeter because the response of the receiver depends only on the product of the gains in the two lines. Correlation polarimetry is therefore a particularly clean method for measuring small polarized signals. However, while a correlation polarimeter is insensitive to the relative gains of the two lines, it is sensitive to relative phase shifts between the lines (a nonzero relative phase shift ϕ between the lines reduces the output voltage by a factor of $\cos\phi$). Consequently, considerable effort goes into insuring the two arms of the receiver are well phase matched (see Section 3.2).

2.2.2 Optics

The operational polarimeter views the sky through a scalar corrugated feed horn, which illuminates a 1.2 meter off axis parabolic dish recycled from the Saskatoon experiment (Figure 2.4; Wollack, 1997). The base consists of a crane bearing which allows the telescope to scan back and forth in azimuth. The elevation of the telescope is fixed, which greatly simplified construction and pointing. Two nested sets of ground screens surround the telescope in order to minimize ground (ambient 300K) pickup. The optics are kept extremely simple in order to minimize spurious polarized signals from either reflected or diffracted radiation.

With the W-Band polarimeter, the telescope has a beam full width at half maximum (FWHM) of about a quarter of a degree on the sky (see Chapters 4 and 6). As shown in Chapter 9, this enables the telescope to probe the angular scales where the polarized signal should be detectable. The telescope could have been even more sensitive to these small angular scales if the beam were smaller, but the beamsize is limited by the size of the readily available dish.

2.3 Observations and Scan Strategy

The PIQUE telescope observed from the roof of Jadwin Hall, Princeton NJ (Latitude 40.346°, West Longitude 74.650°). The W-band polarimeter took data during the entire first observing season (January 1, 2000 to May 5, 2000, referred to PIQUE2000), and the first half of the second observing season (December 15, 2000 to February 28, 2001, called PIQUE2001) before being replaced by the Q-band system. During both of these periods, the telescope measured the polarization of the sky around the North Celestial Pole (NCP) whenever weather (and the telescope) permitted (see Chapter 5).

The finite sensitivity of the radiometer strongly influenced the choice of scan strategy. As shown in Figure 2.5, the telescope measures the polarization on a single ring around the NCP. The radius of this ring is set at 1° as a compromise between improving the signal to noise per pixel (which favors smaller observing regions) and having sufficient pixels to avoid problems with sample variance (which favors larger observing regions). If the telescope had stared at a single point fixed in elevation and azimuth, then the polarized signal everywhere on a ring would be measured each day as the sky rotates around the NCP. The atmosphere and radiometer are not sufficiently stable for this to be a practical observation method (see Section 3.3.2, Chapter 7), so the telescope instead moves back and forth between two points on this ring every few seconds.[3] Taking the difference between the measured signals at these two points not only eliminates long time scale drifts, but also introduces useful systematic checks into the data set (see Chapter 8).

During the first observing season, the telescope pointed slightly above the NCP and the detection axes of the polarimeter were oriented (as shown in Figure 2.5) so that the instrument measured the Q Stokes parameter (as defined by the IAU, see Figure 2.2). The second observing season was dedicated to measuring the other Stokes parameter, U. This could have been done by observing the same points on the sky and rotating the polarimeter by 45°. However, rotating the detection axes in this way would have made the telescope more sensitive to spurious polarized signals from radiation diffracting around the vertical and horizontal edges near the telescope. PIQUE was therefore made sensitive to U by modifying the telescope mount (using a tapering shim to put the optical elements on a slight tilt) so that it observed a different pair of points on the ring (see Figure 2.5).

[3] The two scan positions are at the same elevation because the telescope can only scan in azimuth. This insures the integrated column density of the atmosphere along the line of sight is the same in the two positions.

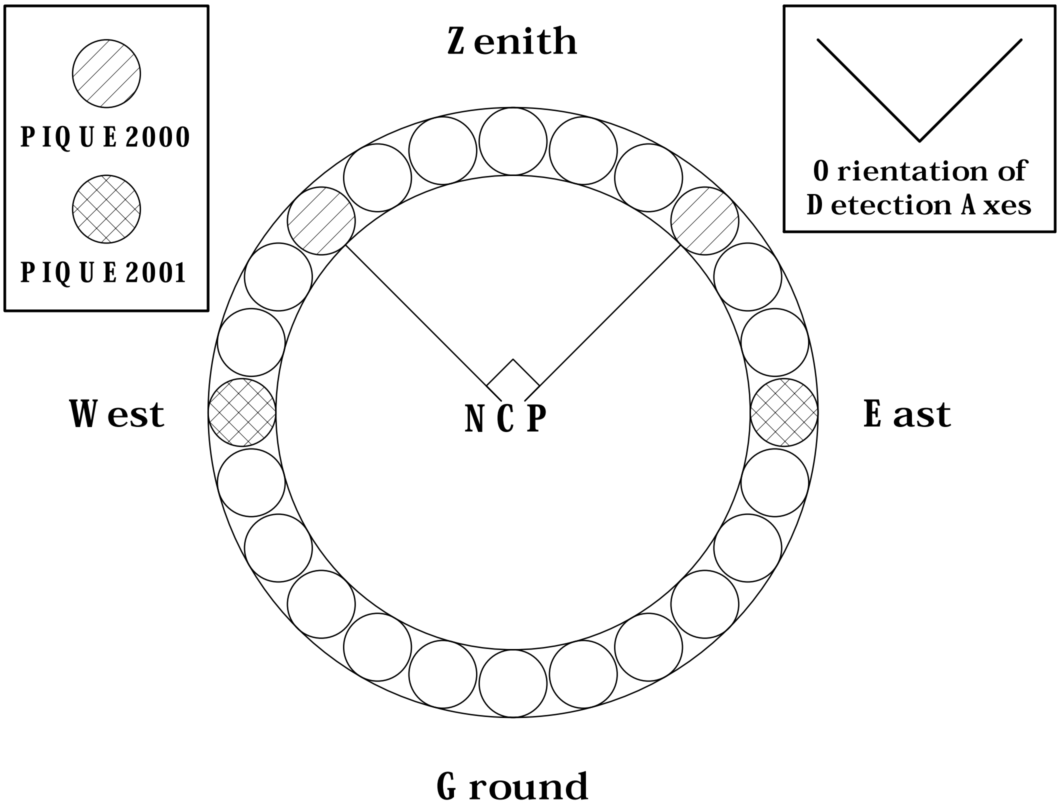

Figure 2.5: Scan Strategy: PIQUE measured the polarization on a ring of 1° radius around the North Celestial Pole, illustrated here as a ring of circles. Each circle is the beam FWHM across, indicating the number of independent pixels around the ring. The orientation of the detection axes are shown in the upper right corner. During PIQUE2000, the telescope chopped between the two diagonally hatched spots, while during PIQUE2001 the telescope chopped between the two cross-hatched spots. Observing different points on the ring allowed both stokes parameters (Q and U) to be measured (see Figure 2.2).

Chapter 3

The Polarimeter

The PIQUE telescope uses correlation polarimeters to search for the polarized component of the CMB (see Section 2.2.1). The W-band polarimeter is described in detail in Section 3.1 below, and a components list is provided in Appendix B (the Q-band polarimeter will be dealt with in Barkats, 2002). Section 3.2 describes how the two lines of the receiver are phase matched so that it can function as a polarimeter. The series of laboratory tests described in Section 3.3, are used to determine the salient characteristics of the radiometer, such as its responsivity, bandpass and sensitivity. During the suite of tests prior to the first observing season, the unexpected features of two components (the cryogenic IF amplifiers and the multipliers) led to significant changes in the operation of the polarimeter, as discussed in Section 3.4.

3.1 Description

The W-band polarimeter illustrated in Figures 3.1 and 3.2 is a heterodyne device, which means the radio (RF) frequency (84-100 GHz) radiation from the sky is mixed down to a lower (IF) frequency (2-18 GHz) before reaching the multipliers. Such a polarimeter is naturally divided into RF, LO and IF sections. The RF and IF portions of the radiometer (described in Sections 3.1.1 and 3.1.3) carry the radiation before and after the frequency conversion, while the LO lines (described in Section 3.1.2) provide a continuous wave 82 GHz signal from a local oscillator (LO) source to the mixers (the devices responsible for

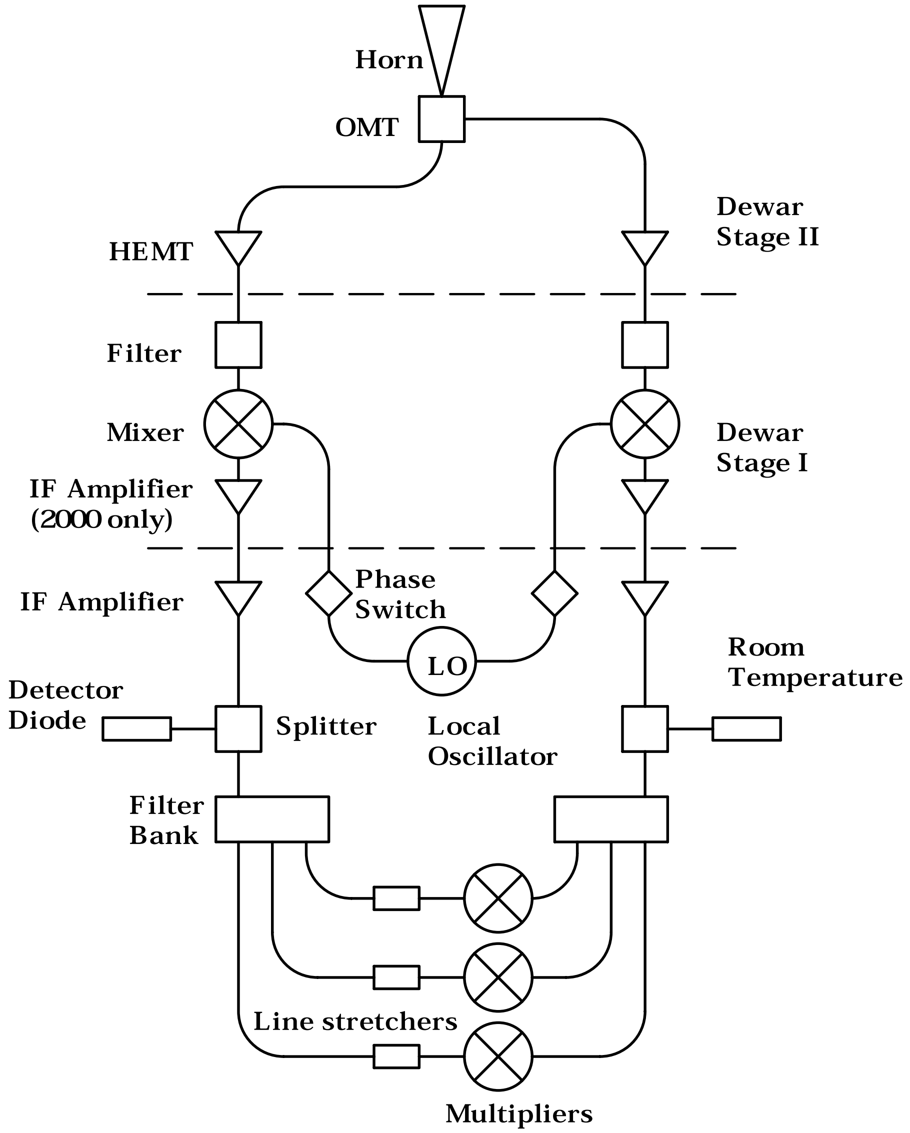

Figure 3.1: Block diagram of the W-band polarimeter

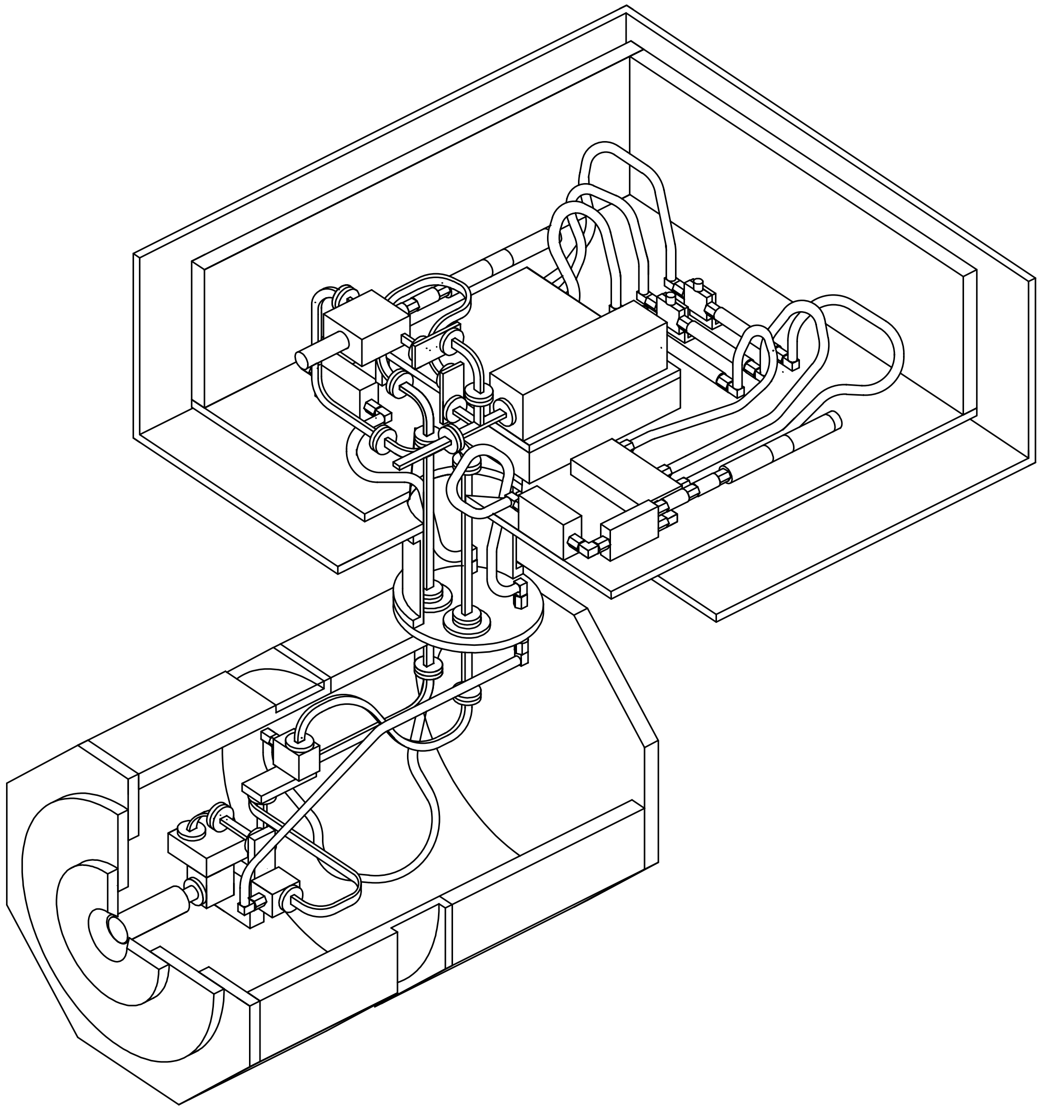

Figure 3.2: Isometric drawing of the W-band polarimeter (2001). Detail views, with specific components indicated, are provided in Figures 3.3, 3.4 and 3.5. The components are also enumerated in the text. The octagonal prism is the dewar that contained the cryogenic portion of the radiometer, while the box on top is the "E-Box" that contains most of the IF and LO components. For scale, the dewar is 9.5″ across.

Table 3.1: Summary of polarimeter responsivity

Component	Physical Temperature[a] (2000/2001)	Gain[b]	Effective Noise Temperature[c]	Signal level (Polarization)[d]	Signal level (Total Power)[d]
Horn	40K/20K	<-0.05 dB	>0.3K	$1.1 * 10^{-12} \mu W$	$1.1 * 10^{-5} \mu W$
OMT	30K/17K	-0.05dB	0.3K	$1.1 * 10^{-12} \mu W$	$1.1 * 10^{-5} \mu W$
HEMT	30K/18K	35 dB	75K	$3.5 * 10^{-9} \mu W$	$8.8 * 10^{-2} \mu W$
E-Bend	30K/18K	-0.3 dB	5 mK	$3.3 * 10^{-9} \mu W$	$8.3 * 10^{-2} \mu W$
Stainless Steel wg.	50K/30K	-0.7dB	10 mK	$2.8 * 10^{-9} \mu W$	$7.0 * 10^{-2} \mu W$
Filter	80K/50K	-0.6dB	10 mK	$2.4 * 10^{-9} \mu W$	$6.0 * 10^{-2} \mu W$
Mixer	80K/50K	-10 dB	72 mK	$2.4 * 10^{-10} \mu W$	$6.0 * 10^{-3} \mu W$
Cold IF Amp (2000)	100 K	35 dB	3K	$7.2 * 10^{-7} \mu W$	$18 \mu W$
Attenuators (2000)	280 K	-30 dB	—	$7.2 * 10^{-10} \mu W$	$0.018 \mu W$
Warm IF Amp (2000)	305 K	30 dB	—	$7.2 * 10^{-7} \mu W$	$18 \mu W$
Warm IF Amp (2001)	305 K	35 dB	3K	$7.2 * 10^{-7} \mu W$	$18 \mu W$
Splitter	305 K	-3dB	—	$3.6 * 10^{-7} \mu W$	$9 \mu W$
Pre-Detector Diode attenuators	305 K	-3/-6dB[e]	—	N/A	$3 \mu W$
Detector Diode	305 K	$1 mV/\mu W$	—	N/A	3 mV
Filter Bank	305 K	-5dB	—	$2.4 * 10^{-7} \mu W$	N/A
Multiplier	305 K	$1.2 mV/\mu W$	—	0.29 nV	N/A
Pre-Amplifier	305 K	100	—	29 nV	300 mV
Programmable Gain Amplifier	305 K	100/40[f]	—	$2.9 \mu V$	12 V
V/F Digitizer	305 K	1638.4 cts/V	—	$4.8 * 10^{-3}$ cts	$2 * 10^4$ cts

[a] The physical temperatures are the typical values during normal operations (excluding warm-ups in 2001, see Chapter 5); the temperatures of cryogenic components differ because the removal of the cryogenic IF amplifiers before the second observing season (see section 3.1.3) greatly reduced the thermal loading on the dewar.

[b] The gains and losses included here are rough, band-averaged estimates based on a combination of actual measurements, typical data, and the specifications of the parts in question. Differences between the components in the two lines are not indicated in this table.

[c] The effective noise temperature for each component is the contribution of that component to the total receiver noise temperature, referenced at the horn. Each component introduces some noise to the system. For amplifiers this is a measurable parameter, while for passive elements it can be calculated from its attenuation using Kirchoffs Laws: $T^{noise} = T^{physical}(1 - L)$, where $T^{physical}$ is the physical temperature of the component and L is the (linear) loss through the element. In order to properly reference this noise component to the signal at the horn, T^{noise} must be divided by the gain and loss of all of the elements above it in the signal chain ($\tilde{G}$): $T_{eff}^{noise} = T^{noise}/\tilde{G}$. After the first IF Amplifier, $\tilde{G}$ becomes so large that subsequent components contribute negligibly small amounts to the total system temperature.

[d] Signal levels assume the polarized signal from the sky is 5 μK, the temperature of the sky is 50 K, and the bandwidth of the radiometer is 16 GHz before the filter bank.

[e] The attenuation before the two total detector diodes are different (see Section 3.1.3) because the gains of the two lines are slightly different. The attenuators insure that the total power channel output voltages are approximately equal.

[f] Nominal settings during both observing seasons were 100 for the three polarization channels and 40 for the total power channels. The actual gains of this stage are closer to 110 and 44 (see Section 3.1.4).

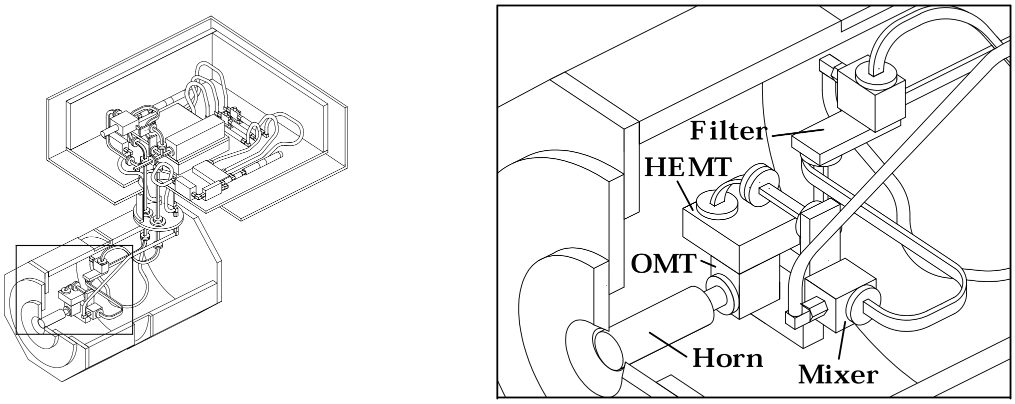

Figure 3.3: The RF portion of the W-band polarimeter (PIQUE2001)

the frequency conversion). The operational radiometer also requires back-end electronics (described in Section 3.1.4) to process the output voltages from the polarimeter into digital data stored by the computer.

3.1.1 RF Section

The RF section of the polarimeter (illustrated in Figure 3.3) conveys W-Band radiation from the feed horn (described in Chapter 4) to the mixers. Waveguide is used to channel this high-frequency radiation to different components. This portion of the radiometer combines low signal power levels with large passive losses and stringent phase-matching requirements, so it includes the minimal number of elements in a rather compact configuration.

The single feed horn couples all components of the input electromagnetic field into the radiometer. An Ortho-Mode Transducer (OMT) attached directly to the horn, decomposes this radiation into two polarized components. With the OMT oriented as shown in Figure 3.3, the radiation entering the arms of the polarimeter is polarized either horizontally or vertically. This insures the detection axes of the polarimeter are oriented $\pm 45^{\circ}$ with respect to vertical (see Chapter 2).

The radiation is amplified immediately upon entering the lines by HEMT (High Electron

Mobility Transistor) amplifiers bolted directly to the OMT. This first stage of amplification is critical because it dominates the receiver noise temperature (see Section 2.2.1), which is due to the Johnson noise of the active components and the finite emissivity of the passive components (see Table 3.1— the latter is particularly important at high frequencies like W-band, where the emissivity of metals is rather high). The HEMTs not only have low noise temperatures themselves ($\sim$ 75 K, see Wollack and Pospiszalski, 1998), but also have such a high gain that the noise contributions from the rest of the radiometer is nearly negligible (see Table 3.1). The total receiver noise temperature is $\sim$ 80 K, so the polarimeter has sufficient sensitivity to effectively search for the cosmological polarization (see Section 3.3).

The HEMTs must be kept as cold as possible in order to achieve the desired low noise temperatures, so the horn, OMT and HEMTs are thermally sunk to Stage II of a closed-cycle Helium refrigerator (CTI Cryogenics model 1020/1020R). These components were cooled to temperatures between 20 K and 40 K (depending on the component) during the first observing season, while during the second observing season temperatures between 10 K and 20 K were achieved (see Table 3.1). Components downstream of the HEMTs do not make significant contributions to the total system temperature, so the thermal emission from significant lengths of lossy waveguide is not a major concern. However, phase-matching requirements (see below) still favor a compact design. Hence a single E-Bend and a straight section of stainless steel waveguide (as a thermal break) connect the horn-OMT-HEMT complex to each of the filters. These filters are connected to stage I of the refrigerator, and operated at temperatures around 80 K during the first observing season and 60 K during the second observing season (see Section 3.4). They define the low-frequency edge of the polarimeter bandpass at 84 GHz, thus preventing the powerful 82 GHz LO signal from propagating through the RF lines.

Since the entire RF lines are housed inside the dewar, there is no way to actively adjust the relative phase shift between them (see Section 3.2). Fortunately, the compact design of the RF section allows the relative phase shift to be accurately calculated based on measurements of the phase shifts through individual components (obtained at room temperature with a vector network analyzer (Hewlett-Packard Model 8510C)). The combination of components in the actual RF lines (including a $0.050''$ long shim in the line carrying the horizontal polarization) were chosen so the total relative phase shifts across the band are small ($< 10°$).

3.1.2 LO section

The mixers are attached directly to the output of the filters and sunk to Stage I of the refrigerator (see Figure 3.3). These devices multiply the input RF radiation at 84-100 GHz by a narrow-band 82 GHz signal from a Local Oscillator (LO) to yield IF radiation at 2-18 GHz (see Appendix A for how this multiplication can be accomplished using diodes). The LO signal has to be strong enough to overdrive the mixer diodes, or else the IF output can be sensitive to variations in the LO signal strength (Elliott, 1993). The manufacturer specifications indicate that an LO input signal of 10 dBm (10 mW) is sufficient to overdrive the mixers. The local oscillator (LO) therefore emits 20 dBm of 82 GHz radiation. In order to generate a signal this strong, the LO requires at least 10W of power. This high power consumption would have put a significant thermal load on the refrigerator, and the LO is designed to work at room temperature anyway, so the LO is located outside the dewar in the same metal box that houses the IF portion of the radiometer (This box, for historical reasons, is called the "E-Box"). The LO is contained inside a smaller box (the "LO Box") within the E-Box, which is thermally isolated with a G10 shim to allow the LO to self-heat to its correct operating temperature (315 K, maintained with an internal heater). The LO Box was also intended to stop RF signals from leaking between the LO and the IF sections of the radiometer, but this function was degraded when it was found necessary to cut large holes in the LO Box to accommodate various pieces of waveguide and coaxial cable.

A rather long and convoluted collection of waveguide is required to get the signal from the LO to the mixers (Figure 3.4). The output from the local oscillator is split between the two LO lines by a waveguide hybrid T attached directly to the LO. Splitting the lines at this early stage, rather than inside the dewar allows the phase switches and the phase tuner to be conveniently located outside the dewar (see below). The main problem with this choice of design is that it requires a long pathlength of W-band waveguide at room

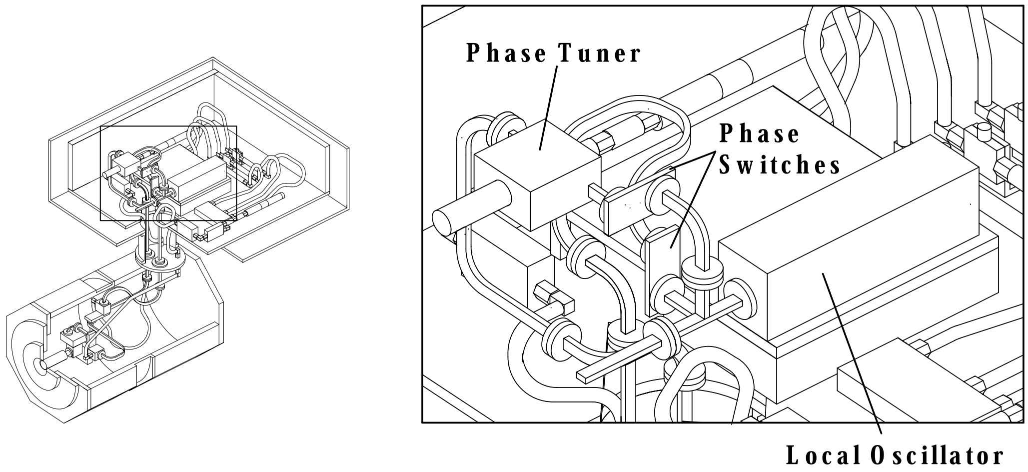

Figure 3.4: The LO section of the W-band polarimeter

temperature. W-band waveguide is rather lossy (0.15 dB/inch for copper, 0.50 dB/inch for stainless steel), and this long pathlength (coupled with the losses incurred through the phase switches, etc.) means that the power actually delivered to the mixers is as low as 9dBm, somewhat below the manufacturer specifications. Laboratory tests demonstrated that such low input powers do not significantly change the loss or noise temperature of the mixers, and there is no evidence for increased gain instabilities (see Chapter 6).

The complicated shape of the LO lines reflects both geometrical and mechanical constraints. In particular, the two lines are approximately the same length so that the mixers receive approximately the same input LO power. With some difficulty, these lines could have been phase matched by making the line lengths exactly identical, but instead a phase tuner is used in one of the LO lines. This device inserts a variable amount of a dielectric into the line, thus changing the phase of the monochromatic radiation passing through it. (For the broad-band radiation in the RF and IF lines, this method would not have worked well because the phase-shift introduced would be a complex function of frequency). The phase tuner is located outside the dewar so the relative phase shift between the LO lines can be tuned while the polarimeter is operational (see Section 3.2).

The (4 kHz) phase switches are also located in the LO lines. Since the mixer takes the product of the input RF and LO fields to produce the output IF radiation, multiplying either the RF or LO field by ± 1 would properly modulate the output signal. It is easier to change the sign of the strong, monochromatic LO radiation than to modulate the phase of the weak, broad-band RF radiation. The phase switches are also active components that require a current source in order to function (see Section 3.1.4). As part of the LO lines, they can be located outside the dewar, which makes it much easier to connect up the appropriate driver circuitry. While only one phase switch is needed to modulate the polarized signal, both LO lines have phase switches to make the lines as symmetric as possible and to allow either line to be phase switched. Normally, only the phase switch in the line carrying the horizontal component of the input field is actually modulated (at 4 kHz); the other phase switch is simply held open (with a constant current) without being switched.

3.1.3 IF section

The IF section of the radiometer carries the phase switched, down-converted (2-18 GHz) radiation from the mixers to the multipliers. At these comparatively low frequencies, passive losses are much smaller, semi-flexible SMA coaxial cable [EZ-141AL-TP (EZ-Form) Aluminum jacketed cable outside the dewar, AS 50141 (Precision Tube) Stainless Steel cable inside the dewar] can be used to transmit the signal, and components are easier to acquire and harder to break. In the PIQUE2000 radiometer, the IF radiation from the mixers first entered amplifiers located inside the dewar (the so-called cryogenic IF amplifiers). These amplifiers were a source of significant problems during the first observing season (see Section 3.4), and due to the higher than expected gain of the HEMTs and the limited dynamic range of the multipliers, they were unnecessary (see Section 3.4). These amplifiers were therefore removed for the second observing season. The IF radiation leaves the dewar via SMA cables and enters the E-Box where the remainder of the radiometer (operating at room temperature) is housed (Figure 3.5). The removal of the cryogenic IF amplifiers required interchanging the two stainless steel coaxial cables (see Figure 3.6), so the SMA line that carried the horizontally polarized component during the first observing season carried the

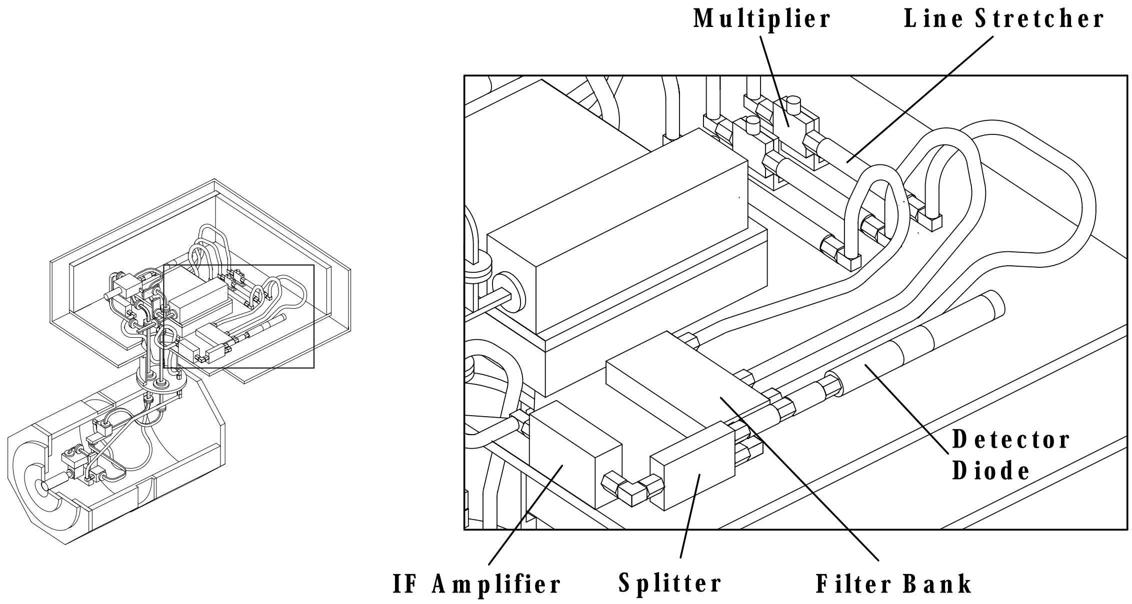

Figure 3.5: The IF section of the (2001) W-band Polarimeter

vertically polarized component during the second observing season, and vice versa.[1]

Upon reaching the room temperature portion of the radiometer, the field in each line is amplified by (another) amplifier (sometimes called the warm IF amplifier), and sent into a (3 dB) splitter, which couples half the radiation out to a detector diode. This radiation is attenuated before reaching the detector diode in order to keep the output voltage on scale on the computer (Section 3.1.4). Different attenuators are used in front of the two detector diodes (-3dB in D0 and -6 dB in D1 during PIQUE2000, - 6 dB in D0 and -3 dB in D1 during PIQUE2001) because of differences in the gains of the two lines. These diodes monitor the total power in the two lines and are used extensively in characterizing the system. The signals from these diodes are denoted D0 and D1. Due to the SMA interchange between the two observing seasons, D0 measured the vertically polarized component of the incident radiation during PIQUE2000 and the horizontally polarized component during PIQUE 2001.

The splitters transmit the remainder of the radiation into a pair of filter banks, which divide it among three frequency sub-bands. This frequency separation both provides spectral resolution and simplifies the phase matching of the system (see Section 3.2). The radiation

[1]The SMA interchange also meant that a given RF line (including the HEMT) was connected to different room temperature IF lines and detector diodes during the two observing seasons.

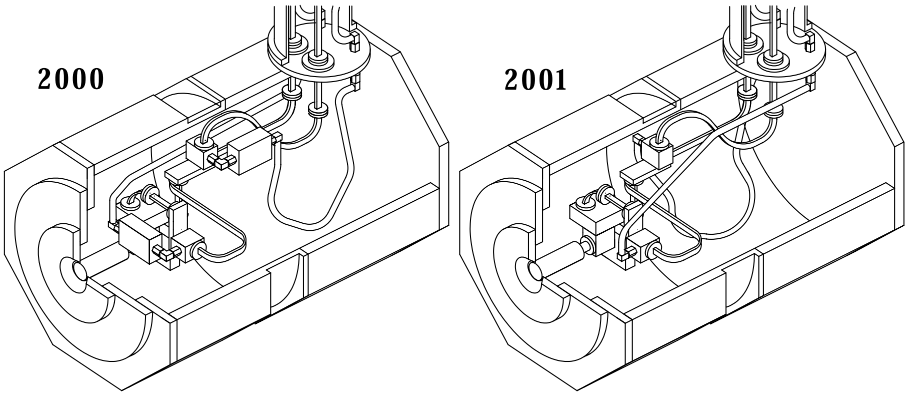

Figure 3.6: The different configurations of the PIQUE2000 and PIQUE2001 polarimeters before and after the removal of the cryogenic IF amplifiers. Note that the coaxial cables in the dewar are crossed in the second figure.

in each sub-band is sent to a separate multiplier. The multiplier channels are sensitive to the frequency ranges 84-89 GHz, 89-94.7 GHz and 94.7-102 GHz, and are denoted s0, s1 and s2 respectively. A line stretcher is attached to one of the input ports of each multiplier in order to facilitate the phase matching of the polarimeter (see Section 3.2).

3.1.4 Back-End Electronics

The back-end electronics (Figure 3.7) process and digitize the voltages from the (three) multipliers and the (two) detector diodes so that they can be read (and stored) by the computer. In addition, these electronics control the phase switch.

The voltages from the various channels are first sent through preamplifiers attached directly to the detector diodes and the multipliers. These are differential amplifiers with gains of 100 (circuit diagram in Appendix C), which strengthen the signals so that they can be sent on short ($4'$ long) cables to the two "thin boxes." Located on the telescope base near the E-Box, these thin boxes contain the remainder of the processing and digitizing circuitry.

The signals from the multipliers and the signals from the detector diodes are sent to different thin boxes, with different processing circuits (diagrams in Appendix C). In both boxes, the signals are first amplified by a computer-controlled factor (typically 44 for the

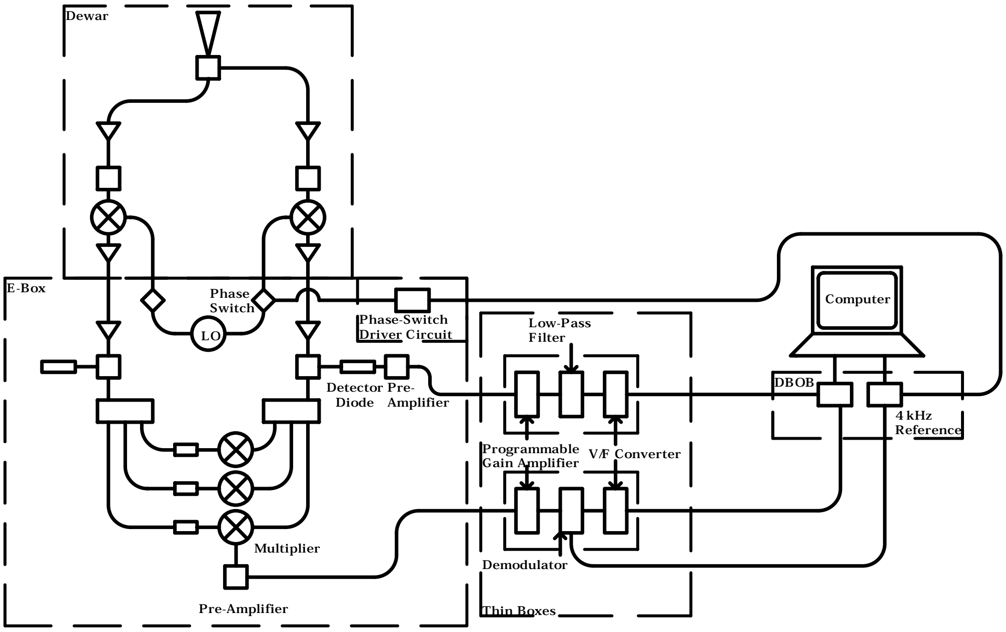

Figure 3.7: Schematic of the electronics which connect the detector with the computer

total power channels, and 110 for the multiplier channels). The amplified signals from the total power channels are then simply low-pass filtered (240 Hz, 8 pole) before being digitized. The signals from the multipliers, however, need more complex processing to extract the signal modulated at 4 kHz by the phase switch. A 4 kHz reference square wave from an oscillator located near the computer (housed in a box called the Digital-Break-out- Box, or DBOB) clocks the circuit that provides the appropriate (square wave) alternating current to one phase switch (see Appendix C, a similar circuit provides a constant current to the other phase switch). This same reference is used by a circuit that demodulates the signals from the multipliers. Thus signals modulated at 4 kHz are converted into a DC voltage and any DC offset is converted into a 4 kHz modulation, which is removed with a low-pass (10 Hz, 2 pole, Bessel) filter. The processed signals from both the detector diodes and the multipliers are therefore analog DC voltages. These voltages are digitized using voltage-to-frequency converters that produce an output 5V (digital) square wave with a frequency proportional to the input (analog) voltage (This device assumes a positive voltage on the input. Since the signal from the multipliers could have either sign, a 5V offset is added to these signals before they enter the V-to-F. Conversely, the signals from the total power channels vary about some large positive number, so a 5V offset is removed from these signals to provide increased resolution).

After being digitized, the five signals are sent on separate lines (twisted, shielded pairs) through long ($30'$) cables to the DBOB. The DBOB in turn passes these signals to the computer, where an RTI827 and an RTI815 card measure the frequency of each digital square wave by counting the number of pulses received within a sampling period (0.016 s). The resulting numbers, or "counts" (0-16384) are stored to file, along with various housekeeping data (see Appendix F), by a C++ program running under DOS. The timing of this data acquisition and storage is controlled by interrupts from the same oscillator that clocks the phase switch.

3.2 Phase Matching

Before it can function as a polarimeter, the radiometer must be tuned so that the relative phase shift between the two arms is as small as possible at all relevant frequencies (Section 2.2.1). This is accomplished by adjusting the phase tuner in the LO lines and the line stretchers in the IF lines[2]. Having two methods for changing the relative phase shift between the two lines is essential to achieve a good phase match. On the one hand, dialing the phase tuner changes the relative phase shift of the essentially monochromatic signals in the LO lines, which in turn introduces a frequency independent relative phase shift into the IF radiation. On the other hand, adjusting the line stretchers changes the relative path-length between the two arms, which alters the relative phase shift in a frequency dependent way (For a given change in relative line length δL, the change in the relative phase shift is $\delta\phi = 2\pi\delta L/\lambda$, where λ is the wavelength of the radiation). Therefore, by adjusting the phase tuner and the line stretchers, an offset and a slope in the relative phase shift between

[2]The RF lines cannot be tuned because they are housed entirely inside the dewar. They were instead assembled from components which yield a small relative phase shift, see Section 3.1.1

the two lines can be removed. Higher order terms could in principle be introduced by the various amplifiers, splitters, and filters, but these components were specifically ordered so that such terms are small (see Figure 3.10).

The procedure for phase-matching the radiometer is as follows. The narrow band test arrangement (shown in Figure 3.8 and described in detail in Section 3.3.1) injects a swept, narrow-band, polarized signal into the polarimeter. Each multiplier channel is phase matched by adjusting the phase tuner and line stretchers until the output of the multiplier is nearly zero at all frequencies. The relative phase shift between the arms is then roughly 90° at all frequencies, and there are no gross frequency dependent phase shifts (i.e. the line stretchers are set correctly). Finally, the residual overall phase shift is removed by adjusting the phase tuner (which has been appropriately calibrated with a vector network analyzer). This procedure is used instead of adjusting the system until the output is maximized at all frequencies because it is much easier to discern the null than the maximum. After adjustments are complete, the line stretchers are locked into position and the setting of the phase tuner is recorded.

3.3 Characterization: Integrated Polarimeter Tests

Once phased up, the polarimeter is characterized using a series of three tests involving different systems for injecting known signals into the polarimeter (Figure 3.8). These three tests provide information on the bandpass (narrow-band test), responsivity (two-load test) and sensitivity (circular load test) of the radiometer. Additional measurements are performed after the horn is installed prior to the observing seasons. This series of tests was performed before both observing seasons, and the narrow-band test and the two-load test were repeated after both observing seasons. Unfortunately, most of the results from the tests before the first observing season are invalid because of a problem with the multipliers (see Section 3.4.2).

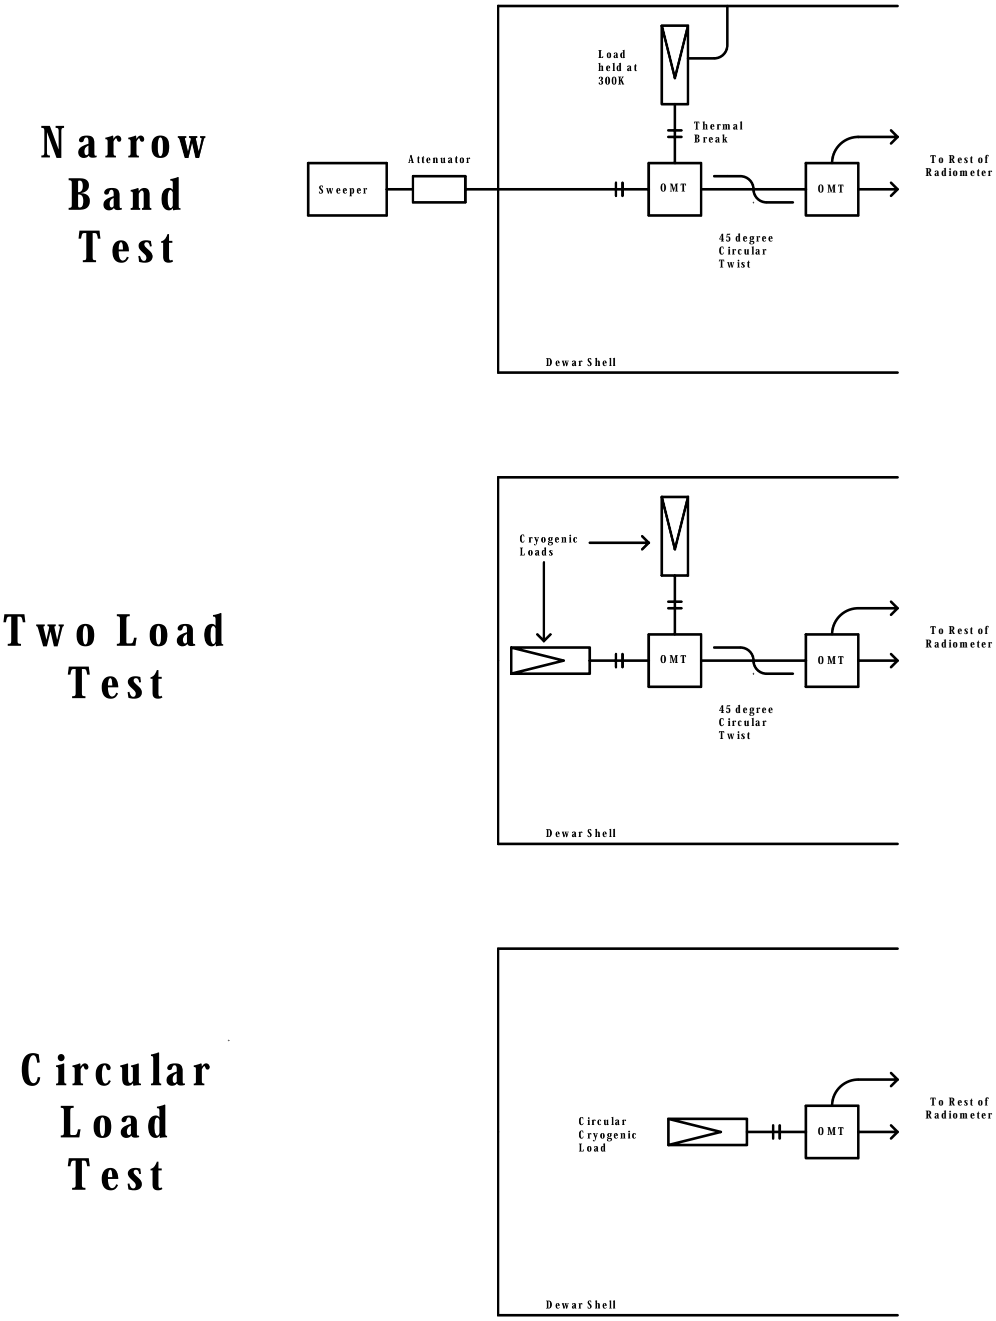

Figure 3.8: Block diagrams showing the set ups of the three integrated polarimeter tests: (top) Narrow-Band test, (middle) Two-Load test, (bottom) Circular Load test

3.3.1 Narrow-Band Test

During this test (and during phase matching, see Section 3.2), a polarized narrow-band signal is injected into the polarimeter using an external sweeper (Figure 3.8). The signal is coupled into the radiometer OMT (the horn is removed for this test) through a combination of another (injector) OMT and a piece of brass circular waveguide which has its two flanges rotated 45° with respect to each other (both these components are housed inside the dewar). This combination of twist and OMT is a handy way to inject polarized signals into the dewar, since a signal entering one of the rectangular ports on the injector OMT is polarized along one of the detection axes of the polarimeter (see Section 2.2.1). For this test, one port of the injector OMT is connected through waveguide to a sweeper (Hewlett-Packard Model 83630L) external to the dewar, while the other port is terminated by a load. In the W-band system this load is held at room temperature inside the dewar. This load insures that 300K thermal radiation is injected along both detection axes of the polarimeter. If this load were not present, then the thermal radiation from the room temperature sweeper and attenuators[3] would introduce a polarized signal an order of magnitude larger than the dynamic range of the polarimeter. With the load in place, the narrow-band signal from the sweeper is the only polarized signal entering the radiometer.

Measurements of the frequency-dependent qualities of the polarimeter are obtained with this arrangement by continuously recording the outputs from the polarimeter while the input signal sweeps from 80 GHz to 105 GHz at a low speed (total sweep time of about 30 s, which prevents the filters in the back-end electronics from smearing the response). The frequency of the input radiation is not included in the data stored by the computer, so it must be reconstructed based on the crossover points between the three multiplier sub-bands at 89 GHz (between s0 and s1) and 94.7 GHz (between s1 and s2). Two data sets are acquired with different (calibrated) LO phase tuner settings. In one case (the "in-phase" condition), the polarimeter is properly phase matched, in the other (the "out-of-phase"

[3]If the attenuators were located inside the dewar, the thermal emission injected into the polarimeter would be much smaller and the room temperature load would be unnecessary. The Q-Band system is tested using cold attenuators and without a room temperature load.

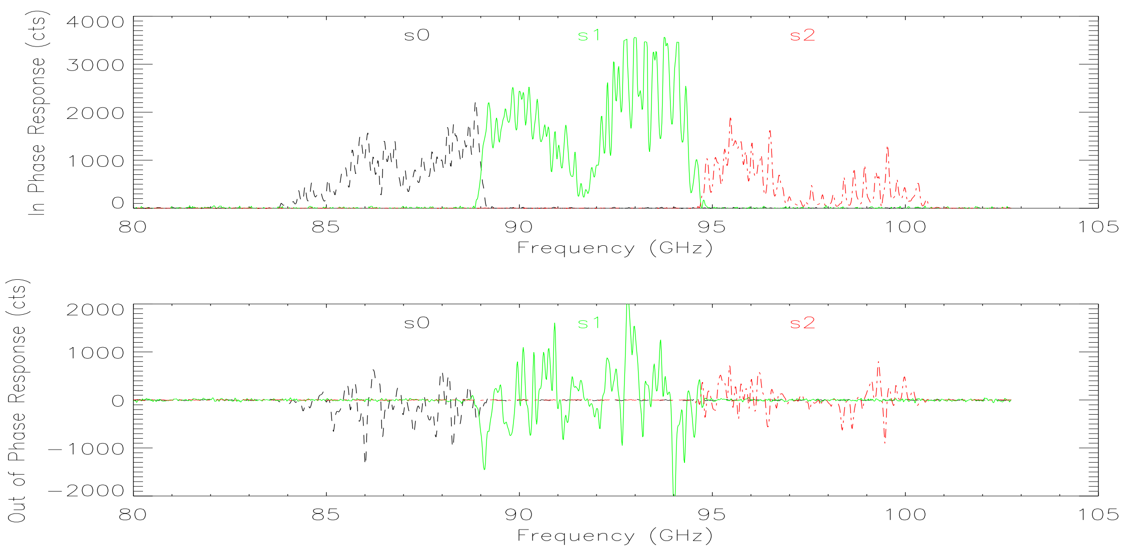

Figure 3.9: The raw data for determining bandpasses and relative phase shifts (PIQUE2001, pre-season), showing the response of the three frequency sub-bands when the system was phased up and when it had 90° relative phase shift between the lines. The different line styles distinguish the three bands (s0,s1,and s2 from left to right).

condition) there is an average of 90° relative phase shift between the lines (Figure 3.9). If the ideal (i.e. no relative phase shift) responsivity of the system is A and the relative phase shift is ϕ, then these two measurements give $a = A\cos\phi$ and $\bar{a} = A\sin\phi$, respectively. The ideal responsivity and relative phase shifts of the radiometer can therefore be easily extracted from these data (Figure 3.10). The effective bandwidth ($\Delta\nu$) and the effective central frequency (ν_{eff}) of each band is then calculated from the (ideal) response $A(\nu)$ using the formulae (Kraus, 1986):

$$\Delta\nu = \frac{(\int A d\nu)^2}{\int A^2 d\nu},$$

and

$$\nu_{eff} = \frac{\int A\nu d\nu}{\int A d\nu}.$$

Similar expressions can be derived from the in-phase response a. The average value of the cosine of the phase shift $< \cos\phi >$ is also calculated for each band. These parameters are provided in Table 3.2

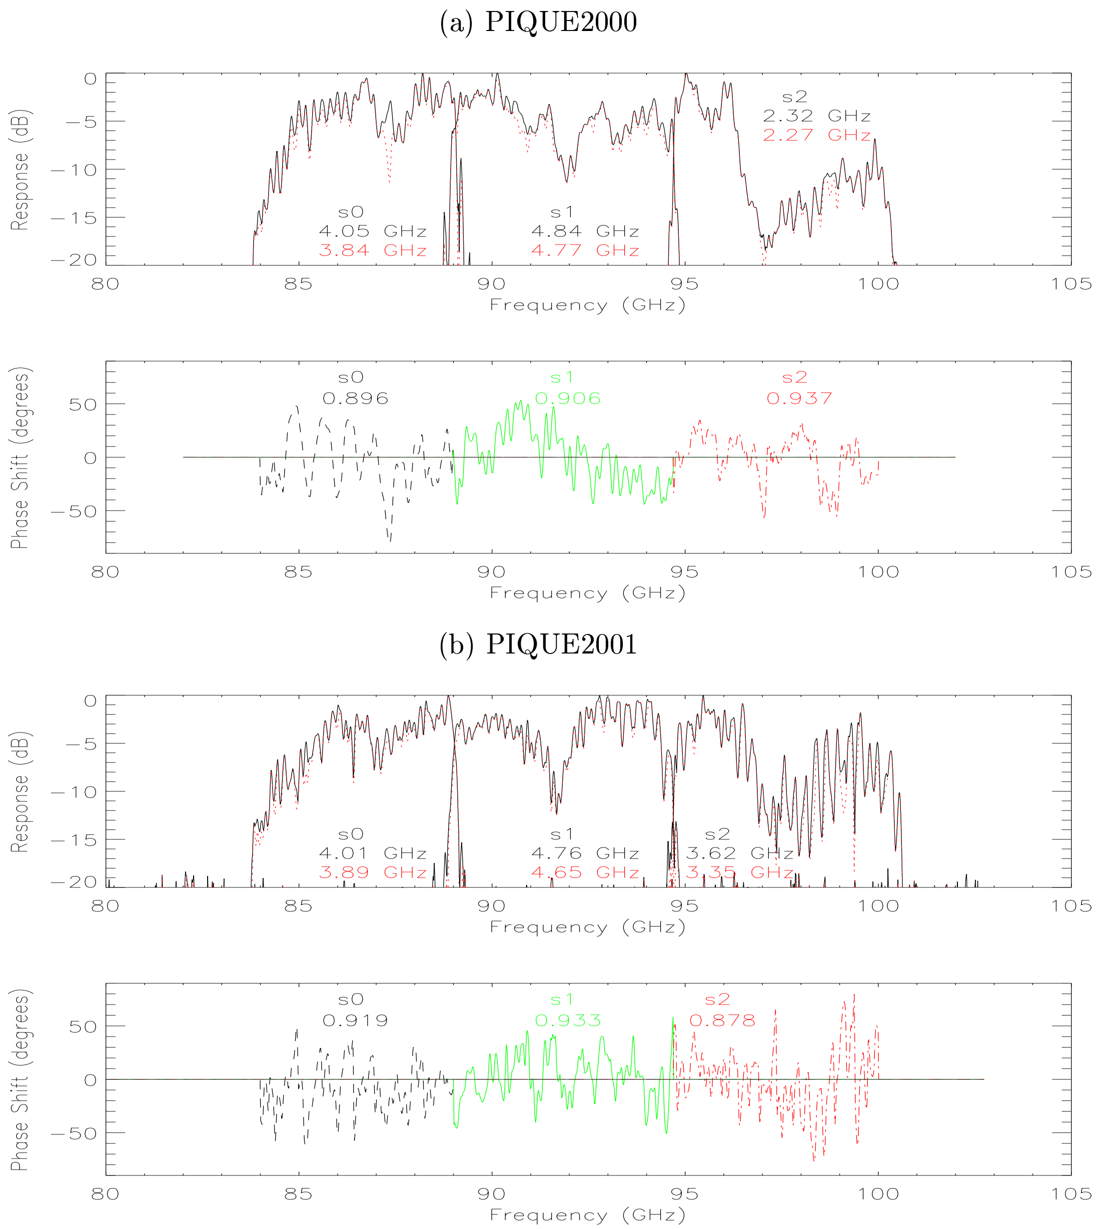

Figure 3.10: The relative responsivities and phase shifts of the W-Band polarimeter as a function of frequency derived from the (a) PIQUE2000 post-season test and the (b) PIQUE2001 pre-season test. In the responsivity plots, the solid lines are the ideal response and the dotted lines are the response from the in-phase measurement. The black and red (gray) numbers are the bandwidths derived from the ideal and in-phase data, respectively. In the phase plots, the different line types correspond to different bands, and the numbers are the value of $\cos\phi$ averaged over the band.

Table 3.2: Frequency-dependent characteristics of the radiometer

Character	s0		s1		s2	
	2000[a]	2001[b]	2000[a]	2001[b]	2000[a]	2001[b]
ν_{eff}	87.0 GHz	87.3 GHz	91.4 GHz	91.9 GHz	96.1 GHz	96.4 GHz
$\Delta\nu_A$	4.05 GHz	4.01 GHz	4.84 GHz	4.76 GHz	2.32 GHz	3.62 GHz
$\Delta\nu_a$	3.84 GHz	3.89 GHz	4.77 GHz	4.65 GHz	2.27 GHz	3.35 GHz
$\Delta\nu_a/\Delta\nu_A$	0.948	0.970	0.985	0.976	0.978	0.918
$< \cos\phi >$	0.896	0.919	0.906	0.933	0.937	0.878

[a] Post-season estimates

[b] Pre-season estimates. The post-season estimates are the same to within 10%, except for s2, for which $\Delta\nu_a$ is 15% smaller.

On average, the relative phase shifts for both observing seasons and all bands are roughly zero, as desired (see Figure 3.10). However, there is significant fine-scale structure in the phase shifts, with fluctuations as high as 45°. These nonzero phase-shifts have two effects on the performance of the polarimeter: a reduction in the responsivity of the polarimeter by a factor off $< \cos\phi >$, and a reduction in the effective signal bandwidth of the polarimeter (the noise bandwidth is not reduced). As shown in Table 3.2, every channel loses $\sim 10\%$ of its bandpass and responsivity to nonzero phase shifts. The polarimeter is therefore very well phase matched and the residual phase shifts have a very minor effect on the polarimeter's performance.

More significant factors for the performance of the polarimeter are the frequency dependent gain variations. These variations significantly reduce the bandwidths of two of the sub-bands (s0 and s2) from the ideal 5 GHz. The bandwidth of s2, in particular, is only about half its ideal value. The combined bandwidth of the three bands is consequently only about 12 GHz instead of 16 GHz, which unfortunately reduces the sensitivity of the polarimeter by 15% (see Section 3.3.3 below).

3.3.2 Two-Load Test

This test, which measures the responsivity (gain) of the polarimeter, uses the same 45° twist and OMT as the narrow band test. However, in this case, both ports of the second

OMT are terminated by cryogenic loads (see Figure 3.8), whose temperature can be controlled and measured using calibrated thermometers (Lakeshore Model DT-470-CU-13) and a servo circuit. The loads are held between 30K and 50K, comparable to the atmospheric noise temperature during the observations (see Chapter 7). Each of the two loads injects radiation polarized along one of the detection axes of the polarimeter, so a difference in the temperatures of the loads produces a known polarized signal at the input to the polarimeter.

The response of the polarimeter to variations in the load temperatures is measured with two different data acquisition procedures. In the so-called "manual" procedure, the loads are set to regulate at a variety of temperatures (each load being set to one of five temperatures, about 3 K apart, for a total of 25 possible combinations) and allowed 10-15 minutes to reach these temperatures before the data are recorded. With this procedure, the load certainly has time to stabilize at the regulated temperature, but collecting data in this way is extremely slow and tedious. The accuracy of this procedure may also be compromised by long time scale drifts (see Section 3.3.3). On the other hand, the "automatic" procedure begins with both loads regulating at low temperatures. One load is then set to regulate at a higher temperature (by 5K- 10K) and allowed to warm up for 10-15 minutes before the heat is turned off so the load can cool back down (The other load remains at a constant temperature). Both the load temperatures and the polarimeter outputs are recorded continuously during this process. This procedure is much faster and less sensitive to long time scale drifts, but may have problems with inaccurate temperature readings due to thermal gradients across the load. Both these procedures give comparable estimates for the responsivity of the polarimeter, so the tedious manual procedure was not used extensively after the first observing season, and only the results of the faster automatic procedure are reported here.

Estimates of the responsivities of both the polarimetry channels and the total power channels can be derived from this test. Say the two loads are at temperatures T_1 and T_2. These loads inject radiation along the detection axes of the radiometer. The radiation from each load is split equally between the two lines, so each total power channel measures the average temperature of the loads, $T_{ave} = (T_1+T_2)/2$. The polarimetry channels, on the other

Table 3.3: Summary of losses in the two-load injector system and the optics chain

Injector system loss		Optics chain loss	
Elements	Transmission	Element	Transmission
Stainless Steel W/G	0.87	Dish Surface Roughness	> 0.96
OMT	0.99	Dewar Window	0.98
Brass Circular Twist	0.97	Horn	> 0.99
Gold Plated Circular Straight Section	> 0.99		
Total	0.84	Total	0.98

hand, indicate the polarization of the input radiation, which is given by $Q_T = (T_1 - T_2)/2$ (see Section 2.1).

The responsivity of each channel is determined by obtaining the linear fit slope for the output signal as a function of T_{ave} or Q_T (The readings from the load thermometers are noisier than the signal outputs, so the fit is actually performed with the temperature difference as the dependent variable). In addition, the receiver noise temperatures for both lines are calculated from the x-intercept of the fit to the total power channel data.

The measured slopes and intercepts are not exactly the desired responsivities and noise temperatures of the operational polarimeter. Due to the losses in the injector, the radiation entering the polarimeter OMT is attenuated by a factor of $L_{injector}$. In the operational telescope, polarized signals from the sky reach the polarimeter through the optics chain, which has loss factor (transmission) L_{optics}. These different losses must be taken into account to determine the response of the radiometer to signals from the sky.

The relevant lossy components and their loss factors are listed in Table 3.3. For $L_{injector}$, the transmission of the OMT is the measured value at room temperature, while the loss factors of the other waveguide elements are calculated assuming the frequency of the radiation is 90 GHz and the published conductivities of stainless steel, brass and gold at 40 K ($50\mu\Omega \cdot cm, 4.6\mu\Omega \cdot cm$ and $0.12\mu\Omega \cdot cm$, respectively; White, 1959). Similar calculations assuming room temperature conductivities predict the transmission through the stainless steel waveguide, OMT, and brass twist should be 0.80 (0.92 dB). This loss factor was measured with a vector network analyzer as roughly 0.70 (1.5 dB). The excess loss is most likely an artifact in the measurement, but may also be due to imperfect waveguide joints in the injector.

For $L_{optical}$, the loss in the horn is less than that of the gold plated section of waveguide (which is approximately the same length), and so is completely negligible. The loss factor for the window is the calculated transmission through the 0.002″ sheet of polypropylene that forms vacuum window (The 0.020″ Goretex^{TM} backing has negligible (<0.05 dB) loss). The loss of the mirror is due to diffuse scattering from the finite surface roughness of the mirror surface. This loss can be calculated from the Ruze expression $L_{diffusive} = exp(-(2\pi d/\lambda)^2)$ (Ruze, 1952), where d is the rms surface roughness and λ is the wavelength of the radiation. Measurements at 40 GHz indicate that d is less than 0.05 mm (Netterfield, 1995), so the diffusive loss factor from the mirror is greater than 0.96.

Based on this information, the loss ratio is: $L_{optical}/L_{injector} = 1.17^{+0.10}_{-0.04}$. The nominal value assumes the calculated losses and no diffusive loss in the dish. The upper error bar assumes the excess loss measured for the injector is real and independent of temperature, while the lower error bar assumes the diffusive loss in the dish is actually 0.96. These error bars should be understood as extreme limits on the possible range for this loss ratio.

For the total power channels, the slope R and the x-intercept T_o of the linear fit are converted into estimates of the operational telescope responsivity $\mathcal{R}$ and receiver noise temperature T_{rec} using the expressions $\mathcal{R} = R{*}L_{optical}/L_{injector}$ and $T_{sys} = T_o L_{injector}/L_{optical} - (1 - L_{injector})T_i$. The factor of $(1 - L_{injector})T_i$ in the expression for T_{sys} accounts for noise introduced by the injector itself, which has a physical temperature $T_i \simeq 30K$[4]. The resulting estimates of $\mathcal{R}$ and T_{sys} are given in Table 3.4. Note that for PIQUE2000 the receiver noise temperature for both lines is about 70K, while during PIQUE2001, one line has a noise temperature of 90K. This is attributed to degradation in the performance of one HEMT with time and exposure to the elements in the lab.

For the polarimetry channels, only the responsivity is of interest (but see Section 3.3.3). The responsivity of these channels in the operational telescope $\mathcal{R}_p$ can be calulated from

[4] A correction accounting for noise injected by the optical system is not included, because loss through the window is due to reflection instead of absorption.

Table 3.4: Total-power responsivities and receiver noise temperatures derived from two-load tests

Year	D0		D1	
	$\mathcal{R}$ (cts/mK)	T_{rec} (K)	$\mathcal{R}$ (cts/mK)	T_{rec} (K)
2000 (Post-Season)	0.14	73	0.11	68
2001 (Pre-Season)	0.14	66	0.12	93
2001 (Post-Season)	0.15	61	0.12	90

The dominant uncertainty in these numbers is the 7% uncertainty in the loss correction factor.

Table 3.5: Responsivities of multiplier channels derived from two-load tests

Year	$\mathcal{R}_p$ (cts/mK)		
	s0	s1	s2
2000 (Post-Season)	2.0	3.6	1.2
2001 (Pre-Season)	2.1	6.0	1.2
2001 (Post-Season)	1.8	6.4	1.1

The dominant uncertainty in these numbers is the 7% uncertainty in the loss correction factor.

the linear fit slope R_p using the formula $\mathcal{R}_p = R_p k(\nu)(L_{optical}/L_{injector})$. The loss ratio factor is the same as for the total power channels, but there is an additional factor of $k(\nu)$ to account for the temperature difference between the (40 K) loads and the CMB (this correction is also applied to the chopper plate data described in Chapter 6). For the loads, the quantity $h\nu/k_BT < 0.1$ is small, so the Rayleigh-Jeans Limit applies and the intensity to temperature conversion factor $K(\nu,T) = \frac{dI}{dT} = K_{RJ}$ is a constant. However, for the CMB (T=2.7K) $h\nu/k_BT \simeq 2$, so the conversion factor K is smaller than K_{RJ} by a factor $k(\nu) = K(\nu,T_{CMB})/K_{RJ} = \frac{dI}{dT}_{CMB}/\frac{dI}{dT}_{RJ} = (e^x - 1)^{-2}e^x x^2$, where $x = h\nu/k_BT_{CMB}$. (Note that the this quantity is independent of the load temperatures). This reduces the responsivity of the polarimeter to polarized signals from the CMB by a comparable factor.[5] Since this factor depends on the frequency of the radiation, it is different for each polarimetry channel, and must be calculated using the bandpasses derived from the narrow-band test. For s0, s1 and s2, $k(\nu)$ is 0.83, 0.81 and 0.79 respectively (Differences in the bandpasses between the two observing seasons do not significantly affect these factors). The resulting

[5]This reduction only occurs when the polarization is measured in units of temperature, if it is measured in units of power, no such correction is necessary.

estimates for $\mathcal{R}_p$ are given in Table 3.5.

This test was performed before and after both observing seasons. However, just before the start of the first observing season, a feature of the multipliers was discovered which required adding attenuators to the lines (see Section 3.4.2). The previously measured values for the responsivity could not be used to estimate the responsivity of the polarimeter with the attenuators in place, and it was too late to repeat the test. Therefore, there are no pre-season estimates of the responsivities for the PIQUE2000 polarimeter.

The uncertainties in the estimates derived from the two-load test are dominated by the $\sim$ 7% uncertainty in the loss correction factor. Statistical errors in the linear fits and systematic errors due to miscalibration of the thermometers are each less than 2%. This large systematic uncertainty is not a significant problem because the responsivity estimates for the polarimetry channels derived from the two-load tests are not directly used to calibrate the operational telescope. These estimates are only used to verify the results from the chopper plate test, which provides a much more direct measure of the telescope responsivity (see Chapter 6).

3.3.3 Circular Load Test

This test provides measurements of the sensitivity, system temperature and the cross-coupling of the polarimeter. In this case the detector is coupled to a circular load (at 30 K) through a section of circular waveguide (Figure 3.8). The circular load injects a stable, essentially unpolarized signal into the radiometer. This load is either held at one temperature or has its temperature changed by the same servo circuitry used in the two-load test. This test is performed before both observing seasons, but the data from the first observing season was affected by an unexpected feature of the multipliers (Section 3.4.2). Hence, the quantitative results given here derive exclusively from the PIQUE2001 polarimeter.

Varying Load Temperature: Noise Temperature and Cross-Coupling

Both manual and automatic data acquisition procedures (see Section 3.3.2) are used to measure the response of the polarimeter to (5-10 K) variations in the temperature of the

Table 3.6: Circular load test results: total power channels

Year	D0		D1	
	$\mathcal{R}$ (cts/mK)	T_{rec} (K)	$\mathcal{R}$ (cts/mK)	T_{rec} (K)
2001 (Pre-Season)	0.15	60	0.12	95

The dominant uncertainty in these numbers is the $\sim 3\%$ uncertainty in the loss correction factor.

load. These data are used to estimate the responsivities of the total power channels and the noise temperatures of the two lines. The responsivity and noise temperatures are obtained from the slope and x-intercept of a linear fit to each total power channel output voltage as a function of load temperature, as for the two load test (see above). These results must also be corrected for the injector loss, which in this case is only 0.93 (due to the 1 1/2$''$ stainless steel circular waveguide), so the correction factor is $L_{optical}/L_{injector} = 1.05^{+0.026}_{-0.04}$.

The estimates for the responsivity and receiver noise temperatures derived from the circular load test are given in Table 3.6. The uncertainty is again dominated by the uncertainty in the loss correction term, which is considerably smaller in this test than it is for the two-load test because of the comparative simplicity of the injection system. The results of this test are reasonably consistent with the results from the two-load test for the 2001 radiometer.

Variations in the load temperature also cause changes in the outputs of the polarimetry channels (see Table 3.7). During the observations, variations in the temperature of the atmosphere produced a similar response (see Chapter 7). A simple test (performed prior to the first observing season) demonstrated that this responsivity of the polarimetry channels to unpolarized radiation is a real effect, and is not simply the result of stray polarized signals. Any polarized signal generated by the load should change sign when the load is rotated by 90°, but the measured responsivities do not simply change sign when the load is rotated. Therefore, the polarimetry channels must also respond to changes in total power. Estimating this unpolarized responsivity is difficult because the response of the polarimetry

[6] The upper limit assumes room temperature loss in the waveguide instead of the loss at 40K, the lower limit assumes 0.96 diffusive loss in the dish, see Section 3.3.2

Table 3.7: Circular load test results: polarimetry channels (2001)

Character	s0	s1	s2
Unpolarized Responsivity Estimate	5.9 cts/K	-14.4 cts/K	-6.3 cts/K
Rejection Ratio[a]	320	430	180

[a]The rejection ratio is the ratio of the measured responsivities here to the responsivities given in Table 3.5.

channels does change when the load is rotated, which implies the load does emit some polarized radiation. This test does not provide sufficient information to separate responses to these small polarized signals from the responses due to true unpolarized responsivity. However, rotating the load has a noticeable effect on the responsivities, so the two contributions to the measured response must be about the same size. The measured responsivities of the polarimetry channels (Table 3.7) therefore provide extremely rough estimates of the actual unpolarized responsivities.

A comparison of the responsivities in Table 3.5 with those in Table 3.7 shows that the multiplier channels are over 200 times less responsive to unpolarized signals than they are to polarized signals. This small but finite responsivity is due to cross-coupling between the two lines of the radiometer. If some small fraction of the signal in one arm of the polarimeter leaks into the other arm above the phase switch, that stray signal correlates with its parent and produces a nonzero output from the multiplier. (Unlike the offsets discussed in Section 2.2.1, phase switching does not remove this false signal because it arises from a true correlation between the lines.) Correlation polarimeters are particularly sensitive to this cross-coupling because the multiplier takes the product of the two fields. Therefore, if the leaked signal has a factor of X less power than its parent in the other line, the response of the polarimeter to this leaked signal will only be a factor of $\sqrt{X}$ smaller than its response to a polarized signal. The fact that the unpolarized responsivity of the polarimeter is as small as it is shows that the polarimeter has very good power isolation between the lines (-50 dB or better).

Cross coupling is most likely to occur where the two arms of the detector are most intimately connected: in the OMT and the HEMTs. The HEMTs have a large reflection coefficient R_{HEMT} (which is on average -5 dB, but gets as bad as -2dB). A large fraction of the radiation coupled through the OMT to the HEMT therefore gets reflected back into the OMT. Most of this radiation exits the radiometer through the horn (or is terminated by the load), but a small portion (I_{OMT} =-40 dB) of the reflected radiation leaks into the other line. The leaked signal is attenuated by $R_{HEMT} + I_{OMT} = -45$ dB, so the response of the polarimeter to this signal should be $-45/2 = -23$ dB smaller than its response to polarized signals. This is consistent which the measurements which indicate that the unpolarized responsivities of the polarimetry channels are a few hundred times smaller than their polarized responsivities (Table 3.7).

The unpolarized responsivity is sufficiently small that it can be dealt with using a reasonable scan strategy (see below, Chapter 7). Even so, various methods for reducing the cross coupling (without redesigning the OMT or the HEMTs) were considered. For example, isolators could be inserted between the OMT and the HEMTs, preventing the radiation reflected from the HEMT input from reaching the OMT. Unfortunately, the reflection coefficient of an isolator is only about a factor of three smaller than that of the HEMTs. The radiation reflected from the isolator would consequently produce a cross coupled signal only three times smaller than that produced by the unisolated HEMTs. The unpolarized responsivity would therefore only be reduced by a factor of 1.5. Furthermore, the isolators would be lossy (1.5 dB) elements in front of the HEMTs, and would significantly increase the receiver noise temperature and degrade the sensitivity of the polarimeter (see Table 3.1). For these reasons, isolators do not seem to be a practical way to reduce the cross-coupling. Another way to reduce the unpolarized responsivities is to insert additional waveguide between the OMT and the HEMTs. This technique exploits the extra pathlength the cross-coupled radiation travels when it reflects off a HEMT, which means the leaked signal is not properly phase matched with its parent in the other line (Figure 3.11, the reflection of the signal off the HEMT also can introduce phases that further degrade the phase match). With sufficient waveguide between the OMT and the HEMTs the phase match will be so badly degraded as to virtually eliminate this source of cross coupling. However, W-band waveguide is quite lossy (0.15 dB/inch for copper), so there is a risk that the necessary amount of waveguide

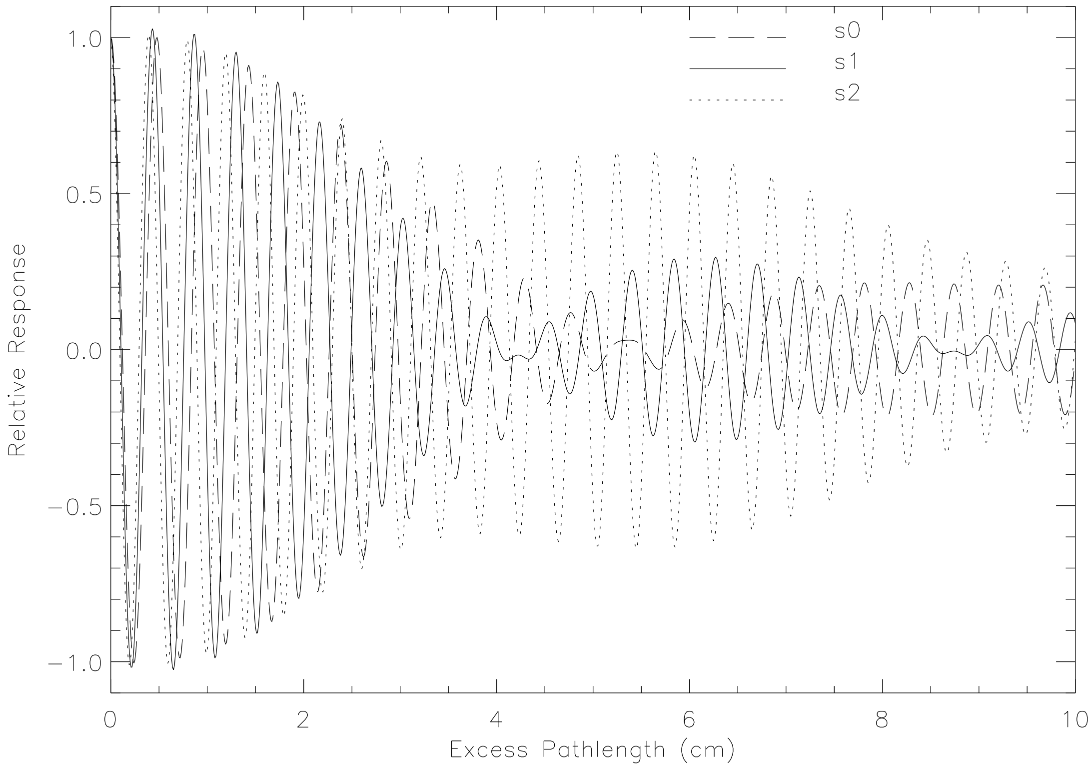

Figure 3.11: The relative responsivities of the PIQUE 2000 polarimetry channels to unpolarized signals (assuming a fixed cross-coupling) as a function of extra line length traversed by the cross coupled signal (the distance in the actual radiometer would be about 2 cm due to the finite size of the OMT). The fast wiggles correspond to the central frequency of each band passing through roughly $n\lambda/2$ (n odd) of line-length, while the long-wavelength damping is due to decoherence across the finite bandwidth

would significantly increase the noise temperature of the polarimeter. In any case, the space constraints inside the dewar do not allow for this much excess waveguide.

Constant Load Temperature: Sensitivity and Stability

Data are also taken over a period of hours with the circular load held fixed at one temperature (30 K). This provides information about the sensitivity of the polarimeter as well as the practical implications of the unpolarized responsivity of the multiplier channels.

The power spectra for the various signal channels, shown in Figure 3.12, are quite informative. The power spectrum for the total power channel shows a strong red slope, or $1/f$ structure. As discussed in Section 2.2.1, this $1/f$ noise is due primarily to gain variations in the HEMTs. The (unchopped) multiplier channel also shows $1/f$ structure on very long time scales, due to the finite responsivity of this channel to fluctuations in the total power (the relative size of the fluctuations in the total power and multiplier channels is consistent with the rejection ratios given in Table 3.7). On times scales shorter than a few minutes, however, the power spectrum of the polarimetry channel becomes a constant as a function of sampling frequency. On these time scales, drifts are sub-dominant to random gaussian fluctuations. During CMB observations, the telescope chops between two points on the sky with chop periods between 10 s and 30 s (see Chapter 5), differencing on these time scales (well below the $1/f$ knee) should completely eliminate any long time scale drifts. This is confirmed by "fake chopping" the data from a multiplier channel (i.e. binning the data into 5 second long segments and differencing adjacent segments). The power spectrum of these chopped data shows no $1/f$ structure at all, even on time scales on the order of a day.

The lack of $1/f$ structure in the power spectrum of the chopped data indicates that there is no excess power on long time scales, and that the system should properly integrate down following the formula $Q_{min} = \mathcal{S}/\sqrt{t_{int}}$ (see Section 2.2.1), where $\mathcal{S}$ is the sensitivity. This is confirmed by calculating the rms of the chopped data averaged over different time periods, and determining that the rms values do scale as the inverse square root of the integration time (see Figure 3.13). Therefore, given sufficient time, the polarimeter should

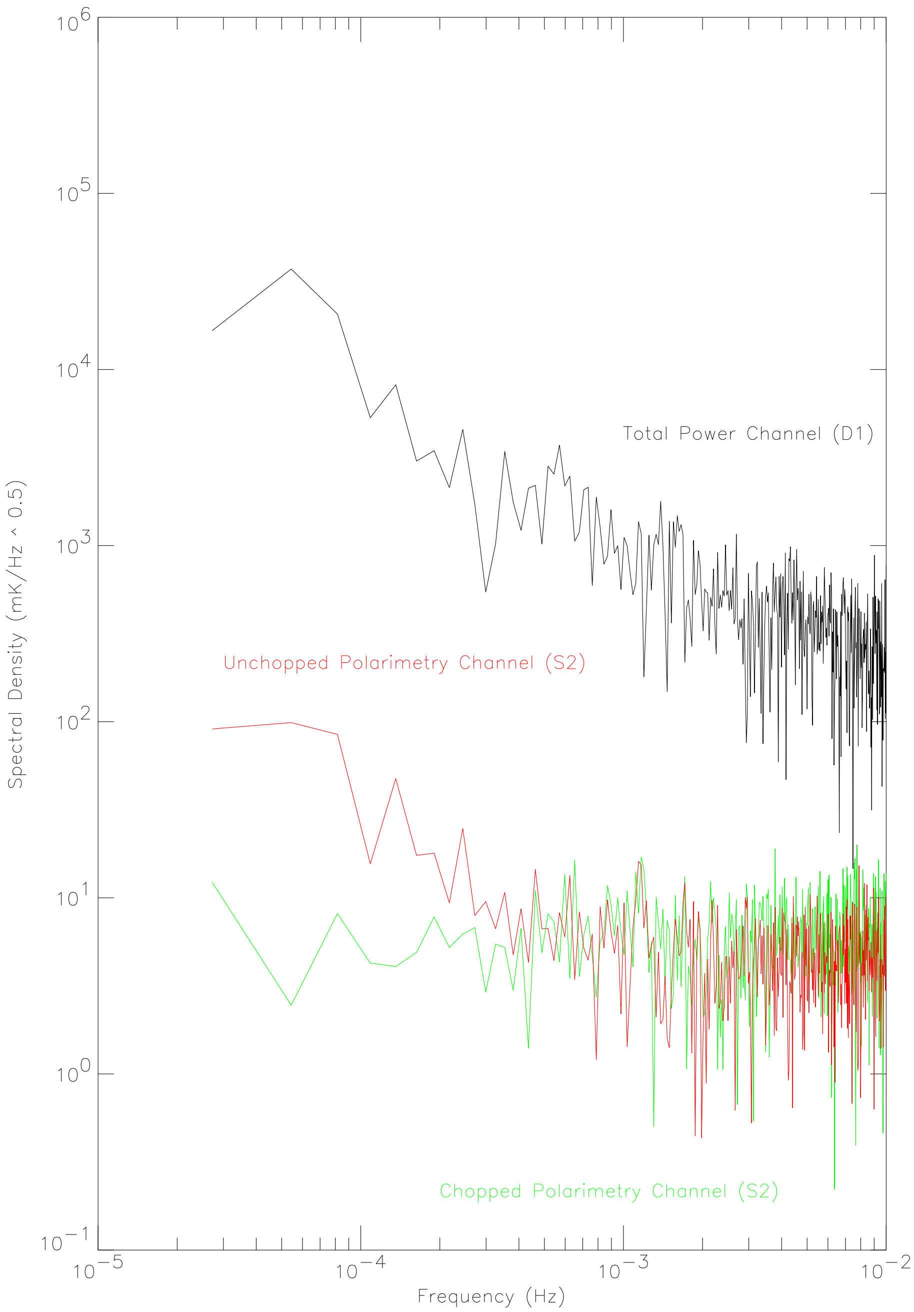

Figure 3.12: Power Spectra of selected radiometer channels with the radiometer viewing a 30 K circular load. The channels displayed are total power channel, a multiplier channel and the same multiplier channel with a fake 10 second chop applied.

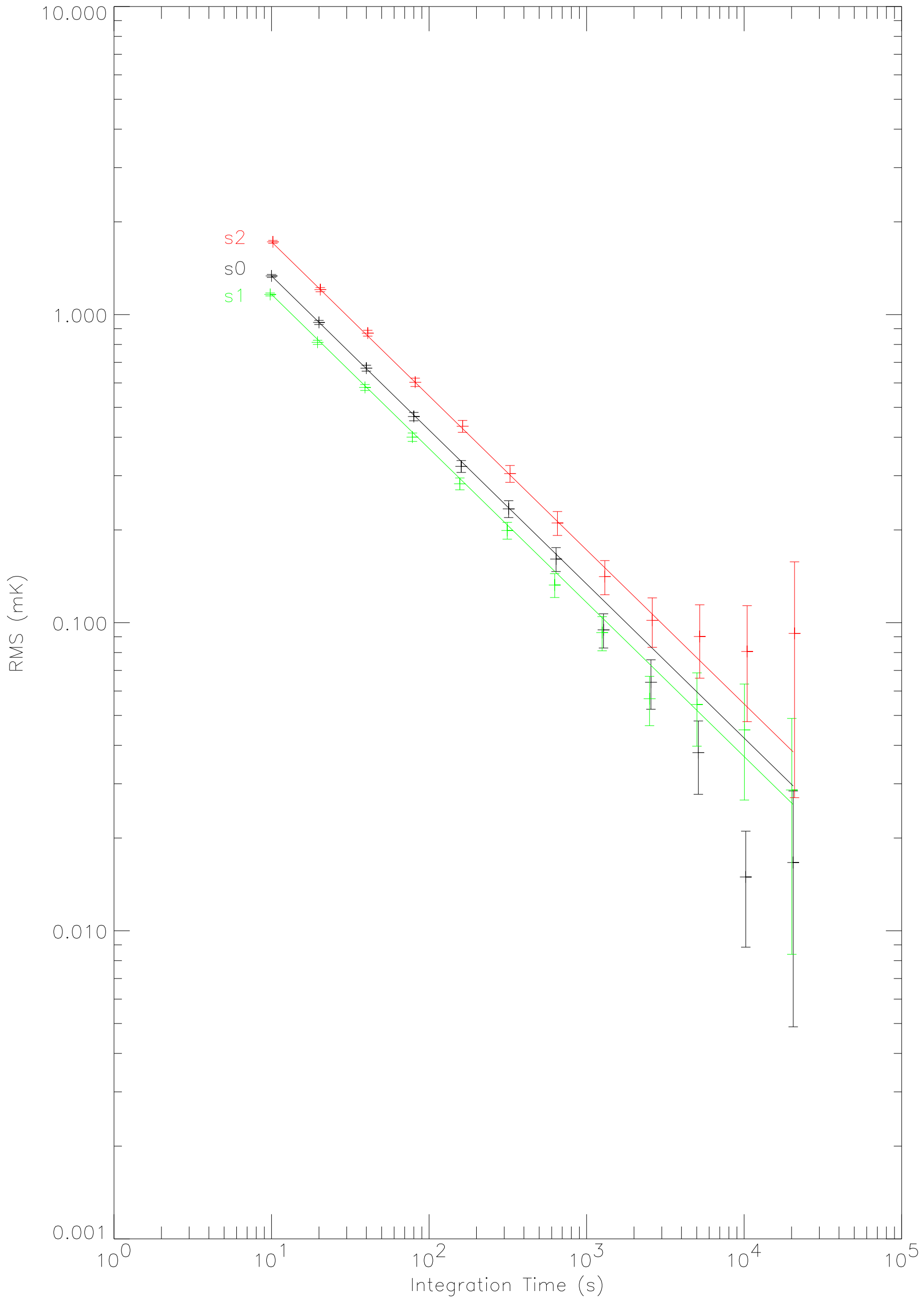

Figure 3.13: Integration down plot. This plot shows the *rms* of the 10 second fake chopped data from the circular load test, averaged over progressively longer time scales. The lines show the expected trend $rms \propto (\text{integration time})^{-1/2}$. The measured *rms* values follow these trends (to within the error bars), thus demonstrating the system integrates down properly.

Table 3.8: Circular load test results: sensitivity of multiplier channels (2001)

Character	s0	s1	s2
Sensitivity[a]	2.1 mK$\sqrt{\mathrm{s}}$	1.8 mK$\sqrt{\mathrm{s}}$	2.7 mK$\sqrt{\mathrm{s}}$
System Temperature[b]	78 K	70 K	96 K[c]

[a] Unchopped sensitivities with load temperature of 30K, calculated with gains from chopper plate tests (see Chapter 6).
[b] System temperature derived from the sensitivities using bandwidths from Table 3.2 and load temperature of 30 K. The factor $k(\nu)$ in the gains is taken into account.
[c] The high noise temperature estimated for s2 may be because the bandwidth in Table 3.2 is an overestimate. The $\sim$ 15% smaller bandwidth measured during the post-season tests yields a noise temperature more consistent with the other estimates.

be able to measure differences in the polarization of the sky with sufficient accuracy to detect the polarization of the CMB.

These rms values can also be used to calculate the sensitivity of the multiplier channels (Table 3.8): $\mathcal{S} = rms\sqrt{t_{int}}/2$. The factor of two arises because the rms (and the sensitivity) of the differenced data is twice that of the unchopped data. Furthermore, these values for the sensitivities can be combined with the bandwidths of the channels (measured with the narrow-band tests) to provide estimates of the receiver noise temperature ($T_{rec} = \mathcal{S} * \sqrt{\Delta\nu} - T_{load}$, see Section 2.2.1). The estimated noise temperatures for two of the channels are around 75 K (the third channel, s2, gives a higher noise figure, but this may be because the bandwidth in Table 3.2 is an overestimate). This agrees very well with the mean noise temperature derived from the total power channels (see Tables 3.4 and 3.6).

Finally, the (unchopped) power spectrum on short time scales is used to look for microphonics (Figure 3.14). The W-band system does show an extremely narrow feature corresponding to the cycle of the cold head of the refrigerator at 1.2 Hz. Fortunately there is very little power in this microphonic (roughly 1μK rms) and it does not affect the performance of the polarimeter.

3.3.4 Horn Test

Prior to both observing seasons, a final test on the integrated polarimeter was performed after the horn was installed but before the telescope was brought to the roof. In this test, the

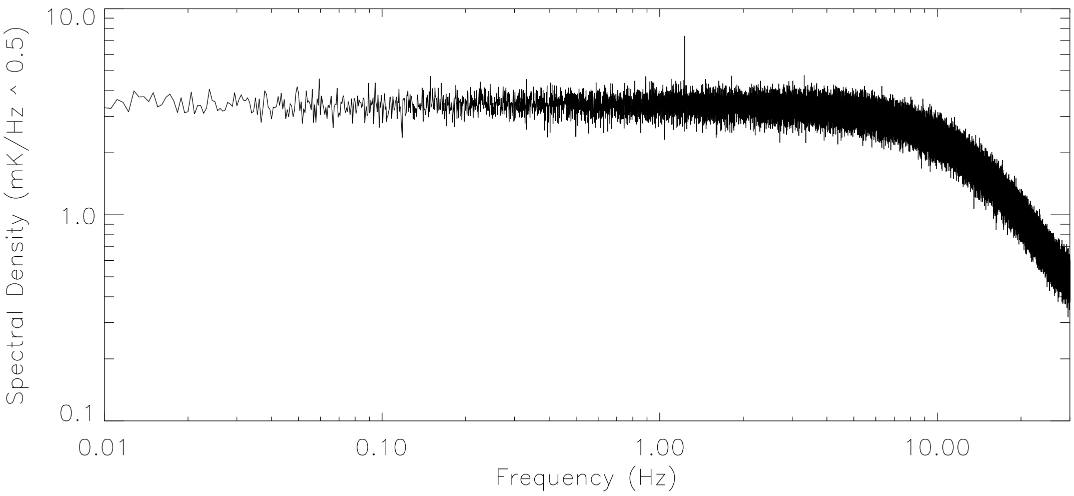

Figure 3.14: Power spectra of one multiplier channel (s2) on short time scales. The roll off around 10 Hz is due to the low-pass filter in the processing circuitry, the narrow spike near 1.2Hz is a microphonic due to the refrigerator cycle

Table 3.9: Estimates of total power responsivities and receiver noise temperatures from the horn test

Year	D0		D1	
	$\mathcal{R}$ (cts/mK)	T_{rec} (K)	$\mathcal{R}$ (cts/mK)	T_{rec} (K)
PIQUE2000	0.14	88	0.10	86
PIQUE2001	0.14	71	0.11	106

Table 3.10: Ratio of observed to expected results from the horn tests

Year	D0		D1	
	77 K	293 K	77 K	293 K
PIQUE2000	1.09	1.03	1.06	0.99
PIQUE2001	1.01	0.97	1.00	0.96

Expected values assume the following values for $\mathcal{R}$ and T_{rec}: PIQUE2000 D0=(0.14 cts/mK, 73K), D1=(0.11 cts/mK, 68 K); PIQUE2001 D0=(0.15 cts/mK, 62 K), D1=(0.12 cts/mK, 93 K).

response of the total power channels is measured with the horn viewing a room temperature load (293K source) and a load saturated in liquid nitrogen (77K source). These data should have provided the cleanest measurement of the total power gains and the receiver noise temperatures. However, the resulting estimates, tabulated in Table 3.9, indicate a problem with this test. The responsivity estimates are comparable to those derived from the load tests, but the estimates of the receiver noise temperatures are uniformly higher than the estimates derived from the load tests.

The source of this inconsistency can be clarified by considering the ratios of the observed output signals from the total power channels to the expected outputs based on the various load tests (for PIQUE2001, the three estimates for $\mathcal{R}$ and T_{rec} from both two-load tests and the circular load test are averaged together to calculate the expected response). These ratios, tabulated in Table 3.10 show a clear pattern. For both channels in both years, the ratio for the 293K source is less than the ratio for the 77 K source by about 5%. This suggests that the response of the detector diodes has become slightly nonlinear at 293K. This is not unreasonable, since in this case the output from the detector diodes is greater than 5 mV. Detector diodes normally go nonlinear when their output exceeds 10 mV (Pozar, 1998), but the diodes used in PIQUE have matching resistors on their outputs that reduce their gain somewhat, and so the linear regime may well be somewhat smaller.

With the data from the 293K source suspect, these tests cannot provide a clear estimate of the responsivities or the noise temperatures of the radiometer. However, the 77K data is almost certainly still within the linear regime of the diodes, and this data alone can provide a useful check on the estimates from the load tests. For PIQUE2001, the observed signal is almost exactly the expected value based on the load tests. This indicates the loss factor is correct for this case. For PIQUE2000, the observed response is between 6% and 9% larger than expected, which indicates that the injector loss is somewhat higher than predicted, perhaps due to a loose waveguide joint in the injector. This discrepancy will be discussed further in Chapter 6.

3.4 Additional Detector Features

Before concluding this chapter on the polarimeter, there are two features of the radiometers which should be briefly mentioned: the temperature sensitivity of the cryogenic IF amplifiers in the PIQUE2000 radiometer, and the limited dynamic range of the multipliers.

3.4.1 Cryogenic IF Amplifiers

During the first observing season, a pair of IF amplifiers were located inside the dewar in order to provide additional gain before the signals left the dewar (see Section 3.1). Due to limitations in the dynamic range of the multipliers, (see below), this additional gain was unnecessary. Furthermore, these amplifiers did not work well cryogenically, so they were removed for the second observing season.

The problem with the cryogenic IF amplifiers is that they only operate in a rather narrow temperature range. If they get colder than about 85K, their gain drops. This is accompanied by a decrease in their current draw, which decreases their self-heating, causing them to cool further and setting up a feedback loop that ultimately results in the amplifiers turning off entirely. Thus we needed to operate the amplifiers at a relatively high temperature, around 90K, which put undue thermal loading on the remainder of the cryogenic radiometer.

On the other hand, if the temperature of the amplifiers rises above 105K, they begin to oscillate, dumping a considerable amount of power at frequencies around 1 GHz. Fortunately, this is outside the bandpasses of the polarization channels. However, there is no filter in front of the detector diodes to block 1 GHz, so the total power channels are sensitive to these oscillations. Worse, the oscillations cause the amplifiers to draw more current, increasing their self-heating and setting up another feedback loop. This quirk of the amplifiers was not much of a problem in the lab, where the thermal loading on the cold head was constant and the refrigerator operated at constant pressure, but became a major headache in the field, where the loading changed significantly and the long-running refrigerator gradually lost pressure due to slow leaks and temperature shifts (see Chapter 5). Fortunately, the responsivities of the polarimetry channels did not change significantly as a result of these problems (see Chapter 6), although certain abrupt shifts in offsets of all five channels during the first observing season may have been due to these oscillations (see Chapter 5).

3.4.2 Multipliers

Relatively late in the testing program prior to the first observing season (after the horn was installed), a limitation in the dynamic range of the multipliers was discovered. It was known that the multipliers' response to correlated (i.e., polarized) signals compressed if these signals were too large (i.e., the responsivity becomes smaller with increasing signal level). Unexpectedly, the response of the multipliers to correlated signals also decreases in the presence of large uncorrelated signals (Appendix A). The latter effect was discovered in the operational polarimeter by inserting attenuators in front of a multiplier and measuring the rms of the (fake chopped) data in that channel while it looked at a 300 K load. The product of these rms values and the attenuation factors is an estimate of the relative responsivity of the multiplier and should be independent of attenuation. In fact, this product is smaller at higher input powers (low attenuations, see Figure 3.15), indicating that the responsivity is reduced when the input power is too high. More direct measurements of the multiplier responsivity (using a separate test set up) confirmed these results (Barkats, 2000). The signal entering the multipliers needed to be attenuated by at least 20 dB (equivalent to the leftmost point in Figure 3.15) to keep the responsivities from being affected by the unpolarized signal level. A 30dB attenuator was therefore inserted in front of the warm IF amplifiers. This extra attenuation entirely canceled out the gain of one pair of amplifiers (see Table 3.1), paving the way for the removal of the cryogenic IF amplifiers after the first observing season.

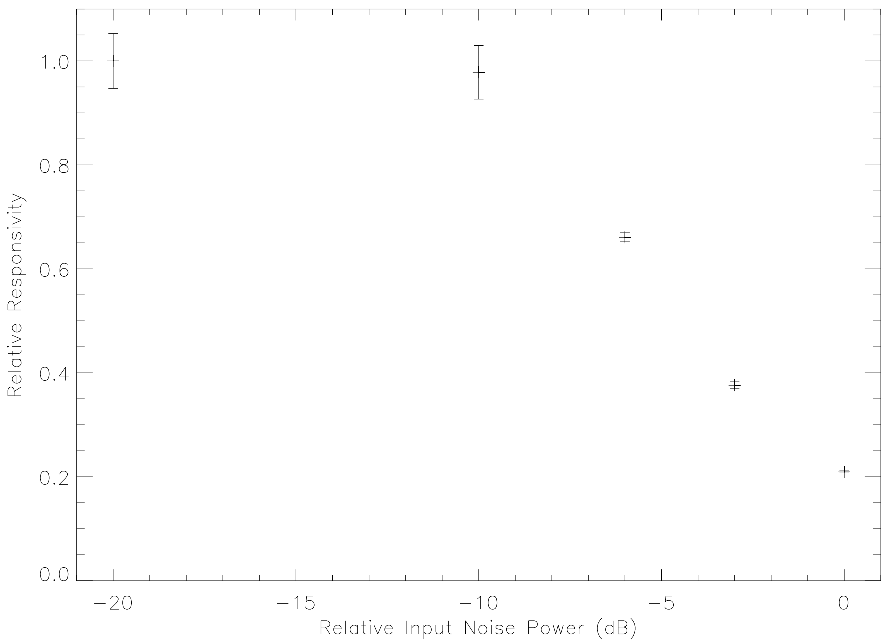

Figure 3.15: Responsivity of the multiplier as a function of uncorrelated input power. Note that the responsivity drops when the uncorrelated power exceeds a certain value

Chapter 4

Optics

The polarimeter described in Chapter 3 views the sky through a scalar feed horn and a 1.2 meter off-axis parabolic mirror. The entire system is mounted on a crane bearing (the only moving part in the entire telescope) which enables the telescope to scan back and forth in azimuth at a fixed elevation. While the mirror is recycled from the Saskatoon experiment (Wollack et al., 1997), the W-band feed horn and the ground screens are original to this experiment (see Section 4.1). Measurements prior to the first observing season establish that the telescope is properly aligned (see Section 4.2, also Chapter 6). Both the feed design and the evaulation of the measured beam patterns rely heavily on predictions from the physical optics programs CCORHRN and DADRA (Rahmat-Samii et al., 1996a,b)[1].

4.1 Description

4.1.1 Parabolic Mirror

The mirror was originally built for Saskatoon's 1995 observing season, and the PIQUE optics system is designed around this dish. A detailed description of this mirror is given in Wollack et al. (1997), so only a summary account of the mirror geometry is provided here (see also Figure 4.1). The mirror surface is a portion of a paraboloid with a focal length of 75 cm. The mirror edge is defined by the intersection of this paraboloid with a 122 cm

[1] See Appendices D and E for the input files for these programs

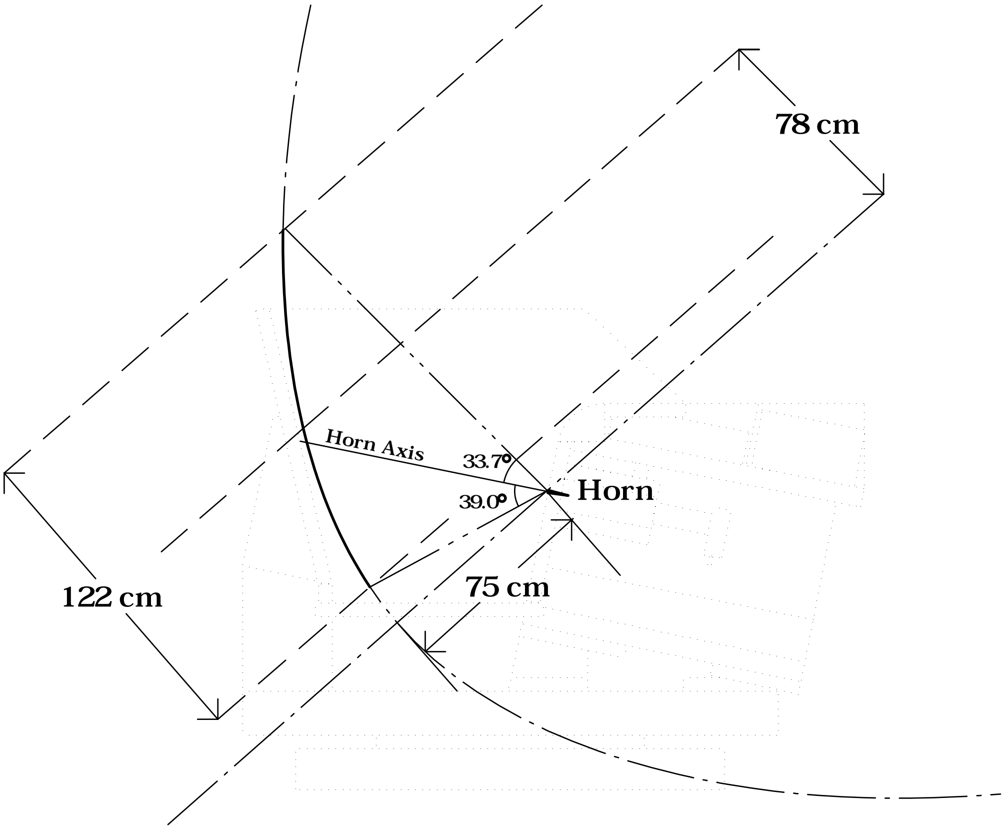

Figure 4.1: Optics geometry. The complete telescope and inner ground screen are shown in dotted lines, while the mirror and the horn are indicated with thick solid lines. The parent paraboloid for the mirror is denoted by the dot-dashed lines. The dashed lines show the axis and edges of the cylinder that defines the edge of the dish.

Table 4.1: Summary of optics characteristics

W-Band Horn

Flare Angle	3.5°
Aperture diameter	0.386 inches
Length	2.955 inches
Beamwidth	25°
Return Loss	-50 dB

Mirror

Type	Parabola
Focal Length	75 cm
Diameter	122 cm
Offset[a]	78 cm
rms surface roughness	<0.05 mm
Edge Illumination	-25 dB
Beam FWHM	0.24°

[a]See text, Figure 4.1

diameter cylinder. The central axes of the cylinder and the paraboloid are parellel to each other and separated by an offset of 78 cm.

The horn is located at the focus of the paraboloid and is oriented as shown in Figure 4.1. Note that the horn axis is not aligned with the geometrical center of the dish (the point where the axis of the cylinder intersects the surface of the dish), instead the horn axis is tilted 3° with respect to the line connecting the focus and the geometrical center of the mirror. This choice of horn orientation yields the most symmetric final beamshape because the edge illumination is nearly a constant around the lip of the mirror.

4.1.2 W-band Horn

The corrugated feed horn for the W-band polarimeter is designed to illuminate the mirror with an edge taper of roughly -25 dB. This edge taper is a compromise between reducing spillover off the dish (in order to preserve the symmetry of the beam and minimize sidelobes), and filling the dish as much as possible (in order to minimize the beamsize). The horn is illustrated in Figure 4.2, and a complete description of the inner surface (radius as function of depth), is provided in Appendix D. This final design has a flare angle of 3.5° and an aperture diameter of 0.386″. The corrugations in the body of the horn have the same

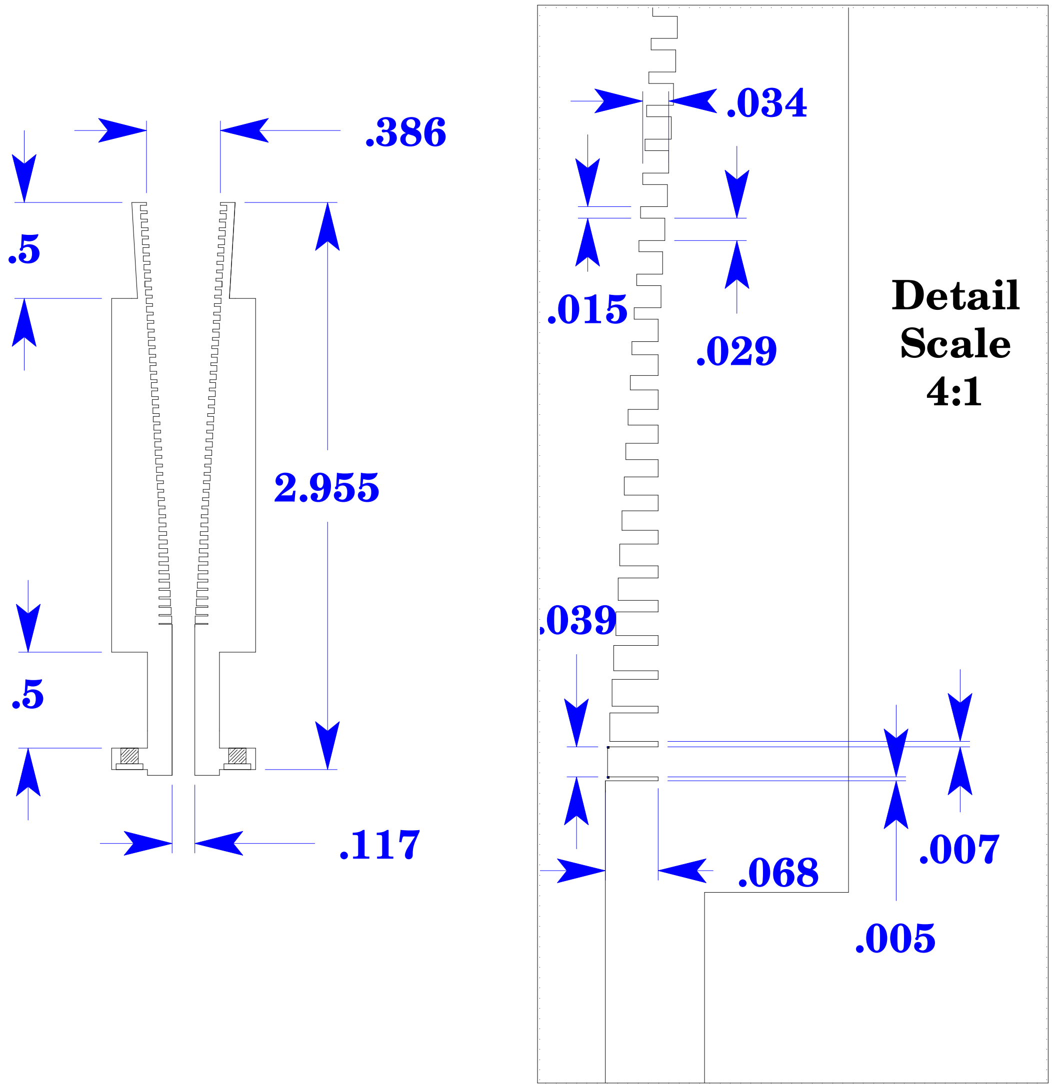

Figure 4.2: The W-band horn

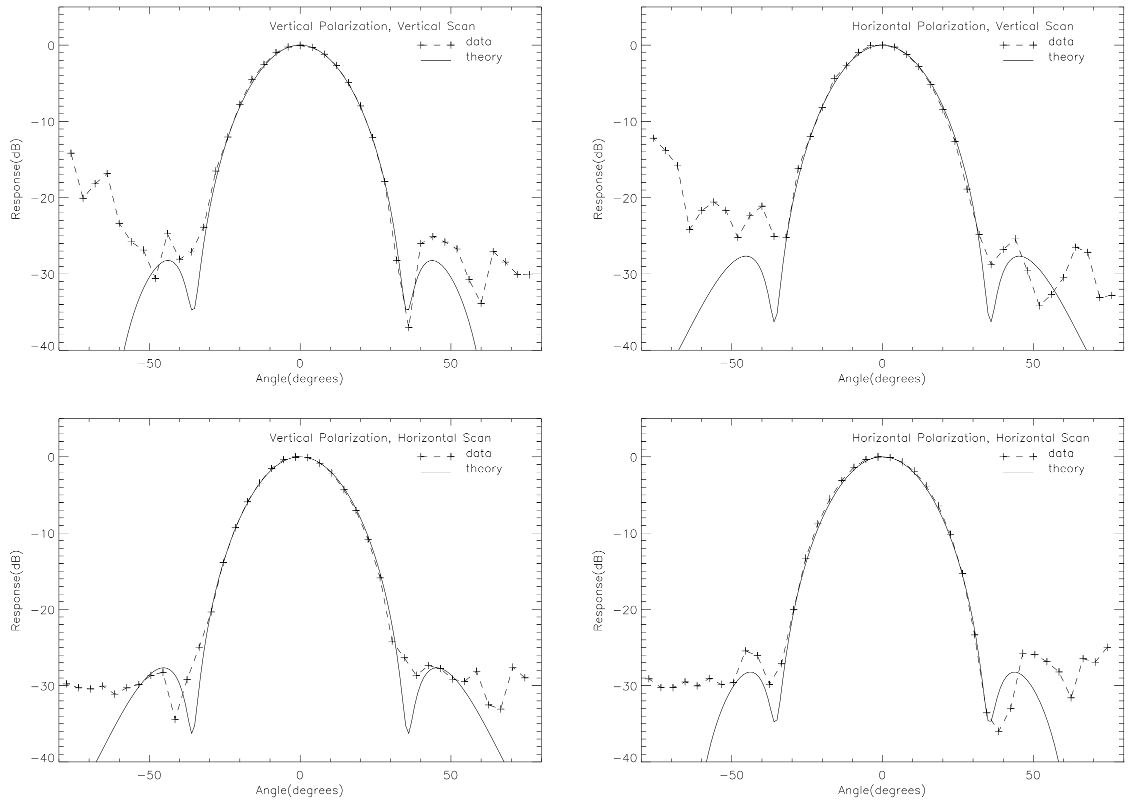

Figure 4.3: Beam maps of the W-Band Horn at 93 GHz, along with predictions from the physical optics program CCORHRN. For the main lobe, theory and data match very well. Measurements of the sidelobe response are limited by the noise floor of the detector, but these data show the sidelobes must be below -25 dB. At negative angles during the vertical scans, where the horn is pointing downward, the data show a rise due to increased pickup from the ground in front of the horn.

depth and spacing as the corrugations in the MAP W-band feeds (i.e. one quarter of a wavelength at 84 GHz; see Barnes, 2002). In the throat section, however, the grooves become progressively deeper and narrower in order to improve the coupling between free space and circular waveguide (Zhang, 1993). Indeed, the return loss of the horn is smaller than -50 dB (measured with a vector network analyzer).

The far-field beam map of the horn was measured using a 93 GHz, 125 mW source modulated at 1 kHz and an automated scanning platform developed for the MAP telescope (Barnes, 2002). During these measuremets the horn is connected through an OMT to a pair of detector diodes. The outputs of the diodes are sent into lock-in amplifiers which use a laser link (a switched laser at the source illuminates a photodiode near the lock-ins) to provide a reference signal. The output from the lock-ins are digitized and stored on a computer running a LABVIEW program. This program takes data as the motorized platform scans either in elevation or in azimuth. The resulting measured beam patterns match the predictions from the physical optics progrm CCORHRN (Figure 4.3).

4.1.3 Ground Screens

The ground screens are designed to prevent microwave radiation from sources like the ground and nearby buildings from reaching the polarimeter. These sources are bright (~ 300 K) and have large angular extents, so if even a small fraction of their emission enters the polarimeter, it can cause a significant response in the multiplier channels (either because it has a true polarized component or because of the unpolarized responsivity of these channels, see Section 3.3.3). These sources also change temperature over time so they can produce confusing time variable signals (such time-variability has caused problems in other polarization experiments, see Nanos (1979), Keating et al. (2001b)). It is extremely difficult to accurately model the responsivity of the telescope to these sources (which are far from the main beam); the relevant physical optics calculations are computationally expensive, and the scattering from elements besides the dish is relevant. The ground screens are consequently designed using simple ray tracing. Ideally, the ground screens should block any ray connecting the horn or mirror with any source of local foreground emission, and reflect these rays to the relatively cold, stable sky.

Telescope Site

The design of the ground screen is strongly influenced by the observing site. For convenience and accessibility, the PIQUE telescope is located on the roof of Jadwin Hall, the physics building for Princeton University (Figure 4.4). This roof has been used for microwave telescopes before, and there are two locations on the roof –the eastern and northern corners– which were specially outfitted for telescopes prior to PIQUE. Neither of these sites is used

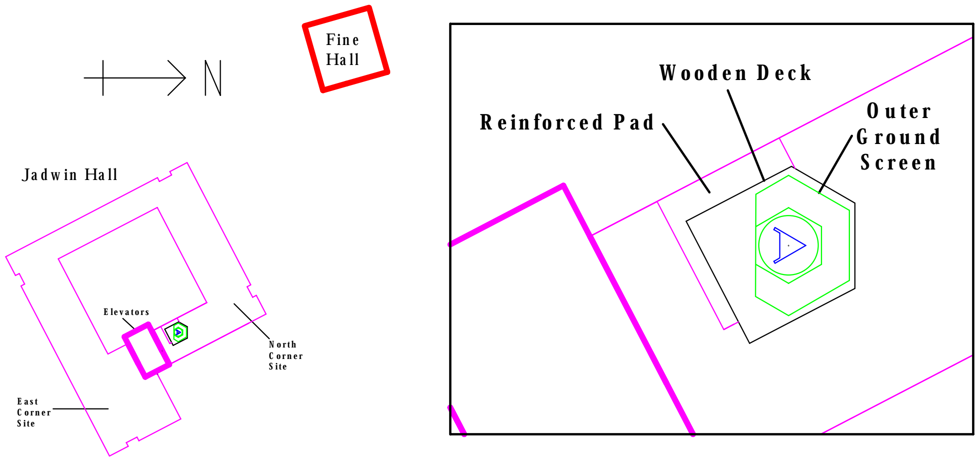

Figure 4.4: Diagram showing the location of the PIQUE telescope on the roof of Jadwin Hall.

for the PIQUE telescope. The eastern corner is not used because a narrow doorway and some ductwork between this site and the elevators make it hard to access. The northern corner of the roof would have been a good location, especially since there is a set of "hard points" in this part of the roof that would connect the telescope directly to the frame of the building. However, this site was not available because it was already being used to conduct tests for the MAP satellite.

The PIQUE telescope site is located along the northeast portion of the roof, between the elevators and the northern corner (see Figure 4.4). The exact location of the telescope is determined by the following considerations: (1) Most of the weight of the telescope needs to sit over a reinforcing pad on this part of the roof, and (2) the telescope needs to be sufficiently far away from the elevator shaft building so that the latter could be blocked with reasonably sized ground screens. The telscope is therefore located as far north on the pad as possible. Three new hard points were constructed for the telescope at this location and a large wooden platform was constructed to support the ground screen (this platform actually extends off the edge of the pad and is appropriated weighted so that the portion of the platform not supported by the pad is cantilevered over the roof.)

Outer Ground Screen

Once the site is chosen, the ground screens can be designed. Two sets of nested ground screens are used: the inner one is attached to the base so that it rotates with the telescope, and the outer one is attached to the ground. Normally, the inner ground screen is designed to block all possible sources of ground pick-up, and the outer ground screen is designed around the inner one. This method was not practical for PIQUE given the horizon at the site. As can be seen in Figures 4.4 and 4.5, there are two objects near the PIQUE site that are quite tall and that cover significant solid angles: (1) the elevators on the roof of Jadwin itself, which are located almost directly south of the telescope, and (2) a very tall nearby building (Fine Hall) to the northwest of the site. The ground screens need to block both these objects from the telescope's view and still allow the telescope to observe the NCP. These constraints require an 11 foot tall ground screen. It seemed unlikely that something this big could be light enough to move with the base and rugged enough to handle high winds. Consequently, it was decided to make the outer ground screen, which is attached to the ground and reinforced with a wood frame, big enough to block these sources, and then design the inner ground screen around the outer one.

The shape and size of the outer ground screen are determined based on the following considerations: (1) the ground screen has to geometrically block radiation from the elevators and Fine Hall from entering either the mirror or the horn; (2) the observed region around the NCP needs to be completely visible to the telescope, and (3) the telescope needs to be able to rotate freely around 360° within the ground screen. The resulting outer ground screen is a truncated hexagonal pyramid, with a vertical height of $11'$ and a footprint $11'$ across. The southern wall is vertical while the other walls slope outward at 20° (Figure 4.6). The edge of the northernmost panel (the edge closest to the main beam) is curved with a radius larger than $6''$ to reduce diffraction (Keller, 1962). The edges of this ground screen, as viewed from the top and bottom of the dish, are shown in Figure 4.5. Fine Hall and the elevators are below the edge of the ground screen as seen from the top of the dish, so these buildings are completely blocked from the telescope (The vertical southern wall also

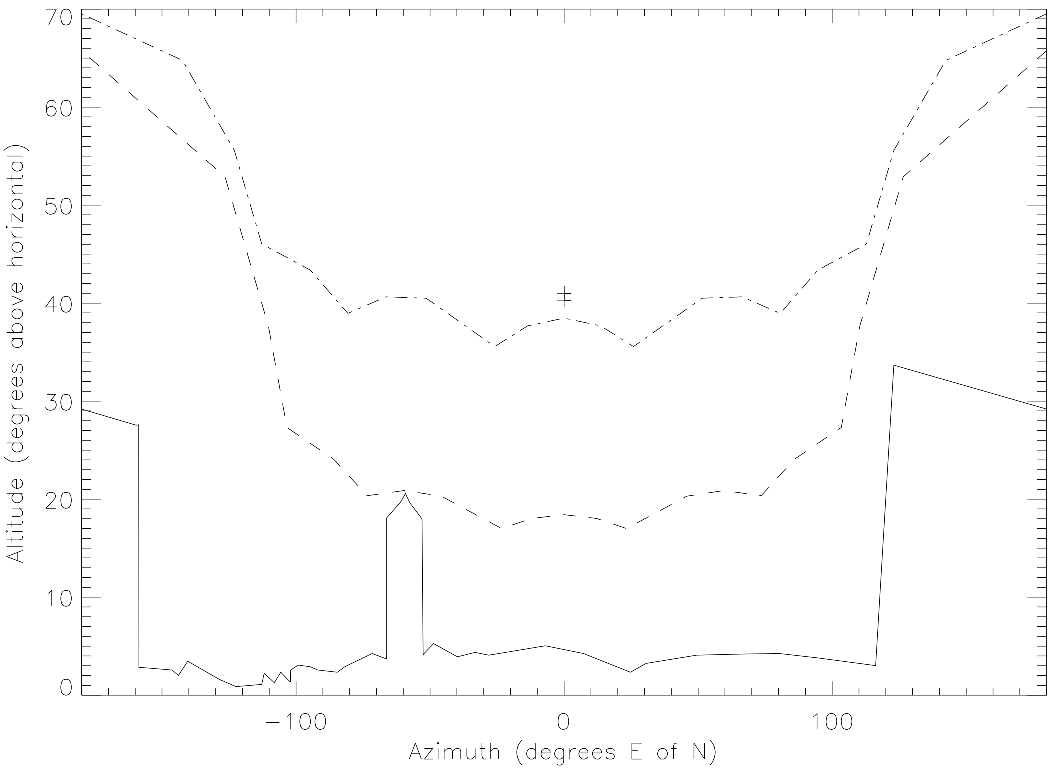

Figure 4.5: The solid line indicates the horizon from the PIQUE observing site, as measured with a theodolite. Note the tall building (Fine Hall) to the northwest and the elevator shaft to the south, The dashed (dot dashed) line shows the edge of the PIQUE2000 outer ground screen viewed from the top (bottom) of the dish. The region below the dashed line is completely blocked from the dish, while the region above the dot-dashed line is completely visible to the telescope. The two crosses mark the observing altitude of the telescope in 2000 and 2001

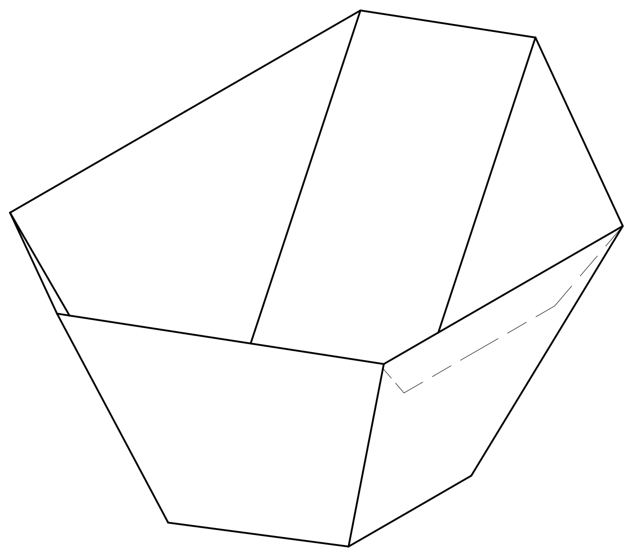

Figure 4.6: The outer ground screen. the dashed line indicates the portion of the front panel that was cut away for the second observing season

prevents radiation from the elevators from reaching the horn). Furthermore, the NCP is well above this edge as viewed from the bottom of the dish, so this region of the sky is clearly visible to the telescope.

A ground screen this large poses certain practical difficulties. It is very difficult to cover up something this big in the event of rain or snow, and an $11'$ tall structure capable of withstanding the full force of winter storms would be an unacceptably large, heavy and dangerous object for the roof of a physics building. The panels are therefore hinged $6'$ above the platform so they can fold down, allowing a tarp to be drawn over both the ground screen and the telescope. The frame of the lower half of the ground screen is constructed of wooden beams, while the upper half uses a lighter aluminum frame, which is held up by a combination of latches and long wooden struts. The panels themselves are constructed of light aluminum- covered styrofoam (Homasote Ultra/R$^{\mathrm{TM}}$ Exterior Sheathing, $1''$ thick). When all the panels are up and locked in place, the outer ground screen is extremely rugged, and easily withstood measured wind speeds in excess of 50 mph.

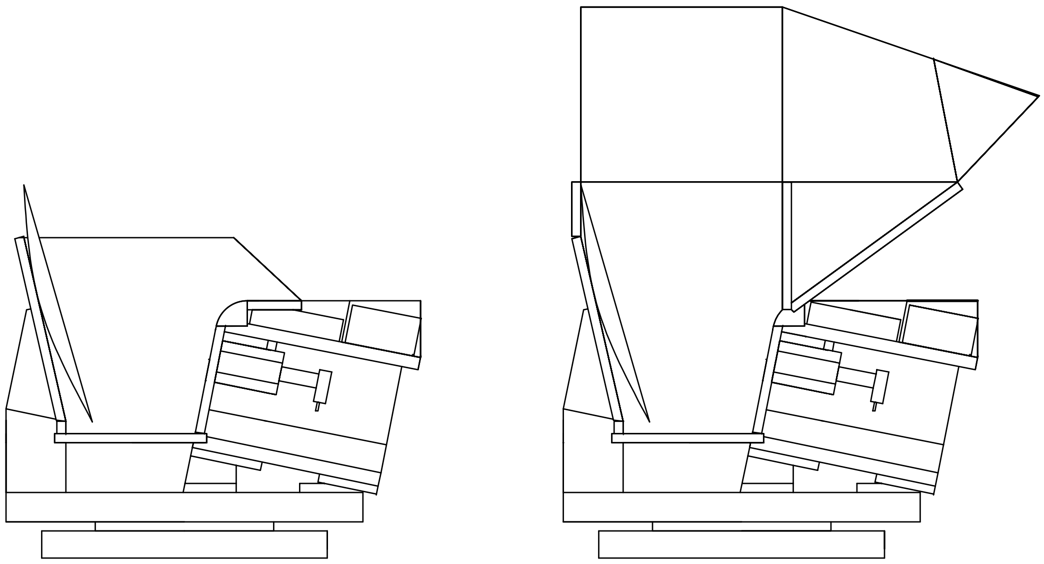

Figure 4.7: The inner ground screens. On the left is the small ground screen used for the 2000 season, while on the right the extended ground screen used for the 2001 season is shown.

Inner Ground Screens

The original design criteria for the inner ground screens were (1) that they should prevent the dish from seeing the floor of the outer ground screen and (2) that they should not extend too far above the folded- down outer ground screen (so that it would be easy to cover up and stow). The resulting ground screen for PIQUE2000 was rather modest (see Figure 4.7) and was constructed of sheet aluminum. The top edge of the side panels were designed such that any rays from the top of the dish that cleared the inner ground screen would be reflected upwards by the outer ground screen. The other important feature of this ground screen is the curved surface located on the front panel over the radiometer. The radius of curvature of this portion of the ground screen is fairly large ($7''$ radius). This not only prevents the ground screen from blocking the NCP[2], but also reduces the spurious polarized signals due to diffraction around this edge (Keller, 1962). The other edges of the

[2]Note in Figure 2.4 how close the ray from the bottom of the dish passes to the curved surface of the IGS.

inner ground screen, which are far from the main beam, are rounded with a smaller radius of curvature (1").

The version of the inner ground screen used during the first observing season was not ideal. As described in Chapter 8, the telescope measured a large (some hundreds of μK) difference between the polarized signals in the two scan positions during the first observing season. This "chopped offset" (so called to distinguish it from the offsets that are removed by chopping, such as that due to the unpolarized responsivity of the polarimetry channels) was the result of the telescope seeing the panels of the outer ground screens. The thermal emission from a metal surface has a small polarized component that depends on the orientation of the surface (see Section 6.3). Since the telescope is oriented at slightly different angles with respect to the panels of the outer ground screens when it is in the two positions, there is a net difference in the polarized signals picked up through the sidelobes of the instrument in these two positions.

At the end of the first observing season, the sources of the chopped offset were mapped out by hanging rugs made of a microwave absorbing material around the inside the outer ground screen. These rugs are much worse reflectors than aluminum (but not so bad as to be near-perfect blackbodies), which means they generate much stronger polarized signals than the ground screen itself. The responsivity of the telescope to signals from different panels can then be explored. These tests showed the telescope responds most strongly (on the order of mK with changes of orientation on the order of degree) to signals from the southernmost (back) panel of the outer ground screen. The northernmost (front) panel is also a source of measureable pickup, although it was an order of magnitude smaller. The other panels did not yield any detectable signals. These results confirm that the chopped offset is due to the parts of the outer ground screen which are located in regions of high sidelobes, either due to spillover past the dish (back panel) or due to proximity to the main beam (front panel).

Further tests done at the end of PIQUE2000 established that the chopped offsets could be significantly reduced if the inner ground screen was extended (with extra aluminum-covered foam panels), and that the azimuth drive motor could handle such tall inner ground screens. The inner ground screen used during PIQUE2001 is therefore tall enough to prevent radiation from any part of the outer ground screen from entering the horn or the mirror (after a part of the northernmost panel of the outer ground screen was removed, see Figure 4.6). Like the outer ground screen, the panels of the enlarged inner ground screen are hinged so they can be folded down when the telescope is stowed. The lower half of these panels are made of sheet aluminum, while the upper half are made of foam panels in an aluminum frame.

The larger ground screens reduced the chopped offsets by a factor of two or more. However, there were still chopped offsets as large as a few hundred μK throughout PIQUE2001 (see Chapter 8). The remaining offsets are almost certainly due to a small gap that was left uncovered between the bottom of the mirror and the bottom of the ground screen. Since this gap is very close to the dish, it is in a region of high sidelobes, and the horn could view the wooden floor through this gap. When the Q-band system replaced the W-band system, this hole was covered with metal tape and the offsets became as small as 50 μK.

4.2 Focusing the telescope with near-field beam maps

In order to achieve the correct final beamwidth and obtain a reasonably symmetric beam, the horn needs to be correctly oriented and located at the focus of the dish. Orienting the horn properly is relatively easy. The horn is rigidly mounted along the central axis of the dewar and passes through a small hole in the dewar behind the window. This hole has about 1 mm clearance, so the horn axis must be aligned with the dewar axis to better than $0.1°$. The dewar in turn is mounted on a cradle. The cradle can be adjusted so the dewar is held parallel with the bottom of the cradle to an accuracy of about $0.1°$, as measured with a digital level. Finally, the cradle sits on to a pair of rails attached rigidly to the telescope base (see Figure 2.4). So long as the slope of these rails is correct ($13.6°$ for PIQUE2000, $14.3°$ for PIQUE2001), which can be verified with levels (to an accuracy of $0.1°$), the focused horn will point in the correct direction, $3°$ below the geometrical center of the dish (see Section 4.1.1). The uncertainty in the orientation is therefore conservatively

0.2°, which is perfectly acceptable (changing the orientation by this much has no signifcant effect on the beam).

Locating the horn at the focus of the dish is not so straightforward. The symmetry of the mount makes it reasonably easy to locate the horn in the horizontal direction perpendicular to the horn axis. However, locating the horn correctly in the vertical plane including the horn axis is difficult due to the large (over 75 cm) distance between the horn and the mirror. Measurements using wires stretched between the dewar and the dish, as well as calculations based on the design of the base, were used to roughly locate the horn before beam mapping began.

Beam maps of the W-Band system were obtained prior to the first observing season. These measurements use the nearly fully assembled telescope.[3] The crane bearing allows the telescope to scan in azimuth, and a special jack (see Figure 4.8) allows the telescope to scan in elevation. The jack's range of motion is limited, so the source has to be located near the fixed observing altitude of the telescope (41°). The source is therefore situated on top of a very tall nearby building (Fine Hall). The source (93 GHz, 125 mW) is modulated at 1 kHz so that the signal can be isolated from other sources of microwaves in the area. Due to the complications involved in extracting a 1 kHz signal from the multiplier channels (which are phase switched at 4 kHz), the beam patterns are measured using the total power channels.

In order to extract the modulated signal from the mapping source, the back-end electronics have to be slightly modified. The low-pass filter in the total-power channel thin boxes is therefore replaced with a lock-in amplifier (Stanford Research Systems Model SR850), which is coupled into the processing circuitry through the BNC ports on the front of the thin box. The signals enter the lock-in after the last stage of (software-programmable) gain. Various attempts to define a reference signal for the lock-in (involving a switched laser illuminating a photodiode) proved unreliable. Fortunately, the internal references of the lock-in and of the signal modulator are both sufficiently stable that the external reference signal is

[3]The polarimeter was only missing the HEMTs, which were unnecessary and could have been damaged as the telescope was manhandled during these tests. A specially-made pair of W-Band waveguide bends replaced the HEMTs and coupled the OMT directly to the filters.

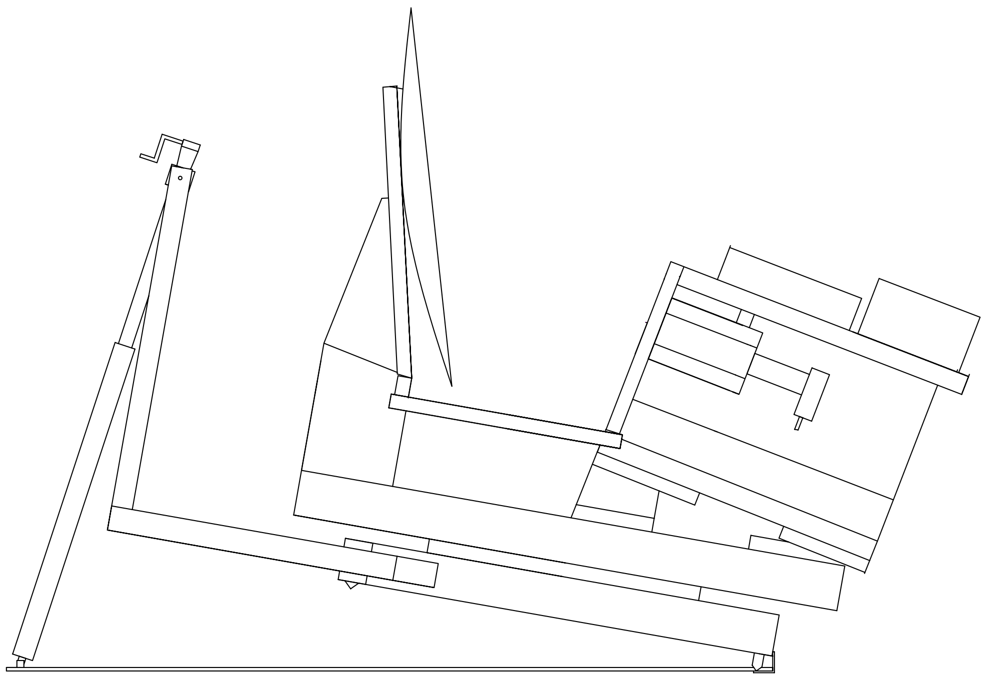

Figure 4.8: The jack used to alter the observing angle (altitude) of the telescope during beam mapping. The base pivots about the point at the front end of the base, while the back end of the base is raised and lowered using the crank. The tilt angle is measured with a clinometer (a surplus gun sight) attached to the base.

unnecessary, and the internal reference of the lock-in is used for most of the measurements. The magnitude of the 1 kHz signal measured by the lock-in (using the $R = \sqrt{(X^2 + Y^2)}$ option) is sent (as an analog voltage) into the V-to-F circuit in the thin box. The signal is then digitized and sent to the computer as usual (see section 3.1.4).

Beam maps are obtained by first orienting the telescope so that it points directly at the source, and then scanning the telescope either in azimuth (using the drive motor and bearing of the telescope mount) or in altitude (using the special jack). Due to the response time of the lock-in, continuous scans are not practical. Instead the orientation of the telescope is stepped through a discrete set of values (roughly $0.05°$ apart, total range $1°$). At each orientation, the lock-in is given enough time to settle down (up to a few minutes, depending

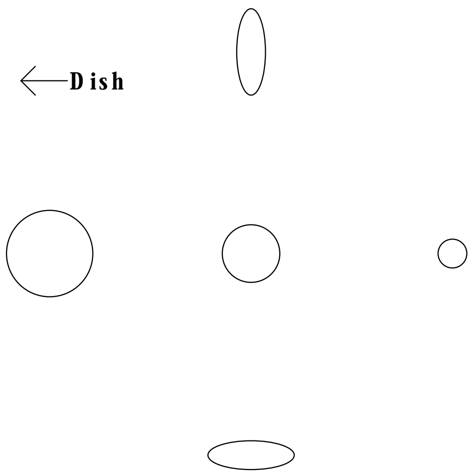

Figure 4.9: The distortions in the size and shape of the near field beam due to the location of the horn relative to the focus. Note that horizontal displacements change the size of the beam, while vertical displacements change the shape of the beam

on the signal strength), and then eight frames worth of data are stored to a file (using a special command in the data-taking program). These frames include 256 samples of each of the values of all five signal channels, as well as the azimuth position of the base measured by the 16 bit encoder. The altitude of the telescope is determined from a clinometer (a surplus gun sight, accurate to about $0.01°$) attached to the base. The clinometer reading is added to the data frames by hand (as part of the command) each time the data are recorded.

For the source to be in the far field of the mirror, it would have to be about 500m from the telescope. In fact, the source and the telescope are only 82m apart. Fortunately, the physical optics program DADRA can calculate the appropriate near field beam pattern. In fact, the near field beam pattern is extremely useful for focusing the telescope because the near field beam pattern is more sensitive to the telescope being off focus than the far field pattern (for example, moving the horn $0.125''$ above focus changes the near field beamwidths by 8% while the far field beam pattern only changes by 2%). Furthermore, the distortions in the near-field beam clearly determine the location of the horn with respect to the focus (see Figure 4.2): If the horn is located above or below the focal point, the average beamwidth is largely unaffected, but the beam becomes oblate (wider in altitude if the horn is above the focus, wider in azimuth if it is below). On the other hand, if the horn is displaced along the horn axis, the beam remains circular, but the size of the beam is too small if the horn is

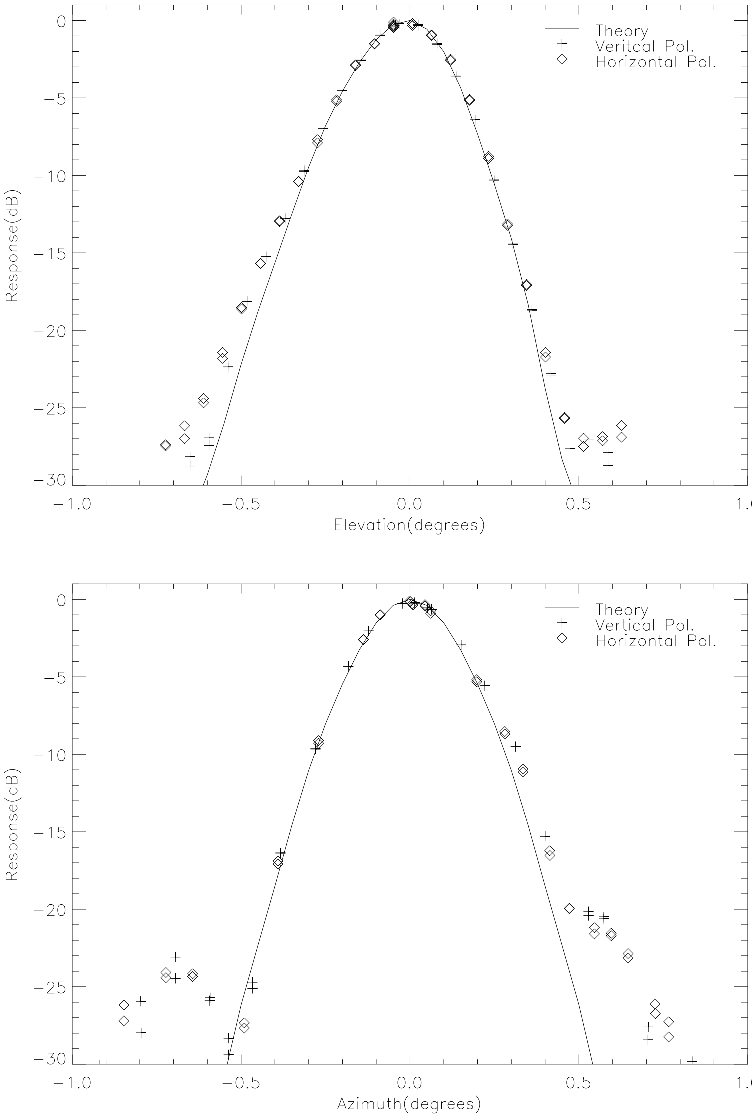

Figure 4.10: The final beam maps obtained with the W-band system viewing a local source, compared with the predictions from the physical optics program DADRA. The agreement is good down to -10 dB off the peak, but below that there are discrepancies which are most likely due to scattered radiation from parts of the roof where the source was located. The asymmetry of the beam in elevation is an artifact of the near field, the far field beam is symetric in this direction.

behind the focus and too large if it is in front of the focus (unlike the far field beam, which becomes broader if the horn is located before or behind the focus). This occurs because locating the horn behind the focal point causes the telescope to focus at a finite ditance instead of infinity. Distortions created when the horn is located to the left or the right of the focus are also distinctive, but the symmetry of the dewar mount makes it difficult to displace the horn in this direction.

The first beam maps of the main lobe gave an average beamsize that agreed with the predicted value, but the beam was noticeably oblate. This was due an error in the mount design which resulted in the horn being located half an inch below the focus. Making the appropriate adjustments yielded beam maps (see Figure 4.10) which agree reasonably well with the predictions down to about -10 dB from the peak. Further from the peak, the data indicate a wider beam and higher sidelobes than predicted. Most of these discrepancies are attributed to radiation scattering from various structures on the roof surrounding the source, although no conclusive tests were made to confirm this theory.

Additional beam maps beyond the main lobe are made by scanning the telescope in altitude such that the main beam points increasingly far above the source (scans in other directions would be corrupted by reflections off the building). These maps (figure 4.11) showed a larger then expected response 1°-2° from the main lobe. This discrepancy is attributed to diffraction over the crate holding the dewar (no ground screen was present during these measurements), although again there are no hard data to prove this supposition.

While there are differences between the measured and expected beam patterns, these discrepancies are relatively small, and the data indicate that the telescope is properly focused. Observations of the planets during the observing seasons confirm that the far-field beam pattern has the predicted shape (see Chapter 8).

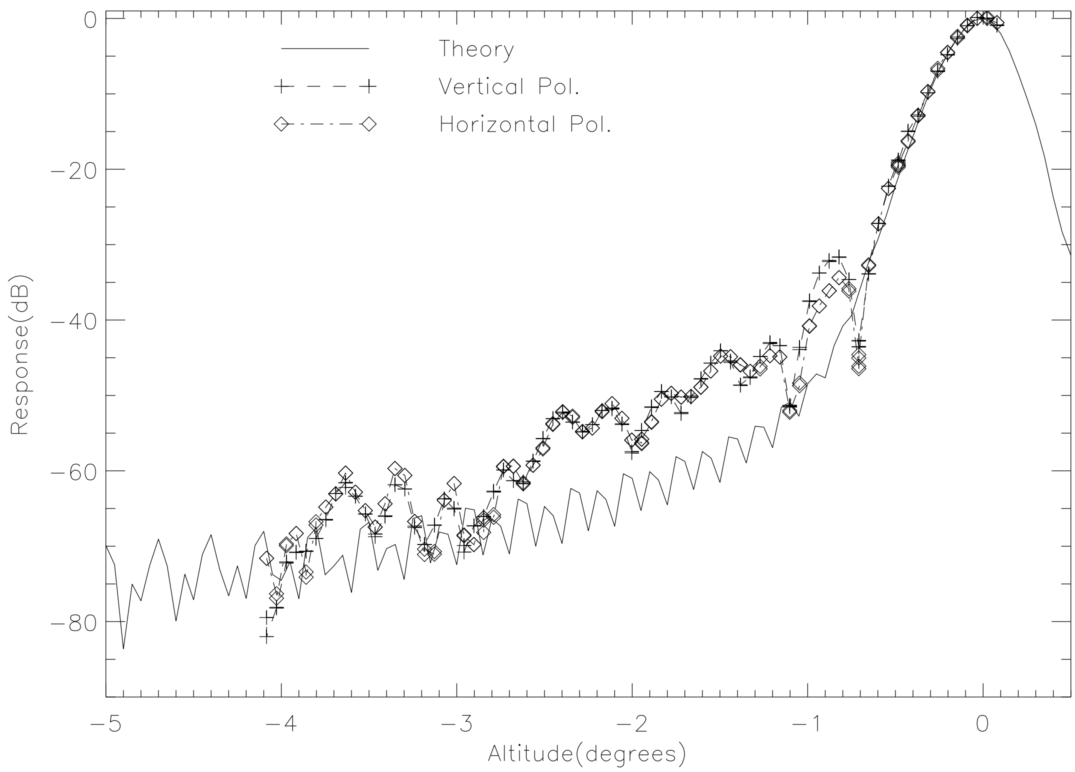

Figure 4.11: Far side-lobe beam maps, compared with theory

Chapter 5

Summary account of the Observing Seasons

The operational W-band telescope sat upon the roof of Jadwin Hall from January 1, 2000 to May 8, 2000 and from December 15, 2000 to February 28, 2001 (The W-band radiometer was then removed and replaced with the Q-band system, which operated until May 8, 2001). During these periods, the telescope measured the polarized microwave emission from the sky around the NCP whenever possible. These observations were interrupted at various times by weather, scheduled calibrations and other tests, and a number of quirks, glitches, problems and catastrophic telescope failures that kept the observing seasons interesting (and the telescope operators confused). Figure 5.1 and Table 5.1 summarize how much time was spent in each of these situations.

5.1 Nominal Observing Program

5.1.1 Deployment

The operational telescope is moved out to the site over the course of about a week. During this process, the base is leveled by attaching a clinometer to the base and measuring variations in the tilt as the base is rotated over a wide range ($> 180°$) in azimuth. The heights of the three legs that support the telescope are adjusted until the clinometer output

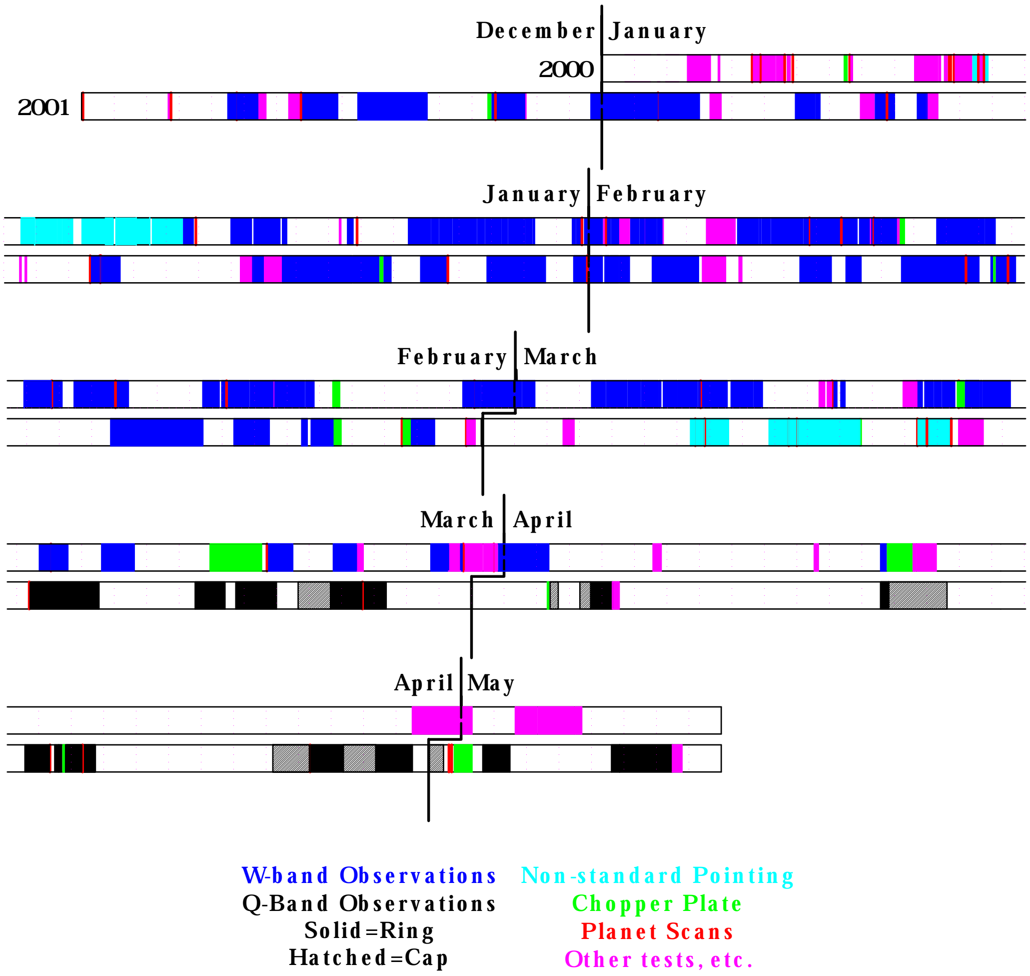

Figure 5.1: Timelines of the two observing seasons. The differently colored bands correspond to various periods when data was being acquired. The blank regions are time when the telescope was stowed or otherwise not taking data.

Table 5.1: Allocation of time during the two observing seasons

Situation	2000	2001
Total Time	3100 hrs	1800 hrs
Initial Tests/Pointing	250 hrs	100 hrs
Post-Season Tests etc.	850 hrs	0 hrs
Active Observing Season	2000 hrs	1700 hrs
Telescope Stowed (Bad Weather)	910 hrs	890 hrs
Electromechanical Problems	50 hrs	80 hrs
Nonstandard Pointing	100 hrs	0 hrs
Planet Scans	20 hrs	10 hrs
Chopper Plate Runs	50 hrs[a]	15 hrs
Other Scheduled Tests	60 hrs	40 hrs
CMB observations	810 hrs	660 hrs

[a]Additional chopper plate runs were performed during the Initial Test and Post-Season Test Periods

is constant as a function of telescope azimuth. The base in then level to the accuracy of the clinometer, which is better than $0.01°$.

After the radiometer is connected up and turned on, the telescope pointing and the azimuth encoder calibration (i.e. the encoder value corresponding to true north) is established with observations of Jupiter (see Chapter 4). However, limiting the search for Jupiter to a reasonable range in azimuth ($\sim 2°$) requires a preliminary rough estimate of north. For the second observing season, this rough estimate was generated based on information from the previous season, and was confirmed with an early observation of the moon (which is an extremely bright source easily visible in real time, even with the dewar warm). Unfortunately, at the beginning of the first observing season the moon was not high enough in the sky to be observable. The rough estimate of north therefore had to be obtained from rather crude methods: aligning the telescope with the hard points, and viewing the Polaris through a small optical telescope attached to the dish. These methods were sufficiently accurate that Jupiter did indeed fall within the scanned region (barely).

Several Jupiter scans were performed early in each observing season and the data from these scans were processed to determine the position of Jupiter (see Section 6.1). This information provides a more precise estimate of the encoder position corresponding to north

(21943 for PIQUE2000, 48540 for PIQUE2001). These encoder calibrations were used to determine the CMB scan positions for the remainder of the observing seasons.

5.1.2 Nominal CMB Observations

Nominal CMB observations occurred whenever weather (and the telescope) permitted. During these observations the telescope chops between two points on the sky (see Section 2.5), while the data acquisition software[1] takes the data from the polarimeter and bundles it together with various housekeeping data into a series of data frames (see Appendix F). The motion of the base is keyed to the rate at which these data frames are written to file (roughly 2 Hz). For each scan period, a certain number of frames (referred to as the dwell) are first acquired with the telescope pointed at one of the scan positions. The base is then instructed to move so the telescope points at the other scan position, where the same number of data frames are acquired. Finally, the base is instructed to return to its original orientation.[2]

The telescope moves in azimuth between the two scan positions using a crane bearing. The inner ring of the bearing is attached to the fixed part of the base. The telescope itself is attached to the outer ring of the bearing, which has gear teeth and is driven by a small gear attached to a drive motor. This motor is controlled by the computer through a Copley 306AP series amplifier. A feedback system between the motor and the Copley maintains the motor velocity at a value requested by the computer (in the form of an analog voltage). The software responsible for sending out this velocity request was different for the two observing seasons.

For the first observing season, a LABVIEW program[3] was ported directly into the C++ data acquisition program (this program therefore exists outside of the interrupt- driven timing structure of the data acquisition program, which caused some problems, described in Section 5.2). When the telescope is supposed to move to a new encoder position, this

[1] written by Jamie Hinderks

[2] Throughout the entire first observing season and the beginning of the second observing season, 10 frames were acquired at each scan position. In an effort to improve the duty cycle of the telescope, the period of the scan was increased on January 23, 2001 such that the telescope acquired 20 frames at each position. The dwell was even increased to 30 frames for a brief period between January 27, 2001 and January 31, 2001.

[3] written by Dave Strozzi

program first requests a finite velocity to overcome the inertia in the base and the motor. Once the base is moving, the requested velocity ramps up to some maximum value. When the base gets within a certain distance of the target value, the requested velocity is ramped down to a minimum velocity so the base slowly coasts towards the target position. Finally, just before reaching the target position, the program requests that the motor be briefly thrown into reverse to halt the base at the desired position.

For the second observing season, the motor control program was better incorporated into the interrupt- driven timing system. This program[4] repeatedly acquires the position of the encoder whenever it is not at the desired location. Depending on the distance between the actual encoder position and the target encoder position, the computer requests either a high velocity (if it is beyond a certain distance from the target), or a low velocity (if it is within a critical distance from the target), until the telescope arrives at the desired position.

5.1.3 Scheduled tests

A variety of tests were performed on the telescope throughout both observing seasons. Observations of Jupiter were performed approximately every week to further clarify the pointing of the telescope and to remeasure the beamwidth (see Section 6.1). The polarimetry channels were calibrated using a chopper plate roughly every two weeks (see Section 6.3). These tests were done, whenever possible, under less-than-perfect atmospheric conditions so that they would not interfere with the search for cosmological polarization.

5.1.4 Late Season Tests

The W-band system was replaced with the Q-band system midway through the second observing season, so there was no real post-season for PIQUE2001. On the other hand, during the first observing season, it became apparent in early April that the weather would only rarely be suitable for regular observations (see Figure 5.1). However, there were a few remaining tests that needed to be done before the telescope was brought off the roof, including one final calibration using the chopper plate. These tests were performed on the

[4] written by Marc Schreiber

few days in April and early May when the sky was sufficiently clear for these tests to yield sensible results.

5.2 Unplanned Situations and Events

Many things happened during the observing seasons that interrupted the search for cosmological polarization. These events had important consequences for the data, so a fairly complete list of all the bugaboos that were encountered with the operational telescope is provided here.

5.2.1 General Problems (Common to both Observing Seasons)

Inclement Weather

Inclement weather was responsible for the largest amount of lost observing time (Table 8.1). A significant fraction of this time was lost because the telescope was stowed due to the imminent threat of precipitation or the presence of thick cloudbanks over the observing site. In addition, data taken when the sky was mostly or even partly cloudy could not be used for astronomical measurements. This is because clouds not only are warm, but also contain water droplets and ice crystals which scatter microwaves in complex ways to produce polarized signals. Any time a cloud drifted through the scan region, the polarimetry channels picked up relatively large signals (visible by eye), which rendered the resulting data useless for measurements of cosmological polarization (see Chapter 8).

High Temperatures

The telescope is built to work outside in a relatively cold environment. Thus the E-Box, the HEMT power supply and the thin boxes are all provided with heaters to maintain them at a constant temperature. Unfortunately, while the W-band system was in operation there was no way to effectively cool these components down if they got too warm. This was never a problem during PIQUE2001, but around March in PIQUE2000, it occasionally became sufficiently hot outside that the heaters turned off and the temperatures of these components stopped being regulated. Without thermal regulation, the LO output power and IF amplifier gains can change significantly ($\sim 0.3\%/\mathrm{K}$), which alters the responsivity of the telescope in ways that are difficult to model or control (see Chapter 8).

Condensation

The polarimeter views the dish through a window made of $0.002''$ polypropylene backed with 0.020" Goretex™. Since the cryogenic ($\sim 50\mathrm{K}$) horn is directly behind this window, its external surface is somewhat colder than the ambient temperature. This means that moisture in the atmosphere tends to condense on the window.

During the first observing season, a small chamber in front of the window was constructed using a foam ring to support a sheet of polypropylene. Heated air flowed into this chamber through a tube connected to a fan and wrapped in heater tape. Warming the air inside the chamber in this way did help reduce condensation during the first observing season. However, the flow through the tube was very low, requiring the heater tape to be held at such a high temperature that it actually melted and fused to the copper tube it was wrapped around.

After several days of the second observing season were lost due to condensation despite the heated air system (because the heater was not cranked to such dangerous levels), that system was abandoned. Instead, a packet of Drierite™ (a material that absorbs water) was placed inside the chamber, which was then sealed. This method works surprisingly well, as it prevents condensation even in foggy conditions with relative humidities approaching 100% (In fact, condensation occasionally formed on the front of the chamber before it formed on the window). Furthermore, the Drierite pack lasted for days, even up to a week, before it needed to be replaced[5].

[5] We could tell when it was time to replace the Drierite by observing the color changes that occur as it absorbs water

Computer Hang-ups

A single C++ computer program is in charge of data acquisition and motor control . This program functioned quite well for long time periods in the lab. However, it was discovered early in the first observing season that the program would occasionally hang (stop taking data) during normal observations. Rebooting the computer and/or the DBOB usually fixed this problem, so it rarely caused significant data loss. However, constant vigilance was necessary to insure the program did not remain in this frozen state for long, especially since the program often crashed in the middle of the night.

There seemed to be multiple factors which influenced the frequency of these hang-ups, including the temperature of the computer. During the first observing season, many crashes occurred when the base was moving at the same time as the computer was closing a file. This may have been because the motor control program operated outside the timing structure of the data acquisition program (see Section 5.1.2). Indeed, there were far fewer crashes during the second observing season (less than one a day), when an interrupt-driven motor control program was being used (other alterations to the program between the two observing seasons may also have had a role in this). Another possible source of these computer problems was discovered during the second observing season, when it was found that one of the wires in the cable connecting the multipliers' thin box to the DBOB shorted out when the cable got wet. However, neither changing the code nor fixing this short completely eliminated the hang-ups

Motor Motion Problems

During much of the first observing season, the motor control program had some trouble getting the telescope to the correct encoder position. Once every 11 chops, the base overshot one of the target positions (usually the eastern position) by roughly 0.1° (see Chapter 8). The regularity of these overshoots indicates a problem in the control system. Most likely, the request to throw the motor into reverse, which brings the base to a complete stop (see Section 5.1.2), was coming late because of the aberrant timing of the motor control program.

Indeed, the interrupt driven program used during the second observing season only rarely overshot the target position (see Chapter 8).

Unfortunately, the interrupt-driven program is not as efficient at finding the target value as its predecessor. The base therefore spends more time moving between the two positions and less time taking interesting data. Furthermore, since this program did not ramp the velocities up and down (this portion of the program was not successfully implemented), the motion is jerkier than before. These facts motivated increasing the dwell time partway through the second observing season (see Section 5.1.2).

LO Failures

For a few days during both observing seasons, the radiometer behaved very oddly. Over the course of an hour, the signals from the total power channels jumped between widely different discrete values, as if some electrical contact was intermittently breaking and being restored. Eventually, the signals from the polarimeter either dropped all the way to zero (during PIQUE2001) or became so large that they suggested that the cryogenic IF amplifiers were oscillating (during PIQUE2000). After a period of a few hours in this state, the radiometer suddenly began functioning normally again. Several such failures occurred over the course of about a week, but then there were no further failures for the remainder of the observing season.

The fact that both total power channels failed at the same time indicates that these problems might be due to the LO shutting off (The cryogenic IF amplifiers were known to oscillate when the LO was not on). However, it is hard to pinpoint the true origin of this problem because it is so infuriatingly intermittent, and often the act of opening the E-Box restored the radiometer to its fully functional state. The LO was successfully accessed during a failure state only two times (once during each observing season). Both times, the power supply lines for the LO were probed. Touching the power leads to the heater or the primary oscillator had no effect and gave the correct voltages. However, when the leads to the secondary oscillator were touched, the radiometer instantly returned to its functional state (and the voltage reading was correct). This suggests that one of these leads might

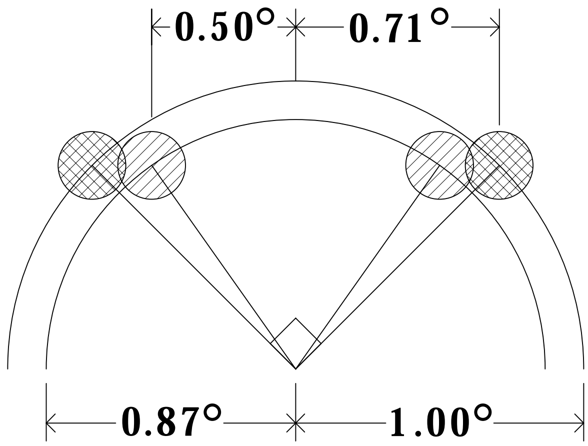

Figure 5.2: Nonstandard pointing during PIQUE2000. During the first 100 hours of data in the first observing season, the base chopped between the hatched points instead of the correct cross hatched points, the resulting data therefore apply to the wrong ring around the NCP.

have had a cold solder joint that cracks open during these failures (these leads were not resoldered because the failures stopped occurring, and the leads are not easily accessible).

Overheated Compressor

The compressor (CTI Cryogenics Model 1020/1020R) which ran the refrigerator is kept outside in a wooden box. Opening and closing hatches in the side of this box provides some control over the temperature of the compressor. Once during each observing season, the hatches were accidentally completely shut after a cold night, and not reopened during the subsequent day. This caused the compressor to overheat and trip itself off. The dewar consequently warmed up and a significant amount of time was required to cool the dewar back down to operational temperatures.

5.2.2 Problems Specific to 2000 season

Nonstandard Pointing

At the beginning of the first observing season, the telescope was incorrectly set to acan between two points 0.5° east and west of the NCP instead of 0.7°. The telescope therefore observed a ring of 0.86° radius instead of the desired 1° ring (see Figure 5.2) and measured a linear combination of Q and U instead of Q alone. Roughly 100 hours of data were taken before this error was detected and corrected for the remainder of the observing season.

Oscillations in the Cryogenic IF Amplifiers

The cryogenic IF amplifiers rarely became warm enough to oscillate during the observing season. However, the D0 signal level does sometimes appear to be somewhat higher than it ought to have been given the signal level in D1. This might be due to low level oscillations in D0, whose cryogenic IF amplifier was warmer and thus closer to the oscillation regime. Fortunately, the chopper plate data indicate that even if the amplifier were oscillating, it should not have significantly affected the performance of the polarimetry channels (see Section 6.3).

Ground loops/Glitches

Rapid shifts or glitches in all five signal levels and in various housekeeping data occasionally occurred throughout the first observing season. These glitches may have been due to some ground loop in the telescope electronics. While this ground loop was never pinpointed, far fewer glitches occurred during the second observing season, which suggests that the cryogenic IF amplifiers might have been involved.

5.2.3 Problems Specific to 2001 season

Complete Motor Control Meltdown

On January 11, 2001, for unknown reasons, the power supply for the motor had a major breakdown. Regardless of what voltage was requested from the supply, it put out its maximum voltage ($> 160V$) and drew huge amounts of current (even when nothing was connected to the supply). Fortunately, this power supply breakdown did not damage the motor itself. Unfortunately, the Copley motor controller was completely destroyed and needed to be replaced. After the Copley and the power supply were replaced, the motor power system posed no major problems.

Refrigerator Problems

The same refrigerator was used for both observing seasons (CTI Cryogenics model 1020/1020R). During the first observing season the only problem with this system was that it slowly leaked Helium, which required recharging the compressor about once a week (otherwise the temperature of the dewar would slowly climb). The system did not leak noticeably during the second observing season, but after the motor broke down and the compressor overheated, the refrigerator had more significant problems. [6] Every so often, the temperatures inside the dewar would increase radically (see Figure 5.3). The thermometers attached most directly to the cold head warmed faster than the other thermometers, and the temperature of Stage I of the cold head increased more than the temperature of Stage II. This indicates a problem in the refrigerator itself. Flushing the lines with Helium and replacing the adsorber in the compressor did not fix this problem. Furthermore, the problem was not correlated with changes in the compressor pressure.

It was eventually discovered that these warm-ups occurred almost exclusively when the wind speed was above 5 mph (see Figure 5.4). In fact, warm-ups could be induced using a large fan directed at the refrigerator lines and the cold head inside the outer ground screen. It is not known how high winds could have caused this problem or why this was not a problem during the first observing season. Nevertheless, the refrigeration system apparently needed to be shielded from the wind. Wrapping the refrigerator lines in blankets and covering the cold head with foam panels did, in fact, reduce the frequency of warm ups, but these actions did not eliminate them altogether. There was some evidence that the warm-ups were tied to the compressor temperature (sometimes the dewar seemed to cool back down when the hatches on the compressor box were shut), and the normal operating temperature of the compressor was therefore raised from 315K to 330K. After this adjustment was made, warm-ups were very rare. The impact of these warm-ups on the data from the second observing season is discussed in Chapter 8.

[6] Although the timing of these events is suspicious, we have no evidence that it is more than a coincidence. Furthermore, the Q-band system, which used a different compressor, had the same problem!

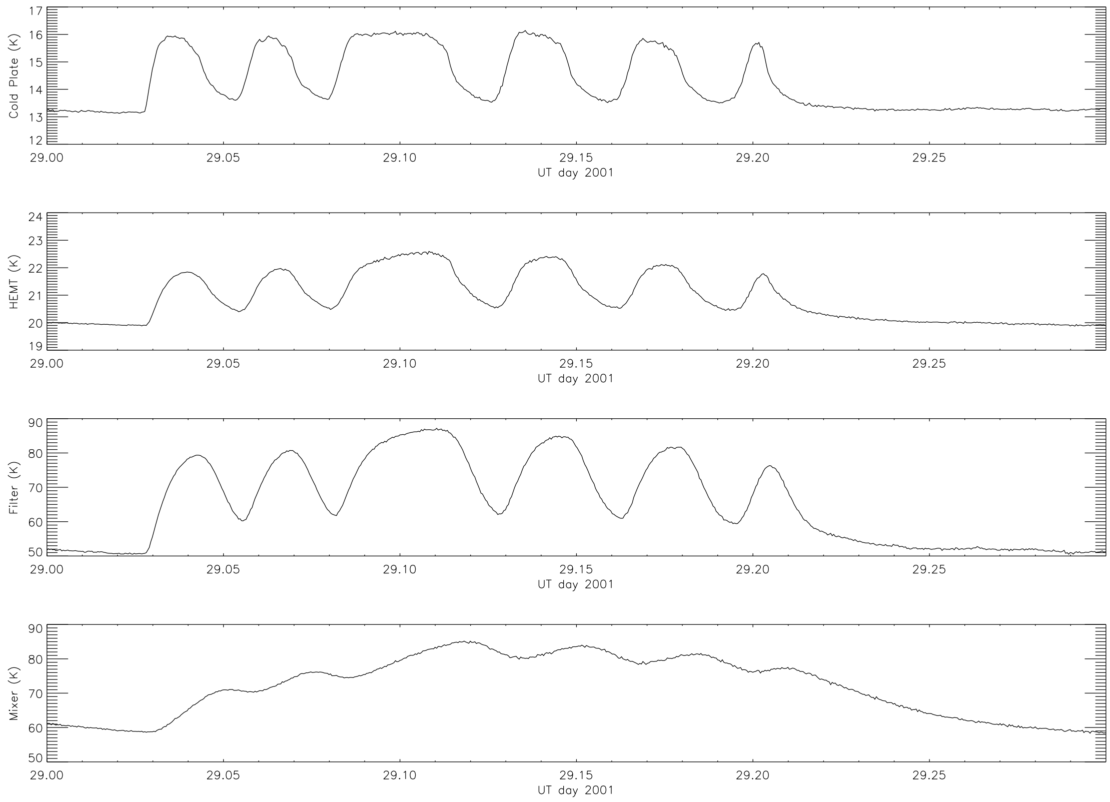

Figure 5.3: An example of the warm-ups that plagued a part of PIQUE2001. The four panels show the time series of four thermometers inside the dewar during a classic warm up. The cold plate and HEMT thermometers are attached to Stage II of the cold head, while the filter and mixer are attached to Stage I. The cold plate thermometer is attached most directly to the cold head, and the HEMT is also fairly robustly connected to the cold plate. The filter and mixer are attached to Stage I through a long post, so they are more remote from the cold head. Note that the time constants of the temperature variations reflect the strength of the connection between the thermometer and the cold head.

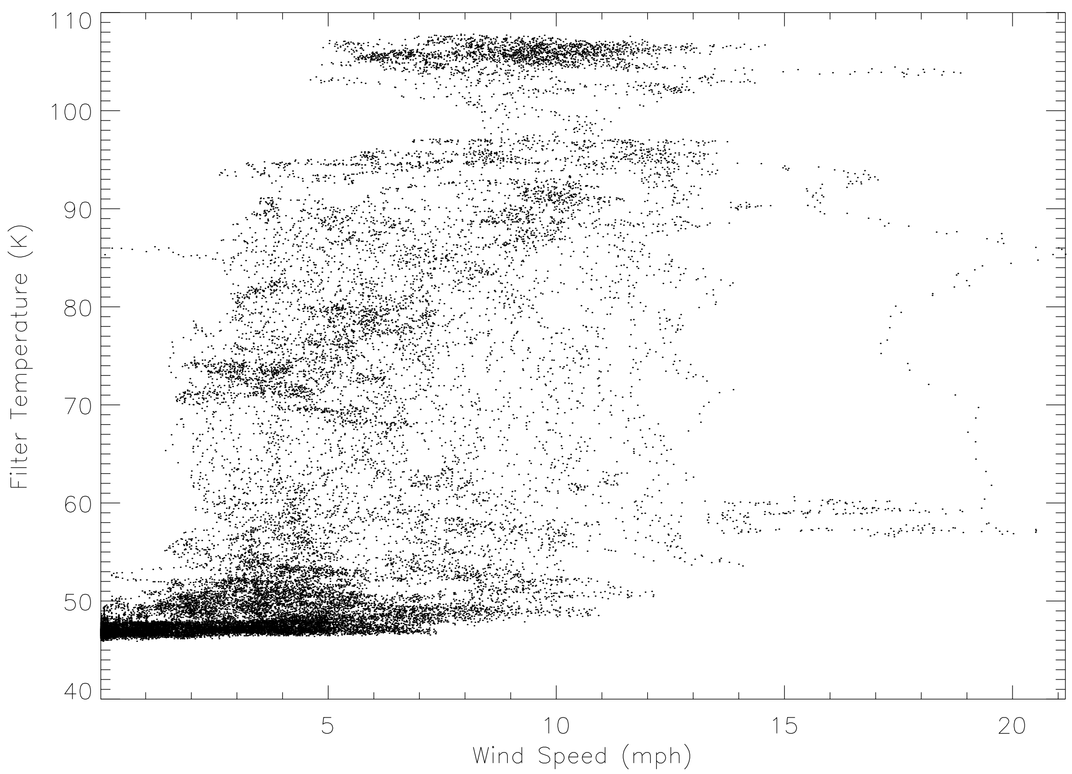

Figure 5.4: Correlation between warm-ups and high winds. Each point corresponds to the wind speed and dewar temperature of a single chop. There are two regions of high density at two discrete temperatures (roughly 50 K and 105 K) and a more diffuse cloud of points between these two regions (representing transitions between the two states). Note that the high temperature region only exists for wind speeds above 5 mph.

Chapter 6

Characterization of the Operational Telescope

As mentioned in Chapter 5, data on the important properties of the operational telescope (e.g. gain, beamshape) are obtained from tests performed during both observing seasons. Observations of the planets (described in Section 6.1) provide estimates of the beamsize and the pointing of the telescope, and observations of the moon (described in Section 6.2) determine the sign of the polarimetry channels' responsivities. The magnitudes of these responsivities, on the other hand, are derived from measurements involving a chopper plate (see Section 6.3).

6.1 Planetary Observations

Observations of the planets are used to establish the optical characteristics (e.g. beamsize) of the telescope[1]. The planets are well suited for this purpose because they are among the brightest sources of 90 GHz radiation in the sky (after the sun and the moon), and they are sufficiently compact that they can be well approximated as point sources. Unfortunately, the radiation from the planets does not have a detectable polarized component (see Section 6.2), so the beamsize and pointing must be extracted from the total power channels (Since the two total power channels have similar beamshapes (see below), the beams for the polarization

[1]based on work by J.O. Gundersen and D. Barkats

channels should also be similar).

6.1.1 Observation Methods and Data Processing

The telescope can only scan in azimuth, so the planets can only be observed when they pass through the correct altitude (which happens at most twice a day). When a planet is to be observed, the telescope rasters in azimuth by one to two degrees (varying speeds and throws were used during different observations) about the predicted location of the planet for about 20-30 minutes while the planet slowly moves through the scanned region. The data recorded during this time include the outputs from the total power channels and the encoder, which are sampled at 62.5 Hz, and the universal time, which is sampled at 0.5Hz (see Appendix F).

In order to extract the signal from the planet, drifts in the total power channels due to atmospheric fluctuations (see Chapter 7) must be removed by subtracting offsets from the time series. For each observation, the data from the total power channels are parsed into segments 180 samples (2.9 seconds) long. The mean value of the signal in each segment is computed and subtracted from the relevant samples. Samples taken when the telescope was pointed within 0.2° of the predicted location of the planet are not used in the calculation of the mean (this prevents the offset subtraction from smearing the beam).

After offset subtraction, the time series are converted into a two-dimensional map by comparing the orientation of the telescope to the expected position of the planet. For each sample, the orientation of the telescope is derived from the encoder, assuming the telescope observes at a single elevation (41.0° for PIQUE2000, 40.3° for PIQUE2001), while the position of the planet on the sky is obtained from an ephemeris program, XEPHEM.[2]. Differencing these two pairs of numbers provides the relative coordinates x and y (relative azimuth and altitude, respectively), which parametrize the orientation of the beam axis relative to the position of the planet. The (offset-subtracted) total power channel data are

[2]Comparisons with two more sophisticated ephemeris programs: the U.S. Naval Observatory Multiyear Interactive Computer Almanac (MICA) and the JPL Ephemeris (JPLEPH) demonstrate that XEPHEM gives the correct azimuth to within $10''$, but the altitude is about $1'$ higher than the other programs. The altitude offsets in Table 6.1 are therefore adjusted upwards by .017 degrees to correct for this discrepancy.

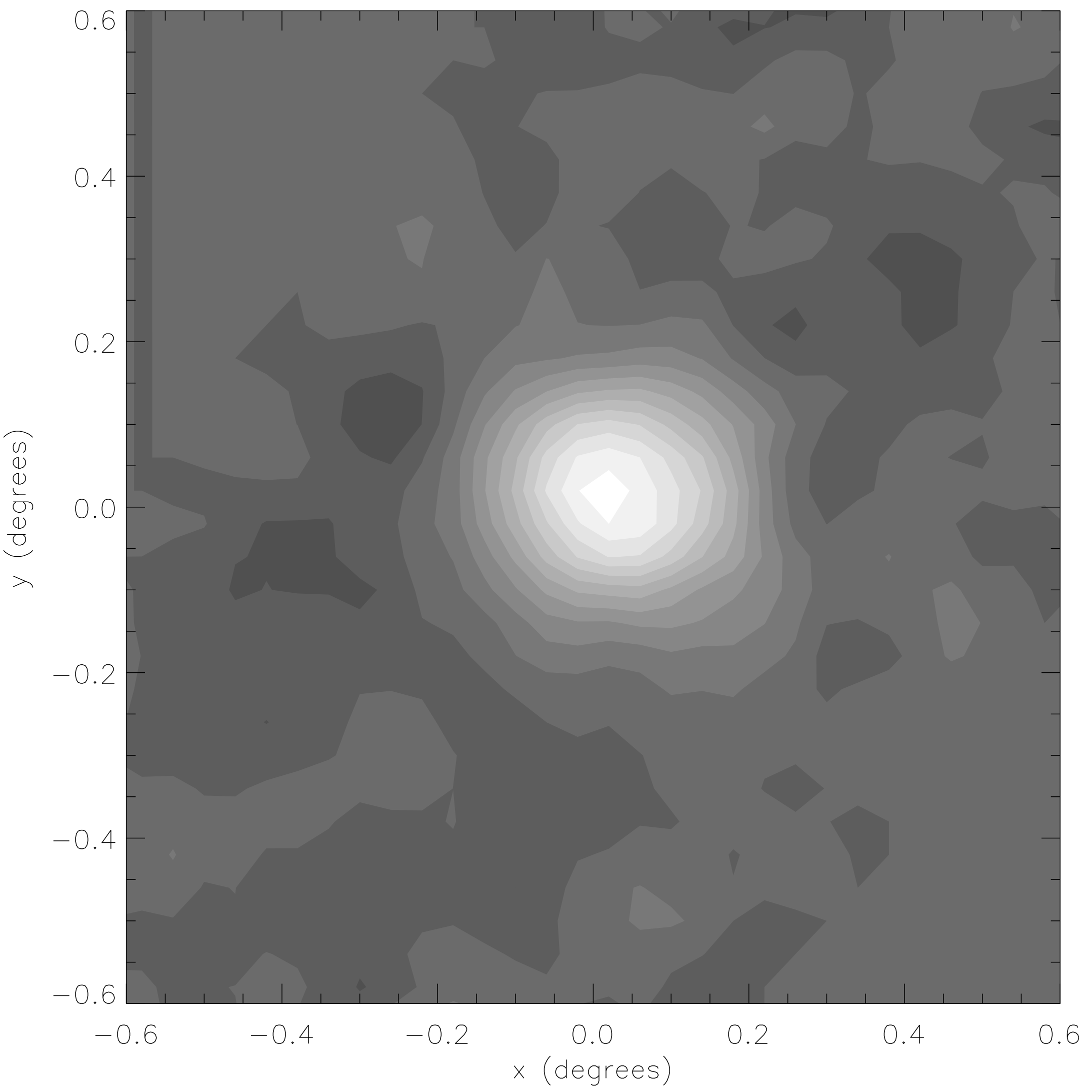

Figure 6.1: Map derived from a single Jupiter scan, taken on December 20, 2000. For presentation purposes, the map has been smoothed over adjacent pixels. The contour levels are 20 mK apart, and the peak is at 232 mK.

binned into square pixels, each 0.4° on a side (roughly 1/6 of the beam FWHM), in these relative coordinates, and the mean and error of the signal in each pixel is computed. An example of the resulting map is shown in Figure 6.1.

Finally, the data for each map is fit to a two-dimensional gaussian

$$A \exp\left(-\frac{(x-x_o)^2}{2\sigma_x^2} - \frac{(y-y_o)^2}{2\sigma_y^2}\right)$$

using the IDL routine MP2DFITPEAK (Consistent results were obtained by J.O. Gundersen using a FORTRAN program). Each successful planetary observation (i.e. each observation with a fit that converges on a reasonable solution) provides estimates of the amplitude A, the offsets x_o and y_o, and the beamwidths σ_x and σ_y.

6.1.2 Results

The fit parameters (with errors) of all of the successful planetary observations are given in Table 6.1. During the 2000 observing season, 12 successful observations of Jupiter were obtained, 3 while Jupiter was rising and 9 while it was setting. The PIQUE2001 data provided another 12 successful Jupiter observations (9 rising, 3 setting) and a single observation of Venus (setting). The data from a single total power channel (D1) are given for PIQUE2000 because the fits for the other channel often failed to converge due to excessive drifts (see Section 5.2).

Beamshape (σ_x and σ_y)

For each observing season, the various estimates for the beamwidths σ_x and σ_y are consistent with each other.[3] The individual estimates can therefore be combined to yield the estimates for σ_x and σ_y given in Table 6.2. For both observing seasons and for both channels, the estimates for σ_x and σ_y are consistent with each other and with the predictions from DADRA ($\sigma = 0.100°$).

[3] In fact, the estimates are too consistent (the reduced χ^2 is always less than one), which perhaps implies that the fitting routine overestimates the error bars.

Table 6.1: Results of gaussian fits to planetary observations

UT Day	UT Time[a]	Azimuth (deg E of N)	r/s[b]	A (cts)	σ_x (deg)	σ_y (deg)	x_o (deg)	y_o^c (deg)
PIQUE2000, Channel = D1, Planet = Jupiter								
1/5/00[d]	20:46:58	118.4	r	21.6 ± 2.7	0.084 ± 0.011	0.101 ± 0.012	0.044 ± 0.011	-0.031 ± 0.012
1/6/00	02:23:50	241.6	s	20.8 ± 2.8	0.089 ± 0.012	0.082 ± 0.010	-0.015 ± 0.012	-0.001 ± 0.011
1/6/00	20:43:08	118.3	r	18.2 ± 2.6	0.118 ± 0.015	0.103 ± 0.013	0.019 ± 0.016	-0.030 ± 0.013
1/8/00	20:35:30	118.2	r	21.6 ± 2.5	0.106 ± 0.014	0.093 ± 0.012	0.025 ± 0.013	-0.020 ± 0.012
1/12/00	02:02:25	242.0	s	18.6 ± 2.4	0.105 ± 0.014	0.117 ± 0.016	-0.001 ± 0.014	0.000 ± 0.016
2/1/00	00:55:04	243.6	s	18.9 ± 3.8	0.096 ± 0.018	0.085 ± 0.018	-0.029 ± 0.018	0.000 ± 0.017
2/9/00	00:29:38	244.4	s	18.7 ± 2.6	0.108 ± 0.015	0.116 ± 0.016	-0.010 ± 0.015	-0.010 ± 0.016
2/10/00	00:26:31	244.5	s	17.6 ± 2.9	0.132 ± 0.023	0.110 ± 0.019	-0.016 ± 0.022	-0.010 ± 0.015
2/15/00	00:07:58	245.2	s	18.7 ± 2.7	0.099 ± 0.015	0.104 ± 0.017	-0.016 ± 0.015	-0.001 ± 0.017
2/17/00	00:01:52	245.4	s	17.1 ± 3.3	0.092 ± 0.017	0.108 ± 0.022	-0.013 ± 0.018	-0.000 ± 0.021
3/6/00	23:08:28	247.6	s	16.1 ± 3.3	0.100 ± 0.021	0.105 ± 0.022	-0.022 ± 0.021	-0.017 ± 0.021
3/31/00	21:57:33	250.9	s	15.3 ± 2.6	0.105 ± 0.018	0.114 ± 0.020	-0.036 ± 0.018	-0.002 ± 0.020
PIQUE2001, Channel = D0, Planet = Jupiter								
UT Day	UT[a] Time	Azimuth (deg E of N)	r/s[b]	A (cts)	σ_x (deg)	σ_y (deg)	x_o (deg)	y_o^c (deg)
12/18/00[d,e]	23:39:12	97.6	r	41.6 ± 4.5	0.102 ± 0.012	0.099 ± 0.010	-0.017 ± 0.012	0.118 ± 0.010
12/20/00	23:30:31	97.6	r	41.7 ± 4.2	0.103 ± 0.009	0.091 ± 0.009	0.030 ± 0.009	0.030 ± 0.009
12/22/00	23:21:52	97.7	r	40.9 ± 3.1	0.101 ± 0.008	0.100 ± 0.008	0.003 ± 0.008	0.019 ± 0.007
12/28/00	22:56:01	97.8	r	35.8 ± 2.5	0.102 ± 0.007	0.099 ± 0.007	-0.006 ± 0.007	0.033 ± 0.007
1/2/01	22:34:57	97.9	r	35.4 ± 2.2	0.105 ± 0.006	0.103 ± 0.007	-0.008 ± 0.007	0.033 ± 0.007
1/9/01	22:05:45	98.0	r	36.4 ± 2.4	0.101 ± 0.007	0.101 ± 0.007	0.042 ± 0.007	0.031 ± 0.007
1/17/01	04:42:41	261.9	s	29.2 ± 2.6	0.095 ± 0.010	0.095 ± 0.0009	0.110 ± 0.010	0.023 ± 0.010
1/26/01	20:57:11	98.0	r	32.7 ± 4.6	0.095 ± 0.014	0.097 ± 0.0014	0.064 ± 0.014	0.024 ± 0.015
2/1/01	03:43:40	262.0	s	27.6 ± 3.2	0.098 ± 0.012	0.105 ± 0.0012	0.097 ± 0.012	0.016 ± 0.012
2/12/01	19:52:04	97.7	r	29.6 ± 5.1	0.089 ± 0.015	0.095 ± 0.0017	0.053 ± 0.015	0.026 ± 0.015
2/14/01	02:59:26	262.3	s	28.4 ± 4.9	0.090 ± 0.009	0.086 ± 0.0014	0.140 ± 0.015	0.038 ± 0.015
2/26/01	19:00:55	97.3	r	25.1 ± 3.2	0.103 ± 0.013	0.094 ± 0.0014	0.060 ± 0.013	0.017 ± 0.014
PIQUE2001, Channel = D0, Planet = Venus								
1/16/01	21:25:40	203.5	s	19.0 ± 3.3	0.105 ± 0.019	0.122 ± 0.026	0.057 ± 0.018	0.005 ± 0.023
PIQUE2001, Channel = D1, Planet = Jupiter								
UT Day	UT[a] Time	Azimuth (deg E of N)	r/s[b]	A (cts)	σ_x (deg)	σ_y (deg)	x_o (deg)	y_o^c (deg)
12/18/00[d,e]	23:39:12	97.6	r	29.7 ± 3.6	0.099 ± 0.012	0.092 ± 0.011	-0.009 ± 0.012	0.106 ± 0.011
12/20/00	23:30:31	97.6	r	32.0 ± 3.3	0.107 ± 0.011	0.089 ± 0.009	0.011 ± 0.011	0.013 ± 0.009
12/22/00	23:21:52	97.7	r	30.7 ± 2.9	0.107 ± 0.010	0.099 ± 0.010	-0.001 ± 0.010	0.017 ± 0.010
12/28/00	22:56:01	97.8	r	28.3 ± 3.0	0.103 ± 0.010	0.111 ± 0.012	-0.004 ± 0.011	0.029 ± 0.012
1/2/01	22:34:57	97.9	r	27.6 ± 2.9	0.098 ± 0.011	0.105 ± 0.012	-0.008 ± 0.011	0.030 ± 0.011
1/9/01	22:05:45	98.0	r	28.1 ± 3.1	0.103 ± 0.011	0.097 ± 0.011	0.043 ± 0.011	0.044 ± 0.011
1/17/01	04:42:41	261.9	s	25.8 ± 3.7	0.100 ± 0.014	0.104 ± 0.015	0.087 ± 0.014	0.048 ± 0.015
1/26/01	20:57:11	98.0	r	29.9 ± 3.4	0.101 ± 0.012	0.099 ± 0.012	0.061 ± 0.012	0.015 ± 0.011
2/1/01	03:43:40	262.0	s	24.5 ± 3.2	0.109 ± 0.015	0.109 ± 0.016	0.087 ± 0.015	0.023 ± 0.016
2/12/01	19:52:04	97.7	r	28.8 ± 4.1	0.102 ± 0.015	0.080 ± 0.012	0.046 ± 0.014	0.030 ± 0.011
2/14/01	02:59:26	262.3	s	23.4 ± 4.1	0.090 ± 0.016	0.104 ± 0.019	0.116 ± 0.015	0.027 ± 0.019
2/26/01	19:00:55	97.3	r	25.2 ± 3.3	0.108 ± 0.015	0.113 ± 0.016	0.054 ± 0.015	0.027 ± 0.016
PIQUE2001, Channel = D1, Planet = Venus								
1/16/01	21:25:40	203.5	s	18.0 ± 3.8	0.113 ± 0.022	0.104 ± 0.024	0.058 ± 0.021	0.024 ± 0.021

[a] Time when XEPHEM predicts the planet should pass through the expected observing altitude

[b] r = planet rising s = planet setting

[c] Reported offsets account for the discrepancy between the XEPHEM ephermis and the ephemeris from MICA and JPLEPH

[d] The first scan from each observing season only contains a partial map of the beam.

[e] The first Jupiter scan during PIQUE2001 was done before the base was properly leveled, so there is a significant offset in y.

Table 6.2: Summary of beamshape

Quantity	PIQUE2000	PIQUE2001	
	D1	D0	D1
σ_x	$0.098° \pm 0.004°$	$0.100° \pm 0.003°$	$0.103° \pm 0.003°$
σ_y	$0.100° \pm 0.004°$	$0.099° \pm 0.003°$	$0.098° \pm 0.003°$

Amplitude (A)

The amplitude A of the gaussian is related to the temperature of the planet. If the temperature of the planets are sufficiently well known, these parameters provide an independent estimate of the gain of the total power channels. However, there are significant systematic uncertainties in the effective temperatures of Venus and Jupiter at 90 GHz (see below), so these amplitudes can only provide a rough consistency check on the calibrations derived in Chapter 3.

Estimates of the planet's brightness temperature T_{planet} can be computed from the amplitude A using the formula:

$$T_{planet} = \frac{A}{\mathcal{R} L_{atmosphere}} \left(\frac{\Omega_A}{\Omega_{planet}} \right).$$

In this expression, $\mathcal{R}$ is the responsivity of the total power channel.[4] The parameter $L_{atmosphere}$ accounts for the loss due to the finite opacity of the atmosphere, and is computed from the atmospheric noise temperature T_{sky} (derived from the total power channel data) and surface temperature $T_{surface}$ (obtained from the National Climate Data Center, `http://lw.ncdc.gov/oa/ncdc.html`) using the equation $L_{atmosphere} = 1 - T_{sky}/(0.95 * T_{surface})$ (see Chapter 7). Finally, the ratio of the beam solid angle $\Omega_A = 2\pi\sigma_x\sigma_y$ to the solid angle subtended by the planet $\Omega_{planet} = \pi D_x D_y/4$ (D_x and D_y are the equatorial and polar angular diameters of the planet, respectively) accounts for the dilution of the signal due to the telescope's large beamwidth. For the purposes of this analysis, we will assume $\sigma_x = \sigma_y = 0.1°$, and that the equatorial radius (eccentricity) of Jupiter and Venus are 71400

[4]Only one total power channel is used in this part of the analysis. For PIQUE2000 only the data from D1 are available. For PIQUE2001 the total power channel D0 is used because it has higher signal-to-noise and better stability than D1 (although the two channels yield roughly the same value for T_J, but there may be excess scatter in the estimates from D1).

Table 6.3: Information relevant to calculating T_{planet}

UT Date	R_{planet} [a] (10^6 km)	Ω_A/Ω_{planet} [b]	T_{sky} (K)	$T_{surface}$ (° C)	$(L_{atmosphere})^{-1}$
PIQUE2000					
1/5/00(r)	701.5	624.7	33.1	3.3	1.14
1/6/00(r)	702.1	625.7	33.4	-1.2	1.15
1/6/00(s)	703.9	628.9	42.5	6.7	1.19
1/8/00(r)	708.6	637.4	41.7	3.9	1.18
1/12/00(s)	716.3	651.5	46.6	4.4	1.21
2/1/00(s)	763.9	740.7	45.1	-1.1	1.21
2/9/00(s)	782.3	776.9	38.5	3.3	1.18
2/10/00(s)	784.5	779.1	58.3	4.1	1.29
2/15/00(s)	795.7	807.9	43.4	1.1	1.20
2/17/00(s)	802.0	816.2	46.0	-1.1	1.22
3/6/00(s)	837.2	889.7	53.6	11.1	1.24
3/31/00(s)	874.0	969.9	55.1	13.9	1.25
PIQUE2001					
UT Date	R_{planet} [a] (10^6 km)	Ω_A/Ω_{planet} [b]	T_{sky} (K)	$T_{surface}$ (° C)	$(L_{atmosphere})^{-1}$
12/18/00(r)	617.4	483.8	43.8	-2.1	1.20
12/20/00(r)	619.5	487.1	38.6	-4.4	1.18
12/22/00(r)	621.8	490.9	34.1	-7.0	1.16
12/28/00(r)	629.6	503.3	40.4	-6.1	1.19
1/2/01(r)	637.2	515.3	34.8	-5.8	1.16
1/9/01(r)	649.1	534.7	48.4	1.7	1.17
1/17/01(s)	663.1	558.1	50.5	1.1	1.29
1/26/01(r)	683.6	593.3	36.4	0.6	1.16
2/1/01(s)	695.5	614.2	45.8	2.8	1.21
2/12/01(r)	723.0	663.6	38.3	1.7	1.17
2/14/01(s)	726.1	669.2	42.0	2.2	1.19
2/26/01(r)	756.8	727.1	38.7	7.8	1.17
UT Date	R_{planet} [a] (10^6 km)	Ω_A/Ω_{planet} [b]	T_{sky} (K)	$T_{surface}$ (°C)	$(L_{atmosphere})^{-1}$
2/16/01(s)	103.1	1768.9	46.2	5.0	1.21

[a] Distance to the planet, provided by the JPL Solar System simulator

[b] Assuming $\sigma_x = \sigma_y = 0.100°$

Table 6.4: Estimates of the planetary brightness temperatures

PIQUE2000 (D1)		PIQUE2001 (D0)	
UT Date	T_{planet} [a] (K)	UT Date	T_{planet} [b] (K)
Jupiter			
1/5/00(r)	129 ± 16	12/18/00(r)	164 ± 18
1/6/00(s)	125 ± 17	12/20/00(r)	154 ± 16
1/6/00(r)	114 ± 16	12/22/00(r)	155 ± 12
1/8/00(r)	137 ± 16	12/28/00(r)	143 ± 10
1/12/00(s)	122 ± 16	1/2/01(r)	141 ± 9
2/1/00(s)	141 ± 29	1/9/01(r)	152 ± 10
2/9/00(s)	135 ± 19	1/17/01(s)	135 ± 12
2/10/00(s)	147 ± 24	1/16/00(r)	151 ± 21
2/15/00(s)	151 ± 21	2/1/01(s)	137 ± 16
2/17/00(s)	142 ± 27	2/12/0(r)	153 ± 27
3/6/00(s)	149 ± 30	2/14/01(s)	150 ± 26
3/31/00(s)	155 ± 26	2/26/01(r)	142 ± 18
Venus			
		2/16/01(s)	289 ± 51

[a] $\mathcal{R} = 0.12$cts/mK assumed
[b] $\mathcal{R} = 0.15$cts/mK assumed

km (0.065) and 6100 km (0) respectively. Table 6.3 provides the values of $L_{atmospheric}$ and Ω_A/Ω_{planet} for each observation, along with the data used to calculate these parameters.

The estimates for the brightness temperatures of the planets are given in Table 6.4. Note that the first five estimates of Jupiter's temperature from PIQUE2000 are noticeably lower than the other estimates from that year. These scans were performed before the beginning of the regular observing season, when the radiometer was not yet thermally regulated. It is therefore not surprising that the system gain was slightly different during these observations.

The estimates of Jupiter's temperature from the PIQUE2000 regular observing season are all consistent with each other. All of the estimates from PIQUE2001 are likewise consistent amongst themselves. However, if we assume the responsivities derived in Section 3.3, the PIQUE2000 estimates for $T_{Jupiter}$ are 6% higher than the estimates from the PIQUE2001 data. A similar 6% discrepancy appears between the expected and observed response to a 77 K load (see Table 3.10). Both these discrepancies are probably the result of excess loss in

Table 6.5: Published estimates for the brightness temperature of Jupiter around 90 GHz

Estimate	Source
179 ± 5 K	Ulich et al. (1980)
183 ± 2 K	Epstein et al. (1980)
172.5 ± 1.4 K	Griffin et al. (1986)
144 K	Ulich et al. (1980), recalibrated[a]
149 K	Epstein et al. (1980), recalibrated[a]
147 ± 4 K	This work

[a]The recalibrated numbers are adjusted so the solar temperatures obtained by Ulich et al. (1980) and assumed by Epstein et al. (1980) agree with the solar model proposed by Fedoseev and Chernyshev (1998).

the PIQUE2000 two load injector system. In order to make the observing seasons' estimates of $T_{Jupiter}$ consistent, the loss in the injector during PIQUE2000 must be 6% higher than the value given in Section 3.3.2.[5] Adjusting the loss correction factors appropriately yields $\mathcal{R} = 0.12$ cts/mK and $T_{rec} = 62.5$ K for D1 during PIQUE2000. The planetary brightness temperature estimates given in Table 6.4 use this corrected responsivity.

The combined estimates for the brightness temperatures of Venus and Jupiter are 289 ± 51 K and 147 ± 4 K, respectively. While the estimate of T_{Venus} has large error bars, the estimate of $T_{Jupiter}$ is sufficiently precise that it can be compared with other measurements (see Table 6.5). While the published data suggest that $T_{Jupiter}$ is significantly higher than our estimate, there are significant systematic uncertainties in the published numbers.[6] The estimate from Epstein et al. (1980) is calibrated using the measured value of the sun's temperature at 90 GHz given in Ulich et al. (1980), so both Epstein's and Ulich's estimates for $T_{Jupiter}$ rely on Ulich's absolute calibration. This calibration may be incorrect because Ulich et al. (1980) say the temperature of the sun at 90 GHz is 7900K, while more recent measurements suggest that it is more like 6400 K (Fedoseev and Chernyshev, 1998). If Ulich's and Epstein's data are recalibrated to make them consistent with this new number, then their estimates of $T_{Jupiter}$ are consistent with our measured value (see Table 6.5). This

[5]This correction is not applied in Section 3.3.4 because the horn test alone does not provide enough information to determine if the problem is in the estimate of the responsivity or the receiver noise temperature.

[6]The data from Griffin et al. (1986) cannot be considered here because they do not provide sufficient detail about their 90 GHz measurement.

Table 6.6: Summary of Pointing

Parameter	PIQUE2000 D1	PIQUE2001 D0	PIQUE2001 D1
y_o (combined)	$0.012° \pm 0.004°$	$0.028° \pm 0.003°$	$0.026° \pm 0.003°$
x_o (Jupiter rising)	$0.032° \pm 0.007°$	$0.002° \pm 0.004°$ [a] $0.049° \pm 0.005°$ [b]	$0.001° \pm 0.005°$ [a] $0.051° \pm 0.006°$ [b]
x_o (Jupiter setting)	$-0.013° \pm 0.005°$	$0.112° \pm 0.007°$	$0.096° \pm 0.009°$
x_o (Venus setting)		$0.057° \pm 0.018°$	$0.058° \pm 0.021°$
x_o (combined estimate)	$0.010° \pm 0.022°$	$0.034° \pm 0.032°$ [a] $0.081° \pm 0.032°$ [b]	$0.023° \pm 0.022°$ [a] $0.073° \pm 0.022°$ [b]

[a] Prior to January 5 2001, see text
[b] After January 5 2001, see text

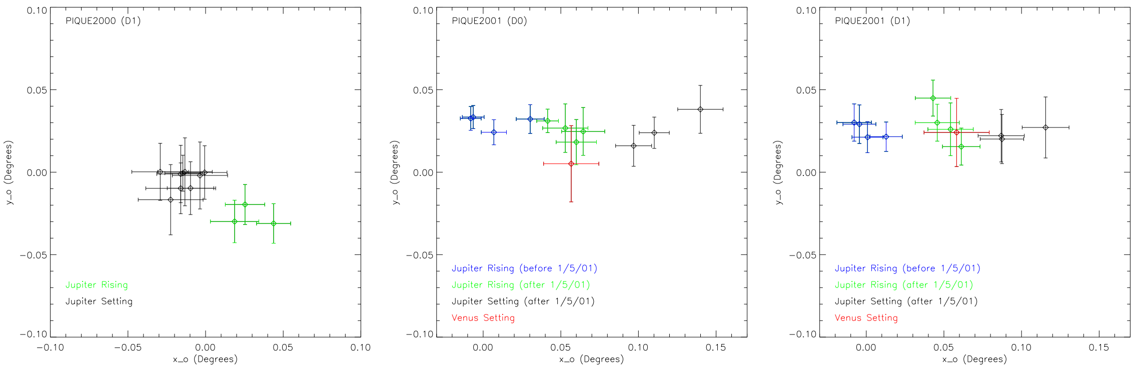

Figure 6.2: The pointing offsets derived from the Jupiter Observations

does not provide definitive evidence that either this recalibration of the published data or the calibration of the PIQUE data is correct. It merely demonstrates that present measurements of Jupiter cannot strongly constrain the calibration of the total power channels.

Pointing (x_o and y_o)

The offset parameters x_o and y_o establish the pointing of the telescope. For both years, the estimates for the altitude offset y_o are consistent and indicate the telescope is observing within 0.03° of the correct elevation (see Table 6.6). On the other hand, the scatter in the azimuth offset x_o is much larger than expected given the error bars (see Table 6.6 and Figure 6.2). For both observing seasons, the observations of Jupiter rising cluster at a different azimuth from the observations of Jupiter setting. Both the rising and setting clusters in PIQUE2000 and the setting cluster in PIQUE2001 have scatters consistent with their error bars, but there is excess scatter in the rising cluster for PIQUE2001.

The excess scatter in the rising cluster of PIQUE2001 occurs because the first four Jupiter scans are significantly to the left of the remaining scans (note all the setting observations occurred late in the season). Between the fourth and fifth Jupiter scans, on January 5, 2001, the replacement DBOB cable was threaded under the base (see Chapter 5). Apparently this operation caused the encoder to slip by about $0.05°$. This problem was not discovered until after the end of the observing season, so the assumed location of north did not change to account for this shift. The data taken after January 5 therefore refers to a slightly different point on the sky than the data taken prior to this date. The separation between the two points is only about one-fifth of a beam FWHM (the overlap integral is about 0.94), so the change can be safely neglected.

The origin of the discrepancy between the rising and setting observations is more of a mystery, especially since the setting points cluster to the left during PIQUE2000 and to the right during PIQUE2001. A $\sim 0.05°$ tilt in the base could account for the spread in the data, but measurements from both observing seasons show the base was level to better than $0.01°$ (see Chapter 5). A more likely explanation is that the shaft connecting the encoder to the base was cocked at a small ($\sim 0.05°$) angle, such that the encoder reading differed from the true azimuth position by a small azimuth-dependent factor. This model naturally explains why only the parameter x_o is affected, but unfortunately the available data cannot be used to determine if this explanation is correct. The pointing estimates come from only two or three narrow ranges in azimuth, which does not provide enough data to either determine if the offset follows the expected sinusoidal form, or constrain the size of the offset at the NCP. For want of a better solution, the average of the offsets estimated from the rising and setting data will be used as the final estimate for x_o, and the separation between the two estimates will be used to estimate the error[7]. These combined estimates are given in Table 6.6. For PIQUE2000 the large uncertainty in x_o is not a problem because

[7]The value of x_o for PIQUE2001 prior to January 5 is calculated by comparing the rising data before and after this date

the entire range of likely offsets is small (less than 0.04°). For PIQUE2001, the offset could be as large as the beamwidth, and some calculations are needed to determine if an offset this large has a significant impact on the data analysis (see Chapter 9)

6.2 Observations of the moon

The moon is the only astronomical source of detectable polarized 90 GHz radiation observed by the PIQUE telescope (we wisely chose never to look directly at the sun). Maps of the moon, shown in Figure 6.3, are generated using the procedures described in the previous section. The moon is just barely resolved in these images, which (for the most part) show a distinct quadrupole pattern. This pattern is consistent with the expected polarized signal from the moon. The thermal radiation from within the moon (or any terrestrial planet) acquires a polarized component when it is refracted at the moon's surface[8], the resulting polarization is radial to the center of the moon and the polarized component is largest near the edge of the moon (Davies and Gardner, 1966; Mitchell and de Pater, 1994; Chapman, 1986). This radial pattern reduces to the observed quadrupole pattern when a single Stokes parameter is measured with the coarse resolution of the PIQUE telescope.

The strength of the polarized radiation is a complex function of phase angle and surface emissivity (Gary et al., 1965), so these data cannot easily be used to calibrate the polarimetry channels. However, the mere fact that PIQUE observes a polarized signal from an astronomical source demonstrates conclusively that the telescope is a functional polarimeter. Furthermore, these measurements nicely determine the sign of the responsivities of the polarimetry channels, in that a polarized signal oriented along the axis rotated 45° clockwise from vertical would produce a negative response from the PIQUE2000 polarimeter and a positive response from the PIQUE2001 polarimeter.

It is also worth noting that in many of the images in Figure 6.3 (particularly PIQUE2001 s2), one pair of lobes is much brighter than the other pair. This is not expected based on the models of the lunar surface. Instead, this asymmetry is an artifact of the telescope and

[8]A similar process generates the polarized signal from the chopper plate, see Section 6.3

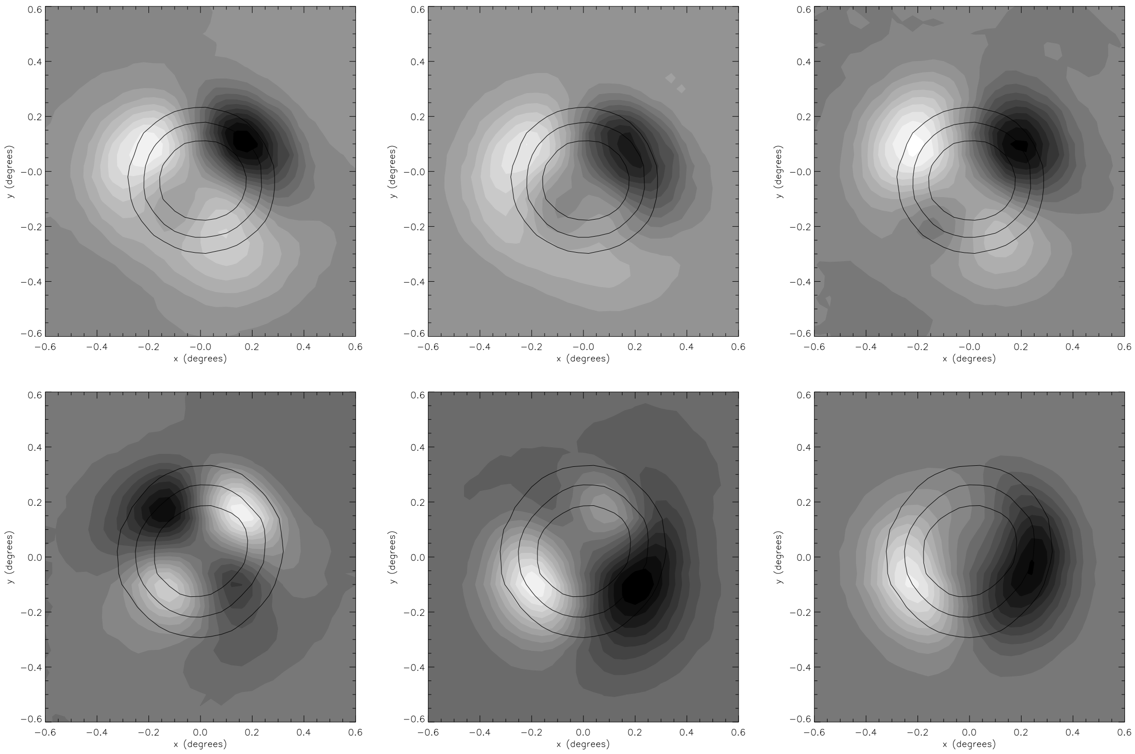

Figure 6.3: The moon as observed on January 12, 2000 (top) and February 28, 2001 (bottom) by the polarimetry channels s0 (left), s1 (middle) and s2 (right). The 2001 maps have been offset by 0.1° to center them on the field. The shading indicates the polarized signal and the greyscale is different for the different channels. The lines are contours from a total power channel's moon map. The quadrupole pattern (clearest in the maps from s0) is due to polarized emission from the moon, while the dipole pattern (strongest in PIQUE2001 s2) is due to common-mode pick up. Note the sign of the signal changes between the two years, which implies that the sign of the responsivity of these channels changed between the two years (as expected due to changes in the dewar configuration).

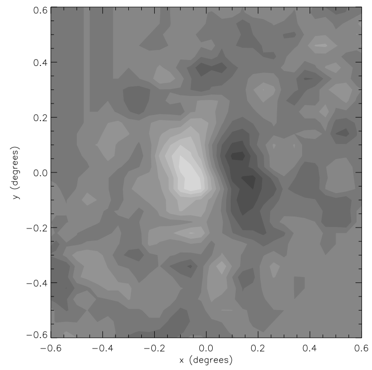

Figure 6.4: Map of the polarimetry channel (s1) data for a Jupiter observation on December 20, 2000. The data have been smoothed over two bins for clarity. The contours are at 1 mK levels. The maximum is at 6.7 mK and the minumum is at -4.5 mK. The signal is probably not due to real polarized signals from the planet. Instead, it most likely represents the cross-polar response of the telescope.

is probably related to some strange results from the observations of Jupiter. As noted in the previous section, the radiation from the planets should not be significantly polarized, and indeed the data from the polarimetry channels do not show any evidence for a point source of polarized radiation. However, in certain planetary observations (those with the highest signal-to-noise in the total power channels) the polarimetry channels do respond to the planet (see Figure 6.4). The origin of the observed two-lobe pattern is not fully understood, but it apparently represents the response of the polarimetry channels to an unpolarized point source. If this pattern is convolved with the lunar disk, the resulting pattern will have two lobes. Combining this dipole pattern with the quadrupole pattern from the true polarized signals can reasonably account for the features in Figure 6.3.

6.3 Calibration with a chopping plate

Without a sufficiently bright, well characterized, astrophysical source of polarized 90 GHz radiation, the responsivities of the polarization channels must be measured using a local calibrator. This calibrator is a large Hexcell^{TM} Aluminum 6061 plate placed directly in the main beam of the telescope (Figure 6.5). Both the radiation from the sky reflected off the plate and the thermal emission from the plate itself have a small polarized component due to the finite conductivity of aluminum. This polarized component is orientation dependent, so when the plate is driven back and forth by a small motor, it injects a small, calculable, time-variable polarized signal into the telescope.

The polarization fraction of (initially unpolarized) radiation reflected off a metal plate can be calculated from first principles using the Maxwell Equations and the low-frequency dispersion equation for metals (see, e.g. Wollack,1994; Cortiglioni, 1994; Strozzi and McDonald, 2000). The polarized component of the thermal emission can then be determined by applying Kirchoff's Law (reflectivity = absorbtivity; Strozzi and McDonald 2000). Finally, the appropriate coordinate transformations yield the polarized signal entering the telescope as a function of the orientation of the plate[9]:

$$q(\alpha) = \kappa \frac{\alpha}{\alpha_o}(T_{plate} - T_{sky}),$$

where T_{plate} is the physical temperature of the plate and T_{sky} is the atmospheric noise temperature. This expression assumes the plate is vertical and the rotation angle α, measured with respect to the vertical plane containing the main beam axis, is small. With the motor running, $\alpha = \alpha_o \cos(2\pi f_n t)$, where the frequency of nutation $f_n \sim 1Hz$ and the amplitude $\alpha_o \sim 5^\circ$ (different values were used in different runs, see Table 6.7). The factors dependent on the composition, location and geometry of the plate are all included in the dimensionless parameter κ:

$$\kappa = 4\pi \frac{\delta}{\lambda} \gamma \frac{\sec\beta - \cos\beta}{\sin\beta} \alpha_o,$$

where λ is the wavelength of the radiation, $\delta = 1/2\pi\sqrt{\lambda/c\mu\sigma_o}$ is the skin depth of the plate (μ and σ_o are the (cgs) permittivity and DC conductivity of aluminum, respectively); β is the elevation angle observed by the telescope ($\beta = 41^\circ$ for PIQUE2000 and $\beta = 40.3^\circ$ for PIQUE2001), and γ is a geometrical factor which depends upon the size of the plate (for a small, uniformly illuminated plate, $\gamma = \sec(\beta) \simeq 1.3$, while for an infinitely large plate, $\gamma = 1$).

[9]These very painful calculations were done by Herzog (1999)

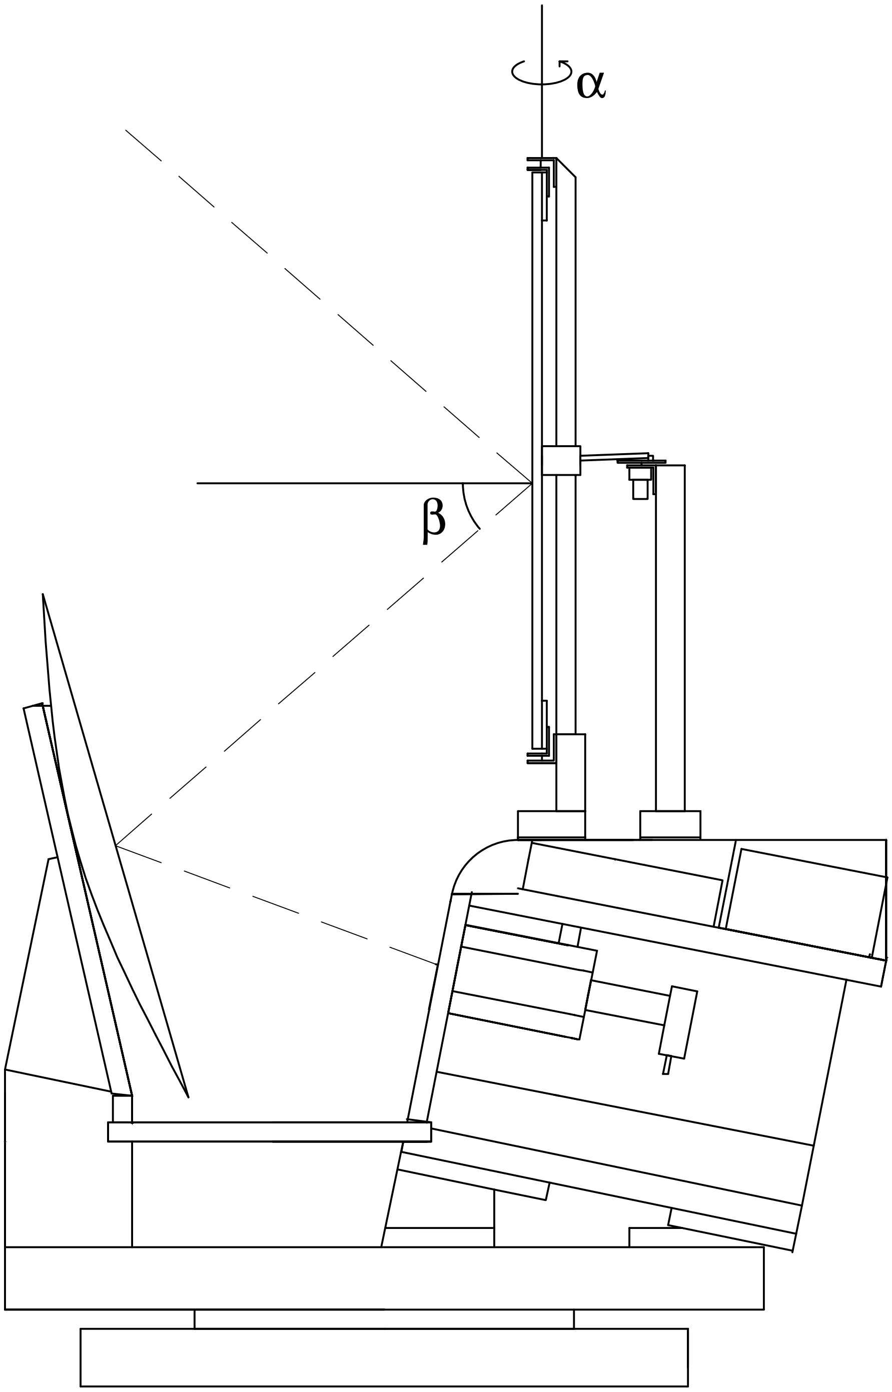

Figure 6.5: The chopper-plate used to calibrate PIQUE. The plate nutates about a vertical axis to produce a time-variable polarized signal.

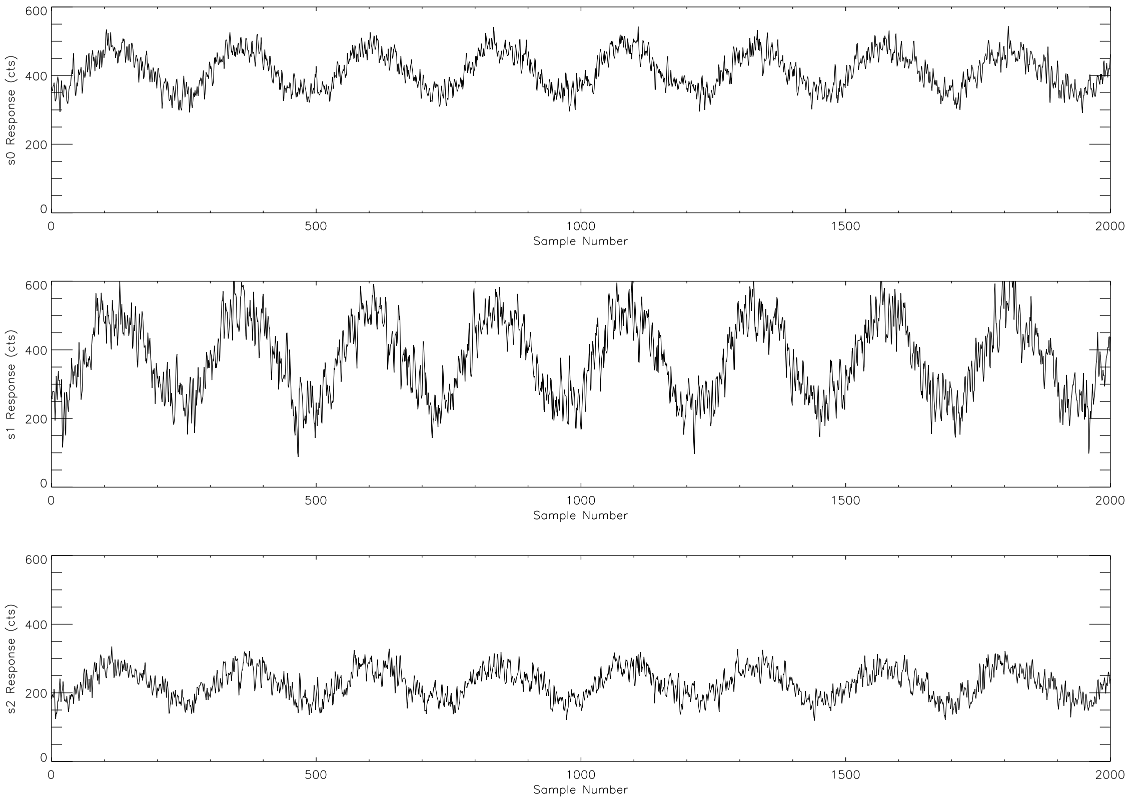

Figure 6.6: Representative sample of the signals from the multiplier channels during a chopper plate calibration. The sampling rate is 62.5 Hz

The parameter κ is calculated using the following data: The angles β and α_o are determined from the mechanics of the telescope and the chopping plate (see Figure 6.5). The geometrical factor γ is assumed to be one (this is a good approximation because only 3% of the beam spills over the edges of the plate). The effective wavelength (λ) for each band is obtained from the results of the narrow band test (see Table 3.2). The skin depth is computed using the appropriate parameters for Aluminum 6061, $\mu = 1$ and $\sigma_o^{-1} = 4.0\mu\Omega \cdot cm$ (Sharp, 1993). For $\alpha_o = 5.6°$, this gives $\kappa = 1.06 * 10^{-4}$, $1.09 * 10^{-4}$ and $1.11 * 10^{-4}$ for s0,s1 and s2, respectively (these values are not significantly different for the two observing seasons).

The temperatures T_{plate} and T_{sky} vary significantly with time, so they must be continuously monitored during the calibration runs. A thermometer attached to the plate measures

Table 6.7: The results of individual chopper plate tests

Date	α_o	$\mathcal{R}_p$ (cts/mK)		
		s0	s1	s2
PIQUE2000				
2/10/00	4.7°	2.39	4.49	1.40
2/24/00	4.7°	2.41	4.47	1.43
3/14/00	9.4°	2.48	4.58	1.47
3/23/00	5.6°	2.33	4.41	1.47
4/13/00	5.6°	2.46	4.73	1.57
PIQUE2001				
12/28/00	5.6°	1.95	6.58	1.19
1/24/01	5.6°	2.05	7.24	1.28
2/13/01	5.6°	1.96	7.08	1.23
2/24/01	5.6°	1.95	7.03	1.15
2/26/01[a]	5.6°	1.88	6.72	1.14

[a]evaluated at Dewar Temperature=50K

T_{plate}, while one of the total power channels (D1 for PIQUE2000, D0 for PIQUE2001) is used to determine T_{sky} ($T_{sky} = T_{sys} - T_{rec}$, where T_{sys} is the total system temperature and T_{rec} is the receiver noise temperature, see Chapter 7).[10]

The amplitude of the sinusoidal variation in the polarized signal is roughly 30 mK, and the response of the polarimetry channels to this signal is clearly visible in the time series (see Figure 6.6). The response of the polarimetry channel is quantified by parsing each time stream from each run into segments 3600 samples (112.5 frames, roughly one minute) long, and fitting the data from each segment to a sine curve (the amplitude, frequency, phase and offset all being allowed to float). The amplitudes derived from these fits are divided by the appropriate values of $\kappa(T_{plate} - T_{sky})$ so that each segment provides a separate responsivity estimate[11]. These estimates are combined to yield a single measurement of the responsivity of each polarimetry channel. Finally, the appropriate finite-temperature correction ($k(\nu) =$ 0.83, 0.81 and 0.79 for s0, s1 and s2 respectively, see Section 3.3), is applied to these estimates.

[10]The 6% correction to PIQUE2000 D1 described in Section 6.1 is included in these calculations

[11]The responsivities are estimated for each segment because the polarized signal can drift significantly over the course of the run as the temperature of the atmosphere and the plate change.

Table 6.8: Comparison of the responsivities of multiplier channels

Year	Test	$\mathcal{R}_p$ (cts/mK)		
		s0	s1	s2
2000	Chopper Plate[a,c]	2.41 ± 0.06	4.53 ± 0.12	1.47 ± 0.06
2000	Two-Load (Post-Season)[b,d]	2.0	3.6	1.2
2001	Two-Load(Pre-Season)[d]	2.1	6.0	1.2
2001	Chopper Plate[c]	1.96 ± 0.06	6.93 ± 0.27	1.20 ± 0.06
2001	Two-Load(Post-Season)[d]	1.8	6.4	1.1

[a]The 6% correction in the gain of the total power channel D1 discussed in Section 6.1 is included in these estimates.

[b] These estimates do not include a correction for the excess loss in the injector system which is implied by the horn test (Section 3.3.4) and the Jupiter data (Section 6.1)

[c]Errors on the Chopper Plate test results are based on the standard deviations of the individual estimates from Table 6.7.

[d]The uncertainty in the two-load test results is dominated by the 7% uncertainty in the loss correction factor (see Section 3.3.2).

The responsivity estimates from the various chopper plate tests are tabulated in Table 6.7. The statistical error of each estimate is calculated based on the distribution of the estimates from the individual segments. These errors are less than 1% and are considerably smaller than the scatter in the results from the different tests. This excess scatter is attributed to a lack of repeatability in the chopper plate test. The chopper plate is a rather unwieldy piece of metal, and it was difficult to mount it in the exact same place and with the exact same orientation every time. Small differences in the location and orientation of the plate result in small changes in the measured responsivity. The excess scatter therefore provides an estimate of the repeatability of this test, and since the different runs do agree with each other to within about 10%, the repeatability of this test is reasonably good. The final estimates of the responsivities for each observing season are calculated by averaging the estimates from the individual runs, and the relevant errors are derived from the scatter of these estimates (see Table 6.8).

The largest systematic uncertainty in these measurements comes from the uncertainty in the conductivity of the chopper plate. The conductivity of commercial Aluminum alloys depends on the concentration of other metals in the alloy. For Aluminum 6061, roughly

1% of the alloy is Magnesium (`http://www.luskmetals.com/chemalum.html`), which has a significant ($\sim 30\%$) impact on the assumed conductivity (Sharp, 1993). Uncertainties in the composition of the plate can therefore lead to significant uncertainties in the conductivity. The emissivity of the chopper plate at 90 GHz has not been measured directly. However, relative measurements of the 90 GHz emissivity of other samples of Aluminum 6061 (Butler, 1998) suggest that the assumed resistivity is approximately correct. A conservative estimate of the systematic uncertainty due to the resistivity of the chopper plate is about 10%.

These estimates can be compared with the responsivities derived from the two-load tests (see Section 3.3, Table 6.8). For PIQUE2001, the two methods agree to within the $\sim$10% systematic uncertainties of the two tests, with the responsivities derived from the chopper plate tests being somewhat higher than the results from the two-load tests. For PIQUE2000, the responsivities derived from the two-load test are significantly ($\sim 20\%$) lower than the estimates from the chopper plate test. This can be attributed to the excess loss in the two load injector system discussed in Sections 3.3.4 and 6.1. The chopper plate tests provide a more direct measurement of the operational telescope gain, and these estimates yield a reasonable estimate of the receiver noise temperature (see Table 3.8), so the responsivities derived from the chopper plate test are used in the data analysis (see Chapter 8).

In addition to establishing the mean responsivity of the polarimetry channels, the chopper plate tests also indicate how variations in the temperature of components inside the dewar affect the gain of the polarimeter. During both observing seasons, there were significant problems with the temperature stability of these cryogenic components (see Chapter 5). For the first observing season, small fluctuations in the dewar temperatures may have occasionally caused one cryogenic IF amplifier to oscillate, while during the second observing season, there were large, rapid changes in the dewar temperatures. The last chopper plate test of each observing seasons was used to directly measure any changes in the responsivities of the radiometer that might have resulted from these temperature variations.

The last chopper plate test of the first observing season began with the dewar in its normal state for 45 minutes. After this baseline data was acquired, a voltage was applied to heater resistors attached to the IF amplifiers. The amplifiers warmed by a few Kelvin

Table 6.9: Changes in responsivity due to changes in dewar temperature[a]

Year	s0 (%/K)	s1 (%/K)	s2 (%/K)
2000[b]	$+0.36 \pm 0.11$	$+0.25 \pm 0.07$	$+0.20 \pm 0.11$
2001[c]	-0.01 ± 0.03	-0.13 ± 0.03	-0.36 ± 0.03

[a]The temperature used to calculate the above slopes are measured with one of the thermometers attached to Stage I of the dewar (see Appendix F).
[b]The errors for PIQUE2000 are derived from the linear fit; the 1% systematic uncertainty in the responsivity of the total power channel is not included.
[c]The errors for PIQUE2001 are the systematic errors based on the difference in the estimated variability when D0 and D1 were used to determine T_{sky}.

and then one of the amplifiers (D0) began to oscillate (indicated by a rapid increase in the signal measured by that channel). The amplifier was allowed to oscillate only briefly (a few minutes) before the heat was turned off. The amplifier was then allowed to cool down to approximately its starting temperature before the voltage was turned on again. The amplifiers were heated a total of three times, and each time the amplifier was allowed to oscillate for a little longer.

For the corresponding chopper plate run during the second observing season, the blankets and pads that protected the helium lines from the wind were removed and a large fan was pointed inside the outer ground screen. At the same time, the compressor hatches were thrown wide open and the compressor was allowed to cool below 310K. This initiated one of the warm-ups characteristic of the season. After the dewar had been warming for about an hour, the fan was turned off and the dewar cooled down to normal. Unfortunately, just before the dewar had reached its peak temperature, it was discovered that the rear panel of the inner ground screen was left up, blocking the reflected beam from the chopper plate, so the baseline data prior to the warm-up cannot be used.

Since the total power channels are used to determine the value of T_{sky}, variations in the gain of the total power channels must be calibrated out before the variations in the multiplier channels can be measured. For PIQUE2000, the episodes of heating are sufficiently short that the gain variation can be estimated by assuming the atmospheric noise temperature is constant (for D1, this variation is $-0.73\%/\mathrm{K}$). The warming episode for PIQUE2001, on the other hand, was too long for this method to be practical. However, it was clear that the two channels responded differently to the changing dewar temperatures. As the dewar approached normal operating temperatures, the signal from one channel (D1) fell while the signal from the other channel (D0) rose. The two estimates of T_{sky} provided by the two channels should therefore bracket the true value.

The variations in the gain are quantified by obtaining the slope of a linear fit to the responsivity as a function of dewar temperature (measured using one of the thermometers attached to Stage I of the refrigerator, see Table 6.3). The gain variations are on the order of a few tenths of a percent per degree Kelvin. For the first observing season, the variations are all positive, so that the gain increases as the dewar temperature increases. This is counterintuitive, and suggests the correction applied to the total power responsivity is inaccurate. In any case, the temperature of the dewar varied by less than 10 K during this season, so the responsivity only changed by a percent or two (at most) over the course of the observing season. On the other hand, during the second observing season the dewar temperatures changed by as much as 50K (see Figure 5.3), so the responsivity of the polarimetry channels changes significantly (by as much as 20%) during a warm-up. These gain variations therefore must be taken into account if the data taken during warm ups are to be useful (see Chapter 8).

Chapter 7

Atmospheric Characteristics at the PIQUE Site (Princeton, NJ)

The atmosphere at Princeton, NJ is not completely featureless and transparent, even during clear weather. This chapter briefly describes the atmospheric characteristics of the observing site, as measured by the PIQUE telescope.

7.1 Atmospheric Noise Temperature

The nonzero atmospheric noise temperature T_{sky} around 90 GHz at Princeton is primarily due to water vapor (Grossman, 1989). The water content of the atmosphere is concentrated near Earth's surface and is extremely variable (Grossman, 1989). The atmospheric noise temperature therefore changes on time scales from hours to months. These variations in T_{sky} are monitored using the total-power channels. Given the system temperature T_{sys} measured by a total power channel[1], and the receiver noise temperature T_{rec} (see Chapter 3), the atmospheric noise temperature can be calculated as $T_{sky} = T_{sys} - T_{rec}$.

Figure 7.1 illustrates the distribution of sky temperatures for different parts of the

[1] Only one total power channel is used for atmospheric measurements in a given observing season. For PIQUE2000, D1 is used because D0 may have been affected by cryogenic IF amplifier oscillations (see Chapter 5). The responsivity and receiver noise temperatures are assumed to be $\mathcal{R} = 0.23$ cts/mK and $T_{rec} = 63$K (see Chapter 6).For PIQUE2001, D0 is used because it has a lower receiver temperature (see Chapter 3), and the assumed responsivity and noise temperature are $\mathcal{R} = 0.15$ cts/mK and $T_{rec} = 62$K (see Chapter 3).

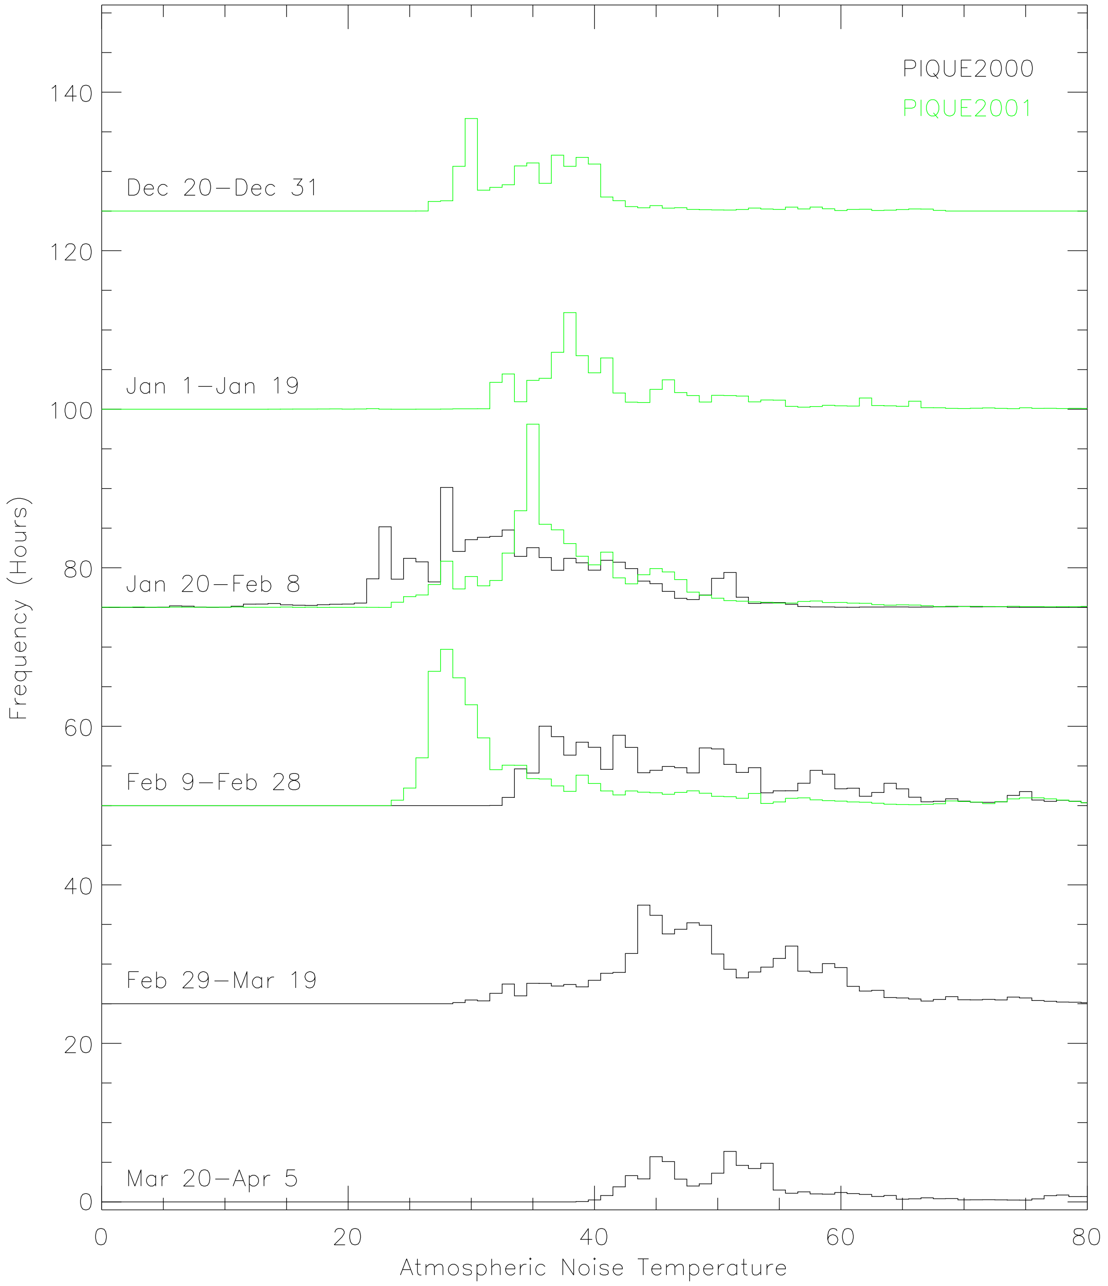

Figure 7.1: Distribution of sky temperature measured during different parts of the two observing seasons

observing seasons. These distributions only reflect times when the telescope was actively taking data (roughly 30% of the observing season, see Chapter 5), so periods of cloudy skies are underrepresented, although some cloudy periods are still evident in the long tails extending out beyond 100 K. The atmospheric noise temperatures corresponding to clear weather cluster around 40K, with some noticeable trends. The average value of T_{sky} increases over the course of the first observing season as one might expect. During the second observing season, however, T_{sky} actually decreases just before the end of the W-band observations. This is the result of unseasonably warm weather during January 2001.

The atmospheric noise temperature has a somewhat broader distribution in PIQUE2000 than in PIQUE2001. This probably reflects real meteorological differences between the two years, but it is also may be influenced by how often the telescope was deployed during the two observing seasons. In an effort to gather as much data as possible during the first observing season, the telescope was deployed whenever even a short stretch of clear weather seemed possible. These small patches of clear weather during otherwise cloudy days were actually nearly useless (see Section 8), so during the second observing season the telescope was only deployed if a reasonably long (6 hours) stretch of clear weather was expected.

The finite emissivity of the atmosphere implies that the atmosphere absorbs some of the astronomical 90 GHz radiation passing through it. This atmospheric loss can be calculated from the atmospheric noise temperature. The effective physical temperature of the water vapor in the atmosphere (the primary absorber) is 0.95 of the temperature at earths surface $T_{surface}$ (Grossman, 1989). Given the measured atmospheric noise temperature T_{sky}, the emissivity can be approximated as $e = T_{sky}/(0.95 * T_{surface})$. The attenuation factor is therefore $a = 1 - e = 1 - T_{sky}/(0.95 * T_{surface})$. Therefore, if $T_{surface} \sim 280$K and $T_{sky} \sim 40K$, astronomical signals at 90 GHz will attenuated by $\sim 20\%$. This attenuation must be taken into account in the data processing (see Chapter 8).

7.2 Atmospheric Stability

In addition to having a finite noise temperature, the atmosphere also has structure. The variations in the atmospheric noise temperature as a function of time and position make it difficult to measure the smaller variations in the CMB. Fortunately, the atmospheric emission should be unpolarized, so the structure of the atmosphere is less of a problem for a polarimeter than it is for a total power receiver.

The effects of atmospheric structure on the PIQUE receiver are best illustrated using power spectra (Figure 7.2). These power spectra are generated using the data from a single clear day (February 22) during the first observing season. (This day was selected because it provides a long, nearly continuous time series.) The power spectrum of the total power channel has a $1/f$ structure, so the fluctuations in this channel are dominated by drifts on long time scales. It is worth noting that the power spectrum displayed here has a steeper slope compared with the power spectrum obtained when the radiometer was viewing a regulated load (Figure 3.12). This implies that the atmosphere is the primary source of long time scale drifts in the total power channels during the observing seasons.

An ideal polarimeter would be completely insensitive to these drifts in atmospheric noise temperature. However, the actual multiplier channels of the PIQUE polarimeter have a finite responsivity to the temperature of the atmosphere (see Section 3.3). Consequently, the (unchopped) multiplier channels do respond to drifts in the atmospheric temperature, and their power spectra have a $1/f$ structure on long time scales. Fortunately, the large rejection ratios of the polarimetry channels suppress these drifts to such an extent that they are subdominant to random detector noise on time scales less than 500 seconds (indicated by the knee in the power spectrum). The telescope chops between two points on the sky on time scales between 10 and 30 seconds (see Chapter 5). The differences in the signal between the two chop positions are therefore drift free (the power spectra are flat) on time scales larger than a day. Therefore, given sufficient integration time, the PIQUE polarimeter can measure these differences accurately enough to detect cosmological polarization (see Section 3.3).

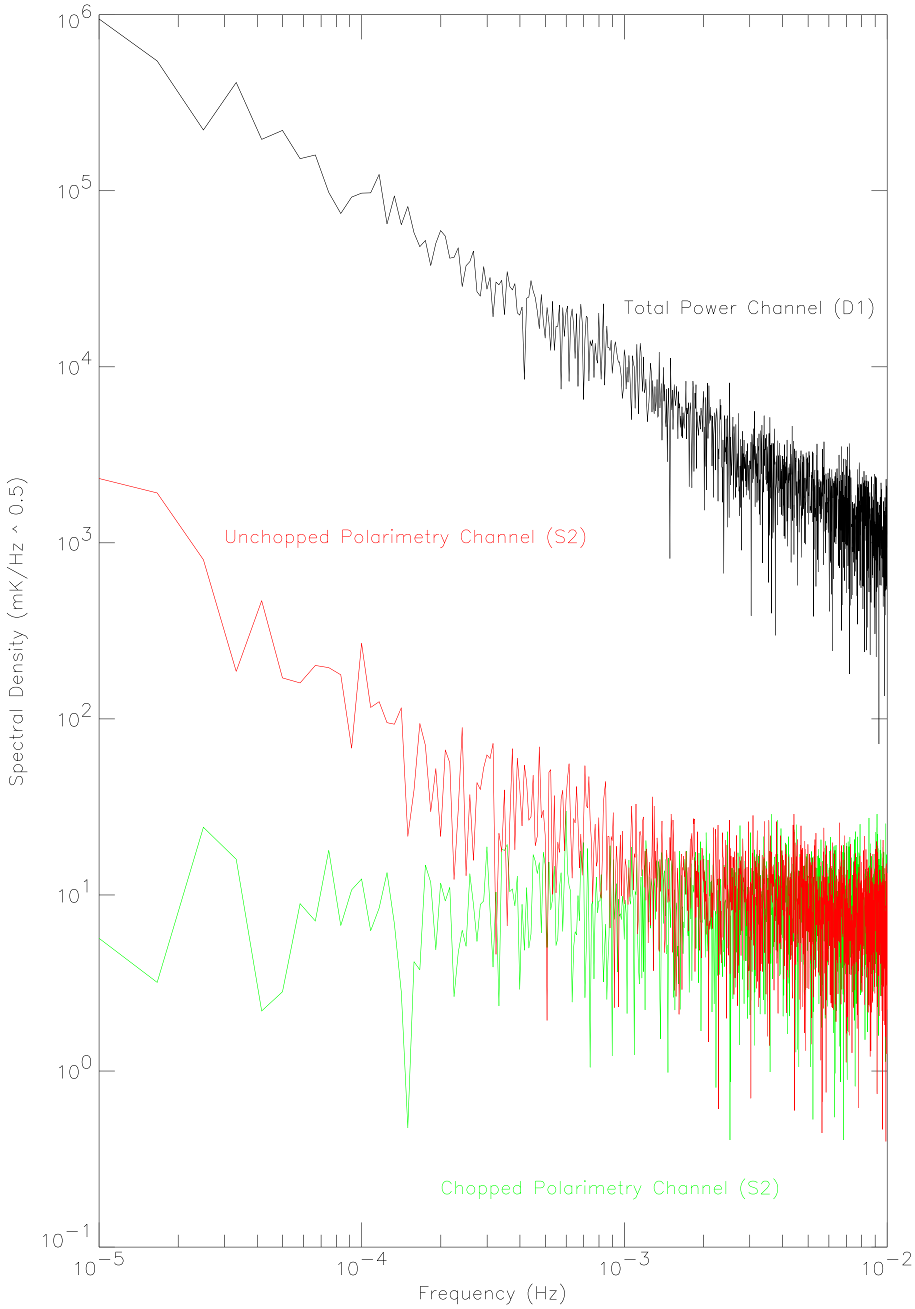

Figure 7.2: Power spectra from a day (February 22, 2000) during the first observing season. The power spectra of a total power channel, an unchopped multiplier channel, and a chopped multiplier channel are shown

Chapter 8

Data Processing

Over the course of two observing seasons, the W-band system observed the sky around the NCP for 1470 hours (see Table 8.1). A variety of selection criteria establish which data are corrupted by unfavorable weather conditions or electromechanical problems and which data contain useful astronomical information (see Section 8.2). The astronomically interesting data then provide estimates (with errors) of the polarization at different points on the sky (see Section 8.1, Figure 8.1). A series of χ^2 tests demonstrate the effectiveness of the selection procedures and the quality of the processed data (see Section 8.3).

8.1 Data Reduction

The data acquisition program writes the signals from the radiometer to file at a sampling rate of 62.5 Hz (see Appendix F); the base chops between the two scan positions with a period between 10 and 30 seconds (see Chapter 5), and the sky rotates around the NCP

Table 8.1: Allocation of time during active observations

Situation	2000	2001
Total Observation Time	810 hrs	660 hrs
Base in Motion	137 hrs	180 hrs
Mispointed scans	66 hrs	0 hrs
Data Removed *a priori*	176 hrs	222 hrs
Quadrature Covariance Cuts	124 hrs	62 hrs
Subsidiary Cuts	0 hrs	3 hrs
Time Remaining After Cuts	307 hrs	192 hrs

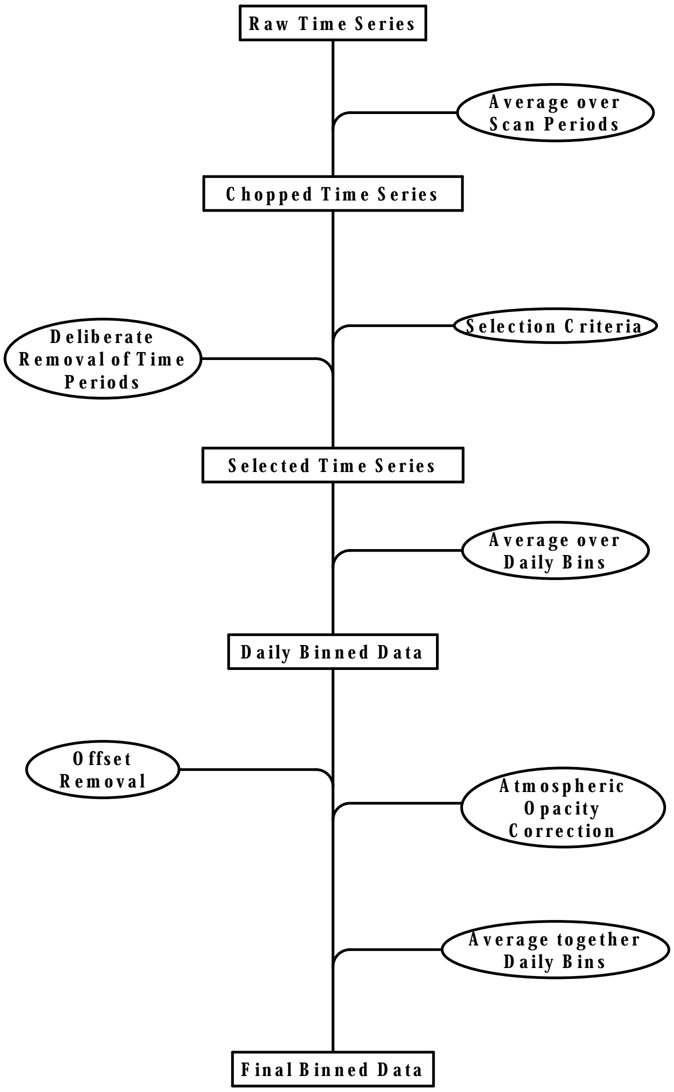

Figure 8.1: Flowchart illustrating the process involved in digesting the data from the raw time series to a final data set

once a day. All of these time scales come into play as part of the data reduction process (see Figure 8.1). The raw data are first converted into averages over each chop period. Various selection criteria (described in Section 8.2) establish which chop periods contain astronomically interesting data. These data are averaged together to produce estimates of the Stokes parameters as a function of position on the sky for each day of observations. After correcting for atmospheric opacity and removing offsets, these daily estimates are combined into the final data set.

8.1.1 From Raw Data to Chopped Data

The data reduction process begins with an IDL routine that takes the raw data from each chop period and computes the mean value of the signal from each radiometer channel at each scan position. (Other programs calculate the mean UT time and date for each chop period, as well as the average values of various housekeeping data.) The differences between the averages at the two scan positions contain the interesting astronomical information and form the "chopped" data set (or "CMB" data set) used in the next step of data processing (another program calculates "quadrature chopped data" for each polarimetry channel, see Section 8.3.1).

Data taken while the base was in motion are not used in these calculations, so 318 hours of raw data are not included in the chopped data set (see Table 8.1). Furthermore, if the encoder reading at one of the two scan positions for a given chop is more than $0.05°$ (10 encoder bits) from the target value (± 170 for PIQUE2000, ± 240 for PIQUE2001), then the data from that chop period are discarded. This removes 66 hours of data from PIQUE2000, but does not remove any data from PIQUE2001 (see Figure 8.2; The mispointed chops from the first observing season are due to irregularities in the timing of the motor control program, see Chapter 5).

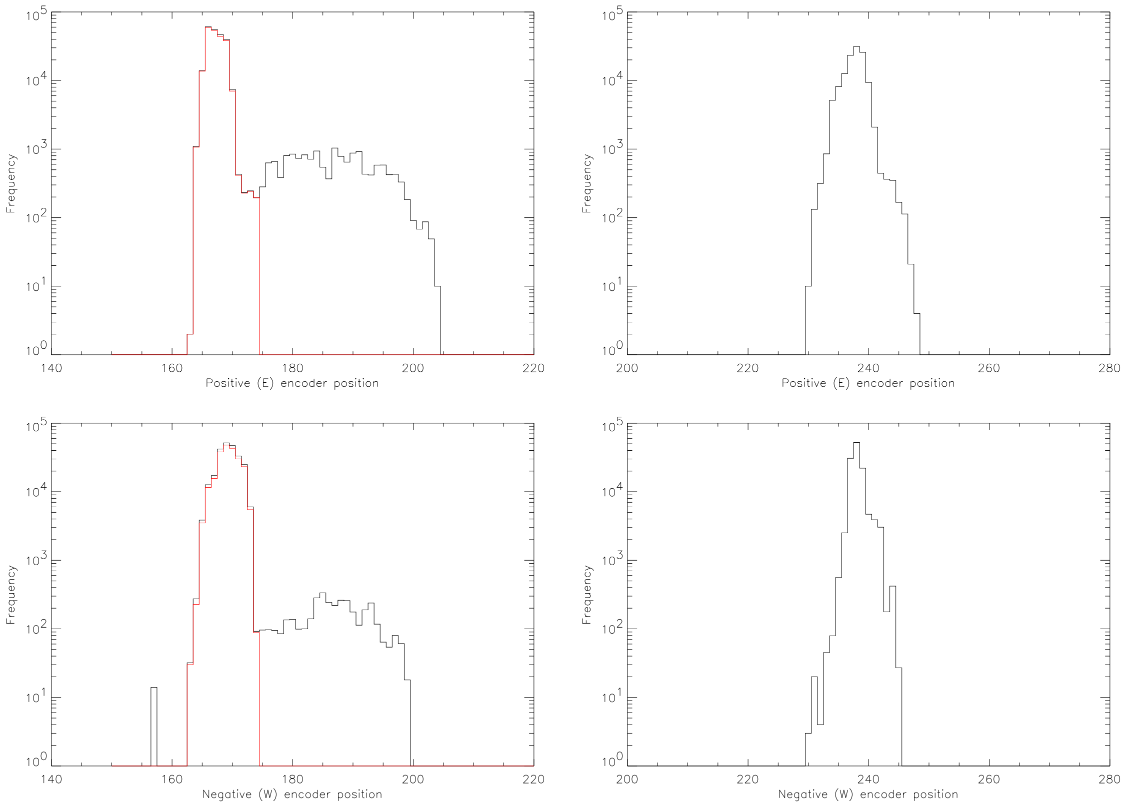

Figure 8.2: Distributions of the encoder values at the two scan positions. The x-axis is the absolute value of the encoder output relative to north, measured in encoder bits (65536 bits=360°). The data for PIQUE2000 is on the left and the data for PIQUE 2001 is on the right. The red curves for PIQUE2000 are the distributions after overshoots are removed.

8.1.2 From Chopped Data to Binned Data

After isolating the astronomically interesting data (see Section 8.2), the Local Sidereal Time (LST, i.e. the right ascension at zenith) of each chopped data point is calculated from the appropriate UT date and time. The LST determines the right ascension of the two scan positions (for PIQUE2000, the two scan positions are at RA = LST $\pm$ 6 hr, while for PIQUE2001, the scan positions are at RA=LST $\pm$ 12 hr), and thus unambiguously establishes the locations on the sky observed by the telescope during each chop.[1]

Next, another IDL routine divides the data among a set of specified time periods, each

[1]Recall the telescope observes a ring of constant declination ($\delta = 89°$).

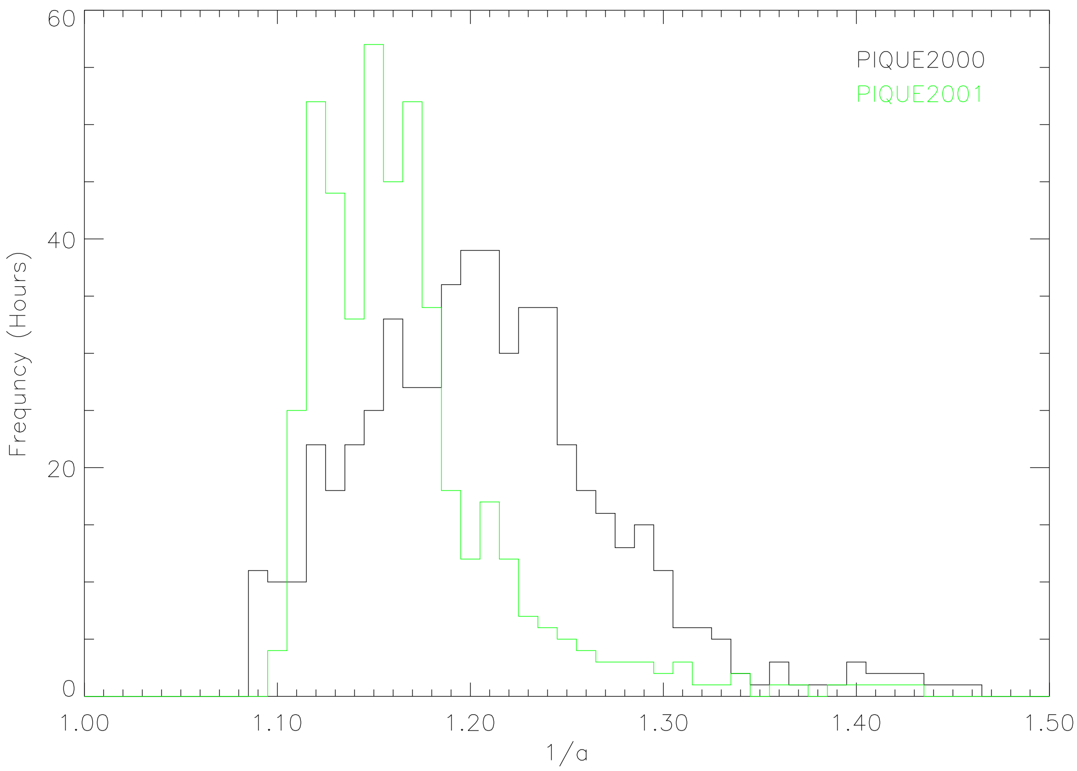

Figure 8.3: Distribution of the atmospheric attenuation corrector factor $1/a$ for the two observing seasons, assuming the physical temperature of the atmosphere is constant at 260K.

less than 24 hours long (see Section 8.2.1). The data from each of these "days" are further divided among either 24 or 144 LST bins.[2] For each of these "daily bins", the program calculates the mean and standard deviation of the mean of the chopped data from each polarimetry channel, along with the correlation coefficients between the three possible pairs of channels. In the 144-bin case, only about 50 chopped data points fall within each daily bin, so these calculations must take into account small number statistical affects (see Appendix G).

[2]With 24 bins, each bin is approximately the FWHM of the beam across in RA, so the signals from the sky in adjacent bins are nearly independent. With 144 bins, the pattern on the sky is oversampled, which preserves information on small angular scales. While studies of the data quality utilize both the 24-bin and the 144-bin data (see Section 8.3), the likelihood analyses use the 144-bin data exclusively (see Chapter 9).

Table 8.2: The deployments for the two observing seasons

PIQUE2000			PIQUE2001		
Deployment index	UT range	Hours of Data	Deployment index	UT range	Hours of Data
0	19.75-20.05	4.51	0	-10.30- -9.30	11.98
1	21.25-23.00	27.57	1	-8.05- -6.90	18.42
20	24.80-25.05	3.66	2	-6.15- -4.10	33.56
2	26.70-30.60	63.31	3	-2.20- -1.10	15.84
3	32.05-34.45	24.31	4	0.85-2.65	29.59
4	36.80-41.75	77.85	5	2.70-4.10	22.98
5	42.95-44.75	32.19	6	7.15-7.95	12.74
6	46.15-47.35	18.84	7	9.60-10.20	9.11
7	47.65-49.40	29.39	8	10.90-11.40	7.62
8	51.60-55.10	54.11	9	16.90-17.55	10.29
9	59.65-61.90	37.38	10	21.75-25.15	46.10
10	63.55-68.90	92.59	11	26.05-26.90	14.41
11	71.00-71.45	6.19	12	28.10-29.90	30.67
12	73.55-76.50	40.20	13	31.75-33.45	23.74
13	77.90- 78.85	13.91	14	34.15-35.60	25.83
14	79.80-80.85	15.78	15	38.70-39.70	16.17
15	84.90-85.75	14.04	16	40.10-40.60	7.99
16	86.95-87.70	12.16	17	41.80-44.10	39.58
17	89.95-90.55	9.18	18	44.55-45.35	9.86
18	92.05-93.65	28.54	19	48.80-51.65	49.72
			20	52.25-53.70	19.28
			21	54.65-55.60	14.04
			22	58.05-58.80	10.77
Total		605.71	Total		480.29

NB: Data were not taken continuously during each deployment

The data in each daily bin come from a short, discrete time period (an hour for 24 bins, ten minutes for 144 bins). The atmospheric noise temperature T_{sky} does not change significantly on these time scales, so an effective atmospheric attenuation factor a is calculated for each daily bin using the mean value of T_{sky} and assuming $T_{surface} = 273$ K (see Chapter 7). In order to correct for atmospheric losses, the daily binned estimates and errors are multiplied by $1/a \sim 1.2$ (see Figure 8.3).[3]

Next, the chopped offsets are removed. As mentioned in Chapter 4, these offsets are most likely due to pick-up from the panels of the outer ground screen. The orientation of these panels may have changed slightly each time the panels were deployed, so nominally a separate offset is calculated for each "deployment" (i.e. each period when the outer ground

[3]This correction is not applied to the chopped data directly because the total power channels used to estimate the atmospheric noise temperature are too noisy on the time scale of a single chop.

screen panels were in their upright and locked positions, see Table 8.2). Another method of partitioning the offsets is discussed in Section 8.3.4. Regardless of the partitioning method, the offsets are the weighted means of the appropriate daily binned data.

After correcting for atmospheric opacity and subtracting off the chopped offsets, the daily binned data are combined to yield the final estimates (with errors) of the polarized signal in each LST bin. For each polarimetry channel, LST bin and observing season, the final estimate is the weighted mean of the appropriate estimates from the different days. The error on this mean is computed taking into account the statistical effects discussed in Appendix G. The correlation coefficients for each bin are simply the averages of the appropriate daily values.

8.2 Data Selection

After generating the chopped data set but before these data are averaged together in the daily bins, the astronomically interesting data are isolated from the data corrupted by meteorological and electromechanical effects. Some data are eliminated *a priori*, while other data are removed based on selection criteria.

8.2.1 Deliberate Elimination of Time Periods

As described in Section 8.1.2, the chopped data are divided among a set of specified time periods prior to binning. Data from time periods outside these "days" are not included in the final data set. A total of 398 hours of data are deliberately omitted from the list of acceptable time periods (see Table 8.3). These omissions eliminate stretches of data that are known *a priori* to be useless for astronomical measurements. Most of these data come from times when the weather was continuously bad (i.e. cloudy), moisture condensed on the window, or various electromechanical problems (such as LO failures) occurred (see Chapter 5). The warm-ups during PIQUE2001 are eliminated from the nominal data sets so the quality of the remaining data (taken with stable dewar temperatures) can be established (see Section 8.3.5). Deployments containing less than 12 hours of data are also excised from

Table 8.3: Time periods deliberately omitted from consideration

Time Period	Hours of Data Removed	Reason For Removal
PIQUE2000		
19.75-20.05	4.51	Short, isolated data fragment
24.80-25.05	3.66	Short, isolated data fragment
26.70-27.00	5.87	Weather (clouds)
32.05-34.45	24.31	LO failures
37.20-37.60	6.41	Condensation
38.20-39.20	14.92	Weather, Condensation
40.50-40.80	3.63	Condensation
41.30-41.80	7.34	Weather, Condensation?
59.65-60.00	5.50	Weather (clouds)
67.00-68.90	34.94	Weather
71.00-71.45	6.19	Short, isolated data fragment
75.00-76.50	21.23	Condensation
89.95-90.55	9.18	Short, isolated data fragment
92.05-93.65	28.54	E-Box out of regulation
Total	176.23	
PIQUE2001		
-10.30- -9.30	11.98	Short isolated data fragment/Condensation
-8.05- -6.90	18.42	Strange quadrature data (s0)
7.15-7.95	12.74	Condensation
9.60-10.20	9.11	Short, isolated data fragment/LO failure?
10.90-11.40	7.62	Short isolated data fragment/Condensation
16.90-17.55	10.29	Short isolated data fragment/Condensation
21.75-22.8	13.13	Condensation/LO failures
26.05-26.90	14.41	Dewar temps unstable
28.10-29.90	30.67	Dewar temps unstable/LO failures
31.75-33.45	23.74	Dewar temps unstable/Weather/LO failure?
34.15-35.60	25.83	Dewar temps unstable
39.50-39.70	2.62	Weather/Dewar temps unstable
40.10-40.60	7.99	Short isolated data fragment/Weather
41.80-42.80	17.46	Dewar temps unstable
48.80-49.50	12.41	Dewar temps unstable
54.65-54.90	3.20	Weather
Total	221.62	

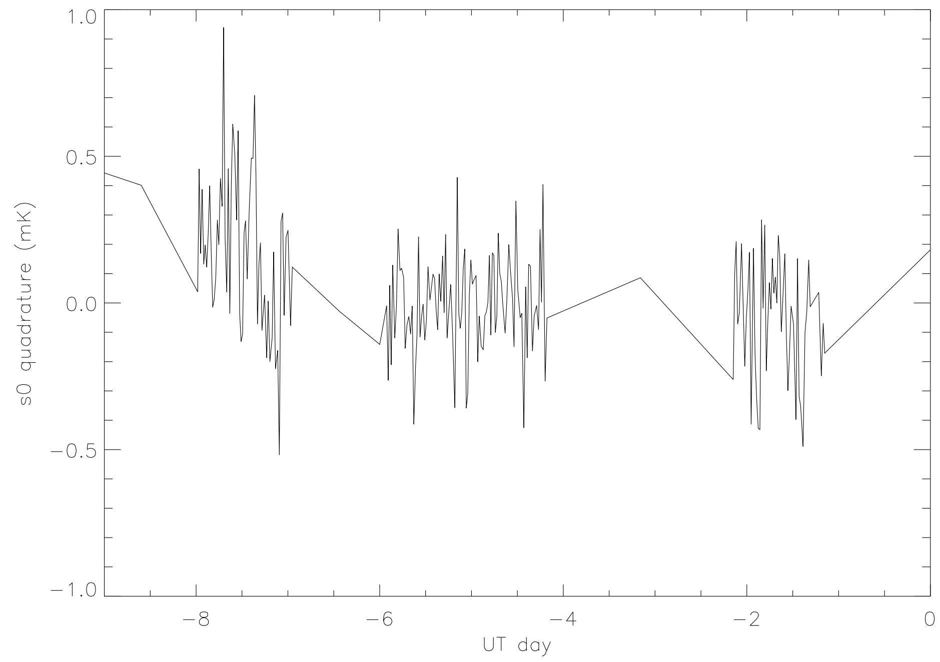

Figure 8.4: Strange behavior in the s0 quadrature data during PIQUE2001. For clarity, the data displayed here are smoothed over hundreds of chops. Between UT day= -8 and UT day = -7, this quadrature data shows a large (~ 0.5 mK) offset not seen in the later deployments. This offset is not seen in the other channels or even in the CMB data.

the nominal data sets to avoid complicating the data analysis[4].

One segment of data (-8.05- -6.90) omitted from the second observing season deserves special mention. During this time, the quadrature data from one multiplier channel (s0) has a large offset. Such an obvious offset is found nowhere else in the quadrature data from either observing season (see Figure 8.4). It is not clear how these data acquired such an offset (the quadrature data are calculated from differences between data taken at the same scan position, see Section 8.3.1), but regardless of its origin, it clearly indicates a problem with the data taken during this time period.

[4]Each such isolated fragment has a separate offset removed, so these fragments do not provide information about the polarized signals on the largest angular scales. This can be taken into account in the likelihood analyses (Keating et al., 2001a), but these fragments contain such a small fraction (16 hours) of the total data that it is not worth the trouble.

8.2.2 Selection Criteria

While electromechanical problems and stretches of bad weather can be removed *a priori*, more subtle meteorological phenomena (e.g. isolated clouds) also affect the data quality. These phenomena are identified using an unbiased selection criteria: the quadrature covariances.

Quadrature Covariances (QC)

In order to apply the quadrature covariance (QC) selection criteria, the quadrature data (see Section 8.3.1) are first divided into units consisting of a set number of chops (called the "range").[5] For each of these units, the correlation coefficient between every possible pair of polarimetry channels is calculated using the formula: $\sum \frac{(\bar{x}-x_i)(\bar{y}-y_i)}{\sigma_x \sigma_y}$ (x_i and y_i are the data from two different channels, and $\bar{x}$ and σ_x are the mean and standard deviation of the x_i). The data from each unit is considered "astronomically interesting" only if the absolute values of all three correlation coefficients are less than a certain value (called the "level").

These criteria remove time periods corrupted by rapid variations in the atmosphere (for example, clouds). When the atmosphere is stable, the quadrature data are random gaussian noise, and the noise in the different channels is uncorrelated. In this case, the correlation coefficients have (approximately) guassian distributions with a mean of zero and a width inversely proportional to the square root of the range (Lupton, 1993). However, when the atmosphere changes rapidly, all three polarimetry channels respond synchronously to the change[6], which generates excess correlations. The actual distributions of the correlation coefficients (shown in Figures 8.5) do indeed have gaussian components peaking near zero and large positive tails. While the distributions for both the quadrature and the CMB data have this same basic shape, the selection criteria are applied to the quadrature data to avoid biasing the CMB data.

Two QC criteria with different ranges are used to extract the nominal data set. The first

[5]Due to the changes in the chop period during PIQUE2001 (Chapter 5), a fixed range does not correspond to a single time span.

[6]Either due to the unpolarized responsivities of these channels or due to real polarized signals

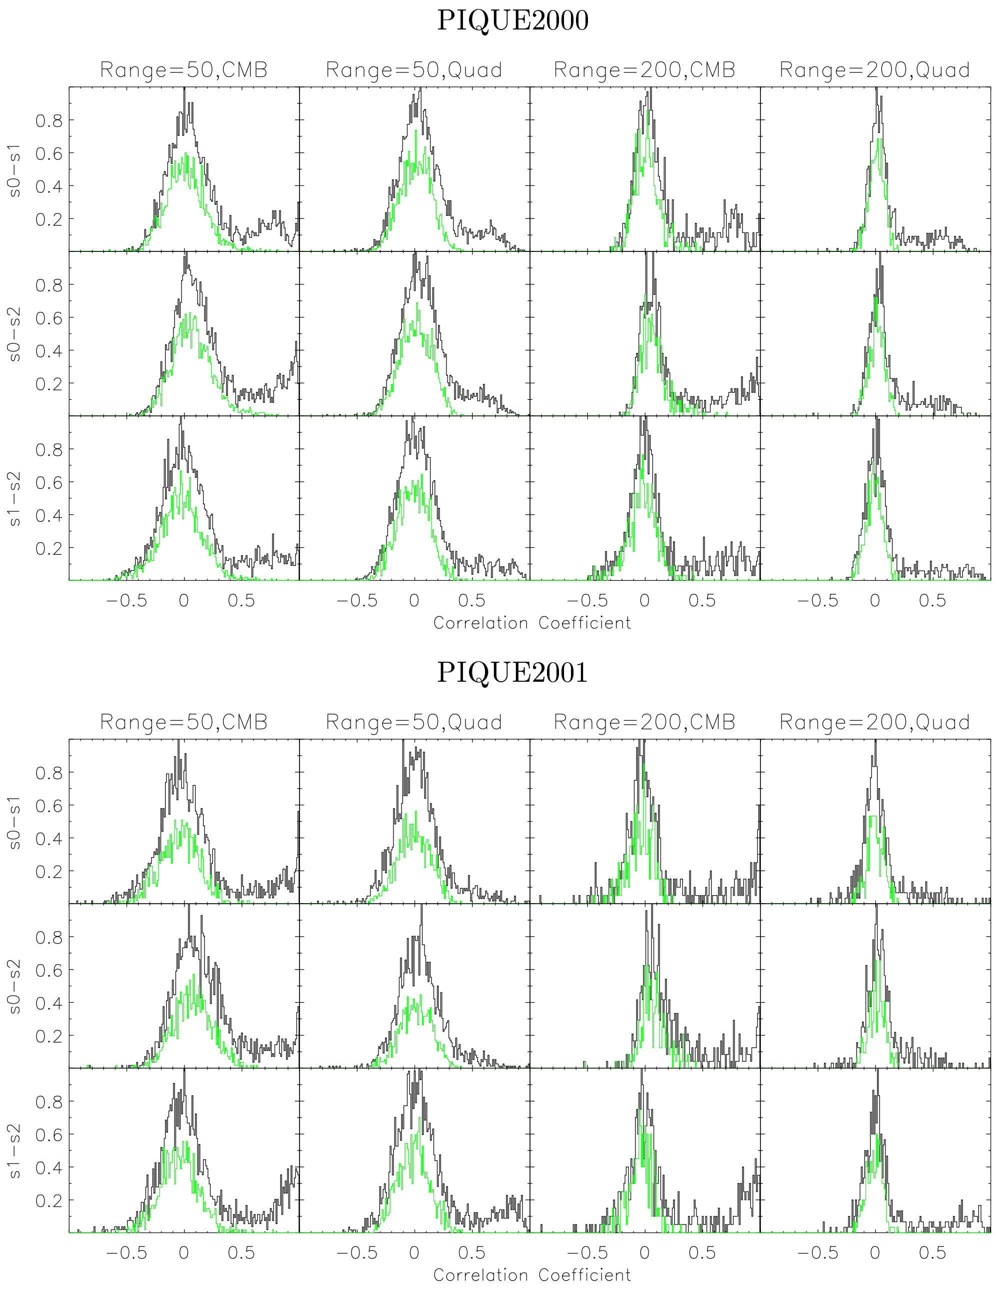

Figure 8.5: Correlation coefficients between the three polarimetry channels. The black curves are the distributions before any cuts, and the green curves show the distributions after all cuts (to scale).

criteria (range=200, level=0.20) efficiently removes long periods of corrupted data, while the second criteria (range=50, level=0.32) is better at eliminating short, isolated patches of contaminated data[7] (the long range is an integral multiple of the short range so that the units for the two selection criteria are commensurate with each other). The nominal levels for these two cuts are between two and three times the width of the gaussian component (see Section 8.3.3). By itself, each of these selection criteria removes about 125 hours of data. Together, they remove 186 hours of data and eliminate the tails in the distributions of both the quadrature and the CMB data (see Figure 8.5).

Subsidiary Cuts for PIQUE2001

Two subsidiary cuts are applied to the data from PIQUE2001. First, a covariance cut similar to the QC criteria is applied to the CMB data with a range of 10 and a level of 0.9. This cut removes data taken while someone was checking the window for frost or condensation (which requires sticking one's head directly in the beam). The CMB data are used to remove the resulting spikes in the time stream because the coherent response of the polarimetry channels to the person's head occurs on a time scale roughly commensurate with the dwell time. Only 6 minutes of data are removed by this cut. (This cut is not applied to PIQUE2000 data because there was no hatch in the inner ground screen during this observing season, and every time the window was checked the data acquisition program had to be stopped so one panel of the inner ground screen could be removed.)

The second cut removes any 10 chop long period where the variance of s0 exceeds the variance of s2 by a factor of 2.5 or more. This cut removes a series of glitches in s0 from early in PIQUE2001, which amount to three hours of data. The origin of these glitches has not been determined, and these data have little effect on the data quality or the results.

[7] Ranges longer than 200 are not practical because there are too few units per observing season, which makes it difficult to identify the guassian component. Ranges shorter than 50 are not practical because the guassian component is too wide, which makes it difficult to cleanly isolate the tails.

8.3 Data Quality

8.3.1 Methods: Nulled Data Sets

The quality of the processed data is established using a set of "nulled" data sets. Every element of these data sets is a combination of the polarimetry data that is identically zero for signals from the CMB. If the nulled data are not consistent with zero (i.e. the χ^2 of zero has low probability to exceed), then the processed data may contain some residual atmospheric contamination. Conversely, if the nulled data are consistent with zero signal, then the processed data are likely free from such contamination. Four different types of nulled data are generated: inter-channel differences, quadrature data, long time scale differences and scan pattern nulls.

Inter-channel Differences

Three nulled data sets are generated by taking the difference between the binned data from every possible pair of multiplier channels (s0-s1, s0-s2, s1-s2). In the observed limit of small cross-correlations between the channels (see Figure 8.5) and in the absence of astrophysical foregrounds with a strong spectral slope (see Section 1.3) and assuming the relative gains of the channels are correct, these differences should be zero.

Quadrature Data

Quadrature chops are calculated for each polarimetry channel during the first stage of data processing (see Section 8.1.1). Unlike the "normal" chops, which are differences between the data taken at different scan positions, the quadrature chops are calculated from differences between data taken at the same scan position. For each chop, the data from each scan position are broken into two halves. If $E_1(W_1)$ and $E_2(W_2)$ are the mean values of the data from the two halves at the east(west) scan position, then the quadrature chop is $\frac{1}{2}((E_1 - E_2) + (W_1 - W_2))$. For any signal fixed on the sky, this quantity should be zero. Processing the quadrature chops using the procedures described in Section 8.1.2 therefore yields a nulled data set for each polarimetry channel.[8]

Long Time Scale Differences (Half-half Test)

Three nulled data sets (one for each polarimetry channel) are created by binning the selected data from the two halves of an observing season separately. The difference between these two data sets should be identically zero because any signal fixed on the sky does not change from day to day.

The point in time separating the two halves is chosen so that the two data sets (1) contain roughly the same amount of useful data (which insures the differenced data set has the smallest possible error bars) and (2) are separated by a natural break in the time series. The division point is between two deployments (5 and 6) for PIQUE2000, while for PIQUE2001 it is when the scan period changes from 10 seconds to 20 seconds (see Chapter 5).[9]

Scan Pattern Nulls

The scan pattern nulls are a consequence of PIQUE's differencing scan strategy. If the differenced signal for a bin at LST x is $D_Q(x)$ for PIQUE 2000 and $D_U(x)$ for PIQUE2001, then $D_Q(x)-D_Q(x+6\text{hrs})+D_Q(x+12\text{hrs})-D_Q(x+18\text{hrs}) = 0$ and $D_U(x)+D_U(x+12\text{hrs}) = 0$ for any signal fixed on the sky (see Figure 8.6). Consequently, there are 36(72) independent combinations of the 144-bin data from each PIQUE2000 (PIQUE2001) polarimetry channel (and 6(12) combinations of the 24-bin data) which should equal zero.

8.3.2 Nominal Data Set

The nominal data set is extracted using the nominal QC selection criteria (range=200, level=0.2; range=50, level=0.32). Separate offsets are removed from each deployment, and the warm-ups are not included in the PIQUE2001 data.

[8] The quadrature data are also used in the QC selection criteria, see Section 8.2.

[9] The results of the χ^2 tests are not sensitive to the location of the division point.

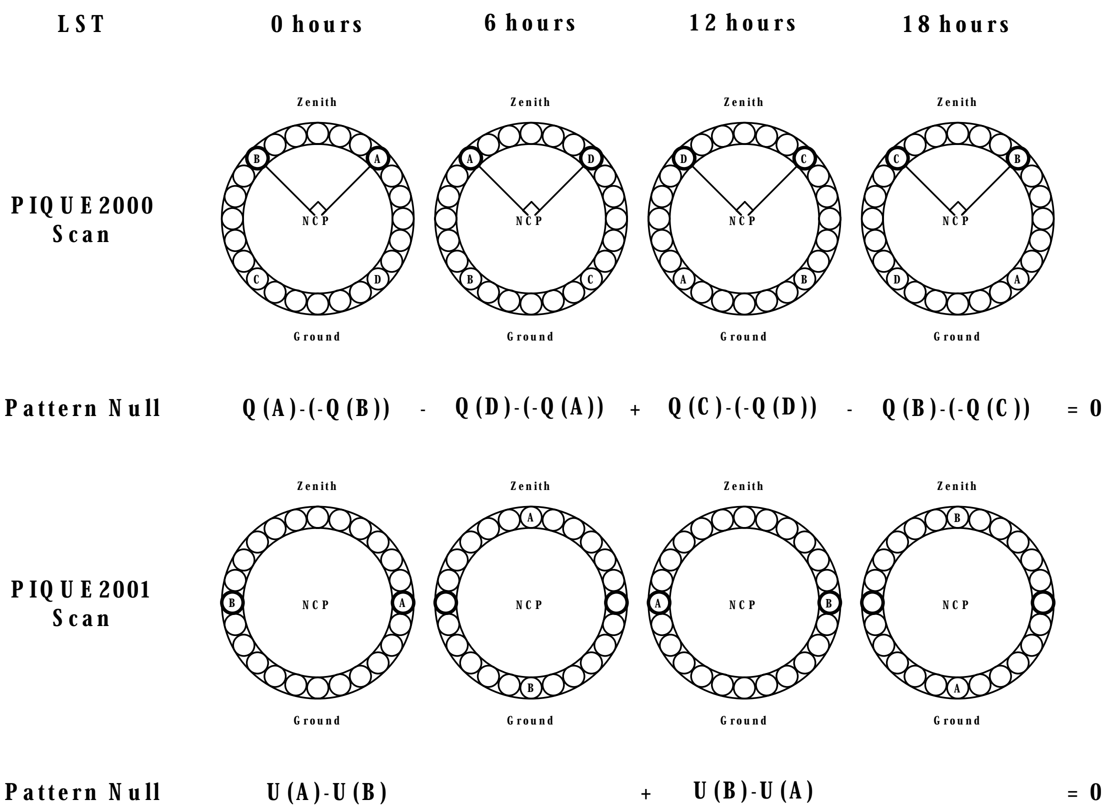

Figure 8.6: Scan pattern nulls. The four columns from left to right show the observed ring around the NCP at a sequence of times separated by 6 hours in LST. The two spots on the sky observed by the telescope are highlighted, and a set of pixels on the sky are marked with letters. Below each figure is the difference in the Stokes parameters measured at the given time (The extra minus sign for PIQUE2000 occurs because telescope measures Q at the east scan position, but $-Q$ at the west scan position). The indicated combinations of these signals equal zero for any signal fixed on the sky.

Table 8.4: χ^2 probabilities to exceed for zero of the null data sets for the nominal data set

PIQUE2000 (307.0 hrs)								
	24 bins				144 bins			
test	DOF	s0	s1	s2	DOF	s0	s1	s2
Pattern Null	6	0.434	0.088	**0.045**	36	0.455	0.264	0.140
Quadrature	23	0.520	0.438	0.173	143	0.960	0.248	0.188
Half-half	23	0.838	0.144	0.472	143	0.183	0.139	0.329
Combined	52	0.767	0.113	0.109	322	0.686	0.091	0.113
		s0-s1	s0-s2	s1-s2		s0-s1	s0-s2	s1-s2
Inter-channel	23	0.593	0.355	0.340	143	0.623	0.832	0.827
Overall	225		.177		1395		0.388	
PIQUE2001 (192.3 hrs)								
	24 bins				144 bins			
test	DOF	s0	s1	s2	DOF	s0	s1	s2
Pattern Null	12	0.589	0.082	0.711	72	0.801	0.120	0.378
Quadrature	23	0.526	0.361	0.083	143	0.383	0.612	0.282
Half-half	23	0.350	0.569	0.138	143	0.173	0.609	0.581
Combined	58	0.505	0.238	0.098	358	0.347	0.436	0.369
		s0-s1	s0-s2	s1-s2		s0-s1	s0-s2	s1-s2
Inter-channel	23	0.173	0.701	0.219	143	0.473	0.658	**0.043**
Overall	243		0.095		1503		0.204	

There are 48 nulled data sets associated with this nominal data set. For each observing season and each binning (24 and 144), there is a scan pattern null, quadrature and half-half data set from each channel, and three inter-channel differences. The χ^2 of zero signal is computed for each of the nulled data sets, and the corresponding probabilities to exceed are given in Table 8.4. Two of these probabilities are below 0.05, which is consistent with the expected number of 2.4. However, only one probability is greater than 0.95. These statistics therefore do not clearly establish whether the overall distribution of the χ^2 statistics is consistent with noise. Consequently, Table 8.4 also includes combined statistics that better reflect the overall quality of the data. Adding the χ^2s of the pattern null, quadrate and half-half data sets together yields a "combined" χ^2 for each channel, and adding these combined numbers together with the χ^2 from the inter-channel data yields a single "overall" statistic for each observing season (and binning). The probabilities to exceed of both the combined and the overall χ^2s are always larger than 0.05. Hence, there is no evidence for residual

Table 8.5: Hours of data remaining for different cut levels

Hours of data remaining in PIQUE2000						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.16
0.90	426.3	392.3	357.4	333.2	317.0	291.4
0.50	391.6	381.9	356.4	333.1	316.8	291.3
0.36	359.5	353.8	337.0	318.6	305.2	282.1
0.32	341.7	336.5	322.6	307.0	294.6	273.0
0.30	327.2	322.4	310.7	297.0	285.4	265.5
0.27	300.4	269.4	286.6	274.5	264.5	246.6

Hours of data remaining in PIQUE2001						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.16
0.90	249.2	231.1	227.7	208.4	198.0	187.2
0.50	235.0	229.4	227.0	208.2	197.9	187.2
0.36	223.8	218.6	216.9	201.1	191.7	182.9
0.32	213.4	208.2	206.5	192.3	183.9	176.0
0.30	206.1	201.2	199.5	186.5	178.9	171.1
0.27	187.6	183.2	182.0	171.4	165.0	158.2

Table 8.6: Overall probabilities to exceed as a function of cut levels

24 bin overall P, PIQUE2000						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.216
0.90	**0.046**	**0.040**	0.083	0.146	0.057	**0.042**
0.50	0.053	**0.036**	0.065	0.132	**0.050**	**0.036**
0.36	**0.046**	**0.050**	0.132	0.249	0.097	0.094
0.32	**0.038**	**0.038**	0.075	0.177	0.069	0.055
0.30	0.070	0.076	0.108	0.200	0.087	0.086
0.27	0.138	0.122	0.121	0.063	0.085	0.098

24 bin overall P, PIQUE2001						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.16
0.90	**0.037**	**0.031**	**0.029**	0.063	0.116	0.135
0.50	**0.015**	**0.014**	**0.022**	**0.050**	0.116	0.135
0.36	**0.028**	**0.039**	**0.043**	0.056	0.079	0.073
0.32	**0.028**	**0.040**	**0.041**	0.095	0.110	0.120
0.30	0.067	0.086	0.086	0.159	0.156	0.154
0.27	0.231	0.249	0.255	0.222	0.268	0.305

144 bin overall P, PIQUE2000						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.16
0.90	0.223	0.206	0.100	0.153	0.077	0.129
0.50	0.145	0.128	0.097	0.145	0.072	0.114
0.36	0.223	0.211	0.153	0.234	0.141	0.145
0.32	0.503	0.502	0.420	0.391	0.212	0.207
0.30	0.348	0.334	0.103	0.245	0.121	0.196
0.27	0.217	0.211	0.148	0.186	0.141	0.234

144 bin overall P, PIQUE2001						
QC level	QC level (range =200)					
(range=50)	0.90	0.50	0.30	0.20	0.18	0.16
0.90	0.423	0.459	0.481	0.295	0.388	0.253
0.50	0.347	0.360	0.437	0.283	0.388	0.253
0.36	0.308	0.388	0.288	0.238	0.314	0.190
0.32	0.175	0.221	0.215	0.205	0.314	0.210
0.30	0.185	0.242	0.231	0.145	0.259	0.166
0.27	0.333	0.281	0.347	0.185	0.301	0.220

atmospheric contamination in the nominal data set.

8.3.3 Variations with cut levels

The levels of the QC criteria are determined based on the distributions of the correlation coefficients. Even so, a variety of cut levels is considered here. These different levels yield criteria that remove widely different amounts of data (see Table 8.5). Table 8.6 gives the overall probabilities to exceed for the 24-bin and the 144-bin data derived from these different cuts. As expected, the quality 24-bin data generally improves (the probabilities to exceed increase) as more data arre removed (the quality of the 144-bin data is good regardless of the cut levels). The quality of the nominal data sets is therefore not a statistical fluke dependent on the exact values of the cut levels.

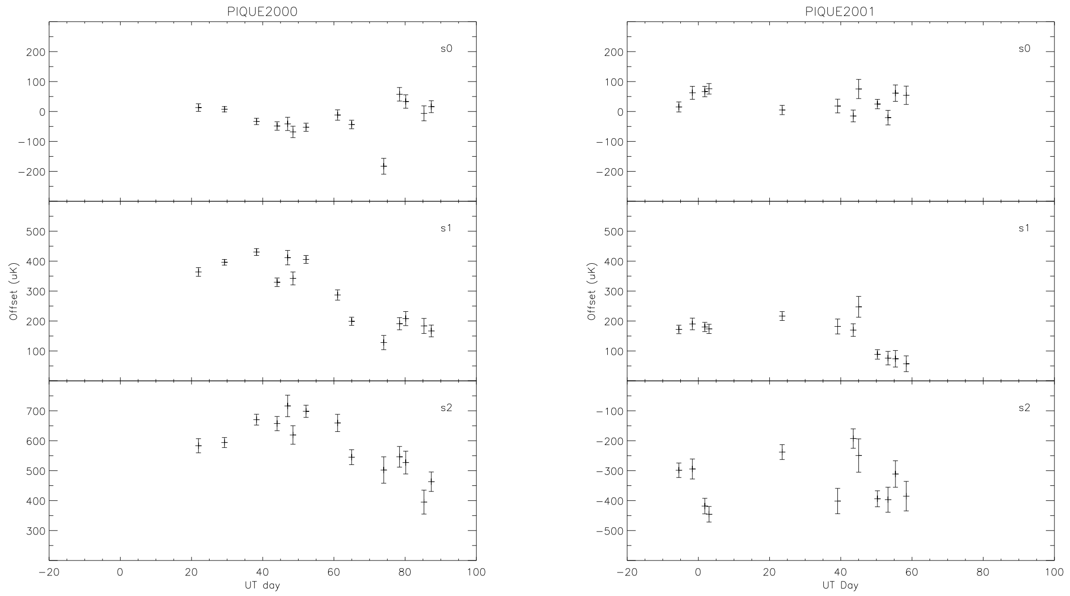

Figure 8.7: The offsets removed from the deployments for PIQUE2000 (left) and PIQUE2001 (right). A data point (with error bar) is given for each deployment.

8.3.4 Variations with number of offsets

Figure 8.7 shows the offsets (with errors) for the different deployments in the nominal data set. The scatter in the offsets is clearly larger than the error bars, and there are some interesting patterns. Most obviously, the offsets for s1 during PIQUE2000 shift between two discrete levels around day 60. Similar shifts appear in the offsets from s2 and even perhaps in s0. More interesting, the offsets in s1 from PIQUE2001 also show a shift around day 50. These discrete shifts cannot be explained by changes in the orientation of the outer ground screen panels (see Section 8.1.2). Perhaps they are due to some critical change in the sun angle and the solar heating of the outer ground screen (although it is hard to see why this would cause the offsets to get smaller later in the year). In any case, these shifts account for most of the excess scatter in the offsets (although the offsets before and after the shifts are not, in general, consistent with each other to within their error bars), so perhaps only two offsets need to be subtracted from each observing season. Table 8.7 compares the nominal data with data using only two offsets per observing season (The offset shift is assumed to occur between deployments 9 and 10 of PIQUE2000 and between deployments

Table 8.7: Combined probabilties to exceed for different offset removals

Condition	24 bins				144 bins			
	s0	s1	s2	overall	s0	s1	s2	overall
PIQUE2000								
Deployment-Based Offsets	0.767	0.113	0.109	0.177	0.686	0.091	0.113	0.388
Two Offsets	0.537	0.077	0.108	0.064	0.541	0.077	0.108	0.232
PIQUE2001								
Deployment-Based Offsets	0.505	0.238	0.098	0.095	0.347	0.436	0.369	0.204
Two Offsets	0.464	0.205	0.060	**0.024**	0.326	0.909	0.367	0.152

Table 8.8: Combined probabilities to exceed for PIQUE2001, including warm-ups

Condition	hours	24 bins				144 bins			
		s0	s1	s2	overall	s0	s1	s2	overall
Nominal	192.3	0.505	0.238	0.098	0.095	0.347	0.436	0.369	0.204
Warm Ups Included	239.6	0.375	0.208	**0.027**	**0.019**	0.245	0.128	0.183	**0.014**

18 and 19 of PIQUE2001). The probabilities to exceed become somewhat smaller with two offsets (and the overall probability to exceed for the PIQUE2001 24-bin data falls below 0.05), but the overall quality of the 144-bin data (which is the only data used in the final analysis, see Chapter 9) is good even with two offsets. The data quality is therefore not particularly sensitive to the number of offsets removed.

8.3.5 Allowing warm-ups back into PIQUE2001

Finally, the quality of the data taken during the warm-ups of PIQUE2001 is explored using data sets that include all of deployments 11 and 14 and portions of deployments 17 and 19 (the data from other warm-ups are affected by LO failures and bad weather). These data are corrected for temperature dependent gain variations (see Chapter 6). Table 8.8 compares the quality of these data with the quality of the nominal data sets. The warm-ups clearly reduce the quality of even the 144-bin data, so these data are not useful for astronomical measurements.

Chapter 9

Preliminary Data Analysis

The end result of the data reduction process is a set of estimates (with errors) of the polarized signal as a function of position on the sky (see Figure 9.1). The scatter in these estimates is consistent with the error bars, so PIQUE cannot claim a detection of the cosmological polarization. However, these measurements do place strong limits the polarized component of the CMB.

9.1 Methods: Likelihood Analysis

Constraints on the cosmological polarization result from comparisons between the observed scatter in the data and the expected scatter due to instrumental noise and any fluctuations in the polarized signals from the sky. This comparison is nontrivial because there are cross-correlations between the various data points. Likelihood analyses account for these cross-correlations in a conceptually straightforward (but computationally expensive) manner.

9.1.1 Bayesian Statistics

Likelihood analysis is a form of Bayesian statistics, which relies on the fundamental probability relation:

$$P(AB) = P(A)P(B|A) = P(B)P(A|B).$$

This relation states that the probability of both A and B occurring is both (1) the probability of A occurring times the conditional probability that B will occur if A occurs, and (2) the

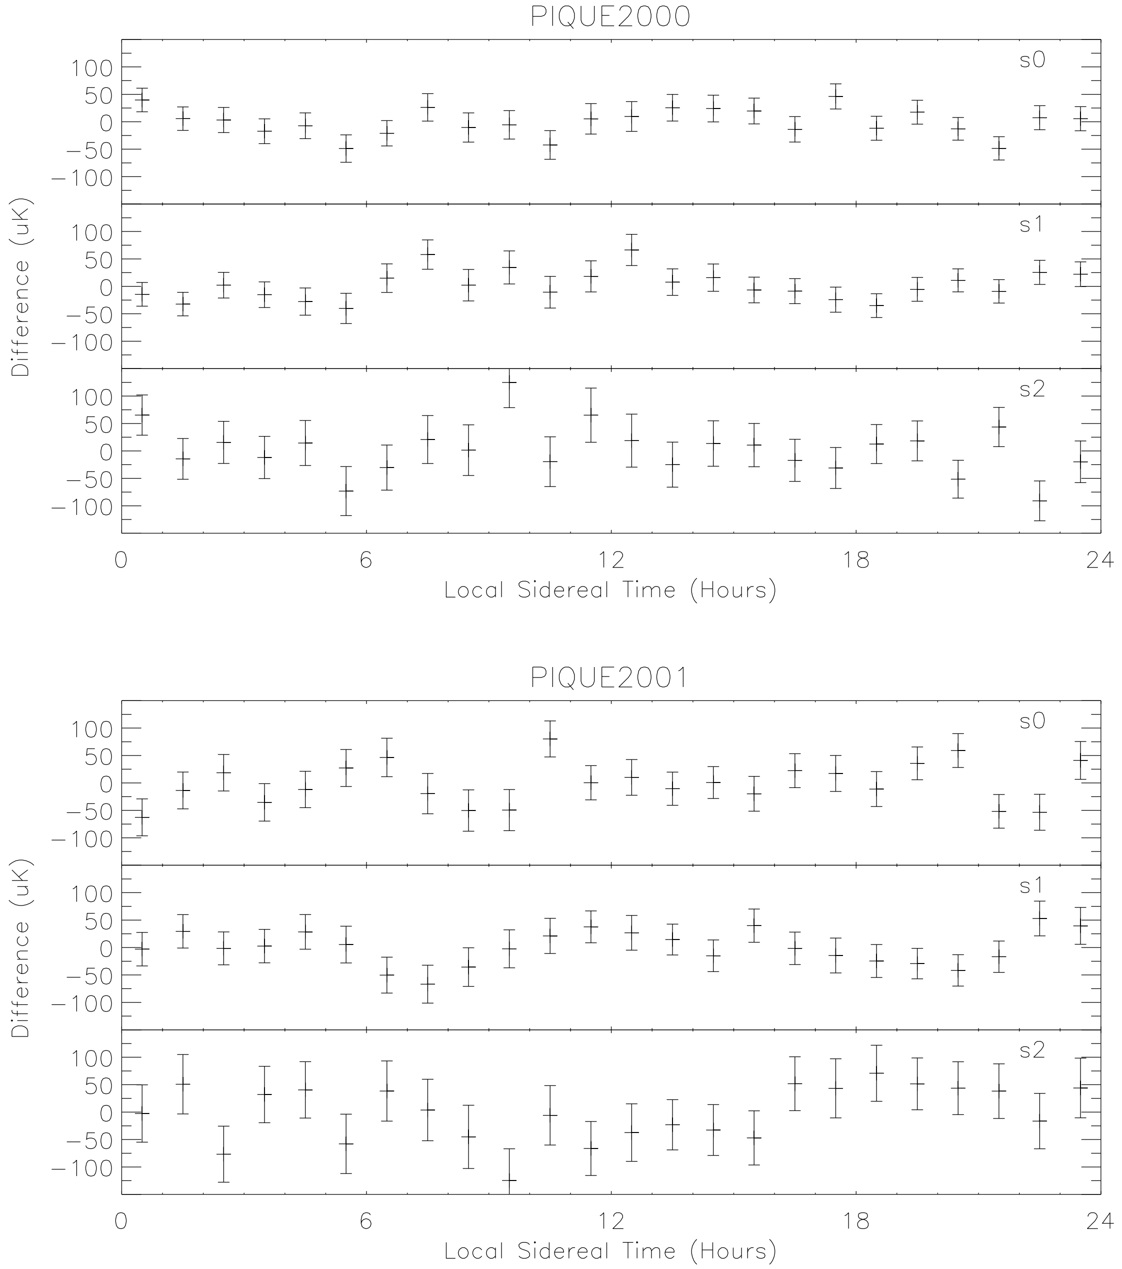

Figure 9.1: The final nominal data set, the three panels show the estimated difference between the signals measured by the three polarimetry channels at the two scan positions as a function of LST. The 24-bin data is shown here for clarity, but the final data analysis uses the 144-bin data.

probability of B occurring times the conditional probability that A will occur if B occurs. If A is a certain experimental data set, and B is some theoretical prediction, then this relation can be rewritten as

$$P(data)P(theory|data) = P(theory)P(data|theory).$$

Rearranging this expression yields the likelihood of the model prediction given the results of an experiment:

$$L = P(theory|data) = \frac{P(theory)}{P(data)} P(data|theory).$$

The results of the experiment are known ($P(data) = 1$), and the conditional probability of obtaining the measured data given a theoretical model is calculable (see below), but the prior probability of the model prediction $P(theory)$ is not a known or even a calculable quantity. Fortunately, there is a well-motivated, if unproven, postulate (known as Bayes' Postulate) which states that in the absence of any other information, $P(theory)$ can be assumed to be a constant (but see below, Lupton, 1993). With this assumption, $L \propto P(data|theory)$ and the likelihood can be computed.

For an illustrative example of a likelihood analysis, consider a simple univariate gaussian distribution. If the mean of this distribution is μ and the width is σ, and if a set of N values are drawn (independently) from this distribution, then the probability of obtaining a given set of values x_i is:

$$P(x_i|\mu, \sigma) \propto \frac{1}{\sigma^N} \prod e^{-\frac{(x_i-\mu)^2}{2\sigma^2}}.$$

What if we know the values of the x_i and we want to estimate the center and the width of the gaussian distribution they came from? In this case, we can use Bayes' postulate to get the likelihood:

$$L \propto \frac{1}{\sigma^N} \prod e^{-\frac{(x_i-\mu)^2}{2\sigma^2}}.$$

This is the exact same formula that gives $P(x_i|\mu\sigma)$, but P is a function of x_i, whereas L is a function of μ and σ. The maximum of L occurs where $\mu = \sum x_i/N$ and $\sigma = \sum(\sum x_i/N - x_i)^2/N$, so the likelihood function provides good estimates of these parameters.[1]

[1] The factor of N instead of $N-1$ in the expression for σ suggests that in order to use this technique to estimate σ, the prior probability should be $1/\sigma$ instead of a constant (Lupton, 1993).

Table 9.1: The ideal structure of the data vector **d**

Indices	Year	Channel
0-143		s0
144-287	PIQUE2000	s1
288-431		s2
432-575		s0
576-719	PIQUE2001	s1
720-863		s2

Furthermore, the errors on these estimates can be deduced from the shape of the likelihood function (Lupton, 1993).

9.1.2 Likelihood analysis of CMB polarization data

The PIQUE data are derived from a multivariate gaussian distribution, and the expectation value of every data point is zero because of the differencing scan strategy and the offset removal procedures (see Section 8). In this case, the probability of obtaining a certain data set is:

$$P(\mathbf{d}|\mathbf{C}) \propto \frac{1}{\sqrt{det(\mathbf{C})}} e^{-\frac{1}{2}\mathbf{d^T C^{-1} d}}.$$

In this expression, the data set is written as a vector $\mathbf{d}$ composed of elements d_i that represent the individual data points. The "covariance matrix" $\mathbf{C}$ is the multivariate generalization of the variance σ^2. The elements of this matrix are expectation values ($C_{i,j} =< d_i d_j >$) that depend on the system noise and the CMB power spectra. The likelihood function of the covariance matrix:

$$L \propto \frac{1}{\sqrt{det(\mathbf{C})}} e^{-\frac{1}{2}\mathbf{d^T C^{-1} d}},$$

therefore yields constraints on the structure of the CMB anisotropies.

Before the likelihood function can provide cosmological information, the elements of the covariance matrix must be written as explicit functions of the CMB power spectra. For PIQUE, the data vector $\mathbf{d}$ contains $144 * 3 * 2$ elements, each being the difference in a Stokes parameter measured at two points on the sky by one of the polarimetry channels (see Table 9.1). Each of these differences d_i includes a contribution from astronomical polarized signals, s_i, and a contribution from instrumental and atmospheric noise, n_i: $d_i = s_i + n_i$. The noise is completely uncorrelated with any signal from the sky, so $C_{i,j} =< d_i d_j >=< (s_i + n_i)(s_j + n_j) >=< s_i s_j > + < n_i n_j >= C^T_{i,j} + C^N_{i,j}$. The covariance matrix $\mathbf{C}$ can therefore be decomposed into two parts: a theory covariance matrix $\mathbf{C^T}$ and a noise covariance matrix $\mathbf{C^N}$.

The noise covariance matrix for PIQUE is extremely sparse. The noise is completely uncorrelated between the two observing seasons, so $< n_i n_j >$ is zero whenever n_i and n_j refer to different observing seasons. Furthermore, studies of the autocorrelation function demonstrate that the correlations between data taken at different times are very small (less than 1%). The expectation value $< n_i n_j >$ is therefore assumed to be zero when n_i and n_j refer to different LST bins. Therefore, the only nonzero elements of the noise covariance matrix are the diagonal elements and the elements which cross-correlate two polarimetry channels in the same LST bin. These elements can be estimated from the errors S_i and the correlation coefficients c_{ij} of the proccessed data set (see Chapter 8): $< n_i n_i >= S_i^2$ and $< n_i n_j >= c_{ij} S_i S_j$.

The theory covariance matrix is not sparse because the polarized component of the CMB has a nontrivial structure. Given the values of the Stokes parameters Q and U as a function of the spherical coordinates θ and ϕ, the s_i are $Q(\theta_E, \phi_E) + Q(\theta_W, \phi_W)$ for PIQUE2000[2] and $U(\theta_E, \phi_E) - U(\theta_W, \phi_W)$ for PIQUE2001 (the E and W refer to the east and west scan positions, respectively). The elements of the theory covariance matrix $< s_i s_j >$ therefore consist of linear combinations of the correlation functions $< Q(\theta_x, \phi_x) Q(\theta_y, \phi_y) >$, $< Q(\theta_x, \phi_x) U(\theta_y, \phi_y) >$ and $< U(\theta_x, \phi_x) U(\theta_y, \phi_y) >$.

The correlation functions of the Stokes parameters are expressed in terms of power spectra by decomposing the polarization anisotropies into spherical harmonics. The polarization has both a magnitude and an orientation, so the relevant harmonic functions are not the familiar Y_{lm}. If Q and U are defined with respect to the poles of the celestial sphere (i.e.

[2]The chopped data for PIQUE2000 is a sum and not a difference due to the orientation of the detection axes with respect to the NCP are different at the two scan positions during this observing season (see Figures 2.2 and 2.5).

the IAU convention, see Section 2.1), then (Zaldarriaga, 1998):

$$Q(\theta,\phi) = -\sum_{lm} a_{E,lm}\sqrt{\frac{(2l+1)}{4\pi}}F_{1,lm}(\theta)e^{im\phi} + ia_{B,lm}\sqrt{\frac{(2l+1)}{4\pi}}F_{2,lm}(\theta)e^{im\phi},$$

$$U(\theta,\phi) = -\sum_{lm} a_{B,lm}\sqrt{\frac{(2l+1)}{4\pi}}F_{1,lm}(\theta)e^{im\phi} - ia_{E,lm}\sqrt{\frac{(2l+1)}{4\pi}}F_{2,lm}(\theta)e^{im\phi}.$$

The harmonic functions $F_{(1,2)lm}$ are:

$$F_{1,lm} = 2\sqrt{\frac{(l-2)!(l-m)!}{(l+2)!(l+m)!}}\left[-\left(\frac{l-m^2}{\sin^2\theta}+\frac{1}{2}l(l-1)\right)P_l^m(\cos\theta)+(l+m)\frac{\cos\theta}{\sin^2\theta}P_{(l-1)}^m(\cos\theta)\right],$$

$$F_{2,lm} = 2\sqrt{\frac{(l-2)!(l-m)!}{(l+2)!(l+m)!}}\frac{m}{\sin^2\theta}\left[-(l-1)\cos\theta P_l^m(\cos\theta)+(l+m)P_{(l-1)}^m(\cos\theta)\right],$$

where P_l^m are the usual Legendre Polynomials.

The harmonic decomposition of the Stokes parameters includes two different harmonic series with two sets of coefficients, $a_{E,lm}$ and $a_{B,lm}$, which correspond to the curl-free and gradient free components of the polarization field, respectively (these are also called the "E-modes" and "B-modes"). The correlation functions of Q and U are linear combinations of the correlation functions of these coefficients. Assuming statistical isotropy, the correlation functions of the a_{lm} are:

$$< a_{E,lm}a_{E,l'm'} >= \delta_{ll'}\delta_{mm'}C_{E,l},$$

$$< a_{B,lm}a_{B,l'm'} >= \delta_{ll'}\delta_{mm'}C_{B,l},$$

$$< a_{E,lm}a_{B,l'm'} >= 0.$$

The parameters $C_{E,l}$ and $C_{B,l}$ specify the size of the variations as a function of angular scale (multipole moment), and they are directly related to the power spectra $\delta P_l^2 = (l(l+1)/2\pi)C_{(E,B)l}$ illustrated in Figure 1.6 (Zaldarriaga, 1998; Seljak and Zaldarriaga, 1996).

The correlation functions of Q and U, as defined by the IAU (see Section 2.1) are linear combinations of $C_{E,l}$ and $C_{B,l}$, but the prefactors are complicated functions of θ and ϕ. However, these Stokes parameters are linear combinations of the "natural" Stokes parameters $\mathcal{Q}$ and $\mathcal{U}$, which have relatively simple correlation functions. The parameters $\mathcal{Q}$ and $\mathcal{U}$ are defined with respect to the great circle connecting the two points under consideration (if the polarization is directed along this great circle, then $\mathcal{Q}$ is positive and $\mathcal{U}$ is zero). The correlation functions of these Stokes parameters depend only the angular distance β between the two points (Zaldarriaga, 1998):

$$< \mathcal{Q}(0,0)\mathcal{Q}(\beta,0) >= \sum_l \frac{2l+1}{4\pi} B_l^2 [C_{E,l} F_{1,l2}(\beta) - C_{B,l} F_{2,l2}(\beta)],$$

$$< \mathcal{U}(0,0)\mathcal{U}(\beta,0) >= \sum_l \frac{2l+1}{4\pi} B_l^2 [C_{B,l} F_{1,l2}(\beta) - C_{E,l} F_{2,l2}(\beta)],$$

$$< \mathcal{Q}(0,0)\mathcal{U}(\beta,0) >= 0.$$

The factor $B_l = exp[-l(l+1)\sigma_b]$ in these expression accounts for the finite beamwidth of the telescope. For the W-band system, $\sigma_b = 0.1°$ (see Chapter 6)

The correlation functions of Q and U, and consequently the entries in the theory covariance matrix, are linear combinations of the above expressions. Every element of $\mathbf{C^T}$ consequently has the form $\frac{1}{4\pi}\sum[(2l+1)(W_{E,l}C_{E,l} + W_{B,l}C_{B,l})]$. The window functions $W_{(E,B),l}$ are calculable functions that determine the range of angular scales probed by the experiment (see Figure 9.2). The effective multipole moments of the $W_{E,l}$ (see below) are $< l >= 211^{+261}_{-146}$ for PIQUE2000 and $< l >= 177^{+219}_{-134}$ for PIQUE2001.[3] The W-Band system on PIQUE is therefore sensitive to angular scales around half a degree.

9.2 Preliminary Results

The preliminary analysis of the W-band data discussed here uses a conmbination of FORTRAN programs and IDL routines to calculate the likelihoods.[4] This program assumes flat power spectra for the E-mode polarization ($\delta P_E = \sqrt{2\pi C_{E,l}/l(l+1)}$ is a constant in l) and zero B-mode polarization ($C_{B,l} = 0$). Nonzero B-modes will be considered before the final analysis is complete. The flat band power assumption simplifies the calculations

[3]The effective multipole moment is computed using the formula $< l >\simeq \sum W_l / \sum W_l/l$ (Steinhardt, 1995). The upper and lower limits indicate where the sum of W_l/l from 2 to l is 10% or 90% of the total sum.

[4]These programs were written by S.T. Staggs, and they have been checked for programming and arithmetic errors using two other, independent FORTRAN programs: one written by J.O. Gundersen and the other managed by B. Winstein

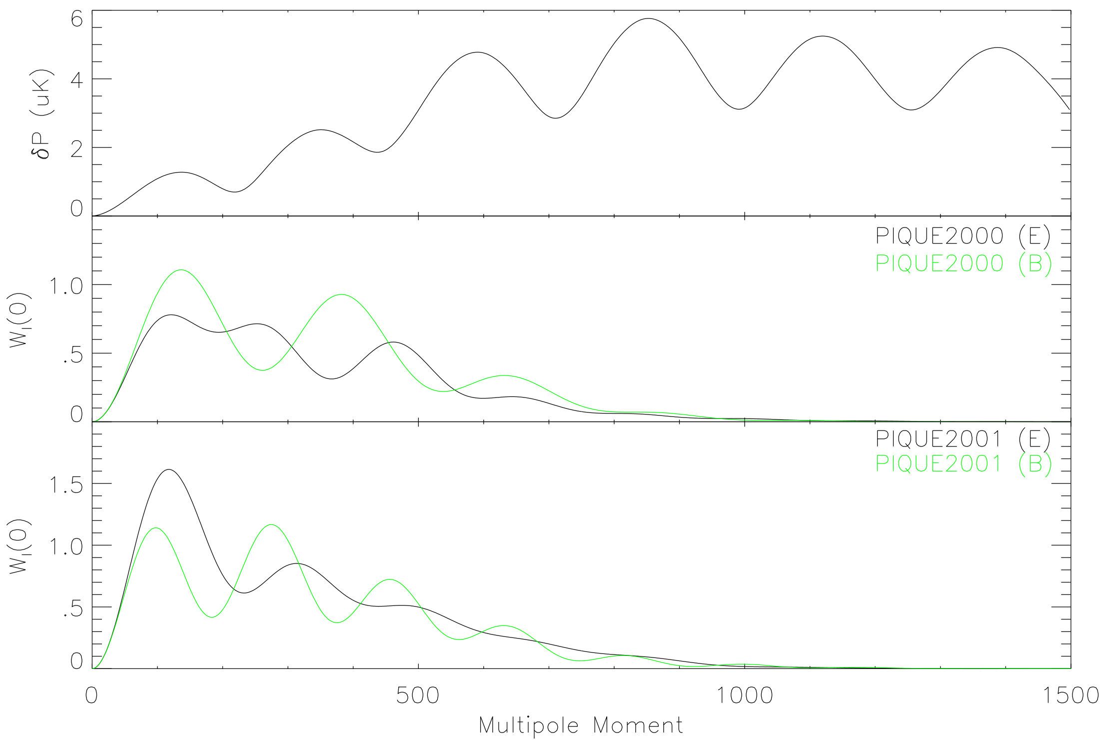

Figure 9.2: The window functions of the PIQUE W-Band system for the two observing seasons, presented with the power spectrum of the polarized signals for comparison. These are the "zero-lag" window functions, i.e. they are the window functions appropriate to the diagonal elements of the covariance matrix

and facilitates comparisons with other experiments and theoretical predictions. Assuming purely E-type polarization is reasonable because the bulk flows in the early universe that are responsible for most of the polarized radiation from the CMB (see Chapter 1)) did not generate B-mode polarization (the relevant gravitational and pressure forces could not generate significant vorticity in the primordial plasma). Primordial gravity waves and various secondary effects like lensing can give rise to B-modes, but the expected signal is an order of magnitude smaller than the signal from the E-modes[5] (Kodama and Sasaki, 1984; Kamnionkowski et al., 1997; Seljak and Zaldarriaga, 1997).

[5]The polarized emission from foregrounds could also have a substantial B-mode component, but foregrounds should be sub-dominant to the CMB at 90 GHz (see Figure 1.10)

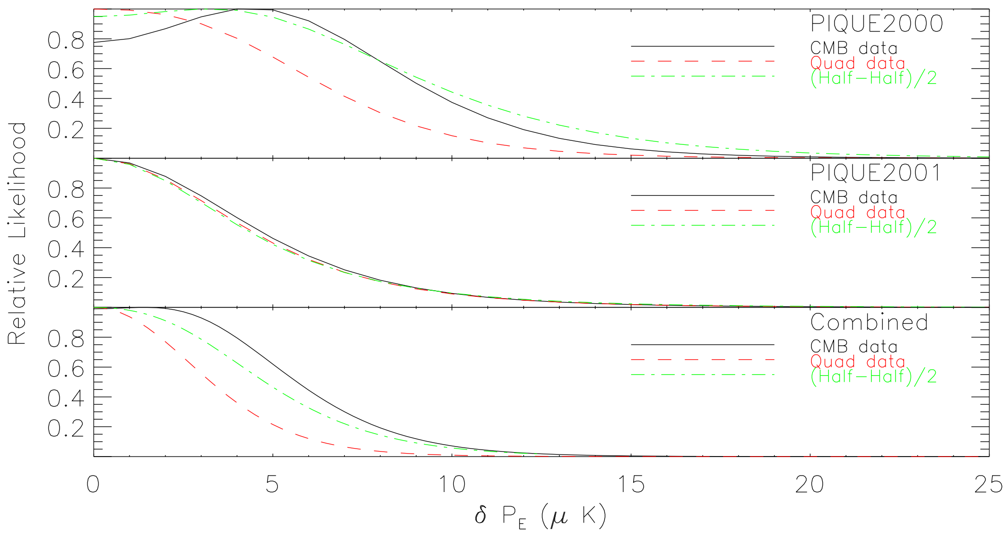

Figure 9.3: Preliminary likelihood functions of δP_E for the nominal data set assuming no B-modes. Note that the half-half data are divided by two so that the error bars are comparable to those of the other data sets.

Figure 9.3 shows the likelihoods of δP_E for the nominal data set. The PIQUE2000 likelihood peaks slightly away from zero, but the relative likelihood of zero is around 0.8, so this cannot be considered evidence for cosmological polarization. In any case, the combined likelihood peaks near zero. The 95% confidence limit on δP_E corresponds to the point on the likelihood curve which bounds 95% of the total (integrated) likelihood, which is 8μK for the combined data set (see Table 9.2). This limit is reasonable given our previous result of 10 μK (Hedman et al., 2001), which was obtained with about half the data[6] Furthermore, the likelihoods of the quadrature and half-half nulled data sets[7] have similar shapes and yield similar limits. These nulled likelihoods not only confirm that the limit for the CMB data is not unusally high or low, but also provide a further check on the data quality.

The systematic errors on these limits are less than 20% (see Table 9.3), and are dominated by the 10% uncertainty in the calibration of the polarimetry channels (see Section 6.3).

[6]Note that the PIQUE2000 limit is higher than the limit given in Hedman et al. (2001), which may indicate that the data was overcut in Hedman et al. (2001). However, the difference between these limits is also consistent with statistical fluctuations.

[7]Similar nulled likelihoods cannot be calculated for the interchannel differences and the pattern nulls because these data sets do not have the same structure as the CMB data, see Section 8.3.1.

Table 9.2: Preliminary 95% Confidence Limits on the curl-free flat-band power δP_E, assuming $\delta P_B = 0$

Data Set	95% Confidence Limits (μK)		
	PIQUE2000	PIQUE2001	Combined
CMB data	13	10	8
Quad. data	10	10	6
(Half-Half)/2[a]	15	11	8

[a] The half-half data are divided by two so the error bars are comparable to those of the other two data sets

Additional contributions come from uncertainties in the beamsize. While the planetary observations yield beamsize measurements accurate to better than 3% (see Section 6.1) the spread in encoder values at the two scan positions (see Figure 8.2) and the encoder slip during PIQUE2001 (see Section 6.1) smear the beam slightly. Assuming, conservatively, that the beamsize is uncertain by 5% yields a comparable error in the limit. Various simplifying assumptions (e.g. neglecting pointing offsets and assuming $T_{surface} = 273$K throughout the observing seasons) in the data analysis routines do not contribute significantly to the systematic uncertainty in the limit.

The expected polarized signal from the CMB is $1 - 2\mu$K, while polarized radiation from foregrounds might contribute another 0.7μK (see Table 9.4). Our limit is consistent with this expected signal level. Furthermore, the limit is only a factor of two or so smaller than the expected signal, so PIQUE is approaching the sensitivity needed to make a clear measurement of the polarized component of the CMB.

Table 9.3: Systematic Uncertainties in the PIQUE results

Systematic	Error
Calibration Uncertainty	10%
Beamsize Uncertainty	4%
Pointing Uncertainty	2%
Variations in $T_{surface}$	1%

Table 9.4: Estimates of Foreground Contamination near the NCP for PIQUE

Foreground	RMS Contamination at 90 GHz (μK)
Free-Free	0.14
Synchrotron[a]	0.36
Vibrating Dust[b]	0.53
Spinning Dust	0.04
Point Sources	0.10

These estimates come from Gundersen (2000)

[a] Mesurements of fluctuations in the synchrotron brightness around the NCP exist at 408 MHz (Haslam et al., 1982) and at 1.411 GHz (Brouw and Spoelestra, 1976), although the latter measurement has considerably coarser resolution than PIQUE. Assuming the synchrotron emission is completely polarized and that the brightness temperature scales with frequency as ν^{β}, with $\beta = -2.9$, the predicted rms polarization due to synchrotron is about 0.36μK at 90 GHz.

[b] The brightness fluctuations of the dust emission are recorded in the IRAS 100 μm map. Convolving these data with the PIQUE beams and using the correlation between the IRAS map and COBE measurements, the brightness fluctuations of dust at 90 GHz are estimated to be about 10μK rms. Assuming 5% polarization (Lazarian, 2000), the polarized signal from dust should be about 0.5μK at 90 GHz.

Chapter 10

Conclusion

The data from the PIQUE W-band polarimeter provide one of the most constraining measurement of the cosmological polarization to date (see Figure 10.1). This is a significant acheivement, especially considering that it comes from a rather modest telescope observing from a less-than-perfect site. It demonstrates that correlation polarimeters have the sensivity, and, more importantly, the stability to measure signals approaching the size of the polarized component of the CMB. Also, this measurement permits continued optimism that the polarized foregrounds will not overwhelm the cosmological signal, at least around 90 GHz. All this bodes well for the future of the CMB polarization experiments, and it should not be long before someone detects the polarized component of the CMB, and opens a new window onto the early universe.

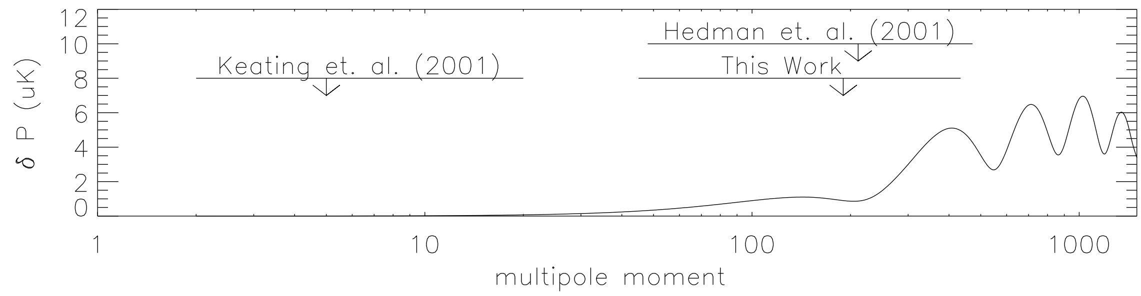

Figure 10.1: Preliminary new PIQUE 95% confidence limit plotted with the expected signal from a ΛCDM model. The limit from POLAR (Keating et al., 2001a) and the previous PIQUE limit (Hedman et al., 2001) are also shown.

Appendix A

Multiplier

The multipliers used in PIQUE are double-balanced mixers. The double-balanced configuration (see Figure A.1) treats the RF and LO ports symmetrically, so the two input fields can be of comparable strength, which is desirable for this application (other mixer designs may be optimized to work with an LO signal much stronger than the RF signal).

A double-balanced mixer multiplies the fields entering its RF and LO ports using a ring of four diodes as and two transformers or stripline couplers (Figure A.1). The oscillating electric fields at the RF and LO inputs produce oscillating voltages V_1 and V_2 across one side of each transformer. These voltages in turn drive currents around the diode ring, producing the output voltage V_o. Assuming the transformers split the signal equally across the output points, then the voltages across each diode are as shown in Figure A.1. The currents flowing through each diode can then be calculated using the current-voltage relation for diodes $I = I_S(e^{AV} - 1)$, (A and I_S are constants). Assuming no current flow at the outputs, the current flow through the different diodes must obey the relation $I_W + I_Y = I_X + I_Z$. This relation yields a complex expression relating V_o to V_1 and V_2.

If the diodes are identical (i.e., have the same values of I_S and A), and the transformers split the signals properly, then, to second order, this relation simplifies to :

$$V_o = gV_1V_2,$$

where $g = A/2$ is the responsivity of the multiplier. Note that terms linear in V_1 and V_2 cancel out due to the symmetry of the double-balanced configuration. Taking the time

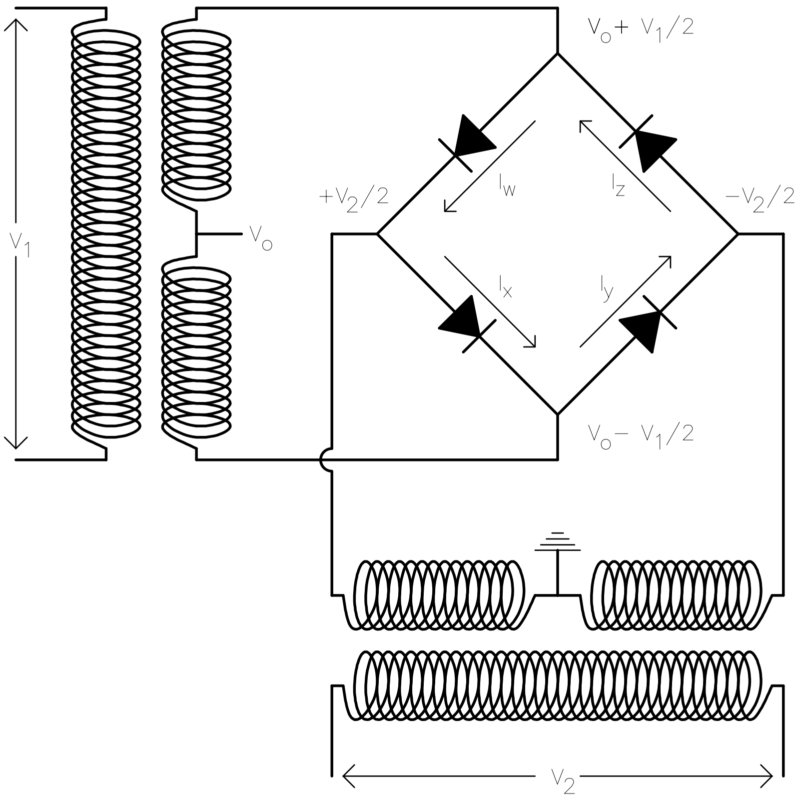

Figure A.1: The circuit of a double-balanced mixer average of both sides of this equation yields the correct expression for the output DC voltage from the multiplier.

Departures from this ideal behavior occur if the diodes are not identical (i.e. have different values of A and I_S). In this case, the cancellation between diodes is not exact, and additional terms linear and quadratic in V_1 and V_2 appear. Upon taking the time average, terms linear in V_1 and V_2 drop out, leaving the following expression for the DC output voltage:

$$< V_o >= g < V_1 V_2 > +a < V_1^2 > +b < V_2^2 >,$$

where a and b are combinations of various diode parameters. Asymmetries in the voltage splitting across the transformers can change the values of a and b, but cannot generate nonzero a and b by themselves. Measured values of a and b (obtained by injecting signal into only one arm of the multiplier) indicate that these parameters are about an order of magnitude smaller than g for the devices used in PIQUE.

For the multipliers used in PIQUE, there is evidence that the nonzero a and b are not due to diode asymmetries, as suggested above, but instead due to internal reflections within the device. Measurements of the response of the multiplier to narrow-band signals injected into one arm of the device show a strong resonance near 10 GHz (Barkats, 2000). Such a resonance is typical of internal reflections, but is inconsistent with the above model of asymmetric diodes, which predicts a frequency independent response.

As noted in Section 3.4.2, the multipliers saturate when the uncorrelated noise power is too high. This was an unexpected result because early tests of the multipliers did not inject sufficient broad-band noise into the multipliers, and because early efforts to model the multipliers had made a serious error in assuming the cubic terms in the expansion average to zero.

The correct third order relationship between V_o, V_1 and V_2 is (assuming identical diodes):

$$V_o(1 + \frac{1}{2}A^2(\frac{4}{3}V_o^2 + V_1^2 + V_2^2)) = \frac{1}{2}AV_1V_2$$

Note that the third order terms are proportional to V_o^3, $V_oV_1^2$ or $V_oV_2^2$. Since V_o is roughly constant, these terms do not average to zero. Taking the time average and rearranging this expression slightly yields:

$$< V_o > \simeq \frac{A}{2 + A^2(\frac{1}{3}A^2 < V_1V_2 >^2 + < V_1^2 > + < V_2^2 >)} < V_1V_2 >$$

(where $< V_o > \simeq \frac{1}{2}A < V_1V_2 >$ has been used in the denominator). If the entire factor in front of $< V_1V_2 >$ is treated as the input-dependent responsivity $g(V_1, V_2)$, the saturation properties of the multiplier become clear. As expected, if the correlated signal $< V_1V_2 >$ gets large, then the responsivity of the multiplier drops. Furthermore, the terms involving V_1^2 and V_2^2 mean that if the raw uncorrelated power entering the radiometer is too large, the responsivity of the device will be affected, which precisely the behavior observed (Figure 3.15).

Appendix B

Components of the W-Band Radiometer

Table B.1: Components in W-Band Radiometer

Component	Manufacturer	Part No.	Serial Numbers[a]
OMT	γ-f	11590	n/a
HEMT	NRAO	n/a	53 (v) 52 (h)
Filter	MRI	FRWS-94	101 (v) 102 (h)
Mixer	Spacek Labs	P94-10-LN	7H11 (v) 7H12 (h)
LO	Millimeter-Wave Oscillator Co.	n/a	n/a
Phase Switch	Pacific Millimeter	W180	001 (v) 003 (h)
Phase Tuner	Millitech	VPS-10-R0000	019 (v)
Cold IF Amp (2000 only)	DBS Microwave	DB97-0421	001 (v) 002 (h)
Warm IF Amp (2000)	Veritech Microwave	VMA S1473	002 (v) 001 (h)
Warm IF Amp (2001)	DBS Microwave	DB97-0421	001 (v) 002 (h)
Splitter	MAC Technology	P8248-2	497003 (D0) 497001 (D1)
Detector Diode	Hewlitt-Pacard	HP 8472B	1822A10358 (D0) 1822A10357 (D1)
Filter Bank	Reactel	3SM-7/12.7-10PM	97-1A (D0) 97-1B (D1)
Line Stretchers	M/A-COM	2054-6002-00	n/a
Multipliers	Miteq	DB0218LW2	475783 (s0) 475784 (s1) 475676 (s2)

[a] The terms in parentheses define the location of the specific components: (v/h) indicates the component is in the line carrying the vertical or horizontal component of the electric field, respectively. (D0/D1) indicates the component is in the line connected to the named detector diode. This differs from the v/h notation because of the shift in topology imposed by the removal of the cold IF amplifiers after the first observing season (see section 3.1.3) (s0/s1/s2) indicates the component is in the named frequency band.

Appendix C

Circuit Diagrams

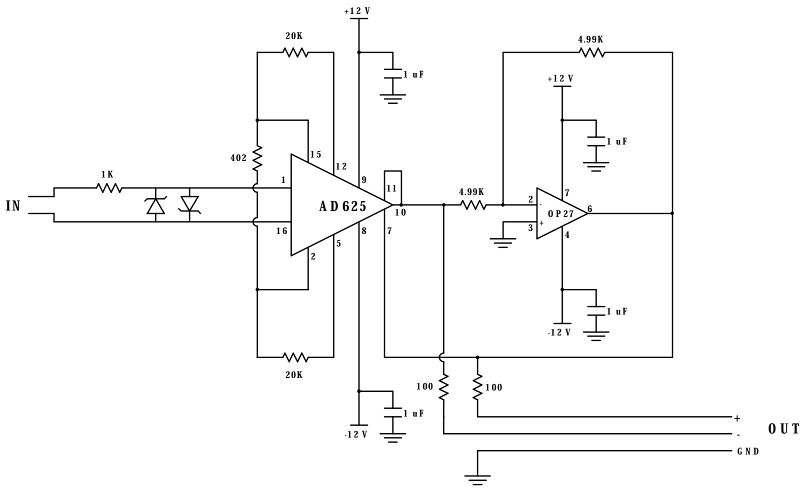

Figure C.1: Pre-Amplifier (attached directly to multipliers and detector diodes.)

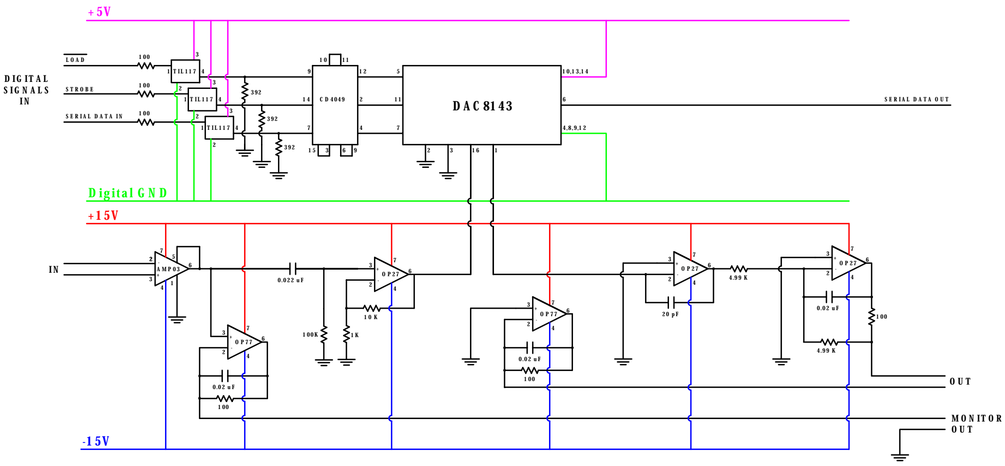

Figure C.2: Programmable-gain amplifier

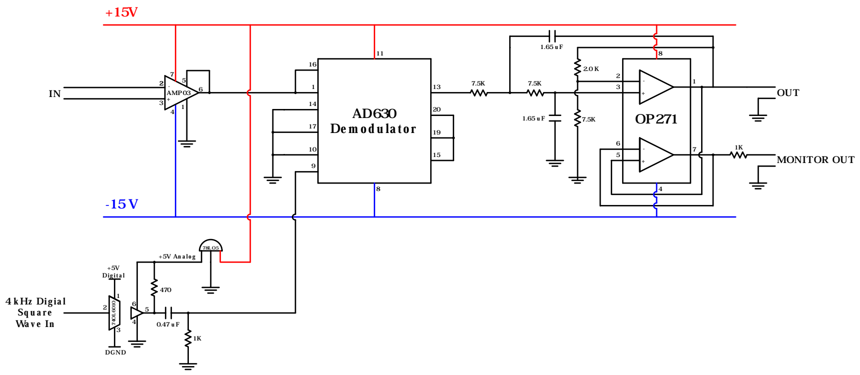

Figure C.3: Demodulator and low-pass filter for multiplier channels

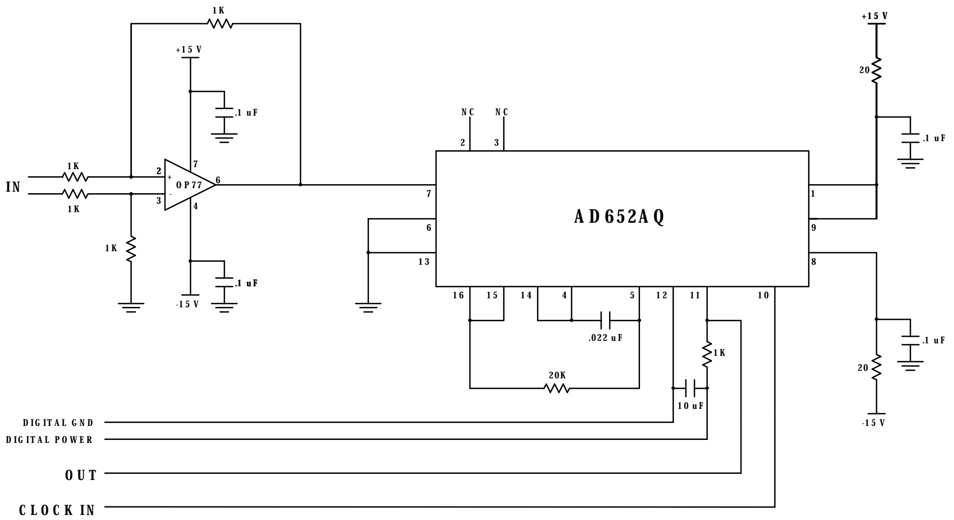

Figure C.4: Voltage-to-frequency converter

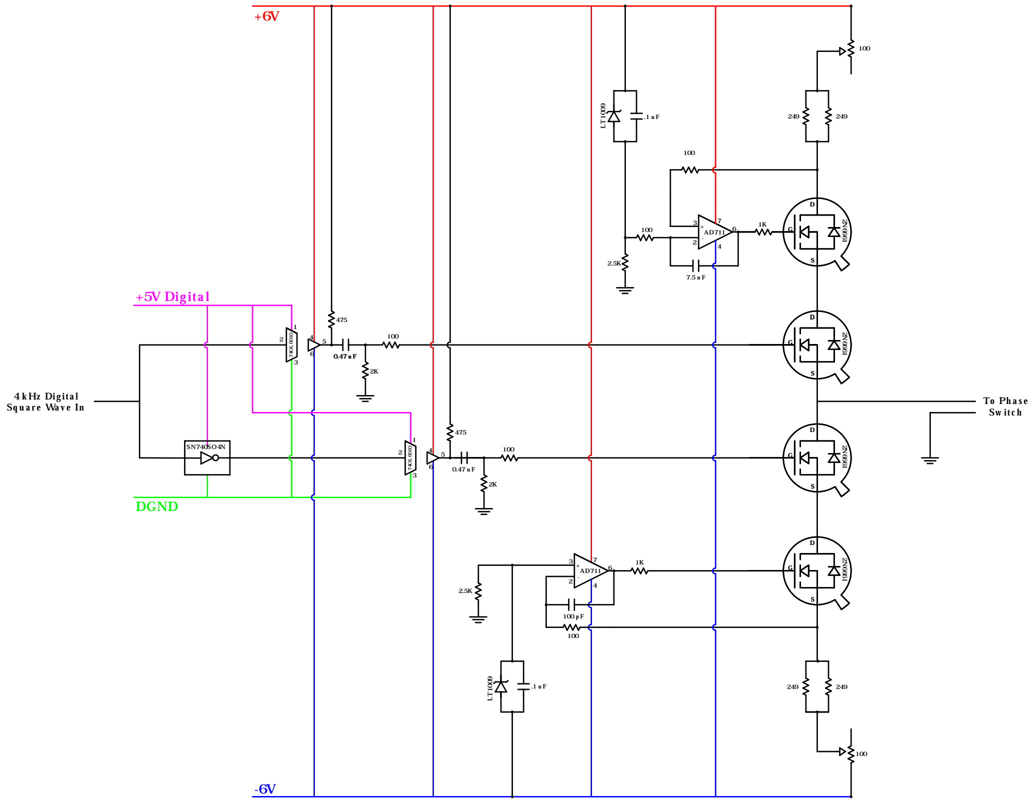

Figure C.5: Phase-switch driver circuit

Appendix D

Complete Description of W-Band Horn

This is a copy of the data sent to Custom Microwave, who built the horn. The interior surface of the horn is a series of cylinders. The first column gives the radius of each cylinder, while the second column specifies the location of the distal edge of each cylinder along the horn axis.

radius (inches)	height (inches)	radius (inches)	height (inches)	radius (inches)	height (inches)	radius (inches)	height (inches)
0.0585	0.7550	0.0935	1.3270	0.1285	1.8990	0.1635	2.4710
0.1265	0.7605	0.1279	1.3560	0.1629	1.9280	0.1979	2.5000
0.0612	0.7990	0.0962	1.3710	0.1312	1.9430	0.1662	2.5150
0.1265	0.8063	0.1306	1.4000	0.1656	1.9720	0.2006	2.5440
0.0639	0.8430	0.0989	1.4150	0.1339	1.9870	0.1688	2.5590
0.1265	0.8521	0.1333	1.4440	0.1683	2.0160	0.2033	2.5880
0.0666	0.8870	0.1016	1.4590	0.1365	2.0310	0.1715	2.6030
0.1265	0.8979	0.1360	1.4880	0.1710	2.0600	0.2060	2.6320
0.0693	0.9310	0.1043	1.5030	0.1392	2.0750	0.1742	2.6470
0.1265	0.9437	0.1387	1.5320	0.1737	2.1040	0.2087	2.6760
0.0720	0.9750	0.1069	1.5470	0.1419	2.1190	0.1769	2.6910
0.1265	0.9895	0.1414	1.5760	0.1764	2.1480	0.2114	2.7200
0.0747	1.0190	0.1096	1.5910	0.1446	2.1630	0.1796	2.7350
0.1265	1.0353	0.1441	1.6200	0.1791	2.1920	0.2141	2.7640
0.0773	1.0630	0.1123	1.6350	0.1473	2.2070	0.1823	2.7790
0.1265	1.0811	0.1468	1.6640	0.1818	2.2360	0.2167	2.8080
0.0800	1.1070	0.1150	1.6790	0.1500	2.2510	0.1850	2.8230
0.1265	1.1269	0.1495	1.7080	0.1845	2.2800	0.2194	2.8520
0.0827	1.1510	0.1177	1.7230	0.1527	2.2950	0.1877	2.8670
0.1265	1.1727	0.1522	1.7520	0.1871	2.3240	0.2221	2.8960
0.0854	1.1950	0.1204	1.7670	0.1554	2.3390	0.1904	2.9110
0.1265	1.2185	0.1549	1.7960	0.1898	2.3680	0.2248	2.9400
0.0881	1.2390	0.1231	1.8110	0.1581	2.3830	0.1931	2.9550
0.1265	1.2643	0.1575	1.8400	0.1925	2.4120		
0.0908	1.2830	0.1258	1.8550	0.1608	2.4270		
0.1265	1.3101	0.1602	1.8840	0.1952	2.4560		

Appendix E

Parametrized Description of Dish for DADRA

The physical optics program DADRA uses two following input files to specify the location and orientation of the dish with respect to the horn. The syntax of these files is given in Rahmat-Samii et al. (1996b).

The dish itself is parametrized in the file REFLECT.DAT:

```
TYPE                        BOUNDARY
PARABOLA                    CIRCLE
NBPTS
0
SURFACE DATA
F              D            X0             Y0             Z0
75.0           122.0        -78.09         0.0            0.0
```

While the orientation and location of the horn are specified in the file DADRA.DAT:

```
FREQ        UNITS
90.0        CM
XS          YS          ZS
0.0         0.0         0.0
SCALEM      BLKRAD
1.0         0.0
MFILE
notused
NDXM        NDYM        XSYMM       YSYMM   POSNZM
2           2           F           F       T
IFLAGM      SIZEM
-1          1.0
READJE      SAVEJE      READJS      SAVEJS
F           F           F           F
XF          YF          ZF
-78.09      0.0         0.0
ALPHAF      BETAF       GAMMAF
90.0        127.9       0.0
SCALES
1.0
SFILE
REFLECT.DAT
NDX         NDYS        XSYMS       YSYMS   POSNZS
100         100         F           F       T
IFLAGS      SIZES       LINPSE
-1          1.0         F
CIRPOL      POLX        LHCP        COSQ    SWE     FILE19
F           F           T           F       T       F
GAIN        QE          QH
0.0         8.75        8.75
HFILE
FILE27.SPW
RR
99999.
PHII        PHIF        DPHI
0.0         180.0       90.0
THETAI      THETAF      DTHETA
-1.0        1.0         0.01
THTBPD      PHIBPD
0.0         0.0
PHSXC       PHSYC       PHSZC
0.0         0.0         0.0
PATNRM
PEAK
POLARZ
LUDWIG
HFIELD      FARFIELD    ROTATE      ADDSUB  NOMAIN  ADDFED
F           F           T           T       T       T
PRTPAT      PLTPAT      PLTPHS      WRIT27  XYCALC  ANGLES
F           T           T           F       F       F
HFMAGN(1)   HFMAGN(2)   HFMAGN(3)
1.0         1.0         1.0
HFPHN(1)    HFPHN(2)    HFPHN(3)
0.0         0.0         0.0
XMAPMIN     XMAPMAX     NMAPX
-2.5        2.5         101
YMAPMIN     YMAPMAX     NMAPY
-2.5        2.5         101
ZMAP
64
```

Appendix F

Data Structure

PIQUE data is stored in files nominally containing 3600 'frames'. Some files were closed early and therefore contain fewer frames. Also, early in the second observing season there was a problem with the file closing algorithm, so some files contain many more frames.

A frame is written to file once every half second, and each frame consisted of 512 bytes or 254 words. The first 192 words are 32 sets of six numbers: the encoder value and the five polarimeter outputs (s0,s1,s2,D0,D1). These numbers were recorded at a rate of 62.5 Hz. The remaining 62 words include a variety of housekeeping data:

- Weather station data (8 words): Humidity, air temperature, wind direction, wind speed, pressure, pressure change, and luminosity

- Buffered voltages (8 words): The current draw and temperatures of various temperature-regulating heater circuits

- E-Box temperatures (2 words): The output of two AD592 thermometers located inside the E-Box (one near the IF amplifiers, the other near the LO)

- External Temperatures (12 words): The outputs of twelve AD592 and LM335 thermometers attached to various points on the frame, dish, and inner ground screen (not used in PIQUE2001)

- Clinometers (2 words): The output of two clinometers attached to the base.

- External Temperatures (6 words): The outputs of an addition six thermometers attached to various external surfaces of the experiment.
- Liquid Nitrogen (Stage I) Thermometers (3 words): The outputs of three Lakeshore diode cryogenic thermometers attached either to both cold IF amplifiers and a filter (PIQUE2000), or to both filters and a mixer (PIQUE2001).
- Liquid Helium (Stage II) Thermometers (4 words): The outputs of four Lakeshore cryogenic thermometers attached to the horn, cold plate, and the HEMTs
- Load Thermometer Data (6 words): The temperatures, set points and current draw through the cryogenic regulated loads used in the polarimeter tests (see section 3.3). During the observing seasons these loads were stowed on Stage II of the refrigerator.
- GPS Date and Time (11 words)

Appendix G

Statistical Effects Relevant to PIQUE Data Reduction

Simulations of the PIQUE data revealed two statistical effects that must be properly taken into account to insure the error bars on the binned data are properly estimated. If these effects are neglected, the error bars for the 144-bin data are underestimated by 4%, which in turn means χ^2 statistics are overestimated by 8%.

G.1 Background

Both effects are relatively subtle small-number phenomena. Therefore, to avoid confusion, the notation used in this analysis is laid out here:

- The data are assumed to be drawn from a guassian distribution with mean μ and variance σ^2.
- The data points drawn from this distribution are denoted x_i.
- The true variance of these data points is σ^2.
- A sample or set of data points is drawn to estimate the mean and variance.
- n is the number of data points in the sample.

- $\bar{x} = \frac{\sum x_i}{n}$ is the sample mean.
- The true variance of $\bar{x}$ is $\bar{\sigma}^2 = \sigma^2/n$.
- $s^2 = \frac{\sum(x_i-\bar{x})^2}{n}$ is here referred to as the **naive sample variance**.
- s is called the **naive standard deviation**.
- $S^2 = \frac{\sum(x_i-\bar{x})^2}{n-1}$ is called the **normal sample variance**.
- S is called the **normal standard deviation**.

G.2 Unbiased Estimators and Weights

If normally distributed data are broken into subsets, and the mean and (normal) standard deviation is calculated for each subset (e.g. when the means and errors of the "daily" bins are calculated), then the χ^2 of those means and standard deviations to a flat line is higher than the expected numbers of degrees of freedom. This occurs because the standard deviation is not the correct unbiased estimator for the weight in the χ^2 statistic.

The subject of biased estimators requires some introduction. The data points x_i are drawn from a gaussian distribution with mean μ and width σ. The parameters μ and σ are usually unknown, and the whole point of measuring x_i is to estimate these parameters. A good estimator should be **unbiased**, which means that the expectation value of the estimator (i.e., the average value of the estimator over many trials) equals the parameter to be estimated.

For example, the sample mean $\bar{x}$ is an unbiased estimator of μ because:

$$<\bar{x}>=<\frac{1}{n}\sum x_i>=\frac{1}{n}\sum<x_i>=\frac{1}{n}\sum\mu=\mu.$$

The classic expample of a biased estimator is s^2, which one might naively guess is an estimator of the true variance σ^2. In fact (Bevington, 1969; Cowden, 1957):

$$<s^2>=\frac{n-1}{n}\sigma^2.$$

Therefore s^2 will, on average underestimate the true variance. On the other hand, $< S^2 >= \sigma^2$, so the normal sample variance is the correct, unbiased estimator of σ^2.

For a χ^2 test, we want an unbiased estimate of $1/\sigma^2$. While S^2 is an unbiased estimator of σ^2, $1/S^2$ is **not** an unbiased estimator of $1/\sigma^2$. (This is not as weird as it might sound, because S^2 and $1/S^2$ have differently shaped distribution functions.)

The correct estimator for arbitrary powers of σ is given on page 78 of Cowden (1957). For convenience, many steps in his derivation are copied here. Note that everything is done in terms of s not S (this keeps factors of $n-1$ from cluttering the equations).

The probability distribution function for s^2 given a sample size n is:

$$f(s) = \frac{n^{(n-1)/2}}{2^{(n-3)/2}\Gamma(\frac{n-1}{2})\sigma^{n-1}} e^{-ns^2/2\sigma^2} s^{n-2}.$$

The expectation value of s^r is therefore:

$$< s^r >= \int_0^\infty s^r f(s) ds = \frac{n^{(n-1)/2}}{2^{(n-3)/2}\Gamma(\frac{n-1}{2})\sigma^{n-1}} \int_0^\infty e^{-ns^2/2\sigma^2} s^{n+r-2} ds.$$

Performing the substituions $a = n/2\sigma^2$ and $y = (n+r-1)/2$ yields:

$$< s^r >= \frac{2a^{(n-1)/2}}{\Gamma(\frac{n-1}{2})} \int_0^\infty e^{-as^2} s^{2y-1} ds.$$

Using the definition for the Gamma Function $\Gamma(y) = \int_0^\infty e^{-z} z^{y-1} dz$, with the substitution $z = as^2$, one can show:

$$\int_0^\infty e^{-as^2} s^{2y-1} ds = \frac{\Gamma(y)}{2a^y}.$$

Substituting this in gives:

$$< s^r >= \frac{\Gamma(y)}{\Gamma(\frac{n-1}{2})} \frac{a^{(n-1)/2}}{a^y}.$$

Finally, subsitiuting in for the values of y and a yields

$$< s^r >= \frac{\Gamma(\frac{n+r-1}{2})}{\Gamma(\frac{n-1}{2})} (\frac{n}{2\sigma^2})^{-r/2}.$$

Using this expression, we can find the unbiased estimator for any power of σ (indeed, using $r = 2$ gives $< s^2 >= \frac{n-1}{n}\sigma^2$). We are concerned here with the case $r = -2$, in which case:

$$\left\langle \frac{1}{s^2} \right\rangle = \left(\frac{n}{n-3}\right)\frac{1}{\sigma^2}.$$

The unbiased estimator of $1/\sigma^2$ is therefore

$$\left(\frac{n-3}{n}\right) 1/s^2 = \frac{n-3}{\sum(x_i - \bar{x})^2} = 1/\varsigma^2,$$

where ς is a **recripocal standard deviation**.

Summary/Moral: In situations where the **reciprocal** of the variance is needed (e.g. for the weight in a χ^2 statistic) , the correct, unbiased estimator is **not** the reciprocal of the normal standard deviation $1/S^2 = \frac{n-1}{\sum(x_i-\bar{x})^2}$; it is instead $1/\varsigma^2 = \frac{n-3}{\sum(x_i-\bar{x})^2}$.

G.3 The Error on a Weighted Mean

A weighted mean is given by the expression (which can be derived as a maximum likelihood estimator, as on page 69 of Bevington (1969)):

$$Z = \frac{\sum x_i/\sigma_i^2}{\sum 1/\sigma_i^2},$$

where x_i are a set of data points and σ_i are the known width of the gaussian each point was pulled from. The error on this quantity is derived using normal error propagation (as is done by Bevington on pg 70-71):

$$\sigma_Z^2 = \sum_i \sigma_i^2 \left(\frac{dZ}{dx_i}\right)^2 = \sum_i \sigma_i^2 \left(\frac{1}{\sum_k \sigma_k^2} \sum_k \frac{dx_k/dx_i}{\sigma_k^2}\right)^2,$$

$$\sigma_Z^2 = \sum_i \sigma_i^2 \left(\frac{1/\sigma_i^2}{1/\sum_k \sigma_k^2}\right)^2 = \frac{\sum_i 1/\sigma_i^2}{(\sum_k 1/\sigma_k^2)^2},$$

$$\sigma_Z^2 = \frac{1}{\sum 1/\sigma_k^2}.$$

This is all well and good, if σ_i is known. However, usually the x_i are means themselves, based on data points that are also used to estimate σ_i. Consequently there are contributions to the derivative $\frac{dZ}{dx_i}$ from the derivative acting on σ_i. The variance of Z is therefore a bit larger than the above formula predicts:

Consider a data set that is composed of N subsets, each subset i contains a set of n_i data points $x_{i,j}$ where $j = 1, 2, 3 ... n_i$. All $x_{i,j}$ are drawn from distributions with mean μ but different sets i may have different widths σ_i. For each subset, we can calculate the mean $\bar{x}_i = \sum_j x_{i,j}/n_i$ and the weight $1/\bar{\varsigma}_i^2 = n_i/\varsigma_i^2 = n_i(n_i-3)/\sum_j(x_{i,j}-\bar{x}_i)^2$. The weight includes a factor of n_i because we want an estimate of the reciprocal of the variance of the mean $1/\bar{\sigma}^2 = n_i/\sigma_i^2$. (Note I am using the correct unbiased estimator for the weight).

The weighted mean in this case is:

$$Z = \frac{\sum_k \bar{x}_k/\bar{\varsigma}_k^2}{\sum_k 1/\bar{\varsigma}_k^2}.$$

The variance of Z is then given by:

$$\sigma_Z^2 = \sum_i \sum_j \sigma_i^2 \left(\frac{dZ}{dx_{i,j}}\right)^2 = \sum_i \sigma_i^2 \sum_j \left(\frac{dZ}{dx_{i,j}}\right)^2.$$

The derivative can be expanded:

$$\frac{dZ}{dx_{i,j}} = \frac{d}{dx_{i,j}}\left(\frac{\sum_k \bar{x}_k/\bar{\varsigma}_k^2}{\sum_k 1/\bar{\varsigma}_k^2}\right) = \frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\frac{d}{dx_{i,j}}\left(\sum_k \frac{\bar{x}_k}{\bar{\varsigma}_k^2}\right) - \frac{\sum_k \bar{x}/\bar{\varsigma}_k^2}{(\sum_k 1/\bar{\varsigma}_k^2)^2}\frac{d}{dx_{i,j}}\left(\sum_k \frac{1}{\bar{\varsigma}_k^2}\right),$$

$$\frac{dZ}{dx_{i,j}} = \frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\frac{d}{dx_{i,j}}\left(\sum_k \frac{\bar{x}_k}{\bar{\varsigma}_k^2}\right) - \frac{Z}{\sum_k 1/\bar{\varsigma}_k^2}\frac{d}{dx_{i,j}}\left(\sum_k \frac{1}{\bar{\varsigma}_k^2}\right).$$

Now the derivatives can rewritten:

$$\frac{d}{dx_{i,j}}\left(\sum_k \frac{1}{\bar{\varsigma}_k^2}\right) = \frac{d}{dx_{i,j}}\left(\frac{1}{\bar{\varsigma}_i^2}\right) = -\left(\frac{1}{\bar{\varsigma}_i^2}\right)^2 \frac{d\bar{\varsigma}_i^2}{dx_{i,j}},$$

$$\frac{d}{dx_{i,j}}\left(\sum_k \frac{\bar{x}_k}{\bar{\varsigma}_k^2}\right) = \frac{d}{dx_{i,j}}\left(\frac{\bar{x}_i}{\bar{\varsigma}_i^2}\right) = \frac{1}{\bar{\varsigma}_i^2}\frac{d\bar{x}_i}{dx_{i,j}} - \bar{x}_i\left(\frac{1}{\bar{\varsigma}_i^2}\right)^2 \frac{d\bar{\varsigma}_i^2}{dx_{i,j}}.$$

Substituting and simplifying yeilds:

$$\frac{dZ}{dx_{i,j}} = \frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\left[\frac{1}{\bar{\varsigma}_i^2}\frac{d\bar{x}_i}{dx_{i,j}} - \bar{x}_i\left(\frac{1}{\bar{\varsigma}_i^2}\right)^2\frac{d\bar{\varsigma}_i^2}{dx_{i,j}}\right] + \frac{Z}{\sum_k 1/\bar{\varsigma}_k^2}\left(\frac{1}{\bar{\varsigma}_i^2}\right)^2\frac{d\bar{\varsigma}_i^2}{dx_{i,j}},$$

$$\frac{dZ}{dx_{i,j}} = \frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\left[\frac{1}{\bar{\varsigma}_i^2}\frac{d\bar{x}_i}{dx_{i,j}} - \left(\frac{1}{\bar{\varsigma}_i^2}\right)^2(\bar{x}_i - Z)\frac{d\bar{\varsigma}_i^2}{dx_{i,j}}\right].$$

Now we need the derivatives of $\bar{x}_i$:

$$\frac{d\bar{x}_i}{dx_{i,j}} = \frac{d}{dx_{i,j}}\frac{\sum_l x_{i,l}}{n_i} = \frac{1}{n_i}\frac{dx_{i,j}}{dx_{i,j}} = \frac{1}{n_i},$$

and $\bar{\varsigma}_i^2$:

$$\frac{d\bar{\varsigma}_i^2}{dx_{i,j}} = \frac{d}{dx_{i,j}}\frac{\sum_l(x_{i,l} - \bar{x}_i)^2}{n_i(n_i-3)} = \frac{2}{n_i(n_i-3)}\sum_l(x_{i,l} - \bar{x}_i)\frac{d}{dx_{i,j}}(x_{i,l} - \bar{x}_i),$$

$$\frac{d\bar{\varsigma}_i^2}{dx_{i,j}} = \frac{2}{n_i(n_i-3)}\sum_l(x_{i,l} - \bar{x}_i)\frac{dx_{i,l}}{dx_{i,j}}(x_{i,l}) - \frac{2}{n_i(n_i-3)}\sum_l(x_{i,l} - \bar{x}_i)\frac{d\bar{x}_i}{dx_{i,j}}.$$

The second term sums to zero since $\sum_l x_{i,l} = n_i\bar{x}_i$, leaving

$$\frac{d\bar{\varsigma}_i^2}{dx_{i,j}} = \frac{2}{n_i(n_i-3)}(x_{i,j} - \bar{x}_i).$$

Substituting in gives:

$$\frac{dZ}{dx_{i,j}} = \frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\left[\frac{1}{\bar{\varsigma}_i^2}\frac{1}{n_i} - \left(\frac{1}{\bar{\varsigma}_i^2}\right)^2(\bar{x}_i - Z)\frac{2}{n_i(n_i-3)}(x_{i,j} - \bar{x}_i)\right],$$

$$\frac{dZ}{dx_{i,j}} = \frac{1}{n_i}\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\left[1 - \frac{2}{n_i-3}\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}(x_{i,j} - \bar{x}_i)\right].$$

Squaring and summing over j

$$\sum_j\left(\frac{dZ}{dx_{i,j}}\right)^2 = \left[\frac{1}{n_i}\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2\sum_j\left[1 - \frac{2}{n_i-3}\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}(x_{i,j} - \bar{x}_i)\right]^2$$

$$= \left[\frac{1}{n_i}\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2 \sum_j \left[1 - \frac{4}{n_i-3}\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}(x_{i,j}-\bar{x}_i) + \frac{4}{(n_i-3)^2}\left(\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}\right)^2 (x_{i,j}-\bar{x}_i)^2\right].$$

Upon summation, the term proprotional to $\sum_j (x_{i,j} - \bar{x}_i)$ becomes zero, leaving

$$\sum_j \left(\frac{dZ}{dx_{i,j}}\right)^2 = \left[\frac{1}{n_i}\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2 \left[n_i + \frac{4n_i}{(n_i-3)}\left(\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}\right)^2 \sum_j \frac{(x_{i,j}-\bar{x}_i)^2}{n_i(n_i-3)}\right],$$

$$\sum_j \left(\frac{dZ}{dx_{i,j}}\right)^2 = \left[\frac{1}{n_i}\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2 \left[n_i + \frac{4n_i}{(n_i-3)}\left(\frac{\bar{x}_i - Z}{\bar{\varsigma}_i^2}\right)^2 \bar{\varsigma}_i^2\right],$$

$$\sum_j \left(\frac{dZ}{dx_{i,j}}\right)^2 = \frac{1}{n_i}\left[\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2 \left[1 + \frac{4}{(n_i-3)}\frac{1}{\bar{\varsigma}_i^2}(\bar{x}_i - Z)^2\right].$$

Now to complete the calculation, multiply by σ_i^2 and sum over i:

$$\sigma_Z^2 = \sum_i \frac{\sigma_i^2}{n_i}\left[\frac{1/\bar{\varsigma}_i^2}{\sum_k 1/\bar{\varsigma}_k^2}\right]^2 \left[1 + \frac{4}{(n_i-3)}\frac{1}{\bar{\varsigma}_i^2}(\bar{x}_i - Z)^2\right].$$

Now we need to convert the unknown σ_i^2 into a known estimate. Since we are trying to estimate the variance , the correct estimate is $S_i^2 = \frac{n_i-3}{n_i-1}\varsigma_i^2 = n_i\frac{n_i-3}{n_i-1}\bar{\varsigma}_i^2$. Plugging this in gives:

$$\sigma_Z^2 = \left(\frac{1}{\sum_k 1/\bar{\varsigma}_k^2}\right)^2 \sum_i \frac{1}{\bar{\varsigma}_i^2}\left(\frac{n_i-3}{n_i-1}\right)\left[1 + \frac{4}{(n_i-3)}\frac{1}{\bar{\varsigma}_i^2}(\bar{x}_i - Z)^2\right].$$

This expression can be simplified if the number of subsets N is large enough to make the dispersion of Z much less than the dispersion in $\bar{x}_i$, becuase in this case $(\bar{x}_i - Z)^2/\bar{\varsigma}_i^2 \simeq 1$. With this approximation:

$$\sigma_Z^2 \simeq \left(\frac{1}{\sum_i 1/\bar{\varsigma}_i^2}\right)^2 \sum_i \frac{n_i+1}{n_i-1}\frac{1}{\bar{\varsigma}_i^2}.$$

If n_i is the same for i, then:

$$\sigma_Z^2 \simeq \frac{1}{\sum_i 1/\bar{\varsigma}_i^2}\left(\frac{n+1}{n-1}\right)$$

The variance of the weighted mean is therefore larger than the naive calculation predicts by a factor of $(n+1)/(n-1)$.

Summary/Moral: Because the means and weights used to calculate the weighted means are derived from the same data, the variance of the weighted mean is slightly larger than one would predict if the weights were derived independently (by a fctor of $\frac{n+1}{n-1}$). Also note that the final expression involves ς, not S. In fact, the two effects discussed in this Appendix contribute roughly equally to the bias in the naive estimation of the error on the weighted mean (while the estimate of the mean itself affected by the choice of estimator alone).

Bibliography

Barkats, D., 2000. Measuring Q and U simultaneously: Advantages and limitations. Princeton U. Experimental Project Report.

Barkats, D., 2002. In preparation.

Barnes, C., 2002. MAP satellite feed horns. Submitted to IEEE Trans, Ant. Prop.

Bevington, P.R., 1969. *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill.

Brouw, W.N. and T.A.T. Spoelestra, 1976. Linear polarization of the Galactic background at frequencues between 408 and 1411 MHz. reductions. *ApJ Letters*, 391:L49.

Butler, N., 1998. Emissivity measurements of MAP satellite optics. Princeton U. Senior Thesis.

Caderni, N., R Fabbri, B. Melchiorri, F. Melchiorri, and V. Natale, 1978. Polarization of the microwave background ii. an infrared survey of the sky. *Physical Review D*, 17(8):1908–1918.

Cartwright, J.K., S. Padin, T.J. Pearson, A.C.S. Readhead, M.C. Shepherd, and G.B. Taylor, 2001. Polaization observations with the Cosmic Background Imager. In *2K1BC Workshop on Experimental Cosmology at Millimeter Wavelengths*.

Chapman, B.D., 1986. *Thermal Radio Emission from the Durfaces of Venus and Mercury*. Ph.D. thesis, M.I.T.

Cortiglioni, S., 1994. Linear polarization and the effects of metal reflectors used to redirect the beam in microwave radiometers. *Rev Sci Instrum*, 65(8):2667–2671.

Cowden, D.J., 1957. *Statistical Methods in Quality Control*. Prentice Hall.

Davies, R.D. and F.F. Gardner, 1966. Linear polarization of the moon at 6,11 and 21 cm wavelengths. *Australian Journal of Physics*, 19:823–836.

de Oliveira-Costa, A., M. Tegmark, D.P. Finkbeiner, R.D. Davies, C.M. Gutierrez, L.M. Haffner, A.W. Jones, A.N. Lasenby, R. Rebolo, R. Reynolds, R.J. Tufte, and R.A. Watson, 2000. A new spin on Galactic dust. `astro-ph/00110527`.

de Oliviera-Costa, A. and M. Tegmark, eds., 1999. *Microwave Foregrounds*. Astonomical Society of the Pacific.

de Oliviera-Costa, A., M. Tegmark, L.A. Page, and S.P. Boughn, 1998. Galactic emission at 19 GHz. *ApJ Letters*, 509:L9.

Duncan, A.R., R.F. Haynes, K.L. Jones, and R.T. Stewart, 1997. Polarized radio emission over the southern Galactic plane at 2.4 GHz. *MNRAS*, 291:279.

Elliott, R. S., 1993. *An Introduction to Guided Waves and Microwave Circuits*. Prentice Hall.

Epstein, E.E., M.A. Janssen, J.N. Cuzzi, W.G. Fogarty, and J. MottMann, 1980. Saturn's rings: 3-mm observations and derived properties. *Icarus*, 41:103–118.

Fedoseev, L.I. and V.I. Chernyshev, 1998. The millimeter spectrum of the quiet sun. *Astronomy Reports*, 42(1):105–109.

Fomalont, E. B., R. B. Partridge, J. D. Lowenthal, and R. A. Windhorst, 1993. Limits to Cosmic Background Radiation fluctuations at 8.44 GHz between angular scales 10" and 200". *ApJ*, 404(1):8–20.

Gary, B., J. Stacey, and F.D. Drake, 1965. Radiometric mapping of the moon at 3 millimeters wavelength. *ApJ Suppl Series*, 12:239–262.

Griffin, M.J., P.A.R. Ade, G.S. Orton, E.I. Robson, W.K. Gear, I.G. Nolt, and J.V. Radositz, 1986. Submillimeter and millimeter observations of Jupiter. *Icarus*, 65:244–256.

Grossman, E., 1989. AT – Atmospheric Transmission Software: User's Manual.

Gundersen, J.O., 2000. Polarization foregrounds near the NCP. `online.itp.ucsb.edu/online/cmb00/gundersen/`.

Halverson, N.W., E.M. Leitch, C. Pryke, J. Kovac, J.E. Carlstrom, W.L. Holzapfel, M. Dragovan, J.K. Cartwright, B.S. Mason, S. Padin, T.J. Pearson, M.C. Shepherd, and A.C.S. Readhead, 2001. DASI first results: A measurement of the Cosmic Microwave Background angular power spectrum. `astro-ph 0104489`.

Hanany, S., 2001. `http://www.physics.umn.edu/cosmology/maxipol`.

Haslam, C.G.T., C.J. Salter, H. Stoffel, and W.E. Wilson, 1982. A 408 MHz all-sky continuum survey. ii- the atlas of contour maps. *Astronomy and Astrophysics*, 47:1.

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. *ApJ Letters*, 548:L111–L114.

Herzog, C., 1999. A calibration of a microwave telescope. Princeton U. Experimental Project Report.

Hu, W. and N. Sugiyama, 1995. Anisotropies in the Cosmic Microwave Background: An analytic approach. *ApJ*, 444:489–506.

Hu, W., N. Sugiyama, and J. Silk, 1997. The physics of microwave background anisotropies. *Nature*, 386:37–43.

Hu, W. and M. White, 1997. A CMB polarization primer. *New Astronomy*, 2:323.

IAU, 1973. Commission 40: Radio astronomy. *Trans Int Astron Union*, 15B:166–168.

Jackson, J. D., 1975. *Classical Electrodynamics*. John Wiley and Sons, 2 edn.

Kamionkowski, M. and A. Kosowky, 1999. The Cosmic Microwave Background and particle physics. *Ann Rev Nucl Part Sci*, 49:77–123.

Kamnionkowski, M., A. Kosowsky, and A. Stebbins, 1997. Statistics of cosmic micrtowave background polarization. *Phys Rev D*, 55:7368.

Keating, B., P. Timbie, A. Polnarev, and J. Steinberger, 1997. Large angular scale polarization and the feasibility of its detection. `astro-ph/9710087`.

Keating, B. G., C. W. O'Dell, A. de Oliviera-Costa, S. Klawikowski, L. Piccirillo, M. Tegmark, and P. T. Timbie, 2001a. A limit on the large angular scale polarization of the Cosmic Microwave Background. *ApJ Letters*, 560:L1–L4.

Keating, B.G., C.W. Odell, J.O. Gundersen, L. Piccirillo, N.C. Steror, and P.T. Timbie, 2001b. An instrument for investigating the large angular scale polarization of the Cosmic Microwave Background. `astro-ph/011127`.

Keller, J.B., 1962. Geometrical theory of diffraction. *Journal of the Optical Society of America*, 52(2):116–129.

Kodama, H. and M. Sasaki, 1984. Cosmological perturbation theory. *Progress of Theoretical Physics Supplement*, (78).

Kolb, E.W. and R.D. Pecceci, eds., 1995. *Particle and nuclear astrophysics in the next millenium. Proceedings, summer study, Snowmass, USA, Jun29-July14 1994*.

Kosowsky, A., 1999. Introduction to microwave background polarization. *New Astronomy Rev*, 43:157.

Kovac, J., 2001. Polarization with DASI. `http://find.uchicago.edu/ john/dasi/efi html/index.html`.

Kraus, John D., 1986. *Radio Astronomy*. Cygnus-Quasar Books, 2 edn.

Lazarian, A., 2000. Polarized microwave radiation from dust. `astro-ph 0101001`.

Lee, A.T., P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P.G. Ferreira, S. Hanany, V.V. Hristov, A.H. Jaffe, P.D. Mauskopf, C.B. Netterfield, E. Pascale, B. Rabii, P.L. Richanrds, G.F. Smoot, R. Stompor, C.D. Winant, and J.H.P. Wu, 2001. A high resolution analysis of the MAXIMA-1 cosmic microwave background anisotropy data. `astro-ph/0105062`.

Liddle, A.R. and D.H. Lyth, 2000. *Cosmological Inflation and Large Scale Structure*. Cambridge University Press.

Lubin, P., P. Melese, and G. Smoot, 1983. Linear and circular polarization of the Cosmic Background Radiation. *ApJ Letters*, 273:L51–L54.

Lubin, Phil M. and George F. Smoot, 1979. Search for linear polarization of the cosmic background radiation. *Physical Review Letters*, 42(2):129–132.

Lubin, Phil M. and George F. Smoot, 1981. Polarization of the Cosmic Background Radiation. *ApJ*, 245(1):1–17.

Lupton, R., 1993. *Statistics in Theory and Practice*. Princeton U. Press.

Miller, A.D., R. Caldwell, M.J. Devlin, W.B. Dorwart, M.R. Nolta, L.A. Page, J. Puchalla, E. Torbet, and H.T. Tran, 1999. A measurement of the angular power spectrum of the CMB from l=100 to 400. *ApJ Letters*, 524:L1–L4.

Mitchell, D.L. and I. de Pater, 1994. Microwave imaging of Mercury's thermal emission at wavelengths from 0.3 to 20.5 cm. *Icarus*, 110:2–32.

Nanos, G.P., 1979. Polarization of the blackbody radiation at 3.2 centimeters. *ApJ*, 232(2):341–347.

Netterfield, C. B., 1995. *A Measurement of the Degree Scale Anisotropy in the Cosmic Microwave Background Radiation*. Ph.D. thesis, Princeton University.

Netterfield, C.B., P.A.R. Ade, J.J. Bock, J.R. Bond, J. Borrill, A. Boscaleri, K. Coble, C.R.Contaldi, B.P. Crill, P. de Bernardis, P. Farese, K. Ganga, M. Giacometti, E. Hivon, V.V. Hristov, A. Iacoamgeli, A.H. Jaffe, W.C. Jones, A.E. Lange, L. Martinis, S. Masi, P. Mason, P.D. Mauskopf, A. Melchiorri, T. Montroy, E. Pascale, F. Piacentini, D. Pogosyan, F. Pongetti, S. Prunet, G. Romeo, J.E. Ruhl, and F. Scaramuzzi, 2001. A measurement by BOOMERANG of multiple peaks in the angular power spectrum of the cosmic microwave background. `astro-ph/0104460` .

Partridge, R. B., J. Nowakowski, and H. M. Martin, 1988. Linear polarized fluctuations in the Cosmic Microwave Background. *Nature*, 331:146–147.

Partridge, R. B., E. A. Richards, E. B. Fomalont, K. I. Kellerman, and R. A. Windhorst, 1997. Small-scale Cosmic Microwave Background observations at 8.4 GHz. *ApJ*, 483(1):38–50.

Penzias, A.A. and R.W. Wilson, 1965. A measurement of excess antenna temperature at 4080 Mc/s. *ApJ*, 142(1):419–421.

Pilhour, B.J., B.G. Keating, P.A.R. Ade, R.S. Bhatia, J.J. Bock, S.E. Church, J. Glenn, J.R. Hinderks, V.V. Hristov, W.C. Jones, M. Kamionkowski, D.E. Kumar, A.E. Lange, J.R. Leong, D.P. Marrone, B.S. Mason, M.M. Shuman, and G. I. Sirbi, 2001. The Polatron: A millimeter-wave cosmic background polarimeter for the OVRO 5.5 m telescope. `astro-ph/0106543`.

Pozar, D.M., 1998. *Microwave Engineering*. J. Wiley and Sons, Inc., 2 edn.

Rahmat-Samii, Y., W. Imbriale, and V. Galindo, 1996a. Conical Corrugated Horn Antenna Analysis (CCORHRN).

Rahmat-Samii, Y., W. Imbriale, and V. Galindo, 1996b. Diffraction Analysis of a Dual Reflector Antenna (DADRA).

Ruhl, J., 2001. `http://www.physics.ucsb.edu/ boomerang/index.html`.

Ruze, J., 1952. Physical limitations on antennas. *MIT Res Lab Electron Tech Report*, 248.

Seljak, U. and M. Zaldarriaga, 1996. A line of sight approach to Cosmic Microwave Background anisotropies. *ApJ*, 469:437–444.

Seljak, U. and M. Zaldarriaga, 1997. Signature of gravity waves in the polarization of the microwave background. *Phys Rev Lett*, 78:2054–2057.

Sharp, M. L., 1993. *Behavior and Design of Aluminum Structures*. McGraw-Hill Inc.

Sironi, G., G. Boella, G. Bonelli, L. Bruneti, F. Cavaliere, M. Gervasi, G. Giardino, and A. Passerini, 1998. A 33 GHz polarimeter for observations of the Cosmic Microwave Background. *New Astronomy*, 3:1–13.

Staggs, S. T., J. O. Gundersen, and S. E. Church, 1999. Cosmic Background Radiation polarization experiments. In de Oliviera-Costa and Tegmark (1999), p. 299.

Steinhardt, P., 1995. Cosmology at the crossroads. In Kolb and Pecceci (1995), p. 51.

Stompor, R., M. Abroe, P. Ade, A. Balbi, D. Barbosa, J. Bock, J. Borill, A. Boscaleri, P. de Bernardis, P.G. Ferreira, S. Hanany, V. Hristov, A.H. Jaffe, A.T. Lee, E. Pascale, B. Rabii, P.L. Richards, G.F. Smoot, C.D. Winant, and J.H.P. Wu, 2001. Cosmological implications of the MAXIMA-1 high resolution Cosmic Microwave Background anisotropy measurement. `astro-ph/0105062`.

Strozzi, D.J. and K.T. McDonald, 2000. Polarization dependence on emissivity. `physics/0005024`.

Subrahmanyan, R., M. J. Kesteven, R. D. Ekers, M. Sinclair, and J. Silk, 2000. An australian telescope survey for CMB anisotropies. `astro-ph/0002467`.

Tegmark, M., D.J. Eisenstein, W. Hu, and A. de Oliveira-Costa, 1999. Overview of foregrounds and their impacts. In de Oliviera-Costa and Tegmark (1999), p. 3.

Timbie, P., 2001. `http://cmb.physics.wisc.edu/compass.html`.

Torbet, E., M. J. Devlin, W. B. Dorwart, T. Herbig, A. D. Miller, M. R. Nolta, L. Page, J. Puchalla, and H. T. Tran, 1999. A measurement of the angular power spectrum of the microwave background radiation made from the high chilean andes. *ApJ Letters*, 524:L1–L4.

Ulich, B.L., J.H. Davis, P.J. Rhodes, and J.M. Hollis, 1980. Absolute brightness temperature measurements at 3.5-mm wavelength. *IEEE Trans Ant Prop*, 28(3):367–377.

Uyaniker, B., E. Furst, W. Reich, P. Reich, and R. Wielebinski, 1998. A 1.4 GHz radio continuum and polarization survey at medium galactic latitude: I. observation and reduction technique. *Astronomy and Astrophysics*, 132:402.

Wang, X., M. Tegmark, and M. Zaldarriaga, 2001. Is cosmology consistent? `astro-ph/0105091`.

White, G.K., 1959. *Experimental Techniques in Low Temperature Physics*. Oxford at the Clarendon.

Wollack, E. J., N. C. Jarosik, C. B. Netterfield, L. A. Page, and D. Wilkinson, 1993. A measurement of the anisotropies in the Cosmic Microwave Background Radiation at degree angular scales. *ApJ Letters*, 419:L49–L52.

Wollack, E.J., 1994. *A Measurement of the Degree Scale Cosmic Background Radiation Anisotropy at 27.5, 30.5 and 33.5 GHz*. Ph.D. thesis, Princeton University.

Wollack, E.J., M.J. Devlin, N. Jarosik, C.B. Netterfield, L. Page, and D. Wilkinson, 1997. An instrument for investigation of the Cosmic Microwave Background at intermediate angular scales. *ApJ*, 476(2):440–457.

Wollack, E.J. and M.W. Pospieszalski, 1998. *IEEE MMT S Int Microwave Symp Digest*, p. 669.

Zaldarriaga, M., 1998. CMB polarization experiments. *ApJ*, 503:1–15.

Zhang, X., 1993. Design of conical corrugated feed horns for wide-band high-frequency applications. *IEEE Trans Microwave Theory and Techniques*, 41(8):1263–1274.