

**Maximum likelihood analysis with the  
Q/U Imaging Experiment**

Sigurd Kirkevold Næss

Institute of Theoretical Astrophysics  
University of Oslo

2012





# Acknowledgements

This thesis would not have seen the light of day without the guidance and inspiration from my friends, family and colleagues.

I would in particular like to thank my advisor Hans Kristian Eriksen and Ingunn Wehus, who have been my main collaborators in building the QUIET ML pipeline. I could not have asked for more inspiring co-workers, and developing the QUIET pipeline with them has been both enjoyable and productive.

I should also thank the rest of the QUIET collaboration for making the experiment work, and for the opportunity to get hands-on experience with a telescope in one of the most exotic locations in the world.

For introducing me to science and astronomy, I would like to thank my parents Petter Næss and Nina Kirkevold; a series of inspiring science teachers through school, high-school, and university; and in particular Øyvind Grøn, who first introduced me to modern cosmology in one of his popular science talks.

I am also grateful to Frode Hansen and Kristin Mikkelsen, who made my teaching duties enjoyable, and to my co-supervisor Per Lilje, who guided me through the red tape of university bureaucracy.

Finally, I would like to thank my fellow Ph.D. students at the institute for thought-provoking discussions and a good social environment. In particular Jostein Kristiansen and Langangen for answering my torrent of questions during the first months of my thesis, and Unni Fuskeland, Eirik Gjerlew, Dag Sverre Seljebotn, Sandro Scodeller, Mikkjel Thorsrud and Hans Winther for useful collaboration and discussion.



# Contents

<b>I</b>	<b>CMB data analysis for QUIET</b>	<b>7</b>
<b>1</b>	<b>Introduction</b>	<b>9</b>
1.1	The CMB anisotropies . . . . .	11
1.2	Polarization . . . . .	14
1.2.1	Stokes parameters . . . . .	14
1.3	Detector technology . . . . .	18
1.4	The Q/U Imaging Experiment . . . . .	19
<b>2</b>	<b>Telescope and calibration</b>	<b>23</b>
2.0.1	CES detection . . . . .	25
2.0.2	Alternative scanning patterns . . . . .	29
2.1	Detector data and noise properties . . . . .	29
2.1.1	Noise estimation . . . . .	32
2.1.2	Filters . . . . .	38
2.1.3	Cuts . . . . .	43
2.2	Pointing . . . . .	46
2.2.1	Verifying the pointing . . . . .	47
2.2.2	Mount model . . . . .	51
2.2.3	Full-season validation of the mount model . . . . .	53
2.3	Gain and beam . . . . .	55
2.3.1	Polarization gains . . . . .	57
<b>3</b>	<b>Null tests</b>	<b>61</b>
3.1	Signal-less validation . . . . .	62
3.2	Map null-tests . . . . .	62
3.3	Pseudo- $C_l$ null-tests . . . . .	63
3.4	Null test suite . . . . .	64
<b>4</b>	<b>Map making</b>	<b>67</b>
4.1	Measuring the sky with a scanning telescope . . . . .	67
4.2	From sky to time-ordered data . . . . .	68
4.3	The noise . . . . .	70
4.4	From time-ordered data to sky . . . . .	71

4.4.1	Binned maps . . . . .	73
4.5	Filters . . . . .	74
4.5.1	Frequency filters . . . . .	74
4.5.2	General filters . . . . .	75
4.6	Practical implementation of the map-making equations . . .	77
4.6.1	Implementing $P^T DV$ . . . . .	77
4.6.2	Implementing $P^T DP$ . . . . .	78
4.6.3	Solving the map-making equation . . . . .	79
4.6.4	Implementation in the QUIET pipeline . . . . .	79
4.7	Biaslessness and deconvolution . . . . .	81
4.8	Results . . . . .	83
4.8.1	CMB maps . . . . .	83
4.8.2	Foreground maps . . . . .	83
<b>5</b>	<b>Power spectrum and parameter estimation</b>	<b>93</b>
5.1	Power spectrum estimation . . . . .	94
5.1.1	Brute force evaluation . . . . .	95
5.1.2	Pseudo- $C_l$ estimation . . . . .	97
5.1.3	Gibbs sampling . . . . .	99
5.2	Application to QUIET . . . . .	102
5.3	Results . . . . .	104
5.3.1	Temperature spectra . . . . .	105
5.3.2	Foregrounds . . . . .	105
5.3.3	Systematic errors . . . . .	114
5.3.4	Parameters . . . . .	116
<b>6</b>	<b>Non-gaussianity</b>	<b>119</b>
<b>7</b>	<b>Summary and outlook</b>	<b>123</b>
<b>8</b>	<b>Bibliography</b>	<b>125</b>
<b>II</b>	<b>Papers</b>	<b>133</b>
	Paper I: First Season QUIET Observations: Measurements of CMB polarization power spectra at 43 GHz in the multipole range $25 \leq l \leq 475$ . . . . .	135
	Paper II: The Q/U Imaging Experiment . . . . .	155
	Paper III: Bayesian noise estimation for non-ideal CMB experiments	185
	Paper IV: Application of the Kolmogorov-Smirnov test to CMB data: Is the universe really weakly random? . . . . .	195

**Part I**

**CMB data analysis for QUIET**



# Chapter 1

## Introduction

Cosmology, the study of the overall composition and history of the universe, is currently in the enviable position of having a well-tested standard model that agrees with all observations, while relying only a low number of adjustable parameters<sup>1</sup>. According to this model, which is known as  $\Lambda$ CDM or the “concordance model”, the universe started out in an extremely hot and dense state called the Big Bang some 13.8 billion years ago. This was almost immediately followed by a phase of extremely rapid expansion, called inflation.




Figure 1.1: A timeline of the universe, courtesy of the WMAP science team. Quantum fluctuations are created during inflations, resulting in the CMB fluctuations (afterglow light pattern) and the later formation of stars, galaxies and other structures.

<sup>1</sup>For a thorough introduction to this topic, see for example [1].

During inflation, the size of the universe increased exponentially by at least a factor of  $10^{26}$ , leaving the universe in an almost perfectly flat and uniform state, with small fluctuations with approximately the same amplitude on all length-scales<sup>2</sup>.

After this, the universe continued expanding much more slowly, eventually cooling down enough for protons and neutrons, then light atomic nuclei, and finally atoms to form. The latter happened about 370 000 years after the Big Bang, when the temperature of the universe had fallen to about 3000 K. The formation of atoms turned the plasma that had filled the universe until then into a neutral gas, making the universe transparent for the first time. Due to the finite speed of light, this event, called “recombination”<sup>3</sup>, is still visible today as an apparent surface, called the *surface of last scattering*, some 13.8 billion light-years away<sup>4</sup>.

Up to this point, the fluctuations from inflation had been gradually growing under the influence of gravity, and at the time of recombination, the universe was inhomogeneous at the level of 1:100 000. In the following billions of years until today, the fluctuations continued to grow, eventually becoming dense enough to form first stars and then galaxies, followed by even larger structures such as clusters of galaxies.

Quantitatively, the model is described by the ten parameters in table 1.1, which can be summarized as the age, expansion speed and density of the universe; the amplitude and scale dependence of the fluctuations; and the time of recombination and another significant event later in the history of the universe called reionization.

This concordance model is supported by a diverse set of observations, including

- the relationship between the redshift and magnitude of supernova explosions, which can be used to map out the expansion history of the universe.
- the distribution of galaxies on large scales, from which information about the fluctuations can be extracted.
- the chemical composition of the universe, which provides information about the baryon density and expansion speed during the early universe.

---

<sup>2</sup>It also had the effect of diluting the particles present in the universe before inflation into irrelevance. At the end of inflation, the universe is re-populated with particles created from the decay of the field(s) responsible for driving inflation.

<sup>3</sup>Though “combination” would have been a more appropriate name, as this is the first time nuclei and electrons combined to form atoms.

<sup>4</sup>In light travel distance, one of several possible ways of measuring distances in the universe.



Parameter	Value	Name
$t_0$	$13.78 \pm 0.11$ Gyr	Age of the universe
$H_0$	$69.9 \pm 1.3$ km/s/Mpc	Hubble parameter
$\Omega_b$	$0.0485 \pm 0.0026$	Baryon density
$\Omega_c$	$0.244 \pm 0.016$	Dark matter density
$\Omega_\Lambda$	$0.708 \pm 0.0016$	Dark energy density
$\sigma_8$	$0.811 \pm 0.023$	Fluctuation amplitude
$n_s$	$0.967 \pm 0.014$	Spectral index
$r$	$< 0.2$	Tensor-to-scalar ratio
$z_*$	$1020.3 \pm 1.3$	Redshift at last decoupling
$\tau$	$0.086 \pm 0.014$	Optical depth of reionization

Table 1.1: The 10 parameters of the  $\Lambda$ CDM model, and current bounds on their values, based on a combination of WMAP 7-year data, supernova observations and galaxy surveys.

However, the most sensitive test of the model currently available is radiation reaching us from the surface of last scattering, called the cosmic microwave background (CMB). We observe the CMB as perfect black-body radiation corresponding to a temperature of 2.725 K, a temperature which is nearly uniform in all directions on the sky<sup>5</sup>. But since the CMB is an image of the surface of last scattering, it cannot be completely uniform; it must contain the faint fluctuations that were present in the universe at that time.

## 1.1 The CMB anisotropies

These anisotropies in the CMB were first detected by the COBE satellite in 1992 [2], and were later mapped out in more detail by several other experiments, including the WMAP satellite [3, 4, 5, 6] (see fig. 1.2). The primordial fluctuations produced during inflation are ultimately sourced by random quantum fluctuations, which are expected to follow a statistically isotropic and homogeneous<sup>6</sup> Gaussian distribution, and these properties are inherited by the temperature fluctuations at the surface of last scattering. The precise position and value of each positive and negative fluctuation is therefore not of cosmological interest, but the statistical properties

<sup>5</sup>The observed temperature is a factor of  $\sim 1000$  lower than the temperature at recombination. This fall in temperature between the time the radiation was emitted and observed is expected from general relativity, which predicts that the wavelength of photons will grow proportionally with the expansion of the universe.

<sup>6</sup>That is, while each realization of the random field is anisotropic and inhomogeneous, they will not systematically prefer any position or direction.




Figure 1.2: The WMAP 7-year ILC map of the CMB temperature anisotropies. At this resolution, the fluctuations have a typical amplitude of  $\sim 100\mu\text{K}$ . The faint horizontal artifacts in the middle of the map are due to contamination from our own galaxy.

of the fluctuation field as a whole are, as they contain information not only about the conditions that sourced them during inflation, but also about the physical conditions during the 370 000 year period between inflation and recombination.

A statistically isotropic Gaussian random field on the sphere has the convenient property that its statistical properties can be completely described in terms of the angular power spectrum,  $C_l$ , which is given by

$$C_l = \langle |a_{lm}|^2 \rangle \quad a_{lm} = \int Y_{lm}^*(\hat{\theta}) \Delta T(\hat{\theta}) d\Omega, \quad (1.1)$$

where  $a_{lm}$  are the coefficients of the decomposition of the fluctuation field  $\Delta T(\hat{\theta})$  in terms of the spherical harmonics<sup>7</sup>  $Y_{lm}(\hat{\theta})$ . The CMB temperature fluctuation power spectrum as measured by the WMAP satellite is compared to the  $\Lambda\text{CDM}$  best fit in figure 1.3. The fit is excellent, despite the much greater number of data points than model parameters. This lends confidence to the idea that  $\Lambda\text{CDM}$  is correct, or at least a very good approximation to reality.

For the most part, the  $\Lambda\text{CDM}$  model builds on the two well-tested theories general relativity and the standard model of particle physics, which together form the status quo of our understanding of the particles and forces

<sup>7</sup>The spherical harmonics are eigenvectors of the angular part of the  $\nabla^2$  operator. They are analogues to the normal harmonic functions, and form an orthogonal basis set on the sphere.




Figure 1.3: The power spectrum of the CMB temperature fluctuations as measured by the WMAP satellite, compared to the best-fit  $\Lambda$ CDM model.

that exist. However,  $\Lambda$ CDM also requires a few ingredients which go beyond these two, and its impressive fit to observations gives good reason to believe that these are actually real. The new ingredients are:

**Dark energy** An unknown form of energy with negative pressure and repulsive gravity, currently making up 71% of the energy density in the universe. It is needed to explain the current accelerated expansion of the universe.

**Dark matter** An unknown form of matter which interacts very weakly with normal matter and radiation, and which currently makes up 24% of the energy density of the universe.

**Inflation** The period of extremely rapid expansion that sets up the fluctuations in the early universe. Possibly related to dark energy.

Of these, the most speculative and poorly measured is inflation, despite its central place in the theory. It is believed to have happened at an energy scale of up to  $\sim 10^{16}$  GeV [7], which puts it far beyond the reach of foreseeable particle experiments on Earth. However, such a high energy scale would also result in the production of a significant amount of tensor fluctuations (gravitational waves) during inflation, which would be detectable in the CMB power spectrum if present in sufficient amounts. Thus, the CMB power spectrum is a unique opportunity for probing physics at these ultra-high energy scales.

The amplitude of the primordial waves is parameterized by the tensor-to-scalar ratio  $r$ , which is defined as the ratio of the primordial tensor and

scalar power at the scale with wavenumber  $k = 0.05/\text{Mpc}$ . So far, no tensor perturbations have been detected, so  $r$  is consistent with zero, but with a quite large confidence interval:  $r < 0.2$  [8]. This bound is based on the effect of tensor perturbations on the CMB temperature power spectrum, and derives most of its significance from the multipoles  $l < 10$ , which are fundamentally limited in sensitivity by cosmic variance<sup>8</sup>. It is therefore not possible to significantly improve on this bound using the CMB temperature fluctuations.

However, the CMB, being electromagnetic radiation, has more degrees of freedom than just its temperature – it can also be polarized, and the CMB polarization has no such fundamental limit on the ability to detect tensor modes.

## 1.2 Polarization

General electromagnetic radiation can be expressed as a linear combination of plane wave solutions of Maxwell's equation, which take the form

$$\vec{E}(\vec{r}, t) = |\vec{E}| \begin{pmatrix} \cos \theta e^{i\alpha_x} \\ \sin \theta e^{i\alpha_y} \\ 0 \end{pmatrix} e^{i(kz - \omega t)} \quad (1.2)$$

$$\vec{B}(\vec{r}, t) = c^{-1} \vec{k} \times \vec{E}(\vec{r}, t), \quad (1.3)$$

for a wave travelling in the  $\vec{z}$  direction, where  $\vec{E}$  and  $\vec{B}$  are the electric and magnetic field strengths,  $\vec{r}$  and  $t$  are the position and time,  $\theta$  is the linear polarization angle and  $\alpha_x, \alpha_y$  are the polarization phase angles [9]. Solutions with  $\alpha_y = \alpha_x$  are called linear polarization;  $\alpha_y = \alpha_x \pm \frac{\pi}{2}$  gives rise to circular polarization, and other choices are called elliptical polarization, and can be described as linear combinations of linear and circular polarization.

No choice of these parameters correspond to unpolarized light. Instead, unpolarized light can be built up as a linear combination of plane waves with different phases and polarization angles. This results in light with polarization changing rapidly over short time scales, adding up to a zero net polarization. In general, it is also possible that the polarization does not completely cancel out, resulting in partially polarized light.

### 1.2.1 Stokes parameters

A useful way of characterizing these possibilities is in terms of the *Stokes parameters*  $\vec{S} = (I, Q, U, V)^T$ , where  $I \equiv |E_x|^2 + |E_y|^2$ ,  $Q \equiv |E_x|^2 - |E_y|^2$ ,

<sup>8</sup>Cosmic variance refers to the uncertainty inherent in only having a single location in a single universe from which to observe. For example, there are only 5 linearly independent quadrupoles on the sky, which sets a lower limit on the uncertainty of  $C_2$ , independently of instrumental noise, etc.




Figure 1.4: Illustration of the linear polarization Stokes parameters  $Q$  and  $U$ , and the circular polarization parameter  $V$ . Stokes parameters are useful due to their linearity when combining light from multiple sources and the ease of measuring them, but they are arbitrary in the sense that they depend on the orientation of the coordinate system. The  $E$  and  $B$  decomposition provide a more natural alternative to  $Q$  and  $U$ . Based on illustration by Dan Moulton.

$U = 2\text{Re}(E_x E_y^*)$  and  $V \equiv 2\text{Im}(E_x E_y^*)$ . These represent respectively the total intensity of the radiation<sup>9</sup>, the linear polarization along the  $x$  (positive) and  $y$  (negative) axes, linear polarization along the  $x+y$  (positive) and  $x-y$  (negative) direction, and the right-handed (positive) and left-handed (negative) circular polarization, as illustrated in figure 1.4.

Under a rotation of the coordinate system,  $I$  and  $V$  are unchanged, while  $(Q, U)$  rotates as a spin 2 quantity. That is, under a rotation  $\psi$  around the  $z$  axis of the local coordinate system,

$$\begin{pmatrix} Q \\ U \end{pmatrix} \rightarrow \begin{pmatrix} Q' \\ U' \end{pmatrix} = \begin{pmatrix} \cos(2\psi)Q - \sin(2\psi)U \\ \cos(2\psi)U + \sin(2\psi)Q \end{pmatrix}. \quad (1.4)$$

Compared to a description in terms of polarization fraction and polarization angles, the Stokes parameters are useful because they add linearly when combining radiation from multiple sources. They are also easy to measure compared to other parameterizations. However, they are not the

<sup>9</sup>The CMB has a black-body spectrum, and its intensity is therefore fully described by its temperature. It is therefore usual to express the Stokes parameters in temperature units in the context of CMB analysis, and when doing this, it is common to denote the total intensity parameter by  $T$  instead of  $I$ .

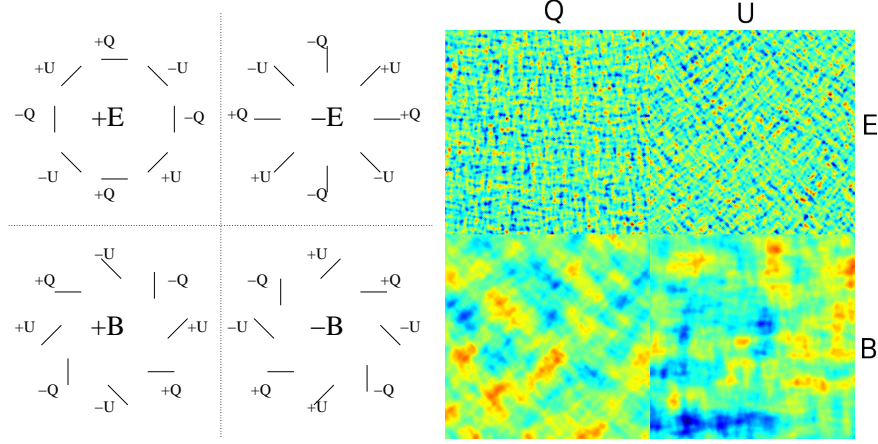


Figure 1.5: Left: Q/U pattern surrounding a point with positive/negative E/B. E-modes are curl-free, while B-modes are divergence-free. Right: Q/U maps corresponding to an E-only (top) and B-only (bottom) CMB realization. The scale difference between the E and B modes shown here is due to the expected behavior of primordial B-modes, and are not intrinsic to the definition of E and B.

most natural description of a statistically isotropic radiation field, as the definition of the linear polarization parameters  $Q$  and  $U$  depends on the arbitrary choice of an “up” direction, which results in  $Q$  rotating into  $U$  and vice versa during a rotation of the coordinate system.

The scalar  $E$  field and the pseudo-scalar field  $B$  are rotationally invariant alternatives to  $Q$  and  $U$ .  $E$  and  $B$  are respectively curl-free and gradient-free, and are named in analogy to the similar properties of the electric and magnetic field. The value of  $E$  and  $B$  in a given point can be defined as radially-weighted averages of  $Q_r$  and  $U_r$ , the Stokes  $Q$  and  $U$  parameters as defined in a polar coordinate system centered on that point<sup>10</sup>.  $E$  and  $B$  modes and their connection to  $Q$  and  $U$  are shown in figure 1.5. A pedagogical explanation can be found in [10], while [11] has an exact full-sky treatment of the topic.

The CMB is partially polarized due to Thomson scattering, in which photons are scattered into random directions by interaction with electrons. As the electromagnetic field of a photon always is orthogonal to its heading, only the component of the photon’s original polarization orthogonal to its new direction survives the scattering. The total radiation leaving any given point in the plasma in any given direction is then the sum of contributions from photons incident from all directions. As illustrated in figure 1.6, this

<sup>10</sup>This makes  $E$  and  $B$  non-local: Their value in one point depends on the  $Q$  and  $U$  over the whole sky, but with the greatest weight from nearby points. This non-locality is the greatest disadvantage of the  $E$ - $B$  parameterization.

will introduce a net polarization in the outgoing radiation if the incoming radiation has a quadrupole moment [12].




Figure 1.6: Local quadrupole anisotropies introduce linear polarization at the surface of last scattering. In this example, unpolarized light (i.e. light with equal amounts of all types of polarization) arrives at an electron from above and the left. The outgoing radiation inherits the horizontal polarization component from the vertically incident radiation and vice versa for the horizontal radiation. If the horizontally and vertically incident radiation have different intensities, the outgoing radiation will be linearly polarized. Based on figure in [12].

There are three qualitatively different types of fluctuations present at the surface of last scattering that can set up quadrupole moments:

**Scalar perturbations**, or simply density perturbations, set up temperature anisotropies by compressing or expanding the gas.

**Vector perturbations** represent vortices in the velocity field, which create a spatially varying Doppler shift in the radiation. Vector perturbations are not expected to be present at detectable levels.

**Tensor perturbations** are quadrupolar distortions in the metric which directly induce a quadrupole moment in the photon field.

These perturbations introduce quadrupole moments with different symmetries as illustrated in figure 1.7. The symmetry of the scalar perturbations ensure that they cannot produce B-mode polarization. Hence, since vector perturbations are unlikely to be present, a detection of B-mode polarization would be evidence for tensor perturbations<sup>11</sup>.

<sup>11</sup>There are some caveats to this. Firstly, E-modes may be turned into B-modes by lensing




Figure 1.7: Illustration of the local quadrupole anisotropy induced by scalar and tensor perturbations. Scalar perturbations correspond to  $m = 0$  modes, while tensor perturbations have  $m = \pm 2$ . The symmetry of scalar modes prevent them from inducing B-mode polarization. Figure borrowed from [12].

The polarization of the CMB was first detected by the DASI experiment [15] in 2002, and has since been measured by several other experiments, including WMAP [4], QUAD [16] and BICEP [17]. So far, only E-modes have been detected. The lack of detection of B-modes implies a bound of  $r < 0.72$  [17] based on the sensitivity of current experiments. This is still not competitive with the bounds from the temperature power spectrum, but as detector technology improves this is very likely to change.

### 1.3 Detector technology

Most of the experiments currently aiming to measure B-modes are based on polarization-sensitive bolometers. These detect radiation by the heat it imparts on the detectors, effectively counting photons while discarding their phase. With sufficient cooling (50 mK - 300 mK) a bolometer can be very sensitive, approaching the limit set by the discrete nature of photons at a few tens of  $\mu K \sqrt{s}$  [18]. However, the time it takes for the detector to reach equilibrium with the incoming radiation means that bolometers measure a slightly delayed and smeared-out signal, which must be taken into account in post-processing.

The alternative to bolometers is coherent amplifiers, which are basically radio antennas that measure both the amplitude and phase of the incoming

---

of the CMB by the matter distribution of the universe [13], and secondly, the presence of large magnetic fields in the early universe may excite vector modes, which also produce a B-mode signal [14]. Luckily, these effects can be decoupled from B-modes from tensor perturbations through their different scale behavior.



radiation. The presence of phase information makes it possible to measure Q and U simultaneously, as well as allowing for interferometry and various techniques for reducing systematics in hardware. However, coherent amplifiers have a minimal noise level set by quantum noise, corresponding to an effective temperature of  $T_q = \frac{h\nu}{k\log 2} \approx \frac{\nu_{\text{GHz}}}{20} \text{ K}$  [18]. This limits the sensitivity of coherent amplifiers at high frequencies, with the quantum noise becoming prohibitively high at  $\nu \gtrsim 300 \text{ GHz}$ .

For comparison, the CMB has greatest intensity at 57 GHz and is dominated by synchrotron at frequencies below about 30 GHz and dust above 200 GHz or so. The quantum limit is thus not a show-stopper for CMB observations with coherent amplifiers. In practice, neither bolometers nor coherent amplifiers reach their theoretical sensitivities. Not only are the detectors themselves imperfect; the signal itself contributes to the overall system temperature, and for ground-based experiments this includes a large contribution from the atmosphere. This puts a rather restrictive limit on the sensitivity achievable with a single detector.

The most effective way of increasing overall sensitivity is therefore to add more detectors. If the noise in each detector is uncorrelated, the sensitivity will fall<sup>12</sup> as the square root of the number of detectors. The challenge is then to fit as many detectors as possible into a focal plane. This is an area where bolometers used to have a big advantage due to being smaller, cheaper and easier to mass produce, which helped make bolometers the norm.

Recent developments in coherent amplifiers have changed this, however. With the development of a new miniaturized polarimeter-on-a-chip design, it is finally feasible to build large arrays of coherent amplifiers with sensitivity competitive with modern bolometer arrays.

## 1.4 The Q/U Imaging Experiment

These new detectors are being fielded for the first time in the Q/U Imaging Experiment (QUIET), making it the only current B-mode experiment built with coherent amplifier detectors. These detectors make QUIET ideally suited for observations at frequencies below  $\sim 100 \text{ GHz}$ , which includes the frequency range where the CMB polarization spectrum attains its greatest intensity relative to the foregrounds. Due to bolometer characteristics, other current B-mode experiments operate above  $\sim 100 \text{ GHz}$ , and this frequency difference implies that QUIET will be exposed to different foregrounds than its competitors. This, together with different and typically lower instrument systematics, makes QUIET's measurements or limits on

---

<sup>12</sup>The sensitivity of a detector is measured in terms of its noise level, which leads to a rather misleading terminology where *lower* sensitivity is better.

B-modes an important cross-check with regards to those of bolometer experiments, and also an important result in their own right.

At a few  $\mu\text{K}$ , the CMB E-modes are extremely faint, and the B-modes are at least several times fainter. Even the most sensitive detectors are therefore strongly noise dominated, and hence a large amount of data reduction is needed to extract a physically relevant signal from the raw telescope data. A radiometer like QUIET is basically a nothing but a fancy radio antenna, and the direct output from the telescope is a time-series of voltages. The process of reducing this to an estimate of  $r$  can be separated into the following steps.

1. Calibration
2. Map making
3. Component separation
4. Power spectrum estimation
5. Parameter estimation

These steps conceptually follow each other, and can be thought of as a processing pipeline, where raw time-ordered data are fed into one side, and parameter estimates issue from the other. Following this analogy, the software and methods responsible for these steps is collectively called an “analysis pipeline”.

This linear model should not be taken too literally: There is usually significant feedback in a realistic pipeline, with for example the calibration step depending on map making, or the latter 3 steps being jointly solved at the same time.

Due to the heavy amount of processing involved in an analysis pipeline, there is a significant risk for error which could bias or otherwise degrade the result. Null-tests and end-to-end simulations are powerful techniques for guarding against this, but in addition to this, it is common to have two or more independent implementations of the pipeline. QUIET employs two such pipelines: A pseudo- $C_l$  based pipeline [19, 20], and a maximum likelihood-based pipeline (see fig. 1.8). The latter was primarily developed by a University of Oslo based team consisting of H. K. K. Eriksen, I. K. Wehus and myself (S. K. Næss), with collaborators from Columbia University. Most of my work in this thesis has been dedicated to the development and application of this pipeline, which will be described in detail in the following chapters.

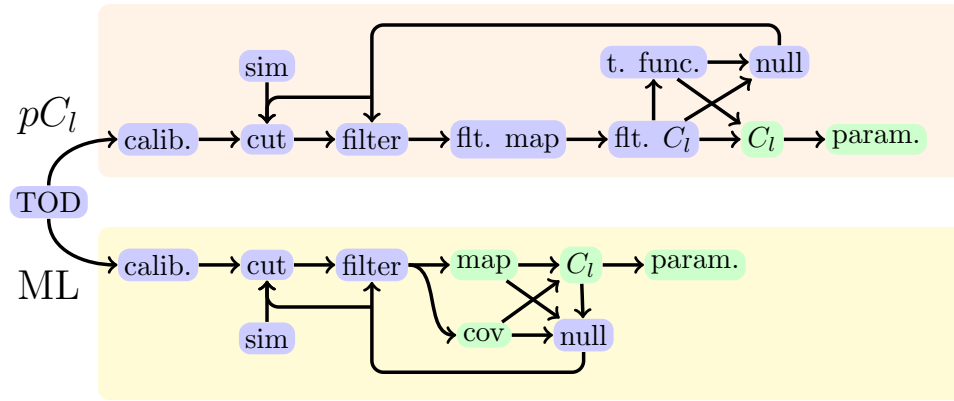


Figure 1.8: Stylized overview of QUIET’s two analysis pipelines. The major internal steps are marked in blue and green, where green indicates a result that is useful in its own right. “flt.” is short for “biased due to filtering”, while “t. func” is short for “transfer function”. Both pipelines share the major steps *data preparation* (calibration, cuts and filtering), *map making*, *validation* and *parameter estimation* (power spectrum and cosmological parameters).



## Chapter 2

# Telescope and calibration




Figure 2.1: Left: QUIET is located at  $67.761667^\circ$  west,  $23.028222^\circ$  north at 5020 m elevation on the Chajnantor plateau in Chile, sharing access and infrastructure with ALMA, ACT, APEX, etc. Right: The QUIET telescope itself, after the installation of the upper ground screen. Map courtesy of Google, telescope image by Joe Zuntz.

QUIET fielded its first array of coherent amplifiers from August 2008 to June 2009, consisting of 19 so-called “modules”, each made up of four individual detectors. This array is sensitive in the Q-band (centered on 43 GHz), and consists of two parts: A polarization sensitive sub-array of 17 modules (of which 15.75 were usable) with a combined sensitivity of  $69\mu K\sqrt{s}$ , and a temperature sensitive sub-array consisting of 2 modules configured to act together as a differential detector.

A second array of 91 W-band (95 GHz) detectors was deployed from July 2009 to December 2010, consisting of 85 polarization modules (of which 77.25 were usable) with a combined sensitivity of  $85\mu K\sqrt{s}$ , and 6 temperature modules (of which 5.5 were usable) arranged in 3 differencing pairs.

Both arrays were arranged in a hexagonal pattern in the telescope’s fo-




Figure 2.2: The location of QUIET’s 4 CMB patches and two galactic foreground patches in galactic coordinates. The bright part of the sky indicates the observable region based on the telescope’s location and elevation limits. Several of QUIET’s patches overlap with those of other current or planned experiments, including ABS [23], EBEX [24], BICEP [17] and SPIDER [25], which are also shown. The overlaid grid is in equatorial coordinates. Figure courtesy of I. Wehus.

cal plane, with each module fed by its own feed horn, which is thus shared by the four detectors within the module. The projected radius of both arrays on the sky was 3.75 degrees, with an average beam size of  $27'$  for the Q-band array and  $12'$  for the W-band array.

These arrays were mounted in a 1.4 m Dragonian [21, 22] telescope located at 5080 m altitude at the Chajnantor plateau in the Atacama desert in Chile. The telescope had three free axes: azimuth, elevation and rotation of the focal-plane about the boresight, called the “deck” axis. In order to reach the necessary sensitivity, observations were focused on 4 CMB-dominated  $35^\circ \times 35^\circ$  patches on the sky plus 2 foreground-dominated patches of similar size (fig. 2.2). These were chosen such that at least one target was available throughout the day, allowing observations to continue 24 hours of the day.

During normal operation of the telescope, the boresight is aimed roughly one focal plane radius ahead of the edge of the target patch. The telescope is then slewed backwards and forwards in azimuth at about 0.1 Hz with an amplitude of 15 degrees at constant elevation (see figure 2.3) while the sky drifts past, until the patch exits the telescope’s field of view. The boresight is then repositioned, and the cycle continues. This scanning pattern ensures that each detector spends as much time as possible looking through the atmosphere at constant optical depth. The temperature intensity of the atmo-




Figure 2.3: The first 400 seconds of data from a Constant Elevation Scan (CES). The top three panels show azimuth, elevation and the deck orientation respectively. The azimuth amplitude is about  $15^\circ$ , while both elevation and deck are almost constant, fluctuating by about  $0.5''$  and  $1'$  respectively. The fourth panel shows the demodulated readout from one of the detectors of the 95 GHz array after decimation to 25 Hz, and the fifth panel contains the corresponding power spectrum and the best fit  $\frac{1}{f}$ -noise profile.

sphere is much greater than the CMB fluctuations QUIET aims to measure, and is proportional the optical depth, which depends strongly on elevation but only varies very weakly with azimuth and time. Scanning at constant elevation means the atmosphere only shows up as a constant or very slow drifts in the data, making it easy to filter out with a high-pass filter.

Thus, the useful data taking of the telescope comes in the form of Constant Elevation Scans (CES), typically lasting about one hour, interspersed with periods of about 5 minutes without data taking.

## 2.0.1 CES detection

For an ideal telescope, the list of the individual scans, their start and end times, their target object, and type of scan would be available as one of

the outputs from the telescope. In the case of the QUIET telescope, part of this information is available, but after problems with miscategorized scans, missing scans and undetected scans, this information was determined to be incomplete and unreliable. We have therefore adopted a heuristic CES detection based on the actual time-ordered data. The two pipelines of the project, the maximum likelihood pipeline described here and the PCL pipeline described in [19] each developed their own CES detection algorithm, which were then refined until they reached an agreement of  $\sim 95\%$  when ignoring small differences in the detected start- and end- times of the CESes. I will here describe the ML pipeline's CES detection algorithm, i.e. the practical definition of "CES" used in this analysis.

Ideally, a CES is simply a continuous time-period where the telescope is scanning in azimuth while keeping the elevation and deck angle constant. However, as figure 2.3 illustrates, the hardware cannot stay totally still, and elevation and deck fluctuate by about half a second of arc and one minute of arc respectively during a typical CES. Furthermore, the telescope does not spend all its time in data-taking mode. One of the output streams of the telescope is the data taking mode, which has the value 3 during normal data taking. Sadly, this data stream has frequent glitches where the value departs from 3 for several seconds while data taking continues as normal. The CES conditions above must therefore be qualified with tolerances. Finally, the Fourier methods used in the map making and filtering depends on a constant sampling rate, so the CES detection must also ensure that the time between samples is very close to constant during the scan.

With these problems in mind, the practical definition of a CES is: A consecutive series of two-second chunks of data, which must each fulfil the following criteria:

1. The deglitched data taking mode is 3
2. The elevation is within the allowed limits of  $[-\frac{\pi}{2}, \frac{\pi}{2}]$
3. No more than 30 consecutive samples<sup>1</sup> have an absolute azimuth change of less than  $10^{-5}$  radians per sample.

The deglitched data taking mode here refers to the stream of data taking modes after removing deviations from 3 of less than 3 seconds in duration. Additionally, the series of chunks as a whole must satisfy

1. No samples must have a time-stamp that deviates by more than 3 ms from the value expected from a constant 100 Hz sampling rate.
2. No samples can have an absolute difference between its elevation and the average elevations of its chunk and the chunks before it of more than  $3 \cdot 10^{-4}$  radians.

---

<sup>1</sup>The raw data at this stage are at 100 Hz. It is reduced to 25 Hz before the main part of the data analysis.



3. Similarly, the deck angle must not deviate by more than  $3 \cdot 10^{-3}$  radians.

Whenever one of these conditions is violated, the current CES candidate is ended and a new one started. The beginning time of each CES candidate is then adjusted so that it is at least 1 minute later than the last phase switch event<sup>2</sup> that happens before the end of that candidate. Finally, the candidate is accepted as a usable constant elevation scan if its duration is longer than 5 minutes.

### Patch detection

Each of these CESes is then classified according to the objects hit during its course. The boresight pointing for each sample of the CES is translated to galactic coordinates<sup>3</sup> and the angular distance  $d$  to a predefined set of candidate objects is calculated. A hit is defined as  $d < R + r_b + r_v + r_f$ , where  $R$  is an approximate radius of the object,  $r_b$  is the beam 5 sigma radius,  $r_v$  is half the distance traveled during a sample and  $r_f$  is the focal plane radius. The objects hit during the CES are then sorted according to their brightness, and the brightest one is assigned to the CES. The list of objects considered can be seen in figure 2.1. This classification scheme ensures that a scan of a CMB patch where the Moon happens to pass in front of the patch registers as a usable Moon scan instead of a contaminated CMB scan.

### Implementation

The CES detection algorithm was implemented in the Fortran 90 program **ces\_detect** through a greedy algorithm – that is, an algorithm which tries to maximize the length of the current CES, without considering how this might affect the length of the next CES. This approach allows the program to iterate through the data set in small chunks, which keeps the memory requirements very low while still requiring only a single, sequential pass through the data. This comes at the cost of not necessarily finding the optimal CES partitioning, but any loss from this is minuscule.

Due to the size of the data set ( $\sim 15$  TB for the W-band analysis), **ces\_detect** benefited greatly from MPI parallelization. This was implemented by splitting the full data set into  $N_{\text{proc}}$  slices of equal size (in number of data files, which corresponds closely to duration and data volume), each of which are processed independently by each MPI process. CESes that span slice

---

<sup>2</sup>A phase switch event is a hardware event that is followed by a period of up to a minute of unstable gain and noise levels due to temperature changes in the electronics.

<sup>3</sup>This conversion assumes an ideal mount model because arcminute precision is not necessary here, and to avoid too many backwards dependencies in the pipeline.

Name	Lon. (°)	Lat. (°)	Rad. (°)	$N_{\text{CES}}^Q$	$N_{\text{CES}}^W$
Sun			0.25	0	0
Moon			0.25	82	107
Venus			0	36	30
Jupiter			0	78	466
Mars			0	0	1
Saturn			0	3	0
Tau A	184.557	-5.784	0	18	0
RCW 38	267.927	-1.050	5	53	124
CMB-1	292.200	22.800	15	873	1584
CMB-2	243.200	-35.300	15	736	1344
CMB-3	304.600	-69.100	15	802	1059
CMB-4	7.000	-62.000	15	322	647
G-1	0.0833	-0.0667	10	189	352
G-2	329.100	0.000	15	295	568

Table 2.1: The objects considered in the automatic CES classification. These are sorted into priority classes by brightness, separated by horizontal lines. The location in galactic coordinates, and an approximate radius is displayed for the stationary objects. Shown in the last two columns are the number of CESes identified for each object for the Q-band and W-band observing seasons.

boundaries are handled by allowing each MPI process to finish processing the last CES in its slice even if it extends outside that slice; the resulting overlapping CESes from this and the next MPI task can then be automatically combined in postprocessing.

### 2.0.2 Alternative scanning patterns

All the CMB data are collected using constant elevation scans, but some calibration measurements employ an alternative scanning pattern called raster scans, which are series of very small amplitude ( $\sim 1^\circ$ ), very short duration ( $\sim 1\text{min}$ ) almost constant elevation ( $\Delta\text{El} \lesssim 0.5^\circ$ ) scans. The primary target of these is the strongly polarized supernova remnant Tau A, of which there are  $\sim 10^4$  such raster scan segments.

## 2.1 Detector data and noise properties

The telescope provides a 100 Hz time-stream modulated at 50 Hz for each detector<sup>4</sup>. After correcting for a well-measured non-linearity in the detector response and demodulating<sup>5</sup>, the result is one 50 Hz time-stream for each detector, for a total of 72 data streams for the Q-band array and 364 for the W-band array.

Though the detectors are quite sensitive, the CMB polarization is weaker still, resulting in a S/N ratio per sample of about  $10^{-3}$ , meaning that the data are thoroughly noise-dominated. It is therefore critical to understand the statistical properties of the noise.

As illustrated in figure 2.4, the noise is Gaussian, but suffers from time correlations in the form of  $1/f$ -noise, which results in a noise power spectrum of the form

$$\phi(f) \equiv \sigma_0^2 (1 + [f/f_{\text{knee}}]^\alpha). \quad (2.1)$$

This profile only holds within a CES; over longer periods the noise is not stationary, and one must therefore estimate  $\sigma_0$ ,  $f_{\text{knee}}$  and  $\alpha$  individually for each CES. The time dependence is illustrated for one of these parameters,  $\sigma_0$ , in figure 2.5.

Furthermore, the noise in each detector is not independent. The detectors within a module are on average about 40% correlated, while inter-module correlations are negligible (see figure 2.7). This means that detectors within a module must be analyzed jointly rather than one-by-one,

<sup>4</sup>Two secondary data channels per detector are also provided: A “TP” stream sensitive to the stokes  $I$  parameter, but with high noise correlations, and a noise channel. These are not used for the main analysis, but are useful for pointing calibration.

<sup>5</sup> $d_i = \frac{1}{2} \sum_{j=0}^1 m_{2i+j} p_{2i+j}$ , where  $d$  is the demodulated signal,  $m$  is the modulated signal and  $p$  alternates between 1 and -1.




Figure 2.4: The noise is well-fit by a Gaussian profile.

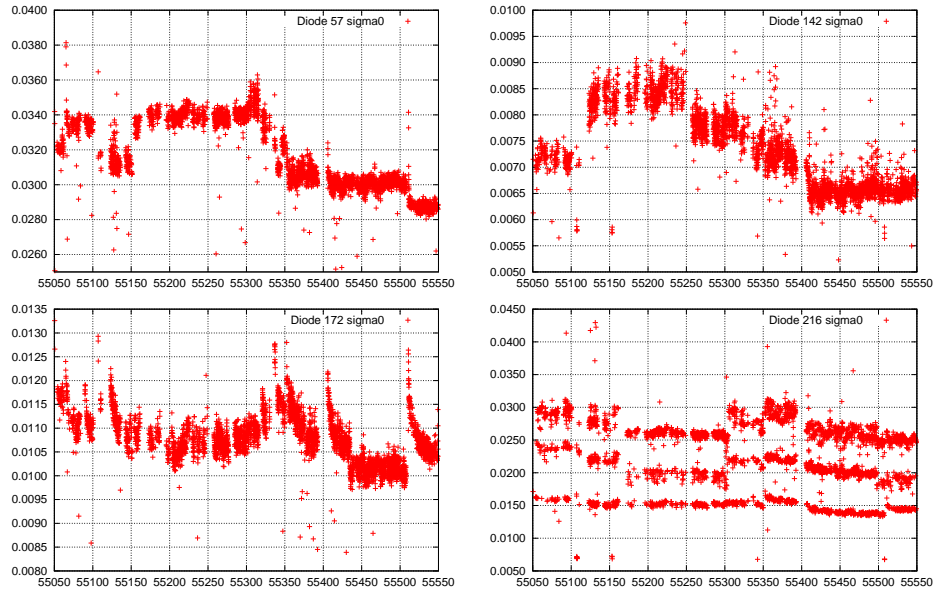


Figure 2.5: The white noise floor  $\sigma_0$  shows a complicated time dependence with a different pattern for each module, and to a lesser degree for the detector within the module. The four panels show the measured  $\sigma_0$  values per CES for four different detectors. Large changes in  $\sigma_0$  happen at discrete events which are common for the detectors. The response to these events differs, but is correlated between the detectors, as illustrated in figure 2.6.




Figure 2.6: The time-dependence of the white noise floor  $\sigma_0$  is strongly correlated between the detectors. This figure shows the W-band season correlations in  $\sigma_0$ .

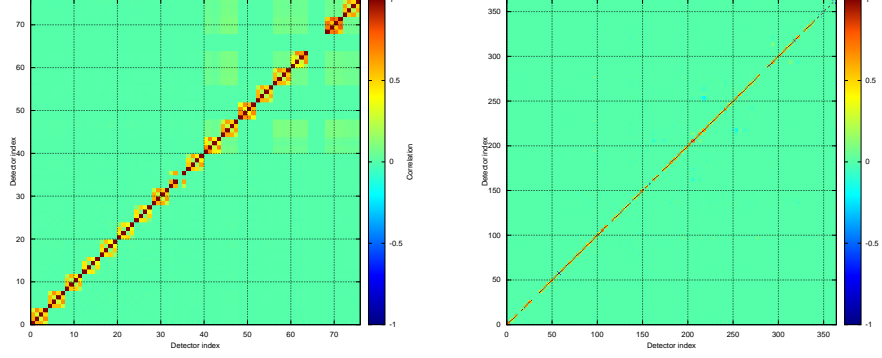


Figure 2.7: Detector white noise correlations for the 43 GHz (left) and 95 GHz (right) arrays. Within modules correlations are strong: on average 40% at 43 GHz and 37% at 95 GHz. Detectors in different modules are on average less than 1% correlated. Only data for detectors that are actually used in the analysis is displayed here. Unused detectors therefore show up as gaps.

which comes as a cost in time and memory use. To complicate matters further, figure 2.8 shows that the degree of correlation is frequency dependent, changing gradually with frequency below 1 Hz.

All in all, these effects result in a final noise model of

$$N_{dd'ff'} = \sqrt{\phi_d(f)\phi_{d'}(f)}C_{dd'f}\delta_{ff'}, \quad (2.2)$$

where  $N_{dd'ff'} \equiv \langle n_{df}n_{d'f'}^\dagger \rangle = \mathcal{F}_{ft}N_{dd'tt'}\mathcal{F}_{t'f'}^{-1}$  is the frequency-domain noise covariance matrix between detectors  $d$  and  $d'$  and frequencies  $f$  and  $f'$ , and  $C_{dd'f}$  is the detector correlations, and where all parameters vary from CES to CES.

### 2.1.1 Noise estimation

We estimate the noise parameters in a two step process: First, we fit a  $1/f$  profile independently for each detector, and then we measure the correlations. This is not as optimal as the full maximum likelihood solution, which would maximize the joint likelihood, but it is much faster, and still unbiased.

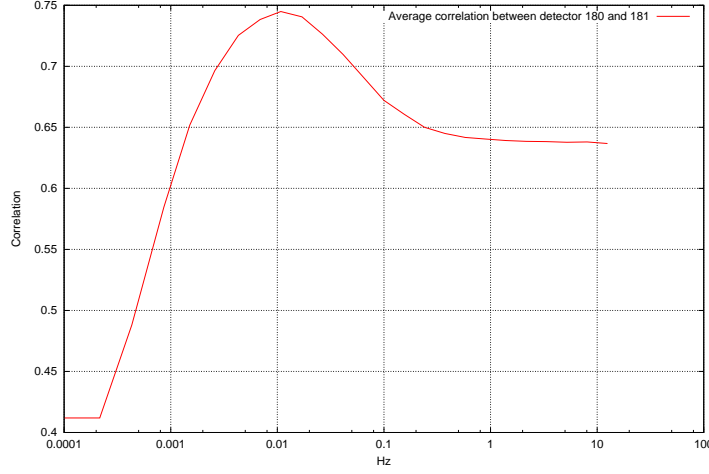


Figure 2.8: The average correlation between detector 180 and 181 in the 95 GHz array as a function of frequency. The correlation becomes frequency dependent in the  $1/f$ -dominated area.

### 1/f profile

Assuming that the time-ordered data are completely noise dominated, and that the noise is Gaussian and diagonal in Fourier space with mean zero and unknown variance  $\langle \tilde{n}_f \tilde{n}_f^* \rangle = \langle P_f \rangle$ , where  $\tilde{n}_f \equiv \mathcal{F}_{f,t} n_t$ , and  $P_f$  is the observed TOD power spectrum with length  $N$ , the likelihood for  $\phi(f)$  is

$$-\log \mathcal{L} = \frac{1}{2} \sum_f \left[ \frac{P_f}{\phi(f)} + \log(2\pi\phi(f)) \right]. \quad (2.3)$$

The maximum likelihood estimator for the parameters  $\sigma_0, f_{\text{knee}}, \alpha$  is given by minimizing equation (2.3) with respect to the model

$$\phi(f) = \sigma_0^2 \left( 1 + \left[ \frac{f}{f_{\text{knee}}} \right]^\alpha \right). \quad (2.4)$$

With respect to some parameter  $X$ , this is done by solving

$$(-\log \mathcal{L})_{,X} = \frac{1}{2} \sum_f \frac{\phi(f)_{,X}}{\phi(f)} \left( 1 - \frac{P_f}{\phi(f)} \right) = 0, \quad (2.5)$$

which for  $X = \sigma_0^2$  results in

$$\sigma_0 = N_{\text{freq}}^{-1} \sum_f \frac{P_f}{1 + \left[ \frac{f}{f_{\text{knee}}} \right]^\alpha}. \quad (2.6)$$

The remaining parameters can be determined with a nonlinear search.

For the QUIET detectors, it turns out that the assumption of complete noise dominance is not realistic, even when observing a weak source like the CMB. Figure 2.9 shows the season average of the observed power spectrum for one W-band detector, based on only constant elevation scans with good observing conditions and observation of CMB patches only. It shows the expected  $1/f$ -profile for most frequencies, but at high frequencies there is a forest of spikes, with the most prominent one being a 10 Hz alias of the 60 Hz mains power at the site. For the Q-band array, the most largest deviation is a broader bump of excess power between 5.4 Hz and 6.35 Hz.

Since these problems are located in a distinct frequency range, they can be dealt with by down-weighting these frequencies in the likelihood,

$$-\log \mathcal{L} = \frac{1}{2} \sum_f w_f \left[ \frac{P_f}{\phi(f)} + \log(2\pi\phi(f)) \right], \quad (2.7)$$

with the weights  $w_f$  being zero in problematic regions and one otherwise. These weights also allow us to exclude the multiples of the scanning frequency from the noise estimate. Any signal which mostly varies with azimuth, which includes ground pickup through sidelobes as well as the actual sky signal, will have the greatest contribution at these frequencies, so we avoid these as a precaution.

The performance of the resulting masked  $1/f$  estimator at recovering the input parameters for noise only simulations is illustrated in figure 2.10, which shows that the estimator is unbiased. However, this unbiasedness only carries over to real data if no significant signal is left unmasked, and if the real noise actually follows a  $1/f$  profile. This is investigated in figure 2.11. It shows the season co-added<sup>6</sup> detection of a deviation between the data and model for one detector, expressed as the number of standard deviations per 1.25 mHz bin in frequency. This can be calculated by noticing that  $\frac{P_f}{\psi(f)}$  follows a scaled chi-squared distribution with 2 degrees of freedom and mean 1. Summing these in bins over the whole season, we arrive at

$$\chi_b^2 = 2 \sum_{\text{CES}} \sum_{f \in b} \frac{P_f}{\psi(f)}, \quad (2.8)$$

which is an unscaled chi-squared with  $N_b = \sum_{\text{CES}} \sum_{f \in b} 1$  degrees of freedom.  $N_b$  will typically be of the order of  $10^5$ , so we can use a Gaussian approximation to express this as the number of standard deviations away from the expectation value:  $\frac{\chi_b^2 - N_b}{\sqrt{2N_b}}$ .

The figure shows that deviations from a pure  $1/f$  profile are detectable at the 2-5 sigma level per bin for most bins, with larger deviations at the

---

<sup>6</sup>Using only CESes from CMB patches that pass the data quality cuts.



lowest frequencies. It is clear that a  $1/f$  profile is only an approximation to the true noise shape, which is more complicated. Still, it is a good enough approximation that deviations are not significant at the single-CES level, and when integrating over all unmasked frequencies, the deviation necessarily averages to zero. Hence, the overall effect of approximating the noise with  $\psi(f)$  is slightly suboptimal noise weighting in the map-making step, which does not lead to any bias in the final map, but may move a small amount of power between small and large scales in the noise covariance matrix.

The discussion above only considered observations with low signal-to-noise per sample, such as CMB observations. When observing strong sources such as the Moon (both polarization and temperature) or Jupiter (temperature), it is unfeasible to mask out their contribution in frequency domain, as every frequency will be contaminated. However, the signal is localized in time-domain, making it possible to perform the masking there instead. The problem can be formulated as one of sampling the noise parameters  $\theta = \{\sigma_0, f_{\text{knee}}, \alpha\}$  given knowledge of only parts of the data stream. That is, we wish to find  $P(\theta|d) = P(P_f(\theta)|d) \propto P(d|P_f(\theta))$ , where  $d$  is the incomplete time ordered data, with the masked samples missing.

Assuming that the unmasked samples in  $d$  have low signal-to-noise,  $d$  is a sample from  $N(0, C)$ , where  $C = M\mathcal{F}^{-1}P\mathcal{F}M^T$  and  $M$  is a projection operator from the full unmasked TOD to the masked TOD. While in principle straightforward, this approach does not scale well with TOD length, requiring memory of the order  $\mathcal{O}(N_{\text{samp}}^2)$  and processing time of the order  $\mathcal{O}(N_{\text{samp}}^3)$ , with  $N_{\text{samp}}$  usually exceeding  $10^5$ .

A much more efficient approach is to use the fact that one can sample from a joint distribution by iteratively sampling from its conditional distributions. This allows one to simplify the problem by *adding* more parameters to the joint distribution. In this case, what makes it difficult to estimate  $\theta$  is the presence of holes in the TOD, which makes the Fourier basis non-orthogonal. We can rectify this by adding the data inside the holes to the set of parameters to be estimated. We then sample  $(\theta, d')$  jointly through the iteration

$$d' \leftarrow P(d'|\theta, d) \quad (\text{gap filling}) \quad (2.9)$$

$$\theta \leftarrow P(\theta|d', d). \quad (2.10)$$

Here,  $d'$  is a gap-filled version of  $d$ . That is, the masked areas of  $d$  have been filled with a noise realization based on  $d$  and the noise parameters  $\theta$ .

The details of this more general estimator can be found in paper III of this thesis, but it is not needed for QUIET's CMB patches due to the low signal-to-noise per sample there.




Figure 2.9: The measured noise profile compared to the model. The green points are the measured power spectrum for a given CES and detector, where the white noise floor has been normalized to 1. The red curve is the season average of such curves for the same detector. The blue curve is an asymptotically correct  $1/f$  profile.

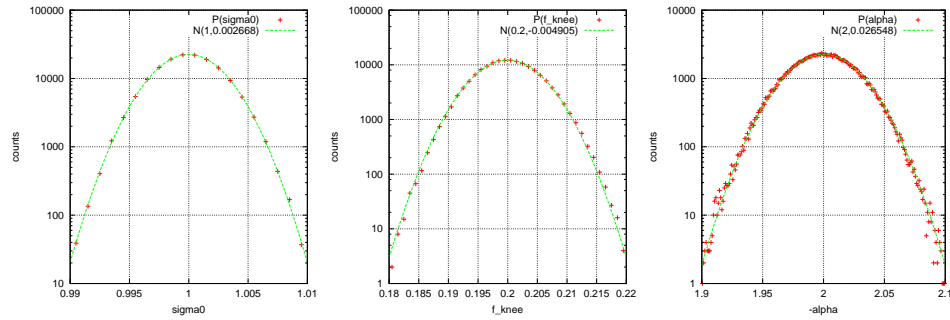


Figure 2.10: The distribution of 120000 noise model fits based on simulations with parameters  $\sigma_0 = 1$ ,  $f_{\text{knee}} = 0.2$ ,  $\alpha = -2$ . The estimator is unbiased with Gaussian errors.




Figure 2.11: Lower panel: The full-season significance of the deviation of the noise power spectrum from the best-fit  $1/f$  profile per 1.25 mHz bin. For a perfect match, the values should be normal distributed with mean 0 and variance 1. Upper panel: logarithmic plot of the square of the significance of the deviation. Two effects are visible in these plots: A trend towards higher than expected noise in the lowest bins, and spikes of varying significance at various frequencies. While these are significant when co-added over the season, they are not detectable for a single CES-detector (except for the 10 Hz spike), and they do not systematically add up during map making.

### Correlations

The covariance between noise streams  $\{n_d\}$  is defined as

$$\text{Cov}(n_d, n_{d'}) \equiv \langle (n_d - \langle n_d \rangle)(n_{d'} - \langle n_{d'} \rangle) \rangle. \quad (2.11)$$

Given a series of samples  $n_{df}$  for each stream in frequency space, and assuming a mean of zero, the covariance can be estimated as

$$\text{Cov}_{dd'} = \frac{1}{N_{\text{freq}}} \text{Re} \left( \sum_f n_{df} n_{d'f}^* \right). \quad (2.12)$$

This assumes that the covariance is the same for all the samples, but we have already seen that both the variance and correlations are frequency dependent for the QUIET detectors. We therefore subdivide the frequencies into bins, within which the covariance is approximately constant. By also allowing frequency weighting, we arrive at

$$\text{Cov}_{dd'b} = \left( \sum_{f \in b} w_f \right)^{-1} \text{Re} \left( \sum_{f \in b} n_{df} n_{d'f}^* w_f \right) \quad (2.13)$$

$$C_{dd'b} = \frac{\text{Cov}_{dd'b}}{\sqrt{\text{Cov}_{ddb} \text{Cov}_{d'd'b}}}. \quad (2.14)$$

The frequency binning results in sharp jumps in correlation when moving from one bin to another, which results in unacceptably long time-correlations in the noise model. This can be avoided by using spline interpolation between the bins.<sup>7</sup>

#### 2.1.2 Filters

The four main contaminants in the QUIET time-ordered data are atmospheric disturbances (i.e. bad weather), ground (fig. 2.14) and sun (fig. 2.13) pickup through sidelobes in the telescope beam, and high frequency spikes from the electronics. The typical shape and magnitude of these in frequency space is illustrated in figure 2.12 in comparison with the typical polarized CMB signal. For a single CES from a single detectors, the contaminants strongly dominate the CMB, and though they will tend to average down when more CES-detectors are co-added, they are strong enough that this will not bring them down to acceptable levels unless they are dealt with by cuts and filtering.

---

<sup>7</sup>This approach, with the spline being done logarithmically in frequency, is what QUIET actually uses. However, a possible improvement would be to not force  $C$  to be a correlation matrix, but instead use it to absorb deviations from a  $1/f$  profile. One would then use

$$C_{dd'b} = \left( \sum_{f \in b} w_f \right)^{-1} \text{Re} \left( \sum_{f \in b} \frac{n_{df} n_{d'f}^* w_f}{\sqrt{\phi_d(f) \phi_{d'}(f)}} \right). \quad (2.15)$$




Figure 2.12: Illustration of the different contributions to QUIET’s time-ordered data. The strongest contribution is the noise, which follows a  $1/f$  pattern with a white noise floor about 6 orders of magnitude (in power) above the polarized CMB itself, which appears as a set of bumps around harmonics of the azimuth scanning frequency. This picture is complicated by several systematic effects. At low frequencies, atmospheric effects (bad weather) may start to dominate, while the high frequencies are polluted by occasional narrow spikes. Neither the very low or very high frequencies have much CMB contributions, so the weather and spikes can be removed by applying a bandpass filter. However, ground pickup from telescope sidelobes enters as a systematic in the same frequency range as the CMB, and must be handled by an azimuth filter instead.

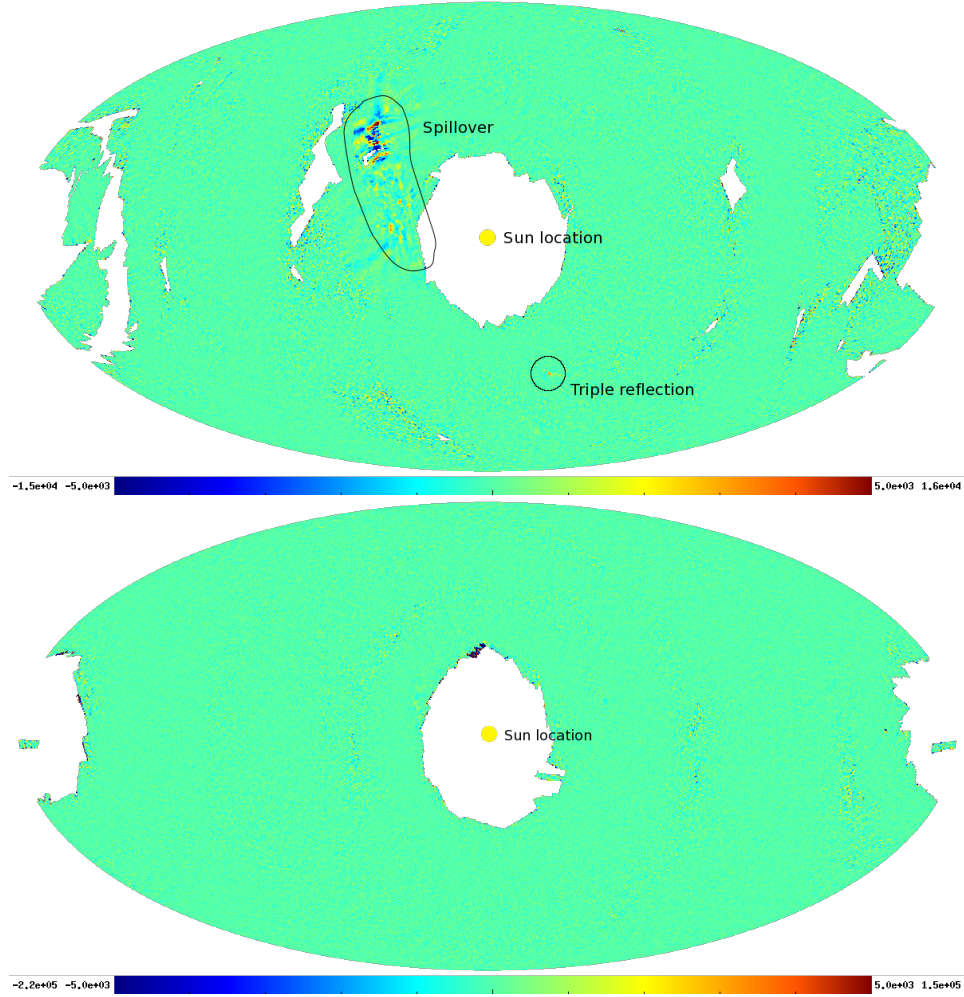


Figure 2.13: Map of the stokes Q component of the sky in Sun-relative coordinates for the detector and deck orientation that shows the strongest contamination from the Sun. The center of the map is the position of the sun. The central  $\sim 30^\circ$  are unexposed due to sun avoidance. The first and second panel cover the part of the W-band observing season before and after the upper ground screen was installed. Before, the triple reflection sidelobe and the spillover sidelobe are both visible, with amplitudes of up to 10mK; after, they are both gone. The strength of both sidelobes and the position of the spillover sidelobe varies with deck angle and position in the focal plane. See [?] for a more comprehensive review of the sidelobes.

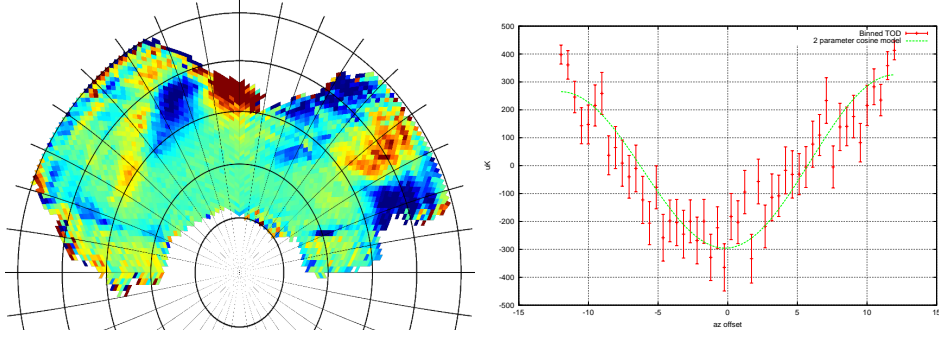


Figure 2.14: Left panel: A stokes Q map in horizontal coordinates based on the CMB scans of the Q-band observations at a deck orientation of  $165^\circ$ . In these coordinates, any signal from the sky will average down, while any function of az and el, such as the ground signal, will add up. Due to their different locations in the focal plane, each detector will see a different ground pattern – this image shows the average of these contributions. Significant ground pickup is visible: The color range in this map goes from -0.1 mK (blue) to 0.1 mK (red). Right panel: Azimuth-binned TOD for a single CES-detector from the Q-band season, compared to a best-fit 2-parameter cosine model.

From the figure it is clear that normal weather and spikes can be effectively dealt with by using an apodized bandpass filter (see section 4.5.1) between roughly 0.1 Hz to 4.5 Hz (Q-band) and 9.5 Hz (W-band), which is indicated with the blue region in the figure. However, bad weather can be much worse than indicated here, and may in the worst cases dominate all the way up to several Hz. These latter cases must simply be cut (see section 2.1.3 for details), but in less severe cases it is still possible to filter it out by using a harsher high-pass filter. But increasing the threshold of the high-pass filter must be done sparingly: The primary goal of QUIET is to put upper bounds on the primordial B-modes, and these are most easily detected in the multipole range  $l \lesssim 70$ , and a significant fraction of the power in this range comes from frequencies below 1 Hz. We therefore use an adaptive scheme which adjusts the filter cutoffs individually for each CES-detector as follows:

**Highpass cutoff** The lowest frequency  $f_{\text{hp}} = (i + \frac{1}{2})f_{\text{scan}}$  for  $i \in \mathbb{Z}^+$  for which the chisquare of the data at all the higher frequencies  $\chi^2 = \sum_{f=f_{\text{hp}}}^{f_{\text{max}}} \frac{P_f}{\phi(f)}$  is within 4 sigma of the expectation value, and for which all scan frequency harmonics above this frequency also have chisquares within 4 sigma of the expectation value, based on the 49 modes centered on each such harmonic. The resulting value is then capped from below to  $f_{\text{hp}} \geq \frac{3}{2}f_{\text{scan}}$  based on the null-tests.

**Lowpass cutoff** The frequency  $f_{\text{lp}} = \min(f_{\text{spike}}, f_{\text{bump}})$ , where  $f_{\text{spike}}$  is the highest frequency such that the range  $[2\text{Hz}, f_{\text{spike}}]$  contains no single modes  $f$  with  $\chi_f^2 = \frac{P_f}{\phi(f)} > 15$ , and  $f_{\text{bump}}$  is the highest frequency such that no consecutive set of 100 Fourier modes within the range  $[2\text{Hz}, f_{\text{bump}} + 50\Delta f]$  has a chisquare more than 5 sigma away from the expectation value. Here  $\Delta f$  is the frequency interval per Fourier mode.

Combining these cutoffs with  $1/f$ -type apodization<sup>8</sup>, we get a total filter profile of

$$F_f = \left[ 1 + \left( \frac{f}{f_{\text{hp}}} \right)^{\alpha_{\text{hp}}} \right]^{-1} \left[ 1 + \left( \frac{f}{f_{\text{lp}}} \right)^{\alpha_{\text{lp}}} \right]^{-1}, \quad (2.17)$$

with  $\alpha_{\text{hp}} = -20$  and  $\alpha_{\text{lp}} = 300$ , which results in typical correlation lengths of 1-2 minutes.

Ground pickup is not localized to a well-defined set of frequencies. They are spread out over the same frequencies as the CMB itself, and hence can't be handled with a bandpass filter. It is formed by sidelobes hitting the ground, and because the ground, unlike the sky, is stationary with respect to the telescope, the observed ground pickup will depend only on the telescope azimuth, elevation, deck orientation, the detector's position and orientation in the focal plane, and the ground signal itself. Hence, during a CES, where elevation and deck orientation are constant, and the ground properties presumably do not change noticeably, the ground pickup will be a function of azimuth only. We therefore employ a filter in azimuth to reduce the impact of the ground. Implementation-wise, this is handled as a generalized filter (section 4.5.2) with a basis choice<sup>9</sup> of

$$b_i(\text{az}) = \cos \left[ i\pi \frac{\text{az} - \text{az}_{\text{min}}}{\text{az}_{\text{max}} - \text{az}_{\text{min}}} \right]. \quad (2.18)$$

As the number of basis functions used for the filter increases, so does its ability to pick up sharp changes in the ground signal, but this comes at the cost of also removing more of the CMB and a large increase in computational cost due to the need to project out these modes from the covariance matrix. The number of basis functions used for the filter should therefore

---

<sup>8</sup>We also evaluated cosine apodization of the form

$$F_f = \begin{cases} 0 & \text{if } f \leq f_c - \Delta \\ \frac{1}{2} \{1 - \cos(\frac{\pi}{\Delta} [f - (f_c - \Delta)])\} & \text{if } f_c - \Delta < f \leq f_c + \Delta \\ 1 & \text{if } f_c + \Delta < f \end{cases}, \quad (2.16)$$

which has the advantage of transitioning from no damping to complete damping in a finite frequency range, but this resulted in much greater correlation lengths than a  $1/f$ -type filter.

<sup>9</sup>Chebyshev or Legendre polynomials would be other good alternatives.



be large enough to make the TOD as a function of azimuth consistent with noise, but no larger. Defining  $a_i$  as the best-fit amplitude of the  $i$ 'th basis function, and the deviation from the expected effective chisquare as

$$\sigma_n \approx \frac{\chi_n^2 - (N_{\text{samp}} - n)}{\sqrt{2(N_{\text{samp}} - n)}} \quad \text{with} \quad \chi_n^2 = \sum_t \left( d_t - \sum_{i=1}^n a_i b_i(\text{az}_t) \right)^2,$$

we choose the number of modes to include,  $n_{\text{basis}}$ , to be the lowest value that fulfills  $\sigma_{n_{\text{basis}}} < \min\{\sigma_1 \dots \sigma_{15}\} + 1$ . Typical values are in the range 0 to 4.

The final major contaminant, the signal from the Sun as seen through the sidelobes, only affects the small fraction of CESes that enter the sidelobe-contaminated region in sun-relative coordinates (see figure 2.13)<sup>10</sup>. The sidelobes are not sufficiently mapped out to allow a reliable filter to be constructed, so this effect must instead be handled through cuts.

### 2.1.3 Cuts

The QUIET maximum-likelihood pipeline employs an array of cuts to safeguard against bad weather, contamination from sidelobes and instrumental defects. These cuts are summarized in table 2.2.

Most of the cuts are tests of consistency between the data and the noise model. Because QUIET is strongly noise-dominated during CMB scans, any deviation of the data from the noise model during a single CES means that either the data have been contaminated by a non-CMB signal, or the noise model does not fit the actual noise properties. In maximum-likelihood map-making both of these cases need to be avoided: Extra signal will bias the map, while wrong noise will result in a suboptimal map and bias the covariance matrix.

The noise consistency cuts can be divided into two subclasses: Direct cuts based on data that will actually be used in the analysis (i.e. data that survive the filter), and indirect cuts which use the filtered data as a proxy for detecting effects that are likely to also contaminate the unfiltered data. For direct cuts, the goal is to detect deviations from noise, and a cut threshold of 4-5 sigma<sup>11</sup> is high enough to avoid significantly truncating the noise distribution through false positives while still being strict enough to detect all but the weakest effects.

<sup>10</sup>These sidelobes are so weak that they are only detected when they hit the sun. A figure of the sky in moon-relative coordinates was consistent with no signal.

<sup>11</sup>This is for fixed-location cuts. For variable location cuts like parts of the spike cut, the threshold should be somewhat higher to compensate for the added opportunities for the increased opportunity for false positives due to thousands of modes being tested. This is called the "look elsewhere effect" in some fields.

Name	Cut Q (%)		Cut W (%)		Description
	P	T	P	T	
Static	12.3	5.3	9.7	8.9	Broken detectors, missing data or precalibration data.
Gain	7.4	0.0	11.5	50.0	Cuts detectors where no gain model has been established due to missing calibration observations.
Noise fit	7.4	0.0	9.1	8.3	CES-detectors where the 1/f-profile fit failed. This usually happens due to bad weather.
Noise outlier	9.3	40.8	12.3	55.4	CES-detectors with $\sigma_0$ or $\alpha$ more than 5 sigma away from that detector's median, or $f_{\text{knee}}$ more than 5 sigma above the median.
ADC residuals	4.0	19.3	6.5	37.6	Cuts if evidence for a non-linear ADC response is greater than 5 sigma. In practice, this is also sensitive to bad weather.
Weather	8.2	6.4	9.9	8.5	Cuts if the rms of the 10 s-rms values of the TOD is more than 5 sigma away from the season median, and similar for the 30 s-rms values. Also cuts if the current PWV > 5 mm.
Fourier $\chi^2$	16.2	39.9	18.5	31.7	Validates the $\chi^2$ of frequency ranges by computing the number of standard deviations $\delta(\chi^2) \approx \left  \frac{\chi^2 - n}{\sqrt{2n}} \right $ , where $n = 2N_{\text{freq}}$ is the number of degrees of freedom. Cuts if $\delta(\chi^2) > 4$ , $\delta(\chi_{f_{\text{scan}} \pm 10\text{mHz}}^2) > 10$ , $\delta(\chi_{0\text{Hz} \dots 0.2\text{Hz}}^2) > 10$ , $\delta(\chi_{10\text{Hz} \dots}^2) > 20$ or $\delta(\chi_{\text{filtered}}^2) > 4$ .
TOD $\chi^2$	6.0	24.3	6.4	25.0	Cuts if a $\chi^2$ based on every 10th sample of the filtered TOD is more than 4 sigma away from the expectation value, or if the sample with the highest absolute value in the filtered TOD is higher than 7 times its rms.
Azimuth $\chi^2$	2.7	9.1	3.5	8.5	Cuts if azimuth structure in the TOD is detected at more than 4 sigma.
Spike	5.4	21.1	7.8	24.3	Cuts if the 1 Hz and 1.2 Hz spikes associated with the cryostat, and their first 6 harmonics are detected at more than 20 sigma, or if any 0.1 Hz interval between 0.2 Hz and 4.5 Hz (Q) or 9.5 Hz (W) has a $\chi^2$ excess of more than 7.5 sigma.
Sun	8.6	9.9	2.2	0.0	Rejects CES-detectors if any of their samples expressed in Sun-relative coordinates hit an area with a season-average detection of the sun of more than 5 sigma. See the text for details.
Conservative	18.3	17.0	16.5	15.1	Rejects all detectors for a given CES if more than 40% of the detectors fail the weather, TOD, Fourier, azimuth, spike or sun cuts.
Total	41.7	65.7	41.7	91.6	Less than the sum because of a large degree of overlap.

Table 2.2: Summary of the QUIET cut criteria and the amount of data cut for each criterion.

For the indirect cuts, however, the goal is not detection of deviations from the noise model, as the filtered regions are expected to be contaminated even for usable data. In this case there are two approaches to choosing cut thresholds: Firstly, one can try to model the relevant contaminations' frequency dependence, and use this to establish how strong a detection in a filtered region is needed for it to significantly affect the unfiltered region. And secondly, one can use trial and error until one finds cut thresholds which cut as little as possible of the data while still passing the null tests.<sup>12</sup> Most of the indirect cuts in table 2.2 are based on the latter approach due to the difficulty of modelling the effect of weather, which is the primary contaminant.

### The Sun cut

As illustrated in figure 2.13, the telescope sidelobes are quite complex, exhibiting both narrow, intense regions like the triple reflection sidelobe and broader but weaker regions like the spillover sidelobe. Given some region  $r$  in sun-relative coordinates<sup>13</sup>, we can determine the strength of detection of the sidelobe in this region by building a map  $m$  and its covariance  $C$  in this region, and calculating  $\chi^2 = m^T C^{-1} m$ .

The Sun cut is based on finding significantly contaminated regions like these, and rejecting CES-detectors which hit them. However, the significance of the detection depends on the choice of region: A small, sharp sidelobe may not be detected if the region is too big, while a weak but extended sidelobe will not be detected if the region is too small.

To get around this, we choose a hierarchical approach: For each HEALPix  $N_{\text{side}}$  [26] from 256 to  $16^{14}$ , we calculate the  $\chi^2$  for each pixel, and use this to produce a map of the detection significance for each resolution. We then consider a point to be contaminated by a Sun sidelobe if any of these maps has a detection of 5 sigma or more for the corresponding pixel. We do this independently for each module, as we observe the shape of the sidelobes to depend on the position in the focal plane.

Contaminated regions are typically detected with a significant margin, with chisquares up to hundreds of sigma away from the expectation value. We can therefore afford to make some approximations in the calculation of the  $\chi^2$ . By ignoring time-correlations in the TOD, and hence pixel-correlations

<sup>12</sup>Normally, being able to tweak cut parameters etc. manually would make the data analysis vulnerable to experimenter bias: If one is free to manually tweak the cut thresholds and other parameters based on how the maps and power spectra look, then it is very easy to end up tweaking them away from the true value towards the expected value. Null-tests eliminate this bias by performing the test on a combination of the data that is independent of the quantity to be measured.

<sup>13</sup>See 2.2 for the definition of this coordinate system

<sup>14</sup>These numbers are somewhat arbitrary, but correspond approximately to the observed minimal and maximal size of the features in the sidelobes.

in the map, we end up with a block-diagonal covariance matrix  $C_{pp'\alpha\alpha'} \propto \delta_{pp'}$  where  $p$  and  $p'$  are pixel indices and  $\alpha$  and  $\alpha'$  are Stokes indices. In this case, the  $\chi^2$  for a low-resolution pixel is simply the sum of the  $\chi^2$  of its constituent high-resolution pixels, meaning that only a single map-making step is necessary per detector.

## 2.2 Pointing

The QUIET telescope can move with three degrees of freedom: Rotations in azimuth ( $\alpha$ ) and elevation ( $\epsilon$ ), which are used to point the boresight at a given point on the sky, and “deck rotations” ( $\delta$ ), which are rotations of the focal plane around the boresight axis. These axes are connected to an encoder, which provides readouts of their orientation as a part of the time-ordered-data.

The translation from these coordinates to the galactic coordinates used in the map-making step can be expressed as a rotation

$$\vec{v}_g = R_{gh}\vec{v}_h, \quad (2.19)$$

where  $\vec{v}_h = \vec{e}(-\alpha, \frac{\pi}{2} - \epsilon)$  and  $\vec{v}_g = \vec{e}(l, \frac{\pi}{2} - b)$  are the unit pointing vectors in horizontal and galactic coordinates respectively, in terms of

$$\vec{e}(\phi, \theta) \equiv \begin{bmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{bmatrix}. \quad (2.20)$$

$R_{gh}$  is the rotation matrix from horizontal to galactic coordinates, which can be decomposed into a rotation from horizontal to apparent equatorial coordinates  $R_{ea}$  followed by a rotation from apparent to astrometric equatorial coordinates, and finally a rotation from equatorial to galactic coordinates  $R_{ge}$  such that  $R_{gh} = R_{ge}R_{ea}R_{ah}$ , with

$$R_{ah} = E(\text{LST}, \frac{\pi}{2} - b_q, 0) \quad (2.21)$$

$$R_{ge} = E(\alpha_g, \beta_g, \gamma_g). \quad (2.22)$$

Here,  $E(\phi, \theta, \psi)$  is the rotation matrix corresponding to the  $zyz$  Euler angles  $\phi, \theta, \psi$ :

$$E(\phi, \theta, \psi) = R_z(\phi)R_y(\theta)R_z(\psi), \quad (2.23)$$

and  $\text{LST} = \text{GMST} + l_q$  is the local sidereal time in radians,  $l_q = -67.76166667^\circ$  and  $b_q = -23.02822222^\circ$  are the telescope’s longitude and latitude respectively, and  $\alpha_g = -57.068351386^\circ$ ,  $\beta_g = -62.871663896^\circ$  and  $\gamma_g = -192.859498564^\circ$  are the J2000 Euler angles for the equatorial to galactic

rotation. The apparent to astrometric rotation  $R_{ea}$  is a nonlinear effect taking into account the precession, nutation and aberration etc. due to the Earth's orbit<sup>15</sup>.

Using equation (2.19), we can calculate the galactic pointing of each individual detector provided that we have a means of finding its horizontal pointing. Each detector has a constant pointing  $\vec{v}_b = \vec{e}(\phi_{0d}, \theta_{0d})$  in boresight-relative coordinates, so

$$\vec{v}_h = R_{hb} \vec{v}_b = E(-\alpha, \frac{\pi}{2} - \epsilon, \delta) \vec{v}_b, \quad (2.24)$$

making the full rotation  $\vec{v}_g = R_{ge} R_{ea} R_{ah} R_{hb} \vec{v}_b$ . This expression is sufficient for temperature detectors, but a crucial component is missing with regards to measuring the linear polarization components  $Q$  and  $U$ .

Each of QUIET's polarization sensitive detector measures a fixed linear combination  $\cos(2\psi_{0d})Q + \sin(2\psi_{0d})U$  of the Stokes  $Q$  and  $U$  parameters in boresight-relative coordinates, but the decomposition of linear polarization into  $Q$  and  $U$  is coordinate system dependent, and when projected on to the sky in galactic coordinates, the detector measures a time-dependent linear combination given by

$$\sigma_{dt} = \begin{bmatrix} \cos(2\psi_{dt}) \\ \sin(2\psi_{dt}) \end{bmatrix}. \quad (2.25)$$

Equation (2.24) does not prescribe how to calculate provide  $\psi$  from  $\psi_0$ .

An elegant way of solving this problem is to do away with the vectors altogether, and work exclusively with rotations. Starting from a detector-relative coordinate system where  $\vec{e}_z = [0, 0, 1]$  is the pointing and  $\vec{e}_x$  is the direction along which the detector is sensitive to linear polarization, this system can be rotated into the boresight system by

$$R_{bd} = E(\phi_{0d}, \theta_{0d}, \psi_{0d}). \quad (2.26)$$

The full rotation from detector-relative to galactic coordinates is then

$$R = R_{ge} R_{ea} R_{ah} R_{hb} R_{bd}. \quad (2.27)$$

This rotation can be factorized into  $R = E(l, \frac{\pi}{2} - b, \psi)$  via  $l = \arctan_2(R_{23}, R_{13})$ ,  $b = \frac{\pi}{2} - \arccos(R_{33})$ ,  $\psi = \arctan_2(R_{32}, -R_{31})$ , which provides both the pointing and the detector orientation we need.

### 2.2.1 Verifying the pointing

The last two terms in equation (2.27), are hardware-dependent, and could therefore be miscalibrated. In part to check for the presence of pointing

---

<sup>15</sup>This is provided by the NOVAS library.

Object	T	P	Remarks
Moon	CD	CD	Hard to model
Venus	CD	N	Few scans
Jupiter	CD	N	
Tau A	CD	CD	Raster scans
RCW 38	CD	N	
G-1	CD	C	Extended
G-2	CD	S	Extended

Table 2.3: The objects used for calibrating the telescope pointing. Columns two and three indicate the strength of the signal in temperature and polarization respectively. The objects can be visible in single CES-detectors (CD), in a single CES if all detectors are coadded (C), only if coadding over an observing season (S), or not at all (N). The Moon is by far the strongest source, but due to its extendedness, time-variability due to phase changes, and its excessive brightness, it is not used for pointing.

errors, our observing strategy includes a set of calibration targets, with the most important ones being Jupiter, RCW 38 and the galactic center patch G-1 (see table 2.3 for a full list). These objects are strong enough to be visible to at least some of the detectors in single constant-elevation scans, and also have known shape, position and amplitude. This allows us to test not only the accuracy of the pointing, but every parameter that enters into the response matrix  $P$ . From section 4.2 we recall that, given a beam-smoothed sky map  $\tilde{m}$ , the signal part of the time-ordered data is given by

$$s_{dt} = \sum_{t'h\alpha} \tau_{dtt'} \psi_{dht'\alpha} \tilde{m}_{p_{dht'\alpha}}. \quad (4.6)$$

By inserting in place of  $\tilde{m}$  a model of the source based on position, shape and amplitude parameters, we can find the optimal parameter values by minimizing the residual of the true TOD  $d_{dt}$  and the predicted signal  $s_{dt}$ :

$$\chi^2 = (d - s)^T N^{-1} (d - s). \quad (2.28)$$

Figure 2.2.1 illustrates the result for this minimization for a scan of the point-like source Jupiter and the extended source G-1, both using data from the temperature detectors. The signal-to-noise is high enough to constrain the model parameters with sufficient accuracy.

For each CES, the result of the fit is a set of apparent positions, shapes and amplitudes per detector. If the pointing model, gain and beam are all correct, these will all scatter tightly around the true values. Otherwise, the apparent position, size and amplitude provide information about errors in pointing, beam and gain respectively.

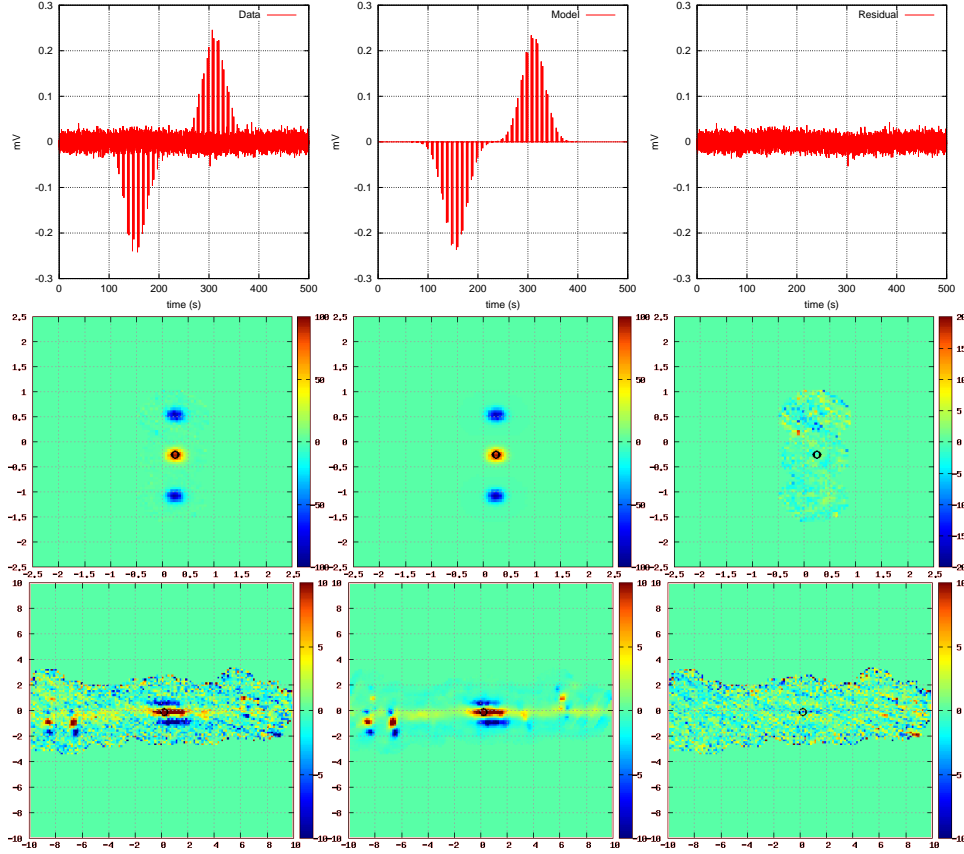


Figure 2.15: Illustrations of the fitting of a source model to observations. Top: time-ordered data for for a scan of Jupiter (left), the best fit model (middle), and the residual (right). Middle: The same data naively projected to pixel domain using a fast, simplified map-making procedure which ignores the differential nature of the map-making equations. These maps are meant only for visually diagnosing the quality of the fit, and use a very fast, simplified map-making procedure which ignores noise correlations and the differential nature of the temperature detectors, resulting in spurious blue shadows (see section 4.4.1). The fit itself happens in time-domain. The full map-making equation does not produce these artifacts. Bottom: The data vs. model for an extended source, the galactic center patch (G-1). The horizontal and vertical axes of the maps measure the offset from the expected position in degrees. The Jupiter model used here is a Gaussian, while the galactic center models is an  $N_{\text{side}} = 1024$  W-band map from the WMAP satellite.

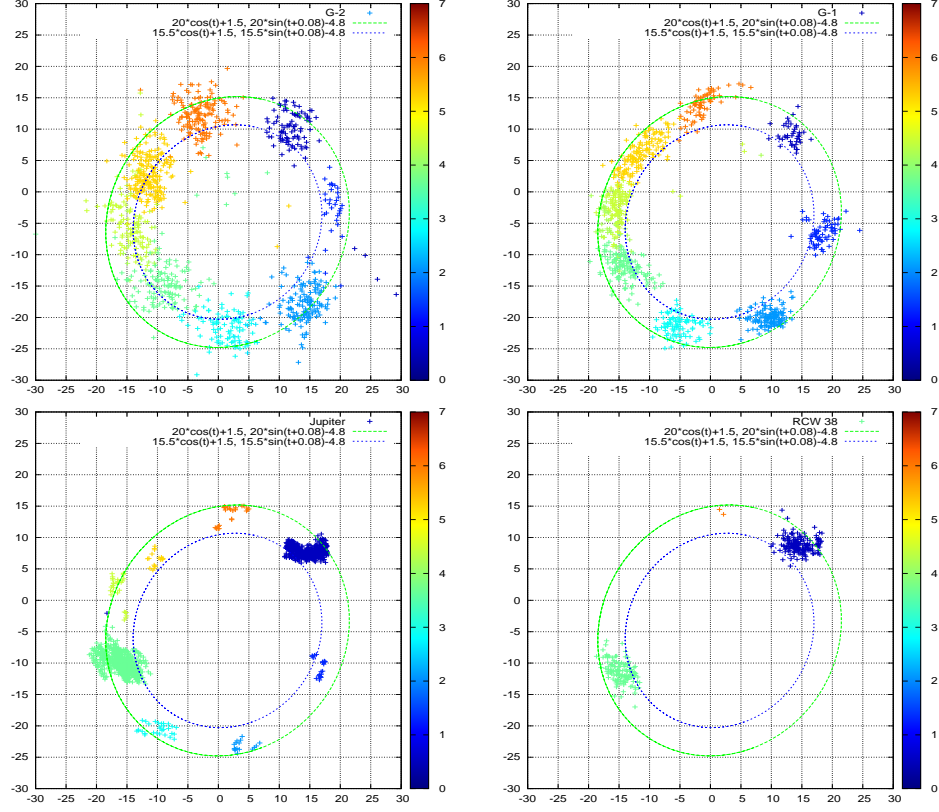


Figure 2.16: The pointing residuals based on single-CES temperature measurements of four objects: the two galactic patches, Jupiter and RCW 38. The horizontal and vertical axes show  $(\alpha_{\text{obs}} - \alpha_{\text{true}}) \cos(\epsilon_{\text{true}})$  and  $\epsilon_{\text{obs}} - \epsilon_{\text{true}}$  respectively, in units of arcminutes. The color scale encodes the deck orientation in radians. The scatter is mostly confined within an ellipse with a semi-major axis of  $\sim 18'$ , most of which is due to a constant collimation offset of  $17.1' \pm 0.1'$ .



Collimation	A constant offset between the true boresight and the fiducial boresight.	$R_{\text{col}} = E(\phi_c, \theta_c, -\phi_c)$
Azimuth tilt	An offset of the azimuth rotation axis from vertical.	$R_{\text{atilt}} = E(\phi_a, \theta_a, -\phi_a)$
Elevation tilt	An offset of the elevation rotation axis from east when the telescope is pointing north.	$R_{\text{etilt}} = E(-\alpha - \frac{\pi}{2}, \theta_e, \alpha + \frac{\pi}{2})$
Flexure	The effect of gravity pulling the boresight lower due to non-rigidity of the telescope.	$R_{\text{flex}} = E(-\alpha, k \cos \epsilon, \alpha)$
Encoder offsets	Errors in the calibration of the encoder may lead to constant offsets in the readouts for each of the azimuth, elevation and deck axes.	$\alpha \rightarrow \alpha - \Delta\alpha, \epsilon \rightarrow \epsilon - \Delta\epsilon,$ $\delta \rightarrow \delta - \Delta\delta$
Acceleration	Constant elevation scans involve regular acceleration in azimuth, which may slightly change the effective direction of the sag.	

Table 2.4: The possible sources of pointing errors that were anticipated before observations started, and the rotations they correspond to.

Figure 2.16 shows the difference between apparent and true position for Jupiter, RCW 38 and the galactic patches G-1 and G-2 based on observations with the temperature-specialized detectors. It is clear that the pointing model in equation (2.27) is far from sufficient. The effect of pointing errors amounts to convolving the maps with an extra beam corresponding to the pattern of residuals, which in this case is a non-Gaussian beam with a diameter of up to  $40'$ . This is several times larger than the W-band beam FWHM of  $11.7'$ , and also larger than the Q-band FWHM of  $27.3'$ , and it is therefore critical to correct for the imperfections in the telescope pointing.

### 2.2.2 Mount model

From a hardware point of view, pointing errors are not unexpected, and even before observations started, the error sources described in table 2.2.2 were anticipated. Of these, the most challenging to correct is the acceleration-induced flexure. The overall effect of this would be extra smearing in the horizontal direction, effectively producing a deck-angle-dependent ellipticity of the beam, which we have not observed. The acceleration can there-

Rotation	Param.	Value	Error
$R_{\text{bore}} = E(-\alpha - [\Delta_{\text{az}} + \theta_{\text{ef,az}} \cos(-\alpha + \phi_{\text{ef}})],$ $\frac{\pi}{2} - \epsilon - [\Delta_{\text{el}} + \theta_{\text{ef,el}} \cos(-\alpha + \phi_{\text{ef}})],$ $\delta - [\Delta_{\text{dk}} + \theta_{\text{ef,dk}} \cos(-\alpha + \phi_{\text{ef}})])$	$\Delta_{\text{az}}$	$-0.0357^\circ$	$0.0041^\circ$
	$\Delta_{\text{el}}$	$-0.0294^\circ$	$0.0010^\circ$
	$\Delta_{\text{dk}}$	$0.1083^\circ$	$0.0263^\circ$
	$\theta_{\text{ef,az}}$	$-0.0225^\circ$	$0.0027^\circ$
	$\theta_{\text{ef,el}}$	$-0.0042^\circ$	$0.0019^\circ$
	$\theta_{\text{ef,dk}}$	$-0.2615^\circ$	$0.0177^\circ$
	$\phi_{\text{ef}}$	$11.3520^\circ$	$2.8992^\circ$
$R_{\text{col}} = E(\phi_c, \theta_c(1 + \theta_{\text{ec}} \cos[\delta + \phi_{\text{ec}}]), -\phi_c)$	$\theta_c$	$0.2948^\circ$	$0.0015^\circ$
	$\phi_c$	$26.6708^\circ$	$0.2302^\circ$
	$\theta_{\text{ec}}$	$0.2126$	$0.0058$
	$\phi_{\text{ec}}$	$253.8268^\circ$	$1.5714^\circ$
$R_{\text{atilt}} = E(\phi_a, \theta_a, -\phi_a)$	$\theta_a$	$-0.0070^\circ$	$0.0013^\circ$
	$\phi_a$	$-49.6202^\circ$	$8.5147^\circ$
$R_{\text{etilt}} = E(-\alpha - \frac{\pi}{2}, \theta_e, \alpha + \frac{\pi}{2})$	$\theta_e$	$-0.0037^\circ$	$0.0024^\circ$

Table 2.5: The four rotations making up the mount model and their parameters. Except for  $R_{\text{atilt}}$ , all models are nonlinear. The elevation axis tilt and the “encoder elevation flexure” are only weakly detected. Notably missing from the model is the telescope flexure correction, which was found to be consistent with zero. All the corrections are small, so the rotation matrices commute to high accuracy.

fore be safely ignored.

However, of the remaining effects, none are capable of producing the elliptical pointing residuals in figure 2.16<sup>16</sup>, nor can they explain the peculiar elongation of the scatter in the angular direction in the plots or the “U” shape observed in the residuals of Jupiter, the highest signal-to-noise source used in the calibration<sup>17</sup>. We therefore adopted a phenomenological mount model which allows for some unanticipated extra degrees of freedom in order to match the observed pointing scatter. The resulting model has 14 parameters, and is summarized in table 2.5. Incorporating these rotations into the full rotation from detector to sky results in

$$R = R_{ge} R_{ea} R_{ah} R_{etilt} R_{atilt} R_{col} R_{bore} R_{bd}. \quad (2.29)$$

A histogram of the pointing residuals from this model is shown in figure 2.17. The best-fit Gaussian to these residuals has major and minor

<sup>16</sup>The collimation error produces circular pointing residuals which can account for the majority of the effect seen in the figure, but it cannot account for the elliptical shape.

<sup>17</sup>Observations of the Moon are also available, but due to hard-to-model temperature variations across its disk we chose not to include these in the pointing calibration.




Figure 2.17: A histogram of the pointing residuals from the mount model described in table 2.5. The horizontal and vertical axes are  $\alpha_{\text{corr}} - \alpha_{\text{true}} \cos(\epsilon_{\text{true}})$  and  $\epsilon_{\text{corr}} - \epsilon_{\text{true}}$  respectively, where  $\alpha_{\text{corr}}, \epsilon_{\text{corr}}$  are the mount model-corrected observations in horizontal coordinates, and  $\alpha_{\text{true}}, \epsilon_{\text{true}}$  are the known coordinates. A circle corresponding to a standard deviation of  $1.91'$  is overplotted (red) and compared with the W-band beam rms (black). The scatter is moderately non-Gaussian, but is not large enough to seriously affect the beam.

semi-axis of  $2.15'$  and  $1.66'$  respectively. We treat this difference as a statistical error, leading to an estimate for the scatter of  $1.91' \pm 0.25'$ , which corresponds to a FWHM of  $4.49' \pm 0.59'$ . These add quadratically to the effective beam size, increasing beams from  $11.7'$  to  $12.5'$  (W) and from  $27.3$  to  $27.7'$  (Q), both of which lead to an acceptable loss of sensitivity of 7% at  $l = 500$  and 27% at  $l = 1000$ .

### 2.2.3 Full-season validation of the mount model

As we have seen, the mount model performs adequately on the observations of the calibration targets. However, this need not be representative for its performance on the CMB data if the calibration observations are systematically different from the CMB observations. This can come about through:

1. Different kinds of detectors: All the calibration observations that en-

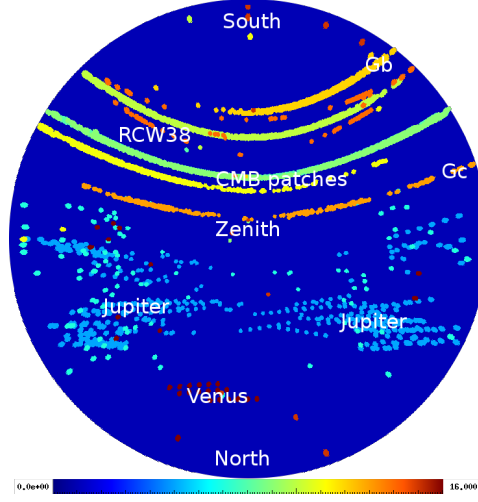


Figure 2.18: The distribution of the calibration observations in horizontal coordinates compared to the CMB observations. Each CES is here indicated with a colored disk, with one color per object. While Jupiter and Venus are measured at systematically different azimuths than the CMB patches, the RCW 38 and the patches G-1 and G-2 (here denoted Gc and Gb) remedy this weakness. Together, the calibration sources cover most of the relevant azimuth/elevation space.

tered into the pointing analysis were based on temperature detectors.

2. Different focal plane coverage: The temperature detectors are concentrated at one edge of the focal plane.
3. Different sampling of horizontal coordinates: If the optimal mount model varies as a function of azimuth, elevation or deck, then the best mount model for the calibration observations may not be the best mount model for the part of the sky covered by the CMB observations (but see figure 2.18 for why this is unlikely to be a problem).

We performed three tests to investigate these possibilities:

1. Comparison of an  $N_{\text{side}} = 2048$  full-season co-added map of Jupiter<sup>18</sup> based on the secondary total power data stream<sup>19</sup> from the polarization detectors, with the expected size of a point source convolved

<sup>18</sup>This requires a coordinate system where Jupiter does not move, for example coordinates centered on Jupiter. Such object-centered coordinates can be constructed by applying an extra rotation  $R_{og} = E(0, b_o, -l_o)$  after the transformation to galactic coordinates. Here  $b_o, l_o$  are the galactic latitude and longitude of the object.

<sup>19</sup>This data stream has much higher  $1/f$  noise than the main data stream, but is sensitive enough to measure Jupiter's temperature signal.

with the beam and pointing scatter. The result can be seen in figure 2.19, where we find a FWHM of  $13.1'$ . The fiducial beam for polarization horns is  $11.7'$ , which is consistent with a scatter FWHM of  $5.9'$ .

2. A similar test for the  $N_{\text{side}} = 2048$ , full-season co-added map of the point source PMN J0538-4405 in patch CMB-2 (fig. 2.20), which finds a best-fit FWHM of  $14.2' \pm 1.0'$ . Comparing this to a temperature detector beam of  $13.1' \pm 0.2'$ , we find a scatter FWHM of  $5.5' \pm_{5.5'}^{2.5'}$ .
3. We compared the full W-band season CMB temperature power spectrum and fit it to the model  $C_l = AC_l^{\Lambda\text{CDM}} e^{-l(l+1)\sigma_{\text{FWHM}}^2(8\log 2)^{-1}}$  to determine the effective beam, and found a FWHM of  $14.1' \pm 0.1'$ , resulting in a scatter FWHM of  $5.2' \pm 0.7'$ .

From these we can conclude that both the temperature and polarization detectors see effective beams consistent with the pointing residuals both when observing CMB patches and Jupiter. The scatter is small enough not to be problematic, but it must be taken into account during power spectrum estimation. We have adopted a value of  $5.1'$  for the rest of the analysis.

## 2.3 Gain and beam

The fitting procedure described in section 2.2.1 produces estimates of the gain and beam as a by-product of the pointing fit, as illustrated in figure 2.3. But there are several caveats one needs to bear in mind before these can be applied in CMB analysis.

Firstly, these numbers were produced by fitting a Gaussian profile to the signal, and while the true beam is approximately Gaussian, it is not exactly so, having significantly heavier tails. This increases the beam's area on the sky compared to the value it would have for a Gaussian beam.

Secondly, the effective wavelength of the observations depends on the spectral index of the target, which varies from 0 for the CMB to  $\sim -3$  for synchrotron-dominated areas like the galaxy. With an average bandwidth of 17% (Q), 11% (W), this can lead to differences in effective frequency of up to 3%, which corresponds to changes of up to 6% in the area of the beam.

Thirdly, when using point sources for calibration, the observed flux depends on the beam area  $\Omega_b$  as  $F \propto \Omega_b^{-1}$ . The gain is defined as the ratio of the detector's response  $s$  to the flux,  $g = \frac{s}{F} \propto \Omega_b$ . Thus, changes in the area of the beam propagate directly into changes in the gain when the latter is based on point source observations. This is not a concern when observing slowly varying, extended sources.

Putting these together, one can easily misestimate the gain by 5% by using the wrong beam area, and this would then translate into 10% errors in

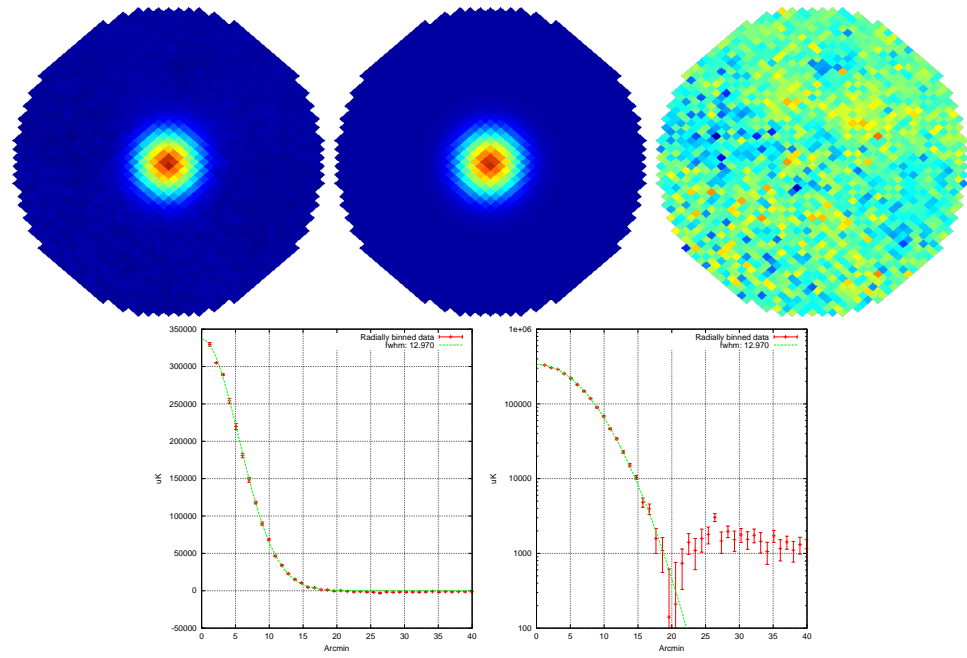


Figure 2.19: Top left: Co-added  $N_{\text{side}} = 2048$  temperature map of Jupiter from the QUIET W-band season in object-centered coordinates based on the secondary total power data stream from the polarization detectors, in arbitrary units. Top middle: Best fit Gaussian with a major FWHM of  $13.2'$  and a minor FWHM of  $13.0'$ , and a bias in the position of  $\Delta l = 0.42'$ ,  $\Delta b = 0.20'$ . Top right: Map of the residuals. Bottom row: Radial average of the map, with a best-fit Gaussian with a FWHM of  $13.0'$ , with linear (left) and logarithmic (right) vertical axis.




Figure 2.20: Left: A full-season co-added  $N_{\text{side}} = 2048$  temperature map of the source PWM J0538-4405 in patch CMB-2, based on the primary data stream from the temperature detectors. Right: WMAP's 7 year W-band  $N_{\text{side}} = 1024$  temperature map of the same source. The object can be seen at consistent position in the two maps. Due to low signal to noise, the effective beam is somewhat poorly determined, at  $14.2' \pm 1.0'$ . Figure courtesy of H. K. K. Eriksen.

the power spectrum. A detailed study of the beam shape was performed in [27], which determined that the relevant beam area for temperature observations of Jupiter is  $15.58 \pm 0.63 \mu\text{Sr}$ , compared to  $16.7 \mu\text{Sr}$  for the Gaussian approximation. The gains from the Gaussian fit must therefore be scaled by 0.93 before being used.

With this done, the Jupiter fit provides  $\sim 500$  gain measurements for each temperature detector, which is sufficient to produce a reliable gain model. The gains show evidence of module-dependent trends at the  $\sim 5\%$  level, or slightly higher than the per-sample scatter of the measurements. Due to the low amplitude of the trends, we chose to adopt a constant gain model for the temperature gains, with the trends entering into the systematic error estimate.

### 2.3.1 Polarization gains

A thorough analysis of the polarization gains for the Q-band season was performed in [28], where it was found that gain measurements based on the Moon, Tau A, elevation dips and a wire-grid<sup>20</sup> were correlated with the measured temperature of the electronics  $T$ , resulting in a gain model of the

<sup>20</sup>An artificial polarized source consisting of a grid of thin wires mounted on a rotating frame in front of the receiver.

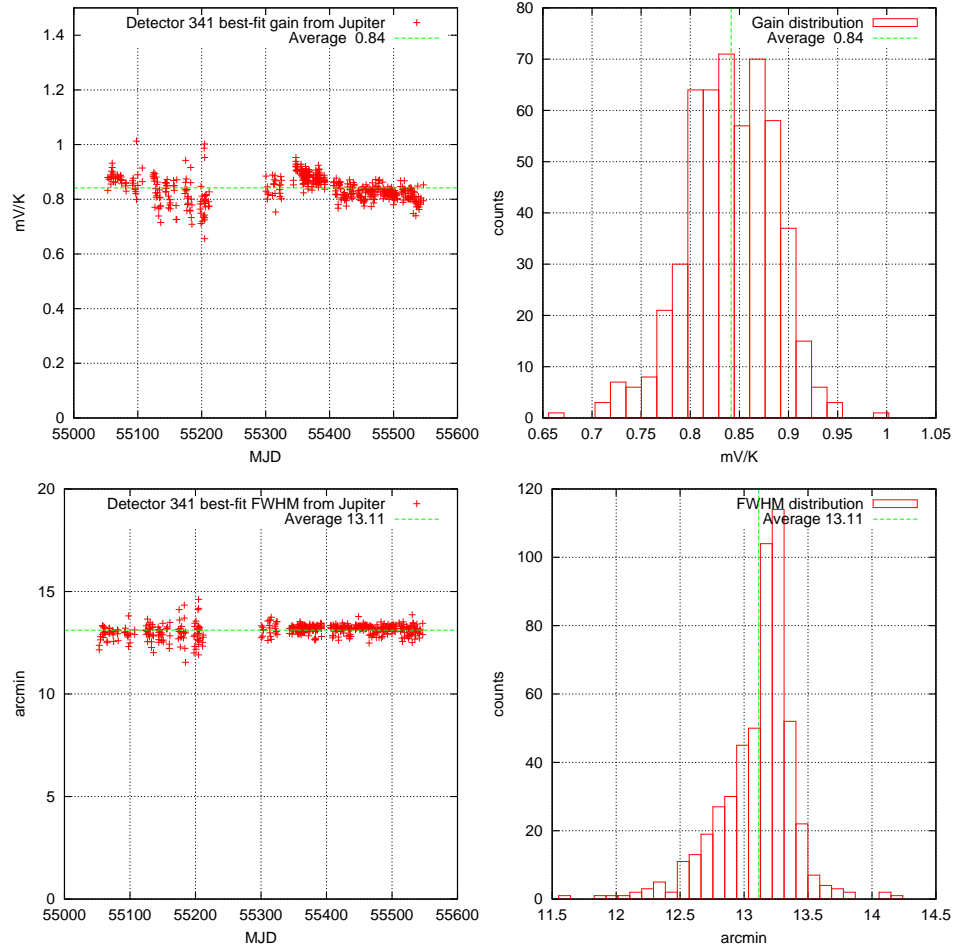


Figure 2.21: An example of gain and beam size estimates produced by fitting a Gaussian profile to Jupiter observations (see section 2.2.1) for a single detector. The top row shows the gain estimates as a function of time (left) and a histogram of these (right). There is some evidence of time-dependence of the gain. The bottom row shows the equivalent graphs for the beam.



form

$$g_{dt} = \alpha_d(1 + \beta_d[T(t) - T_0]), \quad (2.30)$$

with  $T$  mostly varying around  $T_0 = 25.1\text{K}$ .

For the W-band season, the situation is quite different. Firstly, the source of the gain-temperature correlation was eliminated, and secondly, the problem of gain nonlinearity is much more prominent here, making both Moon observations and elevation dips unreliable as calibrators. The W-band gain model must therefore be built based on a much reduced data set, consisting of two Tau A measurements per day, half of which are for the central horn, and of one day of wire-grid measurements at the end of the season, providing relative gains for all the detectors at the same time.

This data set is too small to build a time-dependent gain model for each detector; we therefore assume that the time-dependent part of the gain is shared between detectors, and that these only differ by a constant factor. This picture is supported by figure 2.3.1, which shows that the gain is clearly time-dependent, and that this time-dependence is common for all four detectors for which a significant amount of data is available. The time dependence has the form of a wave with a period of approximately one year, and a slight downwards trend. This motivates the model

$$g_d(t) = \alpha_d [1 + \beta(t - t_0) + \gamma \sin(\omega(t - t_0) + \delta)], \quad (2.31)$$

where only the overall factor  $\alpha_d$  is diode dependent, and the shape parameters  $\beta, \gamma, \omega$  and  $\delta$  are shared between detectors. Two fits of this model to the data are shown in the figure: a maximum likelihood fit based on the central horn only, and a fit based on all the detectors. These have  $\chi^2 = 679$  for 492 data points and 8 parameters, and  $\chi^2 = 2381$  for 1920 data points and 313 parameters respectively. In both cases, there is evidence for extra scatter beyond that expected from the error bars. This is not surprising given the simple and ad-hoc nature of the gain model and the small amount of data available from which to construct it. However, the fit for the full data set is comparatively worse compared to the central horn only case, and as seen in the figure, the other detectors seem to prefer a different phase and amplitude of the modulation than the one exhibited by the central horn. This may indicate that the assumption of a common time dependence for all detectors is inaccurate. Still, the model is a useful approximation to the gain, and performs better than the hypothesis of constant gains or linearly changing gains. We therefore adopt this as the polarization gain model, and interpret the extra scatter in the residuals as a systematic error.

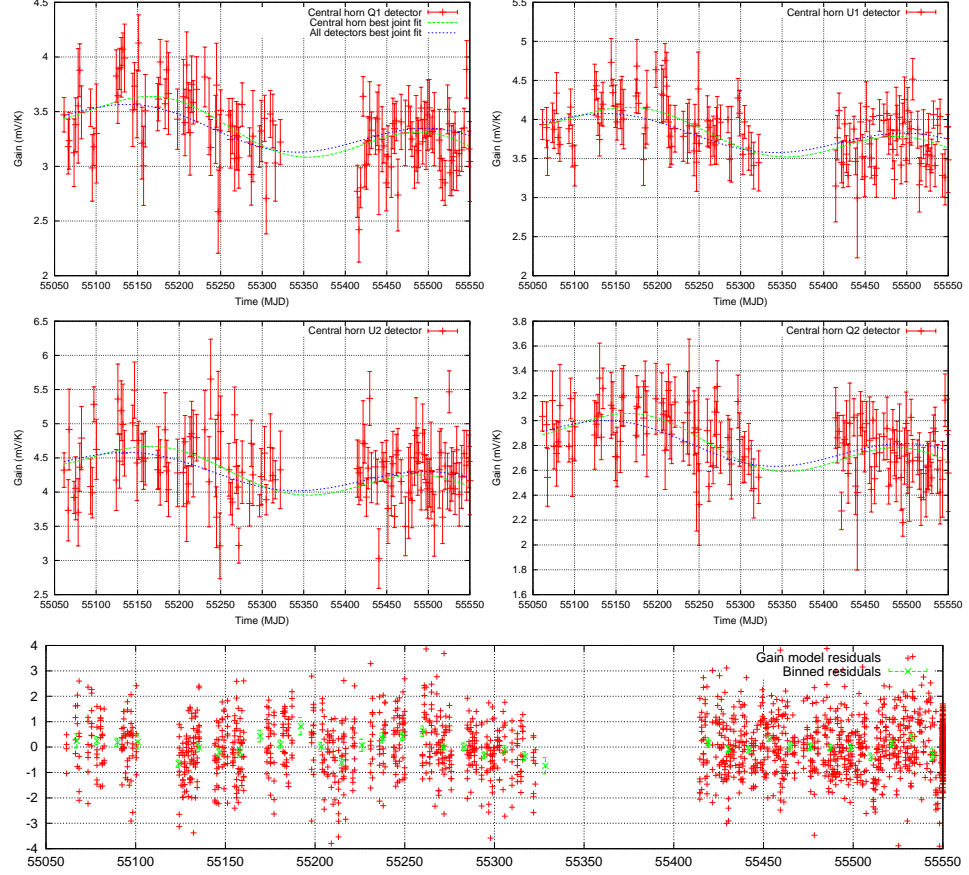


Figure 2.22: Top: Polarization gain measurements from Tau A and the wire-grid for the four detectors in the central module. These are shown as red points with error bars. A roughly seasonal modulation is clearly visible, as well as a slight downwards trend, and the shape of these is the same for the four detectors. The green curve shows the best fit harmonic+linear model based on these four detectors while the blue curve shows the best fit including data from the other 305 usable polarization detectors, which have only on average 4.6 data points each. The blue curve is a somewhat worse fit to the central module data. This may indicate that not all detectors prefer the same time dependence as the central module. Bottom: Residuals of the full model as a function of time.

## Chapter 3

### Null tests

As the calibration chapter shows, the analysis pipeline has a large number of tunable parameters, including the pointing model, gain model, filter parameters and cuts. Both which parameters to use and their value will end up affecting the final result, which thus in a sense inherits the tunability of the parameters: The amplitude of the power spectrum will change depending on the gain model, and a poor pointing model will lead to less power on small scales, to name two examples.

This introduces the danger of *observer bias* – the tendency to observe what one expects to observe because unexpected results will tend to get more scrutiny than expected results. A famous example of this is related by Richard Feynman [29]:

Millikan measured the charge on an electron by an experiment with falling oil drops, and got an answer which we now know not to be quite right. It's a little bit off, because he had the incorrect value for the viscosity of air. It's interesting to look at the history of measurements of the charge of the electron, after Millikan. If you plot them as a function of time, you find that one is a little bigger than Millikan's, and the next one's a little bit bigger than that, and the next one's a little bit bigger than that, until finally they settle down to a number which is higher.

Why didn't they discover that the new number was higher right away? It's a thing that scientists are ashamed of–this history–because it's apparent that people did things like this: When they got a number that was too high above Millikan's, they thought something must be wrong–and they would look for and find a reason why something might be wrong. When they got a number closer to Millikan's value they didn't look so hard. And so they eliminated the numbers that were too far off.

### 3.1 Signal-less validation

The standard way of avoiding observer bias is through *blind analysis*, which is the practice of performing the experiment and data analysis without peeking at the results underway. In CMB data analysis, the analysis pipeline itself can be validated through end-to-end simulations, while most calibration can be tested through calibration-dedicated subsets of the data, such as Jupiter observations in our case. But as we saw in sections 2.1.2 and 2.1.3, some calibration parameters deal with removing variable contaminants such as the atmosphere from the data, and these must be tested on the main data set itself. It is, however, possible to perform this test without exposing oneself to the results, and thus the results to bias. One simply sub-divides the data set into two halves such that each should have the same signal, and considers their difference:

$$\Delta d = d_2 - d_1 = (s + n_2 + c_2) - (s + n_1 + c_1) = \Delta n + \Delta c. \quad (3.1)$$

Here  $d$  is the data,  $s$  is the signal, and  $n$  and  $c$  represent the noise and contaminants respectively. Since  $\Delta d$  is independent of the signal, it can be safely used without introducing bias, and one can ensure the quality of one's cuts and filters by demanding that  $\Delta d$  have no  $\Delta c$  contribution, i.e. that it is consistent with noise. This technique is known as *null-testing*.

### 3.2 Map null-tests

The QUIET maximum likelihood pipeline uses two kinds of null-tests: map-based and pseudo- $C_l$  power spectrum based estimators. The former of these is the most straight-forward: Given two sub-sets of the data, compute the map  $m$  and covariance matrix  $M$  of each of these sets. Then the difference map  $m_2 - m_1$  should consist of Gaussian noise with covariance  $M_2 + M_1$ . We can check this by computing the  $\chi^2$ :

$$\chi^2 = (m_2 - m_1)^T (M_2 + M_1)^{-1} (m_2 - m_1), \quad (3.2)$$

which should be chisquare-distributed with  $N_{\text{dof}} = N_{\text{pix}} N_{\text{stokes}}$  degrees of freedom, with a mean of  $N_{\text{dof}}$  and a standard deviation of  $\sqrt{2N_{\text{dof}}}$ . A significant excess in this number indicates that there remains unfiltered junk in the data<sup>1</sup>, while a significant shortfall only can be produced by an over-estimate of the noise.

Null-tests based on maximum likelihood maps are expensive both time-wise and memory-wise, because they impart the full expense of computing unbiased maps and their covariance matrix. When considering that one will want to perform more than just a single split of the data in order to look

---

<sup>1</sup>Or that the noise has been under-estimated.

for different kinds of pollution, these null-tests can easily come to totally dominate the computational expense of the data analysis. Or put another way, it is only possible to perform a small number of null-tests this way.

### 3.3 Pseudo- $C_l$ null-tests

An alternative, less expensive way of detecting deviations from null is by replacing covariance matrices with simulations. Given, as before, maps  $m_1$  and  $m_2$ , and a set of simulations based on the same split of the data,  $\{\xi_{1i}\}$  and  $\{\xi_{2i}\}$ , we can compute null maps  $\Delta m \equiv m_2 - m_1$  and  $\{\Delta \xi_i \equiv \xi_{2i} - \xi_{1i}\}$ . Given some function of the null maps  $f(\Delta m)$ , we can compute the deviation of the actual null map from the null simulations by computing the probability to exceed the observed value of  $f$ :

$$\text{PTE}(f) \approx \frac{\text{count}_i(f(\Delta \xi_i) > f(\Delta m))}{N_{\text{sim}}}. \quad (3.3)$$

The function  $f$  allows us to choose which parts of the maps we wish to test. We are ultimately interested in computing the binned angular power spectra<sup>2</sup>  $C_b^{\text{EE}}$ ,  $C_b^{\text{BB}}$  and  $C_b^{\text{EB}}$ . It therefore makes sense to make  $f$  simply be the pseudo- $C_l$  approximation of power in these bins,

$$C_b^{AB}(m) = \frac{\sum_{(l \in b)_m} l(l+1) \tilde{m}_{lm}^{A*} \tilde{m}_{lm}^B}{\sum_{(l \in b)_m} l(l+1)}. \quad (3.4)$$

This is more conveniently expressed as a  $\chi$  value, i.e. the normalized deviation from the mean of the simulations:

$$\chi_b^{AB}(m) = \frac{C_b^{AB}(m) - \langle C_b^{AB} \rangle}{\sqrt{\text{var}(C_b^{AB})}}. \quad (3.5)$$

Here uppercase Latin indices run over the polarization component, that is  $E$  or  $B$ , and  $\tilde{m}_{lm}^A$  is the pseudo-harmonic coefficients of the map  $m$ . The choice  $f = \chi$  not only allows us to individually test each of the angular scales we are interested in, it also makes it easier to track down the cause of a failed null-test by its  $l$ -dependence.

The two major causes of null-test failures are contaminations, which result in extra power (high  $\chi$ ) and thus too low PTEs, and noise misestimations which result in either too low or high PTEs. As a catch-all for both these cases, it is useful to consider the  $\chi_b^{2AB}$ , for which both types of deviations end up as high values, and thus low PTEs.

---

<sup>2</sup>The binning is necessary due to the low sky coverage, which strongly correlates nearby multipoles. For patches of  $\sim 1\%$  of the sky, a bin size of 50 is sufficient to decorrelate the bins (see fig. 5.3).

An advantage of this variant of null-tests is that by processing the data and simulations the same way, any bias in the map-making procedure is automatically taken into account. Hence, we do not need to use the full, unbiased maximum likelihood map-making equation; much faster, but biased, binned maps<sup>3</sup> can be used instead. Together with the lack of need for the covariance matrix, this results in the pseudo- $C_l$  null-tests being 1-2 orders of magnitude faster than the maximum likelihood null-tests, greatly increasing the practical number of null-tests.

This does not come completely for free, though. Due to using different map-making and power spectrum estimation than the full analysis would, maximum likelihood and pseudo- $C_l$  null-tests could in theory be sensitive to different things, even when using the same data splits and power spectrum binning. To test for this possibility, we ran a small number of null-tests through both the ML and PCL machinery. This confirmed that the two testing methodologies were consistent, allowing us to choose PCL null-tests as the primary tests for the Q- and W-band analysis.

### 3.4 Null test suite

Table 3.4 shows a summary of the W-band null test suite together with the final probability to exceed (PTE) for each combination of test and CMB patch<sup>4</sup>. The suite consists of 23 individual data splits, each of which results in one test per multipole bin per EB-combination for each patch, for a total of  $N_{\text{test}}N_{\text{bin}}N_{\text{spec}} = 23 \cdot 20 \cdot 3 = 1380$  tests per patch, or 5520 tests in total.

Based on this, we compute 3 summary statistics for each patch:  $\text{PTE}(\max[\chi_b^{2AB}])$ , which exposes the worst individual null-failure;  $\text{PTE}(\sum \chi_b^{2AB})$ , which exposes persistent deviations from expected power independent of the direction; and  $\text{PTE}(\sum \chi_b^{AB})$ , which exposes systematic excess or lack in power. Our requirement for declaring the null-suite as passed is that these 12 summary statistics should have no significant excess of very low  $< 0.05$  or very high  $> 0.95$  values. The Maximum Likelihood pipeline W-band analysis reached this milestone April 2012, 16 months after the data taking finished, making the null-tests the most time-consuming and comprehensive part of the analysis effort. At this point, we finally allowed ourselves to “open the box” and calculate the non-null maps and power spectra, which are described in the next chapters.

---

<sup>3</sup>See section 4.4.1.

<sup>4</sup>The Q-band null test suite is very similar, and is therefore not included here.

Split	Sensitive to	1	2	3	4
Time (MJD)	Gain drifts	0.24	0.01	0.89	0.77
Focal plane radius	Ground, beam ellipticity	0.15	0.51	0.21	0.61
ADC nonlinearity	Uncorrected nonlinearity, weather	0.68	0.95	0.54	0.02
Power at $f_{\text{scan}}$	Ground, weather	0.25	0.35	0.53	0.76
Power at 10 Hz	Electronic problems	0.41	0.56	0.53	0.48
Elevation	Ground, mount	0.58	0.08	0.39	0.55
Deck angle	Ground, mount	0.54	0.12	0.27	0.77
Azimuth	Ground, mount	0.76	0.95	0.66	0.11
$T_{\text{encoder}}$	Weather, electronics	0.08	0.09	0.24	0.06
$\Delta T_{\text{encoder}}$	Weather, electronics	0.13	0.54	0.09	0.51
$T_{\text{cryostat}}$	Weather, electronics	0.17	0.86	0.31	0.26
$\Delta T_{\text{cryostat}}$	Weather, electronics	0.86	0.51	0.07	0.68
Gain	Gain model	0.55	0.81	0.01	0.34
Water vapour (PWV)	Weather	0.52	0.30	0.79	0.94
Wind	Pointing	0.33	0.02	0.57	0.94
$T_{\text{ambient}}$	Weather	0.62	0.16	0.42	0.92
$f_{\text{knee}}$	Weather, noise model	0.21	0.81	0.49	0.03
$\sigma_0$	Weather, noise model	0.25	0.41	0.06	0.52
Module mean $\nu$	Strong foregrounds	0.58	0.27	0.58	0.87
Assembly board	Electronic problems	0.55	0.38	0.60	0.25
TOD rms variability	Weather	0.93	0.18	0.47	0.17
Q vs. U detectors		0.73	0.36	0.99	0.29
$I \rightarrow QU$ leakage	Leakage misestimation	0.16	0.11	0.91	0.43
Total $\chi^2$	Deviation from expected power	0.31	0.08	0.31	0.43
Max $\chi^2$	Individual null-failures	0.63	0.02	0.72	0.80
Mean $\chi$ shift	Power excess	0.20	0.27	0.21	0.97

Table 3.1: Summary of the pseudo- $C_l$  null-tests used in the W-band analysis. The last four columns correspond to each of the 4 CMB patches in the polarization analysis, and specify the probability to exceed (PTE) the observed value in the test, based on 100 simulations. The tests are divided into two types: individual tests and summary tests, corresponding to the top and bottom sections of the table. The main categories of individual tests are splits based on pointing, features in the TOD, temperature of electronics, detector type and weather conditions. Before opening the box, we demanded that there should be no significant excess of very low ( $< 0.05$ ) or very high ( $> 0.95$ ) PTEs in the summary null-tests.





## Chapter 4

# Map making

### 4.1 Measuring the sky with a scanning telescope

Most people's idea of how a telescope takes images of the sky is that it works much like a normal camera: There is an array of detectors (pixels) behind an optical system for focusing the light, and a picture is taken by pointing the instrument at an object and exposing the detectors for a short while. This is a reasonable sketch of how a typical optical telescope works, but when going to lower frequencies, this mode of operation becomes impractical: The practical size of a detector scales in proportion to the wavelength, greatly reducing the number that can be fit in the focal plane. For CMB experiments, the typical number of detectors is from tens to a few thousands, compared to a few millions to hundreds of millions for optical telescopes. Additionally, the signal-to-noise ratio per unit of time is often much lower in radio experiments, and there may not be enough signal to construct a useful image in one continuous exposure before the target rotates below the horizon.

To get around these limitations, telescopes at low frequencies scan the sky, using their small number of detectors as a paint brush and the scanning motion as paint strokes to paint a larger image. The direct output of such a telescope is not a single value per pixel, but instead a time series of values for each detector, called the time-ordered data (TOD). Based on this TOD and information about the telescope's scanning pattern, it is possible to reconstruct an image of the sky in a process called "map making". When low S/N data are involved, it is critical that the resulting map has well-understood statistical properties, which in turn means that one needs to understand the properties of the TOD. Hence, to produce an image with a scanning telescope, one must first build up an accurate model for how the telescope transforms the signal from sky to TOD, and then invert this model to reconstruct the sky from the TOD.

## 4.2 From sky to time-ordered data

Given a sky  $m$ , the most general possible response function give us the noise-free time-ordered data  $s(d, t)$  for the detector labeled  $d$  and the time  $t$  as

$$s(d, t) = P(d, t, m(t)). \quad (4.1)$$

It is not practical to work with infinite-resolution skies or time-streams with arbitrary response functions, so in practice one has to make a few simplifying assumptions:

**Discrete TOD** The telescope provides a set of samples instead of a continuous time-stream.

**Discrete sky** We model the sky as consisting of a set of discrete pixels. This is of course not true, but it is a good approximation as long as the pixels are smaller than our angular resolution.

**Constant sky** We assume that the sky does not change for the duration of our data set. This implies that we are using sky-fixed coordinates such as galactic or equatorial coordinates.

**Linear response** The TOD is a linear function of the sky signal. Nonlinear responses are difficult to work with, so having a linear response is a design requirement for most detectors.

Together, these assumptions allow us to write equation (4.1) as a matrix multiplication:

$$s_{dt} = P_{dti\alpha} m_{i\alpha}. \quad (4.2)$$

Here  $i$  is the pixel index while  $\alpha$  is the signal component within the pixel<sup>1</sup>. Our job, then, is to expand  $P_{dti\alpha}$  into something concrete enough to be implementable. A very general such expansion is

$$s_{dt} = \sum_{t'ijk\alpha\beta\gamma} g_{dt} \tau_{dt't'} \sigma_{dt'\alpha} \delta_{i0} B_{dt'ij\alpha\beta} R_{dt'jk\beta\gamma} m_{k\gamma}. \quad (4.3)$$

Here the indices  $t$  and  $t'$  are time steps,  $d$  is the detector,  $i, j, k$  and  $0$  are pixels, with  $0$  indicating zenith, and Greek indices are Stokes components. So the sky  $m$  is rotated into detector-relative coordinates with the pointing matrix  $R$  and smoothed with the beam matrix  $B$ , after which the Stokes

---

<sup>1</sup>This could be frequency, Stokes parameter or particle type, etc. depending on what the detectors are sensitive to. In the following, I will assume that it is the Stokes parameters only, since it is common for detectors to be sensitive to linear combinations of these (as is QUIET), but not so for different frequency bands. An experiment with multiple frequencies can therefore analyse these frequencies independently.

parameters at zenith are read off and linearly combined to a single value using the detector Stokes sensitivity  $\sigma$ . This value is then convolved with the detector time response function  $\tau$ , and finally scaled with the gain  $g$  to form the signal  $s$  at each time step for each detector.

This daunting expression can be dramatically simplified by adding a few reasonable assumptions. The most important of these comes from assuming a constant, detector-independent, circular beam<sup>2</sup>. A circular  $B$  will commute with  $R$ , and if  $B$  additionally is constant, we can apply it to the map once and for all, avoiding the need to convolve the map with the beam at every time step,

$$\begin{aligned} s_{dt} &= \sum_{t'jk\alpha\beta\gamma} g_{dt} \tau_{dt't'} \sigma_{dt'\alpha} \delta_{i0} R_{dt'ij\alpha\beta} B_{jk\beta\gamma} m_{k\gamma} \\ &= \sum_{t'j\alpha\beta} \tau_{dt't'} g_{dt'} \sigma_{dt'\alpha} R_{dt'0j\alpha\beta} \tilde{m}_{j\beta}. \end{aligned} \quad (4.4)$$

Here  $\tilde{m} = Bm$  is the beam-smoothed map, and I have used the approximation that the  $g$  and  $\tau$  commute<sup>3</sup>.

Each detector points towards a single point on the sky at a given time-step, so to pixel accuracy  $R$  will only have a nonzero contribution from a single pixel  $p_{dt'}$  of the beam-smoothed sky. The spatial part of the rotation therefore becomes a delta function, and we are left with the Stokes parameter part of the rotation,  $\rho$ . Hence,  $R_{dt'0j\alpha\beta} = \delta_{jp_{dt'}} \rho_{dt'\alpha\beta}$ , and

$$\begin{aligned} s_{dt} &= \sum_{t'j\alpha\beta} \tau_{dt't'} g_{dt'} \sigma_{dt'\alpha} \delta_{jp_{dt'}} \rho_{dt'\alpha\beta} \tilde{m}_{j\beta} \\ &= \sum_{t'\alpha\beta} \tau_{dt't'} g_{dt'} \sigma_{dt'\alpha} \rho_{dt'\alpha\beta} \tilde{m}_{p_{dt'}\beta}. \end{aligned} \quad (4.5)$$

This model for the response is much more efficient than equation (4.3) due to the elimination of full-sky convolutions and rotations. However, this simplification came at the cost of assuming a single circular beam for each detector, which leaves out a common class of detectors, namely differential detectors. These measure signal differences between two spots on the sky, and have a beam consisting of two sub-beams, each with the same circular shape, but with opposite signs in their contribution. But since each of these beams are circular, we can still use the formula above, as long as we sum over the contributions from each sub-beam  $h$ , weighted by its amplitude

<sup>2</sup>These are usually good approximations – most experiments are designed to have as round beams as possible, and detectors in the same focal plane will have very similar beams provided that they are sensitive to the same frequencies, which I assume here.

<sup>3</sup>The gain is ideally constant, but detectors are imperfect, and the gain may in practice drift gradually. However, it will never be allowed to drift at anywhere near the time scale of the detector time response. Hence,  $g_{dt} \tau_{dt't'} \approx \tau_{dt't'} g_{dt'}$ .

$a_{dh}$ .

$$\begin{aligned} s_{dt} &= \sum_{t'h\alpha\beta} \tau_{dtt'} a_{dh} g_{dt'} \sigma_{dt'\alpha} \rho_{dt'\alpha\beta} \tilde{m}_{p_{dt'}\beta} \\ &= \sum_{t'h\alpha} \tau_{dtt'} \psi_{dht'\alpha} \tilde{m}_{p_{dht'}\alpha}. \end{aligned} \quad (4.6)$$

Here I have introduced the rotated sensitivity  $\psi_{dht'\alpha} = a_{dh} g_{dt'} \sigma_{dt'\alpha} \rho_{dt'\alpha\beta}$ .

The biggest remaining expense in equation (4.6) is the detector time convolution  $\tau$ . The importance of this factor is strongly dependent on the type of detector used. For amplifier-based radiometers it is practically a delta function, and can be ignored, while for bolometers it takes the form

$\tau_{dtt'} = e^{-\frac{t'-t}{z_d}}$  for  $t' > t$  and 0 otherwise, with  $z_d$  normally being 10 ms or smaller. The effect of  $\tau$  is similar to the effect of sampling at discrete time steps – it acts as a low-pass filter that limits sensitivity at high frequencies and hence, via the scanning pattern, the smallest scales on the sky. The scanning pattern is typically chosen with this in mind, so that  $\tau$  has negligible effect on the scales of interest for the experiment, and it is therefore usually safe to ignore this factor even for bolometers.

Equation (4.6) corresponds to an explicit form of the response matrix of

$$P_{dti\alpha} = \sum_{t'h} \tau_{dtt'} \psi_{dht'\alpha} \delta_{p_{dht'}i}. \quad (4.7)$$

This is a sparse matrix, with the sparsity contained in the Kronecker delta, and this sparsity ensures that multiplications with  $P$  are not prohibitively expensive.

### 4.3 The noise

So far we have only looked at the signal part of the time-ordered data, which are typically a minor constituent of the TOD, with the majority being made up by noise,

$$d_{dt} = s_{dt} + n_{dt}. \quad (4.8)$$

As discussed in section 2.1, the noise is Gaussian and approximately stationary within one CES, with covariance given by equation (2.2), which is repeated here for convenience:

$$N_{dd'ff'} = \sqrt{\phi_d(f)\phi_{d'}(f)} C_{dd'f} \delta_{ff'}. \quad (2.2)$$

A note on notation here: We will need 3 different representations of covariance matrices: the time domain version ( $N_{dd'tt'}$ ), the frequency domain version ( $N_{dd'ff'}$ ), and the power spectrum. The latter is simply the diagonal

of the frequency domain representation  $(N_{dd'}(f) = N_{dd'ff})^4$ . If the detector indices are left out, as in  $N_{tt'}$ , the result is an  $n_{\text{det}} \times n_{\text{det}}$  sub-matrix instead of a single element.

Since equation (2.2) is block-diagonal, we can find its inverse simply by inverting each  $n_{\text{det}} \times n_{\text{det}}$  sub-block. The block-diagonality also makes operating with  $N$  on dense vectors and matrices very efficient in Fourier space, which will come in handy. However, the advantage of working in Fourier space is lost for vectors and matrices which are sparse in time-domain. For these, we will need the time-domain version of  $N$  and  $N^{-1}$ . The discrete Fourier transform  $\mathcal{F}_{tf} = \frac{1}{\sqrt{n_s}} e^{-\frac{2\pi i t f}{n_s}}$ ,  $\mathcal{F}_{tf}^{-1} = \frac{1}{\sqrt{n_s}} e^{\frac{2\pi i t f}{n_s}}$  of a diagonal Fourier-domain matrix  $A_{ff'} = a_f \delta_{ff'}$  is

$$A_{tt'} = \sum_{ff'} \mathcal{F}_{tf}^{-1} A_{ff'} \mathcal{F}_{t'f} = \frac{1}{\sqrt{n_s}} \sum_f \mathcal{F}_{(t-t')f}^{-1} a_f, \quad (4.9)$$

with  $n$  being the number of time samples, so

$$N_{dd'tt'}^{\pm 1} = \frac{1}{\sqrt{n_s}} \sum_f \mathcal{F}_{(t-t')f}^{-1} N_{dd'ff}^{\pm 1}. \quad (4.10)$$

This expression only depends on the time difference  $\Delta t = t - t'$ , so we only need to store the time correlation function<sup>5</sup>  $N_{dd'}^{\pm 1}(\Delta t) = N_{dd'0\Delta t}^{\pm 1}$ .

## 4.4 From time-ordered data to sky

With our signal and noise model in hand, we are ready to reconstruct a map of the sky. The time-ordered data are given by

$$d = Pm + n, \quad (4.11)$$

resulting in  $\text{cov}(d) \equiv \langle dd^T \rangle = \langle nn^T \rangle \equiv N$ . The likelihood of a sky  $m$  given the data stream is

$$\mathcal{L}(m) = \frac{e^{-\frac{1}{2}(d-Pm)^T N^{-1}(d-Pm)}}{\sqrt{|2\pi N|}}, \quad (4.12)$$

which gives a maximum likelihood estimator of

$$\frac{d\mathcal{L}(\hat{m})}{d\hat{m}} = 0 \Rightarrow \hat{m} = (P^T N^{-1} P)^{-1} P^T N^{-1} d. \quad (4.13)$$

<sup>4</sup>Using the same symbol in all these cases is convenient because they all refer to the same underlying quantity. But if we actually insert numbers for the indices, the notation becomes ambiguous. In the rare cases where this happens, I will specify the domain explicitly.

<sup>5</sup>I here use the convention from signal processing, where the correlation function is not normalized by dividing by the variance. This quantity could therefore be more accurately called the time covariance function.

This is an unbiased estimator

$$\langle \hat{m} \rangle = \langle (P^T N^{-1} P)^{-1} P^T N^{-1} (Pm + n) \rangle = m, \quad (4.14)$$

with covariance

$$\begin{aligned} M &= \langle (P^T N^{-1} P)^{-1} P^T N^{-1} d d^T N^{-1} P (P^T N^{-1} P)^{-1} \rangle \\ &= (P^T N^{-1} P)^{-1} P^T N^{-1} N N^{-1} P (P^T N^{-1} P)^{-1} \\ &= (P^T N^{-1} P)^{-1}. \end{aligned} \quad (4.15)$$

For implementation purposes it will be convenient to introduce the quantity

$$r = P^T N^{-1} d, \quad (4.16)$$

such that  $M^{-1} \hat{m} = r$ . Working with  $r$  and  $M^{-1}$  has the advantage that these are linear in the data - if one has several series of time-ordered data, the contributions from each of these can be directly added into the total  $r$  and  $M^{-1}$ . Efficiency-wise, it is convenient that the covariance matrix and the matrix used when solving for the map turn out to be the same, as these matrices can have formidable sizes.

In practical applications, data will be contaminated by various poorly modeled effects, such as signal from the atmosphere, ground pickup, etc., and this makes it necessary to filter the data. It is easy to see that replacing  $N^{-1}$  by a general matrix  $K$  in equation (4.13) still yields an unbiased estimator, and we can use this to implement filtering by choosing  $K = FN^{-1}$

$$\hat{m}_f = (P^T F N^{-1} P)^{-1} P^T F N^{-1} d. \quad (4.17)$$

This filtered estimator has covariance

$$\begin{aligned} M_f &= \langle (P^T F N^{-1} P)^{-1} P^T F N^{-1} d d^T N^{-1} F P (P^T F N^{-1} P)^{-1} \rangle \\ &= (P^T F N^{-1} P)^{-1} P^T F N^{-1} F P (P^T F N^{-1} P)^{-1}. \end{aligned} \quad (4.18)$$

Here the filter  $F$  has been assumed to be symmetric. Implementation-wise,  $\hat{m}_f$  and  $M_f$  must be built up in terms of

$$r_f = P^T F N^{-1} d \quad (4.19)$$

$$A_f^{-1} = P^T F N^{-1} P \quad (4.20)$$

$$A_{ff}^{-1} = P^T F N^{-1} F P, \quad (4.21)$$

such that

$$A_f \hat{m}_f = r_f \quad (4.22)$$

$$M_f = A_f A_{ff}^{-1} A_f. \quad (4.23)$$

We see that introducing the filter means that we now have to deal with three matrices instead of one:  $A_f$ ,  $A_{ff}$  and the covariance matrix itself, making  $\hat{m}_f$  a more computationally expensive estimator compared to  $\hat{m}$ .

Under certain circumstances it is possible to avoid this expense: If  $FN^{-1}F = FN^{-1}$ , equation (4.18) will simplify to

$$M_f = (P^T F N^{-1} P)^{-1}. \quad (4.24)$$

A common case is for both  $N$  and  $F$  to be diagonal in Fourier space, in which case the condition simplifies to  $F^2 = F$ , implying a hard filter with abrupt jumps between the values 0 and 1. As we shall see, hard filters like this correspond to very long correlations in the time domain, limiting the usefulness of this approach.

An informative way of rewriting the filter which we shall need later is

$$FN^{-1} = (N + G)^{-1}, \quad (4.25)$$

so that

$$\hat{m}_f = (P^T (N + G)^{-1} P)^{-1} P^T (N + G)^{-1} d. \quad (4.26)$$

The observant reader will notice that this form of the filtered estimator is equivalent to an unfiltered map estimator based on the data model  $d = Pm + n + g$ , where  $g$  is an extra Gaussian component with covariance  $G$ . That is, an alternative way of implementing filtering is by assuming the presence of extra noise. This leads to the same map estimator, but a simpler form for the covariance:  $M = (P^T (N + G)^{-1} P)^{-1}$ . However, for this to be valid the extra noise must *actually be present* in the data, otherwise the covariance of the final map will be overestimated<sup>6</sup>. If signal contaminants like bad weather and ground pickup could be accurately modeled as Gaussians, this would be an optimal and efficient way of handling them. This is unlikely to be the case in practice.

#### 4.4.1 Binned maps

It is sometimes convenient to be able to make approximate maps of the sky without paying the full computational cost of equation (4.17), for example for diagnostic output of subsets of data going into the full map-making procedure. A simple way of doing this is to ignore the non-diagonal elements of  $P^T F N^{-1} P$ , replacing it with its diagonal  $D$ . This leads to the fast but biased map estimator

$$\hat{m}_b = D^{-1} P^T F N^{-1} d. \quad (4.27)$$

---

<sup>6</sup>Implementing filtering by assuming an extra noise component when it is not actually present is equivalent to assuming that  $FN^{-1} = FN^{-1}F$  when using a soft filter.

In the absence of filtering and differential detectors, this is equivalent to equation (4.13), but filtering inevitably biases the resulting map, producing characteristic stripes and shadows around strong sources.

## 4.5 Filters

The job of the filters is to remove parts of the data that cannot be incorporated into the noise model, either because it is not Gaussian, not stationary, or too poorly known. Common effects that need to be filtered out include atmospheric disturbances, sidelobe pickup from the ground and the sun, glitches in the electronics, etc. (see section 2.1.2). A good filter should remove as much as possible of these effects while removing as little as possible of the actual signal<sup>7</sup> and this requires a rough model for where the effects appear, such as a frequency range, certain time intervals, etc., so that these regions can be specifically targeted.

Since we assume  $N$  to be diagonal in Fourier space, filters with the same property can be much more efficiently implemented than other, more general types. We will therefore decompose the full filter into a frequency filter and a general filter.

### 4.5.1 Frequency filters

A frequency filter can be represented as a power spectrum  $F_{dd'}(f)$  with a value close to zero in contaminated regions and close to one outside. We have already seen that the noise itself can be represented as a power spectrum  $N_{dd'}(f)$ . These are simply the diagonals of the Fourier space covariance matrices, so (leaving out the detector indices in the following)

$$\begin{aligned} (FN^{-1})_{tt'} &= \sum_f \mathcal{F}_{(t-t')f}^{-1} F(f) N(f)^{-1} \\ (FN^{-1}F)_{tt'} &= \sum_f \mathcal{F}_{(t-t')f}^{-1} F(f)^2 N(f)^{-1}. \end{aligned} \quad (4.28)$$

So effectively frequency filters simply act as an update to the noise power spectrum, producing the two new effective spectra  $N_F^{-1} \equiv FN^{-1}$  and  $N_{FF}^{-1} \equiv FN^{-1}F$ .

For example, to completely suppress every frequency below the frequency  $\alpha$ , one would use  $F(f < \alpha) = 0$  and  $F(f \geq \alpha) = 1$ . See figure 4.1 for an illustration of this.

While one is theoretically free to choose any  $F(f)$ , correlation length considerations prevent us from choosing a shape with too sharp transitions.

---

<sup>7</sup>There is a trade-off between statistical and systematic error here. More aggressive filtering will decrease systematic effects but increase error bars by removing more signal.






Figure 4.1: Frequency filters are straightforwardly implemented as mode-wise multiplication with the inverse noise power spectrum. The x and y axes are frequency and inverse power in arbitrary units, respectively.

This is illustrated in figure 4.5.1. We have found a  $1/f$ -type filter  $F(f) = \left(1 + \left[\frac{f}{f_{\text{cut}}}\right]^{\alpha_{\text{cut}}}\right)^{-1}$  to provide a reasonable trade-off between sharpness in the two domains.

#### 4.5.2 General filters

When the filters are not diagonal in frequency domain, we cannot simply merge  $F$  with  $N$ , and must find another way of implementing the filter. The filtered estimator is built up from three sub-expressions given in equations (4.19,4.20,4.21), and using equation (4.25) we can express these as

$$r_f = P^T (N_F + G)^{-1} d \quad (4.29)$$

$$A_f^{-1} = P^T (N_F + G)^{-1} P \quad (4.30)$$

$$A_{ff}^{-1} = P^T (N_{FF} + G)^{-2} N_{FF} P. \quad (4.31)$$

This form is amenable to manipulation with the Woodbury matrix identity [30, 31]

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1}. \quad (4.32)$$

The identity is effectively a prescription for how to find  $(A + UCV)^{-1}$  if you already know  $A^{-1}$  *without needing to invert*  $A^{-1}$ . This comes at the cost of needing to invert the matrix  $C^{-1} + VA^{-1}U$ , so for this expansion to have any point,  $C$  must be a much smaller matrix than  $A$ ,

How can this be used in practice? Consider the case where we have a small number of basis vectors  $\{\vec{U}_i\}$  which span our model for the contamination  $g$ . This could, for example, be a set of Legendre polynomials in terms of the azimuth  $\alpha$ :  $U_{ti} = P_i(2\frac{\alpha_t - \alpha_{\min}}{\alpha_{\max} - \alpha_{\min}} - 1)$ , or perhaps a set of time steps  $\tau_i$  which contain a strong glitch:  $U_{ti} = \delta_{\tau_i t}$ . We can then build up  $g$  as




Figure 4.2: The shape of the frequency filter greatly affects the correlation length, which drives the main computational cost of building the pixel covariance matrix. Shown in the first panel are three different inverse power spectra, corresponding to no filter, a hard filter and a soft filter respectively. In the second panel, the inverse time correlation functions of the same three cases are shown. For the hard filter, the correlation never falls below 1% of the maximal value in this example, meaning that the correlation length is greater than the length of the time-ordered data.

$g_t = U_{ti}\gamma_i$ , with covariance  $\langle \gamma\gamma^T \rangle \equiv C$ , giving  $G \equiv \langle gg^T \rangle = UCU^T$ , which is of the form we need for the Woodbury identity.

Note that we are not assuming that the covariance  $G$  is a good model for the contamination here – if we had a good model of it we could treat it as a noise component instead<sup>8</sup>. Instead, what is being assumed here is that  $G$  is much larger than the variance of the spurious signal, so that to a good approximation the latter no longer matters. A simplifying assumption that guarantees this is to let the nonzero parts of  $G \rightarrow \infty$  by setting  $C^{-1} = 0$ . In addition to removing  $C$  from equation (4.32), it also results in  $(N_{FF} + G)^{-2}N_{FF} = (N_{FF} + G)^{-1}$ , greatly simplifying the application of the Woodbury identity to equation (4.31).

Applying this to equations (4.29,4.30,4.31), we get

$$r_f = P^T N_F^{-1} d - P^T N_F^{-1} U (U^T N_F^{-1} U)^{-1} U^T N_F^{-1} d \quad (4.33)$$

$$A_f^{-1} = P^T N_F^{-1} P - P^T N_F^{-1} U (U^T N_F^{-1} U)^{-1} U^T N_F^{-1} P \quad (4.34)$$

$$A_{ff}^{-1} = P^T N_{FF}^{-1} P - P^T N_{FF}^{-1} U (U^T N_{FF}^{-1} U)^{-1} U^T N_{FF}^{-1} P. \quad (4.35)$$

The first terms here are analogous to those from the filterless estimator (4.15,4.16), while the second term corresponds to a post-hoc update. Defining  $r_{Fu} \equiv P^T N^{-1} U$ ,  $r_{FFu} \equiv P^T N^{-1} U$  and  $N_{Fu}^{-1} \equiv U^T N_F^{-1} U$ ,  $N_{FFu}^{-1} \equiv U^T N_{FF}^{-1} U$ ,

<sup>8</sup>See section 4.4.

this simplifies somewhat to

$$r_f = P^T N_F^{-1} d - r_{Fu} N_{Fu} U^T N_F^{-1} d \quad (4.36)$$

$$A_f^{-1} = P^T N_F^{-1} P - r_{Fu} N_{Fu} r_{Fu}^T \quad (4.37)$$

$$A_{ff}^{-1} = P^T N_{FF}^{-1} P - r_{FFu} N_{FFu} r_{FFu}^T. \quad (4.38)$$

This update is dominated by the cost of calculating the two  $r_u N_u r_u^T$  terms, which will be cheap as long as the number of basis functions is kept low, and has the additional advantage that general filters can be added to an existing program simply by adding a small section for updating  $r_f$ ,  $A_f$  and  $A_{ff}$ .

## 4.6 Practical implementation of the map-making equations

We saw in the previous section that the filtered map estimator  $\hat{m}_f$  in rough steps can be described as

1. Apply the Fourier filters to the noise covariance
2. Build the components of the map-making equations by projecting into pixel space with  $P$
3. Apply the general filters to these
4. Solve for the map and its covariance

Most of these operations involve straightforward matrix operations, but operations involving the response matrix  $P$  need to be tailored to take into account its sparsity in order to avoid prohibitive computational cost.  $P$  appears in two types of products:  $P^T D V$ , where  $D$  is a Fourier-diagonal matrix and  $V$  is a dense matrix; and the heavier operation  $P^T D P$ .

### 4.6.1 Implementing $P^T D V$

By using equation (4.7), we see that

$$\begin{aligned} W &\equiv P^T D V \Rightarrow \\ W_{i\alpha I} &= \sum_{dt} P_{dtp\alpha} [DV]_{dtI} \\ &= \sum_{dth} \psi_{dht\alpha} \delta_{ip_{dht}} [\tau DV]_{dtI}. \end{aligned} \quad (4.39)$$

The detector time response function is stationary, so  $\tau$  is well-approximated as diagonal in frequency space, as is  $D$ . The product in the brackets can therefore be efficiently calculated with fast Fourier transforms:

$$[\tau DV]_{dtI} = \sum_{f't'd'} \mathcal{F}_{tf}^{-1}(\tau_{dff} D_{dd'ff} \mathcal{F}_{ft'} V_{d't'I}). \quad (4.40)$$

This means that support for detector time response functions comes almost for free.

To be useful equation (4.39) needs to be reorganized to exploit the sparsity encoded in the Kronecker delta. This results in the following algorithm

---



---

```

W = 0
for all  $d, t, \alpha, I, h$  do
   $W_{p_{dht}\alpha I} += \psi_{dht\alpha} [\tau DV]_{dtI}$ 
end for

```

---

#### 4.6.2 Implementing $P^T DP$

Starting again from equation (4.7), we find

$$\begin{aligned} W &\equiv P^T DP \\ W_{ij\alpha\beta} &= \sum_{dd'tt'} P_{dti\alpha} D_{dd'tt'} P_{d't'\beta} \\ &= \sum_{dd'tt'hh'} \psi_{dht\alpha} \delta_{ip_{dht}} [\tau^T D \tau]_{dd'tt'} \delta_{jp_{d'h't'}} \psi_{d'h't'\beta}. \end{aligned} \quad (4.41)$$

As before, the product in the brackets can be computed efficiently in Fourier space:

$$[\tau^T D \tau]_{dd'tt'} = \sum_f \mathcal{F}_{(t-t')f}^{-1} \tau_{dff} D_{dd'ff} \tau_{d'ff}. \quad (4.42)$$

Using the sparsity in equation (4.41) and the fact that equation (4.42) only depends on  $\Delta t = t - t'$ , we arrive at

---



---

```

W = 0
for all  $d, d', t, h, h', \alpha, \beta$  do
  for  $-t_{\text{corr}} \leq \Delta t \leq t_{\text{corr}}$  do
     $W_{p_{dht}p_{d'h'(t+\Delta t)}\alpha\beta} += \psi_{dht\alpha} [\tau^T D \tau]_{dd'\Delta t_0} \psi_{d'h'(t+\Delta t)\beta}$ 
  end for
end for

```

---

This algorithm is often the most expensive part of map-making, and since its computational time is proportional to the time correlation length

$t_{\text{corr}}$ , the time difference where  $\tau^T D \tau$  becomes negligible, this puts constraints on the form of  $D$  and  $\tau$  (refer to figure 4.5.1 for an illustration of this).

### 4.6.3 Solving the map-making equation

Equation (4.22) for solving for the final map and its covariance matrix consist of a straightforward dense matrix operation<sup>9</sup>, and can be handled by a linear algebra package like LAPACK<sup>10</sup> [35]. In principle the matrices involved are all symmetric and positive definite, and could be efficiently solved by using Cholesky factorization. Sadly, this is sabotaged by the effect of filtering, which tends to make the covariance matrix poorly conditioned or singular. For example, a high-pass filter which assigns zero weight to low-frequency modes will make the average of the map completely indeterminate. The map covariance matrix will therefore have an infinite eigenmode corresponding to the average of the map. Similar effects apply to a lesser degree to other modes that are damped by filters<sup>11</sup>. One must therefore use slower, more robust methods such as eigenvalue decomposition.

### 4.6.4 Implementation in the QUIET pipeline

Table 4.1 summarizes the steps for implementing the filtered map estimator  $\hat{m}_f$ . These steps are implemented in two components of the QUIET maximum likelihood pipeline.

---

<sup>9</sup>If one is only interested in the map itself, and not its covariance, these equations can be solved without actually needing to store or invert large matrices by using conjugate gradients [32]. But one ignores the covariance matrix at ones peril. As I will describe in the section on postprocessing, a map that is both filtered and unbiased will contain large, correlated noise modes which will be given undue weight if one does not have covariance information.

<sup>10</sup>One of the weaknesses of maximum likelihood map making is the expense of storing and operating on the matrices involved, which will easily grow larger than what is practical to handle with LAPACK. This can be mitigated by using parallel solvers such as those in ScaLAPACK [33] and Elemental [34].

<sup>11</sup>Cross-coupling by scanning the same part of the sky in several different directions will reduce the degeneracy for other modes.

$N_F^{-1}$	$=$	$FN^{-1}$	eq. (4.28)
$N_{FF}^{-1}$	$=$	$FN^{-1}F$	
$r$	$=$	$P^T N_F^{-1} d$	alg. (4.6.1)
$A_f^{-1}$	$=$	$P^T N_F^{-1} P$	alg. (4.6.2)
$A_{ff}^{-1}$	$=$	$P^T N_{FF}^{-1} P$	
$r_{Fu}$	$=$	$P^T N_F^{-1} U$	alg. (4.6.1)
$r_{FFu}$	$=$	$P^T N_{FF}^{-1} U$	
$N_{Fu}^{-1}$	$=$	$U^T N_F^{-1} U$	
$N_{FFu}^{-1}$	$=$	$U^T N_{FF}^{-1} U$	
$r$	$\Leftarrow$	$r_{Fu} N_{Fu} U^T N_F^{-1} d$	
$A_f^{-1}$	$\Leftarrow$	$r_{Fu} N_{Fu} r_{Fu}^T$	
$A_{ff}^{-1}$	$\Leftarrow$	$r_{FFu} N_{FFu} r_{FFu}^T$	
$\hat{m}_f$	$=$	$A_f r$	eq. (4.22)
$M_f$	$=$	$A_f A_{ff}^{-1} A_f$	eq. (4.24)

Table 4.1: Summary of the steps needed for calculating the filtered estimator  $\hat{m}_f$ .

Building  $r$ ,  $A_f^{-1}$  and  $A_{ff}^{-1}$  is done in the program **tod2map**, which consists of about 4000 lines<sup>12</sup> of Fortran 90 code, which processes constant elevation scans in parallel using MPI, and co-adds these into the final  $r$ ,  $A_f^{-1}$  and  $A_{ff}^{-1}$ . The computation time for this step scales linearly with the total length of the time-ordered data involved, and required  $\sim 10^4$  CPU hours to process all the CMB patches in the Q-band analysis, and  $\sim 5 \cdot 10^4$  CPU hours for the W-band analysis due to the larger amount of data involved.

In addition to the use of parallelization, **tod2map** received a large speedup through the use of data decimation. The algorithmic complexity of this part of the map-making algorithm is  $\mathcal{O}(N_{\text{samp}} N_{\text{corr}} N_d f_d)$ , where  $N_{\text{samp}}$  is the total number of samples per detector,  $N_{\text{corr}}$  is the noise correlation length, measured in samples,  $N_d$  is the total number of detectors, and  $f_d$  is the average fraction of these that each detector is correlated with. Both  $N_{\text{samp}}$  and  $N_{\text{corr}}$  scale with the sampling rate of the time-ordered-data, and consequently one can achieve a quadratic speedup simply by reducing the sampling rate.

Given a typical sky scanning speed  $v$  and a Gaussian beam with a standard deviation of  $\sigma = \sigma_{\text{FWHM}}(8 \log(2))^{-\frac{1}{2}}$ , the time-ordered data will be

<sup>12</sup>This number does not include supporting libraries that are shared with other parts of the pipeline.

smoothed by

$$b_{\text{TOD}}(f) = e^{-\frac{1}{2} \frac{f^2 \sigma^2}{v^2}}, \quad (4.43)$$

which has a typical width of  $\sigma_{\text{tod}} = \frac{v}{\sigma}$ . For QUIET,  $v \sim 1^\circ/\text{s}$ <sup>13</sup>, giving  $\sigma_{\text{tod}} = 5.2\text{Hz}$  in the Q-band and 12 Hz in the W-band. The sampling rate must be at least twice as high as the highest frequency one wishes to represent, so these numbers correspond to sampling rates of 10.4 Hz and 24 Hz respectively, which are both significantly lower than QUIET's native sampling rate of 50 Hz. This led us to downsample the TOD to 25 Hz in the W-band analysis, resulting in a factor 4 speedup.

Memory-wise, the cost is driven by the need to store the two  $A$ -matrices, though their symmetry permits us to only store half of each. To avoid artificial loss of resolution due to pixelization, the pixels should be significantly smaller than the beam, which in the Q- and W-bands have a FWHM of 27' and 12' respectively<sup>14</sup>. We therefore require a HEALPix  $N_{\text{side}}$  of at least 256 and 512 corresponding to a pixel diameter of 13.7' and 6.9' for the two frequency bands. For a typical QUIET CMB patch, this results in  $\sim 10^4$  pixels in the Q-band and  $\sim 4 \cdot 10^4$  in the W-band, which coupled with two polarization components per pixel in the case of a polarization only analysis results in  $\sim 2 \cdot 10^4$  and  $\sim 8 \cdot 10^4$  degrees of freedom in the map, and the square of these numbers in the matrices. Using single-precision (4 bytes floating point numbers) for the matrices, this amounts to about 1.5 GB (25 GB) of memory to store one full matrix for the Q-band (W-band) analysis.

## 4.7 Biaslessness and deconvolution

The estimator  $\hat{m}_f$  described in the previous sections was constructed to be unbiased even in the presence of filters. The immediate effect of the filters is to scale parts of the signal down ( $FN^{-1}d$ ), and heavily filtered signal components are effectively eliminated by this. In order to recover an unbiased map, the filters must be subsequently unapplied. In time (or frequency) domain, this could be done without loss of sensitivity by simply multiplying by  $(FN^{-1})^{-1}$ , but the projection operator  $P$  has the effect of mixing noise modes, introducing extra noise in the down-scaled modes. When the filters are unapplied ( $(P^T FN^{-1}P)^{-1}$ ), this noise is scaled up together with the signal, resulting in high-noise modes in the final map.

Unbiased filtered map-making is therefore a form of deconvolution, completely analogous to the process of deconvolving the instrument beam from the map. In the case of a beam, the dampened modes correspond

<sup>13</sup>The telescope scans at  $2^\circ/\text{s}$  in azimuth, which at a typical elevation of  $60^\circ$  corresponds to  $1^\circ/\text{s}$  on the sky.

<sup>14</sup>See section 2.

to the small scales, and deconvolving these produces very high amplitude small-scale noise. In contrast, for unbiased map-making with a high-pass filter the dampened modes correspond to large scales, and deconvolution results in large-scale, i.e. correlated, noise modes. See figure 4.3 for an illustration of these effects.

The full effect of these deconvolution-induced noise modes is described in the covariance matrix, which ensures that they are assigned no more weight than they are due. However, the presence of these modes makes visual inspection of the maps very difficult. When plotting the maps, we therefore apply an *eigenvector*-filter, defined by

$$m_{\text{cut}} = VKV^T m. \quad (4.44)$$

Here  $V$  is the eigenvector matrix of the map  $m$ 's covariance matrix  $M$ ,  $M = VEV^T$ ,  $E$  is a diagonal vector consisting of  $M$ 's eigenvalues, and  $K_{ij} = k_i \delta_{ij}$  with  $k_i = (E_i < \epsilon)$ .<sup>15</sup> The plots in the results section employ  $\epsilon = 100\min(E)$ , which typically results in  $\sim 20 - 40$  large scale modes being cut.<sup>16</sup>

The eigenvector-filter is a filter based purely on the noise properties of the map. If one is willing to make assumptions about the signal itself, one can do better by filtering based on the signal-to-noise ratio using a Wiener-filter [36]. If the signal is taken to be Gaussian with covariance matrix  $S$  (see section 5.1.1), then the Wiener filtered map is given by

$$m_{\text{cut}} = (1 + S^{-1}N)^{-1} m. \quad (4.46)$$

Making assumptions about  $S$ , one of the quantities one is trying to measure, is dangerous. For example, if the fiducial power spectrum does not contain any B-modes, then neither will the filtered map. It is therefore best to estimate the Wiener-filtered map jointly with the covariance matrix.

<sup>15</sup>I here use the C language convention that true/false maps to 1/0. Thus  $k_i$  is 1 if  $E_i < \epsilon$ , and otherwise 0.

<sup>16</sup>While primarily a means of visualization, it is also possible to use this as a cut before further data analysis. This could for example be done if one has evidence that certain eigenmodes are contaminated by ground pickup, atmosphere, or similar<sup>17</sup>. In this case, the corresponding modes in the covariance matrix should be assigned infinite variance, reflecting the fact that the value of these eigenmodes is being ignored:

$$M_{\text{cut}}^{-1} = VKV^T M^{-1} VV^T. \quad (4.45)$$

<sup>17</sup>If there were a 1-1 correspondence between filtered Fourier modes and noise modes in the final map, all the filtered pollution would be completely restored when deconvolving the map (but would still be assigned a huge variance in the covariance matrix). However, at least in the case of more than a single CES being analysed, there is no 1-1 correspondence, and each mode will be a linear combination of different frequencies in the TOD. In this linear combination, the heavily filtered modes weigh little, and may thus be a small part of the deconvolved result.



The Gibbs sampling procedure described in section 5.1.3 produces Wiener-filtered maps based on an unbiased estimate of the power spectrum as a by-product of the map-sampling step.

## 4.8 Results

### 4.8.1 CMB maps

Temperature and polarization maps for QUIET's four CMB patches are shown in figures 4.4-4.7. These have all been subjected to a mild eigenvector-filter in order to remove large, correlated noise modes which mostly correspond to multipoles lower than the  $25 \leq l \leq 1200$  that QUIET can comfortably measure with its small patch size. This typically results in the  $\sim 30$  noisiest modes out of  $\sim 6 \cdot 10^4$  being removed. An example of what these modes look like is shown at the bottom of figure 4.4.

The figures compare the QUIET maps with WMAP at the same frequencies after removing the same modes. In the temperature maps the CMB anisotropies are clearly visible, with both QUIET and WMAP observing the same pattern. QUIET does not quite reach WMAP's noise levels, but as described in the calibration and instrument chapter, the temperature maps are based on only  $\sim 3\%$  of the detectors and a temperature cut efficiency of only  $\sim 30\%$ , so this was to be expected.

In polarization, however, the noise level is low enough to faintly make out the  $\sim 5\mu\text{K}$  E-mode polarization signal from the second and third peak of the EE-power spectrum by eye. This is most clearly visible in the W-band map of the first CMB patch (and in the wiener-filtered version of the patch, which can be seen in figure 4.8). The corresponding WMAP maps are here completely noise dominated.

Several point sources are visible in the maps, with CMB patch 2 being most affected, containing four strong sources and several weaker ones. These are easily visible in the Q-band temperature maps, and one of the sources is even visible in polarization. We therefore applied a point source mask to the maps before further analysis for the Q- and W-band temperature analysis and the Q-band polarization analysis.

At the time of writing, the W-band first results article has still not been released, and these maps should therefore be considered preliminary and may be subject to change.

### 4.8.2 Foreground maps

Figures 4.9-4.10 show polarization maps for QUIET's two foreground patches. In galactic coordinates, the galactic plane is aligned in the +Q direction and its axis of rotation in the -Q direction, so it is not surprising that the galactic emissions are predominantly characterized by Stokes Q. The emissions are




Figure 4.3: The effect of deconvolution on noise. Top left: Input CMB map, based on the power spectrum shown below. Top middle: Map after adding the effects of the beam and noise. This map is biased, as shown by the corresponding power spectrum is shown in green below. Top right: The map after deconvolving the beam is unbiased, but is dominated by noise at small scales, shown in blue below. The bottom row shows the same sequence for filters instead of a beam. The left map is the input, the middle is the biased map  $P^T F N^{-1} d$ , and the right is the unbiased map  $(P^T F N^{-1} P)^{-1} P^T F N^{-1} d$ . Again, the deconvolution has amplified noise modes that now dominate the map, but these are now large-scale modes due to the high-pass and azimuth filters.

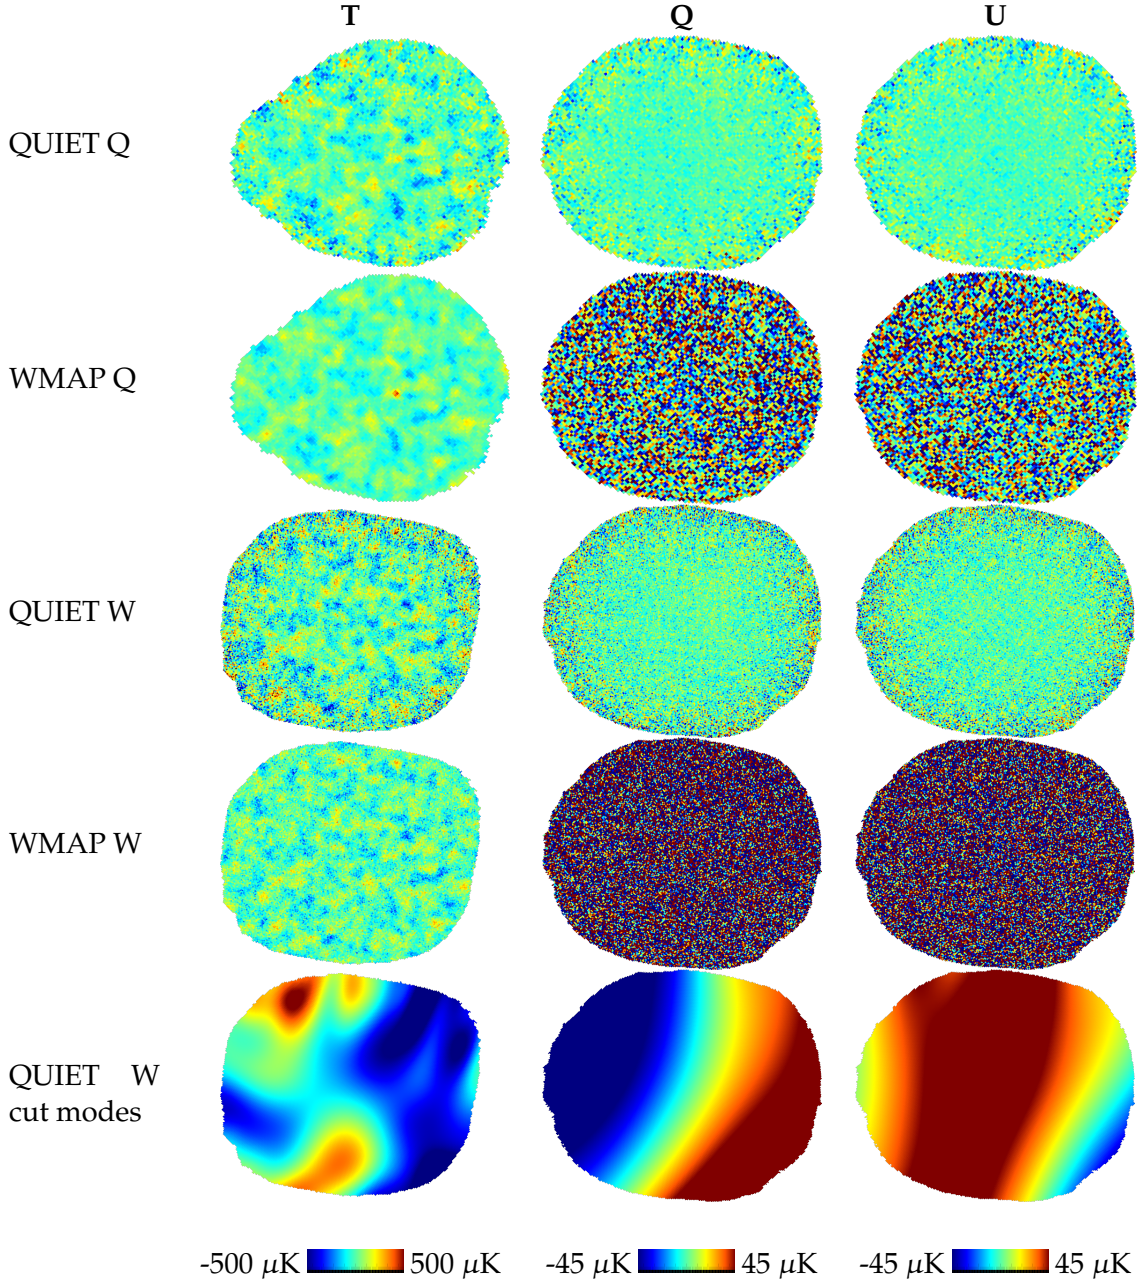


Figure 4.4: Maps for CMB patch 1, centered on galactic  $l = 292.2^\circ, b = 22.8^\circ$ , with a diameter of  $\sim 25^\circ$ . The columns are temperature (left), Stokes Q polarization (middle) and Stokes U polarization (right). Row 1-2: The QUIET Q-band result compared to the WMAP Q-band map, both at  $N_{\text{side}} = 256$ . Row 3-4: The QUIET W-band result compared to the WMAP W-band map, both at  $N_{\text{side}} = 512$ . An E-mode signal is faintly visible in the QUIET maps, particularly in the W-band map at small scales. The temperature map is consistent with WMAP, but a bit more noisy. Row 5: Maps of the high-noise modes removed by the eigenvalue filter for QUIET's W-band maps. While high in amplitude, these maps contain only  $\lesssim 50$  very high variance degrees of freedom.



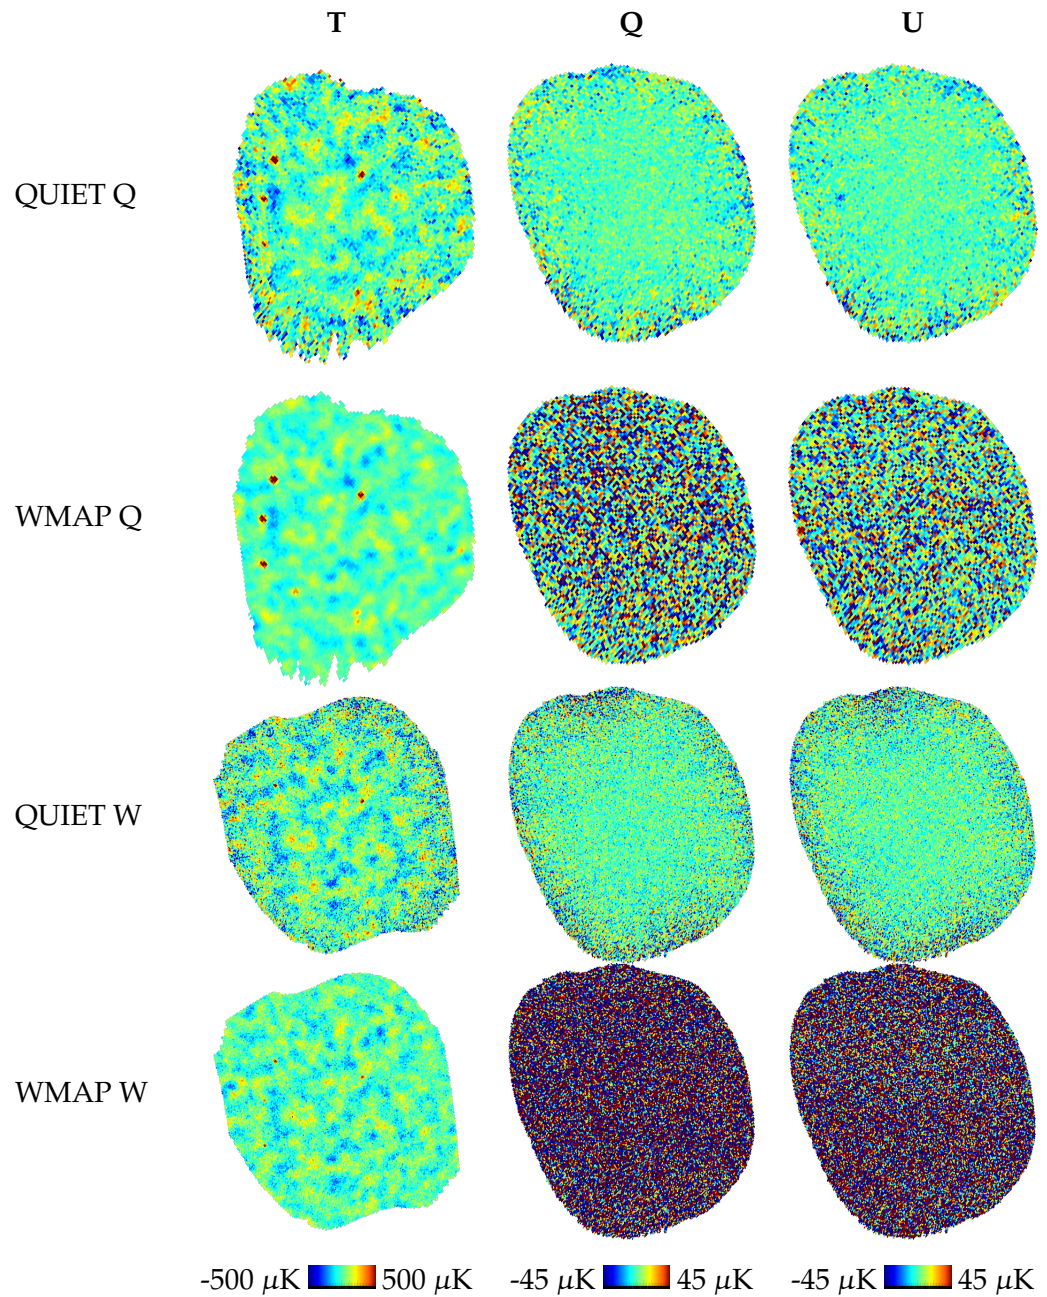
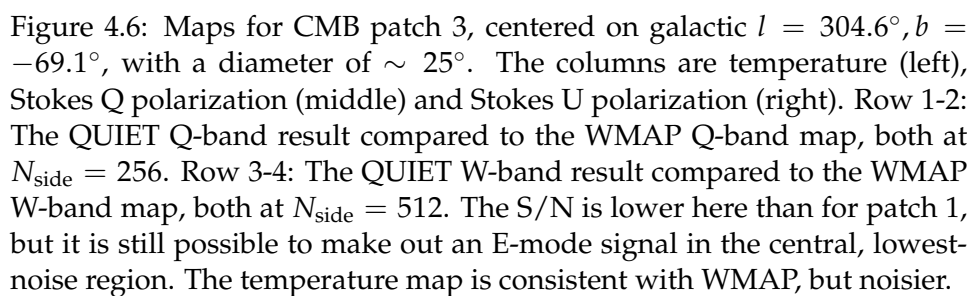


Figure 4.5: Maps for CMB patch 2, centered on galactic  $l = 243.2^\circ, b = -35.3^\circ$ , with a diameter of  $\sim 25^\circ$ . The columns are temperature (left), Stokes Q polarization (middle) and Stokes U polarization (right). Row 1-2: The QUIET Q-band result compared to the WMAP Q-band map, both at  $N_{\text{side}} = 256$ . Row 3-4: The QUIET W-band result compared to the WMAP W-band map, both at  $N_{\text{side}} = 512$ . The S/N is lower here than for patch 1, but it is still possible to make out an E-mode signal. A polarized point source is visible in the Q-band map. This and several other point sources were masked out before the power spectrum was estimated. The temperature map is consistent with WMAP, but noisier.





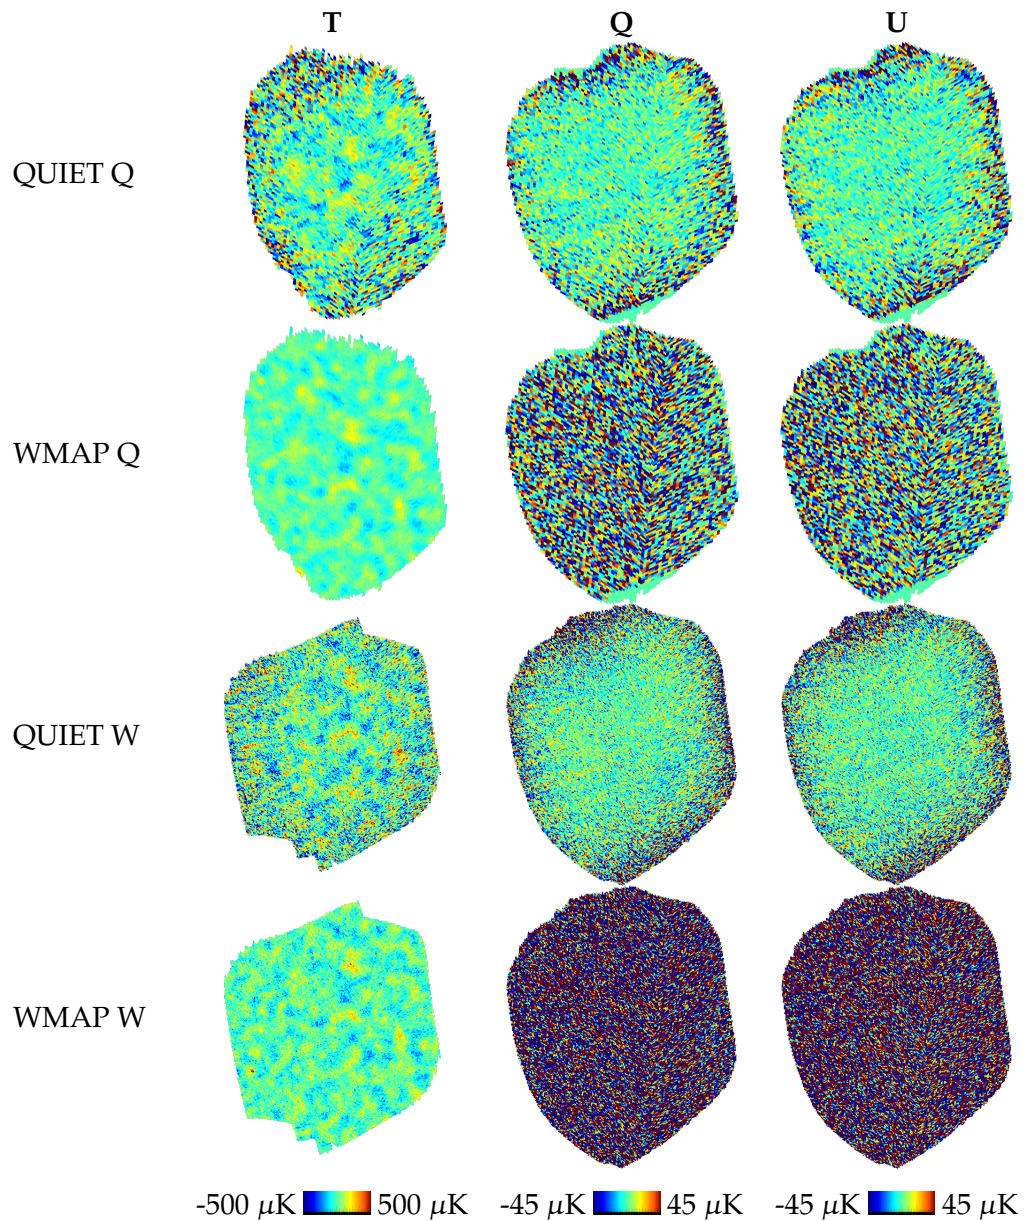


Figure 4.7: Maps for CMB patch 4, centered on galactic  $l = 7.0^\circ, b = -62.0^\circ$ , with a diameter of  $\sim 25^\circ$ . The columns are temperature (left), Stokes Q polarization (middle) and Stokes U polarization (right). Row 1-2: The QUIET Q-band result compared to the WMAP Q-band map, both at  $N_{\text{side}} = 256$ . Row 3-4: The QUIET W-band result compared to the WMAP W-band map, both at  $N_{\text{side}} = 512$ . The S/N very low here due to the limited exposure time. The thin band of zero signal at the bottom of the Q-band maps are pixels that were cut by the eigenvalue filter due to having excessive noise. The temperature map is consistent with WMAP, but much noisier.




Figure 4.8: Wiener-filtered maps of CMB patch 1 from the Gibbs sampling chain. The top row shows the Stokes Q and U components from left to right, with a clear E-mode signature (“+”-shape in Q, “x”-shape in U) visible. The bottom row shows maps of the E and B modes, confirming a strong detection of E-modes, while the B-modes are consistent with noise. Figure courtesy of H. K. K. Eriksen.

several times stronger in the Q-band than W-band, but their shape is different in the two bands, with the emissions being more concentrated towards the plane of the galaxy in the W-band. This is consistent with a transition from synchrotron dominance to a weak dust-dominance when going from Q to W-band, with the W-band being close to the foreground minimum. Most of the galactic emissions are in the +Q-direction, which for synchrotron radiation indicates electrons spiraling in a north-south magnetic field. However, in the galactic center there is a small, vertically elongated region with very strong -Q emissions which indicates a horizontal magnetic field here. This feature is part of a structure called the galactic center lobe [37, 38, 39], and has been previously imaged in the Q-band at higher resolution but lower sensitivity than QUIET [37], revealing it to be a thin, vertical and plume-like feature.

The same features are visible in the WMAP Q and W-maps, but the signal-to-noise is quite low, particularly in the W-band. Due to the steep spectral index ( $s \propto \nu^\alpha$ ,  $\alpha \sim -3$ ) of synchrotron, much higher signal-to-noise synchrotron maps are available at lower frequencies. In the bottom row of the two figures, I therefore plot the WMAP Ka-band polarization maps for comparison. These have a similar signal-to-noise as the QUIET Q-band maps, and agree well with these considering the differences in angular resolution.

Note that the QUIET foreground maps shown here are still preliminary. In particular, QUIET's temperature-to-polarization leakage, which is at  $\sim 1\%$  in the TOD, but much lower in map space due to frequent rotation of the polarization angles, has not been corrected here. Hence, regions of strong, unpolarized emission may show up as a spurious component in these polarization maps. This will be rectified in a future QUIET article on foregrounds.



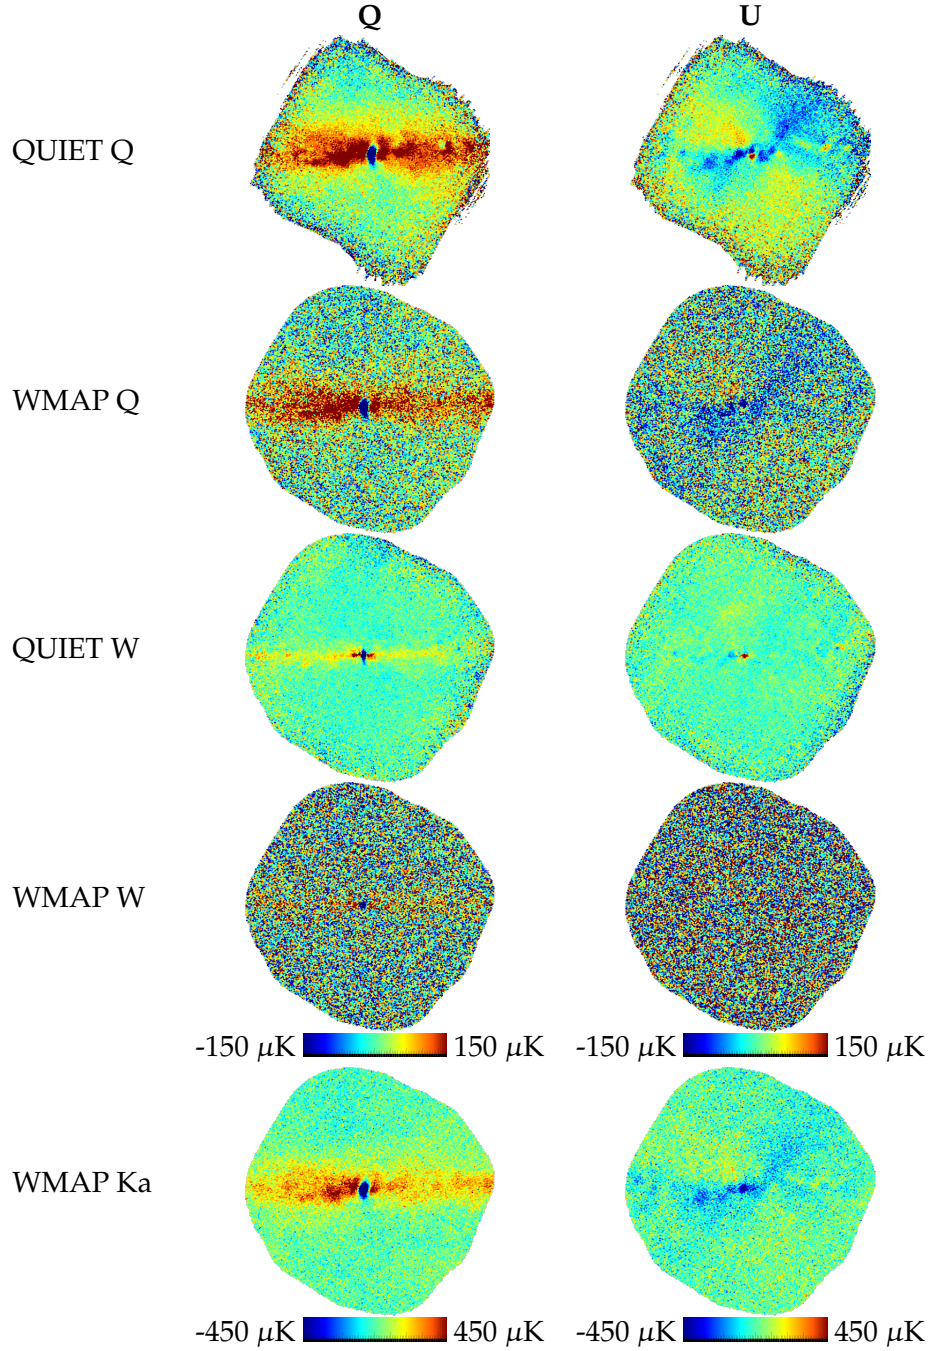


Figure 4.9: Maps of QUIET’s central galactic patch, centered at  $l = 0.1^\circ$ ,  $b = -0.1^\circ$ , with a diameter of  $\sim 20^\circ$ . The QUIET Q- and W-band results are consistent with WMAP, but WMAP’s high noise level makes a visual comparison difficult. I therefore provide the high S/N WMAP Ka-band map in the last row. The QUIET Q-band and WMAP Ka-band maps are both synchrotron-dominated, and show excellent agreement.

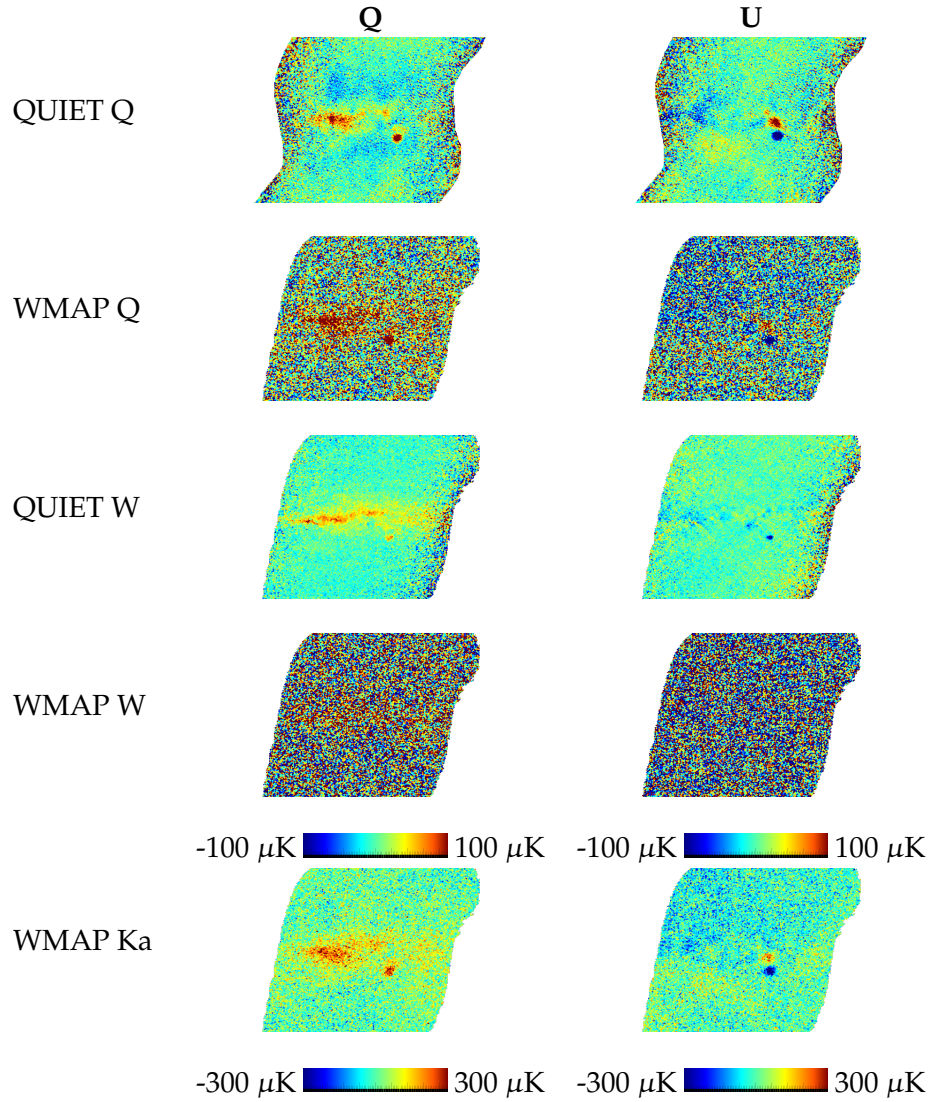


Figure 4.10: Maps of QUIET's other galactic patch, centered at  $l = 329^\circ$ ,  $b = 0^\circ$ , with a diameter of  $\sim 20^\circ$ . The columns are the Q and U Stokes parameters respectively. The top two rows show the QUIET and WMAP maps for the Q-band, while the next two rows have the corresponding maps for the W-band. WMAP is consistent with QUIET, but is very noisy, making comparison difficult. However, as synchrotron dominates below  $\sim 80\text{GHz}$ , and has a steeply falling spectrum, the WMAP Ka-band map is signal dominated, allowing for a straightforward visual comparison.

## Chapter 5

# Power spectrum and parameter estimation

Though map-making is a necessary step of the analysis pipeline, and maps are very useful for e.g. foreground analysis, they are not the final scientific goal of CMB experiments. As discussed in chapter 1, the CMB anisotropies are understood to be ultimately sourced by quantum fluctuations set up during inflation, and while the value of these fluctuations at any given position is random, their statistical distribution is governed by the physics from inflation to the surface of last scattering, providing us with a powerful probe of the physics of the early universe. If these fluctuations are Gaussian, homogeneous and isotropic as indeed they seem to be, then they can be completely described by a covariance matrix which is block-diagonal in spherical harmonic basis<sup>1</sup>:

$$a_{lm} \equiv (a_{T,lm}, a_{E,lm}, a_{B,lm})^T \sim N(0, C_l) \quad (5.1)$$

$$\langle a_{X,lm} a_{Y,l'm'} \rangle = C_l^{XY} \delta_{ll'} \delta_{mm'}, \quad (5.2)$$

The indices  $X, Y$  run over  $T, E, B$ , and the power spectrum  $C_l$  is the  $\{T, E, B\} \times \{T, E, B\}$  covariance matrix for a given multipole.  $a_{T,lm}$  are the normal spherical harmonic coefficients of the temperature map  $m_T$ , and  $a_{E,lm}, a_{B,lm}$  are rotationally invariant linear combinations of the spin-weighted spherical harmonic coefficients of the linear polarization fields  $m_Q$  and  $m_U$ , and

---

<sup>1</sup>A word on notation here: Due to the polarization degrees of freedom,  $a_{lm}$  is not a scalar quantity, but a vector of length  $N_{\text{pol}}$ . Similarly,  $C_l$  is an  $N_{\text{pol}} \times N_{\text{pol}}$  matrix. I will use the notation  $C$  to mean the  $N_l N_{\text{pol}} \times N_l N_{\text{pol}}$  matrix with  $C_l$  on the diagonal, such that  $C_l^{XY} \equiv (C)_{ll}^{XY}$ .

are given by [11]

$$a_{T,lm} = \int_{\Omega} Y_{lm}^*(\hat{n}) m_T(\hat{n}) d\Omega \quad (5.3)$$

$$a_{E,lm} = -\frac{1}{2} \int_{\Omega} [Y_{2,lm}^*(\hat{n}) + Y_{-2,lm}(\hat{n})] [m_Q(\hat{n}) + im_U(\hat{n})] d\Omega \quad (5.4)$$

$$a_{B,lm} = \frac{i}{2} \int_{\Omega} [Y_{2,lm}^*(\hat{n}) - Y_{-2,lm}(\hat{n})] [m_Q(\hat{n}) + im_U(\hat{n})] d\Omega. \quad (5.5)$$

For a given cosmological model one can calculate the predicted power spectrum in terms of that model's parameters<sup>2</sup>, and observational bounds on the power spectrum can therefore be used to constrain the cosmological parameters, or exclude the model altogether. Our primary interest in the sky maps is therefore to use them to estimate the power spectrum and ultimately the cosmological parameters.

## 5.1 Power spectrum estimation

Starting from equation (5.1), we see that given the harmonic coefficients, the likelihood for the power spectrum is

$$\mathcal{L}(C_l) = P(a_{lm}|C_l) = \prod_m \frac{1}{\sqrt{|2\pi C_l|}} e^{-\frac{1}{2} a_{lm}^T C_l^{-1} a_{lm}}, \quad (5.6)$$

corresponding to an independent Wishart distribution for each multipole in the power spectrum:

$$C_l \sim W \left( (2l+1)^{-2} \sum_m a_{lm} a_{lm}^T, 2l+1 \right). \quad (5.7)$$

The spherical harmonics coefficients used in these expressions are defined in equations (5.3-5.5), and while these appear to require  $\mathcal{O}(N_l^2 N_{\text{pix}}) \sim \mathcal{O}(N_{\text{pix}}^2)$  calculations, they can actually be implemented much more efficiently due to the existence of  $\mathcal{O}(N_{\text{pix}}^{\frac{3}{2}})$  algorithms such as those provided in the HEALPix package [26]<sup>3</sup>. Thus, for full-sky noise-less maps (or maps with uniform, uncorrelated noise), the full likelihood of the power spectrum can be very efficiently calculated.

However, in real applications the map, and thus its harmonic coefficients, will not be exactly known. Typically, the noise in the map will be

<sup>2</sup>This can be done by using a numerical Boltzmann solver such as CAMB [40], CMBfast [41] and CMBEasy [42].

<sup>3</sup>Asymptotically faster algorithms are under active research [43, 44], and Fast Legendre Transform-based algorithms promise an asymptotic complexity of  $\mathcal{O}(N_{\text{pix}} \log(N_{\text{pix}})^2)$ . However, current implementations of these algorithms have extremely high memory overhead and large pre-factors, which outweighs their advantages in most cases.

Gaussian with variance varying from pixel to pixel, and as we saw in chapter 4, instrument 1/f-noise and filters result in the noise in different pixels and stokes parameters being correlated, resulting in a dense pixel covariance matrix. Additionally, part of the sky will usually be masked out due to lack of observations or foreground contamination, effectively leading to infinite variance in those pixels, and the telescope optics smooths out the small scales of the map, biasing the higher multipoles low.

### 5.1.1 Brute force evaluation

The most straightforward approach for dealing with these complications is brute-force pixel-space evaluation of the likelihood. The covariance between the signal components at two points on the sky,  $m(\hat{n}_1)$  and  $m(\hat{n}_2)$ , depends on the power spectrum as

$$S(\hat{n}_1, \hat{n}_2) = R(\hat{n}_1, \hat{n}_2) \left\{ \sum_l \frac{2l+1}{4\pi} \lambda_l(\hat{n}_1 \cdot \hat{n}_2, b_l C_l b_l) \right\} R(\hat{n}_2, \hat{n}_1) \quad (5.8)$$

$$\lambda_l(\beta, C_l) = \begin{bmatrix} d_{00}^l C_l^{TT} & d_{20}^l C_l^{TE} & -d_{20}^l C_l^{TB} \\ d_{20}^l C_l^{TE} & \Delta_+^l C_l^{EE} + \Delta_-^l C_l^{BB} & -d_{2-2}^l C_l^{EB} \\ -d_{20}^l C_l^{TB} & -d_{2-2}^l C_l^{EB} & \Delta_-^l C_l^{EE} + \Delta_+^l C_l^{BB} \end{bmatrix}. \quad (5.9)$$

Here,  $R(\hat{n}_1, \hat{n}_2)$  is a matrix representing the rotation of the Stokes parameters when under parallel transport from  $\hat{n}_1$  to  $\hat{n}_2$ ,  $b_l$  is a diagonal matrix with the temperature and polarization beams  $b_{T,l}$ ,  $b_{E,l}$  and  $b_{B,l}$  on the diagonal, and  $d_{mm'}^l(\beta)$  are Wigner d-matrices, and  $\Delta_{\pm}^l \equiv \frac{1}{2}(d_{22}^l \pm d_{2-2}^l)$  [45].

When evaluated at the pixel locations in the map, equation (5.8) provides us with the pixel covariance matrix corresponding to a given power spectrum, which combined with the noisy map data yields the following likelihood

$$\mathcal{L}(C_l) = P(m|C) = \frac{e^{-\frac{1}{2}m^T[S(C)+M]^{-1}m}}{\sqrt{|2\pi[S(C)+M]|}}. \quad (5.10)$$

From this, the maximum likelihood solution can be found with a nonlinear search such as Newton-Raphson iteration, and error bars can be found by grid evaluation, Monte Carlo sampling or through a local Gaussian approximation [46]. These are all quite heavy, as they involve covariance matrices of size  $(N_l N_{\text{comp}})^2$ , which are needed because limited sky coverage, correlated noise and filtering induce correlations between nearby multipoles.

### Binning

The complication of correlated multipoles can be removed by binning several of them together, such that

$$C_b = \sum_{l \in XY} P_{bl}^{XY} C_l^{XY}, \quad (5.11)$$

where  $P_{bl}^{XY}$  is a projection operator down to the binned space, with a typical choice being (following [47])

$$P_{bl}^{XY} = \left( \sum_{l \in b} \right)^{-1} \frac{l(l+1)}{2\pi} \delta_{[XY][XY]_b}, \text{ if } l \in b, \text{ otherwise } 0. \quad (5.12)$$

That is, each bin is the average of the flattened<sup>4</sup> spectrum in multipole ranges, with the different spectra and cross-spectra mapping on to different bins (i.e. not being averaged together). Expanding back from bins to multipoles can be done with the operator

$$Q_{lb}^{XY} = \frac{2\pi}{l(l+1)} \delta_{[XY][XY]_b}, \text{ if } l \in b, \text{ otherwise } 0, \quad (5.13)$$

i.e. piecewise constant interpolation of the flattened spectrum. With sufficient bin size, each bin will be approximately independent, removing the need for evaluating marginals and covariances.

### Direct evaluation of cosmological parameters

Instead of estimating the power spectrum, it is also possible to directly estimate cosmological parameters, and in many ways the latter is easier than the former due to the much smaller number of free parameters, typically of the order of 5 - 10. In this case, no binning is necessary. Since the likelihood for a given power spectrum can be calculated directly, the likelihood for a set of parameters  $\theta$  is simply

$$\mathcal{L}(\theta) = \mathcal{L}(C(\theta)). \quad (5.14)$$

As for the power spectrum itself, the maximum likelihood point can be efficiently found using a non-linear search, and the likelihood can be mapped out using grid evaluation (when the number of parameters is small), or by MCMC sampling such as that provided by CosmoMC [48], or newer approaches such as MultiNest [49].

### Practicality

While straightforward and exact, this method scales poorly with the number of pixels. The cost of an evaluation of the likelihood is driven by the calculation of the determinant, which scales as  $\mathcal{O}(n^3)$ , where  $n = N_{\text{pix}} N_{\text{comp}}$  is the number of degrees of freedom in the map, so each doubling of the angular resolution results in 64 times higher cost.

---

<sup>4</sup>The power spectra have a natural scaling of  $\frac{2\pi}{l(l+1)}$ . Flattering the spectrum refers to dividing by this produces a much more slowly varying spectrum, which is more suited for binning. The flattened spectrum is also the standard when plotting the power spectrum.

The speed of the algorithm can be somewhat improved by performing a change of basis [50]. Let  $R = N^{-1}S_0N^{-1}$ , where  $S_0$  is a fiducial signal covariance matrix, and compute its eigenvalue decomposition  $Q\Lambda Q^T = R$ . This provides an orthogonal basis based on signal-to-noise modes, which allows one to truncate the basis at the signal-to-noise level corresponding to the required accuracy, resulting in a smaller number of degrees of freedom  $n'$ . In the truncated basis,  $S + N \rightarrow S' + N' = Q'^T S Q' + Q'^T N Q'$ , and  $m \rightarrow m' = Q'^T m$ , where  $Q'$  is the  $n \times n'$  truncated eigenvector matrix. This basis only needs to be calculated once, and can provide a substantial performance boost when  $n'$  is significantly smaller than  $n$ . Low signal-to-noise modes are generated by filtering, and may also be present due to the scanning pattern, and with QUIET's filters and scanning strategy a relative eigenvalue threshold of  $10^{-8}$  resulted in  $n/n' \sim 2.5$ , for a speedup of a factor of 15.

Notwithstanding this optimization, the practical limit for this method is of the order  $10^5 - 10^6$  degrees of freedom with current clusters. For comparison, a full-sky polarization map with a HEALPix  $N_{\text{side}}$  parameter of 2048, corresponding to a pixel resolution of 1.7 arcminutes, has  $n \sim 1.5 \cdot 10^8$ . So for experiments which need both high sky coverage and high resolution, like Planck, a cheaper method is needed.

### 5.1.2 Pseudo- $C_l$ estimation

The idea of Pseudo- $C_l$  (PCL) power estimation as implemented in the MASTER algorithm [47], is to divide the estimation of  $C$  into two steps: A fast but biased estimator producing an initial estimate  $C'$ , and a correction step which compensates for the bias of that estimate. In general the true spectrum and the initial estimate can be related through

$$C' = f(C), \quad (5.15)$$

where the function  $f$  represents the effect of creating a realization of the spectrum and the effect of telescope optics, filtering, mapmaking and the biased power spectrum estimator. Most of these steps are linear in the map, and thus quadratic in power. Hence we cannot expect  $f$  to be a linear function. However, it is still possible to approximate it as such, and this approximation turns out to work well in practice,

$$\langle C'_b \rangle \approx \sum_{b'} \alpha_{bb'} C_{b'} + \beta_b. \quad (5.16)$$

Here  $b$  are the power spectrum multipole and polarization component bins, and  $\alpha_{bb'}$  and  $\beta_b$  are coefficients that approximate the effect of  $f$ . Provided estimates of  $\alpha$  and  $\beta$ , which can be determined using signal-only and noise-

only Monte Carlo simulations respectively<sup>5</sup> the unbiased power spectrum can be easily recovered by solving for  $C$  in equation (5.16):

$$\hat{C} = \alpha^{-1}(C' - \beta). \quad (5.17)$$

In principle, this expansion has  $N_{\text{bin}}(N_{\text{bin}} + 1)$  parameters that must be determined through simulations. However, most of the non-diagonal structure in  $\alpha$  is due to limited sky coverage, an effect which can be calculated analytically in the case where  $C'$  is the pseudo-spectrum of the map. The pseudo-spectrum is defined as the observed spectrum of the weighted map  $\{w_i m_i\}$ . The weights are usually chosen to be  $w_i = \sigma_i^{-2}$ , where  $\sigma_i$  is the noise level in pixel  $i$ , and their effect is to make the map well-behaved (but biased) on the whole sky, allowing equations (5.3-5.5 and 5.7) to be used. The result is the parametrisation

$$\alpha_{bb'} = \sum_{lXYX'Y'} P_{bl}^{XY} M_{ll'}^{XX'YY'} b_{l'}^{X'} b_{l'}^{Y'} F_{l'} Q_{l'b'}^{X'Y'}. \quad (5.18)$$

$M$  here models the effect of the partial sky coverage, and is a function only of the PCL map weighting  $w$ ,  $b$  is the instrument beam, and  $F$  is the transfer function, which incorporates the effect of biased map-making, filters, correlated noise, etc. The task of estimating  $\alpha$  is thus reduced to that of estimating the transfer function.

The uncertainty of the resulting power spectrum estimate can be determined using signal + noise Monte Carlo simulations based on a smooth version of the output spectrum. Given a set of such simulations  $\{C_b^i\}$ , the covariance of the estimator is

$$\text{Cov}_{bb'} = \langle (C_b^i - \hat{C}_b)(C_{b'}^i - \hat{C}_{b'}) \rangle. \quad (5.19)$$

The MASTER algorithm sketched above has become the most popular of power spectrum estimators currently in use, and is the only of these that is currently capable of handling the high number of pixels involved in high-resolution full-sky experiments like Planck. It also has the advantage of being very flexible – any effect that can be efficiently simulated is automatically taken care of by the transfer function, making it easy to e.g. quantify the effects of new systematics as they are discovered. And due to not needing unbiased maps and their covariance matrix like the brute force method does, its memory requirements are also much smaller.

This convenience comes at the price of being sub-optimal, due in part to not taking the full noise properties into account when weighting the map,

---

<sup>5</sup>Equation (5.16) is only approximate, which means that it matters slightly which input spectrum is used for the simulations. This spectrum should therefore be a reasonable estimate of the true spectrum. To protect against biasing oneself towards the fiducial spectrum, the process can be iterated by starting over with the output spectrum as the new fiducial spectrum.



and approximate, because  $f(C)$  is not truly linear. The practical result of this is that the uncertainty of the result is larger than for the brute force method, and that the error bars are slightly inaccurate.

### 5.1.3 Gibbs sampling

It is possible to have both the optimality and exactness of the brute force estimator while avoiding its  $\mathcal{O}(n^3)$  complexity. As in the PCL approach above, the key to efficiency is to find a way to use the fast full-sky expressions in equations (5.3-5.5 and 5.7), as these benefit from the speed of  $\mathcal{O}(N_{\text{pix}}^{\frac{3}{2}})$  full-sky spherical harmonics transformations. The PCL approach to this is to fill in the missing parts of the sky with zeroes<sup>6</sup>, a choice which makes it necessary to bias-correct the spectrum later. Gibbs sampling, on the other hand, makes it possible to use the full-sky expressions by producing full-sky samples of what the noiseless CMB could look like, which are then amenable to efficient spherical harmonics transformations [51, 52, 53].

Mathematically, this is based on the observation that

$$P(a, b, \dots) = P(a|b, \dots)P(b, \dots) \propto P(a|b, \dots), \quad (5.20)$$

which implies that given some sample  $(a^i, b^i, \dots)$  from the joint distribution of  $(a, b, \dots)$ , we can produce a new sample from the joint distribution by drawing a new  $a$  from the *conditional* distribution  $P(a|b, \dots)$  and keeping the others constant, producing  $(a^{i+1}, b^i, \dots)$ . This can then be repeated for each other parameter in turn. In the context of power spectrum estimation, this allows us to jointly sample the power spectrum  $C$  and the CMB sky  $s$  from  $P(C, s|m, M)$  through the iterative scheme

$$\begin{aligned} s^{i+1} &\leftarrow P(s|C^i, m, M) \\ C^{i+1} &\leftarrow P(C|s_{i+1}, m, M). \end{aligned} \quad (5.21)$$

The sky sampling step constructs a constrained realization of  $C$  given the observed sky  $m$  and its covariance  $M'$ . Such a realization can be economically produced by solving the system

$$(S^{-1} + M^{-1})s = M^{-1}m + S^{-\frac{1}{2}}\omega_1 + M^{-\frac{1}{2}}\omega_2 \quad (5.22)$$

through Conjugate Gradients iteration [51, 32].  $\omega_1$  and  $\omega_2$  are here vectors of independent samples from the standard normal distribution  $N(0, 1)$ . Conceptually, this amounts to Wiener-filtering  $m$ , and then adding enough fluctuations to compensate for the bias in power introduced by the wiener filter, resulting in a map with covariance  $S$  as required.

---

<sup>6</sup>Unobserved regions of the sky have infinite variance, making  $w = \sigma^{-2} = 0$ .

The resulting map  $s$  is a representative sample of the full, noise-less sky – the full effect of the noise in  $m$  is taken into account through the scatter in  $\{s\}$ , and since subsequent steps are conditional on a given  $s$ , they do not have to worry about the observed map or its noise at all. The second step of the sampling,  $C \leftarrow P(C|s, m, M) = P(C|s)$ , is therefore straightforward. Equations (5.3-5.5 and 5.7) apply directly to  $s$ , allowing us to efficiently sample a new power spectrum within the cosmic variance of the observed map.

By sampling only from the conditional distributions, Gibbs sampling allows a form of separation of responsibilities, letting each step in the sampling chain care about a single effect only. This makes it easy to expand the procedure with new effects, such as foreground separation.

### Parameter estimation

The set of samples  $\{C\}$  contain all the information we have about the power spectrum, together representing  $P(C|\text{data})$ . We can use this for cosmological parameter estimation by testing the power spectrum given by a particular choice of parameters,  $C(\theta)$ , against  $P(C|\text{data})$ . However, this is complicated by  $P(C|\text{data})$  only being available as a set of samples, requiring some binning and interpolation scheme, such as the Blackwell-Rao estimator [54], to be able to evaluate the probability at the required point.

However, as with brute force likelihood evaluation, it is possible to replace the power spectrum  $C$  with the cosmological parameters  $\theta$  as independent variables. This results in the scheme

$$\begin{aligned} s^{i+1} &\leftarrow P(s|\theta_i, m, M) \\ \theta^{i+1} &\leftarrow P(\theta|s_{i+1}, m, M). \end{aligned} \tag{5.23}$$

The result is a set of samples  $\{\theta, s\}$  that directly provide the probability of the parameters  $P(\theta|\text{data})$ . As  $C$  is a function of  $\theta$ , the first step is unchanged from above. The second step differs in that we do not have a simple expression for  $\theta(C)$ , preventing us from simply drawing a  $C$  from the distribution in equation (5.7), and translating it into a set of parameters. But we can combine the probability density function of said Wishart distribution with a short MCMC chain in  $\theta$  to still sample  $\theta \leftarrow P(\theta|s)$  at relatively low cost<sup>7</sup>.

---

<sup>7</sup>This is a planned feature in an upcoming version of H. K. K. Eriksen's cosmological Gibbs sampler *Commander* [55]. Long correlation lengths when sampling high S/N variables (like most cosmological parameters) was a barrier to direct cosmological parameter estimation in Gibbs sampling previously, but this was solved in [52].

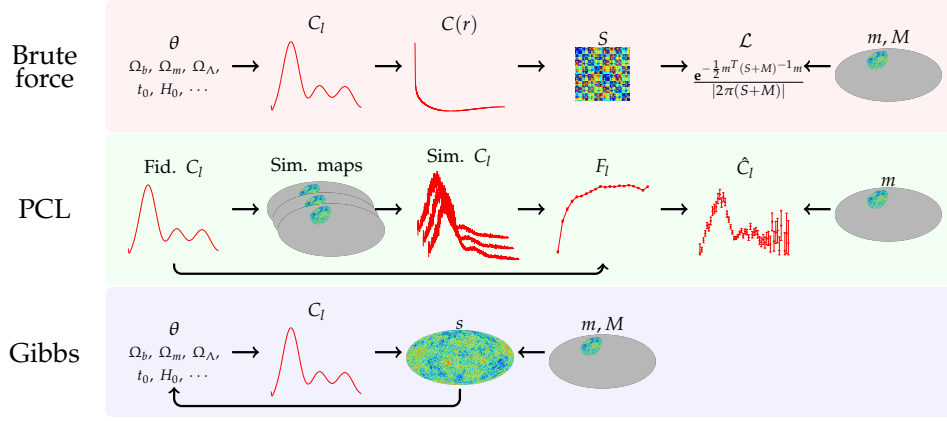


Figure 5.1: A sketch of the three power spectrum estimation methods described in this chapter. *Top*: Brute force maximum likelihood estimation. The likelihood of a set of cosmological parameters is computed by translating them into a signal covariance matrix  $S$  via the power spectrum  $C_l$  and the two-point correlation function  $C(r)$ . The observed map  $m$  should then be  $m \leftarrow N(0, S + M)$ , which is evaluated directly in pixel space. *Middle*: The MASTER pseudo- $C_l$  estimator works by using simulations to compute a transfer function  $F_l$  (and a noise bias, which is neglected here for simplicity) which accounts for the bias introduced by the map-making and filtering. This is then used to derive an unbiased spectrum from the data, and the simulations also provide error bars. *Bottom*: Gibbs sampling jointly samples  $P(\theta, C, s | m, M)$  by iteratively sampling over the conditional distributions. Only the sky sampling step actually involves the data.

## 5.2 Application to QUIET

As described in section 2, the QUIET telescope performs deep scans of relatively small parts of the sky, resulting in a relatively low number of pixels in each patch, typically  $\sim 3 \cdot 10^4$ , and twice that number of degrees of freedom in the case of a polarization only analysis. It is this low number that allows us to use the maximum likelihood map-making described in chapter 4 to obtain both an unbiased map  $m$  and its noise covariance  $M$ . This makes it possible in theory to use any of the three power spectrum/parameter estimation methods described in the previous section, but our limited number of pixels makes the brute force likelihood evaluation a natural choice, having computational costs comparable to the map-making algorithm.

The average area of each of QUIET's four CMB patches is  $\sim 250 \text{ deg}^2$  and this limited sky coverage sets up strong correlations in nearby multipoles. We therefore bin the spectra in bins of 50 multipoles (see section 5.1.1), which makes the bins mostly independent (see figure 5.3). This results in 9 bins for the Q-band analysis and 19 bins for the W-band analysis. Since these bins are only weakly correlated, a good approximation of the likelihood is provided by the conditional distributions for each bin, where the other bins are held fixed at the maximum likelihood point as found through Newton-Raphson iteration. For each frequency, there are 6 spectra to evaluate (TT, TE, TB, EE, EB and BB), resulting in a total of 54 (114) slices being needed for Q (W). Mapping out these slices required a few thousand likelihood evaluations, which is expensive but still less costly than the map-making step.

Based on each slice, we extracted a 68% confidence interval by lowering a likelihood threshold until 68% of the curve is above this likelihood, as illustrated in figure 5.2. The mode as well as the upper and lower limits found this way are what we report as the value and error bars for each bin.

For the tensor-to-scalar ratio we model the power spectrum as

$$C_l = C_l^{\text{scalar}} + C_l^{\text{tensor}} r, \quad (5.24)$$

where we have assumed for simplicity that tensor- and scalar perturbations have the same primordial spectral behavior.  $C_l^{\text{scalar}}$  is here the power spectrum in the absence of tensor modes, and  $C_l^{\text{tensor}}$  is the tensor-only spectrum for the case  $r = 1$ , as provided by CAMB [40]. If we know  $C_l^{\text{scalar}}$ , we can calculate  $C_l$  for each  $r$  and evaluate its likelihood using equation (5.10). But in practice,  $C_l^{\text{scalar}}$  is among the quantities we are trying to estimate, and must therefore be marginalized over to reflect our ignorance of it. In practice, this would involve an integral in as many dimensions as there are bins in our power spectrum, which is unfeasible with brute force evaluation<sup>8</sup>.

<sup>8</sup>It would, however, be an automatic by-product if we used Gibbs sampling. This is one of the major advantages of Gibbs sampling.




Figure 5.2: Illustration of how the reported value and error bars for a power spectrum bin are determined. The red curve shows the conditional distribution for a single bin of the W-band EE spectrum, which we take as an approximation of the marginal distribution for that bin. The reported value for the power in this bin is the mode of the distribution, while the upper and lower error bars are found by lowering a likelihood threshold (shown in blue) until 68% of the likelihood (green shaded area) is above the threshold.




Figure 5.3: The bins used in the Q-band power spectrum analysis. Each bin is 50 multipoles wide, and are centered on every multiple of 50 between 0 and 500. The horizontal bars show the location and extent of each bin, while the curves show which multipoles contribute to each bin. A bin width of 50 is enough to ensure a low level of correlation between the bins. In the W-band analysis, the same bin size is used, but the  $l_{\max}$  used here is 1000. This figure was computed for the PCL pipeline, but the level of correlation between the bins is mostly due to the shape of the sky cut, which is the same for both pipelines.

We therefore approximate the marginal with a partial conditional:

$$C_l^{\text{scalar}} \approx C_l^{\text{fix}} + C_l^{\text{marg}}. \quad (5.25)$$

Here the spectrum has been split into a part which will be held fixed at the maximum likelihood value, and a part which will still be marginalized over. We chose  $C_l^{\text{fix}}$  to be the multipoles  $l > 25$  and  $C_l^{\text{marg}}$  to be a single bin containing the multipoles  $l \leq 25$ , which is enough to make the marginalization manageable.

Though the tensor-to-scalar ratio must necessarily be positive (as it is the ratio of two positive quantities), the properties of the CMB maps only depends on the combined power spectrum, and the likelihood above therefore assigns a nonzero value even to clearly unphysical negative values. We can apply knowledge that  $r \geq 0$  as a prior  $P(r) \propto \Theta(r)$ , where  $\Theta(x)$  is the Heaviside step function, which results in a posterior with these unphysical values removed.

### 5.3 Results

As described in section 2, QUIET is sensitive to both the polarization and overall intensity of the incoming radiation, but due to its focus on measuring B-modes and the tensor-to-scalar ratio, the vast majority of its sensitivity and analysis effort is concentrated on the polarization results. In the absence of temperature information, the three available power spectra are EE, EB and BB, which can be seen in figures 5.4-5.6.

Of these, the EE spectrum is the most easily measured, and is sourced by the same density perturbations that source the TT spectrum. EE is there-

fore well constrained by the  $\Lambda$ CDM best fit to existing temperature experiments, such as WMAP [6]. The EE spectrum has also been directly measured by several polarization experiments such as BICEP [17] and QUAD [16], which were consistent with  $\Lambda$ CDM. As seen in figure 5.4, this is also the case for the QUIET Q-band and W-band EE power spectra. The agreement is especially striking in the W-band spectrum due to its small error bars, and to my knowledge this is the most precise measurement of the second peak of the EE power spectrum at this time.

The BB spectrum is expected to be produced by primordial gravitational waves at low multipoles and by lensing [13] and possibly higher-order effects like non-adiabaticity and magnetic fields at higher multipoles [14]. Neither of these have as yet been detected, and only upper bounds exist for the BB spectrum. As shown in figure 5.6, both QUIET data sets are consistent with a non-detection of B-modes. This was to be expected based BICEP and QUAD's null-results, as these experiments have similar sensitivity to QUIET. QUIET's 95% upper limit on BB power is compared with other current limits in figure 5.7. The QUIET W-band limits are seen to provide the lowest upper bounds in the multipole range  $150 \leq l \leq 400$ .

Like the BB spectrum, EB is expected to have a non-zero contribution from lensed E-modes and possible second-order effects, but is not sensitive to primordial gravitational waves. As seen in figure 5.5, QUIET observes an EB power consistent with zero.

### 5.3.1 Temperature spectra

Figure 5.8-5.10 show the QUIET TT, TE and TB power spectra. These include data from the temperature maps shown in section 4.8.1, which have relatively poor sensitivity compared to the polarization maps, and which have less well understood systematics. They are therefore not a part of the primary QUIET results, but they can still serve as a consistency check for the polarization results. Taken together, the three spectra are consistent with the  $\Lambda$ CDM expectations.

### 5.3.2 Foregrounds

In temperature (total power) emissions, the foregrounds are quite well known, consisting of synchrotron radiation from electrons spiraling in magnetic fields, free-free emission from electrons scattering off ions and dipole radiation and thermal radiation from dust. The spectral behavior of these is shown in figure 5.11.

The situation is less well known for polarization, but free-free and spinning dust are not expected to be significantly polarized. Synchrotron emission, on the other hand, is 3% – 30% polarized, while thermal dust is polarized at the 1% – 6% level, leaving these as the dominant foregrounds




Figure 5.4: The QUIET EE power spectrum based on data from all four CMB patches. The W-band results are indicated using “+”-symbols with 1 sigma (green) and 2 sigma (red) error bars, while the Q-band results use “x”-symbols with 1 sigma (purple) and 2 sigma (blue) statistical error bars. These are compared with the best-fit  $\Lambda$ CDM model based on the WMAP 7-year data (light blue). The overall amplitude of the QUIET W-band spectrum shown here is calibrated based on the WMAP EE spectrum. Hence, the  $\Lambda$ CDM best fit and the W-band spectrum necessarily agree on the overall amplitude. However, the shape of the power spectrum is independently measured by QUIET, and is consistent with  $\Lambda$ CDM. The two QUIET seasons are also consistent with each other. The first three peaks of the EE power spectrum are clearly visible.






Figure 5.5: The QUIET EB power spectrum ( $\frac{l(l+1)}{2\pi} C_l^{EB}$ ) based on data from all four CMB patches. The W-band results are indicated using “+”-symbols with 1 sigma (green) and 2 sigma (red) statistical error bars, while the Q-band results use “x”-symbols with 1 sigma (purple) and 2 sigma (blue) error bars. Due to the large magnitude difference between high and low multipoles, the graph is logarithmic – the upper and lower panels corresponding to positive and negative numbers respectively. The W-band gain is calibrated to WMAP via the EE power spectrum, greatly reducing the systematic errors, which would otherwise be  $\sim 18\%$ . The spectrum is consistent with the expectation of zero EB power except for a 2.7 sigma outlier in the W-band bin centered on  $l = 350$ . The probability of observing one or more such outliers in one of the  $(19 + 9) \cdot 3 = 84$  bins of the EE, EB and BB spectra is 44%.




Figure 5.6: The QUIET BB power spectrum based on data from all four CMB patches. The W-band results are indicated using “+”-symbols with 1 sigma (green) and 2 sigma (red) statistical error bars, while the Q-band results use “x”-symbols with 1 sigma (purple) and 2 sigma (blue) error bars. Due to the large magnitude difference between high and low multipoles, the graph is logarithmic – the upper and lower panels corresponding to positive and negative numbers respectively. The W-band gain is calibrated to WMAP via the EE power spectrum, greatly reducing the systematic errors, which would otherwise be  $\sim 17\%$ . The spectrum is consistent with zero BB power. Though BB cannot physically be negative, we are using the Gaussian error approximation here, which does allow this. Negative power should be interpreted as cases where the noise happened to fluctuate high. Though the truncated, positive part of the distribution is a useful approximation to the likelihood, we do not rely on this when estimating the tensor-to-scalar ratio, which is estimated from the exact likelihood.

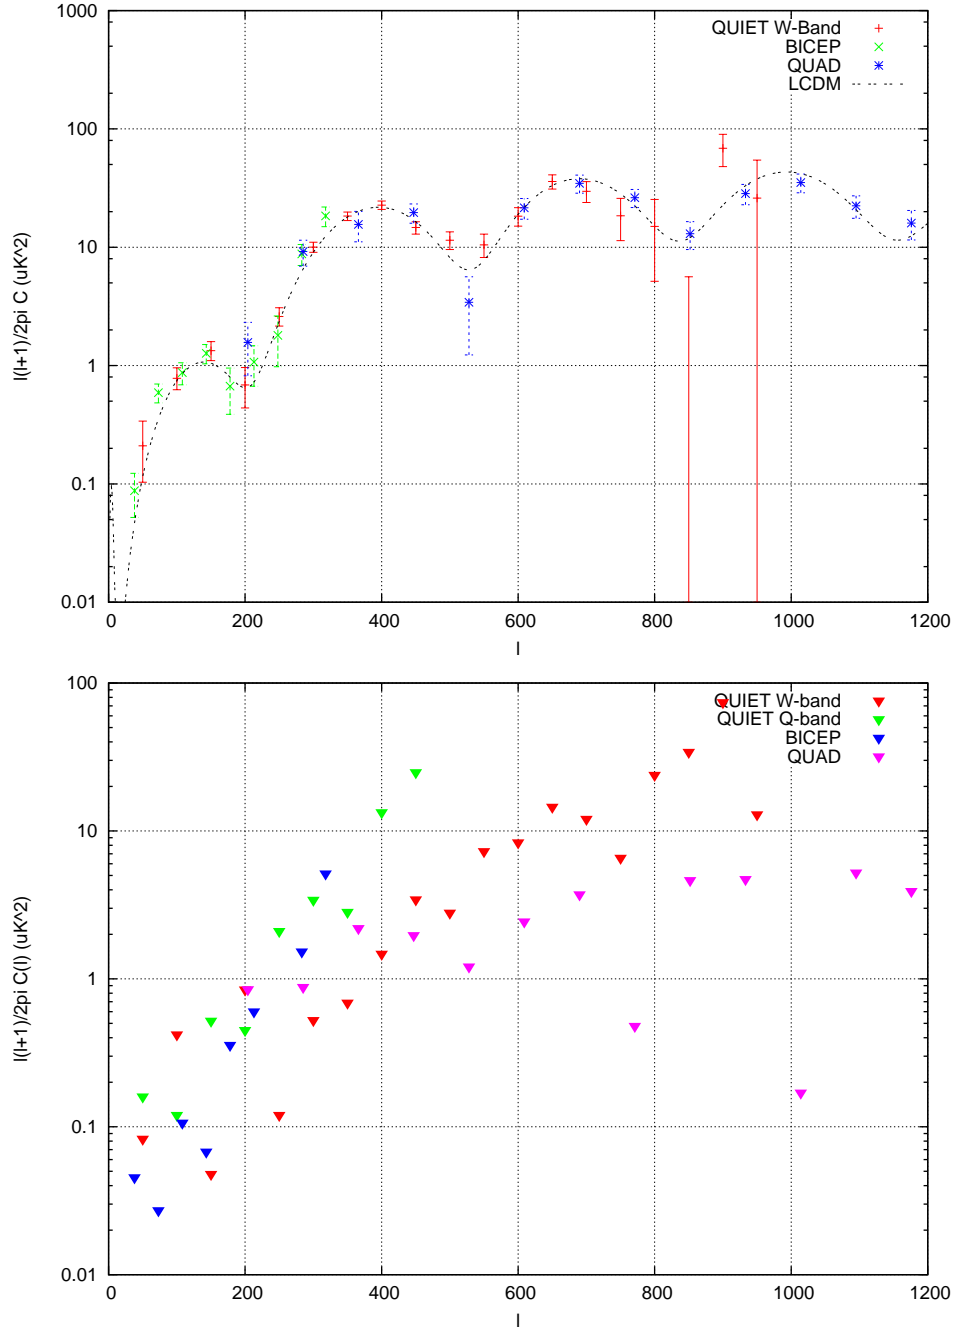


Figure 5.7: Top: The QUIET EE power spectrum compared with the two other most precise measurements to date, BICEP and QUAD, and with the WMAP 7-year best fit  $\Lambda$ CDM spectrum. Bottom: The QUIET BB power 2 sigma upper limits compared with BICEP and QUAD. In both cases, QUIET has the most precise measurements in the multipole range  $150 \leq l \leq 400$ .




Figure 5.8: The QUIET W-band (red) and Q-band (green) TT power spectrum based on data from all four CMB patches for Q-band, and CMB patch 1, 2 and 4 for W-band, compared to the WMAP 7-year best fit  $\Lambda$ CDM model. The temperature results are based on a less comprehensive null test suite than the polarization results, and have not been subject to the same detailed systematics study. Nevertheless, the agreement with  $\Lambda$ CDM is excellent. The  $l = 50$  bin in W-band and CMB patch 3 in Q-band failed the temperature null-tests, and are therefore not included here.




Figure 5.9: The QUIET Q-band TE power spectrum based on data from all four CMB patches, compared to the WMAP 7-year best-fit  $\Lambda$ CDM model. The temperature results are based on a less comprehensive null test suite than the polarization results, and have not been subject to the same detailed systematics study. The results are consistent with the expectation from  $\Lambda$ CDM. The spectrum prefers to shift the top near  $l = 300$  towards higher multipoles, but this is not significant. The W-band temperature cross-spectra were not yet ready at the time of writing.




Figure 5.10: The QUIET Q-band TB power spectrum based on data from all four CMB patches. The temperature results are based on a less comprehensive null test suite than the polarization results, and have not been subject to the same detailed systematics study. Based on our sensitivity, we expect TB to be consistent with zero. This is the case all bins but those at  $l = 150$  and  $l = 350$ , which are nonzero by 2.4 and 2.2 sigma based on the statistical errors. The probability of observing 2 outliers of  $> 2.2$  sigma among the 45 bins in the TT, TE and TB spectra is 35.7%, so this does not amount to a detection. Furthermore, the systematic errors were not included in the error bars, and would further reduce the significance.




Figure 5.11: The spectral behavior of the diffuse foregrounds in total power as a function of frequency. Blue: Synchrotron, green: free-free, pink: spinning dust, red: thermal dust. The horizontal line shows the amplitude of the CMB anisotropies. The corresponding graph for polarization is still poorly determined, but due to synchrotron's high polarization fraction (3% – 30%), it is expected to be the only relevant diffuse component until thermal dust takes over at  $\nu \sim 80\text{GHz}$ . This figure is taken from [56].

[56]. In comparison, the CMB itself is polarized at the 10% level, which is comparable to that of the foregrounds. The CMB polarization is however predominantly in the form of E-modes, and the B-mode amplitude may well be low enough that foregrounds pose a serious problem.

Figure 5.12 shows the expected foreground contamination for QUIET's four CMB patches in the W-band based on the Planck Sky Model [57]. According to this model, dust is by far the most important polarized foreground at 95 GHz, with an amplitude approaching the statistical error level in the  $l = 50$  bin for CMB patch 1, which is the most foreground-contaminated patch due to its proximity to the Milky Way. In the Q-band, the situation is reversed. With a spectral index of  $\sim -3$ , the synchrotron component has  $\sim (\frac{44\text{GHz}}{95\text{GHz}})^{-6} \sim 10^2$  times higher power, which is enough to make it detectable in the lowest multipole bin for CMB patch 1; indeed, the Q-band  $l = 50$  bin EE-power is measured to be  $0.55 \pm 0.14 \mu\text{K}^2$  for this patch, a  $3\sigma$  outlier which is consistent with the expected synchrotron contribution.

Aside from the first bin for CMB patch 1, we do not expect foregrounds to be significant at our sensitivity. However, if the tensor-to-scalar ratio is of the order  $r \lesssim 0.02$  a better understanding of the polarized foregrounds will be needed in order to disentangle these from the CMB.

### 5.3.3 Systematic errors

The error bars in figures 5.4-5.6 do not include the systematic errors, but a summary of these can be seen in figure 5.13. Most of these errors are much smaller than the statistical error, and have no impact on the results. However, the absolute gain uncertainty, which is by far the largest systematic error, is significantly larger than the statistical error in signal-dominated regions. Luckily, the effect of the absolute gain is simply to scale all the polarization spectra up or down by a constant factor, making this error completely correlated between all multipoles. Rather than add large, 100% correlated systematic error bars to the power spectra, which would make it hard to visually judge the impact of this systematic, we instead chose to factorize the power spectra as follows:

$$C_l = AC_l^{\text{norm}}. \quad (5.26)$$

Here  $C_l^{\text{norm}}$  is the spectrum with the amplitude normalized to a reference value, which we choose to be the WMAP 7-year best-fit spectrum, and  $A$  is the relative EE amplitude between  $\Lambda\text{CDM}$  and QUIET. Thus,  $A$  contains the amplitude degree of freedom of the spectrum and its relatively large systematic uncertainty, while  $C_l^{\text{norm}}$  holds all the remaining  $N_{\text{bin}}N_{\text{spect}} - 1$  degrees of freedom of the *shape* of the spectrum, which have only negligible systematics. The spectra in figures 5.4-5.6 are all of the  $C_l^{\text{norm}}$  type, and hence only contain relative amplitude information.



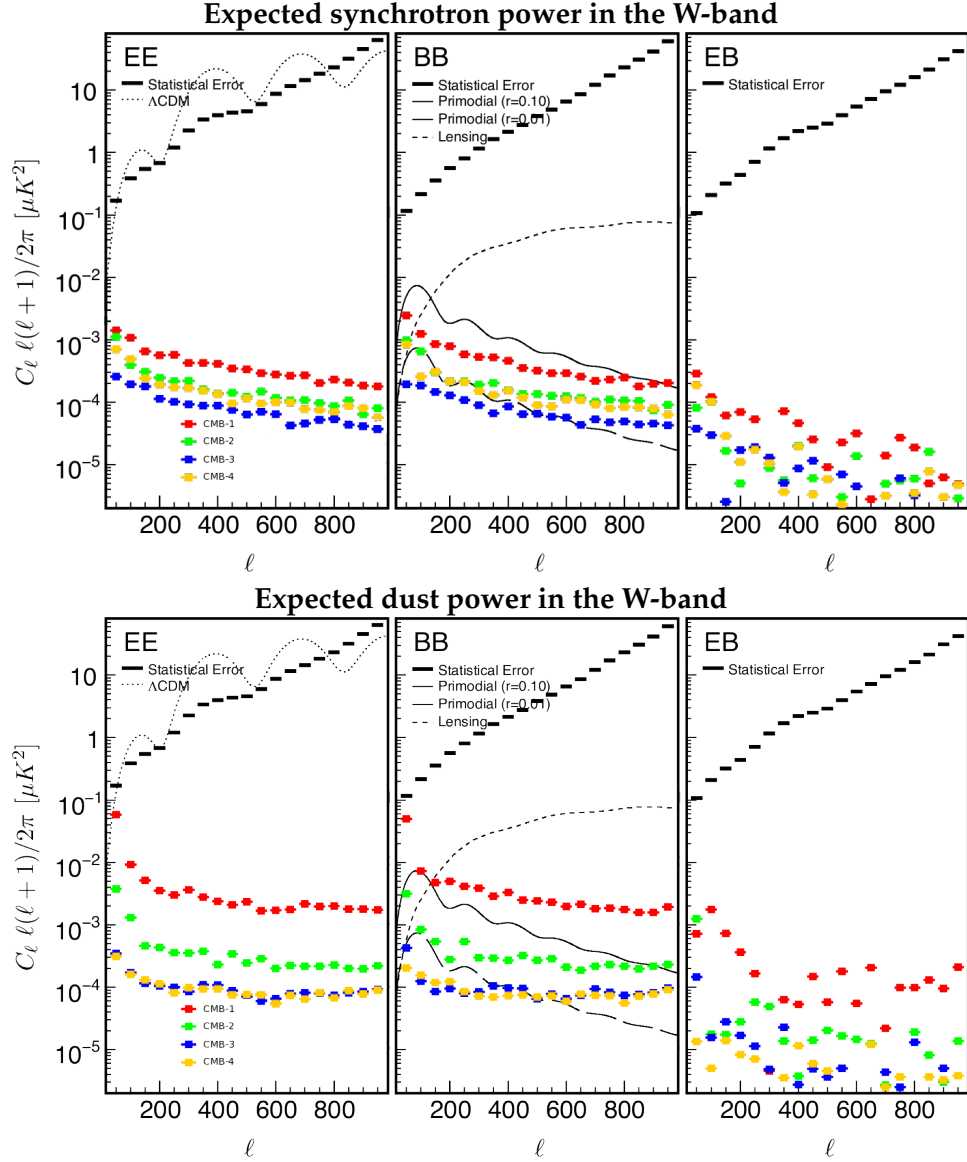


Figure 5.12: The expected power spectrum contamination from synchrotron (top) and dust (bottom) in the W-band, compared to QUIET's statistical error and the  $\Lambda$ CDM expectation for EE modes, lensing B-modes and primordial B-modes for  $r = 0.1$  and  $r = 0.01$ , based on the Planck foreground model v. 1.7.4. The graphs in this figure were produced by Osamu Tajimao as part of his systematics study for the W-band analysis.

The primary contributions to the absolute gain uncertainty are the uncertainty in the polarized intensity of Tau A, beam area, and gain variations between detectors. These are multiplicative errors, so  $A$  can be approximated as having a log-normal distribution, with  $A \sim \log \mathcal{N}(\log(1.0), 0.13)$  in the Q-band and  $A \sim \log \mathcal{N}(\log(1.35), 0.17)$  in the W-band (preliminary). The latter is high by 1.8 sigma compared to the expected value of 1, indicating that this preliminary 17% may be an underestimate.

To avoid propagating this uncertainty into the estimate of the tensor-to-scalar ratio  $r$ , we can instead consider  $A$ , as determined from the EE power spectrum, as a measurement of the absolute gain, and use this instead of the uncertain Tau A-derived estimate. In this case, the uncertainty of  $A$  is reduced to the statistical uncertainty from the EE power spectrum fit, which is only 3.7%. This is similar to the approach of calibrating the gain against the TT power spectrum, which is the standard approach for bolometer-based CMB polarization experiments.

#### 5.3.4 Parameters

The tensor-to-scalar ratio likelihood derived from the Q-band and the preliminary W-band analysis is shown in figure 5.14. Based on these, we find  $r = 0.35^{+1.06}_{-0.87}$  (Q-band) and  $r = 1.08^{+0.88}_{-0.78}$  (W-band), with 95% upper limits of 2.2 (Q) and 2.8 (W). Though the W-band data set has a slightly higher sensitivity than the Q-band set, each central value is noisy due to limited statistics, and this noise fluctuated towards higher values for the W-band analysis, resulting in a poorer limit on  $r$ .

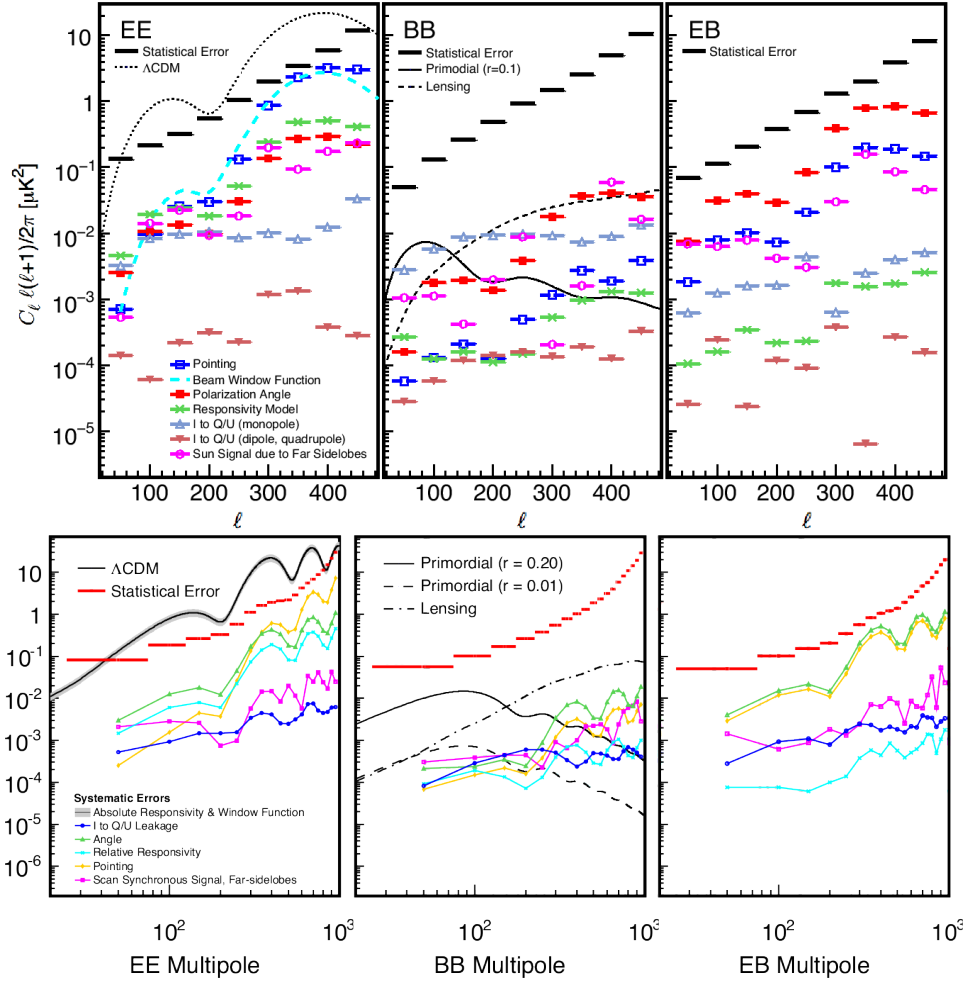


Figure 5.13: The Q-band (top) and W-band (bottom, preliminary) systematic effects compared to the statistical error and the expected signal. With the exception of the absolute gain, the systematic errors are all significantly smaller than the statistical errors. The absolute gain uncertainty is relatively large, but acts as a simple scaling of all spectra by  $\pm 18\%$ . The simplest way of ensuring that this does not affect the B-mode limits is to sacrifice a single EE degree of freedom by calibrating the EE spectrum against  $\Lambda\text{CDM}$ , since the same gain applies to both E and B. Figure courtesy of Osamu Tajima.




Figure 5.14: The Q- and W-band likelihood for the tensor-to-scalar ratio  $r$ . While the W-band measurement is slightly more sensitive than that for the Q-band, the latter results in a lower limit on  $r$  because it happened to fluctuate lower. Both graphs extend into the area of unphysical negative tensor-to-scalar ratios. These regions should be excluded by the prior, but plotting the likelihood directly with these negative regions included makes it easy to see at a glance that these measurements are consistent with zero.

## Chapter 6

# Non-gaussianity

The power spectrum and parameter estimation in the previous chapter is all based on the assumption that the CMB is a Gaussian random field. Theoretically, the primordial perturbations derive from quantum fluctuations during inflation, and these are expected to be in the ground state of the harmonic oscillator, which is a Gaussian [58, 1]. The subsequent evolution towards the surface of last scattering preserves this to linear order in perturbation theory, which is still a good approximation, resulting in a CMB which should be very close to Gaussian<sup>1</sup>.

However, this conclusion can change based on the model of inflation and the importance of second and higher order effects during the subsequent evolution of the universe, neither of which are settled issues [60, 61]. Testing the Gaussianity of the CMB provides an indirect way of probing these issues, and has become an active field of research [62, 63, 59]. As there are infinitely many ways of being non-Gaussian, a large set of tests have been suggested. These can be divided into targeted tests, which test for a specific kind of non-Gaussianity<sup>2</sup>, and general tests, which can find any kind of non-Gaussianity, at the cost of being less sensitive. Examples of the latter are the Kolmogorov-Smirnow test, the Cramér-von Mises statistic and the Anderson-Darling test<sup>3</sup>, which all work by comparing the empirical cumulative distribution function with the theoretically expected one.

In 2010 Gurzadyan et al. [65] published the result of the application of the Kolmogorov-Smirnov test to look for non-Gaussianity in the CMB. Their surprising result was that only 20%(!) of the CMB anisotropies behave as a random Gaussian field, with the remaining 80% being “non-

---

<sup>1</sup>During its travel from the surface of last scattering, small levels of non-Gaussianity enters the CMB due to various effects such as inhomogeneous recombination and gravitational lensing. The latter has been recently been detected [59].

<sup>2</sup>These can be tests for a particular kind of non-Gaussianity predicted by a specific theory, or can be more generic such as testing the higher moments of the empirical distribution.

<sup>3</sup>When applied as a test for normality, the Anderson-Darling test is usually the most sensitive of these [64], but the Kolmogorov-Smirnov test is still popular.




Figure 6.1: Illustration of the Kolmogorov statistic  $K$ , which is the maximal difference between the empirical and theoretical cumulative distribution functions.  $K$  is the maximal vertical difference marked in blue.

random”<sup>4</sup>. This analysis later became the basis of a highly publicised claim for evidence for the Conformal Cyclic Cosmology hypothesis [66], which in turn prompted three independent follow-up studies [67, 68, 69]. Neither of these follow-ups found any significant detection. Subsequent correspondence<sup>5</sup> made it clear that the reason for the difference in significance was due to Gurzadyan and Penrose [66] using simulations with much lower variance than the standard model due to the claim in Gurzadyan et al. [65] that the CMB is only “20% random”, a claim that had gone unnoticed by most of the community, but which would be revolutionary if true.

To test the “20% random” claim, I repeated the analysis of Gurzadyan et al. [65] in paper IV of this thesis. Here the Kolmogorov-Smirnov (K-S) test was applied to  $10^4$  randomly chosen  $1.5^\circ$  radius disks from the WMAP 7-year W-band temperature map after excluding regions of galactic latitude  $|b| < 30^\circ$ .

According to the Kolmogorov theorem, the maximal difference  $K$  between the theoretical and empirical cumulative distribution functions of a set of *independently identically distributed* samples will in the limit of in-

<sup>4</sup>What they mean by this term is not defined in the article, but based on their simulations in later articles, they appear to mean that 80% of the CMB would be the same between different cosmic realizations.

<sup>5</sup>This rapid exchange, with 8 articles following each other in rapid succession in response to each other, all in the space of half a year, and the first 5 within one month, is an example of how preprint-services can drastically reduce the scientific turn-around time compared to traditional journals.

finitely many samples be distributed according to [64]:

$$P(x < K) = F_{\text{KS}}(\sqrt{N_{\text{samp}}}K) \quad (6.1)$$

$$F_{\text{KS}}(x) = 1 - 2 \sum_{i=1}^{\infty} (-1)^{i-1} e^{-2i^2 x^2}. \quad (6.2)$$

This relation can then be used to test the hypothesis that a set of samples were drawn from a given distribution. However, the pixels  $s_i$  within the  $1.5^\circ$  CMB disks do not fulfil the criterion of being independently identically distributed due to the CMB anisotropies' correlatedness<sup>6</sup>, which are described in pixel space by the covariance  $S_{ij} + N_{ij}$  as provided by equation (5.8) and the noise properties of the experiment. But it is possible to *whiten*  $s$  and thus make it eligible for the test: The vector  $s' \equiv S^{-\frac{1}{2}}s$  will follow the distribution  $N(0, 1)$ , and thus be independently identically distributed if  $s \sim N(0, S + N)$ .

After applying the K-S test to the  $s'$  of all  $10^4$  disks with  $N(0, 1)$  as the theoretical distribution I found a result fully consistent with the CMB anisotropies being fully described as a random Gaussian distribution with an angular power spectrum  $C_l$  given by the  $\Lambda$ CDM best fit to the WMAP 7-year data. This result is clearly incompatible with Gurzadyan et al. [65], where the K-S test consistently failed. I was, however, able to reproduce their K-S failure by ignoring the correlations, and instead testing  $s$  against  $N(\mu, \sigma^2)$ , where  $\mu$  and  $\sigma$  are the measured mean and standard deviation of  $s$  respectively. This strongly suggests that Gurzadyan et al. [65]'s claim of a "weakly random" CMB stems from incorrectly ignoring the covariance of the CMB.

---

<sup>6</sup>Ignoring these correlations amounts to assuming a flat angular power spectrum.





## Chapter 7

# Summary and outlook

CMB B-modes provide the most promising avenue for detection of the primordial tensor fluctuations predicted by inflation. Such a detection would in addition to validating inflation itself constitute a probe of physics at far higher energies than those reachable in particle colliders. So far, no experiment has been able to detect B-modes, which implies that the tensor-to-scalar ratio must be less than 0.7. This limit is still worse than the limit of  $r < 0.2$  from T and E-modes, but this is likely to change with future B-mode experiments.

One of the current experiments aiming to measure B-modes is QUIET. QUIET's defining feature is its use of large, compact arrays of coherent amplifier detectors. Until recently, such arrays of microwave polarimeters were only practical with bolometers. However, a recent breakthrough has allowed the miniaturization of coherent amplifier detectors to a polarimeter-on-a-chip format, and QUIET is the first deployment of these in the field. In addition to the measurement of B-modes itself, one of QUIET's goals is therefore to demonstrate the viability of these detectors for high-precision CMB measurements.

After deploying in 2008, QUIET observed in the Q- and W-bands until the end of 2010, resulting in measurements of the CMB E- and B-modes in the multipole range  $25 < l < 1000$ . The W-band measurements provide world record sensitivity between  $l = 150$  and  $l = 400$ . QUIET did not observe any CMB B-modes, resulting in limits of  $r < 2.2$  (Q) and  $r < 2.8$  (W) at 95% confidence. While these are wider than the best limits available, QUIET's results are demonstrated to have very low systematic errors, corresponding to  $r < 0.1$ .

A first step towards this sensitivity will be taken by combining the Q- and W-band data sets into a single estimate on  $r$ , which will be performed in a future article. This will provide an approximate factor of 2 improvement. To go beyond this, more observations are needed. With the viability of the detectors validated, plans are underway for a full deployment of this

technology as QUIET-2.

QUIET 2 would include 500-1000 detector modules, for a total of 2000-4000 concurrent polarization measurements. The majority of these will be W-band detectors with significantly improved sensitivity compared to those used in the current QUIET experiment. This will allow QUIET-2 to reach statistical errors of  $r \lesssim 0.01$  or better, which is enough to rule out all large-field models of inflation and some small-field models [7].

However, to reach this level, it is clear that it will no longer be possible to ignore foregrounds. Polarized component separation will be needed, and that requires sensitive polarization maps at multiple frequencies and a better understanding of the behavior of polarized foregrounds. QUIET 2 will include detectors sensitive in the Ka-low, Ka-high and W-band, which in principle is enough to disentangle CMB, synchrotron and dust. Additionally, Q-band measurements in the same patches will be provided by the FOCUS experiment, a planned re-deployment of the existing QUIET Q-band array, which is currently the most sensitive array in this frequency band.<sup>1</sup>

Though foregrounds are a nuisance for CMB measurements, they are interesting from an astrophysical perspective. The current QUIET experiment includes two foreground-dominated patches in the galactic plane. The maps of these patches, which represent some of the most sensitive microwave polarization measurements of the galaxy, will be the subject of a future QUIET foreground paper.

A strength of the maximum likelihood analysis pipeline described in this thesis is its production of unbiased maps with fully quantified statistical properties. This makes the maps useful in their own right instead of simply being an intermediate step before power spectrum estimation. The maximum likelihood pipeline will therefore be central in the QUIET foreground analysis both for the current QUIET and for the future QUIET-2.

---

<sup>1</sup>Other upcoming experiments will also improve our knowledge of foregrounds:

- C-Bass is currently making polarization maps of the full sky in the C-band (5 GHz). These will essentially be 100% synchrotron radiation with very high signal-to-noise due to synchrotron's steep spectral index.
- Planck will produce full-sky polarization maps at a large number of frequencies between 30 GHz and 857 GHz. While these will have low S/N for the CMB polarizations, Planck will have good S/N for dust and synchrotron.

## Chapter 8

# Bibliography

- [1] Scott Dodelson. *Modern Cosmology*. Academic Press, 2002.
- [2] G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *ApJL*, 396:L1–L5, September 1992.
- [3] C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, E. Wollack, E. L. Wright, C. Barnes, M. R. Greason, R. S. Hill, E. Komatsu, M. R. Nolta, N. Odegard, H. V. Peiris, L. Verde, and J. L. Weiland. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results. *ApJS*, 148:1–27, September 2003, arXiv:astro-ph/0302207.
- [4] M. R. Nolta, J. Dunkley, R. S. Hill, G. Hinshaw, E. Komatsu, D. Larson, L. Page, D. N. Spergel, C. L. Bennett, B. Gold, N. Jarosik, N. Odegard, J. L. Weiland, E. Wollack, M. Halpern, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Angular Power Spectra. *ApJS*, 180: 296–305, February 2009, 0803.0593.
- [5] N. Jarosik, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill, G. Hinshaw, A. Kogut, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. *ApJS*, 192:14, February 2011, 1001.4744.

- [6] D. Larson, J. Dunkley, G. Hinshaw, E. Komatsu, M. R. Nolta, C. L. Bennett, B. Gold, M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year wilkinson microwave anisotropy probe (wmap) observations: Power spectra and wmap-derived parameters. *The Astrophysical Journal Supplement Series*, 192(2):16, 2011. URL <http://stacks.iop.org/0067-0049/192/i=2/a=16>.
- [7] William H Kinney and Antonio Riotto. Theoretical uncertainties in inflationary predictions. *Journal of Cosmology and Astroparticle Physics*, 2006(03):011, 2006, astro-ph/0511127v2. URL <http://stacks.iop.org/1475-7516/2006/i=03/a=011>.
- [8] E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year wilkinson microwave anisotropy probe (wmap) observations: Cosmological interpretation. *The Astrophysical Journal Supplement Series*, 192(2):18, 2011, 1001.4538v3. URL <http://stacks.iop.org/0067-0049/192/i=2/a=18>.
- [9] I. S. Grant and W. R. Phillips. *Electromagnetism*. Wiley, 2 edition, 1990.
- [10] Matias Zaldarriaga. Nature of the  $e - b$  decomposition of cmb polarization. *Phys. Rev. D*, 64:103001, Oct 2001, astro-ph/0106174v1. URL <http://link.aps.org/doi/10.1103/PhysRevD.64.103001>.
- [11] Matias Zaldarriaga and U. Seljak. All-sky analysis of polarization in the microwave background. *Phys. Rev. D*, 55:1830–1840, Feb 1997. URL <http://link.aps.org/doi/10.1103/PhysRevD.55.1830>.
- [12] Wayne Hu and Martin White. A cmb polarization primer. *New Astronomy*, 2(4):323 – 344, 1997, astro-ph/9706147v1. ISSN 1384-1076. URL <http://www.sciencedirect.com/science/article/pii/S1384107697000225>.
- [13] Matias Zaldarriaga and U. Seljak. Gravitational lensing effect on cosmic microwave background polarization. *Phys. Rev. D*, 58:023003, Jun 1998, astro-ph/9803150v1. URL <http://link.aps.org/doi/10.1103/PhysRevD.58.023003>.
- [14] A. Lewis. CMB anisotropies from primordial inhomogeneous magnetic fields. *Phys.rev.D*, 70(4):043011, August 2004, arXiv:astro-ph/0406096.

- [15] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel. Detection of polarization in the cosmic microwave background using dasi. *Nature*, 420:772 – 787, 12 2002, astro-ph/0209478v1. URL <http://dx.doi.org/10.1038/nature01269>.
- [16] M. L. Brown, P. Ade, J. Bock, M. Bowden, G. Cahill, P. G. Castro, S. Church, T. Culverhouse, R. B. Friedman, K. Ganga, W. K. Gear, S. Gupta, J. Hinderks, J. Kovac, A. E. Lange, E. Leitch, S. J. Melhuish, Y. Memari, J. A. Murphy, A. Orlando, C. O’ Sullivan, L. Piccirillo, C. Pryke, N. Rajguru, B. Rusholme, R. Schwarz, A. N. Taylor, K. L. Thompson, A. H. Turner, E. Y. S. Wu, M. Zemcov, and The QUaD collaboration. Improved measurements of the temperature and polarization of the cosmic microwave background from quad. *The Astrophysical Journal*, 705(1):978, 2009, 0906.1003v3. URL <http://stacks.iop.org/0004-637X/705/i=1/a=978>.
- [17] H. C. Chiang, P. A. R. Ade, D. Barkats, J. O. Battle, E. M. Bierman, J. J. Bock, C. D. Dowell, L. Duband, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, W. C. Jones, B. G. Keating, J. M. Kovac, C. L. Kuo, A. E. Lange, E. M. Leitch, P. V. Mason, T. Matsumura, H. T. Nguyen, N. Ponthieu, C. Pryke, S. Richter, G. Rocha, C. Sheehy, Y. D. Takahashi, J. E. Tolan, and K. W. Yoon. Measurement of cosmic microwave background polarization power spectra from two years of bicep data. *The Astrophysical Journal*, 711(2):1123, 2010, 0906.1181v3. URL <http://stacks.iop.org/0004-637X/711/i=2/a=1123>.
- [18] C R Lawrence, S Church, T Gaier, R Lai, C Ruf, and E Wollack. Coherent detectors. *Journal of Physics: Conference Series*, 155(1):012002, 2009. URL <http://stacks.iop.org/1742-6596/155/i=1/a=012002>.
- [19] I. Buder. *Measurement of the CMB Polarization at 95 GHz from QUIET*. PhD thesis, University of Chicago, 2012.
- [20] Y. Chinone. *Measurement of Cosmic Microwave Background Polarization Power Spectra at 43 GHz with Q/U Imaging Experiment*. PhD thesis, Tohoku University, 2010.
- [21] Y. Mizuguchi, M. Akagawa, and H. Yokoi. Offset dual reflector antenna. volume 14 of *IEEE Antennas And Propagation Society International Symposium*, pages 2–5, 1976.
- [22] C. Dragone. Offset multireflector antennas with perfect pattern symmetry and polarization discrimination. *Bell syst. tech. J.*, 57:2663–2684, 1978.
- [23] T. Essinger-Hileman, J. W. Appel, J. A. Beall, H. M. Cho, J. Fowler, M. Halpern, M. Hasselfield, K. D. Irwin, T. A. Marriage, M. D.

- Niemack, L. Page, L. P. Parker, S. Pufu, S. T. Staggs, O. Stryzak, C. Visnjic, K. W. Yoon, and Y. Zhao. The Atacama B-Mode Search: CMB Polarimetry with Transition-Edge-Sensor Bolometers. *ArXiv e-prints*, August 2010, 1008.3915.
- [24] B. Reichborn-Kjennerud, A. M. Aboobaker, P. Ade, F. Aubin, C. Bacigalupi, C. Bao, J. Borrill, C. Cantalupo, D. Chapman, J. Didier, M. Dobbs, J. Grain, W. Grainger, S. Hanany, S. Hillbrand, J. Hubmayr, A. Jaffe, B. Johnson, T. Jones, T. Kisner, J. Klein, A. Korotkov, S. Leach, A. Lee, L. Levinson, M. Limon, K. MacDermid, T. Matsumura, X. Meng, A. Miller, M. Milligan, E. Pascale, D. Polsgrove, N. Ponthieu, K. Raach, I. Sagiv, G. Smecher, F. Stivoli, R. Stompor, H. Tran, M. Tristram, G. S. Tucker, Y. Vinokurov, A. Yadav, M. Zal-darriaga, and K. Zilic. EBEX: a balloon-borne CMB polarization experiment. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010, 1007.3672.
- [25] A. A. Fraisse, P. A. R. Ade, M. Amiri, S. J. Benton, J. J. Bock, J. R. Bond, J. A. Bonetti, S. Bryan, B. Burger, H. C. Chiang, C. N. Clark, C. R. Contaldi, B. P. Crill, G. Davis, O. Doré, M. Farhang, J. P. Filippini, L. M. Fissel, N. N. Gandilo, S. Golwala, J. E. Gudmundsson, M. Hasselfield, G. Hilton, W. Holmes, V. V. Hristov, K. Irwin, W. C. Jones, C. L. Kuo, C. J. MacTavish, P. V. Mason, T. E. Montroy, T. A. Morford, C. B. Netterfield, D. T. O'Dea, A. S. Rahlin, C. Reintsema, J. E. Ruhl, M. C. Runyan, M. A. Schenker, J. A. Shariff, J. D. Soler, A. Trangsrud, C. Tucker, R. S. Tucker, A. D. Turner, and D. Wiebe. SPIDER: Probing the Early Universe with a Suborbital Polarimeter. *ArXiv e-prints*, June 2011, 1106.3087.
- [26] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005, astro-ph/0409513v1. URL <http://stacks.iop.org/0004-637X/622/i=2/a=759>.
- [27] R. Monsalve. *Calibrations and Observations with the QUIET Radiotelescope*. PhD thesis, University of Miami, 2012.
- [28] R. Dumoulin. *Results from the QUIET Q-band Observing Season*. PhD thesis, Columbia University, 2011. URL <http://academiccommons.columbia.edu/item/ac:141910>.
- [29] R. Feynman. Cargo cult science. *Lock Haven University*, Donald Simanek's Pages, rev. August 2008.

- [30] Max A. Woodbury. *Inverting modified matrices*. Statistical Research Group, Memo. Rep. no. 42. Princeton University, Princeton, N. J., 1950.
- [31] William W. Hager. Updating the inverse of a matrix. *SIAM Rev.*, 31(2): 221–239, 1989.
- [32] Magnus R. Hestenes; Eduard Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49(6), December 1952.
- [33] L. S. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users’ Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997. ISBN 0-89871-397-8 (paperback).
- [34] Jack Poulson, Bryan Marker, Robert A. van de Geijn, Jeff R. Hammond, and Nichols A. Romero. Elemental: A new framework for distributed memory dense matrix computations. *ACM Transactions on Mathematical Software*, 2012. To appear.
- [35] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users’ Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999. ISBN 0-89871-447-8 (paperback).
- [36] E. F. Bunn, K. B. Fisher, Y. Hoffman, O. Lahav, J. Silk, and S. Zaroubi. Wiener filtering of the COBE Differential Microwave Radiometer data. *ApJL*, 432:L75–L78, September 1994, arXiv:astro-ph/9404007.
- [37] M. Tsuboi, M. Inoue, T. Handa, H. Tabara, T. Kato, Y. Sofue, and N. Kaifu. Prominent polarized plumes in the galactic center region and their magnetic field. *AJ*, 92:818–824, October 1986.
- [38] R. F. Haynes, R. T. Stewart, A. D. Gray, W. Reich, P. Reich, and U. Mebold. Polarized arcs near the Galactic Centre. *A&A*, 264:500–512, October 1992.
- [39] C. J. Law. A Multiwavelength View of a Mass Outflow from the Galactic Center. *ApJ*, 708:474–484, January 2010, 0911.2061.
- [40] Antony Lewis, Anthony Challinor, and Anthony Lasenby. Efficient computation of CMB anisotropies in closed FRW models. *Astrophys. J.*, 538:473–476, 2000, astro-ph/9911177.
- [41] U. Seljak and M. Zaldarriaga. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *ApJ*, 469:437, October 1996, arXiv:astro-ph/9603033.

- [42] Michael Doran. Cmbeasy: an object oriented code for the cosmic microwave background. *Journal of Cosmology and Astroparticle Physics*, 2005(10):011, 2005, arXiv:astro-ph/0302138v2. URL <http://stacks.iop.org/1475-7516/2005/i=10/a=011>.
- [43] M. Tygert. Fast algorithms for spherical harmonic expansions, III. *Journal of Computational Physics*, 229:6181–6192, September 2010, 0910.5435.
- [44] D. S. Seljebotn. Wavemoth-Fast Spherical Harmonic Transforms by Butterfly Matrix Compression. *ApJS*, 199:5, March 2012, 1110.4874.
- [45] Gayoung Chon, Anthony Challinor, Simon Prunet, Eric Hivon, and István Szapudi. Fast estimation of polarization power spectra using correlation functions. *Monthly Notices of the Royal Astronomical Society*, 350(3):914–926, 2004, astro-ph/0303414v2. ISSN 1365-2966. URL <http://dx.doi.org/10.1111/j.1365-2966.2004.07737.x>.
- [46] W. H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes: The art of scientific computing*. Cambridge University Press, 3 edition, 2009.
- [47] Eric Hivon, Krzysztof M. Górski, C. Barth Netterfield, Brendan P. Crill, Simon Prunet, and Frode Hansen. Master of the cosmic microwave background anisotropy power spectrum: A fast method for statistical analysis of large and complex cosmic microwave background data sets. *The Astrophysical Journal*, 567(1):2, 2002, astro-ph/0105302v1. URL <http://stacks.iop.org/0004-637X/567/i=1/a=2>.
- [48] Antony Lewis and Sarah Bridle. Cosmological parameters from cmb and other data: A monte carlo approach. *Phys. Rev. D*, 66:103511, Nov 2002, astro-ph/0205436v3. URL <http://link.aps.org/doi/10.1103/PhysRevD.66.103511>.
- [49] F. Feroz, M. P. Hobson, and M. Bridges. Multinest: an efficient and robust bayesian inference tool for cosmology and particle physics. *Monthly Notices of the Royal Astronomical Society*, 398(4):1601–1614, 2009, 0809.3437v1. ISSN 1365-2966. URL <http://dx.doi.org/10.1111/j.1365-2966.2009.14548.x>.
- [50] M. Loève. Probability theory. *Graduate Texts in Mathematics*, II, 1978.
- [51] D. L. Larson, H. K. Eriksen, B. D. Wandelt, K. M. Górski, Greg Huey, J. B. Jewell, and I. J. O’Dwyer. Estimation of polarized power spectra by gibbs sampling. *The Astrophysical Journal*, 656(2):653, 2007, astro-ph/0608007v2. URL <http://stacks.iop.org/0004-637X/656/i=2/a=653>.



- [52] J. B. Jewell, H. K. Eriksen, B. D. Wandelt, I. J. O'Dwyer, G. Huey, and K. M. Górski. A Markov Chain Monte Carlo Algorithm for Analysis of Low Signal-To-Noise Cosmic Microwave Background Data. *ApJ*, 697: 258–268, May 2009, 0807.0624.
- [53] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-6(6):721–741, nov. 1984. ISSN 0162-8828.
- [54] Benjamin D. Wandelt, David L. Larson, and Arun Lakshminarayanan. Global, exact cosmic microwave background data analysis using gibbs sampling. *Phys. Rev. D*, 70:083511, Oct 2004, astro-ph/0310080v2. URL <http://link.aps.org/doi/10.1103/PhysRevD.70.083511>.
- [55] H. K. Eriksen, I. J. O'Dwyer, J. B. Jewell, B. D. Wandelt, D. L. Larson, K. M. Górski, S. Levin, A. J. Banday, and P. B. Lilje. Power Spectrum Estimation from High-Resolution Maps by Gibbs Sampling. *ApJS*, 155: 227–241, December 2004, arXiv:astro-ph/0407028.
- [56] R. D. Davies, C. Dickinson, A. J. Banday, T. R. Jaffe, K. M. Górski, and R. J. Davis. A determination of the spectra of galactic components observed by the wilkinson microwave anisotropy probe. *Monthly Notices of the Royal Astronomical Society*, 370(3): 1125–1139, 2006, astro-ph/0511384v2. ISSN 1365-2966. URL <http://dx.doi.org/10.1111/j.1365-2966.2006.10572.x>.
- [57] J. Delabrouille, M. Betoule, and J. B. Melin et al. The pre-launch planck sky model: a model sky emission at submillimetre to centimetre wavelengths. 2012.
- [58] D. H. D. H. Lyth and A. A. Riotto. Particle physics models of inflation and the cosmological density perturbation. *Phys.Rep.*, 314:1–146, June 1999, arXiv:hep-ph/9807278.
- [59] Kendrick M. Smith, Oliver Zahn, and Olivier Doré. Detection of gravitational lensing in the cosmic microwave background. *Phys. Rev. D*, 76:043510, Aug 2007. URL <http://link.aps.org/doi/10.1103/PhysRevD.76.043510>.
- [60] E. Martinez-Gonzalez. Gaussianity. *ArXiv e-prints*, May 2008, 0805.4157.
- [61] K. M. Smith. Beyond the Power Spectrum: Primordial and Secondary Non-Gaussianity in the Microwave Background. In L. M. Stanford, J. D. Green, L. Hao, and Y. Mao, editors, *New Horizons in Astronomy: Frank N. Bash Symposium 2009*, volume 432 of *Astronomical Society of the Pacific Conference Series*, page 147, October 2010, 1111.1783.

- [62] J. Hoftuft, H. K. Eriksen, A. J. Banday, K. M. Górski, F. K. Hansen, and P. B. Lilje. Increasing Evidence for Hemispherical Power Asymmetry in the Five-Year WMAP Data. *ApJ*, 699:985–989, July 2009, 0903.1229.
- [63] G. Rossmannith, H. Modest, C. Räth, A. J. Banday, K. M. Górski, and G. Morfill. Search for Non-Gaussianities in the WMAP Data with the Scaling Index Method. *Advances in Astronomy*, 2011, 2011, 1108.0596.
- [64] M. A. Stephens. Edf statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association*, 69(347):pp. 730–737, 1974. ISSN 01621459. URL <http://www.jstor.org/stable/2286009>.
- [65] V. G. Gurzadyan, A. E. Allahverdyan, T. Ghahramanyan, A. L. Kashin, H. G. Khachatryan, A. A. Kocharyan, S. Mirzoyan, E. Poghossian, D. Vetrugno, and G. Yegorian. A weakly random Universe? *A&A*, 525:L7+, January 2011, 1011.3711.
- [66] V. G. Gurzadyan and R. Penrose. Concentric circles in WMAP data may provide evidence of violent pre-Big-Bang activity. *ArXiv e-prints*, November 2010, 1011.3706.
- [67] I. K. Wehus and H. K. Eriksen. A search for concentric circles in the 7-year WMAP temperature sky maps. *ArXiv e-prints*, December 2010, 1012.1268.
- [68] A. Moss, D. Scott, and J. P. Zibin. No evidence for anomalously low variance circles on the sky. *JCAP*, 4:33–+, April 2011, 1012.1305.
- [69] A. Hajian. Are There Echoes From The Pre-Big Bang Universe? A Search for Low Variance Circles in the CMB Sky. *ArXiv e-prints*, December 2010, 1012.1656.

# **Part II**

# **Papers**



## Paper I

---

*First Season QUIET Observations: Measurements of CMB polarization power spectra at 43 GHz in the multipole range  $25 \leq l \leq 475$*

QUIET Collaboration

ApJ, Volume 690, Issue 2, pp. 111 (2011)

---

# FIRST SEASON QUIET OBSERVATIONS: MEASUREMENTS OF CMB POLARIZATION POWER SPECTRA AT 43 GHz IN THE MULTIPOLE RANGE $25 \leq \ell \leq 475$

QUIET COLLABORATION—C. BISCHOFF<sup>1,22</sup>, A. BRIZIUS<sup>1,2</sup>, I. BUDER<sup>1</sup>, Y. CHINONE<sup>3,4</sup>, K. CLEARY<sup>5</sup>, R. N. DUMOULIN<sup>6</sup>, A. KUSAKA<sup>1</sup>, R. MONSALVE<sup>7</sup>, S. K. NÆSS<sup>8</sup>, L. B. NEWBURGH<sup>6,23</sup>, R. REEVES<sup>5</sup>, K. M. SMITH<sup>1,23</sup>, I. K. WEHUS<sup>9</sup>, J. A. ZUNTZ<sup>10,11,12</sup>, J. T. L. ZWART<sup>6</sup>, L. BRONFMAN<sup>13</sup>, R. BUSTOS<sup>7,13,14</sup>, S. E. CHURCH<sup>15</sup>, C. DICKINSON<sup>16</sup>, H. K. ERIKSEN<sup>8,17</sup>, P. G. FERREIRA<sup>10</sup>, T. GAIER<sup>18</sup>, J. O. GUNDERSEN<sup>7</sup>, M. HASEGAWA<sup>3</sup>, M. HAZUMI<sup>3</sup>, K. M. HUFFENBERGER<sup>7</sup>, M. E. JONES<sup>10</sup>, P. KANGASLAHTI<sup>18</sup>, D. J. KAPNER<sup>1,24</sup>, C. R. LAWRENCE<sup>18</sup>, M. LIMON<sup>6</sup>, J. MAY<sup>13</sup>, J. J. McMAHON<sup>19</sup>, A. D. MILLER<sup>6</sup>, H. NGUYEN<sup>20</sup>, G. W. NIXON<sup>21</sup>, T. J. PEARSON<sup>5</sup>, L. PICCIRILLO<sup>16</sup>, S. J. E. RADFORD<sup>5</sup>, A. C. S. READHEAD<sup>5</sup>, J. L. RICHARDS<sup>5</sup>, D. SAMTLEBEN<sup>2,25</sup>, M. SEIFFERT<sup>18</sup>, M. C. SHEPHERD<sup>5</sup>, S. T. STAGGS<sup>21</sup>, O. TAJIMA<sup>1,3</sup>, K. L. THOMPSON<sup>15</sup>, K. VANDERLINDE<sup>1,26</sup>, R. WILLIAMSON<sup>6,27</sup>, B. WINSTEIN<sup>1</sup>

*Submitted to ApJ—This paper should be cited as “QUIET (2010)”*

## ABSTRACT

The Q/U Imaging ExperimenT (QUIET) employs coherent receivers at 43 GHz and 95 GHz, operating on the Chajnantor plateau in the Atacama Desert in Chile, to measure the anisotropy in the polarization of the CMB. QUIET primarily targets the B modes from primordial gravitational waves. The combination of these frequencies gives sensitivity to foreground contributions from diffuse Galactic synchrotron radiation. Between 2008 October and 2010 December, over 10,000 hours of data were collected, first with the 19-element 43-GHz array (3458 hours) and then with the 90-element 95-GHz array. Each array observes the same four fields, selected for low foregrounds, together covering  $\approx 1000$  square degrees. This paper reports initial results from the 43-GHz receiver which has an array sensitivity to CMB fluctuations of  $69 \mu\text{K}\sqrt{\text{s}}$ . The data were extensively studied with a large suite of null tests before the power spectra, determined with two independent pipelines, were examined. Analysis choices, including data selection, were modified until the null tests passed. Cross correlating maps with different telescope pointings is used to eliminate a bias. This paper reports the EE, BB, and EB power spectra in the multipole range  $\ell = 25\text{--}475$ . With the exception of the lowest multipole bin for one of the fields, where a polarized foreground, consistent with Galactic synchrotron radiation, is detected with  $3\text{-}\sigma$  significance, the E-mode spectrum is consistent with the  $\Lambda$ CDM model, confirming the only previous detection of the first acoustic peak. The B-mode spectrum is consistent with zero, leading to a measurement of the tensor-to-scalar ratio of  $r = 0.35^{+1.06}_{-0.87}$ . The combination of a new time-stream “double-demodulation” technique, Mizuguchi–Dragone optics, natural sky rotation, and frequent boresight rotation leads to the lowest level of systematic contamination in the B-mode power so far reported, below the level of  $r = 0.1$ .

*Subject headings:* cosmic background radiation—Cosmology: observations—Gravitational waves—Inflation—Polarization

<sup>1</sup> Kavli Institute for Cosmological Physics, Department of Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA; send correspondence to A. Kusaka, [ak-ito@kicp.uchicago.edu](mailto:ak-ito@kicp.uchicago.edu)

<sup>2</sup> Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

<sup>3</sup> High Energy Accelerator Research Organization (KEK), 1-1 Oho, Tsukuba, Ibaraki 305-0801, Japan

<sup>4</sup> Astronomical Institute, Graduate School of Science, Tohoku University, Aramaki, Aoba, Sendai 980-8578, Japan

<sup>5</sup> Cahill Center for Astronomy and Astrophysics, California Institute of Technology, 1200 E. California Blvd M/C 249-17, Pasadena, CA 91125, USA

<sup>6</sup> Department of Physics and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

<sup>7</sup> Department of Physics, University of Miami, 1320 Campo Sano Drive, Coral Gables, FL 33146, USA

<sup>8</sup> Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

<sup>9</sup> Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway

<sup>10</sup> Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

<sup>11</sup> Oxford Martin School, 34 Broad Street, Oxford OX1 3BD, UK

<sup>12</sup> Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

<sup>13</sup> Departamento de Astronomía, Universidad de Chile, Casilla 36-D, Santiago, Chile

<sup>14</sup> Departamento de Astronomía, Universidad de Concepción, Casilla 160-C, Concepción, Chile

<sup>15</sup> Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Varian Physics Building, 382 Via Pueblo Mall, Stanford, CA 94305, USA

<sup>16</sup> Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK

<sup>17</sup> Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway

<sup>18</sup> Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

<sup>19</sup> Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA

<sup>20</sup> Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

<sup>21</sup> Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

<sup>22</sup> Current address: Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 43, Cambridge, MA 02138, USA

<sup>23</sup> Current address: Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

<sup>24</sup> Current address: Micro Encoder Inc., Kirkland, WA 98034, USA

## 1. INTRODUCTION

The inflationary paradigm resolves several outstanding issues in cosmology, including the flatness, horizon, and monopole problems, and it provides a compelling explanation for the origin of structure in the Universe (e.g. Liddle & Lyth 2000, and references therein). So far all cosmological data, including measurements of Cosmic Microwave Background (CMB) anisotropies, support this paradigm; still the underlying fundamental physics responsible for inflation is unknown. Inflation produces a stochastic background of gravity waves that induce odd-parity tensor “B modes” at large angular scales in the CMB polarization. If these primordial B modes, parametrized by the tensor-to-scalar ratio  $r$ , are detected, one can learn about the energy scale of inflation. In many attractive slow-roll models, this scale is given approximately by  $r^{1/4} \times 10^{16}$  GeV. For large-field models, the energy scale is near the Grand Unification Scale in particle physics, so that  $r \gtrsim 0.01$ . A new generation of experiments aims for good sensitivity in this range of  $r$ . Establishing the existence of primordial B modes would both verify an important prediction of inflation and provide access to physics at an incredibly high energy scale.

The most stringent limit to date is  $r < 0.20$  at the 95% confidence level (Komatsu et al. 2010) set by a combination of CMB-temperature-anisotropy measurements, baryon acoustic oscillations, and supernova observations, but cosmic variance prohibits improvements using only these measurements.

E-mode polarization has now been detected by many experiments (e.g., Kovac et al. 2002; Leitch et al. 2005; Montroy et al. 2006; Sievers et al. 2007; Wu et al. 2007; Bischoff et al. 2008; Larson et al. 2010). These measurements are consistent with predictions from CMB-temperature-anisotropy measurements, and they provide new information on the epoch of reionization. Only BICEP has accurately measured E-mode polarization in the region of the first acoustic peak (Chiang et al. 2010); that paper also reports the best limit on  $r$  coming from cosmological B modes:  $r < 0.72$  at the 95% confidence level.

Experiments measuring B-mode polarization in the CMB should yield the best information on  $r$ , but this technique is still in its infancy. B modes are expected to be at least an order of magnitude smaller than the E modes so control of systematic errors and foregrounds will be particularly critical. Below  $\approx 90$  GHz, the dominant foreground comes from Galactic synchrotron emission, while at higher frequencies, emission from thermal dust dominates. Most planned or operating CMB polarization experiments employ bolometric detectors observing most comfortably at frequencies  $\gtrsim 90$  GHz, so they cannot estimate synchrotron contamination from their own data.

The Q/U Imaging Experiment (QUIET) is one of two CMB polarization experiments to observe at frequencies suitable for addressing synchrotron contamination, making observations at 43 GHz (Q band) and 95 GHz (W band) and with sufficient sensitivity to begin to probe primordial B modes. The other is *Planck* (Tauber et al. 2010).

QUIET uses compact polarization-sensitive modules based upon High-Electron-Mobility Transistor (HEMT) amplifiers, combined with a new time-stream “double-demodulation” technique, Mizuguchi–Dragone (MD) optics (for the first time in a CMB polarization experiment), natural sky rotation, and frequent rotation about the optical axis to achieve a very low level of contamination in the multipole range where a primordial-B-mode signal is expected.

Between 2008 October and 2010 December, QUIET collected over 10,000 hours of data, split between the Q-band and W-band receivers. Here we report first results from the first season of 3458 hours of Q-band observation. After describing the instrument, observations, and detector calibrations (Sections 2, 3, and 4), we discuss our analysis techniques and consistency checks (5 and 6). CMB power spectra are then presented together with a foreground detection (7). We evaluate our systematic errors (8) and then conclude (9).

## 2. THE INSTRUMENT

The QUIET instrument comprises an array of correlation polarimeters cooled to 20 K and coupled to a dual-reflector telescope, installed on a three-axis mount inside a comoving ground screen. The instrument is illustrated in Figure 1. Further details are given below and in Newburgh et al. (2010), Kusaka et al. (2010), and Buder (2010).




FIG. 1.— Overview of the QUIET instrument. The cryostat and 1.4-m telescope mirrors are enclosed in a rectangular comoving absorbing ground screen; in this figure its walls are transparent. The telescope, cryostat and electronics are mounted on a single platform attached to the deck bearing, which allows rotations around the instrument’s optical axis.

The Q-band QUIET receiver is a 19-element array containing 17 low-noise correlation polarimeters, each simultaneously measuring the Stokes Q, U, and I parameters, and two CMB differential-temperature monitors.

The first element in the QUIET optical chain is a 1.4-m crossed Mizuguchi–Dragone dual-reflective telescope

<sup>25</sup> Current address: Nikhef, Science Park, Amsterdam, The Netherlands

<sup>26</sup> Current address: Department of Physics, McGill University, 3600 Rue University, Montreal, Quebec H3A 2T8, Canada

<sup>27</sup> Current address: Kavli Institute for Cosmological Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

(Mizugutch et al. 1976; Dragone 1978). The crossed MD configuration is very compact, with low cross polarization and a large diffraction-limited field of view. The telescope is described in detail in Imbriale et al. (2010). Light incident on the mirrors is focused into an array of corrugated circular feed horns (Gundersen & Wolack 2009), yielding a full-width half-maximum (FWHM) beam size of  $27'.3$  and a roughly circular field of view of  $7^\circ$  diameter. Radiation from each feed horn enters a septum polarizer (Bornemann & Labay 1995) which separates left and right circularly-polarized components ( $L$  and  $R$ ) into two waveguide ports which mate to a QUIET correlation module, detailed below.

The module array and feed horns are cooled to 20 K in a cryostat to reduce instrumental noise. An electronics enclosure mounted next to the cryostat houses the electronics necessary for biasing the modules and recording their data. The cryostat, electronics, and telescope are installed on the former CBI mount (Padin et al. 2002). This mount provides three-axis motion: azimuth, elevation, and rotation about the optical axis. This last is called “deck” rotation.

The cryostat and telescope are enclosed by an absorbing comoving ground screen. The ground screen was designed to have two parts, but the upper section (not shown in Fig. 1) was not installed until after the Q-band instrument was removed. Its absence was correctly anticipated to result in two far sidelobes, which were mapped with a high-power source by the QUIET W-band instrument in the field and measured to be  $\lesssim -60$  dB with the QUIET Q-band instrument when the Sun passed through them. The effects of these sidelobes are mitigated through filtering and data selection (Sections 5.1.3 and 5.2). Section 8.4 shows that any residual contamination is small.

Each QUIET Q-band correlation module, in a footprint of only  $5.1 \times 5.1 \text{ cm}^2$ , receives the circular polarization modes of the celestial radiation and outputs Stokes  $Q$ ,  $U$  and  $I$  as follows. Each input is independently amplified and passed through a phase switch. One phase switch alternates the sign of the signal voltage at 4 kHz, while the other switches at 50 Hz. The two signals are combined in a  $180^\circ$  hybrid coupler, with outputs proportional to the sum and difference of the inputs. Since the module inputs are proportional to  $(L, R) = (E_x \pm iE_y)/\sqrt{2}$ , where  $E_x$  and  $E_y$  are orthogonal components of the incident electric field, the coupler outputs are amplified versions of  $E_x$  and  $iE_y$ , with the phase switch reversing their roles. Half of each output is bandpass filtered and rectified by a pair of detector diodes, while the other half passes into a  $90^\circ$  hybrid coupler. A second pair of bandpass filters and detector diodes measures the power from this coupler’s outputs (Kangaslahti et al. 2006).

Synchronous demodulation of the 4-kHz phase switching yields measurements of Stokes  $+Q$  and  $-Q$  on the first two diodes and Stokes  $+U$  and  $-U$  on the remaining two. This high-frequency differencing suppresses low-frequency atmospheric fluctuations as well as  $1/f$  noise from the amplifiers, detector diodes, bias electronics, and data-acquisition electronics. Subsequent demodulation of the 50-Hz phase switching removes spurious instrumental polarization generated by unequal transmission

coefficients in the phase-switch circuits. The resulting four “double-demodulated” time streams are the polarization channels.

Averaging the output of each diode rather than demodulating it results in a measurement of Stokes  $I$ , hereafter called total power, denoted “TP.” The TP time streams are useful for monitoring the weather and the stability of the detector responsivities, but suffer too much contamination from  $1/f$  noise to constrain the CMB temperature anisotropy. Therefore, the Q-band instrument includes two correlation modules that are coupled to a pair of neighboring feed horns to measure the temperature difference between them, in a scheme similar to the *WMAP* differencing assemblies (Jarosik et al. 2003). These differential-temperature modules provide calibration data for the telescope pointing, beams, and sidelobes, as well as CMB data. Their feed horns are in the outer ring of the close-packed hexagonal array,  $\approx 3^\circ$  from the center.

Here we summarize several array-wide characteristics of the polarimeters. Bandpass measurements in the lab and at the start of the observing season find that the average center frequency is  $43.1 \pm 0.4$  GHz, and the average bandwidth is  $7.6 \pm 0.5$  GHz. We calculate the noise power spectra of the double-demodulated polarimeter time streams from each 40–90-minute observation to assess their  $1/f$  knee frequencies and white-noise levels (see Section 5.1). The median  $1/f$  knee frequency is 5.5 mHz, well below the telescope scan frequencies of 45–100 mHz.

From the white-noise levels and responsivities (Section 4.1) we find an array sensitivity<sup>28</sup> to CMB fluctuations of  $69 \mu\text{K}\sqrt{\text{s}}$ , such that the mean polarized sensitivity per module is  $280 \mu\text{K}\sqrt{\text{s}}$ .




FIG. 2.— The CMB and Galactic patches, in equatorial coordinates, superimposed on a Q-band all-sky *WMAP* 7-year temperature map (Jarosik et al. 2010). Note that the Galactic-plane temperature signal saturates the color scale. Patch G-2 is the Galactic center.

### 3. OBSERVATIONS

QUIET is located on the Chajnantor plateau in the Atacama Desert of northern Chile ( $67^\circ 45' 42''$  W,  $23^\circ 01' 42''$  S). A combination of high altitude (5080 m) and extreme dryness results in excellent observing conditions for most of the year. During the eight months of QUIET Q-band observations, the median precipitable

<sup>28</sup> This is the sensitivity for 62 polarization channels. Six of 68 polarization channels are non-functional—an array yield of 92%.



water vapor (PWV) measured at the nearby APEX site (Güsten et al. 2006) was 1.2 mm.

We began observations with the Q-band receiver on 2008 October 24, and took 3458 hours of data until 2009 June 13 (when the receiver was replaced on the telescope by the 90-element W-band receiver). Of these data, 77% are for CMB, with 12% of the observing time used for Galactic fields, 7% for calibration sources, and 4% cut due to obvious instrumental problems such as lack of telescope motion. We observe 24 hours a day, except when interrupted. Our full-season operating efficiency is 63%; causes of downtime include occasional snow, power outages, and mechanical failures.

TABLE 1  
PATCH LOCATIONS AND INTEGRATION TIMES

Patch	RA (J2000)	Dec.	Integration Hours
CMB-1	12 <sup>h</sup> 04 <sup>m</sup>	−39°00′	905
CMB-2	05 <sup>h</sup> 12 <sup>m</sup>	−39°00′	703
CMB-3	00 <sup>h</sup> 48 <sup>m</sup>	−48°00′	837
CMB-4	22 <sup>h</sup> 44 <sup>m</sup>	−36°00′	223
G-1	16 <sup>h</sup> 00 <sup>m</sup>	−53°00′	311
G-2	17 <sup>h</sup> 46 <sup>m</sup>	−28°56′	92

NOTE. — The central equatorial coordinates and integration times for each observing patch. G-1 and G-2 are Galactic patches.

### 3.1. Field Selection

We observe four CMB fields, referred to henceforth as “patches.” Table 1 lists their center positions and total integration times, while Figure 2 indicates their positions on the sky. The number of patches is determined by the requirement to always have one patch above the lower elevation limit of the mount (43°). The specific positions of each patch were chosen to minimize foreground emission using *WMAP* 3-year data. The area of each patch is  $\approx 250 \text{ deg}^2$ . In addition to the four CMB patches, we observe two Galactic patches. These allow us to constrain the spectral properties of the polarized low-frequency foregrounds with a high signal-to-noise ratio. The results from the Galactic observations will be presented in a future publication.

### 3.2. Observing Strategy

Scanning the telescope modulates the signal from the sky, converting CMB angular scales into frequencies in the polarimeter time streams. Since QUIET targets large angular scales, fast scanning ( $\approx 5^\circ \text{ s}^{-1}$  in azimuth) is critical to ensuring that the polarization modes of interest appear at higher frequencies than the atmospheric and instrumental  $1/f$  knee frequencies.

So that each module sees a roughly-constant atmospheric signal, each QUIET scan is a constant-elevation scan (CES): periodic motion solely in azimuth with both the elevation and deck-rotation axes fixed. Each CES has an amplitude of  $7.5^\circ$  on the sky, with period 10–22 s. Typical CESes last 40–90 minutes. We repoint the telescope when the patch center has moved by  $15^\circ$  in order to build up data over an area of  $\approx 15^\circ \times 15^\circ$  for each patch. Note that a central region  $\approx 8^\circ$  across is observed by all polarimeters since the instrument’s field of view has a diameter of  $\approx 7^\circ$ . Diurnal sky rotation and weekly

deck rotations provide uniform parallactic-angle coverage of the patch, and ensure that its peripheral regions are also observed by multiple polarimeters.

TABLE 2  
REGULAR CALIBRATION OBSERVATIONS

Source	Schedule	Duration (min.)
sky dips	every 1.5 hours	3
Tau A	every 1–2 days	20
Moon	weekly	60
Jupiter	weekly	20
Venus	weekly	20
RCW38	weekly	20

## 4. CALIBRATION

Four quantities are required to convert polarimeter time streams into polarization power spectra: detector responsivities, a pointing model, detector polarization angles, and beam profiles. To this end, a suite of calibration observations is performed throughout the season using astronomical sources (Taurus A—hereafter Tau A, Jupiter, Venus, RCW38, and the Moon); atmospheric measurements (“sky dips,” which typically consist of three elevation nods of  $\pm 3^\circ$ ); and instrumental sources (a rotating sparse wire grid and a polarized broadband noise source). From these we also measure instrumental polarization, as described below. QUIET’s regular calibration observations are summarized in Table 2.

We typically use two or more methods to determine a calibration constant, taking the spread among the methods as an indication of the uncertainty. We show in Section 8 that aside from the case of absolute responsivity, all calibration uncertainties lead to estimates of systematic effects on the power spectra well below statistical errors. This immunity comes from having a large number of detectors and highly-crosslinked polarization maps.




FIG. 3.— Polarimeter responses from the central feed horn to the polarization of Tau A at four parallactic angles. These data were collected with one correlation module in about 20 minutes. The errors are smaller than the points. From top to bottom, responses are shown for the detector diodes sensitive to the Stokes parameters  $+Q$ ,  $-Q$ ,  $+U$ , and  $-U$ , respectively. For each, the fitted model is plotted as a dashed line.

#### 4.1. Responsivity

The polarized flux from Tau A provides a 5 mK signal which we observe at four parallactic angles. The sinusoidal modulation of the signal induced by the changing parallactic angles is fitted to yield responsivity coefficients for each detector. Figure 3 shows the response of the four polarization channels from the central feed horn to Tau A. A typical responsivity is  $2.3 \text{ mV K}^{-1}$ , with a precision from a single set of observations of 6%. The absolute responsivity from Tau A was measured most frequently for the central feed horn. We choose its +Q diode detector to provide the fiducial absolute responsivity.

The responsivities of other detectors relative to the fiducial detector are determined with the sky dips as described below. We have three independent means of assessing the relative responsivities among polarimeters: from nearly-simultaneous measurements of the Moon, from simultaneous measurements of responses to the rotating sparse wire grid in post-season tests, and from Tau A measurements. The errors from these methods are 4%, 2%, and 6% respectively, while the error from the sky-dip method is 4%. All the methods agree within errors.

Sky dips generate temperature signals of several 100 mK and thus permit measurement of the TP responsivities. The signals vary slightly with PWV. We estimate the slope from the data as  $4\% \text{ mm}^{-1}$  and correct for it. This slope is consistent with the atmospheric model of Pardo et al. (2001). Because the ratios of the responsivities for the TP and polarized signals from each detector diode are stable quantities within a few percent of unity, we use sky dips performed at the beginning of each CES to correct short-term variations in the polarimeter responsivities. The responsivities vary by  $\lesssim 10\%$  over the course of a day, due to changing thermal conditions for the bias electronics. Further post-season tests provide a physical model: the relevant temperatures are varied intentionally while the responsivities are measured with sky dips. We confirm the results with the polarized broadband source.

We bound the uncertainty in the absolute responsivity of the polarimeter array at 6%. The largest contributions to this estimate are uncertainties in (1) the beam solid angle (4%, see below), (2) the response difference between polarized and TP signals for each diode detector (3%), and (3) the Tau A flux (3%, Weiland et al. 2010). The first enters in converting the flux of Tau A into  $\mu\text{K}$ , while the second enters because although one fiducial diode detector is calibrated directly from Tau A, for the rest we find relative responsivities from sky dips and normalize by the fiducial diode's responsivity.

For the differential-temperature modules, all detectors observe the signal from Jupiter simultaneously, providing the absolute responsivity for all channels upon comparison with the Jupiter flux from Weiland et al. (2010). Observations of Venus (Hafez et al. 2008) and RCW38 agree with the Jupiter measurements within errors, and sky dips track short-term variations. We calibrate the absolute responsivity with 5% accuracy.




FIG. 4.— Map of the polarization of the Moon from one detector diode. The amplitude of the quadrupole polarization visible here is  $\approx 400 \text{ mK}$ . Similar maps are produced for all 17 polarization modules in the array with a single  $\approx$  hour-long observation. The dotted line indicates the polarization orientation of the detector. Contours are spaced at intervals of 100 mK, with negative contours indicated by dashed lines.

#### 4.2. Pointing

The global pointing solution derives from a physical model of the 3-axis mount and telescope tied to observations of the Moon with the central feed horn in the array, as well as Jupiter and Venus with the differential-temperature feed horns. Optical observations are taken regularly with a co-aligned star camera and used to monitor the time evolution of the pointing model.

During the first two months in the season, a mechanical problem with the deck-angle encoder results in pointing shifts. The problem was subsequently repaired. Based on pointing observations of the Moon and other astronomical sources, we verify that these encoder shifts are less than  $2^\circ$ . Systematic uncertainties induced by this problem are discussed in Section 8.1.

After the deck-angle problem is fixed, no significant evolution of the pointing model is found. The difference in the mean pointing solution between the start and the end of the season is smaller than  $1'$ . Observations of the Moon and Jupiter also provide the relative pointing among the feed horns. The root mean square (RMS) pointing error in the maps is  $3/5$ .

#### 4.3. Detector Polarization Angles

Our primary measurement of the polarization angle for each detector comes from observing the radial polarization of the Moon, as illustrated in Figure 4. The polarization angles are stable, changing by  $< 0.2^\circ$ , except during the period with the deck-angle-encoder problem mentioned above.

Two other less precise methods also give estimates of the detector angles: fits to the Tau A data, and determination of the phases of the sinusoidal responses of all the detectors to rotation of the sparse wire grid. In each case, the differences between the detector angles determined by the secondary method and the Moon are described by a standard deviation of  $\approx 3^\circ$ . However, we find a mean shift between the Tau A-derived and Moon-derived angles of  $1.7^\circ$ . To estimate the errors in the angles in light of this shift, we use an empirical approach: in Section 8.2 we estimate the impact on the power spectra from using the Tau A results instead of the Moon results, and find it to be small.




FIG. 5.— Top panel: Polarization beam profile from Tau A observations with the central feed horn. The data are overplotted with the expansion in Gauss-Hermite polynomials described in the text. Bottom panel: Beam window function with errors shown by the gray band.

#### 4.4. Beam Profile and Window Function

The polarization and differential-temperature beams are obtained from maps created using the full data sets of Tau A and Jupiter observations respectively, with square pixels of  $1/8$  on a side. For polarization, this process produces the main and leakage beam maps simultaneously, with the latter describing the instrumental polarization. The average FWHM for the beams across the array is  $27.3$ , measured with  $0.1$  precision for the central feed horn and for the differential-temperature feed horns at the edge of the focal plane. The non-central-polarization-horn FWHMs are measured less frequently and thus are less precisely known, with an uncertainty of  $1/5$ . The beam elongation is typically small ( $1\%$ ), and its effect is further reduced by the diurnal sky rotation and weekly deck rotations which result in a symmetrized effective beam in the CMB maps. We compute 1-dimensional symmetrized beam profiles, with a resolution of  $0.6$ . These profiles are modeled as a sum of six even Gauss-Hermite terms (Monsalve 2010). The main-beam solid angles are computed by integrating these models out to  $54'$  (roughly  $-28$  dB), yielding  $78.0 \pm 0.4 \mu\text{sr}$  for the differential-temperature horns and  $74.3 \pm 0.7 \mu\text{sr}$  for the central horn. An average gives  $76 \mu\text{sr}$  for all horns in the array. We also examine alternative estimates such as integrating the raw beam map instead of the analytical fit. We assign a systematic uncertainty of  $4\%$  based on the differences among these different estimates. The systematic error includes possible contributions from sidelobes, which we constrain to  $0.7 \pm 0.4 \mu\text{sr}$  with antenna range measurements carried out before the observation season.

The window functions, encoding the effect of the finite resolution of the instrument on the power spectra, are computed from the central-horn and the temperature-horn-profile models. The central-horn beam profile and window function are shown in Figure 5. The uncertainty accounts for statistical error and differences between polarization and differential-temperature beams,

as described in Section 8.1.

#### 4.5. Instrumental Polarization

Instrumental imperfections can lead to a spurious polarization signal proportional to the unpolarized CMB temperature anisotropy. We call this the I to Q (or U) leakage term. In our instrument, a fraction of the power input on one port of the correlation module is reflected because of a bandpass mismatch to the septum polarizer, and a fraction of the reflected power re-enters the other port. The dominant monopole term comes from this effect. We measure the monopole term from the polarimeter responses to temperature changes, using sky dipoles; Moon, Tau A, and Galactic signals; as well as variations from the weather. The average magnitude is  $1.0\%$  ( $0.2\%$ ) for the Q (U) diodes. Note that the discrepancy in the Q and U averages was predicted from measurements of the properties of the septum polarizers and confirmed in the field.

### 5. DATA ANALYSIS PROCEDURE

QUIET employs two independent analysis pipelines to derive CMB power spectra. We present the methods used for analysis in each pipeline, including data selection, filtering, map making, and power-spectra estimation.

Pipeline A is based on the pseudo- $C_\ell$  analysis framework, first described by Hivon et al. (2002), which is used by numerous experiments (Netterfield et al. 2002; Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010; Lueker et al. 2010). This pipeline made all analysis choices in accordance with a strict (blind) analysis validation policy described in Section 6. An advantage of the pseudo- $C_\ell$  framework is computational efficiency, which is critical for completing the more than 30 iterations of the null-test suite. For the same reason, this pipeline is used for the systematic-error evaluations found in Section 8. Pseudo- $C_\ell$  analysis also enables us to perform cross correlation, making the resultant power spectra immune to possible misestimation of noise bias.

Pipeline B implements a maximum-likelihood framework (e.g., Tegmark 1997; Bond et al. 1998), which has a long history of use by CMB experiments (e.g., Mouskoff et al. 2000; Page et al. 2007; Wu et al. 2007; Bischoff et al. 2008). This framework yields minimum-variance estimates of the power spectra, naturally accounts for E/B mixing, and directly provides the exact CMB likelihood required for estimation of cosmological parameters, without the use of analytical approximations. In addition to power spectra, it produces unbiased maps with full noise-covariance matrices, useful for comparisons with other experiments. On the other hand, this approach is also computationally more expensive than the pseudo- $C_\ell$  framework, and a reduced set of null tests is therefore used to evaluate data consistency.

The processing of the time-ordered data (TOD) and the methodology used for data selection are treated in Sections 5.1 and 5.2, respectively. Brief descriptions of the pseudo- $C_\ell$  and maximum-likelihood techniques are found in Section 5.3. TOD processing, data selection, and analysis for temperature-sensitive modules are discussed in Section 5.4.



### 5.1. Time-Ordered-Data Processing

To prepare the TOD for map making, we execute three steps: pre-processing, noise modeling, and filtering. Of these steps, only the filtering is significantly different between the two pipelines.

#### 5.1.1. Pre-processing

The first data-processing step is to correct for a small non-linearity that was discovered in the analog-to-digital converter (ADC) system. The non-linearities occur every 1024 bits; roughly 14% of the data are affected. Systematic uncertainty from this effect is estimated in Section 8.5. Next, the receiver data are synchronized with the telescope pointing. The double-demodulation step, described in Section 2, is applied, reducing the sample rate from 100 Hz to 50 Hz. A model of the detectors' polarized responsivities converts the data from ADC counts into thermodynamic temperature. The two pipelines use different responsivity models. Pipeline A applies a constant responsivity throughout each CES, addressing possible variability within a CES as part of the systematic error (Section 8); pipeline B updates responsivities on 2-minute timescales (Dumoulin 2010).

#### 5.1.2. Noise Model

After pre-processing, the time streams for each detector diode in each CES are Fourier-transformed and their noise power spectra are fit to a model<sup>29</sup> with three parameters: the amplitude of white noise, the  $1/f$  knee frequency, and the power-law slope of the  $1/f$  noise. We also compute the white-noise correlations among detector diodes in the same module: the most important are between the two Q or the two U detector diodes (with an average coefficient of 0.22). A small fraction of the noise spectra contain features not accounted for in the noise model: beam sidelobes (see Section 2) scanning across features on the ground create a narrow spike at the scan frequency; slowly-changing weather patterns during a CES create a broader peak also at the scan frequency; and there are some narrow spikes at high ( $\gtrsim 6$  Hz) frequencies. To prevent these features from biasing the noise model, the fit excludes a region around the scan frequency as well as frequencies above 4.6 Hz. In addition to the noise-model parameters, several statistics quantifying the agreement between the data and noise model are also used for data selection as described in Section 5.2.

#### 5.1.3. Filtering

In pipeline A, three filters are applied. These were chosen from the results of many runs of the null-test suite (see Section 6). First, to remove the high-frequency narrow spikes, we apply a low-pass filter that cuts signals off sharply above 4.6 Hz<sup>30</sup>. Second, to suppress contamination from atmospheric fluctuations and detector  $1/f$  noise, we subtract a linear function from each telescope half scan (left-going or right-going) removing modes below twice the scan frequency<sup>31</sup>. The third filter, designed

to eliminate signal from ground emission, removes any azimuthal structure that remains after summing over all half scans in the CES.

In pipeline B, an apodized bandpass filter is used that accepts modes from 2.5 times the scan frequency to 4.5 Hz; the highpass component of this filter is designed to suppress scan-synchronous contamination. Further, a time-independent ground-emission model is subtracted. The model of ground emission is generated by building low-resolution and high-signal-to-noise maps in horizon coordinates from the full-season data for each deck angle and module, using large (55') pixels. Only features that are stable in time, azimuth, elevation, and deck angle contribute to this model. The amplitude of the ground correction is  $\lesssim 1 \mu\text{K}$ .

### 5.2. Data Selection

The fundamental unit of data used for analysis is the double-demodulated output of one detector diode for a single CES, referred to as a "CES-diode." Selecting only those CES-diods that correspond to good detector performance and observing conditions is a critical aspect of the data analysis. The data-selection criteria began with a nominal set of cuts and evolved into several distinct configurations, as many as 33 in the case of pipeline A. For each configuration, analysis validation (see Section 6) was performed yielding statistics quantifying the lack of contamination in the data set. The final data set was chosen when these statistics showed negligible contamination and were little affected by changes to the cuts.

Cut efficiencies, defined as the fractions of CES-diods accepted for the analysis, are given for both pipelines in Table 3. While each pipeline applies its own cuts uniformly to all four patches, the efficiencies among patches are non-uniform because of differences in weather quality. Over the course of the eight month observing season, patch CMB-1 is primarily visible at night, when the atmosphere tends to be more stable; patch CMB-3 is mostly observed during the day.

The first step of the data selection is simply to remove known bad data: data from six non-functional detector diodes, data during periods of mount malfunctions, and CESes lasting less than 1000 s. Further, we cut individual CES-diods that show deviation from the expected linear relationship between the demodulated and TP signals. This cut removes data with poor thermal regulation of the electronics or cryostat, or residual ADC non-linearity.

The beam sidelobes, described in Section 2, introduce contamination to the data if the telescope scanning motion causes them to pass over the ground or the Sun. Ground pickup is dealt with by filtering as described in Section 5.1.3. The less frequent cases of Sun contamination are handled by cutting those CES-diods for which the Sun's position overlaps with the measured sidelobe regions for each diode.

Additional cuts are specific to each pipeline. Pipeline A removes data taken during bad weather using a statistic calculated from fluctuations of the TP data during 10-s periods, averaged across the array. This cut removes entire CESes. Several more cuts remove individual CES-diods. While these additional cuts are derived from the noise modeling statistics, they also target residual bad weather. During such marginal weather conditions only some channels need to be cut, since the sensitivity for

<sup>29</sup> At the level of a single CES, the TOD of each detector diode are dominated by noise; the contribution of the CMB is negligible.

<sup>30</sup> For QUIET's beam size and scanning speed a low-pass filter of 4.5–4.6 Hz results in a minimal loss of sensitivity to the CMB.

<sup>31</sup> Typical scan frequencies range from 45 mHz to 100 mHz.

a given detector diode to atmospheric fluctuations depends on its level of instrumental polarization. Next, we reject CES-diodes with poor agreement between the filtered data and the noise model in three frequency ranges: a narrow range (only 40 Fourier modes) about the scan frequency, from twice the scan frequency to 1 Hz, and from 1 Hz to 4.6 Hz. We also cut CES-diodes that have higher than usual  $1/f$  knee frequencies, or large variations during the CES in the azimuthal slopes of the double-demodulated time streams; both these cuts help eliminate bad weather periods. Finally, we also remove any CES-diodes with an outlier greater than  $6\sigma$  in the time domain on three timescales (20 ms, 100 ms, and 1 s).

For pipeline B, the weather cut rejects CESes based on a statistic computed from fluctuations of the double-demodulated signals from the polarization modules on 10-s and 30-s timescales. Three cuts are applied to remove individual CES-diodes. The first is a cut on the  $1/f$  knee frequency, similar to that of pipeline A. Second, a cut is made on the noise model  $\chi^2$  in the frequency range passed by the filter, and third, we reject CES-diodes having a large  $\chi^2$  in the azimuth-binned TOD. This cut rejects data with possible time variation in the ground signal. Finally, an entire CES is removed if more than 40% of its detectors have already been rejected.

### 5.3. Map Making and Power-Spectra Estimation

After filtering, the TOD for all diodes are combined to produce  $Q$  and  $U$  maps for each of the QUIET patches. The maps use a HEALPix  $N_{\text{side}} = 256$  pixelization (Gorski et al. 2005). This section describes the map making and power-spectra estimation from the maps for each of the pipelines.

#### 5.3.1. Pipeline-A Map Making

Polarization maps ( $Q$  and  $U$ ) are made by summing samples into each pixel weighted by their inverse variance, calculated from the white-noise amplitudes. The full covariance matrix is not calculated. Two polarized sources, Centaurus A and Pictor A, are visible in the maps and are removed using circular top-hat masks with radii of  $2^\circ$  and  $1^\circ$ , respectively.

Separate maps are made for each range of telescope azimuth and deck-angle orientations. The coordinates are binned such that there are 10 divisions in azimuth<sup>32</sup> and six distinct ranges of deck-angle orientation. Making separate maps for different telescope pointings enables the cross correlation described in the next section.

TABLE 3  
TOTAL HOURS OBSERVED AND DATA-SELECTION EFFICIENCIES

Patch	Total Hours	A %	B %	Common %
CMB-1	905	81.7	84.3	76.7
CMB-2	703	67.3	70.0	61.2
CMB-3	837	56.0	61.4	51.4
CMB-4	223	70.6	74.2	65.9
All Patches	2668	69.4	72.9	64.2

NOTE. — Selection efficiencies for each pipeline. “Common” gives the efficiencies if both sets of cuts were applied.

<sup>32</sup> The azimuth divisions are the same for all patches, which means that not all divisions are populated for patches CMB-3 and CMB-4.

#### 5.3.2. Power-Spectra Estimation in Pipeline A

The MASTER (Monte Carlo Apodized Spherical Transform Estimator) method is used in pipeline A (Hivon et al. 2002; Hansen & Gorski 2003); it is based on a pseudo- $C_\ell$  technique and takes account of effects induced by the data processing using Monte Carlo (MC) simulations. The pseudo- $C_\ell$  method allows estimation of the underlying  $C_\ell$  using spherical-harmonics transformations when the observations do not cover the full sky uniformly (Wandelt et al. 2001). The pseudo- $C_\ell$  spectrum, designated by  $\tilde{C}_\ell$ , is related to the true spectrum  $C_\ell$  by:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle. \quad (1)$$

There is no term corresponding to noise bias, which would arise if we did not employ a cross-correlation technique. Here  $B_\ell$  is the beam window function, described in Section 4.4, and  $M_{\ell\ell'}$  is a mode-mode-coupling kernel describing the effect of observing only a small fraction of the sky with non-uniform coverage. It is calculable from the pixel weights, which are chosen to maximize the signal-to-noise ratio (Feldman et al. 1994). We bin in  $\ell$  and recover  $C_\ell$  in nine band powers,  $C_b$ , and  $F_\ell$  is the transfer function (displayed in Section 7) due to filtering of the data; its binned estimate,  $F_b$ , is found by processing noiseless CMB simulations through pipeline A and used to obtain  $C_b$ . For the polarization power spectra, equation (1) is generalized for the case where  $\tilde{C}_\ell$  contains both  $\tilde{C}_\ell^{EE}$  and  $\tilde{C}_\ell^{BB}$ .

In the power-spectra estimates, we include only the cross correlations among pointing-division maps, excluding the auto correlations. Because the noise is uncorrelated for different pointing divisions, the cross-correlation technique allows us to eliminate the noise-bias term and thus the possible residual bias due to its misestimate. Cross correlation between different pointing divisions also suppresses possible effects of ground contamination and/or time-varying effects. Dropping the auto correlations creates only a small increase in the statistical errors ( $\approx 3\%$ ) on the power spectra.

The errors estimated for the pipeline-A power spectra are frequentist two-sided 68% confidence intervals. A likelihood function used to compute the confidence intervals is modeled following Hamimeche & Lewis (2008) and calibrated using the MC simulation ensemble of more than 2000 realizations with and without CMB signal. We also use the likelihood function to put constraints on  $r$  and calculate the consistency to  $\Lambda$ CDM.

The partial sky coverage of QUIET generates a small amount of E/B mixing (Challinor & Chon 2005), which contributes an additional variance to the BB power spectrum. We incorporate it as part of the statistical error. This mixing can be corrected (Smith & Zaldarriaga 2007) in future experiments where the effect is not negligible compared to instrumental noise.

#### 5.3.3. Pipeline-B Map Making

In pipeline B, the pixel-space sky map  $\hat{\mathbf{m}}$  ( $N_{\text{side}} = 256$ ) is given by

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{d}, \quad (2)$$

where  $\mathbf{P}$  is the pointing matrix,  $\mathbf{N}$  is the TOD-noise-covariance matrix,  $\mathbf{F}$  corresponds to the apodized band-pass filter discussed in Section 5.1.3, and  $\mathbf{d}$  denotes the TOD. This map is unbiased, and for the case  $\mathbf{F} = \mathbf{1}$  it is additionally the maximum-likelihood map, maximizing

$$\mathcal{L}(\mathbf{m}|\mathbf{d}) = e^{-\frac{1}{2}(\mathbf{d}-\mathbf{Pm})^T\mathbf{N}^{-1}(\mathbf{d}-\mathbf{Pm})}. \quad (3)$$

The corresponding map-noise-covariance matrix (e.g., Tegmark 1997; Keskitalo et al. 2010) is

$$\mathbf{N}_{\hat{\mathbf{m}}} = (\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P})^{-1} (\mathbf{P}^T\mathbf{F}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}) (\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P})^{-1}. \quad (4)$$

Note that one often encounters the simplified expression  $\mathbf{N}_{\hat{\mathbf{m}}} = (\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P})^{-1}$  in the literature. This corresponds effectively to assuming that  $\mathbf{F} = \mathbf{F}^2$  in the Fourier domain, and is strictly valid for top-hat-filter functions only. For our filters, we find that the simplified expression biases the map-domain  $\chi^2$  ( $\equiv \hat{\mathbf{n}}^T\mathbf{N}_{\hat{\mathbf{m}}}^{-1}\hat{\mathbf{n}}$ , where  $\hat{\mathbf{n}}$  is a noise-only map) by  $\approx 3\sigma$ , and we therefore use the full expression, which does lead to an unbiased  $\chi^2$ .

Equations (2–4) apply to both polarization and temperature analysis. The only significant difference lies in the definition of the pointing matrix,  $\mathbf{P}$ . For polarization,  $\mathbf{P}$  encodes the detector orientation, while for temperature it contains two entries per time sample, +1 and −1, corresponding to the two horns in the differential-temperature assembly.

After map making, the maps are post-processed by removing unwanted pixels (i.e., compact sources and low-signal-to-noise edge pixels). All 54 compact sources in the 7-year *WMAP* point source catalog (Gold et al. 2010) present in our four patches are masked out, for a total of 4% of the observed area. We also marginalize over large-scale and unobserved modes by projecting out all modes with  $\ell \leq 5$  ( $\ell \leq 25$  for temperature) from the noise-covariance matrix using the Woodbury formula, assigning infinite variance to these modes.

#### 5.3.4. Power-Spectra Estimation in Pipeline B

Given the unbiased map estimate,  $\hat{\mathbf{m}}$ , and its noise-covariance matrix,  $\mathbf{N}_{\hat{\mathbf{m}}}$ , we estimate the binned CMB power spectra,  $C_b$ , using the Newton–Raphson optimization algorithm described by Bond et al. (1998), generalized to include polarization. In this algorithm one iterates towards the maximum-likelihood spectra by means of a local quadratic approximation to the full likelihood. The iteration scheme in its simplest form is

$$\delta C_b = \frac{1}{2} \sum_{b'} \mathcal{F}_{bb'}^{-1} \text{Tr} [(\hat{\mathbf{m}}\hat{\mathbf{m}}^T - \mathbf{C})(\mathbf{C}^{-1}\mathbf{C}_{,b'}\mathbf{C}^{-1})], \quad (5)$$

where  $b$  denotes a multipole bin,  $\mathbf{C}$  is the signal-plus-noise pixel-space covariance matrix, and  $\mathbf{C}_{,b}$  is the derivative of  $\mathbf{C}$  with respect to  $C_b$ . The signal component of  $\mathbf{C}$  is computed from the binned power spectra,  $C_b$ , and the noise component is based on the noise model described in Section 5.1.2, including diode-diode correlations. Finally,

$$\mathcal{F}_{bb'} = \frac{1}{2} \text{Tr}(\mathbf{C}^{-1}\mathbf{C}_{,b}\mathbf{C}^{-1}\mathbf{C}_{,b'}) \quad (6)$$

is the Fisher matrix. Additionally, we introduce a step length multiplier,  $\alpha$ , such that the actual step taken at

iteration  $i$  is  $\alpha \delta C_b$ , where  $0 < \alpha \leq 1$  guarantees that  $\mathbf{C}$  is positive definite. We adopt the diagonal elements of the Fisher matrix as the uncertainties on the band powers.

We start the Newton–Raphson search at  $C_\ell = 0$ , and iterate until the change in the likelihood value is lower than 0.01 times the number of free parameters, corresponding roughly to a 0.01- $\sigma$  uncertainty in the position of the multivariate peak. Typically we find that 3 to 10 iterations are required for convergence.

Estimation of cosmological parameters,  $\theta$ , is done by brute-force grid evaluation of the pixel-space likelihood,

$$\mathcal{L}(\theta) \propto \frac{-\frac{1}{2}\mathbf{d}^T\mathbf{C}^{-1}(\theta)\mathbf{d}}{\sqrt{|\mathbf{C}(\theta)|}}. \quad (7)$$

Here  $\mathbf{C}(\theta)$  is the covariance matrix evaluated with a smooth spectrum,  $C_\ell$ , parametrized by  $\theta$ . In this paper, we only consider 1-dimensional likelihoods with a parametrized spectrum of the form  $C_\ell = a C_\ell^{\text{fid}}$ ,  $a$  being a scale factor and  $C_\ell^{\text{fid}}$  a reference spectrum; the computational expense is therefore not a limiting factor. Two different cases are considered, with  $a$  being either the tensor-to-scalar ratio,  $r$ , or the amplitude of the EE spectrum,  $q$ , relative to the  $\Lambda$ CDM model.

#### 5.4. Temperature Data Selection and Analysis

As described in Section 2, we dedicate one pair of modules to differential-temperature measurements. While these modules are useful for calibration purposes, when combined with our polarization data they also enable us to make self-contained measurements of the TE and TB power spectra.

For temperature, both pipelines adopt the pipeline-A data-selection criteria used for polarization analysis (see Section 5.2). The temperature-sensitive modules, however, are far more susceptible to atmospheric contamination than the polarization modules. Thus, these cuts result in reduced efficiencies: 12.4%, 6.9%, and 6.8% for patches CMB-1, CMB-2, and CMB-3, respectively<sup>33</sup>. More tailoring of the cuts for these modules would improve efficiencies.

In pipeline A, the analysis proceeds as described in Sections 5.1.3, 5.3.1, and 5.3.2 except for two aspects. First, in the TOD processing a second-order polynomial is fit and removed from each telescope half scan instead of a linear function. This suppresses the increased contamination from atmospheric fluctuations in the temperature data. Second, we employ an iterative map maker based on the algorithm described by Wright et al. (1996). Map making for differential receivers requires that each pixel is measured at multiple array pointings or crosslinked. In order to improve crosslinking we divide the temperature data into only four maps by azimuth and deck angle, rather than the 60 divisions used for polarization analysis. To calculate TE and TB power spectra, polarization maps are made for these four divisions, plus one additional map that contains all polarization data with pointings not represented in the temperature data.

For pipeline B the algorithms for making temperature maps and estimating power spectra are identical to the polarization case, as described in Sections 5.3.3 and 5.3.4.

<sup>33</sup> Patch CMB-4 is excluded due to low data-selection efficiency and a lack of sufficient crosslinking.



## 6. ANALYSIS VALIDATION

The QUIET data analysis follows a policy of not looking at the power spectra until the analysis is validated using a set of predefined tests for possible systematic effects<sup>34</sup>. The validation tests consist of a suite of null tests, comparisons across multiple analysis configurations, and consistency checks among power spectra from different CMB patches. Data-selection criteria, filtering methods, and the division of data into maps for cross correlation in pipeline A are all evaluated based on the test results. We finalize all aspects of the data analysis including calibration and evaluation of the systematic error before unveiling the power spectra (blind analysis). The risk of experimenter bias is thereby eliminated.

Details of tests found in this section describe pipeline A. While the pipeline B analysis follows a similar program of null tests to verify the result, the increased computational requirements of the maximum-likelihood framework limit the number of tests that could be performed and require those tests to be run using lower-resolution maps than for the non-null analysis. The bulk of this section treats validation of the polarization analysis; at the end, we briefly describe the temperature analysis validation.

In a null test, the data are split into two subsets. Maps,  $m_1$  and  $m_2$ , are made from each subset. The power spectra of the difference map,  $m_{\text{diff}} \equiv (m_1 - m_2)/2$ , are analyzed for consistency with the hypothesis of zero signal. The null suite consists of 42 tests<sup>35</sup>, each targeting a possible source of signal contamination or miscalibration. These are highly independent tests; the data divisions for different null tests are correlated at only 8.8% on average. Nine tests divide the data by detector diode based on susceptibility to instrumental effects, such as instrumental polarization. Ten tests target effects that depend on the telescope pointing such as data taken at high or low elevation. Five tests divide based on the proximity of the main or sidelobe beams to known sources such as the Sun and Moon. Eight tests target residual contamination in the TOD using statistics mentioned in Section 5.2. Ten tests divide the data by environmental conditions such as ambient temperature or humidity.

Each null test yields EE and BB power spectra in nine  $\ell$  bins, calculated separately for each CMB patch. Figure 6 shows the power spectra from one null test. Although the EB spectra are also calculated for each null test, they are assigned lesser significance since sources of spurious EB power will also result in the failure of EE and BB null tests. Combining all EE and BB points for all patches and null tests in the null suite yields a total of 3006 null-spectrum points. For each power-spectrum bin  $b$ , we calculate the statistic  $\chi_{\text{null}} \equiv C_b^{\text{null}}/\sigma_b$ , where  $C_b^{\text{null}}$  is the null power and  $\sigma_b$  is the standard deviation of  $C_b^{\text{null}}$  in MC simulations. We evaluate both  $\chi_{\text{null}}$  and its square;  $\chi_{\text{null}}$  is sensitive to systematic biases in the null spectra while  $\chi_{\text{null}}^2$  is more responsive to outliers. We run MC simulations of the full null suite to take into account the small correlation among the null tests and the slight non-

Gaussianity of the  $\chi_{\text{null}}$  distribution. Non-Gaussianity is caused by the small number of modes at low  $\ell$ .




FIG. 6.— EE and BB power spectra for the patch CMB-1 null test between Q and U detector diodes. The inset shows the low- $\ell$  region in detail.

As we refine the data-selection criteria based on the results of the null suite, we use a second test to monitor changes in the non-null power spectra. Using a blind analysis framework, we compute the difference of the power spectra between any two iterations of the data selection without revealing the non-null spectra. Further, we randomize the sign of the difference to hide the direction of the change; knowledge of the direction could allow experimenter bias (e.g. a preference for low BB power). Figure 7 shows the differences in the power spectra between the final configuration and several intermediate iterations of the data selection, starting with data sets that showed significant failures for the null-test suite. Statistically significant differences indicate a change in the level of contamination in the selected data set. Our data-selection criteria are finalized when further iterations only result in statistically expected fluctuations. The sensitivity of this test is demonstrated by the fact that the expected fluctuations are much less than the statistical error of the final result.

Finally, the non-null power spectra are compared among the four CMB patches. A  $\chi^2$  statistic is computed from the deviation of each patch's non-null power spectra from the weighted average over all patches. The total  $\chi^2$  is compared to MC simulations to compute probabilities to exceed (PTE).

When all aspects of the analysis are finalized, the last round of null tests and CMB patch comparisons validates the non-null-power-spectra results. Figure 8 shows the distributions of the  $\chi_{\text{null}}$  statistic and of the PTEs corresponding to all  $\chi_{\text{null}}^2$  values from the full null suite. In pipeline A, the distribution of  $\chi_{\text{null}}$  is consistent with the expectation from MC simulations. The mean of the  $\chi_{\text{null}}$  distribution is  $0.02 \pm 0.02$ ; the mean of the MC-ensemble  $\chi_{\text{null}}$  distribution is also consistent with zero. The distribution of the  $\chi_{\text{null}}^2$  PTEs is uniform as expected. Table 4 lists the PTEs for the sums of the  $\chi_{\text{null}}^2$  statistic over all bins in each patch. Examinations of various subsets of the null suite, such as EE or BB only, do not reveal any anomalies. The EB null spectra do not indicate any failure either. Patch comparison PTEs are 0.16, 0.93, and 0.40 for EE, BB, and EB, respectively, demonstrating no statistically significant difference among the patches.

A similar, but smaller, null suite is run by pipeline B.

<sup>34</sup> Some systematic effects, such as a uniform responsivity-calibration error, cannot be detected by these techniques, and are addressed in Section 8.

<sup>35</sup> Only 41 null tests are performed for patch CMB-4; one test is dropped because there are no data in one of the subsets.




FIG. 7.— Power-spectra differences between the final data selection and six of the 32 earlier data-selection iterations, ordered by date. The lowest- $\ell$  bin of patch CMB-1 is shown. The error bars correspond to the expected fluctuations due to the differences in data selected, which are much smaller than the final statistical errors in this bin ( $\approx 0.10 \mu\text{K}^2$  for BB). Iterations that are closer to the final data selection have smaller errors. The expected EE power in this bin from the  $\Lambda\text{CDM}$  model is also shown for comparison.

Specifically, 21 null tests are made at a HEALPix resolution of  $N_{\text{side}} = 128$ . The results obtained in these calculations are summarized in the bottom panel of Figure 8, and total PTEs for each patch are listed in Table 4. As in pipeline A, no anomalous values are found.

Finally, we make a comment on the usefulness of the  $\chi_{\text{null}}$  distribution (as opposed to the  $\chi^2_{\text{null}}$  distribution) for identifying and quantifying potential contaminants. During the blind stage of the analysis, a positive bias in the  $\chi_{\text{null}}$  distribution of 0.21 (0.19) was identified using pipeline A (B) (corresponding to 21% (19%) of the statistical errors). The number from pipeline A was obtained when including auto correlations in its power-spectra estimator. When excluding auto correlations, and cross-correlating maps made from data divided by time (day by day), the bias decreased to 0.10. Further detailed studies lead to the division of data into maps based on the telescope pointing, as described in Section 5.3; the result is an elimination of the observed bias.

The maximum-likelihood technique employed by pipeline B intrinsically uses auto correlations, and a corresponding shift in the  $\chi_{\text{null}}$  distribution is seen in Figure 8. However, as will be seen in Section 7, the power spectra from the two pipelines are in excellent agreement, thereby confirming that any systematic bias coming from including auto correlations is well below the level of the statistical errors. We close this section by mentioning that we know of no other CMB experiment reporting an examination of the  $\chi_{\text{null}}$  distribution, which is sensitive to problems not detected by examining the  $\chi^2_{\text{null}}$  distribution only.

### 6.1. Validation of the Temperature Analysis

A smaller number of null tests is used for the temperature analysis. Several are not applicable and others are discarded due to lack of data with sufficient crosslinking. Even so, we are able to run suites of 29, 27, and 23 TT null tests on patches CMB-1, CMB-2, and CMB-3, respectively. We calculate the sums of  $\chi^2_{\text{null}}$  statistics, yielding PTEs of 0.26 and 0.11 for patches CMB-1 and CMB-2, respectively. No significant outliers are found for

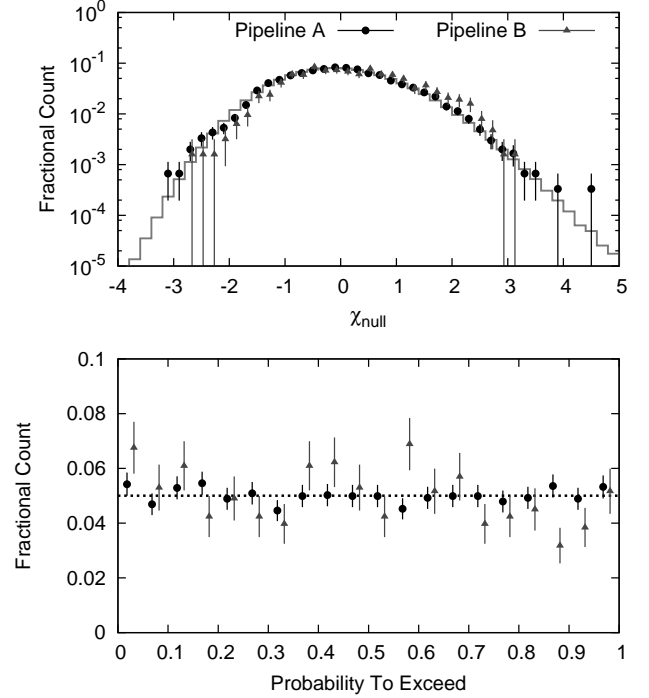


FIG. 8.— Null-Suite Statistics. The upper panel shows a histogram of the  $\chi_{\text{null}}$  values for the pipeline-A null suite (circles), pipeline-B null suite (triangles), and the average of 1024 MC realizations of the pipeline-A null suite (gray histogram). Both data and MC distributions show similar non-Gaussianity in the  $\chi_{\text{null}}$  statistic. The shift in  $\chi_{\text{null}}$  seen for pipeline B, also seen in earlier iterations of pipeline A, is discussed in the text. The lower panel shows a histogram of PTEs calculated from the  $\chi^2_{\text{null}}$  statistic (outliers from either side of the upper distribution manifest as low PTEs).

TABLE 4  
NULL SUITE PROBABILITY TO EXCEED BY PATCH

Patch	Pipeline A %	Pipeline B %
CMB-1	44	7
CMB-2	19	43
CMB-3	16	23
CMB-4	68	28

NOTE. — PTEs calculated from the sums of the  $\chi^2_{\text{null}}$  statistics, for EE and BB spectra points, over the null tests for each patch.

these patches. However, a 5- $\sigma$  outlier in a single test<sup>36</sup> is found in patch CMB-3, implying contamination in its temperature map. CMB-3 is therefore excluded from further analysis. We confirm consistency between the patches CMB-1 and CMB-2 with a PTE of 0.26.

With no significant contamination in TT, EE, or BB spectra, one may be confident that the TE and TB spectra are similarly clean. For confirmation, we calculate TE and TB null spectra for the five null tests that are common to the temperature and polarization analyses. These yield PTEs of 0.61 and 0.82 for TE, and 0.16 and 0.55 for TB, for patches CMB-1 and CMB-2, respectively, with no significant outliers. Patch consistency checks give PTEs of 0.48 for TE and 0.26 for TB. Thus, the TE and TB power spectra, as well as the TT, pass all validation tests that are performed.

<sup>36</sup> This null test divides the data based on array pointing.



## 7. RESULTS

We report results from the first season of QUIET Q-band observations: CMB power spectra, derived foreground estimates, and constraints on the tensor-to-scalar ratio,  $r$ .

## 7.1. Polarization Power Spectra

The CMB power spectra are reported in nine equally-spaced bands with  $\Delta\ell = 50$ , beginning at  $\ell_{\min} = 25$ . Given the patch size, modes with  $\ell < \ell_{\min}$  cannot be measured reliably. The correlation between neighboring bins is typically  $-0.1$ ; it becomes negligible for bins further apart.

The EE, BB, and EB polarization power spectra estimated by both pipelines are shown in Figure 9. The agreement between the results obtained by the two pipelines is excellent, and both are consistent with the  $\Lambda$ CDM concordance cosmology. Our findings and conclusions are thus fully supported by both pipelines. Only the statistical uncertainties are shown here; we treat systematic errors in Section 8. Because the systematic error analysis was only done for pipeline A, we adopt its power-spectra results (tabulated in Table 5) as the official QUIET results.

The bottom sub-panels in Figure 9 show the window and transfer functions for each bin computed by pipeline A. Figure 10 shows the maps for patch CMB-1 computed by pipeline B, and Figure 11 shows the QUIET power spectra in comparison with the most relevant experiments in our multipole range. Additional plots and data files are online<sup>37</sup>.

Fitting only a free amplitude,  $q$ , to the EE spectrum<sup>38</sup> relative to the 7-year best-fit *WMAP*  $\Lambda$ CDM spectrum (Larson et al. 2010), we find  $q = 0.87 \pm 0.10$  for pipeline A and  $q = 0.94 \pm 0.09$  for pipeline B. Taking into account the full non-Gaussian shapes of the likelihood functions, both results correspond to more than a  $10\text{-}\sigma$  detection of EE power. In particular, in the region of the first peak,  $76 \leq \ell \leq 175$ , we detect EE polarization with more than  $6\text{-}\sigma$  significance, confirming the only other detection of this peak made by BICEP at higher frequencies. The  $\chi^2$  relative to the  $\Lambda$ CDM model, with  $C_\ell^{EB} = C_\ell^{BB} = 0$ , is 31.6 (24.3) with 24 degrees of freedom, corresponding to a PTE of 14% (45%) for pipeline A (B).

## 7.2. Foreground Analysis

In order to minimize possible foreground contamination, QUIET's four CMB patches were chosen to be far from the Galactic plane and known Galactic synchrotron spurs. In these regions, contributions from thermal dust emission are negligible in Q band. Spinning dust is expected to be polarized at no more than a few percent in Q band (Battistelli et al. 2006; Lopez-Caraballo et al. 2010), so we expect the contribution to polarized foreground emission in our patches to be small. We therefore consider only two dominant sources of possible foreground contamination, namely compact radio sources and Galactic diffuse synchrotron emission.

<sup>37</sup> <http://quiet.uchicago.edu/results/index.html>

<sup>38</sup> Only  $\ell \geq 76$  are used in the EE fit and the  $\chi^2$  calculation relative to  $\Lambda$ CDM because the first EE bin has a significant foreground contribution; see Section 7.2.

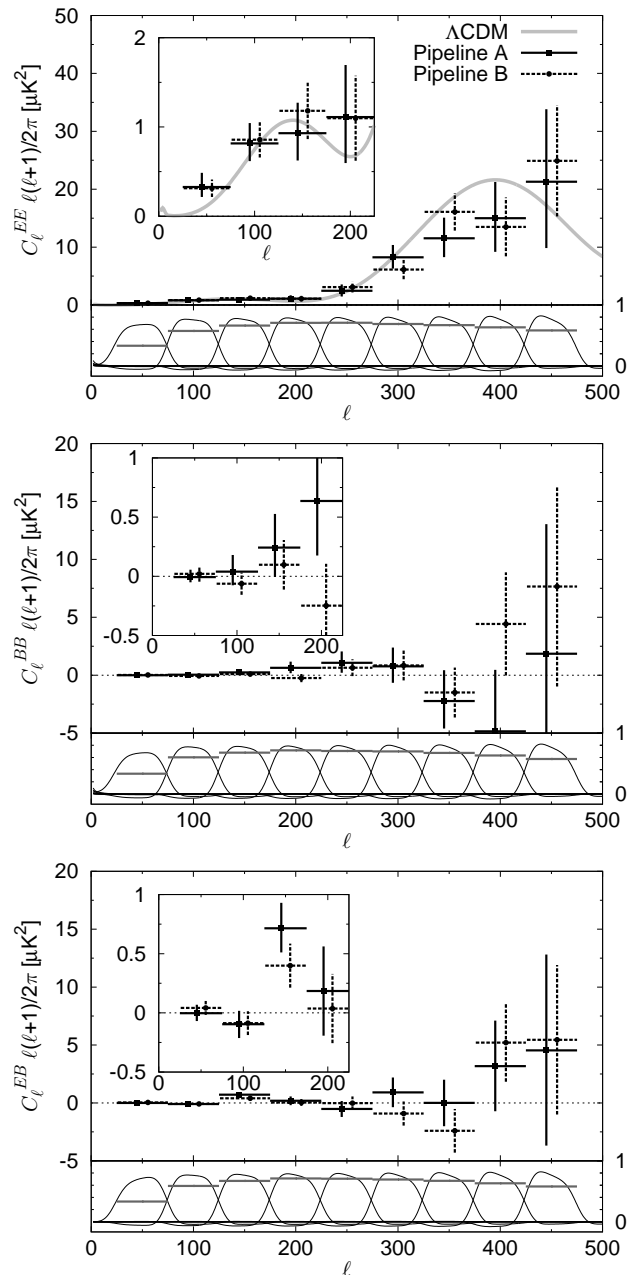


FIG. 9.— EE, BB, and EB power spectra from each QUIET pipeline, all four patches combined. The insets show the low- $\ell$  region in detail. Window and transfer functions for each  $\ell$  bin are shown below the corresponding power spectra in black and gray, respectively. The window function combines the mode-mode-coupling kernel  $M_{\ell\ell'}$  with the beam ( $B_\ell$ ) and represents, in combination with the transfer function ( $F_\ell$ ), the response in each band to the true  $C_\ell$  spectrum. The EE point in the lowest- $\ell$  bin includes foreground contamination from patch CMB-1. For this display, pipeline A shows frequentist 68% confidence intervals while pipeline B uses the diagonal elements of the Fisher matrix; the difference is most pronounced in the lowest- $\ell$  bin where the likelihood is the most non-Gaussian.

To limit the effect of compact radio sources, we apply a compact-source mask to our maps before computing the power spectra, as described in Section 5. We also evaluate the CMB spectra both with and without the full *WMAP* temperature compact-source mask (Gold et al. 2010), and find no statistically significant changes. The possible contribution from compact radio sources with

TABLE 5  
CMB-SPECTRA BAND POWERS FROM QUIET Q-BAND DATA

$\ell$ bin	EE	BB	EB
25-75	$0.33^{+0.16}_{-0.11}$	$-0.01^{+0.06}_{-0.04}$	$0.00^{+0.07}_{-0.07}$
76-125	$0.82^{+0.23}_{-0.20}$	$0.04^{+0.14}_{-0.12}$	$-0.10^{+0.11}_{-0.12}$
126-175	$0.93^{+0.34}_{-0.31}$	$0.24^{+0.28}_{-0.25}$	$0.71^{+0.22}_{-0.20}$
176-225	$1.11^{+0.58}_{-0.52}$	$0.64^{+0.53}_{-0.46}$	$0.18^{+0.38}_{-0.38}$
226-275	$2.46^{+1.10}_{-0.99}$	$1.07^{+0.98}_{-0.86}$	$-0.52^{+0.68}_{-0.69}$
276-325	$8.2^{+2.1}_{-1.9}$	$0.8^{+1.6}_{-1.4}$	$0.9^{+1.3}_{-1.3}$
326-375	$11.5^{+3.6}_{-3.3}$	$-2.2^{+2.7}_{-2.4}$	$0.0^{+2.0}_{-2.0}$
376-425	$15.0^{+6.2}_{-5.8}$	$-4.9^{+5.3}_{-4.9}$	$3.2^{+3.9}_{-3.9}$
426-475	$21^{+13}_{-11}$	$2^{+11}_{-10}$	$4.5^{+8.3}_{-8.2}$

NOTE. — Units are thermodynamic temperatures,  $\mu\text{K}^2$ , scaled as  $C_\ell \ell(\ell+1)/2\pi$ .

<sup>a</sup>Patch CMB-1 has significant foreground contamination in the first EE bin.

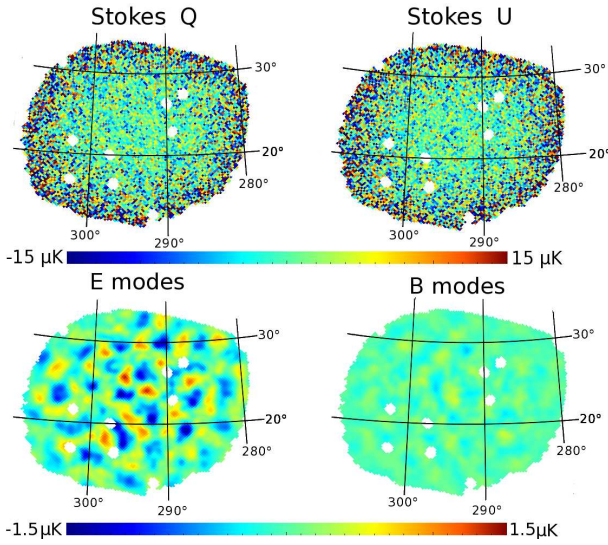


FIG. 10.— Maps of patch CMB-1 in Galactic coordinates. The top row shows our polarization maps with compact sources masked (white disks). The bottom row shows E and B modes decomposed using a generalized Wiener filter technique, implemented through Gibbs sampling (Eriksen et al. 2004; Larson et al. 2007), including only modes for  $\ell \geq 76$  and smoothed to  $1^\circ$  FWHM; lower multipoles are removed due to a significant foreground contribution. Note the clear difference in amplitude: the E modes show a high-signal-to-noise cosmological signal while the B modes are consistent with noise.

fluxes below the *WMAP* detection level (1 Jy) is small:  $0.003 \mu\text{K}^2$  at  $\ell = 50$  and  $0.01 \mu\text{K}^2$  at  $\ell = 100$  (Battye et al. 2010). We therefore conclude that our results are robust with respect to contamination from compact radio sources and that the dominant foreground contribution comes from diffuse synchrotron emission.

In Figure 12 we show the power spectra measured from each patch. The CMB-1 EE band power for the first bin is  $0.55 \pm 0.14 \mu\text{K}^2$ , a  $3\text{-}\sigma$  outlier relative to the expected  $\Lambda\text{CDM}$  band power of  $0.13 \mu\text{K}^2$ ; while not significant enough to spoil the overall agreement among the patches as shown in Section 6, this is a candidate for a bin with foreground contamination.

To estimate the Q-band polarized synchrotron contamination in our CMB patches, we process the *WMAP7* K-band (23-GHz) map through pipeline A and estimate

its band power,  $\hat{C}_b^{\text{KK}}$ , as well as the cross spectra with the QUIET Q-band data,  $\hat{C}_b^{\text{QK}}$ . These results are shown for the first bin ( $25 \leq \ell \leq 75$ ;  $b = 1$ ) in Table 6, together with the corresponding QUIET band powers,  $\hat{C}_b^{\text{QQ}}$ . Since foregrounds do not contribute to the sample variance, the uncertainties for  $\hat{C}_{b=1}^{\text{KK}}$  and  $\hat{C}_{b=1}^{\text{QK}}$  are given by instrumental noise only, including contributions from both *WMAP* and QUIET. For  $\hat{C}_{b=1}^{\text{QQ}}$ , sample variance as predicted by the  $\Lambda\text{CDM}$  model is also included.

There is significant EE power in patch CMB-1 as measured by  $\hat{C}_{b=1}^{\text{KK}}$ . We also find a correspondingly significant cross correlation between the *WMAP* K band and the QUIET Q band, confirming that this excess power is not due to systematic effects in either experiment and is very likely a foreground. No significant power is found in any other case. The non-detection of foreground power at  $\ell > 75$  is consistent with the expected foreground dependence:  $\propto \ell^{-2.5}$  (Carretti et al. 2010), and the low power found in  $\hat{C}_{b=1}^{\text{KK}}$ .

The excess power observed in the first EE bin of CMB-1 is fully consistent with a typical synchrotron frequency spectrum. To see this, we extrapolate  $\hat{C}_{b=1}^{\text{KK}}$  from K band to Q band, assuming a spectral index of  $\beta = -3.1$  (Dunkley et al. 2009), and calculate the expected power in  $C_{b=1}^{\text{QK}}$  and  $C_{b=1}^{\text{QQ}}$ ,

$$C_{b=1}^{\text{QK}} = \frac{1.05}{1.01} \left( \frac{43.1}{23} \right)^\beta \hat{C}_{b=1}^{\text{KK}} = 2.57 \pm 0.69 \mu\text{K}^2, \quad (8)$$

$$C_{b=1}^{\text{QQ}} = \left[ \frac{1.05}{1.01} \left( \frac{43.1}{23} \right)^\beta \right]^2 \hat{C}_{b=1}^{\text{KK}} = 0.38 \pm 0.10 \mu\text{K}^2, \quad (9)$$

where the prefactor accounts for the fact that  $\beta$  is defined in units of antenna temperature, and the uncertainties are scaled from that of  $\hat{C}_{b=1}^{\text{KK}}$ . These predictions are fully consistent with the observed values of  $\hat{C}_{b=1}^{\text{QK}}$  and  $\hat{C}_{b=1}^{\text{QQ}}$ , when combined with the  $\Lambda\text{CDM}$ -expected power. We conclude that the excess power is indeed due to synchrotron emission.

### 7.3. Constraints on Primordial B modes

We constrain the tensor-to-scalar ratio,  $r$ , using the QUIET measurement of the BB power spectrum at low multipoles ( $25 \leq \ell \leq 175$ ). Here  $r$  is defined as the ratio of the primordial-gravitational-wave amplitude to the curvature-perturbation amplitude at a scale  $k_0 = 0.002 \text{ Mpc}^{-1}$ . We then fit our measurement to a BB-spectrum template computed from the  $\Lambda\text{CDM}$  concordance parameters with  $r$  allowed to vary. For simplicity, we fix the tensor spectral index at  $n_t = 0$  in computing the template<sup>39</sup>. This choice makes the BB-power-spectrum amplitude directly proportional to  $r$ .

For pipeline A, we find  $r = 0.35^{+1.06}_{-0.87}$ , corresponding to  $r < 2.2$  at 95% confidence. Pipeline B obtains  $r = 0.52^{+0.97}_{-0.81}$ . The results are consistent; the lower panel of Figure 11 shows our limits on BB power in comparison with those from BICEP, QUaD, and *WMAP*. QUIET lies between BICEP and *WMAP* in significantly limiting  $r$

<sup>39</sup> Our definition of  $r$  agrees with Chiang et al. (2010)

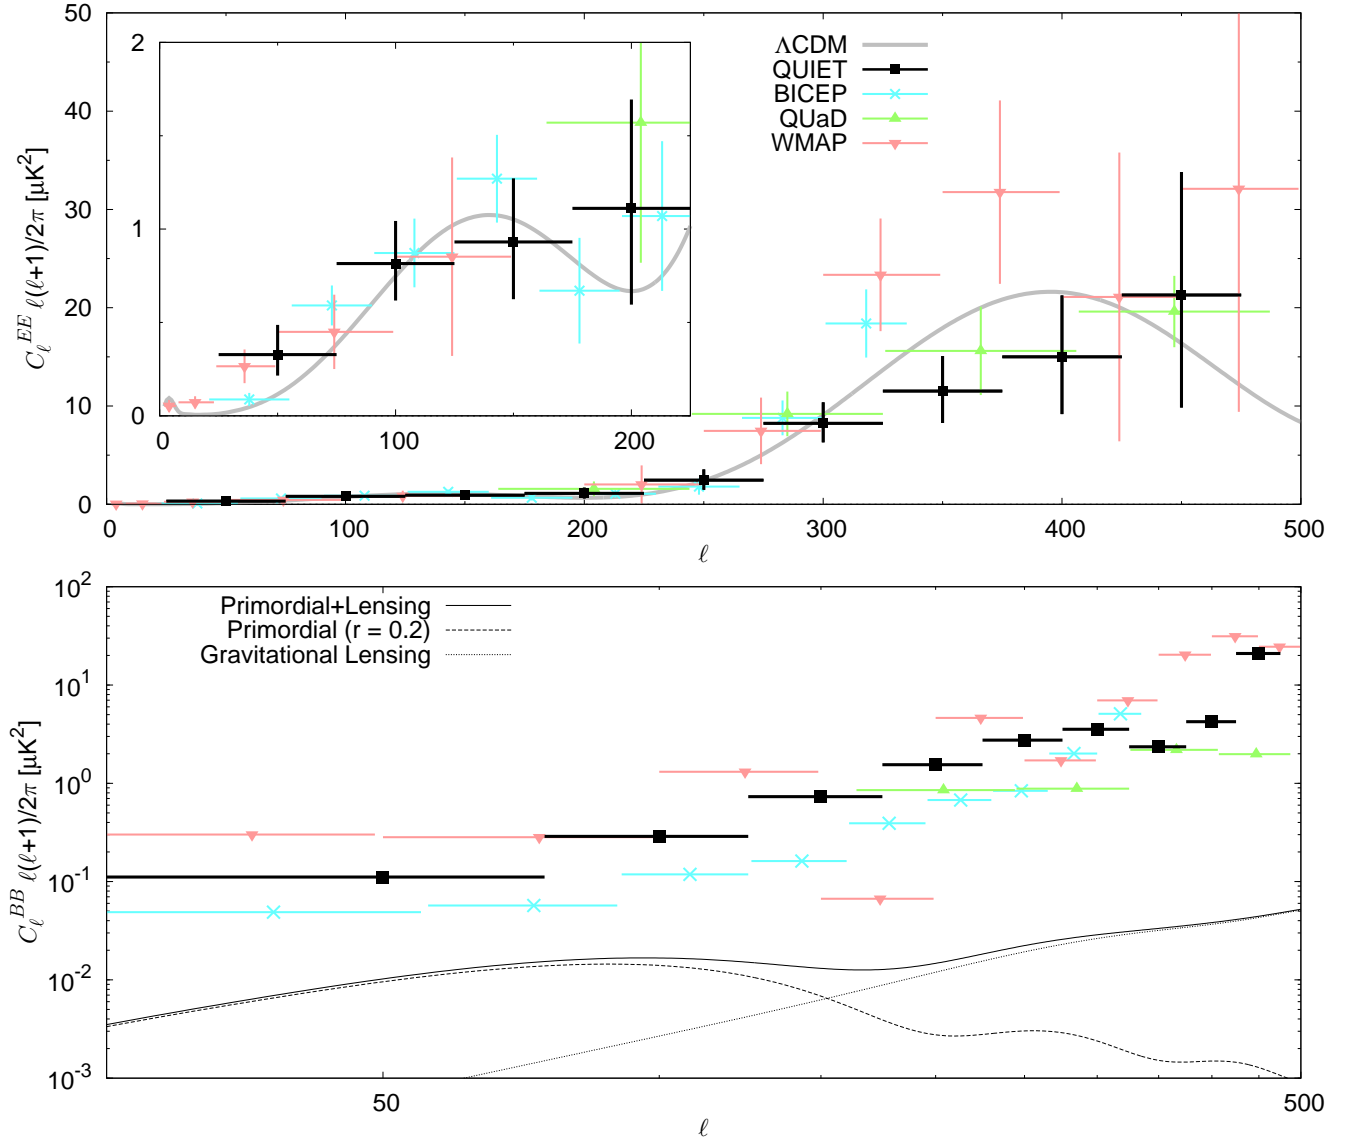


FIG. 11.— The top panel shows EE results with 68% C.L. error bars; the bottom panel shows BB 95% C.L. upper limits. For comparison, we also plot results from previous experiments (Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010) and the  $\Lambda$ CDM model (the value  $r = 0.2$  is currently the best 95% C.L. limit on tensor modes).

from measurements of CMB-B-mode power in our multipole range. Although we neither expected nor detected any BB foreground power, the detection of an EE foreground in patch CMB-1 suggests that BB foregrounds might be present at a smaller level. We emphasize that the upper limit we report is therefore conservative.

#### 7.4. Temperature Power Spectra

Figure 13 compares the QUIET and WMAP Q-band temperature maps and TT, TE, and TB power spectra. Agreement with the  $\Lambda$ CDM model is good. This is a strong demonstration of the raw sensitivity of the QUIET detectors; the single QUIET differential-temperature assembly produces a high-signal-to-noise map using only 189 hours (after selection) of observations. The high sensitivity of these modules makes them very useful for calibration, pointing estimation, and consistency checks (see Section 4).

#### 8. SYSTEMATIC ERRORS

The passing of the null suite itself limits systematic uncertainty, but to get well below the statistical errors, dedicated studies are needed. They are important in gaining confidence in the result and also in evaluating the potential of the methods and techniques we use for future efforts. We pay special attention to effects that can generate false B-mode signals. Our methodology is to simulate and then propagate calibration uncertainties (see Section 4) and other systematic effects through the entire pipeline. The systematic errors in the power spectra are shown in Figure 14. The possible contaminations are well below the statistical errors; in particular, the levels of spurious B modes are less than the signal of  $r = 0.1$ . This is the lowest level of BB contamination yet reported by any CMB experiment. This section describes how each effect in Figure 14 is determined and considers three additional possible sources of contamination.

An uncertainty not shown in Figure 14 is that aris-




FIG. 12.— CMB power spectra are shown for each patch individually. The top and bottom panels show the EE and BB spectra, respectively. The different error bars for each patch mainly reflect the amounts of time each was observed.

TABLE 6  
BAND AND CROSS POWERS FOR  $\ell = 25\text{--}75$

Patch	Spectrum	$\hat{C}_{b=1}^{\text{KK}}$	$\hat{C}_{b=1}^{\text{QK}}$	$\hat{C}_{b=1}^{\text{QQ}}$
CMB-1	EE	<b><math>17.4 \pm 4.7</math></b>	<b><math>3.30 \pm 0.55</math></b>	<b><math>0.55 \pm 0.14</math></b>
	BB	$4.8 \pm 4.5$	$0.40 \pm 0.41$	$0.06 \pm 0.08$
	EB	$-6.2 \pm 3.2$	$0.27 \pm 0.38$	$0.10 \pm 0.08$
CMB-2	EE	$5.5 \pm 3.7$	$0.01 \pm 0.56$	$0.23 \pm 0.19$
	BB	$4.6 \pm 3.4$	$0.18 \pm 0.48$	$-0.11 \pm 0.13$
	EB	$-5.5 \pm 2.8$	$-0.39 \pm 0.41$	$-0.20 \pm 0.12$
CMB-3	EE	$0.2 \pm 1.9$	$0.64 \pm 0.43$	$0.10 \pm 0.18$
	BB	$-0.3 \pm 2.6$	$0.33 \pm 0.35$	$0.01 \pm 0.13$
	EB	$1.4 \pm 1.7$	$-0.34 \pm 0.30$	<b><math>-0.27 \pm 0.11</math></b>
CMB-4	EE	$-5.2 \pm 5.1$	$0.7 \pm 1.2$	$0.65 \pm 0.58$
	BB	$-2.6 \pm 5.2$	$-0.1 \pm 1.1$	$-0.37 \pm 0.52$
	EB	$-1.0 \pm 3.9$	$0.0 \pm 0.9$	$-0.15 \pm 0.47$

NOTE. — Power-spectra estimates for the first multipole bin for each patch, computed from the *WMAP*7 K-band data and the QUIET Q-band data. The units are  $\ell(\ell+1)C_\ell/2\pi$  ( $\mu\text{K}^2$ ) in thermodynamic temperature. Uncertainties for  $\hat{C}_{b=1}^{\text{KK}}$  and  $\hat{C}_{b=1}^{\text{QK}}$  include noise only. For  $\hat{C}_{b=1}^{\text{QQ}}$  they additionally include CMB sample variance as predicted by  $\Lambda\text{CDM}$ . Values in bold are more than  $2\sigma$  away from zero.

ing from the overall responsivity error estimate of 6% (12% in power-spectra units). After including the effect of possible time-dependent responsivity variations (4%, see below), the power-spectra uncertainty is 13%. It is multiplicative, affecting all power-spectra results independent of multipole.

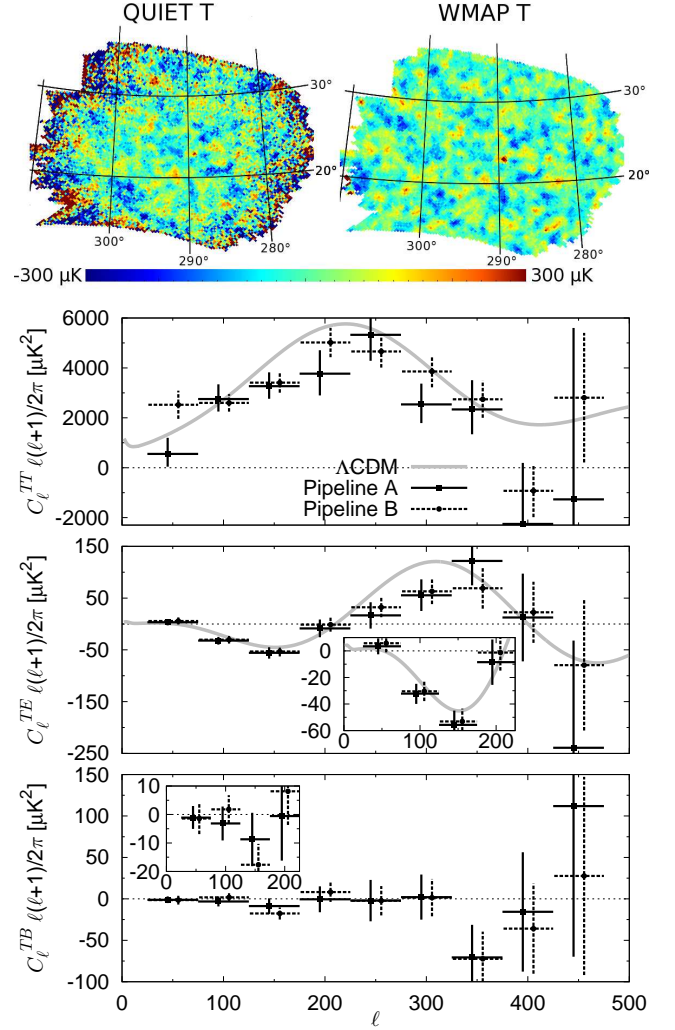


FIG. 13.— The top row compares our temperature map to the *WMAP* 7-year Q-band map (Jarosik et al. 2010) for patch CMB-1 in Galactic coordinates. Lower panels show the CMB temperature power spectra: TT, TE, and TB.

### 8.1. Beam Window Function and Pointing

The uncertainty in the beam window function is another multiplicative factor, one which increases with multipole. We estimate this uncertainty using the difference of the beam window functions measured for the central module and the modules of the differential-temperature assembly, which are at the edge of the array. The difference is statistically significant, coming from the different locations (with respect to the optics) in the focal plane; it is expected from the pre-season antenna range measurements.

Uncertainties in pointing lead to distortions in polarization maps. E power will be underestimated and spurious B power (if the distortions are non-linear) generated (Hu et al. 2003). We quantify these effects by using the differences in pointing solutions from two independent models: the fiducial model used for the analysis and an alternative model based on a different set of calibrating observations. We also modeled and included the effects of the deck-angle-encoder shift which occurred for a portion of the season (Section 4.3).



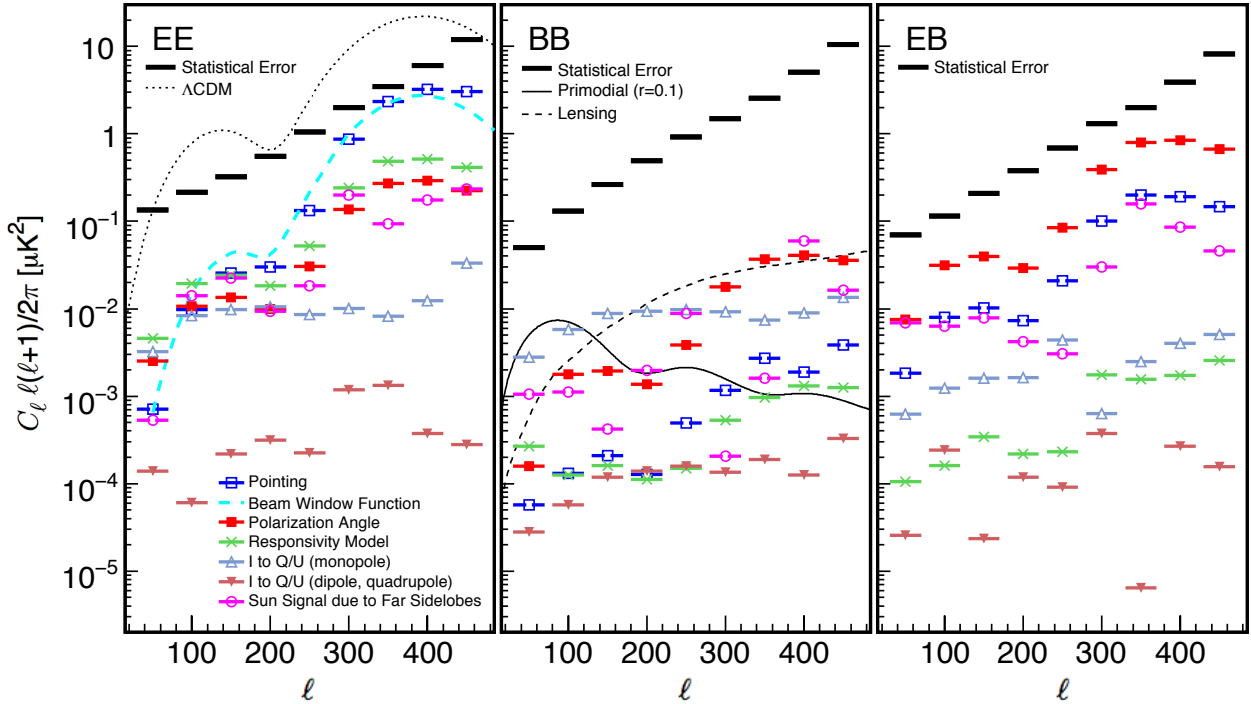


FIG. 14.— Systematic uncertainty estimates for EE, BB, and EB power spectra. Estimates for a variety of effects (see text) are shown for the three power spectra. In all cases, they are well below the statistical errors, which are also shown. In particular, the contaminations to the primordial-B-mode signal, at multipoles below 100, are below the level of  $r = 0.1$ , even though we do not make a correction for the largest contaminant, the monopole leakage.

### 8.2. Responsivity and Polarization Angle

Responsivity shifts, particularly within CESes, lead to distortions in the maps. Full-pipeline simulations quantify the shifts caused by variations in the cryostat or electronics temperatures. Similarly shifts from using responsivities determined from the Moon data, Tau A data, or from the sparse wire grid, rather than those from the sky dips, are determined. We also incorporate the uncertainty in the atmospheric-temperature model used in analyzing the sky-dip data. The largest possible effects on the power spectra are shown in Figure 14.

Uncertainties in the orientation of the polarization axes of the modules can lead to leakage between E and B modes. To quantify this leakage, we use the differences in power spectra where these angles are determined from Moon data, Tau A data, and the sparse-wire-grid data. As expected, the largest effects show up in EB power.

### 8.3. Instrumental Polarization

As described in Section 4.5, the I to Q (U) leakage coefficients for the QUIET detector diodes are small: 1% (0.2%). Except in the case of patch CMB-4, our scanning strategy significantly reduces this effect with the combination of sky and deck-angle rotation.

We estimate spurious  $Q$  and  $U$  in the maps for each CES-diode using the *WMAP* temperature map and our known leakages. Shown in Figure 14 are the estimates of spurious EE, BB, and EB powers from full-pipeline simulations, where for each realization the spurious  $Q$  and  $U$  are added to the  $Q$  and  $U$  from simulated  $\Lambda$ CDM E modes. While this method has an advantage of being able to use the real (not simulated) temperature map, it does not incorporate TE correlation, which only affects the spurious EE power. As a complement, we repeat the

study, but using simulated  $\Lambda$ CDM maps for both temperature and polarization; this only changes the estimate of spurious EE power by 30% at most. Because the spurious power is as small as it is, we have treated it as a systematic rather than correcting for it. Doing so would give us a further order of magnitude suppression.

Differing beam ellipticities can also induce higher multipole polarization signals. We measure these leakages from Tau A and Jupiter observations and find that the higher-order multipoles are at most 0.1% of the main-beam peak amplitude. The corresponding effects on the power spectra, which are seen in Figure 14, are of little concern.

### 8.4. Far Sidelobes Seeing the Sun

While we make cuts to reduce the effects of far sidelobes seeing the Sun (Sections 2 and 5.1.3), small contaminations could remain. We make full-season maps for each diode in Sun-centered coordinates and then use these maps to add contamination to full-pipeline CMB simulations. The excess power found in the simulations is taken as the systematic uncertainty.

### 8.5. Other Possible Sources of Systematic Uncertainty

Here we discuss a few additional potential sources of systematic uncertainty, which are found to be subdominant.

**Ground-Synchronous Signals.** QUIET's far sidelobes do see the ground for some diodes at particular elevations and deck angles. Ground pickup that is constant throughout a CES is removed by our TOD filters; the net effect of this filtering in the full-season maps is a correction of  $\approx 1 \mu\text{K}$ .

The only concern is ground pickup that changes over

the short span of a single CES. We find little evidence for changes even over the entire season, let alone over a single CES. We therefore conservatively place an upper limit on such changes using the statistical errors on the ground-synchronous signal. We start with the CES and module with the largest ground pickup. We then simulate one day's worth of data, inserting a ground-synchronous signal that changes by its statistical error. Given the distribution in the magnitude of the ground-synchronous signal and assuming that changes in this signal are proportional to the size of the signal itself, by considering that the signals from changing pickup add incoherently into the maps made from multiple CES-diodes at a variety of elevations and deck angles, we estimate an upper limit on residual B power from possible changing ground-pickup signals. The result is  $\lesssim 10^{-4} \mu\text{K}^2$  at multipoles below 100.

**ADC Non-linearities.** The possible residual after the correction for the non-linearity in the ADC system results in effects similar to the I to Q (or U) leakage and the variation of the responsivity during the CES. We estimate such effects based on the uncertainty in the correction parameters, confirming that there is at most a 3% additional effect for the leakage bias, and that the responsivity effect is also small, less than half of the systematic error shown for the responsivity in Figure 14.

**Data-Selection Biases.** Cuts can cause biases if they are, for example, too stringent. We expect none but to be sure we apply our selection criteria to 144 CMB + noise simulations. No bias is seen, and in particular we limit any possible spurious B modes from this source to  $\lesssim 10^{-3} \mu\text{K}^2$  at multipoles below 100.

## 9. CONCLUSIONS

QUIET detects polarization in the EE power spectrum at 43 GHz. We confirm with high significance the detection of polarization in the region of the first acoustic peak (Chiang et al. 2010) in the multipole region  $\ell = 76$ –175. We find no significant power in either BB or EB between  $\ell = 25$  and  $\ell = 475$ . We measure the tensor-to-scalar ratio to be  $r = 0.35^{+1.06}_{-0.87}$ .

These results are supported by a very extensive suite of null tests in which 42 divisions of data were used for each of 33 different cut configurations. The selection criteria and systematic errors were determined before the power spectra themselves were examined. Biases were revealed during this process, the last of which was a contamination present in the null spectra at the level of about 20% of the statistical errors, but eliminated when cross-correlating maps with differing telescope pointings. The robustness of the final results is further supported by having two pipelines with results in excellent agreement, even though one uses only cross correlations while the other also uses auto correlations.

Several possible systematic effects are studied with full end-to-end simulations. The possible contaminations in the B-mode power are thereby limited to a level smaller than for any other published experiment: below the level of  $r = 0.1$  for the primordial B modes; simply correcting for the known level of instrumental polarization would reduce this to  $r < 0.03$ . This very low level of systematic uncertainty comes from the combination of several important design features, including a new time-stream “double-demodulation” technique, Mizuguchi–Dragone

optics, natural sky rotation, and frequent deck rotation.

The correlation modules we use have a polarization sensitivity (Q and U combined) of  $280 \mu\text{K}\sqrt{\text{s}}$ , leading to an array sensitivity of  $69 \mu\text{K}\sqrt{\text{s}}$ . Further, the  $1/f$  noise observed in our detectors is small: the median knee frequency is just 5.5 mHz. One important outcome of this work, then, is the demonstration that our detectors, observing from a mid-latitude site, give excellent sensitivity and systematic immunity.

Because of our mid-latitude site, we are driven to collect data in four separate patches. While we lose some sensitivity (compared to going deeper on a single patch), there are a few advantages that we have exploited. The patches are scanned differently, in terms of time of day and the degree of crosslinking, and these differences allow some important systematic checks. Another advantage concerns foregrounds.

Foreground contamination is expected to be one of the main limiting factors in the search for primordial B modes. Indeed we report a  $3\text{-}\sigma$  detection of synchrotron emission in one of our four CMB patches, originally chosen for their expected low foreground levels. Our detection is only in EE but assuming a similar BB level and extrapolating to the foreground minimum of about 95 GHz, we would have synchrotron contamination at the level of  $r = 0.02$ . Neither *WMAP* nor *Planck* will have enough sensitivity (Tauber et al. 2010) to sufficiently constrain the polarized synchrotron amplitude at this level. In fact, our Q-band polarization maps are already as deep or deeper than what *Planck* will achieve at the same frequency. Dedicated low-frequency observations are clearly needed to achieve such constraints. When foreground cleaning becomes important, consistency among separate patches will be an important handle on our understanding.

Further progress must be made through larger arrays and longer integration times. In hand we have data collected by the 90-element W-band array with similar sensitivity to our Q-band array and more than twice the number of observing hours. Results from the analysis of that data set will be reported in future publications. A W-band receiver with the sensitivity to reach below the level of  $r = 0.01$  is under development.

Support for the QUIET instrument and operation comes through the NSF cooperative agreement AST-0506648. Support was also provided by NSF awards PHY-0355328, AST-0448909, AST-1010016, and PHY-0551142; KAKENHI 20244041, 20740158, and 21111002; PRODEX C90284; a KIPAC Enterprise grant; and by the Strategic Alliance for the Implementation of New Technologies (SAINT).

Some work was performed on the Joint Fermilab-KICP Supercomputing Cluster, supported by grants from Fermilab, the Kavli Institute for Cosmological Physics, and the University of Chicago. Some work was performed on the Titan Cluster, owned and maintained by the University of Oslo and NOTUR (the Norwegian High Performance Computing Consortium), and on the Central Computing System, owned and operated by the Computing Research Center at KEK. Portions of this work were performed at the Jet Propulsion Laboratory (JPL) and California Institute of Technology, operating under

a contract with the National Aeronautics and Space Administration. The Q-band polarimeter modules were developed using funding from the JPL R&TD program.

C.D. acknowledges an STFC Advanced Fellowship and an ERC IRG grant under FP7. P.G.F. and J.A.Z. gratefully acknowledge the support of the Beecroft Institute of Particle Astrophysics and Cosmology, the Oxford Martin School, and the Science and Technology Facilities Council. L.B., R.B., and J.M. acknowledge support from CONICYT project Basal PFB-06. A.D.M. acknowledges a Sloan foundation fellowship.

PWV measurements were provided by the Atacama Pathfinder Experiment (APEX). We thank CONICYT for granting permission to operate within the Chajnantor Scientific Preserve in Chile, and ALMA for providing site infrastructure support. We are particularly indebted to the engineers and technician who maintained and operated the telescope: José Cortés, Cristobal Jara, Freddy Muñoz, and Carlos Verdugo.

In addition, we would like to acknowledge the following people for their assistance in the instrument design, construction, commissioning, operation, and in data analysis: Augusto Gutierrez Aitken, Colin Baines, Phil Bannister, Hannah Barker, Matthew R. Becker, Alex Blein, Mircea Bogdan, April Campbell, Anushya Chandra, Sea Moon Cho, Emma Curry, Maire Daly, Richard Davis, Fritz Dejongh, Joy Didier, Greg Dooley, Hans Eide, Will Grainger, Jonathon Goh, Peter Hamlington, Takeo Higuchi, Seth Hillbrand, Christian Holler, Ben Hooberman, Kathryn D. Huff, William Imbriale, Koji Ishidoshiro, Eiichiro Komatsu, Jostein Kristiansen, Richard Lai, Erik Leitch, Kelly Lepo, Martha Malin, Mark McCulloch, Oliver Montes, David Moore, Makoto Nagai, Ian O'Dwyer, Stephen Osborne, Stephen Padin, Felipe Pedreros, Ashley Perko, Alan Robinson, Jacklyn Sanders, Dale Sanford, Mary Soria, Alex Sugarbaker, David Sutton, Matias Vidal, Liza Volkova, Edward Wollock, Stephanie Xenos, and Mark Zaskowski.

## REFERENCES

- Battistelli, E. S., Rebolo, R., Rubiño-Martín, J. A., Hildebrandt, S. R., Watson, R. A., Gutiérrez, C., & Hoyland, R. J. 2006, *ApJ*, 645, L141
- Battye, R. A., Browne, I. W. A., Peel, M. W., Jackson, N. J., & Dickinson, C. 2010, arXiv:1003.5846v1 [astro-ph.CO]
- Bischoff, C. et al. 2008, *ApJ*, 684, 771
- Bond, J. R., Jaffe, A. H., & Knox, L. 1998, *Phys. Rev. D*, 57, 2117
- Bornemann, J., & Labay, V. A. 1995, *IEEE Trans. MTT*, 43, 95
- Brown, M. L., et al. 2009, *ApJ*, 705, 978
- Buder, I. 2010, *Proc. SPIE*, 7741, 77411D
- Carretti, E., et al. 2010, *MNRAS*, 405, 1670
- Challinor, A., & Chon, G. 2005, *MNRAS*, 360, 509
- Chiang, H. C. et al. 2010, *ApJ*, 711, 1123
- Dragone, C. 1978, *Bell Syst. Tech. J.*, 57, 2663
- Dumoulin, R. N. 2010, *Proc. SPIE*, 7741, 77412N
- Dunkley, J. et al. 2009, *ApJ*, 701, 1804
- Eriksen, H. K. et al. 2004, *ApJS*, 155, 227
- Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *ApJ*, 426, 23
- Gold, B. et al. 2010, arXiv:1001.4555v2 [astro-ph.GA]
- Gorski, K. M., et al. 2005, *ApJ*, 622, 759, <http://healpix.jpl.nasa.gov/>
- Gundersen, J., & Wollack, E. 2009, *Journal of Physics: Conf. Ser.*, 155, 012005
- Güsten, R., Nyman, L. Å., Schilke, P., Menten, K., Cesarsky, C., & Booth, R. 2006, *A&A*, 454, L13
- Hafez, Y., et al. 2008, *MNRAS*, 388, 1775
- Hamimeche, S., & Lewis, A. 2008, *Phys. Rev. D*, 77, 103013
- Hansen, F. K., & Gorski, K. M. 2003, *MNRAS*, 343, 559
- Hivon, E., Gorski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Hu, W., Hedman, M., & Zaldarriaga, M. 2003, *Phys. Rev. D*, 67, 043004
- Imbriale, W. A., Gundersen, J., & Thompson, K. L. 2010, To appear in the Special Issue on Radio Astronomy in the IEEE Transactions on Antennas and Propagation
- Jarosik, N. et al. 2003, *ApJS*, 145, 413
- Jarosik, N., et al. 2010, arXiv:1001.4744v1 [astro-ph.CO]
- Kangaslahti, P. et al. 2006, in *Microwave Symposium Digest*, 2006. IEEE MTT-S International, 89–92
- Keskitalo, R. et al. 2010, *A&A*, 522, A94
- Komatsu, E., et al. 2010, arXiv:1001.4538v3 [astro-ph.CO]
- Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel, W. L. 2002, *Nature*, 420, 772
- Kusaka, A., et al. 2010, To appear in Proceedings of the 45th Rencontre de Moriond
- Larson, D. et al. 2010, arXiv:1001.4635v2 [astro-ph.CO]
- Larson, D. L., Eriksen, H. K., Wandelt, B. D., Górski, K. M., Huey, G., Jewell, J. B., & O'Dwyer, I. J. 2007, *ApJ*, 656, 653
- Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, *ApJ*, 624, 10
- Liddle, A. R., & Lyth, D. H. 2000, *Cosmological Inflation and Large-Scale Structure*
- Lopez-Caraballo, C. H., Rubino-Martin, J. A., Rebolo, R., & Genova-Santos, R. 2010, arXiv:1011.1242v1 [astro-ph.CO]
- Lueker, M., et al. 2010, *ApJ*, 719, 1045
- Mauskopf, P. D. et al. 2000, *ApJ*, 536, L59
- Mizugutch, Y., Akagawa, M., & Yokoi, H. 1976, *IEEE Antennas and Propagation Society International Symposium*, 14, 2
- Monsalve, R. A. 2010, *Proc. SPIE*, 7741, 77412M
- Montroy, T. E. et al. 2006, *ApJ*, 647, 813
- Netterfield, C. B., et al. 2002, *ApJ*, 571, 604
- Newburgh, L., et al. 2010, To appear in Proceedings of the Twelfth Marcel Grossmann Meeting on General Relativity
- Padin, S., et al. 2002, *PASP*, 114, 83
- Page, L. et al. 2007, *ApJS*, 170, 335
- Pardo, J., Cernicharo, J., & Serabyn, E. 2001, *IEEE Transactions on Antennas and Propagation*, 49, 1683
- Sievers, J. L. et al. 2007, *ApJ*, 660, 976
- Smith, K. M., & Zaldarriaga, M. 2007, *Phys. Rev. D*, 76, 043001
- Tauber, J. A., et al. 2010, *A&A*, 520, A1
- Tegmark, M. 1997, *ApJ*, 480, L87
- Wandelt, B. D., Hivon, E., & Górski, K. M. 2001, *Phys. Rev. D*, 64, 083003
- Weiland, J. L. et al. 2010, arXiv:1001.4731v1 [astro-ph.CO]
- Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, *ApJ*, 458, L53
- Wu, J. H. P. et al. 2007, *ApJ*, 665, 55





## Paper II

---

*The Q/U Imaging Experiment*  
QUIET Collaboration  
ApJ (to be submitted)

---

# THE Q/U IMAGING EXPERIMENT

QUIET COLLABORATION—C. BISCHOFF<sup>1,22</sup>, A. BRIZIUS<sup>1,2</sup>, I. BUDER<sup>1</sup>, Y. CHINONE<sup>3,4</sup>, K. CLEARY<sup>5</sup>, R. N. DUMOULIN<sup>6</sup>,  
 A. KUSAKA<sup>1</sup>, R. MONSALVE<sup>7</sup>, S. K. NÆSS<sup>8</sup>, L. B. NEWBURGH<sup>6,23</sup>, G. NIXON<sup>23</sup>, R. REEVES<sup>5</sup>, K. VANDERLINDE<sup>1,26</sup>,  
 I. K. WEHUS<sup>10</sup>, M. BOGDAN<sup>1</sup>, R. BUSTOS<sup>7,13,14</sup>, S. E. CHURCH<sup>15</sup>, C. DICKINSON<sup>16</sup>, H. K. ERIKSEN<sup>8,17</sup>, T. GAIER<sup>18</sup>,  
 J. O. GUNDERSEN<sup>7</sup>, M. HASEGAWA<sup>3</sup>, M. HAZUMI<sup>3</sup>, C. HOLLER<sup>10</sup>, K. M. HUFFENBERGER<sup>7</sup>, W. A. IMBRIALE<sup>18</sup>,  
 K. ISHIDOSHIRO<sup>3</sup>, M. E. JONES<sup>10</sup>, P. KANGASLAHTI<sup>18</sup>, D. J. KAPNER<sup>1,24</sup>, C. R. LAWRENCE<sup>18</sup>, E. M. LEITCH<sup>18</sup>, M. LIMON<sup>6</sup>,  
 J. J. MCMAHON<sup>19</sup>, A. D. MILLER<sup>6</sup>, M. NAGAI<sup>3</sup>, H. NGUYEN<sup>20</sup>, T. J. PEARSON<sup>5</sup>, L. PICCIRILLO<sup>16</sup>, S. J. E. RADFORD<sup>5</sup>,  
 A. C. S. READHEAD<sup>5</sup>, J. L. RICHARDS<sup>5</sup>, D. SAMTLEBEN<sup>2,25</sup>, M. SEIFFERT<sup>18</sup>, M. C. SHEPHERD<sup>5</sup>, S. T. STAGGS<sup>21</sup>,  
 O. TAJIMA<sup>1,3</sup>, K. L. THOMPSON<sup>15</sup>, R. WILLIAMSON<sup>6,27</sup>, E. J. WOLLACK<sup>28</sup>, J. T. L. ZWART<sup>6,29</sup>

*Submitted to ApJ—This paper should be cited as “The Q/U Imaging Experiment (2012)”*

## ABSTRACT

The Q/U Imaging Experiment (QUIET) is designed to measure polarization in the Cosmic Microwave Background, targeting the imprint of inflationary gravitational waves at large angular scales ( $1^\circ$ ). Between October 2008 and December 2010, two independent receiver arrays were deployed sequentially on a 1.4 m side-fed Dragonian telescope. The arrays use a highly compact design based on High Electron Mobility Transistors (HEMTs) that provide simultaneous measurements of the Stokes variables Q, U, and I in a single module. The 17-element Q-band polarimeter array, with a central frequency of 43 GHz, has a sensitivity of  $69 \mu\text{Ks}^{1/2}$  and an instrumental systematic error in the tensor-to-scalar ratio of  $r < 0.1$ , the highest sensitivity ever achieved in this band and the lowest systematic error reported so far. The 84-element W-band polarimeter array has a sensitivity of  $78 \mu\text{Ks}^{1/2}$  at a central frequency of 95 GHz. The two arrays cover multipoles in the range  $\ell \approx 25 - 950$ . These are the largest HEMT-based arrays deployed to date. This article describes the design, calibration, performance, and imperfections that can lead to systematic errors for the telescope and instrument.

*Subject headings:* cosmology: cosmic microwave background — cosmology: observations

## 1. INTRODUCTION

The Cosmic Microwave Background (CMB) is a powerful probe of early universe physics. Measurements of the temperature anisotropy power spectrum are critical in establishing the concordance  $\Lambda$ CDM model (e.g. Liddle & Lyth 2000, and references therein), and measurements of CMB polarization currently provide the best prospects for confirming inflation or constraining the level of the primordial gravitational wave background. The CMB is polarized via Thomson scattering off temperature anisotropies. The curl-free component of the polarization field (E-mode polarization) is generated by the same density inhomogeneities responsible for the

<sup>1</sup> Kavli Institute for Cosmological Physics, Department of Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

<sup>2</sup> Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

<sup>3</sup> High Energy Accelerator Research Organization (KEK), 1-1 Oho, Tsukuba, Ibaraki 305-0801, Japan

<sup>4</sup> Astronomical Institute, Graduate School of Science, Tohoku University, Aramaki, Aoba, Sendai 980-8578, Japan

<sup>5</sup> Cahill Center for Astronomy and Astrophysics, California Institute of Technology, 1200 E. California Blvd M/C 249-17, Pasadena, CA 91125, USA

<sup>6</sup> Department of Physics and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

<sup>7</sup> Department of Physics, University of Miami, 1320 Campo Sano Drive, Coral Gables, FL 33146, USA

<sup>8</sup> Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

<sup>9</sup> Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway

<sup>10</sup> Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

<sup>11</sup> Oxford Martin School, 34 Broad Street, Oxford OX1 3BD, UK

<sup>12</sup> Department of Physics and Astronomy, University College London, Gower Street, London WC1E, UK

<sup>13</sup> Departamento de Astronomía, Universidad de Chile, Casilla 36-D, Santiago, Chile

<sup>14</sup> Departamento de Astronomía, Universidad de Concepción, Casilla 160-C, Concepción, Chile

<sup>15</sup> Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Varian Physics Building, 382 Via Pueblo Mall, Stanford, CA 94305, USA

<sup>16</sup> Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK

<sup>17</sup> Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway

<sup>18</sup> Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

<sup>19</sup> Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA

<sup>20</sup> Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

<sup>21</sup> Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA; send correspondence to L. Newburgh, [newburgh@princeton.edu](mailto:newburgh@princeton.edu)

<sup>22</sup> Current address: Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 43, Cambridge, MA 02138, USA

<sup>23</sup> Current address: Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

<sup>24</sup> Current address: Micro Encoder Inc., Kirkland, WA 98034, USA

<sup>25</sup> Current address: Nikhef, Science Park, Amsterdam, The Netherlands

<sup>26</sup> Current address: Department of Physics, McGill University, 3600 Rue University, Montreal, Quebec H3A 2T8, Canada

<sup>27</sup> Current address: Kavli Institute for Cosmological Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

<sup>28</sup> Goddard Space Flight Center, Greenbelt, MD 20771, USA

<sup>29</sup> Current address: Physics Department, University of the Western Cape, Private Bag X17, Bellville 7535, South Africa

measured temperature anisotropy. However a measurement of the E-mode polarization can break degeneracies in cosmological parameters inherent to measurements of the temperature anisotropy spectrum alone. A divergence-free component of the polarization field (B-mode polarization) is generated by lensing of E-mode polarization by intervening large-scale structure along the line-of-sight. It can be used to probe structure formation in the early universe. A large class of inflationary models exists for which each predicts a spectrum of gravity waves generated during inflation which could produce a measureable B-mode amplitude expected to peak around angular multipole of  $\ell \approx 100$  (Seljak & Zaldarriaga 1997; Kamionkowski et al. 1997; Dodelson et al. 2009). The detection of these B-modes, parameterized by the tensor-to-scalar ratio  $r$ , would provide a measurement of the energy scale of inflation. An additional contribution to both E-mode and B-mode polarization spectra is expected from polarized foreground emission. Understanding the spectral dependence and spatial distribution of foregrounds is critical for pushing the limits of B-mode polarization detection or constraint.

The goal of detecting or placing competitive constraints on the inflationary B-mode CMB polarization signature led us to optimize QUIET<sup>30</sup> for both sensitivity and control of systematic errors. We demodulate the signal at two phase-switching rates (“double demodulation”) to reduce both the  $1/f$  noise and instrumental systematic effects. In addition, our scan strategy, consisting of constant elevation scans performed between regular elevation steps, frequent boresight rotations, and natural sky rotation reduces systematic errors. Using arrays with two widely separated central frequencies having good atmospheric transmission allows us to separate a cosmological signal from galactic foreground signals.

This paper describes the QUIET instrument, designed to measure the CMB polarization and the synchrotron foreground. QUIET deployed two arrays of 19 and 90 HEMT-based coherent detector assemblies in the Chajnantor plateau in the Atacama Desert of Northern Chile. The extreme aridity of this region results in excellent observing conditions for most of the year (Radford & Holdaway 1998). The arrays operate at central frequencies of 43 GHz (Q-band) and 95 GHz (W-band) respectively and are the largest HEMT-based arrays used to date. In the focal plane, each assembly contains passive waveguide components and a module, a small interchangeable HEMT-based electronics package. Within these two arrays, 17 (84) of the Q-band (W-band) assemblies are polarimeters, each measuring simultaneously the Q, U, and I Stokes parameters. The remaining 2(6) assemblies measure the CMB temperature anisotropy (“differential-temperature assemblies”). The Q-band (W-band) assemblies are cooled to  $\approx 20\text{K}$  ( $\approx 26\text{K}$ ) in a cryostat and placed at the focus of a 1.4 m side-fed Dragonian telescope enclosed in an absorbing ground screen. The resulting full width at half maximum (FWHM) angular resolution is  $27.3'$  ( $12.2'$ ) for each Q-band (W-band) assembly.

<sup>30</sup> Bruce Winstein, who died in 2011 February soon after observations were completed, was the principal investigator for QUIET. His intellectual and scientific guidance were crucial to the experiment’s success.

Tables 1 and 2 lists the salient characteristics of the QUIET experiment. Figures 1(a) and 1(b) show views of the receiver, telescope, and electronics enclosure. The following sections describe the optics, observing site and strategy, cryogenics and the optical window properties, polarimeter and differential-temperature assemblies, electronics, and calibration tools. Finally, we present a detailed description of the performance of both receivers.

## 2. OBSERVING SITE AND STRATEGY

Observations (Table 3) were performed at the Chajnantor plateau at 5080 m altitude in the Atacama Desert of northern Chile ( $67^\circ 45' 42''\text{W}$   $23^\circ 1' 42''\text{S}$ ). Atmospheric conditions were monitored using data from a 183 GHz line radiometer sited at the APEX telescope (Güsten et al. 2006), located  $\sim 2.5\text{km}$  away from the QUIET site. Typical atmospheric optical depths in our observing bands over all scanning elevations at Chajnantor are 0.02–0.1 (Figure 2). The median precipitable water vapor (PWV) was 1.2 mm (0.94 mm) during the Q-band (W-band) observing season. We found for the Q-band data set we kept 82% of the data below the median PWV, and 59% of the data above the median PWV. For the W-band, we kept 75% below the median PWV and 54% below.

We employed a fixed-elevation, azimuth-scanning technique: a  $\sim 15^\circ \times 15^\circ$  field (the fields are given in Table 2) was scanned in azimuth as it drifted through the  $\sim 7^\circ$  field-of-view. These constant elevation scans (“CES”) typically lasted  $\sim 1.2$ – $1.5$  hours. The telescope then re-tracked the field center and began another CES. By scanning at constant elevation for a given scan, we observed through a constant column density of atmosphere so that only weather variations within a scan contributed an atmospheric signal. Most calibration sources were observed at constant elevation, but occasionally we employed raster scans, changing elevation between azimuth slews to more rapidly observe a calibration source.


The infrastructure and three-axis driving mount previously used for the CBI experiment (Padin et al. 2002) was refurbished for QUIET, in part to enable rapid azimuth scanning. The mount control software is an augmented version of the CBI control system. The principal

TABLE 1  
INSTRUMENT OVERVIEW


Band	Q	W
Frequency (GHz)	43	95
# of Polarization Assemblies	17	84
# of Temperature Anisotropy Assemblies	2	6
FWHM Angular Resolution (arcmin)	27.3	12.2
$\ell$ range	$\approx 25$ – $475$	$\approx 25$ – $950$
Instrument Sensitivity ( $\mu\text{K}\sqrt{\text{s}}$ )	69	78
Telescope (common to Q and W)	Side-fed Dragonian	
Main Reflector Diameter (mm)	1400	
Secondary Reflector Diameter (mm)	1400	

TABLE 2  
FIELDS OF OBSERVATION

Field Centers (J2000 RA, Dec)	$12^{\text{h}} 04^{\text{m}} -39^\circ$ $05^{\text{h}} 12^{\text{m}} -39^\circ$ $00^{\text{h}} 48^{\text{m}} -48^\circ$ $22^{\text{h}} 44^{\text{m}} -36^\circ$
Field Size	$\approx 15^\circ \times 15^\circ$



(a)



(b)

FIG. 1.— *a*: The QUIET instrument before placement upon mount, showing the electronics enclosure, cryostat, and reflectors. *b*: The mounted instrument shown within an absorbing ground screen.

TABLE 3  
PARTITION OF THE QUIET Q-BAND AND W-BAND SEASONS BY OBSERVATION TYPE, WITH NO DATA CUTS. ‘OTHER’ INCLUDES DATA TAKEN DURING ENGINEERING TESTS, ABORTED SCANS, ETC.

Band	Q	W
Season start	2008 Oct 24	2009 August 12
Season end	2009 June 13	2010 December 22
Total Observing Hours	3458	7493
CMB Observing (%)	77	72
Galactic Observing (%)	12	14
Calibration (%)	7	13
Other (%)	4	1

modifications include the addition of support for rapid scanning of the azimuth axis of the mount and for monitoring and archiving of data from the receiver. This software consists of a central control and data collection program, a graphical user interface program, a real-time computer running the VxWorks<sup>2</sup> operating system to control the telescope mount, and a real-time computer running Linux to control the receiver. The mount was operated by a queue of non-interactive observing scripts written in a custom control language. The modifications that were made were aimed at supporting high scanning accelerations without overwhelming the counter-torque in the anti-backlash system of the azimuth drive. Track-

<sup>2</sup> www.windriver.com




FIG. 2.— Zenith optical depth for typical atmospheric conditions at the Chajnantor plateau (left scale) and representative QUIET module bandpass responses (right scale). The atmospheric spectrum is calculated with the ATM model from Pardo et al. (2001).

ing accuracy is sacrificed for high scanning speeds and accelerations. However, accurate pointing information can be reconstructed during the data analysis from frequent readouts of the axis encoders and a dynamic model of the mechanical response of the mount. To facilitate this, the CBI control system was also modified to acquire encoder readouts at 100 Hz. The modified control system supports scans with coasting speeds of up to 6°/s and turnaround accelerations of up to 1.5°/s<sup>2</sup>. The accuracy of the encoder readout timestamps is ~0.5 ms resulting in a worst-case pointing error of ~8′. The achieved pointing accuracy of the entire system is 3′.5 rms for the Q-band array (QUIET Collaboration et al. 2011) and 0.2′ rms for the W-band array (Collaboration 2012). We achieved a mean azimuthal scan speed of ~5°/s. The azimuthal speed on the sky is elevation dependent, it corresponds to ~2°/s on the sky, yielding azimuth scan frequencies of 45–100 mHz. As each 15°×15° observing field rises, its azimuthal extent with respect to the fixed telescope mount increases. Avoiding scanning through the azimuth limit (~440°) of the telescope leads to an effective upper elevation limit of 75°. The lower limit of the elevation range of the mount is 43°.

In addition to the azimuth and elevation axes, the mount provides a third rotation axis through the bore-sight. We rotate this boresight angle (‘deck angle’) once per week in order to separate the polarization on the sky from that induced by systematic errors such as leakage from temperature to polarization.

### 3. OPTICS

The QUIET optical chain consists of a classical side-fed Dragonian antenna (Dragone 1978) coupled to a platelet array of diffusion-bonded corrugated feed horns cooled to  $\simeq 25$  K inside a cryostat. The outputs of these optical elements are directed into the polarimeter and differential-temperature assemblies described in Sections 5.1 and 5.3, respectively. The main reflector (MR) and sub-reflector (SR) as well as the aperture of the cryostat are enclosed by an ambient temperature ( $\simeq 273$  K), absorbing ground shield. The design and characterization of the telescope, feed horns and ground shield are described in Sections 3.1, 3.2, and 3.3, respectively. The optical performance, as measured by the main beam, the sidelobes and the instrumental polarization, is described in Sec-

tions 3.4, 3.5, and 3.6, respectively.

### 3.1. Telescope

The telescope design requirements include: a wide field of view, excellent polarization characteristics with minimal beam distortion and minimal instrumental polarization, and minimal spillover and sidelobes that could otherwise generate spurious polarization. The latter requirements have often been met by CMB experiments by using either classical, dual offset Cassegrain antennas (e.g. Barkats et al. (2005)), Gregorian antennas (e.g. Meinhold et al. (1993)), or shaped reflectors (e.g. Page et al. (2003b)). QUIET is the first CMB polarization experiment to take advantage of the wide field of view enabled by a classical Dragonian antenna (Imbriale et al. 2011). An additional advantage of the classical Dragonian antenna is that it satisfies the Mizuguchi condition (Mizuguchi et al. 1976) which, when combined with the very low cross-polar characteristics of the conical corrugated feed horns, yields very low antenna contribution to the instrumental polarization. As pointed out by Chang & Prata (2004), a classical Dragonian antenna affords two natural geometries, a front-fed design and a side-fed (or crossed) design. QUIET uses the side-fed design because it allows for the use of a larger cryostat, and hence focal plane array, without obstructing the beam.

#### 3.1.1. Telescope Design

The design of the reflectors follows the procedure outlined by Chang & Prata (2004), and it is augmented with physical optics program (Imbriale & Hodges 1991) to predict beam patterns. This procedure relies on the specification of the first five design parameters given in Table 4 and shown in Figure 3. Once these parameters are specified, a number of other useful parameters can be computed, including the MR focal length and the SR eccentricity. The actual MR circular diameter was decreased slightly to 1400 mm, as noted in Table 4. Similarly, the actual SR circular diameter was increased slightly, also to 1400 mm, and this resulted in an increased value of the actual SR edge angle given by  $20^\circ$  in Table 4. The oversized SR reduces feed spillover for the horns on the edge of the array. The design values (not the actual values) shown in the top half of Table 4 were used to establish the calculated values shown in the lower half of Table 4.

#### 3.1.2. Telescope Fabrication and Alignment

The telescope consists of two reflectors, the receiver cryostat (Figure 1(a)) and the structure that supports them (the ‘sled’). The reflectors are each made of a solid piece of 6061-T6 aluminum, light-weighted on the reverse side leaving narrow ribs on a triangular grid, and attached with adjustable hexapod struts to the sled. The sled in turn is mounted on a deck structure, which also supports the ground shield, the receiver electronics enclosure, the telescope drive crates, the uninterruptable power supply, and the expanded steel walkways. The deck is attached directly to the deck bearing.

After the fabrication of the reflectors and sled, the telescope was assembled and pre-aligned using a MetricVision MV200 laser radar. This system was used to measure both the reflector surfaces as well as the absolute

TABLE 4  
TELESCOPE DESIGN PARAMETERS ARE LISTED WHERE NEGATIVE ANGLES ARE MEASURED CLOCKWISE WITH RESPECT TO THE VERTICAL AXIS.

description, parameter	design/actual value
MR circular aperture diameter, $D$	1470/1400 mm
SR edge $\angle$ , $\theta_e$	$17^\circ/20^\circ$
MR-SR separation, $\ell$	1270 mm
MR offset $\angle$ , $\theta_0$	$-53^\circ$
$\angle$ between MR and horn axes, $\theta_p$	$-90^\circ$
calculated value	
MR focal length, $F$	4904.1 mm
SR eccentricity, $e$	2.244
$\angle$ between SR and MR axis, $\beta$	$-63.37^\circ$
SR interfocal distance, $2c$	6516.1 mm
MR offset distance, $d_0$	4890.2 mm




FIG. 3.— This scaled schematic of the QUIET side-fed Dragonian antenna shows a number of the useful design parameters. Table 4 provides a description of each parameter and either the design value or the calculated value.

positions of tooling balls on the perimeter of each reflector once the reflectors are aligned to the focal plane. The root-mean-square (rms) deviations from the MR and SR design surfaces are  $38 \mu\text{m}$  and  $28 \mu\text{m}$ , respectively, once a small fraction ( $< 1\%$ ) of the outlier measurements from the perimeter of each reflector are removed.

In order to align the reflectors after assembling them on the site, an animated 3-D model of the telescope was constructed which also accounts for additional tooling balls on the cryostat face. These tooling balls have well measured displacements from the platelet array. Using the model, a transformation matrix was established that mapped turnbuckle adjustments to tooling ball displacements for each reflector. After assembly at the site, the distances between the tooling balls were measured with a custom-built vernier caliper with a range of 2.4 m. The transformation matrix was inverted and applied to the tooling ball displacements in order to establish the proper turnbuckle adjustments. The turnbuckles were then adjusted to bring the system into alignment. This method enabled convergence to an aligned state after just three iterations. The 17 measurements used to establish the position of the secondary with respect to the cryostat

TABLE 5  
PLATELET ARRAY DESIGN PARAMETERS

Frequency (band/GHz)	# of Feeds	L×W×H (mm × mm × mm)	Mass (kg)	Aperture Diameter (mm)	Throat Diameter (mm)	# of Grooves	Horn Separation (mm)	Semi-flare Angle (degrees)
Q/39–47	19	281.7×427.3×370.1	43.7	71.78	6.69	104	76.20	7.6°
W/89–100	91	129.1×427.8×370.5	20.6	31.62	2.97	103	35.56	7.6°

(for both the Q- and W-band systems) yielded an rms error of  $< 400 \mu\text{m}$  when compared to the ideal positioning. Similarly the 14 measurements used to establish the position of the primary with respect to the secondary yielded an rms error of  $< 500 \mu\text{m}$  when compared to the ideal positioning as established using the laser radar. Tolerance studies allowing for comparable displacements show that there is minimal impact on the optical performance.

### 3.2. Feed Horns

The requirements for the feed horns include high beam symmetry, efficiency, gain and bandwidth, as well as low sidelobes and cross-polarization. These requirements are satisfied by conical, corrugated feed horns (Kay 1962) (Clarricoats & Olver 1984). Standard production techniques for corrugated feed horns (e.g. computer-numerically-controlled lathe machining and electroforming) are prohibitively costly for the large number of feeds for the W-band array. A lower-cost option is described in the next subsection.

#### 3.2.1. Platelet Array Design

A 91-element W-band and a 19-element Q-band platelet array of hexagonally-packed, conical, corrugated feed horns were designed for QUIET (Gundersen & Wolack 2009), (Imbriale et al. 2011). Each array is machined from aluminum 6061-T6 and consists of a number of thin platelets each with a single corrugation, a number of thick plates each with multiple corrugations, and a base plate. The assembly of platelets and plates is then diffusion bonded together. Table 5 provides many of the parameters of each array.

Due to the side-fed geometry of the telescope, the feed horns must have relatively high gain ( $\simeq 27$  to  $28$  dB) in order to provide a low edge taper of  $< -30$  dB for both the Q and W-band systems. This dictates the aperture size of the feed horns and hence the horn-to-horn spacing. For the W-band horns, this spacing is commensurate with the size of the modules. Most of the dimensions of the Q-band horns are scaled by the ratio of the frequencies ( $\sim 90/40 = 2.25$ ) which results in a Q-band horn spacing that is larger than the W-band modules. These horn spacings give rise to angular separations of  $1.75^\circ$  ( $0.82^\circ$ ) between adjacent beams in the Q(W) systems and result in fields of view of  $7.0^\circ$  and  $8.2^\circ$  for the Q and W systems, respectively.

The number of corrugations is fixed at three per wavelength for each horn and a semi-flare angle of  $7.6^\circ$  is chosen using a design procedure (Hoppe 1987), (Hoppe 1988) that ensures both acceptable cross-polar levels and return loss. This optimization procedure also adjusts the depth of the first six corrugations of each horn in order

to reduce the reflection coefficient to better than  $-32$  dB over the full anticipated band of operation.

#### 3.2.2. Platelet Array Testing

A vector network analyzer (VNA) was used to measure the return loss of each horn in each array. Each measurement consisted of attaching one horn in a platelet array to one port of the VNA using a commercially available circular-to-rectangular transition. A sheet of microwave absorber was placed at  $45^\circ$  in front of the horns at a distance of  $\simeq 1$  m. The return losses for five of the 19 Q-band horns are shown in Figure 4 and are similar for the W-band feed horns. Maximum reflection strengths (negative return loss) are listed in Table 6. For comparison, electroformed horns that are identical in design to the Q and W-band horns were fabricated. The array values in Table 6 are comparable to but not quite as good as the electroformed horns or the theoretical predictions both of which were  $< -30$  dB across the band.

Beam patterns were measured for all 91 horns in the W-band array and 13 out of 19 horns in the Q-band array. A synthesizer combined with  $\times 3$  and  $\times 6$  multipliers generated the source signals at 40 and 90 GHz respectively. A standard gain horn was used as a source antenna. The platelet arrays were mounted on an azimuth-elevation mount so that the source was in the far-field of the platelet array horns. The source signals were modulated at 1 kHz and a lock-in amplifier connected to a detector diode on the platelet array detected the signal. A coaligned, alignment laser ensured that the source horn and platelet array horn were parallel and axially aligned to each other. A digital protractor with an accuracy of  $\pm 0.05^\circ$  ensured that the source and receiver horns' polarization axes were coincident for the copolar patterns or perpendicular to each other for the crosspolar patterns. Several measurements were made on each horn including E- and H-plane copolar patterns as well as their corresponding crosspolar patterns. The patterns were taken by keeping the source horn static and rotating the platelet array horn in azimuth about a vertical axis that intersected the horns phase center.

The beam patterns of typical Q-band and W-band horns are shown in Figure 5. This figure shows both E- and H-plane copolar patterns as well as crosspolar patterns for the platelet feeds and for an electroformed feed with identical design parameters. The figure also shows the theoretical model responses. In all cases the E- and H-plane copolar patterns are consistent with both the model and the electroformed feed measurements out to the  $-30$  dB level. Upper limits of  $-33$  dB are set on the E-plane crosspolar levels. The H-plane crosspolar patterns are not in as good agreement with the model, which predicts both E- and H-plane crosspolar levels at the  $< -40$  dB level. The largest discrepancies are similar



TABLE 6  
MEASURED PLATELET ARRAY PERFORMANCE

Frequency (band/GHz)	FWHM (deg)	Gain (dB)	Crosspol E/H (dB)	Reflection Strength (dB)	Insertion Loss (dB)
Q/39–47	8.3–6.9	27.2–28.5	$< -34 / -29$	$< -25$	$< -0.1$
W/89–100	8.3–7.4	27.1–28.0	$< -31 / -29$	$< -24$	$< -0.1$




FIG. 4.— Return loss measurements for five of the 19 Q-band horns.




FIG. 5.— Beam pattern measurements of a typical Q-band (W-band) horn in each platelet array along with an electroformed equivalent horn are shown in the top (bottom) two figures. The left-hand subfigures show the E-plane results and the right-hand subfigures show the H-plane results. The solid line in each case shows the theoretical prediction of the copolar responses. The theoretical predictions of the crosspolar responses are all below -40 dB and are not shown. Upper limits of -33 dB are placed on the E-plane crosspolar responses, while the H-plane crosspolar responses are measured at the -30 to -33 dB level for both Q and W-band platelet array horns as well as for their electroformed equivalents.

in shape to the Q-band H-plane crosspolar measurements shown in Figure 5 and have a non-null crosspolar boresight response. This type of response is typical of angular misalignment between the source and receiver probes. The level of the response is consistent with the precision of the digital level. The W-band H-plane crosspolar response does have a null on boresight and is likely the true crosspolar response. Given that the platelet arrays' crosspolar responses are consistently higher than the corresponding electroformed horns' responses suggests that either the machining or the diffusion bonding process leads to somewhat compromised performance. However, none of the measured feeds has crosspolar levels  $> -29$  dB. Table 6 summarizes the results of the beam pattern

measurements.

Upper limits on the insertion loss were obtained during the return loss measurements of both the W-band and Q-band platelet arrays by placing a flat aluminum plate in front of the horn and generating an effective short. In both cases the measured reflection strength allows one to set a lower bound on the feeds' room temperature transmission efficiency of  $> 99\%$ . Assuming ohmic losses, this transmission efficiency is expected to increase to  $> 99.5\%$  upon cooling to 25 K. This is due to the fact that the ratio of the electrical resistivity for aluminum 6061-T6 operating at 25 K to the resistivity at 293 K is 0.35 (Clark et al. 1970), and this results in reduced ohmic losses.

### 3.3. Ground Shield

The side-fed Dragone design minimizes but does not eliminate sidelobe power. Simulations show that a number of sidelobes are expected. The performance of the ground shields is described in detail in Section 3.5. In order to minimize the radiation from the ground and from celestial sources entering the receiver through sidelobes, an absorbing, comoving ground shield is employed. This shields the instrument from varying ground and Sun pick-up and provides a stable, essentially unpolarized emission source that does not vary during a telescope scan. The ground shield structure (Figure 1(b)) consists of two parts: the lower ground shield is an aluminum box that encloses both reflectors and the front half of the cryostat; the upper ground shield (UGS) is a cylindrical tube that attaches to the lower ground shield directly above the MR. The external surface of the ground shield is coated in white paint in order to reduce diurnal temperature variations and to minimize radiative loading. The interior of the ground shield is coated with a broadband absorber<sup>2</sup> that absorbs radiation and re-emits it at a constant temperature, allowing the ground shield to function as an approximately constant Rayleigh-Jeans source in both Q and W bands. The UGS was not in place for the Q-band measurement. It was installed in January 2010, approximately a third of the way through the W-band measurements.

### 3.4. Main Beam Performance

The main beam profiles are primarily determined from observations of Jupiter. Additional observations of Tau A are performed to check the main lobe response, to measure the polarized responsivity, to determine the polarization angles and to characterize instrumental polarization. Tau A and Jupiter are used for main beam charac-

<sup>2</sup> The absorber is Emerson Cumming HR-10 (www.eccosorb.com) and is covered with Volara. The Volara is transparent at our observing frequencies and acts as weather-proofing.




FIG. 6.— Normalized beam maps of Jupiter are shown on the left for representative differential-temperature assemblies for the Q- and W-band systems with contours at 0.8, 0.5, and 0.2. The corresponding azimuthally-averaged beam profiles for each map are shown on the right in comparison with the theoretical prediction (solid line). Similar maps and profiles of TauA were measured using the polarimeter assemblies but at a reduced signal-to-noise.

terization since they are the brightest polarized and unpolarized, compact sources in the sky, respectively. Figure 6 shows beam patterns of Jupiter for a differential-temperature assembly in each of the Q and W-band arrays. These measurements are consistent with the lower signal-to-noise main beam profiles measured using Tau A once the slightly different instrumental bandpasses, source spectra, and positions in the focal plane are taken into account. The main beam is used to compute the solid angle  $\Omega_B$ , the main beam forward gain,  $G_m = \frac{4\pi}{\Omega_B}$ , and the telescope sensitivity,

$$\Gamma = \frac{10^{-26} c^2}{2k_B \nu_e^2 \Omega_B} \text{ KJy}^{-1} \quad (1)$$

in terms of the effective ,

$$\nu_e = \frac{\int \nu f(\nu) \sigma(\nu) d\nu}{\int f(\nu) \sigma(\nu) d\nu} \quad (2)$$

for a given instrumental bandpass  $f(\nu)$  and source spectrum  $\sigma(\nu)$ . The source spectra of Tau A and Jupiter are based on the WMAP data (Weiland et al. 2011). A Tau A source spectrum with  $\sigma \propto \nu^{-0.302}$  is assumed for the calculation of the effective frequency for the Tau A measurements. An empirical fit to WMAP's measurements of Jupiter's brightness temperatures yields a source spectrum of the form

$$\sigma(\nu_{\text{GHz}}) = \frac{2k_B \nu_{\text{GHz}}^2}{c^2} (96.98 + 2.175 \nu_{\text{GHz}} - 2.219 \times 10^{-2} \nu_{\text{GHz}}^2 + 8.217 \times 10^{-5} \nu_{\text{GHz}}^3). \quad (3)$$

Similarly a source spectrum of the form

$$\sigma \propto \nu^4 \exp(h\nu/k_B T_{\text{CMB}}) / (\exp(h\nu/k_B T_{\text{CMB}}) - 1)^2 \quad (4)$$

is used to compute the effective frequency for unresolved CMB fluctuations. Equations 1 and 2 explicitly

assume  $G_m \propto \nu^2$ .

Table 7 provides a summary of the mean values of these quantities for the Q-band and W-band polarization and total-power modules for a source spectrum of the form 4. The Q-band total power values are for the lone Q-band differential-temperature assembly, while the Q-band polarization values are for the central pixel which is typical for the array. Both the W-band total power and polarization values shown in Table 7 are averaged over the respective differential-temperature and polarization array elements using an inverse-variance weighting.

The shape of the main beam and its uncertainties are used to compute the instrumental window function and its associated uncertainties (Monsalve 2010). Initially, an arbitrarily oriented, 2-D, elliptical gaussian beam is fit to the data shown in Figure 6. If  $\sigma_a$  and  $\sigma_b$  represent the beam widths of the semi-major and semi-minor axes of the elliptical gaussian (with  $\sigma_a \geq \sigma_b$ ), then the elongation is defined by  $\epsilon = \frac{\sigma_a - \sigma_b}{(\sigma_a + \sigma_b)}$ . Typical elongations were found to be  $< 0.02$  and averaged about 0.01. This low elongation, and the fact that the CMB scans use a combination of natural sky rotation and deck angle rotation, implies that the beams are well described by an axially-symmetric beam. As described by Monsalve (2010), a Hermite expansion is applied to the symmetrized beam from which the transfer function and covariance matrix may be calculated (Page et al. 2003a).

### 3.5. Sidelobe Characterization

Two different methods are used to measure sidelobes. These included pre-deployment antenna range measurements and in-situ measurements of a bright, near-field source. In addition, unintentional measurements of the sun in the sidelobes also enabled their characterization. These three measurements and their results are discussed in more detail here.

#### 3.5.1. Antenna Range Measurements of Sidelobes

The telescope was installed on the Jet Propulsion Laboratory's Mesa Antenna Measurement Facility for measurements of both the main lobe and far sidelobes at both 40 and 90 GHz. The telescope was mounted on an elevation-over-azimuth positioner with 4'' pointing accuracy. Individual electroformed versions of the Q and W-band horns, described in Section 3.2.2, were used for the range measurements. The range measurements were conducted before the ground shields were fabricated, so the sidelobe results are only appropriate for the telescope in its bare configuration. The measurements made use of the facility's Scientific Atlanta model 1797 heterodyne receiver system which enabled repeatable measurements down to  $-90$  dB of the peak power level. A combination of a source synthesizer, multiplier and amplifier was used to generate  $\simeq 100 - 200$  mW of power at each frequency. The sources were separately connected to corrugated feeds at the focus of a small cassegrain antenna at a distance of 914 m from the telescope. Due to limitations of the mount, only a simple principal plane cut within  $\pm 90^\circ$  of the telescope boresight (in the plane shown in Figure 3) was performed for a number of arrangements of the source/receiver antennas. These arrangements included moving the receiver horn to a few positions in the focal plane and rotating the source and receiver horns for



TABLE 7  
MEAN EFFECTIVE FREQUENCIES, BEAM SIZES, MAIN BEAM SOLID ANGLES, MAIN BEAM FORWARD GAINS AND TELESCOPE SENSITIVITIES FOR BOTH THE POLARIZATION (SUBSCRIPT P) AND TOTAL POWER (SUBSCRIPT T) ASSEMBLIES.

	$\nu_e$ (GHz)	FWHM (deg)	$\Omega_B$ ( $\mu$ sr)	$G_m$ (dBi)	$\Gamma$ ( $\mu$ K/Jy)
Q <sub>P</sub>	43.0	0.455	74.3	52.3	237
Q <sub>T</sub>	43.4	0.456	78.0	52.1	222
W <sub>P</sub>	94.4	0.195	13.6	59.6	269
W <sub>T</sub>	95.7	0.204	15.6	59.1	228




FIG. 7.— Results from the antenna range measurements. The top measurements are the 40 GHz E-plane results for a horn located in the top row, 20.46 mm above the central horn. The bottom measurements are the 90 GHz H-plane results for a horn located in the top row, 23.87 mm above the central horn. The gap in the measurements from boresight angles of  $+1.5^\circ$  to  $+8.5^\circ$  is due to mount-related elevation angle limitations. The two most prominent far sidelobes are the triple reflection sidelobe and the SR spillover lobe as indicated in each figure. The optical paths associated with these lobes are shown in Figure 9. Top row horns, such as these, are most susceptible to each of these lobes due to their location in the focal plane.

both E- and H-plane cuts.

The results for one feed horn position for each of the Q- and W-band arrays are shown in Figure 7. In each case the feed horn position that was tested corresponds to the top row of the respective platelet array, furthest from the MR and directly above the central feed horn. Cross polar measurements were not made on the antenna range since they are made during routine calibrations. The main lobe beamsizes compare well with initial theoretical predictions (Imbriale et al. 2011); however, the

near-in (i.e. within  $\pm 5^\circ$  of the main lobe) sidelobe levels do not. As described by Imbriale et al. (2011), this is due to the reflector surface imperfections, which were not included in the initial theoretical predictions. As shown in Figure 18 from Imbriale et al. (2011), once the measured reflector surface is incorporated in the theoretical pattern predictions, the predicted envelope of near-in sidelobes matches well with the observations. The surface imperfections caused the near-in sidelobe levels to increase by as much as 15 dB in some regions. The two dominant far sidelobes are the SR spillover lobe and the ‘triple reflection’ lobe. The SR spillover lobe is broad and arises from direct coupling into the feed horn. It is located  $\sim 70^\circ$  from boresight as predicted by Imbriale et al. (2011). The triple reflection lobe is due to an additional reflection off the SR (as indicated in Figure 9) and it is located  $\simeq 50^\circ$  from boresight in the opposite direction from the SR spillover lobe. This position also matches the prediction shown in Figure 10 of Imbriale et al. (2011). The amplitude of each lobe for the W-band case is  $\simeq -61$  dB, while they are both  $\simeq -57$  dB for the Q-band measurement. These amplitudes are both 5–7 dB above the uncorrected predictions of Imbriale et al. (2011). As with the increased near-in sidelobe levels, this increase in the far sidelobes can be attributed to the reflector surface imperfections.

### 3.5.2. Source Measurements of Sidelobes

The performance of the UGS was assessed using the W-band array in 2010 January. For these measurements, a polarized, modulated 92 GHz oscillator was placed in the near field of the telescope at a distance of approximately 15 m. The telescope was scanned over its entire azimuth and elevation range at four different deck angles ( $0^\circ$ ,  $90^\circ$ ,  $-90^\circ$ ,  $-180^\circ$ ). The top and middle panels in Figure 8 show measurements before and after the installation of the UGS, respectively. The main sidelobe feature at the bottom of the top map corresponds to the line-of-sight over the secondary. This feature is clearly removed by the UGS. The remaining sidelobes were tracked down to holes in the floor of the lower ground shield below the SR. A third measurement taken after placing absorber over these holes (bottom panel in Figure 8) verifies this and displays the sidelobe performance in the final ground shield configuration. The UGS was neither in place during the full Q-band observing season nor during the first third of the W-band observing season. Data with the moon or sun in the sidelobes were excised during the Q-band analysis as a result (QUIET Collaboration et al. 2011).

### 3.5.3. Sun Measurements of Sidelobes

Before the installation of the UGS, the sun is occasionally detected in the sidelobes. This is particularly apparent once the data are binned into maps in ‘telescope boresight-centered’ coordinates. The cartesian basis of this coordinate system has  $\hat{i}$  oriented along the feed horn boresight,  $\hat{k}$  oriented along the telescope boresight,  $\hat{j} = \hat{k} \times \hat{i}$ , and  $\hat{s}$  directed toward the sun. The corresponding spherical coordinates of the sun are defined to be  $\theta = \cos^{-1}(\hat{s} \cdot \hat{k})$ , and  $\phi = \tan^{-1}(\frac{\hat{s} \cdot \hat{j}}{\hat{s} \cdot \hat{i}})$ . Figure 9-(a) shows the the optical path of these sidelobes before the installation of the UGS along with the cartesian basis




FIG. 8.— Sidelobe measurements for W-band module 40, located on the edge of the array), with the deck angle set at  $-180$  and the near-field source located at an azimuth of  $\simeq 20^\circ$  and an elevation of  $\simeq -5^\circ$ . *top*: Measurements with only the lower ground shield. The lobe seen at the bottom of the map is from spillover past the secondary. This lobe is removed after the installation of the UGS. *middle*: Measurements with the lower ground shield and UGS installed. The lobe at the top is due to holes in the absorber from the ground shield structure, and is present before the UGS was added as well, but its position has shifted slightly because the source was moved between measurements. *bottom*: Results with the complete ground shield installed and with additional absorber placed over holes in the floor of the lower ground shield. The color scale is the same between all three measurements and has been normalized to match the antenna range measurements. The UGS reduces the far sidelobes by at least an additional 20 dB below the levels shown in Figure 7.

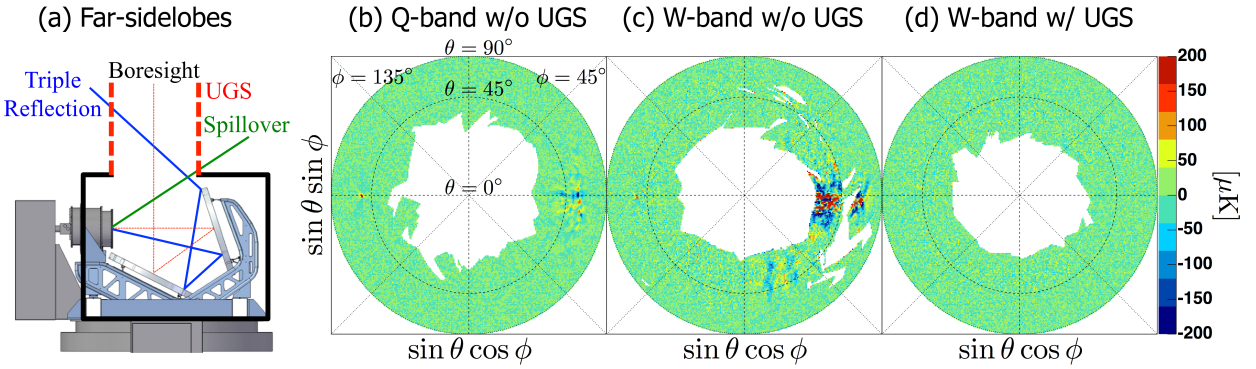


FIG. 9.— Sidelobe characterization using the sun. (a) The optical paths that give rise to the triple reflection and spillover sidelobes before the installation of the UGS are shown as dashed lines. (b) The ‘telescope boresight-centered’ map of the sun before the installation of the UGS for a Q-band feed horn in the top row, nearest to the vertical centerline. The sharp spike induced by the triple reflection is seen at  $(\theta, \phi) \simeq (50^\circ, 180^\circ)$ , while the large-area of sidelobe contamination just under the  $\phi = 0^\circ$  line is induced by the SR spillover. (c) The ‘telescope boresight-centered’ map for a horn in a similar position in the W-band array before the UGS installation. (d) The same map for the same W-band horn after the UGS installation.

of the ‘telescope boresight-centered’ coordinates. Figure 9-(b) shows the ‘telescope boresight-centered’ map for a feed horn on the top row of the Q-band array that is closest to the vertical centerline of the platelet array. The direction of the triple reflection far sidelobe is similar among feed horns. However, the direction of the spillover far sidelobe is different among feed horns because it couples directly to the feed horns and not through the reflectors. Therefore the far sidelobe response is characterized for each feed horn separately. We also measured the far sidelobes for W-band before and after the UGS installation (Figure 9-(c, d)). Figure 9 (d) confirms that both far sidelobes are eliminated by the UGS. The  $\phi = 0^\circ - \phi = 180^\circ$  line in Figure 9 corresponds to the principal plane measurement shown in Figure 7, and both show the SR spillover lobe and triple reflection lobe before the installation of the UGS. The amplitudes of the two far sidelobes measured with the sun are consistent with the  $\sim -60$  dB levels obtained

with the range measurements shown in Figure 7. The addition of the UGS for the W-band data, in combination with azimuth filtering and data rejection used for the Q-band data, makes the spurious polarization signal due to sidelobes a negligible effect on the B-mode measurements.

### 3.6. Leakage Beams

The leakage beams quantify both the Q and U diodes’ responses to an unpolarized source, as well as the leakage that can convert a sky Q into a measured U or a sky U into a measured Q. In order to assess these various forms of leakage, a large number of observations were performed of both Jupiter and Tau A. These produce beam maps that are subsequently decomposed into their respective beam Mueller fields as given by O’Dea et al. (2007). The beam Mueller fields are related to the co- and cross-polar components of the dual, orthogonal polarizations supported by the feed system.




FIG. 10.— The extracted Mueller fields are shown for a Q and U diode pair behind the central horn of the W-band array. An  $\simeq 0.4\%$  quadrupole term is evident in the  $m_{QI}$  and  $m_{UI}$  leakage beams, while no higher order structure is evident in the  $m_{QU}$  or  $m_{UQ}$  leakage beams at the  $\simeq 0.1\%$  level. As described in Section 3.6, the monopole contribution to the  $m_{QU}$  and  $m_{UQ}$  leakage beams can be absorbed into the detector angle which is measured during the calibration procedure. Similar results for the Q-band system are given in Monsalve (2010).

For a linearly polarized source with Stokes parameters ( $I_{\text{src}}, Q_{\text{src}}, U_{\text{src}}, V_{\text{src}} = 0$ ), degree of linear polarization  $p = \sqrt{Q_{\text{src}}^2 + U_{\text{src}}^2} / I_{\text{src}}$ , and position angle  $\gamma_{\text{PA}} = \frac{1}{2} \tan^{-1}(-U_{\text{src}}/Q_{\text{src}})$ , the output voltage  $d_Q$  of a Q diode<sup>2</sup> as a function of instrumental flux density gain  $g_Q$  and instrumental position angle  $\psi$  is given by

$$d_Q = g_Q \cdot I_{\text{src}} \cdot \{m_{QI} + p \cdot m_{QQ} \cdot \cos(2[\gamma_{\text{PA}} - \psi]) + p \cdot m_{QU} \cdot \sin(2[\gamma_{\text{PA}} - \psi])\}, \quad (5)$$

where  $m_{QI}$  and  $m_{QU}$  are the Mueller fields representing the I-to-Q and U-to-Q leakage beams and  $m_{QQ}$  is the extracted Q polarization beam. Similarly, the output voltage of a U diode is given by

$$d_U = g_U \cdot I_{\text{src}} \cdot \{m_{UI} + p \cdot m_{UU} \cdot \sin(2[\gamma_{\text{PA}} - \psi]) + p \cdot m_{UQ} \cdot \cos(2[\gamma_{\text{PA}} - \psi])\}, \quad (6)$$

where  $m_{UI}$  and  $m_{UQ}$  are the corresponding leakage beams, and  $m_{UU}$  is the U polarization beam. In each of these expressions, the factor  $g$  is the product of the receiver responsivity  $R$  and the telescope sensitivity  $\Gamma$ . The instrumental position angle is given by  $\psi = \eta + \phi_d$  where  $\eta$  is the parallactic angle of the beam center and  $\phi_d$  is the deck angle.<sup>3</sup> For a number of sources, Tau A

<sup>2</sup> The diode nomenclature is described in Section 5.2.

<sup>3</sup> For reference, when  $\phi_d = 0^\circ$  or  $\phi_d = 180^\circ$ , the line connecting the feed horn centers of the differential temperature assemblies is parallel to the ground. In the event that the parallactic angle of a given beam is similarly zero (so that the beam is observing the local meridian), then the plane of the feed horn's septum polarizer is perpendicular to the local meridian, yielding an instrumental position angle  $\psi = 0^\circ$ .

in particular, the parallactic angle coverage is not very large, so beam maps at various deck angles are necessary in order to vary the outputs of the Q and U detector diodes. Daily beam maps of Jupiter and/or Tau A are used to extract the various leakage beams as well as the polarized beams. Figure 10 shows the results of this extraction for a Q and U diode pair behind the central W-band horn. A similar figure is shown in (Monsalve 2010) for the Q-band system.

The  $m_{QI}$  and  $m_{UI}$  Mueller fields are of particular importance since they characterize the instrumental polarization. Instrumental polarization can be generated by any of the elements in the optical path including the reflectors, the curved cryostat window, the IR blocker, the feed horns, the septum polarizers and the modules themselves. In the Appendix, specific expressions are derived for these leakage terms for the modules and the septum polarizers. These two elements are the primary cause of the monopole leakage contribution to the  $m_{QI}$  and  $m_{UI}$  Mueller fields. The median W-band monopole leakage is 0.25% and is lower than the Q-band monopole leakage. As reported in QUIET Collaboration et al. (2011), the Q-band monopole leakage is the largest systematic error in the B-mode measurement at  $\ell \sim 100$  where it begins to dominate the constraint on  $r$  at levels of  $r < 0.1$ . A naive estimate of the impact of this leakage would cause it to dominate at a much higher level, however a combination of sky rotation and frequent boresight rotation suppresses this systematic by two orders of magnitude. The cause of the Q-band monopole leakage is described in more detail in section 5.1.

The monopole leakage refers to the  $s_{00}$  term in



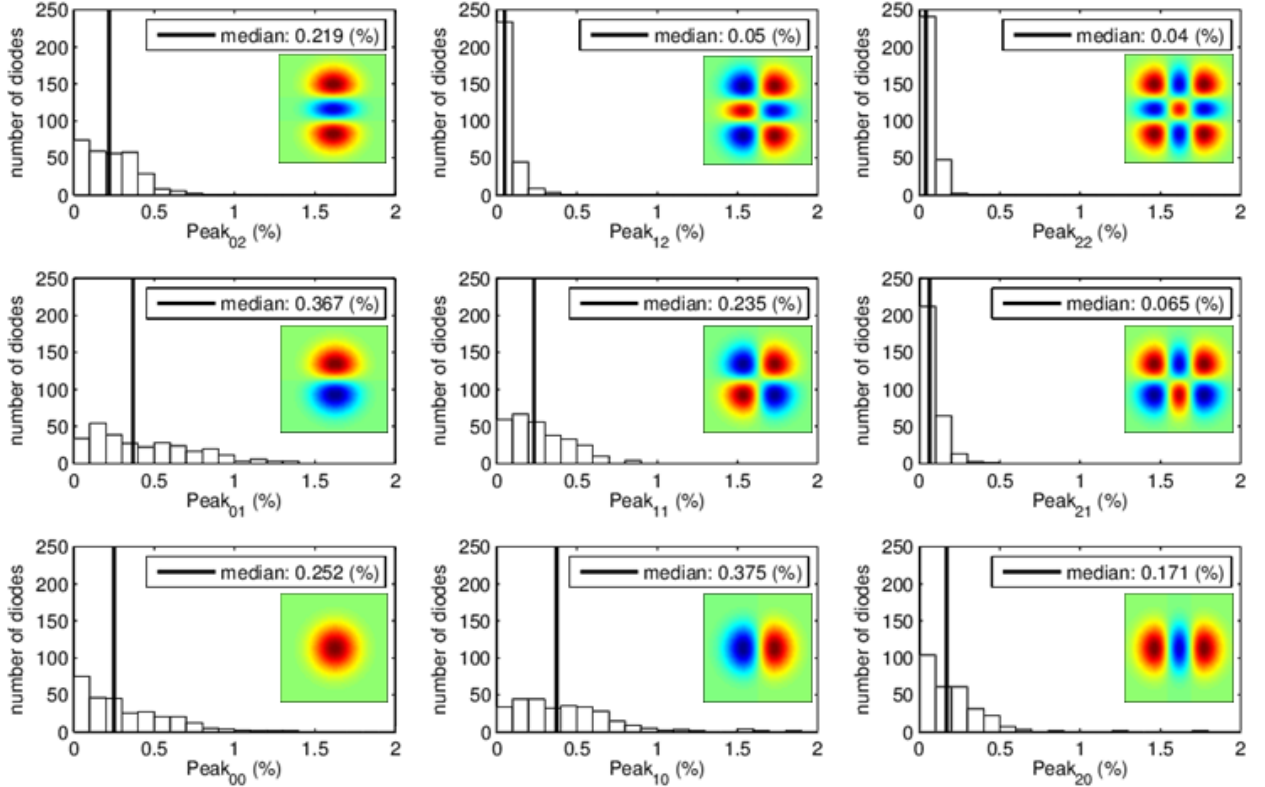


FIG. 11.— These histograms show the number of W-band diodes that have a maximum absolute value of the product  $|s_{ij}f_{ij}|$  (denoted  $\text{Peak}_{ij}$  on the  $x$ -axis) in a given percentile range for both the  $m_{QI}$  and  $m_{UI}$  leakage beams. The Hermite expansion term is also shown in each panel. A median value of all detector diodes is provided in each histogram and indicated with a vertical line. Similar results for the Q-band system are given in Monsalve (2010).

the Gauss-Hermite expansion of these leakage beams  $b_{\text{leak}}(x, y)$  (Monsalve 2010). Here and in Figure 10 the coordinates ( $x = \sin \theta \sin \phi, y = \sin \theta \cos \phi$ ) are ‘telescope boresight-centered’ coordinates defined in Section 3.5.3. The leakage beams can be expressed as

$$b_{\text{leak}}(x, y) = \sum_{j=0}^2 \sum_{i=0}^2 s_{ij} \cdot f_{ij}(x, y), \quad (7)$$

where  $s_{ij}$  are the fit coefficients and the normalized basis functions  $f_{ij}(x, y)$  are

$$f_{ij}(x, y) = \left( \frac{1}{\sqrt{2^{i+j} i! j! \pi \sigma^2}} \right) e^{-\frac{1}{2\sigma^2} [x^2 + y^2]} H_i \left( \frac{x}{\sigma} \right) H_j \left( \frac{y}{\sigma} \right), \quad (8)$$

where  $\sigma$  is the gaussian width of the symmetrized beam described in Section 3.4 and the  $H_i$  and  $H_j$  are Hermite polynomials.

Higher order leakage terms, including dipole ( $s_{01}$  or  $s_{10}$ ) and quadrupole leakages ( $s_{11}$  or  $(s_{20} - s_{02})/2$ ), can also arise due to the off-axis nature of the telescope and the imperfectly matched E and H-plane feed horn patterns. The full array drift scans of Jupiter are particularly useful in measuring these quantities for every diode in the W-band array. Histograms of the peak amplitudes complete to  $i = j = 2$  are shown in Figure 11 for the W-band array. Similar results are provided for the central pixel of the Q-band array in Monsalve (2010). Additional terms in the expansion are also included, but they are consistently less than 0.1%. Leakages above 1%

are quite rare and typical values are in the 0.2 – 0.4% range. The W-band dipole and quadrupole leakages are typically slightly higher than those in Q-band. The systematic effects that these leakage beams generate for power spectrum estimation are provided for the Q-band results (QUIET Collaboration et al. 2011), and similarly they will be provided in the W-band analysis paper (in preparation).

The  $m_{UQ}$  and  $m_{QU}$  Mueller fields measure the leakage of the incident Q Stokes parameter into the measured U Stokes parameter or the incident U Stokes parameter into the measured Q Stokes parameter. Curved reflector surfaces, imperfections in the septum polarizer, and imperfections in the phase switch are potential sources of this leakage. These primarily give rise to monopole leakage and effectively rotate the instrumental position angle. In the case that the ratios  $\frac{m_{QU}}{m_{UU}}$  and  $\frac{m_{UQ}}{m_{UU}}$  are constant over the extent of the beam, the  $m_{UQ}$  and  $m_{QU}$  Mueller fields can be absorbed into the expressions for the two diode outputs with the definition of detector angles  $\psi_Q$  and  $\psi_U$ . The detector angles are defined by replacing the last two terms in each of equations 5 and 6 with a single term as follows:

$$p \cdot m_{QQ} \cdot \cos(2[\gamma_{\text{PA}} - \psi - \psi_Q]) \equiv p \cdot m_{QQ} \cdot \cos(2[\gamma_{\text{PA}} - \psi]) + p \cdot m_{QU} \cdot \sin(2[\gamma_{\text{PA}} - \psi]) \quad (9)$$

and

$$\begin{aligned}
p \cdot m_{UU} \cdot \sin(2[\gamma_{PA} - \psi - \psi_U]) \equiv \\
p \cdot m_{UU} \cdot \sin(2[\gamma_{PA} - \psi]) + p \cdot m_{UQ} \cdot \cos(2[\gamma_{PA} - \psi]),
\end{aligned}
\tag{10}$$

respectively. A Hermite decomposition of the  $m_{QU}$  and  $m_{UQ}$  Mueller fields shown in Figure 10 shows that they are simply related by a multiplicative factor to the  $m_{QQ}$  and  $m_{UU}$  fields. Thus they can be represented in terms of single valued detector angles,  $\psi_Q$  and  $\psi_U$  and are not a source of systematic error. An important feature of the resulting detector angles is that they are separated by nearly integer multiples of  $45^\circ$  for each of the four diodes in a given module. This is shown to be the case in Section 8.5.

#### 4. CRYOSTATS

##### 4.1. Cryostat Design

The Q-band and W-band receiver arrays each have a dedicated cryostat (Figure 12). In each cryostat, cryogenic temperatures are achieved with two Gifford-McMahon dual-stage refrigerators. The first stage of the refrigerators provide cooling power to a radiation shield, maintained at  $\sim 50$  K ( $\sim 80$  K) for the Q-band (W-band) cryostat. The difference in shield temperature between the W-band and Q-band instruments wasn't anticipated from the cryostat design, but ultimately did not greatly impact the module temperatures. Infrared radiation is reduced with 10 cm thick, 3 lb density styrofoam (Table 8) attached to the top of the radiation shield. The first stages of the refrigerators also provide a thermal break for the electrical cables. The second stages of the refrigerators provide cooling power for the feedhorn array and the modules. The two stages are thermally isolated by G-10 rings.

##### 4.2. Cryostat Performance

The cryogenic performance of the Q-band array is consistent with the design goals of 20 K module temperatures and that we maintain a constant temperature during a scan to within  $\pm 0.1$  K. A temperature sensor located on an edge module in the Q-band cryostat had a mean temperature of 20.0 K with a standard deviation of 0.3 K throughout the season and a deviation of 0.02 K




FIG. 12.— The W-band cryostat with the vacuum shell and radiation shields removed.

within a scan.

For the W-band array, we expected the modules to be slightly warmer than the Q-band modules due to additional loading contributions from the active components and conduction through cabling from a factor of five more modules. We found we ultimately had to regulate the W-band modules warmer than we expected by  $\sim 3$  K, likely as a result of both higher shield temperatures and a minor vacuum leak which caused us to vacuum pump the cryostat a few times during the W-band season. A temperature sensor placed directly on the central polarimeter of the W-band array had a mean temperature of 27.4 K with a standard deviation of 1.0 K throughout the season, and a mean variation within a scan of 0.12 K. For both receiver arrays, both the variation of the module temperatures within a scan and throughout the season had a negligible impact on the responsivity (QUIET Collaboration et al. 2011).

##### 4.3. The Cryostat Window

The vacuum windows for both the Q-band and W-band cryostats are  $\sim 56$  cm in diameter, the largest vacuum window to date for any CMB experiment. The vacuum windows must be strong enough to withstand atmospheric pressure, and also should maximize transmission of signal with minimal instrumental polarization.

We chose a window material of ultra-high molecular weight polyethylene (UHMW-PE) after stress-testing a variety of window materials and thicknesses. The index of refraction was expected to be 1.52 (Lamb 1996). To make a well-matched anti-reflection coating for the UHMW-PE in the QUIET frequency bands, we coated the window with expanded teflon, which has an index of refraction of 1.2 (Benford et al. 2003). We adhered the teflon to the UHMW-PE window by placing an intermediate layer of low-density polyethylene (LD-PE) between the teflon and the UHMW-PE. We then heated the plastics above the melting point of LD-PE while applying pressure with a clamping apparatus in a vacuum chamber, allowing us to avoid trapping air bubbles between the material layers (the window material properties are summarized in Table 8). By adding an anti-reflection coating to the window, we expected the band-averaged transmission to improve from 89.8% to 98.8% for the Q-band array and from 91.1% to 98.3% for the W-band array.

An anti-reflection coated sample for the W-band window was measured in a VNA. We fit the envelope of the transmission and reflection response to obtain values for the optical properties and material thicknesses. We computed the expected contributions to the system noise from loss using published loss tangent values (Lamb 1996): 3 K (4 K) for the Q-band (W-band) windows. These values were confirmed by placing a second window over the main receiver window and measuring the change in instrument noise.

The curvature of the window under vacuum pressure could introduce cross-polarization by presenting a variable material thickness to the incoming radiation. A physical optics analysis of the 90 GHz window was performed with the General Reflector Antenna Analysis<sup>6</sup> package to investigate the effect of the curved surface

<sup>6</sup> <http://www.ticra.com>

TABLE 8  
THICKNESSES OF THE WINDOW, ANTI-REFLECTION COATING MATERIAL AND IR BLOCKING STYROFOAM FOR THE W- AND Q-BAND CRYOSTAT WINDOWS. VALUES FOR THE INDEX OF REFRACTION FOR TEFLON AND UHMW-PE COME FROM THE BEST-FIT VALUES TO VNA MEASUREMENTS AT 90 GHz.

Material	Index of refraction	Thickness (mm)		Vendor
		Q-band	W-band	
UHMW-PE	1.52	9.52	6.35	McMaster-Carr
LD-PE	1.52	0.127	0.127	McMaster-Carr
Teflon	1.2	1.59	0.54	Inertech
Styrofoam	—	101.6	101.6	Clark Foam

on the transmission properties of the window. For these simulations we use a window curvature determined from measurements of the deflection of the window under vacuum,  $\sim 7.62$  cm. With a curved window, the central feed-horn has negligible instrumental polarization. The edge pixel has 0.16% additional cross-polarization, where this is defined as leakage from one linear polarization state into the other linear polarization state. This -28 dB cross-polarization is on the same order as expected cross-polarization from the horns alone and would contribute to the cross polarization coefficients  $m_{QU}$  and  $m_{UQ}$  given in Section 3.6.

#### 5. QUIET POLARIMETER AND DIFFERENTIAL-TEMPERATURE ASSEMBLIES

QUIET uses HEMT-based amplifiers with phase sensitive techniques, following the tradition of recent polarization-sensitive experiments such as PIQUE and CAPMAP (Barkats et al. 2005), COMPASS (Farese et al. 2004), WMAP (Jarosik et al. 2003a), DASI (Leitch et al. 2002), and CBI (Padin et al. 2002). Unlike those other experiments, however, QUIET uses a miniaturized design (Lawrence et al. 2004) suitable for large arrays.

The QUIET Q-band (W-band) array contains 19(90) assemblies, where 17(84) are polarization-measuring assemblies. The remaining 2(6) measure the CMB temperature anisotropy (differential-temperature assemblies) and are described in Section 5.3. At the heart of these assemblies are the modules (see Section 5.2), a highly integrated package that replaces many waveguide-block components with strip-line-coupled monolithic microwave integrated circuit (MMIC) devices containing HEMT's. The modules have a footprint of 2.5 cm  $\times$  2.5 cm (W-band) and 5 cm  $\times$  5 cm (Q-band). Figure 5 shows the W-band array assemblies.

##### 5.1. Polarimeter Assemblies

Each QUIET polarimeter assembly consists of (i) a septum polarizer, (ii) a waveguide splitter, and (iii) a module containing the highly integrated package of HEMT-based MMIC devices (see Figure 5.1). The septum polarizer consists of a square waveguide with a septum (a thin aluminum piece with a stepped profile) in the center, which adds a phase lag to one of the propagating modes (Bornemann & Labay 1995). Given an incident electric field with linear orthogonal components  $E_x$  and  $E_y$ , where the  $x$  and  $y$  axis orientations are defined by the septum, the septum polarizer assembly sends a left-circularly polarized component  $L = (E_x + iE_y)/\sqrt{2}$  to one output port, and a right-circularly polarized component  $R = (E_x - iE_y)/\sqrt{2}$  to the other output port. Thus the septum's spatial orientation is used to define




FIG. 13.— The W-band array polarimeter and differential-temperature assemblies. This is the largest HEMT-based array ever assembled to date.

the instrumental position angle. The output ports of the septum polarizer are attached to a waveguide splitter which transitions from the narrow waveguide spacing of the septum-polarizer component to the wider waveguide separation of the module waveguide inputs. A more thorough mathematical description of the septum polarizer is given in Appendix 11.3.

With a Vector Network Analyzer (VNA), we measure the scattering matrices of both the Q-band and W-band septum polarizers, and derive the gain and the total power-to-polarization (monopole) leakage terms. Q-band VNA measurements of prototypes show an enhancement in the return loss near the low frequency end of the module's bandpass. Therefore, we introduce a bandpass mismatch between the septum polarizer and module to prevent oscillations in the module output. The bandpass mismatch leads to an enhancement in the differential loss between the  $E_x$  and  $E_y$  transmissions at 47 GHz, causing a total power-to-Stokes Q leakage of  $\sim 1\%$ , averaged over the module's bandpass. This estimate is consistent with leakage values derived from Tau A measurements (Section 3.6). W-band VNA measurements show no return loss enhancement, and therefore we made no bandpass adjustments. The VNA measurements predict a smaller leakage of  $\sim 0.3\%$ , so that it is subdominant to leakage due to optics. These measurements are consistent with monopole leakage values obtained from on-sky calibrators (see Section 3.6 and figure 11). Note that since the optics leakage has a random direction relative to the polarimeter assembly leakage, the combined leakage averages to a smaller value and is randomly distributed both in sign and amplitude.

##### 5.2. Modules

The QUIET modules are used in the polarimeter and differential-temperature assemblies (see Sections 5.1 and 5.3), functioning as pseudo-correlation receivers so that the output is a product (rather than sum or difference) of gain terms. While the modules employ a high speed switching technique to reduce  $1/f$  noise, they are an improvement on classical Dicke radiometers (Dicke 1946): they do not have an active switch at the amplifier input, and there is an additional improvement of  $\sqrt{2}$  in sensitivity since the modules continually measure the sky signal






FIG. 14.— A W-band polarimeter assembly. The module is more compact than previous generation correlators by an order of magnitude.

(Mennella et al. 2003).

In a polarimeter assembly, the module receives as inputs the left ( $L$ ) and right ( $R$ ) circularly polarized components of the incident radiation, and measures the Stokes parameters  $Q$ ,  $U$  and  $I$ , defined as (we also define Stokes  $V$  for completeness):

$$\begin{aligned} I &= |L|^2 + |R|^2 \\ Q &= 2 \operatorname{Re}(L^* R) \\ U &= -2 \operatorname{Im}(L^* R) \\ V &= |L|^2 - |R|^2. \end{aligned} \quad (11)$$

where the  $*$  denotes complex conjugation.

Figure 15(a) shows a schematic of the QUIET module, in which  $L$  and  $R$  traverse separate amplification “legs” (called legs A and B). A phase switch in each leg allows the phase to be switched between  $0^\circ$  ( $+1$ ) and  $180^\circ$  ( $-1$ )<sup>31</sup>. The outputs of the two amplification legs are combined in a  $180^\circ$ -degree hybrid coupler which, for voltage inputs  $a$  and  $b$ , produces  $(a+b)/\sqrt{2}$  and  $(a-b)/\sqrt{2}$  at its outputs. The hybrid coupler outputs are split, with half of each output power going to detector diodes  $D_1$  and  $D_4$ , respectively. The other halves of the output powers are sent to a  $90^\circ$  coupler which, for voltage inputs  $a$  and  $b$ , produces  $(a+ib)/\sqrt{2}$  and  $(a-ib)/\sqrt{2}$  at its outputs. The outputs of this  $90^\circ$  coupler are each detected in diodes  $D_2$  and  $D_3$ , respectively. The detector diodes are operated in the square-law regime, and so their output voltages are proportional to the squared input magnitudes of the electric fields.

Table 9 shows the idealized detector diode outputs for the two states of leg B, with the leg A state held fixed.

<sup>31</sup> The phase switch acts uniformly across the bandwidth of the module.





FIG. 15.— *a*: Signal processing schematic for an ideal QUIET module in a polarimeter assembly. The diode raw signals are given for the two ( $\pm 1$ ) leg B states, and for the leg A state fixed ( $+1$ ). For simplicity, details of the 3 LNAs and bandpass filters are not shown. *b*: Internal components of a 5 cm  $\times$  5 cm Q-band module.

The diode outputs are averaged and demodulated by additional warm electronics (see Section 6). Given a diode output of  $I \pm Q(U)$ , the averaging and demodulation operations return  $I$  and  $Q(U)$  respectively.<sup>32</sup> We can self-consistently generate the Stokes parameters with units of temperature as follows (Staggs et al. 2003). Let  $T_x$  ( $T_y$ ) be the physical temperature of a black body that emits the observed value of  $\langle E_x^2 \rangle$  ( $\langle E_y^2 \rangle$ ). The Stokes parameters become:

$$\begin{aligned} I &= \frac{1}{2} \cdot (T_x + T_y) \\ Q &= \frac{1}{2} \cdot (T_x - T_y) \end{aligned} \quad (12)$$

For completeness, the voltage  $V_{Q1}$  appearing at the  $Q_1$  diode would measure:

$$V_{Q1} = g \cdot \left( \frac{1}{2}(T_x + T_y) \pm \frac{1}{2}(T_x - T_y) \right) \quad (13)$$

where  $\pm$  represent the states of leg B, and  $g$  is the responsivity constant extracted using calibration tools and procedures described in Sections 7 and 8.

In practice, the phase of leg B changes at 4 kHz, so that the demodulated  $Q$  and  $U$  outputs do not suffer from  $1/f$  noise, which is due primarily to the LNAs. However the phase switches do not reverse the sign of  $I$ , therefore the  $I$  output suffers from significant  $1/f$  noise and so is not used to measure the temperature anisotropy. The choice of circularly-polarized inputs thus allows us to measure both Stokes  $Q$  and  $U$  simultaneously which gives an ad-

<sup>32</sup> When referring to diodes  $D_1$ ,  $D_2$ ,  $D_3$ , and  $D_4$ , we shall use the naming convention  $Q_1$ ,  $U_1$ ,  $U_2$ , and  $Q_2$  diodes respectively.

vantage in sensitivity.

The amplifier gains and transmission coefficients are represented by the proportionality symbols in Table 9. In practice, the transmission through leg B is not exactly identical between the two leg B states, leading to additional free parameters needed to characterize the module. If the leg B transmission differences are not accounted for, they lead to instrumental (i.e. false) polarization. This is resolved by modulating the phase of leg A at 50 Hz during data taking, and performing the “double demodulation” procedure on the offline data. Imperfections in the optics and the septum polarizer introduce additional offsets and terms proportional to  $I$ . These effects are discussed in Appendix 11.2 and 11.3.

In practice, the signal pseudo-correlation is implemented in a single small package as shown in Figure 15(b) (Kangaslahti et al. 2006; Cleary 2010). The low noise amplifiers (LNAs), phase switches and hybrid couplers are all produced using the same Indium-Phosphide (InP) fabrication process. Three LNAs, each with gains  $\sim 25$  dB, are used in each of the two legs. When the input amplifiers are packaged in individual amplifier blocks and cryogenically cooled to  $\sim 20$ K, they exhibit noise temperatures of about 18 K (50–80 K) for the Q-band (W-band). The phase switches operate by sending the signal down one of two paths within the phase switch circuit, one of which has an added length of  $\frac{\lambda}{2}$  (ie.  $180^\circ$  shift). Two InP PiN (p-doped, intrinsic-semiconductor, n-doped) diodes control which path the signal takes. The signals go through band-defining passive filters made from AlN substrates, and then detected by commercially-available Schottky detector diodes downstream of the hybrid couplers. The amplifiers and phase switches are frequency dependent, and hence unique to each array. The detector diodes are capable of functioning at both 40 GHz and 90 GHz, and so are identical between the two arrays.

The module components are packaged into clamshell-style brass housings, precision-machined for accurate component placement and signal routing. Each housing has feedthrough pins connecting the module components via strip-line and wire ribbon bonds for bias and diode readout. Miniature absorbers and an epoxy gasket between the two halves of the clamshell are used to isolate unwanted cross talk between the RF and DC components. All Q-band modules and roughly 30% of W-band modules are assembled by hand. For the remaining W-band modules, the components and substrates are automatically placed in the housings by a commercial con-

tractor using a pick-and-place machine; the wire bonding, absorber and epoxy gasket are then finished by hand.

### 5.3. Differential-Temperature Assemblies

The differential-temperature assemblies are grouped into pairs of assemblies, with waveguide components that mix two neighboring horn signals into two neighboring modules. Figures 16(a) and 16(b) show the schematic and implementation of these assemblies. An orthomode transducer (OMT) located after feedhorn “A” outputs the linear polarizations  $E_{Ax}$  and  $E_{Ay}$ . One of these polarizations,  $E_{Ay}$ , enters a waveguide  $180^\circ$  coupler (a “magic-tee”) and is combined with  $E_{Bx}$  from the adjacent feedhorn. The magic-tee outputs are coupled to a module’s inputs. The OMTs are reused from CAPMAP (Barkats et al. 2005) while the waveguide routing and magic-tees are custom made by Custom Microwave. Note that the differential-temperature assembly design resembles that of WMAP (Jarosik et al. 2003b), with the only significant difference being in the implementation of the LNAs. Where as WMAP uses a conservative design of discrete HEMT LNAs operating at both cold and warm temperatures, advances in MMIC HEMT LNAs allows for QUIET’s more compact design.

For an ideal differential-temperature assembly, the demodulated Q diode (ie.  $D_1$  and  $D_4$ ) outputs measure  $E_{Ax}^2 - E_{By}^2$ , while their counterpart in the adjacent differential-temperature assembly measure  $E_{Ay}^2 - E_{Bx}^2$ . The difference of demodulated Q diode outputs from adjacent differential-temperature assemblies measure the beam-differenced total power  $(E_{Ax}^2 + E_{Ay}^2) - (E_{Bx}^2 + E_{By}^2) = I_A - I_B$  (see Table 10). The demodulated U diode outputs (ie.  $D_2$  and  $D_3$ ) is zero for an ideal assembly. However, unequal path lengths ( $\phi$ ) in the two legs of a module move some of the temperature difference signal from the Q diodes to the U diodes. The Q(U) diode signals vary as  $\cos(\phi)(\sin(\phi))$ . For the differential-temperature assemblies,  $\phi$  is  $\sim 10 - 20$  degrees which moves  $\sim 15\text{-}30\%$  of the signal to the U diodes.

Finally, we note that the sum of demodulated Q diode outputs from adjacent modules is  $Q_A + Q_B$ , where  $Q$  is the Stokes Q parameter seen by the respective horns. Thus one can in principle extract polarization information from the differential-temperature assemblies. However, as these assemblies form a small fraction of the array, the sensitivity gain is marginal and so this option was not explored further in the analyses.

## 6. ELECTRONICS

TABLE 9

RAW, AVERAGED, AND DEMODULATED DETECTOR DIODE OUTPUTS FOR AN IDEAL QUIET MODULE IN A POLARIMETER ASSEMBLY. RESULTS ARE SHOWN FOR THE TWO STATES OF LEG B, WITH THE LEG A STATE HELD FIXED.

Diode	Raw Output	Average	Demodulated
$D_1$	$\propto \frac{1}{4}(I \pm Q)$	$\propto \frac{1}{4}I$	$\propto \frac{1}{2}Q$
$D_2$	$\propto \frac{1}{4}(I \mp U)$	$\propto \frac{1}{4}I$	$\propto -\frac{1}{2}U$
$D_3$	$\propto \frac{1}{4}(I \pm U)$	$\propto \frac{1}{4}I$	$\propto \frac{1}{2}U$
$D_4$	$\propto \frac{1}{4}(I \mp Q)$	$\propto \frac{1}{4}I$	$\propto -\frac{1}{2}Q$

TABLE 10

OUTPUTS OF AN IDEAL DIFFERENTIAL-TEMPERATURE ASSEMBLY AT DETECTOR DIODES  $D_1$  AND  $D_4$  CORRESPONDING TO A LEG B STATE OF  $+1(-1)$ , WITH LEG A FIXED AT  $+1$ . ALSO SHOWN IS THE DIFFERENCE OF THE DEMODULATED  $D_1$  SIGNALS FROM TWO MODULES.

	Mod 1	Mod 2
$D_1$	$\propto E_{Ay}^2(E_{Bx}^2)$	$\propto E_{By}^2(E_{Ax}^2)$
$D_4$	$\propto E_{Bx}^2(E_{Ay}^2)$	$\propto E_{Ax}^2(E_{By}^2)$
$demod(D_{1,Mod1}) - demod(D_{1,Mod2})$	$(E_{Ax}^2 + E_{Ay}^2) - (E_{Bx}^2 + E_{By}^2)$	




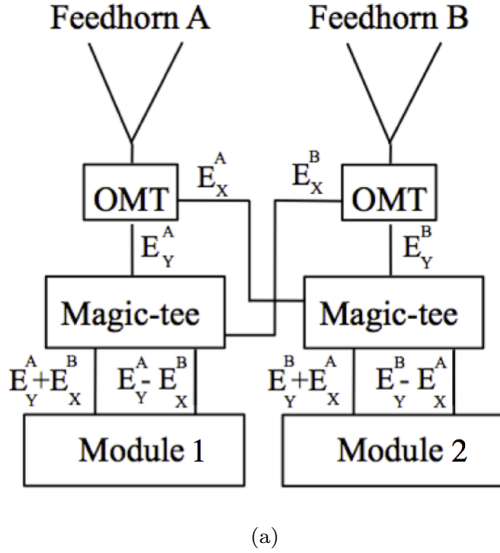


FIG. 16.— *a*: Schematic of the waveguide coupling for the hybrid-Tee assembly. An Orthomode Transducer (OMT) located after feedhorn “A” outputs the linear polarizations,  $E_{Ax}$  and  $E_{Ay}$ . One of these polarizations,  $E_{Ay}$ , enters a waveguide  $180^\circ$  hybrid coupler (a “magic-tee”) and is combined with the orthogonal polarization from an adjacent feedhorn,  $E_{Bx}$ . Proportionality constants such as factors of  $1/\sqrt{2}$  have been omitted for simplicity. *b*: Implementation of a W-band differential-temperature assembly (modules not shown).

Downstream of the modules are electronics for detector biasing, timing, preamplification, digitization, and data collection. These functions are accomplished by four systems: (1) Passive Interfaces, (2) Bias, (3) Readout, and (4) Data Management. The Passive Interfaces system (Section 6.1) creates an interface between the modules, the Bias system, and the Readout system. The Bias system (Section 6.2) provides the necessary bias to each module active’s components. The Readout system (Section 6.3) amplifies and digitizes the module outputs. The Data Management system (Section 6.4) commands the other systems and records the data. The Bias and Readout systems are housed in a weather-proof temperature-controlled enclosure to protect them from the harsh conditions of the Atacama Desert. The enclosure also serves as a Faraday cage to minimize radio-frequency interfer-

ence. Further description of these electronics can be found in Bogdan et al. (2007).

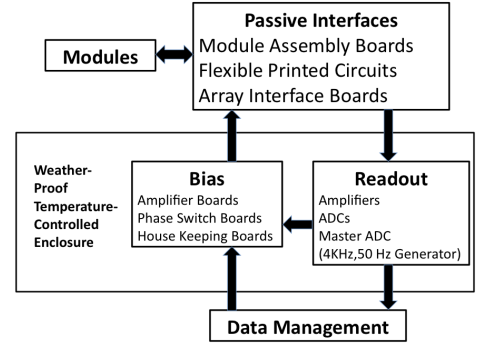


FIG. 17.— Major components of the electronics. Boxes outline the four main systems. Arrows indicate the flow of bias commands and data signals.

### 6.1. Passive Interfaces

The electrical connection to, and protection of, the modules is provided by Module Assembly Boards (MAB). Each MAB is a printed circuit board with pin sockets for seven modules. Voltage clamps and RC low-pass filters protect the sensitive components inside the module from damage. The Q-band(W-band) modules require 28(23) pins for grounding, biasing active components, and measuring the detector diode signals. All of these electrical connections are routed outside the cryostat. After the MAB protection circuitry, these signals travel on high density flexible printed circuits (FPC), which bring them out of the cryostat through Stycast-epoxy-filled hermetic seals. An additional layer of electronic protection circuitry is provided by the Array interface boards (AIB), which also adapt the FPC signals to board-edge connectors and route to the Bias and Readout systems.

### 6.2. Bias System

All biasing is accomplished by custom circuit boards. The amplifier bias boards provide voltage and current to power the amplifiers in the modules. Each of these bias signals is controlled by a 10-bit digital-to-analog converter (DAC), which allows the biases to be tuned for optimal performance of each amplifier. Phase switch boards provide control currents to the PiN diodes in the phase switches. The control current is switched by the board at 4 kHz for one phase switch and 50 Hz for the other phase switch, generating the modulation described in Section 5.2. The data taken during the switch transition time are discarded in the Readout system. A house-keeping board monitors the bias signals at  $\approx 1$  Hz for each item being monitored. The housekeeping board multiplexes between these items, switching only during the phase switch transitions when data will be discarded.

The Q-band amplifier bias boards are designed to operate at  $25^\circ\text{C}$  so the enclosure is thermally regulated at that temperature. The W-Band amplifier bias boards use a different design that is much less temperature sensitive. Therefore, the enclosure regulation temperature for the W-band is varied between  $35^\circ\text{C}$  and  $40^\circ\text{C}$  depending on the season to reduce the power needed for regulation. For

both the Q-band and W-band observing seasons, the enclosure temperature remained within the regulation set-point for  $\approx 90\%$  of the time. For the Q-band system, the excursions will primarily affect the drain-current bias supplied by the amplifier bias boards, which changes the detector responsivity by  $\approx 2\%/^{\circ}\text{C}$ . This effect is taken into account with an enclosure-temperature dependent responsivity model (QUIET Collaboration et al. (2011)).

### 6.3. Readout System

The Readout system first amplifies each module’s detector diode output by  $\approx 130$  in order to match the voltage range of the digitizers. The noise of this warm amplifier circuit does not contribute significantly to the total noise. This is determined *in situ* at the site by selectively turning off the LNA’s in the module and seeing that the total noise decreased by roughly two orders of magnitude. For the W-band array, the amplifier noise contributes less than 2% to the total noise in the quadrature sum. The amplifier chain also low-pass filters the signal at  $\approx 160$  kHz to prevent aliasing in digitization. Each detector diode output is digitized by a separate 18-bit Analog Devices AD7674 (Analog-Digital Converter) ADC with 4V dynamic range at a rate of 800 kHz. Each ADC Board has a field-programmable gate array (FPGA), which accumulates the samples from the 32 ADCs on that Board. The FPGA on one ADC Board, designated the “Master ADC Board,” generates the 4 kHz and 50 Hz signals used by the Bias system to modulate the phase switch control current. This signal is also distributed to all ADC Boards, and the FPGA on each ADC Board uses it to demodulate the detector diode data synchronously with the phase switch modulation.

Figure 18 summarizes the organization of data performed by the FPGA. The FPGA organizes the 800 kHz of detector diode data into continuous 10 msec blocks (i.e. 100 Hz time streams), itself organized into continuous 125  $\mu\text{sec}$  blocks. These 10 msec blocks contain an equal sampling of both 4 kHz clock states. In the “TP” stream, the 800 kHz data within a 10 msec block are averaged, regardless of the 4 kHz clock state. This stream is sensitive to Stokes  $I$  and is used for calibration and monitoring<sup>33</sup>. In the “demodulated” stream, data within a 125  $\mu\text{sec}$  block have the same 4 kHz phase state, and are averaged. Averaged data from sequential 125  $\mu\text{sec}$  blocks are differenced, thus forming the polarization-sensitive data stream. Offline, two adjacent 10 msec blocks in the demodulated stream are differenced to form the “double-demodulated” (50 Hz) stream. The W-band data contains an additional specially demodulated 100 Hz stream, called the “quadrature stream”. Unlike the usual demodulated stream, data within a 125  $\mu\text{sec}$  block populate equally both 4 kHz phase states, and are averaged. When these averaged data are differenced, the result has the same noise as demodulated data but has *no signal*. The quadrature stream are used to monitor potential contamination and to understand the detector noise properties.

We noted in the previous section that data was masked at the phase switch transition. Masking 14% of the sam-

<sup>33</sup> However this stream does not have the benefit of the  $1/f$  noise reduction, and so is not used to compute the temperature anisotropy spectrum.

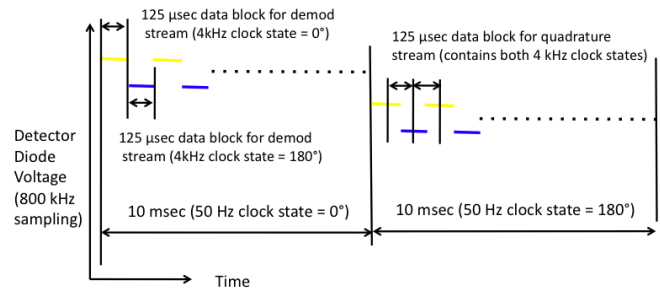


FIG. 18.— Organization of 800 kHz data to form the demodulated and quadrature 100 Hz streams. The blue and yellow lines show detector diode data for the two 4 kHz phase states. Levels are exaggerated for clarity.

ples around the transition was found to be adequate to remove contamination in the data stream.

The ADC Boards have a small non-linearity in their response. At intervals of 1024 counts, the ADC output has a jump discontinuity between 1 and 40 counts, roughly affecting 14% of the data. This jump is shown schematically in Figure 19. When the 800 kHz data stream value falls at a discontinuity, the jump in the output signal will trickle into the 100 Hz stream. We correct for this non-linearity in 100 Hz stream. The correction is statistical in nature, based on the width of the 800 kHz noise and its proximity to the discontinuity (Bischoff 2010b). This nonlinearity, if uncorrected, causes a variation of responsivity during a CES and systematic effect similar to the leakage of temperature to polarization. For the Q-band, the correction reduces the ADC nonlinearity to contribute at most 3% to the leakage bias systematic, and at most 50% to the CES responsivity systematic. For the W-band, the residual ADC nonlinearity adds 40% in quadrature to the leakage bias systematic. The effect from CES responsivity is  $< 1\%$ , negligible compared to other errors in the gain model. These affect  $r$  at a level below 0.01 for the W-band.

The Readout system ensures that the housekeeping data and 100 Hz data from the detectors are synchronized to each other and to the mount motion encoder readout. Synchronization is achieved by distributing the same GPS-derived IRIG-B<sup>34</sup> time code to both the receiver and mount electronics. In the Readout system, the time code is decoded by a Symmetricom TTM635VME-OCXO timing board. One-Hz and 10-MHz clock signals, locked to the IRIG-B time code, synchronize the readout of all ADC Boards. The timing board provides the GPS-derived time to the Data Management system so that each datum is assigned a time stamp.

### 6.4. Data Management

The Data Management system sends commands to the Bias system to prepare for observation, acquires the data from the Readout system, writes them to disk, and creates summary plots of the detector diode signals and housekeeping data for display in real time. The complete data are written to disk and DVD’s in the control room at the observation site at a rate of  $\approx 8$  GB day<sup>-1</sup> for the Q-band array. W-band array data are written to blu-ray optical discs at the rate of  $\approx 35$  GB day<sup>-1</sup>. A subset of  $\approx 10\%$  of the data are transferred by internet every day for more rapid analysis and monitoring.

<sup>34</sup> Inter-range Instrumentation Group Mod B.




FIG. 19.— Non-linear response of the ADC plotted in arbitrary units. The horizontal axis shows the ADC input near the discontinuity. The red and blue distributions show the fluctuations of 800 kHz samples in two different 4 kHz phase switch states. When these distributions cross the ADC response discontinuity, the output distributions are split at the discontinuity. When the red and blue states are differenced to create demodulated data, the split caused by the discontinuity is added to the result.

## 7. ARTIFICIAL CALIBRATORS

We used both astronomical and artificial calibrators to characterize the instrument. Astronomical calibrators are described in Section 3.4 and Section 8. In this section, we focus on the artificial calibrators we developed for QUIET for use in the laboratory and at the observation site.

### 7.1. The Optimizer

We measured the polarized response of the receiver in the laboratory with the ‘optimizer,’ a reflective plate and cryogenic load that rotate around the bore sight of the cryostat (Figure 20).




FIG. 20.— The optimizer consists of a reflective metal plate and a cryogenic load, which co-rotate around the cryostat. We set the plate angle  $\beta$  to be  $45^\circ$ . The reflected signal is polarized (given by equation 14) and the polarized component modulates at twice the angular frequency of the rotating apparatus.

The plate is oriented at angle  $\beta$  from the plane of the

feedhorns and reflects light from the cryogenic load into the window of the cryostat with a Stokes Q in temperature units (Barkats et al. 2005).

$$Q = \frac{1}{2} \cdot \frac{4\pi\delta}{\lambda} (\cos \beta - \sec \beta) (T_{\text{plate}} - T_{\text{load}}) \sin(2\alpha),$$

$$\delta = \sqrt{\frac{\rho}{\mu_0 \pi \nu}}, \quad (14)$$

where  $\rho$  is the bulk resistivity of the metal plate,  $\nu$  and  $\lambda$  correspond to the center frequency and wavelength of the detector bandpass,  $\epsilon_0$  is the vacuum dielectric constant,  $t$  is time, and  $T_{\text{plate}}$  and  $T_{\text{load}}$  are the temperature of the plate and cryogenic load, respectively.  $\delta$  is the skin depth for the frequency  $\nu$ , and  $\mu_0$  is the permeability of free space. The Stokes Q is the difference in temperatures between radiation normal and parallel to the plane of incidence, and scaled by  $\frac{1}{2}$  to maintain consistency as shown in Equation 12. This apparatus rotates at a angular speed  $\alpha$  around the boresight of the cryostat so that the resulting polarized signal will rotate between the Stokes Q and U at a angular speed of  $2\alpha$ . Polarization signals that do not rotate with the system (such as thermal emission from objects in the laboratory) will be detected at a rate of  $\alpha$ , and so can be removed.

The predicted polarized emission from Equation 14 and the measured voltage on the detector diodes allow us to calculate the polarized responsivity for polarimeters whose beam primarily samples the reflected cryogenic load. We used various plate materials (aluminum, stainless steel, and galvanized steel) and two thermal loads (liquid nitrogen and liquid oxygen) to obtain multiple estimates of the polarized responsivity. The loads are too small to fill the entire array beam, so only the measurements from the central polarimeter (Q-band) or inner two rings (W-band) were used. The optimizer was used to verify that the responsivities derived from unpolarized measurements with cryogenic loads were not substantially different from the polarized responsivities, and hence that the projections of instrument sensitivity (which were made from unpolarized measurements) were valid for the Q-band array. For the W-band array, we used the optimizer to select functioning modules for the final array configuration.

### 7.2. The Wire Grid Polarizer

We built a ‘sparse wire grid’ (Tajima et al. 2012): a plane of parallel wires held in a large circular frame with the same diameter as the cryostat window. For the polarization parallel to the grid wires, a fraction of the rays that would ordinarily pass through the telescope to the cold sky are instead scattered to large angles, mostly terminating on the warm ground shield. The grid was placed as close to the cryostat aperture as possible to minimize interference with the telescope optics and to ensure that it covered the field-of-view of each detector (Figure 21).

With this geometry, a  $\sim 2$  K polarized signal is directed parallel to the length of the wire. The circular frame is rotateable via a small motor, allowing us to modulate the injected polarized signal at a constant frequency. We used the wire grid for calibration measurements in the laboratory and three times during the observing season: at the end of the Q-band observing season and at the






FIG. 21.— Sparse wire grid array mounted on the W-band cryostat (right), and the fine wire detail (left).

beginning and end of the W-band observations.

An example of the data taken with the rotating grid is shown in Figure 22. In the ideal case in which the intensity of the reflected radiation is isotropically uniform over the array, the polarized signal from each detector  $P(\theta)$  would exhibit a perfect sinusoidal dependence at twice the frequency of the sparse wire grid angle  $\theta$ . There is an additional component in the measured polarization signal as the grid rotates around the boresight which comes from the rays terminating at different temperatures in the non-uniform ground screen. This variation appears in both the polarized data stream  $P$  as well as the total power data stream  $I$  as a function of  $\theta$ , and so this component can be measured in the total power data stream and removed from the demodulated stream. For each detector, we fit the data to the form

$$P(\theta) = G \cdot I(\theta) \cos(2(\theta - \gamma)) - P_0, \quad (15)$$

to extract  $G$  (the polarized gain),  $\gamma$  (the detector angle as defined within the wire grid system),  $P_0$  (an offset induced by leakage) from the polarized data  $P(\theta)$  and the total power data  $I(\theta)$ . Since the wire grid orientation with respect to the sky coordinates is not measured, the angle  $\gamma$  for each detector provides only the relative direction of detector angles. Similarly, we use only the relative values of  $G$  (i.e. gain ratios amongst the detectors), which agree with gain ratios derived from Tau A observations.




FIG. 22.— Polarization response of a detector as a function of time, where the wiregrid was rotating at  $\sim 8$  RPM. A sinusoidal response is clearly observed at twice the rotation frequency.

## 8. RECEIVER CHARACTERIZATION AND CALIBRATION

Each receiver diode (Table 11) is characterized by its bandpass, noise level, polarization angle, and total power and polarized responsivities. These quantities were mea-

sured for the receiver arrays in the laboratory prior to deployment, and at the site using astronomical calibrators, sky dip measurements, and polarizing grid measurements.

In this section, we discuss our methods of module bias optimization (Section 8.1) and module leakage remediation (Section 8.2). We include discussion of our methods for characterizing module bandpasses (Section 8.3), responsivities (Section 8.4), detector angles (Section 8.5), noise measurements (Sections 8.6 and 8.7), and sensitivity (Section 8.8).

TABLE 11  
DETECTOR YIELD FOR THE Q-BAND AND W-BAND ARRAYS.

Band	Q	W
Number of modules	19	90
Polarization modules	17	84
Polarization diodes	68	336
Working polarization diodes (Stokes Q)	31	153
Working polarization diodes (Stokes U)	31	155
Total power modules	2	6
TT diodes (Stokes Q only)	4	12
Working TT diodes (Stokes Q only)	4	12

### 8.1. Detector Biasing and Optimization

The amplifiers were biased manually for gain balance between the module legs and for adequate signal level for the Q-band array at the beginning of the observing season. This was performed for each module using a room temperature black-body load in front of the cryostat. The phase switches are turned on separately, such that the signal can only propagate through the module leg with the phase switch on. The amplifiers can then be biased one leg at a time such that the first stage amplifier drain current is in the range 0-5 mA, the second stage drain current is in the range 5-15 mA, and the third stage amplifiers are in the range 15-30 mA, and that the signal measured by the detector diodes is  $\sim 5$  mV. We then repeat turning on only the phase switch for the other leg, and adjust again to obtain a signal difference between the two legs of 0.6 mV. This biasing scheme reduces the current through the first stage amplifier to  $\sim 30\%$  that of its maximum value to keep its noise contribution low. The bias values for the phase switches were chosen to equalize the signal measured on the two separate legs of the module. These bias settings were found once at the beginning of the season, and kept fixed during the observing season.

For the W-band array, biasing the modules by hand was not feasible due to the larger number of modules compared to the Q-band array, and so an automatic method was developed. A polarized signal was injected during module biasing by continually rotating the sparse wire grid, generating a sinusoidal polarization signal. Amplifier bias settings were found by maximizing the amplitude of the sinusoid relative to the time-stream noise. The bias settings were sampled via a computer-based downhill simplex algorithm and optimum values were found for all modules within few hours. These settings were also kept fixed during the W-band observing season. Because the settings were chosen for signal-to-noise, balance between the legs was not prioritized.

### 8.2. Module Leakage Remediation

One source of leakage from total power into polarization from the module stems from differential power transmission between the two phase switch states within a given leg (Section 11). We found that double demodulating (described in Section 5.2) typically reduces the root-mean-square of leakage from 0.8% to 0.4% for the W-band modules (Figure 23). The improvement was smaller for the Q-band array,  $< 0.1\%$ , likely because it was dominated by other sources of leakage (Section ??) and because the phase switches were balanced during bias optimization.




FIG. 23.— A histogram of diode leakage values between total power and polarization channels during a large angle sky dip for the W-band array before and after double demodulation. Double demodulating reduces the total power leakage by a factor of  $\sim 2$  for the W-band array.

### 8.3. Bandpasses

Bandpasses were measured for the Q-band array for each diode in the laboratory during the course of array testing and in an end-of-season calibration measurement at the site. The laboratory measurement was performed by injecting a polarized carrier-wave signal from a signal generator with a standard-gain horn over a 35–50 GHz range. The signal was injected into the receiver array through the cryostat window without additional imaging optics, with the horn approximately 3 m away from the window. Sweeps were performed at least eight times. The average bandwidth and central frequency of the polarization modules are given in Table 12. The statistical errors on this measurement are obtained by finding the standard deviation between the eight measurements for a given module, and then averaging that standard deviation for all modules. The systematic error is half of the average of the largest difference between the eight measurements for each module.

Bandpasses were measured at the site for the Q-band array by reflecting the swept signal from a small ( $\sim 1 \text{ cm}^2$ ) plate into the primary mirror. While measurements performed in the laboratory and at the site are consistent with each other, the variation in bandpass shape between the two days of data taking at the site showed that the systematic errors were larger in the experimental setup

at the site, so we used laboratory measurements where available. There was some concern that because the amplifier bias settings were different between the laboratory and the site measurements, the laboratory measurements may not be a good enough proxy for the bandpasses during CMB and calibration observations at the site. Using data taken in the laboratory with different amplifier bias settings, we investigated the relationship between amplifier biasing and the central frequency and bandwidth, and found no relationship between the two. We conclude that changes in amplifier bias over the range of interest do not have a significant effect on the bandpasses.

Bandpasses were measured at the site for the W-band array at the end of the observing season and the central frequency and bandwidth are also given in Table 12. A standard-gain horn was mounted beside the secondary mirror, so it could illuminate the cryostat window from  $\sim 1.5 \text{ m}$  away. The signal generator was swept over 72–120 GHz, while the phase switches were held constant (no switching). In this configuration we can send the signal down each module leg separately. The responses at each frequency bin for each module leg were combined to emulate the power combinations occurring in the module:

$$P_{\text{pol}} = P_A P_B \cos(2(\phi - \gamma)), \quad (16)$$

$$(17)$$

where  $P_A$  and  $P_B$  are the measured bandpasses for the signal travelling through module legs A and B, respectively,  $\phi$  is the assumed detector angle relative to the horn (for example, a Q diode might have  $\phi = 90^\circ$  and a U diode might have  $\phi = 45^\circ$ ), and  $\gamma$  is the horn angle. Thus with the horn oriented at  $\gamma = 90^\circ$ , if the detector angle is aligned such that the Q diode has  $\phi = 90^\circ$ , the polarized bandpass is maximal on the Q diodes, and minimal on the U diodes of the module. This reconstruction is straightforward for the total power bandpass, but requires a well motivated assumption for the angle between the horn and the detector axis, which we took as  $\phi = 90^\circ$  for the Q diodes. Systematic errors have two main sources: the accuracy with which we can detect the spike used to indicate the beginning of a sweep, and from reconstructing the bandpass for both module legs biased from data in which only one leg is biased. The first was computed by noting that the timing was accurate to 1.5 ms, which corresponded to 0.7 GHz during the sweep measurement. The second was computed by comparing the difference between measurements performed with both legs biased and the reconstruction from single-leg bandpasses from the total power stream. Because the total power stream does not have a dependence on detector angle  $\phi$ , the two should be identical and the difference represents the systematic error in the measurement. The systematic error was found to be 0.3 GHz for central frequency and 0.9 GHz for the bandwidth.

Typical bandpasses for the Q-band and W-band arrays are shown with the spectrum of the atmosphere in Figure 2. Central frequencies and bandwidths are computed

from discrete frequency steps as

$$\text{Central frequency} \equiv \frac{\sum_i I_i \nu_i}{\sum_i I_i} \quad (18)$$

$$\text{Bandwidth} \equiv \frac{[\sum_i I_i]^2 \Delta \nu}{\sum_i I_i^2}, \quad (19)$$

where  $I_i$  is the measured intensity from a detector diode for each frequency step,  $\nu_i$  is the corresponding frequency, and  $\Delta \nu$  is the frequency step of the signal generator (100 MHz).

Although the bandwidths and central frequencies in Table 12 are only given for the polarization modules, the values for the differential-temperature modules were consistent within the statistical accuracy of the swept source measurements. While we designed the septum-polarizers and differential-temperature assemblies to match the module bandpass, the consistency between the two was not necessarily expected both because the bandwidth of the septum polarizer is slightly narrower than that of the differential-temperature assemblies, and because the septum polarizer can impart a frequency dependent detector angle.

#### 8.4. Responsivities

We characterize the responsivities for the differential-temperature modules and the polarization modules separately with different calibration sources. Responsivities for the differential-temperature modules are computed from calibration observations of Jupiter, RCW38, and Venus, one of which was observed  $\sim$ once per week for the Q-band receiver, and once a day for the W-band receiver. The average responsivities for the differential-temperature modules was 2.2 mV/K for the Q-band array and 2.3 mV/K for the W-band array.

We found the absolute polarimeter responsivity for the Q-band array central horn from Tau A measurements performed every two days. Relative responsivity values among the polarization modules were found from observations of the Moon (performed once per week). Sky dip measurements (elevation nodes of  $\sim 6^\circ$  for ‘normal’ sky dips, and  $\sim 40^\circ$  for ‘large’ sky dips) are also used to obtain relative total power responsivities of both the differential-temperature and polarized modules before each CES for the Q-band array (“flat fielding”). These frequent (once every  $\sim 1.5$  hours) responsivity measurements provide relative responsivity tracking for the differential-temperature and polarized modules on short time-scales. These relative responsivities were checked with an end-of-season wire-grid measurement and measurements of Tau A with off-center modules.

For the W-band array, the Moon is too bright for relative responsivity measurements, so we used measurements from the wire-grid and measurements of Tau A from off-center modules. The average responsivity for the polarized modules was 2.3 mV/K for the Q-band array and 3.1 mV/K for the W-band array. These responsivities are in terms of antenna temperature and include the gain factor of 130 from the preamplifier boards (Section 6.3).

The responsivity depends on the amplifier bias settings. The bias values for the Q-band array were found

to be dependent on the bias board temperature (typical values are 2% of the average responsivity per K), which was the motivation for thermally regulating the electronics enclosure to 25 C within 1 C. The final responsivity model included a linear term for this temperature dependence, and it was found to be a negligible systematic for scientific analysis (QUIET Collaboration et al. 2011). The bias circuit was upgraded in the W-band array, rendering the temperature dependence of the boards negligible.

One common concern when using amplifiers is signal compression: an input-dependent responsivity where the responsivity is greatly reduced at high input powers. Compression is typically manifested as different responsivity values for different load temperatures, and has important consequences when using responsivities from calibration sources which are all usually much warmer than the CMB itself (for example, the Moon is  $\sim 223$  K, (Ulich et al. 1973)) and from extrapolating total power responsivities to polarization responsivities. For the Q-band array, responsivity measurements in the laboratory and at the site with different calibration sources were all consistent with each other, confirming that the modules were not operating in a compressed regime. Laboratory responsivity studies of the W-band modules using liquid nitrogen as a cold load show some evidence for compression. In the field, the W-band modules exhibited compression during measurements of the Moon. The emission from the  $\sim 1^\circ$  moon varies across its face (Ulich et al. 1973) and we found that the polarized responsivities varied between the brightest and darkest portions of its face by 20% (worst case 50% ).

Compression affects the polarized signal and the total power signal differently (Section 11.1). Since the sky dips measure total power responsivity only, this complicates our use of sky dips to track relative polarized responsivity for the W-band array. As a result we rely on daily Tau A measurements of a single module to measure fast variations, and use the relative responsivities between the central module and the other modules from additional Tau A measurements and an end-of-season polarization grid measurement to extrapolate absolute responsivities for all modules.

Additional laboratory studies performed after deployment explain why the W-band modules were operated in a compressed regime: inactive components in the W-band modules had as much as twice the expected loss. To compensate for this loss, the amplifiers were biased higher than optimal such that the bias power alone is a significant contribution to the power required to compress the signal. We have produced modules with new inactive components with lower loss. These modules exhibit little compression and have noise temperatures closer to the isolated W-band amplifier noise ( $\sim 50$  K).

#### 8.5. Detector Angles

Absolute polarized detector angles are measured for the central module of each array through observations of from Tau A, whose polarization angle is known to  $0.2^\circ$  precision from IRAM measurements (Aumont et al. 2010). For the Q-band array, the absolute angle shifted by as much as  $2^\circ$  due to jumps in the pointing data from a loose encoder during the first half of the Q-band season. The systematic uncertainties related to the en-

TABLE 12  
AVERAGE BANDWIDTHS AND CENTRAL FREQUENCIES FOR THE Q-BAND AND W-BAND ARRAYS.

Band	Bandwidth (GHz)			Central Frequency (GHz)		
	Value	Stat error	System. error	Value	Stat. error	System. error
Q	7.6	0.5	1.0	43.1	0.4	1.0
W	10.7	-	1.1	94.5	-	0.8

coder jumps are discussed in (QUIET Collaboration et al. 2011). The Q-band array relied on weekly Moon observations and an end-of-season sparse wire grid measurement to find the relative polarized angles between the modules. The relative angle between the central module and a given module from all  $\sim 35$  Moon measurements have a variance of  $0.2^\circ$ , indicating that the relative angles remained nearly constant during the season. Relative detector angles are not effected by the encoder jumps.

The W-band array had a smaller, more efficient Tau A scan trajectory and was able to make measurements with all modules over the course of the season to obtain absolute angle calibration. The variance of detector angles for the central module from Tau A is  $0.3^\circ$ . The relative angles between modules were confirmed with end-of-season wire-grid measurements for both arrays to within the  $0.9^\circ$ .

Relative angles for all modules are shown in Figure 8.5 W-band array. Systematic errors in the absolute angle are the largest source of systematic errors for the W-band array, which would limit our measurement of  $r$  to XX at  $\ell \sim 100$ .

### 8.6. Noise Spectra

Noise measurements at the site were obtained from a noise spectrum fit to the Fourier-transform of the double-demodulated time stream for each CES. The measured noise floor should be proportional to the combination of module noise temperature, atmospheric temperature, optical spillover temperature, and CMB temperature. We assume a power law with a flat noise floor for the functional form of the noise spectrum:

$$N(\nu) = \sigma_0 \left[ 1 + \left( \frac{\nu}{\nu_{\text{knee}}} \right)^\alpha \right], \quad (20)$$

where  $N(\nu)$  and  $\sigma_0$  have units  $\text{V}/\sqrt{\text{Hz}}$ ,  $\nu$  is frequency in Hertz,  $\sigma_0$  is the white noise level,  $\alpha$  is the slope of the low frequency end of the spectrum, and  $\nu_{\text{knee}}$  is the knee frequency. A typical noise power spectrum for a W-band module is given in Figure 25, which also shows the effects on the noise of demodulating and double demodulating the time streams. After double demodulation, the median knee frequency is 5.5mHz (10mHz) for the Q-band (W-band) array, thus the noise is white at the scan frequencies of the telescope, 45–100 mHz.

The white noise is correlated among detector diodes within a given module. The correlation between Q and U diodes is expected (Bischoff 2010a); the theoretical expectation and typical measured correlations are given in Table 13. The measured correlation coefficients are larger than theoretically anticipated, and the source is unknown but could come from unequal transmission in the coupling hybrid in the module, or from leakage of the atmosphere causing residual  $1/f$  noise. However, the




FIG. 24.— Detector angle differences among diodes on the same W-band module. **TO BE REVISED**




FIG. 25.— Typical noise spectrum of W-band module with no demodulation, single demodulation, and double demodulation. Double-demodulation reduces the knee frequencies below the telescope scan frequencies.

noise correlation among diodes is easily treated in the data analysis (QUIET Collaboration et al. 2011), and more importantly is not actually present in the polarized signal we measure, which is a difference between diodes:  $(Q_1 - Q_2)$  and  $(U_1 - U_2)$ .

### 8.7. System Noise

TABLE 13  
PREDICTED AND MEASURED CORRELATION COEFFICIENTS AMONG DIODES. THE ERRORS ARE THE STANDARD DEVIATION OF THE CORRELATION COEFFICIENTS BETWEEN MODULES.

Diode $\times$ Diode	Prediction	Typical Measured Value	
		Q-band	W-band
Q $\times$ Q	0	$0.23 \pm 0.09$	$0.06 \pm 0.19$
U $\times$ U	0	$0.22 \pm 0.08$	$0.06 \pm 0.21$
Q $\times$ U	0.5	$0.54 \pm 0.08$	$0.48 \pm 0.11$



The instrument noise for our system is given by:

$$T_{\text{instrument}} = \quad (21)$$

$$T_{\text{CMB}} + T_{\text{atm}} + \frac{T_{\text{M}}}{G_{\text{atm}}} + \frac{T_{\text{W}}}{G_{\text{atm}}G_{\text{M}}} +$$

$$\frac{T_{\text{IR}}}{G_{\text{atm}}G_{\text{M}}G_{\text{W}}} + \frac{T_{\text{H}}}{G_{\text{atm}}G_{\text{M}}G_{\text{W}}G_{\text{IR}}} +$$

$$\frac{T_{\text{SP}}}{G_{\text{atm}}G_{\text{M}}G_{\text{W}}G_{\text{IR}}G_{\text{H}}} +$$

$$\frac{T_{\text{module}}}{G_{\text{atm}}G_{\text{M}}G_{\text{W}}G_{\text{IR}}G_{\text{H}}G_{\text{SP}}},$$

where  $T_{\text{atm}}$  is the effective atmospheric temperature,  $G_{\text{atm}}$  is the transmission through the atmosphere ( $e^{-\tau}$  where  $\tau$  is the opacity from the atmospheric sky model)  $T_{\text{CMB}}$  is the temperature of the CMB in thermodynamic units,  $T_{\text{module}}$  is the noise temperature of a QUIET module,  $\{T_{\text{M}}, G_{\text{M}}\}$ ,  $\{T_{\text{W}}, G_{\text{W}}\}$ ,  $\{T_{\text{IR}}, G_{\text{IR}}\}$ ,  $\{T_{\text{H}}, G_{\text{H}}\}$ ,  $\{T_{\text{SP}}, G_{\text{SP}}\}$  are the temperature and gain for the mirror, window, IR blocker, horns, and septum polarizers, respectively (Table 14).

The instrument noise can be found from the total power time streams taken during sky dips. During a sky dip, the sky temperature as seen by the receiver changes with telescope elevation. Using an atmospheric model, the change in signal with this model-dependent change in sky temperature allows us to estimate the instrument noise. We found noise values of 26 K for the Q-band array, and 106 K for the W-band array. Histograms of instrument noise temperatures for both arrays are shown in Figure 26.

We can estimate the contribution to instrument noise due to the module alone by subtracting assumed or measured values for all other known instrument noise sources (Table 14). All components other than the modules are lossy; thus their noise temperatures are given by  $(\frac{1}{G} - 1) \times T_{\text{phys}}$ , where  $G$  is the gain of the component and  $T_{\text{phys}}$  is its physical temperature. The extrapolated module temperature is 15 K for a Q-band module, and 77 K for a W-band module. Measurements of the Q-band amplifiers give noise values of  $\sim 18$  K; the most likely source of the discrepancy is that we are overestimating the septum polarizer loss. Similar measurements for the W-band module give amplifier noise values of 50 K, although the discrepancy is probably also due to compression, which decreases signal-to-noise and hence noise temperature.

### 8.8. Instrument Sensitivity

The sensitivity for the polarization response,  $S_{\text{pol}}$  ( $\mu\text{K}\sqrt{\text{s}}$ ), is calculated as the ratio of the white noise level to the responsivity. For the Q-band array, after data selection (QUIET Collaboration et al. 2011), the sensitivity obtained from the white noise floor and the responsivity is  $69\mu\text{K}\sqrt{\text{s}}$  corresponding to an average module sensitivity of  $275\mu\text{K}\sqrt{\text{s}}$ . For the W-band array, the array sensitivity is  $78\mu\text{K}\sqrt{\text{s}}$ , corresponding to an average module sensitivity of  $674\mu\text{K}\sqrt{\text{s}}$ . Both values are given in thermodynamic units, such that the power detected by the receiver has been corrected from a Rayleigh-Jeans approximation to correspond to the blackbody temperature of the CMB. Functionally this is performed by dividing by  $C_{\text{RJ}}$ , which is 1.05 (1.26) for the Q-band (W-band)

TABLE 14

ESTIMATED CONTRIBUTIONS TO THE SYSTEM NOISE WHERE THE NOISE FROM EACH COMPONENT HAS BEEN DIVIDED BY THE GAIN OF THE PREVIOUS ELEMENTS IN THE OPTICAL CHAIN. THE VALUES FOR THE GAIN ARE NOT MEASURED AND ARE INCLUDED FOR ILLUSTRATIVE PURPOSES, THUS WE DO NOT ATTACH ERRORS TO THE ESTIMATES. THE ATMOSPHERIC TEMPERATURE AND LOSS WERE COMPUTED FOR AN ELEVATION OF  $66^\circ$  (THE MID POINT OF THE CES ELEVATION RANGE), AND A PWV OF 1.2 MM (Q-BAND) AND 0.94 MM (W-BAND). THE CMB AND ATMOSPHERIC TEMPERATURES ARE GIVEN AS RAYLEIGH-JEANS TEMPERATURES. ALL AMBIENT TEMPERATURES ARE TAKEN AS 270 K.

Description	Q-band		W-band	
	Gain	Noise(K)	Gain	Noise(K)
CMB+sky	0.96	11.1	0.98	5.9
Mirrors	0.99	2.7	0.99	2.7
Window	0.99	2.8	0.98	5.5
Horn	0.99	0.2	0.99	0.2
Septum Polarizer	0.9	2.5	0.9	2.3
TOTAL	0.83	19	0.85	17
Measured $T_{\text{instrument}}$	-	38	-	109
Implied module	-	15	-	77




FIG. 26.— System noise temperatures of W-band (main figure) and Q-band (inset) modules. Elevation-dependent atmospheric temperature and the CMB temperature have been removed. The Q-band noise temperatures are obtained from small sky dips and W-band noise temperatures are obtained from large angle sky dips.

central frequencies. We can compare these to values from the radiometer equation (Krauss 1986) for the sensitivity per module,  $S_{\text{pol}}$ :

$$S_{\text{pol}} = \frac{1}{C_{\text{RJ}}} \times \frac{T_{\text{instrument}}}{\sqrt{2\Delta\nu G_{\text{atm}}(1 - f_{\text{mask}})}}. \quad (22)$$

Using the measured values for  $T_{\text{instrument}}$  and the atmospheric gain,  $G_{\text{atm}}$  (Table 14), the bandwidths  $\Delta\nu$  (Section 8.3), the Rayleigh-Jeans correction  $C_{\text{RJ}}$  for the CMB, and the fraction of the data we mask during the phase switch transitions,  $f_{\text{mask}}$  (14%, Section 6), we find sensitivity values of  $277\mu\text{K}\sqrt{\text{s}}$  for the Q band, and  $772\mu\text{K}\sqrt{\text{s}}$  for the W-band. Errors in bandpasses and the atmospheric temperature contribute directly to the difference between the two methods of computing the sensitivity  $S_{\text{pol}}$ . For the W-band array, compression will significantly impact this estimate as well because the noise temperatures are derived from the total power stream during sky dips. As noted in Section 11, even 1% amplifier compression over a 2 K atmospheric temperature



change can cause as much as 35% compression in the total power responsivity, which can easily explain the  $\sim 15\%$  discrepancy between sensitivities derived from the noise floor of the CMB observations and from noise temperatures obtained from the total power signal from sky dips.

### 9. INSTRUMENTAL SYSTEMATICS

Excellent control of systematic errors is one of the primary design drivers of the QUIET instrument. The most important systematics for the Q-band instrument are temperature-to-polarization leakage, responsivity fluctuations, and sidelobe contamination (9). For the Q-band instrument, we also show the level of the tensor-to-scalar ratio  $r$  corresponding to the spurious B-mode polarization (Ref. QUIET Collaboration et al. (2011)). In this section, we summarize the primary sources of systematic errors.

Temperature to polarization leakage ( $I \rightarrow Q/U$  leakage) is the largest systematic error in the B-mode measurement for the Q-band instrument at  $\ell \sim 100$ , where it begins to dominate the constraint on  $r$  at levels of  $r < 0.1$ . A naive estimate of the impact of this leakage would cause it to dominate at a much higher level, however a combination of sky rotation and frequent boresight rotation suppresses this systematic by two orders of magnitude.

Monopole leakage in the Q-band system is due to two sources: an intentional bandpass mismatch between the septum polarizers and the modules as well as monopole leakage in the optical chain. We expected monopole leakage from the septum polarizers to dominate leakage from the optical chain, giving rise to a monopole leakage value of  $\sim 1\%$ . This is consistent with leakage values derived from Tau A measurements, as discussed in section 3.6. The band mis-match was introduced to prevent module oscillations, which occurred from a return loss spike near the low side of the module bandpass. The leakage from the differential loss spike at 47 GHz (section 5.1) preferentially leaks excess power into the Q diodes. In the W-band system we did not need to introduce the band mismatch between the septum polarizer and the module so the monopole leakage is dominated by the optics. Hence, the leakage is smaller ( $\sim 0.3\%$ ) and randomly distributed both in sign and amplitude (because leakage from the optics has no preferential direction relative to the individual module diode orientations), such that spurious B-mode power will be averaged down. Because the module-based leakage is removed during double-demodulation, the measured leakage is dominated by the optical chain so it is stable over time and can be well calibrated. This enables us to correct for and simulate the effects of the leakage during data analysis (Ref. QUIET Collaboration et al. (2011)).

Optics sidelobes are due to reflector surface imperfections, spillover past the secondary mirror, and a series of reflections between the edges of the secondary and primary mirrors. The antenna range measurement limits all three effects to  $-40$  dB within  $4^\circ$  of the main beam and to  $-60$  dB far from the main beam. The upper ground screen reduces the far sidelobes by at least an additional  $-20$  dB, confirmed by sidelobe measurements with the Sun. The combination of filtering and rejecting contaminated data makes sidelobe pick-up a negligible effect on the B-mode

measurement.

### 10. CONCLUSIONS

QUIET employs the largest HEMT-based receiver arrays to date. The 17-element Q-band array has a polarization sensitivity of  $69 \mu\text{Ks}^{1/2}$ , currently the most sensitive instrument in this band. The 84-element W-band array has a  $78 \mu\text{Ks}^{1/2}$  sensitivity. Together the two arrays give the instrument sensitivity to angular scales  $\ell \sim 25\text{--}950$ .

The QUIET instrument design also achieves extremely low systematic errors. The optical design utilizes high-gain, low-crosspolar, and low-sidelobe corrugated feed horns and septum polarizers. The entire receiver and mirrors are housed in an absorbing ground shield to reduce sidelobe pickup, and is mounted on a 3-axis telescope with boresight rotation. The polarimeter assemblies utilize electronic double modulation to remove both  $1/f$  noise and monopole leakage. Finally, the differential-temperature assemblies and calibration tools provide critical measurements and cross checks of the systematic errors. The dominant systematic errors at  $\ell \sim 100$  are leakage for the Q-band instrument, and detector angle calibration for the W-band instrument. QUIET's Q-band result has a systematic error of  $r < 0.1$  at  $\ell = 100$ , the lowest systematic uncertainty on  $r$  published to date. QUIET's W-band data has excellent prospects for improving upon this uncertainty.

Support for the QUIET instrument and operation comes through the NSF cooperative agreement AST-0506648. Support was also provided by NSF awards PHY-0355328, AST-0448909, AST-1010016, and PHY-0551142; KAKENHI 20244041, 20740158, and 21111002; PRODEX C90284; a KIPAC Enterprise grant; and by the Strategic Alliance for the Implementation of New Technologies (SAINT).

Some work was performed on the Joint Fermilab-KICP Supercomputing Cluster, supported by grants from Fermilab, the Kavli Institute for Cosmological Physics, and the University of Chicago. Some work was performed on the Titan Cluster, owned and maintained by the University of Oslo and NOTUR (the Norwegian High Performance Computing Consortium), and on the Central Computing System, owned and operated by the Computing Research Center at KEK. Portions of this work were performed at the Jet Propulsion Laboratory (JPL) and California Institute of Technology, operating under a contract with the National Aeronautics and Space Administration. The Q-band modules were developed using funding from the JPL R&TD program. We acknowledge the Northrop Grumman Corporation for collaboration in the development and fabrication of HEMT-based cryogenic temperature-compatible MMICs.

C.D. acknowledges an STFC Advanced Fellowship and an ERC IRG grant under FP7. R.B. acknowledges support from CONICYT project Basal PFB-06 and ALMA-Conicyt 31070015. A.D.M. acknowledges a Sloan foundation fellowship. H.K.E. acknowledges an ERC Starting Grant under FP7.

PWV measurements were provided by the Atacama Pathfinder Experiment (APEX). We thank CONICYT for granting permission to operate within the Chajnan-

TABLE 15  
SUMMARY OF INSTRUMENTAL SYSTEMATIC ERRORS

Effect	Primary Source	Excess BB power @ $\ell = 100$	
		Q-band	W-band
$I \rightarrow Q/U$ (monopole)	Septum polarizer and module	$5.8 \times 10^{-3}$	
$I \rightarrow Q/U$ (dipole/quadrupole)	Mirror optics	$5.7 \times 10^{-5}$	
Responsivity fluctuation	Backend electronics	$1.3 \times 10^{-4}$	
Sidelobe contamination	Optics / Ground screen	$< 1 \times 10^{-4}$	

tor Scientific Preserve in Chile, and ALMA for providing site infrastructure support. Field operations were based at the Don Esteban facility run by Astro-Norte. We are particularly indebted to the engineers and technician who maintained and operated the telescope: José Cortés, Cristobal Jara, Freddy Muñoz, and Carlos Verdugo.

In addition, we would like to acknowledge the following people for their assistance in the instrument design, construction, commissioning, operation, and in data analysis: Augusto Gutierrez Aitken, Colin Baines, Phil Banister, Hannah Barker, Matthew R. Becker, Alex Blein, April Campbell, Anushya Chandra, Sea Moon Cho, Joelle Cooperrider, Mike Crofts, Emma Curry, Maire Daly, Richard Davis, Fritz Dejongh, Joy Didier, Greg Dooley, Hans Eide, Will Grainger, Jonathon Goh, Peter Hamlington, Takeo Higuchi, Seth Hillbrand, Ben Hooberman, Kathryn D. Huff, Eiichiro Komatsu, Jostein Kristiansen, Donna Kubik, Richard Lai, David Leibovitch, Kelly Lepo, Siqi Li, Martha Malin, Mark McCulloch, Oliver Montes, David Moore, Ian O'Dwyer, Gustavo Orellana, Stephen Osborne, Stephen Padin, Felipe Pedreros, Ashley Perko, Alan Robinson, Jacklyn Sanders, Dale Sanford, Yunior Savon, Mary Soria, Alex Sugarbaker, David Sutton, Matias Vidal, Liza Volkova, Stephanie Xenos, Octavio Zapata, and Mark Zaskowski.

## 11. APPENDIX

### 11.1. Compression

This section explains some subtleties regarding nonlinearities, and how it affects the polarization and total power measurements differently. This complicates our use of periodic telescope skydips to track the total power responsivity, which we *assume* to also track the polarization responsivity. During CMB operations, the receiver load temperature varies by  $\sim 2\text{K}$  due to changes in the skyloading. Nonlinearities also affect the use of large skydip and moon signals to calibrate the total power responsivity. For HEMT LNAs, the nonlinearities typically encountered in our operating regime is compression.

We now analyze the effect of compression on polarization responsivity. Consider a horn looking at an unpolarized background at temperature  $T_0$ , where  $T_0 = T_{0x} = T_{0y}$ , with axes  $x$  and  $y$  defined by the septum polarizer. Given below are the  $Q_1$  diode measurements for the  $0^\circ$  and  $180^\circ$  leg B states, and the demod output (which is

the polarization measurement):

$$\begin{aligned}
 S_0(0^\circ) &= g_0 \cdot \left( \frac{1}{2}(T_{0x} + T_{0y}) + \frac{1}{2}(T_{0x} - T_{0y}) \right) \\
 S_0(180^\circ) &= g_0 \cdot \left( \frac{1}{2}(T_{0x} + T_{0y}) - \frac{1}{2}(T_{0x} - T_{0y}) \right) \\
 S_0(\text{demod}) &= \frac{1}{2}(S_0(0^\circ) - S_0(180^\circ)) \\
 &= \frac{1}{2} \cdot g_0 \cdot (T_{0x} - T_{0y}) = 0
 \end{aligned} \tag{23}$$

where  $g_0$  is the gain at temperature  $T_0$ .

Consider now the horn looking at a source on top of this background. Without loss of generality, let the source be polarized in the  $x$  direction at temperature  $T_1$  such that  $T_{1x} = T_{0x} + T_{Sx}$ ,  $T_{1y} = T_{0y}$ ,  $T_{avg} = \frac{1}{2} \cdot (T_{0x} + T_{Sx} + T_{0y})$ . Then:

$$\begin{aligned}
 S_1(0^\circ) &= g_1 \cdot (T_{avg} + \frac{1}{2}(T_{0x} + T_{Sx} - T_{0y})) \\
 S_1(180^\circ) &= g_1 \cdot (T_{avg} - \frac{1}{2}(T_{0x} + T_{Sx} - T_{0y})) \\
 S_1(\text{demod}) &= \frac{1}{2} \cdot g_1 \cdot T_{Sx}
 \end{aligned} \tag{24}$$

Note that gain constant  $g_1$  is relevant for the temperature  $T_{avg}$ , for the following reason. Since the incident E-fields at the horn input are linearly polarized, the septum polarizer splits the power equally between legs A and B. And the legs see a constant input power given by  $T_{avg}$ , regardless of the beam's position angle. Within the module, the LNA's are placed upstream of any phase-sensitive circuitry. In our model, compression depends primarily on the input power at the first LNA. Therefore, the first LNA sees power represented by  $T_{avg}$ , so we assign the gain  $g_1$  to that temperature. Thus equation 24 shows that the polarization measurement is compressed by  $(g_0 - g_1)/g_1$ . Using Moon data estimates,  $(g_0 - g_1)/g_1$  changes roughly by 0.1% per Kelvin.

We consider now the effect of compression on the total power responsivity. For an unpolarized background source at temperature  $T_0$ , the  $Q_1$  diode voltages for the leg B  $0^\circ$  and  $180^\circ$  states, and the average (which gives us the total power) are:

$$\begin{aligned}
 S_0(0^\circ) &= g_0 \cdot \left( \frac{1}{2}(T_{0x} + T_{0y}) + \frac{1}{2}(T_{0x} - T_{0y}) \right) \\
 S_0(180^\circ) &= g_0 \cdot \left( \frac{1}{2}(T_{0x} + T_{0y}) - \frac{1}{2}(T_{0x} - T_{0y}) \right) \\
 S_0(\text{avg}) &= \frac{1}{2}(S_0(0^\circ) + S_0(180^\circ)) \\
 &= \frac{1}{2} \cdot g_0 \cdot (T_{0x} + T_{0y}) = g_0 \cdot T_0
 \end{aligned} \tag{25}$$

Similarly, an unpolarized background source at temperature  $T_1$  results in:

$$S_1(avg) = g_1 \cdot T_1 \quad (26)$$

Here,  $g_1$  and  $g_0$  are the gains for temperatures  $T_1$  and  $T_0$  respectively. It can be shown that:

$$S_1(avg) - S_0(avg) = g_1(T_1 - T_0)c \quad (27)$$

$$c = (1 - \frac{g_0 - g_1}{g_1} \frac{T_0}{T_1 - T_0})$$

where  $c$  is ratio between the observed signal difference and the expected difference without compression.

Comparing equation 27 with 24, the total power sensitivity compression is magnified by  $T_0/(T_1 - T_0)$ . Assuming as an example,  $T_1 - T_0 = 2\text{K}$  (typical for a skydip), a system temperature of  $T_0 = 120\text{K}$ , and a typical gain compression of  $(g_0 - g_1)/g_1 = 0.002$  over that range, the resulting ratio is  $c = 93\%$ , or 7% signal loss. Therefore, in the data analysis, the absolute responsivities are derived

from polarized source measurements to avoid systematic biases of this type.

### 11.2. Double demodulation

This section discusses some imperfections in the module, and resolving them using double demodulation. Table 9 shows the detector diode outputs of an ideal module for the two leg B states, with the leg A state held fixed. The idealization (see Figure 15(a)) assumes equal transmission between the two leg B states, the two leg A states, and an ideal septum polarizer (see 11.3). In practice, the transmissions are unequal, thus requiring extra parameters to describe the module. With out loss of generality, let the transmission through the  $0^\circ(\uparrow)$  state of legs A and B be equal to unity, and define  $\beta_A$  and  $\beta_B$  to be the transmissions through these legs for the  $180^\circ(\downarrow)$  state. We define  $g_A$  and  $g_B$  to be the effective gains of the two legs (see fig. 27). In this more general case, the detector diode voltages are given by:

$$V_{Q_1}(V_{Q_2}) = \frac{1}{4} \left\{ \begin{array}{lll} \frac{1}{2}(g_A^2 + g_B^2)I & + \frac{1}{2}(g_A^2 - g_B^2)V & \pm g_A g_B Q \\ \frac{1}{2}(g_A^2 + g_B^2 \beta_B^2)I & + \frac{1}{2}(g_A^2 - g_B^2 \beta_B^2)V & \mp g_A g_B \beta_B Q \\ \frac{1}{2}(g_A^2 \beta_A^2 + g_B^2)I & + \frac{1}{2}(g_A^2 \beta_A^2 - g_B^2)V & \mp g_A g_B \beta_A Q \\ \frac{1}{2}(g_A^2 \beta_A^2 + g_B^2 \beta_B^2)I & + \frac{1}{2}(g_A^2 \beta_A^2 - g_B^2 \beta_B^2)V & \pm g_A g_B \beta_A \beta_B Q \end{array} \right\} \quad (28)$$

$$V_{U_1}(V_{U_2}) = \frac{1}{4} \left\{ \begin{array}{lll} \frac{1}{2}(g_A^2 + g_B^2)I & + \frac{1}{2}(g_A^2 - g_B^2)V & \mp g_A g_B U \\ \frac{1}{2}(g_A^2 + g_B^2 \beta_B^2)I & + \frac{1}{2}(g_A^2 - g_B^2 \beta_B^2)V & \pm g_A g_B \beta_B U \\ \frac{1}{2}(g_A^2 \beta_A^2 + g_B^2)I & + \frac{1}{2}(g_A^2 \beta_A^2 - g_B^2)V & \pm g_A g_B \beta_A U \\ \frac{1}{2}(g_A^2 \beta_A^2 + g_B^2 \beta_B^2)I & + \frac{1}{2}(g_A^2 \beta_A^2 - g_B^2 \beta_B^2)V & \mp g_A g_B \beta_A \beta_B U \end{array} \right\} \quad (29)$$

where the upper (lower) signs correspond to the signal of the diodes  $Q_1$  and  $U_1$  ( $Q_2$  and  $U_2$ ). The four rows for each  $V_i$  correspond to the phase switch states of (A,B) = ( $\uparrow, \uparrow$ ), ( $\uparrow, \downarrow$ ), ( $\downarrow, \uparrow$ ), and ( $\downarrow, \downarrow$ ), from the top to the bottom. Transmission imbalance between the phase switches, signified by the deviation from unity of  $\beta_A$  and  $\beta_B$ , causes  $I \rightarrow Q/U$  leakage. This can be seen in Table 16, showing the demodulated output dependences on  $I \cdot (1 - \beta_B^2)$ . However, the difference between the  $\uparrow$  and  $\downarrow$  demodulated outputs is free from  $I$ -dependence. This operation is referred to as *double-demodulation*. For the W-band, the rms of  $I \rightarrow Q/U$  leakage distribution reduces from roughly 0.8% in the demodulated stream, to 0.4% in the double-demodulated stream. A smaller reduction of  $< 0.1\%$  is found for the Q-band, as the leakage is dominated by other effects.

### 11.3. Polarimeter Assembly Offset and $I \rightarrow Q/U$ Leakage




FIG. 27.— Schematic of signal processing in a QUIET polarimeter assembly.

As shown in Section 11.2, we can assume that the module does not generate any instrumental polarization on its own, since the double demodulation procedure nulls out this effect. However, the interaction between the module and septum polarizer can cause irreducible instrumental polarization and offsets; this section derives these couplings. Since the module does not generate instrumental

TABLE 16

DEMOMULATED SIGNAL FOR EACH LEG A PHASE STATE, WITH THE LEG B PHASE STATE SWITCHING AT 4 KHz. A FACTOR OF 1/4 HAS BEEN OMITTED FROM EACH EXPRESSION. THE TERMS INVOLVING STOKES V ARE ALSO OMITTED FOR SIMPLICITY. THE UPPER (LOWER) SIGNS CORRESPOND TO DETECTOR DIODES  $Q_1$  AND  $U_1$  ( $Q_2$  AND  $U_2$ ).

Leg A phase state	Diode	Demodulated output
$\uparrow$	$Q_1$ ( $Q_2$ )	$\frac{1-\beta_B^2}{4}g_B^2I \pm \frac{1+\beta_B}{2}g_Ag_BQ$
$\uparrow$	$U_1$ ( $U_2$ )	$\frac{1-\beta_B^2}{4}g_B^2I \mp \frac{1+\beta_B}{2}g_Ag_BU$
$\downarrow$	$Q_1$ ( $Q_2$ )	$\frac{1-\beta_B^2}{4}g_B^2I \mp \beta_A \frac{1+\beta_B}{2}g_Ag_BQ$
$\downarrow$	$U_1$ ( $U_2$ )	$\frac{1-\beta_B^2}{4}g_B^2I \pm \beta_A \frac{1+\beta_B}{2}g_Ag_BU$

polarization on its own, the module measures:

$$Q_m = 2\Re(L_m^* R_m) \quad (30)$$

and

$$U_m = -2\Im(L_m^* R_m), \quad (31)$$

where  $L_m$  and  $R_m$  are the signals transmitted into the module inputs. With out loss of generality, we absorb all constant factors into the responsivity, and set it to unity. The signals transmitted into the module inputs need not be the same as the  $L$  and  $R$  components at the septum polarizer input; this difference is a cause of instrumental polarization.

The effect of the septum polarizer is described by a 4x4 complex scattering matrix  $S$ :

$$\begin{pmatrix} E'_x \\ L' \\ R' \\ E'_y \end{pmatrix} = S \cdot \begin{pmatrix} E_x \\ L_r \\ R_r \\ E_y \end{pmatrix} \quad (32)$$

$$S = \begin{pmatrix} r_1 & \frac{e^{i\gamma}}{\sqrt{2}}\tau_{21} & \frac{e^{i\gamma}}{\sqrt{2}}\tau_{31} & r_{41} \\ \frac{e^{i\gamma}}{\sqrt{2}}\tau_{21} & r_2 & c & i\frac{e^{i\gamma}}{\sqrt{2}}\tau_{24} \\ \frac{e^{i\gamma}}{\sqrt{2}}\tau_{31} & c & r_3 & -i\frac{e^{i\gamma}}{\sqrt{2}}\tau_{34} \\ r_{41} & i\frac{e^{i\gamma}}{\sqrt{2}}\tau_{24} & -i\frac{e^{i\gamma}}{\sqrt{2}}\tau_{34} & r_4 \end{pmatrix} \quad (33)$$

where  $E_x$  and  $E_y$  are electric field components at the septum polarizer input port,  $L'$  and  $R'$  are the fields at the two septum polarizer output ports.  $E'_x$  and  $E'_y$  are the electric fields emitted from the septum polarizer back toward the feed horn.  $L_r$  and  $R_r$  are signals reflected (or emitted) from the module inputs travelling back toward the septum polarizer output ports. Here,  $e^{i\gamma}$  is the propagation phase shift,  $\tau_{ij}$  and  $r$  are transmission and the reflection coefficients respectively.  $c$  is the cross-polarization coefficient for the backward travelling fields. For an ideal septum polarizer,  $\tau_{ij} = 1$  and  $r_i = r_{41} = c = 0$ . Symmetry across the septum implies  $\tau_{21} = \tau_{31}$ ,  $\tau_{24} = \tau_{34}$ ,  $r_2 = r_3$ , and  $r_{41} = 0$ ; although manufacturing errors can cause these conditions to be violated. As described in Sections 5.1 and 8 there are small departures from ideality; in this section we consider such effects up to second order. Note that the scattering matrix is frequency-dependent. The analysis given here is strictly for a single frequency; in practice, the result should be averaged with the effective bandpass.

We shall derive perturbatively  $L_m$  and  $R_m$ , the fields transmitted into the module inputs due to a sky source consisting of fields  $E_x$  and  $E_y$ . We also assume a noiseless module. The case of a noise signal from the module is described later. To lowest order, the  $S$  matrix applied to the column vector  $(E_x, 0, 0, E_y)$  yields  $(0, L', R', 0)$ , where

$$\begin{aligned} L' &= \frac{e^{i\gamma}}{\sqrt{2}} \left( \tau_{21} \frac{L+R}{\sqrt{2}} + \tau_{24} \frac{L-R}{\sqrt{2}} \right) \\ &= \frac{e^{i\gamma}}{2} [(\tau_{21} + \tau_{24})L + (\tau_{21} - \tau_{24})R]. \end{aligned} \quad (34)$$

Similarly,

$$R' = \frac{e^{i\gamma}}{2} [(\tau_{31} + \tau_{34})R + (\tau_{31} - \tau_{34})L]. \quad (35)$$

where we use  $E_x = (L+R)/\sqrt{2}$  and  $E_y = (L-R)/(i\sqrt{2})$ . However,  $L_m$  and  $R_m$  differ from  $L'$  and  $R'$  due to reflection at the module input. Let  $r_L$  ( $r_R$ ) be the reflection coefficient at the module's  $L$  ( $R$ ) input. Then the  $S$  matrix applied to  $(E_x, r_L L', r_R R', E_y)$  yields  $(-, L_m, R_m, -)$  where

$$L_m = (1 + r_2 r_L) L' + c r_R R'. \quad (36)$$

$$R_m = (1 + r_3 r_R) R' + c r_L L'. \quad (37)$$

and for simplicity, we omit giving the expression for the first and fourth component. The module output is

$$L_m^* R_m = L'^* R' (1 + r_3 r_R + r_2^* r_L^*) + L'^* L' c r_L + R'^* R' c^* r_R^* \quad (38)$$

where we assume that  $r_i$ ,  $c$ ,  $r_R$  and  $r_L$  are all small, and we drop terms above second order.

We now simplify the RHS of equation 38 into the underlying physics parameters  $Q$  and  $U$ , in order to identify the sources of instrumental polarization. The terms  $L'^* L'$  and  $R'^* R'$  need only be calculated to leading order, since in they appear in equation 38 multiplied by the second order terms  $c r_L$  and  $c^* r_R^*$ . To leading order,  $L'^* L' = L^* L$  and  $R'^* R' = R^* R$  since  $\tau_{ij} \approx 1$ .

The first term in Eq. 38 is expanded by substituting equation 34 and 35 and using  $L^* R = (Q - iU)/2$ ,  $LL^* = (I + V)/2$ , and  $RR^* = (I - V)/2$  to obtain

$$\begin{aligned} L'^* R' &= \frac{1}{4} ((\tau_{21}^* \tau_{31} + \tau_{24}^* \tau_{34})Q - i(\tau_{21}^* \tau_{34} + \tau_{24}^* \tau_{31})U \\ &\quad + (\tau_{21}^* \tau_{31} - \tau_{24}^* \tau_{34})I + (\tau_{24}^* \tau_{31} - \tau_{21}^* \tau_{34})V). \end{aligned} \quad (39)$$

The first two terms are the expected response to  $Q$  and  $U$ . The presence of  $\tau_{ij}$  in these terms parameterizes the imperfections in the septum polarizer transmissions. These terms reduce the gain to  $Q$  and  $U$ , and in general cause mixing between  $Q$  and  $U$ . In practice, we absorb the gain into the calibration of the total system responsivity<sup>35</sup>, and absorb the induced  $Q/U$  mixing into a redefinition of the detector angle. Therefore, these two terms do not cause instrumental polarization, and we can neglect these imperfections in the following discussion. By the same argument, the terms  $r_3 r_R$  and  $r_2^* r_L^*$  in equation 38 can be ignored since their only effect is to change the gain.

<sup>35</sup> The effect of the gain difference on the leakage terms is a neglected higher order effect.

The third and fourth terms represents  $I \rightarrow Q/U$  and  $V \rightarrow Q/U$  leakage respectively. Since  $V \ll I$  for reasonable sources and the coefficients have the same order, we neglect these circular polarization terms. Combining these observations, we finally arrive at the desired expression for the RHS of equation 38:

$$L_m^* R_m = \frac{1}{4} [2Q - 2iU + (\tau_{21}^* \tau_{31} - \tau_{24}^* \tau_{34})I] + \frac{1}{2} (c r_L + c^* r_R^*) I. \quad (40)$$

The module output is:

$$Q_m = Q + \frac{1}{2} \Re(\tau_{21}^* \tau_{31} - \tau_{24}^* \tau_{34}) I + \Re(c r_L + c^* r_R^*) I, \quad (41)$$

where the first term is the expected response, the second term is  $I \rightarrow Q$  leakage due to differential loss, and the third term is leakage caused by reflections at the module inputs coupling with the septum polarizer crosstalk. Similarly:

$$U_m = U - \frac{1}{2} \Im(\tau_{21}^* \tau_{31} - \tau_{24}^* \tau_{34}) I - \Im(c r_L + c^* r_R^*) I. \quad (42)$$

In summary, the two equations above describe the measurements of a sky signal, in the absence of noise from the module.

We now consider the case of noise emitted from the module inputs, reflecting from the septum polarizer and returning into the module. Module noise stems primarily from the HEMT-based first stage LNA's. Since the sky signal and module noise are relatively incoherent, they decouple and we can neglect the sky signal in the following. Let the module noise fields be given by the column vector  $(0, L_r, R_r, 0)$ . Applying the  $S$  matrix, we get the vector  $(-, L_m, R_m, -)$  where

$$L_m = L' = r_2 L_r + c R_r \quad (43)$$

$$R_m = c L_r + r_3 R_r. \quad (44)$$

The output is

$$L_m^* R_m = r_2^* L_r^* c L_r + c^* R_r^* r_3 R_r \quad (45)$$

because the  $L_r R_r$  terms average to zero because the two amplifier noises are uncorrelated. Thus each output acquires an offset

$$Q_m = 2L_r^* L_r \Re(r_2^* c) + 2R_r^* R_r \Re(c^* r_3). \quad (46)$$

$$U_m = -2L_r^* L_r \Im(r_2^* c) - 2R_r^* R_r \Im(c^* r_3) \quad (47)$$

The offset is independent of the input  $I$ ; however, it is modulated by gain fluctuations so the offset also contributes to 1/f noise.

## REFERENCES

- Aumont, J., et al. 2010, *Astronomy and Astrophysics*, 514
- Barkats, D. et al. 2005, *ApJS*, 159, 1
- Benford, D. J., Gaidis, M. C., & Kooi, J. W. 2003, *Appl. Opt.*, 42, 5118
- Bischoff, C. 2010a, Ph.D. Dissertation, University of Chicago
- . 2010b, Ph.D. Thesis, University of Chicago
- Bogdan, M., Kapner, D., Samtleben, D., & Vanderlinde, K. 2007, *Nuclear Instruments and Methods in Physics Research A*, 572, 338
- Bornemann, J., & Labay, V. A. 1995, *IEEE Trans. MTT*, 43, 95
- Chang, S., & Prata, A. 2004, *IEEE Transactions on Antennas and Propagation*, 52, 12
- Clark, A. F., Childs, G. E., & Wallace, G. H. 1970, *Cryogenics*, 10, 295
- Clarricoats, P. J. B., & Olver, A. D. 1984, *Corrugated horns for microwave antennas* (London, UK: Peter Peregrinus Ltd.)
- Cleary, K. A. 2010, *Proc. SPIE*, 7741, 77412H
- QUIET Collaboration, T. Q. 2012, Submitted to *ApJ*
- Dicke, R. H. 1946, *Rev. Sci. Instr.*, 17, 268
- Dodelson, S., et al. 2009, *arXiv:0902.3796*
- Dragone, C. 1978, *Bell Syst. Tech. J.*, 57, 2663
- Farese, P. C. et al. 2004, *ApJS*, 610, 625
- Gundersen, J., & Wollack, E. 2009, *Journal of Physics Conference Series*, 155, 012005
- Güsten, R., Nyman, L. Å., Schilke, P., Menten, K., Cesarsky, C., & Booth, R. 2006, *A&A*, 454, L13
- Hoppe, D. 1987, *Scattering Matrix Program for Circular Waveguide Junctions*, Tech. Rep. NASA-CR-179669, NTO-17245, National Aeronautics and Space Administration, Georgia: NASA's Computer Software Management and Information Center, in *Cosmic Software Catalog*
- Hoppe, D. J. 1988, *Modal Analysis Applied to Circular, Rectangular and Coaxial Waveguides*, Tech. Rep. Telecommunications and Data Acquisition Progress Report 42-95, Jet Propulsion Laboratory, Pasadena, California
- Imbriale, W. A., Gundersen, J., & Thompson, K. L. 2011, *IEEE Transactions on Antennas and Propagation*, 59, 1972
- Imbriale, W. A., & Hodges, R. E. 1991, *Applied Computational Electromagnetics Society*, 6, 74
- Jarosik, N. et al. 2003a, *ApJS*, 148, 29
- . 2003b, *ApJS*, 145, 413
- Kamionkowski, M., Kosowsky, M., & Stebbins, A. 1997, *Phys. Rev. D*, 55, 7368
- Kangaslahti, P. et al. 2006, in *Microwave Symposium Digest*, 2006. IEEE MTT-S International, 89–92
- Kay, A. F. 1962, *A wide flare angle horn. A novel feed for low noise broadband and high aperture efficiency antennas*, Tech. Rep. US Air Force Cambridge Research Labs Report 62-757, US Air Force Cambridge Research Lab, Cambridge, Massachusetts
- Krauss, J. D. 1986, *Radio Astronomy (Cygnus-Quasar)*
- Lamb, J. W. 1996, *International Journal of Infrared and Millimeter Waves*, 17, 1997
- Lawrence, C. R., Gaier, T. C., & Seiffert, M. 2004, in *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, Vol. 5498, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. C. M. Bradford, P. A. R. Ade, J. E. Aguirre, J. J. Bock, M. Dragovan, L. Duband, L. Earle, J. Glenn, H. Matsuhara, B. J. Naylor, H. T. Nguyen, M. Yun, & J. Zmuidzinas, 220–231
- Leitch, E. M. et al. 2002, *Nature*, 420, 763
- Liddle, A. R., & Lyth, D. H. 2000, *Cosmological Inflation and Large-Scale Structure* (UK: Cambridge University Press)
- Meinhold, P. R., Chingcuanco, A. O., Gundersen, J. O., Schuster, J. A., Seiffert, M. D., Lubin, P. M., Morris, D., & Villela, T. 1993, *ApJ*, 406, 12
- Mennella, A., Bersanelli, M., Seiffert, M., Kettle, D., Roddis, N., Wilkinson, A., & Meinhold, P. 2003, *Astronomy and Astrophysics*, 410, 1089
- Mizugutch, Y., Akagawa, M., & Yokoi, H. 1976, *IEEE Antennas and Propagation Society International Symposium*, 14, 2
- Monsalve, R. A. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7741, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- O'Dea, D., Challinor, A., & Johnson, B. R. 2007, *MNRAS*, 376, 1767
- Padin, S. et al. 2002, *PASP*, 114, 83
- Page, L. et al. 2003a, *ApJS*, 148, 39
- . 2003b, *ApJ*, 585, 566
- Pardo, J., Cernicharo, J., & Serabyn, E. 2001, *IEEE Transactions on Antennas and Propagation*, 49, 1683
- QUIET Collaboration et al. 2011, *ApJ*, 741, 111
- Radford, S. J., & Holdaway, M. A. 1998, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 3357, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. T. G. Phillips, 486–494
- Seljak, U., & Zaldarriaga, M. 1997, *Phys. Rev. Lett.*, 78, 2054

- Staggs, S. T., Barkats, D., Gundersen, J. O., Hedman, M. M., Herzog, C. P., McMahon, J. J., Vanderlinde, K., & Winstein, B. 2003, AIP Conference Proceedings, 609, 183
- Tajima, O., Nguyen, H., Bischoff, C., Brizius, A., Buder, I., & Kusaka, A. 2012, Journal of Low Temperature Physics, 167, 936
- Ulich, B. L., Cogdell, J. R., & Davis, J. H. 1973, Icarus, 19, 59
- Weiland, J. L. et al. 2011, ApJS, 192, 19

### **Paper III**

---

*Bayesian noise estimation for non-ideal CMB experiments*

I. K. Wehus, S. K. Næss, H. K. Eriksen

ApJS, Volume 199, Issue 1, article id. 15 (2012)

---

## BAYESIAN NOISE ESTIMATION FOR NON-IDEAL CMB EXPERIMENTS

I. K. WEHUS<sup>1,2</sup>, S. K. NÆSS<sup>3</sup> AND H. K. ERIKSEN<sup>3,4</sup>*Draft version October 7, 2011*

## ABSTRACT

We describe a Bayesian framework for estimating the time-domain noise covariance of CMB observations, typically parametrized in terms of a  $1/f$  frequency profile. This framework is based on the Gibbs sampling algorithm, which allows for exact marginalization over nuisance parameters through conditional probability distributions. In this paper we implement support for gaps in the data streams and marginalization over fixed time-domain templates, and also outline how to marginalize over confusion from CMB fluctuations, which may be important for high signal-to-noise experiments. As a by-product of the method, we obtain proper constrained realizations, which themselves can be useful for map making. To validate the algorithm, we demonstrate that the reconstructed noise parameters and corresponding uncertainties are unbiased using simulated data. The CPU time required to process a single data stream of 100 000 samples with 1000 samples removed by gaps is 3 seconds if only the maximum posterior parameters are required, and 21 seconds if one also want to obtain the corresponding uncertainties by Gibbs sampling.

*Subject headings:* cosmic microwave background — cosmology: observations — methods: statistical

## 1. INTRODUCTION

Detailed observations of the cosmic microwave background (CMB) during the last two decades have revolutionized cosmology. Through detailed measurements of the angular CMB power spectrum, a highly successful cosmological concordance model has been established, stating that the universe is statistically isotropic and homogeneous, filled with Gaussian random fluctuations drawn from a  $\Lambda$ CDM spectrum, and consists of 4% baryonic matter, 23% dark matter and 73% dark energy (e.g., Komatsu et al. 2011, and references therein). Using this model, millions of data points from many different types of cosmological observations can be fitted with only six free parameters.

This success has been driven primarily by rapid progress in CMB detector technology, allowing experimentalists to make more and more detailed maps of the CMB fluctuations. However, such maps are imperfect, in the sense that they typically are contaminated by various instrumental effects. For instance, the optics of a given experiment can be asymmetric; the detector gain may be unknown and time-dependent; the data may exhibit resonant frequencies due to electronics or cooling non-idealities; and the observations are invariably noisy. All these non-idealities must be properly understood before one can attempt to extract cosmology from the observations.

In this paper we consider one specific component within this global calibration problem, namely how to

estimate the statistical properties of the instrumental noise in light of real-world complications. This problem has of course already been addressed repeatedly in the literature (e.g., Prunet et al. 2001; Hinshaw et al. 2003), and our method is in principle similar to that of Ferreira & Jaffe (2000), taking a Bayesian approach to the problem. The main difference is that we formulate the algorithm explicitly in terms of a Gibbs sampler including both the time stream and the noise parameters as unknown variables, and this has several distinct advantages. First, it allows us to obtain proper uncertainties on all derived quantities. Second, gap filling is directly supported through built-in proper constrained realizations. This can for instance be used to account for instrumental glitches in the time stream, or to exclude point sources and other bright sources from the analysis. Third, it is straightforward to add support for additional nuisance parameters, due to the conditional nature of the Gibbs sampler. In this paper we implement template marginalization, which may for instance be useful for removing cosmic ray glitches in the Planck HFI data (Planck 2011b) or ground pickup for ground based experiments (e.g., QUIET 2011). We also outline the formalism for marginalization over CMB fluctuations, which may be relevant for experiments with high signal-to-noise ratio.

The method presented here is mathematically identical to the CMB Gibbs sampling framework developed by Jewell et al. (2004); Wandelt et al. (2004); Eriksen et al. (2004, 2008), and the main difference is simply that the object under consideration is a one-dimensional time stream instead of a two-dimensional field on the sphere. This makes the implementation considerably simpler, and the run times correspondingly faster. As a demonstration of the practicality of the method, we apply it to simulated data with properties typical for current ground-based experiments, and demonstrate explicitly that the computational costs of the method are tractable. The experiment of choice will be QUIET

i.k.wehus@fys.uio.no  
sigurdkn@astro.uio.no  
h.k.k.eriksen@astro.uio.no

<sup>1</sup> Theoretical Physics, Imperial College London, London SW7 2AZ, UK

<sup>2</sup> Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway

<sup>3</sup> Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

<sup>4</sup> Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway



(2011), for which this method was initially developed.

## 2. DATA MODEL

The first step of any Bayesian analysis is to write down an explicit parametric model for the observations in question. In this paper, we start with the assumption that the output,  $\mathbf{d}$ , from a given detector can be written in terms of the following sum,

$$\mathbf{d} = \mathbf{n} + \mathbf{P}\mathbf{s} + \mathbf{T}\mathbf{a} + \mathbf{m}. \quad (1)$$

Here each term indicates a vector of  $n$  values sampled regularly in time in steps of  $\Delta t$ , that is,  $\mathbf{d} = \{d_i\}$  with  $i = 1, \dots, N$ .

The first term on the right-hand side,  $\mathbf{n}$ , indicates the instrumental noise, which is our primary target in this paper. All the other components are only nuisance variables that we want to marginalize over.

We assume that the noise is Gaussian distributed and stationary over the full time range considered. In practice this means that the full data set of a given experiment should be segmented into parts which are individually piecewise stationary. For QUIET this corresponds to division into so-called “constant elevation scans” (QUIET 2011), while for Planck it corresponds to division into so-called “rings”, which are one-hour observation periods with a fixed satellite spin axis (Planck 2011a). Because the noise is assumed stationary, the time-domain noise covariance matrix,  $\mathbf{N}$ , depends only on the time lag between two observations,  $N_{tt'} = N(t - t')$ : It is a Toeplitz matrix, and may therefore be well approximated in Fourier domain with a simple diagonal matrix,  $N_{\nu\nu'} = N_{\nu}\delta_{\nu\nu'}$ . Here  $N_{\nu}$  is the Fourier-domain noise power spectrum, which is given by the Fourier transform of  $N(t - t')$ .

Our main task is to estimate  $N_{\nu}$ , and we do so in terms of a parametrized function. For many experiments this function is well approximated by a so-called  $1/f$  profile,

$$N_{\nu} = \sigma_0^2 \left[ 1 + \left( \frac{\nu}{f_{\text{knee}}} \right)^{\alpha} \right] \quad (2)$$

which describes a sum of a correlated and an uncorrelated noise component in terms of three free parameters. The white-noise RMS level,  $\sigma_0$ , defines the overall amplitude of the noise; the knee frequency,  $f_{\text{knee}}$ , indicates where the correlated and the uncorrelated components are equally strong, and  $\alpha$  is the spectral index of the correlated component. Collectively, we denote  $\{\sigma_0, \alpha, f_{\text{knee}}\}$  by  $\theta$ . Of course, other parametrizations may easily be implemented if necessary.

The second term on the right-hand side,  $\mathbf{P}\mathbf{s}$ , indicates the contribution from the CMB sky, with  $\mathbf{P}$  being a pointing matrix, typically equal to zero everywhere except at  $P_{ip}$  if the detector points towards pixel  $p$  at time  $i$ , and  $s_p$  is the true (beam convolved) CMB signal. We make the usual assumption that  $\mathbf{s}$  is isotropic and Gaussian distributed with a given angular power spectrum,  $C_{\ell}$ . In this paper, we will simply outline the formalism for how to deal with this term, and leave the implementation for a future paper dedicated to Planck analysis; as mentioned in the introduction, this machinery was initially developed QUIET, which is strongly noise dominated for a single data segment, and the CMB component is therefore not important, as will be explicitly

demonstrated in this paper.

The third term is a sum over  $n_{\text{temp}}$  time-domain templates. These can be used to model several different types of nuisance components. Three examples are diffuse foregrounds and cosmic ray glitches for Planck, and ground pick-up for QUIET. In either case, we assume in this paper that the template itself is perfectly known, and the only free parameter is an overall unknown multiplicative amplitude  $a$ . This is a vector of length  $n_{\text{temp}}$ , and  $\mathbf{T}$  is the two-dimensional  $n \times n_{\text{temp}}$  matrix listing all templates column-wise.

Finally, the fourth term on the right-hand side of Equation 1 denotes a time-domain mask,  $\mathbf{m}$ . This is implemented by a “Gaussian” component having zero variance for samples that are not masked, and infinite variance for samples that are masked. In order to make analytic calculations more transparent, we write the corresponding covariance matrix as a diagonal matrix with elements

$$M_{ii} = \begin{cases} a, & i \text{ not masked} \\ \epsilon, & i \text{ masked} \end{cases}, \quad (3)$$

where  $a \rightarrow \infty$  and  $\epsilon \rightarrow 0$ . A typical application of this component is to remove periods of instrumental glitching, or to discard particularly bright observations when the telescope points towards bright astrophysical sources, such as point sources or the Galactic plane.

## 3. GIBBS SAMPLING AND THE POSTERIOR

Our primary goal is now to map out  $P(\theta|\mathbf{d})$ , the noise spectrum posterior distribution marginalized over all nuisance components. By Bayes’ theorem this distribution reads

$$P(\theta|\mathbf{d}) = \frac{P(\mathbf{d}|\theta)P(\theta)}{P(\mathbf{d})} \propto \mathcal{L}(\theta)P(\theta), \quad (4)$$

where  $\mathcal{L}(\theta) = P(\mathbf{d}|\theta)$  is the likelihood,  $P(\theta)$  is a prior on  $\theta$ , and  $P(\mathbf{d})$  is an irrelevant normalization constant.

In this paper we adopt for simplicity uniform priors on  $\sigma_0$ ,  $\alpha$  and  $f_{\text{knee}}$ . For typical relevant time series which contain  $\sim 10^5$  samples, these parameters are usually strongly data-driven, and the choice of priors is therefore irrelevant. However, if an informative prior (or the Jeffreys’ prior) is desired for a given application, it is straightforward to include this as indicated by Equation 4.

Since we assume that the noise is Gaussian distributed with covariance  $\mathbf{N}(\theta)$ , the likelihood is given by

$$\mathcal{L}(\theta) \propto \frac{e^{-\frac{1}{2}\mathbf{n}^T\mathbf{N}^{-1}(\theta)\mathbf{n}}}{\sqrt{|\mathbf{N}(\theta)|}}, \quad (5)$$

where  $\mathbf{n} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{F}\mathbf{a} - \mathbf{m}$  is the noise component of the data stream. The goal is to compute this distribution, marginalized over  $\mathbf{s}$  and  $\mathbf{a}$ , while at the same time taking into account possible gaps in the data.

The latter point touches on an important computational issue. If there are no gaps in the data, then  $\mathbf{N}$  is a Toeplitz matrix, and multiplication with  $\mathbf{N}$  is performed most efficiently in Fourier space. However, the same does not hold if there are gaps in  $\mathbf{d}$ , since the symmetry of  $\mathbf{N}$  is broken. The well-known solution to this problem is to fill the gap with a constrained noise realization with the appropriate spectrum (e.g., Hoffman & Ribak 1991). In

our formulation, this is equivalent to estimating  $\mathbf{n}$  jointly with  $\theta$ .

More generally, we want to estimate the joint density  $P(\mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \theta | \mathbf{d})$ , from which any desired marginal may be obtained. At first sight, this appears like a formidable computational problem, involving more than  $10^5$  free parameters. However, this is also a problem that may be tackled by means of the statistical technique called Gibbs sampling, which has already been described in detail for computing the CMB angular power spectrum with contaminated data by Jewell et al. (2004); Wandelt et al. (2004); Eriksen et al. (2004, 2008).

According to the theory of Gibbs sampling, samples from a joint distribution may be obtained by iteratively sampling from each corresponding conditional distribution. For our case, this leads to the following sampling scheme,

$$\mathbf{m}, \mathbf{n} \leftarrow P(\mathbf{m}, \mathbf{n} | \mathbf{s}, \mathbf{a}, \theta, \mathbf{d}) \quad (6)$$

$$\mathbf{s}, \mathbf{n} \leftarrow P(\mathbf{s}, \mathbf{n} | \mathbf{a}, \theta, \mathbf{m}, \mathbf{d}) \quad (7)$$

$$\mathbf{a}, \mathbf{n} \leftarrow P(\mathbf{a}, \mathbf{n} | \theta, \mathbf{m}, \mathbf{s}, \mathbf{d}) \quad (8)$$

$$\theta \leftarrow P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d}) \quad (9)$$

The symbol  $\leftarrow$  indicates sampling from the distribution on the right-hand side. With this algorithm,  $(\mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \theta)^i$  will be drawn from the correct joint distribution.

Note that each of the sampling steps that involve time-domain vectors are joint steps including the noise component itself. This approach is highly computationally advantageous as it allows for fast multiplication with  $\mathbf{N}$  in Fourier domain; conditional algorithms for sampling each component separately would require slow convolutions in time domain. Of course, it is fully acceptable within the Gibbs sampling machinery to sample some components more often than others.

Note also that if we are only interested in the joint maximum-posterior parameters, we can replace the relevant steps in the above algorithm by a maximization operation, such that we maximize the conditional instead of sampling from it. The algorithm then reduces to a typical iterative approach, but formulated in a convenient and unified statistical language. The advantage of this approach is computational speed, while the disadvantage is the loss of information about uncertainties. Both versions of the algorithm will be implemented and demonstrated in the following.

#### 4. SAMPLING ALGORITHMS

Equations 6–9 defines the high-level algorithm in terms of conditional sampling steps. To complete the method, we have to establish efficient sampling algorithms for each conditional distribution.

##### 4.1. Noise estimation with ideal data

Perhaps the most fundamental conditional distribution in the sampling scheme outlined above is  $P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d})$ . This describes the distribution of the noise parameters given perfect knowledge about all components of the data. To obtain an explicit expression for this distribution, we first note that  $P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d}) = P(\theta | \mathbf{n})$ ; if we know the true noise component,  $\mathbf{n}$ , no further information about either the CMB signal, the tem-




FIG. 1.— Correlation function for  $\alpha$  of the Metropolis sampler employed to sample from  $P(\theta | \mathbf{n})$  with a curvature matrix based proposal density. Similar plots for  $\sigma_0$  and  $f_{\text{knee}}$  look visually the same. The correlations fall below 10% at a lag of  $\sim 20$  samples, and we adopt a thinning factor of  $\sim 20$  samples to suppress correlations, given that the computational cost of this sampling step is lower than the constrained realization sampler.

plate amplitudes, or, indeed, the actual data is needed in order to estimate the noise parameters.

The expression for the conditional distribution  $P(\theta | \mathbf{n})$  is then formally the same as that for  $P(\theta | \mathbf{d})$  given by Equations 4 and 5. Writing this out in Fourier space for the  $1/f$  profile discussed in Section 2, one finds the following explicit distribution for  $\theta = \{\sigma_0, \alpha, f_{\text{knee}}\}$ ,

$$-\ln P(\sigma_0, \alpha, f_{\text{knee}} | \mathbf{n}) = -\ln P(\sigma_0, \alpha, f_{\text{knee}}) + \sum_{\nu > 0} \left[ \frac{p_\nu}{N_\nu} + \ln N_\nu \right]. \quad (10)$$

Here  $p_\nu$  are the power spectrum components of the data  $\mathbf{n}$ , while  $N_\nu = N_\nu(\sigma_0, \alpha, f_{\text{knee}})$  is the covariance matrix which in Fourier space is diagonal and given by Equation 2. The first term on the right-hand side is a user-defined prior.

To sample from this distribution, we use a standard Metropolis sampler with a Gaussian proposal density (e.g., Liu 2001). Each chain is initialized at the maximum-posterior point, which is found by a non-linear quasi-Newton search, and the covariance matrix of the Gaussian proposal density is taken to be the square root of the curvature matrix, evaluated at the maximum-posterior point. The elements of the inverse curvature matrix,  $\mathcal{C}^{-1} = -\partial^2 \log P(\theta | \mathbf{n}) / \partial \theta_i \partial \theta_j$ , read

$$\mathcal{C}_{\theta_i \theta_j}^{-1} = \sum_{\nu > 0} \left[ \left( \frac{1}{N_\nu} - \frac{p_\nu}{N_\nu^2} \right) \frac{\partial^2 N_\nu}{\partial \theta_i \partial \theta_j} + \left( \frac{2p_\nu}{N_\nu^3} - \frac{1}{N_\nu^2} \right) \frac{\partial N_\nu}{\partial \theta_i} \frac{\partial N_\nu}{\partial \theta_j} \right]. \quad (11)$$

We have found that this proposal density leads to a Markov chain correlation length of about 20 samples for typical parameter values, and we therefore thin our chains by this amount. In addition, we remove a few post-thinned samples at the beginning of the chain to remove potential burn-in, although we have never seen evidence of any such effects. Finally, if we only want to find the maximum-posterior point, and not run a full-blown




FIG. 2.— Constrained realizations through a gap in the data. The extent of the gap is indicated by the two vertical dashed lines; the solid black curve shows the input data, which is dominated by a point source (thin spikes) in the masked out region. Five different constrained realizations are shown by colored lines. These depend only on the data outside the mask, and are therefore not affected by the point source.

Gibbs chain, as discussed in Section 3, we terminate the process directly after the initial quasi-Newton search.

#### 4.2. Gap filling by constrained realizations

Next, we need to establish a sampling algorithm for  $P(\mathbf{n}, \mathbf{m} | \mathbf{s}, \mathbf{a}, \theta, \mathbf{d}) = P(\mathbf{n}, \mathbf{m} | \mathbf{r}, \theta)$ , where the residual  $\mathbf{r} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{T}\mathbf{a} = \mathbf{n} + \mathbf{m}$ . Note that  $\mathbf{r}$  and  $\mathbf{n}$  differ only by  $\mathbf{m}$ , which represents contaminated segments of data that are masked out. The problem is therefore reduced to sampling the components of a sum of Gaussians given the sum itself. This may be achieved efficiently by solving the equation

$$[\mathbf{N}^{-1} + \mathbf{M}^{-1}] \mathbf{n} = \mathbf{M}^{-1} \mathbf{r} + \mathbf{N}^{-\frac{1}{2}} \mathbf{v} + \mathbf{M}^{-\frac{1}{2}} \mathbf{w} \quad (12)$$

for  $\mathbf{n}$ , where  $\mathbf{v}$  and  $\mathbf{w}$  are vectors of standard  $N(0, 1)$  Gaussian random variates (e.g., Jewell et al. 2004; Wandelt et al. 2004; Eriksen et al. 2004).

The matrices involved here are typically of the order  $10^5 \times 10^5$ , so the equation cannot be solved by brute force. But since  $\mathbf{M}$  is diagonal in time domain and  $\mathbf{N}$  is diagonal in frequency domain (due to its Toeplitz nature), all multiplications are cheap in either time or Fourier domain, and the equation can be efficiently solved with the Conjugate Gradients method, properly changing basis as needed.

However, there is one practical complication involved in Equation 12: Since we intend to let  $\epsilon \rightarrow 0$  and  $a \rightarrow \infty$ , the matrix  $\mathbf{N}^{-1} + \mathbf{M}^{-1}$  becomes infinitely poorly conditioned. We can solve this problem, as well as significantly simplifying the equation, by splitting it into one equation for the masked region, and one for the unmasked region. Introducing the notation  $\mathbf{x}_1$  and  $\mathbf{x}_2$  for the unmasked and masked subsets of a vector  $\mathbf{x}$  respectively, we find

$$(\mathbf{N}^{-1} \mathbf{n})_1 + \epsilon^{-1} \mathbf{n}_1 = \epsilon^{-1} \mathbf{r}_1 + (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_1 + \epsilon^{-\frac{1}{2}} \mathbf{w}_1 \quad (13)$$

$$(\mathbf{N}^{-1} \mathbf{n})_2 + a^{-1} \mathbf{n}_2 = a^{-1} \mathbf{r}_2 + (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 + a^{-\frac{1}{2}} \mathbf{w}_2 \quad (14)$$

which in the limit  $\epsilon \rightarrow 0$ ,  $a \rightarrow \infty$  simplifies to

$$\mathbf{n}_1 = \mathbf{r}_1 \quad (15)$$

$$(\mathbf{N}^{-1} \mathbf{n})_2 = (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 \quad (16)$$

Note that equations 15–16 form an asymmetric equation system<sup>5</sup>. The Conjugate Gradients method is therefore not directly applicable, and the more general Biconjugate Gradients method must be used instead.

We found a simple diagonal preconditioner with value  $\text{Var}(\mathbf{n})$  inside the mask and value 1 outside it to be sufficient.

#### 4.3. Marginalization over fixed templates

The sampling algorithms described in Sections 4.1 and 4.2 defines the core noise Gibbs sampler, and together form a well-defined and complete noise estimation algorithm for low signal-to-noise data with gaps. However, one of the main advantages of the Gibbs sampling algorithm compared to other alternatives is its natural support for marginalization over nuisance parameters. In this section we describe the sampling algorithm for marginalization over fixed time-domain templates, describing for instance ground pickup, diffuse foregrounds or cosmic ray glitches.

First, we note that while the original data stream,  $\mathbf{d}$ , may contain gaps, and the Toeplitz nature of the noise covariance matrix is in that case broken, the constrained realization produced in Section 4.2 restores the Toeplitz symmetry. It is therefore computationally advantageous to use  $\mathbf{q} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{m} = \mathbf{n} + \mathbf{T}\mathbf{a}$  as the data in this step<sup>6</sup>, so that  $P(\mathbf{a}, \mathbf{n} | \theta, \mathbf{m}, \mathbf{s}, \mathbf{d}) = P(\mathbf{a}, \mathbf{n} | \mathbf{q}, \theta)$ .

Starting with Equation 5, solving for  $\mathbf{a}$  and completing the square in the exponential, the appropriate conditional distribution for  $\mathbf{a}$  is found to be the well-known distribution

$$P(\mathbf{a} | \mathbf{q}, \theta) \propto e^{-\frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})(\mathbf{a} - \hat{\mathbf{a}})}, \quad (17)$$

where  $\hat{\mathbf{a}} = (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})^{-1} \mathbf{T}^t \mathbf{N}^{-1} \mathbf{q}$ ; that is,  $P(\mathbf{a} | \mathbf{q}, \theta)$  is a Gaussian distribution with mean  $\hat{\mathbf{a}}$  and covariance  $\mathbf{C}_a = (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})^{-1}$ . The same result has been derived for numerous other applications, one of which was described by Eriksen et al. (2004), outlining template amplitude sampling with CMB sky map data.

Sampling from this distribution is once again straightforward. If one only wish to marginalize over a small number of templates, the easiest solution is simply to compute both  $\hat{\mathbf{a}}$  and  $\mathbf{C}_a$  by brute-force, and let  $\mathbf{a}^{i+1} = \hat{\mathbf{a}} + \mathbf{L}_a \eta$ , where  $\mathbf{L}_a$  is the Cholesky factor of  $\mathbf{C}_a = \mathbf{L}_a \mathbf{L}_a^t$ , and  $\eta$  is a vector of uncorrelated standard Gaussian  $N(0, 1)$  variates. On the other hand, if there are more than, say, 1000 templates involved, it may be faster to

<sup>5</sup> Equations 15–16 can be rewritten as

$$\begin{bmatrix} 1 & 0 \\ N_{21}^{-1} & N_{22}^{-1} \end{bmatrix} = \begin{bmatrix} \mathbf{r}_1 \\ (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 \end{bmatrix}$$

This is asymmetric because we multiplied the unmasked part of the equation by  $\epsilon$  in order to get a finite result.

<sup>6</sup> Of course, one could write down a sampling algorithm for  $\mathbf{a}$  that only uses the non-masked data directly, but this would require heavy time-domain convolutions, and not take advantage of the symmetries inherent in the noise covariance matrix.




FIG. 3.— Recovered noise parameters from 10 000 noise-only simulations with gaps. The dashed horizontal lines indicate the true input value.

solve the following equation by Conjugate Gradients,

$$(\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T}) \mathbf{a} = \mathbf{T}^t \mathbf{N}^{-1} \mathbf{q} + \mathbf{T}^t \mathbf{N}^{-\frac{1}{2}} \omega, \quad (18)$$

where  $\omega$  is a full time-stream of  $N(0, 1)$  random variates. In this paper, which has the QUIET experiment as its main application, we only use the former algorithm; for Planck the latter algorithm may be useful in order to account for frequent and partially overlapping cosmic ray glitches efficiently.

#### 4.4. CMB signal marginalization

Most CMB experiments are strongly noise-dominated within relatively short time periods, and need to integrate over the sky for a long time in order to produce high-sensitivity maps. For instance, for the first-season QUIET experiment the mean polarized sensitivity of a detector was  $280 \mu\text{K}\sqrt{\text{s}}$  (QUIET 2011), while the polarized CMB sky has an RMS of  $\sim 1 \mu\text{K}$  on the relevant angular scales. In such cases, it is an excellent approximation simply to ignore the CMB contribution when estimating the noise parameters. However, this approximation does not hold for all experiments, and one particularly important counterexample is Planck.

In order to marginalize over the CMB signal we need to be able to sample from  $P(\mathbf{s}, \mathbf{n} | \mathbf{a}, \theta, \mathbf{m}, \mathbf{d}) = P(\mathbf{s}, \mathbf{n} | \mathbf{u}, \theta)$ , where the residual is  $\mathbf{u} = \mathbf{d} - \mathbf{m} - \mathbf{F}\mathbf{a} = \mathbf{P}\mathbf{s} + \mathbf{n}$ . This is again a case that involves sampling terms of a Gaussian sum given the sum itself. The only difference from Section 4.2 is that a projection operator is involved in this case,

$$(\mathbf{S}^{-1} + \mathbf{P}^T \mathbf{N}^{-1} \mathbf{P}) \mathbf{s} = \mathbf{P}^T \mathbf{N}^{-1} \mathbf{u} + \mathbf{S}^{-\frac{1}{2}} \mathbf{v} + \mathbf{P}^T \mathbf{N}^{-\frac{1}{2}} \mathbf{w} \quad (19)$$

Here  $\mathbf{v}$  and  $\mathbf{w}$  are vectors of standard normal samples in pixel and time domain respectively. Note that  $\mathbf{M}$  is not involved this time, so the expression can be used as it is. As before, multiplications involving  $\mathbf{N}$  are efficient in Fourier space due to the Toeplitz nature of  $\mathbf{N}$ , but the CMB signal covariance,  $\mathbf{S}$ , will in general be dense, depending on the true CMB power spectrum. This makes this method prohibitively expensive for general asymmetric scanning strategies. There are, however, circumstances for which also this multiplication becomes efficient. One obvious case is that of full sky coverage,

where a change to spherical harmonic basis makes  $\mathbf{S}$  diagonal.

More interestingly,  $\mathbf{S}$  also becomes diagonal when expressed in Fourier basis on a circle on the sky. To see this, consider a circle with radius  $\theta$  on the sphere, parametrized by the angle  $\phi$ . The covariance between the points  $p_1 = (\theta, \phi_1)$  and  $p_2 = (\theta, \phi_2)$  with angular distance  $r$  is given by the two-point correlation function:

$$S(p_1, p_2) = C(r) = \frac{1}{4\pi} \sum_{l=0}^{\infty} (2l+1) C_l P_l(\cos(r)) \quad (20)$$

Since  $r(p_1, p_2)$  for the case of a circle only depends on  $\Delta\phi$ , and is independent of  $\phi$  itself,  $S(\phi_1, \phi_2)$  is a Toeplitz matrix, and is therefore diagonal in Fourier space. This is highly relevant for Planck, since the Planck scanning strategy (Planck 2011a) naturally divides into scans of circles. Therefore, in this case sampling  $\mathbf{s} \leftarrow P(\mathbf{s} | \mathbf{u})$  can be done at very low extra cost.

#### 5. APPLICATION TO SIMULATED DATA

In this section we demonstrate the noise Gibbs sampler as described in Section 4 on a particular type of QUIET simulations. QUIET is a radiometer-based CMB B-mode polarization experiment located in the Atacama desert (QUIET 2011), which took observations from August 2008 to December 2010. The first results were based on only nine months worth of data, and yet already provided the second most stringent upper limit on the tensor-to-scalar ratio,  $r$ , to date based on CMB polarization measurements.

In its normal mode of operation, QUIET observed four separate CMB fields on the sky which were each chosen because of their low foreground levels. In this mode, the experiment is totally noise dominated on time scales of an hour or less, with a mean polarized sensitivity per detector diode of  $280 \mu\text{K}\sqrt{\text{s}}$  (QUIET 2011).

However, QUIET also observed two Galactic patches, one of which was the Galactic center, as well as several bright calibration objects, such as the Moon, Tau Alpha and RCW38. These sources are bright enough to be seen visually in each detector time stream, and they can therefore bias any noise estimates unless properly accounted for. Such objects also complicate automated data selection processes, since it is difficult to distinguish between

TABLE 1  
VALIDATION BY SIMULATIONS

Simulation	$\sigma_0$ ( $10^{-5}\text{V}$ )	$\alpha$	$f_{\text{knee}}$ ( $10^{-1}\text{Hz}$ )
Posterior maximization (absolute parameter values)			
Gaps only	$1.000 \pm 0.003$	$-1.80 \pm 0.06$	$1.00 \pm 0.05$
Gaps + uncorrected CMB	$1.000 \pm 0.003$	$-1.80 \pm 0.06$	$1.00 \pm 0.05$
Gaps + uncorrected ground pickup	$1.006 \pm 0.003$	$-1.52 \pm 0.05$	$1.51 \pm 0.08$
Gaps + corrected ground pickup	$1.000 \pm 0.003$	$-1.80 \pm 0.06$	$1.00 \pm 0.05$
Gibbs sampling (normalized parameter values)			
Gaps only	$-0.02 \pm 1.00$	$0.01 \pm 1.01$	$0.03 \pm 1.01$
Gaps + uncorrected CMB	$-0.02 \pm 1.00$	$0.00 \pm 1.01$	$0.02 \pm 1.01$
Gaps + uncorrected ground pickup	$1.78 \pm 1.03$	$6.04 \pm 1.36$	$6.63 \pm 0.77$
Gaps + corrected ground pickup	$-0.02 \pm 1.00$	$0.00 \pm 1.00$	$0.03 \pm 1.00$

NOTE. — Summary of recovered noise parameters from various simulated ensembles. Each column indicates the mean and standard deviation of the resulting parameter distribution. The top section shows results obtained when simply maximizing the posterior, while the bottom section shows the results for a full Gibbs sampling analysis. Each run in the top section consists of 10000 simulations, while each run in the bottom section consists of 5000 simulations. All runs have been started with the same seed, to ensure directly comparable results.

an astrophysical object and an instrumental glitch.

In this section we show how the algorithm developed in Section 4 may be applied to such situations. Specifically, we consider a observing session lasting for about 40 minutes of a field including a bright source with known location, and assume that the data may be modeled as  $\mathbf{d} = \mathbf{n} + \mathbf{T}\mathbf{a} + \mathbf{m}$ . Here  $\mathbf{T}$  is a single template describing possible sidelobe pick-up from the ground, constructed from the full observing season as described by QUIET (2011), and  $\mathbf{m}$  is a time-domain mask that removes any samples that happen to fall closer than  $1^\circ$  from the source center. The total number of samples in the data stream is 60949, and the total number of masked samples is 651.

The simulations used in this section are constructed as follows. We set up an ensemble of  $10^4$  time streams containing correlated Gaussian random noise with  $\sigma_0 = 10^{-5}\text{V}$ ,  $\alpha = -1.8$  and  $f_k = 0.1\text{Hz}$ ; the white noise and spectral index are representative for a QUIET detector, while the knee frequency is grossly exaggerated to push the algorithm into a difficult region of parameter space, as well as to more clearly visualize the outputs of the algorithm. A far more reasonable value for QUIET is  $f_k = 10\text{mHz}$ , and we have of course verified that the algorithm also works for such cases. Further, it reaches convergence faster in that case than with the extreme value of  $f_k$  used in the present simulations.

### 5.1. Visual inspection of constrained realizations

Before considering the statistical properties of the resulting posterior distributions, it is useful to look visually at a few constrained realizations in order to build up intuition about the algorithm. In order to highlight the behavior of the constrained realizations, we make two adjustments to the above simulation procedure for this case alone: First, we replace the tuned mask with a wide 6000-sample mask, covering the entire time range in which the source is visible, and second, we make the correlated noise component stronger by setting  $\sigma_0 = 10^{-6}\text{V}$ ,  $f_{\text{knee}} = 1\text{Hz}$  and  $\alpha = -2.3$ .

The results are shown in Figure 2. The raw data are shown in the solid black line, and the vertical dashed

lines indicate the extent of the gap. The colored curves within the gap shows 5 difference constrained realizations; note that together with the black solid curve outside the mask, any of these form a valid noise realization with the appropriate noise power spectrum as defined by  $\sigma_0$ ,  $\alpha$  and  $f_{\text{knee}}$ . They are each a valid sample drawn from  $P(\mathbf{n}, \mathbf{m}|\mathbf{d})$ . However, if one had tried to estimate the noise spectrum also using the data inside the gap, the source signal (seen as sharp spikes in Figure 2) would bias the resulting noise parameters.

In this paper, we consider the constrained realizations primarily to be a useful tool that allows for fast noise covariance matrix multiplications in Fourier space. However, these constrained realizations can of course also be useful in their own right, for instance for deglitching a time stream before map making.

### 5.2. Validation and statistical characterization

We now seek to statistically validate our algorithms and codes. Both the posterior maximization and the Gibbs sampling algorithms are considered. The number of simulations are 10000 for posterior maximization and 5000 for Gibbs sampling, with properties as described above. In each case, we consider four different models. First, we analyze simulations including only noise and gaps. Second, we add a CMB signal to each realization, but do not attempt to correct for it. Third, we add a strong ground template to each realization, and do also not attempt to correct for it. Fourth, we analyze the same ground-contaminated simulations as above, but this time do marginalize over an appropriate template. The same random seeds were used in each of the four simulation and analysis results, in order to allow for direct comparison of results between runs.

The results from this exercise are summarized in Table 1, and histograms for the first case are shown in Figure 3. First, the the upper section in Table 1 lists the mean and standard deviation of the recovered parameters for the posterior maximization algorithm. Recall that the input parameters were  $(\sigma_0, \alpha, f_{\text{knee}}) = (10^{-5}\text{V}, -1.8, 0.1\text{Hz})$ , and these are recovered perfectly




FIG. 4.— *Left:* CPU time as a function of total number of samples in the time stream, keeping the number of masked samples constant. *Middle:* CPU time as a function of number of masked samples. The size of each gap is kept constant at 100 samples, and only the number of gaps vary. *Right:* CPU time as a function of number of masked samples, but in this case there is only one gap of varying width.

in all cases, except for the one involving an uncorrected ground template, as expected.

Second, in the bottom half we show the results from the Gibbs sampling analyses, but this time in terms of normalized parameters on the form  $r = (\theta_{\text{est}} - \theta_{\text{in}})/\sigma_{\text{est}}$ , where  $\theta_{\text{est}}$  and  $\sigma_{\text{est}}$  are the mean and standard deviation of the Gibbs chain for a given parameter (removing the first 10% of the samples for burn-in), and  $\sigma_{\text{in}}$  is the true input value. If the Gibbs chain is both unbiased and has the correct dispersion,  $r$  should be Gaussian distributed with zero mean and unit variance. As seen in Table 1, this is indeed the case.

We also note that adding a CMB component to these simulations do not bias the noise estimates, simply because the CMB is too weak to be detected on the time scales considered here. This confirms the assumption made by the QUIET team when estimating their noise properties: The QUIET observations are sufficiently noise dominated on a one-hour time scale that the CMB can be safely neglected for noise estimation purposes.

### 5.3. Resource requirements

To be practical, it is not sufficient that a method is robust and accurate, but it must also be computationally efficient. For the present algorithm the two most important parameters for computational speed are 1) the total number of samples in the time stream,  $n$ , and 2) the number of masked samples,  $m$ , while also the relative position of the masked samples play an important role.

In Figure 4 we show the scaling of each of the two algorithms (posterior maximization and Gibbs sampling) as a function of both  $n$  (left panel) and  $m$  (right panel). In the left plot,  $m$  was fixed at 1000, divided into ten gaps of 100 samples each, and only the total length of the

data stream was varied. In this case, we should expect the scaling of the overall algorithm to be dominated by Fourier transforms, suggesting an overall behavior given by  $\mathcal{O}(n \log n)$ . As seen in Figure 4, this approximation holds to a very high degree, both for posterior maximization and Gibbs sampling. Further, we see that the CPU time required to analyze a single 100 000 sample data set with 1000 samples removed is 3 seconds for posterior maximization and 21 seconds for Gibbs sampling.

In the middle panel, we fix  $n$  at 100 000, and increase  $m$  by varying the number of gaps, each extending 100 samples. Perhaps somewhat surprisingly, we see that the computing time in this case is nearly independent of  $m$ . The reason for this is simply that the number of conjugate gradient iterations required for the gap filling procedure is largely determined by condition number (i.e., the ratio between the highest and smallest eigenvalue) of the covariance matrix of a single gap. Having more gaps separated by more than one time-domain correlation length effectively corresponds to performing multiple matrix inversions in parallel, and the net cost therefore do not increase significantly.

In the third panel, we increase  $m$  by making one gap larger, as opposed to adding many small gaps. In this case the CPU time does increase dramatically, because it becomes increasingly hard for the algorithm to fill the missing pieces of the data stream. In this case, the noise covariance matrix condition becomes larger and larger.

## 6. SUMMARY

We have described and implemented a Bayesian framework for estimating the time-domain noise power spectrum for non-ideal CMB experiments. This framework is conceptually identical to a previously described method

for estimating the angular CMB power spectrum from CMB sky maps (Jewell et al. 2004; Wandelt et al. 2004; Eriksen et al. 2004), and relies heavily on the Gibbs sampling algorithm. The single most important advantage of this method over existing competitors in the literature derives from the conditional nature of the Gibbs sampler: Additional parameters may be introduced *conditionally* into the algorithm. This allows for seamless marginalization over nuisance parameters, which otherwise may be difficult to integrate. A second important advantage of the method is the fact that it provides proper uncertainties on all estimated quantities, which at least in principle later may be propagated into final estimates of the uncertainties of the CMB sky map and angular power spectra.

In this paper we implemented support for two general features that are useful for analysis of realistic data, namely constrained realizations and template sampling. The former is useful whenever there are gaps in the data, for instance due to an instrumental glitch, or there are strong localized sources in the sky that may bias the noise estimate: In these cases, the gaps are refilled with a constrained noise realization with the appropriate noise parameters, such that the full time stream represents a proper sample from a Gaussian distribution with a noise covariance matrix,  $\mathbf{N}$ . Since the time stream no longer contains gaps, the Toeplitz symmetry of the noise covari-

ance matrix is restored, and matrix multiplications may be performed quickly in Fourier space.

The second operation, template sampling, is also a powerful and versatile technique for mitigating systematic errors. In this paper we mostly focused on data from the QUIET experiment, for which ground pickup from sidelobes is one significant source of systematics (QUIET 2011). In a future publication we will apply the same method to simulations of the Planck experiment, for which cosmic ray glitches is an important source of systematic errors. As detailed by Planck (2011b), these cosmic rays may be modeled in terms of a limited set of time domain templates, and the algorithms presented in this paper should therefore prove useful for mitigating the effects of these glitches, as well as for propagating the corresponding uncertainties into the final noise spectrum parameters.

We thank the QUIET collaboration for stimulating discussions. The computations presented in this paper were carried out on Titan, a cluster owned and maintained by the University of Oslo and NOTUR. This project was supported by the ERC Starting Grant StG2010-257080. Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005) software and analysis package.

#### REFERENCES

- Eriksen, H. K., et al. 2004, ApJS, 155, 227  
 Eriksen, H. K., Jewell, J. B., Dickinson, C., Banday, A. J., Górski, K. M., & Lawrence, C. R. 2008a, ApJ, 676, 10  
 Ferreira, P. G., & Jaffe, A. H. 2000, MNRAS, 312, 89  
 Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., Bartelman, M. 2005, ApJ, 622, 759  
 Hinshaw, G., et al. 2003, ApJS, 148, 63  
 Hoffman, Y., & Ribak, E. 1991, ApJ, 380, L5  
 Komatsu, E., et al. 2011, ApJS, 192, 18  
 Jewell, J., Levin, S., & Anderson, C. H. 2004, ApJ, 609, 1  
 Liu, J. S., Monte Carlo Strategies in Scientific Computing, Cambridge, USA: Springer, 2001,  
 Planck Collaboration 2011, A&A, submitted, [1101.2022]  
 Planck HFI Core Team 2011, A&A, submitted, [1101.2048]  
 Prunet, S., et al. 2001, in MPA/ESO/MPA conf. Proc. "Mining the Sky", [astro-ph/0101073]  
 QUIET Collaboration 2011, ApJ, in press, [1012.3191]  
 Tauber, J. A., et al. 2010, A&A, 520, A1  
 Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, Phys. Rev. D, 70, 083511





## Paper IV

---

*Application of the Kolmogorov-Smirnov test to CMB data: Is the universe really weakly random?*

S. K. Næss

A&A, Volume 538, pp. A17 (2012)

---

# Application of the Kolmogorov-Smirnov test to CMB data: Is the universe really weakly random?

Sigurd K. Næss

Institute of theoretical Astrophysics, University of Oslo, P.O.Box 1029 Blindern, 0315 Oslo, Norway

Preprint online version: May 26, 2011

## Abstract

A recent application of the Kolmogorov-Smirnov test to the WMAP 7 year W-band maps claims evidence that the CMB is “weakly random”, and that only 20% of the signal can be explained as a random Gaussian field. I here repeat this analysis, and in contrast to the original result find no evidence for deviation from the standard  $\Lambda$ CDM model. Instead, the results of the original analysis are consistent with not properly taking into account the correlations of the  $\Lambda$ CDM power spectrum.

**Key words.** cosmic background radiation

## 1. Introduction

In astronomical data analysis, it is often useful to be able to test whether a set of data points follows a given distribution or not. For example, many analysis techniques depend on instrument noise being Gaussian, and to avoid bias, one must check that this actually is the case. There are many different ways in which two distributions can differ, and correspondingly many different ways to test them for equality. The simplest ones, such as comparing the means or variances of the distributions, suffer from the problem that there are many ways in which distributions can differ that they cannot detect no matter how many samples are available. For example, samples from a uniform distribution can easily pass as Gaussian if one only considers the mean and variance.

The popular Kolmogorov-Smirnov test (K-S test) resolves this problem by considering the cumulative distribution functions (CDF) instead: Construct the empirical CDF of the data points and find its maximum absolute difference  $K$  from the theoretical CDF. Due to the limited number of samples, the empirical CDF will be noisy, and  $K$  will therefore be a random variable with its own CDF, which in the limit where the number of samples goes to infinity is given by

$$P(x < K) = F_{KS}(\sqrt{N_{\text{obs}}}K) \quad (1)$$

with

$$F_{KS}(x) = 1 - 2 \sum_{i=1}^{\infty} (-1)^{i-1} e^{-2i^2 x^2}. \quad (2)$$

In contrast with the simplest tests, this test can detect any deviation in the distributions, but may require a large number of samples to do so, especially in the tails of the distribution.

Recently, a series of papers (Gurzadyan et al. 2011; Gurzadyan & Kocharyan 2008; Gurzadyan et al. 2010)

has applied this test to WMAP’s cosmic microwave background (CMB) maps, resulting in the remarkable claim that the CMB is “weakly random”, with only 20% of the CMB signal behaving as one would expect from a random Gaussian field. This result went on to be used in a much discussed series of papers (Gurzadyan & Penrose 2010a,b, 2011) claiming a strong detection of concentric low-variance circles in the CMB, which was taken as evidence for Conformal Cyclic Cosmology. Other groups failed to significantly detect the circles (Wehus & Eriksen 2010; Moss et al. 2011; Hajian 2010). The difference in significance was due different CMB models: Wehus & Eriksen (2010); Moss et al. (2011); Hajian (2010) used realizations of the best-fit  $\Lambda$ CDM power spectrum, while Gurzadyan & Penrose (2010b) used a “weakly random” CMB model.


Both in order to resolve this issue, and because a weakly random universe would be a strong blow against the  $\Lambda$ CDM model in its own right, it is important to test this result.

## 2. Method

Before applying the K-S test, one must be aware of its limited area of validity: Equation (1) requires an infinite number of independently identically distributed samples, while CMB maps actually consist of a limited number of correlated samples. However, both the correlations and number of samples can be compensated for, as we shall see.

### 2.1. Application of the K-S test to correlated data

Though the K-S test is not immediately applicable to a correlated data set, it is possible to perform an equivalent test on a transformed set of samples. The question we are trying to answer with the K-S test is “Do the samples follow the theoretical distribution?”. The truth or falseness of this is preserved if we apply the same transformation to both the samples and the distribution we test them against, and to be able to use the K-S test, the logical transformation to use is a *whitening* transformation, which results in an



**Figure 1.**  $\Lambda$ CDM two point correlation function after applying the WMAP W-band beam and the HEALPix (Górski et al. 2005) nsidc 512 pixel window.

independent, identical distribution for the samples. With original samples  $\mathbf{d}$  with covariance matrix  $\mathbf{C}$ , the whitened (uncorrelated with unit variance) samples  $\mathbf{r}$  are given by:

$$\mathbf{r} = \mathbf{C}^{-\frac{1}{2}} \mathbf{d} \quad (3)$$

Thus, to test whether the data points  $\mathbf{d} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{C})$ , we can test the equivalent hypothesis  $\mathbf{r} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{1})$ .

In the case of CMB maps, both the data itself and the noise is expected to be Gaussian, so the obvious theoretical distribution here is  $N(0, \mathbf{S} + \mathbf{N})$ , where the CMB signal covariance matrix  $\mathbf{S}$  is given by the two-point correlation function:


$$S_{ij} = \sum_{l=0}^{\infty} \sqrt{\frac{2l+1}{4\pi}} C_l B_l P_l(\cos(|\mathbf{p}_i - \mathbf{p}_j|)) \quad (4)$$

$P_l(x)$  are the Legendre polynomials normalized to  $\frac{1}{2\pi}$ , and  $\mathbf{p}_i$  and  $\mathbf{p}_j$  are the direction vectors for pixel  $i$  and  $j$  in the disk.  $C_l$  is the  $\Lambda$ CDM angular power spectrum, while  $B_l$  accounts for the beam and pixel window.  $\mathbf{N}$  is instrument dependent, but for the WMAP W-band CMB map we will use here, the noise is nearly diagonal, and given by the corresponding W-band RMS map.

## 2.2. Application of the K-S test with few samples

The other problem we need to account for is our finite number of samples. In this case equation (1) is only approximate. For most uses of the test, this approximation is good enough, especially when employing analytical expressions for improving the quality of the approximation for low numbers of samples (von Mises 1964). For example, when performing a single test to accept or reject a test distribution, a bias of a few percent in the confidence with which the hypothesis is rejected is not important.

However, when making statistics for a large number of such test results, such a bias may make the results ambiguous. Given a set of experiments with a corresponding set of maximum deviations  $\{K_i\}$ , the corresponding probabilities  $\{p_i = P(x < K_i)\}$  should be uniformly distributed if the



**Figure 2.** When applying the K-S test to samples known to come from the correct distribution, the resulting values  $\{p_i = P(x < K_i)\}$  should be uniformly distributed, but when working with a limited number of samples, the Kolmogorov distribution is only approximate, and the actual CDF of the results,  $G(p)$ , differs from the ideal  $G_{\infty}(p) = p$ . This is shown in the upper panel for the case of 540 samples per experiment, where  $G(p)$  is the solid line and  $G_{\infty}(p)$  is dashed. The lower panel shows the deviation between the two, which is of the order of 1% in this case (but larger with fewer samples).

samples actually follow the theoretical distribution, and a histogram of  $\{p_i\}$  should therefore be flat. Deviations from this indicate that the theoretical distribution does not accurately describe the samples. However, the approximate equation (1) also introduces a small non-uniformity in  $\{p_i\}$  even if the samples actually do follow the distribution. To avoid the ambiguity this causes, we will instead compute a numerical correction function mapping the approximate  $p$  to the true  $p'$ .<sup>1</sup>


To build up the correction, we simulate a large number<sup>2</sup> of experiments, each with the same number of samples as the actual data set, but drawn directly from the theoretical distribution. Thus, for these,  $\{p_i\}$  should be uniform, with a CDF of  $G_{\infty}(p) = p$ . However, since equation (1) is inexact, for small numbers of samples, the actual CDF is  $G(p) \neq G_{\infty}(p)$ . The mapping between the approximate  $p$  and true  $p'$  is given by  $G(p) = G_{\infty}(p') \Rightarrow p' = G_{\infty}^{-1}(G(p)) = G(p)$ . Thus, for a limited number of samples

$$P(x < K) = G(F_{\text{KS}}(\sqrt{N_{\text{obs}}}K)). \quad (5)$$

Figure 2 illustrates the correction function for  $5 \cdot 10^6$  simulations of 540 each. For this many samples, the correction is only of the order of 1%.

<sup>1</sup> What we do here is essentially replacing the analytical Kolmogorov distribution (equation (1)) with a numerical distribution. This could also be done without using the analytical distribution as a basis, at a small cost in clarity.

<sup>2</sup> The number necessary depends on the level of accuracy desired. The noise in the estimate of  $G(p)$  propagates to the final results. To make this a subdominant noise contribution, the number of simulations should be at least as large as the number of actual experiments, preferably much higher.



**Figure 3.** A randomly selected disk before (left) and after (right) the whitening operation. The samples are strongly correlated and thus unsuitable for the K-S test before the transformation, but afterwards no correlations are visible and the variance is 1. Note that whitening the data does not mean that we are “forcing” the K-S test to pass. The whitened data will only end up matching  $N(0, 1)$  after whitening if they followed our theoretical distribution  $N(0, C)$  before.

### 3. Does $\Lambda$ CDM fail the K-S test?


With this in hand we can finally apply the K-S test on CMB data. Following Gurzadyan et al. (2011), we randomly pick 10000 disks with a radius of 1.5 degrees from the WMAP 7 year W-band map (Jarosik et al. 2011), with the region within 30 degrees from the galactic equator excluded. Each disk contains on average 540 pixels, which are whitened using equation (3). A typical disk before and after the whitening operation can be seen in Fig. 3. After whitening, the values should follow the distribution  $N(0, 1)$  if our model is correct.

The histogram of resulting probabilities  $\{p_i = P(x < K_i)\}$  from of applying equation (5) to the hypothesis  $\mathbf{r} \leftarrow N(\mathbf{0}, \mathbf{1})$  is shown in Fig. 4, together with the 68% and 95% intervals from 300 simulations. The data and simulations are consistent, and follow a uniform distribution as expected<sup>3</sup>: The CMB map is fully consistent with  $\Lambda$ CDM + WMAP noise as far as the K-S test is concerned.

This is dramatically different from the curve found by Gurzadyan et al. (2011), which was strongly biased towards *low* values. Low values of  $P(x < K)$  would mean that the empirical CDF of the samples matches the theoretical one *too well*, i.e. even better than samples drawn directly from the theoretical distribution.

What could cause Gurzadyan et al. to get results so different from ours? One way biasing  $P(x < K)$  low is by basing the parameters of your test distribution on the values themselves. However, even without doing this, it is possible to get low values if the values used in the K-S test are *correlated*. This is also consistent with the presentation given by Gurzadyan et al. (2011) who apparently applied the K-S test directly to the raw samples  $\mathbf{d}$ , or equivalently, that they model the pixel values as coming from a 1-dimensional dis-

<sup>3</sup> It should be noted that the histogram bins are not completely independent for two reasons: Firstly, some disks are going to overlap, meaning that the same samples enter into several different K-S tests, and secondly, while our transformation has made the samples within each disk independent, the correlation between different disks is still present.



**Figure 4.** Histogram of results of the K-S test. Each panel compares the results from properly taking the correlations into account (solid line) with those one gets from ignoring them (dashed line), together with 68% and 95% intervals (dotted lines) from simulations. The upper panel corresponds to using samples further than 30 degrees away from the galactic equator, while the lower panel instead uses the WMAP KQ85 analysis mask. In both cases, both the map and the simulations pass the K-S test when taking the correlations into account, while if the are ignored, the K-S test fails in the same way Gurzadyan et al. (2011) reported.

tribution. To check this, I repeated the analysis, this time using the theoretical distribution  $d \leftarrow N(\mu, \sigma^2)$ , where  $\mu$  and  $\sigma^2$  are the measured mean and variance of the samples in the disk. The result is also shown in Fig. 4. This time, the bias towards low values is clearly recreated.

It therefore seems likely that Gurzadyan et al.’s reported “weak randomness” is the result of not properly taking the CMB’s correlations into account. One is, of course, free to use whatever distribution one wants as the theoretical distribution in a K-S test, even a model where the CMB pixels are independently identically distributed, with no correlations at all. The problem lies in the interpretation of the test results. For Gurzadyan et al. (2011), the K-S test results are clearly not uniform, indicating that the chosen theoretical distribution has been disproven. However, Gurzadyan et al. then go on to create a set of simulations (linear combinations of 20% Gaussian and 80% static signal) that fail the test in the same way as the WMAP map does. But having two sets of samples fail the K-S test the

same way does not prove that they have the same properties. It simply means that the chosen test distribution was a poor choice.

#### 4. Kolmogorov maps

While Gurzadyan et al.'s Kolmogorov statistics are biased by not taking the correlations into account, the approach of making sky maps of K-S test results introduced in Gurzadyan et al. (2009) is still an interesting way to search for regions of the sky that do not follow the expected distribution. Making an unbiased Kolmogorov map straightforwardly follows the procedure in Sect. 3, with the main difference being the selection of pixels. Instead of randomly selecting disks, we now systematically go through nside 16 pixels, using the 1024 nside 512 subpixels inside each one as the samples. These are then tested against  $N(0, \mathbf{C})$  by whitening them via equation (3) and then comparing the whitened samples to  $N(0, 1)$ .

The result is the nside 16 map of  $P(x < K)$  shown in Fig. 5. Regions that pass the test have a value uniformly distributed between 0 and 1, and we see that this applies to the CMB-dominated areas of the sky, while areas dominated by the galaxy fail the test as expected.

For comparison, Fig. 5 also includes the result of making the same map while ignoring correlations. In this case, the whole sky fails the test: The CMB-dominated areas are biased low, while the galaxy is biased high. This map is similar to the map in Gurzadyan et al. (2009), which is also too low outside the galaxy, and too high inside, which is, again, consistent with Gurzadyan et al. applying the K-S test directly to the raw samples.


#### 5. Summary

The Kolmogorov-Smirnov test is a useful and general way of testing whether a data set follows a given distribution or not. However, it only applies to independently identically distributed samples. The CMB is strongly correlated, and thus not immediately compatible with the test. However, this can be resolved by the application of a whitening transformation, replacing the hypothesis  $\mathbf{d} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{C})$  with the equivalent  $\mathbf{C}^{-\frac{1}{2}}\mathbf{d} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{1})$ . With this, we find that the  $\Lambda$ CDM passes the K-S test. This is incompatible with the original analysis by Gurzadyan et al. (2011), which claimed detection of an unknown non-random component making up 80% of the CMB based on the CMB failing the K-S test there. It turns out that this analysis did not take the CMB correlations into account, which we confirm by producing the same failure of the K-S test when we skip the whitening step. When the correlations are handled properly, there is no need for a weakly random universe.

*Acknowledgements.* The author would like to thank Hans Kristian Eriksen for useful discussion and comments.

#### References

- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759, arXiv:astro-ph/0409513  
 Gurzadyan, V. G., Allahverdyan, A. E., Ghahramanyan, T., et al. 2009, *A&A*, 497, 343, 0811.2732  
 Gurzadyan, V. G., Allahverdyan, A. E., Ghahramanyan, T., et al. 2011, *A&A*, 525, L7+, 1011.3711



**Figure 5.** nside 16 map of  $P(x < K)$  based on pixels from the WMAP 7 year nside 512 W-band map. This time, the K-S test is performed on the 1024 nside 512 pixels inside each nside 16 pixel instead of on disks. The upper panel uses  $N(0, \mathbf{S} + \mathbf{N})$  as the theoretical distribution (by testing the whitened data against  $N(0, 1)$ ), while the lower panel ignores the correlations, instead using  $N(\mu, \sigma^2)$ , where  $\mu$  and  $\sigma$  are the measured mean and standard deviation of the samples. The former passes the test outside the galactic plane, while the latter fails everywhere, being biased low outside the galaxy.

- Gurzadyan, V. G., Kashin, A. L., Khachatryan, H. G., et al. 2010, *Europhysics Letters*, 91, 19001, 1002.2155  
 Gurzadyan, V. G. & Kocharyan, A. A. 2008, *A&A*, 492, L33, 0810.3289  
 Gurzadyan, V. G. & Penrose, R. 2010a, ArXiv e-prints, 1011.3706  
 Gurzadyan, V. G. & Penrose, R. 2010b, ArXiv e-prints, 1012.1486  
 Gurzadyan, V. G. & Penrose, R. 2011, ArXiv e-prints, 1104.5675  
 Hajian, A. 2010, ArXiv e-prints, 1012.1656  
 Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, *ApJS*, 192, 14, 1001.4744  
 Moss, A., Scott, D., & Zibin, J. P. 2011, *JCAP*, 4, 33, 1012.1305  
 von Mises, R. 1964, *Mathematical Theory of Probability and Statistics*, ed. von Mises, R.  
 Wehus, I. K. & Eriksen, H. K. 2010, ArXiv e-prints, 1012.1268