

THE UNIVERSITY OF CHICAGO

OBSERVING THE COSMIC MICROWAVE BACKGROUND POLARIZATION
ANISOTROPY AT 40 GHZ WITH QUIET

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
COLIN A. BISCHOFF

CHICAGO, ILLINOIS
AUGUST 2010

Copyright © 2010 by Colin A. Bischoff

All rights reserved

*To my mom,
whose enthusiasm for cosmology
might even exceed my own*

TABLE OF CONTENTS

LIST OF FIGURES	ix
LIST OF TABLES	xii
ACKNOWLEDGMENTS	xiii
ABSTRACT	xv
1 INTRODUCTION	1
1.1 Thesis Overview	1
1.2 The Expanding Universe	1
1.3 The Cosmic Microwave Background	4
1.3.1 Spectrum of the Cosmic Microwave Background	4
1.3.2 Cosmic Microwave Background Temperature Anisotropies	6
1.3.3 The Horizon Problem	9
1.4 Polarization of the Cosmic Microwave Background	10
1.4.1 Descriptions of a Polarization Field	10
1.4.2 Polarization Generated by Thomson Scattering	13
1.4.3 Sources of B -mode Polarization	14
1.5 Foregrounds	18
2 CAPMAP AND QUIET	22
2.1 The Cosmic Anisotropy Polarization Mapper	23
2.1.1 CAPMAP Site, Telescope, and Observation Strategy	23
2.1.2 Calibration	25
2.1.3 Data Analysis	26
2.2 The Q/U Imaging Experiment	28
3 QUIET POLARIMETER MODULES	30
3.1 Components of a QUIET Module	31
3.1.1 Feedhorn	32
3.1.2 Septum Polarizer	32
3.1.3 Low Noise Amplifiers	35
3.1.4 Phase Switches	36
3.1.5 Hybrid Coupler	36
3.1.6 Bandpass Filters	37
3.1.7 Detector Diodes	37
3.2 Basic Theory of Operation	39
3.3 QUIET Module Radiometer Equation	42
3.3.1 Noise Distributions for RF Signals	42
3.3.2 Detected Signal at the $Q1$ Diode	44
3.3.3 Sensitivity of the $Q1$ Diode	47

3.4	Diode Noise Correlation	49
3.4.1	Diode Noise Correlation Without Hybrid Imperfection	50
3.4.2	Diode Noise Correlation Including Coupler Imperfections	51
3.4.3	Correlation of Difference Streams	57
3.5	Phase Switching and Demodulation	58
3.5.1	Phase Switch Imbalance and Double Demodulation	60
3.6	Temperature Differencing Module Pairs	61
3.7	Comparison Between CAPMAP and QUIET Polarimeters	62
3.7.1	CAPMAP Polarimeter Overview	63
3.7.2	Comparison of Correlation Schemes	65
3.7.3	Heterodyne Detection	66
4	THE QUIET INSTRUMENT	67
4.1	QUIET Cryostat	67
4.2	QUIET Electronics Overview	69
4.3	Module Protection and Interface	70
4.3.1	Module Attachment Boards	70
4.3.2	Connecting the Module Array to the Bias Electronics	74
4.4	Bias Electronics	75
4.4.1	Amplifier Bias Board	76
4.4.2	Phase Switch Bias Board	77
4.4.3	Preamplifier Board	78
4.4.4	Housekeeping Board	80
4.5	Analog to Digital Converter Board	81
4.5.1	Data Acquisition	82
4.5.2	Digital Control Signals	83
4.5.3	Timing Synchronization	85
4.6	Electronics Enclosure	86
4.7	QUIET Receiver Software	87
4.7.1	ADC Server	88
4.7.2	Data Compilation	89
4.7.3	Bias Server	90
4.7.4	Receiver Control	91
4.7.5	Peripheral Server	91
4.8	Optimization of QUIET Modules	92
4.8.1	Receiver Temperature Measurement	92
4.8.2	Isolation and Bandwidth Measurement	93
4.8.3	Sensitivity Optimization	94
4.9	Telescope	95
4.10	Groundscreen	96
4.11	Telescope Mount	97

5	OBSERVATION WITH THE 40 GHZ ARRAY	100
5.1	Observing Site	100
5.1.1	Weather and Atmospheric Properties of the Atacama Desert	101
5.2	Scan Strategy	103
5.3	Patches Observed by QUIET	105
5.4	Events of the Q-band Observing Season	108
6	INSTRUMENTAL PROBLEMS AND IMPERFECTIONS	110
6.1	Broken Modules and Diodes	111
6.2	Temperature to Polarization Leakage	112
6.2.1	Suppression of the Temperature Leakage Signal Through Uniform Par- allactic Angle Coverage	113
6.3	Type-A Glitching	115
6.4	Type-B Glitching	117
6.4.1	Correcting Type-B Glitching	119
6.5	High Frequency Noise	121
6.6	Timing Offset Between Telescope and Receiver Data	123
6.7	Beam Sidelobes	124
6.8	Deck Encoder Slips	127
7	CALIBRATION	129
7.1	Pointing Model	129
7.1.1	Definition of Pointing and Orientation Vectors	130
7.1.2	Mount Model	130
7.1.3	Deck Encoder Offset	131
7.1.4	Deck Flexure Correction	131
7.1.5	Azimuth Tilt Correction	132
7.1.6	Elevation Axis Non-perpendicularity	132
7.1.7	Array Collimation	133
7.1.8	Coordinate Transformations	134
7.1.9	Array Pointing Model	135
7.1.10	Pointing Model Parameters	136
7.1.11	Evaluation of Pointing Error	138
7.2	Gain Model	139
7.2.1	Sky Dips	140
7.2.2	Tau A Observations	141
7.3	Noise Model	143
7.3.1	Noise Power Estimation	143
7.3.2	Noise Correlation Estimation	146
7.4	Beam Measurement	146
7.5	Calibration from Moon Observations	147
7.5.1	Why Emission Instead of Reflection?	148
7.5.2	Physical Model of Polarized Emission from the Moon	149

7.5.3	Moon Calibration Observations	152
7.5.4	Fitting Procedure for Moon Observations	153
7.5.5	Understanding the Shortcomings of the Moon Model	156
7.5.6	Results from Moon Calibrations	158
8	DATA SELECTION	162
8.1	Baseline Data Set	163
8.2	Weather Cut	164
8.3	Cut on Noise Model Goodness-of-Fit	167
8.3.1	Low Frequency (10 to 200 mHz) χ^2	170
8.3.2	Medium Frequency (200 mHz to 2 Hz) χ^2	172
8.3.3	High Frequency (2 to 4.6 Hz) χ^2	172
8.4	Knee Frequency Cut	173
8.5	Sun Position Cut	176
8.6	Glitch Cut	178
8.7	Type-B Cut	178
8.8	Summary of Cut Efficiency	179
9	DATA ANALYSIS	181
9.1	Time-Ordered Data Processing	182
9.1.1	High Frequency Cut	184
9.1.2	Linear Azimuth Filter	184
9.1.3	Binned Azimuthal Structure Removal	186
9.1.4	Effect of Combined Filters on CMB Signal	186
9.2	Map Making	187
9.2.1	Point Source Masks	189
9.3	Power Spectrum Estimation	192
9.3.1	Using Day-to-Day Cross-Correlation to Avoid Noise Bias	194
9.3.2	Overview of Analysis Procedure	195
10	ASSESSING DATA QUALITY	196
10.1	Calculating Null Power Spectra	197
10.2	The QUIET Null Suite	198
10.2.1	Module Divisions	199
10.2.2	Pointing Divisions	200
10.2.3	Temporal Divisions	202
10.2.4	Performance Divisions	203
10.2.5	Weather Divisions	205
10.2.6	Housekeeping Divisions	205
10.3	Evaluating Null Suite Results	206
10.4	Null Suite Results	207
10.5	Prospects for the Analysis	210
A	MOON CALIBRATION RESULTS	214

B	QUIET DATA SELECTION USING MYSQL	224
B.1	Overview of Database Tables	224
B.2	Standard Data Selection Query	226
B.3	Nulltest Selection Criteria	229
B.3.1	Module Divisions	229
B.3.2	Pointing Divisions	230
B.3.3	Temporal Divisions	233
B.3.4	Performance Divisions	234
B.3.5	Weather Divisions	236
B.3.6	Housekeeping Divisions	237
C	NULL SUITE TOY MODEL	238
	REFERENCES	243

LIST OF FIGURES

1.1	Spectrum of the Cosmic Microwave Background	5
1.2	Full sky CMB temperature map, as measured by WMAP	7
1.3	WMAP CMB temperature spectrum	8
1.4	Illustration of E - and B -modes	12
1.5	Illustration of linear polarization from Thomson scattering	14
1.6	Measurements of the CMB EE spectrum	15
1.7	Illustration of B -mode polarization generated by a gravitational wave	16
1.8	Measurements of the CMB BB spectrum	18
1.9	Foreground maps for synchrotron, free-free, and dust emission from WMAP	20
2.1	The Crawford Hill 7-meter Antenna	24
2.2	Distribution of χ^2 values from the CAPMAP null suite	28
2.3	EE and BB power spectra from CAPMAP third season	29
3.1	Schematic of a QUIET module	33
3.2	Photograph of the interior of a QUIET module	34
3.3	Diode noise correlation with power imbalance	52
3.4	Q – U diode noise correlation with power imbalance and hybrid coupler imperfection	54
3.5	$U1$ – $U2$ diode noise correlation with power imbalance and hybrid coupler imperfection	56
3.6	Effect of 4 kHz phase switching on $1/f$ noise power.	59
3.7	Schematic of a CAPMAP polarimeter	64
4.1	Schematic diagram of QUIET bias and readout electronics	71
4.2	Modules installed in a Module Attachment Board	73
4.3	Flexible Printed Circuits sealed into a hermetic feedthrough	75
4.4	Blanking and demodulation of 800 kHz data	83
4.5	Illustration of a serial bias command	84
4.6	Histogram of enclosure temperature and maximum variation	87
4.7	Schematic of QUIET receiver software	88
4.8	Photograph showing the lower groundscreen on the QUIET mount	97
4.9	Diagram of the axes of rotation for the QUIET mount	98
5.1	Precipital water vapor levels at the QUIET site	102
5.2	Coverage map for one observing run	104
5.3	Patches of the sky observed by QUIET	106
5.4	Patch availability during an observation day	107
6.1	Parallactic angle coverage for patch 2a.	114
6.2	Example of Type-A glitching	116
6.3	Illustration of “Type-B” glitching.	118
6.4	Example of Type-B glitching and correction	121

6.5	Common mode high frequency noise	122
6.6	Timing offset between pointing and radiometer data	124
6.7	Illustration of the timing offset correction	125
6.8	Illustration of QUIET sidelobes	126
6.9	Sun contamination due to sidelobe pickup	127
6.10	Deck offset measurements	128
7.1	Array pointing relative to module RQ-09	138
7.2	Magnitude of mount model correction	139
7.3	Residual pointing errors for 18 moon observations	140
7.4	Noise power spectrum for RQ-09 diode $Q1$ from CES 128.0	145
7.5	Total power and polarization intensity as a function of radial distance for emission from the moon.	152
7.6	Simulated \mathbf{I} , \mathbf{Q} , and \mathbf{U} maps of the moon	153
7.7	Bias in detector angles recovered from moon simulations with temperature non-uniformity.	157
7.8	Detector angles measured from moon observations	159
7.9	$\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients measured from moon observations	160
7.10	Relative gain between RQ-05 diode $Q1$ and RQ-09 diode $Q1$, measured from moon observations	161
8.1	Fraction of CES passing weather cuts, as a function of time of day and time of year	168
8.2	Distribution of time of day for observations of the four CMB patches	168
8.3	Filtered power spectrum and noise model used for data selection	169
8.4	Scan synchronous signal as a function of sidelobe elevation	171
8.5	Diode knee frequency histograms	175
8.6	Cut on sun position relative to the telescope boresight	177
9.1	Filters employed for time-ordered data processing	183
9.2	Amplitude of azimuth polynomials of up to fifth order for CES with various weather conditions	185
9.3	Effect of TOD filtering on CMB signal	187
9.4	Inverse noise weighted polarization maps for Patch 2a	190
9.5	Inverse noise weighted polarization maps for Patch 4a	190
9.6	Inverse noise weighted polarization maps for Patch 6a	191
9.7	Inverse noise weighted polarization maps for Patch 7b	191
10.1	Null suite results for Patch 2a	208
10.2	Null suite results for Patch 4a	208
10.3	Null suite results for Patch 6a	209
10.4	Null suite results for Patch 7b	209
10.5	Projected results for QUIET measurements of the EE and BB spectra at 40 GHz	213

C.1	Comparison between the null suite toy model and Monte Carlo results	241
-----	---	-----

LIST OF TABLES

4.1	Correspondence between modules and their support electronics	72
5.1	Patches observed by QUIET	106
6.1	Parallactic angle of the array orientation averaged over all CES for each of the QUIET patches.	115
7.1	Collimation offset parameters for each module, relative to RQ-09	136
7.2	Mount model parameters	137
8.1	Total number and duration for all CES by patch	163
8.2	Data excluded from the baseline data set	165
8.3	Knee frequency cut thresholds for all detector diodes, in mHz	174
8.4	Number of CES-diodes rejected by various cuts	180
10.1	Null Suite Status by Patch	211
A.1	Moon drift scans used for calibration.	215
A.2	Detector angles for Q diodes, in degrees, as measured from moon observations.	216
A.3	Detector angles for U diodes, in degrees, as measured from moon observations.	217
A.4	Fractional $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients for Q diodes, as measured from moon observations.	218
A.5	Fractional $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients for U diodes, as measured from moon observations.	219
A.6	Average (total power) gains of Q diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.	220
A.7	Average (total power) gains of U diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.	221
A.8	Demodulated (polarized) gains of Q diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.	222
A.9	Demodulated (polarized) gains of U diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.	223

ACKNOWLEDGMENTS

QUIET is a huge group effort. While I only have room to single out a few people, it was a pleasure to work with you all (even the telecons).

The first person to thank is my advisor, Bruce, who convinced me years ago that QUIET would be a great project to work on (which is true!). From oscilloscopes to χ^2 distributions, he has always pushed me to better understand the instrument and the data.

The Chicago QUIET group has been both productive and a lot of fun to work with. Dorothea has always provided a great example of insightful data analysis, dating back to my early days with CAPMAP. Dan carried the receiver development through the toughest stretches, when hundreds of bugs needed ironing out, and he showed me that all you need to build a good calibrator is a machine shop, epoxy, and plenty of liquid nitrogen. Akito taught me most of what I know about CMB analysis and still makes sure that I'm keeping everything rigorous. Ali always amazes me with her ability to get things done, which is something that we always needed more of. Immanuel showed up and immediately became a critical member of the team; QUIET will be in very good hands with him. Kendrick never fails to provide astonishing contributions whenever he appears. Osamu has handled every job that might otherwise slip through the cracks, in lab, at the site, or in the data, always with incredible craftsmanship. Mircea gave me some great lessons in electronics design and met every deadline that he set.

The rest of the QUIET collaboration, at institutions across the world, was instrumental in making the project happen. I especially want to thank Suzanne, who has been like a second advisor for me on both CAPMAP and QUIET. Josh, who also worked on both projects, has done an excellent job leading the QUIET optics design and calibration. The Columbia QUIET group — Amber, Ross, Laura, Robert, Jon, and Michele — assembled the Q-band receiver that plays a prominent role in this thesis. The work that Todd and Kieran have done at JPL on module design and performance made the idea of correlation polarimeter

arrays into a reality.

Keith was a great friend and colleague on CAPMAP but then he escaped to the South Pole just when things got ultra-fun. While we never quite made it so far as actually writing a song, Joaquin, Jay, and I did craft a few big balls of sound. Jason and Eric always kept life lively, especially on Wednesday nights. Tuesday nights belonged to Team Suki, though our triumphs were few and far between. Thanks to Cora for a thousand office conversations on hundreds of topics, and for sharpening my idl skills one question at a time.

Trying to get QUIET down on paper was an occasionally rocky process, so I am very grateful for all the support I had while writing this thesis. Bob and Lynn were incredibly hospitable and I enjoyed my stay in Sturbridge, even though I was working the whole time. Lilly has provided the best entertainment. Thanks also to my mom and Steve who helped me escape the winter and provided an idyllic spot to get work done. The upside of spending a few months away from Chicago was the chance to see a lot more of my family than I have in years. Thanks to my dad, mom, Steve, Kara, Ben, George, Michael, Melissa, David, Wende, Zoe, Casey, and Christine for the love and encouragement.

Most of all, thank you Jenn for your total support.

ABSTRACT

The Cosmic Microwave Background has proven to be an invaluable tool for cosmology, but much of the scientific information contained in its polarization anisotropy is still to be discovered. Of particular interest is the angular power spectrum of the tiny B -mode polarization signal, the divergence free component which is geometrically distinct from the dominant E -modes. The theory of inflation, which has been proposed to explain several curious features of our universe, predicts the existence of B -mode fluctuations on large angular scales and at a signal level that may be achievable in the near future.

The Q/U Imaging Experiment (QUIET) is designed to make sensitive measurements of the CMB polarization using arrays of integrated “detectors on a chip” that can be scaled to reach the sensitivity requirements for detecting B -modes.

This thesis covers the design of QUIET, which consists of separate instruments for observation at 40 and 90 GHz. After the discussion of the detectors, which generally applies to both instruments, the remaining sections focus on observations, calibration, and data analysis for the 40 GHz detector. Tests performed on the data indicate some not-yet-identified source of contamination that must be understood and corrected before a reliable result can be obtained.

CHAPTER 1

INTRODUCTION

1.1 Thesis Overview

The focus of this thesis is the Q/U Imaging Experiment (QUIET), an instrument designed to measure the angular power spectrum of the polarized component of the Cosmic Microwave Background (CMB). Also included is some discussion of the Cosmic Anisotropy Polarization Mapper (CAPMAP) experiment, which preceded QUIET. While the two projects share many similarities, the areas where they contrast highlight developments in the detector technology and their differing science goals.

This chapter describes the Cosmic Microwave Background and the motivation for measuring its polarization. An overview of CAPMAP is included in §2. The design of the QUIET instrument is described in §3 and §4. A summary of the observations is included in §5. The actual operation of the experiment did not always go smoothly, of course; descriptions of several problems that arose, and the solutions that were found, are discussed in §6. The remaining chapters are devoted to the analysis of QUIET data. Calibration is covered in §7, data selection in §8, analysis methods in §9, and tests used to verify the data quality in §10. The test results point to some issues in the data that are currently unsolved, but prospects for addressing the problem are discussed.

1.2 The Expanding Universe

The cosmological principle states that, on sufficiently large scales, the universe is homogeneous and isotropic. Put simply, this means that there are no preferred locations or preferred directions in space. Enforcing these conditions on the space-time metric yields the

Friedmann–Robertson–Walker (FRW) metric

$$g_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & a^2(t) & 0 & 0 \\ 0 & 0 & a^2(t) & 0 \\ 0 & 0 & 0 & a^2(t) \end{bmatrix} \quad (1.1)$$

The FRW metric is uniform in space but features a time-dependent scale factor, $a(t)$, that determines physical distances. The scale factor is defined to be equal to 1 at the present time. Edwin Hubble’s pioneering measurement of the redshift–distance relation (Hubble’s Law) in 1929 found that distant galaxies recede with velocity proportional to their distance. This demonstrated that the universe is expanding, or that a is increasing with time. That the size of the universe is increasing now, and was increasing in the past, implies that the early universe was a smaller, denser, and hotter place. This is the basis of the Big Bang theory of cosmology, which was dramatically different from previously held assumptions that the universe exists in a steady state.

The dynamics of the scale factor depend on the energy density of the universe. The Hubble rate, $H(t) \equiv (\dot{a}/a)$, is defined to describe the change in the scale factor, with present day value H_0 . The Hubble rate is calculated as

$$\frac{H^2(t)}{H_0^2} = \frac{\rho}{\rho_{cr}} \quad (1.2)$$

where $\rho_{cr} \equiv \frac{3H_0^2}{8\pi G}$ is known as the critical density.

The energy density, ρ , also varies with the scale factor, with different dependence for different types of matter or energy. For non-relativistic massive particles, rest mass accounts for nearly all of the energy, with a negligible kinetic energy contribution. For temperatures significantly below the chemical potential of the particle species, number count is preserved,

so the energy density of that species scales inversely with volume, like a^{-3} . Substituting this dependence into equation (1.2), the time dependence of the scale factor is found to be $a(t) \propto t^{2/3}$.

The energy in photons depends on their wavelength, which redshifts as the scale factor increases. This causes their energy density to decrease like a^{-4} , faster than it would due to the volume increase alone. For a universe dominated by photons, or highly relativistic matter, the scale factor varies in time proportional to $t^{1/2}$.

Evidence for a third type of energy, known as dark energy, has been found from measurements of Type 1A supernovas (Riess et al. (1998), Perlmutter et al. (1999)), which are used to extend the reach of the redshift–distance relation. These data provide evidence that the expansion of the universe is accelerating. The simplest model proposed to explain this phenomena is the cosmological constant, Λ , which has constant energy density, independent of the scale factor. Solving equation 1.2 for this case results in a scale factor that increases exponentially with time.

The FRW metric, as written in equation (1.1), is for the case of flat space-time, which requires that the total energy density, from matter, radiation, and dark energy, is equal to the critical density. If the actual density is higher or lower than this value, then the universe has positive or negative curvature, with consequences for the scale factor evolution. Observations indicate that the universe is flat, or very close to it.

As the scale factor varies, the dominant component of the energy density changes. At very early times, when $a \ll 1$, photons were the largest component. At the epoch of matter–radiation equality, the photons have redshifted sufficiently that their energy density falls below the density of matter. Eventually, the matter also spreads out to the point where dark energy becomes the most significant component. Observations suggest that we are currently in the midst of this second transition, entering an era dominated by dark energy, with dark energy density equal to $0.73\rho_{cr}$ and matter density only equal to $0.27\rho_{cr}$. This picture is

known as Λ CDM cosmology, with Λ indicating the presence of a cosmological constant and CDM representing cold dark matter, which makes up about 80% of the matter density in the theory, with the rest mostly comprised of the more familiar protons and neutrons.

1.3 The Cosmic Microwave Background

Looking out into the universe at light from very distant sources, we observe back into the past. Those earlier times correspond to smaller values of a , larger densities, and higher temperatures. When the temperature is sufficiently high – about 3000K or a quarter of an electron-Volt – all atoms in the universe (mostly Hydrogen) are ionized. The resulting plasma has a small mean free path for photons, so the universe is opaque beyond this point. This screen, which can be seen in all directions, is known as “the surface of last scattering”. The epoch in which the universe transitioned from an ionized to a neutral state is referred to as “recombination”¹.

The emission from this primordial plasma was first detected by Arno Penzias and Robert Wilson in 1965 using a horn antenna at the Crawford Hill Laboratory in Holmdel, New Jersey (Penzias & Wilson, 1965). Their measurement was performed at a frequency of 4.08 GHz, in the microwave region of the electromagnetic spectrum. Though it had been 36 years since Hubble’s discovery of cosmic expansion, the Big Bang theory was not fully accepted. The significance of Penzias’ and Wilson’s discovery was quickly appreciated, leaving little doubt that the universe today is descended from an early hot and dense state.

1.3.1 *Spectrum of the Cosmic Microwave Background*

While Penzias and Wilson took great care to rule out any possibility that the signal they measured was due to a local source, they only made measurements at a single frequency.

1. This is a misnomer, since the electrons and protons were never combined at any earlier stage.

Much stronger evidence for the cosmological origin of the microwave background radiation is found from its spectrum. If the signal originated from the primordial plasma, which is made up of electrons and protons in equilibrium with a photon gas, then it should have a characteristic blackbody spectrum. The alternatives – signals of terrestrial or galactic origin – would have other types of spectra.

The Far Infrared Absolute Spectrometer (FIRAS) instrument on board the Cosmic Background Explorer (COBE) satellite made a definitive measurement of the CMB frequency spectrum, shown in Figure 1.1 (Mather et al., 1999). The spectrum is a near perfect match to a blackbody at $2.725 \pm 0.002\text{K}$. A measurement of the CMB temperature also pins down the redshift at which it was emitted. A temperature of 3 degrees Kelvin is a factor of ~ 1100 less than what is necessary to fully ionize the Hydrogen in the early universe, indicating that the scale factor at the time of decoupling was $1/1100 = 0.0009$.

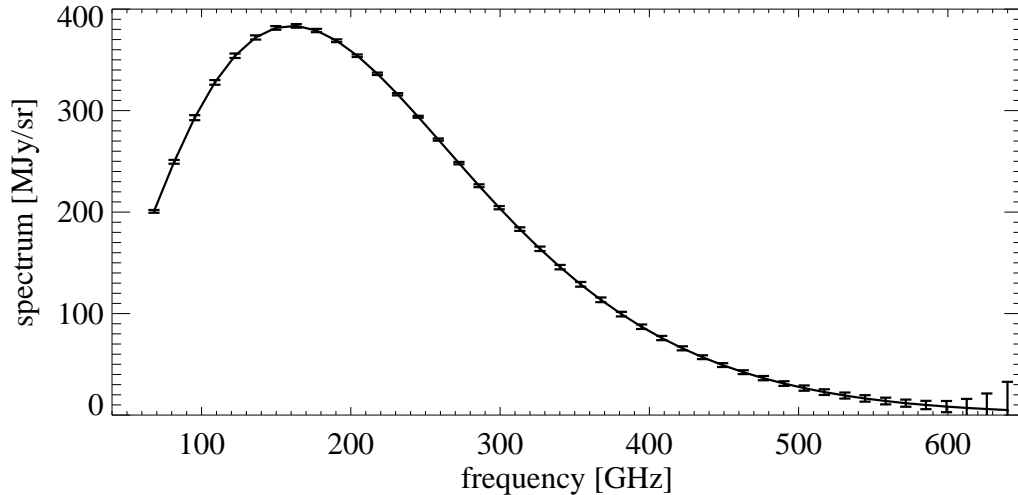


Figure 1.1: Spectrum of the Cosmic Microwave Background as measured by the Far Infrared Absolute Spectrometer (FIRAS) on the Cosmic Background Explorer (COBE) satellite. Units on the y axis are flux density per unit solid angle ($1 \text{ MJy} = 10^{-20} \text{ W/m}^2/\text{Hz}$). The solid line is an ideal 2.725K blackbody spectrum; error bars indicate the 100σ interval for the FIRAS data points.

As an aside, the amplitude of the CMB and its anisotropies are usually expressed in units of temperature. This is appropriate, given the blackbody nature of the source. However, microwave detectors do not measure temperature; they measure the electromagnetic flux in a frequency band. Fortunately, these two quantities are closely related. In the Rayleigh-Jeans (low frequency) limit, the blackbody spectrum can be approximated as

$$I_\nu = \frac{2\nu^2 k_B T}{c^2} \quad (1.3)$$

In this limit, the observed flux is proportional to temperature T . Throughout this thesis, instrumental response will frequently be quoted in Kelvin, using a simple gain factor to convert from the raw detector output. In fact, a small correction factor to the gain, calculated using the frequency response of the detector, is needed to account for the true shape of the CMB blackbody spectrum.

1.3.2 Cosmic Microwave Background Temperature Anisotropies

One of the most notable features of the Cosmic Microwave Background is that it has an extremely uniform temperature across the sky. Small temperature variations, at the level of tens of μK , or a part in 10^5 , were first detected by the Differential Microwave Radiometer (DMR) carried by COBE (Smoot et al., 1992). Since then, many ground and balloon based experiments have refined the measurements, with finer angular resolution and higher sensitivity.

Currently, the best measurement of the CMB temperature anisotropies at large to moderate angular scales comes from the Wilkinson Microwave Anisotropy Probe (WMAP) satellite, which launched in 2001 and released results from its first year of observation in 2003 (Hinshaw et al., 2003). A map of the microwave sky from the most recent WMAP data release (7 year) is shown in Figure 1.2.




Figure 1.2: Full sky temperature map from the WMAP 7 year data release (Jarosik et al., 2010a). Galactic foreground emission has been subtracted using the Internal Linear Combination method described in Hinshaw et al. (2007).

The fluctuation power of the anisotropies is measured as a function of scale using the angular power spectrum. The temperature field over the celestial sphere is decomposed using the spherical harmonics, $Y_{\ell m}$, as an orthonormal basis. For each value of ℓ , which corresponds to a particular angular scale, there are $2\ell + 1$ independent functions with m ranging from $-\ell$ to ℓ , in integer steps. If amplitude values $a_{\ell m}$ are measured for each value of m , the angular power spectrum C_ℓ is defined to be their variance:

$$C_\ell = \langle a_{\ell m}^* a_{\ell m} \rangle \quad (1.4)$$

If the temperature field is Gaussian, then all information about its statistical properties is contained in the power spectrum (there are no correlations between multipole moments with different values of ℓ or m). Note that, because there are a finite number of modes for each value of ℓ , C_ℓ is an imperfect estimate of the true variance of the distribution from which $a_{\ell m}$ values are drawn. This limitation, known as cosmic variance, means that even a perfect

measurement of the CMB across the entire sky has finite power to constrain theory.

A close inspection of Figure 1.2 might reveal that the CMB tends to feature hot and cold spots at characteristic scales, but it becomes very clear once the angular power spectrum is plotted. Figure 1.3 shows the spectrum of temperature fluctuations (TT spectrum) reported in the WMAP 7 year data release (Larson et al., 2010).




Figure 1.3: CMB temperature (TT) spectrum reported in the WMAP 7 year data release (Larson et al., 2010). Vertical error bars include contributions from cosmic variance and instrumental noise. The red curve is calculated for Λ CDM cosmology using CAMB (Lewis et al., 2000).

The rich structure of the CMB angular power spectrum can be understood from a simple description of the hot plasma in the early universe. We start with initial conditions that consist of a mostly smooth density field plus small random fluctuations with a flat power spectrum (no characteristic size). The causal horizon grows with time since the Big Bang, setting the scale for the largest fluctuations which are able to evolve at a given time. Once the causal horizon exceeds a particular scale, gravity acts to increase overdense regions, pulling

in surrounding matter. This process is opposed by photon pressure from light which is coupled to the plasma. These two forces create conditions for oscillation of the initial density perturbations. Eventually, the universe expands and cools to the point of recombination and the oscillations cease. At this point, each mode has had a particular length of time to evolve, based on the time when it entered the causal horizon. For the scale corresponding to $\ell \simeq 200$, or an angular diameter on the sky of about 1° , there was just enough time for initial overdensities to collapse to their point of maximum compression, resulting in larger fluctuation power. The second peak, near $\ell = 500$, represents the scale that collapsed, expanded, and had just reached a point of maximum rarefaction at recombination. Continuing to smaller scales, there is a harmonic series of spectral peaks (continuing beyond those shown in Figure 1.3). The modes with multipole less than ~ 20 are larger than the size of the causal horizon at recombination, so they are imprinted on the CMB at their primordial value.

1.3.3 The Horizon Problem

A paradox arises from the features of the Cosmic Microwave Background at large angular scales. If, by the time of decoupling, causal contact had only been established over scales that corresponds to angles of several degrees, then it is hard to understand why the CMB temperature is uniform to a part in 10^5 across the entire sky. Moreover, the existence of CMB fluctuations with $\ell < 20$ requires some mechanism to create initial perturbations with wavelength much larger than the horizon.

The theory of inflation has been proposed as a solution to this problem (Guth, 1981). Inflation refers to a brief period that occurs just a fraction of a second after the Big Bang. During this instant, the scale factor undergoes exponential expansion, increasing by roughly 60 e -foldings ($\sim 10^{26}$). This rapid stretching of space-time means that our entire observable universe, including points that are still out of causal contact at the present day, was once a very small, causally connected region.

The most common mechanism used to drive inflation is a homogeneous scalar field in a state of high potential energy. This results in a constant energy density throughout space, which acts like the cosmological constant described in §1.2 to make the scale factor increase exponentially. In the “slow roll” inflation scenario, the potential is chosen so that the scalar field slowly returns to its potential minimum, at which point inflation ceases.

Inflation makes an excellent solution to the horizon problem – our entire universe is at a uniform temperature because there was a time in the past when it was able to equilibrate. A flat spectrum of initial density perturbations are generated from quantum fluctuations blown up to large size. Inflation also explains the observed spatial flatness of the universe. The downside is that there is no mechanism in the Standard Model that could drive inflation. While they are a common feature of quantum field theories, scalar fields have never been observed. Direct evidence for inflation would be an astonishing discovery with implications for both cosmology and high energy physics.

1.4 Polarization of the Cosmic Microwave Background

In addition to the temperature anisotropies, the Cosmic Microwave Background also carries polarization information. The power in the polarization spectrum is about two orders of magnitude smaller than the TT spectrum, so the first detection was not made until 2002 by the Degree Angular Scale Interferometer (Kovac et al., 2002). Since that time, numerous other experiments, including CAPMAP, have reported measurements with increasing precision. The CMB polarization is one of the most promising sources of information about inflation, which accounts for the high level of enthusiasm in the field.

1.4.1 Descriptions of a Polarization Field

The CMB polarization field is often represented using the Stokes parameters. The parameter \mathbf{I} specifies the unpolarized intensity (or temperature), \mathbf{Q} and \mathbf{U} specify the linear polariza-

tion, and \mathbf{V} measures circular polarization. In a two-dimensional coordinate system, with two sets of Cartesian coordinates, (x, y) and (a, b) , defined at 45° with respect to each other, the Stokes parameters are calculated as follows.

$$\mathbf{I} = \frac{1}{2} (|E_x|^2 + |E_y|^2) \quad (1.5)$$

$$\mathbf{Q} = \frac{1}{2} (|E_x|^2 - |E_y|^2) \quad (1.6)$$

$$\mathbf{U} = \frac{1}{2} (|E_a|^2 - |E_b|^2) = \frac{1}{2} (E_x^* E_y + E_x E_y^*) \quad (1.7)$$

$$\mathbf{V} = \frac{1}{2} (|E_L|^2 - |E_R|^2) = i \frac{1}{2} (E_x^* E_y - E_x E_y^*) \quad (1.8)$$

Circular polarization is not expected to be found in the CMB, so we will focus on linear polarization only.

When defining Stokes parameters on a sphere, the convention is for the x axis to point North–South while the y axis points East–West². This definition depends on the choice of spherical coordinate systems, usually equatorial or galactic.

In the discussion of detectors in §3, \mathbf{Q} and \mathbf{U} are often referred to in the instrument basis, where \mathbf{Q} represents the difference in polarization between the perpendicular symmetry axes of the detector, and \mathbf{U} is defined to be 45° away.

An alternate description of the polarization field is to express it in terms of gradient-type and curl-type components, known as E - and B -fields in analogy to the electric and magnetic fields. Simple examples of E - and B -modes can be seen in Figure 1.4. For the E -mode case, the polarization is parallel or perpendicular to the wave vector of the fluctuation, so a reflection about that axis leaves the pattern unchanged. In the B -mode case, the polarization is oriented at 45° to the wave vector and a reflection causes the pattern to flip sign.

To describe the CMB polarization across the entire celestial sphere, the spin-2 spherical

2. There is still room for ambiguity in defining the $\pm\mathbf{U}$ axes (NW–SE or NE–SW?). Naturally, the convention defined by the IAU does not match the one used in the Healpix pixelization software.




Figure 1.4: Example E -mode (left) and B -mode (right). For both cases, the direction of the wave vector points upward. For the E -mode, the polarization is aligned parallel or perpendicular to the wave vector; for the B -mode, the polarization direction is rotated by 45° from \vec{k} .

harmonics $\pm 2Y_{\ell m}$ are used as basis functions (Zaldarriaga & Seljak (1997), Kamionkowski et al. (1997), Hu & White (1997)). From the amplitudes $a_{\pm 2, \ell m}$ in this basis, E -mode and B -mode amplitudes can be calculated as

$$a_{E, \ell m} = -(a_{2, \ell m} + a_{-2, \ell m})/2 \quad (1.9)$$

$$a_{B, \ell m} = \imath(a_{2, \ell m} - a_{-2, \ell m})/2 \quad (1.10)$$

and the polarization power spectra are calculated as

$$C_\ell^{EE} = \langle a_{E, \ell m}^* a_{E, \ell m} \rangle \quad (1.11)$$

$$C_\ell^{BB} = \langle a_{B, \ell m}^* a_{B, \ell m} \rangle \quad (1.12)$$

Besides these autocorrelation spectra, one can also measure the cross spectra, C_ℓ^{TE} , C_ℓ^{TB} , and C_ℓ^{EB} . However, because the B -modes flip sign under reflection, C_ℓ^{TB} and C_ℓ^{EB} must be zero by symmetry.

1.4.2 Polarization Generated by Thomson Scattering

CMB polarization is sourced by Thomson scattering of photons off of free electrons during the era of recombination. An unpolarized photon with wavevector \vec{k}_0 parallel to the \hat{x} axis has transverse electric fields that can drive an electron in the \hat{y} or \hat{z} directions. Scattered radiation with \vec{k}_1 along the \hat{y} axis will be due only to oscillations of the electron along \hat{z} , so it is linearly polarized.

Of course, for an electron in the surface of last scattering, photons are incident from all directions, so the polarization generally averages away to zero. A net linear polarization can occur if the temperature field has a quadrupolar pattern in the plane perpendicular to the line of sight. For this case, the y component of the scattered light could have a different temperature than the z component, as illustrated in Figure 1.5. This description only considers incoming radiation perpendicular to the line of sight, which would be the case if the surface of last scattering was a sharp transition (i.e. the universe transitioned instantaneously from a state where the mean free path of a photon was zero to one where the mean free path is infinite). In fact, the visibility function, which measures the probability that a CMB photon last scattered at a particular redshift, has support over the range $900 \lesssim z \lesssim 1200$, so the “surface” of last scattering is really an era. This significantly reduces the polarization fraction of the CMB, accounting for why the polarization spectrum is ~ 100 times smaller than the temperature spectrum.

A critical feature of the polarization pattern created by quadrupole temperature anisotropies is that it can produce E -modes only. This can be understood intuitively due to the fact that a mode in the temperature field has only one direction associated with it – the direction of the wave vector – which provides both the orientation of the polarization and the direction along which it varies. This is the case for the polarization pattern produced by any scalar perturbation, but B -modes can be produced by tensor perturbations, as discussed in §1.4.3.

Figure 1.6 shows the angular power spectrum of E -mode polarization and measurements




Figure 1.5: Linear polarization due to Thomson scattering off an electron located in a quadrupole of the temperature field. Higher temperature photons incident from the $\pm\hat{x}$ direction and lower temperature photons incident from the $\pm\hat{z}$ direction result in a net polarization for light scattered in the \hat{y} direction. Figure adapted from Hu & White (1997).

reported by several experiments since the original detection in 2002. Because the signal is a direct consequence of CMB temperature fluctuations, the theoretical curve is tightly constrained by the precise measurements of the TT spectrum from WMAP and others.

1.4.3 Sources of B -mode Polarization

While temperature anisotropies produce only E -modes, B -mode polarization patterns can be imprinted on the CMB by tensor perturbations such as gravitational waves. Besides its wave vector, a gravitational wave has an additional associated direction due to its polarization. Figure 1.7 illustrates a B -mode pattern created by a gravitational wave passing through the last scattering surface at a 45° angle. Seen in projection, the polarization amplitude varies along the horizontal direction, but the orientation is rotated by 45° from that direction, resulting in a pure B -mode. For other geometries, gravity waves can create E -mode patterns.

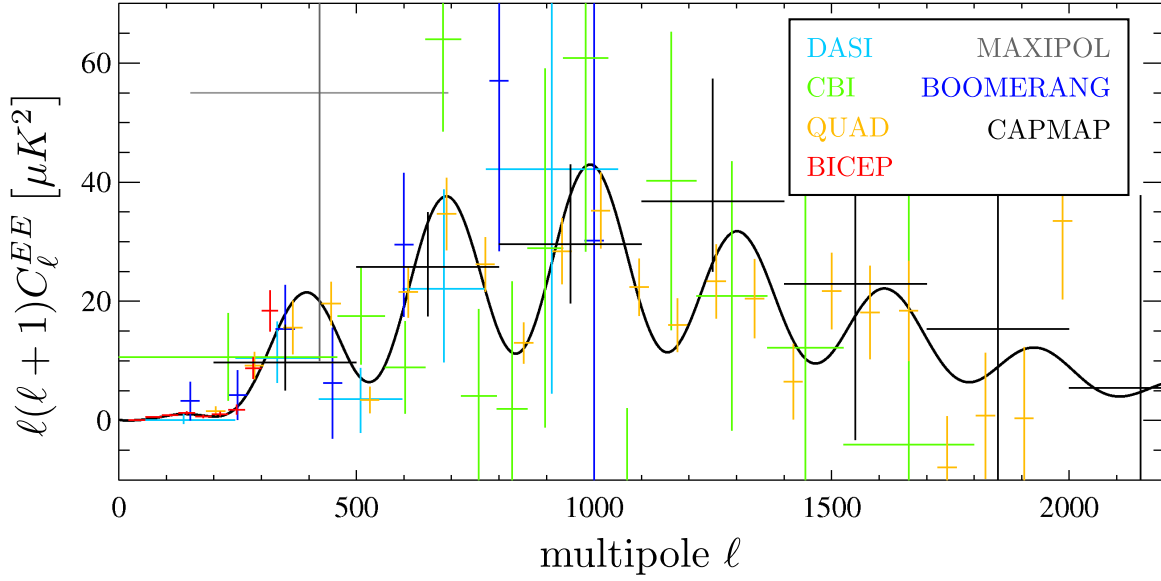


Figure 1.6: Measurements of the CMB EE spectrum from DASI (Leitch et al., 2005), BOOMERANG (Montroy et al., 2006), MAXIPOL (Wu et al., 2007), CBI (Sievers et al., 2005), CAPMAP (Bischoff et al., 2008), QUAD (Brown et al., 2009), and BICEP (Chiang et al., 2010). The black curve is calculated for Λ CDM cosmology using CAMB (Lewis et al., 2000). The E -mode polarization has been detected at high significance, including confirmation of the acoustic peak structure seen in the theoretical curve (focus on the orange points from QUAD to see this).

For a general population of tensor fluctuations, E - and B -mode power will be generated in roughly equal amounts (Hu & White, 1997).




Figure 1.7: On the left, a gravitational wave passes through the surface of last scattering, which is represented by the blue rectangular frame. Projected onto that surface, the induced polarization shows the B -mode pattern seen on the right.

The existence of a background of primordial gravity waves is a signature prediction of inflation. These tensor perturbations are generated in the same way as the initial spectrum of scalar density perturbations, through quantum fluctuations blown up to cosmological scales³. At the time of recombination, the gravity waves leave an imprint on the temperature, E -mode polarization, and B -mode polarization anisotropies of the CMB. The fact that only tensor fluctuations are capable of creating B -mode polarization means that measurements of B -modes are an extremely powerful technique to search for evidence of inflation.

The magnitude of the B -mode signal from gravity waves is parametrized by r , the ratio of tensor to scalar mode power generated by inflation. The energy scale, $V^{1/4}$, at which inflation occurred is proportional to $r^{1/4}$ (Baumann et al., 2009).

$$V^{1/4} = (3.35 \times 10^{16} \text{GeV}) r^{1/4} \quad (1.13)$$

3. Vector perturbations should be produced as well, but these are not expected to persist for long enough to affect the CMB.

The current best upper limit on r is 0.24, at 95% confidence, from WMAP in combination with measurements of Baryon Acoustic Oscillations (BAO) and a prior on the current value of Hubble’s constant. Most of the power in this constraint is due to the high precision measurement of the temperature power spectrum, not B mode polarization. The BICEP experiment constrains r to be less than 0.72, which is the tightest constraint directly from B -mode upper limits. However, with many new polarimeters (like QUIET!) nearing deployment or already taking data, polarization will soon be the primary constraint. A detection of primordial tensor fluctuations in the range $0.001 < r < 0.1$ would provide strong confirmation of inflation. Additionally, this measurement would provide a glimpse of physics at energy scales of 10^{16} GeV, the scale of Grand Unified Theories in particle physics, which is many orders of magnitude beyond what can ever be achieved in accelerators.

However, there are other sources of B -mode polarization. The CMB photons pass through the entire observable universe as they travel to us from the surface of last scattering. Large scale structure in the universe, such as galaxy clusters, will gravitationally lens the light as it passes through deep potential wells. This distortion mixes E and B modes, contaminating the clean B mode channel with a small fraction of the much larger E mode power. Fortunately, the power spectrum for the lensing B modes peaks at small angular scales ($\ell \sim 1000$), so it can be differentiated from the gravity wave spectrum, which is concentrated at low ℓ . Also, the lensing power spectrum contains interesting physics itself. For example, the growth of structure at the largest scales is sensitive to the average mass of the neutrino species (Smith et al., 2008).

Figure 1.8 shows the BB spectrum from inflationary gravity waves for r equal to 0.1 (upper black curve) or 0.01 (lower black curve), and the spectrum of the lensing B modes. No detections of B mode polarization have been reported to date, and the plotted 95% confidence upper limits are all substantially higher than the expected signals. To understand the difficulty of this measurement, simply compare the y axis scale between Figures 1.3, 1.6,

and 1.8.




Figure 1.8: Upper limits, at 95% confidence, for the CMB BB spectrum from DASI (Leitch et al., 2005), BOOMERANG (Montroy et al., 2006), MAXIPOL (Wu et al., 2007), CBI (Sievers et al., 2005), CAPMAP (Bischoff et al., 2008), QUAD (Brown et al., 2009), and BICEP (Chiang et al., 2010). The black curves show the expected BB spectrum from gravitational lensing and for primordial gravity waves with tensor to scalar ratio $r = 0.1$ (upper curve) or 0.01 (lower curve).

1.5 Foregrounds

Microwave emission from foregrounds, primarily in our own galaxy, represents a significant potential difficulty for measurements of the Cosmic Microwave Background. Foreground sources can be categorized into three classes, which have different spectra generally characterized by a power law spectral index. This is key to the identification and removal of foreground signals, because none of the foregrounds match the 2.7K blackbody spectrum of the CMB.

High energy electrons moving in the galactic magnetic field produce synchrotron emission.

The spectrum of synchrotron emission depends on the population of free electrons and the distribution of magnetic fields in the galaxy. These both vary significantly between different locations, leading to spectral indexes in the range $-1 \lesssim \alpha \lesssim -0.5$ for the observed flux density, or spectral indexes in the range $-3 \lesssim \beta \lesssim -2.5$ when compared to a Rayleigh-Jeans spectrum⁴ (Bennett et al., 2003).

Scattering of electrons off ions in the galaxy leads to a foreground signal known as free-free emission. Like synchrotron emission, the spectrum of free-free emission falls off relative to the Rayleigh-Jeans spectrum for increasing frequency, with $\beta \sim -2.15$. Free-free emission is subdominant to other foregrounds across the microwave frequency range.

The third foreground component is thermal emission from dust at temperatures around 20K. The dust has a rising spectrum, with typical spectral indexes of $1.5 \lesssim \beta \lesssim 2$ and emission into the far infrared. Again, the spectral index can vary across the sky, with dependence on temperature and dust grain composition.

Figure 1.9 shows temperature maps of the three foreground components, estimated from WMAP observations at five frequencies ranging from 20 to 90 GHz. The synchrotron and free-free maps are shown at K-band (23 GHz) while the dust map is shown at W-band (94 GHz). Averaged across the sky, but not including the galactic plane, the total foreground emission is minimized at frequencies near 65 GHz. However, the dust emission tends to be more tightly localized than the other components, so the foreground minimum is higher for selected small patches of the sky.

The presence of galactic magnetic fields means that synchrotron and dust emission will be partially polarized. The polarization of these signals varies for different regions of the sky (Kogut et al., 2007). For synchrotron emission, the polarization fraction has a value of $\sim 5\%$ near the galactic plane, but it climbs to 20% at higher latitudes. Similarly, the dust polarization fraction is smallest in the galactic plane and reaches $\sim 6\%$ near the poles.

4. The Rayleigh-Jeans spectrum has flux density proportional to ν^2 , so a source with flux density spectral index α will have a Rayleigh-Jeans temperature that depends on frequency as $T_{RJ} \propto \nu^\beta$, where $\beta = \alpha - 2$.

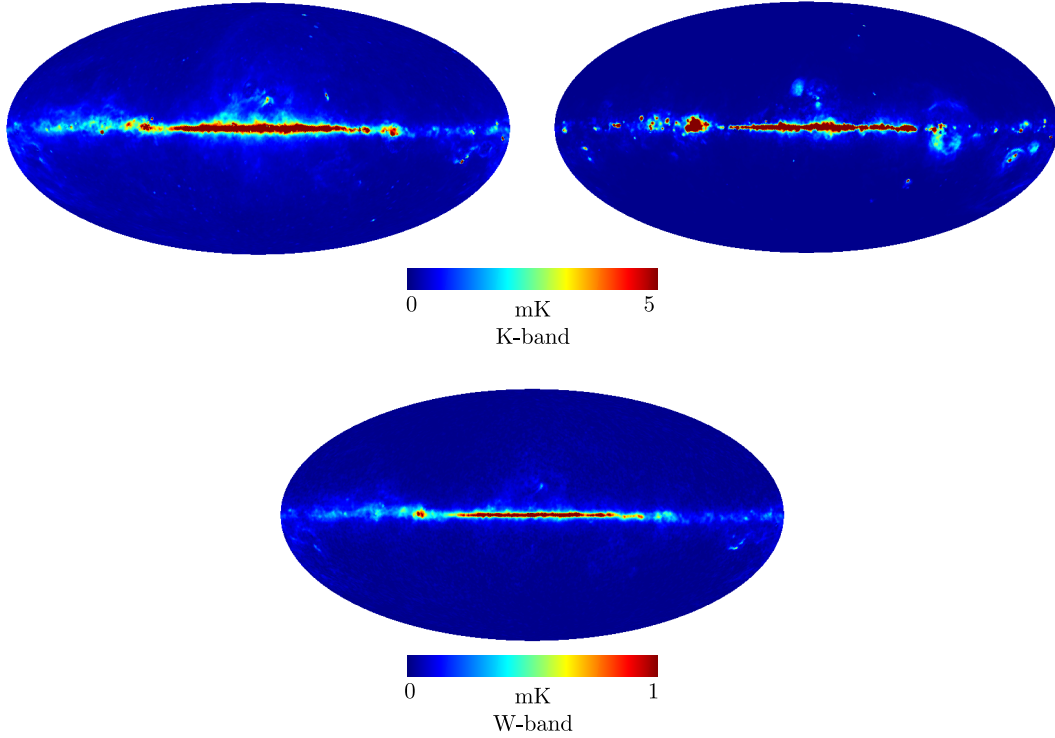


Figure 1.9: Foreground temperature maps for synchrotron (top left), free-free (top right), and dust emission (bottom) from the WMAP 7 year data release (Gold et al., 2010). To maximize the foreground signals, the synchrotron and free-free maps are shown for K-band (23 GHz) while the dust emission is shown for W-band (94 GHz). In fact, these maps were derived using all frequencies observed by WMAP, as well as template maps from other experiments, to differentiate between CMB and the various foreground types.

Besides the diffuse foregrounds described above, the sky is populated with point sources, both galactic and extra-galactic, that emit in the microwave range. Depending on the physics of emission, these sources can have rising or falling spectra. Extremely bright objects can be easily detected and masked out of maps for CMB analysis, but the contributed power from unresolved point sources can lead to biased results.

With significant uncertainties in the spectral index and polarization fraction of foregrounds, cleaning them from maps is a very challenging problem. Ground based polarization experiments, which only observe a small fraction of the sky, have managed to date to avoid foregrounds by observing in only the cleanest patches. As instrumental sensitivities improve,

this method will eventually become inadequate. Accurate foreground removal will require observations at a larger range of frequencies than is found in a typical ground based experiment, or else collaboration between multiple instruments operating at different frequencies to identify the various foreground components.

CHAPTER 2

CAPMAP AND QUIET

PIQUE (the Princeton IQU Experiment), CAPMAP (the Cosmic Anisotropy Polarization Mapper), and QUIET (the Q/U Imaging Experiment) represent three generations of CMB polarimeters using coherent detection techniques. At each step, advances in design and production allow for increasing numbers of detectors, which leads to improvements in sensitivity to the power spectrum proportional to N .

PIQUE, which predates the author’s involvement, consisted of single correlation polarimeters operating at 40 and 90 GHz (Q-band and W-band). PIQUE produced the tightest upper limits on the CMB EE power spectrum (Hedman et al., 2001) before it was detected by DASI. CAPMAP used a similar polarimeter design, but increased the number of detectors to 12 at W-band and 4 at Q-band. The key to this improvement was the use of High Electron Mobility Transistor (HEMT) amplifiers that are manufactured as Monolithic Microwave Integrated Circuits (MMICs). The HEMTs used in PIQUE provided similar performance but required painstaking hand assembly by a highly skilled technician while the CAPMAP MMICs could be produced in large batches using chip fabrication techniques. For QUIET, the polarimeter design was radically changed to create a highly integrated polarimeter module. Automated assembly techniques were used to build the modules and batch optimization allowed for entire arrays to be characterized simultaneously. With these advances, QUIET has been able to field two receivers consisting of 19 Q-band polarimeters and 90 W-band polarimeters, respectively. Looking forward, QUIET is designed to scale up to 1500 W-band polarimeters, plus additional Q-band and Ka-band (30 GHz) modules, with the goal of achieving limits on B mode polarization of $r < 0.01$.

This chapter provides an overview of the key features, goals, and scientific results of CAPMAP, in §2.1. A description of QUIET is included in §2.2, but this is kept very brief since the rest of the thesis focuses on QUIET in great detail.

2.1 The Cosmic Anisotropy Polarization Mapper

CAPMAP, a collaboration between groups at University of Chicago, Princeton University, University of Miami, and the Jet Propulsion Laboratory, operated for three observing seasons in the winters of 2003, 2004, and 2005. For the first season, only four polarimeters were built, all operating at W-band. First season results, including a 2σ detection of the EE spectrum, were published in Barkats et al. (2005a). Second season observations occurred concurrently with the effort to upgrade from the original four detectors to the full complement of sixteen. With attention and effort divided between deployment and observation, neither task was done well and CAPMAP second season results were never published, though a useful result to come out of season two was the verification that the ring scan (described below) was a viable strategy. In the third season, CAPMAP operated with all 16 polarimeters, at two frequencies. The third season results, published in Bischoff et al. (2008), featured measurements of the EE spectrum in seven bins of ℓ with a total detection significance of 11σ . A detailed description of CAPMAP and PIQUE instrumentation was provided in Barkats et al. (2005b). The theory of operation for CAPMAP polarimeters is also discussed here, in §3.7.

2.1.1 CAPMAP Site, Telescope, and Observation Strategy

The primary science goal of CAPMAP was to measure the CMB EE spectrum, which had not been detected until around the time of CAPMAP's initial deployment. This signal peaks for multipoles of $\ell \sim 1000$ (see Figure 1.6), so a large telescope was needed to achieve sufficiently good angular resolution. Fortunately, the Crawford Hill 7-meter Antenna (Chu et al., 1978) was unused and available for this purpose. The location holds significant historical interest, as the 7-meter antenna is located about 200 meters from the horn antenna that Penzias and Wilson used to first detect the CMB. A photograph of the Crawford Hill 7-meter antenna can be seen in Figure 2.1.



Figure 2.1: Photograph of the Crawford Hill 7-meter Antenna located at Lucent Technologies in Holmdel, New Jersey.

Located in coastal New Jersey¹, Crawford Hill is not an ideal site for microwave observations. Only during the winter months (late November to early April) does the level of atmospheric water vapor dip to acceptably low levels, and even then entire weeks are often lost due to clouds, rain, and snow. To counter these disadvantages, CAPMAP focused its sensitivity into a small region of sky, just 3° in diameter, centered on the North Celestial Pole (NCP). The telescope was scanned by moving in both azimuth and elevation to inscribe a 0.725° radius circle about the NCP every 21 seconds. This “ring scan” is quite unusual among CMB experiments, which typically avoid changes in elevation that introduce large atmospheric signals to the timestreams. For CAPMAP, the sinusoidal motion in elevation proved to be easily filtered out of the time ordered data, and the ring scan contributed highly uniform and symmetric coverage of the observing region, plus a constant monitor of sky temperature.

2.1.2 Calibration

The calibration requirements for CAPMAP and QUIET are quite similar, so many of the same techniques were used. QUIET calibration is covered in detail in §7, but some of the differences are highlighted here.

To target fine angular scales, CAPMAP required a much more accurate pointing model than is needed for QUIET. However, the task of deriving the model was simplified because the CAPMAP telescope has only two degrees of freedom (azimuth and elevation) rather than the three axes of motion provided by the QUIET mount. Also, the small beam provided by the 7-meter antenna means that CAPMAP was much more sensitive to point sources, so there were many more targets that could be used to constrain model parameters. The final CAPMAP pointing model had *rms* residual errors of ~ 0.5 arc-minutes and was found to be consistent across all three seasons of observation.

1. Latitude: $40^\circ 23' 31'' N$, longitude: $74^\circ 11' 10'' W$, elevation: 119 m

While CAPMAP did perform calibration observations of Tau A (Crab nebula supernova remnant, see §7.2) to measure polarized responsivities, the primary gain calibration involved a flat aluminum mirror (“chopper plate”), installed temporarily between the detectors and the secondary mirror. The chopper plate was nutated by 13° , so that the detectors observed the plate across a range of incidence angles. Reflection off the chopper plate induces a known polarized signal that depends on the resistivity of the aluminum, and the temperature difference between the plate and the sky. QUIET does not use this type of hardware calibrator during observations, but the CAPMAP chopper plate functions on similar principles to the rotating wire grid optimizer, described in §4.8.3.

2.1.3 Data Analysis

The CAPMAP time ordered data was filtered by removing five Fourier modes from each 21 second scan orbit. For data parametrized by the angle θ of the telescope relative to the NCP, the functional form of the subtracted modes is

$$M(\theta) = A_0 + A_1 \sin \theta + A_2 \cos \theta + A_3 \sin 2\theta + A_4 \cos 2\theta \quad (2.1)$$

The $\sin \theta$ term has large amplitude due to the ring scan elevation motion. The A_2 , A_3 , and A_4 terms were not significantly different from zero, on average, but their removal functioned as a high pass filter, effectively whitening the data. Data selection was performed based on the χ^2 goodness-of-fit between the data and equation (2.1) at timescales of 21 seconds and ~ 4 minutes. After removing the Fourier modes, the entire observing season was binned by θ and the resulting scan-synchronous structure was projected out as well.

Two independent pipelines were used for mapmaking and power spectrum estimation. Pipeline 1 created conventional maps, using the Healpix pixelization (Górski et al., 2005). Pipeline 2 used a map-making strategy based on the symmetry of the scan strategy. The

scan area was binned by right ascension and declination (resulting in smaller pixels near the NCP), but the analysis was performed in a half Fourier space obtained by taking the Fourier transform in the direction of right ascension. This change of coordinates resulted in a block diagonal noise covariance matrix, greatly speeding up map making operations. In contrast to the QUIET analysis described in §9, both CAPMAP pipelines calculated the full map covariance and used maximum likelihood techniques for power spectrum estimation.

CAPMAP relied on an extensive set of null tests to verify the quality of the selected data. For each null test, the full data set was divided into two subsets based on some criteria. Maps were made from each subset and then differenced to produce a null map that should be consistent with noise only. A total of 72 null tests were performed, using data divisions based on time, detector, or instrument performance. The null suite was used to test and tune the data selection algorithm, avoiding possible bias that might result from adjusting parameters of the analysis after viewing the actual science result. This blind analysis strategy was enthusiastically adopted for QUIET, as discussed in §10. Figure 2.2 shows a histogram of the χ^2 values obtained by comparing the null suite results to the hypothesis of zero signal. The distribution matches expectations, indicating that the underlying data is well understood.

The CAPMAP third season results, from pipeline 1, are shown in Figure 2.3. The 95% confidence upper limits on the BB spectrum can be seen in Figure 1.8. Other results from the third season include the EB power spectrum, which was consistent with zero, as well as power spectra calculated from a map constructed by differencing the two frequencies. That the Q/W difference map contained no signal provides evidence that CAPMAP results were not contaminated by foreground emission.




Figure 2.2: Distribution of χ^2 values obtained for the hypothesis that each null map contains zero sky signal. The red curve is a χ^2 distribution with one degree of freedom. The total χ^2 of the black histogram is 77.8 for 72 degrees of freedom, which corresponds to a probability to exceed of 30%. Figure from Bischoff et al. (2008).

2.2 The Q/U Imaging Experiment

With increased sensitivity made possible by large arrays of integrated polarimeters, the science goal for QUIET is shifted from the EE spectrum to the challenge of detecting or limiting the B -mode signature of inflation. This signal peaks at large angular scales ($\ell \sim 90$), making a 7-meter telescope unnecessary. In the quest for more sensitivity, a site was chosen for QUIET in the Atacama desert of northern Chile, which features observing conditions that are far superior to New Jersey. The QUIET collaboration is greatly expanded from CAPMAP, including groups at University of Chicago, Caltech, Columbia University, University of Manchester, Max-Planck-Institut für Radioastronomie in Bonn, University of Oslo, the Jet Propulsion Laboratory, University of Miami, KEK, Princeton University, Fermilab, University of Michigan, Oxford University, and Stanford University.

Like CAPMAP, QUIET observes at both Q and W bands, but now with the two frequencies split into separate receivers. The QUIET Q-band array, consisting of 19 detectors, was deployed first, observing from late October 2008 until June 2009. The QUIET W-band array, consisting of 90 detectors, was then installed and it continues to make observations as




Figure 2.3: The EE (left) and BB (right) power spectra reported for the third season of CAPMAP. Only results from pipeline 1 are shown.

of this writing. The measured total sensitivity for the QUIET arrays is 64 (57) $\mu K\sqrt{s}$ for Q (W) bands, which improves on CAPMAP by factors of 6.25 and 5.25, respectively (factors of 30–40 in integration time).

The remainder of this thesis focuses in detail on QUIET, particularly the completed Q-band observations. The design and operation of QUIET polarimeter modules is covered in §3. The non-module components that make up the rest of the instrument are described in §4. An overview of the Q-band observing season is found in §5, while §6 contains examples of various non-ideal behavior from the detectors and the solutions found for these problems. Calibration is addressed in §7. The remaining three chapters, §8, §9, and §10, describe the current state of the Q-band data analysis. Much progress has been made, while following a philosophy of blind analysis, but some issues, discussed in §10, remain.

CHAPTER 3

QUIET POLARIMETER MODULES

This chapter describes the inner workings of a QUIET module. The components that make up the module are covered in §3.1. The basic theory of operation is described in §3.2 while §3.3 goes into greater detail for the calculation of the polarimeter sensitivity. A calculation of the noise correlation between the various signal outputs of the module is found in §3.4. The benefits of fast modulation of the polarization signal are described in §3.5. An alternate configuration of QUIET modules, described in §3.6, results in paired devices with sensitivity to temperature differences, rather than polarization. Finally, §3.7 gives a comparison of the CAPMAP and QUIET polarimeters.

Correlation polarimeters detect linear polarization of an input microwave signal as a correlation between orthogonal components of the electric field. This technique depends on the existence of low noise amplifiers that are capable of amplifying GHz frequency signals with low noise while preserving their phase. In the past, correlation polarimeters have been successfully used for a variety of CMB polarization experiments, including CAPMAP (Barkats et al., 2005b), PIQUE (Hedman et al., 2002), COMPASS (Farese et al., 2003), and POLAR (Keating et al., 2003). Important advantages of correlation polarimeters, compared to direct differencing of the detected power along two orthogonal axes, are good immunity to systematic effects, particularly relative gain fluctuations between the arms of the detector, and the ability to quickly modulate the polarization signal.

QUIET uses a novel technique, which differs from normal correlation polarimeters in that the signal is detected by adding left circular and right circular components of the radiation, rather than multiplying linearly polarized components. The key breakthrough for QUIET is in miniaturization and integration of the detectors into a “polarimeter on a chip” (Gaier et al., 2003). All radio-frequency (RF) electrical components are contained in a compact package, which is coupled to waveguide for signal input and DC electronics for biasing and

readout. For comparison, a polarimeter from the CAPMAP experiment, operating at 90 GHz, is approximately 18 inches long, involving discrete microwave components and many waveguide pieces. The equivalent QUIET module is packaged in a 1 square inch housing that interfaces via twenty-seven pins.

A unique advantage of the QUIET correlation scheme is that a single module can simultaneously measure both Stokes parameters for linear polarization. A photon that enters a differencing polarimeter will be absorbed by either a device that is sensitive to X polarization or Y polarization, which means that linear polarization oriented at 45° is indistinguishable from unpolarized light. By measuring both Stokes parameters, QUIET modules effectively double the number of detectors for each feed and introduce additional systematics checks, since the measurements of \mathbf{Q} and \mathbf{U} share the same microwave circuitry.

Two QUIET detector arrays have been built and operated in the field, one with modules operating at 40 GHz, or Q-band, and the other with modules operating at 90 GHz, or W-band. While the components making up these modules are built to operate at different frequencies, they function in the same way. The relations derived in this chapter are applicable to both types of QUIET module, with a few differences between the two types noted.

3.1 Components of a QUIET Module

Figure 3.1 shows a schematic of a QUIET polarimeter module. An interior photograph, with the same parts labeled, is found in Figure 3.2. The module body consists of a brass housing with waveguide inputs on the top and pins on the bottom, evocative of an integrated circuit – hence the “polarimeter on a chip” description. In fact, the housing contains many separate circuits operating at microwave frequencies and connected by stripline (a flat electrical conductor sandwiched between two ground planes that operates like coax cable to propagate RF signals) that carries signals from the output of one device to the input of the next. A key

feature of the module schematic is the division of the signal into two legs for amplification. The signals in each leg are later combined to perform the correlation. The components of the module, and the input waveguide elements, are described here.

3.1.1 Feedhorn

The purpose of the feedhorn is to couple microwave radiation from free space into the input of the septum polarizer. The overall horn profile consists of a wide circular opening that tapers gradually down until reaching the waveguide diameter. Corrugations in the horn – alternating small and large diameter segments – improve the impedance match and reduce diffractive sidelobes.

With the large number of detectors in QUIET, it is necessary to use a feedhorn design that can be mass produced. The solution developed, referred to as a platelet array, is to assemble an entire array of feedhorns by drilling hole patterns in thin plates, stacking the plates up, and then fusing them together in a diffusion bonding process (Gundersen & Wollack, 2009). The resulting feedhorns have excellent performance, indicating that the mechanical precision of this process is sufficient.

The 40 GHz and 90 GHz feedhorns are scaled versions of each other, which means that the 40 GHz array is significantly longer. Both arrays are quite heavy, despite holes for lightweighting. Since the horns have to be cooled to cryogenic temperatures, their weight is one downside of the platelet array approach.

3.1.2 Septum Polarizer

The septum polarizer is a waveguide device that divides an input signal into its left circularly polarized and right circularly polarized components.

The input to the device is a square waveguide, which couples to the feedhorn via an interface plate, and supports all polarization states. Inside the polarizer, the waveguide

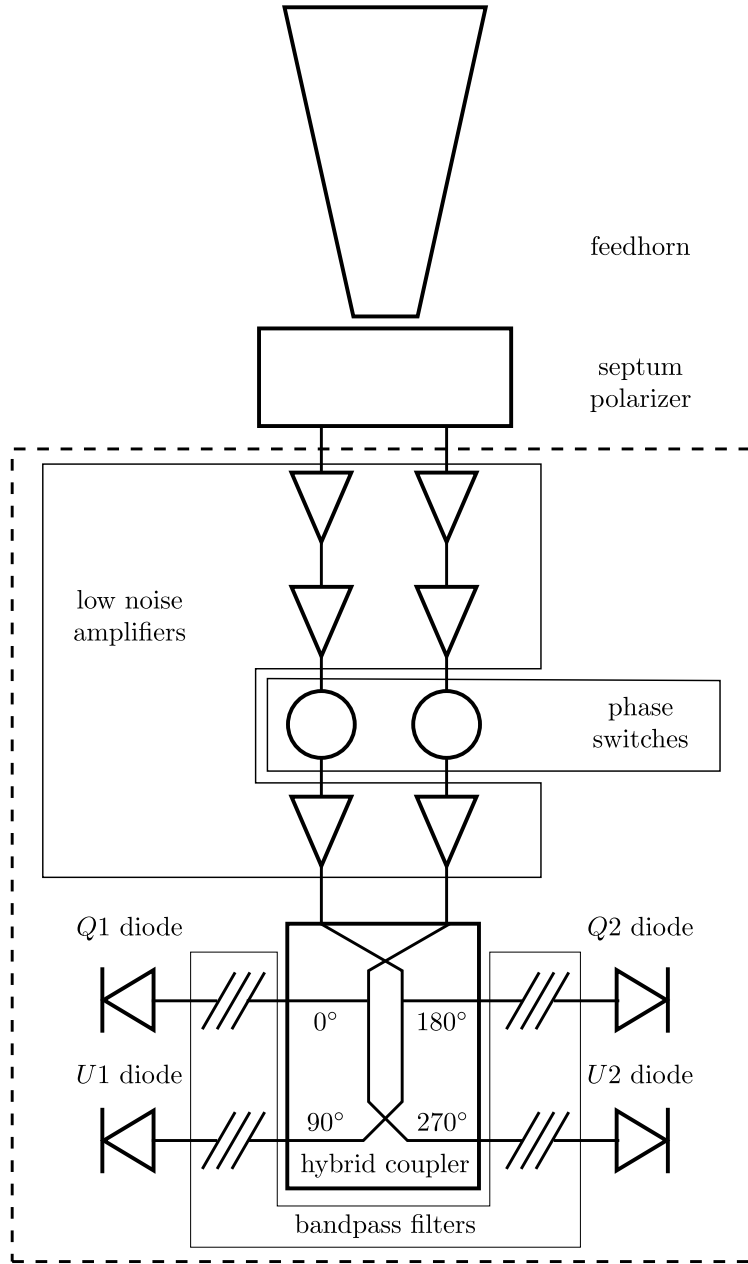


Figure 3.1: Schematic diagram of a QUIET polarimeter module. The components labeled on the diagram are described in Sections 3.1.1 through 3.1.7. The dashed line encloses the components which are physically integrated in the module package to form a “polarimeter on a chip”. Not shown in the schematic are bandpass filters between the phase switches and third amplification stage.




Figure 3.2: Photograph of a 90 GHz QUIET module, with the lid removed and components labeled. The signal path runs from the waveguide inputs on the left to the detector diodes on the right (unlike Figure 3.1, where the signal path runs from top to bottom). The component labeled “phase discriminator” is the same as the hybrid coupler, described in §3.1.5. Three horizontal rows of pins are visible; these pins pass through the bottom of the module to interface with the mating circuit board. Photo credit: JPL/California Institute of Technology.

is split by a thin metal septum with a stepped profile, based on the design described in Bornemann & Labay (1995), that induces a 90° phase shift along one axis. This phase shift transforms left and right circular polarizations into orthogonal linear polarizations, which are piped into separate rectangular waveguides. The net behavior of the device is that one output port emits a signal with amplitude proportional to the amplitude of the input left circular component while the other port emits a signal with amplitude proportional to the amplitude of the input right circular component. The polarizer is bolted to the top of the QUIET module, so these output ports are connected to the module inputs.

3.1.3 Low Noise Amplifiers

Low-noise amplifiers (LNAs) used for QUIET are InP High Electron Mobility Transistors (HEMTs) manufactured using Monolithic Microwave Integrated Circuit (MMIC) technology by Northrup-Grumman (Weinreb et al., 1999). Each of the six low-noise amplifiers included in the module schematic (Figure 3.1) consists of four transistors arranged in series to make an open-loop amplifier. The low-noise amplifiers can provide gain of ~ 25 dB over a bandwidth of $\sim 20\%$.

Since they are active components, the amplifiers will add noise to the microwave signal. In order to minimize this noise, the entire QUIET module is operated at a cryogenic temperature of 20K. The amplifier noise is measured in units of temperature referred to the device input, based on the linear relationship between RF power in a given bandwidth and blackbody temperature. Measured in this sense, the LNAs used for the 40 GHz QUIET modules have typical noise temperatures around 30K and amplifiers for the 90 GHz modules have noise temperatures around 60K. The amplifier gain is sufficiently large that only the first amplifier in each leg contributes significantly to the overall module noise figure.

To turn on the amplifiers, it is necessary to supply bias current to the transistor drains and bias voltage to the gates. The amplifier biases must be carefully tuned for each module,

to achieve sufficient gain, lowest possible noise, and a good phase match between the two legs of the module. These biases are provided by the LNA bias boards, described in Section 4.4.1. While the best performance would be achieved by independently adjusting the drain current and gate voltage to each transistor, such a large quantity of bias connections multiplied by the number of QUIET modules in an array becomes unfeasible. The compromise adopted for QUIET is to gang together the second and third stage amplifiers in a module leg, while providing more independent control for the first stage amplifiers, which are most important in determining the module sensitivity.

3.1.4 *Phase Switches*

Phase switches are located in each arm of the module, between the second and third amplifier stages. The phase switch consists of a fork that splits the stripline signal path into two branches. The branches merge back together, but one has additional length that adds 180° of phase to the signal. If an input signal splits equally between the two branches and is recombined, it will be strongly attenuated at the phase switch output due to destructive interference. However, in normal operation, the input signal passes through only one leg or the other. PIN diode switches¹ on each branch control the flow of signal.

Each phase switch device has four possible states, depending on the bias applied to the two PIN diodes. The $\uparrow\uparrow$ and $\downarrow\downarrow$ states both result in a strongly attenuated signal. The $\uparrow\downarrow$ and $\downarrow\uparrow$ states both transmit signal, but with a phase length that differs by 180° .

3.1.5 *Hybrid Coupler*

The QUIET hybrid coupler consists of a 180° coupler and a 90° coupler arranged in series. The amplified outputs of the two legs of the module each connect to an input of the 180°

1. PIN diodes, so called because of a region of Intrinsic semiconductor between the P-type and N-type regions, are commonly used for microwave switches. (Pojar, 2005, Chapter 10.3)

coupler. The two outputs of that coupler are the sum and difference of the input signals (Pozar, 2005, Chapter 7.8), which can also be thought of as sums of the inputs with an added relative phase of 0° or 180° . Each of these signals is divided by a 3dB power splitter, which diverts half the RF power to the $Q1$ and $Q2$ detector diodes. The remaining parts of each signal are used for the two inputs of the 90° coupler (Pozar, 2005, Chapter 7.5). This coupler again sums the input signals, but with relative phase shifts of 90° and 270° . The resulting outputs go to the $U1$ and $U2$ detector diodes.

3.1.6 Bandpass Filters

Bandpass filters are located between the hybrid coupler outputs and detector diodes and also between the phase switches and the third stage amplifiers (these filters are not include in Figure 3.1) to limit the range of frequencies detected. The waveguide components (feedhorn, septum polarizer, module inputs) feature a sharp low frequency cutoff and the performance of the low-noise amplifiers degrades outside of their optimal band. Bandpass filters are chosen to allow the maximum range of frequencies where the module has good sensitivity to polarization, but to remove signals outside that range which would only contribute extra noise. Also, attenuation of out of band power helps to stabilize the module, blocking standing waves that might develop between components.

3.1.7 Detector Diodes

After being filtered, each of the four outputs of the hybrid coupler is terminated on a detector diode. Detector diodes are used to measure signal power through rectification. The diode is able to convert RF power to DC by virtue of its non-linear I-V curve. The following derivation of the diode voltage sensitivity follows Pozar (2005), Chapter 10.3.

The detector diode I-V curve can be approximated by

$$I(V) = I_s \left(e^{\alpha V} - 1 \right) \quad (3.1)$$

where I_s is the saturation current and α depends on the temperature and properties of the diode. The voltage, V , applied to the diode consists of a large DC bias plus a small RF signal.

$$V = V_0 + V_1 \cos \omega t \quad (3.2)$$

Expanding equation (3.1) about V_0 , we obtain the following second order expression for I .

$$I(V) = I_0 + V_1 \alpha I_s e^{\alpha V_0} \cos \omega t + \frac{V_1^2}{2} \alpha^2 I_s e^{\alpha V_0} \cos^2 \omega t + \dots \quad (3.3)$$

$$= I_0 + \frac{1}{R_j} V_1 \cos \omega t + \frac{\alpha}{2R_j} V_1^2 \cos^2 \omega t + \dots \quad (3.4)$$

Here R_j is the junction resistance of the diode, given by dV/dI , evaluated at $V = V_0$.

The RF power deposited in the diode (which should be equal to the RF power incident, if there is a good impedance match) is given by the product of V and I (ignoring the DC bias terms):

$$P = \langle VI \rangle = \frac{V_1^2}{2R_j} \quad (3.5)$$

The rectified voltage measured by the diode is given by the junction resistance times the induced current. In the small signal limit, the rectified voltage is proportional to the square of the amplitude of the RF signal (which is temperature, in the Rayleigh-Jeans limit of the blackbody spectrum):

$$\Delta V = R_j (\langle \Delta I \rangle - I_0) = \alpha \frac{V_1^2}{4} \quad (3.6)$$

Taking a ratio of equations (3.5) and (3.6), we find the diode voltage sensitivity to RF power,

β_v :

$$\beta_v = \frac{\alpha R_j}{2} \quad (3.7)$$

The detector diodes used in QUIET modules are Agilent HSCH-9161 GaAs beamlead detector diodes. The diode specifications describe typical sensitivities of 1 mV/ μ W for zero bias operation at room temperature. In QUIET, these diodes are incorporated into the module body, which is cooled to cryogenic temperatures. The diode I-V curve changes dramatically when it is cooled. Diodes measured at 13K were found to have typical $\alpha^{-1} \sim 6$ mV and $I_s \sim 10^{-19}$ mA (Kapner (2006), Richardson (2007)). Using these values, we find that the resistance is astronomical at zero bias ($R_j \sim 10^{19}\Omega$) and it becomes necessary to apply bias to the diode in order to read it out. At typical bias currents of 30 μ A, R_j is decreased to only 200 Ω , but the diode sensitivity to RF signals, $\beta_v \sim 17$ mV/ μ W, is still much greater than the room temperature value.

The electrical connections to the detector diode are different for radio-frequency and DC signals. Capacitive couplings between the signal stripline and the diode anode and between the diode cathode and the module body create a path for RF signals to terminate on the detector diode. However, since both of these connections are capacitive, the DC voltage on the diode floats free from the module ground. Application of DC bias, and readout of the rectified voltages, occurs through a pair of pins for each diode, which connect to the preamplifier board described in Section 4.4.3.

3.2 Basic Theory of Operation

This section contains a simple description of how a QUIET module measures polarization. A more detailed description, including a calculation of the module sensitivity to polarized signals, is found in §3.3.

The inputs to legs A and B of the module are the circular polarization components,

E_L and E_R , of the input radiation. The amplifiers in each leg multiply the electric field amplitudes by gain factors, labeled g_A and g_B (here, we are lumping together the three amplifiers in each leg). Also, the phase switch in one leg of the module is modulated, alternately multiplying the signal in that leg by ± 1 . Equations (3.8) and (3.9) give the electric field amplitudes for each leg at the input of the hybrid coupler, where they are combined:

$$E_{legA} = g_A E_L \quad (3.8)$$

$$E_{legB} = \pm g_B E_R \quad (3.9)$$

The coupler adds the signals from the two legs with phase shifts of 0° , -90° , $+90^\circ$, and 180° , and the resulting signals are rectified by the $Q1$, $U1$, $U2$, and $Q2$ detector diodes. Calculating the rectified signal, and changing the electric field basis from left-right to X - Y , we find the following expressions.

$$|E_{Q1}|^2 = |E_{legA} + E_{legB}|^2 \quad (3.10)$$

$$= |g_A E_L \pm g_B E_R|^2 \quad (3.11)$$

$$= g_A^2 |E_L|^2 + g_B^2 |E_R|^2 \pm g_A g_B (E_L^* E_R + E_L E_R^*) \quad (3.12)$$

$$= (g_A^2 + g_B^2) (E_X^2 + E_Y^2) \pm 2g_A g_B (E_X^2 - E_Y^2) \quad (3.13)$$

$$= 2 (g_A^2 + g_B^2) \mathbf{I} \pm 4g_A g_B \mathbf{Q} \quad (3.14)$$

$$|E_{U1}|^2 = |E_{legA} + e^{-i\pi/2} E_{legB}|^2 \quad (3.15)$$

$$= 2 (g_A^2 + g_B^2) \mathbf{I} \mp 4g_A g_B \mathbf{U} \quad (3.16)$$

$$|E_{U2}|^2 = |E_{legA} + e^{i\pi/2} E_{legB}|^2 \quad (3.17)$$

$$= 2 (g_A^2 + g_B^2) \mathbf{I} \pm 4g_A g_B \mathbf{U} \quad (3.18)$$

$$|E_{Q2}|^2 = |E_{legA} + e^{i\pi} E_{legB}|^2 \quad (3.19)$$

$$= 2 \left(g_A^2 + g_B^2 \right) \mathbf{I} \mp 4g_A g_B \mathbf{Q} \quad (3.20)$$

The signals measured at the four detector diodes, shown in equations (3.14), (3.16), (3.18), and (3.20), all have the same first term. This term is proportional to the total intensity of the input signal, or Stokes parameter \mathbf{I} , and is multiplied by a gain factor $G_{av} = (g_A^2 + g_B^2)$. It is also important to note that the first term does not change sign when the phase switch state flips.

The second terms in those equations all do change sign with the phase switch state. The $Q1$ and $Q2$ diodes both measure $(E_X^2 - E_Y^2)/2$, or Stokes parameter \mathbf{Q} . The $U1$ and $U2$ diodes measure $E_X E_Y$, or Stokes parameter \mathbf{U} . For all four diodes, the gain factor multiplying these polarization signals is given by $G_{de} = 2g_A g_B$.

To measure polarized signals, the phase switch is toggled rapidly at 4 kHz. The voltages read out from each detector diode consist of an offset component proportional to \mathbf{I} and a switching component proportional to either \mathbf{Q} or \mathbf{U} . This signal is demodulated by differencing the phase switch \uparrow state from the phase switch \downarrow state to retain only the polarized signal. The benefits of 4 kHz demodulation are discussed in more detail in §3.5.

3.3 QUIET Module Radiometer Equation

The radiometer equation describes the sensitivity of a generic detector, defined as the minimum detectable signal, $\Delta\mathbf{T}$ in integration time² τ (Kraus, 1986).

$$\Delta\mathbf{T} = \frac{k(\mathbf{T}_{\text{rec}} + \mathbf{T}_{\text{load}})}{\sqrt{(\tau \times BW)}} \quad (3.21)$$

$$BW = \frac{(\int G d\nu)^2}{\int G^2 d\nu} \quad (3.22)$$

The receiver temperature, \mathbf{T}_{rec} , gives the noise, in temperature units, added by active devices (amplifiers) or lossy elements (waveguide, stripline, etc) in the detector. The load temperature, \mathbf{T}_{load} , gives the temperature of the observed source, which is the sum of the atmosphere (8–10K) and the 3K CMB. The bandwidth, BW , is a measure of the range of frequencies measured by the radiometer; the fact that sensitivity improves with larger bandwidth emphasizes that the radiometer equation is valid for observations of broad band sources, like the CMB. The factor k is a number of order unity that depends on the details of the radiometer.

This section contains a derivation of a radiometer equation for the specific case of a QUIET polarimeter module. The derivation is performed for the Q1 diode, but it is applicable for the other detector diodes as well. Included in this derivation is the impact of the relative gain and phase between the two legs of the QUIET module, which must be carefully optimized to achieve good performance.

3.3.1 Noise Distributions for RF Signals

All radio-frequency signals in the QUIET module will be treated as Gaussian noise with zero mean. Because of this, they can be described just by their variance, which can generally be

2. Note that the effective integration is less than the total observing time due to blanking that occurs at each phase switch transition, described in §4.5.

a function of frequency. The detector diodes will be treated as perfect square law detectors. In the case of non-Gaussian noise or non-square law detectors, the modules will exhibit non-linear effects, such as saturation.

Before starting the derivation, it is useful to establish some notation and calculate some expectation values that will come up frequently. For a single frequency, a RF signal can be written in terms of its electric field as

$$E = (a_0 + ia_1) e^{i\omega t} \quad (3.23)$$

The $e^{i\omega t}$, which gives a sinusoidal time dependence, can be safely dropped from further expressions, as the phase of the signal is contained in the complex amplitude.

If E is Gaussian noise with zero mean, then a_0 and a_1 are independent random numbers drawn from the probability distribution

$$P(z) = \frac{1}{\sqrt{2\pi\sigma_a^2}} \exp \left[-\frac{a^2}{2\sigma_a^2} \right] \quad (3.24)$$

Any odd power of E has zero expectation value. The expectation values for $|E|^2$ and $|E|^4$ are given by

$$\langle |E|^2 \rangle = \langle |a_0|^2 \rangle + \langle |a_1|^2 \rangle = 2\sigma_a^2 \quad (3.25)$$

$$\langle |E|^4 \rangle = \langle |a_0|^4 \rangle + 2\langle |a_0|^2 \rangle \langle |a_1|^2 \rangle + \langle |a_1|^4 \rangle = 8\sigma_a^4 \quad (3.26)$$

Also for future use, the expressions for the Stokes parameters, in terms of the electric

field in the X - Y basis, are included here.

$$\mathbf{I} = \frac{1}{2} \left(\langle |E_X|^2 \rangle + \langle |E_Y|^2 \rangle \right) = \sigma_X^2 + \sigma_Y^2 \quad (3.27)$$

$$\mathbf{Q} = \frac{1}{2} \left(\langle |E_X|^2 \rangle - \langle |E_Y|^2 \rangle \right) = \sigma_X^2 - \sigma_Y^2 \quad (3.28)$$

$$\mathbf{U} = \frac{1}{2} \left(\langle E_X^* E_Y \rangle + \langle E_X E_Y^* \rangle \right) = \langle x_0 y_0 \rangle + \langle x_1 y_1 \rangle \quad (3.29)$$

$$\mathbf{V} = \frac{1}{2} \left(\langle E_X^* E_Y \rangle - \langle E_X E_Y^* \rangle \right) = \langle x_1 y_0 \rangle - \langle x_0 y_1 \rangle \quad (3.30)$$

3.3.2 Detected Signal at the Q1 Diode

For a single frequency, ω , the input signal can be written in the instrumental X - Y basis as

$$E_{in} = E_X + E_Y \quad (3.31)$$

$$E_X = (x_0 + ix_1) e^{i\omega t} \quad (3.32)$$

$$E_Y = (y_0 + iy_1) e^{i\omega t} \quad (3.33)$$

The septum polarizer converts E_{in} into the circular polarization basis. The left and right circular components are amplified separately, which also introduces noise from the leg A and leg B amplifiers, written as E_A and E_B , respectively. Both E_A and E_B are scaled so that the noise levels are referred to the module input. For each leg, all amplifiers are lumped together and the combined gains are written as g_A and $g_B e^{i\phi}$, where ϕ is the difference in phase length between the two legs. The effect of the phase switch is represented as a \pm sign for the leg B signal. Combining these terms, the leg A and leg B signals input to the hybrid coupler can be written as

$$E_{legA} = g_A \left(\frac{1}{\sqrt{2}} E_X + i \frac{1}{\sqrt{2}} E_Y + E_A \right) \quad (3.34)$$

$$E_{legB} = \pm g_B e^{i\phi} \left(\frac{1}{\sqrt{2}} E_X - i \frac{1}{\sqrt{2}} E_Y + E_B \right) \quad (3.35)$$

While the $e^{i\omega t}$ terms have been dropped, it is important to keep in mind that equations (3.34) and (3.35) are expressions for the electric field at a single frequency. In particular, g_A , g_B , and ϕ can all vary considerably across the module response band.

The outputs of the first section of the hybrid coupler are the 0° and 180° combinations of the leg A and B signals. Half of the RF power from the 0° output is incident on the $Q1$ detector diode (the other half is used for the 90° and 270° combinations, which are rectified by the $U1$ and $U2$ diodes). The electric field signal at the $Q1$ diode is given by

$$E_{Q1} = \frac{1}{2\sqrt{2}} \left(g_A \pm g_B e^{i\phi} \right) E_X + i \frac{1}{2\sqrt{2}} \left(g_A \mp g_B e^{i\phi} \right) E_Y + \frac{1}{2} g_A E_A \pm \frac{1}{2} g_B e^{i\phi} E_B \quad (3.36)$$

The full expression for the rectified signal is

$$\begin{aligned} |E_{Q1}|^2 = & \frac{1}{8} \left(g_A^2 + g_B^2 \right) \left(|E_X|^2 + |E_Y|^2 \right) \pm \frac{1}{8} (2g_A g_B \cos \phi) \left(|E_X|^2 - |E_Y|^2 \right) \\ & \pm \frac{1}{8} (2g_A g_B \sin \phi) (2x_0 y_0 + 2x_1 y_1) + \frac{1}{8} \left(g_A^2 - g_B^2 \right) (2x_1 y_0 - 2x_0 y_1) \\ & + \frac{1}{4} g_A^2 |E_A|^2 + \frac{1}{4} g_B^2 |E_B|^2 \\ & + \frac{1}{4\sqrt{2}} \left(2g_A^2 \pm 2g_A g_B \cos \phi \right) (x_0 a_0 + x_1 a_1 - y_0 a_1 + y_1 a_0) \\ & + \frac{1}{4\sqrt{2}} \left(2g_B^2 \pm 2g_A g_B \cos \phi \right) (x_0 b_0 + x_1 b_1 + y_0 b_1 - y_1 b_0) \\ & \pm \frac{1}{4\sqrt{2}} (2g_A g_B \sin \phi) (x_0 a_1 - x_1 a_0 + y_0 a_0 + y_1 a_1) \\ & \pm \frac{1}{4\sqrt{2}} (2g_A g_B \sin \phi) (x_1 b_0 - x_0 b_1 + y_0 b_0 + y_1 b_1) \\ & \pm \frac{1}{4} (2g_A g_B \cos \phi) (a_0 b_0 + a_1 b_1) \pm \frac{1}{4} (2g_A g_B \sin \phi) (a_1 b_0 - a_0 b_1) \end{aligned} \quad (3.37)$$

However, only the first six terms of equation (3.37) have non-zero expectation value. Taking the expectation value, and applying the identities from equations (3.27) through

(3.30), the $Q1$ diode signal is given by

$$\begin{aligned} \langle |E_{Q1}|^2 \rangle = & \frac{1}{4} (g_A^2 + g_B^2) \mathbf{I} \pm \frac{1}{4} (2g_A g_B \cos \phi) \mathbf{Q} \pm \frac{1}{4} (2g_A g_B \sin \phi) \mathbf{U} \\ & + \frac{1}{4} (g_A^2 - g_B^2) \mathbf{V} + \frac{1}{4} g_A^2 \mathbf{T}_A + \frac{1}{4} g_B^2 \mathbf{T}_B \end{aligned} \quad (3.38)$$

The terms \mathbf{T}_A and \mathbf{T}_B are the noise temperature of amplifiers in the two legs of the module, equal to $|E_A|^2$ and $|E_B|^2$ respectively.

As was the case with equation (3.14), this expression has terms that are common mode or differential mode with respect to the two phase switch states. The common mode terms include the total power signal, \mathbf{I} , noise power from the amplifiers, and the circular polarization signal, \mathbf{V} , but only for the case where there is a mismatch between the leg A and B gains.

The Stokes parameters for linear polarization are found in the differential signal, which is obtained by demodulation. The effect of the phase difference, ϕ , is to rotate the axis along which the diode detects polarization. Notably, this means that a phase offset which is constant across the module frequency band does not adversely affect sensitivity (as it would for a CAPMAP polarimeter, see §3.7.2). However, there are usually enough degrees of freedom in selecting amplifier bias values to tune the mean value of ϕ near zero.

When measuring module performance, a statistic known as “isolation” is often used.

$$\text{isolation} [dB] = 10 \log_{10} \left(\frac{g_A^2 + g_B^2 - 2g_A g_B \cos \phi}{g_A^2 + g_B^2} \right) \quad (3.39)$$

Optimizing the isolation of a module involves balancing the gain and phase of the two legs, which improves the polarimeter performance.

3.3.3 Sensitivity of the Q1 Diode

The variance of the rectified Q1 signal is given by

$$\sigma_{Q1}^2 = \langle |E_{Q1}|^4 \rangle - \langle |E_{Q1}|^2 \rangle^2 \quad (3.40)$$

The terms in equation (3.40) can be calculated from equation (3.37). To simplify the result, we will assume that the input signal is unpolarized, so $\mathbf{Q} = \mathbf{U} = \mathbf{V} = 0$. This assumption is valid at the percent level or better for all atmospheric and astronomical signals. With this simplification, the expression for σ_{Q1}^2 is

$$\sigma_{Q1}^2 = \frac{1}{16} \left[(g_A^2 + g_B^2) \mathbf{I} + g_A^2 \mathbf{T}_A + g_B^2 \mathbf{T}_B \right]^2. \quad (3.41)$$

The factor of 1/16 accounts for the input power being divided in the septum polarizer and in the hybrid coupler. In practice, it is absorbed into the gain factors.

To calculate the detector diode sensitivity, it is necessary to integrate over the module frequency band. In an integration time τ , independent measurements are made at frequencies $\nu = \frac{n}{\tau}$, for $n = 0, 1, 2, \dots$. Replacing the sum over modes with an integral over frequencies, and defining lower and upper cutoff frequencies, ν_1 and ν_2 , that bracket the detector response, the expressions for the integrated signal and variance are

$$\text{integrated signal} = \tau \int_{\nu_1}^{\nu_2} \text{signal } d\nu \quad (3.42)$$

$$\text{integrated variance} = \tau \int_{\nu_1}^{\nu_2} \text{variance } d\nu \quad (3.43)$$

The noise is calculated by integrating the variance rather than the standard deviation, because each frequency mode provides an independent measurement of a broad band source.

Plugging the demodulated signal from equation (3.38) into equation (3.42) and the variance from equation (3.41) into equation (3.43), the minimum detectable polarization signal,

$\Delta \mathbf{Q}$, can be found by setting signal equal to noise. Here, we will assume that the input source has a blackbody spectrum, so the Stokes parameters are constant across the frequency range:

$$\frac{1}{4}\tau \Delta \mathbf{Q} \int_{\nu_1}^{\nu_2} (2g_A g_B \cos \phi) d\nu = \frac{1}{4}\sqrt{\tau} \left[\int_{\nu_1}^{\nu_2} \left((g_A^2 + g_B^2) \mathbf{I} + g_A^2 \mathbf{T}_A + g_B^2 \mathbf{T}_B \right)^2 d\nu \right]^{1/2} \quad (3.44)$$

$$\Delta \mathbf{Q} = \frac{1}{\sqrt{\tau}} \left[\frac{\int ((g_A^2 + g_B^2) \mathbf{I} + g_A^2 \mathbf{T}_A + g_B^2 \mathbf{T}_B)^2 d\nu}{\left(\int (2g_A g_B \cos \phi) d\nu \right)^2} \right]^{1/2} \quad (3.45)$$

If we assume that $\mathbf{T}_A = \mathbf{T}_B \equiv \mathbf{T}_{\text{rec}}$ and that \mathbf{T}_{rec} does not depend on frequency, then the expression simplifies further:

$$\Delta \mathbf{Q} = \frac{\mathbf{I} + \mathbf{T}_{\text{rec}}}{\sqrt{\tau}} \left[\frac{\int (g_A^2 + g_B^2)^2 d\nu}{\left(\int (2g_A g_B \cos \phi) d\nu \right)^2} \right]^{1/2}. \quad (3.46)$$

Comparing equations (3.21) and (3.46), we see that the proportionality factor k is equal to one and the bandwidth is given by

$$BW = \frac{\left(\int (2g_A g_B \cos \phi) d\nu \right)^2}{\int (g_A^2 + g_B^2)^2 d\nu} \quad (3.47)$$

Note that, for a given frequency response profile, bandwidth is maximized when $2g_A g_B \cos \phi = (g_A^2 + g_B^2)$, or $g_A = g_B$ and $\phi = 0$, which is the condition of good isolation. Equation (3.46) considers only the sensitivity to \mathbf{Q} polarization, but it can be seen from equation (3.38) that there is also some response to \mathbf{U} polarization when $\phi \neq 0$. In fact, an offset in ϕ that is constant across the detector band merely rotates the sensitivity axis of the detector diode and doesn't actually affect its sensitivity to general linear polarization.

3.4 Diode Noise Correlation

The hybrid coupler outputs four linear combinations of the signals from the two legs of the module, which results in noise correlation between the rectified signals. In practice, diode noise correlations are measured from the CMB data, as part of the noise model described in §7.3. Here, we derive the correlation coefficients, including effects due to gain imbalance or non-ideal behavior from the hybrid coupler.

For this calculation, we will assume that the input signal is unpolarized. It is split into uncorrelated left and right circular components with equal amplitudes. After applying amplifier gain and noise, the signals in each leg are still uncorrelated, but can now have different amplitude. At one particular frequency, ω , the RF signals in legs A and B are given, with all signals lumped together, by

$$E_A = (a_0 + \imath a_1)e^{i\omega t} \quad (3.48)$$

$$E_B = (b_0 + \imath b_1)e^{i\omega t} \quad (3.49)$$

Since there are no correlations, it is not even necessary to keep track of the relative phase state, switching at 4 kHz. As in §3.3, variables a_0 , a_1 , b_0 , and b_1 are Gaussian random numbers, drawn from distributions with zero mean variance σ_A^2 or σ_B^2 . If $\sigma_A^2 \neq \sigma_B^2$, then the RF power is not evenly balanced between the two arms of the module. This could be due to a difference in gain between the two legs or a difference in amplifier noise.

Between two diodes, X and Y , the correlation coefficient is given by equation (3.50).

$$C(X, Y) = \frac{\langle XY \rangle - \langle X \rangle \langle Y \rangle}{\sqrt{(\langle X^2 \rangle - \langle X \rangle^2)(\langle Y^2 \rangle - \langle Y \rangle^2)}} \quad (3.50)$$

3.4.1 Diode Noise Correlation Without Hybrid Imperfection

We start by calculating the noise correlations between diodes under the assumption that the hybrid coupler works perfectly (i.e. Q_1 sees $|E_A + E_B|^2$, Q_2 sees $|E_A - E_B|^2$, U_1 sees $|E_A - \imath E_B|^2$, and U_2 sees $|E_A + \imath E_B|^2$). Factors of $\frac{1}{2}$ that reflect division of power in the hybrid coupler are left out here, to simplify the equations; the correlation coefficient is a ratio, so any factors that are constant for all diodes are unimportant.

The rectified signals at diodes Q_1 and Q_2 are given by

$$|E_{Q1}|^2 = a_0^2 + a_1^2 + b_0^2 + b_1^2 + 2a_0b_0 + 2a_1b_1 \quad (3.51)$$

$$|E_{Q2}|^2 = a_0^2 + a_1^2 + b_0^2 + b_1^2 - 2a_0b_0 - 2a_1b_1 \quad (3.52)$$

$$\langle |E_{Q1}|^2 |E_{Q2}|^2 \rangle = 8 \left(\sigma_A^4 + \sigma_B^4 \right) \quad (3.53)$$

$$\langle |E_{Q1}|^2 \rangle = \langle |E_{Q2}|^2 \rangle = 2 \left(\sigma_A^2 + \sigma_B^2 \right) \quad (3.54)$$

$$\langle |E_{Q1}|^4 \rangle - \langle |E_{Q1}|^2 \rangle^2 = \langle |E_{Q2}|^4 \rangle - \langle |E_{Q2}|^2 \rangle^2 = 4 \left(\sigma_A^2 + \sigma_B^2 \right)^2 \quad (3.55)$$

Plugging these values into equation (3.50), we find the following expression.

$$C \left(|E_{Q1}|^2, |E_{Q2}|^2 \right) = \left[\frac{\sigma_A^2 - \sigma_B^2}{\sigma_A^2 + \sigma_B^2} \right]^2 \quad (3.56)$$

From equation (3.56), we see that the white noise is uncorrelated between the Q_1 and Q_2 diodes only if the power level into the hybrid coupler from each leg is the same. In the limit that one leg is turned off completely, the correlation goes to 1. In this limit, the module no longer operates as a polarimeter. Instead, all four diodes see the same signal proportional to total power.

The $U1$ diode signal is given by

$$|E_{U1}|^2 = a_0^2 + a_1^2 + b_0^2 + b_1^2 + 2a_0b_1 - 2a_1b_0 \quad (3.57)$$

$$\langle |E_{Q1}|^2 |E_{U1}|^2 \rangle = 8 \left(\sigma_A^4 + \sigma_A^2 \sigma_B^2 + \sigma_B^4 \right) \quad (3.58)$$

$$\langle |E_{U1}|^2 \rangle = \langle |E_{Q1}|^2 \rangle \quad (3.59)$$

$$\langle |E_{U1}|^4 \rangle - \langle |E_{U1}|^2 \rangle^2 = \langle |E_{Q1}|^4 \rangle - \langle |E_{Q1}|^2 \rangle^2 \quad (3.60)$$

The correlation between $Q1$ and $U1$ is

$$C \left(|E_{Q1}|^2, |E_{U1}|^2 \right) = \frac{\sigma_a^4 + \sigma_b^4}{(\sigma_a^2 + \sigma_b^2)^2} \quad (3.61)$$

Unlike the $Q1$ – $Q2$ case, the correlation between $Q1$ and $U1$ is non-zero, even if $\sigma_a^2 = \sigma_b^2$. In fact, the value starts at 0.5 and increases if the power in the two legs is imbalanced.

Figure 3.3 shows the quadratic dependence of the noise correlation on power imbalance, for the parametrization $\sigma_A^2 = 1 + \epsilon$ and $\sigma_B^2 = 1 - \epsilon$. While not every diode combination is worked out here, it can be shown that $C \left(|E_{U1}|^2, |E_{U2}|^2 \right) = C \left(|E_{Q1}|^2, |E_{Q2}|^2 \right)$ and all correlations between Q and U diodes are given by equation (3.61).

3.4.2 Diode Noise Correlation Including Coupler Imperfections

The white noise correlation between detector diodes in a module will also change if the hybrid coupler doesn't combine the input signals with exactly the right phase. To get a general result, we will define ϕ as phase of the $Q2$ output of the 180° coupler, and ψ and θ as the phases of the $U1$ and $U2$ outputs of the 90° coupler. The nominal values of ϕ , ψ , and θ are 180° , -90° , and 90° , respectively. The freedom to choose an overall phase means that we do not lose any generality by assuming that the $Q1$ output of the 180° coupler is exactly




Figure 3.3: Diode noise correlation as a function of power imbalance in the two legs of a QUIET module. The black curve is for correlations between diodes of the same type ($Q1-Q2$ or $U1-U2$). The red curve gives the correlation between a Q diode and a U diode. The imbalance is parametrized by $\sigma_A^2 = 1 + \epsilon$ and $\sigma_B^2 = 1 - \epsilon$.

0, so the $Q1$ signal is still given by equation (3.51).

With the added complication from the hybrid coupler imperfection, the $Q2$ signal is given by

$$|E_{Q2}|^2 = \left| (a_0 + ia_1) + e^{i\phi} (b_0 + ib_1) \right|^2 \quad (3.62)$$

$$= a_0^2 + a_1^2 + b_0^2 + b_1^2 + 2(a_0b_0 + a_1b_1)\cos\phi + 2(a_1b_0 - a_0b_1)\sin\phi \quad (3.63)$$

$$\langle |E_{Q2}|^2 \rangle = 2(\sigma_A^2 + \sigma_B^2) \quad (3.64)$$

$$\langle |E_{Q2}|^4 \rangle - \langle |E_{Q2}|^2 \rangle^2 = 4(\sigma_A^2 + \sigma_B^2)^2 \quad (3.65)$$

Plugging equations (3.51), (3.54), (3.55), and (3.63) through (3.65) into the correlation formula yields the following expression.

$$C(|E_{Q1}|^2, |E_{Q2}|^2) = \frac{\sigma_A^4 + 2\sigma_A^2\sigma_B^2\cos\phi + \sigma_B^4}{(\sigma_A^2 + \sigma_B^2)^2} \quad (3.66)$$

If we set $\phi = 180^\circ$, then we recover equation (3.56). For all other values of ϕ , the correlation is larger than it would be from the power imbalance alone. Note that it is impossible to obtain a negative correlation value from equation (3.66).

The U diodes are more difficult to understand, since any effect that induces correlation between the Q diodes also implies that the inputs to the 90° coupler are already correlated at some level. The signal at the $U1$ diode is given in equation (3.68), where ψ is the phase shift of the $U1$ output of the 90° coupler and is nominally equal to -90° . The $U2$ signal is given by the same formula, but with ψ replaced by θ , which has a nominal value of $+90^\circ$.

$$|E_{U1}|^2 = \frac{1}{2} \left| a_0 + ia_1 + b_0 + ib_1 + e^{i\psi} \left(a_0 + ia_1 + e^{i\phi} (b_0 + ib_1) \right) \right|^2 \quad (3.67)$$

$$\begin{aligned} &= a_0^2 (1 + \cos \psi) + a_1^2 (1 + \cos \psi) \\ &\quad + b_0^2 (1 + \cos (\phi + \psi)) + b_1^2 (1 + \cos (\phi + \psi)) \\ &\quad + (a_0 b_0 + a_1 b_1) (1 + \cos \phi + \cos \psi + \cos (\phi + \psi)) \\ &\quad + (a_0 b_1 - a_1 b_0) (\sin \psi - \sin \phi - \sin (\phi + \psi)) \end{aligned} \quad (3.68)$$

$$\langle |E_{U1}|^2 \rangle = 2 \left[\sigma_A^2 (1 + \cos \psi) + \sigma_B^2 (1 + \cos (\phi + \psi)) \right] \quad (3.69)$$

$$\langle |E_{U1}|^4 \rangle - \langle |E_{U1}|^2 \rangle^2 = 4 \left[\sigma_A^2 (1 + \cos \psi) + \sigma_B^2 (1 + \cos (\phi + \psi)) \right]^2 \quad (3.70)$$

The correlation between $Q1$ and $U1$ is shown in equation (3.71). The correlation between $Q1$ and $U2$ can be obtained by replacing ψ with θ . It can be shown that the correlations between the U diodes and $Q2$ are given by equation (3.71) as well.

$$\begin{aligned} C(|E_{Q1}|^2, |E_{U1}|^2) &= \frac{1}{Z} \left[\sigma_A^4 (1 + \cos \psi) + \sigma_B^4 (1 + \cos (\phi + \psi)) \right. \\ &\quad \left. + \sigma_A^2 \sigma_B^2 (1 + \cos \phi + \cos \psi + \cos (\phi + \psi)) \right] \end{aligned} \quad (3.71)$$

where

$$Z = \sigma_A^4 (1 + \cos \psi) + \sigma_B^4 (1 + \cos (\phi + \psi)) + \sigma_A^2 \sigma_B^2 (2 + \cos \psi + \cos (\phi + \psi)) \quad (3.72)$$

Figure 3.4 shows $C(|E_{Q1}|^2, |E_{U1}|^2)$ as a function of power imbalance for several values of ϕ and ψ .




Figure 3.4: Noise correlation between Q and U diodes as a function of power imbalance, parametrized as $\sigma_A^2 = 1 + \epsilon$ and $\sigma_B^2 = 1 - \epsilon$. For the left-hand panel, $\psi = -\frac{\pi}{2}$ and $\phi = \pi, \frac{7\pi}{8}, \frac{6\pi}{8}$ for the black, red, and blue curves, respectively. For the right-hand panel, $\phi = \pi$ and $\psi = -\frac{\pi}{2}, -\frac{5\pi}{8}, -\frac{6\pi}{8}$ for the black, red, and blue curves, respectively.

The correlation between U diodes is the most complicated. One might guess that it would be the same as the $Q1$ – $Q2$ correlation, but with ϕ replaced with the difference of ψ and θ . However, if there is any power imbalance or imperfection in the 180° hybrid, then the input signals to the 90° hybrid are correlated and it is necessary to keep track of all three phase angles, ϕ , ψ , and θ .

The full expression for the correlation between U diodes is given by

$$\begin{aligned}
C(|E_{U1}|^2, |E_{U2}|^2) = \frac{1}{Z} & \left[\sigma_A^4 (1 + \cos \psi) (1 + \cos \theta) \right. \\
& + \sigma_B^4 (1 + \cos(\phi + \psi)) (1 + \cos(\phi + \theta)) \\
& + \sigma_A^2 \sigma_B^2 (1 + \cos \phi + \cos \psi + \cos \theta + \cos(\phi + \psi) \\
& + \cos(\phi + \theta) + \cos(\psi - \theta) \\
& \left. + \cos(\phi + \psi + \theta)) \right] \quad (3.73)
\end{aligned}$$

where

$$\begin{aligned}
Z = & \left[\sigma_A^2 (1 + \cos \psi) + \sigma_B^2 (1 + \cos(\phi + \psi)) \right] \\
& \times \left[\sigma_A^2 (1 + \cos \theta) + \sigma_B^2 (1 + \cos(\phi + \theta)) \right] \quad (3.74)
\end{aligned}$$

If the power is balanced in the two legs of the module and the 180° coupler has a phase shift of exactly 180° , then equation (3.73) reduces to

$$C(|E_{U1}|^2, |E_{U2}|^2) = \frac{1}{2} (1 + \cos(\psi - \theta)) \quad (3.75)$$

On the other hand, if the 90° coupler is working perfectly and the only noise correlations come from further upstream, then the correlation coefficient is

$$C(|E_{U1}|^2, |E_{U2}|^2) = \frac{(\sigma_A^2 + \sigma_B^2 \cos \phi)^2}{\sigma_A^4 + \sigma_B^4 \cos^2 \phi + 2\sigma_A^2 \sigma_B^2} \quad (3.76)$$

Figure 3.5 shows the rich dependence of the $U1$ – $U2$ noise correlation on power imbalance and phases of the couplers.




Figure 3.5: Noise correlation between $U1$ and $U2$ diodes as a function of power imbalance, parametrized as $\sigma_A^2 = 1 + \epsilon$ and $\sigma_B^2 = 1 - \epsilon$. For the upper-right panel, $\psi = -\frac{\pi}{2}$, $\theta = +\frac{\pi}{2}$ and $\phi = \pi, \frac{7\pi}{8}, \frac{6\pi}{8}$ for the black, red, and blue curves, respectively. For the lower-left panel, $\phi = \pi$, $(\theta - \psi) = \pi$ and $(\theta + \psi) = 0, \frac{\pi}{4}, \frac{\pi}{2}$ for the black, red, and blue curves, respectively. For the lower-right panel, $\phi = \pi$, $(\theta + \psi) = 0$, and $(\theta - \psi) = \pi, \frac{9\pi}{8}, \frac{10\pi}{8}$ for the black, red, and blue curves, respectively.

3.4.3 Correlation of Difference Streams

While calculating the Q – U correlation, we found that

$$\langle |Q_1|^2 |U_i|^2 \rangle = \langle |Q_2|^2 |U_i|^2 \rangle \quad (3.77)$$

and

$$\langle |Q_1|^2 \rangle = \langle |Q_2|^2 \rangle \quad (3.78)$$

for arbitrary ϕ , ψ , and θ . Applying these identities to the correlation coefficient of the Q diode difference stream with the U diode difference stream, we obtain the result

$$\begin{aligned} C \left(|E_{Q1}|^2 - |E_{Q2}|^2, |E_{U1}|^2 - |E_{U2}|^2 \right) \\ \propto \langle (|E_{Q1}|^2 - |E_{Q2}|^2) (|E_{U1}|^2 - |E_{U2}|^2) \rangle \end{aligned} \quad (3.79)$$

$$\begin{aligned} &= \langle |E_{Q1}|^2 - |E_{Q2}|^2 \rangle \langle |E_{U1}|^2 - |E_{U2}|^2 \rangle \\ &\propto \langle |E_{Q1}|^2 |E_{U1}|^2 \rangle - \langle |E_{Q2}|^2 |E_{U1}|^2 \rangle \\ &\quad - \langle |E_{Q1}|^2 |E_{U2}|^2 \rangle + \langle |E_{Q2}|^2 |E_{U2}|^2 \rangle \end{aligned} \quad (3.80)$$

$$\begin{aligned} &= \left(\langle |E_{Q1}|^2 \rangle - \langle |E_{Q2}|^2 \rangle \right) \langle |E_{U1}|^2 \rangle \\ &\quad + \left(\langle |E_{Q1}|^2 \rangle - \langle |Q_2|^2 \rangle \right) \langle |E_{U2}|^2 \rangle \\ &= 0 \end{aligned} \quad (3.81)$$

Regardless of power imbalance or coupler imperfections, the module combined measurements of Q and U are totally uncorrelated. This cancellation is imperfect in an actual module, however, due to differences in bandpass between detector diodes.

3.5 Phase Switching and Demodulation

While §3.2 shows that phase switching is important for measuring polarization, because it toggles the detector diodes between measurements along the X and Y axes, it did not address the reason why the switching speed is so high. The amplifiers used in QUIET are prone to gain drifts with a $1/f$ type spectrum. Under normal operating conditions, the power output by the LNAs is dominated by this noise for frequencies less than a few hundred Hz for 40 GHz amplifiers and up to 1 kHz for 90 GHz amplifiers.

By demodulating at 4 kHz, drifts in the total power level at longer timescales are differenced away. Residual $1/f$ noise in the demodulated signal comes from gain fluctuations multiplying small offsets in the demodulated signal due to $\mathbf{I} \rightarrow \mathbf{Q}$ leakage. Since this leakage is typically 1% or less, $1/f$ noise power in the demodulated data is suppressed by a factor of 10^4 or more. Figure 3.6 illustrates the vast difference in noise between the averaged and demodulated data.

The PIN diodes used in the phase switches are capable of extremely fast (10 ns or less) switching times between states. In practice, there are obstacles to increasing the demodulation rate much higher than 4 kHz (and no obvious benefits). The PIN diodes are fragile and can be broken by large electrical transients. Protection circuitry designed to protect the diodes limits the speed of the transitions. During the transition, the phase switch is in an undetermined state, typically blocking all RF power in the module leg. The transition periods must be blanked out during demodulation and this blanking does reduce the fraction of time that the detector is integrating signal. In operation, the QUIET Q-band array blanks out $17.5 \mu\text{s}$ at each transition, which adds up to 14% of total time.




Figure 3.6: Noise power spectra for averaged (black) and demodulated (red) data, taken from CMB observations with the $Q1$ diode of Q -band module RQ-09. The blue lines, drawn to guide the eye, show that phase switching suppresses the $1/f$ noise power by a factor of nearly 10^5 .

3.5.1 Phase Switch Imbalance and Double Demodulation

Ideally, switching the phase switch state should change the phase of the signal in that arm by exactly 180° but leave the amplitude unchanged. In practice, the phase difference might not be exactly 180° , which would adversely affect the module isolation. A more serious problem occurs if the phase switch transmission amplitude is different in the two states.

Suppose that we are modulating the phase switch in leg B, which has transmission amplitude of 1 in its up (\uparrow) state but transmission amplitude of $(1 - \epsilon)$ in its down (\downarrow) state. The phase switch in leg A is in the \uparrow state, with transmission amplitude 1. Following the simple signal derivation of Section 3.2, the signal at the Q1 diode is given by equations (3.82) and (3.83).

$$Q1_{\uparrow\uparrow} = 2 \left(g_A^2 + g_B^2 \right) \mathbf{I} + 4g_A g_B \mathbf{Q} \quad (3.82)$$

$$Q1_{\uparrow\downarrow} = 2 \left(g_A^2 + (1 - \epsilon)^2 g_B^2 \right) \mathbf{I} - 4(1 - \epsilon) g_A g_B \mathbf{Q} \quad (3.83)$$

$$\frac{1}{2} (Q1_{\uparrow\uparrow} - Q1_{\uparrow\downarrow}) = g_B^2 (\epsilon^2 - 2\epsilon) \mathbf{I} + 4g_A g_B (1 - \epsilon/2) \mathbf{Q} \quad (3.84)$$

The demodulated signal, shown in equation (3.84), has a term proportional to Stokes parameter \mathbf{Q} , with a gain term that is slightly modified by the phase switch imbalance. More critical is the appearance of a term proportional to Stokes parameter \mathbf{I} , leaking the temperature signal into the polarization measurement.

This $\mathbf{I} \rightarrow \mathbf{Q}$ leakage can be eliminated, however, by modulating the leg A phase switch – a process referred to as “double demodulation”. Equation (3.85) describes a similar signal to equation (3.84), measured by demodulating across the leg B phase switch states, but now with the leg A phase switch in the \downarrow state. An additional phase switch imbalance term, η , has been introduced to account for any imbalance in the leg A phase switch.

$$\frac{1}{2} (Q1_{\downarrow\uparrow} - Q1_{\downarrow\downarrow}) = g_B^2 (\epsilon^2 - 2\epsilon) \mathbf{I} + 4g_A g_B (1 - \epsilon/2) (1 - \eta) \mathbf{Q} \quad (3.85)$$

The harmful leakage term is common mode to equations (3.84) and (3.85), while the polarized signal is differential. By switching the leg A phase switch at a slower rate (only 50 Hz, compared to 4 kHz for the leg B phase switch) and performing an extra demodulation step, the problem of $\mathbf{I} \rightarrow \mathbf{Q}$ leakage due to phase switch imbalance is effectively eliminated. QUIET modules still do suffer from $\mathbf{I} \rightarrow \mathbf{Q}$ leakage, with the dominant contribution due to the septum polarizers, as discussed in §6.2.

3.6 Temperature Differencing Module Pairs

While averaging (rather than demodulating) the detector diode signals provides some sensitivity to \mathbf{I} , large $1/f$ noise (see Figure 3.6) make these data useless for precision measurements. QUIET is primarily a polarization experiment, but some sensitivity to the CMB temperature spectrum is preferable. For example, while it would be possible to cross-correlate QUIET polarization maps with CMB temperature maps from WMAP to extract a \mathbf{TE} spectrum, there are numerous practical difficulties, such as different beam sizes and scan strategies.

An alternate detector configuration, known as “TT modules”, uses a pair of QUIET modules with a different waveguide input section to give QUIET the ability to make precise differential measurements of the CMB temperature field. Out of the 19 element Q-band array, two modules are configured for temperature measurements; the 91 elements W-band array contains six such modules. The CMB temperature data is not covered in this thesis, but a simple explanation of the TT module operation is included here.

For a pair of TT modules, the septum polarizers are replaced by a differencing assembly, which couples signals from two feedhorns into both modules. The differencing assembly can be broken down into two sections. First, the signal from each feedhorn passes through a separate ortho-mode transducer (OMT), which divides the signal into its X and Y linear polarization components. These OMTs are identical to those used in CAPMAP polarimeters

(see §3.7). Next, the X polarization from feedhorn 1 and the Y polarization from feedhorn 2 are used as inputs to a magic-T, which is a waveguide device equivalent to a 180° coupler (Pozar, 2005, Chapter 7.8). The outputs of the magic-T, $E_{X1} + E_{Y2}$ and $E_{X1} - E_{Y2}$, are inputs to one module of the pair. The other module is also attached to a magic-T and sees input signals $E_{X2} + E_{Y1}$ and $E_{X2} - E_{Y1}$.

Applying amplifier gain and phase switching, the rectified signal on the $Q1$ diode of the first module is given by

$$|E_{Q1,1}|^2 = |g_{A,1} (E_{X,1} + E_{Y,2}) \pm g_{B,1} (E_{X,1} - E_{Y,2})|^2 \quad (3.86)$$

$$\begin{aligned} &= \left(g_{A,1}^2 + g_{B,1}^2 \right) \left(|E_{X,1}|^2 + |E_{Y,2}|^2 \right) \\ &\quad \pm 2g_{A,1}g_{B,1} \left(|E_{X,1}|^2 - |E_{Y,2}|^2 \right) \end{aligned} \quad (3.87)$$

The first term in equation (3.87) gives the average temperature between the separate spots on the sky pointed to by feedhorns 1 and 2. The second term gives the temperature difference between the two spots³. Because the temperature difference term is obtained by demodulation, it benefits from greatly improved $1/f$ noise, compared to average data, and can be used for CMB temperature measurements.

3.7 Comparison Between CAPMAP and QUIET Polarimeters

While QUIET modules contain many of the same components as CAPMAP polarimeters, particularly the low noise amplifiers and phase switches, there are substantial differences. Some of these differences, such as the ability to simultaneously measure both Stokes parameters, represent improvements in the QUIET modules. Many of the differences are compromises that are necessary in order to make QUIET modules scalable to large arrays.

3. While the terms in equation (3.87) involve only single linear polarizations, they can still be thought of as the temperatures for their respective spots on the sky, because the linear polarization fraction of the CMB is $\leq 1\%$ and the linear polarization fraction of the atmosphere is zero

3.7.1 CAPMAP Polarimeter Overview

A schematic diagram of a CAPMAP polarimeter is shown in Figure 3.7. The two dashed boxes on the diagram indicate the cryogenic (20K) and room temperature portions of the cryostat housing the detector. The remaining components at the bottom of the diagram are housed in the “IF box”, which mounts on the outside.

Signals are coupled into the polarimeter via a lens (not shown) and a feedhorn with similar optical properties to the QUIET feedhorns, but fabricated individually, instead of as part of a platelet array. Instead of the septum polarizer used in QUIET, the CAPMAP polarimeter features an ortho-mode transducer (OMT), which divides the input signal into X and Y linearly polarized components. These signals are piped using rectangular waveguide to low noise amplifiers of the same type as used in QUIET.

After amplification, the signals in each arm are converted from the input frequencies (either 40 or 90 GHz) to intermediate frequencies (IF) in the 2–18 GHz range. This down-conversion is performed using a mixer in each arm of the polarimeter and a local oscillator running at 30.5 or 82 GHz. The same local oscillator feeds both mixers, through waveguide sections of identical length, so the downconversion preserves the relative phase of the two arms. However, one leg of the local oscillator signal passes through a phase switch, which can modulate this phase by 180° .

The IF signals are carried on coaxial cable (semi-rigid SMA cable) through hermetic feedthroughs in the cryostat wall. Outside, there is an additional stage of amplification, and some of the signal power is diverted and rectified by a detector diode, to serve as a noisy monitor of the unpolarized signal.

The remaining signal passes into a triplexer, which divides it into three frequency sub-bands (or a duplexer, with two sub-bands, for a 40 GHz CAPMAP detector). Since the downconversion to intermediate frequency simply shifts frequencies by the local oscillator frequency, these sub-bands are equivalent to divisions of the polarimeter’s RF bandwidth.

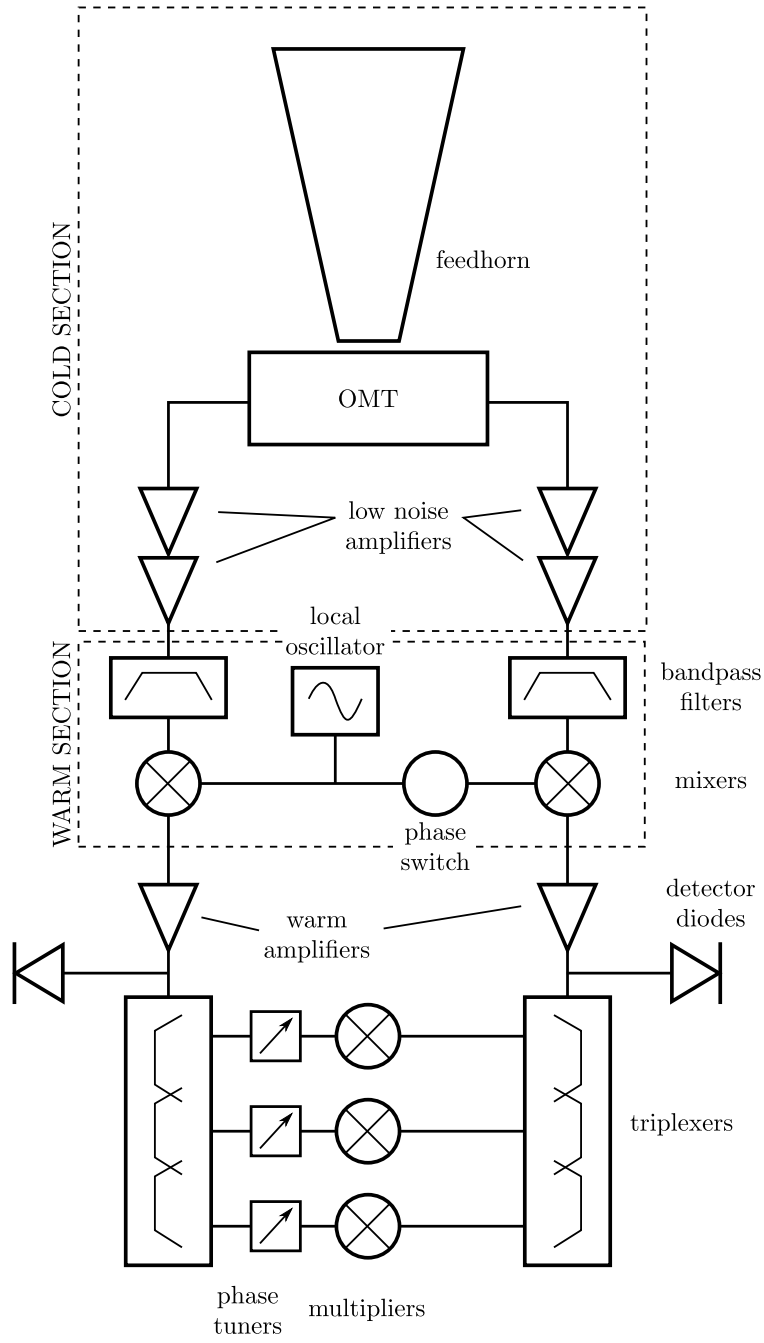


Figure 3.7: Schematic diagram of a CAPMAP polarimeter. The dashed boxes indicate the parts which are in the cryogenic (20K) and room temperature portions of the cryostat.

Finally, the two arms of the polarimeter are correlated by multipliers (one for each sub-band). The multiplier output is demodulated to measure the polarization signal.

3.7.2 Comparison of Correlation Schemes

The correlation method used in QUIET modules is described in §3.2. For CAPMAP polarimeters, the septum polarizer is replaced with an X – Y OMT, so the signals in the two legs of the device are proportional to E_X and E_Y , respectively. The whole chain of amplifiers, filters, and the mixer can be lumped together into a single gain term for each leg, g_A and g_B , as well as the relative phase, ϕ , between the two legs.

The multipliers for each sub-band output the product the two legs, which will have a DC component if they are correlated. With the phase switch toggling the sign of leg B, the multiplier output is shown in equation (3.89):

$$S_{\pm} = \Re\left(\pm g_A g_B e^{i\phi} E_X E_Y\right) \quad (3.88)$$

$$= \pm g_A g_B \cos \phi \mathbf{U} \quad (3.89)$$

Demodulating this signal gives a measure of Stokes parameter \mathbf{U} , but the CAPMAP polarimeters are insensitive to \mathbf{Q} polarization, as defined in the instrumental frame. To measure both components of the sky polarization, CAPMAP took advantage of a scan strategy centered near the North Celestial Pole. While the telescope scans this patch in a uniform pattern, the rotation of the sky causes the sensitivity axis of the polarimeters to change from \mathbf{U} to \mathbf{Q} and back again, with a 12 hour period.

While it is an obvious advantage to be able to simultaneously measure both Stokes parameters, there is one downside to the QUIET scheme. The systematics associated with the septum polarizer, described in §6.2, have significantly higher magnitude than similar systematics of the CAPMAP OMTs. This means that QUIET modules have higher $\mathbf{I} \rightarrow \mathbf{Q}$

leakage than CAPMAP polarimeters. Well designed scans for both experiments reduce the risk of the CMB temperature anisotropy contaminating the polarization maps, but higher leakage contributes to increased $1/f$ noise in the demodulated data. The 40 GHz CAPMAP polarimeters typically have knee frequencies (the frequency where the noise power spectrum reaches twice the value of the white noise level) of less than 1 mHz. For QUIET modules, the U diode knee frequencies are a few mHz, and Q diodes are often in the 10 to 30 mHz range.

3.7.3 Heterodyne Detection

Heterodyne detection refers to the process of mixing down to intermediate frequencies before the signals are correlated by the multipliers. There are several advantages to a heterodyne scheme. Intermediate frequencies signals can be carried on coaxial cable (CAPMAP uses semi-rigid SMA cable) and are amplified with small commercial amplifiers. A triplexer in each arm of the CAPMAP polarimeter divides the full bandwidth of the detector into three sub-bands, providing some (very small) amount of leverage in determining the spectrum of a source. Phase tuners at one input of each multiplier can be used to adjust the relative phase of the polarimeter legs, making it easier to achieve a good phase match.

The reason why heterodyne detection is implausible for QUIET is simply that 90 polarimeter modules would require 180 sets of intermediate frequency electronics. If something similar to the CAPMAP “IF box” were used for the QUIET W-band array, then it would be necessary to route 180 coaxial cables on hermetic connections through the cryostat wall. By rectifying the RF signal with detector diodes inside the module, the QUIET design ensures that only DC signals are routed in and out of the cryostat.

CHAPTER 4

THE QUIET INSTRUMENT

While the QUIET module is designed to integrate all of the radiofrequency electronics necessary for detecting polarization into a small package, many other components are needed to make the array function. Like the modules, the other parts making up QUIET are designed with the goals of achieving good detector performance in a compact and scalable design.

The cryostat, which houses the modules and cools them to their operating temperature of 20K, is described in §4.1. Electronics to bias and read out signals from the modules are described in §4.2 through §4.6. Details about the software used to control the array and record data are found in §4.7. Various methods used to optimize bias values for entire arrays of modules in an automated fashion are described in §4.8. The QUIET telescope and groundscreen are covered in §4.9 and §4.10. Finally, §4.11 describes the telescope mount and its control system.

4.1 QUIET Cryostat

To minimize amplifier noise, it is necessary to cool the modules to an operating temperature of 20–30K. While this is a modest temperature requirement compared to the sub-Kelvin focal planes used for bolometric radiometers, the QUIET cryogenic design still poses challenges due to its scale.

The QUIET Q-band cryostat has a cylindrical design, with a height of 29.5 inches and radius of 14.25 inches at the widest point (Newburgh et al., 2006b). It is mechanically supported by a thick steel plate that serves as the middle section of the cryostat and includes mounting flanges. Inside, the cryogenic stages are built above the steel plate, with sections of G-10 used to provide mechanical support with little thermal conduction.

The modules and septum polarizers are bolted to the horn array via an interface plate.

The interface plate is a single piece of metal with waveguide transitions to couple the narrow end of each horn to the input of the corresponding septum polarizer. An inevitable byproduct of the tight joints needed to mate together waveguide pieces is the fact that the modules, polarizers, interface plate, and horn array are all in extremely good thermal contact. This means that, in order to cool the modules, one must also cool the much larger mass (about 20 kg) of the horn array¹. Sources of thermal loading on the cold stage include radiation through the cryostat window, thermal conduction along the module bias lines, and heat dissipated by the low noise amplifiers.

Cooling power is provided by two CTI-Cryogenics Cryodyne 1020 two-stage refrigerators connected to CTI model 9600 compressors. The refrigerators operate as Gifford-McMahon coolers with high purity Helium used for the coolant gas. The first stages of the refrigerators reach temperatures of 80–110K and are used to cool a shield layer that intercepts most of the radiative loading from the outer shell of the cryostat. Wiring for the module bias lines is strapped to the 80K stage to sink heat. The second stages of the refrigerators, which reach temperatures of 15–20K, attach to the interface plate, on either side of the array. Activated charcoal epoxied to the refrigerator heads absorbs any residual gas left after evacuation, allowing the cryostat to remain cold and under vacuum for months at a time without pumping.

Eight silicon diode thermometers (Lakeshore model DT-470) are used to monitor cryogenic temperatures. Two thermometers are located on the interface plate at 20K, three are attached to modules in the array, and the last three are mounted to locations on the 80K stage. These thermometers are biased and read out by the housekeeping board (see §4.4.4). Two more thermometers are used by a Cryo-con Model 32 temperature controller to monitor and regulate the cold stage temperature. The controller supplies current to heaters mounted near the refrigerator heads to hold the cold stage temperature fixed.

1. Cooling the horns and other waveguide components also provides the benefit of improving the conductance of the metal, which reduces loss.

The compressors are physically separated from the cryostat, mounted below the elevation axis of the telescope mount so that they remain upright regardless of the telescope orientation. Long gas lines and power cables pass through the elevation and deck cable wraps to connect the compressors and refrigerators. Air cooling is inefficient at high altitude, so cooling water is used to remove heat from the compressors. The cooling water flows in a closed loop system to chillers located outside the telescope grounds.

Due to the size of the module array, a large window is necessary to ensure that all modules have an unobstructed view of the sky. The radius of the cryostat window is 10.5 inches. The window material is 3/8 inch thick Ultra High Molecular Weight Poly-Ethylene (Newburgh et al., 2006a), which is strong enough to safely hold vacuum, though the window does bow by several inches. A Teflon layer is applied to both sides of the window as an anti-reflection coating (index of refraction is 1.52 for UHMW Poly-Ethylene and 1.2 for Teflon). Besides microwaves, the window transmits infrared radiation, which would create a significant loading on the cold stage. A thick piece of styrofoam mounted inside the cryostat, between the window and the horn array, acts as an infrared filter but is transparent at microwave frequencies.

4.2 QUIET Electronics Overview

Figure 4.1 shows a block diagram of the electronics used to bias an array of QUIET modules and read the detector diode timestreams. Components shown in the diagram are located inside the cryostat, mounted to the outside of the cryostat, and in two 6U VME crates housed in a separate enclosure. The correspondence between modules and electronics boards is also listed in Table 4.1.

Modules are named according to their physical position in the focal plane, seen in the upper left of Figure 4.1, with the prefix “RQ”. Modules RQ-00 through RQ-16 are configured as polarimeters. Modules RQ-17 and RQ-18 are coupled together to form a TT module pair.

The components shown in Figure 4.1 are described in §4.3, which covers passive circuitry for module protection and interfacing, §4.4, which covers electronics for module biasing, and §4.5, which describes the Analog to Digital Converter boards.

4.3 Module Protection and Interface

The amplifiers and phase switches used in QUIET modules are sensitive devices that can be easily harmed by voltage transients, even if they are very brief. Several stages of passive protection circuitry are used to prevent catastrophic events that might damage a module or even the entire array. This section also discusses the challenge of routing a large number of bias and data lines from the modules inside the cryostat to an electronics crate mounted in a separate enclosure.

4.3.1 Module Attachment Boards

Module Attachment Boards (MABs) are printed circuit boards that mate to the module pins, making electrical connections and also providing protection for sensitive RF components. Three MABs, each with a different shape, are used for the Q-band array, holding seven, six, and six modules, respectively (Dumoulin, 2007). Figure 4.2 is a photograph of seven Q-band modules mounted in an MAB.

Each module plugs in to 34 spring socket connectors on the MAB, arranged in two rows of 17. Friction between the pins and sockets for 34 connections times six or seven modules is sufficient to hold the MAB tightly to the modules, which are attached to the interface plate via the septum polarizers. Traces from the sockets are routed to six 40-contact Zero Insertion Force (ZIF) connectors, manufactured by Hirose, that attach to Flexible Printed Circuits, described in §4.3.2. The ZIF connectors are organized by function, rather than module. Two connectors are used for detector diode biasing and readout, three connectors are used for LNA bias lines, and the last connector carries the phase switch bias signals.

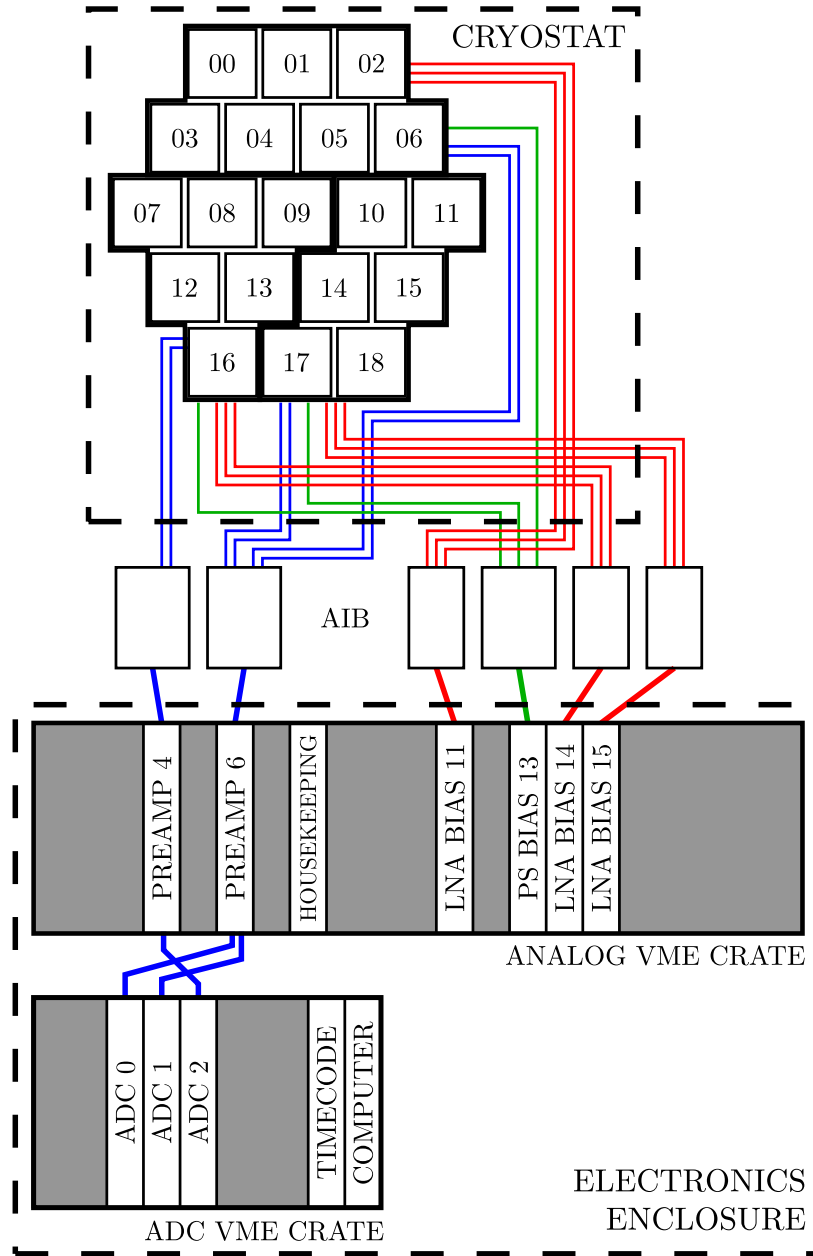


Figure 4.1: Schematic diagram of QUIET analog electronics. The arrangement of modules in the focal plane is shown in the upper left, with the MABs indicated by heavy outlines. Six FPCs run from each MAB to the AIBs located outside the cryostat. Shielded ribbon cables connect the AIBs to preamp, LNA bias, and phase switch bias boards in the analog VME crate, located inside the electronics enclosure. Signals out from the preamps are digitized by ADC boards in a separate VME crate. Connections are color coded by type: LNA bias lines are red, phase switch bias lines are green, and detector diode signal lines are blue.

Table 4.1. Correspondence between modules and their support electronics

Module	MAB	LNA Bias ^a	PS Bias ^a	Preamp ^a	ADC ^b	Q1 ^c	U1 ^c	U2 ^c	Q2 ^c
RQ-00	0	11	13C	6B	1	33	37	41	45
RQ-01	0	11	13C	6B	1	49	53	57	61
RQ-02	0	11	13C	6B	1	32	36	40	44
RQ-03	0	11	13C	6B	1	35	39	43	47
RQ-04	0	11	13C	6B	1	51	55	59	63
RQ-05	0	11	13C	6B	1	34	38	42	46
RQ-06	0	11	13C	6B	1	50	54	58	62
RQ-07	1	14	13A	4B	2	82	86	90	94
RQ-08	1	14	13A	4B	2	65	69	73	77
RQ-09	1	14	13A	4B	2	81	85	89	93
RQ-10	2	15	13B	6A	0	1	5	9	13
RQ-11	2	15	13B	6A	0	17	21	25	29
RQ-12	1	14	13A	4B	2	83	87	91	95
RQ-13	1	14	13A	4B	2	66	70	74	78
RQ-14	2	15	13B	6A	0	2	6	10	14
RQ-15	2	15	13B	6A	0	18	22	26	30
RQ-16	1	14	13A	4B	2	67	71	75	79
RQ-17	2	15	13B	6A	0	3	7	11	15
RQ-18	2	15	13B	6A	0	19	23	27	31

^aLNA bias boards, phase switch (PS) bias boards, and preamplifier boards are labeled by the slot that they occupy in the analog electronics crate. A single phase switch bias board contains circuitry to bias three MABs, with separate power supplies for each; the three divisions of the board are labeled “A”, “B”, and “C”. Likewise, each preamp board contains circuitry to bias two MABs; the preamp divisions are labeled “A” and “B”.

^bADC boards are labeled according to their ordering, from left to right, in the ADC VME crate.

^cThe values listed for each detector diode correspond to the position where samples from that diode are recorded in data files. Data file indexes run from 0 to 95, with 20 unused channels.



Figure 4.2: Photograph of seven Q-band modules, with septum polarizers attached, installed on an MAB. Photo credit: Ross Williamson.

Protection circuitry for LNA gate bias lines consists of a voltage divider in parallel with a $1\ \mu\text{F}$ capacitor, attenuating the input voltage by one-third at DC and by two-thirds at 19 kHz. Switching diodes are used to clamp gate voltages between ± 0.6 Volts before the voltage divider, so gate voltages at the module pin are limited to the range $\pm 0.4\text{V}$.

The LNA drain bias lines are voltage clamped between a yellow LED and a switching diode, limiting voltages to the range $[-0.4, +1.5]\text{ V}$. A $1\ \mu\text{F}$ shorting capacitor is used to protect against voltage transients.

Phase switch bias lines are filtered by a one-pole RC filter with $R = 200\Omega$ and $C = 1000\text{pF}$, resulting a 3dB attenuation frequency of about 800 kHz. A voltage clamp consisting of two switching diodes in one orientation and two red LEDs in the other orientation limits voltages to the range $[-3.0, +1.43]\text{ V}$.

No protection circuitry is used on the detector diode anode or cathode lines.

The MAB protection circuitry is designed to prevent voltage transients from damaging

the module during lab operations at room temperature. When the cryostat cools down, the turn on voltages of various diodes increase dramatically, causing the voltage clamps to stop functioning. These voltage clamps are duplicated on the Auxiliary Interface Boards (see §4.3.2) mounted to the outside of the cryostat so that the modules are still protected when they are cold, though that protection circuitry is not as physically proximal.

4.3.2 Connecting the Module Array to the Bias Electronics

Connections to the MABs are made using 40 conductor Flexible Printed Circuits (FPCs) with 0.5 mm pitch. The use of FPCs solves one of the major design challenges for QUIET – how to route 228 LNA bias lines, 76 phase switch bias lines, and 152 detector diode signal lines out of the cryostat. The problem is even more dire for the W-band array, which features five times as many modules. The solution developed for QUIET is to use the shell of a hermetic connector with a slotted G-10 insert. A set of six FPCs for one MAB are threaded through the slots to the appropriate depth and then the shell, insert, and FPCs are all sealed together with Stycast epoxy to make the connector vacuum tight. The hermetic connector shell includes a jam nut o-ring seal so that it can be mounted in an appropriately sized hole in the cryostat wall. Figure 4.3 shows a photograph of a potted FPC.

Outside the cryostat, the FPCs plug into Auxiliary Interface Boards (AIBs). The AIBs are mounted onto the cryostat and feature protection circuitry that duplicates some of the MAB circuitry (see §4.3.1). By potting the FPCs to make hermetic feedthroughs, the link between the MABs and AIBs is made with only a single connector at either end. Unfortunately, the ZIF connectors used for the FPCs were a common failure mode. They degrade after a fairly small number of connect–disconnect cycles and are easy to break if they are not handled carefully. Also, the connectors are quite sensitive to the thickness of the FPC itself, so extra material was added near the FPC ends.

From the AIBs, bias and signal lines connect to the backplane of a 6U VME crate




Figure 4.3: Photograph of Flexible Printed Circuits threaded through a hermetic connector shell and sealed with epoxy. Photo credit: Osamu Tajima.

housing the analog bias and readout electronics (preamplifiers, LNA bias boards, and the phase switch bias board). The analog VME crate is housed in a enclosure (see §4.6) located behind and to the side of the cryostat. Six foot shielded ribbon cables, referred to as “AIB cables”, are sufficiently long to make the connection through a hole cut in the back door of the enclosure. Durable weatherproofing is added to shield the AIBs and cables from sun, wind, dust, and snow.

Longer, unshielded AIB cables were made for use in lab. These cables were reinstalled for a brief period during the observing season, as part of a debugging effort (see §5.4).

4.4 Bias Electronics

The QUIET bias electronics supply the voltages and currents necessary to operate the module array, and also amplify the detector diode voltages that carry the data signal. Two preamp boards, three LNA bias boards, and one phase switch bias board are housed in the analog VME crate, along with a housekeeping board that monitors the biases. All bias settings are

controlled digitally, with command signals sent by the master ADC board (see §4.5.2 for more details). While the biases remain fixed during observations, the ability to quickly set and adjust biases by computer command, rather than by adjusting potentiometers, as was done in CAPMAP, is crucial for testing and operating a large array of modules.

The analog crate is a 6U VME crate with a custom backplane. The P2 connectors for each board are through pins that connect to the modules via the AIB cables (see §4.3.2). The P1 connectors interface to traces on the backplane that carry digital bias commands, clock signals for the phase switch board, and multiplexer addresses used in bias monitoring. The backplane is configured to accept preamp boards in slots 1 through 7, the housekeeping board in slot 9, and LNA bias or phase switch bias boards interchangeably in slots 10 through 21.

Power for the analog electronics is provided by a tray of $\pm 5\text{V}$ linear power supplies. A separate power supply is provided for each LNA bias board, each half of a preamp board (one preamp board operates two MABs worth of modules), and each third of the phase switch board (one phase switch board operates 3 MABs worth of modules). All of the module case grounds are unavoidably shorted together through the platelet array, so achieving a true star grounding scheme would require independent power supplies for each module. This would result in an unfeasibly large number of supplies, so a compromise solution of providing separate power supplies for each MAB was adopted.

4.4.1 Amplifier Bias Board

Three LNA bias boards are used to supply gate voltages and drain currents to modules in the array. Each board was originally designed to bias two W-band MABs (fourteen modules). However, Q-band modules require 1–2 extra independent drain current lines (depending on the module type). Rather than redesigning the board for the Q-band array, a single board is used for each MAB, with the additional drain lines taken from the otherwise unused circuits

on the second half of the board.

The gate voltage circuit is simple, consisting of a digitally controlled voltage source² and a voltage divider. The output voltage range is -0.53V to +1.68V, with an output impedance of 500 Ω .

The drain bias circuit is an adjustable current source that uses the collector of a PNP transistor to supply the module amplifiers. Two variants of the drain circuit are used. First stage LNAs are individually biased, so a circuit with maximum current of 25 mA is used. The second and third stage LNAs in each module arm have their drain inputs biased in parallel, so the current requirements are twice as large, resulting in a drain bias circuit capable of sourcing up to 50 mA.

The transistors used for the drain bias circuit suffer from large temperature coefficients, which cause the output current to vary by up to 1.5% per degree change of the bias board temperature (Bischoff et al., 2006). For the bias boards used in the Q-band observing season, a compensation scheme was implemented using thermistors with matching coefficients (Seiffert et al., 2006). This fix reduced the drain current temperature coefficient by a factor of 5–10 over a limited temperature range, centered at 25°C. Later versions of the LNA bias board abandoned the current source circuit for drain biasing and switched to a voltage bias configuration.

4.4.2 Phase Switch Bias Board

The phase switches in each leg of a QUIET module are biased by applying forward voltage to one input while simultaneously reverse biasing the other. When the forward and reverse biased inputs are swapped, the phase length for RF signals passing through the switch changes by 180°.

²Linear Technologies LTC1660 Digital to Analog Converters are used for digital bias control on the LNA bias (gate and drain circuits), phase switch bias, and preamp boards.

The phase switch bias circuit uses a double pole, double throw (DPDT) switch to select between an adjustable forward bias voltage, ranging from 0 to 2.5V, and a constant reverse bias voltage of -1.8V. Each output passes through an RC filter with $R = 470\Omega$ and $C = 1000$ pF, which prevents sudden voltage spikes that could damage the delicate PIN diodes found in the phase switches. This filter also has the effect of slowing the transition time between states. In mid-transition, the phase switch blocks RF power, so the rectified signal goes to zero. This results in large spikes in the signal out of the modules, which must be blanked out before demodulation. Filtering on the phase switch lines balances the need for reliable protection with the desire to only blank a minimum fraction of data.

One phase switch bias board has enough circuits to operate all modules in the array. Separate clock signals are used for A-leg and B-leg phase switches, so the legs can be set independently to switch at 4 kHz or 50 Hz, or to a DC state. For normal operation, one leg is switched at 4 kHz while the other is switched at 50 Hz, as described in §3.5.1. A separate pair of clocks is used for even and odd modules, so they could be switched with different phase. In practice, these clock signals are always the same for all modules.

4.4.3 Preamplifier Board

The preamplifier board biases the detector diodes and amplifies their output signals, with an adjustable offset on the output voltage.

The diode bias circuit consists of a digitally controlled voltage source with $10k\Omega$ of series resistance. QUIET detector diodes are not grounded to the module case, so the preamp board both sources and sinks the bias current.

Voltage across the detector diode is amplified by an instrumentation amplifier with gain of 3.8. While the input impedance of the instrumentation amplifier is very large, the actual input impedance of the preamp circuit is set by the series resistance of the bias circuit. This means that it is important to apply enough bias to the detector diodes to bring their

impedance significantly below $10k\Omega$. For cryogenic operation, the typical bias setting is about 260 mV of bias on the diodes, which corresponds to roughly $35\ \mu\text{A}$ of current, or half scale on the bias circuit.

After the first amplification stage, an offset removal circuit, with gain of ~ 2 , is used to subtract the large DC level in the signal left over from the diode bias. As opposed to simply AC coupling the signal, this offset subtraction scheme retains some information about the DC voltage on the diode, which provides a noisy monitor of the total input power to the module. The usual detector diode biasing procedure is to first apply the desired bias to the diodes and then to adjust the offset level to approximately zero the output, all before turning on the LNAs and phase switches, so that the DC level out of the preamp reflects the RF power rectified by the detector diodes, but not the applied bias voltage.

The output stage of the preamp is a differential line driver with gain of 20. The differential output signal is balanced about a DC level of +2.5V, in order to correctly interface with the ADC input. To achieve good noise performance from the ADC, it is necessary to limit the bandwidth of the preamp output signal so that it doesn't significantly exceed the sampling Nyquist frequency (400 kHz). However, if the signal is filtered too aggressively, the result will be to smear out the transitions between phase states, greatly increasing the amount of data that needs to be blanked before demodulation. The QUIET preamp uses three filter poles, with 3dB frequencies of 340 kHz, 160 kHz, and 32 MHz, in combination with a one pole filter on the ADC board that has a 3dB frequency of 600 kHz.

Combining the amplification stages, the preamp has gain of ~ 130 . Referred to the preamp input, the noise of the back end electronics (preamp plus ADC), with a cold, biased detector diode at the preamp input, is about $8\ \text{nV}/\sqrt{\text{Hz}}$. This noise adds in quadrature with the noise due to any rectified RF signal. The RF noise is just given by the size of the rectified signal, in Volts, divided by the square root of the module bandpass³. For a typical rectified

3. This result can be obtained by multiplying the radiometer equation by a factor of the gain

signal level of 1 mV at the detector diode and 7.5 GHz of bandwidth, the front end noise is about $12 \text{ nV}/\sqrt{\text{Hz}}$, so the back end noise is subdominant, but only by a small margin.

Both the diode bias and amplifier circuits have some sensitivity to temperature. The primary effect of changing the preamp temperature is to cause an offset drift in the preamp output, which has no effect on demodulated data. Of the terms that do affect the polarization signal, the most significant is due to changes in the diode bias level leading to variations in gain of 0.13% per degree (Bischoff & Kusaka, 2007).

4.4.4 Housekeeping Board

The housekeeping board monitors module biases and temperatures both inside the cryostat and in the electronics enclosure. Multiplexing on the housekeeping board and the bias boards is used to cycle through values, which are then sampled by an analog to digital converter on the board.

Twenty-four thermometer bias circuits are included on the housekeeping board, with half configured for the Silicon diode thermometers used in the cryostat and the other half configured to bias $100k\Omega$ thermistors that are used to monitor temperatures in the electronics enclosure. Thermometer bias currents are applied constantly, rather than multiplexing along with the readout circuit.

The full multiplexer address is specified by 12 parallel bits. The first eight bits are shared between the housekeeping board, LNA bias boards, and the phase switch board. On the bias boards, the bit values are used to select a quantity to monitor. On the LNA bias boards, monitored quantities include drain currents, drain voltages, gate voltages, and a ground sense voltage for each MAB. On the phase switch bias board, forward bias current is monitored for each phase switch. On the housekeeping board, the first eight bits of the multiplexer address select between the 24 thermometers. The last four bits of the multiplexer address are used by the housekeeping board to select between the values supplied from each bias board or the

thermometer. The single monitoring channel chosen after all multiplexing is then digitized by an analog to digital converter chip on the housekeeping board.

Care is taken to isolate the bias boards from the noisy digital signals found on the housekeeping board. Monitoring outputs of the LNA and phase switch bias boards use analog optoisolators to avoid sharing a ground connection with the housekeeping board. Additionally, steps in the multiplier address and digitization of the current monitored value both occur during the transitions between 4 kHz phase states, when the data will be blanked anyway.

The multiplexer address is advanced every two milliseconds, with the analog to digital conversion occurring one millisecond later. For the Q-band array, a total of 491 values are monitored, resulting in a sampling rate of about 1 Hz for each quantity.

4.5 Analog to Digital Converter Board

The QUIET ADC boards have a flexible design, based around an Altera Field Programmable Gate Array (FPGA). The FPGA handles blanking, averaging, and demodulation of polarimeter data. One ADC board is designated the master; this board generates clock signals to synchronize sampling across all ADC boards and issues digital command signals to set bias values and control the housekeeping multiplexer.

ADC boards are housed in a Wiener Series 6000 powered 6U VME crate. The ADC crate is located inside the electronics enclosure, along with the analog VME crate. The ADC crate also contains a computer that runs the ADC Server software to control the ADC boards and a time code reader that keeps the receiver synchronized to a GPS receiver located in the control room.

4.5.1 Data Acquisition

Data signals from the preamp are digitized using Analog Devices model AD7674 Analog to Digital Converters. These devices provide 18 bits of precision with a sampling rate of 800 kHz. The digitization is carried out using a Successive Approximation Register (SAR), which samples and holds the input data while running a binary search to find the appropriate bit value. A consequence of using this type of ADC is that the actual sampling occurs almost instantaneously, so it is important to filter the data to remove noise above the Nyquist frequency which could otherwise alias in.

The FPGA processes the raw 800 kHz data to produce average and demodulated data streams. First, the phase switch transition spikes are blanked out, using an adjustable mask. For Q-band observations, 14 samples ($17.5 \mu\text{s}$) are masked at each phase switch transition, which removes 14% of the raw data samples. Averaged data, proportional to the total power signal, is calculated by summing up all unmasked samples for an integration period of 10 ms. Demodulated data, which carries the polarization signal, is also produced by summing samples, but with a minus sign applied to samples in one of the phase states. Blanking and demodulation of 800 kHz data is illustrated in Figure 4.4. Since the ADC board also generates the phase switch clock signal (see §4.5.2), alignment of the data with the blanking mask and demodulation clock is guaranteed.

Each ADC board can digitize 32 analog input channels, enough for all detector diodes on an MAB, with channels to spare. One 100 Hz averaged or demodulated sample consists of up to 8000 raw samples (if the blanking mask is set to zero), each with 18 bits of precision. A 32 bit integer is large enough to record a single sample of average or demodulated data, with no loss of dynamic range. Average and demodulated data for the 32 channels on each board are accumulated into 0.25 second (25 sample) frames. When the frame is complete, it is transferred to the crate computer via the VME backplane.

The 800 kHz raw data, 100 Hz averaged or demodulated data, and 4 Hz frames are




Figure 4.4: Raw 800 kHz data from module RQ-09, diode Q1. Phase switch transitions cause the large upward going spikes every 100 samples. The blanking region is shaded gray and accounts for 14% of the samples. The demodulated signal is calculated by summing the samples in each phase state (red and blue) and then differencing the results.

synchronized between all ADC boards. The master ADC board uses phase locked 1 Hz and 10 MHz clock signals from the time code reader to synthesize the sampling clocks, which are distributed across the VME backplane. These clocks do not contain any absolute time information, they only ensure that frames begin simultaneously at the start of each second and individual samples for various ADCs line up correctly.

4.5.2 Digital Control Signals

In addition to digitizing data, the master ADC board also sends digital signals to adjust bias values, set the housekeeping multiplexer address, and toggle the phase switches. Differential line drivers on the ADC board use Low-Voltage Differential Signaling (LVDS) logic

to transmit signals over ribbon cables that connect to a small LVDS receiver board plugged into the backplane of the analog VME crate. The LVDS receiver board distributes signals across the backplane to bias boards in the crate.

The LNA bias boards, phase switch board, and preamp all use LTC1660 10-bit Digital to Analog Converters (DACs) to control their bias circuits. A single LTC1660 contains eight separately addressable output channels, and on each board multiple DAC chips are connected together in a daisy chain configuration (28 DACs per LNA bias board, 12 on each phase switch board, and 14 on each preamp). Four digital lines are used to control the DACs: card select (CS), serial clock (SCK), data in (DIN), and an asynchronous clear line (CLR). Three of the lines, SCK, DIN, and CLR, are common to all bias boards, but there are separate CS lines for each board.




Figure 4.5: Example of a serial bias command. With card select (CS) low, bit values from data in (DIN) are registered on each rising edge of the serial clock (SCK).

Figure 4.5 illustrates the bias update command. To set a bias value on a particular board, the first step is to switch the logic level of the appropriate CS line to low. Next, serial bits are pushed into the DAC register, one at a time, on each rising edge of SCK. A bias update command is 16 bits long, with the first four bits specifying an address on the DAC chip, the next 10 bits specifying the command value, and the last 2 bits unused. In the daisy chain configuration, older bits get pushed from one DAC to the next as newer bits are added. So,

for an LNA bias board with 28 DAC chips, a total of 448 bits are sent. Once that is finished, the CS line is switched back to high, updating the analog biases. The CLR line instantly zeros the bit setting for all addresses on all DACs, but it is never used.

Bit settings for all DACs on each board are stored in memory on the master ADC board. These values can be updated, but the actual DAC settings are not changed until an update command is sent. On an update command, the master ADC sequentially sends out commands to update every DAC address on every bias board.

Twelve more LVDS output lines are used for the housekeeping multiplex address, as described in §4.4.4. The total number of monitored quantities (491) is far less than the 2^{12} possible addresses. Instead of blindly stepping through all address values, the ADC cycles through an address list stored in memory. Shorter address lists can be used to achieve more frequent monitoring of fewer values.

Additionally, the master ADC generates the four clock signals used by the phase switch bias board. The A and B leg phase switches can be independently set to DC mode, 50 Hz switching, or 4 kHz switching. The first pair of clocks applies to modules 0, 2, 4, and 6 on each MAB while the second pair of clocks applies to modules 1, 3, and 5. In practice, all modules on all MABs are phase switched in the same way.

4.5.3 Timing Synchronization

A Symmetricom TTM635VME Time and Frequency Processor (timecode reader) is used to keep the QUIET receiver synchronized to a GPS receiver located in the control room. The GPS transmits timing information formatted as an IRIG-B timecode. The timecode reader interprets this signal to determine the absolute date and time at 1 second intervals. This absolute timing information can be read on demand by the crate computer via the VME backplane, along with possible error conditions, such as loss of the IRIG-B signal.

The timecode reader is also programmed to output the synchronized 1 Hz and 10 MHz

clock signals that are used by the master ADC board to generate clocks for phase switching, sampling, masking, and demodulation. A buffer circuit is used to convert the logic levels of the clock signal from TTL (used by the timecode reader) to LVDS (used by the ADC boards).

4.6 Electronics Enclosure

The analog and ADC VME crates are housed in a weatherproof enclosure that provides protection from the elements. The other important purpose of the enclosure is to regulate the temperature of the electronics boards inside (see §4.4.1 and §4.4.3 for some discussion of temperature coefficients for the QUIET electronics). A PID controller (Omega model CNi16D22-EI) is used for regulation. Cooling the enclosure is accomplished by running a large fan to draw in filtered air from outside. Heaters are mounted near the analog VME crate, with additional fans designed to circulate air throughout the enclosure.

Figure 4.6 shows histograms of the mean temperature (left) and temperature range (right) for every constant elevation scan during the observing season, divided into day (red) and night (blue). The enclosure temperature is regulated at 25°C and usually remains within one degree of the target temperature. On warm days, the cooling power provided by drawing in air from outside is sometimes insufficient to keep the enclosure cold, leading to the tail at high temperatures of the red distribution. The right hand panel shows the maximum range of temperatures within each scan (usually about 70 minutes). This range is less than one degree for 76% of daytime scans and 94% of nighttime scans.

Besides the two VME crates, the electronics enclosure also contains a variety of other equipment used to operate the QUIET array.

- Surge protection and power conditioning units
- Linear power supplies for the analog electronics




Figure 4.6: On the left, a histogram of the mean enclosure temperature for every CMB scan in the observing season, divided into daytime (red) and nighttime (blue). On the right, a histogram of the temperature range for each scan.

- PID temperature controllers for the cryostat and the enclosure
- Commercial analog I/O boards (Sensoray model 2608) used to monitor telescope mirror and groundscreen temperatures, and to operate the calibration noise source
- The computer used to command the array and record data (see §4.7)
- The computer used to operate the optical star tracker telescope
- Ethernet hub

4.7 QUIET Receiver Software

Software to operate the QUIET receiver runs on two computers, one located in the ADC VME crate and the other mounted separately in the electronics enclosure. The crate computer is limited in memory, disk space, and processing power, but communicates directly with the ADC boards via the VME backplane. The other computer, known as the receiver computer, records data and runs scripts for biasing operations.

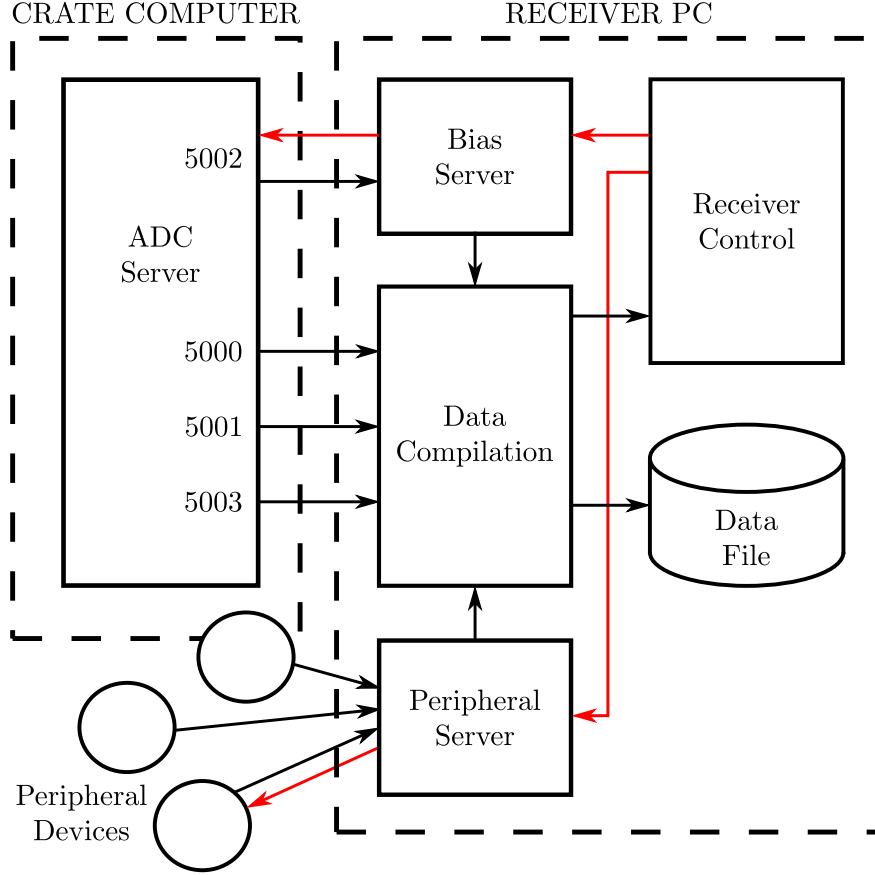


Figure 4.7: Schematic diagram of the QUIET receiver software. Dashed boxes indicate which programs run on the crate computer or the receiver computer. Black arrows indicate the flow of radiometer, housekeeping, and peripheral data. Red arrows indicate the propagation of user commands.

4.7.1 ADC Server

The ADC Server program, running on the crate computer, acts as the software interface to the ADCs. Four ports are made available by ADC Server for receiving commands and transmitting data. The data port (5000) streams frames every quarter of a second containing average and demodulated data from all ADC boards. Each data frame is timestamped by ADC Server, using the absolute date and time from the time code reader. A third type of data, called “quadrature data”, was not implemented in time for the Q-band observations, though space for it was still reserved in the data frames. Quadrature data is generated by

demodulating the raw timestreams 90 degrees out of phase from the phase switching, to create a timestream with zero signal but similar noise properties to the demodulated stream.

The housekeeping port (5003) streams frames every quarter of a second containing the 500 Hz data stream from the housekeeping board and the multiplexer addresses needed to decode it. Housekeeping data is timestamped in the same fashion as the radiometer data.

The snapshot port (5001) can be used intermittently to acquire a short stretch (1024 samples) of 800 kHz raw data for all ADC channels. These data are very useful for diagnosing module problems. Snapshot files were recorded every 10 minutes during the Q-band observing season.

The bias port (5002) is used to send commands to the ADC boards. The available commands are

- Set the bias bit value for a specified channel
- Update biases on all boards to the current bit values
- Set the mode (DC, 50 Hz, or 4 kHz) of the phase switch clocks
- Set a list of addresses to use for the housekeeping multiplexer
- Update the length and offset of the blanking mask

In addition to receiving commands, the bias port can be used to request the block of memory on the master ADC containing the blanking mask definition, the phase switch clock states, and the bit values of all bias channels.

4.7.2 Data Compilation

The Data Compilation program assembles data from several sources and writes to files on the receiver computer. The bulk of these data are the average and demodulated samples from the modules, but it also includes the housekeeping data, records of the blanking mask, phase

switch clock states, and bias bit values, plus assorted peripheral data (see §4.7.5). Frames in the data file are one second long, so they each contain four sub-frames of radiometer and housekeeping data.

The files written on the receiver computer consist of data from the QUIET receiver only. Notably they lack any information about the telescope pointing. This state of affairs came about in the rush to deploy the QUIET Q-band array. The telescope mount, inherited from the Cosmic Background Imager (CBI) project, uses a program called CBI Control to control the telescope and record its position. While updates were made to CBI Control to accommodate QUIET, particularly improvements in the telescope motion control to allow for fast scanning, the planned integration of the receiver data into the program did not finish in time. This system, with separate data files for the receiver and pointing data, continued for the full Q-band observing season. During the transition from the Q-band to W-band array, both Data Compilation and CBI Control were updated so that the receiver data is recorded alongside the telescope encoder coordinates.

To simplify data analysis, the two sets of data files are reprocessed after being transferred to Chicago. The results of this reprocessing are known as “Level-1” data files (the files written by Data Compilation and CBI Control are Level-0). In addition to lining up the radiometer data with the telescope pointing, the Level-0 to Level-1 conversion organizes the data in a more intuitive fashion – demultiplexing the housekeeping data, for example. The use of Level-1 files for analysis was so successful that the format persists even after the switch to W-band observations, with Level-0 data files that integrate receiver and mount data.

4.7.3 Bias Server

The Bias Server program handles the problem of competing commands and data requests on the ADC Server bias port. On startup, the program queries ADC Server to get the blanking mask, phase switch clock states, and bias bit values. These data are provided to

Data Compilation for inclusion in each data frame. After a successful command to update any of the receiver settings, Bias Server updates its internal record, so the change is reflected in the data files.

4.7.4 Receiver Control

Receiver Control consists of a set of routines that perform common receiver tasks. The most commonly used functions in Receiver Control are turning on the array by setting all biases to values specified in a file, or turning off the array by setting all biases to zero.

More complicated bias setting processes are possible by feeding back on radiometer or housekeeping samples read from Data Compilation. For example, the LNA bias board provides control of gate voltages and drain currents. However, when turning on an amplifier, it is often convenient to specify the drain voltage and current. Receiver Control carries this out by adjusting the gate voltages and feeding back on housekeeping measurements until the drain voltage reaches the desired bias point.

4.7.5 Peripheral Server

The Peripheral Server program creates a unified interface for a variety of other devices related to the receiver. Its main function is the same as Data Compilation – to arrange data from several sources into 1 Hz frames. Some of the peripheral devices were added shortly before deployment, so defining Peripheral Server as a separate program meant that Data Compilation would not be subjected to as much last minute software development.

The devices producing data compiled by Peripheral Server include

- Cryostat pressure sensor
- Temperature controllers for the cryostat and electronics enclosure
- Thermometers mounted on the telescope mirrors and groundscreen

- Calibration noise source
- Status information from the ADC VME crate

4.8 Optimization of QUIET Modules

Electronic control of bias values is critical for efficient optimization of arrays of modules. Putting aside the phase switch and detector diode biases, which are typically set to fiducial values, the 19 element Q-band array takes 228 voltages and currents to operate the low noise amplifiers; for the 91 element W-band array, there are 910 values to set. Details of the biasing affect noise temperature, gain, and phase for each amplifier, so the goal of module optimization is achieve the best possible sensitivity, as described in §3.3.

4.8.1 Receiver Temperature Measurement

The noise contribution from the amplifiers, \mathbf{T}_{rec} , can be obtained by measuring the rectified power for two values of load temperature and extrapolating to zero load. Keeping the non-switching terms from equation (3.38), and ignoring circular polarization, the response to load temperature \mathbf{T}_{load} is given by

$$\langle |E|^2 \rangle = \frac{1}{4} \left[(g_A^2 + g_B^2) \mathbf{T}_{\text{load}} + g_A^2 \mathbf{T}_{\mathbf{A}} + g_B^2 \mathbf{T}_{\mathbf{B}} \right] \quad (4.1)$$

$$= \frac{1}{4} (g_A^2 + g_B^2) (\mathbf{T}_{\text{load}} + \mathbf{T}_{\text{rec}}) \quad (4.2)$$

where the receiver temperature, \mathbf{T}_{rec} , is given by

$$\mathbf{T}_{\text{rec}} = \frac{g_A^2 \mathbf{T}_{\mathbf{A}} + g_B^2 \mathbf{T}_{\mathbf{B}}}{g_A^2 + g_B^2} \quad (4.3)$$

Blackbody loads are constructed from buckets that are transparent to microwaves⁴ filled with Eccosorb, which is a material that acts as a blackbody absorber at microwave frequencies. The bucket is filled with a liquid cryogen, such as liquid Nitrogen (77K) or liquid Argon (87K), to set the load temperature. If the loads are sufficiently large, then they will fill the beams for multiple modules in the array, parallelizing the measurement. Switching between LN2 and LAr loads, the receiver temperature can be calculated for each detector diode as

$$\mathbf{T}_{\text{rec}} = \frac{T_{\text{LAr}} \langle |E|^2 \rangle_{\text{LN2}} - T_{\text{LN2}} \langle |E|^2 \rangle_{\text{LAr}}}{\langle |E|^2 \rangle_{\text{LAr}} - \langle |E|^2 \rangle_{\text{LN2}}} \quad (4.4)$$

This measurement assumes that the detector response is linear with temperature. However, as load temperature increases, the low noise amplifiers will eventually exhibit saturation behavior. Ideally, the receiver temperature measurement would be performed for load temperatures that match the observing conditions of QUIET (atmospheric load temperatures are typically $\sim 10\text{K}$), but this would require a colder cryogen (such as liquid Helium), which is difficult to handle and much more expensive. A saturation type effect biases the measured receiver temperature high, so this test can still provide an upper limit on \mathbf{T}_{rec} .

4.8.2 Isolation and Bandwidth Measurement

A narrowband source that sweeps through the module frequency range can be used to directly measure isolation and bandwidth. The signal is generated at intermediate frequencies of 0–20 GHz and upconverted to the RF band using a frequency multiplier⁵. The multiplier output is broadcast with linear polarization from a feedhorn pointed at the array. If the linear polarization is aligned with the instrumental X axis, then the Q and U diode demodulated

4. Various types of foam, particularly Zotefoam, are quite transparent at microwave frequencies while also providing the structural strength and thermal insulation necessary to hold liquid cryogens.

5. Frequency multipliers use a non-linear element to generate harmonics of the input signal (Pozar, 2005, Chapter 12.4). The desired harmonic is then filtered and amplified.

signals are given by

$$\langle |E_Q|^2 \rangle_{de} \propto 2g_A(\nu)g_B(\nu) \cos \phi(\nu) \quad (4.5)$$

$$\langle |E_U|^2 \rangle_{de} \propto 2g_A(\nu)g_B(\nu) \sin \phi(\nu) \quad (4.6)$$

while the averaged signal for all four diodes is given by

$$\langle |E|^2 \rangle_{av} \propto \left(g_A^2(\nu) + g_B^2(\nu) \right) \quad (4.7)$$

Isolation can be calculated at each frequency as $1 - \langle |E_Q|^2 \rangle_{de} / \langle |E|^2 \rangle_{av}$ for Q diodes or as $\langle |E_U|^2 \rangle_{de} / \langle |E|^2 \rangle_{av}$ for U diodes. The module bandwidth can be calculated by integrating these signals over the frequency range, as shown in equation (3.47). However, for the bandwidth measurement, it is important to take into account any frequency dependence in the output power of the swept source.

4.8.3 Sensitivity Optimization

The most direct method of optimization is to simply search for the bias settings that provide the best sensitivity to polarization. A modulated blackbody polarization source is created by installing a rotating array of thin parallel wires between the modules and one of the cryogenic loads. A thin wire scatters radiation that is polarized with its electric field parallel to the wire while allowing the orthogonal polarization to pass (Houde et al., 2001). For parallel wires with spacing much less than the wavelength, the wire array fully reflects one linear polarization while transmitting the other. In practice, sparse arrays are used, as smaller polarization fractions are quite adequate for optimization.

It is very difficult to calculate the exact size of the polarized signal induced by such a wire grid, because it would be necessary to determine where all of the reflected rays

terminate. However, as long as the signal is stable, a search algorithm can be used to find bias settings that maximize its signal to noise. As the wires rotate, the true polarization signal is modulated at twice the rotation frequency. Changes in effective load temperature, caused by stray reflections scanning across objects in the lab environment are periodic at the rotation frequency, so they can be effectively removed from the analysis.

With a sufficiently large wire array and cryogenic load, the bias settings for many modules can be optimized simultaneously. This method was used to tune the 84 polarimeters in the W-band array and will be refined for larger arrays of QUIET modules planned in the future.

4.9 Telescope

The QUIET telescope focuses the detector response pattern from the parallel feeds of the platelet array to produce beams on the sky with full width at half maximum (FWHM) of 28 arc-minutes. The 19 elements of the array produce a hexagonal field of view measuring 7° across, with 1.75° spacing between neighboring beams.

The telescope employs a crossed-Dragone optical design with 1.4 meter primary and secondary reflectors. The crossed-Dragone is based on an off-axis Gregorian, where the optical axes of the parabolic primary mirror and hyperbolic secondary mirror are not aligned. The benefits of this telescope design include a large diffraction limited focal plane, even for a flat focal plane like QUIET, and low cross-polar response (Hanany & Marrone (2002), Tran (2003)). Performance of the QUIET optics is discussed in §7.4.

The mirrors are each machined from single pieces of aluminum and mounted on a rigid hexapod frame. Alignment of the mirrors and cryostat is accomplished by measuring the distances between tooling balls installed on the sides of the mirrors and adjusting turnbuckles in the hexapod until the distances match the desired geometry. This alignment process was first performed during a test integration before the receiver was shipped to Chile, and then repeated during deployment. The turnbuckles are supported by a rigid sled structure, which

is in turn bolted to two large parallel rails which form the deck of the telescope mount.

4.10 Groundscreen

A small fraction of the detector response spills past the telescope mirrors to create sidelobe response at large angles from the main beams. Depending on the orientation of these sidelobes, they could terminate on the sky (radiometric temperature $\sim 10\text{K}$) or the ground ($\sim 270\text{K}$). For QUIET, which only measures anisotropies, the magnitude of this contribution is less important than its stability.

The QUIET groundscreen consists of two sections, co-moving with the telescope, designed to intercept sidelobe beams. The lower groundscreen is a large box, which can be seen in Figure 4.8, that surrounds the mirrors and telescope sled, with the cryostat projecting through a circular hole in the side. A large square opening in the roof, above the primary mirror, allows the main beam to escape. Hinged panels on the sides allow access to the telescope and the front of the cryostat. The interior of the lower groundscreen is covered with Eccosorb. A layer of polyethylene foam (Volara) is added on top of the Eccosorb for weatherproofing. No attempt is made to regulate the temperature of the groundscreen, so it varies based on the ambient temperature and direct sunlight. However, thermometers are installed under the Eccosorb at several locations, to monitor changes.

Simulations of the QUIET optics show that the lower groundscreen is insufficient to block all sidelobes, as there are some which pass near the main beam and escape through the lower groundscreen aperture (Holler, 2007). The upper groundscreen, designed to address this problem, consists of a cylindrical tube that attaches to the lower groundscreen, restricting the aperture so that only rays aligned with the main beam can pass. Like the lower groundscreen, the inside surfaces of the cylinder are lined with Eccosorb. The design of the upper groundscreen is complicated by the fact that it is too tall to fit inside the telescope dome when the mount is parked. This issue is solved by building the cylinder from



Figure 4.8: Photograph of QUIET, with the lower groundscreen indicated. Also labeled are the electronics enclosure and the cryostat (which is mostly obscured by the enclosure).

telescoping segments, which retract when the telescope is parked with the dome closed and extend for observing. This added design requirement meant that the upper groundscreen was not finished in time for the Q-band observing season. The consequences of the sidelobes that evade the lower groundscreen are discussed in §6.7.

4.11 Telescope Mount

QUIET uses a mount originally built for the CBI experiment, which operated at the site now occupied by QUIET from 1999 through 2008.

The mount features three degrees of freedom, as shown in Figure 4.9. The azimuth gear is mounted on top of a conical base and is driven by pinion gears attached to two large electric motors. The azimuth motors are capable of driving the telescope at speeds up to 6 degrees per second. The scan strategy used by QUIET is based on high speed azimuth scans, as described in §5.2. The azimuth cable wrap is located inside the conical base. Lines for power, signals, or coolant water all make several turns around the cable wrap, and a tensioner spring is used to take up slack. The wrap allows the telescope to rotate up to 440° in azimuth, from 175° West of North to 265° East of North.




Figure 4.9: Diagram of the QUIET mount, with the axes of rotation indicated.

The elevation drive consists of a rigid deck hinged on one side, above the azimuth gear. Opposite the hinge is a long ball screw that turns to tip the deck up and down. In the rest position, with the deck parallel to the ground, the telescope points straight up towards zenith. As the deck is tilted (telescope pointing elevation is decreased), the center of mass for objects mounted on the deck (the telescope, cryostat, and electronics enclosure) shifts, increasing the amount of torque necessary to hold the deck stationary. This imbalance results in an effective lower elevation limit for the telescope pointing, which depends on the deck loading. For QUIET Q-band operations, the lower elevation limit was 43 degrees.

Above the elevation axis is another bearing, which provides for rotations about the telescope line of sight. This third degree of freedom is very useful for a polarization experiment, as deck rotations can be used to shift the instrumental detector axis from \mathbf{Q} to \mathbf{U} at will. The deck drive operates by three friction drives, spaced equally around the bearing, that push against rollers. This arrangement is not nearly as powerful as the gears and motors

powering the azimuth rotations, so changes in deck position are only performed between scans with the telescope halted. The deck cable wrap provides for a similar range of motion as the azimuth wrap.

Control of the telescope is provided by commercial motion control boards (PMAC VME board, manufactured by Delta Tau). This board servos the torque applied by each motor to achieve the desired motion. Commanding the PMAC board is a program called Tracker, running on a VME computer in the PMAC crate. Tracker computes the trajectories needed for each axis to move to the desired position and supplies the moves to the PMAC board. Tracker can compute trajectories to slew the telescope to a specified position, follow the side-real motion of a source, or execute scans specified by the position, velocity, and acceleration at discrete steps in time. This last mode of operation is used for QUIET observing.

Running on the same computer as Tracker, other programs collect data from the telescope encoders, which specify the positions of the three axes, the motors and servos, and a weather station, which monitors temperature, wind, and humidity at the site. These data are packaged into frames and streamed to CBI Control, which records them to data files.

A large dome, made of fabric stretched between metal bars, can be closed over the entire mount for protection during bad weather. The dome is also deployed to provide shelter while repair work is performed on the QUIET mount or receiver.

CHAPTER 5

OBSERVATION WITH THE 40 GHZ ARRAY

The QUIET telescope and Q-band receiver were deployed to the observing site, located in Northern Chile, in August 2008. After a period of installation and testing, science observations commenced in late October. Q-band operations continued for approximately nine months, until early June 2009, when the receiver was replaced on the telescope with the 90 element W-band receiver.

Details about the QUIET telescope site are found in §5.1. The scan strategy employed by QUIET to observe the CMB is described in §5.2. The regions of sky observed are discussed in §5.3. Finally, §5.4 details the events of the Q-band observing season, from deployment through decommissioning.

5.1 Observing Site

QUIET is located at the Chajnantor Observatory in the Atacama desert of Northern Chile. The site longitude and latitude are $67^{\circ}45'42.0''$ West, $23^{\circ}1'41.6''$ South. The site elevation is 5020 meters.

Before QUIET, the Chajnantor Observatory was the location for the Cosmic Background Imager (CBI, Padin et al. (2002)), which was used to make measurements of the CMB temperature and polarization anisotropy from 1999 through 2008. Much of the site infrastructure was inherited from CBI, including the telescope control room (heated and oxygen enriched), warm lab (heated and oxygen enriched), machine shop, as well as the telescope mount and dome. Support staff, some of whom were originally employed by CBI, handle maintenance and repairs of the telescope mount and Chilean logistics.

The Chajnantor Observatory is named after a nearby mountain, Cerro Chajnantor, and the Llano de Chajnantor, the large plain at 5000 meters elevation on which the observatory

is located. The flatness of the Llano de Chajnantor makes it a convenient place for access and construction, so it is also the home of the Atacama Pathfinder EXperiment (APEX) and the Atacama Large Millimeter Array (ALMA), currently under construction. Other nearby telescopes include the Atacama Cosmology Telescope (ACT), located on the mountain Cerro Toco, the Atacama Submillimeter Telescope Experiment (ASTE), and the NANTEN2 Observatory, both located in the valley Pampa la Bola. ALMA is a particularly large project and QUIET makes use of some of their infrastructure, including dirt roads leading down from the mountains and an internet connection.

The nearest town is San Pedro de Atacama, approximately a one hour drive down the mountain. Food and lodging in San Pedro are shared with observers for ACT and ASTE.

5.1.1 *Weather and Atmospheric Properties of the Atacama Desert*

The Atacama desert is a nearly ideal location for microwave observations. A combination of high altitude and extreme dryness result in excellent observing weather for most of the year. Figure 5.1 shows a histogram of the precipital water vapor (PWV) during the observing season, as measured by the APEX weather station¹. The lower panel of the figure shows the relationship between PWV and atmospheric transmittance, integrated over a typical bandpass for a QUIET Q-band polarimeter. The atmospheric transmittance as a function of frequency is calculated using the *am* Atmospheric Model software (Paine, 2004). Even more important for observing are the amplitudes of atmospheric fluctuations. These have been found to be acceptably low, within a reasonable range of PWV and wind speed.

Daily temperatures range from 0 to 5° C during the day and drop below freezing at night. A strong diurnal feature of the weather on the Chajnantor plain are winds that pick up early in the afternoon and continue until dusk. The afternoon wind frequently reaches speeds of 10 m/s and sometimes blows much harder. Due to concerns about damage to the

1. APEX weather data can be found at <http://www.apex-telescope.org/weather>

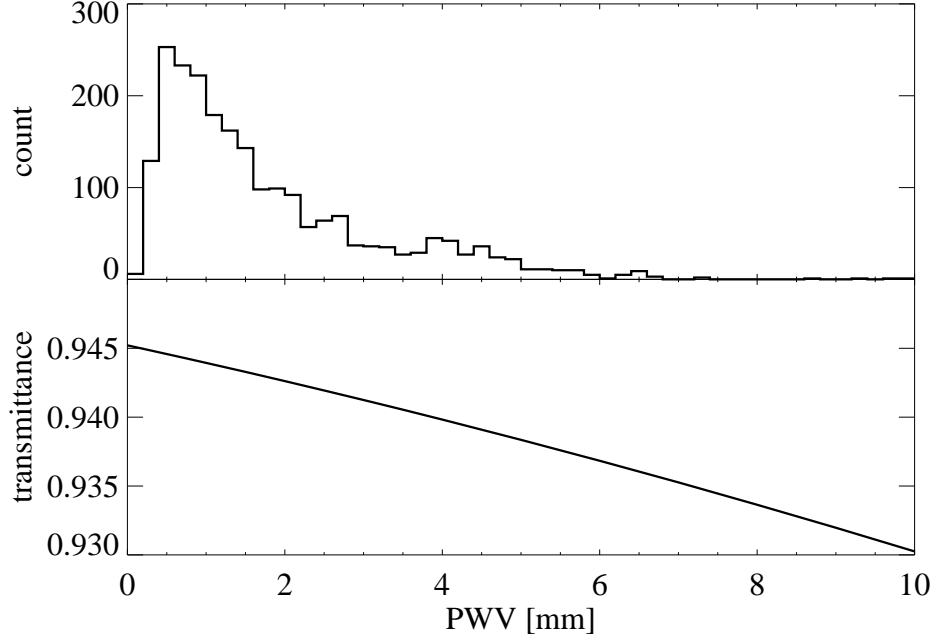


Figure 5.1: A histogram of the precipital water vapor (PWV) for the QUIET observing season is shown in the top panel. PWV values are measured by the APEX weather station. The lower panel shows the dependence of atmospheric transmission on PWV, averaged across a typical bandpass for a QUIET Q-band polarimeter. Atmospheric transmission is calculated using the *am* Atmospheric Model.

experiment, QUIET stops observing when wind speeds rise significantly above 15 m/s.

The weather degrades during January and February, which is the wet season in Atacama. This period is known as the Bolivian winter, despite being in the middle of southern hemisphere summer, because of the storms that come from the northeast (direction of Bolivia) and deposit snow at high altitudes. QUIET was able to continue observing through much of the Bolivian winter, but scans were frequently cut short so that the telescope could be stowed and the dome closed before the arrival of snow.

5.2 Scan Strategy

Scanning the telescope modulates the sky signal, converting the angular scales of the CMB to frequencies in the time ordered data. Since QUIET targets large angular scales, fast scanning is critical to ensure that the polarization modes of interest are measured on time scales shorter than the atmospheric and instrumental $1/f$ knee frequencies.

The scan used for QUIET involves periodic motion in the azimuth direction only. Both the elevation and deck axes of the telescope are fixed during a normal CMB scan. By scanning at fixed elevation, each module observes a roughly constant atmospheric signal. Despite the atmosphere being unpolarized, the $\mathbf{I} \rightarrow \mathbf{Q}$ leakage signal resulting from a change in elevation is easily large enough to overwhelm any CMB signal². Also, while the QUIET mount is capable of fast slewing in azimuth, it moves quite slowly in elevation or deck angle.

The full width of the azimuth scan is 15 degrees on the sky. For the 7 degree field of view of the array, this means that a large region (about 8 degrees across) in the center of the scan range will be observed by all modules. The azimuth range required to cover that distance depends on elevation, given by $15^\circ \cos el$. For a given azimuth range, there is a minimum scan period due to the azimuth velocity limit of 6 degrees per second and a requirement in the motor control software that the period be an even integer value in seconds. For scans near the 43° lower elevation limit the scan period is 10 seconds, while for scans over 65° the scan period is 20 seconds or higher.

As the telescope repeats its azimuthal pattern, the observed patch of sky moves through the field of view under sidereal motion. To observe a roughly square patch, the elevation and central azimuth of the scan are chosen to be at a location that is 7.5° (half a scan width) ahead of the patch center, in the direction of increasing right ascension. The scan duration is the amount of time it takes for the patch center to move by 15 degrees, which depends on

2. CAPMAP was a rare example of a CMB experiment with a scan that involved elevation changes. In that case, the scan motion was a circle about the North Celestial Pole, and the sinusoidal modes corresponding to elevation were filtered from the data.

its declination as $\text{duration} = \frac{15^\circ}{\cos \text{dec}} \left(\frac{24^h}{360^\circ} \right)$. Scan durations for QUIET CMB patches are included in Table 5.1.




Figure 5.2: Maps of the number of 50 Hz double demodulated samples per pixel for five CES in run 128 (small images, with a vertical scale running from 0 to 3400). The sum of all five CES is shown in the large image, with a scale from 0 to 6500. The area of each pixel is approximately 0.05 square degrees. All maps are centered at $\text{RA} = 12^h04^m$, $\text{Dec} = -39^\circ$, the center location of Patch 2a. Grid lines are drawn at 10 degree intervals in RA and Dec.

A series of scans tracking a single patch as it rises and sets is referred to as a run. The four to six segments of a run are called constant elevation scans, or CES. The CES, which form the fundamental unit of data for the analysis, are indexed by their run number and a sub-id that starts at 0 for the first CES of each run. Figure 5.2 shows the number of 50 Hz samples that fall into each 0.05 square degree pixel for 5 CES making up run 128. After the

entire run, an area of 240 square degrees is observed with at least $N_{max}/4$ samples, where N_{max} is the number of hits for the most frequently visited pixel in the map.

In addition to the azimuthal scan, each CES starts with a few brief calibrations. After the telescope has slewed to the position where the scanning will occur, it is slowly tipped in elevation three times, with an amplitude of 4 degrees. These elevation scans are used to track changes in the amplifier gains over long timescales (see §7.2). Next, the phase switches in all modules are turned off³ and 30 seconds of data is recorded to measure the detector diode offsets. These calibrations take 4–5 minutes at the start of each CES.

5.3 Patches Observed by QUIET

Table 5.1 lists the locations of the four CMB patches observed by QUIET, as well as two galactic patches that are scanned less frequently. The locations can also be seen in Figure 5.3, superimposed on a Q-band temperature map, as measured by WMAP (Jarosik et al., 2010b). The patches are selected to be in regions of low foreground emission, and distributed so that QUIET can spend as much time as possible on a small, deep region. The area observed for each patch is about 250 square degrees, or 1000 square degrees for the four CMB patches combined.

The useful range of patch declinations is set by telescope elevation limits. With a lower elevation limit of 43 degrees, it is necessary to wait until the patch center rises to about 37 degrees elevation before scanning begins. Above 70 degrees, the azimuth range of the scan is so large (22 degrees), that it can cause trouble for the telescope. Because of these two constraints, patches are selected with declinations such that they transit at elevations near 70 degrees. This ensures that they spend the maximum amount of time above our lower elevation limit, but at sufficiently low elevation that they can be scanned efficiently.

3. The phase switches are turned off rather than the LNAs, because the LNAs have a significant settling time after turning on.

Table 5.1. Patches observed by QUIET

Patch	R.A.	Dec	CES Duration	Type
Patch 2a	$12^h04^m00.00^s$	$-39^\circ00'00.0''$	1:17	CMB
Patch 4a	$05^h12^m00.00^s$	$-39^\circ00'00.0''$	1:17	CMB
Patch 6a	$00^h48^m00.00^s$	$-48^\circ00'00.0''$	1:30	CMB
Patch 7b	$22^h44^m00.00^s$	$-36^\circ00'00.0''$	1:14	CMB
Patch Gb	$16^h00^m00.00^s$	$-53^\circ00'00.0''$	1:40 ^a	Galactic
Patch Gc	$17^h45^m37.20^s$	$-28^\circ56'10.2''$	1:09 ^a	Galactic

^aFor galactic patch observations, various scan sizes were used, so many scans have different duration than the value listed here.




Figure 5.3: The four CMB patches and two galactic patches observed by QUIET, shown in equatorial coordinates and superimposed over a full sky temperature map at Q-band from the WMAP seven-year data release (Jarosik et al., 2010b). Note that the color scale is highly saturated in the galactic plane where the temperature signal can be as large as 100 mK.

Since QUIET can observe day and night, patches should be spaced uniformly in right ascension so that there is one patch available at all times. Patches with declination near -39° spend about seven hours per day above the lower elevation limit, so four patches is the minimum number needed to meet this criteria. After adjusting patches to avoid galactic foregrounds, there is a gap of time between when Patch 2a sets and Patch 7b rises. Two galactic patches, labeled Gb and Gc, can be observed during this period. Patch Gc targets the galactic center; Patch Gb lies in the galactic plane slightly away from the galactic center. Figure 5.4 shows the pattern of patch availability for a typical observing day. Also shown in Figure 5.4 is the path of Tau A (Crab Nebula), which serves as an important calibration for QUIET (see §7.2). With a declination of $+22^\circ$, Tau A spends less than two hours per day above the lower elevation limit. Tau A calibration scans are scheduled for times when Patch 4a is transiting at elevations above 70 degrees.




Figure 5.4: Elevation of QUIET observation targets as a function of time for an example 24 hour period. CMB patches shown are Patch 2a (dark blue), Patch 4a (light blue), Patch 6a (green), and Patch 7b (red). Galactic patches Gb (dotted line) and Gc (dashed line) are available for observation during the gap between Patches 2a and 7b. Tau A, an important calibration source, is denoted by the dot-dashed line.

5.4 Events of the Q-band Observing Season

The Q-band observing season officially started with run 125 on October 28, 2008. At this point, the decision had been made to use double-demodulation for general observing, scanning scripts had been tested, and the electronics enclosure temperature regulation was functioning well. The final CMB run of the season was run 1853 on June 13, 2009. Between those two dates, QUIET carried out 905 runs (2430 CES) targeting the four CMB patches, as well as 276 runs focused on galactic patches and numerous calibration observations.

Significant events from the observing season, such as equipment failures, are listed below, in chronological order.

- Only three days after the start of the season, a failure of one of the azimuth drives shut down operations for over a week (October 31 through November 7). The problem was eventually traced to a broken motor encoder that is used by the servo circuit to control the motor torque.
- Patch 7b observations did not begin until November 13. Before then, QUIET was observing a patch at a slightly different location. The choice to move to the patch 7b location was made to improve the overlap in sky coverage between QUIET and other CMB polarization experiments.
- For runs 329 through 341, the electronics configuration was altered by swapping the positions of the two preamp boards. The purpose of this switch was to study Type-B glitching (see §6.4), the cause of which was not understood at the time. When the test failed to eliminate the glitching, the preamps were restored to their original configuration.
- A new generator was installed at the site and first activated on December 19. Before that installation, QUIET had been running for several weeks without a backup.

- QUIET was shut down from December 29, 2008 through January 4, 2009, so that observers and staff could enjoy the holiday.
- Shortly after restarting in January, the telescope mount failed, suspending observations from January 4 through 10. The problem was traced to a bad signal line in the motor control system. Near the end of this period of downtime, the AIB cables connecting the bias electronics to the cryostat were replaced with a set of unshielded spare cables. As with the preamp board swap, the goal of this test was to understand the source of Type-B glitching. The alternate AIB cables remained installed after observations resumed, for runs 563 through 593. This failed to solve the glitching problem, so the cables were switched back. After reverting to the original AIB cables, one of the gate bias lines for module RQ-16 was no longer connected. This connection was unreliable for earlier times in the season as well, and RQ-16 was ultimately dropped from the data analysis entirely.
- On January 28, it was discovered, during maintenance on the mount, that the wrong sized bolts were being used to hold the deck axis encoder in place. The bad bolts meant that the encoder was not screwed down tightly and could slip, causing non-physical jumps in deck position to show up in the data. This problem is covered in more detail in §6.8. At this time, the bolts were replaced and the deck encoder was screwed down tightly.
- On February 19, heaters were installed in the box enclosing the AIBs. This change was made to address occasional failures of the phase switch protection circuitry (particularly for RQ-12) that occurred in high humidity conditions. Additional heaters were added on March 13.

CHAPTER 6

INSTRUMENTAL PROBLEMS AND IMPERFECTIONS

This chapter covers problems that arose in the operation of QUIET and the solutions developed to address them. None of the topics discussed here resulted in a serious degradation of performance.

Instrumental leakage, which couples unpolarized signals into the demodulated data, is described in §6.2. This effect is strongly mitigated by the excellent parallactic angle uniformity of the QUIET scan strategy, but it still contributes to $1/f$ noise in the demodulated data.

Two categories of glitching were identified in the timestream data after the start of observations. “Type-A” glitching, described in §6.3, is still poorly understood, but it has no significant impact on the polarization data. A more serious source of glitches, labeled “Type-B” glitching, was eventually traced down to non-linear behavior in the QUIET ADC boards. The cause of Type-B glitching, and a method for correcting it, are described in §6.4. Another type of unaccounted for noise, distinct from the glitches, is described in §6.5.

Having separate data acquisition systems for the QUIET receiver and the mount encoders led to a small, stable timing offset between the radiometer and pointing data, described in §6.6. Beam sidelobes that are not intercepted by the lower groundscreen result in contamination signals, described in §6.7. Finally, a mechanical problem in the encoder for the deck axis of the mount caused random slips in the deck position. This was discovered and fixed during the observing season, but the deck slips contribute additional pointing uncertainty to the early season data, as described in §6.8.

6.1 Broken Modules and Diodes

The Q-band array contains 17 modules configured as polarimeters, each with four detector diodes. Of the 17 modules, one entire module, plus single detector diodes on two others, did not function acceptably during the observing season and are not included in the data analysis.

Module RQ-16 stopped functioning on January 15, 2009, after the AIB cables were reverted back to the original configuration (see §5.4). The cause of the failure was a broken connection on one of the gate bias lines, which shut off an amplifier in one leg of the module. While this failure was permanent after the cable change, it also appears intermittently earlier in the season. Due to concerns about unreliable biasing, no RQ-16 data is used for the analysis.

The U2 detector diode on module RQ-08 did not work at any point during the observing season. When the preamp is used to apply bias voltage to this channel, it behaves like an open circuit, suggesting that one of the bond wires connecting the diode anode or cathode to the module pins may have broken.

The Q1 detector diode on module RQ-04 can be biased and read out normally, and it detects polarization signals such as the moon, but it suffers from large spurious noise at frequencies of a few Hz. This noise is different and much more severe than the narrow lines that are seen in the noise spectrum of many detector diodes at frequencies of 6 Hz or higher (see §6.5). The most probable cause for the excess noise in RQ-04 Q1 is if the diode cathode is shorted to the module case ground, which occurred occasionally in early lab testing of W-band modules.

All together, six detector diodes out of 68, or 9% of the polarization channels, are unusable for the Q-band array. This failure rate was not high enough to require any drastic measures, such as attempting to replace modules after deployment.

6.2 Temperature to Polarization Leakage

Double demodulation, described in §3.5.1, effectively eliminates $\mathbf{I} \rightarrow \mathbf{Q}$ leakage caused by phase switch imbalance. The most significant remaining causes of leakage for QUIET are two systematic effects of the septum polarizers (Bannister et al., 2006).

The first effect, known as differential loss, occurs if the polarizer has different transmission amplitudes for two orthogonal linear polarizations. Input radiation that is linearly polarized along the X direction consists of equal parts left and right circular polarization, and the same is true for an input signal polarized along the Y direction. Differential loss describes the situation where input signals of equal amplitude, but orthogonal linear polarization, result in output signals with different amplitude. In the QUIET polarizers, the symmetry of the device is broken by the metal septum, creating a differential loss between input signals polarized parallel and perpendicular to the septum. Since the symmetry is unbroken for signals linearly polarized along the axes oriented at 45° with respect to the septum, the differential loss is significant only for Stokes parameter \mathbf{Q} , in the instrument basis.

The second important systematic, known as crosstalk, occurs because the two output ports of the septum polarizer are imperfectly isolated. A fraction of the signal input to a module is reflected off of the first stage amplifier back into the polarizer’s output port. Most of this reflected radiation is sent harmlessly out the feedhorn, but part of it couples into the module’s other input port. Since the QUIET module detects polarization by measuring the correlations between the two inputs, this systematic can fake a polarized signal. Unlike the differential loss, which primarily affects only one of the Stokes parameters, crosstalk generally affects both, with dependence on the phase shift of the crosstalk signal.

Additionally, the low noise amplifiers have a tendency to broadcast noise out of their input. The polarizer crosstalk systematic turns this noise into a polarized offset, but one that is not proportional to the input temperature.

Vector Network Analyzer (VNA) measurements have been made for both of these system-

atics (Nixon, 2008). For the 40 GHz septum polarizers, the typical leakage due to differential loss and crosstalk are -20 dB and -35 dB, respectively. The much stronger differential loss primarily affects Q diodes, which results in measurably larger leakage coefficients and $1/f$ noise for Q diodes than for U diodes. At 90 GHz, the crosstalk systematic is dominant, and the asymmetry between Q and U diodes is much less pronounced.

6.2.1 Suppression of the Temperature Leakage Signal Through Uniform Parallactic Angle Coverage

Leakage coefficients measured from elevation nods (see §7.2) or source observations are about 1% for Q diodes and 0.5% for U diodes, but some modules (RQ-05 and RQ-14) have Q diode leakage coefficients as high as 2%. Based on the relative amplitudes of the power spectra, contamination in the polarization signal equal to 1–2% of the temperature anisotropy could be disastrous for any attempts to measure B -modes. Fortunately, the QUIET scan strategy is highly effective for suppressing the leakage signal, when polarization maps are summed over the entire observing season.

The type of $\mathbf{I} \rightarrow \mathbf{Q}$ leakage caused by OMT systematics is known as monopole leakage, because it creates a fake polarization signal that is proportional to the observed temperature, as opposed to a higher order moment of the temperature field. The orientation of this polarization signal is fixed in the instrument frame of reference; for QUIET, it is closely aligned with the Q diodes, since they suffer more from the effect. When the same patch of sky is observed with the instrument at different orientations, the monopole leakage signal will tend to cancel out.

For every CES in the observing season, the parallactic angle of the orientation axis (defined in §7.1.1) is calculated and the patch 2a results are histogrammed in Figure 6.1. Since polarization is periodic in 2ψ , two times the parallactic angle is used. The uniformity of the distribution, with no strongly favored orientations, is thanks to the rotation of the patch

on the sky as it rises, transits, and sets, and the many deck angles used for observation throughout the season. For an example, run 1618, which consists of 6 CES targeted on patch 2a, covers a range of 155° of parallactic angle from the first CES to the last, due to sky rotation alone.




Figure 6.1: Histogram of two times the parallactic angle of the array orientation for patch 2a CES. For each CES, the orientation is calculated based on the mean azimuth, elevation, and deck angle during the CES.

A rough estimate of the monopole leakage suppression can be obtained by averaging $\cos 2\psi$ or $\sin 2\psi$ over all CES. These averages are found in Table 6.1. Even for patch 7b, which was infrequently observed, the scan strategy has sufficiently uniform coverage to cancel the leakage signal to 5–6% of its original value. More accurate estimates can be made using simulated data (Kusaka & Bischoff, 2007), which confirm this suppression.

For the sensitivity level of QUIET, monopole $\mathbf{I} \rightarrow \mathbf{Q}$ leakage has a negligible effect on the EE or BB power spectra, and it is ignored in the analysis. The effect of leakage still exists in the form of increased $1/f$ noise. Quadrupole leakage, which transforms like polarization under changes of the array orientation and therefore doesn’t cancel out over the observing season, could be sourced at a low level by the QUIET telescope. Efforts to measure this very small effect are inconclusive so far, but simulation studies indicate that the effects would be

Table 6.1. Parallactic angle of the array orientation averaged over all CES for each of the QUIET patches.

Patch	$\langle \cos 2\psi \rangle$	$\langle \sin 2\psi \rangle$
Patch 2a	-0.0282	-0.0092
Patch 4a	0.0198	-0.0097
Patch 6a	0.0113	0.0481
Patch 7b	0.0527	-0.0636

negligible.

6.3 Type-A Glitching

A source of noise in average data, known as “Type-A glitching”, was discovered in QUIET data during analysis of sky dip calibrations (see §7.2). This noise manifests as random jumps in average data between two or more distinct states, separated by ~ 1 mV offsets. An example of Type-A glitching can be seen in Figure 6.2.

After being discovered in the Q-band data, this noise was also found in lab tests of the W-band array (Bischoff & Tajima, 2008). The noise appears only for particular detector diodes and persists even when the LNAs and phase switches are turned off. Adjusting the detector diode bias does not change the amplitude of the jumps between states, but it does change the fraction of time spent in each state. If the preamp from a glitching channel is used to bias a resistor instead of the detector diode, then the noise disappears, suggesting that the problem originates on the detector diode or as an interaction between the diode and the bias circuit.

Fortunately, Type-A glitching has almost no impact on demodulated data, as can be seen in the bottom panel of Figure 6.2. If the glitch seen in average data corresponds to a simple offset in the raw data stream, without any corresponding change in gain or noise, then




Figure 6.2: Averaged (top panel) and demodulated (bottom panel) data from RQ-12 diode U2, recorded with the phase switches unbiased so that little RF power is incident on the detector diode. Type-A glitching is clearly visible in the average data, but there is no sign of it in the demodulated data. Note that the vertical scale is zoomed in by a factor of ~ 100 for the demodulated data. When the module is turned on, the typical offset for the average data is 100–200 mV, with much larger noise, so the effect is less noticeable.

it can affect the demodulated data only if the change is sufficiently abrupt that the offset changes significantly from one half of a 4 kHz demodulation period to the other. There is no correlation between average and demodulated data for channels with Type-A glitching, but the average data does correlate with the derivative of demodulated data in channels with strong Type-A glitching, confirming this hypothesis.

The net effect of Type-A glitching on QUIET polarization data is to contribute a small, symmetric non-Gaussian tail to the instrumental noise. Since this tail is barely noticeable with the LNAs and phase switches turned off, it is safe to treat it as totally negligible once noise from the RF circuitry is added. It is important to account for Type-A glitching when using the average data for calibration (see §7.2) or data selection (see §8.2).

6.4 Type-B Glitching

A second type of spurious noise, known as “Type-B” glitching, was also discovered early in the Q-band observing season. Type-B glitching, which affects roughly 10% of the data, manifests as greatly increased noise in demodulated data for individual detector diodes that appears at the start of a CES and persists for the entire CES. This noise is correlated to the averaged data, as if the $\mathbf{I} \rightarrow \mathbf{Q}$ leakage of that channel suddenly increased by a factor of 10 or more. Since each CES is preceded by an elevation change, initial theories about the cause of Type-B glitching were based on the idea that a loose cable was making intermittent contact that depended on the telescope elevation. However, these theories struggled to explain how the noise could affect one detector diode in a module but not the others.

Eventually, the cause of the Type-B glitching was traced to a very different problem, originating on the ADC board. The transfer function of the analog-to-digital converters, specifying the output bit value for each input voltage level, should be purely linear. Instead, the transfer function contains step-like discontinuities, as seen in Figure 6.3, which are spaced at regular intervals but have different heights, that can be positive or negative. The signal into the ADC can be thought of as two data streams, one for each phase state, which have slightly different mean values, corresponding to the size of the demodulated signal. Figure 6.3 demonstrates that, when the average level of the two distributions crosses the discontinuity, the apparent size of the demodulated signal increases (or decreases, if the discontinuity has negative step height). If the average level for a detector diode is near the discontinuity, then any drifts in that average level induce large changes in the demodulated signal, creating the appearance of large $\mathbf{I} \rightarrow \mathbf{Q}$ leakage. Since the atmospheric load temperature is a function of elevation, a move to new elevation changes the average level on each detector diode, moving some diodes near a glitch location while moving others out of the way.




Figure 6.3: Illustration of the mechanism for Type-B glitching. The central panel shows the ADC transfer function between input voltage and output bit value, featuring a step discontinuity. Noise distributions for three input signals are shown in the bottom panel, with the two 4 kHz phase switch states represented in red and blue. The output signals shown in the left panel result from applying the ADC transfer function to the inputs. Barred points along the bottom and left of the central panel show the results of averaging and demodulating the input and output data. The demodulated signal is boosted when the averaged signal approaches the glitch location, leading to characteristic Type-B behavior.

6.4.1 Correcting Type-B Glitching

Type-B glitching affects a sizable fraction of the data, but it is possible to correct it. An iterative method is used, based on an understanding of the cause of the effect and assumptions about the underlying statistical distribution of the input data. The derivation for the correction presented here is based on Kusaka (2010).

Given an ADC transfer function of the form

$$y = x + \Delta f(x) \tag{6.1}$$

$$\Delta f(x) = \begin{cases} -\frac{A}{2} & \text{for } x < x_0 \\ 0 & \text{for } x = x_0 \\ +\frac{A}{2} & \text{for } x > x_0 \end{cases} \tag{6.2}$$

and an input Gaussian distribution with mean $\langle x \rangle$ and standard deviation σ , we can calculate the mean of the output distribution.

$$\langle y \rangle = \langle x \rangle + \frac{A}{2} \operatorname{erf} \left(\frac{\langle x \rangle - x_0}{\sqrt{2}\sigma} \right) \tag{6.3}$$

From the measured average and demodulated samples, we can construct the output mean values, $\langle y_0 \rangle$ and $\langle y_1 \rangle$, of the two 4 kHz states. For each state, we can iteratively solve for the mean of the input distribution using Newton's method. Defining $s \equiv \langle x \rangle - \langle y \rangle$, equation (6.3) becomes

$$h(s) \equiv s + \frac{A}{2} \operatorname{erf} \left(\frac{s + \langle y \rangle - x_0}{\sqrt{2}\sigma} \right) = 0 \tag{6.4}$$

Starting with the initial guess $s = 0$, the first three steps in the iteration are

$$0^{th} \text{ order:} \quad s_0 = 0 \quad (6.5)$$

$$1^{st} \text{ order:} \quad s_1 = 0 - \frac{0 + \frac{A}{2} \operatorname{erf} \left(\frac{\langle y \rangle - x_0}{\sqrt{2}\sigma} \right)}{1 + \frac{A}{\sqrt{2\pi}\sigma} \exp \left(-\frac{(\langle y \rangle - x_0)^2}{2\sigma^2} \right)} \quad (6.6)$$

$$2^{nd} \text{ order:} \quad s_2 = s_1 - \frac{s_1 + \frac{A}{2} \operatorname{erf} \left(\frac{s_1 + \langle y \rangle - x_0}{\sqrt{2}\sigma} \right)}{1 + \frac{A}{\sqrt{2\pi}\sigma} \exp \left(-\frac{(s_1 + \langle y \rangle - x_0)^2}{2\sigma^2} \right)} \quad (6.7)$$

This process converges rapidly, achieving a sufficiently accurate result at 3^{rd} order. Applying the iterative procedure to $\langle y_0 \rangle$ and $\langle y_1 \rangle$ yields $\langle x_0 \rangle$ and $\langle x_1 \rangle$, which are summed and differenced to recover the corrected average and demodulated samples.

Several parameters are needed to implement the Type-B correction. The height, A , and location, x_0 , for each step discontinuity on each ADC channel are measured directly by inputting a sine wave source to the ADC and adjusting its offset voltage to explore the full range of the transfer function. Also necessary for the correction is the standard deviation, σ , of the input signal. This noise level is difficult to estimate for two reasons. First, the data available are 100 Hz demodulated samples, but the correction requires the standard deviation of the raw 800 kHz samples used for demodulation. Accounting for the blanking mask, each 100 Hz sample accumulates 6880 samples of 800 kHz data. However, low-pass filters on the preamp filters correlate these samples, so the noise at 800 kHz is somewhat less than $\sqrt{6880}$ times the 100 Hz noise. This exact factor can be calculated from the preamp filter function, but instead the factor is found by applying the Type-B correction to a sample of glitching data and searching for the value that results in a linear correlation between average and demodulated data across the discontinuity. For an example, see the left hand panel of Figure 6.4, which has a highly non-linear correlation in the uncorrected case. The second complication is that the measured data samples have excess noise, due to

the Type-B glitching itself, so they provide an overestimate to the input noise; this problem is solved by employing another iterative process, like the one described above, where the noise level is recalculated after applying the correction and then used to determine a better correction.

Figure 6.4 shows an example of the Type-B correction applied to glitching data. The uncorrected (black) data points in the left hand panel show a non-linear relationship between average and demodulated data, which is the signature of Type-B glitching. After the correction is applied (red points), the relationship is linear, with a slope equal to the normal $\mathbf{I} \rightarrow \mathbf{Q}$ leakage for this detector diode.




Figure 6.4: The right hand panel shows double demodulated data for module RQ-12 diode Q2 from CES 477.0 before (black) and after (red) applying the Type-B correction. The left hand panel shows the correlation between the double demodulated data and averaged data. The peak shape seen in the left hand panel indicates the location of the discontinuity in the ADC transfer function.

6.5 High Frequency Noise

Many detector diodes suffer from excess noise visible as narrow lines in their power spectra at frequencies $\gtrsim 6$ Hz. The source of this noise has not been conclusively found. One possibility is that it is due to pickup of the electromagnetic environment by the bias cables.

In particular, the detector diode bias lines are unfiltered and the diodes themselves act as rectifiers across a wide range of frequencies.

Fortunately, this noise is almost perfectly common mode between paired diodes in a module. Figure 6.5 shows an example of a spike in the power spectrum of module RQ-11, diode *Q1*. When that diode is differenced with the *Q2* diode, weighted by their inverse variances as would be done in the map making process, the spike disappears entirely. This indicates that the spurious noise will automatically subtract out of the polarization maps, as long as both diodes from the pair are included.




Figure 6.5: Zoom in on a small region of a noise power spectrum from CES 128.0. The spectrum for module RQ-11, diode *Q1* is shown in red, featuring a large spike around 8.645 Hz. The spectrum for the RQ-11 *Q1* – *Q2* diode difference stream (weighted by inverse variance) is shown in black. No spike is seen for the difference stream, indicating that the source of the spike is common mode between the two diodes.

There are two diodes in the array, RQ-04 *Q2* and RQ-08 *U1*, that never have a paired diode included in the analysis. Furthermore, the data selection, described in §8, sometimes accepts only one diode of a pair for certain CES. Rather than requiring all cuts to act on diode pairs, it is simpler to just filter out the affected frequencies. A strong filter is applied to the time ordered data to remove all power above 4.7 Hz, as described in §9.1.1.

This filter has an acceptably small effect on the CMB signal. For the fastest azimuth scans, which have a full period of 10 seconds, each horn moves across 7.5° on the sky in 5 seconds. At an average velocity of $1.5^\circ/\text{second}$, the high frequency cutoff removes power on scales less than a third of a degree, which corresponds to multipoles $\gtrsim 550$. Due to its large beam size, the QUIET Q-band array has little sensitivity to these scales anyway, so it is not a big loss to filter them out entirely.

6.6 Timing Offset Between Telescope and Receiver Data

Telescope encoder data and radiometer data are acquired by separate hardware and software systems. Both systems timestamp frames of data using an IRIG-B timing signal from the GPS receiver in the control room and the clocking for the receiver ADC boards is generated from that timing signal as well.

However, scanned source observations, such as the moon observation shown in Figure 6.6, reveal that there is a timing offset between the two data streams. The value of this offset can be found by lining up the left-going and right-going scans. For Q-band observations, the receiver data lags behind the telescope pointing data by 2.5 samples (25 milliseconds). The half sample component of the lag can be explained by the definition of sampling time used in the two data acquisition systems. The telescope encoders are sampled instantaneously and stamped with that time. The radiometer data is integrated over a 10 millisecond period and the time at the end of the sampling interval is recorded as the timestamp. The origin of the remaining timing offset is unknown.

This lag is very stable throughout the observing season and can be corrected in the data analysis. The correction, shown in Figure 6.7, involves shifting the radiometer data relative to the timestamps and telescope encoder data. Because the timing offset is a non-integer number of samples, the timestamps and encoders are interpolated to line up with the radiometer samples.




Figure 6.6: Moon signal as a function of azimuth for adjacent right-going (red) and left-going (blue) slews. With no timing offset correction (solid lines), the apparent moon center is shifted for each direction. After applying the timing offset correction (dashed lines), the two halves of the scan line up well.

6.7 Beam Sidelobes

The QUIET groundscreen is designed to intercept rays that miss the primary or secondary mirrors, ensuring that they terminate on a stable absorbing surface. However, two particular sidelobes of the radiometer beams do escape the groundscreen and can introduce contamination signals when the scan across the ground or other hot objects.

The significant sidelobes for QUIET are illustrated in Figure 6.8. The first sidelobe to be discovered is a triple reflection sidelobe. Some rays that reflect from very low on the secondary mirror are directed to the primary, but then graze the top of the secondary mirror for a third reflection. The geometry of this effect is such that the affected rays pass through the groundscreen aperture at an angle 50° from the main beam. The triple reflection sidelobe was expected from physical optics simulations of the QUIET telescope, with a response about 50 dB less than the main beam (Holler, 2007). To block this sidelobe, an additional section of groundscreen was designed, consisting of a telescoping tube that extends from the aperture of the groundscreen box, shown in gray in Figure 6.8. Unfortunately, the upper groundscreen

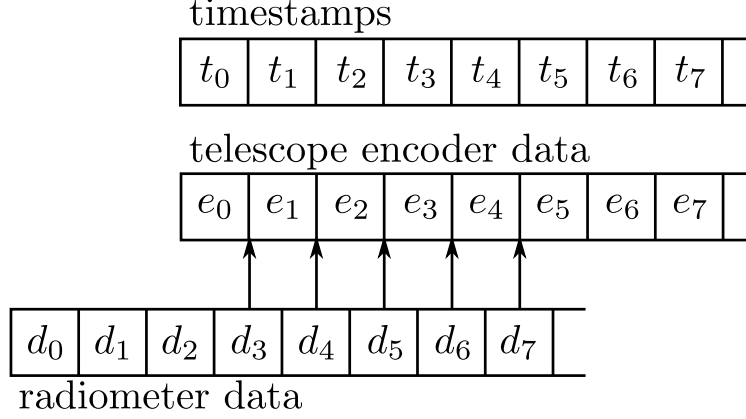


Figure 6.7: Illustration of the timing offset correction.

tube was not installed in time for Q-band observations.

The second sidelobe is oriented 60° from the main beam in the opposite direction of the triple reflection sidelobe. The source of the second sidelobe has not been definitively shown, but dedicated measurements made with the W-band array suggest that it could be due to rays passing through the groundscreen aperture and directly into the receiver window without reflecting off the telescope mirrors (Sanders, 2010). This sidelobe appears for polarimeter modules RQ-00, RQ-01, and RQ-02, as well as the TT modules, RQ-17 and RQ-18, but it has not been observed for the other modules in the array.

Both sidelobes can be observed from CMB observations where the sidelobe position is near the sun, which is bright enough to create a signal that is easily visible over single CES timescales. As the telescope scans in azimuth, the sidelobe sweeps back and forth across the sun. Over the course of the CES, the motion of the sun through the sky will eventually cause it to pass out of the region probed by the sidelobe, resulting in a distinct features in the maps made from that data. Figure 6.9 shows examples of maps with sun contamination from each sidelobe. The triple reflection sidelobe shows up as a sharp spot in the left hand map. The second sidelobe produces the broad striped structure in the right hand map.

Constant elevation scans where the sun happens to line up with the sidelobes are fortu-




Figure 6.8: Not to scale illustration of the QUIET telescope, groundscreen box, and sidelobes. The main beam is indicated with a solid line. The triple reflection sidelobe is shown by the dot-dashed line. The second sidelobe is shown as a dotted line. The upper groundscreen, which was not installed for Q-band observations is shown in gray.

nately rare and can be cut during data selection, as described in §8.5. The moon, which has similar angular size to the sun, is dimmer by two orders of magnitude and doesn't pose a risk of contamination via the sidelobes. However, a significant signal is generated when the telescope is aligned so that one of the sidelobes points to the ground¹. As in the case of the sun contamination, the ground signal is modulated as the telescope scans in azimuth. Unlike the sun, the ground does not move during the course of the CES, so the contamination

1. The back-to-back geometry of the two sidelobes ensures that at most one of them will hit the ground for any observation.

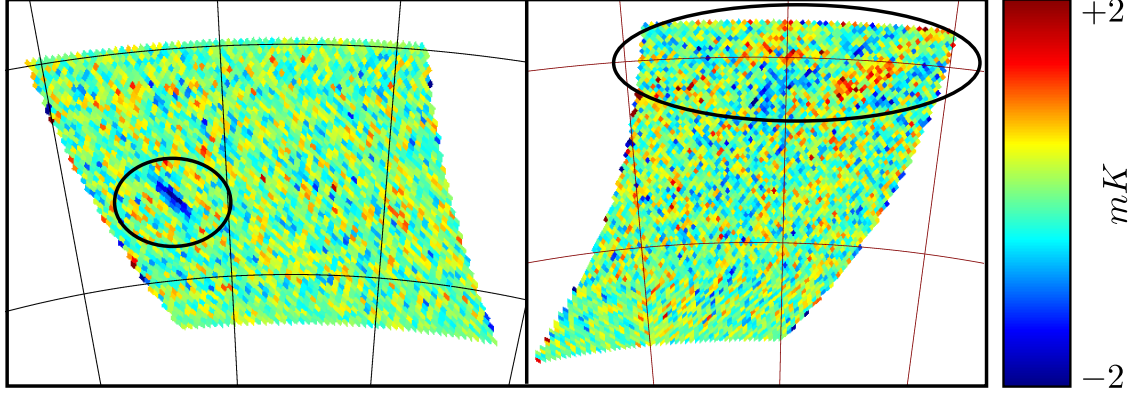


Figure 6.9: On the left is a map of demodulated data from module RQ-02, diode Q2 for CES 1060.0. The circled spot is an image of the triple reflection sidelobe illuminated by the sun. On the right is a map of demodulated data from module RQ-02, diode U1, for CES 1531.0. The red and blue structure along the top edge of the map is the result of the sun passing through the second sidelobe.

signal is an approximately constant function of azimuth and can be removed by filtering, as described in §9.1.3.

6.8 Deck Encoder Slips

As noted in §5.4, a problem with the deck encoder was discovered in January 2009. The telescope had been stalling frequently, with errors related to the deck axis. During the inspection, it was discovered that the encoder used to read the deck position had been fastened with the wrong size of screws, which allowed it to work its way loose. This led to telescope stalls, as motions of the encoder were interpreted by the telescope control software to be sudden dangerous jumps in the deck position.

Besides causing telescope stalls, the loose encoder also means that the deck position data is unreliable for scans occurring before the fix. The magnitude of the deck offsets can be estimated for moon observations where the moon is viewed by many modules. The relative pointing of different modules in the array gives a true measure of the array orientation, which can be used to check the reported deck angle. In practice, the average (total power) data for

a moon observation is analyzed for each module using the standard QUIET pointing model (see §7.1). A pointing offset is calculated between the observed moon center position and the value predicted by ephemeris. More details on the fitting procedure can be found in §7.5.4. Other terms in the pointing model contribute offsets that are, to first order, common to all modules but deck angle offsets create scatter in the pointing offsets from horn to horn. The true deck angle for the scan is calculated by finding the deck offset value that minimizes this scatter.

Figure 6.10 shows the deck offset as a function of time, measured from the 18 moon observations listed in Table A.1. Error bars are estimated roughly based on the typical scatter in recovered moon center positions (0.01°) and the number of modules used to find the deck offset in each case. The offset value is constant after late January, when the encoder was fixed. The dashed line indicates the deck offset correction applied in the QUIET pointing model. Instead of attempting to perfectly track the observed deck error, the model uses four discrete correction values for different blocks of time (it is difficult to see from the figure, but the mount model line has a value of zero for dates in October).




Figure 6.10: Measurements of the deck offset from 18 moon observations during the Q-band season. The dashed line indicates the deck correction used in the mount model.

CHAPTER 7

CALIBRATION

The raw data used in the CMB polarization analysis are the outputs of the telescope axis encoders and the demodulated detector diode timestreams. For the encoder samples, a model of the telescope pointing, inspired by physical and mechanical considerations, is used to calculate the correct astrometric coordinates. The pointing model is described in §7.1. The demodulated voltages are converted to measurements of \mathbf{Q} and \mathbf{U} , with units of μK , using the gain model, which is covered in §7.2.

Additional calibration tasks include modeling the atmospheric and instrumental noise, which is described in §7.3, and measuring the detector beams, which smooth out the sky signal (see §7.4). Section 7.5 contains a detailed discussion of calibration observations targeting the moon. This source is useful due to its extreme brightness, but it must be treated carefully to minimize systematic errors in the derived parameters.

7.1 Pointing Model

The pointing model converts the azimuth, elevation, and deck angle samples recorded by encoders on the telescope axes into accurate pointings for each module in the array. The model involves three steps. First, the pointing is calculated for module RQ-09, the central module in the array. This calculation involves the mount model, which describes imperfections in the mount and telescope alignment. Next, the RQ-09 pointing is transformed to the desired coordinate system, usually equatorial or galactic coordinates. Finally, the positions of the other modules in the array are calculated, using the array model.

Since QUIET measures polarization, it is important to also know the orientation of the detectors on the sky. For this purpose, an orientation vector is defined, perpendicular to the boresight pointing. Mount model corrections, coordinate transformations, and array offsets

are applied to the orientation vector as well as the pointing vector, providing a reference orientation on the sky relative to which the polarization sensitivity axis for each detector diode is defined.

7.1.1 *Definition of Pointing and Orientation Vectors*

The pointing and orientation for any module can be represented by a pair of orthogonal 3-vectors with unit length. The pointing vector, \vec{p} , indicates the spot on the sky being observed while the orientation vector, \vec{s} , defines a reference orientation in the plane perpendicular to \vec{p} . Sometimes it is useful to specify the third vector in the basis, $\vec{r} \equiv \vec{p} \times \vec{s}$.

Calculation of the array pointing starts with uncorrected pointing and orientation vectors, derived from the encoder values

$$\vec{p} = \begin{bmatrix} \cos az \cos el \\ -\sin az \cos el \\ \sin el \end{bmatrix} \quad (7.1)$$

The convention used to define \vec{s} is that, when the deck encoder reads zero, the orientation vector points in the direction of increasing elevation. For increasing values of the deck encoder coordinate, the orientation vector rotates about \vec{p} in a clockwise sense:

$$\vec{s} = \begin{bmatrix} -\cos az \sin el \cos dk - \sin az \sin dk \\ \sin az \sin el \cos dk - \cos az \sin dk \\ \cos el \cos dk \end{bmatrix} \quad (7.2)$$

7.1.2 *Mount Model*

After calculating \vec{p} and \vec{s} from the encoder values, the mount model is used to transform those vectors into the pointing and orientation vectors for module RQ-09. The five corrections

applied are, in order, the deck encoder offset, deck flexure correction, correction for the tilt of the azimuth axis, correction for non-perpendicularity of the elevation axis, and the array collimation offset. The mount model corrections are native to the horizon coordinate system, so they are applied before converting to equatorial or galactic coordinates.

7.1.3 Deck Encoder Offset

Mechanical problems with the deck axis encoder are discussed in §6.8. To correct the pointing, different offsets are applied to the deck encoder values for four distinct periods during the season. Unlike the other mount model corrections, which are implemented as rotations of \vec{p} and \vec{s} , the deck offset is simply subtracted from the encoder values before those vectors are calculated.

These offsets were determined by observing the pattern of array pointing on the sky from moon observations. While there is evidence that the deck encoder slips might have been more common, the moon scans are not frequent enough to locate every shift. Instead, the mount model only attempts to correct the long time scale behavior of the deck offset, which varies by 1.88° from its most positive to most negative values.

Note that the model of array pointing offsets is derived from early season data, so the deck encoder offset at those times is zero by definition. Any offset to the deck encoder position that was constant throughout the entire season would be absorbed into the array pointing model. It is only because this offset changes throughout the observing season that it has to be treated separately.

7.1.4 Deck Flexure Correction

The QUIET instrument is mounted on rails which attach to the deck axis bearing of the mount. Below this bearing is the elevation axis, which consists of a hinge on one side and a long ball-screw on the other. When the telescope is pointed at zenith, the deck is parallel

to the ground, but as the elevation is decreased, it tips further and further. While the deck is quite rigid, the large weight of the QUIET telescope, cryostat, and electronics cause some flexure.

The effect of the flexure is that the actual elevation, el' , of the boresight pointing vector is somewhat lower than the value reported by the encoder. The magnitude of the effect is given by the parameter k_f :

$$el' = el - k_f \cos el \quad (7.3)$$

The flexure correction goes to zero at $el = 90^\circ$, where the deck is not tipped, and increases for lower elevations.

7.1.5 Azimuth Tilt Correction

The azimuth axis of the mount is not perfectly aligned with zenith. This means that rotating the telescope in azimuth will also lead to small changes in the actual elevation for the boresight pointing. The tilt is parametrized by θ , the magnitude of the tilt, and Ω , the direction of the tilt expressed in azimuth (i.e. degrees measured east from north). Given values of θ and Ω , we can correct \vec{p} and \vec{s} by applying the rotation matrix

$$\mathbf{R}_{az} = \begin{bmatrix} \cos^2 \Omega \cos \theta + \sin^2 \Omega & \cos \Omega \sin \Omega (1 - \cos \theta) & \cos \Omega \sin \theta \\ \cos \Omega \sin \Omega (1 - \cos \theta) & \sin^2 \Omega \cos \theta + \cos^2 \Omega & -\sin \Omega \sin \theta \\ -\cos \Omega \sin \theta & \sin \Omega \sin \theta & \cos \theta \end{bmatrix} \quad (7.4)$$

7.1.6 Elevation Axis Non-perpendicularity

The elevation axis should be oriented parallel to the horizon and perpendicular to the azimuth direction of the telescope. If the elevation axis does not lie in the horizon plane, then the true azimuth of the telescope will change as it slews in elevation. The elevation axis non-perpendicularity is parametrized by θ_e , which is the angle between the actual and ideal

elevation axes. Given a value of θ_e , we can correct \vec{p} and \vec{s} by applying the rotation matrix

$$\mathbf{R}_{el} = \begin{bmatrix} \sin^2 az \cos \theta_e + \cos^2 az & \cos az \sin az (\cos \theta_e - 1) & -\sin az \sin \theta_e \\ \cos az \sin az (\cos \theta_e - 1) & \cos^2 az \cos \theta_e + \sin^2 az & -\cos az \sin \theta_e \\ \sin az \sin \theta_e & \cos az \sin \theta_e & \cos \theta_e \end{bmatrix} \quad (7.5)$$

7.1.7 Array Collimation

After applying the above corrections, one obtains the boresight pointing vector, \vec{p}_b , which is the vector that is invariant under deck rotation. The central module in the array, RQ-09, is not perfectly aligned with the deck axis, so, as the deck rotates, its pointing vector inscribes a small circle on the sky.

Module RQ-09 pointing is parametrized by θ_c and ψ_c , the magnitude and orientation of the collimation offset. The orientation angle, ψ_c , is defined clockwise from \vec{s}_b . The pointing and orientation vectors for RQ-09 can be calculated from \vec{p}_b and \vec{s}_b by applying the following transformations

1. Rotate \vec{s}_b by angle ψ_c about \vec{p}_b to obtain \vec{s}' .
2. Rotate \vec{p}_b and \vec{s}' by angle θ_c about $\vec{p}_b \times \vec{s}'$ to obtain \vec{p}_{09} and \vec{s}'' .
3. Rotate \vec{s}'' by angle $-\psi_c$ about \vec{p}_{09} to obtain \vec{s}_{09} .

An equivalent description of this transformation is to use \vec{p}_b , \vec{s}_b , and \vec{r}_b as an orthonormal basis. Then \vec{p}_{09} and \vec{s}_{09} can be calculated directly as

$$\vec{p}_{09} = \vec{p}_b \cos \theta_c + \vec{s}_b \cos \psi_c \sin \theta_c + \vec{r}_b \sin \psi_c \sin \theta_c \quad (7.6)$$

$$\begin{aligned} \vec{s}_{09} = & \vec{s}_b \left(\cos^2 \psi_c \cos \theta_c + \sin^2 \psi_c \right) - \vec{p}_b \cos \psi_c \sin \theta_c \\ & - \vec{r}_b \cos \psi_c \sin \psi_c (1 - \cos \theta_c) \end{aligned} \quad (7.7)$$

7.1.8 Coordinate Transformations

All of the preceding corrections are defined in the horizon coordinate system, which is fixed to the earth with the \hat{x} axis pointing north and the \hat{z} axis pointing toward zenith. For determining pointing and orientation on the sky, these vectors are transformed to equatorial (aligned with the earth's equatorial plane) or galactic (aligned with the plane of the Milky Way) coordinates.

The transformation from horizon to equatorial coordinates is given by a rotation matrix.

$$\mathbf{R}_{hor \rightarrow equ} = \begin{bmatrix} -\cos lst \sin lat & \sin lst & \cos lst \cos lat \\ -\sin lst \sin lat & -\cos lst & \sin lst \cos lat \\ \cos lat & 0 & \sin lat \end{bmatrix} \quad (7.8)$$

This rotation involves the latitude of the telescope, lat , which is $-23^\circ 01' 41.6''$ and the local apparent sidereal time, lst .

$$lst = lon + \left(\frac{360^\circ}{86400 sec} \right) \left(UT1 + 24110.54841 + 8640184.812866 T_u + 0.093104 T_u^2 - 6.2 \times 10^{-6} T_u^3 \right) \quad (7.9)$$

For QUIET, telescope longitude, lon , is $-67^\circ 45' 42.0''$. Time is specified by $UT1$, which is the universal time in units of seconds since midnight, and T_u , which is the time in centuries since January 1, 2000. The higher order terms in equation (7.9) correct for the precession of the earth's orbit.

Once it has been converted to equatorial coordinates, the coordinates of a pointing vector

or parallactic angle of an orientation vector can be calculated as

$$R.A. = \arctan \left[\frac{\vec{p} \cdot \hat{y}}{\vec{p} \cdot \hat{x}} \right] \quad (7.10)$$

$$Dec = \arcsin [\vec{p} \cdot \hat{z}] \quad (7.11)$$

$$Par = \arctan \left[\frac{\vec{s} \cdot (\hat{z} \times \vec{p})}{\vec{s} \cdot (\vec{p} \times (\hat{z} \times \vec{p}))} \right] \quad (7.12)$$

The rotation from equatorial to galactic coordinates is effectively constant for the QUIET observing season. The transformation matrix, calculated for the J2000 epoch, is

$$\mathbf{R}_{equ \rightarrow gal} = \begin{bmatrix} -0.05487 & -0.87344 & -0.48383 \\ 0.49411 & -0.44483 & 0.74698 \\ -0.86767 & -0.19808 & 0.45599 \end{bmatrix} \quad (7.13)$$

For pointing or orientation vectors in galactic coordinates, calculation of latitude, longitude, or parallactic angle is the same as shown in equations (7.10) through (7.12).

7.1.9 Array Pointing Model

After the corrected position and orientation of module RQ-09 have been calculated and converted to the desired coordinate system, the pointing for the rest of the modules in the array is obtained by applying additional collimation offsets. Mathematically, this process is the same as the one described in equations (7.6) and (7.7) except with \vec{p}_{09} , \vec{s}_{09} , and \vec{r}_{09} replacing \vec{p}_b , \vec{s}_b , and \vec{r}_b . The offset parameters, ψ_i and θ_i , for each module are listed in Table 7.1. Figure 7.1 shows the relative pointing of the 19 modules in the array.

Table 7.1. Collimation offset parameters for each module, relative to RQ-09

Module	ψ	θ	Module	ψ	θ
RQ-00	-179.664	3.4658	RQ-01	150.052	2.9968
RQ-02	119.739	3.4653	RQ-03	-149.602	3.0162
RQ-04	-179.773	1.7413	RQ-05	119.821	1.7406
RQ-06	89.679	3.0159	RQ-07	-119.628	3.4914
RQ-08	-119.860	1.7491	RQ-09	0.000	0.0000
RQ-10	59.901	1.7490	RQ-11	59.644	3.4889
RQ-12	-89.798	3.0284	RQ-13	-60.000	1.7419
RQ-14	0.163	1.7472	RQ-15	24.843	3.0352
RQ-16	-59.802	3.5019			
RQ-17	-29.926	3.0356	RQ-18	-0.060	3.5089

7.1.10 Pointing Model Parameters

The mount model parameters are derived from moon observations throughout the season (Thompson, 2009). Table 7.2 lists the parameters, which have a stepwise dependence on time. The initial motivation for allowing the parameters to vary is that the deck encoder correction is known to change (see Figure 6.10). However, the variation was expanded to all mount model parameters, as it was found to improve the fit. The dates chosen as steps are selected based on the availability of observations used to determine pointing, rather than any events which might be expected to cause pointing changes. The exception is January 29, 2009, which is the date when the deck encoder was fixed. Figure 7.2 shows the magnitude of the pointing correction, compared to a model where only the collimation offset is applied.

Table 7.2. Mount model parameters

start date	end date	Δdk	k_f^a	Ω	θ	θ_e	ψ_c	θ_c
—	2008-10-31	0.00000	0.07129	−10.8853	0.01265	−0.01457	−28.4665	0.28462
2008-11-01	2008-11-30	1.29000	0.07131	−85.6323	0.02876	−0.01749	−23.0846	0.28449
2008-12-01	2009-01-19	−0.59000	0.07131	−85.6323	0.02876	−0.01749	−23.0846	0.28449
2009-01-20	2009-01-28	−0.59000	0.07926	−114.1854	0.02117	−0.03630	−20.2166	0.31003
2009-01-29	2009-04-29	0.66000	0.07926	−114.1854	0.02117	−0.03630	−20.2166	0.31003
2009-04-30	—	0.66000	0.06500	−50.0699	0.00801	0.00823	−20.3033	0.29974

^aParameter k_f has units of degrees / $\cos el$. All other parameters have units of degrees.




Figure 7.1: Array pointing relative to module RQ-09. The x and y axes are both in units of degrees. For $dk = 0$, the x axis corresponds to the direction of increasing azimuth and the y axis corresponds to the direction of increasing elevation.

7.1.11 Evaluation of Pointing Error

Figure 7.3 shows the residual pointing offsets after applying the mount model for the moon observations listed in Table A.1¹. The absolute pointing offset (solid lines) for each observation is the mean difference between the true moon position and the centroid. This quantity is sensitive to pointing errors that affect the entire array, such as deck flexure, tilt of the azimuth or elevation axes, and the RQ-09 collimation offset. The relative pointing offset (dashed lines) is the *rms* scatter from module to module in the centroid locations. This is

1. The model was derived from these and other moon observations, but used a totally independent analysis.




Figure 7.2: Magnitude of the mount model correction, in arc-minutes, compared to applying a collimation offset only. The solid black lines show the paths of QUIET CMB patches. The outermost ring is the path of Patch 7b, the middle ring is the path of Patch 2a or 4a, and the inner ring is the path of Patch 6a.

particularly sensitive to errors in the deck encoder value.

Simulations have been used to evaluate systematic error using a significantly more pessimistic estimate for the size of the pointing offsets (Tajima, 2010b). The conclusion of that study was that uncertainty in the pointing model has a noticeable effect for the EE spectrum only at small angular scales ($\ell > 300$) and the effect on the BB spectrum is negligible.

7.2 Gain Model

The purpose of the gain model is to convert the average and demodulated signals for each detector diode from voltage to temperature units. The primary gain calibration comes from the small sky dips (scans in elevation) performed at the start of each CES. These provide a frequent monitor, but only for the response to unpolarized signals. Observations of the polarized source Tau A are used to find the scaling from unpolarized to polarized responsivity.




Figure 7.3: Residual pointing errors, in arc-minutes or as a fraction of the beam size, for the 18 moon observations listed in Table A.1. The solid line indicates the absolute pointing offset, averaged over all modules in the array. The dashed line indicates the *rms* of the relative pointing offset between modules in the array.

7.2.1 Sky Dips

Before every constant elevation scan, a sky dip calibration is made, consisting of three slow scans in elevation with amplitude of four degrees. The atmospheric depth changes with elevation, which causes the input load temperature to vary as $(\sin el)^{-1}$. The elevation is calculated for each module according to the pointing model and the averaged data samples, V_i , for each detector diode are fit to the following model (Tajima, 2010a)

$$V_i = G_{av} \left(\frac{T_z}{\sin el_i} + \kappa \right) + \gamma t_i \quad (7.14)$$

where fit parameter G_{av} is the average gain, in units of mV/K , κ is an offset that combines the amplifier noise power as well as the 3K contribution from the CMB, and γ is the slope for a linear drift in time (averaged data suffers from large $1/f$ noise, frequently leading to offset drifts during the sky dips).

The sky temperature at zenith is modeled as a linear function of the precipital water vapor (PWV)

$$T_z = T_z^0 (1 + \alpha \text{PWV}) \quad (7.15)$$

Values for T_z^0 ($8.5K$) and α ($3.8\%/mm$) are obtained from the data and found to be consistent with atmospheric models.

An additional linear correction is applied based on the enclosure temperature, using a thermometer mounted on one of the LNA bias boards. While the sky temperature coefficient, α , is the same for all modules and diodes, the gain temperature coefficients are measured separately for each channel. The values are typically 1% / degree, but can be as high as 4% / degree for certain detector diodes (module RQ-04 has the largest temperature coefficients).

The result of these corrections is a set of measured gain values over time that are all translated to the same levels of precipital water vapor and enclosure temperature. To obtain the gain model coefficient for a given CES and diode, the median is selected from all sky dips with ± 12 hours, subject to some quality criteria. While the precipital water vapor dependence affects only the temperature of the calibration source (atmosphere), not the actual detector gain, the enclosure temperature effect does represent a true gain change. Accordingly, the median gain value is adjusted from the fiducial enclosure temperature to the measured value for the CES in question.

7.2.2 *Tau A Observations*

To convert the sky dip model for G_{av} to a model for responsivity to polarized signals, it is necessary to find a calibration source with known polarized flux². QUIET uses observations of the supernova remnant Tau A, also known as the Crab Nebula, as an absolute polarization

2. The difference between total power and polarized responsivity is covered in §3.2. In summary, the polarized gain, G_{de} , is less than or equal to the total power gain, G_{av} . The inequality is saturated for the case where the two legs of the module have equal gain and phase.

calibrator. One of the brightest polarized sources in the microwave sky, Tau A attracts interest from many CMB experiments (Aumont et al., 2009).

The polarized flux of Tau A is obtained from measurements by WMAP (Weiland et al., 2010). The reported value is 22.12 ± 0.60 Jy, at a frequency of 40.64 GHz. To calculate the signal observed by QUIET, it is necessary to scale the flux to the QUIET Q-band observing frequency of 43 GHz, using a spectral index of -0.3476 . Next, the observed polarization temperature is calculated using the measured beam solid angle. This yields a value of 5.06 mK for the Tau A polarized signal as observed by QUIET.

At a declination of $+22^\circ$, Tau A is only viewable for a short period of time (~ 1.5 hours) each day. The times when it is above the telescope lower elevation limit correspond to transit, so the usual azimuth scans do not work well. Instead, raster scans are used, consisting of back and forth slews in azimuth between elevation steps. To maximize the signal to noise, the size of the raster scan is kept small, limiting the observation to one module at a time. The raster scan is repeated at several deck angles, so that the measured amplitudes can be fit by the function

$$A_i = (5.06 \text{ mK}) G_{de,i} \cos [2(\psi_i - \psi_{\text{taua}})] \quad (7.16)$$

where ψ_i is the parallactic angle of the sensitivity axis for detector diode i and ψ_{taua} is the polarization angle of Tau A, -87.3 ± 0.2 degrees in galactic coordinates (nearly parallel to the galactic plane). With several observations at a range of deck angles, equation (7.16) can be used to solve for both the absolute polarized responsivity and the detector angle for all diodes in the target module.

Roughly half of the Tau A measurements, which were performed almost daily, targeted module RQ-09. As a result, the ratio of the polarized to total power responsivity for this module is measured very well. The other modules in the array, with only one to three Tau A observations each, are consistent with module RQ-09, to the limit of the measurement uncertainty. To avoid introducing extra noise to the gain model, the well measured polarized

to total power gain ratio from RQ-09 is used for all modules, and potential module-to-module fluctuations in this value are treated as a source of systematic error. An overall error in the absolute calibration, such as uncertainty in the polarized flux of Tau A, directly scales the CMB power spectrum result like the square of the calibration error. But relative gain errors from diode to diode have a smaller effect, which is estimated from simulations to be much less than the statistical error of the CMB measurement (Tajima, 2010b).

7.3 Noise Model

Noise in the demodulated data must be modeled in order to generate realistic simulations for the Monte Carlo based analysis (see §9.3). The noise is measured for each detector diode in every CES by fitting the power spectrum to a three parameter model. Additional parameters encode the noise correlations between diodes in a module.

7.3.1 Noise Power Estimation

The function used to model the diode noise power spectra is

$$\bar{P}(\nu) = A^2 \left[1 + \left(\frac{f_{\text{knee}}}{\nu} \right)^\alpha \right] \quad (7.17)$$

The parameter A represents the white noise level, in units of $\mu K/\sqrt{s}$. Parameter f_{knee} is the knee frequency, where the power spectrum reaches two times the white noise level. The slope of the $1/f$ part of the spectrum is given by the parameter α .

Before calculating the spectrum for a diode, the 100 Hz samples are corrected for Type-B glitching and differenced to produce a 50 Hz double-demodulated stream. The full CES timestream is apodized by subtracting a mean and multiplying the first and last 5% by an envelope that goes smoothly to zero at the ends as $\sin^2 x$. The gain model is applied to convert the data from voltage units to μK .

The power spectrum is then calculated and binned as follows:

- First 20 Fourier modes are left unbinned.
- Modes 20–1000 are averaged into bins with increasing width proportional to mode number.
- Higher modes are average into bins with increasing width proportional to (mode number)^{1/2}.

For a typical CES with duration of 4000 seconds and scan period of 14 seconds, bins near the scan frequency are ~ 7 mHz wide and bins near the 25 Hz Nyquist frequency have widths approaching 250 mHz.

The best fit model is found by minimizing the log-likelihood function

$$\log \mathcal{L} = \sum_k \left[N_k \left(\log \bar{P}_k + \frac{P_k}{\bar{P}_k} \right) \right] \quad (7.18)$$

where P_k and \bar{P}_k are the power spectrum and model, respectively, and N_k is the number of modes per frequency bin. The summation excludes frequency bins within 20% of the scan frequency, which often contain excess power due to ground pickup or atmospheric structure (the actual CMB signal is so much smaller than the noise for a single CES and diode that it can be safely ignored). For a CES with 14 second scan period, this masked region would extend from 57 to 86 mHz. Frequency bins above 4.7 Hz are also ignored, due to excess high frequency noise.

A couple of tricks are used to handle cases where the $1/f$ part of the spectrum is not well measured (CES with short duration and/or very low knee frequencies). Before fitting, the white noise level is estimated by averaging all power spectrum bins with frequencies between 2 and 4.6 Hz. The significance of the $1/f$ noise in the power spectrum is estimated from the excess power at frequencies up to 80% of the scan frequency. If that significance is less than 2σ , then no fit is performed. The white noise level for that CES and diode is set to the

value found by averaging and the knee frequency is set to zero. This occurs for about 1.4% of data used in the CMB analysis. If the significance is between 2 and 3σ , then the noise model fit is performed, but with the slope parameter α fixed to 2. This situation occurs in 5.7% of the data.

Figure 7.4 shows the binned noise power spectrum and best fit model for module RQ-09, diode $Q1$, from CES 128.0. The grayed regions are not used to determine the noise model parameters. This particular CES has no significant structure near the scan frequency (0.1 Hz), but excess high frequency noise is visible near 6 Hz.




Figure 7.4: Noise power spectrum for RQ-09 diode $Q1$ from CES 128.0. The noise model spectrum is shown in red. The gray frequency regions, near the 0.1 Hz scan frequency and above 4.6 Hz, are not used for fitting the noise model.

7.3.2 Noise Correlation Estimation

The correlations between detector diodes in a module are another important component of the noise model. Different correlation coefficients are used for the two components of noise in the model: white noise and $1/f$ noise.

White noise correlations are due to the effect described in §3.4, which results in correlation coefficients ≥ 0 for like diodes (i.e. $Q1$ and $Q2$ or $U1$ and $U2$) and coefficients $\simeq 0.5$ for unlike diodes. Given the Fourier transforms, $D_i(\nu)$ and $D_j(\nu)$, of demodulated data for diodes i and j , their correlation coefficient is calculated as

$$C_{ij}(\nu) = \Re \left(\frac{D_i^*(\nu) D_j(\nu)}{|D_i| |D_j|} \right) \quad (7.19)$$

The white noise correlation coefficient is estimated by averaging $\arccos(C_{ij})$ (which has a uniform distribution for the case of zero correlation) over the frequency range of 2 to 4.6 Hz and then taking the cosine of the result.

The $1/f$ noise in demodulated data can be caused by leakage of atmospheric noise or by amplifier gain fluctuations multiplying a polarized offset. For either case, the noise is highly correlated between detector diodes in a module. It is difficult to measure the low frequency correlation coefficients from the data, so they are just assumed to be ± 1 , with sign given by the relative sign of the $\mathbf{I} \rightarrow \mathbf{Q}$ leakage between the two diodes.

7.4 Beam Measurement

Measurement of the detector beam profiles is important because they determine the transfer of CMB fluctuation power on the sky to observed power, as a function of angular scale.

Deep maps of the detector response to a point source are constructed by stacking several Jupiter observations (for the TT modules, RQ-17 and RQ-18) or many Tau A observations (for polarization module RQ-09). The beam ellipticity is small, and fairly unimportant

because the QUIET scan strategy results in observations of each pixel with many detector orientations, so the two dimensional beam maps can be averaged down to radial profiles. These profiles are fit using Gauss-Hermite functions³ (Monsalve, 2010). The even terms up to 18th order are used for the fit.

After determining the Gauss-Hermite function that describes the beam, the window function is calculated by the Legendre transform. The beam window functions are used in the analysis to correct the measured angular power spectrum for the effects of beam smoothing. For module RQ-09, the window function is measured to better than 5% over the full range of angular scales that QUIET is sensitive to. For other polarization modules, which do not have enough Tau A observations to build a deep map for this purpose, the consistency of the RQ-09, RQ-17, and RQ-18 beams is used to estimate the error introduced by assuming that all beam profiles are identical.

7.5 Calibration from Moon Observations

The moon is a useful calibration source because of its extreme brightness. While other astronomical sources require long observations and careful map making, QUIET detects the moon's polarized emission with signal-to-noise greater than 1000 in one second. However, extracting calibration parameters from moon observations is made difficult because of the complicated nature of the source. The moon has an angular diameter on the sky of about half a degree, which is slightly larger than the beam size. This means that internal details like the temperature profile across the moon's face show up in the data and cause confusion in the analysis.

The primary purpose of moon observations for QUIET is to measure the detector angle for each diode. However, other calibration information that can be extracted includes relative

3. The Gauss-Hermite functions are a polynomial expansion used to describe departures from a pure Gaussian.

gains (both polarized and total power gains) between various modules and instrumental $\mathbf{I} \rightarrow \mathbf{Q}$ leakage. A physical model of the moon’s microwave emission is derived in §7.5.2. The set of Q-band moon observations is described in §7.5.3. The methods used for the data analysis are found in §7.5.4, and some discussion of how the shortcomings in the model are understood is included in §7.5.5. Section 7.5.6 contains a overview of the moon calibration results, which are tabulated in full detail in Appendix A.

7.5.1 *Why Emission Instead of Reflection?*

At optical frequencies, the dominant signal from the moon is reflected sunlight. However, in the microwave frequency range, the moon’s thermal emission is stronger than reflected power from the sun, even though the sun is much hotter.

We will treat the sun and the moon as blackbodies with temperatures of 10,000K and 150–300K, respectively. For either of these temperatures, observations at 40 or 90 GHz are firmly in the Rayleigh-Jeans limit of the blackbody spectrum, so the power emitted over a limited band varies only linearly with temperature (as opposed to the power integrated over the entire blackbody spectrum, which is proportional to T^4).

If the temperature at the surface of the sun is 10,000K, then the power in a microwave band from the sun’s radiation at the moon’s surface is given by equation (7.20).

$$P_{\text{incident}} \propto T_{\text{sun}} \left(\frac{R_{\text{sun}}}{R_{\text{orbit}}} \right)^2 \quad (7.20)$$

The ratio of the radius of the sun to the radius of the moon’s orbit around the sun is about 0.0046, so the Rayleigh-Jeans temperature of the sun’s emitted microwaves is only 200mK or so when incident on the moon. The sun’s total radiation, integrated over all frequencies, does contain enough power to heat the moon up to temperatures of 300K, but most of that power is at much higher frequencies that are far outside of QUIET’s observing

band.

7.5.2 *Physical Model of Polarized Emission from the Moon*

The simple physical model used for the moon is a dielectric sphere at a uniform temperature. This model overlooks many of the more complicated features of the source, particularly temperature variations across the moon's surface. These details are treated as sources of systematic error for the extracted calibration, as discussed in §7.5.4.

We start by considering a plane wave traveling through a dielectric medium with index of refraction n_i , and incident on the boundary to a region with index of refraction n_t . For thermal radiation from the moon being emitted into space, we have $n_i \sim 1.8$ and $n_t = 1$. In addition to the incident wave, there are also transmitted and reflected waves with the same frequency. Matching phase factors along the boundary yields Snell's Law (7.21), specifying the angle of the transmitted wave.

$$n_i \sin \theta_i = n_t \sin \theta_t \tag{7.21}$$

The full boundary value problem can be solved separately for the waves polarized perpendicular to the plane of incidence and for waves polarized in the plane, yielding the Fresnel equations for the transmitted wave. For the details of this derivation, see Jackson (1998), Section 7.3. The results from Jackson are given in terms of θ_i , the angle of incidence; here, Snell's Law is used to convert them into functions of θ_t , the angle between the surface normal

and the observer's line of sight. In practice, n_t , μ_i , and μ_t are all set equal to 1.

$$\left\{ \frac{E_t}{E_i} \right\}_{\perp} = \frac{2\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t}}{\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t} + (\mu_i/\mu_t)n_t \cos \theta_t} \quad (7.22)$$

$$\left\{ \frac{E_t}{E_i} \right\}_{\parallel} = \frac{2n_i\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t}}{(\mu_i/\mu_t)n_t\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t} + n_i^2 \cos \theta_t} \quad (7.23)$$

Equations (7.22) and (7.23) give the transmission amplitudes for waves as a function of angle, but refraction at the interface means that the amount of incident power is not equal for all values of θ_t . For a point just below the surface of the medium, the energy flux in the plane of incidence is uniform for all directions that point down into the medium, $-\pi/2 < \theta_i < \pi/2$. After passing through the interface, the flux lines are not uniformly distributed in θ_t . Their angular distribution can be calculated by differentiating both sides of Snell's Law:

$$n_i \sin \theta_i = n_t \sin \theta_t \quad (7.24)$$

$$n_i \cos \theta_i d\theta_i = n_t \cos \theta_t d\theta_t \quad (7.25)$$

$$d\theta_i = d\theta_t \frac{n_t \cos \theta_t}{n_i \cos \theta_i} \quad (7.26)$$

$$= d\theta_t \frac{n_t \cos \theta_t}{\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t}} \quad (7.27)$$

The factor in equation (7.27) can also be thought of as the ratio of an area element on the surface as seen from angles θ_i and θ_t .

Combining equations (7.22), (7.23), and (7.27), we can write expressions for the observed

temperature and polarization of the dielectric surface as a function of n_i , n_t , θ_t , and T :

$$\mathbf{I} = \frac{T}{2} \left(\left\{ \frac{E_t}{E_i} \right\}_{\parallel}^2 + \left\{ \frac{E_t}{E_i} \right\}_{\perp}^2 \right) \frac{n_t \cos \theta_t}{\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t}} \quad (7.28)$$

$$\mathbf{P} = \frac{T}{2} \left(\left\{ \frac{E_t}{E_i} \right\}_{\parallel}^2 - \left\{ \frac{E_t}{E_i} \right\}_{\perp}^2 \right) \frac{n_t \cos \theta_t}{\sqrt{n_i^2 - n_t^2 \sin^2 \theta_t}} \quad (7.29)$$

For the case where $n_i > 1$ and $n_t = 1$, inspection of the Fresnel equations reveals that the transmission amplitude is larger for the waves polarized parallel to the plane of incidence than for waves polarized perpendicular to the plane of incidence, so \mathbf{P} , as defined above, is positive.

Modeling the moon as a dielectric sphere, seen as a disk in projection, the plane of incidence and θ_t depend on position on the disk. If we define a system of polar coordinates, (r, ϕ) , centered at the moon, then θ_t is given by

$$\theta_t = \arcsin \left(\frac{r}{R_{\text{moon}}} \right) \quad (7.30)$$

where $R_{\text{moon}} \simeq 0.25^\circ$. The plane of incidence forms a radial line in the (r, ϕ) plane, so the two relevant polarization directions, E_{\parallel} and E_{\perp} , correspond to polarization aligned radially and tangentially, respectively. Since the parallel component of the emission is larger, light emitted from the moon is polarized along the radial direction. Visualized as \mathbf{Q} and \mathbf{U} , the radial pattern becomes a pair of quadrupoles, oriented 45° apart.

Figure 7.5 shows the total power and polarized intensity emitted by the moon, as measured by a hypothetical experiment with perfect angular resolution. The total power signal is fairly flat across the moon's disk and is inversely related to n_i . The polarization signal is strongly peaked near the edge of the moon's disk and increases with n_i .

To account for QUIET's large beam, a two-dimensional map of the moon is smoothed using a Gaussian kernel. Figure 7.6 shows simulated I , Q , and U signals for a Q-band

detector, with beam full-width at half maximum (FWHM) of 0.471° .




Figure 7.5: Emitted total power (left panel) and polarization (right panel) intensity, in units of the moon’s physical temperature T_{moon} , as a function of radial distance r , in units of the moon radius. Values of n_i are 1.2, 1.4, 1.6, 1.8, and 2.0, shown in black, dark blue, light blue, green, and red, respectively.

7.5.3 Moon Calibration Observations

During the observing season, there were a total of 91 runs targeting the moon, adding up to 118 distinct scans (sometimes multiple moon scans would be lumped into one run when logged by the observer).

Two scan patterns were used for moon observations. The drift scan is very similar to the standard QUIET CMB scan. The telescope scans in azimuth at fixed elevation while the moon moves through the scan region by sidereal motion. The drift scan works well when the moon is rising or setting, and its motion is mostly perpendicular to the scan direction. However, when the moon is transiting, its elevation changes very little over the course of the scan and it would be observed by only one row of the array. For these cases, the raster scan, described in §7.2.2, is used. For scans over the full array, which require about 45 minutes, the




Figure 7.6: Maps of **I**, **Q**, and **U** for the moon model, smoothed by a Gaussian beam with FWHM of 0.471 degrees. The moon is modeled with a uniform physical temperature of 200K and index of refraction of 1.8. The vertical scales for temperature and polarization are in Kelvin; the temperature map peaks at $\sim 92K$ while the polarization maps peak at $\sim 0.84K$. Each map is 1.5 degrees on a side.

drift scan produces better results. The disadvantage of the raster scan is that the elevation steps are quite slow and result in different atmospheric signals at each step.

Of the 118 moon scans, 56 are drift scans and 62 are raster scans. Moon scans from October 2008 were taken before the double demodulation was implemented for data taking. These early scans account for 18 of the 56 drift scans and none of the raster scans.

For the moon calibration analysis, only a subset of the scans are used. The scans selected are drift scans where all, or nearly all, of the detectors in the array observed the moon's disk in its entirety. Details of these scans can be found in Table A.1.

7.5.4 *Fitting Procedure for Moon Observations*

Observations of the moon are fit to the model described in §7.5.2. Important parameters extracted from the fit are the polarization angle for each detector diode, relative signal amplitudes between channels, and **I** \rightarrow **Q** leakage coefficients. The fit is done in two parts. First the average data is analyzed to determine pointing offsets and average signal amplitudes. Next, the demodulated data is analyzed to find polarization angles, leakage coefficients, and demodulated signal amplitudes.

The first step is to convert the pointing from telescope azimuth, elevation, and deck angle to a coordinate system fixed to the moon. The moon's position in equatorial coordinates is calculated using ephemerides from the Jet Propulsion Laboratory's Solar System Dynamics Group⁴. The `aephem` software library⁵ is used for astronomical reduction of ephemerides and to calculate topocentric corrections to the moon's apparent position. The pointing vectors for the module being analyzed and for the moon position, \vec{p}_i and \vec{p}_m respectively, are calculated in equatorial coordinates. At each sample, a flat, two-dimensional coordinate system is defined, centered on the moon position. The basis vectors, \hat{x} and \hat{y} , of this coordinate system are defined as

$$\hat{x} = \begin{bmatrix} \cos RA_m \sin Dec_m \\ \sin RA_m \sin Dec_m \\ -\cos Dec_m \end{bmatrix} \quad (7.31)$$

$$\hat{y} = \begin{bmatrix} -\sin RA_m \\ \cos RA_m \\ 0 \end{bmatrix} \quad (7.32)$$

and the module pointing in the coordinate system is calculated as

$$\begin{bmatrix} x \\ y \end{bmatrix}_i = \begin{bmatrix} \arcsin(\vec{p}_i \cdot \hat{x}) \\ \arcsin(\vec{p}_i \cdot \hat{y}) \end{bmatrix} \quad (7.33)$$

The convention for \hat{x} and \hat{y} is chosen such that \hat{x} corresponds to the direction of decreasing declination and \hat{y} corresponds to the direction of increasing right ascension.

The fitting procedure is the same for both average and demodulated data, but with different model functions. Working with only one 50 Hz state at a time, data points within

4. <http://ssd.jpl.nasa.gov/>

5. <http://aephem.sourceforge.net>

0.8° of the moon center are selected for fitting. The standard deviation of data points more than 2° from the moon is used as an estimate of the uncertainty on each sample.

For average data, the fit function is given by

$$\mathcal{F}_{av} = A_{av}\mathbf{I}_m + B_{av} \quad (7.34)$$

where A_{av} and B_{av} are fit parameters for the amplitude and offset of the average data and the moon total power signal, \mathbf{I}_m , is calculated based on additional parameters specifying pointing offsets in the x and y directions, beam size, moon radius, and moon index of refraction. The amplitude and index of refraction parameters are strongly degenerate, so index of refraction is held fixed at a value of 1.8 to improve fit convergence. There is a similar degeneracy for the beam size and moon radius parameters, so the beam size is fixed at 0.451° while the moon radius is allowed to vary.

For demodulated data, the fit function is given by

$$\mathcal{F}_{de} = A_{de} \{ \mathbf{Q}_m \cos [2 (\psi + \psi_i)] + \mathbf{U}_m \sin [2 (\psi + \psi_i)] \} + B_{de} + \eta v_{av} \quad (7.35)$$

where A_{de} and B_{de} are amplitude and offset parameters. The polarized moon signal, \mathbf{Q}_m and \mathbf{U}_m , is calculated with parameters for pointing offsets, beam size, moon radius, and moon index of refraction fixed to the values found from the fit to average data for this observation, diode, and 50 Hz state. The detector angle, $\psi + \psi_i$, consists of two terms. The first term, ψ , is the parallactic angle calculated for the orientation vector of this module in equatorial coordinates.

$$\psi = \arctan \left(\frac{\vec{s} \cdot \hat{y}}{\vec{s} \cdot \hat{x}} \right) \quad (7.36)$$

The second term, ψ_i , is a free parameter in the fit, representing the actual orientation of the sensitivity axis for the diode. For the two different 50 Hz states, values of ψ_i differ by 90°.

The final fit parameter, η , represents $\mathbf{I} \rightarrow \mathbf{Q}$ leakage. In the fit function, it multiplies v_{av} , which are the corresponding samples of mean-subtracted average data⁶.

Best-fit parameters are found by minimizing the χ^2 statistic comparing time-ordered data samples calculated from the model to the actual data in the fit region. The MINUIT function minimization library⁷ (MIGRAD minimizer) is used to find the best-fit model. The combination of large signal-to-noise for the measurement and known inaccuracies in the model function means that the minimized χ^2 values are still much larger than the number of degrees of freedom, typically by factors of 10 or more. Likewise, the statistical errors on the parameters calculated in the fitting procedure are not good estimates of the calibration uncertainty.

7.5.5 *Understanding the Shortcomings of the Moon Model*

The assumption of temperature uniformity for the model specified in equation (7.35) leads to systematic bias in the recovered fit parameters. Simulated moon observations with added temperature non-uniformities are used to understand the size and behavior of these errors.

The simulations use the telescope encoder information from an actual moon observation, run 1058.0. The moon signal model includes a dipole temperature variation over the sphere. Detailed models of the moon surface temperature exist, such as Gary et al. (1965) or Linsky (1973), but once they have been smoothed out by the QUIET beam, these more complicated models are indistinguishable from the simple one used here. Detector noise is added, but the typical noise level is much smaller than the signal. The signal plus noise timestreams are then analyzed using the fitting function described above, that assumes a uniform moon temperature.

6. One could use the moon total power model, \mathbf{I}_m , instead of the average data, but this would cause the fit value of η to depend on the value used for the moon index of refraction. Using the average data, while it adds noise into the model, ensures that this convention for the leakage fraction is consistent with measurements from other sources.

7. <http://lcgapp.cern.ch/project/cls/work-packages/mathlibs/minuit/index.html>

Figure 7.7 demonstrates the bias in the recovered detector angles for three simulations ranging from uniform temperature to a dipole variation with 33% amplitude. Each simulation was analyzed for a range of detector diode sensitivity axes, to determine how the temperature non-uniformity couples to the measured polarization quadrupole. The result, seen in the upper panel of Figure 7.7, is a sinusoidal variation in the recovered detector angle that increases in amplitude with the size of the temperature dipole. The period of the variation is 90° , so diode pairs (i.e. $Q1$ and $Q2$ or $U1$ and $U2$) are affected identically.

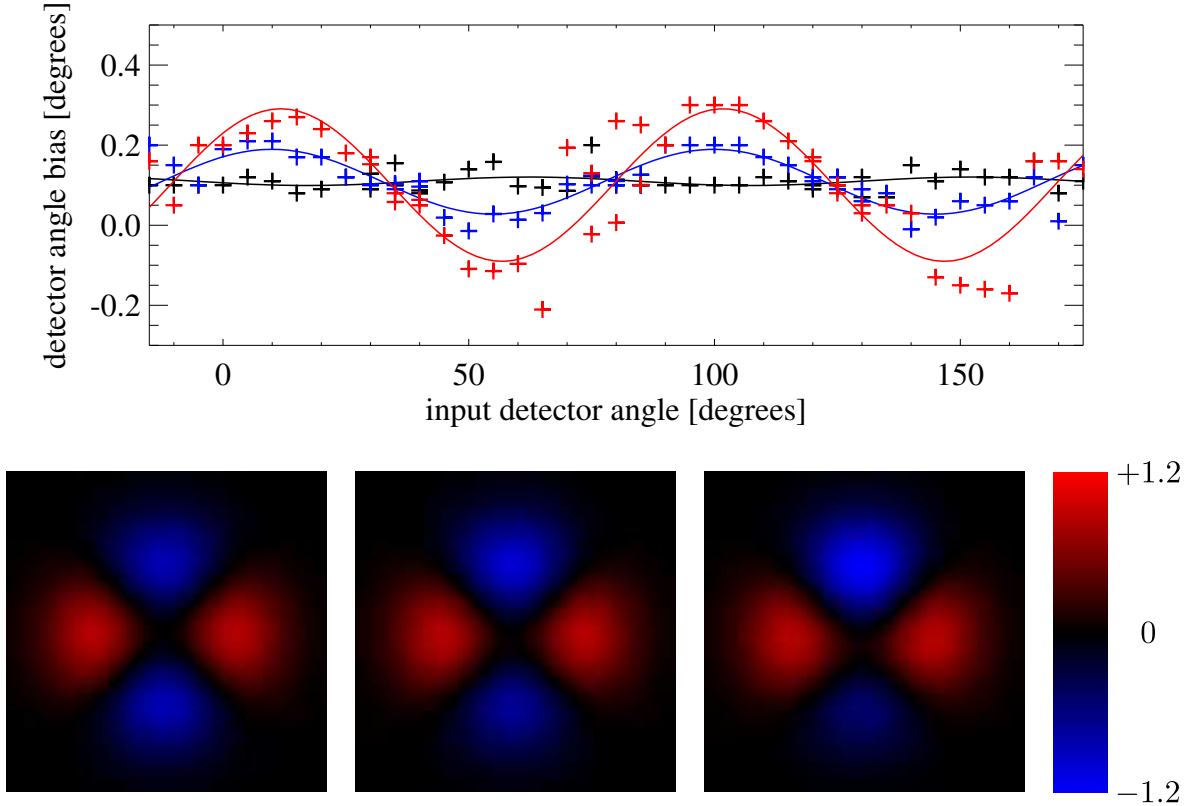


Figure 7.7: The top panel shows bias in detector angles recovered from moon simulations as a function of true detector angle. Black points correspond to the moon simulation shown at bottom left, with no temperature non-uniformity. Blue points correspond to the simulation shown at bottom center, with a temperature dipole with 13% amplitude. Red points correspond to the simulation shown at bottom right, with a temperature dipole with 33% amplitude. The maps show the simulated Q Stokes parameter; detector diodes analyzed in the simulation measure various combinations of Q and U (not shown). Solid lines in the upper panel are fit curves with the three parameter model $y = A + B \sin(4x + C)$.

While this effect introduces the possibility of systematic bias in the moon calibration, it gets averaged out when many observations are combined. The 18 scans used to measure the detector angles include the full range of moon phase, observations while the moon is rising and setting, and four distinct deck angles. This effectively randomizes the relative orientation of the polarization sensitivity axes and the direction of thermal gradients on the moon, so the potential bias turns out to be just a source of noise.

7.5.6 Results from Moon Calibrations

Appendix A contains detector angles, relative total power and polarized gains, and $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients measured from the moon observations. An overview of the results is included here.

Figure 7.8 shows the detector angles for all diodes of modules RQ-00 through RQ-15 (except for broken diode RQ-08 *U2*). As a cross check, the angles are calculated separately for the two 50 Hz states, yielding consistent results⁸. Also shown are detector angles derived from observations of Tau A (Chinone & Tajima, 2010). While the moon and Tau A results are in generally good agreement for *U* diodes, there does seem to be an overall offset in the measurements for *Q* diodes. This suggests that there is some residual bias in the analysis for the case of larger $\mathbf{I} \rightarrow \mathbf{Q}$ leakage.

The measured leakage coefficients for all diodes are shown in Figure 7.9. The characteristic pattern of large leakage in *Q* diodes and smaller leakage in *U* diodes can be clearly seen. Also shown are the leakage coefficients measured from the response of the demodulated signal to sky dips. The sky dip leakage coefficients are uniformly higher than the ones derived from moon observations. The probable source of this discrepancy is the difference between the spectrum of the moon (rising spectrum) and the atmosphere (falling spectrum at 43

8. Really, the sensitivity axis flips by 90° with the 50 Hz state, but this angle has been subtracted out for Figure 7.8.




Figure 7.8: Detector angles derived by averaging 18 moon observations. The black and red points are for independent measurements in the two 50 Hz states. The gray lines indicate the detector angle values measured from observations of Tau A.




Figure 7.9: $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients measured from moon observations. For each module, the four detector diodes are ordered $Q1$, $U1$, $U2$, $Q2$, from left to right. The Q diodes, shown in red, have much higher leakage than the U diodes, shown in blue. The gray line indicates the values of the leakage coefficients as measured from sky dips.

GHz). This difference means that the measurement of leakage coefficients from the moon is weighted towards the high end of the module frequency band while the sky dip result is weighted towards the low end.

Because the moon analysis is performed separately for each 50 Hz state, it is possible to construct the leakage coefficients for both double demodulated data and also for the data obtained by taking the average of the two states. The former are shown in the figure; the latter are the result of phase switch imbalance (see §3.5.1) and are found to be quite small.

Figure 7.10 shows a comparison of the gain model, described in §7.2, and relative polarized gains derived from the amplitude of the moon signal in demodulated data. The temperature of the moon changes dramatically with its phase, so no attempt is made to use it as a source for absolute calibration. Instead, RQ-09 is designated as the reference used to calculate relative gains for all other modules. In this case, the gain ratio between RQ-05 $Q1$ and RQ-09 $Q1$ is shown. For this pair of channels, there is good agreement between the gain




Figure 7.10: The small black points show the ratio of the gain model polarized responsivity for RQ-05 diode $Q1$ to RQ-09 diode $Q1$ for the entire observing season. The red boxes show the same ratio, as calculated from 18 separate moon observations.

model and all moon observations except for a few outliers.

CHAPTER 8

DATA SELECTION

Over nine months of operation, QUIET observed in a variety of conditions. While scans were halted during periods of extremely bad weather, there was still a significant amount of data taken during marginal periods. In addition to selecting for weather quality, it is also important to identify times where one or more modules might be operating abnormally.

This chapter describes the criteria used for data selection and the motivations for the choices. The baseline data set, which excludes only data that are known to be defective, is detailed in §8.1. Cuts are applied based on weather monitoring statistics, described in §8.2, and the results of the noise model fit, discussed in §8.3 and §8.4. Additional selection criteria, which remove small amounts of data, are based on the position of the sun (§8.5), glitches in the timestream (§8.6), or residual Type-B noise (§8.7). The data selection efficiency, defined as the percentage of the baseline data set that is accepted for use in the analysis, is given in §8.8.

The fundamental unit of data is a single detector diode for one full constant elevation scan, referred to as a “CES-diode”. While it is undoubtedly true that some CES have only short periods of bad data and are otherwise usable, trying to account for and keep track of smaller periods of data would require large additional effort for a small increase in the fraction of data retained. Selecting or rejecting individual detector diodes, rather than requiring entire modules to pass cuts, does save significant amounts of useful data.

The data selection is implemented as a MySQL database, which contains several tables of pre-calculated statistics for all CES-diodes. The criteria described below are all combined into a database query that returns the list of CES-diodes to use. Details of the database tables and the selection query are found in Appendix B.

Table 8.1. Total number and duration for all CES by patch

Patch	# of CES	Total duration [h:m:s]
Patch 2a	823	947:20:28
Patch 4a	677	739:58:16
Patch 6a	676	875:12:34
Patch 7b	254	234:30:29
All	2430	2797:01:47

8.1 Baseline Data Set

Data selection begins with a master list of CES for each of the four CMB patches. Table 8.1 shows the total number and duration of CES by patch. The average duration for a Patch 6a CES is substantially longer than for the others, because it is at a lower declination than the other patches, so it moves more slowly through the sky.

Before proceeding with data selection based on weather or detector performance criteria, some data are removed from consideration due to known problems. The problems fall into three categories: broken detector channels, bad or abnormal periods of observation, and very short CES. The data set obtained after cutting these data is referred to as the “baseline data set”. Table 8.2 lists the CES and diodes that are not included in the baseline data set.

One module, RQ-16, and two diodes from other modules, RQ-04 *Q1* and RQ-08 *U2*, are cut entirely from the data set. The malfunctions of these channels are described in §6.1. This cut removes 6 out of 68 total detector diodes from polarimeter modules, or 9% of all channels.

CMB scans were often stopped early for maintenance, approaching weather, or other issues. Estimation of $1/f$ noise, described in §7.3, is less reliable for short stretches of data. Also, while there are a fairly large number of short CES, they don’t collectively add up to a very large fraction of the data. Rather than expending a lot of extra effort for very

little added sensitivity, all CES with duration less than 1000 seconds are dropped from the baseline data set.

Several CES are cut from the baseline data set because of known problems or configuration changes to the QUIET instrument. Many of these events are listed in §5.4.

8.2 Weather Cut

When the weather at the site worsens, QUIET data is affected in two ways. First, the sky temperature increases, due to higher atmospheric opacity. Second, atmospheric fluctuations worsen, as cloud-like structures begin to form. Atmospheric emission is unpolarized, so the first effect only contributes additional white noise to the polarization data. The second effect is more serious, introducing unstable scan synchronous signals with a $1/f$ type power spectrum. Despite being attenuated by a factor of the instrumental $\mathbf{I} \rightarrow \mathbf{Q}$ leakage, the bad weather signal can strongly corrupt the polarization maps.

The module averaged data, though too noisy to be used for CMB temperature measurements, serves as an excellent weather monitor. Selecting based on the averaged, rather than demodulated, data avoids introducing bias in the polarization maps, which would happen if data were selected based on fluctuation power in the polarized timestreams.

Two statistics, calculated using the averaged data from many modules in the array, quantify the atmospheric noise during each CES. First, the gain model is applied and the averaged data is downsampled to 1 Hz for all channels. No distinction is made between the two 50 Hz states used for double demodulation (in averaged data, the 50 Hz signal is small compared to the overall magnitude of the total power signal). Next, the *rms* of the 1 Hz samples is calculated for every ten second period. Ten seconds matches the timescale of the azimuth scan, so these *rms* values are very sensitive to structure on the sky that is picked up by scanning. The *rms* values are summed in quadrature across all detector diodes on ADC

Table 8.2. Data excluded from the baseline data set

Excluded Data	Reason	Patch	Cut Size
RQ-04, diode $Q1$	Excess high frequency noise	all	1 diode
RQ-08, diode $U2$	Broken diode (short)	all	1 diode
RQ-16	Unreliable gate bias connection	all	4 diodes
CES < 1000 seconds	Short duration CES	all	98 CES
Runs 329–341	Alternate preamp configuration ^a	all	32 CES
Run 398	Groundscreen door came open during observation	4a	2 CES
CES 529.1	Receiver was off	4a	1 CES
Runs 563–593	Alternate AIB cables ^b	all	50 CES
CES 654.0	Mount stalls	4a	1 CES
CES 759.0	Dome maintenance during observation	6a	1 CES
Run 937	Mount stalls	6a	3 CES
CES 951.0	Mount stalls	4a	1 CES
Runs 953–955	Mount stalls	2a	7 CES
CES 980.0	Receiver cryogenic temperature was still stabilizing	7b	1 CES
CES 1413.0	Mount stalls	6a	1 CES
CES 632.4–Run 635	Blown fuse on ADC 2	2a, 6a, 7b	11 CES
RQ-07, 08, 09, 12, 13, 16 only			

^aPreamp boards were swapped, as part of Type-B debugging effort.

^bSpare AIB cables were installed for these runs, as part of Type-B debugging effort.

1 (modules RQ-00 through RQ-06)¹ The following diodes with significant Type-A glitching are also excluded from the calculation: RQ-00 *U*2, RQ-02 *U*2, RQ-03 *U*1, RQ-06 *Q*1, and RQ-06 *Q*2.

The result is a single aggregate *rms* value for every 10 second interval in the CES. A rough calculation can be made for the expected noise level of averaged data. For demodulated data, an individual diode has a white noise level of $\sim 600 \mu\text{K}/\sqrt{s}$, so the *rms* on 10 second scales is $\sim 190 \mu\text{K}$. For averaged data, the noise spectrum is dominated by $1/f$ noise for frequencies below ~ 100 Hz, with a typical power law slope of 1.6, so the demodulated data sensitivity value should be scaled by $\left(1 + \left(\frac{100}{0.1}\right)^{1.6}\right)^{-1/2}$. This yields an expected 10 second *rms* of 48 mK, for good weather conditions. The standard deviation of the *rms* values is approximately 11 mK, though the small number statistics involved lead to long tails on the distribution.

The statistic R_{10} is defined as the maximum value of the aggregate 10 second *rms* over the entire CES. A second statistic, R_{10}^{95} , is defined to be 95th percentile value (i.e. the maximum value after excluding the highest 5%). Cut thresholds are set at 150 mK for R_{10} and at 80 mK for R_{10}^{95} , significantly higher than the expectation value calculated above. If either of R_{10} or R_{10}^{95} exceeds the defined threshold, then the entire CES is excluded from the analysis. Because they are based on outliers of a statistical distribution, both statistics are extremely sensitive to changes in the fluctuations of the averaged data. R_{10} is focused particularly on brief flares in the noise level, but a longer stretch of bad weather is required to trigger the cut on R_{10}^{95} .

Good observing weather is not distributed uniformly throughout the Q-band observing season. Figure 8.1 shows the fraction of CES that pass the two cut criteria as a function of time of day or time of year. A strong diurnal variation is visible, with the worst weather occurring between UTC of 16:00 and 24:00, which corresponds to the afternoon and evening

1. Only ADC 1 is used because ADCs 0 and 2 show anomalously large noise in their averaged data after approximately run 1000. The cause is not understood, but it does not affect demodulated data.

by Chilean local time. This matches the intuition about daily weather pattern at the site, which features wind and higher temperatures in the afternoon but calm and cold nights. Also, it explains the difference in cut efficiency between the four CMB patches. Because the Q-band observing season lasted slightly more than seven months, the patches were observed at substantially different times of day (see Figure 8.2). Patch 2a was almost never observed between UTC 16:00 and 21:00, when the weather is worst, while Patch 6a missed all of the excellent weather between UTC 4:00 and 9:00.

8.3 Cut on Noise Model Goodness-of-Fit

The QUIET noise model involves fitting the power spectrum for every CES-diode to a three parameter function, as described in §7.3. CES-diodes that are not fit well by the model indicate problems with the data, often caused by bad weather. Also, the parametrized function is used to generate simulated timestreams for the Monte Carlo based analysis, so including data which are not well described by the noise model will lead to systematic differences between the actual and simulated data.

Two of the most common power spectrum features that are not accounted for in the noise model, scan synchronous signals and high frequency noise, are removed from the time ordered data in the filtering step of the analysis (see §3.1.6). To avoid cutting large amounts of data based on features that will be removed in the analysis, filters are applied before calculating the noise model goodness-of-fit.

- A low pass filter is applied to the Fourier transform of the data that sharply cuts off power above 4.6 Hz. The filter function has a value of 1 for frequencies below 4.6 Hz, value $F(\nu) = \cos^2 \left[\frac{\pi}{2} \left(\frac{\nu - \nu_1}{\nu_2 - \nu_1} \right) \right]$ for frequencies between 4.6 and 4.7 Hz, with $\nu_1 = 4.6$ Hz and $\nu_2 = 4.7$ Hz, and a value of 0 above 4.7 Hz.
- A high pass filter is applied with a shape given by the inverse of the noise power. For




Figure 8.1: Fraction of CES passing cuts based on R_{10} , in black, and R_{10}^{95} , in red, as a function of time of day and time of year.




Figure 8.2: Number of CES as a function of time of day for Patch 2a (upper left), Patch 4a (upper right), Patch 6a (lower left), and Patch 7b (lower right). Comparing this plot to Figure 8.1, we see that Patch 2a was observed only during parts of the day with good weather, while Patch 6a was mostly observed during the worst fraction of the day.

noise model parameters, f_{knee} , and α , defined in §7.3, the filter function is given by

$$F(\nu) = 1 + \left(\frac{f_{\text{knee}}}{\nu} \right)^\alpha.$$

- After applying the low and high pass filters, the time ordered data are averaged into 20 bins by azimuth. After the azimuthal structure has been calculated, the value in the appropriate bin is subtracted from each timestream sample, removing any such mode from the data.

Figure 8.3 shows the filtered power spectrum and noise model for an example CES-diode.




Figure 8.3: The binned power spectrum, calculated after filtering, for RQ-09 diode $Q1$ from CES 128.0 is shown in black. The noise model, with low and high pass filters applied is shown in red. The dashed lines delimit the frequency regions used for data selection statistics (10 to 200 mHz, 200 mHz to 2 Hz, and 2 to 4.6 Hz). The gray regions are not used for data selection.

A χ^2 statistic is calculated to determine the goodness-of-fit between the power spectrum

of the filtered data and the noise model (which is modified by applying the low and high pass filters). In the frequency range $[\nu_1, \nu_2]$, the goodness-of-fit between the binned power spectrum and the similarly binned noise model is given by

$$\chi^2 = \frac{1}{N} \sum_{\nu=\nu_1}^{\nu_2} n_\nu \left(1 - \frac{P_\nu}{\tilde{P}_\nu} \right)^2 \quad (8.1)$$

where N is the number of binned points in the frequency range, n_ν is the number of Fourier modes in each bin, P_ν is the value of the power spectrum, and \tilde{P}_ν is the value of the noise model. The expectation value of χ^2 is 1, if the noise model represents the true power spectrum of the data. This statistic is calculated for three frequency ranges, which are sensitive to different features in the data. The frequency ranges used are 10 to 200 mHz, 200 mHz to 2 Hz, and 2 to 4.6 Hz.

Note that the high pass filter used for this calculation does not match the filter used in the data analysis (see §9.1). That filter, which removes a linear function of azimuth from each left or right going scan segment, does not have such a simple description for its effect on the noise model. However, for the purposes of calculating the χ^2 statistics for data selection, the important effect of high pass filtering is to removing long time scale drifts which would contaminate the calculation of the binned azimuthal structure. The Fourier space filter, shaped as the inverse of noise power spectrum, accomplishes this goal.

8.3.1 Low Frequency (10 to 200 mHz) χ^2

Any CES-diode with a χ^2 statistic greater than 2.5, for the frequency range 10 to 200 mHz, is cut from the data analysis. The period of the QUIET CMB scan is between 10 and 20 seconds, so this range includes the scan frequency. Excess power near the scan frequency can generally be classified into two groups: sharp spikes or broad bump features.

The sharp scan synchronous feature consists of a very narrow spike – often just a single

Fourier mode. The narrow width indicates that it is stable over duration of the CES. This type of signal can be sourced by pickup of ground features through beam sidelobes. As the telescope scans in azimuth, the sidelobe sweeps across different regions of the area surrounding the telescope. The timescale for variations in the ground temperature is several hours at least, so ground pickup can create stable signals at the scan frequency and its harmonics. Figure 8.4 shows how a measure of the sharp scan synchronous signal varies with the calculated elevation of the triple reflection sidelobe, described in §6.7. The average value increases for low sidelobe elevations, demonstrating that ground pickup is responsible for at least part of the signal.




Figure 8.4: Scan synchronous signal statistic as a function of the calculated elevation of the triple reflection sidelobe for module RQ-09, diode *Q1*. Points are plotted in gray for all 1768 CES that pass cuts for this diode. Red points are calculated by binning the individual CES by sidelobe elevation. The scan synchronous signal statistic measures the significance of the power in a single Fourier mode at the scan frequency. Excess scan synchronous signal at low sidelobe elevations provides evidence that ground pickup is contributing to the sharp scan synchronous feature in the noise spectra.

If the contamination signal is truly constant throughout the entire scan, then it will be removed when the binned azimuthal structure is subtracted from the data. More dangerous

are signals near the scan frequency that vary slowly during the CES. These result in “broad bump” features in the power spectrum, as the excess power is spread across a region of Fourier space with width set by the coherence length of the signal. For example, in a CES with 10 second scan period, a contamination signal that is coherent over one minute timescales would contribute to the power spectrum in a range from 92 to 108 mHz. Typical broad bump features in the data have widths of up to 50 mHz, which corresponds to a coherence time of as little as 20 seconds. This type of excess power is exactly what one would expect from bad weather. As it scans, the telescope crosses patches of non-uniform sky temperature. Unlike the temperature pattern of the ground, these sky signals will drift, causing them to move in and out of the scan region. Since a time varying scan synchronous signal will not be removed by the binned azimuth filter, it is important to cut the CES-diodes affected by bad weather.

8.3.2 Medium Frequency (200 mHz to 2 Hz) χ^2

The medium frequency range, from 200 mHz to 2 Hz, is usually a good match to the noise model, with no systematic sources of excess power. A cut is still applied for any CES-diode with a χ^2 statistic greater than 2.5 in this range. This criteria removes very little data, as can be seen from Table 8.4.

8.3.3 High Frequency (2 to 4.6 Hz) χ^2

The high frequency range is defined from 2 Hz up to 4.6 Hz, where the signal is rolled off sharply by the low pass filter. Many of the unfiltered diode spectra have significant features at higher frequencies, but problems below 4.6 Hz are quite rare. Again, a threshold of 2.5 is used. Very few CES-diodes are removed by this cut, as for the medium frequency range.

8.4 Knee Frequency Cut

The knee frequency, as determined by the noise model, is sensitive to several issues in the data which should be cut. Residual Type-B glitching, left over from an imperfect correction, will greatly increase the knee frequency. If the temperature of the enclosure drifts significantly during a CES, the data will show a similar drift, due to the temperature susceptibility of the electronics boards, which adds low frequency power. Finally, atmospheric fluctuations during periods of bad weather show up in the demodulated data due to $\mathbf{I} \rightarrow \mathbf{Q}$ leakage, with a $1/f$ type spectrum.

The nominal knee frequency, for operation in good conditions, varies from detector diode to detector diode. While there is some module to module variation in amplifier $1/f$ noise, $\mathbf{I} \rightarrow \mathbf{Q}$ leakage is the main factor that determines the knee frequency of each diode. At the same time, the high leakage diodes are more susceptible to atmospheric fluctuations. This leads to a common distribution for knee frequencies, which scales from one diode to the next. The CES with an additional contribution to $1/f$ noise from imperfect Type-B removal or varying enclosure temperature show up as outliers from the main distribution.

The cut on knee frequency uses a threshold that is defined separately for each detector diode. The threshold is calculated using a procedure that takes two parameters, ϵ and κ , as inputs. The cutoff value for a particular diode is calculated by starting with the set of knee frequencies for that diode from all CES in the baseline data set. The mean, μ , and standard deviation, σ , of the distribution are calculated, after excluding the highest n outliers, where the parameter ϵ specifies n as a fraction of the full data set. The cut threshold is then defined to be $\mu + \kappa\sigma$.

For QUIET data, the parameters used for the knee frequency cut are $\epsilon = 0.15$ and $\kappa = 5$. Table 8.3 lists the threshold values. Figure 8.5 shows the knee frequency histograms for all detector diodes; the common shape of the distribution is quite evident here.

Table 8.3. Knee frequency cut thresholds for all detector diodes, in mHz

Module	$Q1$	$U1$	$U2$	$Q2$
RQ-00	26.61	6.88	7.69	21.68
RQ-01	22.56	16.26	40.67	26.31
RQ-02	58.38	17.00	13.71	35.78
RQ-03	36.59	7.45	7.29	40.97
RQ-04	—	5.68	7.26	24.12
RQ-05	69.38	22.79	26.23	57.61
RQ-06	40.80	8.42	14.87	35.05
RQ-07	72.83	12.63	10.73	64.24
RQ-08	31.49	9.57	—	25.67
RQ-09	52.00	18.59	19.49	66.86
RQ-10	30.74	7.33	5.56	29.10
RQ-11	18.73	9.92	12.01	16.47
RQ-12	21.60	10.27	13.53	24.87
RQ-13	106.26	32.86	25.50	83.52
RQ-14	47.68	11.06	9.33	48.92
RQ-15	18.85	10.45	20.44	42.99
RQ-16	—	—	—	—

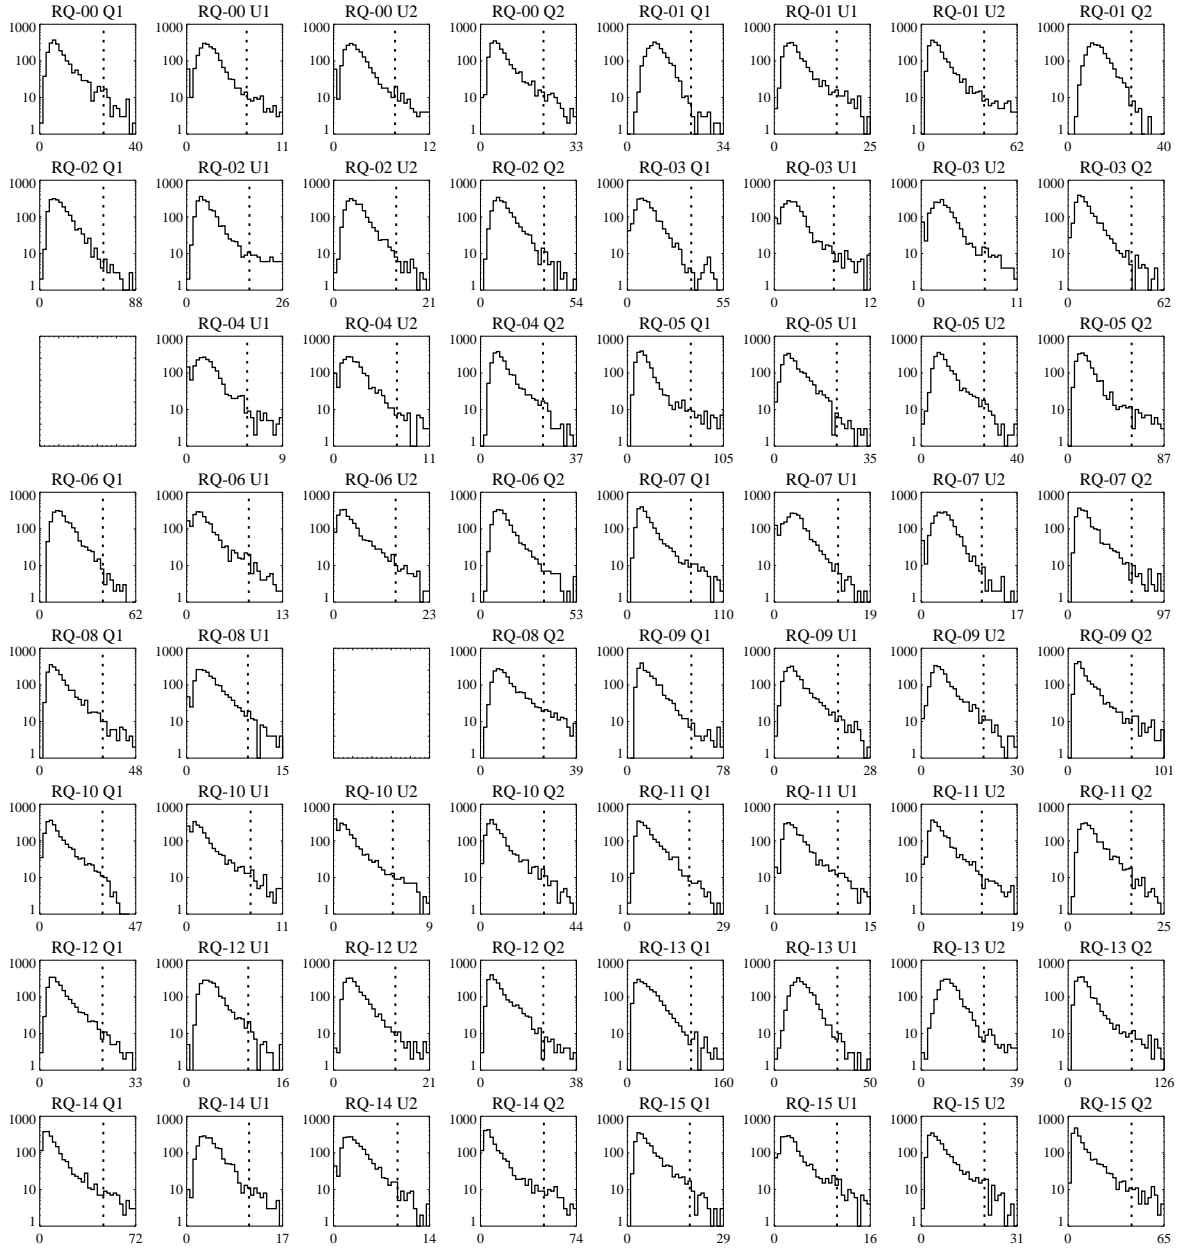


Figure 8.5: Knee frequency histograms for each detector diode, including all CES from the baseline data set. The cut thresholds listed in Table 8.3 are indicated by vertical lines.

8.5 Sun Position Cut

The sidelobes described in §6.7 can create contamination signals in the data when they scan across the ground or the sun. For the case where a sidelobe hits the ground, the signal is static over a CES time scale and is removed by filtering, as described in §9.1.3. Sun contamination, on the other hand, varies during the CES, as the sun drifts through the sidelobe region under sidereal motion. This effect can not be removed by filtering, so cuts are applied to remove the affected data.

The position of the sidelobe during every CES is calculated as follows. First, the mean az , el , and dk are defined for the scan; el and dk are constants, and the scan az is calculated as the mean of the minimum and maximum values during the CES. The telescope pointing and orientation vectors, \vec{p} , \vec{s} , and \vec{r} , are calculated from these values, as described in §7.1.1. An additional pointing vector, \vec{p}_s , specifies the position of the sun at the midpoint of the CES. Now, the sun position relative to the telescope boresight can be expressed by calculating the position of \vec{p}_s in a spherical coordinate system defined by the telescope basis.

$$\theta_s = \arccos[\vec{p}_s \cdot \vec{p}_0] \quad (8.2)$$

$$\phi_s = \arctan \left[\frac{\vec{p}_s \cdot \vec{r}_0}{\vec{p}_s \cdot \vec{s}_0} \right] \quad (8.3)$$

Regions of (θ_s, ϕ_s) that suffer from sun contamination are identified by visual inspection and χ^2 statistics for individual CES maps. In the case of the triple reflection sidelobe, the region where spots were noticed in the maps matches exactly with the predicted sidelobe location from physical optics simulations. Figure 8.6 shows the regions of sun coordinates where CES are removed from the data. The cut region for the triple reflection sidelobe is $44^\circ < \theta_s < 60^\circ$ and $120^\circ < \phi_s < 180^\circ$. Two regions are cut for the second sidelobe: $60^\circ < \theta_s < 65^\circ$ and $-90^\circ < \phi_s < 0^\circ$ plus $44^\circ < \theta_s < 60^\circ$ and $-75^\circ < \phi_s < -15^\circ$. Note that the second sidelobe cut is applied only to modules RQ-00, RQ-01, and RQ-02, which clearly

show the contamination.




Figure 8.6: The θ_s and ϕ_s coordinates for all CES where the sun is less than 90° from the boresight are shown by the black points. The region cut due to the triple reflection sidelobe is outlined in red. The region cut due to the second sidelobe is outlined in blue (but this cut is applied to modules RQ-00, RQ-01, and RQ-02 only).

Because the sun position for each CES is calculated at only one value of azimuth and one moment in time, this statistic is a somewhat fuzzy measure of which ones are contaminated. Additionally, the sidelobe positions could vary somewhat from module to module. To compensate for these factors, the cut regions are defined to be broader than is strictly necessary, creating a buffer around the affected zones.

8.6 Glitch Cut

A glitch cut is used to remove CES-diodes with spikes remaining in the filtered timestreams. This cut is applied at four timescales: 1 sample (50 Hz), 5 samples (10 Hz), 50 samples (1 Hz), and 500 samples (0.1 Hz). The statistic is calculated by taking a boxcar sum over every interval of the specified length (so for a timestream with N samples, the number of 500 sample intervals is $N - 500$, not $N/500$). For each CES-diode, the glitch statistic at each timescale is defined to be the absolute value of the most discrepant interval, in units of the standard deviation of all intervals at that timescale.

A CES-diode is cut if any of the four glitch statistics exceeds a threshold of 6 standard deviations. A 4000 second CES contains 200000 samples, so the statistical probability of a 6σ glitch is only 0.04%. While the glitch cut does directly cut on the demodulated data timestreams, the extremely high threshold means that the events being cut are almost certainly spurious, and the true noise distribution of the data is not being significantly biased. It is reassuring to see, in §8.8, that the glitch cut removes very little data.

8.7 Type-B Cut

Type-B glitching, and the method for correcting it, is described in §6.4. For a small number of CES-diodes, the correction does not manage to completely remove the effect (perhaps due to an inaccurate noise estimate). These data are usually removed by the cut on knee frequency, as the main effect of Type-B glitching is to add $1/f$ noise. However, a more direct way to detect any residual effect is to examine the correlation between demodulated and averaged data. This correlation can be seen in the left hand panel of Figure 6.4. When Type-B glitching is present, the graph of demodulated vs averaged samples has a characteristic non-linear shape. In the absence of glitching, the graph is linear, with slope equal to the diode $\mathbf{I} \rightarrow \mathbf{Q}$ leakage.

The Type-B cut is applied by calculating a reduced χ^2 for fitting the demodulated vs averaged data graph to a line. The cut is very loosely applied, with a threshold of 10. There are some CES-diodes that fail this cut due to reasons other than Type-B glitching. Other potential causes for a non-linear relationship between demodulated and averaged data are cases where the DC level for a channel has a sudden shift, which would affect averaged but not demodulated data. This type of irregular operation is removed from the analysis by the Type-B cut.

8.8 Summary of Cut Efficiency

Table 8.4 lists the number of CES-diodes rejected by each cut. The “Total Cut” column lists the number of CES-diodes cut when the criteria is applied directly to the baseline data set. The “Uniquely Cut” column lists the number of CES-diodes cut if the criteria is applied after all others. For example, there are 1076 CES-diodes targeted on Patch 2a with single sample glitches that exceed the 6σ threshold. However, nearly all of these data are also rejected by other cuts (the same glitch could still be present at longer time scales or it could be an artifact of bad weather, which is detected by the weather cut), so the number of CES-diodes uniquely cut by the one sample glitch criteria is only 5. Since the glitch cuts at various timescales overlap significantly in terms of the data rejected, an additional row is included in Table 8.4 describing the effect of the combined glitch cut at four timescales. Similar rows are included for the combined weather cut and for the χ^2 cut in three frequency regions.

It can be clearly seen that the most significant cut is based on the low frequency χ^2 statistic, with the weather and knee frequency cuts having the next largest impact. All three of these criteria are sensitive to atmospheric structure, so the overall effect of the data selection is primarily to reject bad weather. This has the biggest effect on data from Patch 6a and the least effect on Patch 2a data, for the reasons discussed in §8.2.

Table 8.4. Number of CES-diodes rejected by various cuts

Criteria	Total Cut ^a				Uniquely Cut ^b			
	2a	4a	6a	7b	2a	4a	6a	7b
Baseline data set	46915	38006	38922	14718	—	—	—	—
Weather cut R_{10}	1550	4154	7860	1388	272	173	249	73
Weather cut R_{10}^{95}	1054	3720	7550	1574	0	5	136	140
Weather cut total	1550	4216	8294	1760	419	462	830	377
Low frequency χ^2	4133	6902	12027	2667	2631	2989	3380	1056
Medium frequency χ^2	422	1537	4521	529	15	5	3	4
High frequency χ^2	413	658	2364	226	203	149	212	46
χ^2 cut total	4452	7087	12299	2730	2910	3203	3688	1125
Knee frequency	1614	3958	7655	1469	363	433	361	153
Sun cut	1572	768	2480	168	1284	650	1380	78
Glitch, 1 sample	1076	1040	3502	292	5	0	7	0
Glitch, 5 samples	1108	1254	3882	360	9	3	4	1
Glitch, 50 samples	1245	2926	5983	733	7	10	9	4
Glitch, 500 samples	854	3086	5991	701	26	36	48	15
Glitch cut total	1725	3461	6545	871	745	166	149	304
Type-B cut	1342	1746	3355	638	585	356	253	100
All cuts	8491	9702	15852	3690	—	—	—	—
Percentage passing cuts	81.9%	74.5%	59.3%	74.9%	—	—	—	—

^aNumber of CES-diodes, out of the entire baseline data set, cut by each criteria.

^bNumber of CES-diodes that are cut by this criteria *but no others*.

CHAPTER 9

DATA ANALYSIS

The analysis of CMB polarization data consists of two main steps. First, the time-ordered data is made into a sky map, accounting for the scan pattern used in the observation and applying any necessary filters. Once the map has been made, it is used to estimate the angular power spectrum C_ℓ values. This process must be done very carefully, since the signal-to-noise is small (especially for B -modes!). Also, the signal itself is a Gaussian random field, making it very difficult to distinguish from the instrumental or atmospheric contributions.

A method often used to separate the noise and signal contributions to the map is to maximize the following likelihood function (Bond et al., 1998):

$$\mathcal{L}(a_p) = \frac{1}{(2\pi)^{N/2} |C_T(a_p) + C_N|^{1/2}} \exp \left[-\frac{1}{2} \Delta^T (C_T(a_p) + C_N)^{-1} \Delta \right] \quad (9.1)$$

where Δ is the length- N vector of map pixel values, C_T is the pixel–pixel correlation matrix for the calculated theory parametrized by a_p , and C_N is a correlation matrix that accounts for detector noise as well as any filtering that may have been applied to the data in the map making process. The parametrization, a_p , could be anything — cosmological parameters, for example — but it is usually chosen to be $\ell(\ell+1)C_\ell/2\pi$, averaged across bins in ℓ .

The maximum likelihood method treats the data optimally, but it is computationally difficult for large maps. In particular, the step of inverting the matrix $C_T(a_p) + C_N$ scales as $\mathcal{O}(N^3)$.

An alternate method, known as pseudo- C_ℓ , produces nearly optimal results with much faster computation times (Wandelt et al., 2001). In the pseudo- C_ℓ method, the filtered time-ordered data is binned into a map as usual, but the full correlation matrix C_N is not calculated. Instead, a spherical harmonic transformation is applied to the naive map, resulting in a biased “pseudo-power spectrum” \tilde{C}_ℓ . Monte Carlo (MC) simulations are used

to correct the bias and estimate the uncertainty of the final result.

Two independent pipelines have been developed for QUIET data analysis, one using the maximum likelihood technique and the other using pseudo- C_ℓ . Only the pseudo- C_ℓ pipeline will be discussed here. Filtering and other processing of the time-ordered data is described in §9.1. Details of map making are found in §9.2. The power spectrum estimator is discussed in §9.3.

9.1 Time-Ordered Data Processing

An advantage of the pseudo- C_ℓ method is the freedom to employ a wide range of filtering schemes. Instead of worrying about the effect of an operation on the noise correlation matrix, it is simply accounted for in the Monte Carlo. As a result, the QUIET pseudo- C_ℓ pipeline relies heavily on processing of the time-ordered data (TOD).

Besides the usual filtering tasks – attenuating $1/f$ noise, cutting out contaminated frequency ranges – other corrections to the data are performed in this step. Figure 9.1 illustrates the operations applied to the TOD to go from files on disk to cleaned data ready for map making.

The first several TOD processing actions are more akin to calibration and are discussed elsewhere. The Type-B glitch correction is covered in §6.4. The timing offset, which is applied to align the radiometer and telescope data, is found in §6.6. The pointing and gain models are described in §7.1 and §7.2, respectively. The double demodulation step reduces the sampling rate from 100 Hz to 50 Hz, by simply averaging the pointing samples and differencing the demodulated samples from \uparrow and \downarrow clock states. The advantages of double demodulation are discussed in §3.5.1. After these steps, three filtering operations are applied to the polarization data, as described below.

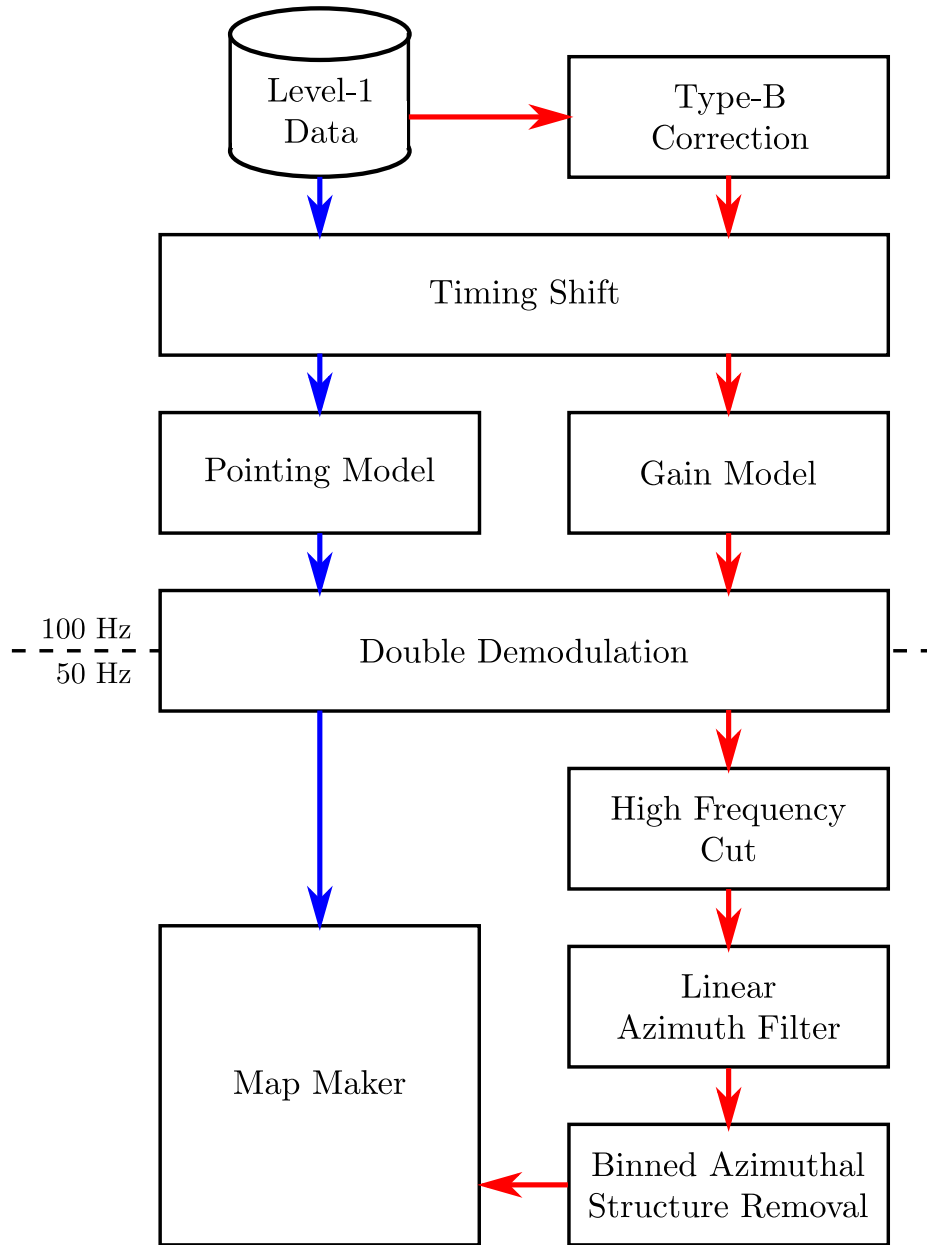


Figure 9.1: Filtering process for time-ordered data. The flow of polarization data from the modules is shown by red arrows; the blue arrows represent the telescope pointing data. Above the dashed line, both timestreams have a sample rate of 100 Hz; after the double demodulation step, the sample rate is reduced to 50 Hz.

9.1.1 High Frequency Cut

The high frequency cut is a very sharp frequency domain filter that starts to roll off at $\nu_1 = 4.6$ Hz and eliminates all power above $\nu_2 = 4.7$ Hz. The filter shape is given by

$$F(\nu) = \begin{cases} 1 & \nu < \nu_1 \\ \cos^2 \left[\frac{\pi}{2} \left(\frac{\nu - \nu_1}{\nu_2 - \nu_1} \right) \right] & \nu_1 < \nu < \nu_2 \\ 0 & \nu > \nu_2 \end{cases} \quad (9.2)$$

This extremely abrupt filter is chosen to totally remove the range of frequencies that often suffer from spikes of narrow band noise, as discussed in §6.5, but to avoid cutting into the CMB signal band, which is $\lesssim 5$ Hz.

9.1.2 Linear Azimuth Filter

The linear azimuth filter subtracts a slope and offset from each left going or right going segment of the azimuth scan. In addition to removing scan synchronous signals, which might be due to atmospheric or ground pickup, the slope subtraction effectively acts as a high pass filter, attenuating $1/f$ noise in the data. While the CMB does feature fluctuations with very long wavelengths, which would appear as slopes across the particular patch in question, QUIET is incapable of measuring these modes, so discarding them is acceptable.

The line to subtract is calculated by linear regression with the azimuth coordinate, not time, as the independent variable. Since the telescope motion does not have constant velocity — it slows down to turn around at the edges of the scan range — the azimuth filter does not have a simple representation in the frequency domain, though signals slower than twice the scan frequency will clearly be diminished.

The decision to remove only a linear slope, instead of a higher order polynomial, can be examined by studying the amplitudes of higher order polynomials in the time-ordered data.




Figure 9.2: Mean values of the coefficients for Legendre polynomials of order zero through five, calculated for module RQ-06 $Q1$ for all scan segments of five CES and averaged. The CES shown, which all pass the data selection, are approximately ordered from worst to best weather conditions. Their run ids are 1127.1 (red), 290.3 (orange), 1492.2 (green), 1582.1 (light blue), and 1373.2 (dark blue).

For several CES, spanning a variety of weather conditions, the time-ordered data is processed by removing Legendre polynomials up to 5th order from each left or right-going scan segment. The mean amplitude for each polynomial is calculated, as well as the error on the mean to determine whether that polynomial was significantly non-zero for the CES. The results are shown in Figure 9.2. Module RQ-06, diode $Q1$ was chosen for illustrative purposes because it has relatively large $\mathbf{I} \rightarrow \mathbf{Q}$ leakage, which leads to increased sensitivity to weather and higher $1/f$ noise. For two of the five scans shown, the first order coefficients are significantly different from zero, indicating that linear structure is present. Higher order coefficients show less significant amplitude so, while removing them would result in harder filtering, it would not target any contamination signal. The zero order coefficients are suppressed in Figure 9.2 because the mean of the entire CES was subtracted before solving for the polynomials.

9.1.3 *Binned Azimuthal Structure Removal*

While the linear azimuth filter is very effective at removing low frequency noise and large spatial gradients, pickup of the ground emission via beam sidelobes can cause contamination with more complicated structure but that is constant over the timescale of one CES. After all other filters have been applied, the data are binned by azimuth into 20 bins, for each detector diode separately. Any structure remaining after averaging over the entire CES is subtracted out of the timestreams.

The amplitude of the azimuthal structures can be quite large, up to several hundred μK , and their removal is perfect only if the contamination signal is constant. If the azimuthal structure varies during the CES, it will manifest in the noise spectrum as a broadened peak around the scan frequency. Because the data selection uses the goodness of fit for the noise spectrum to the noise model after applying filters, those CES-diodes whose azimuthal structure is fully removed by this process are accepted, but the CES-diodes with residual excess near the scan frequency are cut.

9.1.4 *Effect of Combined Filters on CMB Signal*

Figure 9.3 shows the power spectrum for a single CES, simulated with CMB signal only (no noise), before and after the filters are applied. The scan periodicity creates a strong harmonic structure in the signal. The linear azimuth filter and binned structure removal primarily affect the fundamental harmonic. The part of the spectrum removed by the high frequency cutoff is already significantly attenuated in the unfiltered data, due to beam smoothing of the CMB signal.

The filtered spectrum contains about 76% as much total power as the unfiltered spectrum, though the exact amount of signal lost will vary from realization to realization.




Figure 9.3: Power spectra for a simulated CMB observation with no detector noise are shown before filtering (red) and after filtering (black). The scan period for this observation is 10 seconds, which creates harmonic structure in the spectra. The linear azimuth filter and binned structure removal greatly reduce the peak at the fundamental scan harmonic, but have only a small effect on higher harmonics. The high frequency cutoff can be seen clearly. For reference, the white noise level of a single diode is typically $3.6 \times 10^5 \mu\text{K}^2/\text{Hz}$ and the white noise level for the entire array is $4096 \mu\text{K}^2/\text{Hz}$.

9.2 Map Making

After filtering, the time-ordered data for all diodes are combined to produce **Q** and **U** maps of the QUIET patches. The samples summed into each pixel are weighted by their inverse variance, calculated from the white noise level found by the noise model (see §7.3). For demodulated data values d_i , with white noise level σ_i , the inverse noise weighted Stokes parameters, $\tilde{\mathbf{Q}}$ and $\tilde{\mathbf{U}}$ for a pixel are calculated by simply summing over all samples that hit

the pixel.

$$\tilde{\mathbf{Q}} = \sum_i \frac{1}{\sigma_i^2} d_i \cos(2\psi_i) \quad (9.3)$$

$$\tilde{\mathbf{U}} = \sum_i \frac{1}{\sigma_i^2} d_i \sin(2\psi_i) \quad (9.4)$$

The parallactic angle values, ψ_i , are the calculated orientations of the detector diode polarization sensitivity axis for each sample.

As the $\tilde{\mathbf{Q}}$ and $\tilde{\mathbf{U}}$ maps accumulate, the inverse variances of the pixels are calculated as

$$(\mathbf{Q}\mathbf{Q})^{-1} = \sum_i \frac{1}{\sigma_i^2} \cos^2(2\psi_i) \quad (9.5)$$

$$(\mathbf{U}\mathbf{U})^{-1} = \sum_i \frac{1}{\sigma_i^2} \sin^2(2\psi_i) \quad (9.6)$$

$$(\mathbf{Q}\mathbf{U})^{-1} = \sum_i \frac{1}{\sigma_i^2} \cos(2\psi_i) \sin(2\psi_i) \quad (9.7)$$

Calculating the pixel noise levels by equations (9.5) through (9.7) implies that the samples are all independent. This is clearly not true, due to noise correlations between detector diodes and filtering applied to the timestreams. This leads to bias in the unweighted \mathbf{Q} and \mathbf{U} maps, which are calculated by inverting the matrix described in equations (9.5) through (9.7) and applying it to the inverse noise weighted maps. This bias is corrected in the power spectrum estimation step.

The pixelization used for the maps is known as Healpix – Hierarchical, Equal Area, and isoLatitude Pixelization (Górski et al., 2005). A key advantage of using Healpix for CMB analysis is that the spherical harmonic transform can be calculated more efficiently in a system where pixels are arranged along rings of constant latitude. Maps are made with Healpix $N_{side} = 256$, which corresponds to pixel size of ~ 0.05 square degrees, about one fifth of the beam solid angle.

Maps of the inverse noise weighted Stokes parameters, $\tilde{\mathbf{Q}}$ and $\tilde{\mathbf{U}}$, are shown in Figures 9.4 through 9.7, with point source masks indicated, as discussed in §9.2.1. Pixels near the edge of the patch have higher variance, due to less integration time, so features in those regions are de-emphasized.

9.2.1 Point Source Masks

QUIET is fairly insensitive to point sources, due to its large beam and also because most sources are unpolarized or weakly polarized. While this can make calibration difficult, it is useful for the CMB analysis, where point sources represent an unwanted contaminant. The QUIET CMB patches, which are located in very clean patches of the sky, still contain two sources which are sufficiently bright and polarized that they must be masked out.

Centaurus A, or Cen A for short, is a bright galaxy with two large lobes that feature strong polarized emission. The Cen A polarization signal is ~ 1 mK, as seen by QUIET – so bright that it was considered for use as a calibration source (Huff, 2008). Cen A is located at the edge of Patch 2a, where it is occasionally observed by modules at the array periphery. A circular region centered on Cen A, with 2 degree radius, is masked out of the Patch 2a maps by setting the inverse variances to zero, blowing up the error values for those pixels. Because Cen A is so far out from the patch center, masking a large region is not very costly in terms of lost sensitivity.

The second source detected in the QUIET maps is Pictor A, a jet galaxy. Pictor A is fainter than Cen A by almost an order of magnitude, but its location near the center of Patch 4a means that much more observing time is spent on it. Pictor A is polarized at a significant level (Wright et al., 2009), making it more easily detectable for QUIET. A smaller mask, with radius of 1 degree, is used for at this location.




Figure 9.4: Inverse noise weighted \tilde{Q} (left panel) and \tilde{U} (right panel) maps for Patch 2a. Thin black lines are drawn at 10 degree intervals in R.A. and Dec. Centaurus A can be clearly seen near the edge of the patch; the circle indicates the 2 degree radius region that is masked for power spectrum estimation.




Figure 9.5: Inverse noise weighted \tilde{Q} (left panel) and \tilde{U} (right panel) maps for Patch 4a. Thin black lines are drawn at 10 degree intervals in R.A. and Dec. Pictor A can be seen in the Patch 4a maps; the circle indicates the 1 degree radius region that is masked for power spectrum estimation.




Figure 9.6: Inverse noise weighted \tilde{Q} (left panel) and \tilde{U} (right panel) maps for Patch 6a. Thin black lines are drawn at 10 degree intervals in R.A. and Dec. No point sources are masked in Patch 6a.




Figure 9.7: Inverse noise weighted \tilde{Q} (left panel) and \tilde{U} (right panel) maps for Patch 7b. Note that the color scale for the Patch 7b maps has a smaller range than the other patches (Figures 9.4, 9.5, and 9.6), because there is significantly less sensitivity for Patch 7b. Thin black lines are drawn at 10 degree intervals in R.A. and Dec. No point sources are masked in Patch 7b.

9.3 Power Spectrum Estimation

The MASTER (Monte Carlo Apodized Spherical Transform Estimator) method is a pseudo- C_ℓ technique designed to analyze maps that cover only a fraction of the full sky (Hivon et al., 2002). For an overview of the MASTER method, we will start by describing how the method is applied to CMB temperature data.

A CMB temperature map covering only a patch of the sky can be written as the measured temperature values for each pixel, T_i , and a weighting function, W_i , which is chosen heuristically¹. The spherical harmonics, which are orthogonal over the full sky, are not all orthogonal on a small patch, but their coefficients can be calculated anyway.

$$\tilde{a}_{\ell m} = \Omega \sum_i T_i W_i Y_{\ell m}^*(i) \quad (9.8)$$

where Ω is the solid angle of a pixel. The angular power spectrum, \tilde{C}_ℓ , calculated as the variance of the $\tilde{a}_{\ell m}$ values is referred to as the pseudo-power spectrum.

In addition to the CMB signal, the map also contains filtered instrumental noise, which contributes to the pseudo-power spectrum. The \tilde{C}_ℓ values are related on average to the full sky power spectrum by

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle + \langle \tilde{N}_{\ell'} \rangle \quad (9.9)$$

where B_ℓ is the beam window function, F_ℓ is the transfer function due to filtering of the data, $M_{\ell\ell'}$ is a mode-mode coupling caused by only observing a small fraction of the sky with non-uniform coverage, and $\langle \tilde{N}_\ell \rangle$ is the noise bias. With the pseudo-power spectrum measured from the data, equation 9.9 can be solved to obtain an estimate of the true power spectrum.

1. For QUIET, the weighting function used is the inverse of the sum of the variances from noise and the expected signal. For B -modes, this is just the inverse noise variance, but for E -modes the signal contribution can be significant.

The calculation of B_ℓ is described briefly in §7.4. The filter transfer function is found by analyzing noise-free simulations of the CMB signal, with the usual filtering applied (like Figure 9.3, but with the analysis continued through to power spectrum estimation). The noise bias is also calculated from simulation, but using noise only data. The final piece is $M_{\ell\ell'}$, which is computed from the pixel weighting function W_i . This term couples power between different values of ℓ , matching the intuition that different spherical harmonics can not be cleanly distinguished over a small region.

For the case of polarization, the single pseudo-power spectrum \tilde{C}_ℓ is replaced by a pair of pseudo-power spectra, \tilde{C}_ℓ^{EE} and \tilde{C}_ℓ^{BB} . The relation between pseudo-power spectra and true power spectra includes not only coupling between different multipoles, but also coupling between E -modes and B -modes:

$$\begin{pmatrix} \langle \tilde{C}_\ell^{EE} \rangle \\ \langle \tilde{C}_\ell^{BB} \rangle \end{pmatrix} = \sum_{\ell'} \begin{pmatrix} M_{\ell\ell'}^+ & M_{\ell\ell'}^- \\ M_{\ell\ell'}^- & M_{\ell\ell'}^+ \end{pmatrix} \begin{pmatrix} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'}^{EE} \rangle \\ F_{\ell'} B_{\ell'}^2 \langle C_{\ell'}^{BB} \rangle \end{pmatrix} + \begin{pmatrix} \langle N_\ell^{EE} \rangle \\ \langle N_\ell^{BB} \rangle \end{pmatrix} \quad (9.10)$$

Leakage between E -modes and B -modes can lead to a serious problem. Calculating and inverting the mode coupling matrix from equation (9.10) removes the leakage of E to B power on average, but a pure E -mode signal will still contribute variance to the estimated BB spectrum. Since the EE spectrum is so much larger, this added noise could dominate the real B -mode signal. A solution to this problem, known as a “pure” B -mode estimator, has been developed (Smith (2006), Smith & Zaldarriaga (2007)). This method classifies polarization modes in the observed region into three categories: pure E -modes, pure B -modes, and ambiguous modes. Additional terms included in the pseudo-power spectrum estimator can compensate for the ambiguous modes, resulting in \tilde{C}_ℓ values that contains contributions only from the pure B -modes. However, the entire E -to- B leakage effect is small enough that it will not be detected at the current sensitivity level of QUIET. So,

though the analysis code does make use of Kendrick Smith’s `ppc1`² code for pseudo-power spectrum estimation, the method of using only pure B -modes is not activated.

With all of the components of equation (9.10) calculated, the issue still remains that the relation applies only for the average over many identical analyses. To address this, additional end-to-end simulations are performed, now containing both signal and noise. Analysis of these simulations provides the distribution of pseudo-power spectrum outcomes for a given input power spectrum. The width of this distribution can serve as an estimate of the uncertainty on the result. A more rigorous evaluation of the error can be obtained by varying the input power spectrum to determine the point at which the output distribution becomes statistically inconsistent with the \tilde{C}_ℓ values obtained from the data.

9.3.1 Using Day-to-Day Cross-Correlation to Avoid Noise Bias

A useful trick is employed in the QUIET analysis to avoid the danger of leftover bias due to imperfect noise characterization. For a given patch, the cumulative map for the entire season involves many CES gathered over the course of seven months. This map can be thought of as the sum of a large number of daily maps, which all contain a common CMB signal plus a noise contribution that should be uncorrelated from day to day. Measuring the power spectrum of the cumulative map is equivalent to summing up all possible auto and cross-correlations of the daily maps. However, because the noise is independent in each of the daily maps, the only contribution to the noise bias comes from the autocorrelations.

There are roughly 150 daily maps for each of the QUIET patches, so the autocorrelations account for only a percent or so of the CMB signal. Subtracting the auto-correlation contribution to the pseudo-power spectrum makes only a small increase in the error bar of the result but it ensures that errors in the noise model will not contribute bias, except for improbable cases where the instrumental noise is correlated from one day to the next.

2. Stands for “pure pseudo- C_ℓ ”

9.3.2 Overview of Analysis Procedure

The QUIET analysis procedure for one patch is as follows:

- A data selection query produces the list of CES-diodes to include in the analysis.
- The time-ordered data is filtered and binned into daily maps, which are stored separately but also combined to make the full season map.
- The transfer matrix is calculated from the pixel weight functions.
- Signal-only simulations are performed to calculate the filter transfer function. Noise-only simulations are used to calculate the noise bias, though this step is not necessary if only the daily auto-correlations are subtracted, as described above.
- The pseudo-power spectrum is calculated for the full season map, as well as for all the daily maps. The daily map spectra are subtracted from the full season map spectrum to eliminate the noise bias. The remaining terms from equation 9.10 are applied to recover the estimated C_ℓ values.
- Full signal plus noise simulations are analyzed to determine the width of the \tilde{C}_ℓ distribution. For a final result, additional sets of simulations would be generated and analyzed to find the range of input power spectra that are consistent with the data.

CHAPTER 10

ASSESSING DATA QUALITY

This chapter describes current efforts to assess the quality of the QUIET Q-band data.

Measuring CMB polarization, which has an amplitude orders of magnitude below the instantaneous sensitivity of the detector, requires long, careful observation. The key feature that makes this feasible is that the CMB signal itself remains constant while noise averages away. This property can be turned around by constructing and analyzing maps made from the difference of two data subsets scanning the same patch. Because the CMB signal is identical in each subset, it cancels out in the difference. If all sources of noise are treated correctly, the power spectrum of the difference map will be consistent with zero input CMB power. This type of check, called a null test, is the main technique used to assess data quality for QUIET.

Of course, many things can go wrong in the observations or analysis that would add power to the null map. For example, calibration error in one subset of the data would lead to CMB signal contributions that do not difference away perfectly. Incorrect reconstruction of the telescope pointing can shift or otherwise distort the signal maps for each half of the null test. Non-stationary noise, such as bad weather, can contribute fluctuations to the maps, beyond what is predicted by the Monte Carlo.

For each example of an effect that can cause a null test to fail, there will be a particular division of the data that maximizes the failure. To take advantage of this, QUIET employs a suite of 40 null tests, with divisions chosen based on concerns regarding the instrument, observations, or calibration. There do exist problems, such as a constant overall calibration error, that will not be detected by a null test, but statistically passing the full suite provides very strong evidence that the data set is clean and well understood. Focusing the analysis effort on achieving a successful null suite result provides the benefits of a blind analysis – problems in the data can be discovered and addressed without danger of bias.

Section 10.1 describes some technical details necessary for evaluating the power spectrum of null maps in a pseudo- C_ℓ pipeline. The various null tests making up the null suite are listed in §10.2. The criteria used to evaluate the results of the null suite are discussed in §10.3. The results of the full null suite, which indicate that there is some residual contamination in one region of the QUIET data, are found in §10.4.

10.1 Calculating Null Power Spectra

Analysis using the pseudo- C_ℓ framework described in §9 leads to some intricacies that must be addressed to correctly calculate the power spectra of null maps. The problem is that the sub-maps, which are differenced to create a null map, will not necessarily have the same filter transfer function or pixel weighting function. A procedure for correctly handling these issues, so that the null spectrum will have an expectation value of zero, is described in Smith (2009), and is summarized here.

The first detail to address is what weight function should be used for the null map. Ideally, the same function would be used for null tests as for the non-null data analysis, but there will be cases where a pixel is observed only in one of the two data subsets and can not be used. Instead, the weight function is chosen to be

$$W_i = \frac{1}{\sigma_{1,i}^2 + \sigma_{2,i}^2} \quad (10.1)$$

For the case where a pixel is not observed in one sub-map, the variance of that pixel is infinite and W_i is set to zero, as desired.

The second issue has to do with the filter transfer function. If the two sub-maps have different transfer functions, they will contain different amounts of power for the same CMB signal input. Furthermore, even in the case where the transfer functions are identical, the filtering might affect each sub-map in a different way. As an example, consider a typical

high pass filter, used to remove $1/f$ noise. When applied to an azimuth scan, this filter will not affect the modes which are parallel to the scan direction (because they are measured quickly by the back and forth telescope motion), but it will remove the modes in the sky drift direction (because they are measured very slowly as the sky rotates through the scan region). If the division of the data results in sub-maps that have considerably different scan directions, then different modes will be missing from each map and affecting the signal cancellation, even though the transfer function is the same for both cases.

The solution to this problem is to calculate three transfer functions, using signal only simulations as usual. F_ℓ^{11} and F_ℓ^{22} are the usual transfer functions, calculated separately for each sub-map. The third transfer function, F_ℓ^{12} , is calculated by measuring the cross pseudo-power spectrum between the two sub-maps which are made from identical signal only simulations. With these three transfer functions, the null power estimator can be calculated as

$$\tilde{C}_\ell^{null} = \frac{\tilde{C}_\ell(m_1)}{F_\ell^{11}} - 2\frac{\tilde{C}_\ell(m_1, m_2)}{F_\ell^{12}} + \frac{\tilde{C}_\ell(m_2)}{F_\ell^{22}} \quad (10.2)$$

where $\tilde{C}_\ell(m_{1,2})$ are the usual pseudo-power spectrum estimators applied to the sub-maps and $\tilde{C}_\ell(m_1, m_2)$ is the cross pseudo-power spectrum of the two maps. Note that the differencing of the two data subsets occurs here at the pseudo- C_ℓ stage, rather than the map stage, because it becomes impossible to create a true null map in this case. While the expectation value of the estimator is zero, its variance will include some dependence on the signal power, so it is necessary to use a realistic signal in the null test Monte Carlo.

10.2 The QUIET Null Suite

The QUIET null suite consists of a total of 40 separate null tests, applied to each of the four CMB patches, and each analyzed for 10 bins of EE power and 10 bins of BB power. The tests are chosen to target potential issues in the data, due to changes in the instrument

configuration, possible calibration errors, or biased data selection. When it is feasible, the division into subsets is chosen so that they each contain approximately equal amounts of data, to maximize the sensitivity of the result.

The data analysis procedure, which involves processing the time ordered data one CES at a time and building up maps, makes it very easy to run many null tests in parallel. Instead of a single map, all 40 null maps are built up simultaneously, with only a single pass through the TOD. The same benefit is realized in the Monte Carlo steps.

A description and rationale for each null test is provided here. The tests are organized into the following categories: divisions by module, divisions based on pointing, temporal divisions, divisions based on detector performance, weather divisions, and housekeeping divisions. The database queries used to select the data subsets for each null test can be found in Appendix B.

10.2.1 Module Divisions

Divisions by module check whether all modules observe a consistent sky signal. Many instrumental problems would be expected to affect only certain modules, or to affect some modules more than others, and would be detected by a null test that segregates the badly behaving detectors from the ones operating normally.

MAB A vs B + C

Modules RQ-00 through 06 vs others.

MAB B vs A + C

Modules RQ-07, 08, 09, 12, and 13 vs others.

MAB C vs A + B

Modules RQ-10, 11, 14, and 15 vs others.

Central vs peripheral modules

Modules RQ-03, 04, 08, 09, 10, 13, and 14 vs others. This null test was first proposed after the discovery of the deck encoder slipping behavior. Errors in the deck position will result in larger pointing offsets for the peripheral modules than for the central modules.

$\mathbf{I} \rightarrow \mathbf{Q}$ leakage

Division between modules with high or low $\mathbf{I} \rightarrow \mathbf{Q}$ leakage on average over all four diodes. Note that, if the null test division was simply between high and low leakage diodes, then it would be effectively the same as the Q vs U diode null test. The high leakage modules are RQ-05, 07, 08, 09, 10, 13, 14, and 15.

\mathbf{Q} vs \mathbf{U} diodes

This is not actually a division between modules. The Q and U diodes observe different Stokes parameters for each CES and are systematically different in terms of $\mathbf{I} \rightarrow \mathbf{Q}$ leakage, so this null test provides an important check.

10.2.2 Pointing Divisions

Divisions by telescope position are sensitive to errors in the pointing model, which would cause sky signals to cancel imperfectly, as well as sun, moon, or ground pickup. The pointing statistics used in these null tests are calculated at the midpoint of the azimuth scan and a time halfway through the CES. This means that many of the statistics are a bit fuzzy. For instance, the sun’s angular distance from the boresight will vary across each azimuth scan and throughout the CES, but only a single value is used to select which subset a particular CES should belong to. However, for all divisions, the ambiguous border cases represent a small amount of data in comparison to the CES which clearly belong in one subset or the other.

Scan elevation

CES divided by scan elevation, with the threshold set to the median value for each patch.

Patch rising vs setting

All four patches transit to the south, so this division is equivalent to a division by azimuth with a threshold of 180° .

Deck angle 30

CES made at deck angle of 30 degrees (the most common position during the observing season) vs others.

Deck division 90°

CES made at deck angles of 30, 75, or -150 degrees vs others.

Deck division 45°

CES made at deck angles of 30, 120, -150, or -240 degrees vs others.

Array orientation 180°

CES where the calculated parallactic angle (in equatorial coordinates) of the boresight orientation vector is in the range $[-90, 90]$ vs others.

Array orientation 90°

CES where the calculated parallactic angle of the boresight orientation vector is either in the range $[-90, 0]$ or $[90, 180]$ vs others.

Array orientation 45°

CES where the calculated parallactic angle of the boresight orientation vector is in one of the following ranges – $[-180, -135]$, $[-90, -45]$, $[0, 45]$, $[90, 135]$ – vs others.

Sun proximity

CES divided based on the distance between the sun and the scan center, with the threshold set to the median value for each patch.

Moon proximity

CES divided based on the distance between the moon and the scan center, with the threshold set to the median value for each patch.

Sidelobe 1 elevation

CES divided based on the calculated elevation of the triple reflection sidelobe, with the threshold set to the median value for each patch.

Sidelobe 1 galactic latitude

CES divided based on the absolute value of the calculated galactic latitude for the triple reflection sidelobe (i.e. is the sidelobe pointing near or far from the galactic plane), with the threshold set to the median value for each patch.

Sidelobe 2 galactic latitude

CES divided based on the absolute value of the calculated galactic latitude for the second sidelobe, with the threshold set to the median value for each patch.

Scan period 1.2 Hz commensurate

CES with scan periods of 10 or 20 seconds vs others. This null test was introduced after 1.2 Hz noise (the frequency of the refrigerator cycle) was seen in data from the W-band array. Scans with period of 10 or 20 seconds are commensurate with the fridge cycle, so there is always the possibility that noise induced by that source could add coherently.

10.2.3 Temporal Divisions

Over the course of a seven month observing season, problems occurred, equipment broke, and various small changes were made to the instrument. The temporal divisions target four of those events to test whether the observed signal was consistent from before to after.

New generator installation

CES before run 478 vs others. While the new generator worked well (most of the time), the switch could have resulted in changes to the instrument grounding or noise on the power lines.

AIB cable change

CES before run 579 vs others. The AIB cables are not extremely robust, so changing between two sets of cables, and then later changing back, could cause unnoticed problems with the modules.

Deck encoder fix

CES before run 688 vs others. Before this time, the deck encoder was not properly screwed down, resulting in jumps that increase the pointing error. After the fix, the pointing model is more reliable.

Heater installed for the AIBs

CES before run 980 vs others. The AIBs, which feature module protection circuitry, began to act irregularly in cold or snowy weather, so small heaters were installed to keep them warm and dry.

10.2.4 Performance Divisions

These nulltests divide the CES-diodes based on various measure of the data quality. All data in both the “good” and “bad” subsets has passed the data selection, but failure of the null test could indicate that the data selection criteria is too permissive.

Scan synchronous signal

CES-diodes divided based on the scan synchronous signal statistic, with the threshold set to the median value for each patch. The SSS statistic measures the value of the single Fourier mode at the scan frequency as an excess above the noise model prediction for the CES-diode.

High frequency χ^2

CES-diodes divided based on the χ^2 of the fit to the noise model, in the range 2 to 15 Hz, for the noise power spectrum of the unfiltered data. This statistic measures the level of high frequency noise in each CES-diode. The null test threshold is set to the median value for each patch.

Knee frequency

Modules divided based on the average knee frequencies for their detector diodes. The modules with high knee frequencies are RQ-01, 02, 03, 05, 06, 07, 09, 13.

CES gain

CES divided based on the average gain of all diodes, with the threshold set to the median value for each patch.

Diode gain

CES-diodes divided based on whether the gain is above or below the average gain of that diode over the season. This null test splits data up between individual CES-diodes with higher or lower gain than normal while the previous test only considers whether the average array gain is high or low.

Type-A glitching

Division between diodes that frequently exhibit Type-A glitching and those that do not. The glitching diodes are RQ-00 *U2*, RQ-02 *U2*, RQ-03 *U1*, RQ-06 *Q1*, RQ-06 *Q2*, RQ-08 *Q2*, RQ-09 *U1*, RQ-10 *U2*, RQ-12 *U2*, all diodes on RQ-14, RQ-16 *Q1*, and RQ-16 *Q2*.

Type-B glitching

CES-diodes that suffer from Type-B glitching vs those that don't. Whether or not there was Type-B glitching for a given CES-diode is determined based on whether the glitch correction algorithm induces a change in the Type-B χ^2 statistic, which is

described in §8.7.

10.2.5 Weather Divisions

Like the performance divisions, null tests based on weather divisions test check that the data selection eliminates any data which might corrupt the result.

Weather statistic R_{10}

CES divided based on the value of the R_{10} weather statistic (see §8.2), with the threshold set to the median value for each patch.

Weather statistic R_{10}^{95}

CES divided based on the value of the R_{10}^{95} weather statistic, with the threshold set to the median value for each patch.

Ambient temperature

CES divided based on the ambient temperature at the site, with the threshold set to the median value for each patch.

Humidity

CES divided based on the mean value of the humidity during the CES, with the threshold set to the median value for each patch.

Humidity change

CES divided based on the difference between the maximum and minimum humidity values recorded during the CES, with the threshold set to the median value for each patch.

10.2.6 Housekeeping Divisions

The housekeeping division null tests check for problems caused by temperature changes in the electronics enclosure. The gain model does account for the average enclosure temperature

during each CES, but not for drifts in that temperature over sub-CES timescales.

Enclosure temperature

CES divided based on the average enclosure temperature during the CES, with the threshold set to the median value for each patch.

Enclosure temperature change

CES divided based on the difference between the maximum and minimum enclosure temperature during the CES, with the threshold set to the median value for each patch.

Enclosure temperature zig-zag fluctuation

Division between CES that exhibit the “zig-zag” temperature fluctuation (a quick temperature variation with characteristic shape) and those that don’t.

10.3 Evaluating Null Suite Results

A run of the null suite yields 40 EE and BB spectra, each evaluated in 10 ℓ bins, starting with $\ell \in [25, 74]$ and going up to $\ell \in [475, 524]$. This adds up to a grand total of 800 data points for each patch, which are expected to be consistent with no sky signal. There will be outliers in any data set of that size, so it is of course unrealistic to expect that every point will lie within 1σ of zero. Instead, it is necessary to make a statistical evaluation to determine whether the null suite results in a passing grade.

A χ^2 -like statistic is calculated from the null spectra C_ℓ values and their error bars

$$\chi_{null}^2 \equiv \left(\frac{C_\ell}{\sigma_{C_\ell}} \right)^2 \quad (10.3)$$

However, the χ_{null}^2 values will not have the same distribution as a true χ^2 statistic, for two reasons. First, that would require that the C_ℓ values obey Gaussian statistics, which is a good approximation for the bins at larger values of ℓ , but not for the $\ell \in [25, 74]$ bin. The

second reason is that the forty null tests do not constitute forty independent measurements. Every null test is performed on the same data set, divided up in different ways, and some of the divisions are rather similar. An obvious example are the four temporal divisions, which all put the CES from the beginning of the season into sub-map 1 and the CES from the end of the season into sub-map 2, only disagreeing about a small amount of data in the middle. These two effects are explored further in Appendix C.

The signal plus noise Monte Carlo simulations, used to find the error estimates for the C_ℓ values, provide a natural tool to evaluate the null suite results. These simulations are analyzed in the same way and provide hundreds of examples for the \tilde{C}_ℓ values obtained for the case where the data contain no unknown contaminations.

10.4 Null Suite Results

Histograms of the χ_{null}^2 values for each patch are shown in Figures 10.1 through 10.4, with values from the EE spectra on the left and values from the BB spectra on the right. The gray shaded regions are calculated by constructing the same histogram for each Monte Carlo realization of the null suite, and then determining the 1, 2, and 3σ intervals in each histogram bin.

The most striking feature in the null suite results is seen in the histogram for the Patch 2a EE results (left-hand panel of Figure 10.1). There is a single outlier point with a χ_{null}^2 value of 33.5. That particular null spectrum point comes from the measurement of EE power in the bin $\ell \in [25, 74]$ (largest angular scales) with the data divided based on the elevation of the triple reflection sidelobe. Only one out of the 640 Monte Carlo simulations contains a single point that is as discrepant. That such a large χ_{null}^2 value shows up in this particular division of the data immediately suggests that the source of the problem is ground pickup through the sidelobe, but there are several factors that cast doubt on that theory.

The strongest argument against a sidelobe explanation is that the scanning pattern is




Figure 10.1: Histogram of χ^2_{null} values obtained from running the full null suite for Patch 2a. Points from EE spectra are included in the left-hand panel. Points from BB spectra are included on the right. The gray shaded regions designate the 1, 2, and 3σ intervals for each histogram bin, as determined from Monte Carlo realizations of the null suite.




Figure 10.2: Histogram of χ^2_{null} values obtained from running the full null suite for Patch 4a. Points from EE spectra are included in the left-hand panel. Points from BB spectra are included on the right. The gray shaded regions designate the 1, 2, and 3σ intervals for each histogram bin, as determined from Monte Carlo realizations of the null suite.




Figure 10.3: Histogram of χ^2_{null} values obtained from running the full null suite for Patch 6a. Points from EE spectra are included in the left-hand panel. Points from BB spectra are included on the right. The gray shaded regions designate the 1, 2, and 3σ intervals for each histogram bin, as determined from Monte Carlo realizations of the null suite.




Figure 10.4: Histogram of χ^2_{null} values obtained from running the full null suite for Patch 7b. Points from EE spectra are included in the left-hand panel. Points from BB spectra are included on the right. The gray shaded regions designate the 1, 2, and 3σ intervals for each histogram bin, as determined from Monte Carlo realizations of the null suite.

essentially the same for all patches, yet the other patches show no sign of trouble for this particular null spectrum point (χ_{null}^2 values of 0.61, 0.56, and 1.55, for patches 4a, 6a, and 7b, respectively). Another problem is that the ground pickup must be devised in such a way that it contributes only to the EE spectrum, and only at large angular scales.

Less extreme outliers are found among the BB power spectra of the Patch 4a null suite. One of the divisions based on the parallactic angle of the array orientation vector (array orientation 45° , described in §10.2.2) features a null spectrum point with $\chi_{null}^2 = 26.1$ for the bin $\ell \in [25, 74]$. The division based on ambient temperature fails for the bin $\ell \in [125, 174]$, with $\chi_{null}^2 = 19.0$.

The Patch 6a and 7b null suite results are generally free of individual outlier points, but closer inspection of the EE results (left-hand plots) in Figures 10.3 and 10.4 reveals a deficit of counts with $\chi_{null}^2 < 2$ and an excess in the range 2 to 5, which leads to overall statistical incompatibility with the simulations. For Patch 6a, the sum of χ_{null}^2 values for all 800 points in the null suite is 877, which exceeds the calculated total for all but 7.5% of the Monte Carlo realizations – a marginal result.

Table 10.1 summarizes the status of the null tests for each patch.

10.5 Prospects for the Analysis

There are several paths forward for the goal of resolving the failures of the null suite. A new study of the data has shown that the cut on sun position due to the second sidelobe should be applied to more of the modules, in order to truly eliminate that effect. New filtering schemes have been proposed and tested, and the option remains to increase the order of the polynomials subtracted from the data on each left or right-going scan segment.

The most promising recent discovery is that the contamination leading to the Patch 2a extreme outlier point can be seen directly in the Fourier transform of the affected null map. It shows up as a strong but very narrow signal in frequency space, perhaps only a single

Table 10.1. Null Suite Status by Patch

Patch	worst outlier				total χ^2_{null}		
	χ^2_{null} ^a	PTE ^c	EE or BB	ℓ bin	χ^2_{null} ^b	d.o.f. ^d	PTE ^c
Patch 2a	33.5	0.16%	EE	[25, 74]	738	760	71.6%
Patch 4a	26.1	1.81%	BB	[25, 74]	825	760	10.4%
Patch 6a	13.6	39.60%	EE	[125, 174]	877	800	7.5%
Patch 7b	16.9	17.80%	EE	[274, 324]	885	780	3.8%

^aThe χ^2_{null} value for the single worst outlier point in the entire null suite.

^bThe sum of χ^2_{null} values for all points in the null suite.

^cProbabilities to exceed are calculated by determining what percentage of the Monte Carlo null suite simulations meet or exceed the criteria in question.

^dThe number of degrees of freedom for the null suite is just equal to the number of data points. Here, it is different for the different patches because not all of the null tests described in §10.2 were implemented for these null test runs.

contaminated mode of the sky, which should provide strong evidence for the origin of the problem and also suggests a way that it could be fixed, by simply subtracting out that mode in the map domain.

It is also possible that the null test result could be more sensitive to the contamination than the power spectrum itself. Some support for this idea is provided by a blind comparison that was performed between the non-null power spectra obtained from the different patches and found them to be consistent (Buder, 2010).

Once the residual problems in the data are fixed, and it is finally time to calculate the observed EE and BB power spectra, the QUIET Q-band result will have sensitivity similar to what is shown in Figure 10.5. The error bars used for this rough projection are calculated from the errors for the null test based on ambient temperature. This null test is selected because it features an even division of data and no cause for a strong difference in the transfer or weighting functions between the two subsets.

In addition to QUIET, Figure 10.5 also shows data points from BICEP (Chiang et al., 2010) and QUAD (Brown et al., 2009), the two most recently released results from ground based CMB polarization experiments. As projected, the QUIET result will be competitive with BICEP, which observes on similarly large scales. The target B -mode signals from inflationary gravity waves ($r = 0.1$) and gravitational lensing, are still well below the sensitivity level of currently published results.




Figure 10.5: The projected QUIET 40 GHz measurement of the EE spectrum is shown as the blue points in the upper panel. The values for the data points are fixed to the Λ CDM spectrum, but the size of the error bars are calculated from null test result. Also included in the figure are the published results from BICEP (Chiang et al., 2010), in red, and QUAD (Brown et al., 2009), in orange. The projected upper limits to the BB spectrum from QUIET are shown as the blue points in the lower panel, along with the BICEP and QUAD results. The thick black lines show the predicted BB spectra for primordial gravity waves with $r = 0.1$ and for gravitational lensing.

APPENDIX A

MOON CALIBRATION RESULTS

This appendix tabulates the results of the moon calibration analysis described in §7.5.

- Table A.1 lists the 18 full array drift scan observations used for the analysis.
- Tables A.2 and A.3 list the detector angles obtained by averaging over the 18 drift scans, for each diode and 50 Hz state. These values match the ones shown in Figure 7.8.
- Tables A.4 and A.5 list the leakage coefficients obtained by averaging over the 18 drift scans, for each diode and 50 Hz state. These values match the ones shown in Figure 7.9.
- Tables A.6 through A.9 list relative averaged and demodulated gains obtained by comparing each diode against the detector diode of the same type from RQ-09 (i.e. all $Q1$ diode gains are specified relative to RQ-09 $Q1$). Figure 7.10 shows the 18 measurements of relative demodulated gain between RQ-05 $Q1$ and RQ-09 $Q1$ as a function of time. Here, only the mean value over the season is listed, along with the standard deviation of the values, which gives an indication of the level of gain variations.

Table A.1. Moon drift scans used for calibration.

Run id	Date	Moon phase ^a	Deck angle	Notes
213.0	2008-11-21	287	30	Part coverage for RQ-00, 01, 02, 15
213.1	2008-11-21	287	30	
219.0	2008-11-21	289	30	
367.0	2008-12-03	66	165	RQ-11 works in 50Hz state 1 only
505.0	2008-12-23	317	75	
765.0	2009-02-11	199	165	
773.0	2009-02-12	214	120	
779.0	2009-02-13	224	120	
836.1	2009-02-21	320	75	
987.0	2009-03-15	228	120	
1058.0	2009-03-23	324	165	Part coverage for RQ-12, 13
1266.0	2009-04-14	232	30	
1334.0	2009-04-21	319	75	
1522.1	2009-05-12	213	120	
1704.0	2009-06-02	113	−195	
1763.0	2009-06-06	159	−195	
1811.0	2009-06-10	207	−195	
1838.0	2009-06-12	229	−195	

^aMoon phase is specified in degrees, with 0° corresponding to a new moon and 180° corresponding to a full moon. SOURCE United States Naval Observatory, <http://www.usno.navy.mil/>

Table A.2. Detector angles for Q diodes, in degrees, as measured from moon observations.

Module	$Q1$		$Q2$	
	50Hz = 0	50Hz = 1	50Hz = 0	50Hz = 1
RQ-00	59.00 ± 0.25	59.44 ± 0.22	141.23 ± 0.24	141.74 ± 0.21
RQ-01	72.20 ± 0.44	72.65 ± 0.34	155.38 ± 0.41	155.60 ± 0.42
RQ-02	67.66 ± 0.83	67.80 ± 0.77	155.50 ± 0.81	155.58 ± 0.79
RQ-03	60.50 ± 0.29	60.98 ± 0.32	148.02 ± 0.28	148.32 ± 0.20
RQ-04	55.18 ± 0.24	55.52 ± 0.24	139.67 ± 0.24	139.97 ± 0.22
RQ-05	64.99 ± 0.23	65.25 ± 0.23	150.42 ± 0.19	150.69 ± 0.24
RQ-06	62.64 ± 0.24	62.90 ± 0.31	145.64 ± 0.26	146.08 ± 0.30
RQ-07	55.25 ± 0.35	55.76 ± 0.38	140.09 ± 0.30	140.54 ± 0.34
RQ-08	62.97 ± 0.29	63.34 ± 0.31	146.70 ± 0.22	147.13 ± 0.23
RQ-09	68.25 ± 0.22	68.59 ± 0.25	153.67 ± 0.20	153.77 ± 0.20
RQ-10	60.92 ± 0.19	61.39 ± 0.21	147.06 ± 0.17	147.53 ± 0.20
RQ-11	63.87 ± 0.30	64.18 ± 0.21	151.27 ± 0.26	151.85 ± 0.26
RQ-12	62.75 ± 0.21	63.12 ± 0.27	149.15 ± 0.19	149.74 ± 0.18
RQ-13	45.66 ± 0.22	46.02 ± 0.22	129.96 ± 0.21	130.11 ± 0.27
RQ-14	59.73 ± 0.25	60.11 ± 0.20	142.86 ± 0.25	143.24 ± 0.23
RQ-15	39.58 ± 0.53	40.14 ± 0.59	131.72 ± 0.50	132.34 ± 0.42

Table A.3. Detector angles for U diodes, in degrees, as measured from moon observations.

Module	$U1$		$U2$	
	50Hz = 0	50Hz = 1	50Hz = 0	50Hz = 1
RQ-00	107.26 ± 0.19	107.50 ± 0.20	17.61 ± 0.27	17.78 ± 0.27
RQ-01	119.66 ± 0.39	120.15 ± 0.45	34.62 ± 0.53	35.02 ± 0.53
RQ-02	117.02 ± 0.90	116.88 ± 0.78	27.35 ± 0.69	27.47 ± 0.64
RQ-03	112.47 ± 0.20	112.84 ± 0.18	22.43 ± 0.20	22.91 ± 0.15
RQ-04	104.53 ± 0.24	104.68 ± 0.17	14.26 ± 0.20	14.66 ± 0.22
RQ-05	111.34 ± 0.21	111.68 ± 0.23	24.40 ± 0.17	24.76 ± 0.22
RQ-06	110.14 ± 0.21	110.17 ± 0.24	21.62 ± 0.30	21.61 ± 0.20
RQ-07	103.54 ± 0.35	103.95 ± 0.35	14.87 ± 0.38	15.52 ± 0.35
RQ-08	110.79 ± 0.21	111.25 ± 0.20	—	—
RQ-09	118.05 ± 0.22	118.19 ± 0.24	25.40 ± 0.20	25.77 ± 0.22
RQ-10	106.88 ± 0.23	107.33 ± 0.23	17.65 ± 0.25	18.01 ± 0.17
RQ-11	112.71 ± 0.50	113.66 ± 0.22	24.78 ± 0.25	24.71 ± 0.25
RQ-12	112.47 ± 0.21	112.40 ± 0.26	23.35 ± 0.21	23.48 ± 0.22
RQ-13	94.69 ± 0.25	95.05 ± 0.21	7.34 ± 0.22	7.54 ± 0.26
RQ-14	105.75 ± 0.40	105.74 ± 0.40	12.47 ± 0.38	12.60 ± 0.38
RQ-15	108.76 ± 0.60	109.39 ± 0.49	-9.01 ± 0.31	-10.08 ± 0.25

Table A.4. Fractional $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients for Q diodes, as measured from moon observations.

Module	$Q1$		$Q2$	
	50Hz = 0	50Hz = 1	50Hz = 0	50Hz = 1
RQ-00	-0.01064 ± 0.00016	0.00988 ± 0.00017	0.00881 ± 0.00016	-0.00951 ± 0.00016
RQ-01	0.00113 ± 0.00024	-0.00164 ± 0.00018	0.00013 ± 0.00018	-0.00047 ± 0.00017
RQ-02	-0.00926 ± 0.00018	0.00903 ± 0.00018	0.00679 ± 0.00017	-0.00676 ± 0.00017
RQ-03	-0.00327 ± 0.00017	0.00243 ± 0.00017	0.00242 ± 0.00016	-0.00368 ± 0.00016
RQ-04	-0.01005 ± 0.00017	0.00905 ± 0.00016	0.00659 ± 0.00016	-0.00766 ± 0.00016
RQ-05	-0.01914 ± 0.00018	0.01893 ± 0.00018	0.01842 ± 0.00017	-0.01835 ± 0.00017
RQ-06	-0.00833 ± 0.00019	0.00889 ± 0.00018	0.00843 ± 0.00017	-0.00840 ± 0.00017
RQ-07	-0.01509 ± 0.00017	0.01476 ± 0.00017	0.01351 ± 0.00015	-0.01419 ± 0.00016
RQ-08	-0.00747 ± 0.00019	0.00778 ± 0.00019	0.01001 ± 0.00017	-0.01004 ± 0.00018
RQ-09	-0.00929 ± 0.00020	0.00934 ± 0.00021	0.00892 ± 0.00018	-0.01010 ± 0.00020
RQ-10	-0.01022 ± 0.00018	0.01009 ± 0.00018	0.01159 ± 0.00018	-0.01182 ± 0.00018
RQ-11	-0.00493 ± 0.00020	0.00462 ± 0.00020	0.00341 ± 0.00020	-0.00472 ± 0.00019
RQ-12	-0.00659 ± 0.00019	0.00702 ± 0.00019	0.00575 ± 0.00017	-0.00619 ± 0.00018
RQ-13	-0.01389 ± 0.00016	0.01356 ± 0.00017	0.01127 ± 0.00017	-0.01193 ± 0.00017
RQ-14	-0.01872 ± 0.00018	0.01816 ± 0.00019	0.01882 ± 0.00017	-0.01947 ± 0.00018
RQ-15	-0.00426 ± 0.00024	0.00656 ± 0.00019	0.01095 ± 0.00021	-0.00782 ± 0.00025

Table A.5. Fractional $\mathbf{I} \rightarrow \mathbf{Q}$ leakage coefficients for U diodes, as measured from moon observations.

Module	$U1$		$U2$	
	50Hz = 0	50Hz = 1	50Hz = 0	50Hz = 1
RQ-00	0.00026 ± 0.00024	-0.00112 ± 0.00024	-0.00191 ± 0.00024	0.00101 ± 0.00025
RQ-01	0.00538 ± 0.00021	-0.00561 ± 0.00021	-0.00804 ± 0.00022	0.00760 ± 0.00022
RQ-02	0.00269 ± 0.00023	-0.00225 ± 0.00021	-0.00260 ± 0.00022	0.00320 ± 0.00023
RQ-03	-0.00074 ± 0.00023	-0.00024 ± 0.00023	-0.00131 ± 0.00024	0.00044 ± 0.00024
RQ-04	-0.00043 ± 0.00026	-0.00017 ± 0.00024	-0.00072 ± 0.00024	0.00002 ± 0.00026
RQ-05	0.00367 ± 0.00025	-0.00423 ± 0.00026	-0.00706 ± 0.00025	0.00705 ± 0.00025
RQ-06	0.00131 ± 0.00023	-0.00116 ± 0.00023	-0.00197 ± 0.00023	0.00228 ± 0.00024
RQ-07	-0.00190 ± 0.00026	0.00125 ± 0.00025	-0.00006 ± 0.00025	-0.00035 ± 0.00026
RQ-08	0.00103 ± 0.00025	-0.00092 ± 0.00025	—	—
RQ-09	0.00413 ± 0.00021	-0.00481 ± 0.00020	-0.00415 ± 0.00024	0.00477 ± 0.00024
RQ-10	0.00134 ± 0.00023	-0.00161 ± 0.00022	-0.00159 ± 0.00023	0.00169 ± 0.00024
RQ-11	0.00182 ± 0.00024	-0.00265 ± 0.00021	-0.00315 ± 0.00023	0.00218 ± 0.00024
RQ-12	0.00121 ± 0.00025	-0.00120 ± 0.00025	-0.00171 ± 0.00026	0.00214 ± 0.00025
RQ-13	-0.00417 ± 0.00027	0.00379 ± 0.00027	0.00204 ± 0.00027	-0.00218 ± 0.00027
RQ-14	-0.00021 ± 0.00023	-0.00094 ± 0.00022	-0.00019 ± 0.00025	-0.00001 ± 0.00025
RQ-15	0.00350 ± 0.00025	-0.00177 ± 0.00028	0.00718 ± 0.00028	-0.00718 ± 0.00025

Table A.6. Average (total power) gains of Q diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.

Module	$Q1$				$Q2$			
	50Hz = 0		50Hz = 1		50Hz = 0		50Hz = 1	
	G_{av}	σ^a	G_{av}	σ^a	G_{av}	σ^a	G_{av}	σ^a
RQ-00	0.671	0.033	0.668	0.034	0.723	0.038	0.725	0.040
RQ-01	0.772	0.051	0.746	0.041	0.970	0.062	0.975	0.059
RQ-02	0.621	0.030	0.612	0.040	0.648	0.034	0.650	0.030
RQ-03	0.705	0.027	0.702	0.030	0.581	0.029	0.584	0.038
RQ-04	0.760	0.052	0.739	0.047	0.778	0.048	0.770	0.046
RQ-05	0.842	0.033	0.835	0.034	0.743	0.027	0.740	0.029
RQ-06	0.607	0.027	0.604	0.032	0.590	0.023	0.592	0.028
RQ-07	0.799	0.022	0.790	0.024	0.901	0.026	0.903	0.033
RQ-08	0.897	0.030	0.895	0.026	0.662	0.028	0.665	0.024
RQ-09	1.000	—	1.000	—	1.000	—	1.000	—
RQ-10	0.758	0.027	0.747	0.017	0.735	0.039	0.742	0.038
RQ-11	0.877	0.022	0.863	0.029	0.747	0.022	0.758	0.024
RQ-12	0.764	0.016	0.751	0.034	0.709	0.038	0.710	0.037
RQ-13	0.802	0.024	0.797	0.026	0.846	0.021	0.845	0.019
RQ-14	0.947	0.031	0.941	0.021	0.954	0.025	0.957	0.029
RQ-15	0.944	0.026	0.943	0.026	0.800	0.017	0.797	0.019

^aStandard deviation of the relative gain values obtained from each of the 18 moon observations. The standard deviation is given, rather than an error on the mean, because the relative gains are presumed to vary during the observing season.

Table A.7. Average (total power) gains of U diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.

Module	$Q1$				$Q2$			
	50Hz = 0		50Hz = 1		50Hz = 0		50Hz = 1	
	G_{av}	σ^a	G_{av}	σ^a	G_{av}	σ^a	G_{av}	σ^a
RQ-00	0.863	0.055	0.867	0.033	0.892	0.074	0.902	0.088
RQ-01	1.055	0.074	1.040	0.063	0.713	0.066	0.729	0.072
RQ-02	0.764	0.058	0.757	0.048	0.581	0.055	0.595	0.059
RQ-03	0.784	0.031	0.790	0.035	0.922	0.066	0.938	0.085
RQ-04	0.864	0.054	0.843	0.053	0.819	0.090	0.831	0.092
RQ-05	0.559	0.025	0.549	0.026	1.011	0.075	1.030	0.080
RQ-06	0.981	0.072	0.974	0.059	1.095	0.096	1.124	0.107
RQ-07	1.020	0.034	1.013	0.042	0.823	0.054	0.843	0.063
RQ-08	1.003	0.048	1.006	0.035	—	—	—	—
RQ-09	1.000	—	1.000	—	1.000	—	1.000	—
RQ-10	1.070	0.030	1.060	0.032	1.115	0.055	1.127	0.066
RQ-11	0.840	0.031	0.843	0.042	0.701	0.045	0.696	0.041
RQ-12	1.072	0.040	1.070	0.047	1.038	0.072	1.063	0.080
RQ-13	1.040	0.038	1.027	0.041	1.006	0.064	1.022	0.092
RQ-14	0.959	0.042	0.961	0.032	0.737	0.043	0.754	0.053
RQ-15	0.925	0.034	0.920	0.027	1.336	0.092	1.347	0.108

^aStandard deviation of the relative gain values obtained from each of the 18 moon observations. The standard deviation is given, rather than an error on the mean, because the relative gains are presumed to vary during the observing season.

Table A.8. Demodulated (polarized) gains of Q diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.

Module	$Q1$				$Q2$			
	50Hz = 0		50Hz = 1		50Hz = 0		50Hz = 1	
	G_{de}	σ^a	G_{de}	σ^a	G_{de}	σ^a	G_{de}	σ^a
RQ-00	0.594	0.059	0.591	0.038	0.625	0.068	0.629	0.046
RQ-01	0.629	0.073	0.599	0.049	0.758	0.070	0.763	0.061
RQ-02	0.589	0.089	0.571	0.063	0.536	0.064	0.539	0.044
RQ-03	0.580	0.050	0.588	0.051	0.490	0.034	0.498	0.048
RQ-04	0.643	0.076	0.615	0.071	0.615	0.071	0.592	0.051
RQ-05	0.838	0.080	0.832	0.079	0.731	0.063	0.723	0.055
RQ-06	0.518	0.059	0.505	0.064	0.483	0.040	0.469	0.045
RQ-07	0.686	0.050	0.678	0.064	0.755	0.064	0.741	0.074
RQ-08	0.786	0.063	0.784	0.049	0.579	0.037	0.580	0.024
RQ-09	1.000	—	1.000	—	1.000	—	1.000	—
RQ-10	0.760	0.061	0.752	0.037	0.774	0.069	0.773	0.058
RQ-11	0.827	0.073	0.785	0.077	0.657	0.080	0.671	0.039
RQ-12	0.628	0.036	0.593	0.071	0.583	0.058	0.569	0.043
RQ-13	0.776	0.060	0.750	0.079	0.764	0.065	0.769	0.047
RQ-14	0.810	0.067	0.807	0.048	0.818	0.068	0.810	0.043
RQ-15	0.610	0.057	0.598	0.067	0.584	0.052	0.585	0.053

^aStandard deviation of the relative gain values obtained from each of the 18 moon observations. The standard deviation is given, rather than an error on the mean, because the relative gains are presumed to vary during the observing season.

Table A.9. Demodulated (polarized) gains of U diodes, relative to the corresponding diode from module RQ-09, as measured from moon observations.

Module	$Q1$				$Q2$			
	50Hz = 0		50Hz = 1		50Hz = 0		50Hz = 1	
	G_{de}	σ^a	G_{de}	σ^a	G_{de}	σ^a	G_{de}	σ^a
RQ-00	0.750	0.093	0.771	0.046	0.789	0.075	0.826	0.111
RQ-01	0.881	0.067	0.855	0.070	0.571	0.066	0.592	0.081
RQ-02	0.634	0.071	0.628	0.080	0.481	0.074	0.507	0.076
RQ-03	0.626	0.030	0.638	0.050	0.801	0.060	0.829	0.112
RQ-04	0.691	0.074	0.659	0.062	0.632	0.076	0.652	0.098
RQ-05	0.559	0.049	0.552	0.039	0.966	0.091	1.012	0.122
RQ-06	0.768	0.083	0.751	0.070	0.826	0.096	0.860	0.112
RQ-07	0.781	0.065	0.761	0.070	0.643	0.051	0.665	0.091
RQ-08	0.857	0.072	0.861	0.052	—	—	—	—
RQ-09	1.000	—	1.000	—	1.000	—	1.000	—
RQ-10	0.996	0.068	0.964	0.043	0.988	0.075	1.000	0.108
RQ-11	0.670	0.084	0.691	0.061	0.568	0.050	0.570	0.042
RQ-12	0.829	0.065	0.815	0.069	0.751	0.091	0.786	0.096
RQ-13	0.867	0.079	0.841	0.075	0.850	0.051	0.877	0.114
RQ-14	0.762	0.077	0.767	0.046	0.711	0.052	0.755	0.117
RQ-15	0.758	0.070	0.753	0.055	0.944	0.095	0.969	0.137

^aStandard deviation of the relative gain values obtained from each of the 18 moon observations. The standard deviation is given, rather than an error on the mean, because the relative gains are presumed to vary during the observing season.

APPENDIX B

QUIET DATA SELECTION USING MYSQL

Analysis of QUIET data starts with the selection of CES-diodes to combine into maps. The criteria used for this selection are discussed in §8. The actual list of data to use is obtained by querying a MySQL database containing the pre-computed values of the selection statistics. This method proved extremely useful not only for the full blown data analysis, which was run a relatively small number of times, but also for studies of data selection, which could be performed quickly using the available statistics rather than by laboriously processing the full time ordered data.

This appendix includes description of the database tables and the values contained in each. The standard data selection query, containing all of the criteria from §8, is found in §B.2. Division of the selected data for null tests is accomplished by simply appending one additional criteria or its inverse to the query. The null test selections, with thresholds that are defined separately for each patch in some cases, are found in §B.3.

B.1 Overview of Database Tables

The database tables are organized by their purpose or by the source of the data composing the table. For example, the `q19_scan` table contains information about the telescope scan motion for each CES, including its duration, target pointing, and deck orientation, while the `q19_timestream` table contains the results obtained from processing the time ordered data, such as the noise model parameters and goodness-of-fit.

The primary keys used to organize each table are the run id and sub-id, which uniquely specify each CES, and the module number and diode type, which specify the data channel. Some tables use only two of these four primary keys, such as the `q19_scan` table, which is indexed only by run id and sub-id. Common primary keys allow one to easily combine

multiple tables in a query by using the natural join predicate.

The database tables used for data selection and null test divisions are:

q19_scan

This table contains information about the telescope scan for each CES. Fields include the start time for the scan, target object (i.e. which CMB patch is being observed), scan duration, and the mean *az*, *el*, and *dk* coordinates. From these basic results, additional statistics are derived, including the position of the sun and moon (both in equatorial coordinates and relative to the telescope boresight), the hour angle of the patch (used to determine if it is rising or setting), approximate horizon and galactic coordinates calculated for the two telescope sidelobes, and the parallactic angles of the array orientation and the scan direction.

q19_typeb

This table contains statistics describing the effects of Type-B glitching for each CES-diode. Fields include the Type-B χ^2 statistic, described in §8.7, for both uncorrected and corrected data, the noise level in the data before and after applying the correction, and the closest approach for the average data to a Type-B glitch location.

q19_timestream

This table contains the results of timestream processing and the noise model fit. Fields include the gain coefficient calculated for each CES-diode, noise model coefficients, and the frequency domain χ^2 values used for data selection (see §8.3. Also included in the table are measurements of the largest time domain glitches at four time scales, the statistical significance of the linear slopes fit at each scan segment, and two statistics used to measure excess power at or near the scan frequency.

q19_weather

This table contains the weather statistics, R_{10} and R_{10}^{95} , which are calculated from

averaged data. Additional fields include humidity measurements for each CES, as reported by the telescope weather station.

`q19_diodes`

This table contains the mean and standard deviation of the gain model and noise model parameters across the entire season for each detector diode.

`q19_housekeeping`

This table contains the temperature mean and range in each CES for the modules (cryogenic temperature), the electronics enclosure, and ambient (i.e. outside the electronic enclosure).

`q19_ces_usable`

This table contains a flag to indicate whether or not a CES should be included in the analysis. It is mostly redundant with the selection for the baseline data set.

`q19_atm`

This table contains detailed atmospheric information obtained from the APEX weather station, including temperature, precipital water vapor, wind speed, and direction.

B.2 Standard Data Selection Query

The MySQL query corresponding to the data selection described in §8 is shown below. This set of criteria is known as “cuts8”, to distinguish it from earlier versions. Explanatory comments have been added, **in red**.

```
# CES-diodes specified by run id, run subid, module, and diode
```

```
SELECT DISTINCT q19_scan.run_id, q19_scan.run_subid,  
                q19_typeb.module, q19_typeb.diode
```

```
# Database tables used for data selection
```

```

FROM q19_scan NATURAL JOIN q19_typeb NATURAL JOIN
    q19_timestream NATURAL JOIN q19_weather NATURAL JOIN
    q19_diodes NATURAL JOIN q19_housekeeping NATURAL JOIN
    q19_ces_usable, q19_atm

# Selection criteria

WHERE

# Join q19_atm table
(q19_atm.run_id = q19_scan.run_id AND
    q19_atm.run_subid = q19_scan.run_subid)

# Select CMB patch using one of the following lines
AND (q19_scan.object = 'patch2a')
# AND (q19_scan.object = 'patch4a')
# AND (q19_scan.object = 'patch6a')
# AND (q19_scan.object = 'patch7b')

# Baseline data set
AND (q19_ces_usable.usable)
AND (q19_scan.duration >= 1000)
AND (q19_typeb.module < 16)
AND NOT (q19_typeb.module = 4 AND q19_typeb.diode = 'Q1')
AND NOT (q19_typeb.module = 8 AND q19_typeb.diode = 'U2')
AND NOT (q19_scan.run_id BETWEEN 329 AND 341)
AND NOT (q19_scan.run_id = 398)
AND NOT (q19_scan.run_id = 529 AND q19_scan.run_subid = 1)
AND NOT (q19_scan.run_id BETWEEN 563 AND 593)
AND NOT (q19_scan.run_id = 654 AND q19_scan.run_subid = 0)
AND NOT (q19_scan.run_id = 759 AND q19_scan.run_subid = 0)

```

```

AND NOT (q19_scan.run_id = 937)
AND NOT (q19_scan.run_id = 951 AND q19_scan.run_subid = 0)
AND NOT (q19_scan.run_id BETWEEN 953 AND 955)
AND NOT (q19_scan.run_id = 980 AND q19_scan.run_subid = 0)
AND NOT (q19_scan.run_id = 1413 AND q19_scan.run_subid = 0)

# Weather cut
AND (q19_weather.tp_rms_10sec <= 0.15)
AND (q19_weather.tp_rms_10sec_95 <= 0.08)

# Cut on noise model goodness-of-fit
AND (q19_timestream.filtered_chisquare_10mHz_200mHz <= 2.5)
AND (q19_timestream.filtered_chisquare_200mHz_2Hz <= 2.5)
AND (q19_timestream.filtered_chisquare_2Hz_4_6Hz <= 2.5)

# Cut on knee frequency, uses a pre-defined function containing
# individual diode cut thresholds
AND FKNEE_CUT_15_5(q19_timestream.module, q19_timestream.diode,
                    q19_timestream.knee_frequency)

# Sun position cut
AND (NOT (q19_scan.sun_phi BETWEEN 120 AND 180 AND
          q19_scan.sun_theta BETWEEN 44 AND 60))
AND (NOT (((q19_scan.sun_phi BETWEEN -75 AND -15 AND
            q19_scan.sun_theta BETWEEN 44 AND 60) OR
            (q19_scan.sun_phi BETWEEN -90 AND 0 AND
             q19_scan.sun_theta BETWEEN 60 AND 65)) AND
        (q19_typeb.module BETWEEN 0 AND 2)))

# Glitch cut
AND (q19_timestream.glitch_1sample <= 6)

```

```

AND (q19_timestream.glitch_100msec <= 6)
AND (q19_timestream.glitch_1sec <= 6)
AND (q19_timestream.glitch_10sec <= 6)
# Type-B cut
AND (q19_typeb.chisquare_demodav_corrected <= 10);

```

B.3 Nulltest Selection Criteria

The CES-diodes for the two subsets of each null test are selected by applying the usual data selection query, shown above, plus on additional criteria, which defines the first subset. The second subset is obtained by applying the inverse of that criteria.

Some of the null tests have a clear criteria for splitting the data, such as a run number which is used to divide the observing season into halves. Others are based on statistics measured for the CES-diodes, and require a defined threshold. For these cases, the median value of the statistic is chosen separately for each patch, so that the null test is as close as possible to an even split of the data, which leads to the smallest error bars for the null power spectrum.

The MySQL selection criteria for each of the null tests described in §10 are included here, with the patch dependent thresholds, where applicable.

B.3.1 Module Divisions

MAB A vs B + C

```
(q19_timestream.module BETWEEN 0 AND 6)
```

MAB B vs A + C

```
((q19_timestream.module BETWEEN 7 AND 9) OR
(q19_timestream.module BETWEEN 12 AND 13) OR
```



```
(q19_timestream.module = 16))
```

MAB C vs A + B

```
((q19_timestream.module BETWEEN 10 AND 11) OR  
(q19_timestream.module BETWEEN 14 AND 15) OR  
(q19_timestream.module BETWEEN 17 AND 18))
```

Central vs peripheral modules

```
((q19_timestream.module BETWEEN 3 AND 4) OR  
(q19_timestream.module BETWEEN 8 AND 10) OR  
(q19_timestream.module BETWEEN 13 AND 14))
```

I → Q leakage

```
((q19_timestream.module = 5) OR (q19_timestream.module = 7) OR  
(q19_timestream.module = 8) OR (q19_timestream.module = 9) OR  
(q19_timestream.module = 10) OR (q19_timestream.module = 13) OR  
(q19_timestream.module = 14) OR (q19_timestream.module = 15))
```

Q vs U diodes

```
((q19_timestream.diode = 'Q1') OR (q19_timestream.diode = 'Q2'))
```

B.3.2 Pointing Divisions

Scan elevation

```
(q19_scan.elevation > 62.6143) # Patch 2a  
(q19_scan.elevation > 63.1633) # Patch 4a  
(q19_scan.elevation > 58.8852) # Patch 6a  
(q19_scan.elevation > 59.4408) # Patch 7b
```

Patch rising vs setting

```
(q19_scan.hour_angle < 0)
```

Deck angle 30

```
(ROUND(q19_scan.deck_angle, 0) = 30)
```

Deck division 90°

```
((ROUND(q19_scan.deck_angle, 0) = 30) OR  
(ROUND(q19_scan.deck_angle, 0) = 75) OR  
(ROUND(q19_scan.deck_angle, 0) = -150))
```

Deck division 45°

```
((ROUND(q19_scan.deck_angle, 0) = 30) OR  
(ROUND(q19_scan.deck_angle, 0) = 120) OR  
(ROUND(q19_scan.deck_angle, 0) = -150) OR  
(ROUND(q19_scan.deck_angle, 0) = -240))
```

Array orientation 180°

```
(q19_scan.array_orientation_equatorial >= -90. AND  
q19_scan.array_orientation_equatorial < 90.)
```

Array orientation 90°

```
((q19_scan.array_orientation_equatorial >= -90. AND  
q19_scan.array_orientation_equatorial < 0.) OR  
(q19_scan.array_orientation_equatorial >= 90. AND  
q19_scan.array_orientation_equatorial < 180.))
```

Array orientation 45°

```
((q19_scan.array_orientation_equatorial >= -180. AND  
  q19_scan.array_orientation_equatorial < -135.) OR  
(q19_scan.array_orientation_equatorial >= -90. AND  
  q19_scan.array_orientation_equatorial < -45.) OR  
(q19_scan.array_orientation_equatorial >= 0. AND  
  q19_scan.array_orientation_equatorial < 45.) OR  
(q19_scan.array_orientation_equatorial >= 90. AND  
  q19_scan.array_orientation_equatorial < 135.))
```

Sun proximity

```
(q19_scan.sun_theta < 125.231) # Patch 2a  
(q19_scan.sun_theta < 81.6184) # Patch 4a  
(q19_scan.sun_theta < 79.0609) # Patch 6a  
(q19_scan.sun_theta < 76.1315) # Patch 7b
```

Moon proximity

```
(q19_scan.moon_theta < 89.5357) # Patch 2a  
(q19_scan.moon_theta < 94.2536) # Patch 4a  
(q19_scan.moon_theta < 93.5700) # Patch 6a  
(q19_scan.moon_theta < 88.3389) # Patch 7b
```

Sidelobe 1 elevation

```
(q19_scan.sidelobe_el > 29.1416) # Patch 2a  
(q19_scan.sidelobe_el > 25.8815) # Patch 4a  
(q19_scan.sidelobe_el > 25.9742) # Patch 6a  
(q19_scan.sidelobe_el > 30.0007) # Patch 7b
```

Sidelobe 1 galactic latitude

```
(ABS(q19_scan.sidelobe_lat) <= 24.827) # Patch 2a  
(ABS(q19_scan.sidelobe_lat) <= 16.614) # Patch 4a  
(ABS(q19_scan.sidelobe_lat) <= 34.720) # Patch 6a  
(ABS(q19_scan.sidelobe_lat) <= 43.717) # Patch 7b
```

Sidelobe 2 galactic latitude

```
(ABS(q19_scan.sidelobe2_lat <= 34.724) # Patch 2a  
(ABS(q19_scan.sidelobe2_lat <= 32.587) # Patch 4a  
(ABS(q19_scan.sidelobe2_lat <= 28.296) # Patch 6a  
(ABS(q19_scan.sidelobe2_lat <= 14.631) # Patch 7b
```

Scan period 1.2 Hz commensurate

```
((q19_scan.scan_period = 10) OR (q19_scan.scan_period = 20) OR  
(q19_scan.scan_period = 30))
```

B.3.3 Temporal Divisions

New generator installation

```
(q19_scan.run_id < 478)
```

AIB cable change

```
(q19_scan.run_id < 579)
```

Deck encoder fix

```
(q19_scan.run_id < 688)
```

Heater installed for the AIBs

```
(q19_scan.run_id < 980)
```

B.3.4 Performance Divisions

Scan synchronous signal

```
(q19_timestream.scan_synchronous_signal > 0.695684) # Patch 2a  
(q19_timestream.scan_synchronous_signal > 0.687522) # Patch 4a  
(q19_timestream.scan_synchronous_signal > 0.766133) # Patch 6a  
(q19_timestream.scan_synchronous_signal > 0.830864) # Patch 7b
```

High frequency χ^2

```
(q19_timestream.chisquare_2Hz_15Hz > 3.08600) # Patch 2a  
(q19_timestream.chisquare_2Hz_15Hz > 2.60758) # Patch 4a  
(q19_timestream.chisquare_2Hz_15Hz > 2.82151) # Patch 6a  
(q19_timestream.chisquare_2Hz_15Hz > 2.57911) # Patch 7b
```

Knee frequency

```
((q19_typeb.module = 1) OR (q19_typeb.module = 2) OR  
(q19_typeb.module = 3) OR (q19_typeb.module = 5) OR  
(q19_typeb.module = 6) OR (q19_typeb.module = 7) OR  
(q19_typeb.module = 9) OR (q19_typeb.module = 13))
```

CES gain

Patch 2a

```
((SELECT AVG(fiducial_gain) FROM q19_timestream  
AS timestream_gain_temp  
WHERE q19_timestream.run_id = timestream_gain_temp.run_id AND  
q19_timestream.run_subid = timestream_gain_temp.run_subid)  
> 2.32975)
```

```
# Patch 4a threshold = 2.33826
# Patch 6a threshold = 2.34350
# Patch 7b threshold = 2.34282
```

Diode gain

```
(q19_timestream.fiducial_gain >
 (SELECT gain_mean FROM q19_diodes AS diodes_gain_temp
  WHERE diodes_gain_temp.module = q19_timestream.module AND
        diodes_gain_temp.diode = q19_timestream.diode))
```

Type-A glitching

```
((q19_timestream.module = 0 AND q19_timestream.diode = 'U2') OR
 (q19_timestream.module = 2 AND q19_timestream.diode = 'U2') OR
 (q19_timestream.module = 3 AND q19_timestream.diode = 'U1') OR
 (q19_timestream.module = 6 AND q19_timestream.diode = 'Q1') OR
 (q19_timestream.module = 6 AND q19_timestream.diode = 'Q2') OR
 (q19_timestream.module = 8 AND q19_timestream.diode = 'Q2') OR
 (q19_timestream.module = 9 AND q19_timestream.diode = 'U1') OR
 (q19_timestream.module = 10 AND q19_timestream.diode = 'Q2') OR
 (q19_timestream.module = 12 AND q19_timestream.diode = 'U2') OR
 (q19_timestream.module = 14) OR
 (q19_timestream.module = 16 AND q19_timestream.diode = 'Q1') OR
 (q19_timestream.module = 16 AND q19_timestream.diode = 'Q2'))
```

Type-B glitching

```
(ABS(q19_typeb.chisquare_demodav_uncorrected -
      q19_typeb.chisquare_demodav_corrected) > 0.1)
```

B.3.5 Weather Divisions

Weather statistic R_{10}

```
(q19_weather.tp_rms_10sec > 0.055304) # Patch 2a  
(q19_weather.tp_rms_10sec > 0.057137) # Patch 4a  
(q19_weather.tp_rms_10sec > 0.058370) # Patch 6a  
(q19_weather.tp_rms_10sec > 0.055825) # Patch 7b
```

Weather statistic R_{10}^{95}

```
(q19_weather.tp_rms_10sec_95 > 0.048427) # Patch 2a  
(q19_weather.tp_rms_10sec_95 > 0.048947) # Patch 4a  
(q19_weather.tp_rms_10sec_95 > 0.049342) # Patch 6a  
(q19_weather.tp_rms_10sec_95 > 0.048535) # Patch 7b
```

Ambient temperature

```
(q19_atm.air_temp_med > -3.4800) # Patch 2a  
(q19_atm.air_temp_med > -1.2910) # Patch 4a  
(q19_atm.air_temp_med > -1.3290) # Patch 6a  
(q19_atm.air_temp_med > -4.0305) # Patch 7b
```

Humidity

```
(q19_weather.humidity_max > 36.60) # Patch 2a  
(q19_weather.humidity_max > 32.08) # Patch 4a  
(q19_weather.humidity_max > 27.06) # Patch 6a  
(q19_weather.humidity_max > 25.18) # Patch 7b
```

Humidity change

```
(q19_weather.humidity_diff > 6.78) # Patch 2a  
(q19_weather.humidity_diff > 6.00) # Patch 4a  
(q19_weather.humidity_diff > 6.44) # Patch 6a  
(q19_weather.humidity_diff > 3.76) # Patch 7b
```

B.3.6 Housekeeping Divisions

Enclosure temperature

```
(q19_housekeeping.mmic1_temp_p3t6_mean > 24.5529) # Patch 2a  
(q19_housekeeping.mmic1_temp_p3t6_mean > 24.6367) # Patch 4a  
(q19_housekeeping.mmic1_temp_p3t6_mean > 24.7654) # Patch 6a  
(q19_housekeeping.mmic1_temp_p3t6_mean > 24.5279) # Patch 7b
```

Enclosure temperature change

```
(q19_housekeeping.mmic1_temp_p3t6_range > 0.268059) # Patch 2a  
(q19_housekeeping.mmic1_temp_p3t6_range > 0.362078) # Patch 4a  
(q19_housekeeping.mmic1_temp_p3t6_range > 0.402503) # Patch 6a  
(q19_housekeeping.mmic1_temp_p3t6_range > 0.284849) # Patch 7b
```

Enclosure temperature zig-zag fluctuation

```
(q19_housekeeping.enclosure_temp_zigzag_fom > 0)
```


APPENDIX C

NULL SUITE TOY MODEL

While the Monte Carlo simulations provide the best reference for the expected distribution of χ_{null}^2 values, it is a useful exercise to study the reasons why that distribution differs from a true χ^2 distribution. The null suite toy model, derived here, does a fairly good job at capturing the statistical properties of the null suite, by accounting for the non-Gaussianity of the C_ℓ values and also the correlations between different null tests.

We start by considering a single mode in a map derived from a single CES. The amplitude of this mode for the i^{th} CES is described by a Gaussian random number, x_i , with mean μ and standard deviation σ . For a single CES, the instrumental noise, σ , is much larger than the underlying CMB signal, μ . Because of this fact, and also because we will be summing the CES into null maps, it is fine to assume $\mu = 0$.

Each CES that passes cuts is evaluated according to some null criteria, labeled by α , and added into one of the two sub-maps, $X^{\alpha 1}$ and $X^{\alpha 2}$, which will eventually be differenced to make the null map. The amplitude of a mode in each of the sub-maps is calculated as a simple average, as shown in equation (C.1). In this model, we will assume that every CES is weighted equally. That isn't exactly right, since the duration and noise level vary between CES, but it shouldn't make much difference.

$$X^{\alpha 1} = \frac{1}{N^{\alpha 1}} \sum_{i=1}^N \epsilon_i^{\alpha 1} x_i \quad , \quad X^{\alpha 2} = \frac{1}{N^{\alpha 2}} \sum_{i=1}^N \epsilon_i^{\alpha 2} x_i \quad (C.1)$$

The values $\epsilon_i^{\alpha 1,2}$ are equal to 1 if the CES is selected for this sub-map or 0 otherwise. The total number of CES is given by N , and $N^{\alpha 1,2}$ give the number of CES in each sub-map. Ideally, the CES would be evenly divided for each null test, so $N^{\alpha 1} = N^{\alpha 2} = N/2$, but this is not always the case.

The null map, X^α , is given by the difference of the two sub-maps. The variance of this

map is given by

$$X^\alpha = X^{\alpha 1} - X^{\alpha 2} \quad (\text{C.2})$$

$$\langle (X^\alpha)^2 \rangle = \left(\frac{N}{N^{\alpha 1} N^{\alpha 2}} \right) \sigma^2 \quad (\text{C.3})$$

which is smallest for the case where the data are divided equally.

The correlation between two null tests, α and β , can be calculated as follows.

$$C(\alpha, \beta) = \frac{\langle X^\alpha X^\beta \rangle}{\sqrt{\langle (X^\alpha)^2 \rangle \langle (X^\beta)^2 \rangle}} \quad (\text{C.4})$$

$$= \frac{\sqrt{N^{\alpha 1} N^{\alpha 2} N^{\beta 1} N^{\beta 2}}}{N \sigma^2} \left[\langle X^{\alpha 1} X^{\beta 1} \rangle - \langle X^{\alpha 1} X^{\beta 2} \rangle \right. \\ \left. - \langle X^{\alpha 2} X^{\beta 1} \rangle + \langle X^{\alpha 2} X^{\beta 2} \rangle \right] \quad (\text{C.5})$$

The four sub-map correlation terms are all calculated similarly, making use of the identity $\langle x_i x_j \rangle = \delta_{ij} \sigma^2$, which is true if the noise is independent for each CES.

$$\langle X^{\alpha 1} X^{\beta 1} \rangle = \frac{1}{N^{\alpha 1} N^{\beta 1}} \sum_{i=1}^N \sum_{j=1}^N \epsilon_i^{\alpha 1} \epsilon_j^{\beta 1} \langle x_i x_j \rangle \quad (\text{C.6})$$

$$= \frac{\sigma^2}{N^{\alpha 1} N^{\beta 1}} \sum_{i=1}^N \epsilon_i^{\alpha 1} \epsilon_i^{\beta 1} \quad (\text{C.7})$$

$$= \left(\frac{N^{\alpha 1 \beta 1}}{N^{\alpha 1} N^{\beta 1}} \right) \sigma^2 \quad (\text{C.8})$$

The summation in equation (C.7) is equal to the number of CES that are included in sub-map 1 for null test α and null test β . The full expression for the null test correlation contains

these terms for all four combinations of sub-maps.

$$C(\alpha, \beta) = \frac{\sqrt{N^{\alpha 1} N^{\alpha 2} N^{\beta 1} N^{\beta 2}}}{N} \left[\frac{N^{\alpha 1 \beta 1}}{N^{\alpha 1} N^{\beta 1}} - \frac{N^{\alpha 1 \beta 2}}{N^{\alpha 1} N^{\beta 2}} - \frac{N^{\alpha 2 \beta 1}}{N^{\alpha 2} N^{\beta 1}} + \frac{N^{\alpha 2 \beta 2}}{N^{\alpha 2} N^{\beta 2}} \right] \quad (\text{C.9})$$

It can be verified that the statistic in equation (C.9) is equal to +1 for the case where $\epsilon_i^{\alpha 1} = \epsilon_i^{\beta 1}$ and is equal to -1 for the case where $\epsilon_i^{\alpha 1} = \epsilon_i^{\beta 2}$.

Once we have the correlation matrix for amplitudes of single modes, we can generate random values for the mode amplitudes in each null test with the proper correlations by calculating

$$X^\alpha = P_\beta^\alpha \left(D_\gamma^\beta \right)^{1/2} Y^\gamma \quad (\text{C.10})$$

where Y^γ is a vector of independent random values, D_γ^β is the diagonal matrix made up of the eigenvalues of the null test correlation matrix, and P_β^α is a matrix whose columns are the eigenvectors of the correlation matrix.

Each ℓ bin contains many modes, with the approximate number given by

$$\text{number of modes} \simeq f_{\text{sky}} \sum_{\ell_{\min}}^{\ell_{\max}} (2\ell + 1) \quad (\text{C.11})$$

For QUIET, the patch area is roughly 225 square degrees, so $f_{\text{sky}} \simeq 0.00545$. The number of modes for each bin ranges from 27 for $\ell \in [25, 74]$ to 272 for $\ell \in [475, 524]$.

Now we can write the full procedure for a null suite simulation. For each ℓ bin in EE and BB , calculate n mode amplitudes, X^α , for each null test using equation (C.10), where n is the number of modes in that bin. The power in each mode is given by the square of the mode amplitude. To find the null C_ℓ value, we sum up all the modes in an ℓ bin and subtract the expected power, as shown in equation (C.12). The error bar for each point is

calculated according to equation (C.13).

$$C_\ell^\alpha = \left(\sum_{n_\ell} (X^\alpha)^2 \right) - n_\ell \quad (\text{C.12})$$

$$\sigma_{C_\ell} = \sqrt{2n_\ell} \quad (\text{C.13})$$

Figure C.1 compares the cumulative distribution function of the χ_{null}^2 values for the Monte Carlo results from a null suite run on patch 2a to the same function for the toy model. The actual distribution functions match so well that it would be hard to tell them apart, so they have all been subtracted from the CDF for an ideal χ^2 distribution.




Figure C.1: All three curves are the difference between the cumulative distribution function for a data sample and the cumulative distribution function for an ideal χ^2 distribution. For the black and blue curves, the CDF is calculated for the χ_{null}^2 values from a 832 Monte Carlo simulations of a null suite run on Patch 2a. For the red curve, the CDF is calculated by the toy model, with the equivalent of 10,000 null suite simulations.

The agreement between the toy model (red curve) and the null suite Monte Carlo (black

and blue curves) is not perfect, but they do exhibit the same general shape. Compared to a true χ^2 distribution, both samples have a deficit of points with $\chi^2 < 4$ and an excess in the tail of the distribution. It might be possible to improve the agreement by making a better estimate of the number of independent modes for each ℓ bin or by calculating the null test correlation matrix in a way that takes into account the different weights assigned to various CES-diodes, but agreement at this level is good evidence that the understanding of the χ^2_{null} statistic distribution is basically correct.

REFERENCES

- Aumont, J. et al. 2009, ArXiv e-prints
- Bannister, P., Nixon, G., & Staggs, S. 2006, QUIET internal memo
- Barkats, D. et al. 2005a, *The Astrophysical Journal Letters*, 619, L127
- . 2005b, *The Astrophysical Journal Supplement Series*, 159, 1
- Baumann, D. et al. 2009, in *American Institute of Physics Conference Series*, Vol. 1141, American Institute of Physics Conference Series, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith, 10–120
- Bennett, C. L. et al. 2003, *ApJS*, 148, 97
- Bischoff, C. et al. 2008, *The Astrophysical Journal*, 684, 771
- Bischoff, C., Kapner, D., & Sugarbaker, A. 2006, QUIET internal memo
- Bischoff, C., & Kusaka, A. 2007, QUIET internal memo
- Bischoff, C., & Tajima, O. 2008, QUIET internal memo
- Bond, J. R., Jaffe, A. H., & Knox, L. 1998, *Phys. Rev. D*, 57, 2117
- Bornemann, J., & Labay, V. 1995, *Microwave Theory and Techniques, IEEE Transactions on*, 43, 1782
- Brown, M. L. et al. 2009, *ApJ*, 705, 978
- Buder, I. 2010, QUIET internal memo
- Chiang, H. C. et al. 2010, *ApJ*, 711, 1123
- Chinone, Y., & Tajima, O. 2010, private correspondence
- Chu, T. S., Wilson, R. W., England, R. W., Gray, D. A., & Legg, W. E. 1978, *AT& T Technical Journal*, 57, 1257
- Dumoulin, R. 2007, QUIET internal memo
- Farese, P. C. et al. 2003, *New Astronomy Reviews*, 47, 1033 , proceedings of the Workshop on The Cosmic Microwave Background Radiation and its Polarization
- Gaier, T., Lawrence, C. R., Seiffert, M. D., Wells, M. M., Kangaslahti, P., & Dawson, D. 2003, *New Astronomy Reviews*, 47, 1167 , proceedings of the Workshop on The Cosmic Microwave Background Radiation and its Polarization

- Gary, B., Stacey, J., & Drake, F. D. 1965, *ApJS*, 12, 239
- Gold, B. et al. 2010, ArXiv e-prints
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Gundersen, J., & Wollack, E. 2009, *Journal of Physics: Conference Series*, 155, 012005
- Guth, A. H. 1981, *Phys. Rev. D*, 23, 347
- Hanany, S., & Marrone, D. P. 2002, *Appl. Opt.*, 41, 4666
- Hedman, M. M., Barkats, D., Gundersen, J. O., McMahon, J. J., Staggs, S. T., & Winstein, B. 2002, *The Astrophysical Journal Letters*, 573, L73
- Hedman, M. M., Barkats, D., Gundersen, J. O., Staggs, S. T., & Winstein, B. 2001, *The Astrophysical Journal Letters*, 548, L111
- Hinshaw, G. et al. 2007, *ApJS*, 170, 288
- . 2003, *ApJS*, 148, 135
- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Holler, C. 2007, QUIET internal memo
- Houde, M., Akeson, R. L., Carlstrom, J. E., Lamb, J. W., Schleuning, D. A., & Woody, D. P. 2001, *PASP*, 113, 622
- Hu, W., & White, M. 1997, *Phys. Rev. D*, 56, 596
- Hu, W., & White, M. 1997, *New Astronomy*, 2, 323
- Huff, K. 2008, QUIET Celestial Gain Calibrations, University of Chicago Senior Thesis, undergraduate Thesis
- Jackson, J. D. 1998, *Classical Electrodynamics*, 3rd edn. (Wiley)
- Jarosik, N. et al. 2010a, ArXiv e-prints
- . 2010b, ArXiv e-prints
- Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, *Phys. Rev. D*, 55, 7368
- Kapner, D. 2006, QUIET internal memo
- Keating, B. G., O'Dell, C. W., Gundersen, J. O., Piccirillo, L., Stebor, N. C., & Timbie, P. T. 2003, *The Astrophysical Journal Supplement Series*, 144, 1

- Kogut, A. et al. 2007, ApJ, 665, 355
- Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzzapfel, W. L. 2002, Nature, 420, 772
- Kraus, J. D. 1986, Radio Astronomy, 2nd edn. (Cygnus-Quasar Books)
- Kusaka, A. 2010, QUIET internal memo
- Kusaka, A., & Bischoff, C. 2007, QUIET internal memo
- Larson, D. et al. 2010, ArXiv e-prints
- Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, ApJ, 624, 10
- Lewis, A., Challinor, A., & Lasenby, A. 2000, Astrophys. J., 538, 473
- Linsky, J. L. 1973, ApJS, 25, 163
- Mather, J. C., Fixsen, D. J., Shafer, R. A., Mosier, C., & Wilkinson, D. T. 1999, ApJ, 512, 511
- Monsalve, R. 2010, QUIET internal memo
- Montroy, T. E. et al. 2006, ApJ, 647, 813
- Newburgh, L., Miller, A., Grainger, W., & Williamson, R. 2006a, QUIET internal memo
- Newburgh, L., Miller, A., & Williamson, R. 2006b, QUIET internal memo
- Nixon, G. 2008, QUIET internal memo
- Padin, S. et al. 2002, PASP, 114, 83
- Paine, S. 2004, SMA Technical Memo #152
- Penzias, A. A., & Wilson, R. W. 1965, ApJ, 142, 419
- Perlmutter, S. et al. 1999, The Astrophysical Journal, 517, 565
- Pozar, D. M. 2005, Microwave Engineering, 3rd edn. (John Wiley & Sons)
- Richardson, J. 2007, QUIET internal memo
- Riess, A. G. et al. 1998, The Astronomical Journal, 116, 1009
- Sanders, J. R. 2010, QUIET internal memo
- Seiffert, M., Gaier, T., & Smith, S. 2006, QUIET internal memo

- Sievers, J. L. et al. 2005, ArXiv Astrophysics e-prints
- Smith, K. 2009, QUIET internal memo
- Smith, K. M. 2006, Phys. Rev. D, 74, 083002
- Smith, K. M. et al. 2008, ArXiv e-prints
- Smith, K. M., & Zaldarriaga, M. 2007, Phys. Rev. D, 76, 043001
- Smoot, G. F. et al. 1992, ApJ, 396, L1
- Tajima, O. 2010a, QUIET internal memo
- . 2010b, QUIET internal memo
- Thompson, K. L. 2009, QUIET internal memo
- Tran, H. T. 2003, New Astronomy Reviews, 47, 1091 , proceedings of the Workshop on The Cosmic Microwave Background Radiation and its Polarization
- Wandelt, B. D., Hivon, E., & Górski, K. M. 2001, Phys. Rev. D, 64, 083003
- Weiland, J. L. et al. 2010, ArXiv e-prints
- Weinreb, S., Lai, R., Erickson, N., Gaier, T., & Wielgus, J. 1999, in Microwave Symposium Digest, 1999 IEEE MTT-S International, Vol. 1, 101 – 104
- Wright, E. L. et al. 2009, ApJS, 180, 283
- Wu, J. H. P. et al. 2007, ApJ, 665, 55
- Zaldarriaga, M., & Seljak, U. c. v. 1997, Phys. Rev. D, 55, 1830