MEASUREMENT OF THE TEMPERATURE AND POLARIZATION ANISOTROPIES IN THE COSMIC MICROWAVE BACKGROUND WITH QUAD

Michael B. Zemcov

A THESIS SUBMITTED TO CARDIFF UNIVERSITY FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

September 15, 2006

UMI Number: U584847

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U584847 Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author. Microform Edition © ProQuest LLC. All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code.



ProQuest LLC 789 East Eisenhower Parkway P.O. Box 1346 Ann Arbor, MI 48106-1346

Author: Michael B. Zemcov Title: Measurement of the Temperature and Polarization Anisotropies in the Cosmic Microwave Background with QUaD Date of submission: September 15, 2006

Permission is granted to Cardiff University to circulate and to have copied for non-commercial purposes, at its discretion, the above title upon the request of individuals or institutions. The author reservs other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

Copyright © 2006 by Michael B. Zemcov



To my parents

vi

ABSTRACT

This thesis paper describes QUaD, a millimetric polarimeter designed to precisely measure the temperature and polarization anisotropies in the cosmic microwave background (CMB) radiation. QUaD was built and integrated between 2000 and 2004, and commissioned at the south pole in the austral summer of 2004/2005. First light was achieved in February 2005, and QUaD acquired a full austral winter of data during that year. Modifications to the instrument were performed in the austral summer of 2005/2006, and QUaD is currently taking a second season of data. A third season is planned.

QUaD is composed of a cryogenically cooled receiver mounted on a Cassegrain telescope using a 2.6 m primary mirror. The detectors are polarization sensitive Neutron Transmutation Doped germanium bolometers which are coupled to the optics via single-moded corrugated feed horns. These detectors are split between two frequency bands, 100 and 150 GHz, and the optics allow angular resolutions of 6.0 and 4.0 arcminutes in these two bands. The performance of the instrument has been characterized using extensive laboratory data, a battery of tests performed during commissioning, and specially designed measurements during the observing season. QUaD uses constant elevation azimuth scans with 15 degree lengths, allowing reconstruction of all 6 CMB power spectra in the range $50 < \ell < 2000$. Customized software and analysis methods are required to transform data time series from the telescope to refined data products.

QUaD data have provided the most accurate measurements of the E-mode CMB power spectrum to date, as well as measurements of T-mode power competitive with the best high angular resolution experiments. These measurements provide a wealth of information, and both confirm our model for the early Universe as well as constrain cosmological parameters. Furthermore, QUaD measures the formation of structure after recombination, and can constrain the physics of inflation.

ACKNOWLEDGEMENTS

I suspect that debts of gratitude are probably the most difficult to discharge. In the case of this thesis, a substantial part of that difficulty is due to the sheer number of people involved in inserting the necessary knowledge into my head, and subsequently dragging it back out into the light. However, through sheer, blind stubbornness, I am going to include a ham–fisted attempt at thanking you, even if that is an impossible feat to begin with. Just so you're aware, I'm probably going to mess it all up, so prepare yourself. Here's the summary statement for those with better things to do: my sincerest, deepest thanks to you all.

There, you can stop reading now.

It says here in my build-your-own-thesis guidebook that the first people to acknowledge for their efforts are my ever patient, rarely present supervisors, Walter Gear and Phil Mauskopf. So, my most heartfelt thanks for the vigilant and understanding way you have shepherded me through this moderately dull, occasionally painful, but always educational process. I certainly couldn't have done it without you.

My deepest thanks are also due to Douglas Scott, Colin Borys, and especially Mark Halpern, for their tireless attempts to try to get some sense out of me. I am amazed, flattered, frustrated, and incredibly thankful that they even deign to look in my direction four years after any formal responsibility for any part of my well-being passed from their hands. Would that I had a tenth of your patience, or a hundredth of your skill.

An enumeration of people who have furthered my education with no formal reward throughout my graduate work would reach a very high number indeed. So, many, many thanks to you all. However, in the spirit of personalizing things, here's a go at listing some of the ones that I can recall right this minute. Many thanks also to Sarah Church, who has taught me a great deal over the few years I've been working with her, and has been very kind to me while I've thrashed around, trying to find my place. Thanks also to Clem Pryke, whose lessons about care and vigilance in the great scientific endeavour have not fallen on deaf ears, or so one would hope. Also, my profound thanks are due to Brad Johnson, for being a fantastic person to bounce ideas off, and a friend and comrade in exile. Thanks to John Kovac for his constant demand that I do things better, more carefully, and with elegance: his strident exhortations have certainly made me a better scientist and stand as a shining beacon dedicated to that old truism: pride cometh before a fall. Thanks also to Erik Leitch for profound advice that I'm never sure I quite understand, but I'm sure - were I able to - would be pure gold. Ken Ganga also deserves a huge thanks for always being there to talk things through, no matter what the topic; from career advice to optimal mapping algorithms to the exact pronunciation of 'un beignet framboise', nothing is out of bounds or a waste of time.

Many thanks to the QUaD crew, in all of their forms. But especially to Jamie, for putting up with me against all odds or sense (by the way, I got a cow bell the other day dude, she's a real hummer), Mel for being such a good friend and a hell of a lot tougher than she looks (shut up, yes you are), Angiola for the same but with an macchiato flavour, Ben for his patient teaching of too many lessons to list and the heads up to procure a goat ahead of time (I owe you one for that!), Ed for lending me his house (sucka!), Gary, Michael, Patricia, and all the rest of you. Even when I wanted to strangle you to death with my bare hands, it's been a ball. And a super south pole thanks to Robert Schwartz and Allan Day for running the damn telescope so well. And, of course, all the Ritter Sport.

Thanks to all my friends both near and far (in no particular order): James, Douglas, Philippe, Simon, Mitri, Alison, Cara, Ryan, Rhys, Kio, Luke, Alex, Geoff, Ty, and many more besides. I'll end it there, but if any of you ever read this and notice I didn't mention you, don't take it personally, I'm pressed for space and time here. So let's do this: my biggest thanks go to (insert name here): we've come a long way, baby! I feel I should mention separately the people I spent the most time with in Cardiff: Bruce, Gustav and Robbie, who put up with me in ways nobody should have to. Frankly, I'm amazed we're still on speaking terms.

Thanks also to my extended family around the world, the people who *have* to put up with me, no matter what I do or say. I won't try to list you all, but just know you all mean the world to me. At some level, this is the most difficult thanks, since it's impossible to encapsulate just how much these people mean to me, and how much we've been through together. How do you do a proper job of thanking somebody who's been a part of your life, well, forever?

I must separate out my uncle Bill, whose aid and understanding during my time in Britain is something I will not forget. Thank you for spawning a love of Kentish pubs and an appreciation for my ancestors that will last a lifetime.

To my sister Clare, thanks for sharing so much of my life with me; it's been, well, it's been growing up together. And perhaps the biggest debt is to my parents, to whom this work is dedicated: I just don't even know where I would start, or where starting would finish. So: thanks.

And to Andrea, the person who is my home. Over the past few years I've had to divide my life into two parts for the sake of sanity: the kind of things in this book, and everything else. Thank you for giving, being, sharing the everything else with me.

CONTENTS

A	Abstract vi		
A	cknow	viedgements	iii
Li	st of 🛛	Tables xi	i ii
Li	st of l	Tigures X	iv
1	Intr	oduction	1
	1.1	Cosmology and the Cosmic Microwave Background Radiation1.1.1History1.1.2The Modern CMB Radiation Model1.1.3Primary Temperature Anisotropies1.1.4Secondary Temperature Anisotropies1.1.4Polarization of the Cosmic Microwave Background1.2.1Primary Polarized Anisotropies1.2.2Secondary Polarized Anisotropies1.2.3The Experimental State of the Art	2 2 4 7 11 5 16 23 25
	1.3	QUaD Science	27
2	QUa	D Instrument Design	!9
	2.1	Mount and Telescope	29
	2.2	Cryogenics	52 20
	2.3		59 • •
	2.4	Detectors	14
		2.4.1 A Bolometer Model	+U 10
	2.5	Electronics	51 52 53
3	Rece	iver Characterization	59
	3.1	Bolometer & Electronics Characterization	59
		3.1.1 Load Curves	50
		3.1.2 Microphonics	53
		3.1.3 Radio Frequency Interference	55
	3.2	Optical Characterization	56
		3.2.1 Band Determination	56

		3.2.2 Optical Efficiency	70
		3.2.3 Responsivity	71
		3.2.4 Impulse Response & Time Constants	73
	3.3	Polarization Properties	74
		3.3.1 Measurements & Results	78
4	Insti	rument Performance	85
	4.1	Pointing	86
		4.1.1 Radio Pointing Sources	86
		4.1.2 Cross Observations	88
	4.2	Naive Raster Maps	93
	4.3	Feed Offsets	100
	4.4	Beams	105
		4.4.1 Beam Ellipticity & Temperature Variation	108
		4.4.2 Beam Differences	115
	4.5	Sidelobes	117
	4.6	Sensitivity & Noise Performance	122
		4.6.1 Optical Loading	122
		4.6.2 NET, NEQ and Atmospheric Behaviour	123
	4.7	Signal Filtering Rectification	125
		4.7.1 Time Constants	127
		4.7.2 Filter Deconvolution & Decimation	128
5	Cali	bration	131
	5.1	Atmospheric Optical Depth	131
	5.2	Elevation Dips	137
		5.2.1 The El Dip – τ Relation	145
	5.3	Instrumental Calibration Source	146
	5.4	Row Calibrations	157
	5.5	Absolute Calibration	165
(Mar	Malin - Derver Sneetman	
U	Ana	p Making, rower spectrum	169
	61	Survey Description	169
	6.2	Man Making	172
	0.2	6.2.1 Temperature Maps	173
		6.2.2 Polarization Maps	179
	6.3	Power Spectra Estimation & Cosmological Parameters	182
	6.4	Other Astronomical Results	188
		6.4.1 The Moon	190
		 6.4.1 The Moon	190 192
		6.4.1 The Moon	190 192 192
		6.4.1 The Moon	190 192 192 195

A	QUa	D at MAPO	199
	A.1	Commissioning	199
	A.2	Refocus	200
	A.3	Work During the 2005/2006 Summer Season	201

LIST OF TABLES

2.1	Readout Channel Breakdown
3.1	QUaD Band Definitions
3.2	QUaD Polarization Parameters
4.1	QUaD's UC HII Region Pointing Sources
4.2	Radio Pointing Summary for 2005 Observing Season
4.3	RCW 38 Mapping Run Dates & Parameters
4.4	Beam Statistics for 2005 Observing Season
4.5	Beam Model Coefficients for the 2005 Observing Season
5.1	Absolute Calibration Factors for the 2005 Observing Season
6.1	CMB Survey Observation Block Summary
6.2	CMB Survey 24-h Duty Cycle Summary
6.3	Preliminary Cosmological Parameters for the 2005 Season Data

LIST OF FIGURES

1.1	A Model CMB Temperature Spectrum
1.2	Model CMB TT , TE , EE and BB Spectra $\ldots \ldots \ldots$
1.3	Summary of Previous Measurements of C_{ℓ}^{EE}
2.1	QUaD Cross Sectional View
2.2	QUaD Panorama
2.3	The QUaD Cryostat
2.4	The Science Core
2.5	Schematic of a Sorption Refrigerator
2.6	Photograph of the QUaD Refrigerator
2.7	Typical QUaD Refrigerator Cycle 40
2.8	QUaD's Optical Chain
2.9	QUaD's Optical Filtering
2.10	Photograph of the QUaD Focal Plane
2.11	Horn Assemblies
2.12	A Composite Bolometer Model
2.13	Bolometers in QUaD
2.14	Overview of the QUaD Electronics Chain
2.15	The Back Side of the QUaD Focal Plane
3.1	Load Curves
3.2	R versus T for QUaD Bolometers $\ldots \ldots \ldots$
3.3	Microphonic Test Results
3.4	QUaD Bands for Individual Horns
3.5	Average QUaD Bands with Atmospheric Transmission Spectrum 69
3.6	Optical Efficiency Histograms
3.7	DC Responsivity
3.8	Typical Impulse Response
3.9	Polarized Source Near Field Beam Maps
3.10	Sinusoidal Modulation of QUaD's Near Field Polarized Beams 81
3.11	Polarization Angles α from Polarized Source Measurements
3.12	Polarization Efficiencies ϵ from Polarized Source Measurements 83
4.1	Dust Maps of QUaD's UC HII Pointing Sources
4.2	A Single Pointing Cross Observation
4.3	Cross Observation Atmospheric Removal
4.4	QUaD Pointing Performance for Winter 2005

4.5	RCW 38 Raster Map with Median Filtering	97
4.6	Raster Time Series Filter Function	99
4.7	RCW 38 Raster Map with Advanced Filtering	100
4.8	Predicted Focal Plane Projection onto the Sky	101
4.9	Raw Array Offset Data for the 2005 Season	103
4.10	Array Offsets for the 2005 Observing Season	104
4.11	Array Offset Radial Structure	106
4.12	Example Beam Maps, Pre- and Post- 2005 Refocus	109
4.13	Beam Width Time Variability	110
4.14	Calculated Beam Widths and Primary–Secondary Mirror Distance	111
4.15	Anti-Correlation of Beam Width and External Temperature	112
4.16	Typical Beam Performance for Entire 2005 Observing Season	114
4.17	Example Beam Difference Maps	116
4.18	Ground Scans with Sidelobe Pick–Up	119
4.19	2005 Season Sidelobe Maps	120
4.20	Fully Baffled Sidelobe Maps	121
4.21	Optical Loading Time Series for the 2005 Observing Season	124
4.22	Power Spectrum of Noise Data	126
4.23	Histogram of Bolometer Time Constants, 2005 Season	128
4.24	Decimation	130
5 1	Ortical Darths for the 2005 Observing Sessor	124
5.1	Completions for the 2005 Observing Season	134
5.2	τ Correlations for the 2005 Observing Season	120
5.5 5.4	τ Scaling Comparison Over Entire 2005 Season	120
5.4	τ Cumulative Statistics for the 2005 Observing Season \ldots	139
5.5 5.6	Elevation Dip Massurements for the 2005 Observing Sesson for Horn 150.2	140
5.0 5 7	Elevation Dip Measurements for the 2005 Observing Season for Horn 100.2	142
5.1	Elevation Dip Measurements for the 2005 Observing Season for Horn 100-2 El Dip Goins and Sky Dip σ for 2005 Observation Season	145
J.0	Er Dip Gains and Sky Dip 7 for 2005 Observation Season $\ldots \ldots \ldots \ldots$	140
5.9	Contract Contract $\gamma_{150 \text{ GHz}}$	14/
5.10	Sinusoidal Models for Calibration Source Data	140
5.12	Calibration Source Measurements for the 2005 Observing Season for Horn	147
J.12	150-2	151
5 13	Calibration Source Measurements for the 2005 Observing Season for Horn	151
5.15	100-2	152
5 14	Intrinsic Gain Measurements for the 2005 Observing Season for Horn 150-2	152
5 1 5	Intrinsic Gain Measurements for the 2005 Observing Season for Horn 100-2	155
5 16	Calibration Source Temperature for the 2005 CMB Observing Season	157
5 17	Row Calibration Measurements for the 2005 Observing Season for Horn	101
5.17	150-2	159
5.18	Row Calibration Measurements for the 2005 Observing Season for Horn	
2.20	100-2	160
5.19	Row Calibration Flux Conversion Factors for the 2005 Observing Season	
	for Horn 150-2	163

Row Calibration Flux Conversion Factors for the 2005 Observing Season
for Horn 100-2
2005 Season Temperature Anisotropy Maps
Point Sources in the CMB Survey Field
Detail of Radio Point Sources
2005 Season Q and U Polarization Anisotropy Maps $\ldots \ldots \ldots$
Preliminary 2005 Season CMB Power Spectra
All CMB <i>EE</i> Power Spectra Data to Date
The Polarized Moon at 100 and 150 GHz
Carina Nebula at 100 and 150 GHz
Super Nova Remnant G 260.4 at 100 and 150 GHz
The Galactic Centre Region at 100 and 150 GHz
Radio Galaxy Centarus A at 100 and 150 GHz
Refocusing the Telescope, July 10 th , 2005
Secondary Mirror Modifications for 2006 Season
Sidelobe Maps, 2005 Season Baffle
Sidelobe Maps, Blackened Flat Baffle
Sidelobe Maps, Full Top Hat Baffle
Polarization Measurement Apparatus

1 INTRODUCTION

Physical cosmology is the branch of physics and astronomy which deals with the origins of the Universe and its nature on the largest scales. A number of separate lines of inquiry contribute to our understanding of cosmology, from studies of large scale structure in the X-ray, optical, infra-red (IR) and radio wavelengths, to measurements of the cosmic microwave background (CMB) radiation, to more esoteric sources like neutrino telescopes and (potentially) gravitational radiation detectors. These separate probes of the physical nature of our Universe have converged to a very strange model; for example, less than 5 % of the mass-energy of the Universe is in the form of types of matter we have measured in the laboratory. Some 25 % of the mass-energy is in the form of 'dark matter', a proposed entity which acts almost exclusively via gravitation, and for which a number of candidate particles exist, none of which have been directly detected. The largest part of the Universe's density is in the form of 'dark energy', a negative pressure which been a serious component of the standard cosmological model for less than ten years. Because of all of these unknown parameters, physical cosmology is currently an extremely active and productive part of fundamental science. Moreover, measurement of the CMB has become one of the most exciting arenas of experimental research within physical cosmology itself.

This chapter deals with the production, measurement, and scientific information accessible using the CMB radiation and its polarization. Firstly, a broad overview of the CMB radiation and its measurement is presented. The physics which produce the observed fluctuations in the CMB are then discussed. This is followed by an introduction to the physics of the polarization of the CMB radiation, including a discussion of current results in the field. Finally, the scientific motivation of the QUaD experiment is discussed.

1.1 COSMOLOGY AND THE COSMIC MICROWAVE BACK-GROUND RADIATION

Studies of the CMB are one of the most valuable tools in the cosmologist's arsenal. This is because this radiation is comprised of the oldest photons which can be observed. As such, measurements of the CMB provide unique insight into the conditions in the Universe at the time of their emission, some 13.7 Gyrs in the past. In fact, the CMB radiation can provide information about the evolution of the equation of state of the Universe before and after recombination, the evolution of matter and metric perturbations, and the characteristics of the photon–baryon fluid present in the early Universe. The advent of precision measurements of anisotropies in the CMB has allowed the construction of a consistent model for the evolution of the Universe from seconds after the big bang to the beginning of non-linear collapse of structure in the Universe.

For frequencies between about ten and several hundred GHz, the CMB is the dominant form of extragalactic electromagnetic radiation. However, instrumentation and observation in this frequency range is technically challenging, and has been one of the last parts of the electromagnetic spectrum to open to astrophysicists. The level of sensitivity required to measure the anisotropies in the CMB is truly impressive, and although great strides have been made in the last 40 years, a great deal of information remains to be gathered from studies of the ancient light from the beginning of the Universe.

1.1.1 HISTORY

The history of modern observational cosmology really begins with the work of Edwin Hubble in the $1920s^1$. Using careful measurements of Cepheid variable stars in local galaxies and combining them with redshift data from the work of Slipher, Hubble was able to construct a relation between the distance d and recessional velocity of the galaxies v according to:

$$v = H_0 d \tag{1.1}$$

where H_0 is a constant of proportionality called Hubble's constant. This came as a great vindication to Lemaître, who had originated the idea of a Big Bang (under a different name) a few years previously based on his studies in general relativity and the metastable nature of a static Universe. The essential logic behind the Big Bang theory is simple; if the Universe

¹This is in contrast to modern theoretical cosmology, whose roots lie firmly in Einstein's general relativity, and the theoretical explorations of such luminaries as De Sitter, Schwarzschild and Eddington. Their work first introduced such concepts as using Λ as a modification to the flat space metric to explain an apparently expanding Universe.

has always been expanding, there must have been a time in the past when all of the space and matter in it was condensed to a single $point^2$. These ideas are the direct ancestors of the cosmological model we find ourselves with today.

Based on arguments involving the consequences of the Hubble expansion on nucleosynthesis in the early Universe, Alpher & Herman first predicted the existence of a thermal background radiation with a temperature of about 5 K in two papers in 1948 & 1949 (although the theory is most concisely presented in Alpher & Herman 1950). This, and work by Gamow, Dicke, and others, set the theoretical stage for an isotropic Planckian background radiation in the Universe, peaking somewhere in the far infra-red (FIR). In 1965, using a 4.08 GHz system, Penzias and Wilson first detected a diffuse radiation with a thermal temperature 3.5 ± 1.0 K (Penzias & Wilson, 1965). This was correctly identified with the radiation predicted by Alpher and Herman by Dicke, Peebles, Roll and Wilkinson in a companion paper to the Penzias and Wilson article (Dicke et al., 1965). This single observation began what has become over 40 years of CMB observation, and ushered in a new understanding of the nature of the Universe.

The assignment of the extraneous antenna temperature Penzias and Wilson measured with their system was not ubiquitously thought to be caused by the CMB; it took until the 1970s before it became widely accepted that this radiation was, in fact, a relic of the big bang. This agreement primarily came about because of measurements at a variety of frequencies showing that the CMB had a blackbody spectrum with $T \approx 3$ K. At the same time as these measurements were taking place, a number of workers realized that the early Universe would have required inhomogeneities at the 10^{-4} - 10^{-5} level to explain the formation of structure in the later Universe (Harrison 1970, Peebles & Yu 1970, Zeldovich 1972). The goal of CMB radiometry thus became to measure this anisotropy in the CMB.

After 20 years of null results, COBE provided the first measurement of the CMB temperature anisotropy in 1992 (Smoot et al., 1992). The Differential Radiometer Experiment on COBE had fairly poor resolution, but was sensitive enough to measure anisotropies of $\Delta T = 30 \pm 5 \,\mu\text{K}$ at 10 degree angular resolution. This initial measurement spurred a series of ground and balloon based experiments, culminating in TOCO's tentative detection of the first peak of the CMB (Miller et al., 1999), and resolved maps of the structure in the CMB anisotropy around its peak angular scale by BOOMERanG (Mauskopf et al., 2000) and MAXIMA (Hanany et al., 2000).

These successes were followed by detections of the second and third acoustic peaks, and evidence for Silk damping by a number of experiments including BOOMERanG (Ruhl et al., 2003), the VSA (Grainge et al., 2003), CBI (Readhead et al., 2004), and ACBAR (Kuo et al., 2004). These successes culminated in the first WMAP data release, which

²Broadly, this is the content referred to by the popular term 'Big Bang'.

measured the CMB power spectrum to unprecedented precision based on a year's observation of the whole sky (Hinshaw et al., 2003). Together, these experiments have turned cosmology based on the CMB from a speculative venture to a precision science.

The current frontier of CMB science and instrumentation is the measurement of the polarization of the background microwave light. Because Thompson scattering (the physical mechanism that produces the anisotropies in the CMB radiation) polarizes light, it was known from fairly early in the study of the CMB that the CMB should be linearly polarized at approximately 10 % (Bond & Efstathiou, 1984). However, as the CMB temperature anisotropies are at the 10^{-5} level, polarization detections need sensitivity to *polarized* CMB anisotropies of at least 10^{-6} . Furthermore, to fully exploit the potential scientific information available from the polarized CMB signal, one needs sensitivities perhaps a factor of a million better than this limit (see Section 1.2). The detection of a polarized signal in the power spectrum of the CMB was first achieved by the DASI experiment in 2002 (Kovac et al., 2002). In the few years following this, a number of experiments have detected polarized CMB anisotropies, although precise measurement has required the development of instruments like QUaD.

1.1.2 THE MODERN CMB RADIATION MODEL

The modern understanding of the observed properties of the CMB radiation is rooted in the study of a perturbed matter-energy fluid acting under the effect of gravity as described by Einstein's general theory of relativity³. Here we quickly present the foundations of the model and define ubiquitous terms used later.

Observations of CMB, as well as large scale structure, tell us that the part of the Universe we can observe is, to a very good approximation, homogeneous and isotropic on the largest scales. The gravitational metric describing this type of space-time is called the Friedmann Robertson Walker (FRW) metric, and is given by

$$ds^{2} = -dt^{2} + R^{2}(t) \left[\frac{dr^{2}}{1 - \kappa r^{2}} + r^{2}d\theta^{2} + r^{2}\sin^{2}\theta d\phi^{2} \right]$$
(1.2)

where the metric is written using standard spherically symmetric co-moving coordinates and proper time t, R(t) is the cosmological scale factor, κ is the curvature of the spacetime⁴, and we have set the speed of light c = 1. Furthermore, let the current value of R be

³Fairly complete discussions of the theory behind our current understanding of the CMB can be found in, for example, Hu & Sugiyama (1995). The general relativity theory required to follow the arguments presented here can be found in a number of places, for example Weinberg (1972) or Peebles (1993).

⁴Which, with the proper choice of the coordinate r, need only ever take on the values $k = \{-1, 0, 1\}$.

 R_0 ; since the wavelength of a photon scales as the expansion parameter,

$$\frac{R_0}{R(t)} = \frac{\lambda_{\text{obs}}}{\lambda_{\text{emit}}}.$$
(1.3)

Using the result

$$\frac{\lambda_{\rm obs}}{\lambda_{\rm emit}} = 1 + z \tag{1.4}$$

yields the scale factor related to redshift⁵ z via

$$\frac{R(t)}{R_0} = a(t) = \frac{1}{z(t) + 1}.$$
(1.5)

Here, a(t) is a rescaled version of the cosmological scale factor such that a = 1 today. Equation 1.5 allows one to measure the relative size of the Universe using a single observable, z. This can be converted to a comoving distance (that is, something which can be measured in metres) by using

$$D = \frac{c}{H_0} \int_0^z \frac{dz'}{\sqrt{\Omega_M (1 + z')^3 + \Omega_\kappa (1 + z')^2 + \Omega_\Lambda}}.$$
 (1.6)

Here, $\Omega_{\{M,\kappa,\Lambda\}}$ are the relative densities of matter, curvature and cosmological constant, compared to the critical density $\rho_{\rm crit} = 3H_0^2/8\pi G$, respectively. If the Universe is homogeneous, isotropic and matter dominated, these parameters satisfy

$$\Omega_{\kappa} = \Omega_{\rm M} + \Omega_{\Lambda} - 1 \tag{1.7}$$

where $\rho_{\rm crit}$ is scaled to 1.0. There is further good evidence that the Universe is flat (k = 0) meaning that $\sum_i \Omega_i = 1$, and that $\Omega_{\Lambda} > \Omega_{\rm M}$, as discussed in Section 1.1.3.

Using the result in Equation 1.5 combined with the thermodynamics of a photon fluid, the average temperature of a bath of photons with temperature T_0 today at a given z is simply

$$T = T_0/a(t) = T_0(1+z).$$
(1.8)

This is the mathematical expression governing the intuitive thermodynamic assertion that if the Universe has always been expanding, and there has always been a bath of photons in it, then the average temperature of the Universe was higher in the past. It follows from this that at some point in the past, the average temperature was higher than the binding energy of hydrogen, $kT \gtrsim 13.6 \,\mathrm{eV^6}$. Before this, the average path length of a photon was

⁵The redshift of light is due to the fact that the Universe is expanding: as the Universe expands, a photon's wavelength must increase proportionally. Redshift is thus a proxy for the distance to an object.

⁶As usual, reality is a bit more complex than this simple model; because only the highest energy photons

low enough that the photons and baryons in the Universe were coupled as a single fluid⁷, and conversely, after this time, the photons had an essentially infinite path length. The time in between at which hydrogen, helium, and a few heavier elements formed out of nucleons and electrons is called the epoch of recombination (or alternatively the epoch of last scattering). This epoch was almost instantaneous so the photons which we observe as the CMB radiation originate in a sphere with a radius equal to the time since recombination. This sphere is called the last scattering surface.

CMB experiments measure the intensity of radiation per unit frequency per polarization at each point in the sky. In the absence of other sources of photons, these instruments are directly measuring the photons from the last scattering sphere. Instruments capable of precisely measuring the intensity of the CMB as a function of frequency show that it is incredibly well fit by a Planck spectrum with $T_0 = 2.725 \pm 0.002$ K (Mather et al., 1999). The derivative of the Planck function with respect to temperature can be used to relate measurements of intensity to a physical temperature of the photons.

This being said, the interesting observation about the CMB is not necessarily how close it is to isotropic. Because we see so much structure in the Universe around us (galaxies, star, planets, people), we must conclude that the Universe was never completely homogeneous and isotropic. Therefore, it is desirable to study the anisotropy of the CMB to understand the structure present in the later Universe. The 'interesting' CMB experiment therefore compares different parts of the sky to determine the point to point variation in its temperature field. One can represent the temperature at a given point $\hat{\mathbf{n}}$ by using the standard Stokes parameters $T = \{I, Q, U, V\}$ (Jackson, 1998). One then maps the various measurements of T as a function of $\hat{\mathbf{n}}$ and studies the statistics of these maps. This is the fundamental program of most CMB experiments.

Because the sky is best projected onto a sphere rather than Cartesian coordinates, these kind of measurements can be given as a spherical harmonic decomposition

$$\frac{T(\hat{\mathbf{n}})}{T_0} = 1 + \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^T Y_{\ell m}(\hat{\mathbf{n}})$$
(1.9)

where

$$a_{\ell m}^{T} = \frac{1}{T_0} \int d(\hat{\mathbf{n}}) T(\hat{\mathbf{n}}) Y_{\ell m}^*(\hat{\mathbf{n}})$$
(1.10)

in a black-body population are required to keep a proton/electron fluid dissociated, the Universe's temperature at recombination was lower than this. Detailed calculations give $kT \gtrsim 0.3 \,\text{eV}$ as the cosmic temperature at recombination.

⁷The study of the Universe very much before recombination is an interesting field, but beyond the scope of the present discussion, which will focus mainly on the physics which impact the observed CMB. The reader is referred to, for example, Peacock (1999) for a presentation of the particle physics of the big bang.

and * denotes complex conjugation (see e.g. Kamionkowski et al. (1997)). The $\ell = 0$ term is excluded from the expansion because it is represented by 1; this is the CMB monopole, and, as discussed above, dominates the other terms by approximately 10^5 . The $\ell = 1$ term is the 'dipole', and is mostly caused by the Doppler shift due to the observer's motion with respect to the monopole. It is impossible to separate the Doppler induced signal from the CMB dipole, so this term is usually neglected. The $\ell \geq 2$ terms are thus what really interest us.

The various $a_{\ell m}^T$ in Equation 1.10 completely describe the radiation pattern on the sky. Multipole moments C_{ℓ}^T of these $a_{\ell m}^T$ can be formed using

$$\langle a_{\ell m}^{T*} a_{\ell' m'}^{T} \rangle = C_{\ell}^{T} \delta_{\ell \ell'} \delta_{m m'} \tag{1.11}$$

where the averages are over all realizations of the sky which can be produced by the statistical processes that yield the T fields (this is called an ensemble average power). If the fluctuations in the CMB field are Gaussian distributed, the statistical properties of a map of T are fully specified by the multipole moments. The C_{ℓ}^{T} are the most commonly used description of the statistics of a CMB map (or simulated map realization).

1.1.3 PRIMARY TEMPERATURE ANISOTROPIES

Armed with the tools developed in Section 1.1.2, we can now turn to the physics which produces the anisotropies in the CMB. There are a great number of treatments of this subject in the literature, but two particularly lucid and enlightening versions are Kosowsky (2002), Kosowsky (2002), while Hu & Sugiyama (1995) is a seminal paper on the subject and a classic in CMB theory. Here we outline only enough to give an idea of the physics behind the observed features in the CMB. A full treatment uses numerical simulations to evolve the fluid and Boltzmann equations for the different particles in the Universe before recombination, and computes the various effects the photons experience at the sphere of last scattering to determine the CMB we would observe today⁸.

The Universe immediately before recombination was a simple place, composed of baryons⁹, neutrinos, photons, and dark matter particles. The neutrinos and dark matter particles only interacted via the gravitational force, as cross sections to the weak interaction are very small. The baryons and photons interacted electromagnetically, predominantly via Thompson scattering. Baryons were also influenced by gravity which, due to the tight coupling of the photons and baryons, coupled to the photon fluid as well.

⁸Whole theses have been devoted to the details and subtleties involved in this one sentence.

⁹In CMB theory, this term encompasses protons, fully ionized helium and heavier elements, and electrons, even though electrons aren't strictly baryons.

Today, the $a_{\ell m}$ are a set of complex coefficients which describe the temperature of each point in the Universe at recombination (at least in the foreground-free limit). However, in the Universe before recombination, an observer would see a very uniform radiation field (in fact, one can show that only the $\ell = 0$ and $\ell = 1$ modes can exist in the tightly coupled regime, Hu & White (1997)) near thermal equilibrium governed my Thompson scattering of photons off free electrons. The differential cross section of Thompson scattering is described by

$$\frac{d\sigma}{d\Omega} = \frac{3\sigma_{\rm T}}{8\pi} |\hat{\mathbf{E}} \cdot \hat{\mathbf{E}}'|^2 \tag{1.12}$$

where $\hat{\mathbf{E}}$ and $\hat{\mathbf{E}}'$ are the electric fields of the incoming and outgoing photons and σ_{τ} is the Thompson cross section. The mean free path of the photons, λ_{τ} , is given by

$$\lambda_{\rm T} = \frac{1}{\sigma_{\rm T} n_{\rm e}(t) a(t)} \tag{1.13}$$

where $n_{\rm e}(t)$ is the electron density around the scatterer as a function of time. Before recombination the photons would have originated (on average) a distance $\lambda_{\rm T}$ away from the scatterer; this is quite a short distance by cosmological standards, on order Mpc at recombination. Due to the symmetries of Thompson scattering, only the monopole, dipole and (small) quadrupole terms of the radiation field are important before recombination; the rest cancel on average. To understand CMB anisotropy, it is sufficient to understand how these three multipole moments evolve in the early Universe.

If the early Universe were perfectly smooth, the physics described above would operate, but every multipole moment but the monopole would be zero. However, because the present Universe is not perfectly smooth, we know that the pre-recombination Universe could not have been, even though the early anisotropy was quite small ($\sim 10^{-5}$ as alluded to above). In order to produce perturbations in the photon-baryon fluid in the early Universe which can evolve to structure in the CMB or today, we need to perturb the metric describing the space–time (Equation 1.2). For small perturbations, we can write

$$g_{\mu\nu} = \gamma_{\mu\nu} + h_{\mu\nu} \tag{1.14}$$

where $\gamma_{\mu\nu}$ is the original, unperturbed FRW metric, and $h_{\mu\nu}$ is the perturbation to the metric. For simplicity, we chose the conformal Newtonian representation; in this form, the metric's normal ten degrees of freedom are reduced to two via coordinate gauge symmetry and the assumption that shear and vorticity terms (*i.e.* off-diagonal terms in the metric) are not present. These two terms are Ψ , the gravitational redshift or h_{00} , and Φ , the Newtonian potential or h_{ii} perturbation:

diag
$$(h_{\mu\nu}) = [-(1+2\Psi), 1-2\Phi, 1-2\Phi, 1-2\Phi].$$
 (1.15)

All other terms in $h_{\mu\nu}$ are zero in this gauge.

Now, with these perturbations in place, the evolution of the photon and baryon fluids satisfy different continuity equations:

$$\dot{\Theta} = -\frac{kv_{\rm T}}{3} - \dot{\Phi} \tag{1.16}$$

$$\dot{\delta}_{\rm b} = -kv_{\rm b} - 3\dot{\Phi} \tag{1.17}$$

but coupled Euler equations:

$$\dot{v}_{\rm T} = k \left(\Theta - Q_{\gamma}\right) + k\Psi + a n_{\rm e} \sigma_{\rm T} (v_{\rm b} - v_{\gamma}) \tag{1.18}$$

$$\dot{v}_{\rm b} = -Hv_{\rm b} + c_s^2 k \delta_{\rm b} + k \Psi^2 + Ran_{\rm e} \sigma_{\rm T} (v_{\gamma} - v_{\rm b}).$$
 (1.19)

In these equations, k is an arbitrary wave vector, Θ is the temperature field, δ_b is the baryon density field, v_i is the velocity field of the i^{th} component, and Ψ , Φ , a, n_e and σ_T are defined above. In addition, c_s is the sound speed in the fluid, $R = 3\rho_b/4\rho_\gamma$ is the ratio of baryons to photons, and Q_γ is the anisotropic stress which is only important near recombination.

Normally, these equations are evolved using a computer to produce the expected fluctuations in the photon-baryon fluid at a given time. However, it is instructive to assume that the coupling between the photons and baryons is strong so that the their fluid velocities and temperatures are equal (the 'tight-binding' regime: $v_{\gamma} = v_{\rm b}$, and $\dot{v}_{\gamma} = \dot{v}_{\rm b}$). Combining the two Euler equations, we obtain the (in)famous oscillator equation:

$$\ddot{\Theta} + \ddot{\Psi} - \frac{R}{1+R}H\dot{\Theta} - \frac{R}{1+R}H\dot{\Phi} + \frac{k^2}{3}(\frac{\Theta}{1+R} + \Psi) - \frac{R}{1+R}(1+3c_s^2)\Theta = 0. \quad (1.20)$$

This equation represents damped, driven oscillations of the radiation density, and says that the force acting on the photon-baryon fluid is determined by a competition between the pressure (restoring) and gravitational (driving) forces which causes the system to oscillate about its potential minimum. The frequency of this oscillation is given by

$$\omega^2 = \frac{k^2}{3(1+R)} = c_{\rm s}^2 k^2. \tag{1.21}$$

The oscillator equation is the fundamental relation which governs the evolution of anisotropies in the early Universe.

Given this knowledge of the behaviour of the photon and baryon fluids, all that is required to predict the evolution of the perturbations in the temperature, radiation and matter fields are some initial condition constraints. This is a complex topic; here, we assume initial conditions which produce phase coherent oscillations¹⁰. Using these initial conditions and our knowledge of how gravitation and Thompson scattering combine to produce acoustic oscillations, we can now qualitatively understand the power spectrum Figure 1.1.



Figure 1.1: A Model CMB Temperature Spectrum

This figure shows a CMB temperature power spectrum calculated using CMBFAST (Seljak & Zaldarriaga 1996, Zaldarriaga et al. 1998, Zaldarriaga & Seljak 2000). The independent variable is ℓ , which is a measure of angular scale on the sky. The dependent variable is $\ell(1 + \ell)C_{\ell}/2\pi$, which is a measure of the strength of the temperature fluctuations. The text describes the physical mechanisms which create the various features in this plot; this particular spectrum was created using the best-fitting cosmological model measured by three years of WMAP data (Hinshaw, 2006).

Beginning at low ℓ , the flat, low part of the spectrum in the range $2 \le \ell \le 50$ is due to super-horizon fluctuations which did not have time to begin oscillating before recombination. The multipoles in this range best reflect the initial power spectrum of the Universe, and highlight the famous causality problems associated with the CMB. The first peak is due to the largest-scale oscillation, which was maximally compressed at recombination and is extremely sensitive to the curvature of the Universe κ . The second peak, caught after one full oscillation, is relatively lower than those around it because of 'baryon loading'. This

¹⁰This is not necessarily the simplest situation, but arises naturally from inflation, discussed below.

arises from the expansion of the Universe: as the Universe expands, densities fall by $a^{3}(t)$. Once the baryons in a potential have reached their maximal pressure, the decay of the potential has lessened the difficulty of re-expanding. This serves to enhance the compressional (i.e. first, third, fifth...) peaks, or alternatively, it serves to suppress the even peaks compared to them. The third peak, caught after one and a half oscillations, is somewhat high because of the presence of dark matter, which deepens the effective potential. Dark matter also shifts the first peak to lower ℓ than would be expected in a baryon-only Universe. Finally, imperfections in the photon-baryon fluid, namely heat conduction and shear viscosity, tend to damp the acoustic oscillations at small scales. They basically add a damping term to the oscillator equation at large ℓ , which has amplitude $e^{-k^2 t \lambda_T}$. The damping scale is therefore on order $1/\sqrt{t\lambda_{\rm T}}$, which is most important at small distances. This leads to a dramatic suppression of the power spectrum beyond the third peak, as the fluctuations on these scales have been damped away. This is generally called 'Silk damping' and has been observed by CMB experiments in recent years. These are the major features of the CMB temperature power spectrum imprinted at recombination, and together tell us a great deal about the Universe at that epoch.

1.1.4 SECONDARY TEMPERATURE ANISOTROPIES

Secondary CMB anisotropies are defined to be effects which change the energy or spatial distribution of the CMB photons between the sphere of last scattering and the observer. They can be broadly broken into three different categories:

- (i) foreground emission and absorption,
- (ii) gravitational interactions, and
- (iii) electromagnetic interactions.

Each are discussed in turn in this section.

CMB foregrounds can be subdivided into two categories; local foregrounds, and astrophysical foregrounds. Both of these foregrounds are a nuisance for the CMB experimenter, and a great deal of time and effort is expended in trying to minimize their effects. Local foregrounds are both things like the sun and moon, which move around the sky in a predictable way over time, and other sources of emission and absorption, like the ground or atmosphere around the telescope. Mitigation of objects like the sun and moon is simple; one merely observes in a direction where these objects aren't visible. The ground and atmosphere (and particularly the latter) require slightly different techniques. The ground can be effectively shielded from view using an isotropic reflector to block photons originating

there. The atmosphere is a much more complex subject, as its characteristics vary over time (as discussed in *e.g.* Archibald et al. (2002) and Peterson et al. (2003)). Essentially, this particular foreground requires chopping between a target and reference signal, although how this is implemented experimentally varies considerably. Ground-based experiments working above approximately 100 GHz are dominated by atmospheric noise.

Astrophysical foregrounds are those sources of emission and absorption fixed in the sky reference frame. Examples include synchrotron, free-free and thermal dust emission from our galaxy, and extragalactic sources like active galactic nuclei (in various forms) and star bursting galaxies. These sources are often studied in their own right, and a number of methods of removing their effect from CMB data exist. These include source emission templates (whether in the form of a list of point sources, an extrapolated dust map, or a map of synchrotron emission), filtering based on the spatial power spectrum of the source, or even frequency based weighting. The simplest method of minimizing the effects of astrophysical foregrounds is merely to observe the places on the sky where they are smallest.

Gravitational interactions are, as the name suggests, due to concentrations of mass along the direction of CMB photon propagation. Although photons are massless, gravity can affect their paths by warping space-time according to general relativity. This warping is the cause of two basic affects: gravitational lensing which is due to warps orthogonal to the direction of CMB photon propagation, and the Rees-Sciama (RS) effect which is due to time varying potentials parallel to the photon propagation direction (Rees & Sciama, 1968).

The study of gravitational lensing is a complex subject, but broadly, one attempts to measure the tiny deflections in the CMB photon distribution caused by lensing due to assemblies of mass (*e.g.* Stompor & Efstathiou (1999), Peacock (1999)). These deflections are most efficiently observed using some statistical measure, like the change in the CMB power spectrum that gravitational lensing would cause¹¹. Because large scale structure has some characteristic lensing deflection as a function of lens mass, one can compute the expected effect of this foreground using models for large scale structure. The photons are achromatically deflected according to the gradient of the gravitational potential projected along the line of sight, $\nabla \Phi$. Surface brightness is conserved in gravitational lensing, so the observed anisotropies are remapped as

$$T(\hat{\mathbf{n}}) \longrightarrow T(\hat{\mathbf{n}} + \nabla\Phi) \tag{1.22}$$

¹¹In fact, lensing affects both temperature and polarization anisotropies. As discussed later, it is expected to affect the *BB* polarization spectrum most strongly.

where $\hat{\mathbf{n}}$ is normal to the plane of the sky and T can represent polarized quantities as well. While individual deflections are small, the power of this effect lies in measuring the coherence of the deflections across the sky. In the temperature power spectrum, gravitational lensing tends to smooth the peaks slightly.

The Rees-Sciama effect is schematically simple. When a photon traverses a time varying potential, it accumulates a gravitational redshift which causes a temperature perturbation

$$\frac{\delta T}{T} = 2 \int dl \,\dot{\Phi} \tag{1.23}$$

where dl is the infinitesimal path length along the line of sight, and Φ is the potential's time derivative. The Φ one expects from large scale structure in the Universe is largely due to the gravitational collapse of the potential well; as the photon 'falls' into the potential, it accumulates a blueshift, and as it 'climbs' back out, is accumulates a redshift. If the potential has become deeper in the intervening time, there is a net redshift, or energy loss, of the photon. As such, the RS effect is simply a non-linear, secondary version of the integrated Sachs-Wolfe (ISW) effect. Generally, the RS effect is canceled as photons traverse many crests and troughs between the CMB and the observer. However, at low ℓ this effect should not cancel completely, and so should provide a good way to measure the equation of state of dark energy and the large scale structure of the Universe. Unfortunately, this measurement is difficult for a number of reasons. Firstly, because it requires precise comparison of huge patches of sky, this measurement is prone to many systematic errors. Secondly, because there are only a limited number of large patches of sky to compare in the observable Universe, this measurement is subject to a large sample-limited variance, which restricts the amount of information that can be gathered. Precise measurement of the RS effect will therefore require more time and instrumental development to fully exploit.

Secondary temperature anisotropies due to electromagnetic interaction come in two varieties, the Ostriker-Vishniac (OV) effect, and the Sunyaev–Zel'dovich (SZ) effect. The OV effect is scattering from large scale bulk motions of partially ionized gas in the early Universe (Ostriker & Vishniac (1986), Vishniac (1987)). After the first stars began to shine, their UV radiation dissociated the atomic species in the early Universe. During this epoch, the dissociated electrons could act as scatterers for CMB photons. Although only a few per cent of the electrons acted as scatterers, they had a dramatic effect on the CMB by re-isotropizing the photons, thus reducing the heights of the peaks in the CMB power spectrum. This effect is difficult to disentangle from the primordial heights of the peaks, and as such any such effort needs to be treated with care. The OV contribution to the CMB power spectrum is mainly at very high ℓ , and has not yet been detected in temperature alone (although the correlation of temperature with polarization is a different matter, as discussed

below).

The other important electromagnetic interaction CMB photons experience on their way to the observer is from inverse Compton scattering off hot electrons in galaxy clusters, called the Sunyaev–Zel'dovich effect (originally discussed in Zeldovich & Sunyaev (1969), and recent reviews by Birkinshaw (1999) and Carlstrom et al. (2002)). The ionized gas in the cluster potential interacts with CMB photons via inverse Compton scattering. The probability of this occurring to any given photon is given by

$$\tau = \int dl \,\sigma_{\rm T} \,n_{\rm e}(l) \tag{1.24}$$

where $n_{\rm e}(l)$ is the electron density as a function of the distance along the line of sight and τ is the optical depth through the cluster. Since the total probability of scattering is low $(\tau \sim 0.01$ through the densest parts of very massive clusters), photon number is conserved in this process. Although in any given scattering event, energy can be gained or lost, the photons gain energy on average. This is equivalent to a net temperature increase, which gives rise to a spectral distortion of the CMB as seen through the cluster.

The canonical result for the SZ effect distortion shape is given by

$$\Delta I(x) = I_0 y \frac{x^4 e^x}{(e^x - 1)^2} \left[x \coth(x/2) - 4 \right] \left(1 + \delta_{\text{SZE}}(x, T_{\text{e}}) \right)$$
(1.25)

where $x = h\nu/kT$ is the dimensionless frequency, I(x) is the intensity at x, I_0 is the intensity through the centre of the cluster, y is the Compton parameter given by

$$y = \tau \frac{kT_{\rm e}}{m_{\rm e}c^2} \tag{1.26}$$

and δ_{sze} gives corrections to the analytic function due to a relativistic electron population (which is only important for very large electron temperatures T_e). Equation 1.25 is a wellbehaved analytic function yielding a cluster's SZ effect based only on a single parameter y, which is simply the electron gas' pressure integrated along the line of sight. The relativistic correction δ_{sze} depends only on x and the temperature of the electron gas, albeit in a fairly complex way (Rephaeli & Yankovitch, 1997). This function is generally calculated via a numerical solution to the Kompaneets equation, as discussed in (for example) Itoh & Nozawa (2004).

The SZ effect can be used effectively for a variety of cosmological observations. Of great interest is its use as a tracer of mass and structure formation in the Universe. SZ measurements are more robust than X-ray observations because, while cosmological dimming affects X-ray photons as z^4 , the SZ distortion is virtually independent of redshift.

This comes about because the (1 + z) effect of cosmological dimming is exactly canceled by the (1 + z) increase in CMB temperature with redshift. The experiment of interest is to find SZ clusters in a blank field of sky. This sample should be a mass-limited, unbiased sample of clusters in the field. The redshifts of these clusters are found via some other means, and the volume density of clusters is determined. Given enough clusters, this type of experiment can constrain a variety of cosmological parameters, like the density measure σ_8 , the dark energy Λ as a function of redshift, and w, which parametrizes the equation of state of the dark energy. Currently, this type of measurement is attracting a huge amount of effort in experimental cosmology. Furthermore, the SZ effect can be used in combination with the X-ray brightness distribution of a given cluster to measure a variety of parameters. Of great interest is the measurement of distance through the cluster; if one assumes that the cluster is spherically symmetric, the angular diameter distance of the cluster gives the proper distance to it. This is completely independent of the cosmological distance ladder, which makes it very attractive. One can go further and measure the Hubble constant with this method, although this type of experiment has traditionally been severely hampered by systematic effects. QUaD is scheduled to attempt a measurement of the SZ effect in bright, southern clusters during the 2006-2007 summer season.

1.2 POLARIZATION OF THE COSMIC MICROWAVE BACK-GROUND

The primary polarization anisotropies in the CMB radiation are imprinted because Thompson scattering is a function of the polarization of the incoming and scattered photons. The scattered radiation intensity peaks in the direction normal to the incident radiation polarization. Thus, if the radiation field in the photon-baryon fluid is anisotropic before recombination, this can imprint a polarization on the scattered radiation. The polarization fraction is directly related to the size of these anisotropies, and we expect a polarized signal about 10 % of the total temperature of the CMB (on order μ K).

In this section, the physics of both primary and secondary polarization CMB anisotropy are presented. Furthermore, the science which can be extracted from the polarization, and the methods used to achieve this, are discussed. The original work on CMB polarization is given in Rees (1968), while a modern review can be found in Hu & White (1997), and the mathematical details of the subject can be found in Zaldarriaga & Harari (1995) or Kamionkowski et al. (1997).

1.2.1 PRIMARY POLARIZED ANISOTROPIES

The primary polarized anisotropies in the CMB radiation were laid down by local quadrupolar moments during recombination. The quadrupolar dependence comes from the Thompson scattering cross section, given in Equation 1.12; by the physics of the process and the orthogonality of spherical harmonics, no other scattering symmetry is possible. We know that the polarization anisotropies must have been imprinted at reionization because a non-negligible quadrupolar moment in the CMB radiation field only existed during last scattering as the optical depth of the matter–radiation field changed from very large values to zero in a short time¹². Based on this, we can predict that the angular size at which we expect the CMB polarization signal to peak is the horizon scale at recombination, corresponding to ~ 1° in a flat Universe.

To proceed further, it is useful to define the Stokes parameters, given by

$$I = \langle E_x^2 \rangle + \langle E_y^2 \rangle \tag{1.27}$$

$$Q = \langle E_x^2 \rangle - \langle E_y^2 \rangle \tag{1.28}$$

$$U = 2 < E_x > < E_y > \cos\delta \tag{1.29}$$

$$V = 2 < E_x > < E_y > \sin \delta. \tag{1.30}$$

These four variables -I, Q, U, V - completely define the radiation field. In this formulation, the photon is propagating in the z direction, so E_x and E_y are the electric fields in the x and y directions, respectively, and δ measures the phase between these two components (this is explained in, for example, Hecht (1998)). In a north, east, south, west coordinate system, a N-S (E-W) polarization is Q > 0 (Q < 0), while a NE-SW (NW-SE) polarization is U > 0 (U < 0).

The first observation to make here is that Compton scattering *cannot* create circular polarization V; we will therefore ignore this parameter hereafter¹³. Furthermore, I is identified with T, which was discussed above; we will therefore restrict our attention to Q and U in this section.

Q and U can be rotated by an angle ϕ and describe the same polarization field under

¹²The epoch of recombination is measured to have lasted a FWHM redshift interval of $\Delta z = 195 \pm 2$ by WMAP (Bennett et al., 2003).

¹³In passing it may be mentioned that CMB instruments are generally not built to measure V; this design feature is essentially driven by a combination of expediency and this non-existence argument. If a non-zero V were present, it would be attributable to a foreground, either astronomical or atmospheric, and thus be uninteresting from the point of view of CMB science. By the same argument, a non-zero V in the CMB radiation would point to a gross misunderstanding of the physics involved in the early Universe.

the transformation

$$Q' = Q\cos 2\phi + U\sin 2\phi$$
$$U' = -Q\cos 2\phi + U\cos 2\phi$$
(1.31)

with Q' and U' being the new coordinates defining the field (Kosowsky, 2002). It is easy to see the total polarization $P^2 = Q^2 + U^2 = Q'^2 + U'^2$ remains constant under this transformation. We can define a tensor which behaves correctly under these transformation properties by

$$P_{ab}(\hat{\mathbf{n}}) = \frac{1}{2} \begin{pmatrix} Q(\hat{\mathbf{n}}) & -U(\hat{\mathbf{n}})\sin\theta \\ -U(\hat{\mathbf{n}})\sin\theta & -Q(\hat{\mathbf{n}})\sin^2\theta \end{pmatrix}$$
(1.32)

where θ is the standard spherical coordinate. This tensor quantity can be expanded in terms of tensor spherical harmonics, which are just generalizations of standard spherical harmonics. This can be formulated via

$$\frac{P_{ab}(\hat{\mathbf{n}})}{T_0} = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} \left[a_{\ell m}^E Y_{\ell m, ab}^E(\hat{\mathbf{n}}) + a_{\ell m}^B Y_{\ell m, ab}^B(\hat{\mathbf{n}}) \right]$$
(1.33)

where

$$a_{\ell m}^{E} = \frac{1}{T_{0}} \int d\hat{\mathbf{n}} P_{ab}(\hat{\mathbf{n}}) (Y_{\ell m}^{E,ab}(\hat{\mathbf{n}}))^{*}, \qquad (1.34)$$

$$a_{\ell m}^{B} = \frac{1}{T_{0}} \int d\hat{\mathbf{n}} P_{ab}(\hat{\mathbf{n}}) (Y_{\ell m}^{B, ab}(\hat{\mathbf{n}}))^{*}.$$
 (1.35)

These equations are quite complex; the reader is referred to Kamionkowski et al. (1997) for a construction of them explicitly in the case of CMB radiation fields.

The two sets of basis functions, denoted E and B, arise from the fact that any symmetric and traceless 2×2 tensor can be decomposed into 'gradient only' and 'curl only' components. The names E and B come from analogy to electromagnetism - the gradient only component is electric-field like (an E-mode) and the curl component is magnetic-field like (a B-mode). This distinction is actually a meaningful one; in this basis, E-modes contain no curl component, and B-modes contain no gradient component.

Armed with E and B, we can define the set of two point correlation functions (or, via the cross-correlation theorem, power spectra) analogous to Equation 1.11 via

$$\langle a_{\ell m}^{T*} a_{\ell' m'}^{T} \rangle = C_{\ell}^{TT} \delta_{\ell \ell'} \delta_{mm'}, \quad \langle a_{\ell m}^{T*} a_{\ell' m'}^{E} \rangle = C_{\ell}^{TE} \delta_{\ell \ell'} \delta_{mm'}, \langle a_{\ell m}^{E*} a_{\ell' m'}^{E} \rangle = C_{\ell}^{EE} \delta_{\ell \ell'} \delta_{mm'}, \quad \langle a_{\ell m}^{T*} a_{\ell' m'}^{B} \rangle = C_{\ell}^{TB} \delta_{\ell \ell'} \delta_{mm'}, \langle a_{\ell m}^{E*} a_{\ell' m'}^{B} \rangle = C_{\ell}^{EB} \delta_{\ell \ell'} \delta_{mm'}, \quad \langle a_{\ell m}^{B*} a_{\ell' m'}^{B} \rangle = C_{\ell}^{BB} \delta_{\ell \ell'} \delta_{mm'}.$$

$$(1.36)$$

The same conditions applied to Equation 1.11 apply here.

Estimators of the different power spectra can be created from temperature and polarization maps of the CMB radiation, making these six spectra (TT, EE, BB, TE, TBand EB) the *de facto* standard representation of CMB information. By parity arguments, TB and EB power must be zero¹⁴, meaning that the other four power spectra are enough to study the CMB radiation.

Turning back to physics in the early Universe, it is a well-known (but slightly mindtwisting) result that the temperature perturbations that cause quadrupole moments in the photon-baryon fluid can arise from three distinct perturbations: scalar, vector, and tensor (corresponding to the $m = 0, \pm 1, \pm 2$ modes of $\ell = 2$: Hu & White 1997 gives a very good treatment of the subject). Considering these in turn for an adiabatic perturbation scenario, the scalar perturbations are due to two effects: (i) quadrupoles produced at the peaks/troughs of density anisotropies and (ii) quadrupoles produced by velocity perturbations perpendicular to temperature gradients in the fluid. Vector perturbations are due to vortical motions of the matter (*i.e.* are velocity rather than density perturbations), and do not have any gradient component. Tensor fluctuations are transverse traceless perturbations to the metric, and as such represent a stretching of space-time along one axis, and a compression of space time along the other. This is exactly how a spin-2 gravitational particle would affect the metric (see, for example, Weinberg 1972).

Given these general considerations, we can immediately simplify the model of the polarization we expect based on the physics. Compton scattering can only produce a local E-mode, since primordial density perturbations can only produce scalar fluctuations¹⁵ and these must be rotationally symmetric about the k of the perturbation. Therefore, of the spectra listed in Equation 1.36, we only expect TT, TE and EE to be non-zero in an adiabatically perturbed Universe at recombination.

A complexity is introduced when we consider rather uncomfortable questions like:

- Why is the Universe measured to be so close to flat ($\Omega_{tot} = 1$)? This value can be shown to be meta-stable, *i.e.* Ω_{tot} diverges from this value very quickly unless tuned to be *exactly* unitary.
- Why is the measured CMB radiation so close to isothermal, when it is impossible for it to have been in causal contact at any epoch near to recombination (unless the speed of light were some 100 times greater than it is today)?

¹⁴However, these spectra are usually generated when considering real CMB data to act as a sanity check. If either is inconsistent with zero, it points to an error somewhere in the analysis.

¹⁵Vector and tensor perturbations may have been present as an initial condition of the CMB, but these are washed out due to the very short mean free path of photons in the early Universe.

• By the same token, what laid the perturbations which caused the CMB anisotropies, later allowing the gravitational collapse of complex structures, in the first place? What set them to be so small (1 part in 10⁵)?

Two classes of models try to answer these questions, and others: inflation models, and cosmological defect models. Defect models deal with phenomena due to symmetry breaking at phase transitions in the early Universe, and while not ruled out, are becoming increasingly constrained in model parameter space. We therefore restrict our discussion to inflationary models, which postulate that the Universe went through a brief but critical super-luminal expansion in its very early history. Good reviews of the subject can be found in, for example, Guth (2004) and Kinney (1998); we briefly sketch the general ideas behind the subject here¹⁶.

Essentially, inflation postulates the existence of an epoch long before recombination when $\ddot{a}(t) > 0$. This can be phrased in terms of the Hubble parameter H as

$$a \propto \exp\left(\int dt H\right) = e^N$$
 (1.37)

where N is defined to be the number of e-folds of the expansion. During inflation H is nearly a constant value, and the expansion of the Universe is parametrized by a quasiexponential. This leads to comoving distances rapidly increasing in size relative to the horizon distance, which is also approximately constant ($d_H \simeq H^{-1}$). Thus, regions initially in causal contact expand away from one another so quickly that, by the end of inflation, they are completely non-causal. This explains the near isotropy of the CMB; parts of the Universe which appear to never have been in causal contact could have been before or during inflation. To match the observed flatness and homogeneity of the Universe, between 50 and 70 expansion e-foldings are necessary, depending on the precise model we are considering.

Not only do inflationary models solve the problem of flatness, but they provide a mechanism by which the anisotropies of the CMB are seeded, and can provide predictions for their amplitudes. The process of inflation smooths the Universe almost completely by washing out structures of any appreciable size which occurred in the very early Universe. However, during inflation quantum mechanical fluctuations occurring on very small scales are rapidly shifted to scales much larger than the horizon size, where they are frozen in as the initial conditions for the background metric. These then evolve under the model

¹⁶As can be found in these discussions, inflation attempts to explain far more than the few questions discussed above and as such is quite an ambitious intellectual endeavour.

discussed in Section 1.1.3 to produce the observed CMB radiation anisotropy. These fluctuations are generically predicted to be adiabatic, Gaussian distributed, and nearly scale– invariant, which are all features found in the observed CMB.

To put this in the frame of reference of the polarization mechanisms discussed above, the metric perturbations laid down by inflation are of two types: scalar fluctuations, which provide the seeds of the standard CMB anisotropy, and tensor fluctuations, which take the form of gravitational radiation. The gravitational radiation causes some amount of B-mode power to exist at recombination, and for high enough gravitational wave energy densities, this leads to observable B-mode power in the resulting CMB polarization power spectrum.

Calculating a CMB prediction from an inflationary model requires defining a 'inflaton' potential function, which describes the evolution of the Universal inflation over time. This inflaton leads to a prediction for the initial power spectrum of both the scalar and tensor fluctuations, which can be used as initial conditions to CMB evolution code. This leads to the six power spectra in Equation 1.36, which can be parametrized in terms of a minimum set of model parameters:

- 1. the tensor/scalar ratio r,
- 2. the spectral index of the power spectrum n,
- 3. the normalization of the power spectrum A,
- 4. the baryon density of the Universe $\Omega_{\rm b}$,
- 5. the total matter density of the Universe Ω_m , and
- 6. the Hubble constant h.

To this minimum set a number of other model parameters can be added which describe secondary effects that change the observed anisotropy. Figure 1.2 shows power spectra resulting from assuming the WMAP 3 year data as a fiducial model.

All this being said, what are the practical implications of a measurement of CMB polarization? Essentially, there are three major purposes for CMB polarization measurement:

1. Measurement of the partial polarization of the CMB is a fundamental check that our basic model for the production of the fluctuations in the CMB and the physics at work in the early Universe is correct.


Figure 1.2: Model CMB TT, TE, EE and BB Spectra

This figure shows the 4 non-zero CMB power spectra calculated using CMBFAST, as in Figure 1.1. These spectra were calculated using the fiducial 3 year WMAP model, assuming a simple inflationary model with adiabatic initial conditions. There is a clear hierarchy of power here, with TT dominating the EE power, and BB below both. TE is plotted separately because it can be negative (that is, TT and EE can be anti-correlated at various ℓ).

- 2. Because the temperature spectrum can be distorted and changed by a variety of effects in the post-recombination Universe, polarization allows a clean measurement of the conditions *at* recombination. By including polarization information, the ingredients of the cosmological model (*e.g.* whether inflation or defect models are a better description of the initial power spectrum) can be distinguished in a way not possible with temperature alone.
- 3. The polarization power spectrum allows one to break degeneracies in the standard

scalar perturbation paradigm by providing information complementary to the temperature power spectrum (e.g. a change in the temperature power spectrum's normalization versus 'late' reionization).

Thus, a measurement of the polarization power spectrum provides a number of unique checks and measurements not available by any other means.

Before moving on, item (2) above is worth a brief discussion in the context of current and future CMB polarization experiments. Currently, a phrase bandied around quite frequently is that a detection of *B*-mode power is 'smoking gun' evidence in favour of inflation. But why is such a measurement definite proof that some type of inflation happened? The logic is somewhat convoluted: we attempt a simple explanation here, based on Spergel & Zaldarriaga (1997) and Kinney (1998).

Of the two classes of models attempting to explain the cause of the initial perturbations which evolve into structure in the Universe (*i.e.* inflation and cosmological defects), only inflation can produce correlations in the power spectrum larger than the horizon size at recombination. This is due to the physics discussed above: quantum mechanical perturbations in the extremely early Universe are inflated by at least $N \sim 50$ to produce the seed cosmological fluctuations. In this picture, there are correlated fluctuations on all scales, even ones which inflate to larger than the horizon size at recombination. Defect models have no mechanism to cause the same effect; thus, if super-horizon power exists in the measured power spectrum, it must be due to some type of inflation having occurred.

This comment is true for all of the TT, EE or BB power spectra. However, in TT the ISW effect can mimic these large scale correlations, and with the TT spectrum alone there is no way to determine whether apparent super-horizon power is due to the ISW effect or correlations due to inflation. Furthermore, there is a similar degeneracy in the EE power spectrum from the signature of reionization, discussed below. Reionization also damps the TT power, complicating the picture as well. Both of these effects change the power spectrum at low ℓ , the only part of the power spectra where the original signature of inflation is not obscured by evolution before recombination.

The only power spectrum which is not affected by some foreground at these scales is BB. Therefore, detection of cosmological BB power at low ℓ constitutes a unique proof of inflation, simply because correlations at these ℓ cannot be caused by any other means. This simple picture is, of course, complicated by the existence of foregrounds, but the fact remains that measurement of B-mode power is the only way to achieve a clear detection of the large-scale correlations induced by inflation. The other proviso here is that, because the quadrupolar moment which causes primordial B-mode power is due to gravitational radiation in the early Universe, a non-detection of B-modes does not mean inflation did not occur. It could simply be that the energy density in gravitation radiation in the early Universe was too small to make an appreciable impact on the CMB. However, because of the unique ability of B-mode polarization to constrain inflation, the next generation of CMB polarimeters will be designed with the 'holy grail' of BB power spectrum detection as their major driver.

1.2.2 SECONDARY POLARIZED ANISOTROPIES

There are two major secondary polarized CMB anisotropies: those due to reionization, and those due to gravitational lensing of background CMB photons by mass in the Universe between the observer and the CMB¹⁷. Furthermore, polarized foregrounds might be a major contaminant for CMB polarization observations; these subjects are briefly discussed in this section.

Reionization – the ionization of the inter-galactic medium by UV-bright stars in the very early Universe – produces a distinctive signature in the CMB power spectra by rescattering CMB photons from free electrons at $z \gtrsim 5$. This effect erases fluctuations below the horizon scale and regenerates them only weakly, so large scale power in the TT spectrum is reduced. The opposite is true for polarization: because the quadrupolar fluctuations at recombination are large, the power in EE due to Thompson scattering is increased. However, as the optical depth at reionization is quite low, this is only a small (but measurable) effect in TT and EE; it shows up quite strongly in TE as well (*viz.* the first bump in TE in Figure 1.2). These spectra can be used to determine the mean redshift of reionization due to the first generations of stars in the Universe by fitting the three for the parameters τ , the optical depth to reionization, and z_{reion} , the mean redshift of the process.

Gravitational lensing of CMB photons by mass between the surface of last scattering and the observer (discussed above for temperature anisotropies) deflects the path of the photons perpendicular to their direction of travel, thus changing the observed brightness distribution of the CMB. This random deflection smears out CMB structure in temperature, but has slightly different effects on polarization (Zaldarriaga & Seljak, 1998). The first effect of gravitational lensing is to change the heights of the oscillation peaks in *E*-mode power, which is similar to but fractionally larger than its effect on *TT*. The high ℓ end of the spectrum especially is affected due to the typical size scale of the lensing structures. The other, slightly more interesting, effect is to turn primordial *E*-mode power into *B*-mode power. This works because lensing can induce an overall curl-like 'shear' in the path of the CMB photons, which creates apparent *B*-mode power from pure *E*-mode power. This effect peaks at around $\ell = 1000$, and is always quite small, typically $\lesssim 1\%$ of the *EE*

¹⁷In addition, there are a number of very minor polarized secondary anisotropies like the polarized SZ effect, the Ostriker-Vishniac effect, et cetera. At times it seems the bestiary of nK–level polarized secondary effects and foregrounds grows larger on a daily basis!

power. Because it is sensitive to structure formation, this *B*-mode power can constrain parameters like the amount of dark energy in the Universe, the mass of the neutrino(s), and other variables which affect the formation of structure. This lensing signal is thus of great interest to cosmologists, although the instrumentation required to robustly measure such an effect is almost certainly a few years away.

Because the amplitude of the interesting CMB polarization signals are so small, foreground contamination is a major issue. This is especially true of present and future all-sky surveys employing satellites, but QUaD also needs to be concerned that these might impact the final power spectra. Unfortunately, the amplitude and spatial distribution of these foregrounds is largely unknown at this point, and QUaD is very much a path finding experiment in the regard.

Since atmospheric emission is thought to be negligibly polarized in the mm regime at QUaD's sensitivity (Hanany & Rosenkranz, 2003), there are two basic classes of polarized foreground emission: diffuse galactic emission, and polarized extra-galactic point sources. The latter are expected to arise from polarized radio point sources known at other wavelengths, which are generally synchrotron emitting AGN galaxies with polarization fractions < 20 %. In general, these sources are not thought to be of major importance, although the very brightest can appear in QUaD maps (see Section 6.2.1). Fortunately, extensive catalogues for QUaD's survey region exist at lower frequencies; these highlight the positions of bright radio sources. Extrapolating the synchrotron spectrum to QUaD's bands, these catalogues list sources down to levels much dimmer than QUaD can detect; we are therefore confident that dim extra-galactic radio sources contribute a negligible amount of signal to the QUaD data. Those radio sources that are visible are excised from the map before power spectra are calculated.

Diffuse galactic emission is more problematic. Free-free emission is intrinsically unpolarized, but can be partially repolarized by Thompson scattering in HII regions. The polarization fraction of this emission is expected to be at most 10%, and very dim at QUaD's frequencies (but becoming increasingly important below 100 GHz). Galactic synchrotron emission is potentially a major contaminant. It can be highly polarized, and is expected to have a polarization fraction between 10 and 75% of a total signal strength as large as $50 \,\mu\text{K}$ at $30 \,\text{GHz}$. Assuming $T \propto \nu^{-3}$, this gives a signal strength of $\sim 1 \,\mu\text{K}$ at 100 GHz, which is large enough to be a contaminant for QUaD. Dust emission is perhaps the biggest potential problem, but also the least studied of these foregrounds. The polarization fraction of dust is not well known, but may be present at the $\sim 2\%$ level in diffuse interstellar clouds in the galaxy, and as high as > 10% near star forming regions. Although the brightness of thermal dust emission can vary greatly across the sky, typical values are large enough to make this a major contaminant for QUaD, particularly at 150 GHz. QUaD mitigates the effect of these foregrounds with three strategies. The first is completely passive: because QUaD is not a full-sky survey experiment, the regions where the very worst of these foregrounds occur (near the galactic plane) need not be observed. The second mitigation strategy is to choose bands which minimize the expected foreground contamination; the region between 90 and 200 GHz forms a minimum between the (falling) synchrotron emission and the (rising) dust emission in the galaxy. The two QUaD bands are therefore designed to inhabit these regions¹⁸. The third mitigation technique is to observe far from the galactic plane, in a region of the sky known to have very little of either kind of contamination from studies at other wavelengths (this is discussed further in Chapter 6).

1.2.3 THE EXPERIMENTAL STATE OF THE ART

Prior to 2002, the strongest limits on the polarization of the CMB were made by the PO-LAR experiment at small ℓ (Keating et al., 2001), and the PIQUE (Hedman et al., 2001), COMPASS (Farese et al., 2004) and CBI (Cartwright et al., 2005) experiments at larger ℓ . These set an upper limit on CMB polarization of approximately 10 μ K. The first convincing detection of the polarization signal was made by the DASI experiment in 2002 (Kovac et al., 2002), which convincingly detected *E*-mode power and made a statistical detection of the *TE* correlation. This detection 'set the bar' for the next generation of CMB polarization experiments, and has since led to a flurry of measurements.

There are three major styles of instrument used in CMB polarization work:

- 1. Interferometers, which measure the Fourier modes of the CMB on the sky and are intrinsically polarized devices (like DASI or CBI),
- 2. Pseudo-correlating radiometer based instruments, which detect CMB photons using standard square-law detectors in conjunction with millimetre waveguide, phase switches and HEMT amplifiers (like WMAP), and
- 3. Quasi-optical systems employing bolometric detectors (like QUaD).

There is no particular pattern to the environments in which these instruments operate: for example, WMAP is a satellite orbiting the second Lagrange point of the earth-sun system, BOOMERanG is a balloon based instrument, and QUaD is a ground based telescope situated at the south pole. None of these types of instrument has proven to be intrinsically superior to the others at this point: these designs all have advantages and drawbacks, and

¹⁸The corollary to this argument is, of course, that *nothing* astronomical is particularly bright at these wavelengths. This makes finding sources for calibration a major problem for QUaD and other instruments working at these wavelengths.

hybridized systems seeking to maximize sensitivity while minimizing systematic errors and cost are becoming popular.

Following DASI, the WMAP collaboration published a precise measurement of the TE correlation (Kogut et al., 2003). Other experiments like CAPMAP (Barkats et al., 2005), CBI (Readhead et al., 2004), DASI (Leitch et al., 2005) and BOOMERANG (Montroy et al., 2005) have since published increasingly better measurements, leading to the current set of data plotted in Figure 1.3.





This plot shows a summary of previous measurements of $\ell(\ell + 1)\tilde{C}_{\ell}^{EE}/2\pi$ versus ℓ from various experiments. The coloured points show the measurements and their reported errors, and the thin black line shows a fiducial Λ CDM model prediction for the same model plotted in Figure 1.2.

The results plotted in Figure 1.3 provide support for the paradigm developed to explain the CMB temperature spectrum. The TE correlation provides evidence that the perturbations were adiabatic, and measures the thickness of the last scattering surface. Furthermore, the super-horizon correlations measured in TT, TE and EE together lend support to (although by no means prove) inflationary models. Thus far, among the most interesting results provided by these polarization measurements is the detection of the signature of reionization at $\ell < 10$. This signal is consistent with the first generation of stars just beginning to reionize the inter-galactic medium at $z \sim 20$.

The milieu embodied by Figure 1.3 is the one in which QUaD finds itself. Although some measurements of E-mode power exist, they do not constrain the precise shape of the power spectrum, particularly at large ℓ . In addition, constraints on B-mode power exist, but allow a great deal of parameter space for models of inflation and structure formation. Based on these observations of the state of the art in the field, the next section outlines QUaD's science goals for a three year observation campaign from the south pole.

1.3 QUAD SCIENCE

Based on the sheer amount of information available in CMB anisotropy polarization, an experiment like QUaD was the natural next step in the progression of CMB polarimeters after DASI. In many ways, QUaD is a 'pathfinder' instrument: the level of *E*-mode polarization may be known, but the exact shape and amplitude of the signal remains unmeasured. A precise measurement of C_{ℓ}^{EE} would allow a definitive check of predictions of the standard cosmological model against the real data. Furthermore, *B*-mode polarization is a total unknown; although, based on predictions of the strength of this signal, QUaD will only place upper limits on C_{ℓ}^{BB} power, these will constrain *B*-mode parameter space to a much higher degree than ever before. Because of the unknown properties of the polarized foregrounds, the precise shape of the C_{ℓ}^{EE} spectrum, and the level of *B*-mode power at large ℓ , QUaD's data will also be instrumental in pointing the way forward for the next generation of CMB experiments.

QUaD's science goals are the following:

- Q and U polarization maps with signal to noise ratios unrivaled by any other experiment of QUaD's generation. These would permit unparalleled consistency checks of the standard CMB paradigm and a measurement of the foregrounds at high galactic latitudes.
- A measurement of C_{ℓ}^{TT} competitive with the best current CMB experiments for $100 \le \ell \le 2000$.
- An exquisite measurement of the E-mode power spectrum. Due to QUaD's capacity for fast surveys, the minimum inter-ℓ distance in the power spectrum is very small this allows a full characterization of the *shape* of the power, not just the amplitudes of the peaks. The constraints on E-mode polarization for ℓ > 500 will be significantly better than for any other experiment in the near future and QUaD data will go to very small angular scales, probably ℓ ~ 2000 in its final data release.
- The detection of *B*-mode polarization of gravitational lensing, should the signature of structure in the Universe be large enough. If not, the best upper limit yet available will be placed on this power. These measurements will constrain parameters like the neutrino mass(es) and dark energy.

• Understanding how *B*-mode power from lensing will affect primordial *B*-mode power from recombination. This will be particularly powerful in conjunction with experiments like BICEP, a complementary sister experiment to QUaD (Keating et al., 2003).

These goals will be met using a survey area of about 100 deg^2 , integrating over 3 austral winter seasons (from about March to October). At the end of this time, the next generation of CMB polarimeters is due to come on line, and presumably QUaD will make way for the next upgrade to the DASI mount.

The work presented in this thesis deals with QUaD's construction, integration, and observations at the south pole site, leading to a discussion of results from the first year of observations, and expected future direction for the second and third years. I joined the project in October 2003; at this point, the designs for the instrument had been finalized, and components were undergoing fabrication and testing (the final system is described in Chapter 2). My first effort on QUaD was to write and run instrument simulations to understand the mapping algorithm and effects of noise on the observation strategy. Following this, I helped commission the cryostat and cryogenic systems in Cardiff, and helped with the telescope and optical systems. In the spring of 2004, I followed the cryostat to Stanford University, where I spent 6 months helping to integrate and commission the receiver; the results of this work comprise the bulk of Chapter 3. I did not follow the receiver to the south pole in the austral summer of 2004/2005, but stayed in the northern hemisphere, analyzing commissioning data and planning observations. Following commissioning, I played a crucial role in characterizing the performance of the fielded instrument (Chapter 4) and developing and testing calibration schemes (Chapter 5). This lead to my assisting our collaborators around the world with the work discussed in Chapter 6, including mapping algorithms, observation strategies, and analysis techniques. This thesis ends with a discussion of power spectrum estimation and science extraction, which forms the centre piece of the QUaD experiment.

2 QUAD INSTRUMENT DESIGN

QUaD is a millimetre wave polarimeter using cryogenically cooled polarization sensitive bolometers and a 2.6 m diameter primary mirror Cassegrain telescope. Because QUaD is the combination of two experiments (the telescope and receiver QUEST, which stands for the Q and U Extragalactic Sub-mm Telescope, and the mount from DASI, the Degree Angular Scale Interferometer), the acronym QUaD itself stands for QUEST at DASI. This serendipitous arrangement (*i.e.* combining a pre-existing mount, control system and electronics in an excellent site with a receiver and telescope in need of such components) was arrived at by the members of the two projects sometime in the Spring of 2003. QUEST was built over the next two years, and QUaD was finally commissioned in the winter of 2004/2005. QUaD has now taken a full season of data, is in the middle of a second austral winter of observation with a third year planned.

This chapter describes the QUaD instrument, and details its design. The mount, telescope, and control electronics/software are described, followed by details of the receiver's cryogenic systems. This is followed by an explanation of the system's optical design. A description of the bolometric detectors and readout electronics closes the chapter.

2.1 MOUNT AND TELESCOPE

Of primary concern to ground based millimetric experiments is the quality of the site at which the instrument is located. Atmospheric stability on order a few seconds or more is required for high quality CMB observations, and only a few high, dry sites achieve these conditions. The south pole is possibly the best site for millimetric observations in the world, even though its altitude of 2800 m is lower than that of the Atacama desert or Mauna Kea (Peterson et al., 2003). The low atmospheric opacity and ambient temperature mean that conditions which are stable over long time scales are common at the pole. As the precipitable water vapour is exceptionally low even in poor polar conditions, the pole's low

sky noise can be used to great advantage. The ACBAR experiment measured a Kolmogorov median fluctuation power amplitude of $10 \text{ mK}^2 \text{ rad}^{-5/3}$ in the austral winter, which is at least 30 times better than the next best astronomical site in the world (the Atacama desert) (Bussmann et al., 2005). The south pole does not experience diurnal fluctuation *per se*, and the sun is completely absent for half of the year, which makes for very stable observing conditions.

QUaD is situated atop a tower which is attached to the Martin A. Pomerantz Observatory (MAPO) at the United States Antarctic Program's South Pole station. The QUaD collaborators at the University of Chicago were responsible for the construction of the infrastructure at MAPO long before QUEST was even proposed, and provided most of the equipment available there from the computer systems, to the external electronics and testing apparatus, to the telescope mount itself. MAPO is approximately 0.7 km from the pole itself, in a section of the station designated as a low light pollution and radio quiet zone. An outer tower surrounds the inner tower to provide reasonable working temperatures inside the telescope, and the inner and outer towers are mechanically isolated from one another to minimize the transmission of vibrations from the main building. The legs of the inner tower are thermally insulated to prevent contraction under different ambient temperature conditions. The interior of the telescope is connected to MAPO via a 'compressor room', so access to the drive systems, electronics and receiver is available year round (see Fig. 2.1). This area houses the telescope control computer, data acquisition cards, and interfaces to external computers in a VME-card mainframe. An insulated fabric bellow insulates the telescope cab area from the outside while allowing full movement of the telescope's elevation drive. This makes it possible for the drives and electronics to operate above freezing temperatures using only waste heat from the building, even when the ambient outside temperature is below -70 C. The telescope mount is an altitude-azimuth design, and uses a counter weighted elevation drive to minimize the power consumed while driving and pointing the telescope. QUaD has a further degree of freedom, the 'deck', z, or θ axis (the latter is used throughout this work). This axis is situated just below the point where the primary mirror and cryostat are attached to the mount, and allows the telescope to completely rotate its polarization characteristics with respect to the sky. Heavy box steel construction makes the mount very stable and immune to flexure under load.

Reflective ground shields are required to prevent the instrument from picking up thermal emission from the ground and polar station structures (Fig. 2.2). These are attached to the outer tower, and extend to an angle of approximately 29° above the horizon as seen from the receiver. The geometry of these shields is designed so that any part of the illumination pattern of the telescope which falls on them is reflected to the sky. The snow that accumulates in the joints between panels is swept clear every few days in the winter.





This plot shows a cross section of the QUaD instrument. Marked are the primary and secondary mirrors, the secondary support structure, the receiver and electronics, and the various axes of motion.

The telescope itself is composed of a 2640 mm diameter primary mirror and a 450 mm diameter secondary mirror. The design calls for these to be configured according to the parabolic primary/hyperbolic secondary arrangement of the classical Cassegrain. After delivery of the primary mirror from the manufacturer, it became obvious that the mirror's surface was warped. When the nominal shape of the mirror is subtracted from its measured shape, the residual exhibits an obvious 'saddle' shape with a peak to peak variation of $300 \,\mu\text{m}^1$. Although it is unclear what caused this warp, it certainly has an effect on the telescope's performance, and causes many of the effects discussed in Section 4.4. The solution to this problem has been to remake the secondary mirror with a customized

¹Because this is a significant fraction of a wavelength, it will significantly warp the phase front of the incoming radiation.

surface designed to counteract the effect of the warp in the primary mirror, although this was not implemented in the 2005 observing season.

The secondary mirror is supported by a conical foam structure, similar to that used on the COMPASS experiment (Farese et al., 2004). The 35 mm propozote PPA 30 foam should be up to 99.5 % transparent in the millimetric range; this solution is chosen so that the polarization symmetry of the instrument is not broken. The primary is backed with an aluminium frame welded on at strategic positions, and the whole primary mirror assembly is mounted on the telescope with a ring that adapts from the DASI mounting points to the ones on the primary. The foam cone is mounted into a fibre glass flange and secured with bolts. Outside this mounting ring, a skirt of aluminum approximately 40 cm wide is used to ensure any possible beam spill over from the secondary mirror is projected onto the sky and not the ground shield surrounding the telescope.

A small optical camera is attached to the mount in close proximity to the primary mirror, outside of the skirt. This camera is used to image stars in the field of the main telescope for pointing measurements. This 'star camera' is coupled to the telescope control system and saves the offsets of the pointing guide stars during each optical pointing measurement.

The secondary mirror is held in place by a support structure mounted on the foam cone with a fiber glass mounting ring. This structure is attached to the mirror itself with rods that permit the adjustment of the z distance from the primary mirror surface so that the instrument can be focused. The secondary mirror itself is made of an aluminum sheet covering a carbon fiber backing, and has a 40 mm diameter hole in its centre; this allows standing modes between the secondary and detector block to escape the optical system. A thermal source with a rotating polarizer is mounted behind the secondary mirror and can be used as an external calibration source. A flip mirror is used to couple the emission from this source to the telescope on demand. During a calibration observation, a command is sent via an IR link to rotate the polarizing grid and insert the flip mirror into the telescope's beam. This modulated signal is observed for some length of time, and then the flip mirror is parked in its off position. The analysis of this kind of data is discussed in Chapter 5.

2.2 CRYOGENICS

QUaD uses a combination of liquid cryogens and a ${}^{3}\text{He}{-}^{4}\text{He}$ sorption refrigerator to cool its bolometric detectors to around 250 mK². The cryogenic system (or *cryostat*) is arranged

 $^{^{2}}$ Excellent reviews of the thermodynamic principles employed in QUaD can be found in Pobell (1996) and White & Meeson (2002).



Figure 2.2: QUaD Panorama

A panoramic photograph of the QUaD instrument *in situ* at the south pole. Visible in the centre of the image is the telescope, with its prominent foam cone and secondary mirror support structure. The small cylinder on the lower left of the telescope is the pointing camera. The ground shield is composed of the aluminum–fronted plywood panels surrounding the telescope. MAPO itself can just be seen in the extreme lower right of the photograph. During CMB observations the ladders, snow, and people are removed from the inside of the ground shield. The geographic south pole itself is 700 m to the left of this image.

in layers, each colder than the last, to minimize the heat load on each component. Thermodynamically speaking, heat is transferred by three processes: conduction, convection, and radiation. Convection is not a concern in QUaD's cryostat, and the radiative load from a cryogenic layer at a given temperature can be calculated from the Stefan-Boltzmann equation. The best way to minimize the effect of this radiation is to use a system of thermally isolated layers, where each layer is cooler than the one on its outer side, and warmer than the layer on its inner side. This system is employed in QUaD, and is quite effective.

Heat conduction is one of the biggest problems in any cryogenic system. The first major heat conductor is air; this can be dealt with by making the cryostat vacuum tight and pumping the air out. Gas pressures below 10^{-4} mbar are sufficient to ensure negligible heat conduction through otherwise thermally isolated components in a cryostat. A second important heat conduction path is through the materials that join the different temperature

stages in a cryostat. Choosing materials that have low heat conduction at cryogenic temperatures (usually various kinds of plastics like G10 or Vespel) and maximizing their length while minimizing their cross section is a good solution to this consideration³. Finally, the wires that service the detectors, sorption refrigerator unit, thermometry, et cetera, can be an important heat conduction path, particularly since good electrical conductors tend to be good heat conductors. Again, the solution here is to use thin gauge wires with long lengths between stages at different temperatures, and to ensure they are made of materials with low heat conduction and electrical resistance, like manganin or constantan.



Figure 2.3: The QUaD Cryostat

This figure shows the QUaD cryostat in detail. The image on the left is the cryostat in the lab. The drawing on the right is a cross sectional view of the cryostat, with the various stages and components discussed in the text highlighted. The different stages are colour-coded; blue shows the LN_2 stage, red the LHe stage, solid orange the intermediate stage; the cryogen tanks are highlighted in hatched orange fill.

QUaD's cryogenic strategy and an image of the cryostat itself are shown in Figure 2.3. From warmest to coldest, the cryostat's thermal stages are: an outer shell (the 300 K stage), the liquid nitrogen stage, the liquid helium stage, the 400 mK (intermediate) stage, and the 250 mK (cold) stage. The QUaD cryostat has a feature called a 'snout'; this structure allows cold optics to be positioned at the correct distances with respect to the detectors but does not require a large increase in the volume of the cryostat.

³This comes from the heat conduction equation $\partial T/\partial t = \kappa \partial^2 T/\partial \mathbf{x}^2$ with suitable boundary conditions.

The 300 K stage is essentially composed of an aluminum cylinder, with the aluminum 'snout' attached via an aluminum plate; the 300 K snout contains the cryostat window at its apex and juts above the surface of the primary mirror by about 40 cm when installed on the telescope. This upper plate also has the liquid nitrogen (LN_2) and liquid helium (LHe) fill and vent line holes. The bottom plate is solid stainless steel, and has the vacuum flange attachment and holes for the electronic connectors in it. It also has a 'man hole' plate that allows easy access to the science core; this allows the electronic cabling to remain installed when the detector block is removed from the cryostat.

The next thermal stage is the LN_2 stage. This consists of an annular liquid nitrogen tank with a 35ℓ volume that extends the height of the main body of the cryostat, with an extension into the snout. A set of filters is positioned at the top of the 77 K snout. The LN_2 tank is filled via one of two lines that extend to the outside (the interior of the tank is kept at or very slightly above ambient pressure). In operation, this stage runs at $\sim 77 \text{ K}^4$; it is attached to the 300 K stage via a set of G10 legs near the top of the main body.

Inside of this is the helium stage; this and the LN_2 stage are also connected by G10 struts. The LHe stage extends into the snout, which contains filters and lenses held at temperatures around 10 K. The helium tank has a 21 ℓ volume, and is filled/vented via pipes very similar to the ones for the LN_2 stage (these are longer for the LHe stage and are heat sunk to 77 K approximately half way down their length). The bottom end of the LHe stage is adapted for connecting to the science core base surface (called the 4 K plate). The 4 K plate has a number of functions; firstly, it is the attachment point for the Vespel legs which join the science core to the cryostat. It also hosts the JFET amplifier modules, as well as the electronic connections for the detector cables. The sorption cooler, discussed in detail below, is attached to the upper surface of the inner part of the LHe stage.

Attached to the 4 K plate is the science core, which is composed of the intermediate and ultra cold stages, including the focal plane. The science section is shown in Fig. 2.4. The intermediate stage is attached to the science core base plate via hollow Vespel legs, which are chosen because they have a very low heat conduction at T < 4 K. The sorption cooler allows access to a component which operated around 420 mK; this is connected to the intermediate stage base plate via a heat strap. Below 1 K, effects which aren't obvious at higher temperatures become important; for example, the dominant heat conduction path to this stage is actually the wires that service the detectors. It is thus desirable to have this buffer stage to reduce the thermal load on the cold stage. Attached to the intermediate stage base plate is the intermediate stage shield; this consists of a aluminum cylinder with a filter window at its top. This exists to reduce the radiative load on the cold stage, which is not a

⁴Due to the reduced ambient pressure at the south pole, the operating temperature of this stage at the pole is closer to 72 K.



Figure 2.4: The Science Core

This figure shows a schematic view and photograph of the QUaD science core. The schematic is labeled, showing the various components in the science core. Here, the component called the 400 mK ring is referred to as the intermediate stage base plate in the text. The photograph was taken during commissioning (June 28, 2004), and shows the science core in a 'blanked' state, meaning that the filters at the top of the feed horns were replaced with caps so that photons external to the feed horns were not incident on the detectors. The load resistor boxes are not visible in the photograph as they are shielded with an aluminum cylinder. Also not shown is the 400 mK baffle which attaches to the intermediate stage base plate. (Figure courtesy J. Hinderks.)

negligible effect at these temperatures.

Finally, attached to the intermediate stage via 30 mm long, hollow, carbon-loaded Vespel legs is the cold stage. This stage contains the load resistors for the detectors and the focal plane itself. It is connected to the cold head of the sorption refrigerator, and runs around 260 mK during operation.

The sorption cooler itself was manufactured by Chase Research Cryogenics; detailed discussions of similar systems can be found in Bhatia et al. (2000) or Teleberg (2006). The fundamental physical principle behind these refrigerators is that when one lowers the vapour pressure of a liquid/gas mixture by removing the gas, the temperature of the liquid falls. In a sorption cooler of this design, this is achieved by removing ³He or ⁴He atoms from the gas state by absorbing them into activated charcoal. A schematic of how this is achieved in practice is shown in Fig. 2.5.

A sorption cooler is operated by alternatively heating and cooling the absorbing material so that the pressure in the refrigerator varies depending on whether it is desirable to condense or evaporate the cryogenic liquid. Consider the case of a ⁴He stage. To begin with, the ⁴He is completely absorbed in the charcoal. To start the cycle, the charcoal



Figure 2.5: Schematic of a Sorption Refrigerator

This figure shows a schematic representation of a single stage sorption cooler (see text for full details of the use of this device). The cycle begins with the helium completely absorbed into the charcoal in the pump. This charcoal is heated, and the gas is driven off. It condenses in the condenser, and the liquid falls into the evaporator under the force of gravity. When the gas in the pump is completely driven off the charcoal and either in the evaporator or in the gas state, the gas gap heat switch is closed. This cools the charcoal, which then begins to absorb the helium in the gas state. This lowers the pressure in the chamber, which drives the temperature of the liquid in the evaporator down. The refrigerator stays cold until the liquid in the evaporator is completely exhausted.

is heated, which drives off the ⁴He gas: this process begins above about 10 K and the absorbing material is usually heated above about 25 K to ensure that most of the ⁴He is in the gas state. The pressure in the cooler vessel increases until liquid begins to condense in the condenser. This liquid falls into the evaporator under the force of gravity and remains liquefied there. The process continues until the pressure in the vessel is such that it is no longer energetically preferable to condense liquid. At this point, a gas-gap heat switch is used to thermally connect the absorber to the 4 K base plate, which cools the absorbing material back below 10 K. The ⁴He remaining in the gas state then absorbs into the material; this is why the main absorber is referred to as the 'pump'. As ⁴He evaporates from the liquid in the evaporator, it is absorbed, and the temperature of the evaporator stage (and anything attached to it) falls. This process continues until the liquid remaining in the evaporator is exhausted.

The lowest achievable temperature of this type of refrigerator is dictated by four

factors: (i) the total heat load attached to the evaporator, (ii) the pumping speed of the system, (iii) the Kapitza resistance⁵ of the evaporator and (iv) the vapour pressure of the liquid. Item (i) depends on what the experimental requirements are, and items (ii) and (iii) are addressed at the refrigerator design stage⁶. Item (iv) can be changed by using a cryogen with a lower vapour pressure, like ³He, which can reach base temperatures < 800 mK. However, because ³He does not liquefy at 4.2 K, the condenser for a ³He refrigerator must be kept at a lower base temperature. The standard ways to address this are to either pump on the helium tank thermally connected to the 4 K base plate thus cooling the heat sink to ~ 2.6 K, or to use a ⁴He stage to cool the ³He condenser before cycling the ³He stage. This leads to double, triple, and possibly higher stage fridges like those employed in QUaD.

The QUaD triple stage sorption cooler uses a ⁴He stage, a primary ³He (or intermediate) stage, and a secondary ³He (or ultra–cold) stage, and is shown in Fig. 2.6. Using the basic cycle described above chain–wise, the ⁴He stage is used to cool the condenser of the intermediate stage to below 2 K, and this in turn is used to cool the ultra–cold condenser below 1 K. The advantage of this, in the case of QUaD, is two–fold. First, it increases the efficiency of ³He condensation in the ultra–cold stage, thereby increasing the refrigerator's hold time (*i.e.* the time between cycles). Secondly, the intermediate stage reaches a base temperature of about 420 mK, and this is used to cool a baffle and filter around the focal plane, which reduces the radiative load on the cold stage. In practice, the ³He in the inter-cooler evaporator runs out before that in the ultra–cold stage. Fig. 2.7 shows a typical QUaD cooler cycle as employed during winter observations.

The intermediate and ultra-cold stages are attached to the thermal mass of their science core counterparts by means of copper heat straps. In the case of the intermediate stage, this thermal strap is composed of braided copper wires with a total cross section of approximately 1 cm^2 . For the ultra-cold stage, the thermal strap has two parts; a multiply-braided copper strap with cross section of about 1.5 cm^2 and a copper bar with cross section of about 0.75 cm^2 . The copper of both heat straps is annealed to improve its thermal conduction properties. Because in metals at low temperatures heat transfer occurs primarily via electronic conduction, it is important to clamp the various metal-metal interfaces together tightly, and to keep the contact surfaces free of impurities. The first of these considerations

⁵The Kapitza, or boundary, resistance is the thermal resistance between liquid helium and the solid wall of the heat exchanger. At very low temperatures, the dominant thermal energy transfer carriers are phonons. Kapitza resistance arises because an acoustic mismatch between the helium and exchanger material causes energy transfer over the phonon boundary to be inefficient (for a great deal more on the subject, see Pobell (1996)).

⁶A great many other considerations go into sorption cooler design; impurities in the gas can block the capillary lines, super fluid He can leak out of the evaporator, Kapitza resistance can be very high for poor evaporator geometries, et cetera. In practice, this leads to somewhat arcane cooler designs, as exemplified in Fig. 2.6.



Figure 2.6: Photograph of the QUaD Refrigerator

This is an image of the actual sorption cooler used in QUaD, which was manufactured by Chase Research Cryogenics. This model is called 'Helium-10', because it employs three stages (${}^{4}\text{He}-{}^{3}\text{He}-{}^{3}\text{He}$), each cooling the next, to achieve temperatures below 300 mK. Also shown in this image are the pre-cooling gas gap heat switches (long steel cylinders at extreme left and horizontally attached to the intermediate stage block), and the pre-cooling heat links (copper straps running along the top of the system). This photograph captures the state of the refrigerator as used at the south pole.

is achieved my means of a large number of screws, and the second by gold plating the thermally interacting components before installation. Thermometry installed on the science core shows that this strategy is quite successful at thermally connecting the focal plane and intermediate stage to the fridge.

2.3 Optics

QUaD's optics are designed to couple the radiation incident on the telescope primary mirror to the detectors as efficiently as possible whist retaining the polarization state of the incoming photons. The first two components in the optical chain, the primary and secondary mirrors, have already been discussed in Section 2.1. In this section, we describe the optics between the secondary mirror and the detectors; a very detailed discussion can be

CHAPTER 2. QUAD INSTRUMENT DESIGN



Figure 2.7: Typical QUaD Refrigerator Cycle

This plot shows an actual cooling cycle of the sorption cooler used in QUaD (in this case from July 20, 2005, *in situ* at the south pole). The upper plot shows the temperature of the various components attached to the 4 K stage; these are used to drive the fridge pumping cycles. The lower plot shows the temperature of the cold components of the fridge that result from its operation. During regular observation the fridge is either cold or being cycled, so the temperatures just after t = 0 on the plot reflect the final thermal state reached using the cycle which occurred the day before. The cycle QUaD uses is actually composed of two separate coolings; the first two pump cycles are used to pre-cool the fridge and detector block, and the second three actually cool the components to their operating temperatures. In total, the cycle takes approximately 5 hours to go from ultra cold stage base temperature through the cycle back to ultra cold base temperature. The fridge hold time is approximately 24 hours under real operating conditions.

40

found in O'Sullivan (2006). The optics are schematically shown in Fig. 2.8, and the frequency filtering scheme is shown in Fig. 2.9. The optical system was designed by QUaD collaborators at Cardiff University, NUI Maynooth, and Stanford University. Cardiff was responsible for overseeing the manufacture of the large optical elements, while Stanford collaborators manufactured the focal plane and its associated components. The group lead by P. Ade at Cardiff University designed and manufactured the filters as well.



Figure 2.8: QUaD's Optical Chain

This is a schematic of QUaD's optical chain, showing the large elements before the focal plane block. The components are artificially colour-coded for clarity. The primary mirror is shown in green, the secondary in blue, and the baffle is shown as grey. The other components are inside the cryostat, and are the pink camera lens, the yellow field lens, and the orange cold stop. The focal plane, which is composed of additional optical elements (not shown), is white.

Following the secondary mirror in the optical chain is the window-baffle system. The window is made of approximately 5 mm thick anti-reflection coated high density polyethylene (HDPE), which is transparent at mm wavelengths, yet holds vacuum even at large diameters ($d \approx 210$ mm in QUaD). The baffling system is a somewhat more complex issue as discussed in Section 4.5 and Appendix A.3, but its function can be simply described. The hole in the secondary mirror, which is used to reject standing-mode rays between the secondary and the detectors, is not large enough to ensure all rays from the central pixels map onto the sky rather than the edge of the cryostat⁷. The baffle's function is to reject rays that, in the baffle's absence, would fall on the cryostat by reflecting them out of the optical path to the sky. Following the window on the inside of the cryostat is the first of the optical filters, which block some of the IR light in the optical path.



Figure 2.9: QUaD's Optical Filtering

This figure schematizes the optical elements and filters in the QUaD cryostat. The various thermal stages are colour-coded, with blue as the warmest components, and red as the coldest.

Following this are the optical elements in the 77 K snout. Because these components are far from the cryogenic tanks and a great deal of optical power is dissipated in them, they actually operate closer to 90 K. These elements are composed of a set of three filters; an IR blocker, the 12 cm^{-1} edge, and a second IR blocker (see Fig. 2.9). This system rejects all of the IR power which would otherwise be incident on the detectors. The next part of the optical chain are the 4 K snout optics. At the top of the 4 K snout are a 9 cm^{-1} edge filter

⁷Were the hole in the secondary large enough to reject all these rays for the central pixels, as much as 1/4 of the power falling on the central pixel would be lost. The current secondary mirror design loses about 15% of the power through the hole in the secondary mirror.

and an 8 cm^{-1} edge filter. These are closely followed by the camera lens, which is a highly curved, anti-reflection coated lens composed of HDPE. It actually operates at a temperature close to 7 K due to its distance from the LHe tank. At the bottom of the snout is the field lens, which is again a strongly curved piece of anti-reflection coated HDPE. Below this is the cold stop, which is a blackened⁸, shaped optical stop. The purpose of this component is to cleanly truncate the point spread function of the telescope at its edge, around the -20 dB level. This results in a clean beam, and reduced side lobe levels.

In the original plans for QUaD, a wideband rotating waveplate was positioned behind the cold stop. This position is chosen because the stop is at a pupil in the optical system, so the rays from the primary mirror are constricted to a minimum breadth near it. Unfortunately, it proved technically difficult to anti-reflection coat the bonded sapphire that composes such a waveplate, and the design was abandoned. However, the electronics and motors that would service such a component remain installed in the cryostat, and could be used should a waveplate become available.

Behind the stop is the 400 mK stage. As discussed above, this is baffled with an aluminum 'can' surrounding the ultra cold stage, with a 7 cm edge filter at the top of the baffle. This section operates around 425 mK. Behind this baffle is the focal plane itself, shown in Fig. 2.10.

The focal plane base is a gold plated aluminum bowl which is configured to accept attachments from the feed horns on the front (concave) side and the detectors on the back (convex) side. The feed horns are an integral part of the optical system, and are discussed in great detail in Murphy et al. (2005). QUaD uses corrugated feeds; these are chosen because they satisfy the strict constraints on frequency bandwidth (up to 30 %), cross polar leakage, and sidelobe levels QUaD requires. Specifically, the requirement for QUaD is that the beam illuminate the cold stop (and thus primary mirror) with an edge taper of -20 dB at 12.5° from the centre of the beam. The corrugations in the horns preserve the polarization characteristics of the incoming light, while the flare at the top of the horns allows them to have a shorter physical length than conical horns working at the same wavelength. The internal diameter of the horn reaches a minimum called the throat; this acts as a section of waveguide and has the effect of a high pass filter. Coupled with the band-defining low pass filters installed in the cap of each of the horns, only a single waveguide mode is allowed to propagate past the throat. The end of the horn flares out, and allows the radiation to couple to the bolometers efficiently. These horns were designed by collaborators at NUI Maynooth and manufactured using electroformed copper on aluminum mandrels by Thomas Keating Ltd. Figure 2.11 shows the design and filtering employed in a QUaD feed horn.

⁸The term 'blackened', in this context, refers to a coating of carbon–loaded stycast, which acts as a good approximation to a black–body at mm wavelengths (Bock, 1994).



Figure 2.10: Photograph of the QUaD Focal Plane

This is a photograph looking down on the focal plane assembly. The focal plane itself is the spherical section to which the feed horns and handles are attached. The uppermost band defining filters are visible at the top of the horns. The detector modules are attached to the back side of the focal plane, and are not visible in this image. (Photograph courtesy J. Hinderks.)

2.4 DETECTORS

Detector technology for light with wavelengths around 1 mm is a complex field; an excellent review of the technologies at play and their various attributes can be found in Richards & McCreight (2005). The traditional combination of heterodyne mixers and correlators used in interferometry can be used effectively up to about 100 GHz before sensitivity and correlator speed become prohibitive for CMB measurements. However, projects like ALMA are pushing this branch of technology to higher sensitivities and frequencies. Stressed photoconductors, like those used in the Spitzer Space Telescope, can be pushed to wavelengths as long as $200 \,\mu m$ (Ward-Thompson et al., 2005). The gap between, from about one hundred microns to a few millimetres, is the realm of the bolometer ((Richards, 1994)). QUaD uses an advanced generation of intrinsically polarization sensitive bolometers to detect CMB photons.



Figure 2.11: Horn Assemblies

This figure shows cross sectional schematics of the QUaD feed horn assemblies at both 100 and 150 GHz. The band defining filters at the top of the horn are labeled; in reality, these are fitted into a cap with an indium light–tight seal at the cap end, and a beryllium spring washer on the bottom end. These caps are then fitted onto the horn and sealed with aluminum tape; the beryllium washer ensures that, as the copper of the feed shrinks during cooling, the filters remain tightly fitted to the cap. The horns themselves are made of gold plated electroformed copper. The constriction toward the left of both horns is the throat of the feed, and acts as a high pass filter. The throat then flares out into the detector integrating cavity at far left. The vertical flanges shown in the schematic are used to attach the horn to the focal plane via three screws. (Drawing courtesy J. Hinderks.)

A bolometer is a thermal radiation detector that measures, in principle, all of the power incident on it, irrespective of frequency (thus the name 'bolometer'). All thermal detectors – including bolometers – use an absorber with heat capacity C whose function is to convert incoming photons to heat. The absorber is attached to a heat sink at temperature T_s via a link with thermal conductance G. If power P is incident on the thermal detector, the temperature of the absorber T_B increases as $dT_B/dt = P/C$ until the limiting temperature $T_B = T_s + P/G$ is reached. The time constant of the heating interaction is $\tau = C/G$; when the power is turned off, the absorber relaxes back to T_s with the same τ .

The bolometer was originally developed by Langley (Langley, 1881), and is composed of a thermal detector which uses an electrical resistance thermometer to measure the temperature of the absorber. The absorber must have a size similar to that of the wavelength of the photons of interest, have good absorption characteristics at these frequencies, and a low heat capacity. The substrate to which the absorber is attached must have good thermal conductivity and low heat capacity, and the thermal link to the heat sink must have a low heat capacity and a thermal conductance suitable for the experiment. The electrical resistor (called a *thermistor* in this application) must also have a low heat capacity, low electrical noise characteristics, and a well understood R-T relationship. All of these elements must be combined in such a way as to match the required sensitivities, instrumental set ups, and frequencies for a given experiment. In a classical bolometer the thermistor is also the absorber: QUaD actually utilizes 'composite' bolometers in which the absorber and thermistor are separate entities joined by a tight thermal contact.

2.4.1 A BOLOMETER MODEL

The practical operation of a bolometer is best understood by developing a simple model based on the prescriptions given above; such a model – which is very common in the literature – is briefly reviewed here. This model was developed in work like Jones (1953), and given its modern, non-equilibrium incarnation in the work of Mather (1982) and Mather (1984). Sudiwala et al. (2002) gives details of how such a detector can be characterized in practice. Figure 2.12 shows a schematic representation of the type of bolometer commonly found in astrophysical applications.

The model we will develop can describe Neutron Transmutation Doped (NTD) composite bolometers; this process produces large quantities of very homogeneous thermistor material (Lange et al., 1983). Given the definitions of heat capacity, conductance, et cetera discussed above, a bolometer in thermal equilibrium is described by

$$P + Q = \int_{T_{\rm s}}^{T_{\rm B}} G(T) \, dT \tag{2.1}$$

where G(T) is called the differential thermal conductivity, and reflects the fact that Gs in real materials change with temperature. Also, P is the electrical power dispersed in the detector (P = IV), and Q is the radiative power incident on the absorber. For semiconductor bolometers, G(T) generally follows a power law

$$G(T) = G_0 \left(\frac{T}{T_0}\right)^{\beta}$$
(2.2)



Figure 2.12: A Composite Bolometer Model

These two schematics show a simple model for the thermal (left) and electrical (right) operation of a bolometer. The thermal model involves an absorber coupled to a heat sink via a thermal link, while the electrical model is composed of three resistors in series, with the central resistor being the thermistor in the bolometer.

where T_0 is an arbitrary reference temperature. Using this model, Equation 2.1 becomes

$$P + Q = \frac{G_0}{T_0^{\beta}(\beta + 1)} \left(T_{\rm B}^{\beta + 1} - T_{\rm s}^{\beta + 1} \right).$$
(2.3)

Now, using the bolometer circuit shown in Fig. 2.12 and Ohm's law, we get the following equations:

$$I_{\rm bias} = \frac{V_{\rm bias} - V_{\rm bolo}}{2R_{\rm L}} \tag{2.4}$$

$$V_{\text{bolo}} = \frac{R(T)}{R(T) + 2R_{\text{L}}} V_{\text{bias}}$$
(2.5)

where I_{bias} is the bolometer bias current, V_{bias} is the associated bias voltage, R(T) is the resistance of the thermistor on the bolometer, R_{L} is the value of the load resistors, and V_{bolo} is the voltage loss over the bolometer. Finally, we need to relate thermistor resistance to temperature, which can be achieved using

$$R(T) = R_0 e^{\sqrt{\Delta/T}}.$$
(2.6)

This equation gives very good results for NTD bolometers. Here, R_0 and Δ are just parameters of the model.

We can also compute the noise performance of these devices. The noise equivalent

power (NEP) is defined as being the power required to give a unit signal to noise ratio for a 1 Hz detector bandwidth. Bolometer noise is usually composed of two uncorrelated components, photon noise and detector noise:

$$NEP^2 = NEP_{\rm pht}^2 + NEP_{\rm det}^2 \tag{2.7}$$

Modern bolometers are generally engineered so that the noise in the system is background limited, meaning that the detector noise is slightly less than the background photon noise. At millimetre wavelengths, photon noise comprises two roughly equal contributions

- shot noise from the Poissonian nature of the photon flux, and
- Bose noise from the fact that photons are Bose-Einstein particles and tend to bunch together.

Following Mather (1982), these effects sum as:

$$NEP_{\rm tot}^2 = \sum_i NEP_i^2 \tag{2.8}$$

where the sum runs over every component at a given temperature T_i , and

$$NEP_i^2 = \frac{A\Omega(kT_i)^5}{c^2h^3} \int_{\text{band}} dx \, \frac{\eta_i \epsilon_i x^4}{e^x - 1} \left(1 + \frac{\eta_i \epsilon_i}{e^x - 1} \right). \tag{2.9}$$

The integral runs over the band (this is usually defined by the filters in the optical system), x is again the dimensionless frequency $h\nu/kT_i$, η_i and ϵ_i are the optical efficiency and the emissivity of element i, and $A\Omega$ is the instrumental throughput. Under the assumption that the system is single-moded and in the Rayleigh–Jeans limit, this can be simplified to

$$NEP_{\rm pht}^2 \approx 2h\nu P_{\rm tot} + \frac{P_{\rm tot}^2}{\Delta\nu}$$
 (2.10)

where P_{tot} is the total incident power on the detector and $\Delta \nu$ is the optical bandwidth of the system (Lamarre, 1986).

The detector noise is usually dominated by two components;

- phonon noise from the fluctuation of phonons carrying heat out of the bolometer, and
- Johnson noise from the random motion of electrons in the current through the thermistor.

If the responsivity of the detector S is known (or assumed), the Johnson noise contribution is

$$NEP_{\rm Johnson}^2 = 4kTR/S^2 \tag{2.11}$$

and the phonon noise is

$$NEP_{\rm phonon}^2 = 4kGT_{\rm B}^2 \tag{2.12}$$

Together, these equations allow one to calculate most bolometer characteristics given a few model parameters.

2.4.2 POLARIZATION SENSITIVE BOLOMETERS AND QUAD

Polarimetry using bolometers is commonly performed using some variety of waveplate (for example, SCUBA (Greaves et al., 2003), or MAXIPOL (Johnson et al., 2003)), or a polarizing quasi-optical beam splitter working at the frequencies of interest (as in Benoît et al. 2004). These approaches both have limitations, namely a loss of sensitivity due to increased loading, and limited bandwidths due to limitations in the frequency transmission spectrum of available materials. A solution is to make an intrinsically polarization sensitive bolometer (PSB). These systems enjoy all of the best traits of bolometers, combined with polarization capability usually associated with coherent detectors (Jones et al., 2003).

The path to these detectors began with the development of 'Spider Web' composite bolometers (Bock et al., 1995), which have become the *de facto* standard bolometer in the millimetre community over the last decade⁹ (Yun et al., 2003). Based on an optimization of standard bolometric principles, these devices have NTD thermistors in the centre of silicon nitride legs, which give the device a very small filling factor and heat capacity. This design gives a very small cross section to cosmic rays, meaning that spurious spikes in the time series are minimized. Typical time constants for these bolometers are a few milliseconds.

The next step in the progression to the PSB is to create a grid of silicon nitride supports, and to metalize only one direction (*e.g.* north-south or east-west) of these supports. The thermistor is also moved from the centre of the grid to its edge; this allows unimpeded coupling of the incident radiation to the metalized grid. This system essentially forms a PSB; these have already been used successfully in BOOMERanG. Feed horn coupled bolometers, as in QUaD, are usually placed a distance $\lambda/4$ away from the back short; this ensures they are coupled to the incoming radiation most efficiently. Because a single PSB only interacts with one polarization vector at the end of a horn, two PSB modules can be placed at the end of a single horn. The difference in their signals gives the polarized component of the incoming radiation, and their sum gives the total power measurement. These devices are thus an elegant solution to the problem of combining intrinsic polarization sensitivity and advanced bolometer designs, and are used in QUaD due to their numerous advantages over other methods.

⁹These are due to be supplanted by Transition Edge Superconducting bolometers in the near future.



Figure 2.13: Bolometers in QUaD

Left: a photograph of a QUaD PSB module. Right: a schematic of one of QUaD's PSBs. The central grid shows the silicon nitride absorber material, metalized in one direction. The red rectangle represents the thermistor, which is serviced by the green wires. The edge of the grid is also thermally coupled to the thermistor to provide faster response, and to the blue casing material through sparsely spaced silicon nitride legs.

Figure 2.13 shows a schematic and image of the PSB modules used in QUaD. These devices are fabricated using photolithography from silicon nitride at the JPL Micro Devices Laboratory. The absorber is metalized by depositing a layer of titanium (~ 20 Å thick), followed by a layer of gold (~ 120 Å thick). The geometry and thickness of the grid was optimized by simulation and experiment at JPL. Non-metalized silicon nitride legs run orthogonally to the metalized ones for support. A thermally conducting ring surrounding the metalized grid helps couple the NTD thermistor to the absorber. This assembly is thermally connected to the bath by the electrical leads to the thermistor. Two additional leads are manufactured in these devices to allow adjustments to the thermal properties of the PSBs to be made; in QUaD this was never necessary. Two PSBs are placed into brass housings $60 \,\mu$ m apart so that they sample the same electric field; this arrangement maximizes the efficiency of QUaD's detectors.

A PSB's response to incoming radiation is usually modeled as a function of the gain of the detector g (*i.e.* the factor that converts from Kelvin to Volts), the polarization leakage of the detector ϵ , the polarization angle of the system α , and the incident Stokes parameters T, Q, U (Jones et al., 2003). For detector d (*i.e.* A or B) in horn h at time sample t, the voltage is modeled by

$$V_{h,d,t} = \frac{g_{h,d,t}}{2} \left[(1 + \epsilon_{h,d}) T(\theta_{h,d,t}, \phi_{h,d,t}) + (1 - \epsilon_{h,d}) \{ Q(\theta_{h,d,t}, \phi_{h,d,t}) \cos 2\alpha_{h,d,t} + U(\theta_{h,d,t}, \phi_{h,d,t}) \sin \alpha_{h,d,t} \} \right].$$
(2.13)

Here, θ and ϕ are generic positions on the projection of a sphere's inner surface; these are usually identified with the azimuthal or equatorial coordinate systems. This generic result becomes important later during the map making process, when determination of the true T, Q and U from measurements of V are desired. The map making also drives an interest in determining the parameters g and ϵ both absolutely and to as low an error as possible; measurements designed to characterize these parameters are the subject of later chapters.

For a single horn, it is common to form the sum and difference of the V(t) from the A and B detectors; these time series correspond to the (unpolarized) temperature and polarized signal, respectively. In the case of polarization, this correspondence can be easily seen if we assume that the g and ϵ for both of the detectors are equal, and that the two devices are perfectly orthogonal ($\alpha = \alpha_{\rm A} = \alpha_{\rm B} + \pi/2$):

$$V_{A,t} - V_{B,t} = \frac{g_{A,B,t}(1 - \epsilon_{A,B})}{2} [Q\cos(2\alpha_t) - Q\cos(2\alpha_t + \pi) + U\sin(2\alpha_t) - U\sin(2\alpha_t + \pi)]$$

= $g_{A,B,t}(1 - \epsilon_{A,B}) [Q\cos(2\alpha_t) + U\sin(2\alpha_t)]$
(2.14)

where the horn and angular dependencies have been suppressed. Since we can measure g, ϵ and α , Q and U can be reconstructed by using a variety of different α . Under the same assumptions as for Equation 2.14, the sum of the voltages of two horns becomes

$$V_{\rm A} + V_{\rm B} = g_{\rm A,B} (1 - \epsilon_{\rm A,B}) T \tag{2.15}$$

meaning that a clean measurement of unpolarized T alone is also possible using these devices. The assumptions made above can be relaxed so maps can be made using bolometers with different gains, polarization leakages, et cetera; this just leads to slightly more complex equations. However, T, Q and U can indeed be separated using this model given adequate knowledge of the instrumental parameters.

2.5 ELECTRONICS

The electronic systems in QUaD are divided up into three major groups; telescope control, data acquisition, and thermometry. There is a great deal of shared information between

these systems, and because a single software package deals with all three, they are all fed through the VME control computer at some point. The telescope control systems were discussed in Section 2.1; here we turn our attention to the other two systems, starting with thermometry.

2.5.1 THERMOMETRY & REFRIGERATOR CONTROL

There are three basic forms of thermometer on QUaD: diodes, germanium resistance thermometers (GRTs), and thermistors. Although the difference between them is transparent on an end-user basis, these each require separate electronics chains to read out properly. The simplest are the diodes. It is found that the voltage over many standard diodes increases as temperature decreases, and this can be calibrated using a known thermometer. Measuring the voltage drop at constant current across such a system gives a good thermometer that works between about 3 and 300 K. About 15 such devices are used in QUaD, from measuring the temperature at the camera lens to the temperature of the focal plane as it cools to liquid helium temperatures.

The second type of thermometer used on QUaD are GRTs (Lake Shore GT-200 series). These are again resistors whose temperature dependence can be calibrated. GRTs are useful between about 40 K and 100 mK, and as such are used to measure temperatures on the intermediate stage and focal plane itself. As heat loading is a real concern at the temperatures of these stages, GRTs are biased using a constant, low current in a bridge circuit. The resistance R of such devices is then found via

$$R = \sqrt{V^2 + V_{\phi}^2} / |I|$$
 (2.16)

where I is the (constant) current, V is the measured voltage drop over the GRT, and V_{ϕ} is the phase lag voltage. This R is then converted to K using a calibration look up table in software. Both the diodes and GRTs on QUaD are read out using a custom built cryogenic controller, which interfaces with the telescope control system via an Ethernet connection. The signal from the GRTs needs to be amplified; these amplifiers are contained as part of the cryostat's warm electronics package on its back end.

The third and most complex type of thermometer used on QUaD are the thermistors. These are, in essence, the same thermistors as are used in the bolometers themselves, and both the bolometers and thermistors share the same bias and readout (see Section 2.5.2 for a description). The thermistors used on QUaD are commercial products¹⁰ and are placed at strategic places around the focal plane. These devices are used to measure the temperature

¹⁰http://www.haller-beeman.com/thermis.htm.

Table 2.1: A breakdown of the uses of the 96 ADC channels in QUaD.		
Channel service	Number	Purpose
Light detectors	62	Channels containing the science data.
Dark detectors	8	Blanked off detectors for monitoring focal plane conditions.
Resistors	4	$10 \mathrm{M}\Omega$ resistors located on the focal plane for electronics characterization.
Shorts	8	Low impedance shorts; used to monitor JFET and warm amplifier noise.
Thermistors	3	Used for focal plane temperature monitoring and control.
Unused	11	

below about 500 mK. Thermistors are incredibly sensitive, and at these temperatures give resolutions of fractions of a mK.

The thermistors on QUaD are not actually used in an absolute sense, and were never calibrated to Kelvin. Rather, they are used as part of a feedback control loop designed to keep the focal plane at a constant temperature¹¹. The output of the thermistors is fed through to a SRS-SIM960 PID controller, and this is interfaced to the control system. The operator can define a set point about which the PID attempts to control using heaters on the focal plane (see Appendix A of Hinderks (2005) for the details of this system).

2.5.2 DETECTOR BIAS & READOUT

The electronics that control and read out QUaD's detectors are a complex topic and only a cursory review is performed here; the reader is directed to Hinderks (2005) for an exhaustive discussion from the QUaD collaboration's resident expert. The basic outline of the electronics chain is shown in Figure 2.14 below.

The electronics chain for QUaD begins with the bias generator; this system runs current (either AC or DC, depending on user input) through the bolometer. Each bolometer is connected in series with two load resistors, forming a bridge circuit. The signals from these bridge circuits are first amplified using a cold JFET system, which reduces the effects of pickup and interference. These signals then exit the cryostat and are amplified and demodulated by lockin amplifiers on the back end of the cryostat. The signals from these are then read into data acquisition cards in the VME computer that digitize the signal and pass it to the control computer for storage. In all, QUaD has 96 readout channels, which are used for servicing different components are per Table 2.5.2.

The bias generation circuit is used to pass either a constant DC or AC current

¹¹He-10 sorption coolers tend to drift by a few mK as the helium evaporates or under conditions of increased loading. This, in turn, affects their responsivity, which is extremely undesirable.



Reference freq.

Figure 2.14: Overview of the QUaD Electronics Chain

This is a highly simplified schematic of the QUaD electronics. The bias generator is used to run a current across the bridge circuit formed by the load resistors and the bolometer. This is then turned to a low impedance signal and amplified by the JFET buffers, and sent out of the cryostat to the lockin amplifier. This then demodulates the signal based on the reference signal from the bias generator, and amplifies the result. This signal is then digitized and stored using the ADC and control computer. (Figure courtesy J. Hinderks.)

through the load resistors and detectors themselves. In AC (science) mode, this circuit generates two sine waves $V_+(t) = V_0 \sin(\omega_0 t)$ and $V_-(t) = -V_0 \sin(\omega_0 t)$, where both the amplitude V_0 and frequency ω_0 are computer controlled. Because bias amplitude variations look like optical signals, the sine wave output must have an exceedingly stable V_0 . In QUaD's bias system, ω_0 can range from about 40 to 250 Hz to better than 0.1 Hz resolution and is set by a standard quartz crystal oscillator. The system first generates a 20 V peak to peak square wave of the desired frequency. This is then filtered using a bandpass filter to create a pure sine wave (nearly -40 dB at the third harmonic), still with a 20 V peak to peak amplitude. This is then scaled using multiplying DACs to provide waves with amplitudes in the range 0 to 15 V adjustable with 12-bit resolution. In theory, the bias for the 150 and 100 GHz channels can be adjusted separately, but in practice they are slaved together.

The load resistors are housed in boxes just below the focal plane and are kept at the same temperature as the detectors to reduce Johnson noise in the system. The wiring servicing the light-tight load resistor boxes (LRBs) is twisted in pairs, which reduces cross talk. The load resistors themselves have a value of $20 \pm 0.4 \text{ M}\Omega$ (J. Hinderks, private communication), and convert the constant voltage bias into a quasi-constant current. Each

LRB services 24 of the 96 channels, and are connected to the detectors via custom built twisted pair copper wires. This high impedance section of wiring is particularly sensitive to microphonic pickup, and is rigidly fixed to the back side of the focal plane using nylon thread and copper mesh. The twisted pairs are themselves grouped into pairs, and one of these quad units services a single detector pair. The back side of the focal plane is shown in Figure 2.15.

The cold JFET boxes are situated on the 4 K base plate beneath the focal plane, and the high impedance bolometer signals are carried through manganin twisted pair cables from the focal plane to these via Vespel struts. The JFETs buffer these high impedance signals so that the roughly 2 m of wiring downstream does not need to be tied down so carefully. The JFETs themselves do not work below about 50 K, and have a noise minimum closer to 120 K; they thus require heaters to keep them near their optimum operating temperature. There is therefore a trade off between situating the JFET modules as close to the detectors as possible, and minimizing the heat load on the colder sections of the cryostat. This design difficulty is solved by supporting the JFET membranes on their input/output wiring, which have a low thermal conductivity (~ 0.5μ W/K) due to their small cross section. The power dissipated in the modules is in fact sufficient to heat the membranes to the correct temperature, and the heaters are only used on start-up to 'jump start' the process. The connections to the modules are RF filtered to help reject this kind of interference. The modules used in QUaD were designed and manufactured at NASA JPL.

Following the JFET buffers, the signals are fed through the various temperature stages out of the cryostat via twisted pair cabling. The goal in this wiring is to maximize the wire length between the different cryogenic stages to reduce heat transmission. The signals exit the cryostat and enter the lockin amplifiers, which are attached directly to the back side of the cryostat. These amplifiers both amplify and demodulate the incoming signals. The first stage in this process is to amplify the incoming signals by a factor of 100. The signals are then passed through a relatively wide bandpass filter to suppress noise far from the bias frequency. This bandpass is composed of a 4-pole Butterworth low pass with a cutoff frequency of 475 Hz, followed by a 2-pole Butterworth high pass filter with a 2.8 Hz cutoff.

The signal is then rectified using a square wave demodulator. The reference for this comes from the bias circuit, which is connected directly to each lockin via an external BNC cable. Stray capacitance in the cryostat causes a phase delay in the output compared with this reference; the computer control system initially varies this delay to find where the maximum signal occurs. The system then locks itself to this frequency delay, and based on it the demodulator multiplies this signal output from the cryostat by +1 when the reference is positive and -1 when it is negative.

Following rectification, another 6-pole Butterworth low pass with $f_{3dB} = 20$ Hz is applied to the signals. This filter, which suppresses high frequency noise in the system, is the harshest in the system, and must be undone in later processing (see Section 4.7). At this point, the output could be stored; however, it would require ADC cards with greater than 16-bit resolution, which are relatively uncommon and expensive. QUaD's solution is to remove a DC offset from each of the bolometer signals, which is calculated and controlled by the computer system¹². This is possible because the sky signal modulation is less than a few tens of mV on top of a few V DC component. Performing this step provides room for an extra factor of 100 in gain. A final 30 Hz 2-pole Butterworth low pass filter is applied to remove any residual high frequency noise.

The lockin cards are also capable of determining which of three bias modes (DC bias, low gain, or high gain) is being used at a given time. DC bias mode uses a constant bias through the bolometers, and is used for characterization, as in load curves (discussed in Section 3.1.1). Low gain mode applies a post-demodulation amplification of 500, and is used both when determining the correct DC offset to subtract from the signals, and occasionally during science observation (see Section 6.4.1). Finally, high gain mode applies an amplification of 50,000, and is used for general science observation and instrument characterization, which is the subject of the next chapter.

¹²This DC component is in fact stored in the analysis pipeline, but has yet to be used for any practical data analysis.
3 RECEIVER CHARACTERIZATION

QUaD's receiver was built and characterized between the autumn of 2003 and the autumn of 2004. The cryostat was commissioned and cryogenically tested at Cardiff University from November 2003 until April 2004, at which time it was shipped to Stanford University. At Stanford, the electronics, detector block, and optical systems were installed in the cryostat and tested in both dark and light configurations until the system was shipped to the South Pole for field commissioning at the end of September 2004. The testing discussed in this chapter was performed in the laboratory both in Stanford and at the south pole and forms the basis of our current understanding of the instrument's properties. In this chapter, first bolometer and electronics testing are discussed, then the optical characterization is presented, and finally the polarization characterization which occurred in the lab is detailed¹. A full discussion of the characterization of the electronics chain *without* the detectors is not treated here, but can be found in Hinderks (2005).

3.1 BOLOMETER & ELECTRONICS CHARACTERIZATION

Once the cryogenic system was commissioned (that is, the refrigerator could cool the focal plane to about 260 mK for more than a few hours), testing of the bolometers and associated electronic components was performed. Because these tests check the intrinsic properties of the bolometric system, the cryostat aperture (and, in some cases, the feed horn apertures) may have been blanked off to reduce the incoming radiation load. Note that references to the bolometer model derived in Section 2.4.1 are made throughout this chapter.

¹It should be noted that there is no attempt at chronology here; the order of presentation of topics does not necessarily follow the order of the tests themselves.

3.1.1 LOAD CURVES

The load curve is the fundamental measurement used to derive the model parameters for a particular detector as presented in Section 2.4.1. This test involves stepping through a range of input bias voltages and measuring the output bolometer voltages. Load curves can be performed either with the bolometer blanked off (*i.e.* the only loading being a thermal photon background with $T \sim 250$ mK), or open to some external load. These different loading conditions allow measurement of different quantities of the system. Using Equation 2.4, we can plot output voltage versus bias current, and Equation 2.5 allows us to calculate the electrical resistance of the device. Both of these quantities are plotted in Figure 3.1 for both the optically blanked and open cases.



Figure 3.1: Load Curves

This figure shows load curves for dark and light detectors taken June 1, 2005. Upper left shows the I–V curve for light and dark bolometers at both of QUaD's frequencies using a linear scale. Although the devices are physically different so there is some intrinsic scatter in their parameters, we can see the effects of different loading temperatures on the detectors. The two dark detectors shown have the same incident radiation power, and because this Q is very low their turn–over points are much more pronounced than the light bolometers'. The Rayleigh-Jeans temperature from the atmosphere was approximately 5 K at 100 GHz and 14 K at 150 GHz when these data were taken; the effect of the difference in optical loading can be seen in the shape of the I–V curves for the light bolometers. The plot on the lower left is the same as the one above it, but uses a logarithmic scaling for the bias current axis. This highlights the I–V behaviour of the bolometers at low bias. The plot on the right is the I–R curve for these same data. They show the typical range of the QUaD bolometer resistances at a base temperature of about 260 mK and these bias currents.

The behaviour of the bolometers under different bias currents exhibit characteristic features. For small bias currents, the bolometers behave like ordinary electrical resistors. Increasing the bias causes self-heating, which causes deviations from Ohmic behaviour. For bolometers under large optical loading, a large amount of bias current must be applied to measure this effect (and in fact it may disappear entirely for sufficiently loaded detectors).

In the lab, load curves were taken under a variety of conditions. Primarily, dark (optically blocked) load curves were taken to determine the G_0 , Δ , R_0 and β parameters in the detector model. Other load curves were taken with different optical loads, for example, with eccosorb CV-3 covering the cryostat window. Another common load was a Styrofoam cooler with the same type of eccosorb at its bottom, filled with LN₂. Because the Styrofoam was found to be better than 99% transmissive, this acted as beam-filling source with temperature 77 K. A fan placed under the container prevented water condensation between the cold eccosorb and cryostat window (which would change the apparent loading). A liquid helium load of similar design was also used, although using a smaller container. These were used to calibrate the behaviour of the bolometers under different optical loading conditions.

The following method is used to calibrate the bolometer model parameters using the dark load curve data. The points at low bias current $(P \sim 0)$ are used to determine the bolometer's resistance at the operating temperature $(T \sim 260 \text{ mK})$. Equation 2.3 shows that under dark conditions (Q = 0 and P small), $T_B \approx T_s$. This means that measuring $R(T_B)$ at different T_s can yield R_0 and Δ . This is most easily achieved by transforming the bolometer model given in Equation 2.6 according to

$$x = 1/\sqrt{T} \tag{3.1}$$

$$y = \ln R. \tag{3.2}$$

This linearizes the model, allowing a simple linear fit to determine R_0 (via the y offset of the fit) and Δ (via the slope). An example of this method is shown in Figure 3.2. In this type of plot, deviation from the linear model indicates stray optical loading, particularly for the coldest (largest x) points.

Once R_0 and Δ are known, any load curve can be used to determine G_0 and β . Equations 2.5 and 2.6 are used to convert V_{bolo} into T_{bolo} and the power dissipated in the thermistor as a function of bias current can be found using

$$P = I_{\text{bias}} V_{\text{bolo}}.$$

$$P = I_{\text{bias}} V_{\text{bolo}}.$$
(3.3)

discussed in Section 4.6.1). This being said, measurement of the *absolute* incident power at a given time is very difficult because it requires very accurate knowledge of the base plate temperature and parameters of the bolometer.

3.1.2 MICROPHONICS

Microphonic noise is a catch-all term which refers to the noise induced in an electronic system due to vibration. In the case of QUaD, microphonic noise is caused by internal vibrations, like cryogens boiling, or external noise, like heavy machinery passing by the laboratory or vibration from the telescope mount as the telescope moves. This noise arises from two major contributions: current noise in the high impedance wiring, and fluctuations in the heat sink temperature (*i.e.* the thermal stability of the refrigerator and its coupling to the focal plane). The first of these is dealt with by securing the high impedance sections of the wiring to stable parts of the focal plane as discussed in Section 2.2. The second is difficult to combat, except by ensuring a stable fridge, thermal link to the focal plane, and employing active temperature control with a fast response time.

Nevertheless, some level of microphonic noise is an unavoidable demon. We tested for the level of this type of noise by tapping on the Dewar with a rubber tipped mallet to incite resonant response in the wiring. The output signals were analyzed and recorded using a spectral analyzer. During these tests, the bolometers were DC biased, which allowed the full bandwidth of the spectral analyzer to be utilized rather than the standard 20 Hz of bandwidth available to the demodulated signal modes. Figure 3.3 shows the result of these tests for four bolometers.

In this plot, correlated spikes in the power spectra are attributable to vibrations of the entire focal plane assembly rather than individual wires servicing the detectors. The prominent spikes at harmonics of 60 Hz are due to pickup from the AC power used at the south pole station. Similar tests performed with the cryostat installed on the telescope showed no significant deviation from the baseline response, proving that the DASI mount is very stable. Two major microphonic–free windows were identified: that around 200 Hz, and that near 100 Hz. Because of stray capacitance between the bolometers and JFET amplifiers, frequencies around 200 Hz are too high for stable biasing. A bias frequency of 110 Hz was therefore chosen for the 2005 and 2006 seasons, as this window was relatively free from microphonic noise but still far faster than either the bolometers' time constants or the 1/f knee of the system.





These plots show the results of the microphonics testing performed in the lab at the south pole on Dec. 17, 2004. The topmost plot (panel A) shows the baseline noise spectrum for four devices. The middle plot (panel B) shows the effect of banging on the side of the cryostat on these spectra *i.e.* the microphonic response of the system to an approximately δ -function impulse. The sampling frequency and bias frequencies are set to approximately 100 Hz as a result of tests like these (observations in the 2005 season were performed at $f_{\text{bias}} = 110$ Hz and $f_{\text{sampling}} = 100$ Hz). The lower plot (panel C) shows a closer view of the area around this frequency (as marked in yellow in panel B), showing the small scale features in this microphonic window. The large spikes at multiples of 60 Hz are pick-up from the AC power supply in MAPO.

64

3.1.3 RADIO FREQUENCY INTERFERENCE

The wires running into and out of the cryostat which service the detectors can act as an antenna, picking up radio frequency (RF) signals from the environment². Some of this RF power can dissipate in the bolometer's thermistor, heating the detector and producing noise. NTD germanium thermistors have a negative thermal coefficient, so this heating always appears as a signal with negative voltage; this is the same direction as real signal, and so is quite problematic. Given typical wire lengths of 1 m and knowing that half or full-wave antennæ couple to radiation most efficiently, we expect that frequencies around 1 GHz are the most problematic for the QUaD system. Unfortunately, a great deal of modern telecommunications equipment works near these frequencies.

To reduce the effect of RF pickup, all of the lines entering or leaving the cryostat are filtered using RF-suppressing D connectors. Another filtering occurs at the 4K stage, where all of the detector block service cables are passed through filtered micro-D connectors. Unfortunately, the section of wiring connecting the JFET box to the focal plane (approximately 10 cm in length) is still sensitive to RF entering through the cryostat optics. Because of concerns about the total impedance of the wiring at this stage, filtered connectors are not used.

The RF testing apparatus consisted of a dipolar radiator attached to a function generator working up to 15 GHz via a standard BNC cable. The function generator was set to a relatively high power, and the dipole antenna was swept around the outside of the cryostat. With the window blanked off, the system did not pick up any interference between $\sim 1 < \nu < 10$ GHz, proving that the cable filtering scheme worked well. However, with the window open, a great deal of pick-up was observed at these frequencies, particularly when the radiator was in the π steradians above the lip of the window. At lower frequencies (tens of MHz), no pick-up occurred. This frequency dependent behaviour is due to two effects:

- the snout acts as a waveguide for frequencies near 1 GHz, but severely attenuates longer wavelength radiation, and
- wavelengths longer than a few GHz do not efficiently couple to the relatively short length of wire between the focal plane and JFET box.

Because the south pole is relatively radio quiet, and the ground shield blocks a great deal of the solid angle from which RF interference could originate, the RF filtering used in QUaD is found to be adequate.

²Although the dark sector of the south pole station where MAPO is located is meant to be a radio-quiet sector, there is evidence for periodic RF interference in our data, making these RFI tests quite important.

3.2 OPTICAL CHARACTERIZATION

The battery of tests designed to characterize the optical properties of QUaD were extremely important in the later understanding of the data from the instrument. These tests are discussed here; each of them were performed with the detectors open to external light, although a variety of photon sources were used.

3.2.1 BAND DETERMINATION

Determining the band shapes (that is, the overall effect of the optical filtering discussed in Section 2.3) is important in determining the overall calibration of the instrument, as well as the relative contribution of both the CMB signal and foregrounds in either band. This becomes critical during the final science extraction, when the data from the two wavelengths are combined in various ways to remove, reduce or enhance either the signal or foreground emission (Bowden et al., 2004).

The instrument used to measure QUaD's bands is a millimetre wavelength Fourier transform spectrometer (FTS). This measurement was performed both in the Stanford lab and at the south pole prior to receiver installation; the data presented here are from the south pole runs, although the Stanford experimental set up and results are very similar. The spectrometer was of the interferometric polarizing Martin–Puplett type; Martin (1982) gives an exhaustive discussion of this type of instrument, which is a variation of the famous Michelson interferometer³.

Schematically, this type of interferometer is composed of a beam splitter and moving mirror arranged in such a way as to produce interference fringes at the instrument's output. Although one of the attractive features of the Martin–Puplett design is its dual input and output ports (Martin, 1982), only one of each was used for these measurements. The input source was a Styrofoam bucket containing eccosorb and liquid nitrogen which was chopped using a rotating chopper wheel. The wheel is divided into equal quadrants, which alternate between clear and aluminum–backed room temperature eccosorb. The wheel's motor is connected to a lockin amplifier which allows demodulation of the chopped signal. The cold source itself requires careful monitoring to ensure that it remains isothermal over a set of runs (this requirement is met by ensuring that the bucket contains a reasonable amount of cryogen). The interferometer itself is optically coupled to the receiver via flat mirrors which transfer the optical signal from the optical bench in through the cryostat's window.

³Martin (1982) includes a rather interesting discussion of using this design of experiment to measure the *frequency* spectrum of the CMB, which was a topical issue at the time and intimately related to the work presented here.

A Michelson-type interferometer works by scanning the central beam splitter back and forth, thus changing the optical path length of the two arms of the instrument. When the beams are recombined, this produces interference fringes which are imaged at detectors (in this case, the QUaD focal plane). The interferogram thus produced is the autocorrelation function of the amplitude of the electromagnetic field incident on the interferometer. This means that the power spectrum of the incident field can be determined by computing the Fourier transform of the interferogram. If the incident field is uniform, and the optical elements of the interferometer are chosen to have no effect on the spectrum of the radiation, then the spectral response thus computed gives the spectral response of the instrument. This response can be normalized to transmission ($0 \le T \le 1$) using knowledge of the optical efficiency of the instrument.

The procedure for this measurement is thus to record a series of interferograms using a stable source, and to take Fourier transforms of these interferograms to determine the summed effect of the optical filters (discussed in Section 2.3) on the spectral response of the detectors⁴. This is performed for the QUaD receiver, and the results for each of the detectors are plotted in Figure 3.4.

Although there is slight variation in the measured spectral responses between each feed horn, it is useful to compute an average bandpass for use in the analysis pipeline. This is formed simply by computing the average spectral response of the detectors at either frequency. These are both plotted in Figure 3.5 with a representative atmospheric transmission spectrum; this highlights how the choice of filtration is engineered to match windows in the atmosphere's opacity in the millimetre regime. The average bandpass is used whenever the frequency response of the system is required in a calculation during data analysis. As a rough approximation to the full spectral response, the band centre ν_0 and band width $\Delta \nu$ are calculated. The band centre is defined by

$$\nu_{0} = \frac{\int_{0}^{\infty} d\nu \,\nu f(\nu)}{\int_{0}^{\infty} d\nu \,f(\nu)}$$
(3.7)

where $f(\nu)$ is the effective bandpass computed above, and the band width is defined by the full width at half maximum (FWHM) of the average band. These quantities are both given in Table 3.2.1, and are useful when discussing the general characteristics of the receiver.

⁴In principle the detectors have their own spectral response, but this is assumed to be negligible compared with that of the filters.



Figure 3.4: QUaD Bands for Individual Horns

These plots show the spectral response for the individual detectors in each feed horn. Because the filtration is common to both detectors at the bottom of a feed, any difference between them is due to some difference in the spectral response of the detectors; it is reassuring that this is not apparent in any of the detectors. The data plotted in dashes are for horns with only one working detector in the 2005 season; these are not used in the polarization analysis. (Figure courtesy J. Hinderks.)



Figure 3.5: Average QUaD Bands with Atmospheric Transmission Spectrum These plots show QUaD's average bandpass in terms of absolute transmission per polarization at 100 in red and 150 GHz in blue (they have been multiplied by two for clarity). The black line shows the atmospheric spectrum for the south pole with $\tau_{150 \text{ GHz}} = 0.40$ using the model of P. Ade (private communication). The upper plot highlights the region of the spectrum immediately around the QUaD bands, while the lower plot shows the bands in the context of the entire mm/sub-mm regime.

Table 3.1: QUaD Band Definitions			
$ u_0$	(GHz)	Δu	
Nominal	Measured	(GHz)	(%)
100	94.5 ± 0.6	26.2 ± 2	28
150	149.5 ± 0.6	40.5 ± 2	27

3.2.2 OPTICAL EFFICIENCY

In the context of astronomical instruments, optical efficiency is defined to be the ratio of inband power detected by the bolometer to in-band power incident on the receiver. Because of losses in the optical chain (due to reflection and absorption in the optical elements), this ratio is always less than one. Today's millimetric instruments typically achieve about 40 % optical efficiency.

The total optical power (Q) falling on a PSB for a beam filling source is given by

$$Q = \frac{1}{2} \int_{\text{band}} d\nu f_{n}(\nu) B(T,\nu) A\Omega$$
(3.8)

where f_n is the normalized spectral response, $B(T, \nu)$ is the spectrum of the source (which is generally assumed to be Planckian), the factor $A\Omega$ is called the *throughput*⁵ of the system, and the integral is over the band pass. The factor 1/2 comes from the fact that a single PSB only measures one polarization, and thus for unpolarized light one half of the total incident power. For single moded feed horns such as QUaD uses, $A\Omega = \lambda^2$, so in the Rayleigh-Jeans limit we can write

$$Q = kT \int_{\text{band}} d\nu f_{n}(\nu) \tag{3.9}$$

Under this simplification, the optical efficiency can then be written as the transmission averaged over the band;

$$\eta = \frac{1}{\nu_2 - \nu_1} \int_{\nu_1}^{\nu_2} d\nu f_{\rm n}(\nu) \tag{3.10}$$

This is a general result; however, for PSBs, this definition makes η the optical efficiency of the detector to radiation polarized along its axis of sensitivity.

The best tool for measurement of optical efficiency is to use load curves; these can be used to calibrate the band sensitivities in power units. Equation 3.10 allows measurement of η from a single load curve taken using a black body load of known temperature; however, due to emission, absorption and reflection in the optical system, this is extremely difficult in practice. Using two load curves at different external loadings provides a way around these complexities, and can be used even without knowledge of the model parameters of the bolometer.

⁵Or alternatively, *etendue*.

This method requires two load curves using beam-filling blackbody external loads. The loads are held at temperatures T_1 and T_2 . It is useful to think of the load curves in terms of electrical resistance versus input bias power; this way, lines of constant resistance correspond to lines of constant bolometer temperature, and thus lines of constant total power P + Q. At one of these constants, the external load at lower temperature requires more electrical power to heat the bolometer to the same temperature. This means that the difference in optical power must be equal to the difference in electrical power

$$Q_2 - Q_1 = P_1 - P_2. (3.11)$$

Using Equation 3.8, we can relate the difference in power to difference in load temperatures via

$$\Delta Q_2 - \Delta Q_1 = k \int_{\text{band}} d\nu \left(B(T_2, \nu) - B(T_1, \nu) \right) f(\nu) A\Omega.$$
 (3.12)

This way, the constant k can be determined, and the transmission spectrum of the bolometer can be written as

$$f_{\rm n}(\nu) = k f(\nu) \tag{3.13}$$

which quite simply leads to η .

These tests were performed on QUaD. The two loads were beam-filling room temperature eccosorb, and eccosorb at nitrogen temperature. The surfaces between the blackbody eccosorb and the inside of the cryostat were kept free of water condensation. The results of this analysis are shown in Figure 3.6; typically, η around 30 % are achieved at 100 GHz and 35 % at 150 GHz.

3.2.3 **Responsivity**

The responsivity S of a bolometer is defined to be the change in bolometer voltage occurring for a (small) change in incident optical power under the assumption that the bolometer is in thermal equilibrium during the measurement. During real observations, this is determined from some calibration (as discussed in Chapter 5). However, the DC responsivity S_{∞} can also be determined in the lab. From the definition

$$S_{\rm DC} = \frac{dV_{\rm bolo}}{dQ} \tag{3.14}$$

and the bolometer model, the expression

$$S_{\rm DC} = \frac{R_{\rm L}}{2V_{\rm bolo}} \frac{z - R}{z + 2R_{\rm L}} \tag{3.15}$$

CHAPTER 3. RECEIVER CHARACTERIZATION



Figure 3.6: Optical Efficiency Histograms

This plot gives histograms of the measured value of η for the detectors used in QUaD during the 2005 observing season. The left hand plot shows the results for the 100 GHz detectors, while the right hand plot shows those for the 150 GHz detectors. In the right hand plot, the shaded histogram is for the central ring of bolometers, while the unshaded histogram are for the outer ring of 150 GHz bolometers. (Figure courtesy J. Hinderks.)

can be derived, as given in (Jones, 1953); Sudiwala et al. (2002) gives this expression in a variety of different forms. Here, we have defined the dynamical impedance by

$$z = \frac{dV_{\text{bolo}}}{dI_{\text{bias}}},\tag{3.16}$$

which is the derivative of the I–V curve at each data point for a particular device. Richards (1994) discusses a variety of methods to measure the responsivity of bolometers in the lab; we have used the method of Jones (1953). First, the DC I–V curve of a given device is measured for a range of bias voltages and currents around the operating point of the bolometer. Secondly, R and z are measured at a variety of bias points, and Equation 3.15 is used to determine S_{DC} at each point. This yields plots like that shown in Figure 3.7.

For QUaD, the S defined in 3.15 is always negative because the bolometers have a negative temperature coefficient of resistance. This means that greater input optical power warms the detector, which decreases the responsivity. In fact, very bright sources can saturate the detector, forcing the responsivity to very low values. In addition, it is not desirable to bias the bolometers near the peak in the responsivity curve because the derivative of the responsivity is maximized there. That is, small changes in the background loading affect



Figure 3.7: DC Responsivity

This plot shows the |S| as a function of I_{Bias} for the same four devices shown in Figure 3.1 computed using Equation 3.15. The effects of optical loading are evident; the dark devices, which have only 250 mK of optical loading incident on them, are almost twice as responsive as the light detectors, which in this case see the ≈ 10 K sky. In fact, it is clear that the 150 GHz light detector is slightly more loaded than the 100 GHz light detector. The QUaD detectors are biased at a current slightly greater than that at the peak of these curves (1.25 nA in observation mode).

the responsivity quite strongly near the responsivity maximum, making the system difficult to calibrate. To alleviate the effects of this, QUaD's bolometers are slightly 'over biased' during observations, meaning that the current bias is set slightly higher than the peak value.

3.2.4 IMPULSE RESPONSE & TIME CONSTANTS

Bolometers are thermal detectors, and as such they respond to any incident energy input, whether from photons or not. Cosmic rays can travel to the detectors, interact with a bolometer's web or thermistor, and thus heat the detector. This appears as signal in the data stream, and because cosmic rays are quite energetic, the signal thus produced is typically quite large. Because PSBs are designed to have a low cross section to cosmic rays, the event rate is about 0.5 events per channel per hour; because of this low rate, during observations the scans containing these events are simply rejected from the map making process. However, as cosmic ray events are essentially instantaneous, they are a very good approximation to a delta function impulse, and can be used to measure the transfer function of the instrument.

73

The expected transfer function comprises two separate contributions, the bolometer response function and the electronics response function. In Fourier space, the effective filter of a bolometer can be modeled as

$$H_{\rm b}(s) = \frac{1}{\tau_{\rm b}s + 1} \tag{3.17}$$

where τ_b is the bolometer's time constant, and $s = -i\omega$. The second term, that due to the electronics, is dominated by the 6-pole Butterworth filter:

$$H_{\rm f}(s) = \frac{1}{(\tau_{\rm f}^2 s^2 + A\tau_{\rm f} s + 1)(\tau_{\rm f}^2 s^2 + B\tau_{\rm f} s + 1)(\tau_{\rm f}^2 s^2 + C\tau_{\rm f} s + 1)}$$
(3.18)

where $\tau_{\rm f}$ is the time constant of the electronics filter (set to have a 3 dB point of 20 Hz) and the parameters A = 0.5176387, $B = \sqrt{2}$ and C = 1/A come from the definition of a 6-pole Butterworth filter. The overall transfer function of the system is the product of these two terms

$$H_{\rm tot} = H_{\rm b} H_{\rm f} \tag{3.19}$$

To determine the impulse response function, the inverse Laplace transform of H_{tot} is calculated. This can be performed either analytically or digitally, and leads to a series of sums of damped sine and cosine terms (Hinderks, 2005).

Cosmic ray hits can be fit to this function with free τ_b (and τ_f if desired) to check that the impulse response of the system matches that expected from theory. The result of such a test is shown in Figure 3.8; the model and data agree well. This is suggestive that QUaD's bolometer and electronic systems are working as anticipated.

In principle, it is possible to measure the bolometer time constants by determining τ_b based on a number of cosmic ray events. However, the impulse response is a relatively weak function of the bolometer time constant, and in practice this method and others developed to measure the bolometer time constant do not agree well (see Section 4.7.1). Only very slow bolometers can be reliably identified using this method.

3.3 POLARIZATION PROPERTIES

As QUaD's fundamental purpose is to measure polarization, characterization of the instrument's polarization properties is extremely important. To this end, a series of careful measurements were performed to measure these characteristics, both in the laboratory at Stanford University, and in the field at the south pole. This section first presents a discussion of the type of effects one expects based on our understanding of the instrument and the results of measurements made to characterize these effects in the real instrument.



Figure 3.8: Typical Impulse Response

This plot shows data from bolometer D100-2-A (a dark bolometer) from March 9, 2006, as well as a series of models derived using the invariant impulse response method on the QUaD transfer function Equation 3.19. The models have been calculated with $\tau_{\rm b} \in \{0, 10, 20, 30\}$ ms, as indicated on the plot. The data are best fit by a model with $\tau_{\rm b} = 8$ ms; this is typical of a 'good' bolometer.

In CMB polarimetry, there are four major systematic errors induced by the instrument which must be understood and characterized in order to properly convert the detected signal back to the radiation which was initially incident on the telescope. These effects are the following (see Johnson, B. R. et al. 2006 for a detailed discussion):

- 1. Instrumental Polarization in Emission⁶ (I. P. E.). This is polarization due to emission from the telescope. For example, if the primary mirror is constructed of aluminum, the material itself emits slightly ($\sim 1 \%$) polarized radiation due to correlations in the metal's structure emplaced when it cooled at the factory. I. P. E. has the effect of systematically increasing the observed intensity of one polarization state (say Q) over the other (U).
- 2. Instrumental Polarization in Transmission (I. P. T.). This effect is due to diffraction, oblique reflection, and birefringence in the telescope and optics, and converts unpolarized incident radiation into (slightly) polarized radiation at the detectors. This effect is particularly problematic because it converts T on the sky into Q and U (and

⁶Here, we differentiate between instrumental polarization in emission and instrumental polarization in transmission: although these effects both a generally considered under the broad term 'instrumental polarization', they actually have different effects on the observed signal in the time stream.

possibly V depending on whether the interaction with the instrument is phase preserving) in a way that makes it appear as real signal on the sky.

- 3. Cross Polarization. This is an effect where power from one polarization state is converted to the orthogonal polarization, *i.e.* Q ↔ U; this is extremely problematic because Q ↔ U mixing leads to E ↔ B mixing, which corrupts the final power spectra⁷. Typical sources of this kind of effect are the horns and optics.
- 4. Depolarization⁸. This effect causes a loss in the measured Q and U intensity, which means that the polarimeter has slightly depolarized the incident radiation.

Although these effects can have insidious repercussions on the polarization measurement, they can be dealt with given the proper measurements and instrumental characterizations.

The first effect, I. P. E., can be simply dealt with by insisting that the telescope and optics have constant emission. In this case, the signal from I. P. E. appears as a constant DC term in the time series, and is simply modulated away as part of the AC coupling of the system. For QUaD, thermometers placed on the back of the primary, behind the secondary, and throughout the cryostat show that the instrument itself exhibits very little temperature variation on long time scales (about 10 K peak to peak over a year). This, coupled with the low emissivity of the optical components, guarantees that the variation in the I. P. E. signal is low.

The second effect, I. P. T., is more problematic. Because it turns T into Q and U in a scan-synchronous way, it is impossible to disentangle from true signal on the sky, and thus will corrupt measurements of E and B. To characterize this effect, one would take a perfectly unpolarized source in the far field, and observe it with the telescope at various θ angles. This measurement would allow one to disentangle the polarization (in transmission) of the telescope and optics from the other effects mentioned above. In QUaD's case, this measurement is difficult because no perfectly unpolarized far field source exists (for example, RCW 38 may be slightly polarized at the $\sim 1 \%$ level). Based on the results presented in Section 6.2, it appears as though this effect is not an obvious contaminant, even though it may indeed adversely affect our measurements of E and B.

In PSB-based systems, items (3) and (4) above can be shown to be equivalent. That is, the algebra of cross polarization response leads to an apparent depolarization effect in the time series. To show this, one introduces the Jones formalism, the Stokes formalism,

⁷Again, a detailed treatment shows that these polarizations can mix to V as well if phase is not preserved in the instrument (Johnson, B. R. et al., 2006).

⁸This effect is what a number of other authors have called 'Cross Polar Response' (Jones, 2005), 'Cross Polar Leakage' (Hinderks, 2005), and other variants. Here, we have settled on the term given in *e.g.* Johnson, B. R. et al. (2006), although (as shown below) for PSBs alone this term is the same as cross polarization under certain assumptions.

and the Meuller matrix. These algebraic formalisms and first principles analyses are given in Jones (2005) and Hinderks (2005); here, we merely point out that applying these to a system like that used in QUaD leads to Equation 2.13. The next logical step required is to introduce the concept of a polarized beam. This is the angular response of the instrument for the co- and cross-polarizations, denoted by $P_{\parallel}(r, \theta, \phi)$ and $P_{\perp}(r, \theta, \phi)$. In single moded systems, the $P_{\parallel,\perp}$ are not a function of the detectors, but only the optics and feed horns (both devices at the bottom of a feed share the polarized beams; they differ only in their coupling to them). O'Sullivan (2006) presents a thorough analysis of the simulated co- and cross-polar beams for QUaD; the reader is referred to this work for details on the method and results of such a calculation.

A PSB detects the polarized sky convolved with the polarized beam, integrated over the bandpass and with the efficiency dictated by the polarization efficiency of the detector. The general expression for the signal from a single PSB in the flat sky approximation and not including noise is (see Jones 2005):

$$V(\mathbf{x}) = \frac{g'}{2} \int d\nu \frac{\lambda^2}{\Omega_{\text{beam}}} F_{\nu} \int d\Omega$$

$$P_{\parallel}(\nu, \mathbf{x}) \left[(1+\epsilon)I(\nu, \mathbf{x}) + (1-\epsilon) \left(Q(\nu, \mathbf{x}) \cos 2\alpha + U(\nu, \mathbf{x}) \sin 2\alpha \right) \right] +$$

$$P_{\perp}(\nu, \mathbf{x}) \left[(1+\epsilon)I(\nu, \mathbf{x}) - (1-\epsilon) \left(Q(\nu, \mathbf{x}) \cos 2\alpha + U(\nu, \mathbf{x}) \sin 2\alpha \right) \right] \quad (3.20)$$

where variables are as in Equation 2.13 except we introduce x to represent Cartesian coordinates and define $\Omega_{\text{beam}} = \int d\Omega (P_{\parallel} + P_{\perp})$. If the explicit angular and spectral dependencies are dropped, this equation can be re-written as

$$V = \frac{g}{2} \int d\nu \lambda^2 F_{\nu} \int d\Omega \{ I + \gamma \mathcal{P}(Q \cos 2\alpha + U \sin 2\alpha) \}$$
(3.21)

where we define parameters as in Jones (2005):

$$\gamma = \frac{(1-\epsilon)}{(1+\epsilon)} \tag{3.22}$$

and

$$\mathcal{P} = \frac{(P_{\parallel} - P_{\perp})}{(P_{\parallel} + P_{\perp})}.$$
(3.23)

A few assumptions have been used in this definition. First, it should be noted that Equation 3.23 is only valid in the flat sky approximation, which means that signal due to, for example, far side lobes creates a far more complex equation. Secondly, it is assumed that the beam and polarization efficiencies can be removed from the integral over the sky $(\int d\Omega)$; this assumption functionally means that neither of these two terms depend on the pointing of

the telescope. The beam and polarization efficiency prefactors are then thought of as part of a redefined calibration constant, $s = s' \int d\nu (1 + \epsilon)$. This means that the detector cross polar response results in a loss of sensitivity to the polarized part of the incoming radiation, but *not a mixing of I, Q or U*. We can thus conclude that cross polarization is degenerate with depolarization in PSB-based systems due to the way the polarized detectors and optics couple. This depolarization can in turn be absorbed into the overall calibration if a suitable calibration source is available and the polarization efficiencies of the detectors are known.

3.3.1 MEASUREMENTS & RESULTS

QUaD's polarization characteristics were measured both in the laboratory and the field. The experimental philosophy in these measurements is to used a chopped thermal source coupled with a rotating polarizer to measure the polarization response of the instrument. A description of both the lab and field experiments and their results follows.

In the Stanford lab, the experimental apparatus was to mount a large bucket filled with LN_2 on the top of the cryostat (a more thorough description of this experiment can be found in Hinderks 2005). A chopper wheel with alternating transparent and eccosorb quadrants was used to AC couple the signal to the detectors at ~ 2 Hz, which reduces the effects of 1/f noise in the experiment. A rotating polarizer⁹ was placed on the cryostat window, and the grid was rotated using a 'step and integrate' procedure. This method consists of stepping the polarizer angle 1° and holding it steady for several chopping periods, then stepping to the next position. An eccosorb mask was used to restrict the aperture of the bolometers so that only a single pair of devices 'saw' signal during a given measurement. This mask had two functions: i) to limit the loading on the detectors to keep them in the linear regime and ii) to limit highly off-axis rays which artificially increase the measured cross polar response.

In the first stage of analysis, the data are demodulated only for those periods when the gird is not rotating. The average of each polarizer step yields a sinusoidal variation in a bolometer's complete time series, which can be fit to a sine wave model. The ratio of the minimum signal of the model wave to the maximum signal of the model wave yields ϵ , while the phase of the model wave yields the α of the detector. Based on this test, most of the detectors appeared to be within $\pm 2^{\circ}$ of their nominal position; bolometer pairs are orthogonal to a similar error. The variation in α is due to errors in the machining of the detector housings and focal plane. The cross polar leakage, defined in Equation 3.22, appeared to be as high as 10 % from these measurements. This is much higher than the

⁹The polarizer is constructed from a $15 \,\mu$ m thick sheet of polypropylene with parallel lines of copper deposited every $10 \,\mu$ m along its surface; it is virtually the same as the polarizing grid used in the secondary mirror calibration source.

3% determined for the detectors alone: it is thought that a large part of the error is due to the feed horns or the horn-bolometer interface. In any case, the Stanford measurements were not ideal, as they did not involve the full optical chain and suffered from complicated effects like those due to the aperture reduction.

During the 2005/2006 summer season at the south pole, these measurements were performed again using a different apparatus (shown in Appendix A.3). Essentially, the source consists of a box with no back mounted on a tower above the lip of the ground shield. A chopper wheel with alternating transparent and eccosorb quadrants is rotated at a few Hz, and a wire grid polarizer at the aperture of the box polarizes the radiation (because the box has no back, the sky is the cold load and the ~ 250 K eccosorb on the chopper wheel is the warm load). The telescope was rastered over this source to build up an image of the (near-field) beam of the telescope (this is similar to making a map of the sky, only for a az/el stationary source). The grid was rotated to a number of polarization angles and the experiment was repeated; this allowed a measurement of the polarized beam at a variety of input polarization angles. The data are demodulated using a reference signal originating from the source's chopper, and binned into square pixels. Near field beam maps like those shown in Figure 3.9 result from this process.

The sum of the pixels in maps like these gives the integral over the near field beam of the instrument. This should give us a reasonable estimate of the polarization performance of the instrument for each grid angle¹⁰. Plotting the integrals over the beam maps for each grid angle should result in a sinusoidal modulation whose ratio of minimum to maximum will yield ϵ and phase will yield α for each bolometer. Examples of this type of analysis are shown in Figure 3.10.

To perform the fits, a non-linear fitting procedure employing a sinusoid model with variable amplitude, phase and offset is used. The error estimates for the data are based on the variance in the highest elevation row of pixels in each map; this row should be free of signal from the source, and should gives the random error per pixel from the atmosphere and instrument in a clean way. The fitting procedure is automated; each fit is tested for χ^2 , and if this statistic is larger than some threshold, the fit is performed again with a slightly different set of initial parameter estimates. This process is repeated until the χ^2 statistic comes below the threshold, or 5 iterations have occurred.

The threshold statistic in this procedure is problematic: the noise estimate from the highest row of pixels yields a number approximately an order of magnitude too small to give a well-behaved χ^2 . However, in looking at plots like Figure 3.10, it is obvious that

¹⁰In practice, this is a very difficult measurement to relate to the far field performance of the instrument. A variety of effects could be coming into play that may have different effects in the far field beam, so individual estimates of the polarization performance of the instrument should be taken with a grain of salt.



Figure 3.9: Polarized Source Near Field Beam Maps

This figure shows beam maps for bolometers 150-1-A and 150-1-B made using the polarized thermal source on March 16, 2006. The image on the left is the near field beam for the A bolometer when the grid alignment is expected to produce a maximally polarized signal. The image on the right is made from data taken simultaneously to that on the left, but for the B bolometer. This should represent a minimum for the B bolometer, as its polarization sensitivity is orthogonal to the the A device. The images are on a logarithmic colour scale with 1 square degree pixels; the B device's map exhibits small but non-zero signal, meaning that some cross polarization is occurring. The dip in the middle of these maps is due to the secondary mirror occulting the thermal source; the integral over these beams should give the average polarization performance for each device in the focal plane.

many points which have these small error bars lie many σ from the best fit curve. This behaviour is certainly due to some systematic effect, although it is not clear what that is at this point. By testing three different types of aperture at the source box, it was found that the aperture does appear to have some sort of depolarizing effect on the radiation coming from the polarized thermal source. This type of effect, coupled with variations in the atmospheric loading and ambient conditions at the telescope, probably contribute to the variation of the data above that explained by random noise¹¹. In any case, the fit threshold is set to be ten times double the number of degrees of freedom of the fit to account for the difference in the size of the error bars and the systematic variation between them.

It is not necessary to perform an integral over the beam to determine the polarization parameters of the instrument; each pixel in the beam map can be treated as a separate measurement of the telescope's polarization properties. Sine waves can be fit to each of the pixel values (using some estimate for the noise, like the per pixel noise estimate based

¹¹This type of argument also makes this whole measurement somewhat suspicious.



Figure 3.10: Sinusoidal Modulation of QUaD's Near Field Polarized Beams This plot shows the result of integrating over the type of maps shown in Figure 3.9 for each grid angle. The left panel shows results for bolometer 150-1-A (red) and B (blue), while the right panel shows the results for 100-1-A (red) and B (blue). The data are plotted as black points with colour bars matching the A-red & B-blue convention. The dashed lines show the model fits to these data, and the vertical lines show the phases of the fits. The data are well described by a sinusoidal model, although the error bars on the points are slightly smaller than expected (see the discussion in the text).

on that measured for the top row of the map above), and the parameters ϵ and α can be plotted for each pixel. This gives an idea of how uniform the parameters are over the beam. Subtleties in this analysis are that the individual images from each of the detectors and grid angles must be moved to a common centroid and rotated to a common telescope angle (*e.g.* $\theta = 0.0$). Once this is done, the pixels can be fit to the sine model in the same way as above and the parameters found from the fit can be plotted in map space. The results of this type of analysis are shown in Figure 3.11 for α and Figure 3.12 for ϵ .

A close inspection of Figure 3.12 shows that a clear systematic trend exists in the parameter ϵ . As one goes to higher radius in the beam, the value of ϵ increases. This means that the radius to which one performs the integral over the near field beam has a strong effect on the resulting beam-averaged value of ϵ . The cause of this radius-dependent beam shape is unknown, but is possibly caused by the fact that these measurements were made in the telescope's near field. This effect is therefore quoted as part of our systematic uncertainty in the measured ϵ .

81



Figure 3.11: Polarization Angles α from Polarized Source Measurements This figure shows the results of analyzing each pixel in the map as an independent measurement of the polarization characteristics of the instrument. This figure is for the polarization angles of the bolometers, α . The whole array is shown: left are the A-side bolometers, and right are the B-side bolometers. No systematic effects over the beam are visible in these images. (Figure courtesy J. Hinderks).

and the second second	Tabl	e 3.2: QUal	Polarization Param	eters	
ν	$\langle \epsilon \rangle$		σ_{ϵ}	$\langle \Delta lpha angle$	$\sigma_{\Delta lpha}$
(GHz)		Random	Systematic	(Deg	rees)
150	0.07	± 0.010	+0.012, -0.003	0.0	± 1.2
100	0.07	± 0.015	+0.012, -0.003	0.0	±1.2

Once the bolometer angles and polarization efficiencies are determined, their statistics can be studied. In the final analysis of these thermal polarization source data, the parameter values listed in Table 3.3.1 are found. Because α can take on four values ($\alpha \in \{0, 45, 90, 135\}$), the bolometer angle result is quoted as the deviation from the nominal bolometer angle $\Delta \alpha = \alpha_{nominal} - \alpha_{measured}$.

Due to the difficulties involved in interpreting the parameters measured in the 'source on a tower' experiment in terms of the far field beam, they are taken to be statistical measurements of the telescope's performance rather than measurements of the parameters for individual bolometers. Furthermore, these measurements were only performed on the 2006 season bolometer arrangement; no such data exists for the 2005 season, as the thermal



Figure 3.12: Polarization Efficiencies ϵ from Polarized Source Measurements This figure shows the results of analyzing each pixel in the map as an independent measurement of the polarization efficiencies of the detectors, ϵ . The whole array is shown: left are the A-side bolometers, and right are the B-side bolometers. There is a clear trend to larger ϵ with increasing beam radius in these plots; this constitutes a serious systematic effect. (Figure courtesy J. Hinderks).

source apparatus was not available during commissioning¹². Fortunately, since the bolometers in the focal plane during the 2006 season are substantially the same as those during the 2005 season, interpretation of these measurements as statistical rather than per bolometer allows us to use the statistical measurements for the analysis of the 2005 data. Although this is not the ideal situation, this assumption allows us to assign realistic errors to the polarization parameters of the telescope as inputs to instrumental simulations designed to estimate the final errors on our measured power spectra. However, systematic errors in our measurements of these parameters *will* propagate through to the final result, and we must be careful to both minimize these during later data analysis and ensure that the simulations capture all of the salient features of the polarization characteristics of the instrument.

¹²This is unfortunate as the data taken during the austral winter of 2005 forms the bulk of that presented in this thesis, making the polarization characteristics of the telescope one of the dominant uncertainties in the final analysis.

4 INSTRUMENT PERFORMANCE

Of major importance to the QUaD project is the characterization of the instrument, both in the laboratory in preparation for fielding (discussed in the last chapter), and – possibly more importantly – in the field itself. This is because understanding the *in situ* performance of the instrument directly affects the analysis and interpretation of the data it produces which is, after all, the object of the experiment. A great deal of time and effort has therefore been expended on measuring the characteristics and performance of the instrument on the telescope, both to facilitate data analysis, as well as to guide necessary changes to the system's hardware when necessary.

This chapter presents the various tools and methods used to characterize the performance of the system. The order of presentation is not necessarily based on the order in which the tests were performed. In fact, many of the tests discussed here are performed daily to ensure a running record of the instrument's performance exists for later data analysis. On the other hand, a number of tests were performed only at specific times during the experiment's lifetime, or are only possible at a specific time of year. Issues surrounding the frequency and duration of each of these tests are discussed on a specific basis.

Roughly speaking, this chapter's organization mirrors the order in which tests must be performed in the field to understand the instrument's performance. That is, the first test required once the receiver is on the telescope is to measure some photons from an astronomical source. Then a map of a point source should be made, followed by measurements leading to an understanding of the focal plane's mapping to the sky. Then, the beams and sidelobes should be characterized. A measurement of the sensitivity and polarization characteristics achieved in the system in the presence of the true atmospheric column should then be undertaken. Finally, secondary effects like electronic filtering should be understood. A noteworthy topic conspicuously missing from this chapter is a discussion of calibration, which is treated separately in Chapter 5.

4.1 POINTING

The term 'pointing' refers to measurements of the positional accuracy of the telescope with respect to the sky. In QUaD's case, there are two kinds of pointing: optical pointing, and radio pointing. These two modes use separate instruments which are tied together during data analysis to provide a unified model of how the telescope's axes of motion are working. This model is modified heavily at the beginning of an observing season, and approximately once per month subsequently. During the summer, the pointing model is considered to be at its worst, since the combination of large temperature fluctuations and shifting of snow outside the observatory serve to make the absolute position of the telescope and its track unstable.

Of prime importance to the telescope's function is the radio pointing model, since this relates how the telescope's position encoders map to the sky during observations. Unfortunately, bright, compact radio sources are rare on the sky and those that exist are mostly distributed along the galactic plane, making the determination of the pointing model based on radio sources alone extremely difficult. Thus, a class of pointing sources – at any wavelength — which are uniformly distributed across the sky is required. The solution employed in QUaD is to mount a small optical telescope to the side of the main mirror which can use stars as pointing targets. This telescope projects onto the sky in the same way as the main dish, with some known offset. During optical pointing measurements, the centroids of a (large) set of bright stars uniformly distributed across the sky are measured, and pointing models for the three axes (azimuth, elevation, and rotation) are derived. This process is automated, and usually requires a whole day to measure most or all of the telescope's tracks. The optical pointing model derived from these measurements can then be tied to the radio pointing model to give an accurate absolute coordinate system for the telescope across the entire southern sky.

That said, it is absolutely necessary to check the radio pointing model separately, and then to tie this to the optical model. This is performed using cross observations on a selection of mm-bright astronomical sources. The known characteristics of these sources are discussed next.

4.1.1 RADIO POINTING SOURCES

The ideal pointing sources for a telescope are bright in the band of observation and pointlike, so that they have an easily identifiable peak (for example, bright quasars make excellent pointing sources for optical telescopes). In the millimetre waveband, sources considered acceptable for these criteria are either planets, approximately thermal sources with very low effective temperatures (for example, star–forming regions), or bright non-thermal sources (like flat–spectrum quasars). Unfortunately, planets are always too low in the southern sky to be observed with QUaD. Furthermore, typical fluxes for bright quasars or local, compact active galactic nuclei (AGN) are a few Janskys at 100 GHz, which is too close to the atmospheric noise to provide a high instantaneous signal to noise level from the ground¹. QUaD is therefore restricted to cold thermal sources in the galactic plane for the purposes of measuring the instantaneous pointing.

The brightest mm sources in the sky beyond the solar system are cold star-forming regions, which are popular astronomical targets in their own right. Of these, Ultra-Compact HII (UC HII) regions form the brightest and most compact population (one of many extensive reviews of the subject can be found in Churchwell (2002)). These regions are the locations of newly formed massive stars still embedded in their natal molecular clouds. Because these regions are still heavily dust enshrouded, they are only observable at millimetre and far infra-red wavelengths thanks to their thermal emission. The central star in these sources are thought to have spectral types of B3 or hotter, and have just settled onto the main sequence. This means that over a relatively short time scale, the UC HII region will evolve into a classical nebula due to photoionization from the hot central star(s). The spectral energy distribution (SED) of these sources typically peaks around 200 μ m, and their fluxes can reach a few Jy per square arcmin at 150 GHz, with typical extents up to ten square arcmin.

Because the QUaD beam widths are comparable to or larger than the typical UC HII angular extent (meaning they can be considered 'point-like'), we have chosen to use five of these as quick-look sources for pointing tests. Although it would be ideal if they were point sources, the UC HII regions are the smallest, brightest mm sources visible from the south pole. The QUaD UC HII sources and their rough characteristics are given in Table 4.1.1. In this table, $S_{ap}(\nu)$ is the total flux at frequency ν inside an aperture defined to have a radius of $2\sigma_b^2$. For MAT 6a, $\sigma_b = 5.1$ arcmin at 150 GHz (derived from the MAT/TOCO experiment; Puchalla et al. (2002)), and for the other four sources, $\sigma_b = 4.0$ arcmin at 150 GHz and 7.6 arcmin at 100 GHz (from the BOOMERanG experiment; Coble et al. 2003). Also, dust maps of these sources are available in public archives³ based on IRAS and DIRBE data (Schlegel et al., 1998); these are shown in Figure 4.1.

As can be seen in Figure 4.1, these UC HII sources are actually quite complex, with

¹Note that this isn't a concern for mm telescopes with the ability to chop out the atmospheric temperature variation, like the JCMT, CSO, or IRAM. Likewise for interferometers, which can reject a great deal of the atmospheric noise at these wavelengths. Only scanning telescopes like QUaD are subject to this kind of requirement.

 $^{^{2}\}sigma_{b}$ is defined in terms of a Gaussian model for an experiment's beam. Using this model, σ_{b} is equivalent to the Gaussian's standard deviation.

³We acknowledge the use of NASA's SkyView facility (http://skyview.gsfc.nasa.gov) located at NASA Goddard Space Flight Center (McGlynn et al., 1997).

Source	$lpha_{ m J2000}$	$\delta_{ m J2000}$	$S_{\rm ap}(90{\rm GHz})$ (Jy)	$S_{\rm ap}(150{\rm GHz})$ (Jy)
RCW 38 ^{<i>a</i>,<i>b</i>}	134.778°	-47.522°	137 ± 16	122 ± 12
IRAS 08576 ^b	134.86°	-43.76°	20.6 ± 2.6	31.9 ± 2.2
IRAS 1022 ^b	156.05°	-57.80°	143 ± 15	119 ± 11
RCW $57^{a,b}$	167.983°	-61.335°	74.6 ± 9.1	70.1 ± 6.0
MAT 6a ^{<i>a</i>,<i>c</i>}	182.515°	-62.846°	-	34.2 ± 8.5

Table 4.1: QUaD's UC HII Region Pointing Sources

References are: ^aGoldstein (2004), ^bCoble et al. (2003), ^cPuchalla et al. (2002).

multiple peaks and complicated background emission. The typical scale of the emission is such that these sources are adequate for the purposes of pointing measurements, but cannot be used for deep measurements of the beam, or any other test requiring extremely detailed knowledge of the telescope's performance. Of the five sources, RCW 38 appears to have the best combination of compactness (~ 50 arcsec, Coble et al. (2003)) and brightness; we have therefore adopted this source as our default standard for quick measurements of the system, such as gross characterization of the beams and astronomical calibrations. Therefore, throughout this chapter data taken on RCW 38 will be presented unless specifically noted otherwise.

4.1.2 CROSS OBSERVATIONS

The standard quick–look observation employed with QUaD is the 'pointing cross', which can be used to measure the radio pointing model. The pointing cross observation simply consists of centering the telescope on the nominal source position, nodding the telescope up and down, re-centering, and nodding the telescope left and right. This sequence is generally performed twice to provide data in the case that one of the individual cross observations is corrupted. A typical single pointing cross observation is shown in Figure 4.2.

In fact, the very first measurement performed whenever the receiver is installed on the telescope is a simple pointing cross observation. This has two functions:

- 1. it ensures that the receiver is set up to detect astronomical photons (approximately correct focusing, correct bias and gain modes, *et cetera*), and
- 2. it ensures that the telescope's motion control systems and pointing models are approximately correct.

Both of these items are required for any other form of observation to be successful. If the pointing model is near to true, the blips will be large enough that they can be picked out of the time series by eye, or using a relatively simple algorithm.



Figure 4.1: Dust Maps of QUaD's UC HII Pointing Sources

These are images of the dust emission in QUaD's five UC HII region pointing sources. They are based on the work of Schlegel et al. (1998), and, because they are thermal sources, should strongly correlate with the 100 and 150 GHz emission in these regions. The colour stretch is logarithmic and the same in all of these images, so that the relative brightness and background level in each can be easily assessed by eye (the colour is actually the logarithm of the optical reddening E(B - V) in magnitudes). The effective beam size in these maps is 6.0 arcmin FWHM, which is similar to the QUaD beam sizes. Clearly, the emission in these regions is actually quite complex, with multiple peaks and bright background emission present.

During regular CMB observations, pointing crosses are performed four times per daily cycle as part of the standard calibration block (see Section 6.1). Two crosses are performed back to back for redundancy. To compare the blip peak positions to the pointing model, the timings of the peak blips in the two central bolometer's time series are compared with the reported positions in the azimuth and elevation time series. The difference between reported position of the peak and the known position of the peak gives the offset measurement at that time.

Although the pointing offset measurements are given by a simple enough procedure, the sheer volume of data makes performing the search of peak blips by eye impossible. Therefore, an automated procedure for identifying peaks in the time series is required; the algorithm we have developed is as follows. First, all of the cross observations on a given day are identified by searching through the bit mask register (discussed in Section 6.1).



Figure 4.2: A Single Pointing Cross Observation

These plots show the right ascension $(\Delta \alpha)$ and declination $(\Delta \delta)$ offsets for a single cross observation, as well as the resulting time series for the central co-polar bolometer from an observation taken on June 6, 2005. The central plot shows the offsets, the lower plot shows the bolometer time series for the forward (black) and backward (grey) azimuth scan, and the right plot shows the bolometer time series for the upward (black) and downward (grey) elevation scans. The nominal source position is shown as a red diamond in the central plot. The locations of the blips in these time series give the pointing offsets of the telescope at the time of the observation. The large slope in the elevation scan is due to increasing sky temperature (and thus signal) as the airmass through which the telescope's beam is scanning increases; the time series do not line up exactly due to temporal variation in atmospheric emission. Typically two of these observations are performed back to back to provide redundancy. Because individual cross observations (*i.e.* single azimuth scan, or single elevation scan) are separated by a reset of this bit mask, each cross observation is identified by forming the derivative of the bit mask register and flagging steps (the first data point in the cross observation) followed by regions of zero derivative (the observation itself). This is then used to strip out the data corresponding to each cross observation. Variants of this method are used to separate out the different types of observations based on the bit mask time series.

The atmosphere is a major contaminant in these observations, and can cause false blip identifications, fits to the blips biased by a large slope or higher order structure in the data, or in very bad weather, even completely mask the blip under noise⁴. To deal with this, four data channels are used in the pointing cross analysis: the channels corresponding to the central two bolometers (these 150 GHz devices contain the astronomical signal), and two 150 GHz channels to the top and bottom (in elevation) of the central channel, called the 'atmospheric monitors'. It is well known that atmospheric noise is highly correlated across the typical angular distances of the QUaD focal plane, so the philosophy here is to use the two atmospheric monitor channels – which are chosen because they do not contain any appreciable astronomical signal – to remove the atmosphere from the signal channels. This is done by first fitting a fifth order polynomial to the monitor channels, and checking that no large, time correlated blips are contained in the residuals. Residual structure is identified by setting a threshold and looking for points higher than the threshold away from the mean (because the pointing sources are so bright, more complex algorithms are not required). Typically, the threshold is 5σ , although this can be varied in particularly poor conditions.

If strong simultaneous blips are found in both of the monitor time series, they are almost certainly due to some strange correlated atmospheric variation that can masquerade as a real astrophysical signal, so that cross observation is discarded. These type of signals occur in both monitor time series in fewer than 0.5 % of the cross observations, but they *must* be identified and removed. On the other hand, if a blip above the threshold is found in only one of the two monitor time series, the user is warned, and the monitor time series containing the large blip is removed from analysis. The monitor time series without the blip is used as the new atmospheric monitor time series. Finally, in the case that neither monitor channel contains large blips (which is by far the case for most of these cross observations), the atmospheric monitor time series is formed from their mean. Figure 4.3 shows an example of this process.

Once the atmosphere has been removed, it is necessary to find the centre of each blip to determine the pointing of the telescope. In the absence of noise, it would suffice to simply pick off the minimum value in the atmosphere subtracted time series. However,

⁴Which is impressive given that in good weather the blips are typically greater than 15σ signals!



Figure 4.3: Cross Observation Atmospheric Removal

This plot shows the results of our atmospheric removal process in the some of the same data featured in Figure 4.2. The upper plots show the time series for the central co-polar device for half of an azimuth scan (left) and elevation scan (right) in black. The blip due to the astronomical source is clearly visible near the centres of both. The two monitor channels are shown in grey in both cases; clearly, a great deal of structure is common to all three channels. The red line shows the fifth order polynomial fit to the monitor channels; no strong blips are contained in these (this would show up as a large feature in the residual). The mean of the monitor channels is formed, this is fit to the central bolometer channel, and the fit is subtracted. This residual is shown for both cases in the lower plots. Once the atmospheric emission is removed, the blips stand out strongly against the baseline level, which is necessary for an accurate fit to find the blip's centroid.

$\langle x \rangle$	σ_x
-1.04'	2.72'
0.49'	1.74'
	< x > -1.04' 0.49'

Table 4.2: Radio Pointing Summary for 2005 Observing Season

a more complex approach must be adopted when the shape of the source is noisy. In this case, a Gaussian function based on the telescope's point spread function (PSF) is used:

$$V(t) = V_{\rm A} \exp\left(\frac{-(t-t_0)^2}{2\sigma_t^2}\right) + V_0$$
(4.1)

where t is time, V is the signal, and V_A , V_0 , t_0 and σ_t are free parameters in the fit. A gradient-expansion algorithm is used to compute a non-linear least squares fit to the time series (this combination of function and fitting algorithm is a recurring theme in this Chapter). In the case of the pointing analysis, the function's centre, given by the telescope's right ascension and declination at time t_0 , is stored.

This analysis can be performed over the whole observing year; these data act as a record of the telescope's pointing quality over time. Assuming a Gaussian probability density function (PDF) for the pointing offsets⁵, the means and standard deviations in right ascension and declination for the 2005 season pointing measurements are given in Table 4.1.2. These corrections, while smaller than a beam width, are nevertheless non-negligible.

These data are most useful for making (small) corrections to the pointing model; the mean correction over, for example, a single day can be determined and applied to the telescope pointing registers for that day. Performing this correction improves both the spatial width and centroid position of point sources (such as quasars) by up to 10% in our full year maps. The pointing measurements for the entire 2005 season are shown in Figure 4.4.

4.2 NAIVE RASTER MAPS

After assessing the ability of the telescope to point accurately, the first step toward instrument characterization is to make a simple raster map of an astronomical source. The term 'raster' refers to scanning back and forth over the source in a grid pattern to build up an image of the source and its environs. This technique is used with QUaD to make maps of sources like RCW 38 to characterize the feed offsets, beams, *et cetera*.

⁵This is not necessarily the best description of the data; rather, it is better to think of random Gaussian scatter about some slow polynomial-like variation, see Figure 4.4.



Figure 4.4: QUaD Pointing Performance for Winter 2005

This plot shows the calculated offsets for the QUaD pointing in right ascension (α) and declination (δ) for the winter 2005 CMB observing season. The points show the results of each measurement, while the coloured regions are: known bad pointing (purple), fixing the pointing model (grey), and refocusing (green; discussed in more detail in Section 4.4). The gaps between day 90 and day 110 are due to combinations of poor weather, moon contamination, and a bad pointing model.

The typical QUaD rastered time series for calibration observations involves constantelevation azimuth scans of duration 48 s with an azimuthal length of $3.2^{\circ 6}$. For RCW 38 at an elevation of -47.5° , this corresponds to an angular distance on the sky of $3.2^{\circ} \times \cos(-47.5) = 2.16^{\circ}$. After scanning the full throw in one direction and back in the other direction, the telescope steps up in elevation by 0.02° . This pattern is repeated N times to complete a calibration map (usually, N = 101). The angular size of an RCW 38 map is thus 2.16° in right ascension and $0.02^{\circ} * 101 = 2.02^{\circ}$ in declination. This map size is chosen so that every bolometer in the focal plane maps the RCW 38 region completely.

The time series of a single scan looks very similar to that of the azimuth scan featured in Figure 4.3. When a bolometer is off-source, its time series is indistinguishable from that of the monitor channels, and when it is on-source, it looks similar to that for the signal channel. The full set of time series can be binned into a grid, which produces a map of the astronomical source region.

In the absence of noise, a telescope can be thought of as a machine that performs a linear transformation on a vector of sky brightnesses⁷ to create a vector of signals at its output (as discussed by *e.g.* Lineweaver et al. (1994)). Let m_i be a vector of the true sky brightnesses, and d_j be the data stream at the output of the telescope. The transformation can then be written as:

$$\mathbf{A}_{ij}m_i = d_j. \tag{4.2}$$

In this equation, A_{ij} is called the 'pointing matrix', and it encodes the way the telescope scanned the sky brightnesses and turned them into the signal vector. Here, *i* is the number of sky pixels, and *j* is the number of time samples, so in general A_{ij} is not square. The numerical entries in A_{ij} depend on the instrument, scanning strategy, and sky map pixelization scheme. At its simplest, which is the case for QUaD temperature mapping, A_{ij} simply contains ones and zeros⁸ (provided that the telescope time series is calibrated in terms of sky brightness units). The simplest way of thinking about constructing a pointing matrix is that it contains a representation of the sky map at each time step as its row entries. So for each row, one enters '1' at pixels which were observed, and '0' for all others. This procedure is repeated at each time step to build up the full A_{ij} .

⁸Much more complex formulations are possible. COBE or WMAP have pointing matrices containing 1, 0, and -1 to encode the fact that differential radiometers are used (Lineweaver et al. (1994), Wright et al. (1996), Hinshaw et al. (2003)). An experiment like SCUBA, which employs a double sky difference, contains 1, 0, and two -1/2 at each time sample (Zemcov et al., 2003). Instrument parameters can even be folded into the linear solution, like Fixsen et al. (2000) developed for the IRAC camera on the Spitzer space telescope. It seems that the only limit to this way of thinking about sky mapping is one's imagination!



⁶These numbers refer to the constant velocity section of the scan, and do not include the components at the end of the scan which suffer from acceleration noise.

⁷This may seem to be a slightly odd way of thinking about it, but recall that a pixelized map - which is just a matrix of sky brightnesses - can be converted into a vector under a suitable transformation.

In practice, the time series is given, and once A_{ij} is built from the pointing data, one would like to invert the linear equation to solve for m_i . Using Equation 4.2 we can write

$$(\mathbf{A}^{\mathsf{T}}\mathbf{A})m = \mathbf{A}^{\mathsf{T}}d \tag{4.3}$$

where we have suppressed the pixel and time indices for clarity. This representation can be multiplied by $(\mathbf{A}^{\mathsf{T}}\mathbf{A})^{-1}$ to yield the least-squares estimate for the true map brightnesses m. In the case of simple raster temperature data, the matrix $(\mathbf{A}^{\mathsf{T}}\mathbf{A})^{-1}$ is simple; it is diagonal, with size $n_{\text{pix}} \times n_{\text{pix}}$. The diagonal entries are simply the inverse of the number of times the sky pixel was observed. Moreover, because the term

$$\mathbf{A}_{ij}^{\mathsf{T}}d_j = \left(\sum_j d_j\right)_i,\tag{4.4}$$

we can write

$$m_i = (\mathbf{A}^{\mathsf{T}} \mathbf{A})^{-1} \mathbf{A}^{\mathsf{T}} d = \left(\sum_j d_j\right)_i / N_i$$
(4.5)

where N_i is the number of times pixel *i* was observed. That is, the least squares map simply boils down to being the average of the observations in each pixel! Although this is a simple result, it is interesting that it can be derived from the full map making formalism employed in the most complex CMB experiments.

One thing neglected in the treatment above is noise, both correlated and uncorrelated. In the presence of noise, the solution represented by Equation 4.5 is no longer equivalent to the optimal solution, and must be treated more carefully. A simple next step is to add the effect of white noise; if the variance⁹ of the time series is σ^2 , then the optimal map is given by

$$m_{i} = \left(\sum_{j} \frac{d_{j}}{\sigma_{j}^{2}}\right)_{i} / \text{diag}\left(\frac{\mathbf{A}_{ij}^{\mathsf{T}} \mathbf{A}_{ij}}{\sigma_{j}^{2}}\right)$$
(4.6)

which is (once all of the subscripts are sorted out) just a weighted average of the data contributing to each pixel.

This being said, the noise affecting the QUaD receiver is not necessarily white, and is certainly not uncorrelated. In fact, the dominant source of noise is the atmosphere, which induces large, polynomial-like drifts in the bolometer time series which are highly correlated across the array. In the face of these types of noise, more *ad hoc* methods than the optimal mapping algorithm are employed to make maps. The very simplest method is to use Equation 4.5 as it stands. This makes maps that have a great deal of striping,

⁹The variance of data is the optimal weight function for uncorrelated white noise; this is not true of more complex noise distributions.
where the source is at best barely visible, and so is not generally used. The next level of complexity is to use median filtering, which simply involves computing the median of a scan and subtracting it from the time series. This works quite well, and a source like RCW 38 can be easily discerned in such a map, an example of which is shown in Figure 4.5. Beyond this, there are a number of different techniques that can be used to clean raster scan data. However, all of these filter the distribution of flux in the map at some level, in some cases in very complex ways. We will discuss three of these methods here.



Figure 4.5: RCW 38 Raster Map with Median Filtering

This is a raster map of RCW 38 made with the central co-polar bolometer and median filtering (from data taken on March 9, 2006). The binning was performed according to Equation 4.5. In this image, a logarithmic scaling has been used to bring out the dimmer structure in the map (RCW 38's central lobe dominates if this is not done). Clearly, although the median filtering has made the source visible, a great deal of striping is evident. This can be removed using more advanced algorithms. The discontinuous structure at the extreme right is due to the gain reset at the end of a scan and is not real astronomical signal.

The first method is called 'polynomial filtering'. This technique consists of fitting a polynomial between first and fifth order to the data from a single scan before binning it to form a map. This has the effect of removing the polynomial-like atmospheric noise from the time series, at least to the order of the polynomial used in the fit. Subtracting a first order polynomial (*i.e.* a slope and an offset) removes a great deal of the power in the atmospheric drift, but does not adversely affect the astronomical signal in the data. However, the drawback of this technique is that any polynomial higher than first order will remove some of the source flux, and a considerable amount of the source's background structure, from each scan. Because the polynomial fit results in a different solution for each scan, it is difficult to estimate the size of this effect based on the polynomial's shape.

Another method that has been used with some success is the array average subtraction technique. This involves computing the average of the (single frequency) array at each time step to create an atmospheric monitor channel; this is more powerful than using a subset of the channels because uncorrelated noise is reduced drastically by using the full array. Furthermore, because the whole array contains a large number of bolometers, the requirements on an individual channel's behaviour when it is included in the mean are not so stringent as in the cross observation analysis. This array average time series is subtracted from each channel before it is binned into the map, which removes the correlated atmospheric contribution but not the astronomical signal from the time series.

The drawback to this method is that if the whole array is used to calculate the atmospheric monitor time series, the source occasionally falls into the monitor channels. This causes negative images of the source to be distributed around the map at positions corresponding to the offset of a monitor channel bolometer from the signal monitor. Assuming that both bolometers at the bottom of a feed horn are used in the calculation, these sources have amplitude $1/n_{\text{horns}}$ that of the positive source, where n_{horns} is the number of horns used to calculate the atmospheric monitor time series. Since $n_{\text{horns}} \sim 15$, these negative ghosts are not negligible compared to the central source's amplitude. In fact, a similar issue arises if the source's background flux distribution is complex: the map made in this way is a complex linear combination of the background subtracted from itself. A way around this is to excise the bolometers near the source from the atmospheric monitor calculation, so that flux from the source and its environs doesn't get included in the monitor channel. However, this has the effect of increasing the noise around the source itself. In this case, it is more desirable to use the full array and ignore the negative ghosts in the map.

A third method which can be used to clean the time series prior to binning is to use a low-pass filter to remove the high frequency components of the time series. The filter function we apply is defined by

$$H(k) = \cos\left(\frac{\pi}{2}\frac{(k-k_1)}{(k_2-k_1)}\right), \quad k_1 \le k < k_2$$

$$0, \quad k \ge k_2$$
(4.7)

where k is the frequency variable. Figure 4.6 shows this function in detail. The low pass

filter method is quite successful, but the quality of the results depends strongly on the values of k_1 and k_2 . These two parameters define when the filter begins to affect the power and when zero power passes through the filter; the cosine term is used as a smooth cut off to avoid ringing in the filtered time series. A problem arises when k_1 is too low - it can cut into the power from the astronomical source. The lowest acceptable value for k_1 is given by the effective beam size in time, which in turn comes from the angular speed of the scan divided by the beam size, or $(2\pi \times 2.2^{\circ})/(48 \text{ s} \times 0.083^{\circ}) = 3.5 \text{ Hz}$. Any filtering occurring around or below this frequency actually reduces the astronomical signal in the final map, which affects its use as a source of absolute calibration.



Figure 4.6: Raster Time Series Filter Function

This plot shows the low-pass filter applied to each scan of the raster map data. The filter is defined by Equation 4.7, where $k_1 = 3$ Hz, and $k_2 = 7$ Hz. The inset shows the same function on a logarithmic scale.

In terms of final calibration raster maps, the simple median filtered map is used in preference to the other methods. This gives the highest fidelity on the source characteristics, despite the presence of atmospheric noise. However, for the purposes of presentation or when a 'very clean' map is required, some combination of the methods discussed above is applied to the maps, despite their possibly detrimental effects on the astronomical signal in the maps. As an example of this, Figure 4.7 shows a map made using a low pass filter, followed by the subtraction of a third order polynomial fit, followed by a linear fit removal on the individual map rows.



Figure 4.7: RCW 38 Raster Map with Advanced Filtering

This is a raster map of the same single-bolometer data shown in Figure 4.5, but with advanced filtering techniques applied. First, the low pass filter defined in Equation 4.7 is applied and then a third order polynomial fit is subtracted from each scan. The data is then binned according to Equation 4.5, and following this a linear fit is subtracted from each row, excluding the data from the gain reset at the end of each scan and central source in the fit. This produces the logarithmic scale map shown here, which shows a great deal more structure than median filtering brings out.

4.3 FEED OFFSETS

Noise theory tells us that a single detector observing for time t has an error that reduces as $1/\sqrt{t}$. Thus, observing a region of the sky for twice as long only improves the error in the measurement by $1/\sqrt{2}$. A way around this is to add more detectors which make independent measurements of the sky; if the number of detectors added is n, the improvement in the error goes as $1/\sqrt{n}$ (as adding independent detectors is the same as an increase in the integration time). This effect was used with QUaD; instead of a single detector, QUaD uses ~ 15 detectors at either frequency. Because the detectors are feed-horn coupled to the optics, the number of detectors is dictated by the physical space available on the focal plane. Also because of these physical constraints, QUaD's focal plane is laid out in the regular, close-packed pattern shown in Figure 4.8, and photographed in Figure 2.10. Based on quasi-optical simulations, the expected projection of the focal plane onto the sky is known

4.3. FEED OFFSETS



to high accuracy. However, these positions must be measured in the field.

Figure 4.8: Predicted Focal Plane Projection onto the Sky

This plot shows the predicted projection of the QUaD focal plane onto the sky. The blue circles mark the 150 GHz bolometers and the red circles mark the 100 GHz bolometers; both populations are drawn with radius equal to the predicted FWHM of the beams at either frequency. The black crosses on each horn position show the design orientations of the co- and cross-polar devices at the bottom of each feed horn, each of which is labeled with its number. The beam centres are actually separated by ~ 6 beam widths due to the physical size of the feeds compared with the available space on the focal plane. This means that QUaD is not exactly a 'camera' class instrument, as some form of dithering or cross-scanning is required to fully spatially sample the array. (Plot courtesy M. Bowden.)

The method we employ to measure the focal plane's image on the sky is to make a map of RCW 38 for each detector using one of the methods discussed above. In this case,

the absolute flux from the source is not particularly a concern so virtually any of the filtering methods discussed above is applicable; rather, what matters is that the source has a clearly defined peak. Coble et al. (2003) notes that the angular extent of RCW 38's main lobe is about 50", which is sufficiently smaller than the typical QUaD beam size to be considered point–like. Once the map is made, the apparent position of the source is archived; this is the array offset for each bolometer.

Although one day-long mapping measurement is sufficient to determine the array offsets, a single measurement is prone to large errors. We therefore perform this measurement a number of times (\sim 5) at the beginning of each observing season, and again approximately once per month throughout. The exact dates and rotation angles of the telescope at each 'total coverage' RCW 38 map during the 2005 observing season are listed in Table 4.3. These total coverage observations provide a set of data for the entire observing season which can be used to determine the array offsets, beam performance, *et cetera*. The array offset data are kept in look-up tables, and are plotted in Figure 4.9.

Table 4.3: RCW	' 38 Map	ping Run	Dates	& Parameters
----------------	----------	----------	-------	--------------

Observation Date	ϕ (degrees)	$T_{\rm ext}$ (C)			
March 12, 2005	+30	-55.0			
March 13, 2005	-30	-54.3			
March 17, 2005	0	-59.7			
March 19, 2005	15	-64.6			
March 20, 2005	-15	-60.5			
April 10, 2005	0	-60.2			
April 11, 2005	+30.0	-57.8			
June 11, 2005	0	-65.0			
June 12, 2005	0	-62.2			
Re-focus Occurred Morning of July 10, 2005					
July 10, 2005	+30	-52.6			
July 12, 2005	0	-71.0			
July 13, 2005	-30	-68.6			
August 14, 2005	0	-70.8			
September 9, 2005	+30	-61.2			
September 27, 2005	0	-57.8			

The data plotted in Figure 4.9 require a great deal of processing before they can be used to determine the best array offsets. The first step, which has already been applied to the data in Figure 4.9, is to fit a sine wave with variable amplitude, frequency, phase and offset to both the α and δ raw data points. In this model, the amplitude corresponds to the radius of each ring of bolometers, the frequency naturally falls to 2π in the fit, and the phase is used to find the rotation angle ϕ of the array with respect to the global coordinate system. However, the first parameter of interest is the offset of the fits in α and δ . The mean



Figure 4.9: Raw Array Offset Data for the 2005 Season

This plot shows the raw array offsets for the 2005 observing season corrected for pointing error, but not rotation *et cetera*. A great deal of processing is required to get to this point, and from this point to the best fit array offsets.

of these offsets over all bolometers is subtracted from each offset for each bolometer. This corrects the data for random pointing errors (in principle, one could use pointing crosses for this, but the circular fits to the offsets yield higher signal to noise ratios on the offsets for that observation day).

Because the RCW 38 maps are made at a variety of telescope rotation angles, the offsets must be put into a common frame before they can be averaged. This is done by converting the α, δ offsets into spherical coordinates, reading the telescope rotation angle encoder values and adding them to the raw ϕ values, and converting back to α, δ offsets.

The data points then overlap each other to the level of the noise, as shown in Figure 4.10. The average of the offsets can than be calculated; these averages (green triangles in Figure 4.10) are then the nominal offsets for each bolometer.



Figure 4.10: Array Offsets for the 2005 Observing Season

This plot shows the array offsets for the 2005 season (which are also featured in Figure 4.10). The pre-refocus offsets are shown in blue, while the post-refocus offsets are shown in red (see Section 4.4.1 for a discussion of the refocus). Triangles denote the co-polar bolometer and diamonds the cross-polar bolometer in each horn, which are labeled. The blue-green triangles are the average position for each bolometer for the pre-refocus data, while the yellow-green triangles are those for the post-refocus data.

One subtlety is that for the 2005 season, a re-focusing was required to improve the performance of the telescope (see Appendix A.2 and below). During analysis, it was noted that the array offsets actually changed at the time of the focus. We have therefore broken up the offset data into pre-refocusing and post-refocusing epochs; these are plotted separately in Figures 4.10 and 4.11. In addition, the 2006 season data will have to be analyzed with this issue in mind. This behaviour comes as somewhat of a surprise from the quasi-optical simulations, which suggest that mapping of the focal plane to the sky is not a strong function of the primary–secondary mirror distance. However, as shown in Figure 4.11, this variation is certainly present in the data. This figure also shows the mean radii of the bolometer populations. Looking at the data for the inner and outer rings of 150 GHz bolometers, it appears as though the projection of the array onto the sky is actually elliptical at a small level (this comes from noticing the small sinusoidal variations in the bolometer positions with respect to the mean radius of a given ring). The size of the ellipticity is approximately 0.6 % for the inner ring and 6 % for the outer ring of 150 GHz bolometers.

A final use for these data is to calculate the array rotation with respect to the global coordinate system. This is given by the fitted ϕ in the sinusoidal α , δ fits discussed above. It is found that the average array rotation between the line joining horn 150-1 with 150-2 and the line $\Delta \alpha = 0.0$ is $+122^{\circ}$ based on these data.

4.4 BEAMS

A telescope's beam is defined to be the output of the optical transfer function of the system when a delta function is input, which is to say, it is the (two dimensional) optical impulse response function of the optical system. Measurement of this quantity is very important in telescopes, as the beam dictates how astronomical flux distributions are 'seen' by the system. In CMB experiments such as QUaD, this is doubly important because the fidelity of CMB maps produced by a telescope directly affect the power spectrum which results (see Section 6.3). Only a precise understanding of the transfer function allows the appropriate beam filter corrections to the power spectra to be made.

As discussed in Section 4.1.1, sources available in the mm regime tend to be poor for precise beam measurements. An example of an ideal beam measurement source is a bright quasar; these are certainly point sources and could therefore provide high fidelity measurement of the beam to very low levels¹⁰. However, even the brightest quasars (\sim 1 Jy at 100 GHz) are too dim to provide a good measurement of the beams in a single

$$P_{\rm dB} = 10 \log_{10} P/P_0,$$

¹⁰In this context, the beam response is generally measured in deci-Bels, or dB, which is defined to be

where P_0 is the peak response of the beam function (0 dB). A reasonable measurement of the beams is to ~ -20 dB, or 1 % of the peak response, and a deep measurement of the beams would be to, say, < -30 dB, or 0.1 % of peak response.



Figure 4.11: Array Offset Radial Structure

This plot shows an 'unwound' version of the QUaD focal plane which highlights structure in the radial direction. Bolometer number refers to the scaling of bolometers in which the central two are 0 & 1, the inner ring of 150's are 2–13, the outer ring of 150s are 14–37, and the 100s are 38–61. Co-polar and cross-polar devices sharing a horn are even and odd in this numbering scheme. In this plot, the blue crosses are the pre-refocus array offset points and the red crosses are the post-refocus array offset points in terms of radial distance - any ϕ information is lost in this type of plot. The average for each bolometer is plotted as a blue (pre-refocus) or red (post-refocus) diamond. The mean radial distance for each population is plotted as a blue or red dashed line. Two structures are evident in this plot: (i) the refocus (see Appendix A.2) changed the average radial distance, more so for the outer bolometers than for the inner and (ii) both of the 50 GHz bolometer rings are mildly elliptical (this comes from the sinusoidal structure about the mean for these two populations). observation. This means that to use these sources for beam measurements, a series of observations must be stacked. This measurement has been performed, but because the beam can be time-varying, the average over a season should not be trusted (see Section 6.2.1). We are therefore driven to consider other options for beam measurement.

One alternative to beam measurements using dim point sources is to perform this type of measurement in the lab. Indeed, the QUaD horn response functions were measured in the lab, and agree with modeling to beyond the $-35 \,dB$ level O'Sullivan (2006). However, even measurement of the beams of only the optics contained in the cryostat is unwieldy in the lab; measurement of the whole telescope's impulse response function using a laboratory set up is virtually impossible¹¹. One could also use a source in the far field limit, such as a Gunn diode on a balloon at a suitable distance, but this type of measurement is too difficult given QUaD's available manpower and economic constraints. We are therefore forced to use astronomical sources which are not ideal for beam mapping.

As in other calibration observations, the object we have settled on as our beam mapping source is RCW 38. The small angular extent of the main lobe of the source (~ 50") coupled with its brightness make it a reasonable choice for measurements down to the ~ $-15 \, dB$ level. Beyond this, the diffuse background structure in the field severely affects the beam measurement. Because of this, the position we find ourselves forced into is one where no measurement of the main beam exists below about $-15 \, dB$, which is certainly not ideal. This uncertainty becomes an issue for $\ell > 1500$, and for our *B*-mode polarization measurement at all ℓ s.

Because both array offset determination and beam maps require a full map of the source for each bolometer, total coverage RCW 38 maps are used for both types of measurement. The reader is referred to Section 4.3 for a discussion of both the observation strategy and the frequency of these types of observations during the austral winter; Table 4.3 gives a list of dates, telescope ϕ angles, and average external temperatures for each RCW 38 mapping observation performed over the 2005 winter season.

These RCW 38 beam maps produce data typified by those shown in Figure 4.12. The central lobe in these maps is fit to a Gaussian function with variable offset I_0 , amplitude I_0 and ellipticity (both amplitude and angle):

$$I(\alpha, \delta) = I_0 + I_1 \exp\left(\frac{(\alpha'/\sigma_\alpha)^2 + (\delta'/\sigma_\delta)^2}{2}\right)$$
(4.8)

¹¹In principle, one can perform this experiment using a source imaged with the identical optical system at a relatively short distance; this result can be derived from the Rayleigh–Carson Reciprocity Theorem. However, this experiment would require at the very least a copy of the primary and secondary mirrors, and as such is difficult for purely economic and logistical reasons. That being said, the BOOMERanG experiment successfully performed this measurement at CalTech (K. Ganga, private communication).

and the primed coordinates are related to the true α and δ by a rotation:

$$\begin{aligned} \alpha' &= (\alpha - \alpha_0)\cos(\theta - \phi) - (\delta - \delta_0)\sin(\theta - \phi) \\ \delta' &= (\alpha - \alpha_0)\sin(\theta - \phi) + (\delta - \delta_0)\cos(\theta - \phi). \end{aligned}$$
(4.9)

In these equations, $2\sigma_{\alpha}$ and $2\sigma_{\delta}$ are the lengths of the two axes, α_0 and δ_0 are the centroids of the point spread functions (PSFs), θ is the angle of the ellipse with respect to zero telescope rotation ($\phi = 0$), and ϕ is the rotation angle of the telescope. Note that throughout we switch between σ and full width at half maximum (FWHM) via

$$FWHM = 2\sigma\sqrt{2\ln 2}.$$
 (4.10)

The centroid positions are actually those used for the 'raw' measurements discussed in Section 4.3, although a slightly different statistic could be used in their place. As they parametrize the beam, the σ_i and θ are of interest for characterization purposes. In the next section, we present the characteristics of the beam these measurements revealed, and discuss the evolution of the QUaD beam which occurred as our understanding of the instrument grew.

4.4.1 BEAM ELLIPTICITY & TEMPERATURE VARIATION

Using the RCW 38 maps, it was noted as early as March 2005 (that is, just after the commissioning summer ended) that the beams as measured on the telescope might be significantly asymmetric. This hypothesis was confirmed as the season drew on, and measurements taken in constant conditions showed that the ellipticity result was repeatable, time varying, and large compared to the beam size. Unfortunately, full focusing characterization maps are very time consuming to create (1 observation day per full array map), so a proxy for these measurements which takes very little time was of interest. It so happens that part of the daily QUaD calibration observation set are 4 'row calibrations', discussed in Section 5.4. These involve scanning each row of the QUaD array over RCW 38 to produce two blips in each bolometer's time series, in the same way as is performed for the central bolometers in a pointing cross observation. These blips can be extracted using the same algorithm as is applied to the cross observations (again adaptively choosing the atmospheric monitor channels to ensure they themselves don't contain significant astronomical signal). Once the blips are found, they are fit to a one dimensional Gaussian distribution as in Equation 4.1, and the σ statistic is stored. This statistic can be plotted as a function of time, as in Figure 4.13.

Figure 4.13 exhibits a great deal of information, and three observations about these

12111



Figure 4.12: Example Beam Maps, Pre- and Post- 2005 Refocus

These images show beam maps for bolometers 150-1-A, 150-5-A, 150-14-A and 100-7, in order from left to right. The upper row shows pre-refocus maps from data taken June 12, 2006, and the lower row shows post-refocus maps from data taken July 12, 2006. The colour range is the same in all plots (from 0 to ~ 1.5 Jy per $1.2' \times 1.2'$ pixel). The array shows ellipticity, much worse at 150 GHz than 100 GHz, which is correlated along the same angle. Because of the warped primary mirror, this points to a focusing error; after refocusing the telescope, the ellipticity was dramatically reduced. In addition, the amplitude of the signal from RCW 38 was increased after the telescope refocus, as shown by the brightening in the lower maps.

plots are of immediate and obvious interest. The first of these is the fact that there is a great deal of time variability in the beam widths. This is strange, and somewhat worrying, as it requires time variation in the optical elements themselves. The second point of interest is that, for any ϕ angle, the measured beam widths are significantly broader than expected based on simulations of our optical system and the angular size of RCW 38. The third strange feature in these plots is that the FWHM for the two ϕ angles used in these measurements are significantly separated. This points to elliptical beams; worryingly, the numbers in these plots suggest a factor of two in the length of the semi-major to semi-minor axes¹².

Interestingly, a single fault with the system can explain all of these features. The first hint as to the cause of the problem is that the beam widths are most likely to be much broader than expected if the telescope were out of focus. This was a plausible situation for QUaD at the beginning of the 2005 observing season: the method used to focus the telescope was to adjust the secondary mirror's position to a static value during commissioning in the summer. This was done using some predefined calibrated rule, which did



Figure 4.13: Beam Width Time Variability

The plot on the left shows the variation in the QUaD beam over the first four weeks of the 2005 CMB observation season. The blue-hued points show the beam width for the co- (light blue) and cross- (dark blue) bolometers in horn 150-2 for $\phi = +57^{\circ}$. The red-hued points show the same for the co- (orange) and cross- (red) polar bolometers for $\phi = -3^{\circ}$. Triangles denote forward-going scans, and diamonds backward-going scans. The two dashed lines show the predicted FWHM measurements based on two measurements of RCW 38's width at these frequencies. Clearly, during the first part of the season, the beam performance was nowhere near the optimal prediction. The beams are also highly elliptical, as shown by the strong ϕ angle dependence of the widths. The plot on the right shows the telescope's external temperature as a solid black line. Also shown is the external temperature multiplied by -1 and scaled to show that it is strongly anti-correlated with the measured beam widths.

not necessarily reflect the as-built parameters of the telescope¹³. In any case, the precise primary-secondary mirror distance of the telescope was not well known, at least to the sub-mm tolerances required by the system.

Also, simulations showed that, due to the warp in the primary mirror caused during the manufacturing process, a focus error would also cause ellipticity in the beam. This ellipticity was quantitatively calculated using our optical models, and found to be in good agreement with the observed value if the secondary were too close to the primary. Figure

¹³Combined with the failure of the active secondary mirror control and lack of knowledge of how much the foam cone would shrink as the ambient temperature on the outside of the telescope dropped, it is incredible that the telescope was so close to the correct focus.



4.14 shows the results of the calculation of beam width versus primary-secondary separation for the two beam axes; these are in remarkably good agreement with the data.

Figure 4.14: Calculated Beam Widths and Primary–Secondary Mirror Distance This plot shows the results of quasi-optical calculations to determine the effect of focus on the QUaD beam widths during the 2005 season. Shown are the results for the central 150 GHz bolometers, the outer 150 GHz ring, and the 100 GHz bolometers. The solid lines show the FWHM in the axis parallel to the warp in the primary, while the dashed lines show the FWHM in the axis perpendicular to the warp. Without a warp, these two lines would overlap, that is, the beams would not be elliptical for any primary–secondary separation. However, the warp in the primary induces a preferred direction in the optical system which leads to the behaviour in Figure 4.13. Before the refocus in July 2005, the primary–secondary separation was at the position shown as a black dotted line; following the refocus it moved to the position shown as the black dashed line. The total change in distance was ~ 1.7 mm.

In addition, the time variability of the beam width itself can be understood in terms of the combination of focusing error and the expansion or contraction of the foam cone under different ambient external temperatures. That is, the foam cone should contract as the external temperature drops, and expand as the external temperature rises¹⁴. This would have the effect of bringing the telescope into better focus in warm temperatures, and pushing it out of focus in colder temperatures, *i.e.* that the beam width and external temperature should be anti-correlated. Figure 4.15 shows the points plotted in Figure 4.13 with a fit to the external temperature for both ϕ angles; this is an extremely good description

¹⁴This makes the dotted line in Figure 4.14 'slide' back and forth along the separation axis by up to 1 mm in either direction. It was not known *a priori* how much the cone would shrink under the ambient conditions at the pole.

CHAPTER 4. INSTRUMENT PERFORMANCE



Figure 4.15: Anti-Correlation of Beam Width and External Temperature This plot shows the same data featured in Figure 4.13 (symbols and coloured lines are the same as in that figure). Over plotted is the fit to the external temperature for both ϕ angles; these two quantities are clearly strongly anti-correlated. This time variability complicates the analysis of the CMB data from this period.

Based on these pieces of evidence, it was decided to refocus the telescope in July 2005 (as discussed in Appendix A.2). This involved moving the secondary mirror 1.7 mm further from the primary to put as many beams as possible near the minimum in the focus

function¹⁵. Following this refocus, the beam widths behaved as predicted from our modeling. The widths were far closer to their optimum size and zero ellipticity, although they were still somewhat larger then predicted in the symmetric primary case. Unfortunately, the ellipticity did not vanish, but reduced by up to 40 % in the 150 GHz horns, which were the most affected by the warp in the primary mirror. Table 4.4.1 gives the FWHM of the semi-minor axis and the ratio of the axes of the beams both before and after the 2005 refocus.

Table 4.4: Beam Statistics for 2005 Observing Season					
	Pre-refocus		Post-refocus		
	Semi-minor	$\sigma_{ m b}/\sigma_{ m a}$	Semi-minor	$\sigma_{ m b}/\sigma_{ m a}$	
	FWHM (arcmin)		FWHM (arcmin)		
Central Pixel	4.21 ± 0.005	1.76 ± 0.05	4.02 ± 0.000	1.27 ± 0.04	
Inner 150 GHz	4.30 ± 0.006	1.81 ± 0.06	4.09 ± 0.006	1.30 ± 0.03	
Outer 150 GHz	4.16 ± 0.01	1.49 ± 0.10	4.16 ± 0.008	1.13 ± 0.04	
100 GHz Ring	5.31 ± 0.008	1.26 ± 0.04	5.43 ± 0.008	1.08 ± 0.02	

The most important effect of the re-focus was to set the primary-secondary distance near its minimum in both beam axes. This makes the derivative of the mirror separation-FWHM function small, so that the time variation of the beams is negligible. This makes the CMB data much simpler to analyze; Figure 4.16 shows an example result of the refocus on the measured RCW 38 blips widths.

In order to apply proper correction to the power spectra during the final analysis, a beam model must be constructed which faithfully reproduces the day to day variation of the beam width as a function of θ angle and the ambient temperature. This is performed by using time variation data like those shown in Figure 4.16 for each bolometer, separately for the pre- and post-refocus data. Based on the θ angle of the telescope during the observations and assuming that the ellipticity of the beam is constant, the beam widths are extrapolated to the widths of the semi-major and semi-minor axes which would have been measured at the time of observation via a Cartesian rotation of the measured widths. This extrapolation is checked against the full map observations, and found to reproduce the full beam shape within 10 %. These semi-minor axes measurements are then fit to the ambient temperature using a linear model

$$FWHM_b = B_1 T_{ext} + B_0 \tag{4.11}$$

where FWHM_b is the FWHM of the semi-minor axis, T_{ext} is the ambient external temperature, and the B_i are the parameters of the fit. This procedure yields an estimate of the beam

¹⁵For a classical Cassegrain telescope, the beam width variation as a function of the primary-secondary mirror separation describe a parabola in FWHM-separation space. The minimum of this parabola is the best focus.



Figure 4.16: Typical Beam Performance for Entire 2005 Observing Season This plot shows the beam performance for the bolometers in horn 150-2 for the entire 2005 CMB observing season (the data shown in Figures 4.13 and 4.15 are the same as these, but the time axis in this plot is much longer). The days around the refocus are delineated by the two vertical dashed lines. Clearly, before the refocus the beams were very poor, but following this, the minimum size is much closer to that predicted by simulations, and the ellipticity is much reduced. More importantly, the post-refocus data are uncorrelated with the external temperature. This particular horn had one of the smallest improvements in the whole focal plane, so these particular data represent a worst case scenario improvement.

shape as a function of the ambient external temperature. The fit parameters for the four different bolometer populations are given in Table 4.4.1. This table shows that the beams were a very strong function of the ambient temperature before the refocus. However, after the refocus they become a weaker function of the temperature, and are much stabler throughout. The data plotted in Figure 4.16 exhibit an example of this.

Table 4.5: Beam Model Coefficients for the 2005 Observing Season						
and the second states in some	Pre-refocus		Post-refocus			
	B_1	B_0	B_1	B_0		
	arcmin / C	arcmin	arcmin / C	arcmin		
Central Pixel	-0.0411 ± 0.0044	2.10 ± 0.28	-0.0059 ± 0.0014	3.63 ± 0.03		
Inner 150 GHz	-0.0402 ± 0.0056	2.28 ± 0.35	-0.0067 ± 0.0057	3.77 ± 0.39		
Outer 150 GHz	-0.0176 ± 0.0353	3.63 ± 2.27	-0.0095 ± 0.0075	3.92 ± 0.45		
100 GHz Ring	-0.0041 ± 0.0068	5.51 ± 0.51	-0.0016 ± 0.0010	5.70 ± 0.12		

The beam model any any given time is produced by reversing this procedure. First, the ambient external temperature is used to predict the semi-minor axes of the beam, and this is used to produce an elliptical Gaussian kernel assuming that the ellipticity is not a function of θ^{16} . This kernel is then rotated based on the θ angle of the telescope, and the beam model for each bolometer is fed back to the user. Unfortunately, due to the large variation in the row calibration determined beam widths, these beam models are somewhat uncertain, particularly in the pre-refocus data where the variation is the strongest.

To remove this focusing issue entirely, two modifications were made to the secondary mirror unit during austral summer 2005/2006. The first was to implement remote computer control of the primary–secondary mirror distance to allow immediate refocusing. The second was to replace the original hyperbolic secondary with a shaped mirror designed to counteract the effect of the warp in the primary mirror. Thus far, this system has performed very well during the 2006 observing season; this system is discussed in detail in Appendix A.3.

4.4.2 **BEAM DIFFERENCES**

CMB polarization measurements require precise knowledge of the beams to separate cleanly the T, E and B mode power. This is because, if the two beams do not perfectly match one another, the power in their residual will be interpreted as polarized power on the sky. This effect is particularly insidious at high ℓ , where the power spectrum is most sensitive to the precise shape of the beams. The effect of beam differences on the actual power spectra is discussed in Section 6.3; here we show the data which defines the level of the beam mismatch.

To measure the beam mismatch, an unpolarized far-field source is required. As discussed in Section 4.1.1, bright, compact and completely unpolarized sources are vanishingly rare at mm wavelengths (in reality, all astronomical sources will be polarized at some level in any case). Fortunately, work like that presented in Kane et al. (1993) shows that compact HII regions are not highly polarized, particularly for instruments like QUaD where the beam size can serve to depolarize the source¹⁷. RCW 38 itself was measured by BOOMERanG in 2003, and is found to be less than 6 % polarized at 2σ (Masi et al., 2005). We can therefore difference the standard QUaD beam maps with some confidence that, if the maps are properly calibrated, the residual should be less than 6 % of the unpolarized

¹⁶This assumption was checked, and the results were inconclusive. This assumption may therefore be a source of systematic error in the final power spectrum analysis.

¹⁷This occurs because the beam integrates over the polarization incident on it, so even if the compact HII regions have complex polarization structure on small scales, this would tend to average to zero in a QUaD beam.

flux.

The co-polar/cross-polar difference maps for the data shown in Figure 4.12 are shown in Figure 4.17. The structure observed in these maps is perplexing: polarization in the source should appear as an overall residual over one beam; the type of substructure seen here cannot be due to RCW 38 itself. In fact, the dipolar and quadrupolar shapes in the residuals point to some asymmetry in the optical system. The typical peak amplitude of the residual signal is $\sim 5\%$ of the peak temperature signal, so the differences between beams are not small. We therefore conclude there is some fault in the instrument itself.



Figure 4.17: Example Beam Difference Maps

These maps show the A-B bolometer pair differences for horns 150-1, 150-5, 150-14 and 100-7 (in order, left to right). The upper plots show pre-refocus data (taken June 12, 2006), and the lower plots show post-refocus data (taken July 12, 2006; these are from the same sets as are imaged in Figure 4.12). The colour range is $\pm 4\%$ of that in Figure 4.12; clearly, the beam residuals are much larger than the noise, and are in fact not small compared to the amplitude of the undifferenced beam. The fact that these differences are repeatable before and after the refocus suggests that they are caused by an optical element in the cryostat itself.

Further evidence for this mismatch comes from the model fits. Interestingly, elliptical Gaussians fit to the RCW 38 maps, when differenced, show similar features to those in these (data) maps. In fact, the ratio of major axis to minor axis of these fits is systematically bigger for the co-polar devices than the cross-polar devices at the $10^{-3} - 10^{-4}$ level. This suggests that one polarization's field is affected more than another as it propagates through the instrument. While this may be obvious when integrated over the beam (this is a type of instrumental polarization), this polarized *sub-structure* in the beam is much harder to explain. There are two immediate suspects for the cause of these beam mismatches. The first is the warp in the primary, which certainly causes an asymmetry in QUaD's optics (see Section 2.1). However, feed horns should ensure that the radiation patterns at the detectors are not sensitive to this type of optical aberration (B. Maffei, private communication). The second possibility is birefringence in the polyethylene lenses. This asymmetry is most likely a characteristic of the plastic imprinted as it cooled, and as such is very difficult to quantify or properly account for without careful laboratory measurements of the optical components in question. The combination of the high curvature of the lenses and their birefringence could affect the polarization characteristics of the radiation fields propagating through the lenses at a low (but non-negligible) level. At this point, it is not possible to fix this beam mismatch at the hardware level; ultimately, this affects QUaD's ability to measure faint B modes. This type of optical effect is going to be a major concern for the next generation of CMB polarimeters.

4.5 SIDELOBES

The term 'sidelobe' is used in antenna theory to describe any part of the antenna pattern of a receiver which is not in the main beam. The power density in sidelobes is generally much less than that in the main lobe, but may still have an effect on output signal from the receiver, particularly if a bright source is situated in the sidelobe. Sidelobes can be caused by a number of different optical effects, the most common being diffraction rings from the telescope's Airy pattern, spill over from an over–illuminated secondary, or stray reflection from telescope components in the main beam (the secondary mirror support structure, for example). In radio/mm telescopes, measurement and characterization of the sidelobes is crucial to minimizing systematic effects in the data.

In QUaD, a great deal of time and attention has been devoted to minimizing the existence and effect of sidelobes. For example, the optical path from bolometers to sky contains the following sidelobe-reducing elements:

- The feed horns apodize the illumination pattern of the primary and secondary mirrors.
- A (4 K) Lyot stop is used to restrict the effective aperture of the feed horns.
- The mechanisms behind the secondary mirror are completely hidden from the main beam.
- The cryostat snout is hidden behind a reflecting baffle which rejects rays falling between the cryostat window and the primary mirror out of the optical system to the

sky.

- The primary mirror is ringed by a 30 cm width radiation skirt to reduce the effects of secondary spill over.
- The secondary support structure is made of uniform foam which does not reflect in the mm.
- The entire telescope is surrounded by a reflecting ground shield which rejects rays falling onto it to the sky.

All of these elements help to create an optical system designed to be clean to below at least -30 dB.

Unfortunately, during QUaD's first observing season, it became clear that some type of scan-synchronous pick-up was affecting the data. This pick-up manifested itself as a strong signal associated with the azimuthal position of the telescope. Typical correlation coefficients between the time series and the azimuth encoder approached 1.0, although the amplitude of the signal appeared to be time-varying in a completely random way. Scans like those shown in Figure 4.18 solidified our picture of this pick-up signal.

The simplest explanation for these signals is pick–up from the ground. Indeed, the signal plotted in Figure 4.18 seems to comprise two major components: a single-cycle signal due to a tilt in the telescope mount¹⁸, and a twelve-cycle signal which in some bolometers completely dominates the single-cycle signal. Twelve cycles is an interesting number, as it corresponds to the number of panels in the ground shield. However, it is not clear what aspect of the ground shield could contribute such a large signal to the data.

This signal's apparent connection to the ground shield suggests that a stray sidelobe is actually sensitive to the shield during normal telescope operation. The amplitude of the pick-up signal can be as large as 0.2 K at 100 GHz and 0.9 K at 150 GHz; this level of pick-up requires extremely large sidelobes. For example, ~ 250 K snow builds up in the joints between the ground shield panels, and may be the cause of the stray radiation¹⁹. We can perform a back of the envelope calculation to assess the level of signal we expect from this sidelobe/snow combination. Assuming the snow is 20 % emissive, and that the typical

¹⁸This causes a sinusoidal signal because of optical loading from the atmosphere: as the telescope turns on a tilted track, the signal is modulated by the airmass through with the receiver observes the background. Unfortunately, tests at the south pole were unable to account for the full amplitude of this sinusoid based on the atmospheric loading and the known size of the mount's tilt. This either means the canonical track tilt is erroneous, or some other source is contributing to this signal.

¹⁹Although we have embarked on a program to determine if this is, in fact, the source of the pick-up, the results are thus far inconclusive.



Figure 4.18: Ground Scans with Sidelobe Pick-Up

This plot shows the results of ground scans used to characterize the shape and amplitude of the ground pick–up during the 2005 CMB season. To produce these plots, the telescope was scanned 360° in azimuth at constant elevation. The plot on the left shows the time series for the 100 GHz bolometers, and the plot on the right the time series for the 150 GHz bolometers. Bolometer pairs are plotted in the same colour, and the colour scale runs over all the bolometers at a single frequency (arbitrary DC offsets have been applied to the various time series to correct for the random voltage offsets between channels). The *y* axis on these plots are in Volts; to convert to K, these values need to be multiplied by 0.52 K/V at 100 GHz, and 0.44 K/V at 150 GHz. The single-cycle sinusoid is thought to be caused by the tilt of the telescope mount, while the short wavelength 12-cycle structure is thought to be caused by some type of pick–up originating from the ground shield. (Figure courtesy C. Pryke.)

pick-up feature is 0.09 K at 150 GHz, this leads to sidelobes at a level of

$$P = \frac{0.09 \,\mathrm{K}}{250 \,\mathrm{K} \times 0.2} = 0.0018 \tag{4.12}$$

or $-27 \, dB$, integrating over the entire sidelobe. For the worst cases, this signal is actually as much as ten times as large as this estimate. In any case, it completely dominates the CMB anisotropy signal, and must be understood to minimize systematic errors in the data.

During the 2005/2006 summer work at the south pole, a Gunn diode tuned to 115 GHz or 150 GHz was used to map the sidelobes of the telescope. The experimental procedure for this measurement is discussed in Appendix A.3. Essentially, the measurement involved placing the Gunn source high on the ground shield, and mapping the region around the source with the telescope. The resulting data yields a map of the sidelobes for all of the bolometers; Figure 4.19 shows these maps for one bolometer at either frequency



(a more detailed presentation of these kind of maps can be found in Appendix A.3).

Figure 4.19: 2005 Season Sidelobe Maps

This figure shows maps of the sidelobes as measured by bolometers 150-10A (left) and 100-4A (right). The source was positioned at $el = 26^{\circ}$ during these measurements. The colour scale is essentially arbitrary; because the main lobe (the large blob near (-52, 26)) was heavily saturated, it is difficult to determine the calibration from Volts to Power in this measurement. In any case, 0 dB corresponds to a signal of 1 V per square degree pixel. The structure in the radial spike is not due to the sidelobe, but rather to the fact that a Gunn diode is a monochromatic source.

A great deal of structure is apparent in these sidelobe maps. The (saturated) main lobe is clearly visible at the bottom centre of the plots at both frequencies. Because these maps are made from near field measurements, extrapolating the far field response of the main lobe from them is fraught with difficulty. Therefore, little importance is given to the shape of the main lobe in these maps. However, structures outside of the main lobe are more meaningful. Considering the 150 GHz channels first, there appears to be a ring around the main lobe with a radius of about 25° in all of the 150 GHz bolometers, which is present but much reduced in the 100 GHz bolometers. Secondly, the outermost ring of 150 GHz bolometers appears to have a radial spike extending from the centre of the source out to beyond zenith (which is approximately 65° away from the source at these elevations). The structures in these maps are calculated to have amplitudes of approximately -60 dB / QUaD beam, which when integrated over the area covered by the sidelobes, yields numbers in the range -20 to -30 dB. It appears as though these sidelobes are indeed the cause of the ground shield pick-up in the QUaD data.

4.5. SIDELOBES

The physical causes of these two sidelobes are in one case simple, in the other subtle. The simply explained sidelobe is the ring, which is caused by errors in designing and manufacturing the baffle. As the baffle is an optical element whose characteristics strongly affect the sidelobes, these problems conspired to create an unexpected sidelobe.

More subtly, the radial sidelobe in the outer ring of bolometers is caused by diffraction off the edges of optical elements in the cryostat. Although the optics were manufactured to specifications, those specifications were too small, and some rays intercept the filter rings on the 77 K stage. This causes some of the main lobe power to be diffracted, which in turn causes the long radial spikes in the power pattern of the telescope. Fortunately, it is possible to restrict the angles available to these sidelobes, and thereby remove their effect on the receiver. This was performed by using an extended, blackened baffle; Appendix A.3 gives a full discussion. Upon performing sidelobe maps with this extended baffle in place, the sidelobes were reduced or absent, as shown in Figure 4.20. This gives good evidence that the cause and nature of the sidelobes are understood.



Figure 4.20: Fully Baffled Sidelobe Maps

These plots show sidelobe maps for the same bolometers featured in Figure 4.19, but with extensive baffling present on the receiver's snout. The colour stretch and Gunn source position remain the same as in Figure 4.19. Clearly, all of the structure has been eradicated from the 100 GHz channel, while the radial sidelobe is still present in the 150 GHz outermost ring. However, because the allowed incident angles have been reduced, this radial spike only extends to $\sim 14^{\circ}$ from the main lobe (any further restriction and rays propagating from the primary to the secondary mirrors would be intercepted by the baffle).

Unfortunately, the loading from the blackening on the new baffle reduced the telescope's sensitivity by a factor of 2, and the decision was made to remove it for the 2006 CMB observing season²⁰. Instead, a reflective flat baffle was installed in the place of the 2005 version. This has brought the annular sidelobe in from a radius of $\sim 25^{\circ}$ to $\sim 10^{\circ}$, and concentrated its width somewhat. The CMB observation strategy has been tuned to ensure that this sidelobe never observes the ground shield or galactic plane, which means that the pick-up from this feature is no longer an issue. The radial sidelobe is more problematic. Because it only affects the outermost ring of pixels, we have chosen to do nothing and suffer from an unchanged pick-up amplitude but retain full sensitivity in these pixels. Various hardware solutions are being discussed for the 2007 observing season including changing the optics at the top of the cryostat's snout, although it remains unclear if a hardware solution will ever be implemented. In any case, the ground pick-up from the telescope's sidelobes remains in the data stream at some level. The preferred method of reducing the effect of these signals during data analysis is a combination of field-differencing and time series polynomial subtraction; discussion of these techniques is differed to Section 6.2.

4.6 SENSITIVITY & NOISE PERFORMANCE

The QUaD bolometers are designed to be dominated by background photon noise, which is directly related to the number of photons falling on the detectors (see Section 2.4.1). This, in turn, is proportional to the optical loading the telescope experiences in the field. This means that the ultimate sensitivity of the telescope, and thus the error on our measurement of the CMB, is directly related to the optical loading on the telescope at the south pole. In this section, the optical loading in the field and the projected noise level in the data are discussed, and compared to estimates of the noise derived from the atmospheric signal in the time series.

4.6.1 OPTICAL LOADING

The total optical loading incident on the bolometers can be determined by solving Equation 2.3 for Q, although as previously mentioned, converting this measurement to an absolute standard is fraught with difficulty. However, it is instructive to study the variation of the instrumental loading over time to understand how the sensitivity and NEP of the

²⁰The cause of this extra loading is simple; the sidelobes which had been falling on a structured source with temperature ~ 20 K now fall on a uniform source with temperature ~ 300 K. Because the total solid angle in the sidelobes is constant, the optical loading increases by a factor 300/20 = 15; although the sidelobes are now unstructured, they cause a great deal more power to fall on the focal plane, thus reducing the detectors' sensitivity by a factor $\gtrsim 2$.

instrument behave during the observing season. Hinderks (2005) presents a breakdown of QUaD's expected optical loading based on the properties of the materials used in the optics and telescope; the sum over all components including the atmosphere is expected to be $5.6 \pm 1.0 \,\text{pW}$.

The standard experimental tool used to measure the optical loading on the bolometers are load curves, which QUaD performs several times per day as part of the standard calibration block. The loading inferred from these load curves can be extracted and plotted as a function of time. Because many of the bolometers in the array have model parameters available, we can use the average of 20 bolometers at 100 GHz and 29 bolometers at 150 GHz to track the loading, as shown in Figure 4.21. This loading combined with known instrumental parameters can then be used to estimate the *NEP* of the instrument, as discussed in Hinderks (2005).

As this figure shows, the total loading variation is dominated by the variation in the atmospheric emission. When this variation is removed, a stable instrumental loading at a level of 4.0 pW (corresponding to 19 K) is found at 150 GHz. Tests performed during QUaD's commissioning while the foam cone and secondary were not in place showed that the loading from the optics alone (excluding the primary and secondary mirrors) was 1.7 pW, or 8 K at this frequency (as noted in Hinderks (2005)). Apparently, the foam cone contributes an extraneous loading of approximately 2.3 pW; based on the quoted transmissivity of 99.5 %, we would expect only about 0.25 pW from this source. This suggests that the foam used to construct the cone is optically poor; ideally, it should be replaced with a more transparent material. Fortunately, there is no evidence that this level varies over time due to water absorption into the foam or other degradation of the material.

4.6.2 NET, NEQ AND ATMOSPHERIC BEHAVIOUR

A complete model for the noise arising in the instrument alone is beyond the scope of this presentation; Hinderks (2005) gives a thorough analysis of this subject with respect to QUaD's performance. Here, we merely note that the instrument is designed to be atmosphere noise dominated, and discuss the implications of this fact.

Section 2.4.1 develops the expected NEP of the instrument, which is defined in terms of the power absorbed at the detector. Briefly, the NEP expected from optical loading of the size presented in the preceding section is given by $\sqrt{2Qh\nu}$; this is $\sim 3 \times 10^{-17}$ W Hz^{-1/2} at 150 GHz. This level makes it larger than Johnson and Phonon noise, and approximately equal to the amplifier noise in the system, giving a predicted total NEP around 6×10^{-17} W Hz^{-1/2} overall. As shown in Hinderks (2005), this prediction matches the system's performance very well.



Figure 4.21: Optical Loading Time Series for the 2005 Observing Season This plot shows the average optical loading over the 2005 CMB observation season based on load curves from 29 150 GHz bolometers. The black curve shows the raw loading, and the equivalent $T_{\rm RJ}$. The blue curve shows the loading from the atmosphere as measured by QUaD's once-daily sky dip measurements (see Section 5.1). When this is subtracted from the total loading, the instrumental loading (shown in red) results. This loading is constant at 4.0 pW over the entire season, which suggests that the instrument itself is very stable over long time periods. The dashed green line shows the result of loading measurements performed in the absence of the foam cone and secondary mirror; these suggest that the cone and telescope are adding some 2.3 pW of extra optical power to the system.

This being said, it is useful to characterize the noise performance in terms of the brightness of a source; to this end, one can define the noise equivalent temperature or NET. Similarly, for polarization measurements, one can define the NEQ (or alternatively and generally equally, the NEU) of the system. Both of these have units of $\mu K Hz^{-1/2}$, and take into account the throughput, band width and shape, and detector noise performance to yield a statistic which can be used in sensitivity and integration time calculations.

The first step in the determination of an NET for the system is to take a power spectrum of the time series to measure the NEP under real observing and noise conditions. Such a spectrum is shown in Figure 4.22 for time series data taken at zenith during testing in January 2006. This measurement is simple: with the telescope parked, the bolometers are run in their normal bias modes for some length of time. The white noise level in such

power spectra yield the NEP of the system in W Hz^{-1/2}. For the 2005 season, the median NEPs were 4.4×10^{-17} W Hz^{-1/2} at 100 GHz and 5.6×10^{-17} W Hz^{-1/2} at 150 GHz.

These *NEPs* can be converted to *NEQ* using an absolute calibration factor (discussed in Section 5.5) to change from Volts to Kelvin. The temporal unit conversion (from μ K Hz^{-1/2} to μ K sec^{1/2}) requires division by $\sqrt{2}$. The *NET* of a single detector gives the amplitude of the temperature fluctuation which could be measured in an integration time of 1 s with unitary signal to noise ratio. With a PSB pair, the integration time on temperature is effectively doubled, so $NET_{pair} = NET_A/\sqrt{2} = NET_B/\sqrt{2}$. Applying this to the measured bolometer power spectra gives an average NET of 551 μ K s^{-1/2} at 100 GHz and 530 μ K s^{-1/2} at 150 GHz for the 2005 observing season.

The NET can related to the NEQ using the following relation

$$NEQ^{2} = \left[\left(\frac{\partial Q}{\partial T_{\rm A}} \right)^{2} + \left(\frac{\partial Q}{\partial T_{\rm B}} \right)^{2} \right] NET^{2}$$
(4.13)

where T_A and T_B are the temperatures measured by the two bolometers, that is:

$$Q = \frac{S}{2(1-\epsilon)}(P_{\rm A} - P_{\rm B}) = \frac{1}{2(1-\epsilon)}(T_{\rm A} - T_{\rm B})$$
(4.14)

where S is the calibration conversion. With these definitions, we get a simple expression for the NEQ:

$$NEQ = \frac{NET}{\sqrt{2}(1-\epsilon)}.$$
(4.15)

This means that the average NEQs for the 2005 season were $\sim 420 \,\mu\text{K s}^{-1/2}$ at 100 GHz and $\sim 390 \,\mu\text{K s}^{-1/2}$ at 150 GHz assuming ϵ of 8 % and 5 %, respectively.

4.7 SIGNAL FILTERING RECTIFICATION

A subtle but necessary step in the proper analysis of these data is to account for the effects of the filters discussed in Section 3.2.4 on the time series. This constitutes rectifying the effects of the non-zero bolometer time constant and low pass Butterworth filter. Furthermore, the QUaD detectors are sampled quickly enough that the Butterworth filter correlates adjacent samples. This means that in a second but related step, a 'decimation' procedure can be applied to reduce the size of the time series, thus reducing the total computation time required during subsequent analysis. These techniques were first applied to the QUaD data by collaborators at Stanford University and the University of Chicago: their application to these data is now instituted as part of a standardized pipeline maintained by both institutions.





These plots show the noise power spectrum for the sum (black) and difference (gold) of the calibrated time series for bolometers 150-1-A and 150-1-B from noise data taken on Jan. 28, 2006 (the telescope was stationary and pointed at zenith during data acquisition). The upper plot shows the full power spectrum up to the Nyquist frequency for both the sum and difference time series, including the roll-off from the Butterworth filter in the electronics and the noise floor level of the ADC system near 50 Hz. The lower plot shows the same data, but with a restricted frequency range. Over plotted is the best fit 1/f spectrum for both sum (blue) and difference (red); these have a white noise level of 4.9×10^{-17} W Hz^{-1/2} and knee frequencies of 3.0 and 0.28 Hz, respectively.

Because the Butterworth filter has a known and verified spectral shape, a simple deconvolution in Fourier space is required to rectify the signal from its effects. However, the bolometer time constants must be measured before they can be properly rectified: in this section, we first present and discuss the time constant measurement method, and then the filter deconvolution and decimation algorithm.

4.7.1 TIME CONSTANTS

The bolometers act as low pass filters on incoming signals as defined in Equation 3.17; in order to account for this filtration, their time constants must be determined. As discussed in Section 3.2.4, it is technically possible to make this measurement based on cosmic ray transients. However, this is difficult because the impulse response of the system is only a weak function of the bolometer time constant. Furthermore, the loading due to the (unavoidable) brightness of a laboratory source requires a higher bias current through the bolometer, which acts to change the time constants somewhat. Therefore, we adopt a different approach.

A scan of a relatively bright astronomical source, taken in QUaD's standard CMBobservation bias mode, should produce a blip whose shape is the true shape of the source on the sky convolved with the beam and filtered by both the low pass Butterworth and the bolometer response function. If the telescope is scanned over the source quickly enough that the bolometer time constant has an effect on the shape of the blip, this can be exploited to deduce the time constant for each bolometer. This deduction is performed by taking a fiducial model for the blip (say, the ideal beam shape) and producing a time-domain model blip with a width and height which vary with scan speed in the correct way. This model blip is then passed through the correct Butterworth filter to model the effect of the electronic low pass on the shape of the blip.

This blip can then be filtered with a set of bolometer filter functions (which in the time domain are simply exponential decays with different time constants). The resulting blips are fit to the data, and the time constant for the model which minimizes the χ^2 is chosen as the value for the bolometer in question. This procedure can be applied to each bolometer with different telescope scan speeds to check for consistency. For QUaD, this test was performed using azimuth scan speeds of 0.266 deg/sec and 0.533 deg/sec (as compared to 0.25 deg/sec during CMB scans) at $\theta \in \{-3^\circ, +57^\circ\}$. The model reproduced the data well, although there was some ambiguity due to the fact that RCW 38 has dim structure surrounding the main lobe. Figure 4.23 shows the histogram of time constants for the 2005 season array.

As a check on these values, laboratory measurements using a chopped 77/300 K



Figure 4.23: Histogram of Bolometer Time Constants, 2005 Season This plot shows the histogram of bolometer time constants for every (working and open to light) bolometer in the array for the 2005 season. These follow a roughly Poisson distribution with a median value of 21.6 ms. Also plotted are the data for the 150 GHz and 100 GHz bolometers separately; these should show no statistical difference, as the time constant is not a frequency dependent quantity in these devices.

load were performed at the south pole in December 2005. These values (and particularly their bolometer to bolometer variation) agreed well with those determined from RCW 38, although there was a systematic decrease in the time constants due to the (necessarily) increased bias current.

4.7.2 FILTER DECONVOLUTION & DECIMATION

Once the filter functions of the instrument are known, the time stream must be corrected for their effects (we have previously shown that the theoretical filter function is a good description of that affecting the data; see Section 3.2.4). This procedure involves dividing the time series by the filter function in Fourier space, according to the following steps:

- 1. The relatively-calibrated data are divided up scan-wise; this is necessary since the gain offset of the bolometers is reset after each scan, thereby causing discontinuities in the time series.
- 2. The Fourier transform of each scan is taken, and this is multiplied by $H^{-1}(s)$, where H(s) is given in Equation 3.19 and is different for each bolometer.

128

3. The inverse Fourier transform of the result is taken; this is the deconvolved time series for each scan.

In reality, a few modifications to this simple procedure must be implemented. Firstly, if the analytic invariant impulse method is used (as in Section 3.2.4), the frequencies of the data stream and filter are exactly comparable. However, when using a digital transform to create the filter (which is preferable as a digital filter is a much simpler and faster to implement), the Fourier transformation procedure itself warps the digital filter to the incorrect frequencies. To correct this, the filter's cutoff must be *pre-warped* according to:

$$\omega' = \frac{2}{T}\tan(\omega T/2) \tag{4.16}$$

where ω is the old frequency variable, ω' is the new frequency variable, and T is the sampling period. This is sufficient to modify the frequency's characteristic of the filter to account for its digitization.

A second complication is that the filter function H(s) is small at frequencies \gtrsim 30 Hz, so that the deconvolution procedure can *increase* the power at high frequencies in the deconvolved time series by the simple expedient of dividing the (non-zero) time stream power by a small number. A way to minimize the effect of this is to apply a low pass finite impulse response filter to the deconvolved time series, which suppresses power at these higher frequencies. For the QUaD data, the filter is a compromise between a sharp cut-off at 6 Hz and ringing in the filtered time series due to the sharpness of the cut-off. The sharpness of a finite impulse response filter is defined by the Gibbs parameter G, where the lower the numerical value of G, the sharper the filter. The standard value used in the analysis of the 2005 observations is G = 100, which begins cutting off at 4 Hz and is effectively 0 above 6 Hz. Due to QUaD's standard scan speed and beam size, this cut-off does not attenuate astronomical signals, which appear at $\nu \leq 3$ Hz.

A second, and very useful, aspect of the low pass filtration is that the data can be 'decimated' before being used in the final analysis. Decimation is essentially a resampling the (archived) data at a frequency lower than the hardware sample rate, and has the effect of reducing the amount of RAM and drive space that are required for subsequent data analysis steps. It can be safely applied here because the QUaD sampling of 100 Hz is very much greater than the dominant cut-off filter applied to the data at 20 Hz²¹. The decimation is performed by dividing the deconvolved, low pass filtered time series into sets containing 5 samples. The decimated data is produced merely by taking the middle data point of each

²¹This cut-off filter is the 6-pole Butterworth, which has the effect of correlating data at frequencies greater than 20 Hz. This means that any subset of the data resampled above this frequency reflects the statistical properties of the original time series.

CHAPTER 4. INSTRUMENT PERFORMANCE



Figure 4.24: Decimation

This plot shows the result of applying the low pass filter and decimation procedure to 5 seconds of data from bolometer 150-14-A taken on May 30, 2005. In this plot, the blue points are the raw data, the red points are the deconvolved and low pass filtered data, and the black points are the final decimated points. Clearly, the decimation procedure, if properly implemented, can greatly decrease the computation time required in later analysis steps. (Figure courtesy M. Bowden).

set, thus producing a time series sampled at 20 Hz. Figure 4.24 shows an example of this procedure applied to real data.

The decimated final data are produced at Stanford and sent to the rest of the collaboration for further analysis steps. However, the work in this thesis has generally not utilized these data, as it is useful to have access to the raw values from the telescope when characterizing the instrument. As a general rule of thumb, the data presented in Chapters 4 and 5 do not use decimated data, while those in Chapter 6 do, unless otherwise indicated.

130

5 CALIBRATION

It is critical to the success of an experiment that data taken at different times can be compared in a meaningful way. Furthermore, it is just as critical that the results determined by one instrument can be compared to those determined by another. This is the purpose of calibration: to ensure that data taken throughout the lifetime of an experiment can be compared, preferably in a way that ensures that no systematic effect is introduced into the result, and to place the measurements on an absolute scale for comparison to measurements of the same quantities made by independent experiments.

Although calibration can, in principle, be performed in one step, in QUaD we divide calibration into relative calibration, which ensures that the data can be inter-compared correctly, and absolute calibration, which puts our results in terms of a real thermodynamic temperature. We use a number of different methods to perform relative calibration, some of which can be used to derive the absolute calibration. Absolute calibration is performed by comparing our results to the results of other experiments as a final stage in the map making process. Each of these calibration methods is discussed in this chapter; as a prelude, we discuss the effects of the atmospheric optical depth on the data.

5.1 ATMOSPHERIC OPTICAL DEPTH

The optical depth, commonly denoted by the symbol τ , is a measure of how opaque a medium is to radiation. The differential optical depth is defined as

$$d\tau = \kappa \rho dz \tag{5.1}$$

where κ is the opacity of the medium, ρ is its density, and dz is the infinitesimal optical path length through the medium. One integrates along the line of sight to determine the optical depth of the medium. Using a simple model, it can be shown that, at a given frequency, the transmission of the atmosphere varies as

$$\mathcal{T}_{\nu} = e^{-\tau_{\nu}A} \tag{5.2}$$

where the airmass $A = \sec z$, and the zenith angle $z = 90^{\circ} - \text{el}$. This means the emissivity of the atmosphere is given by

$$\epsilon = 1 - \mathcal{T}_{\nu} = 1 - e^{-\tau_{\nu}A}.$$
(5.3)

As the telescope tips down from zenith, the optical loading increases, until at the horizon it is infinite for this model¹. In addition, the variation in transmission means that the the brightness of astronomical sources is diminished the larger the airmass they are viewed through. This can be phrased mathematically: if the total temperature load on the telescope is T_{load} , then

$$T_{\text{load}}(z) = T_{\text{fixed}} + \epsilon T_{\text{atm}} + (1 - \epsilon)T_{\text{CMB}}$$
 (5.4)

$$= T_{\text{fixed}} + (1 - e^{-\tau_{\nu} \sec z}) T_{\text{atm}} + e^{-\tau_{\nu} \sec z} T_{\text{CMB}}$$
(5.5)

$$= T_{\text{fixed}} + T_{\text{atm}} - e^{-\tau_{\nu} \sec z} \left(T_{\text{atm}} - T_{\text{CMB}} \right).$$
 (5.6)

where T_{atm} is the effective temperature of the atmosphere. Because the different radiation sources (*i.e.* the atmosphere and astronomical sources) have different spectra, this expression should be viewed as most applicable in the Rayleigh-Jeans (RJ) approximation. Here, T_{fixed} represents the loading from the telescope's optics, which should be a constant.

Next, we can assume that the effective temperature of the atmosphere $T_{\rm atm} > 200 \,\text{K}$ is sufficiently large that its physical temperature is the same as its RJ temperature, and that the CMB's temperature is smaller than the other two terms. Also, because $\tau \sec z \ll 1$, we can say $\epsilon \approx \tau \sec z$. This makes Equation 5.4 approximate to

$$T_{\text{load}} \approx T_{\text{tot}} + \tau T_{\text{atm}} \sec z$$
 (5.7)

where T_{tot} is now a constant total temperature offset that combines the constant loading from the atmosphere and telescope. Unfortunately, Equation 5.7 implies that it is impossible to fit for τ and T_{atm} separately. Fortunately, a 350 μ m τ tipper is operated at the south pole (Peterson et al., 2003) on a building adjacent to MAPO². This tipper - as its name suggests - tips up and down every 12.7 minutes, measuring the atmosphere. It *can* separate

¹Practically speaking, this exponential law breaks down below about $z = 60^{\circ}$, even in otherwise ideal conditions.

²Actually, during the 2005 observing season, it was operated at the SPARO building, and during the 2006 season, it was operated on the DSL building.
τ and $T_{\rm atm}$ by virtue of the fact that the typical values of τ are extremely large at 350 μ m, so that the exponential dependence in Equation 5.4 is very strong. This tipper has built-in software that autonomously analyzes the data and emails it to a list of recipients around the world every two hours³. These data are used for two purposes: firstly, the tipper reports a $\tau_{350\,\mu\text{m}}$, which can be correlated with the τ measurements at other wavelengths and secondly, the tipper extracts and reports $T_{\rm atm}$ independently of any QUaD data. This $T_{\rm atm}$ can be used to break the degeneracy in Equation 5.7 for τ at QUaD's wavelengths.

For QUaD, the measurement of τ is performed via the skydip technique (Dicke et al., 1946). This consists of tipping the telescope to a known set of elevations and integrating on the sky. A measurement of the power incident on the detectors at each elevation is fit to Equation 5.7 yielding the factor $\tau T_{\rm atm}$. From this, τ is extracted using the most recent $T_{\rm atm}$ measurement from the 350 μ m tipper.

Unfortunately, calibrated measurements of the absolute power incident on the detectors are difficult with QUaD (see Section 2.4.1). In the interest of being as precise as possible given our instrumental limitations, we have adopted a conservative technique in our skydip measurements. First, a load curve is performed before and after the skydip observation itself. This load curve is analyzed to extract the bolometer parameters G_0 and β for each device, since the ambient temperature conditions vary enough to cause a slight apparent change in their values from day to day. The sky dip consists of a set of DC-biased data streams at known elevations. Beginning at $el = 85^\circ$, the telescope tips down, integrating every 5° in el until $el = 45^{\circ}$ is reached. Then the telescope's motion is reversed, and it integrates in the same fashion back upward to the initial elevation. To get the signal at each elevation, the median voltage at each step is calculated, removing 100 points (1 s) at the beginning and end of the data stream for each elevation step. This statistic is sufficient to remove the (low-level) noise in skydip observations. The voltages are then converted to power using the formalism discussed in Section 2.4.1; first voltage is converted to resistance, which is then converted to T_{bolo} . This, combined with G_0 , β and the electrical power dissipated in the thermistor, yields Q, the optical power incident on the detector at each elevation. Using a look up table for dT/dP determined in the laboratory, the equivalent Rayleigh-Jeans temperature of the power can be calculated. This is then equated to $\tau T_{\rm atm}$ to solve for τ . Figure 5.1 shows the results of these calculations for each day averaged over detectors for the 2005 observing season.

The characteristics of these plots are interesting. First, variation in τ at QUaD's frequencies is never that strong, leading to a $\sim 1 - \exp(-0.01) = 1\%$ variation in the

³P. Calisse has shown that the 350 μ m τ values reported are actually slightly erroneous due to the specifics of the instrument (Calisse, 2002). Under most conditions, simply adding 0.3 to the reported τ is sufficient to account for this subtlety. All of the 350 μ m optical depths given here have been corrected for the Calisse term.

CHAPTER 5. CALIBRATION



Figure 5.1: Optical Depths for the 2005 Observing Season

This plot shows $\tau_{100 \text{ GHz}}$, $\tau_{150 \text{ GHz}}$ and $\tau_{350 \,\mu\text{m}}$ for the entire 2005 season (upper, middle and lower panels, respectively). These data points were determined by computing the τ for each bolometer with model parameters, and calculating the mean and standard deviation of the results (of course, $\tau_{350 \,\mu\text{m}}$ comes from the tipper rather than QUaD). The $\tau_{100 \text{ GHz}}$ show very little structure, due to the fact that the atmospheric emission lines bounding this window are not due to water. However, $\tau_{150 \text{ GHz}}$ and $\tau_{350 \,\mu\text{m}}$ do vary with the water vapour content of the atmosphere, and are highly correlated as a result. The 350 μ m tipper samples the optical depth far more frequently than does QUaD (every ~ 10 minutes compared to once per day), and thus exhibits a great deal more detail than the QUaD optical depth time series do. extinction factor⁴. Therefore, neglecting the atmospheric extinction should only affect the calibration of the final data by about this amount. This has indeed been shown with the QUaD data using the final data maps (M. Bowden, private communication), which is reassuring. Secondly, the 100 GHz optical depths exhibit very little time variation compared with the other channels: this can be understood in terms of the physics of the atmosphere. The two lines bounding the classical 100 GHz atmospheric window are actually O₂ lines rather than water lines; because O₂ is fairly evenly distributed through the atmosphere, there is very little in–band variation at 100 GHz. The wings of the 182 GHz H₂0 line are occasionally strong enough to come into this band, but this is very rare at the south pole. However, the wings of this same line can be quite strong at 150 GHz, leading to a correlation between the 150 GHz optical depth and the water vapour content of the atmosphere. Furthermore, the water vapour content is extremely important at 350 μ m, leading to a correlation between $\tau_{150 \text{ GHz}}$ and $\tau_{350 \,\mu\text{m}}$.

It is informative to investigate these inter-wavelength correlations to understand how well a measurement of τ at one frequency tracks to other frequencies⁵. This type of correlation analysis is straightforward: for each skydip measurement, find the corresponding 350 μ m tipper measurement, and then fit the three frequency permutations to one another (this same type of analysis was performed with a similar tipper at the CalTech Sub-Millimetre Observatory and skydips from the James Clerk Maxwell Telescope, as presented by Archibald et al. (2002)). Because each tipper measurement has an approximately 30 % random error, the relevant $\tau_{350\,\mu m}$ measurement is constructed by using a quadratic interpolation at the time of each QUaD skydip. This yields a low-noise time series which can be correlated with the QUaD τ measurements⁶. These time series are plotted against each other in Figure 5.2, including two models to quantify the correlation between them as over-plots.

The models fit to the τ data are based on physical knowledge of the atmosphere, and we can predict the outcome of these fits based on this knowledge. First, $\tau_{100 \text{ GHz}}$ should not be correlated with either 150 GHz or 350 μ m, since water vapour does not affect this channel. This hypothesis is proved by Figure 5.2; in all cases, the model is virtually a

⁴This is a well-known truism of mm/sub-mm ground-based astronomy: the raw τ value is not particularly important, but rather the *variability* of the atmospheric emission on ~ 1 Hz time scales is the quantity of importance. While it is true that a higher τ value generally indicates higher variability on these short time scales, this is not necessarily the case, and observations performed in high τ conditions must be checked rather than thrown away outright.

⁵One could argue that by doing so, should the correlation be sufficiently well understood and well behaved, we can do away with QUaD skydips altogether and thereby increase our observing efficiency. This idea is discussed in more detail in Section 5.2.

⁶In fact, close inspection of the 350 μ m tipper data shows that the correlation time scale of the 'DC' component of the atmospheric water vapour content is ~ 1 h, so information could be lost in this procedure. The only remedy would be to perform skydips every hour, which is not a feasible solution.





These three panels show the data shown in Figure 5.1 plotted against one another to show the relationship between them (in this case the data from the $350 \,\mu\text{m}$ tipper has been interpolated at the QUaD skydip points to facilitate the comparison). The error bars plotted in the previous figure have been suppressed for clarity. Also plotted are fits to these data in each panel; blue is a linear model, and red a quadratic model. Since there is no physical expectation that a quadratic model is a good description of the relationship between $\tau_{150 \text{ GHz}}$ and $\tau_{100 \text{ GHz}}$, this model has not been fit to that combination. It is clear from these plots that only in the $\tau_{150 \text{ GHz}} - \tau_{350 \,\mu\text{m}}$ case is there a strong correlation.

straight line. This means there is no correlation between these variables, and that $\tau_{100 \text{ GHz}}$ is best described as a simple offset ($\tau_{100 \text{ GHz}} \approx 0.046$) throughout the QUaD winter season. Secondly, although we expect the $\tau_{150 \text{ GHz}}$ to be correlated with the precipitable water vapour in the atmosphere, this correlation is much stronger at 350 μ m, since water emission is much stronger at this frequency. In fact, the $\tau_{350 \mu m}$ may be as strong a function of the precipitable water vapour as its square (P. Ade, private communication); we therefore fit both a linear and quadratic model in the case of the tipper data. Again, the 100 GHz optical depth appears to be uncorrelated with that at 350 μ m. However, evidence for a strong correlation between $\tau_{150 \text{ GHz}}$ and $\tau_{350 \mu m}$ does exist, as shown in the lower panel of Figure 5.2. Both quadratic and linear models have similar χ^2 values for this relationship, meaning that it is difficult to determine which is a better description of the data. We therefore report both fits:

$$\tau_{150 \text{ GHz}} = -0.00559(\tau_{350 \,\mu\text{m}} - 3)^2 + 0.00297(\tau_{350 \,\mu\text{m}} - 3) + 0.0597 \qquad (5.8)$$

$$\tau_{150 \text{ GHz}} = 0.0181\tau_{350 \,\mu\text{m}} + 0.0148. \qquad (5.9)$$

Applying the linear fit to the true (*i.e.* pre-interpolation) tipper time series and plotting it with the $\tau_{150 \text{ GHz}}$ time series as measured by QUaD yields Figure 5.3; it is clear from this plot that the correctly scaled tipper data can indeed be used as a proxy for the sky dips for the 2005 season.

Of interest are the statistics of the τ distributions at the QUaD wavelengths. The histograms for both, as well as their cumulative distributions (which can be interpreted as the τ cumulative probability density functions) are plotted in Figure 5.4. From these plots, it can be seen that the 50 % quartile values are $\tau_{100 \text{ GHz}}^{50} = 0.045$ and $\tau_{150 \text{ GHz}}^{50} = 0.035$. Based on these values, atmospheric models predict that the median precipitable water vapour over the south pole for the 2005 observing season was 0.5 mm, which is the lowest for any ground based astronomical site in the world (see Peterson et al. 2003).

5.2 **ELEVATION DIPS**

A method of determining the relative gain of the bolometers frequently is of primary importance in this experiment. This requirement is fundamentally due to the fact that in instruments like QUaD, the two bolometers at the bottom of a feed horn are differenced to produce the polarized time series. This means that a gain mismatch between the two devices can appear as a spurious polarized signal, which would be disastrous. Therefore, knowledge of g(t) at all t – or at least as frequently as possible – is an experimental requirement.



Figure 5.3: τ Scaling Comparison Over Entire 2005 Season

This plot compares the scaled 350 μ m tipper data (black points) to the QUaD $\tau_{150 \text{ GHz}}$ data (red points with error bars) for the 2005 CMB observing season. This is achieved simply by using Equation 5.8 on the tipper data plotted in Figure 5.1, and over plotting this on the 150 GHz points from the same figure. Clearly, a great deal of structure is common to both wavelengths.

One method of measuring the gain of the instrument simultaneously in each bolometer is to tip the telescope a small angular distance in elevation. Since, under the same assumptions as used in Equation 5.7, the loading from the atmosphere goes as $\sec z$, we can write

$$\frac{\partial T_{\text{load}}}{\partial z} \propto \frac{\partial \sec z}{\partial z} = \sec z \tan z$$
 (5.10)

which, if δz is small and $z \leq 50^{\circ}$, is well approximated by a linear function so that $\delta T_{\text{load}} \approx g_{\text{ed}} \sec z$. Here, we have introduced the term g_{ed} , which is the constant of proportionality between the change in airmass and the apparent change in optical loading.

In terms of experimental procedure, this means that we can tip the telescope a small amount in elevation, and fit a line to the resulting bolometer time series expressed as a linear function of airmass. Provided the elevation is sufficiently high to prevent nonlinearities in the detectors and model, the slope of the fit yields the relative responsivity of the bolometers. This method has been used successfully with other experiments (*e.g.* SuZIE, S. Church, private communication) to measure relative calibrations as frequently as every few minutes if required.

In the case of QUaD, an elevation dip (hereafter 'el dip') consists of slewing the telescope up in elevation $\approx 1.35^{\circ}$, and then back down by the same amount. The el dip is

5.2. ELEVATION DIPS



Figure 5.4: τ Cumulative Statistics for the 2005 Observing Season

This plot shows the histograms and cumulative distributions for the optical depths at QUaD's two wavelengths. The red-black line is the histogram for $\tau_{150 \text{ GHz}}$, while the redyellow line is the cumulative distribution for this histogram (this is formed by summing the previous points at each τ and normalizing such that the probability of τ smaller than the largest τ is 1). Similarly, the blue-black line is the histogram for $\tau_{100 \text{ GHz}}$, while the blue-purple line is the cumulative distribution for this histogram.

begun at the same elevation as the source under observation to increase total observation efficiency by minimizing slew time. During CMB observations, el dips are performed between every 5 full azimuth scans at the same elevation as the azimuth strip being observed as part of the scan calibration block. The linear fits to airmass are performed using a simple linear least–squares algorithm; the ~ 5 s of data at both the beginning and end of the dip are not used as they may suffer from acceleration effects. These fits can then be applied as a multiplicative factor to each set of 5 scans to provide the relative calibration. Figure 5.5 shows typical time traces for el dips, including the resulting fits to the airmass.

One minor subtlety in this analysis is that some of the bolometer channels 'ramp' over time. Specifically, these channels experience a decrease in responsivity which follows an exponential decay law after each bias current reset. These ramping channels, such as 150-1-B shown in Figure 5.5, are therefore difficult to characterize, and are generally excluded from our final data analysis. The cause of this ramping is unknown, but seems connected to the thermal relaxation time of the bolometer. If the thermal connection of the thermistor to the heat sink is poor or abnormal, behaviour like that seen in these bolometers

139



Figure 5.5: Elevation Dip Time Series

These plots show the time series for the elevation dips for bolometers 150-1-A (left panels) and 150-1-B (right panels) for data taken at the beginning of the July 1, 2005 observation set. In the upper panels, the black line shows the bolometer voltages in high gain mode, while purple shows the elevation of the telescope at each time (a full el dip takes 51 s). The lower panel shows the same bolometer in black, but plotted as a function of airmass. According to Equation 5.10, the loading should be an approximately linear function of the airmass. This is indeed found to be the case; the fits to the upward–going scan (red) and downward–going scan (blue) are over plotted. It is obvious from these plots that this method works better for 150-1-A than 150-1-B; this is because 150-1-B is one of the 'ramping' channels (discussed below) which suffers from some thermal lag causing its output to vary over short time scales. However, the ramping channels are by far the exception, and most of the bolometers on the focal plane behave as well as 150-1-A.

may occur. The bolometers affected by this problem were in positions 150-1-A, 150-1-B, 150-16-A, 150-16-B, 100-8-A, and 100-8-B during the 2005 observing season⁷. These bolometers were thoroughly cleaned under a microscope during the 2005-2006 summer season and replaced in different positions on the focal plane.

To properly determine their effectiveness in their capacity as our relative calibration source, it is necessary to study the behaviour of the el dip derived g_{ed} over time. This shows general calibration trends, and helps determine when and if this calibration method might fail. Figures 5.6 and 5.7 give example time series of the el dip measurements for the 2005 season for bolometers 150-2-A/B and 100-2-A/B, respectively.

Figure 5.6 exhibits a great deal of structure. First, gaps in the time series are visible; these are times at which the telescope required maintenance, for example, during refocusing or optical pointing measurements near day 75. Furthermore, the magnitude of the measured g_{ed} is strongly correlated between the two bolometers, and between all of the 150 GHz bolometers as well. This correlation suggests that the gain depends on something common to all of the devices either within the cryostat, or external to it. None of the cryostat house keeping data suggests any variation of this type in the detector's environment on the relevant time scales. Interestingly, a quick glance at Figure 5.1 suggests a strong relationship between the atmospheric conditions and the measured g_{ed} (this correlation is investigated in detail in Section 5.2.1). The cause of the temporal variation becomes obvious when this is taken into account: g_{ed} depends most strongly on the atmospheric loading at any given time.

In fact, the structure in Figure 5.6 correlates quite well with the observed external weather conditions. Typically, the weather at the south pole is somewhat bimodal: clear and cold (good observing conditions) or warmer and overcast (poor observing conditions). The typical pattern is a week or two of good conditions, followed by a storm for ~ 5 days. Typically, the weather deteriorates during September into October. This pattern is clearly visible in these plots. The grossly abnormal g_{ed} always occur during storm conditions, and are due to the atmospheric loading behaving so erratically as to preclude observations. Observing days identified as suffering from poor weather conditions – based on these el dip values, the reported weather at the south pole, and a visual examination of the bolometer data – are removed from analysis, as the atmosphere's noise characteristics during these times can be strongly non-Gaussian⁸.

The ratio of the g_{ed} (bottom panel of Figure 5.6) also provides an interesting check.

⁷Interestingly, the 'B' side bolometers exhibited this effect more strongly than the 'A' side bolometers, which in many cases only suffered weakly and intermittently from it. This is perhaps a clue to the cause of the effect. However, although ideas have been suggested, no conclusive evidence pointing to the cause of the fault has been found.

⁸a.k.a. 'Completely Crazy.'



Figure 5.6: Elevation Dip Measurements for the 2005 Observing Season for Horn 150-2 This plot shows the total set of elevation dip $g_{\rm ed}$ from the beginning of CMB observations to the time the helium supply ran out in 2005. The upper panel shows the results for bolometer 150-2-A, the middle for bolometer 150-2-B, and the bottom panel shows the ratio of the two. Here, red shows data taken in the $\theta = +57^{\circ}$ telescope rotation, and blue shows data taken in the $\theta = -3^{\circ}$ rotation. The grey band shows the period at which the telescope was refocused.



Figure 5.7: Elevation Dip Measurements for the 2005 Observing Season for Horn 100-2 This plot shows the total set of elevation dip g_{ed} from the beginning of CMB observations to the time the helium supply ran out in 2005. The upper panel shows the results for bolometer 100-2-A, the middle for bolometer 100-2-B, and the bottom panel shows the ratio of the two. Colour allocation is the same as in Figure 5.6.

For the elevation dips to provide a reliable calibration source, the ratios of g_{ed} must be stable over long time scales. Instabilities in the ratio indicate strong non-linearities in the detectors under loading changes, which will cause models like Equation 2.13 to fail during later analysis. This plot shows that the ratio is, in fact, quite stable over the relevant time scales. The only major features in these plots are during the poor observing conditions discussed above; at these times, the atmosphere may provide loading large enough to drive the detectors slightly non-linear, causing a change in the ratio. Again, the very abnormal ratios occur during the very worst storms.

A slightly more worrying feature of these plots is the difference due to telescope rotation. This issue is not immediately apparent in the single-bolometer plots, but becomes quite obvious in the ratio. A cause for this effect has not been identified, but is thought to be related in some way to the behaviour of the refrigerator in different telescope rotation angles⁹. This issue leads to some complexities during the later analysis. Fortunately, the statistics of the g_{ed} , when the abnormal outliers are excluded and split into the two θ angle populations are well described by Gaussian distributions. This means that, although our relative calibrations may vary with time, the important quantity (*i.e.* how they vary jointly) is noisy at the ~ 1 % level but will not inject a bias into the final data set.

The g_{ed} for the 100 GHz detectors show similar, but slightly less obvious structure, as shown in Figure 5.7. Again the very strong storms appear in these plots, but in general the measured g_{ed} are more constant than those at 150 GHz. This is simple to understand based on the model derived in Section 5.1; water vapour variations dominate the optical loading from the sky, and as very little water vapour emits in the 100 GHz window, this channel is much less affected by the polar weather. Presumably, the variation that does exist (*i.e.* during strong storms) is due to the wing of the 183 GHz water line bleeding inband. Again, the very abnormal fits are synchronous with the very worst scan data. The ratio for the 100 GHz g_{ed} is again very stable, but shows a disparity between $\theta = +57^{\circ}$ and $\theta = -3^{\circ}$. In fact, other 100 GHz bolometer pairs exhibit this problem more strongly than the 100-2 pair. Why this problem depends on wavelength is unknown, but may be related to the overall responsivities of the bolometers being higher at 100 GHz. Once again with these data, the ratios are Gaussian distributed, meaning that they provide an unbiased relative calibration.

⁹The orientation of the fridge with respect to the gravity vector is a major contributor to its base temperature and hold time. This has to do with the geometry of the evaporation chamber and still (see Figure 2.5), and is not something we have a great deal of control over with this type of cooler. This problem is what prevents us from observing at arbitrary θ angles, and as such is a limitation of the instrument design.

5.2.1 The EL DIP – τ Relation

A simplistic way of thinking about an elevation dip is that it is a 'tiny' sky dip: that is, one is observing the atmosphere at various elevations while archiving data. This being the case, since both the sky dip and elevation dip sample the same loading conditions, the major component of which is a constant offset at any given time, the time variability of the measurements should be correlated at some level. This can be seen more formally by observing that *both* sky dips and the slope of an elevation dip measure the term τT_{atm} in Equation 5.7; the only major difference is that for elevation dips Δz is small so sec z behaves approximately linearly.

As discussed above, this relationship can be seen by comparing the temporal structure in Figures 5.1 and 5.6; they are clearly highly correlated. While this may seem like an academic issue, an immediate use for this correlation becomes obvious: optical depth determination. That is, since the elevation dips are so strongly tied to τ , is it possible to model the measured optical depths using the el dip values? This would allow a virtually on-the-fly determination of the optical depth at 150 GHz (since we have already determined that the 100 GHz optical depths are not particularly variable or interesting), which would in principle negate the need for sky dips.

Figure 5.8 shows a scatter plot of the g_{ed} and sky dip determined $\tau_{150 \text{ GHz}}$; the latter are precisely the same as those presented in Section 5.1 for the 2005 season. In this case the g_{ed} used are the mean of those for bolometers 150-2-A & B, as plotted in Figure 5.6; this quantity is denoted $\langle g_{ed} \rangle$. The choice of bolometers is arbitrary; this same analysis could be performed with any combination of the bolometers at a single wavelength (recalling that the full set of $\tau_{100 \text{ GHz}}$ are consistent with a constant value and thus trivial). The choice of 150-2-A/B is simply driven by expediency; this is one of the best-behaved bolometer pairs on the focal plane, so using some larger sample would not provide a significant increase in signal to noise ratio on the mean. It would, however, represent a significant increase in the complexity of the analysis and calculation time, and so the simpler choice of the two 150-2 bolometers is made here¹⁰.

Once the choice of $\langle g_{\rm ed} \rangle$ is made, a median filter of width 10 points is applied to the resulting el dip gain time series. This has the effect of reducing the noise in the time series, while still providing time sampling of the gain below the typical correlation length of τ ($\gtrsim 1$ hr). This time series is then interpolated at the times of the sky dip observations to provide a directly comparable data set; the interpolated $\langle g_{\rm ed} \rangle$ are plotted versus their associated sky dip values in Figure 5.8.

¹⁰Other combinations have been used in this analysis, including the mean of all, the median of all, the mean of seven random bolometers, et cetera. The results were found to be consistent for all, although only numbers specific to 150-2 are presented here.



Figure 5.8: El Dip Gains and Sky Dip τ for 2005 Observation Season This plot shows the interpolated, median filtered $\langle g_{\rm ed} \rangle$ plotted against the $\tau_{150 \, \rm GHz}$ determined from sky dips for the 2005 season. The blue-gold line shows the best fit linear model relating the two variables; this can be used to predict τ based on gain measurements.

Also plotted in Figure 5.8 is a linear fit of the two parameters; this can be used to predict $\tau_{150 \text{ GHz}}$ based on the measured $\langle g_{\text{ed}} \rangle$. The model is given by

$$\tau_{150\,\text{GHz}} = \frac{0.06412 \langle g_{\text{ed}} \rangle + 1.209}{100}.$$
(5.11)

This will, of course, change when different bolometer combinations are used to determine $\langle g_{\rm ed} \rangle$. Using this model, and that given by Equation 5.8 for the $\tau_{350\,\mu\rm m}$ scaling to 150 GHz produces a plot like Figure 5.9, which shows how the independent tracers of the optical depth perform over the 2005 CMB observing season.

5.3 INSTRUMENTAL CALIBRATION SOURCE

The QUaD optical system contains a calibration source located behind the secondary mirror, coupled to the main beam via a flip mirror; the hardware aspects of this device are discussed in Section 2.1. In this section, the analysis of data archived during the calibration source's operation is discussed.

During calibration source operation, the bolometers are in low gain mode, as the signal from the source is bright enough to saturate the range of QUaD's ADC system. The flip mirror is flipped into the optical path, and a polarizer in front of a thermal source



Figure 5.9: Grand Unified Model of $\tau_{150 \text{ GHz}}$

This plot shows $\tau_{150 \text{ GHz}}$ calculated from scaling relations for $\tau_{350 \,\mu\text{m}}$ (Equation 5.8) and $\langle g_{\text{ed}} \rangle$ (Equation 5.11) for the 2005 CMB observing season. The black points with red centres and error bars are the sky dip determined $\tau_{150 \text{ GHz}}$ as plotted in Figure 5.1, the yellow points are the scaled $\tau_{350 \,\mu\text{m}}$, and the blue points are the scaled $\langle g_{\text{ed}} \rangle$. Clearly, these variables are all measuring the same fundamental quantity and can be used as proxies for one another.

is physically rotated at approximately 1.7 Hz to produce a polarization vector rotating at twice this frequency. The signal from the source is shown in Figure 5.10; the signal has a sinusoidal modulation to due to the rotating polarizer.

In order to analyze this data, we must fit it to a sinusoidal model:

$$V(t) = V_{\rm A}\sin(\omega t + \phi) + V_0 \tag{5.12}$$



Figure 5.10: Raw Calibration Source Signal

This plot shows the raw signal due to the calibration source on the telescope for data taken on July 1, 2005. The left panel is for bolometer 150-1-A, and the right panel is for bolometer 100-1-A. The signal is clearly sinusoidal, and smaller at 100 GHz due to the thermal nature of the calibration source.

where V_A, ω, ϕ and V_0 are free parameters. Since non-linear sine fits are notoriously sensitive to initial conditions, we must ensure that the fit settles on the correct region of parameter space. This is performed in an automated way, since checking each fit by eye would be extremely time consuming (there are on order 10⁴ such observations for each bolometer in a single season). To reduce erroneous fits, the initial values in the fit are set to be very close to their true values: V_A is set to the mean rms of the data set multiplied by $\sqrt{2}$, V_0 is set to the mean of the data, ω is determined by finding the first peak of the power spectrum of the data, and ϕ is set to a value dictated by setting the derivative of the function to 0. The χ^2 statistic of the fit is then determined: fits which have settled on harmonics of the fundamental tone of the wave are rejected based on this. These fits are refit using the same initial conditions, except with $\phi_{new} = \phi_{old} + \delta \phi$, where $\delta \phi$ is drawn from a random Gaussian distribution with variance 0.05ω . This 'kicks' the fit slightly in parameter space, and this fit is tested using χ^2 again. If it fails, the process is repeated until the fit passes the χ^2 test, or 10 iterations have been performed. If 10 iterations are reached, the data are not used in the fit¹¹. No examples of a bad fit passing through this filter into the final set of calibration source results have been found in the 2005 season; we are therefore confident that all of the data included in this set are genuine parameter fits. Figure 5.11 shows an

¹¹This occurs for very poor weather, as discussed for the abnormal fits in the preceding section.



example of the results of this procedure for data from a bolometer at either wavelength.



These plots show the sine wave fit results for the same data as is shown in Figure 5.10 (note the y-axes of these plots are different in this figure). The upper left panel shows the data for bolometer 150-1-A in black, and its fit in red, while the upper right panel shows the same for bolometer 100-1-A. Clearly, the models are very well matched to the data; the lower panels show the residuals of the fits. These residuals have amplitudes approximately 1% that of the original sine wave, meaning that the single sinusoid model is a reasonable one. A slight complexity is the residual for bolometer 100-1-A; this shows a residual harmonic signal. This is thought to be due to the time constant of the slow bolometers (a few of which exist at either QUaD band) affecting the shape of the signal, and could be mitigated by slowing the calibration source rotation rate (were that possible).

The parameters for fits which pass the χ^2 test are stored for later analysis. The parameter of interest for the purposes of calibration is V_A , which we will investigate further. Of the other parameters, only ϕ may be of interest, as this can yield the relative polarization angles of the devices (the other model parameters are due to the arbitrary dc offset of the bolometer signal and the rotation rate of the calibration source). However, a careful and extremely thorough analysis of the data shows that these data suffer from some systematic effects which preclude it for this use.

There are two major contaminants to the signal which cause the calibration source to be untrustworthy for an absolute calibration based on the known optical properties of the calibration source and polarizer. The first of these is that the polarizer used in the calibration source is up to 40 % reflecting, meaning that standing modes exist between it and the detector block. This makes the absolute power incident on the detectors, as well as their absolute polarization properties, somewhat ambiguous (*i.e.* the detectors are 'seeing themselves' at some non-negligible level). The second complication is that, under careful scrutiny, it appears stray light external to the calibration source's optical path is contaminating the signal. That is, some complex reflection of an unknown number of ~ 300 K components somewhere in the secondary mirror support assembly is also present in the calibration source signal. This makes the source of the photons seen at the detectors even more ambiguous, and unfortunately it appears that this reflected signal is much stronger in some bolometers than others. This is why the calibration source is thought to be only useful for calibration of a single detector over time (since these two contaminants are hoped to be constant over time), rather than between detectors (which was its original, but failed, purpose).

Given that the contaminating signals are constant over time, the parameter V_A should be a robust measurement of the intrinsic responsivity of the bolometers. A tensecond calibration source observation is performed in conjunction with the elevation dips every 5 azimuth scans; together, these make a scan calibration block. In the same way as for the elevation dips, it is instructive to examine the behaviour of V_A over time. Figures 5.12 and 5.13 show the raw V_A values over the 2005 observing season.

The first aspect of these two plots one notices is that they exhibit similar structure to that exhibited in Figures 5.6 and 5.7 for the elevation dip calibration values, but anticorrelated with them. Once again, the variation is much stronger at 150 GHz than 100 GHz, while the ratio of the quantities is flat except during storms when either the atmospheric variation is extremely large on short time scales, or in the worst cases, the assumption of detector linearity is invalid. This anti-correlation can be understood quite simply in terms of the behaviour of bolometers: under higher loading conditions, the intrinsic responsivity of the detectors is decreased, leading to a lower response to a constant source. Vice versa, in low optical loading conditions, the responsivity is greater, and the signal from the constant source is larger. Another feature of these plots one notices is that there is, once again, a systematic difference between the two telescope rotation angles in the V_A ratio. This is presumably caused by the same effect as for the elevation dip data, although the root cause of this rotation-dependent difference is unknown.

Given that the atmosphere-dependent responsivity change is caused by the loading conditions, it should be possible to correct for them. Assuming that the responsivity is a linear function of loading, all that is required is an independent tracer of loading with which we can correct V_A . This is provided by the elevation dip data themselves: we can simply



Figure 5.12: Calibration Source Measurements for the 2005 Observing Season for Horn 150-2

This plot shows the total set of calibration source derived V_A from the beginning of CMB observations to the time the helium supply ran out in 2005 (this plot can be directly compared with Figure 5.6 for the set of elevation dip calibrations). The upper panel shows the results for bolometer 150-2-A, the middle for bolometer 150-2-B, and the bottom panel shows the ratio of the two. Here, red shows data taken in the $\theta = +57^{\circ}$ telescope rotation, and blue shows data taken in the $\theta = -3^{\circ}$ rotation. The grey band shows the period at which the telescope was refocused.



Figure 5.13: Calibration Source Measurements for the 2005 Observing Season for Horn 100-2

This plot shows the total set of calibration source derived V_A from the beginning of CMB observations to the time the helium supply ran out in 2005 (this plot can be directly compared with Figure 5.7 for the set of elevation dip calibrations). The upper panel shows the results for bolometer 100-2-A, the middle for bolometer 100-2-B, and the bottom panel shows the ratio of the two. Here, red shows data taken in the $\theta = +57^{\circ}$ telescope rotation, and blue shows data taken in the $\theta = -3^{\circ}$ rotation. The grey band shows the period at which the telescope was refocused.

use them as a monitor of the optical loading, fit to the calibration source gain time stream, and remove the effect. In practice, this is performed by removing the outliers from both time series, and calculating the mean of the A/B bolometer g_{ed} to get $\langle g_{ed} \rangle$, as above. This mean elevation dip time series is then passed through a median filter with width 10 points. This new time series is then interpolated at the calibration source observation times, and the two time series are fit to each other using a linear model. A subtlety is that the different telescope rotations are treated independently; this is due to the slight difference between the two rotations in both calibration time series. If this separation is not performed, the fit results do a poor job describing the data for either telescope rotation.

Once the linear relation between the interpolated, median filtered $\langle g_{ed} \rangle$ and V_A has been determined, this fit is subtracted from V_A . Because this would produce data with approximately zero mean, the median of the V_A are preserved during this process by adding them back to the fit-subtracted V_A . This is acceptable because, over long time scales, the mean (which is equivalent to the median by the central limit theorem) should be dominated by the constant term; if this is not true, there is something pathological occurring in the instrument. The resulting time series is plotted in Figure 5.14 for a 150 GHz bolometer pair and 5.15 for a 100 GHz bolometer pair.

A number of features of these plots are interesting; consider first the results at 150 GHz (Figure 5.14; the features discussed here are present in *all* of the detectors, not just 150-2-A/B). First, a great deal of the time variation evident in Figure 5.12 has been removed; this shows that the bolometer responsivities are constant through the season. To quantify this, we divide the CMB observing season into two epochs:

- 1. $T \leq 117$ Days: During this time, good CMB data was being taken, and the calibration source appeared to be operating well. These calibration data only show evidence of non-constant performance during storms, and even then this only increases the scatter of the loading corrected V_A values.
- 2. T > 117 Days: This epoch experienced two events in quick succession that change the loading corrected V_A . The first was two strong storm events back to back between days 105 and 120; this was a bad stretch of time for data quality. The second event is the jump evident in the top and middle panels of Figure 5.14; at the time the data were being taken, the cause of this structure was a complete mystery (although it was seen in all light channels). When the calibrator unit was inspected during the 2005/2006 summer telescope maintenance, it was found that the polarizer (which is composed of $\sim 15 \,\mu$ m plastic) had ripped. Presumably, the jump evident in the calibration source data is due to the rip, which must have occurred in high winds during the second strong storm. The rip changed the polarization properties of the



Figure 5.14: Intrinsic Gain Measurements for the 2005 Observing Season for Horn 150-2 This plot shows the V_A measured using the calibration source in the secondary mirror assembly corrected for optical loading responsivity changes using the elevation dips for horn 150-2 over the entire 2005 CMB observing season. The colour allocation of the points are as described in Figure 5.12. The two dashed lines show the best fit line to these data for $0 \le T \le 100$ days; these show that the intrinsic responsivities of the detectors are remarkably stable over long time scales (the text discusses the obvious changes in structure outside of this T range).



Figure 5.15: Intrinsic Gain Measurements for the 2005 Observing Season for Horn 100-2 This plot shows the V_A measured using the calibration source in the secondary mirror assembly corrected for optical loading responsivity changes using the elevation dips for horn 100-2 over the entire 2005 CMB observing season. The colour allocation of the points are as described in Figure 5.13, and the two dashed lines are as in Figure 5.14. Again, the intrinsic responsivities of the detectors are remarkably stable over the observing season.

calibration source, thereby changing the amplitude of the sinusoidal signal¹².

The data during the normal operation (epoch '1' above) should be constant over the whole season if the intrinsic responsivities of the bolometers remain constant. To quantify this, a line is fit to the data for each bolometer in both of the telescope rotations for $T \leq 117$. The derivative of this model is never more than a tiny fraction of a per cent of the total amplitude of the signal, and never significant. We therefore believe that the change in the intrinsic responsivity of the 150 GHz bolometers over time is consistent with zero. Likewise, the instrumental conditions – that is, some type of degradation or systematic change of an instrumental component – exhibit no trends over long time scales. This is quite an exciting result, and allows us to confidently perform some kind of average calibration as discussed in Section 5.5. In particular, the lack of any identifiable structure in the loading corrected gain ratio suggests that the instrument is highly stable once the data are corrected for changes in the ambient conditions.

The situation for the 100 GHz detectors is much the same, although as for the elevation dips, structures due to the atmosphere are suppressed at this wavelength. The rip in the polarizer is less obvious at this wavelength as well. There is again a clear separation between telescope rotation angles which is more obvious at this band than 150 GHz, particularly in the ratio of gains. Again, a line has been fit to the CMB observation epoch, and the results at 100 GHz are consistent with no change over the season.

A final effect that may cause temporal variation of V_A on long time scales is the temperature behaviour of the thermal calibration source. This source is a 50 mm diameter, 5 mm thick disc of silicon carbide granules embedded in carbon–loaded stycast. This recipe has been shown to emit thermal radiation with a spectrum very close to that of a blackbody at an equivalent temperature by members of the SPIRE consortium (P. Hargrave, private communication); this assertion was checked in the lab by the QUaD team and shown to be correct. The black disc calibration source includes a thermometer mounted to the back of the disk; the temperature data plotted in Figure 5.16 were recorded using this device.

The calibration source temperature hovers around 285 K during the winter of 2005, with a maximum variation of about 5 K over the entire observing season. This means that, in principle, the power can vary by about $5 \text{ K}/285 \text{ K} \times 100 = 1.7 \%$ over the season, leading to a similar size signal variation in the data. However, the Pearson correlation coefficient between these two quantities is less than 0.1 for any bolometer, suggesting that the temperature variation of the source and the V_A time series are not connected. Evidently, the effect of the variation of the calibration source's temperature is subordinate to other sources of error in these data.

¹²It is not obvious why the signal *increased* rather than decreased, but the optical environment behind the secondary mirror following the rip event was quite complex.

5.4. ROW CALIBRATIONS



Figure 5.16: Calibration Source Temperature for the 2005 CMB Observing Season This plot shows the calibration source's temperature over the 2005 observing season; the grey band shows the refocus.

A final mystery here is that a comparison of g_{ed} and V_A ratios from horn to horn shows there is no correlation between the two calibration methods (this is discussed further in the next section). This is somewhat disconcerting; although the calibration values for a given bolometer can take any value, we expect the calibration ratios to track one another relatively well. The obvious solution to this is that the calibration source does not uniformly illuminate the focal plane, meaning that the signal one bolometer pair 'sees' is different than that another bolometer pair 'sees' during a calibration source run. If the signal were dominated by the illumination pattern differences rather than the intrinsic differences between bolometers, the elevation dips and calibration source runs would not have correlated structure over the focal plane. This hypothesis has been borne out by extensive testing, and the cause of the illumination difference has been traced to various faults in the design of the calibration source. Due to this, the calibration source can only be used on a horn by horn basis, and is not useful for inter-horn calibrations.

5.4 ROW CALIBRATIONS

Row calibrations are another type of calibration observation performed on a regular basis with QUaD. These involve exploiting the symmetry of the QUaD array by lining a single row of the array up at the elevation of RCW 38, and scanning back and forth over the source once in azimuth. This yields two blips in each bolometer's time series (which were used to analyze the beam shape over time in Section 4.4). The telescope is then stepped up in elevation to the next row of bolometers, and the azimuth scan is repeated; each row (of

which there are seven in all) is scanned over RCW 38 in this way. These row calibration observations are performed 4 times per CMB observing day as part of the main calibration block (see Section 6.1).

In the absence of the atmosphere and with a perfect source and telescope, this would comprise an ideal astronomical calibration, since the integral under each blip should be constant. However, a number of factors complicate the use of this type of observation as a perfect calibration. These can be broken into three issues: the atmosphere, the telescope and the source. The first of these, the atmosphere, has two major effects on the observed blip (besides adding noise to the time series). Firstly, it causes a variable optical load in the bolometers, which changes their responsivity. Secondly, it reduces the observed flux of the source via atmospheric extinction. The second two issues are intimately tied together; firstly, RCW 38 (the row calibration source) is known to have complex structure surrounding it. It is therefore not perfectly modeled by the telescope's PSF, and so any reasonably simple model for the blip shape cannot be entirely correct. Furthermore, instrumental effects like a pointing error or variability in the beam shape can cause the model to be particularly erroneous. This is all tied to the telescope's performance: the telescope needs to have accurate absolute pointing in declination to pass right over the middle of the source. Otherwise, the brightness of the main lobe of the source is suppressed while the background flux contamination becomes a bigger problem. Likewise, knowledge of the telescope's beam needs to be precise, or it is impossible to relate the measured beam to the shape of the blip in the time series. All of these factors cause the row calibrations to be non-ideal in terms of providing a consistent, time-independent check on the performance of the instrument's gain. In an attempt to correct for these effects and make the row calibrations a useful measurement of the system's performance, we have adopted a somewhat complex analysis method for this type of calibration observation. It is composed of three major phases: blip extraction, blip correction, and source modeling.

Blip extraction is actually the simplest of the three phases. First, the blips are located via the method introduced in Section 4.1.2. These are then fit to an one dimensional Gaussian function with variable mean, amplitude, width and offset, as in Equation 4.1. Because the row calibrations only involve constant elevation scans, the blip finding algorithm is quite accurate (provided the telescope is pointed accurately enough to give a high instantaneous signal to noise ratio on the source). The height V_{blip} and width σ_{blip} of the blip are stored, along with parameters like the time of observation, telescope rotation angle, et cetera. The time series of uncorrected V_{blip} for two bolometer pairs, one at each wavelength, is shown in Figures 5.17 and 5.18¹³.

¹³The time series of σ_{blip} for an example bolometer pair is actually shown in Figure 4.16: although there is some variation in this parameter, it is necessary to account for it only during the source modeling stage (see



Figure 5.17: Row Calibration Measurements for the 2005 Observing Season for Horn 150-2 This plot shows the time series of row calibration derived V_{blip} for bolometer pair 150-2-A (top panel) and 150-2-B (middle panel). The bottom panel shows the ratio of the two (A/B). In this plot, red points are for $\theta = +57^{\circ}$ and blue points are for $\theta = -3^{\circ}$. The dark grey band shows a period of poor pointing (meaning that the source model cannot be reliably applied) and the light grey band shows the telescope refocus. The discontinuity at the telescope refocus is due to the refocus itself: a telescope in better focus has a smaller beam width, but a larger peak signal for a given (time constant) point source.



Figure 5.18: Row Calibration Measurements for the 2005 Observing Season for Horn 100-2 This plot shows the time series of row calibration derived $V_{\rm blip}$ for bolometer pair 100-2-A (top panel) and 100-2-B (middle panel). The bottom panel shows the ratio of the two (A/B), and the colour allocation is the same as in Figure 5.17. Again, the discontinuity at the refocus time is due to the improvement of the telescope's beam shape, although at this wavelength the change is less pronounced (due to the larger size of the beams at 100 GHz relative to that at 150 GHz).

Once the full set of blip parameters is found, they must be corrected for the effects mentioned above. Logically, these effects must be undone in reverse order to that in which they affected the blip size and shape. Thus, the first step in this process is to correct for the variation in detector responsivities due to optical loading variation, as this is the last effect on the measurement of photons from the astronomical calibration source. This is performed in a similar manner to that used with the hardware calibration source; the elevation dip g_{ed} for a given bolometer pair is median filtered with box width 10 points and the resulting time series is interpolated at each of the row calibration observations. This interpolated time series is then fit to the row calibration V_{blip} values using a linear model; the preand post-refocus data are considered separately in this case. The resulting fits are then subtracted from the V_{blip} data; because this would cause the data to scatter about zero, the median of the pre- and post-refocus V_{blip} is preserved in this process. This provides the correct offset for the V_{blip} since they should be independent of the loading responsivity changes over long time scales. This procedure flattens the points in Figures 5.17 and 5.17 slightly.

The next step in the correction algorithm is to multiply the time series by the atmospheric extinction correction factor. To do this, the relation Equation 5.7 is applied to the $350 \,\mu\text{m}$ tipper data to determine the τ values at 150 GHz. This time series is then interpolated at each row calibration observation to yield a time series of $\tau_{150 \,\text{GHz}}$. The situation for 100 GHz is simpler: the constant value assumption $\tau_{100 \,\text{GHz}} = 0.046$ is used. The loading corrected V_{blip} time series is then multiplied by the vector $e^{\tau_{\nu} \sec z}$ to correct for the effect of atmospheric extinction. At this stage, the blips have been corrected to have the amplitude that would have been measured by an instrument with constant responsivity situated above the atmosphere.

Next, we need to adjust for the 'observational' effects on the blip amplitude: beam shape and pointing correction. The first step in this process is to create a model beam based on our knowledge of the beam's temperature and θ angle dependence for each bolometer individually. To achieve this, the beam is parametrized in these two variables based on the model presented in Section 4.4. For each row calibration observation, the external temperature and telescope rotation angle are fed into this model, which returns a beam shape. Then, the elevation pointing offsets discussed in Section 4.1.2 are used to cut across the model along the path the telescope took; this is typically a correction of up to an arc minute away from the peak of the source (which, based on the beam size, is equivalent to a fairly large decrease in blip height). The loading and extinction corrected blip height is then used to normalize this cut along the beam model, which normalizes the model itself.

At this stage, we have a model of the image the telescope would have made of

below).

RCW 38 normalized to telescope units (Volts/pixel). Since the flux of RCW 38 is known at QUaD's two wavelengths (see Table 4.1.1) this estimate can be used to find a flux conversion factor (FCF), which is the number one multiplies the time series by to calibrate it in Janskys (which can in turn be converted to Kelvin using an estimate of the source's frequency spectrum, if desired). This is done by summing the number of Volts in the same aperture as defined in the sum to determine the total flux of RCW 38. The flux sum is then divided by the sum of QUaD observed Volts. In the case of the fluxes listed in Table 4.1.1, this aperture is defined to be 2 BOOMERanG beam widths across, or d = 8.0' at 150 GHz and d = 15.2' at 100 GHz. The FCFs derived in this way give the number of Jy/V in an astronomical point source for each bolometer; plots of the time series of these quantities are shown in Figure 5.19 for a 150 GHz bolometer pair and Figure 5.20 for a 100 GHz

These plots exhibit a number of interesting features common to all of the light bolometers at a given wavelength. All of the bolometers show the discontinuity occurring as a result of the refocus; this discontinuity is due to the fact that the beam width changed drastically after the refocus. At 150 GHz, the beam width decreased by almost a factor of $\sqrt{2}$, leading to a factor of two increase in the peak height (and thus reducing the calculated calibration factor for the time series). At 100 GHz the difference is less marked, but certainly apparent. Points around the poor pointing period ($26 \le T \le 36$ days) are missing because poor pointing causes the blip height to be so small as to be indistinguishable from the noise in the time series. Clearly, this method is more successful at 100 GHz where the required corrections are not nearly as large as at 150 GHz, although this method can provide an FCF for any light bolometer in the array.

The ratio of FCFs can be compared with the ratio of g_{ed} and V_A for each pair of bolometers. As with the comparison of g_{ed} and V_A discussed above, the FCF ratios do not agree with the loading corrected V_A at all. This is again presumably due to the variation of the hardware calibration source's illumination pattern across the array, and is not a major concern as a result. However, when the FCF and g_{ed} ratios of the bolometer pairs are compared, a slightly more worrying result is found. There appears to be a systematic difference from unity in these two ratios of about 5 % at 100 GHz and up to 15 % at 150 GHz. Because of this wavelength dependence, it seems as though the difference may have to do with some mismatch in the A versus B beams, although the physical cause of such an effect is not clear (see Section 4.4.2). Better models based on more data may lead to an understanding of the effect. In any case, because of this, the row calibrations are not thought to provide very good relative calibrations, although they *do* yield a reasonable estimate of the absolute calibration for each bolometer individually.



Figure 5.19: Row Calibration Flux Conversion Factors for the 2005 Observing Season for Horn 150-2

This plot shows the FCFs for bolometers 150-2-A (top panel), 150-2-B (middle panel) and their ratio (bottom panel) for the 2005 observing season. The dark grey band shows a period of poor pointing during which the FCF model fails, and the light grey band shows the telescope refocus. The dashed red lines show the mean of the FCF values/ratios for preand post-refocus data. These values are constant with time, lending weight to the assertion that the gain of the telescope is well modeled as a constant.



Figure 5.20: Row Calibration Flux Conversion Factors for the 2005 Observing Season for Horn 100-2

This plot shows the FCFs for bolometers 100-2-A (top panel), 100-2-B (middle panel) and their ratio (bottom panel) for the 2005 observing season. The dark grey band shows a period of poor pointing during which the FCF model fails, and the light grey band shows the telescope refocus. The dashed red lines show the mean of the FCF values/ratios for pre- and post-refocus data. Again, these values are constant with time, and have even less scatter than their 150 GHz equivalents.

5.5 ABSOLUTE CALIBRATION

A standard way to calibrate FIR/mm instruments is to map some source with a known flux (or alternatively, known effective temperature and frequency spectrum), and to find the FCF associated with the measurement. This procedure will allow a good measurement of the overall calibration of the instrument, usually good to 5 or 10 % per observation. Using the measurements of RCW 38, we can perform such a procedure with QUaD, particularly using something like the blip-determined FCF, which provides a relatively frequent measurement of the calibration of the instrument.

Sadly, a number of complexities make this a poor calibration method. Firstly, the actual flux of RCW 38 is unknown at the level of 10%. This causes it to be a poor absolute calibration source, since in this case QUaD's final calibration would be limited to the uncertainty in RCW 38's flux at best. Furthermore, the mm spectral energy distribution of RCW 38 is not well known, and so a great deal of uncertainty in the final calibration in Kelvin would be caused by the uncertainty in F_{ν} of the source. In addition, it is unknown how polarized RCW 38 is - estimates allow an upper limit of about 6%, but this would also add to our final calibration uncertainty. Finally, the CMB is a beam–filling source; thus corrections must be applied to calibration factors determined from point–like sources such as RCW 38.

Fortunately, a different calibration method for CMB data exists. As the sun moves around the galaxy, this induces a dipole in the observed CMB due to the Doppler shift from the sun's motion. This signal has amplitude T_0v/c , where $T_0 = 2.728 \pm 0.010$ K (Mather et al., 1999), v is the velocity of the motion, and c is the speed of light. Fixsen et al. (1996) measure this dipolar term to be $T_D = 3.369 \pm 0.004$ mK using the FIRAS instrument on COBE. For an arbitrary velocity relative to the sun's rest frame, the v mentioned above is just the vectorial sum of the two velocities; this makes it possible to determine the dipole for arbitrary velocities (that is, that of a satellite, or Earth itself). One can use this simple physics to calibrate a CMB telescope based on the dipolar signal itself.

Unfortunately, QUaD's direct measurement of the dipolar radiation field is limited by the fact that only a small patch of the sky is ever observed, and that signals on the very largest scales are heavily filtered by the data analysis strategy. However, other experiments, such as WMAP, use the dipole to calibrate their all sky maps. Using a very complex and precise model of their instrument and data, they conservatively estimate a final absolute calibration uncertainty of 0.5 % (Hinshaw et al., 2003). This is based on an iterative solution to the map and relative gains, combined with the 0.1 % uncertainty on T_0 and accuracy of their velocity with respect to the sun of $< 1 \text{ cm s}^{-1}$. This means that a comparison of the WMAP maps of the QUaD survey region and the QUaD maps should yield an absolute

	Calibration Factor S	
ν	$\theta = +57^{\circ}$	$\theta = -3^{\circ}$
100 GHz	0.523 K/V	0.511 K/V
150 GHz	0.449 K/V	$0.445\mathrm{K/V}$

Table 5.1: Absolute Calibration Factors for the 2005 Observing Season

calibration for QUaD to very high accuracy.

However, yet another impediment exists for this calibration scheme. As WMAP creates maps of the whole sky with an angular resolution more than a factor of 2 larger than QUaD's, the overlap of the data in ℓ space is very poor, meaning that such a comparison is strongly influenced by the spatial filter imposed on the sky brightnesses by QUaD's scanning strategy. Nevertheless, the comparison between the QUaD map and WMAP map of the same region has been attempted, and is found to be very poor due to the small region of ℓ overlap of the two experiments (N. Rajguru, private communication). Instead of comparing directly to WMAP, a proxy with intermediate angular resolution and survey size has been found: BOOMERanG. The B03 flight's 145 GHz temperature maps, absolutely calibrated to an error of 1.8 % (Masi et al., 2005), can be used to measure the calibration factor $S = V_{\text{OUAD}}/T_{\text{CMB}}$ via

$$\frac{1}{S} = \frac{a_{\ell m}^{\text{QUaD}} F_{\ell}^{\text{QUaD}} \times a_{\ell m}^{\text{B03,Ref}} F_{\ell}^{\text{B03,Ref}}}{a_{\ell m}^{\text{B03,Cal}} F_{\ell}^{\text{B03,Cal}} \times a_{\ell m}^{\text{B03,Ref}} F_{\ell}^{\text{B03,Ref}}}.$$
(5.13)

This equation defines the cross power spectrum between the calibration maps and the QUaD maps. Here, F_{ℓ} is a weight function describing the filtration of the sky brightnesses by the telescope scanning strategy and data analysis pipeline (see Section 6.3 for further details). The superscripts 'Cal' and 'Ref' refer to subsets of different 145 GHz channels of the B03 data¹⁴. This method works well with the B03 data; because the ℓ range overlap between the instruments is similar, the calibration is well modeled by a single scalar multiple for $100 \leq \ell \leq 1000$. For $\ell > 1000$, the same scalar is assumed, since there is no way to determine the proper calibration factor using available data.

It is found that the calibration factors change slightly as a function of telescope rotation angle. This is unfortunate, but not totally surprising based on the fact that all of the other calibration numbers (as well as the effective beam's shape on the sky) vary with θ as well. The cross spectra are therefore calculated for the two θ angles separately; these are listed in Table 5.5 for the 2005 season.

¹⁴Map 'Ref' is composed of the calibrated, noise weighted co-added map of channels W1, W2, X2 and Z2, while 'Cal' is composed of the same for channels X2, Y1, Y2 and Z1. This choice is made based on the noise performance of the B03 detectors, and is designed to give equal statistical noise in the 'Ref' and 'Cal' maps.

These factors allow one to convert a map in V to one in K. Because the CMB is Planckian to a very high degree and our band shapes are well known, performing the calibration at either of our bands is a relatively simple procedure. Ideally, these calibration numbers should be checked against ones determined from another method – for example, the FCFs determined from RCW 38 – for consistency. If they agree, this will not only affirm that these factors are correct, but also that the telescope's beams are behaving in a way that is well modeled by the analysis pipeline. However, this is a complex exercise, and will be left as work for the future.

6 MAP MAKING, POWER SPECTRUM ANALYSIS & RESULTS

Once the instrument's performance has been characterized in detail, analysis of the CMB data can proceed. The first step in this process is the construction of maps of the CMB anisotropies as detected by the telescope (both in temperature and polarization), and the second is power spectrum estimation from the resulting maps. This second step includes complexities like instrument simulations to understand the effect of the telescope's scan pattern and other characteristics of the instrument on the calculated power spectra and their error bars, as well as final parameter estimation.

In this chapter, the QUaD CMB survey philosophy and implementation is described first. This leads to a description of the map making algorithms and presentation of first year temperature and polarization maps. Then the power spectrum estimation method including a brief description of the instrument simulations is described, and results from the first year of CMB observation are presented. Finally, other astronomical results from the first and second seasons are presented to highlight the fact that QUaD is versatile enough to allow studies of a variety of interesting astrophysical sources.

6.1 SURVEY DESCRIPTION

QUaD's survey philosophy is determined by the following considerations:

- 1. The desire for a constant-elevation scan to eliminate the effects of atmospheric loading variation and focal plane temperature changes on the calibration of the detectors.
- 2. The requirement of keeping above $el = 45^{\circ}$ and the telescope rotation in the range $-60^{\circ} \le \theta \le 60^{\circ}$ during CMB observation scans in order to keep the sorption cooler cold for a full day.
- 3. Negligible astrophysical foregrounds in the survey region to reduce contamination.
- 4. The ability to reject ground signal from the data stream & maps.
- 5. Maximization of sensitivity to *E*-mode and *B*-mode power in the ℓ range which addresses our scientific goals.
- 6. The physical scanning speed of the instrument (that is, how fast can the telescope be driven, and for how long), coupled with the desire to bring the ℓ modes of interest into the measurement band.
- 7. Consideration for the loss of CMB scan efficiency to calibration, telescope rotation, and down time for cryogenic servicing.

Clearly, the survey design which takes all of these issues into account is complex indeed. In some sense, the overriding constraint on the survey strategy is due to time: it is extremely desirable to have at least a 24-hour cryogenic duty cycle, and to continue to run the telescope for as long as possible during the austral winter. The solution to the latter is dictated by the length of time the liquid helium supply at the south pole can last; the task of the QUaD team (and particularly the winter-over members at the south pole) is to maximize the total number of days of CMB observation by maximizing the available helium. Because the survey design is a function of many variables, it took several months in the beginning of 2005 to settle on the strategy presented below; this is why the telescope was commissioned in February 2005, but did not begin final CMB scans until May 17, 2005. The final survey strategy represents a conservative balance between the need to maintain and characterize the instrument, and total number of hours spent surveying the CMB.

The specific strategy is composed of a number of general elements which are designed to address the considerations listed above. Firstly, it is designed to be modular, in that the full daily (24 h) duty cycle is composed of a set of multiply repeated observation blocks. These include frequent calibrations to track the relative gains of the bolometers as well as sky dips and optical pointing checks, which occur once per day. The second strategy is to divide the day into two telescope rotation angles; this allows the comparison of data taken in either rotation on a frequent basis. If the schedule were changed to measure, say, in only a single θ every day, the instrumental parameters will have changed between the two θ angles much more than if both angles are performed on a single day.

Another major consideration is ground pickup; as shown in Section 4.5, ground contamination is an issue for QUaD. A conservative observation strategy is the 'lead/trail' technique. In this technique, CMB scans are performed for some length of time in the lead field. After a section of lead field is completed, the telescope is moved back the same

Table 6.1: CMB Survey Observation Block Summary				
Observation Block	Observation Type	Duration	Bit Mask	
Setup Block	Initialization	$\sim 10 \mathrm{m}$	-	
	Sky Dip	9 m	2^{7}	
	Optical Pointing Check	$\sim 35{ m m}$	-	
Calibration Block	Elevation Dip	50 s	2^{3}	
	Calibration Source	19 s	2^5	
	Load Curve	254 s	2^{6}	
	Pointing Cross	325 s	2^4	
	Row Calibration	1010 s	2^{2}	
CMB Scans	Calibration Source	19 s	2^5	
	Elevation Dip	50 s	2^3	
	8 Constant-el Half Scans	$8 \times 40 \mathrm{s}$	2^{1}	
	Step down in Elevation 1.2'	$\sim 10~{ m s}$	-	
CMB Scan Block	$4 \times \text{CMB Scans}$	26.6 m		
	Change to Trail Field	3.4 m		
	$4 \times \text{CMB}$ Scans	26.6 m		
	Change to Lead Field	$3.4\mathrm{m}$		

Table 6.2: CMB Survey	24-h Duty Cycle	Summary
Telescope Activity	Duration	Fraction of

relescope meanly	Durution	
		Day Spent
Cryogen Service	4.65 h	19.4%
Setup Block	56 m	3.9%
Calibration Block	29 m	2.0~%
$8 \times$ CMB Scan Block, $\theta = -3^{\circ}$	8 h	33.3%
Calibration Block	29 m	2.0~%
Telescope Rotation,		
Settle & 'Unwind' [†]	$\sim 30{ m m}$	2.1~%
Calibration Block	29 m	2.0~%
$8 \times$ CMB Scan Block, $\theta = +57^{\circ}$	8 h	33.3%
Calibration Block	29 m	2.0~%

[†]The unwind occurs at a random time of day, which depends on the starting azimuth of the telescope and the az range rail.

amount in LST as had elapsed while observing the lead field (ΔT). This move puts the telescope in the trail field, which is separated by ΔT from the lead field in right ascension but is *exactly the same* on an azimuth scan by azimuth scan basis. This puts the same ground signal (provided it is not time-varying on the offset time scales or shorter) in the telescope's side lobes, which means that maps of the two fields can be subtracted to remove the ground contamination. Of course, the drawback of this method is that the map thus created is not a simple representation of the true sky brightnesses, but this is not important for power

spectrum analysis.

In QUaD's case, $\Delta T = 0.5$ hour, meaning that a single lead/trail scan block requires one hour to complete (see Table 6.1). Eight of these scan blocks comprise a complete scan set for a single telescope rotation angle, and one block for both telescope angles are completed in a day. The field moves by a step of 1.2' down in elevation after every four full CMB azimuth scans, so that 38.4' in elevation are covered daily. The elevation overlap between days is 75 %, meaning that 40 days are required to cover the entire elevation range $-51^{\circ} < \delta < -43^{\circ}$ (of which only 6.4° is covered completely by any given bolometer pair). The survey coverage begins at the top of the field (in elevation) and works down until the entire field is measured. This is repeated as many times as can fit into the observing year (slightly over twice for the 2005 season, and projected to be more than 3 times for the 2006 season).

Field selection is another free parameter in the survey design. As discussed previously, a field away from the galactic plane which minimizes the foreground emission is of primary importance. Secondarily, a field containing a few bright quasars and AGN which can be used as sources for pointing performance checks, beam mapping, et cetera is desirable. Finally, a field which has been measured by other CMB experiments is attractive, as it allows comparison of our results directly to theirs (for the purposes of absolute calibration, for example). For all these reasons, the QUaD survey has settled on the region discussed in Bowden et al. (2004), which has all of the features listed above. It is predicted to be the lowest-foreground region of the southern sky, is high in elevation from the south pole, contains a number of known sources, and has been mapped by other CMB experiments including ACBAR and BOOMERanG. This field is essentially the natural choice for an extragalactic survey of ~ 100 square degrees from the south pole.

The final survey region for the 2005 season covers approximately 80 square degrees, and comprises 96 days of CMB observation. However, due to bad weather, contamination from the moon, high noise levels in the electronics and other instrumental errors, this list is culled to 67 full days for the final CMB map. All of the bolometers are used in the analysis, except for dead bolometers, or those with large ramping, noise or cross talk (the list of *excluded* pairs is: 100-3, 100-10, 150-12, 150-16 and 150-17). The number of good days and bolometer pairs is expected to increase by an appreciable factor for the 2006 season.

6.2 MAP MAKING

Map making begins with the raw time series, which must be preprocessed before they are useful in a standardized data analysis pipeline. First, corrections due to the instrument's pointing model are applied to the pointing data, which are interpolated to 100 Hz sampling. Then, the data are decimated to 20 Hz sampling using the procedure discussed in Section $4.7.2^1$, and reduced to 1024 constant velocity half scans of 30 s duration for each bolometer. The mean of each of these half scans is removed and the resulting time series are multiplied by -1 to account for the instrument's voltage parity. The relative calibration is applied by determining the nearest (in time) elevation dip measurement for each bolometer, with bolometer 150-14-A acting as reference for the 150 GHz bolometers and 100-2-A the reference for the 100 GHz bolometers (that is, the bolometers are all multiplied by the ratio of their gain to the reference channel's gain). Finally, the time series are all multiplied by the function of τ (usually calculated based on Equation 5.11 at 150 GHz and constant at 0.046 at 100 GHz) which accounts for atmospheric extinction. In addition to this, a spikechecking algorithm scans the data for many- σ events due to cosmic ray hits. Half scans containing spikes are flagged in a scan mask and rejected from further analysis. At this stage, the data are considered to be 'preprocessed'.

The first stage of the true map making pipeline is to subtract a polynomial from each half scan for each bolometer. This serves to remove a great deal of the medium-scale atmospheric variation in the time series which heavily masks the CMB signal. During polynomial subtraction, bright point-like sources are masked out of the fit based on a 'source mask'². Typically, a third order polynomial is used in the fit, although this can be varied. A current topic of study is the effect of raising this polynomial to fifth order on the final QUaD results; although this may clean the atmosphere out of the data more effectively, it can also heavily filter the CMB signal. After the polynomial has been subtracted, the pipeline diverges for the temperature and polarization analysis. We discuss both of these in turn below.

6.2.1 **TEMPERATURE MAPS**

The first step in the temperature map analysis pipeline is to sum the A and B bolometers in each pair. Assuming the polarization efficiencies of both bolometers are equal, this yields the solution given in Equation 2.15, meaning that the sum of the bolometer pair is a representation of the (scalar) temperature on the sky. With this identification, we can employ the map making formalism developed in Section 4.2, culminating in Equation 4.3. However, simply applying this equation to the QUaD data set will not work: the data set is so large that this formalism is intractable to implement on many, if not all, computers.

¹This is done to all the time series sampled at 100 Hz, not just the bolometer data, since they must all match temporally.

²This mask is developed by making the maps without any mask, and determining the locations of bright sources.

Instead, a hybridized approach is required. In this approach, the following assumptions are made:

- That each day can be treated as an independent measurement, and that each singleday map can be co-added using some method to produce the whole year map.
- That the bolometer pairs can be treated similarly *i.e.* individual maps can be made and co-added over the whole array.
- That the noise is white and described by the variance of each half scan for each bolometer. This allows one to construct a diagonal N with entries σ²_{i,b}, where i denotes the time sample and b the bolometer sum time series in question. The σ²_{i,b} for each half scan are set to be the variance of scan; they are thus different for each scan and bolometer. This assumption simplifies N to be a number of time samples long vector, which simplifies its storage and the required arithmetic.

Under these assumptions, Equation 4.3 effectively simplifies to a weighted average of all the data for a single bolometer. Although this does not include optimal weighting as the cross correlations are ignored, in the limit that these cross correlations are negligible, this reproduces the optimal map. For such a map, the noise vector is given by

$$\boldsymbol{\Sigma}_{i}^{-1} = (\mathbf{A}_{i}^{\mathrm{T}}\mathbf{A}_{i})^{-1}\mathbf{A}_{i}^{\mathrm{T}}\mathbf{N}_{i}\mathbf{A}_{i}(\mathbf{A}_{i}^{\mathrm{T}}\mathbf{A}_{i})$$
(6.1)

Now, to co-add these individual per bolometer per day maps (denoted by m_i), we can define a new data vector s as

$$\mathbf{s} = \begin{pmatrix} \mathbf{m}_0 \\ \mathbf{m}_1 \\ \mathbf{m}_2 \\ \vdots \\ \mathbf{m}_{n-1} \end{pmatrix}. \tag{6.2}$$

In a similar fashion, define B to be a vector made up of pointing matrices as

$$\mathbf{B} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2, \cdots, \mathbf{A}_{n-1}) \tag{6.3}$$

and S in an analogous way for N_i . The final co-added map c can be found via

$$\mathbf{c} = (\mathbf{B}^{\mathrm{T}} \mathbf{S}^{-1} \mathbf{B})^{-1} \mathbf{B}^{\mathrm{T}} \mathbf{S}^{-1} \mathbf{s}.$$
 (6.4)

Of course, things become more complex if the instrument and data are not well

6.2. MAP MAKING

behaved. For example, the elliptical beam may have an effect on the final maps due to the two telescope rotations. In addition, their time variation can present an issue. Formally, one must degrade the resolution of the final maps to match that of the *worst* beam performance, and then co-add them³. In fact, co-adding subsets of days to test the effect of various systematic errors is a requirement; due to this, a great number of maps can be created by the various day and horn combinations available.

However, working under the assumption that the data are well behaved, maps like those shown in Figure 6.1 result. These maps represent QUaD's cleanest estimate of the CMB temperature anisotropies based on the data taken in 2005.



Figure 6.1: 2005 Season Temperature Anisotropy Maps

These images are maps of the CMB temperature anisotropy at 100 and 150 GHz (left and right panels, respectively). The colour stretch is set to saturate at $300 \,\mu\text{K}$ in each. The striping effects of heavier atmospheric contamination can be seen by comparing the 150 GHz map with the 100 GHz map. The coverage of the 150 GHz map is slightly larger than that of the 100 GHz map as the array is wider and taller at the higher frequency. Point sources are also visible in these maps; these are discussed further below. (Maps courtesy M. Brown).

A couple of features of these images are worth noting. The first is that these maps are not apodized by the pixel hit function, meaning that the noise increases nearer to the edge of the map due to the decreased integration time there. Furthermore, the striping in the azimuthal direction due to the effect of the atmosphere is obvious, particularly at 150 GHz. Despite these noise effects, the temperature anisotropy of the CMB is clear in these maps, and is obviously correlated between frequencies. Because of the telescope's scanning pattern, there is no cross-linking in the elevation direction in these maps, meaning that structure in the vertical direction is affected by assumptions about the mean of each

³In actual fact, for QUaD the MASTER scheme discussed in Section 6.3 negates this requirement by simulating the effect of a time varying beam on the final power spectra.

scan imposed during data analysis. However, this does not pose a problem for power spectrum estimate, and only inflates the error bars by $\sqrt{2}$.

Variance maps can also be created using these data. These are formed by calculating the variance of the data going into each pixel, weighted by the estimated bolometer noise. This procedure yields estimates of the weight for each pixel in the power spectrum calculation. The typical per pixel variances in the temperature maps are $1100 \,\mu$ K at $100 \,\text{GHz}$ and $950 \,\mu$ K at $150 \,\text{GHz}$.

A few bright point sources are obvious in these maps. Figure 6.2 highlights the positions of known radio point sources brighter than 0.4 Jy in the 5 GHz Parkes survey⁴; there is a great deal of correlation between the Parkes catalog and our map. These source images can themselves be used for a variety of purposes, including checking our pointing and beam models, filter deconvolution and decimation procedure, and for scientific study of the sources themselves. Figure 6.3 shows the six brightest of these sources in some detail at both frequencies.

Testing the pointing and beam models using these sources is simple: if the models and instrument are well-behaved, then the point sources will appear to be well matched to the telescope's nominal point spread function (PSF). On the other hand if, for example, the pointing jitter were much higher than accounted for in the map making process, or the temperature variation of the beams were very large, then the image of the point source would be smeared from day to day variations in their characteristics. In addition, the axis of the ellipticity in the beams varies as a function of the telescope rotation angle; if the time average of the beams is sufficiently non-Gaussian, this will show up as a mismatch between a Gaussian beam model and the point source images. Figure 6.3 shows the residual images when an elliptical Gaussian is fit to each point source map (in the case of PKS 0454-46, two Gaussians are used as this is known to be a double-lobed source). For most of the source images, the residual does indeed exhibit structure at the ~ 5 % level, meaning that an elliptical Gaussian is a poor match to the time averaged beam around 13 dB. Also, the widths of the source imaged are also ~ 10 % larger than the nominal values at either frequency. However, we expect these results based on the time variation in the beam and the jitter in the pointing of the telescope. More importantly than either of these points, the residuals do not exhibit telescope rotation angle dependent structure, which suggests that the beam fluctuations can be treated as random. This result simplifies power spectrum estimation.

The filter deconvolution procedure can also be checked using these maps. If one makes maps for the positive δ azimuth direction only and also for the negative δ azimuth

⁴See http://www.parkes.atnf.csiro.au/research/surveys/pmn/pmn.html for references.

6.2. MAP MAKING



Figure 6.2: Point Sources in the CMB Survey Field

This figure shows the 2005 100 GHz CMB temperature map with the positions of known radio point sources superimposed as yellow boxes. Clearly, there is a high degree of correlation between the known source positions and the point sources in the maps, although the relation is by no means 1:1. However, all of the point sources detected in the QUaD survey field have counterparts at 5 GHz, implying that we understand the bright point source foreground well.

direction only, the difference between such maps will exhibit a dipolar-like structure at the positions of the bright point sources if the deconvolution is failing. This test has been performed, and the filter function required slight modification to satisfy the criterion that this residual lies below the noise in the map.

Finally, these images can be used to study the sources themselves. Most of these sources are known to be quasars at moderate redshifts. One could, in principle, study per day maps to understand the time variation in the sources themselves. Of course, this would be complicated by the time variability of the instrument, but we believe we understand this issue. The source PKS 0454-46 is interesting; it is a local double-lobed AGN; QUaD resolves some of this structure, and allows estimation of the spectral energy distribution





This figure shows the six brightest sources in the map plotted in Figure 6.2 in detail. The sources are the same for each row of the plot (the name of each appears on the far left of each row). In each of these are shown the source image at 100 GHz, the 100 GHz residual after a beam model is removed, the 150 GHz image, and finally the 150 GHz residual after a beam model is removed. The beam model is determined for each source at each frequency by fitting an elliptical Gaussian to the map image; the residual plots show how well the sources match this model. Note that in the case of PKS 0454-46, two beams are used, since this is known to be a double-lobed source. (Figure courtesy C. Pryke).

of the lobes in combination with other data. A final use of these sources is for polarization measurements; this class of moderately polarized radio source is known to be highly variable on the time scale of weeks. However, measurements of the behaviour of their polarization properties at high radio frequencies are rare: this type of study will be performed with the QUaD data as well.

6.2.2 POLARIZATION MAPS

Polarized map making is a slightly more involved prospect than temperature mapping. As a starting point, Equation 4.3 is still applicable, although it requires modification to handle polarization data. In the simplest case, one could make a map using the difference of the detector pairs separately for the different telescope rotation angles and bolometer orientations (making 4 maps in total). Since these are just linear combinations of 'absolute' Q and U, where the combination prefactors are sines and cosines of the relative polarization angles, one can reconstruct the Q and U maps in absolute coordinates by inverting the linear system. This procedure was used to produce the very first QUaD polarization maps. One could even go a step further, and form the linear combination in the time series rather than map space to produce Q and U maps directly. This process was also used with QUaD data for a time.

However, if we don't want to make the assumption that the PSBs are well aligned or the cross polarization is negligible, a slightly more complex apparatus is required. First, let the angle between an A bolometer and its B partner be $\pi/2 + \delta \alpha_h$ for horn h. Also, given a cross polarization for detector d in horn h of $\epsilon_{h,d}$, define the cross polar leakage term as

$$\gamma_{h,d} = \frac{1 - \epsilon_{h,d}}{1 + \epsilon_{h,d}}.\tag{6.5}$$

With this relation, the effective angle of a given bolometer is defined as

$$\tan(2\beta_h) = \frac{(\gamma_{h,A} - \gamma_{h,B})\sin(\delta\alpha_h)}{(\gamma_{h,A} + \gamma_{h,B})\cos(\delta\alpha_h)}.$$
(6.6)

This new definition means that the nominal polarization angle of a given detector $\alpha_{h,d,t}$ defined for Equation 2.13 is, in reality, equivalent to the sum of this nominal value and the effective angle written above, or $a_h = \alpha_h + \beta_h$. Now, we can define a χ^2 statistic for each pixel as:

$$\chi^2 = \sum_{h,\alpha_h} w_{h,\alpha_h} (M_{h,\alpha_h} - Q\cos(2a_h) - U\sin(2a_h))^2$$
(6.7)

where M_{h,α_h} is a per horn and polarization angle map of the difference data. This equation is essentially asserting that the model $Q \cos 2a + U \sin 2a$ is a good description of the raw differenced data binned into pixels with the appropriate weight. The weight function w_{h,α_h} can be uniform (to weight the maps equally), zero (to disregard a map in the sum), or $1/\sigma^2$ for Gaussian noise weighting; the normalization of Equation 6.7 dictates that the sum of the weights should be unitary.

The next step in this procedure is to differentiate the χ^2 function with respect to Qand U, and set these two functions to zero. This operation allows us to find the values of Qand U that minimize the difference between the data and the model, which are given by

$$\sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) M_{h,\alpha_h,p} = \left[\sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \sin(2a_h) \right] \begin{pmatrix} Q_p \\ U_p \end{pmatrix}$$
(6.8)

for Q and

$$\sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) M_{h,\alpha_h,p} = \left[\sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \sin(2a_h) \right] \begin{pmatrix} Q_p \\ U_p \end{pmatrix}$$
(6.9)

for U. Combining these yields

$$\begin{pmatrix} \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) M_{h,\alpha_h,p} \\ \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) M_{h,\alpha_h,p} \end{pmatrix} = \begin{pmatrix} \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \sin(2a_h) \\ \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \sin(2a_h) \end{pmatrix} \begin{pmatrix} Q_p \\ U_p \end{pmatrix}.$$
 (6.10)

To determine Q and U, all that is required is that we invert this equation via

$$\begin{pmatrix} Q_p \\ U_p \end{pmatrix} = \begin{pmatrix} \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \sin(2a_h) \\ \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \sin(2a_h) \end{pmatrix}^{-1} \times \begin{pmatrix} \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) M_{h,\alpha_h,p} \\ \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) M_{h,\alpha_h,p} \end{pmatrix}$$
(6.11)

which is simple, as only a 2×2 matrix needs to be inverted for each pixel.

Maps made using this method are shown in Figure 6.4 for the QUaD 2005 season. These are directly comparable with those shown in Figure 6.1; although no CMB structure is visible to the eye, these maps contain the Q and U power on the sky. In fact, these are

6.2. MAP MAKING



the maps used to calculate the CMB power spectrum, as described in the next section.

Figure 6.4: 2005 Season Q and U Polarization Anisotropy Maps

These images show the measured Q and U maps at 100 and 150 GHz from the 2005 season data set. The left hand panels show the 100 GHz maps, and the right hand panels the 150 GHz maps. Similarly, the two upper panels show the Q polarization maps, and the bottom panels the U polarization maps. The colour stretch is $\pm 150 \,\mu\text{K}$ in these images. Although no structure is visible by naked eye due to noise, these polarization maps actually do contain a great deal of polarized power from the CMB.

The calculation of the noise in polarized maps is performed by a similar method; we outline the procedure here. Define a pixel polarization vector by

$$\mathbf{P}_p = \begin{pmatrix} Q_p \\ U_p \end{pmatrix} \tag{6.12}$$

and let

$$\mathbf{\Gamma} = \left(\begin{array}{c} \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \cos(2a_h) \sin(2a_h) \\ \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \cos(2a_h) \sum_{h,\alpha_h} w_{h,\alpha_h} \sin(2a_h) \sin(2a_h) \end{array} \right)^{-1}$$
(6.13)

The object of the exercise is to calculate the quantity

$$\Sigma_{p} = \langle (\mathbf{P}_{p} - \langle \mathbf{P}_{p} \rangle) (\mathbf{P}_{p'} - \langle \mathbf{P}_{p'} \rangle)^{\mathrm{T}} \rangle.$$
(6.14)

Under the assumption that the noise is uncorrelated between detectors

$$\Sigma_{p} = \sum_{h,\alpha_{h}} w_{h,\alpha_{h}}^{2} \begin{pmatrix} T_{00}\cos(2a_{h}) + T_{01}\sin(2a_{h}) \\ T_{10}\cos(2a_{h}) + T_{11}\sin(2a_{h}) \end{pmatrix} \times (T_{00}\cos(2a_{h}) + T_{01}\sin(2a_{h}), T_{10}\cos(2a_{h}) + T_{11}\sin(2a_{h}))\sigma_{h,\alpha_{h},p}^{2}, \quad (6.15)$$

results. Here $\sigma_{h,\alpha_h,p}^2$ is the variance of the data going into a single pixel.

6.3 POWER SPECTRA ESTIMATION & COSMOLOGICAL PA-RAMETERS

The concept of power spectra of the CMB radiation as a measurement of the statistical properties of its anisotropies is presented in Chapter 1; here, we discuss the estimation of the power spectra based on our measurements using QUaD. More complete discussions of the formalism and philosophy presented only briefly here can be found in Hivon et al. (2002) and Brown et al. (2005).

If the spatial anisotropies of the CMB are Gaussian distributed, all of the information contained in a map of the CMB can be encapsulated in its angular power spectrum C_{ℓ} , which can be estimated from the $a_{\ell m}$ of a full sky map by

$$C_{\ell} = \frac{1}{2\ell + 1} \sum_{m = -\ell}^{\ell} |a_{\ell m}|^2.$$
(6.16)

There are two classes of tools generally used in estimating C_{ℓ} from a map of the CMB: maximum likelihood and Monte Carlo simulation methods. Maximum likelihood methods are computationally prohibitive for $N_{\text{pix}} > 10^4$, so the QUaD analysis has settled on the MASTER Monte Carlo method described by Hivon et al. (2002). This method is attractive as it is based on the direct spherical harmonic transform of a limited region of the sky, and can incorporate the characteristics of a specific experiment including survey coverage, realistic instrumental noise, beam performance, and non-Gaussianity/stationarity in the noise. In principle, the limit on the accuracy of this method for measuring the C_{ℓ} and their associated errors is imposed by how realistic the instrument simulation is.

In ground-based CMB experiments, it is common for the sky coverage to be some

(relatively small) fraction of the sky with a non-uniform scan pattern. This means that we must employ a 'window function' $W(\hat{\mathbf{n}})$ to describe the position-dependent measurement weight of the true sky brightnesses (for example, a good choice of W would be the inverse of the noise estimate of each map discussed in the previous section). In the presence of W, we are really measuring

$$\tilde{a}_{\ell m} = \int d\hat{\mathbf{n}} \Theta(\hat{\mathbf{n}}) W(\hat{\mathbf{n}}) Y_{\ell m}^*(\hat{\mathbf{n}})$$
(6.17)

where $\Theta \in \{T, E, B\}$. This being the case, we define the pseudo-power spectrum – that is, the spectrum that actually results from applying this window function to the sky – as

$$\tilde{C}_{\ell} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |\tilde{a}_{\ell m}|^2.$$
(6.18)

The relationship between \tilde{C}_{ℓ} and C_{ell} is given by

$$\langle \tilde{C}_{\ell} \rangle = \sum_{\ell'} M_{\ell\ell'} \langle C_{\ell'} \rangle, \tag{6.19}$$

where the average is ensemble-wise (*i.e.* over many realizations of the Universe, or many simulations of the Universe). The mode-mode coupling matrix $M_{\ell\ell'}$ describes the effect of inter-mode coupling which affects \tilde{C}_{ℓ} ; this is purely dependent on the geometry of the sky coverage (Hivon et al., 2002).

Brown et al. (2005) develop a formalism for determining the mode-mode coupling matrix in the case of polarized maps. This generalization deals with the fact that, in polarization measurements, not only are different ℓ modes coupled in the raw \tilde{C}_{ℓ} , but also couplings between (independent) polarization bases can occur. Thus, a family of $M_{\ell\ell'}$ must be used, one for each polarization combination TT, TE, TB, et cetera. Brown et al. (2005) go on to determine expressions for these $M_{\ell\ell'}$, and show that only E and B mode power are affected by this polarization mixing. Unfortunately, this means that one must be particularly careful in constructing the $M_{\ell\ell'}$ for mixing between E and B.

To fully describe the \tilde{C}_{ℓ} which would come from the maps shown in the previous section, we must add the effect of the instrument's beam, experimental noise and filtering of the time ordered data (TOD) during previous analysis steps. These can be encapsulated using the following equation⁵:

$$\langle \tilde{C}_{\ell} \rangle = \sum_{\ell'} M_{\ell\ell'} \mathcal{F}_{\ell'} \langle C_{\ell'} \rangle + \langle \tilde{N}_{\ell} \rangle.$$
(6.20)

⁵Note that, in this expression, we combine the variables F_{ℓ} and B_{ℓ}^2 used in Hivon et al. (2002) into one variable \mathcal{F}_{ℓ} .

Determining our estimate of the power spectrum thus merely boils down to determining the functions $M_{\ell\ell'}$, \mathcal{F}_{ℓ} and $\langle \tilde{N}_{\ell} \rangle$ given our scan strategy and instrumental characteristics⁶.

For QUaD, the MASTER simulations determine both \mathcal{F}_{ℓ} and the error bars on the measurements of C_{ℓ} (given by both $\langle \tilde{N}_{\ell} \rangle$ and the effects of sample and cosmic variance). In contrast, $M_{\ell\ell'}$ is not determined by simulation, but rather merely by applying an algorithm to the (known) window function. The simulations themselves are performed in the following way:

- 1. Maps of the CMB sky in T, Q and U are simulated using a combination of CMBFAST (Zaldarriaga & Seljak, 2000) and HEALPIX (Górski et al., 2005) and the best-fitting Λ CDM cosmological model to the WMAP 3 year data (Hinshaw, 2006).
- 2. Using the pointing data and time-dependent beam models determined from the telescope data, the maps are scanned to determine a set of TOD for the 2005 season. It is very important that the beam models and pointing data are accurate, or the transfer function and noise estimate will be incorrect.
- 3. Correlated noise is introduced into the TOD based on the measured noise properties of the instrument. This is performed in Fourier space to simplify the calculation.
- 4. Maps are made from the noisy, instrument-modeled data using the method discussed in the previous section.
- 5. The TT, TE, TB, EE, EB and BB power spectra of these maps are then computed.

In this procedure, the addition of CMB signal or instrumental noise is optional, depending on the parameter of interest.

To determine the measured power spectrum, the first step is to settle on the survey strategy and weight function to be used. In QUaD's case, this is relatively simple, as these are both dictated directly by the data. Following this, these steps are performed:

- 1. Calculate $M_{\ell\ell'}$ (using the recipe given in, for example, Brown et al. (2005)).
- 2. Produce a number of noise free simulations to determine \mathcal{F}_{ℓ} . In the signal only case, the ratio of the output power spectra to the input power spectra yields the value of this function at each ℓ ; a number of simulations are performed to ensure that a consistent solution is reached (the error on the transfer function decreases as $\delta F_{\ell}/F_{\ell} \propto N_{\rm MC}^{-1/2}$).

⁶There is also a correction which must be applied to the power spectra because of the fact that the map is pixelized. Fortunately, this correction is included in the HEALPIX software, and is transparent to the user (see Górski et al. 2005 for details).

- 3. Produce a number of noise only simulations in the same way; this yields the angular power spectrum of the noise projected onto the sky.
- 4. Use the measured TOD to make a map, and then \tilde{C}_{ℓ} from the data itself.
- 5. Compute the covariance matrix of the experiment using the variation of the simulated power spectra due to the injected instrumental noise. The square root of the covariance matrix's diagonal elements plus the effects of sample and cosmic variance yields the error bars on the measured power spectra.

In actual fact, it is desirable to bin ℓ values, since the signal from several ℓ can be averaged to produce a smaller error bar on a multi- ℓ bin. This procedure requires defining the number of bins n_b , each of which is labeled with index b. Next, define an operator:

$$P_{b\ell} = \begin{cases} \frac{1}{2\pi} \frac{\ell(\ell+1)}{\ell_{\text{low}}^{b+1} - \ell_{\text{low}}^{b}} & \text{if } \ell_{\text{low}}^{b} \le \ell < \ell_{\text{low}}^{b+1} \\ 0 & \text{otherwise.} \end{cases}$$
(6.21)

and its reciprocal operator

$$Q_{\ell b} = \begin{cases} \frac{2\pi}{\ell(\ell+1)} & \text{if } \ell_{\text{low}}^{b} \leq \ell < \ell_{\text{low}}^{b+1} \\ 0 & \text{otherwise.} \end{cases}$$
(6.22)

In this case, the binned power spectrum becomes $C_b = P_{b\ell}C_{\ell}$. Now, write Equation 6.20 in the following form:

$$\langle \tilde{C}_{\ell} \rangle = K_{\ell \ell'} \langle C_{\ell'} \rangle + \langle \tilde{N}_{\ell} \rangle.$$
(6.23)

A small amount of algebra shows that we are now seeking a solution to the equation

$$P_{b\ell}K_{\ell\ell'}\langle C_{\ell'}\rangle = P_{b\ell}(\langle C_{\ell}\rangle - \langle N_{\ell}\rangle).$$
(6.24)

Next, we define

$$K_{bb'} = P_{b\ell} K_{\ell\ell'} Q_{\ell'b'}$$

$$= P_{b\ell} M_{\ell\ell'} \mathcal{F}_{\ell'} Q_{\ell'b'}$$
(6.25)

This leads to our final unbiased estimator of the whole sky spectrum based on the experimental measurements:

$$\mathcal{C}_{b} = K_{bb'}^{-1} P_{b'\ell} (\tilde{C}_{\ell} - \langle \tilde{N}_{\ell} \rangle)$$
(6.26)

where $\langle \tilde{N}_{\ell} \rangle$ is measured from simulations. This estimator can be considered QUaD's best measurement of the true C_{ℓ} . This procedure has been applied to the QUaD data set for

2005 by QUaD collaborators at the University of Edinburgh, and the power spectra shown in Figure 6.5 and 6.6 result⁷.

These power spectra can be combined using an extension of inverse variance weighting by defining the total power spectrum for each polarization combination according to:

$$C_b^{\text{comb}} = \frac{\sum_i C_b^i w_b^i}{\sum_i w_b^i}$$
(6.27)

where the *i* run over three indices: the 100 GHz spectra, the 150 GHz spectra, and their cross spectra. The w_i are the weight functions for each of the spectra, defined by summing the columns of the inverse covariance matrix for each of the spectra, *i.e.* $w_b^i = \sum_{b'} C_{bb'}^{-1.i}$. These spectra define the 'best' measurement of the CMB anisotropies in a multi-frequency instrument like QUaD; the combined *EE* spectrum is shown in Figure 6.6 with all other published measurements of this spectrum for comparison.

The method of determining cosmological parameters from power spectra like these can be quite complex; we review the salient details here (a detailed treatment of this subject can be found in, for example, Verde et al. 2003). Essentially, cosmological parameter estimation uses Bayes' theorem to evaluate the likelihood of a particular cosmological model given the data. The model corresponding to the maximum likelihood is that favoured by the data, and thus the model most consistent with the measurement. The likelihood function is derived by generating the six power spectra corresponding to a given cosmological model, and the χ^2 of the model given the data is calculated using the covariance matrix of the power spectrum. This χ^2 can then be related to the likelihood function. The model parameter space is searched for its maximum by repeating this procedure for a variety of cosmological models. Integrals over the likelihood function also yield the 68 % and 95 % confidence intervals, which correspond to 1 and 2σ error bars in the limit of Gaussian statistics.

In modern CMB models, more than 15 parameters can be varied to produce the set of 6 power spectra for each model. Because fully sampling 15 dimensional parameter space using many thousands of models can be quite time consuming, a Markov Chain approach is generally used to isolate the areas in parameter space where the likelihood function is varying quickly. This reduces the computation time significantly, and is used with the QUaD data.

The preliminary likelihood function used to describe the QUaD data is a function of only six variables; these, with their best fit values for the 2005 season, are listed in Table

⁷At this point, these power spectra are preliminary; the spectra are failing various jack-knife tests which should yield results consistent with zero signal: we thus suspect unaccounted systematic effects are affecting the analysis pipeline. This issue must be resolved before these data can be published.



Figure 6.5: Preliminary 2005 Season CMB Power Spectra

These plots show the preliminary TT, TE, TB, EE, EB and BB power spectra measured by QUaD based on the first year of integration. The blue points show the power spectra based on the 100 GHz maps, and the red points show the power spectra based on the 150 GHz maps. The range is restricted to $\ell \leq 1500$ because complex beam effects in the first year data become important above this. At the time of writing, these data fail null- χ^2 tests, and so will be modified slightly as the understanding of temperature and polarization systematic effects increases.



Figure 6.6: All CMB EE Power Spectra Data to Date

800

l

1000

1200

1400

600

This figure shows the preliminary QUaD combined EE spectrum as well as all other measurements of the EE spectrum in the literature (the latter points are also shown in Figure 1.3; references are given in Section 1.2.3). This plot summarizes QUaD's first year performance quite well; over three years, we hope to improve the errors on this measurement by a factor $> \sqrt{3}$ because of instrument improvements during the second and third years.

Parameter	Fiducial Value	100 GHz QUaD	150 GHz QUaD
h	0.71	$0.67^{+0.14}_{-0.12}$	$0.71^{+0.16}_{-0.14}$
τ	0.09	$0.13^{+0.093}_{-0.087}$	$0.14_{-0.091}^{+0.095}$
n_s	1.00	$1.08^{+0.19}_{-0.18}$	$1.02^{+0.19}_{-0.19}$
$\Omega_{ m b}h^2$	0.022	$0.028^{+0.008}_{-0.008}$	$0.029^{+0.009}_{-0.008}$
$\Omega_{\rm m}h^2$	0.14	$0.19^{+0.05}_{-0.04}$	$0.16^{+0.04}_{-0.04}$
A	0.84	$0.81^{+0.20}_{-0.18}$	$0.82^{+0.16}_{-0.19}$

6.3. These values are computed using the TT, TE and EE spectra only. As the QUaD analysis pipeline matures and the 2006 and 2007 data is combined into the power spectra, these constraints will improve, and new parameters will be added.

6.4 OTHER ASTRONOMICAL RESULTS

0

200

400

As a very sensitive polarimeter at the south pole, QUaD can be used for measurements other than those of the CMB. There are a great many galactic and extragalactic sources that

are thought to be polarized at these wavelengths; this section reviews observations of a few astronomical objects of interest.

The analysis pipeline for these raster map data is similar to that for the CMB data or RCW 38 maps⁸. For these astronomical targets, the starting point in the mapping analysis is the scan by scan, filter-deconvolved time series. Each scan set for each bolometer is relatively calibrated using an elevation dip which occurs at the beginning of each set scaled to the mean over the array at a single frequency. This is then low pass filtered using the simple filter function discussed in Section 4.2 to remove the high frequency noise contribution. A linear model is fit to each scan, and this is then subtracted from each bolometer's time stream for each scan.

At this point, the temperature and polarization mapping analysis chains diverge. For temperature mapping, the sum of each scan for each bolometer is formed. This is then binned into a map using the formalism discussed in Section 4.2. A linear fit is then subtracted from each row in the resulting map to remove scan-synchronous noise and calibrated using an flux conversion factor (FCF) derived using the method discussed below. In the case of dim sources or those with dim, diffuse structure in their environs (which is common for galactic sources), the map can be convolved with the normalized point spread function K given by:

$$K(x,y) = \frac{1}{\pi\sigma^2} e^{\frac{x^2 + y^2}{2\sigma^2}}$$
(6.28)

where x and y are standard Cartesian coordinates, and σ^2 is the nominal width of the symmetric Gaussian beam at each frequency⁹. The convolution draws out extended structure for the eye.

In the case of the polarization maps, the difference of each relatively calibrated scan for each bolometer pair is formed. The formalism developed in Section 6.2.2 is then applied to the data vectors, positional information, and telescope rotation angles. This process yields Q and U maps at each frequency. Again, a linear fit is removed from each row in both maps to remove scan-synchronous artifacts. These maps are then calibrated using the same FCF as for the temperature maps, and convolved with the beam kernel K if desired. To form standard polarization vector plots, the transformations

$$I_{\rm p} = \sqrt{Q^2 + U^2} \tag{6.29}$$

$$I_{\theta} = \frac{1}{2}\arctan(U/Q) \tag{6.30}$$

⁸The exception to this are the data for the Moon; these are sufficiently different to require special consideration (see Section 6.4.1).

⁹This K is actually a matched filter for point sources, and can serve to warp or smear structures larger than the beam. However, it makes for better presentation of low signal to noise ratio data.

are applied to the Q and U maps. Here, I_p is the polarized intensity (length of the vectors) and I_{θ} is the polarization angle. The fractional polarization is given by I_p/T in each pixel.

The calibration of these maps is reported in Janskys (Jy); this is a standard unit in radio astronomy corresponding to 10^{-26} W Hz⁻¹ m⁻². The FCF for the system is found by running RCW 38 calibration runs through the temperature mapping pipeline discussed above. As discussed in Section 4.1.1, the sum of the flux in Volts over the apertures defined for Table 4.1.1 is equated to the flux reported. For high-gain mode, this yields FCFs of 4.2 Jy V⁻¹ arcmin⁻² at 100 GHz and 6.1 Jy V⁻¹ arcmin⁻² at 150 GHz for the 2005 season, with errors of ~ 10 % on both factors.

It should be noted that the radio sources presented in this section are extremely well-studied and famous southern sky objects. No attempt will be made at a review of the (extensive) literature regarding these objects; we will only discuss the salient points with respect to their ~ 100 GHz polarized emission characteristics.

6.4.1 THE MOON

The QUaD observations of the moon are special for a number of reasons. Firstly, because the moon is so bright, these data are taken in low-gain mode. This has a few effects:

- causes the time stream noise to be much lower than in standard observations,
- makes the FCFs derived above inapplicable, and
- gives very high signal to noise ratios for the polarization, which can be used to calibrate the analysis pipeline.

All of this means that a modified analysis pipeline is required for these data. Because the moon is such a high signal to noise ratio object, the relative calibration of the channels can be determined from single-channel maps alone. Functionally, this means that the first step in the analysis pipeline is to make simple temperature maps for each channel using the filter-deconvolved, per scan, median-subtracted data. The sum of these maps in an aperture centred on the image of the moon (in this instance, a circular aperture with a radius of 4 pixels, where each pixel is 3×3 square arcmin) then gives a calibration factor. The mean calibration factor at each frequency can be calculated, and each channel can be scaled according to this.

Once the time streams are calibrated, the next step is to again separate the temperature and polarization mapping algorithms. For the temperature maps, the median subtracted, calibrated, per scan data is simply binned into pixels. This procedure applied to every bolometer in the array yields the two temperature maps shown in Figure 6.7.



unusately, this fast can be used to calibrate the QUAD polynametion a

Figure 6.7: The Polarized Moon at 100 and 150 GHz

This figure shows QUAD images of the moon at 100 GHz (left) and 150 GHz (right). The grey scale component in each is a temperature map of the moon, while the red vectors show the polarized emission fraction in each pixel. Because these data were taken in low gain mode, the calibration is somewhat problematic; the grey scale is essentially arbitrary. However, the polarization fraction can be consistently calculated, and corresponds to $\sim 1 \%$. Because the moon is such a large, bright source, no beam kernel convolution need be applied to these data. Interestingly, the phase of the moon can be seen in both temperature and polarization; the dominant source of radio emission from the moon is in fact reflected solar light, which also explains the radial polarization pattern.

The polarization mapping algorithm again uses the calibrated, median subtracted per scan time series to form the difference between each bolometer pair's time series. These data are then converted into Q and U maps using the formalism discussed in Section 6.2.2, which can then be converted into I_p and I_{θ} maps. These are over plotted as vectors in Figure 6.7.

The polarized properties of the moon at radio wavelengths are relatively well known. As presented recently in Poppi et al. (2002), the dominant contribution to the moon's radio signal is scattered solar radiation, which explains why the lunar phase is evident in Figure 6.7 (these data were taken slightly before full moon). A radial polarization pattern which is zero at the centre of the moon and strong near its limb arises because the Fresnel reflection coefficients have a gradient in the radial direction but are constant in the azimuthal direction. This leads to a preferential change in the polarization in the radial direction which is dependent on the angle of incidence of the solar radiation at the 1 % level or so.

Fortunately, this fact can be used to calibrate the QUaD polarization mapping algorithm. There are a number of complex conventions required to make polarization maps which are calibrated properly compared to the absolute definitions of polarization on the celestial sphere. It is a requirement of any algorithm used in QUaD analysis that it be able to reproduce the polarization properties of the moon. This is a powerful test of our analysis code, and perhaps the best use of the moon data. That these maps reproduce the polarized properties of the moon is an affirmation that the conventions have been applied properly in the polarization mapping algorithm used here.

6.4.2 THE CARINA NEBULA

The Carina nebula is a star forming region in the southern sky, famous for η Carinæ, one of the most massive stars known (having $T \sim 30,000$ K and $M \sim 100 M_{\odot}$). Cox et al. (1995) gives an excellent review of the millimetre properties of this region, which is dominated by the central stellar core and its outflow region. η Carinæ is actually known to be quite variable at millimetre wavelengths: this is thought to arise because the central source is actually a binary system containing *two* massive stars in orbit about one another (Duncan & White, 2003). The variability is caused by the elliptical nature of the orbit, which has a period of about 5.5 years. Because of this, the millimetre flux of η Carinæ can vary by factors of 5 or more over the same time scale.

The QUaD observations of the Carina Nebula are centred on η Carinæ. The data are reduced in the manner discussed in the introduction to this section; the maps produced by this procedure are shown in Figure 6.8.

This figure clearly shows η Carinæ (the leftmost bright blob), and the slightly extended and bright emission to its west. Also visible in the temperature maps is flocculent structure hitherto unknown at these wavelengths in this source. These structures, including the western 'blob' of emission, are typical of HII regions, of which the Carina nebula itself is an example. The dominant emission mechanism in the source corresponding to η Carinæ itself is free-free emission, which causes the observed polarization at the level of a few percent.

6.4.3 G 240.6

Supernovæ remnants (SNR) are the debris of catastrophic explosions which occur at the end points in the lives of the most massive ($\gtrsim 8 M_{\odot}$) stars (see Longair (1994) for an excellent review of the physics at work in these objects). These stars swell to gigantic proportions before exploding; their outermost layers are cast off with large initial velocities, while their



Figure 6.8: Carina Nebula at 100 and 150 GHz

This figure shows images of the Carina Nebula made with QUaD at 100 (left) and 150 GHz (right). The images show the temperature map of the source, while the vectors in red show the polarization amplitude and direction scaled to the line shown at lower left. This source's peak brightness is ~ 1.2 Jy per square arcmin pixel (similar to that of RCW 38).

cores collapse under gravity to produce neutron stars, or in the most extreme cases, black holes. The SNR themselves are composed of the outermost layers of the star, which expand away from the initial object and remain hot due to their high initial velocity. As these layers spread out from the star, they 'plow' into the surrounding interstellar medium, which can cause the creation or destruction of dust at the shock interface. The (sub-)mm emission in these objects is thus due to a competition between synchrotron and dust emission, the relative contribution of each component depending strongly on the exact composition of the progenitor star, the local interstellar medium, and the speed of the SNR shock front.

G 240.6 (also known as Puppis A) was observed with QUaD 16-17 March, 2006 using the standard raster scan observation strategy. The data were processed as described in the beginning of this section and convolved with the PSF of the telescope; the maps shown in Figure 6.9 are result of this processing.

From the QUaD images, it is obvious that the major emission mechanism in this particular SNR is synchrotron radiation. This is because the remnant is clearly visible at 100 GHz, where synchrotron is expected to be dominant, and virtually absent at 150 GHz, where an approximately equal amount of flux is expected to be contributed from synchrotron and thermal dust emission. The results reported in Arendt et al. (1990) and Saken et al. (1992) tend to support this: neither of them find strong dust emission at 100 μ m.

193



Figure 6.9: Super Nova Remnant G 260.4 at 100 and 150 GHz

This figure shows images of Puppis A made with QUaD at 100 (left) and 150 GHz (right). The images show the temperature map of the source, and are scaled according to the colour bar at the bottom of each map. Although the ring-shaped remnant is obvious at 100 GHz, it is not at all clear at 150 GHz; this is due to the fact that very little dust is present in this source. Because of the lack of a clear detection at 150 GHz, the polarization vectors are only drawn for the 100 GHz map. These are scaled to the line shown at the lower left of the 100 GHz map. Also evident in these images as the bright, diffuse structure towards the top (north) of the map is IRAS 081199-4153, a little–studied galactic infrared source.

Polarimetric measurements of SNR in the millimetre are non-existent; these QUaD data represent a first in this regard. Unfortunately, this SNR is so dim at 150 GHz that no polarimetric measurement is possible using these data alone; fortunately, this measurement is possible at 100 GHz. This is due to the synchrotron origin of the 100 GHz emission, which can yield polarization fractions greater than 70 % at these frequencies. The closest (in frequency) measurement available in the literature is Milne et al. (1993), who present a polarized radio measurement at 4.75 GHz and 8.4 GHz. These measurements show that the expected radio polarization in this remnant is not particularly strong, with a maximum of 10 % or so. Interestingly, the QUaD data support a much higher polarization fraction, approaching 50 % or more at various points in the remnant. However, this is a low signal to noise ratio source, and as such further measurements should be performed before hard conclusions are drawn.

194

6.4.4 THE GALACTIC CENTRE

The galactic centre region is important because it offers the most detailed view of the dynamics, physical conditions and processes at the centres of normal galaxies. That being said, the role of magnetic fields in this region are not understood; polarimetric observations can help build our understanding of their effects. Chuss et al. (2005) report on measurements of this region at 450 μ m using SPARO at the south pole, and Chuss et al. (2003) report on similar measurements at 350 μ m using the Hertz polarimeter. Similarly, Aitken et al. (2000) presents polarization measurements of the galactic centre regions at 750, 850, 1300 and 2000 μ m using SCUBA on the JCMT. These papers detail the polarization properties due to dust emission from Sgr A* and theories about the origin of the dust emission.

QUaD observed the galactic centre region on July 18, 2006, again using a simple raster scan observation strategy. The analysis methods discussed earlier in this section were employed to produce the maps shown in Figure 6.10. Although a detailed understanding of the physics behind the features in these maps has not been reached, we present preliminary results here. As the QUaD channels span the regime where synchrotron and free-free emission give way to dust emission, these maps provide insight into both the highest energy regions (which were mapped in the radio years ago), and the very coldest dusty regions, which are at the forefront of current research.

The QUaD maps highlight the complex, extended structures visible in the central part of the galaxy. They have lower angular resolution than the maps presented in the works listed above, but cover a much larger area. Both the 100 and 150 GHz QUaD bands are dominated by free-free and synchrotron emission, meaning that they are intrinsically highly polarized (while the dust emission is not). The features picked out in Figure 6.10 are thus due to the most ionized regions in the galactic centre. These QUaD maps are consistent with previous work showing that the emission in the galactic centre region is mainly due to synchrotron and free-free radiation and is polarized at the $\sim 10\%$ level. Highlighted by these data is the result of Bower et al. (1999), showing that the most central regions of this field are relatively unpolarized at 86 GHz. This result holds over a much broader area in the QUaD maps: why this should be is a mystery. The real strength of these QUaD maps lies in the fact that they cover a large area; work like that of Bower et al. (1999) can be put into the broader context of the dynamics of the entire galactic centre region using them.

6.4.5 CENTARUS A

Centaurus A is the closet powerful active galaxy, and contains a variety of well-known features which make it an interesting object for study across the entire electromagnetic





Figure 6.10: The Galactic Centre Region at 100 and 150 GHz These maps show the galactic centre region at 100 (left) and 150 GHz (right). Again, the colour table applies to temperature, while the red vectors show the amplitude and direction of the magnetic fields in this source. The elongated structure visible in both maps is aligned with the galactic plane; Sgr A^{*} (the radio source at the centre of the galaxy) is located at $\alpha = 17^{h}45^{m}12^{s}, \delta = -28^{\circ}48.3'$.

spectrum, and particularly at radio wavelengths where synchrotron radiation from the central source is dominant. At high resolution, the centre of Centarus A is resolved into two radio lobes joined by a compact source between them, which presumably corresponds to the black hole that powers the radio emission from the galaxy. Burns et al. (1983) provide a detailed report on 1.4 and 4.9 GHz polarimetry of this source using the VLA.

Centaurus A was observed with QUaD on 7 June, 2006 using the standard raster scan observation strategy. The data were processed as described in the beginning of this section; the maps shown in Figure 6.11 show the result.

At QUaD's resolution, very little of the detailed structure of this source can be discerned. However, the source is visibly extended along the same line as the radio jets discussed in Burns et al. (1983); QUaD is therefore resolving this feature, albeit with relatively poor resolution. This structure corresponds to the so-called 'inner radio lobe', which is roughly 18 kpc long (about the size of the host galaxy itself) assuming a distance to Centarus A of 5 Mpc. The polarization characteristics revealed by QUaD are consistent with those determined at lower frequencies; the northern lobe of the source is more polarized than the southern lobe, with polarization fractions approaching 20 %. These data are analogous to those for the moon, in that they are extremely useful for calibrating conventions

196



Figure 6.11: Radio Galaxy Centarus A at 100 and 150 GHz

These are maps of temperature and polarization in Centarus A at 100 GHz (left) and 150 GHz (right). The images show the source flux scaled according to the key at the bottom of the maps. Again, the red vectors show the polarization angle and amplitude according to the scaling shown in the upper left of the maps.

and algorithms in the high gain mode of the telescope. Based on our success with this source, the QUaD collaboration may embark on a survey of radio-loud galaxies in the local Universe during the austral summer of 2006/07 to study the polarization properties and temporal variation in a large set of such sources.



APPENDIX A: QUAD AT MAPO

This appendix gives a brief review and discussion of work performed to make QUaD observation-ready during commissioning, during the 2005 observation season, and in the austral summer 2005/2006.

A.1 COMMISSIONING

A great deal of time and man power was required to commission QUaD during the austral summer of 2004/2005. Hinderks (2005) gives a detailed account of this work which show-cases the huge amount of effort required to install and test the new receiver and telescope on the existing mount; here, we review the salient details.

Although the south pole station is extremely remote, a remarkable amount of infrastructure exists, and many experimental needs can be met with minimal difficulty. This is fortunate, as even the most urgent deliveries from the north require days to arrive at MAPO. The station has access to:

- a cryogenic facility which stores liquid helium and produces liquid nitrogen with an on-call technician,
- a fully equipped (if minimally redundant) machine shop,
- a large store of electronics components and equipment, and
- data connections to the outside world 8 hours per day.

The first stage in the commissioning process was to unpack and test the cryostat, and to ensure that the fridge would cycle under no load in case it or the cryostat was damaged during shipping. Following this, the focal plane was tested with the cryostat's window open, so that the detectors were under full optical loading. Various laboratory measurements were performed at this time, including band determination, polarization characterization, optical efficiency measurement, and noise characterization for each detector (some of the results of these tests are presented in Chapter 3). The cryostat was then lifted into position in a very laborious and moderately dangerous procedure using a winch, jack cart and lifting straps. The cryostat was aligned with the telescope according to positions calculated by the QUaD collaborators at Maynooth. Following installation, the receiver was tested for noise performance and optical loading in the absence of the foam cone, and then again with the cone and secondary mirror installed.

Simultaneously to the receiver testing and installation in the lab and on the mount, the telescope systems were being installed. The primary mirror and a test foam cone had been installed in the summer of 2003/2004; this foam cone and its testing apparatus were removed, leaving the primary mirror open to the sky. A crane was used to install the new foam cone, which was secured to the outer edge of the primary mirror using a fibre-glass rail and bolts. In addition, due to QUaD's height when compared to DASI, the ground shield required an extension of approximately 3 m to prevent the telescope from 'seeing' the ground while at low elevations. This extension was built by south pole station staff carpenters. At the same time as this work was being performed, the secondary mirror and its support structure were assembled and installed in the lab. Unfortunately, mechanical faults with the secondary mirror position adjustment system caused a drastic redesign at the pole; this involved removing the ability to adjust the mirror's position using remote electronic commands. Instead, a threaded-rod and bolt assembly was constructed which, while allowing adjustment of the secondary mirror's position on the telescope, made such a change very difficult. This system was installed into the secondary mirror support assembly, which was then installed on the top of the foam cone. These alterations were completed, and the mirror's position was set to be in the expected optimal focus position for winter conditions $(T_{\rm ext} \sim -70 \,{\rm C}).$

QUaD's astronomical first light was achieved on RCW 38 on February 8th, 2005. Due to the limited length of the polar summer, the QUaD team was scheduled to leave only a few days after first light; the remaining instrument testing and characterization was performed remotely after the south pole station closed for winter 2005 (these tests are presented in Chapter 4).

A.2 REFOCUS

As early as April 2005 it became apparent that the telescope was out of focus by a considerable amount. As discussed in Section 4.4, the QUaD collaborators at Maynooth calculated the distance between the primary and secondary mirrors to be 1.7 mm too small. Upon compiling compelling evidence that this was indeed the cause of the error, the decision was made to ask the winter-overing technicians to move the secondary mirror.

As discussed above, due to the *ad hoc* design of the secondary support structure, this proved to be a very difficult change. Using head lamps in -70 C weather, the winter

over technicians undid the nuts clamping the secondary mirror in place on July 10th, 2005 (the peak of the austral winter). Using a micrometer, they moved the mirror the required amount (as imaged in Figure A.1), and then re-clamped the mirror in place. An RCW 38 mapping run was performed right after the change to ensure that the refocus had had the desired effect. As shown in Figure 4.12, the refocus was remarkably successful.



Figure A.1: Refocusing the Telescope, July 10th, 2005

This is a photograph of Allan Day checking the distance the secondary mirror moved during the refocus. It is clearly miserably cold and dark at the south pole during this time of year. (Photograph courtesy R. Schwarz).

A.3 WORK DURING THE 2005/2006 SUMMER SEASON

Beginning in November 2005, the instrument was modified to improve its performance over the 2005 observation season. The first change to the system was to bring the receiver down from the telescope cab and to warm it to room temperature. The focal plane block was removed, and the detector configuration was changed. Detectors known to be extremely noisy or dead were removed and replaced with fresh stock from JPL/CalTech. Similarly, detectors with large thermal ramps were moved to dark positions, and channels experiencing cross talk problems were moved to new cables. This gave the array a full complement of detectors; in contrast, during the 2005 winter only about 80% of the detectors were considered 'good'. At this point, there was ample time to cool down the cryostat in a dark configuration to perform dark characterization of the bolometers. The cryostat was then warmed, the detectors re-opened to light, and then cooled again.

At this point, the refrigerator experienced hold time problems. This was thought to be due to poor condensation efficiency, which occurs when the condenser plate is not in sufficiently good thermal contact with the 4 K plate. The cryostat was warmed, and the bolts clamping the cooler condensation plate to the work surface of the LHe tank were retightened with a torque wrench. The wrench showed that the screws were tightened to an average of about 7 foot pounds with a minimum of about 4 and maximum of 9 foot pounds *before* the re-tightening; the screws were all tightened to a uniform 20 foot pounds and the cryostat was closed again. Upon cooling the cryostat and testing the refrigerator, the hold time was found to have improved to more than 30 hours, which is ample for QUaD's needs. At this point, the receiver was reinstalled into the telescope cab and the telescope was ready for testing.

At the same time as these modifications were occurring, changes to the secondary mirror and its support assembly were being performed. QUaD had been experiencing large beam ellipticities due to the warp in the primary mirror (see Chapter 2). To counter this, a new secondary mirror was constructed. In the interest of keeping the weight of the mirror low, a carbon-fibre backed aluminum mirror design was used. Based on ray tracing in our optical model for the instrument, this mirror has a three dimensional shape designed to cancel the effect of the warp from the primary (as shown in Figure A.2). This new secondary mirror was installed on the secondary mirror backing assembly; extra care was taken to ensure it aligned with the primary mirror in the correct configuration to cancel rather than double the beam ellipticities.

The telescope backing assembly was itself modified to allow simpler changes to the focus of the telescope during the 2006 season. Instead of the threaded-rod and bolt set up, a system involving linear actuators and bearings was installed. This system allows remote changes to the z-axis focus of the telescope via motors controlled by commands sent from the main telescope control software. This system should negate the need for difficult and dangerous manual refocusing of the telescope (although, in the event of a failure of the automatic system, the focus can be changed manually). This system was installed on the telescope, and tests were performed to verify its performance (as shown in Figure A.2).

Focus tests are performed in the following manner. The mirror is commanded to go to a set of offsets (in mm) from its nominal zero point, and a pointing-cross like observation is performed. The offsets are generally set to be 1 mm apart, and there are 11 steps in all, yielding 22 data points for each of the central bolometers per test (the other bolometers



Figure A.2: Secondary Mirror Modifications for 2006 Season

These plots highlight the secondary mirror and its support system installed in the austral summer of 2006. The left hand panel shows the complex three-dimensional shape of the new secondary mirror. The fiducial secondary mirror shape has been subtracted from actual measurements of the as-built mirror using a Romer arm, and this plot shows the residuals. Clearly, the new mirror exhibits a strong deviation from the nominal hyperbolic shape used with a properly shaped primary mirror; this strong high-order structure is required to counteract the warp in the actual QUaD primary mirror. The plot on the left shows the result of the new secondary mirror support system using cross observations to measure the beam width. The diamonds show the widths for constant-declination scans, and the triangles the widths for constant-right ascension scans. As a quadratic model is expected to describe the locus of points; the red and blue curves (for azimuth scans and elevation scans, respectively) show the best fit parabolas to these data. The vertical dashed lines show the optimal focus position at the time of this measurement (February 1, 2006); one can see that the secondary mirror was set approximately 3.3 mm too far from the primary mirror for an optimal focus at these ambient temperatures.

are not used). The resulting RCW 38 time series blips are fit as in Section 4.1.2, and the resulting FWHM of the blips for the central bolometers are archived. It can be shown that as the primary-secondary distance in a Cassegrain telescope is varied about the best focus position, the locus traced by the beam's FWHM is a paraboloid. A quadratic model is fit to the blip widths as a function of primary-secondary distance; the minima of these curves give the best focus of the telescope. The mirror was initially set to its nominal position for the winter season, which causes the telescope to be out of focus in the warmer temperature of the summer. However, as the ambient temperature drops, the foam cone shrinks, and the

telescope comes into better focus. Because these focus tests can be performed at virtually any time, the focus slew measurement is used throughout the season (and particularly at the beginning of CMB observations) to determine and reset the best focus position for the telescope. Small raster maps centered on the central horn can also give quick-look capability for the full beam shape. These methods have been utilized throughout the winter season to keep the telescope at an optimal focus.

When it was verified that the receiver and secondary mirror system were working correctly, a campaign to measure the telescope's sidelobes began. Section 4.5 gives a discussion of the causes of these sidelobes and their effects on the QUaD data; here we present the tests used to determine their characteristics. These involved rastering the telescope over a source mounted to the edge of the ground shield (located near an azimuth of -50° and an elevation of $+21^{\circ}$). The source consisted of a tunable Gunn diode (actually, one working at the third harmonic of 50 GHz for the 150 GHz tests, and another working at the second harmonic of 57 GHz for the 100 GHz tests). Because the bandwidth of a Gunn diode is extremely small, the Gunn was biased in such a way as to quickly slew through a few hundred MHz about the nominal zero frequency in a cyclic manner. This gave an effective bandwidth of approximately 1% for the source system (although this is essentially monochromatic to bolometric systems, it is enough to help reduce standing mode features). The Gunn source was fitted with a feed horn projecting into a beam of about 1 square degree. This whole assembly was mounted on the rim of the ground shield using a pana-vice aimed at the centre of the telescope. The telescope was rastered over the source with an azimuth throw of $\pm 55^{\circ}$ from the nominal source position in the elevation range $26^{\circ} < el < 65^{\circ}$. The Gunn source is chopped with a square wave at 3 Hz, the reference signal for which is fed into the QUaD DAC to enable demodulation. Once the data are demodulated, they can be binned into maps to study the angular structure of the sidelobes. The raster map measurement was intended to draw out structure at angular scales a few tens of degrees from the beam. Examples of the results of this type of test can be seen in Figure A.3.

A second type of raster map was performed by setting the telescope to zenith, and rotating 360° in θ . This draws out structure far from the main beam. An example of this type of test is shown in Figure A.4 where it is compared with the raster map experiment discussed above.

As discussed briefly in Section 4.5, these tests incited the polar team to attempt a series of modifications to the snout baffling to see if the sidelobes could be reduced. The starting point for this process were maps like those shown in Figure A.3. After a great deal of thinking, it was hypothesized that the radial feature in these plots was due to the reflecting baffle. The first modification was therefore to replace the reflecting baffle with an



Figure A.3: Sidelobe Maps, 2005 Season Baffle

This plot shows sidelobe raster maps for bolometers 150-1-A, 150-3-A, 150-10-A and 100-4-A (these were chosen as they lie on the same radial line from the centre of the array). In these plots, the colour scale is the same in all, but uncalibrated; 0 dB corresponds to a demodulated signal strength of 1 Volt. However, to give a general idea of the calibration in these images, a back of the envelope calculation suggests these structures have amplitudes about -60 dB per beam size. The bright spot near -50° , $+26^\circ$ is the main lobe response for the system, and is not of interest in these plots. However, the annulus visible in all of these plots is more relevant; this is a sidelobe due to the reflecting baffle. It is visible in all of the detectors because they are all sensitive to rays which fall on this baffle. The radial spike in the map for 150-10-A is also of interest; this appears in all of the outermost 150 GHz bolometers, and is thought to be due to diffraction off optical elements in the 77 K stage. The apparent oscillating signal along this spike is due to the monochromatic nature of the source. Furthermore, the cause of the diffuse structure in the 100 GHz sidelobe maps is unknown.

eccosorb-coated one of the same radius. When this was performed, data like those shown in Figure A.4 resulted.

These maps show a distinct improvement over those shown in Figure A.3. As expected, the annular ring is gone. Furthermore, the diffuse structure in the 100 GHz raster maps is absent. However, the radial spikes in the outermost ring of 150 GHz bolometers were not affected by this change. This spike is thought to be due to diffraction from the 77 K filters in the cryostat; at the point in the season at which these tests were performed, any modification of the optics would have been ill advised. Therefore, the decision was reached to but a eccosorb-coated baffle extension on the snout. This extended baffle was approximately 30 cm in diameter and 60 cm high; it was designed to block the most off–axis diffracted rays (this device was euphemistically called the 'top hat baffle' due to its resemblance to a formal article of clothing). The height was carefully estimated from drawings of the telescope; if too low, it would not reduce the length of the radial spikes sufficiently to improve the optical performance of the system, and if too high, it would block rays in the main lobe of the beam traveling from the primary to the secondary mirrors. This baffle



Figure A.4: Sidelobe Maps, Blackened Flat Baffle

These plots show the raster maps for the same bolometers featured in Figure A.3 (upper panels) and the θ rotation maps for the same devices (lower panels) for data taken with a flat, eccosorb coated baffle. The colour scale is the same in all of these images, and in Figure A.3. As discussed in the text, the purpose of this baffle was to absorb the rays causing the annular sidelobe; it succeeded in this task.

was installed, and the resulting sidelobe maps are shown in Figure $A.5^1$.

Figure A.5 shows that the top hat baffle was very successful at reducing the sidelobe contamination of the telescope. Unfortunately, because this baffle was coated in ~ 300 K eccosorb, and the solid angle in the radial sidelobes is not small, this baffle also increased the loading on the detectors. Load curves taken at each stage in this process showed that the loading increase due to the full top hat baffle was large enough to reduce the overall sensitivity of the instrument by more than a factor of 2. This sensitivity reduction would increase the error bars on our final measurement of the CMB power spectra by the same factor. Because methods exist to reduce the effect of the sidelobes directly in the time series, it was thought that the advantage afforded by the baffling was not large enough to justify such a penalty in sensitivity. Therefore, only a minor modification to the snout baffling is present in the final 2006 instrument; the curved baffle from 2005 has been removed,

¹One may notice that the 100 GHz detector is missing from this plot - this is due to the fact that a major electrical accident destroyed the telescope control computer near the end of the season. It took approximately a week to bring the telescope back on line, but at the point these data were taken only one data acquisition card was available. This card was installed in the 150 GHz position, meaning that the 100 GHz data during this epoch were not recorded; a replacement card was rushed to the south pole and was installed on the telescope just before the station closed.


Figure A.5: Sidelobe Maps, Full Top Hat Baffle

These sidelobe maps are for the full 'top hat' baffle epoch (*i.e.* flat ring with a tall cylinder attached to its upper side, completely coated in eccosorb) arranged in the same manner as in Figure A.4 for devices 150-1-A, 150-3-A, and 150-10-A¹. The colour stretch is again the same as defined in Figure A.3. As intended, this baffle absorbs both the rays which cause the annular ring, but also a great proportion of the diffracted rays causing the radial spike in the outer 150 GHz bolometers.

and a flat reflecting baffle installed in its place. This has the effect of drawing the annular sidelobes in radially by a factor of 2; as long as the telescope observes above an elevation of 40° , only the spike sidelobes in the outermost 150 GHz bolometers will cause a large signal in the time series. This solution was thought to afford the best compromise between sidelobe contamination and sensitivity for the 2006 CMB season.

At the completion of these tests, the polarized source discussed in Section 3.3.1 became available. This source, which is shown in Figure A.6, consists of a box with an aperture and no back. The source, which is a chopper wheel with alternating clear and eccosorb-coated quadrants, rotates at a few Hz to produce a polarized signal. The telescope was rastered over this source in a way very similar to that discussed above for the Gunn source, albeit with a much smaller scan throw of only $\sim \pm 6^{\circ}$. This experiment produced the data discussed in Section 3.3.1. The experiment was repeated with both the full top hat

baffle and the 2006 flat black baffle in place; the results did not change, meaning that the baffling does not affect the polarization properties of the instrument.



Figure A.6: Polarization Measurement Apparatus

A photograph of the thermal polarized source on a tower in place at MAPO. This photograph was taken March 15th, 2006, four weeks after the station was closed but before the sun set for the winter. (Photograph courtesy R. Schwarz).

The aperture of the box was changed to three variations: an aluminum sheet with a d = 7.6 cm circular aperture cut into it, the same with a d = 3.8 cm circular aperture cut into it, and an eccosorb sheet with a d = 7.6 cm circular aperture in it. The measured polarization properties of the instrument changed as a function of the aperture; this means that some level of depolarization of the signal is occurring as a result of diffraction off the edges of the aperture. This issue increases the size of the systematic error assigned to the polarization parameters of the telescope.

BIBLIOGRAPHY

- Aitken, D. K., Greaves, J., Chrysostomou, A., Jenness, T., Holland, W., Hough, J. H., Pierce-Price, D., & Richer, J. 2000, Astrophys. J. Letters, 534, L173
- Alpher, R. A. & Herman, R. C. 1950, Reviews of Modern Physics, 22, 153
- Archibald, E. N., Jenness, T., Holland, W. S., Coulson, I. M., Jessop, N. E., Stevens, J. A., Robson, E. I., Tilanus, R. P. J., Duncan, W. D., & Lightfoot, J. F. 2002, Mon. Not. Royal Astro. Society, 336, 1
- Arendt, R. G., Dwek, E., Petre, R., Dickel, J. R., Roger, R. S., Milne, D. K., & Kesteven, M. J. 1990, Astrophys. J., 350, 266
- Barkats, D., Bischoff, C., Farese, P., Fitzpatrick, L., Gaier, T., Gundersen, J. O., Hedman, M. M., Hyatt, L., McMahon, J. J., Samtleben, D., Staggs, S. T., Vanderlinde, K., & Winstein, B. 2005, Astrophys. J. Letters, 619, L127
- Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S.,
 Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., Wright, E. L., Barnes, C., Greason,
 M. R., Hill, R. S., Komatsu, E., Nolta, M. R., Odegard, N., Peiris, H. V., Verde, L., &
 Weiland, J. L. 2003, Astrophys. J. Suppl., 148, 1
- Benoît, A., Ade, P., Amblard, A., Ansari, R., Aubourg, É., Bargot, S., Bartlett, J. G., Bernard, J.-P., Bhatia, R. S., Blanchard, A., Bock, J. J., Boscaleri, A., Bouchet, F. R., Bourrachot, A., Camus, P., Couchot, F., de Bernardis, P., Delabrouille, J., Désert, F.-X., Doré, O., Douspis, M., Dumoulin, L., Dupac, X., Filliatre, P., Fosalba, P., Ganga, K., Gannaway, F., Gautier, B., Giard, M., Giraud-Héraud, Y., Gispert, R., Guglielmi, L., Hamilton, J.-C., Hanany, S., Henrot-Versillé, S., Kaplan, J., Lagache, G., Lamarre, J.-M., Lange, A. E., Macías-Pérez, J. F., Madet, K., Maffei, B., Magneville, C., Marrone, D. P., Masi, S., Mayet, F., Murphy, A., Naraghi, F., Nati, F., Patanchon, G., Perrin, G., Piat, M., Ponthieu, N., Prunet, S., Puget, J.-L., Renault, C., Rosset, C., Santos, D., Starobinsky, A., Strukov, I., Sudiwala, R. V., Teyssier, R., Tristram, M., Tucker, C., Vanel, J.-C., Vibert, D., Wakui, E., & Yvon, D. 2004, Astron. & Astrophys., 424, 571

- Bhatia, R. S., Chase, S. T., Edgington, S. F., Glenn, J., Jones, W. C., Lange, A. E., Maffei,B., Mainzer, A. K., Mauskopf, P. D., & Rownd, B. K. 2000, Cryogenics, 40, 685
- Birkinshaw, M. 1999, Phys. Reports, 310, 97
- Bock, J. J. 1994, Ph.D. Thesis
- Bock, J. J., Chen, D., Mauskopf, P. D., & Lange, A. E. 1995, Space Science Reviews, 74, 229
- Bond, J. R. & Efstathiou, G. 1984, Astrophys. J. Letters, 285, L45
- Bowden, M., Taylor, A. N., Ganga, K. M., Ade, P. A. R., Bock, J. J., Cahill, G., Carlstrom, J. E., Church, S. E., Gear, W. K., Hinderks, J. R., Hu, W., Keating, B. G., Kovac, J., Lange, A. E., Leitch, E. M., Maffei, B., Mallie, O. E., Melhuish, S. J., Murphy, J. A., Pisano, G., Piccirillo, L., Pryke, C., Rusholme, B. A., O'Sullivan, C., & Thompson, K. 2004, Mon. Not. Royal Astro. Society, 349, 321
- Bower, G. C., Wright, M. C. H., Backer, D. C., & Falcke, H. 1999, Astrophys. J., 527, 851
- Brown, M. L., Castro, P. G., & Taylor, A. N. 2005, Mon. Not. Royal Astro. Society, 360, 1262
- Burns, J. O., Feigelson, E. D., & Schreier, E. J. 1983, Astrophys. J., 273, 128
- Bussmann, R. S., Holzapfel, W. L., & Kuo, C. L. 2005, Astrophys. J., 622, 1343
- Calisse, P. G. 2002, ArXiv Astrophysics e-prints, astro-ph/0210456
- Carlstrom, J. E., Holder, G. P., & Reese, E. D. 2002, Ann. Rev. Astron. Astrophys., 40, 643
- Cartwright, J. K., Pearson, T. J., Readhead, A. C. S., Shepherd, M. C., Sievers, J. L., & Taylor, G. B. 2005, Astrophys. J., 623, 11
- Churchwell, E. 2002, Ann. Rev. Astron. Astrophys., 40, 27
- Chuss, D. T., Davidson, J. A., Dotson, J. L., Dowell, C. D., Hildebrand, R. H., Novak, G., & Vaillancourt, J. E. 2003, Astrophys. J., 599, 1116
- Chuss, D. T., Dowell, C. D., Hildebrand, R. H., & Novak, G. 2005, in Astronomical Society of the Pacific Conference Series, ed. A. Adamson, C. Aspin, & C. Davis, 311-+

- Coble, K., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Contaldi, C. R., Crill, B. P., de Bernardis, P., Farese, P., Ganga, K., Giacometti, M., Hivon, E., Hristov, V. V., Iacoangeli, A., Jaffe, A. H., Jones, W. C., Lange, A. E., Martinis, L., Masi, S., Mason, P., Mauskopf, P. D., Melchiorri, A., Montroy, T., Netterfield, C. B., Nyman, L., Pascale, E., Piacentini, F., Pogosyan, D., Polenta, G., Pongetti, F., Prunet, S., Romeo, G., Ruhl, J. E., & Scaramuzzi, F. 2003, ArXiv Astrophysics e-prints, astro-ph/0301599
- Cox, P., Mezger, P. G., Sievers, A., Najarro, F., Bronfman, L., Kreysa, E., & Haslam, G. 1995, Astron. & Astrophys., 297, 168
- Dicke, R. H., Beringer, R., Kyhl, R. L., & Vane, A. B. 1946, Physical Review, 70, 340
- Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, Astrophys. J., 142, 414
- Duncan, R. A. & White, S. M. 2003, Mon. Not. Royal Astro. Society, 338, 425
- Farese, P. C., Dall'Oglio, G., Gundersen, J. O., Keating, B. G., Klawikowski, S., Knox, L., Levy, A., Lubin, P. M., O'Dell, C. W., Peel, A., Piccirillo, L., Ruhl, J., & Timbie, P. T. 2004, Astrophys. J., 610, 625
- Fixsen, D. J., Cheng, E. S., Gales, J. M., Mather, J. C., Shafer, R. A., & Wright, E. L. 1996, Astrophys. J., 473, 576
- Fixsen, D. J., Moseley, S. H., & Arendt, R. G. 2000, Astrophys. J. Suppl., 128, 651
- Goldstein, J. H. 2004, Ph.D. Thesis
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, Astrophys. J., 622, 759
- Grainge, K., Carreira, P., Cleary, K., Davies, R. D., Davis, R. J., Dickinson, C., Genova-Santos, R., Gutiérrez, C. M., Hafez, Y. A., Hobson, M. P., Jones, M. E., Kneissl, R., Lancaster, K., Lasenby, A., Leahy, J. P., Maisinger, K., Pooley, G. G., Rebolo, R., Rubiño-Martin, J. A., Sosa Molina, P. J., Ödman, C., Rusholme, B., Saunders, R. D. E., Savage, R., Scott, P. F., Slosar, A., Taylor, A. C., Titterington, D., Waldram, E., Watson, R. A., & Wilkinson, A. 2003, Mon. Not. Royal Astro. Society, 341, L23
- Greaves, J. S., Holland, W. S., Jenness, T., Chrysostomou, A., Berry, D. S., Murray, A. G., Tamura, M., Robson, E. I., Ade, P. A. R., Nartallo, R., Stevens, J. A., Momose, M., Morino, J.-I., Moriarty-Schieven, G., Gannaway, F., & Haynes, C. V. 2003, Mon. Not. Royal Astro. Society, 340, 353
- Guth, A. H. 2004, in Measuring and Modeling the Universe, ed. W. L. Freedman, 31-+

- Hanany, S., Ade, P., Balbi, A., Bock, J., Borrill, J., Boscaleri, A., de Bernardis, P., Ferreira,
 P. G., Hristov, V. V., Jaffe, A. H., Lange, A. E., Lee, A. T., Mauskopf, P. D., Netterfield,
 C. B., Oh, S., Pascale, E., Rabii, B., Richards, P. L., Smoot, G. F., Stompor, R., Winant,
 C. D., & Wu, J. H. P. 2000, Astrophys. J. Letters, 545, L5
- Hanany, S. & Rosenkranz, P. 2003, New Astronomy Review, 47, 1159
- Harrison, E. R. 1970, Phys. Rev. D, 1, 2726
- Hecht, E. 1998, Optics 3nd edition (Optics 3nd edition by Eugene Hecht Reading, MA: Addison-Wesley Publishing Company, 1998)
- Hedman, M. M., Barkats, D., Gundersen, J. O., Staggs, S. T., & Winstein, B. 2001, Astrophys. J. Letters, 548, L111
- Hinderks, J. R. 2005, Ph.D. Thesis
- Hinshaw, G., Barnes, C., Bennett, C. L., Greason, M. R., Halpern, M., Hill, R. S., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Odegard, N., Page, L., Spergel, D. N., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003, Astrophys. J. Suppl., 148, 63
- Hinshaw, G. e. a. 2006, Submitted to ApJ, astro-ph/0603451
- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, Astrophys. J., 567, 2
- Hu, W. & Sugiyama, N. 1995, Astrophys. J., 444, 489
- Hu, W. & White, M. 1997, New Astronomy, 2, 323
- Itoh, N. & Nozawa, S. 2004, Astron. & Astrophys., 417, 827
- Jackson, J. D. 1998, Classical Electrodynamics, 3rd Edition (Classical Electrodynamics, 3rd Edition, by John David Jackson, pp. 832. ISBN 0-471-30932-X. Wiley-VCH, July 1998.)
- Johnson, B. R., Abroe, M. E., Ade, P., Bock, J., Borrill, J., Collins, J. S., Ferreira, P., Hanany, S., Jaffe, A. H., Jones, T., Lee, A. T., Levinson, L., Matsumura, T., Rabii, B., Renbarger, T., Richards, P. L., Smoot, G. F., Stompor, R., Tran, H. T., & Winant, C. D. 2003, New Astronomy Review, 47, 1067
- Johnson, B. R. et al. 2006, in "Half-Wave Plate Polarimetry with MAIXPOL", to appear in the Proceedings of the 2006 Rencontres de Moriond, "Contents and Structures of the Universe"

- Jones, R. C. 1953, Journal of the Optical Society of America (1917-1983), 43, 1
- Jones, W. C. 2005, Ph.D. Thesis
- Jones, W. C., Bhatia, R., Bock, J. J., & Lange, A. E. 2003, in Millimeter and Submillimeter Detectors for Astronomy. Edited by Phillips, Thomas G.; Zmuidzinas, Jonas. Proceedings of the SPIE, Volume 4855, pp. 227-238 (2003).
- Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. D, 55, 7368
- Kane, B. D., Clemens, D. P., Barvainis, R., & Leach, R. W. 1993, Astrophys. J., 411, 708
- Keating, B. G., Ade, P. A. R., Bock, J. J., Hivon, E., Holzapfel, W. L., Lange, A. E., Nguyen, H., & Yoon, K. W. 2003, in Polarimetry in Astronomy. Edited by Silvano Fineschi. Proceedings of the SPIE, Volume 4843, pp. 284-295, ed. S. Fineschi
- Keating, B. G., O'Dell, C. W., de Oliveira-Costa, A., Klawikowski, S., Stebor, N., Piccirillo, L., Tegmark, M., & Timbie, P. T. 2001, Astrophys. J. Letters, 560, L1
- Kinney, W. H. 1998, Phys. Rev. D, 58, 123506
- Kogut, A., Spergel, D. N., Barnes, C., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Wollack, E., & Wright, E. L. 2003, Astrophys. J. Suppl., 148, 161
- Kosowsky, A. 2002, in Modern Cosmology, ed. S. Bonometto, V. Gorini, & U. Moschella, 219–263
- Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel,W. L. 2002, Nature, 420, 772
- Kuo, C. L., Ade, P. A. R., Bock, J. J., Cantalupo, C., Daub, M. D., Goldstein, J., Holzapfel,
 W. L., Lange, A. E., Lueker, M., Newcomb, M., Peterson, J. B., Ruhl, J., Runyan, M. C.,
 & Torbet, E. 2004, Astrophys. J., 600, 32
- Lamarre, J. M. 1986, Appl. Opt., 25, 870
- Lange, A. E., Kreysa, E., McBride, S. E., Richards, P. L., & Haller, E. E. 1983, Int. J. Infrared and Millimeter Waves, 4, 689+
- Langley, S. P. 1881, Nature, 25, 14+
- Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, Astrophys. J., 624, 10

- Lineweaver, C. H., Smoot, G. F., Bennett, C. L., Wright, E. L., Tenorio, L., Kogut, A., Keegstra, P. B., Hinshaw, G., & Banday, A. J. 1994, Astrophys. J., 436, 452
- Longair, M. S. 1994, High energy astrophysics. Vol.2: Stars, the galaxy and the interstellar medium (Cambridge: Cambridge University Press, —c1994, 2nd ed.)
- Martin, D. H. 1982, in Infrared and Millimeter Waves, Vol. 6, Systems and Components, Academic Press, New York, ed. K. J. Button, pp. 65–148
- Masi, S., Ade, P., Bock, J., Bond, J., Borrill, J., Boscaleri, A., Cabella, P., Contaldi, C., Crill, B., de Bernardis, P., De Gasperis, G., de Oliveira-Costa, A., De Troia, G., Di Stefano, G., Ehlers, P., Hivon, E., Hristov, V., Iacoangeli, A., Jaffe, A., Jones, W., Kisner, T., Lange, A., MacTavish, C., Marini-Bettolo, C., Mason, P., Mauskopf, P., Montroy, T., Nati, F., Nati, L., Natoli, P., Netterfield, C., Pascale, E., Piacentini, F., Pogosyan, D., Polenta, G., Prunet, S., Ricciardi, S., Romeo, G., Ruhl, J., Santini, P., Tegmark, M., Torbet, E., Veneziani, M., & Vittorio, N. 2005, ArXiv Astrophysics e-prints, astroph/0507509
- Mather, J. C. 1982, Appl. Opt., 21, 1125
- Mather, J. C., Fixsen, D. J., Shafer, R. A., Mosier, C., & Wilkinson, D. T. 1999, Astrophys. J., 512, 511
- Mauskopf, P. D., Ade, P. A. R., de Bernardis, P., Bock, J. J., Borrill, J., Boscaleri, A., Crill, B. P., DeGasperis, G., De Troia, G., Farese, P., Ferreira, P. G., Ganga, K., Giacometti, M., Hanany, S., Hristov, V. V., Iacoangeli, A., Jaffe, A. H., Lange, A. E., Lee, A. T., Masi, S., Melchiorri, A., Melchiorri, F., Miglio, L., Montroy, T., Netterfield, C. B., Pascale, E., Piacentini, F., Richards, P. L., Romeo, G., Ruhl, J. E., Scannapieco, E., Scaramuzzi, F., Stompor, R., & Vittorio, N. 2000, Astrophys. J. Letters, 536, L59

McGlynn, T., Scollick, K., & White, N. 1997, in IAU Symposium, 465-466

- Miller, A. D., Caldwell, R., Devlin, M. J., Dorwart, W. B., Herbig, T., Nolta, M. R., Page, L. A., Puchalla, J., Torbet, E., & Tran, H. T. 1999, Bulletin of the American Astronomical Society, 31, 1458
- Milne, D. K., Stewart, R. T., & Haynes, R. F. 1993, Mon. Not. Royal Astro. Society, 261, 366

- Montroy, T. E., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Cabella, P., Contaldi, C. R., Crill, B. P., de Bernardis, P., De Gasperis, G., de Oliveira-Costa, A., De Troia, G., di Stefano, G., Hivon, E., Jaffe, A. H., Kisner, T. S., Jones, W. C., Lange, A. E., Masi, S., Mauskopf, P. D., MacTavish, C. J., Melchiorri, A., Natoli, P., Netterfield, C. B., Pascale, E., Piacentini, F., Pogosyan, D., Polenta, G., Prunet, S., Ricciardi, S., Romeo, G., Ruhl, J. E., Santini, P., Tegmark, M., Veneziani, M., & Vittorio, N. 2005, ArXiv Astrophysics e-prints, astro-ph/0507514
- Murphy, J. A., Gleeson, E., Cahill, G., Lanigan, W., OSullivan, C., Cartwright, E., Church, S. E., Hinderks, J., Kirby, E., Thompson, K., Rusholme, B., Gear, W. K., Maffei, B., Ade, P. A. R., Tucker, C., & Jones, B. 2005, Int. J. Infrared and Millimeter Waves, 26, 505
- Ostriker, J. P. & Vishniac, E. T. 1986, Astrophys. J. Letters, 306, L51
- O'Sullivan, C. et al.. 2006, In. Prep.
- Peacock, J. A. 1999, Cosmological Physics (Cosmological Physics, by John A. Peacock, pp. 704. ISBN 052141072X. Cambridge, UK: Cambridge University Press, January 1999.)
- Peebles, P. J. E. 1993, Principles of physical cosmology (Princeton Series in Physics, Princeton, NJ: Princeton University Press, —c1993)
- Peebles, P. J. E. & Yu, J. T. 1970, Astrophys. J., 162, 815
- Penzias, A. A. & Wilson, R. W. 1965, Astrophys. J., 142, 419
- Peterson, J. B., Radford, S. J. E., Ade, P. A. R., Chamberlin, R. A., O'Kelly, M. J., Peterson, K. M., & Schartman, E. 2003, Pub. Astron. Soc. Pacific, 115, 383
- Pobell, F. 1996, Matter and Methods at Low Temperatures 2nd edition (Berlin: Springer-Verlag, 1996)
- Poppi, S., Carretti, E., Cortiglioni, S., Krotikov, V. D., & Vinyajkin, E. N. 2002, in AIP Conf. Proc. 609: Astrophysical Polarized Backgrounds, ed. S. Cecchini, S. Cortiglioni, R. Sault, & C. Sbarra, 187–192
- Puchalla, J. L., Caldwell, R., Cruz, K. L., Devlin, M. J., Dorwart, W. B., Herbig, T., Miller, A. D., Nolta, M. R., Page, L. A., Torbet, E., & Tran, H. T. 2002, Astron. J., 123, 1978

- Readhead, A. C. S., Myers, S. T., Pearson, T. J., Sievers, J. L., Mason, B. S., Contaldi, C. R., Bond, J. R., Bustos, R., Altamirano, P., Achermann, C., Bronfman, L., Carlstrom, J. E., Cartwright, J. K., Casassus, S., Dickinson, C., Holzapfel, W. L., Kovac, J. M., Leitch, E. M., May, J., Padin, S., Pogosyan, D., Pospieszalski, M., Pryke, C., Reeves, R., Shepherd, M. C., & Torres, S. 2004, Science, 306, 836
- Rees, M. J. 1968, Astrophys. J. Letters, 153, L1+
- Rees, M. J. & Sciama, D. W. 1968, Nature, 217, 511
- Rephaeli, Y. & Yankovitch, D. 1997, Astrophys. J. Letters, 481, L55+
- Richards, P. L. 1994, Journal of Applied Physics, 76, 1
- Richards, P. L. & McCreight, C. R. 2005, Physics Today, 58, 41
- Ruhl, J. E., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Contaldi,
 C. R., Crill, B. P., de Bernardis, P., De Troia, G., Ganga, K., Giacometti, M., Hivon,
 E., Hristov, V. V., Iacoangeli, A., Jaffe, A. H., Jones, W. C., Lange, A. E., Masi, S.,
 Mason, P., Mauskopf, P. D., Melchiorri, A., Montroy, T., Netterfield, C. B., Pascale, E.,
 Piacentini, F., Pogosyan, D., Polenta, G., Prunet, S., & Romeo, G. 2003, Astrophys. J.,
 599, 786
- Saken, J. M., Fesen, R. A., & Shull, J. M. 1992, Astrophys. J. Suppl., 81, 715
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, Astrophys. J., 500, 525
- Seljak, U. & Zaldarriaga, M. 1996, Astrophys. J., 469, 437
- Smoot, G. F., Bennett, C. L., Kogut, A., Wright, E. L., Aymon, J., Boggess, N. W., Cheng,
 E. S., de Amici, G., Gulkis, S., Hauser, M. G., Hinshaw, G., Jackson, P. D., Janssen, M.,
 Kaita, E., Kelsall, T., Keegstra, P., Lineweaver, C., Loewenstein, K., Lubin, P., Mather,
 J., Meyer, S. S., Moseley, S. H., Murdock, T., Rokke, L., Silverberg, R. F., Tenorio, L.,
 Weiss, R., & Wilkinson, D. T. 1992, Astrophys. J. Letters, 396, L1
- Spergel, D. N. & Zaldarriaga, M. 1997, Physical Review Letters, 79, 2180
- Stompor, R. & Efstathiou, G. 1999, Mon. Not. Royal Astro. Society, 302, 735
- Sudiwala, R. V., Griffin, M. J., & Woodcraft, A. L. 2002, Int. J. Infrared and Millimeter Waves, 23, 545+
- Teleberg, G. 2006, Ph.D. Thesis

- Verde, L., Peiris, H. V., Spergel, D. N., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Wollack, E., & Wright, E. L. 2003, Astrophys. J. Suppl., 148, 195
- Vishniac, E. T. 1987, Astrophys. J., 322, 597
- Ward-Thompson, D., Ade, P. A. R., Araujo, H., Coulson, I., Cox, J., Davis, G. R., Evans, R., Griffin, M. J., Gear, W. K., Hargrave, P., Hargreaves, P., Hayton, D., Kiernan, B. J., Leeks, S. J., Mauskopf, P., Naylor, D., Potter, N., Rinehart, S. A., Sudiwala, R., Tucker, C. R., Walker, R. J., & Watkin, S. L. 2005, Mon. Not. Royal Astro. Society, 364, 843
- Weinberg, S. 1972, Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity (Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity, by Steven Weinberg, pp. 688. ISBN 0-471-92567-5. Wiley-VCH, July 1972.)
- White, G. K. & Meeson, P. J. 2002, Experimental Techniques in Low–Temperature Physics 4th edition (Oxford, UK: Oxford University Press, 2002)
- Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, Astrophys. J. Letters, 458, L53+
- Yun, M., Beeman, J. W., Bhatia, R., Bock, J. J., Holmes, W., Hustead, L., Koch, T., Mulder, J. L., Lange, A. E., Turner, A. D., & Wild, L. 2003, in Millimeter and Submillimeter Detectors for Astronomy. Edited by Phillips, Thomas G.; Zmuidzinas, Jonas. Proceedings of the SPIE, Volume 4855, pp. 136-147 (2003).
- Zaldarriaga, M. & Harari, D. D. 1995, Phys. Rev. D, 52, 3276
- Zaldarriaga, M. & Seljak, U. 1998, Phys. Rev. D, 58, 023003
- ---. 2000, Astrophys. J. Suppl., 129, 431
- Zaldarriaga, M., Seljak, U., & Bertschinger, E. 1998, Astrophys. J., 494, 491
- Zeldovich, Y. B. 1972, Mon. Not. Royal Astro. Society, 160, 1P
- Zeldovich, Y. B. & Sunyaev, R. A. 1969, Astrophys. Space Sci., 4, 301
- Zemcov, M., Halpern, M., Borys, C., Chapman, S., Holland, W., Pierpaoli, E., & Scott, D. 2003, Mon. Not. Royal Astro. Society, 346, 1179

