

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

Sigurd Kirevold Næss

Institute of Theoretical Astrophysics
University of Oslo

2012

© **Sigurd Kirevold Næss, 2012**

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 1216*

ISSN 1501-7710

All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, without permission.

Cover: Inger Sandved Anfinsen.
Printed in Norway: AIT Oslo AS.

Produced in co-operation with Akademika publishing.
The thesis is produced by Unipub merely in connection with the thesis defence. Kindly direct all inquiries regarding the thesis to the copyright holder or the unit which grants the doctorate.

Acknowledgements

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

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

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

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

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

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

Contents

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

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

Part I

CMB data analysis for QUIET

Chapter 1

Introduction

Cosmology, the study of the overall composition and history of the universe, is currently in the enviable position of having a well-tested standard model that agrees with all observations, while relying only a low number of adjustable parameters¹. According to this model, which is known as Λ CDM or the “concordance model”, the universe started out in an extremely hot and dense state called the Big Bang some 13.8 billion years ago. This was almost immediately followed by a phase of extremely rapid expansion, called inflation.




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

¹For a thorough introduction to this topic, see for example [1].

During inflation, the size of the universe increased exponentially by at least a factor of 10^{26} , leaving the universe in an almost perfectly flat and uniform state, with small fluctuations with approximately the same amplitude on all length-scales².

After this, the universe continued expanding much more slowly, eventually cooling down enough for protons and neutrons, then light atomic nuclei, and finally atoms to form. The latter happened about 370 000 years after the Big Bang, when the temperature of the universe had fallen to about 3000 K. The formation of atoms turned the plasma that had filled the universe until then into a neutral gas, making the universe transparent for the first time. Due to the finite speed of light, this event, called “recombination”³, is still visible today as an apparent surface, called the *surface of last scattering*, some 13.8 billion light-years away⁴.

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

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

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

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

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

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

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

Parameter	Value	Name
t_0	13.78 ± 0.11 Gyr	Age of the universe
H_0	69.9 ± 1.3 km/s/Mpc	Hubble parameter
Ω_b	0.0485 ± 0.0026	Baryon density
Ω_c	0.244 ± 0.016	Dark matter density
Ω_Λ	0.708 ± 0.0016	Dark energy density
σ_8	0.811 ± 0.023	Fluctuation amplitude
n_s	0.967 ± 0.014	Spectral index
r	< 0.2	Tensor-to-scalar ratio
z_*	1020.3 ± 1.3	Redshift at last decoupling
τ	0.086 ± 0.014	Optical depth of reionization

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

However, the most sensitive test of the model currently available is radiation reaching us from the surface of last scattering, called the cosmic microwave background (CMB). We observe the CMB as perfect black-body radiation corresponding to a temperature of 2.725 K, a temperature which is nearly uniform in all directions on the sky⁵. But since the CMB is an image of the surface of last scattering, it cannot be completely uniform; it must contain the faint fluctuations that were present in the universe at that time.

1.1 The CMB anisotropies

These anisotropies in the CMB were first detected by the COBE satellite in 1992 [2], and were later mapped out in more detail by several other experiments, including the WMAP satellite [3, 4, 5, 6] (see fig. 1.2). The primordial fluctuations produced during inflation are ultimately sourced by random quantum fluctuations, which are expected to follow a statistically isotropic and homogeneous⁶ Gaussian distribution, and these properties are inherited by the temperature fluctuations at the surface of last scattering. The precise position and value of each positive and negative fluctuation is therefore not of cosmological interest, but the statistical properties

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

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




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

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

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

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

where a_{lm} are the coefficients of the decomposition of the fluctuation field $\Delta T(\hat{\theta})$ in terms of the spherical harmonics⁷ $Y_{lm}(\hat{\theta})$. The CMB temperature fluctuation power spectrum as measured by the WMAP satellite is compared to the ΛCDM best fit in figure 1.3. The fit is excellent, despite the much greater number of data points than model parameters. This lends confidence to the idea that ΛCDM is correct, or at least a very good approximation to reality.

For the most part, the ΛCDM model builds on the two well-tested theories general relativity and the standard model of particle physics, which together form the status quo of our understanding of the particles and forces

⁷The spherical harmonics are eigenvectors of the angular part of the ∇^2 operator. They are analogous to the normal harmonic functions, and form an orthogonal basis set on the sphere.

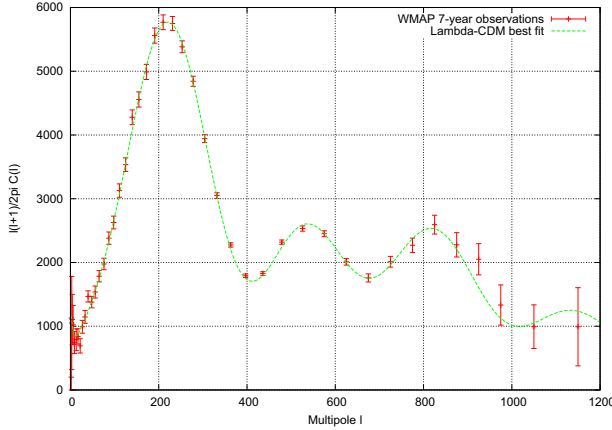


Figure 1.3: The power spectrum of the CMB temperature fluctuations as measured by the WMAP satellite, compared to the best-fit Λ CDM model.

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

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

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

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

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

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

scalar power at the scale with wavenumber $k = 0.05/\text{Mpc}$. So far, no tensor perturbations have been detected, so r is consistent with zero, but with a quite large confidence interval: $r < 0.2$ [8]. This bound is based on the effect of tensor perturbations on the CMB temperature power spectrum, and derives most of its significance from the multipoles $l < 10$, which are fundamentally limited in sensitivity by cosmic variance⁸. It is therefore not possible to significantly improve on this bound using the CMB temperature fluctuations.

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

1.2 Polarization

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

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

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

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

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

1.2.1 Stokes parameters

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

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




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

$U = 2\text{Re}(E_x E_y^*)$ and $V \equiv 2\text{Im}(E_x E_y^*)$. These represent respectively the total intensity of the radiation⁹, the linear polarization along the x (positive) and y (negative) axes, linear polarization along the $x+y$ (positive) and $x-y$ (negative) direction, and the right-handed (positive) and left-handed (negative) circular polarization, as illustrated in figure 1.4.

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

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

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

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

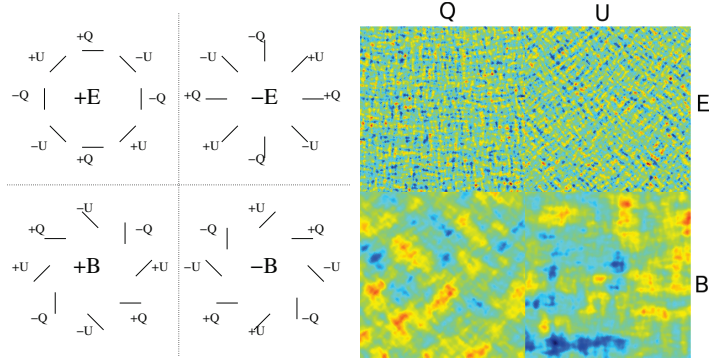


Figure 1.5: Left: Q/U pattern surrounding a point with positive/negative E/B. E-modes are curl-free, while B-modes are divergence-free. Right: Q/U maps corresponding to an E-only (top) and B-only (bottom) CMB realization. The scale difference between the E and B modes shown here is due to the expected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

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

The scalar E field and the pseudo-scalar field B are rotationally invariant alternatives to Q and U. E and B are respectively curl-free and gradient-free, and are named in analogy to the similar properties of the electric and magnetic field. The value of E and B in a given point can be defined as radially-weighted averages of Q_r and U_r , the Stokes Q and U parameters as defined in a polar coordinate system centered on that point¹⁰. E and B modes and their connection to Q and U are shown in figure 1.5. A pedagogical explanation can be found in [10], while [11] has an exact full-sky treatment of the topic.

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

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

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

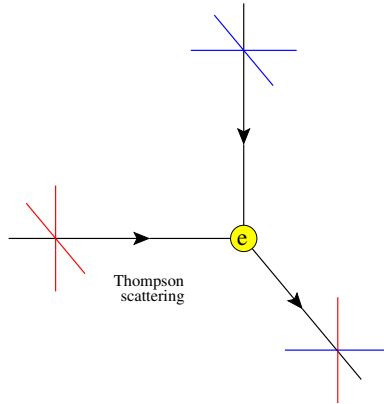


Figure 1.6: Local quadrupole anisotropies introduce linear polarization at the surface of last scattering. In this example, unpolarized light (i.e. light with equal amounts of all types of polarization) arrives at an electron from above and the left. The outgoing radiation inherits the horizontal polarization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

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

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

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

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

These perturbations introduce quadrupole moments with different symmetries as illustrated in figure 1.7. The symmetry of the scalar perturbations ensure that they cannot produce B-mode polarization. Hence, since vector perturbations are unlikely to be present, a detection of B-mode polarization would be evidence for tensor perturbations¹¹.

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

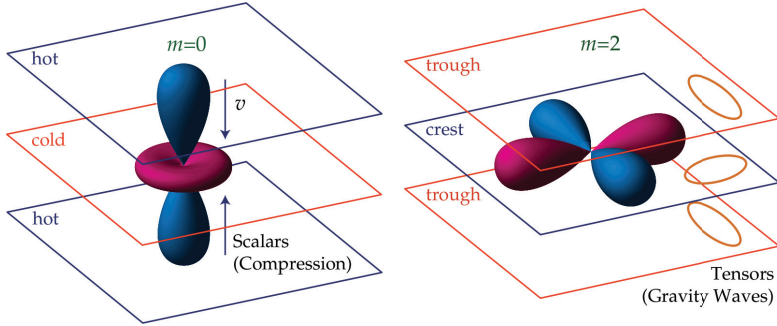


Figure 1.7: Illustration of the local quadrupole anisotropy induced by scalar and tensor perturbations. Scalar perturbations correspond to $m = 0$ modes, while tensor perturbations have $m = \pm 2$. The symmetry of scalar modes prevent them from inducing B-mode polarization. Figure borrowed from [12].

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

1.3 Detector technology

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

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

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

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

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

The most effective way of increasing overall sensitivity is therefore to add more detectors. If the noise in each detector is uncorrelated, the sensitivity will fall¹² as the square root of the number of detectors. The challenge is then to fit as many detectors as possible into a focal plane. This is an area where bolometers used to have a big advantage due to being smaller, cheaper and easier to mass produce, which helped make bolometers the norm.

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

1.4 The Q/U Imaging Experiment

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

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

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

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

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

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

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

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

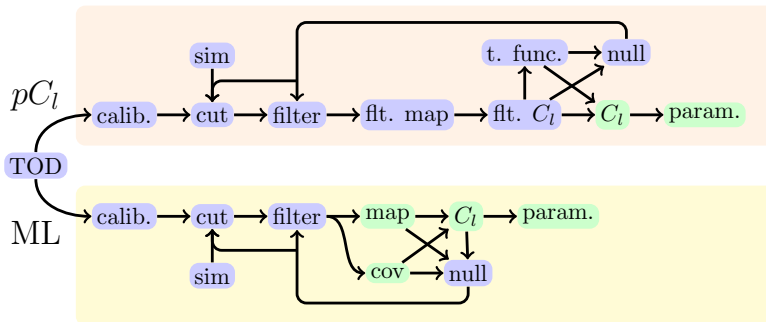


Figure 1.8: Stylized overview of QUIET’s two analysis pipelines. The major internal steps are marked in blue and green, where green indicates a result that is useful in its own right. “flt.” is short for “biased due to filtering”, while “t. func.” is short for “transfer function”. Both pipelines share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

Chapter 2

Telescope and calibration




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

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

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

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

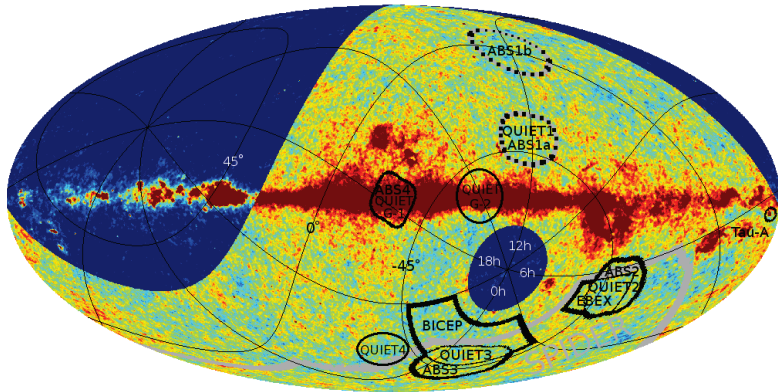


Figure 2.2: The location of QUIET's 4 CMB patches and two galactic foreground patches in galactic coordinates. The bright part of the sky indicates the observable region based on the telescope's location and elevation limits. Several of QUIET's patches overlap with those of other current or planned experiments,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

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

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

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




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

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

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

2.0.1 CES detection

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

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

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

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

1. The deglitched data taking mode is 3
2. The elevation is within the allowed limits of $[-\frac{\pi}{2}, \frac{\pi}{2}]$
3. No more than 30 consecutive samples¹ have an absolute azimuth change of less than 10^{-5} radians per sample.

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

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

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

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

Whenever one of these conditions is violated, the current CES candidate is ended and a new one started. The beginning time of each CES candidate is then adjusted so that it is at least 1 minute later than the last phase switch event² that happens before the end of that candidate. Finally, the candidate is accepted as a usable constant elevation scan if its duration is longer than 5 minutes.

Patch detection

Each of these CESes is then classified according to the objects hit during its course. The boresight pointing for each sample of the CES is translated to galactic coordinates³ and the angular distance d to a predefined set of candidate objects is calculated. A hit is defined as $d < R + r_b + r_v + r_f$, where R is an approximate radius of the object, r_b is the beam 5 sigma radius, r_v is half the distance traveled during a sample and r_f is the focal plane radius. The objects hit during the CES are then sorted according to their brightness, and the brightest one is assigned to the CES. The list of objects considered can be seen in figure 2.1. This classification scheme ensures that a scan of a CMB patch where the Moon happens to pass in front of the patch registers as a usable Moon scan instead of a contaminated CMB scan.

Implementation

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

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

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

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

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

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

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

2.0.2 Alternative scanning patterns

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

2.1 Detector data and noise properties

The telescope provides a 100 Hz time-stream modulated at 50 Hz for each detector⁴. After correcting for a well-measured non-linearity in the detector response and demodulating⁵, the result is one 50 Hz time-stream for each detector, for a total of 72 data streams for the Q-band array and 364 for the W-band array.

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

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

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

This profile only holds within a CES; over longer periods the noise is not stationary, and one must therefore estimate σ_0 , f_{knee} and α individually for each CES. The time dependence is illustrated for one of these parameters, σ_0 , in figure 2.5.

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

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

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




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

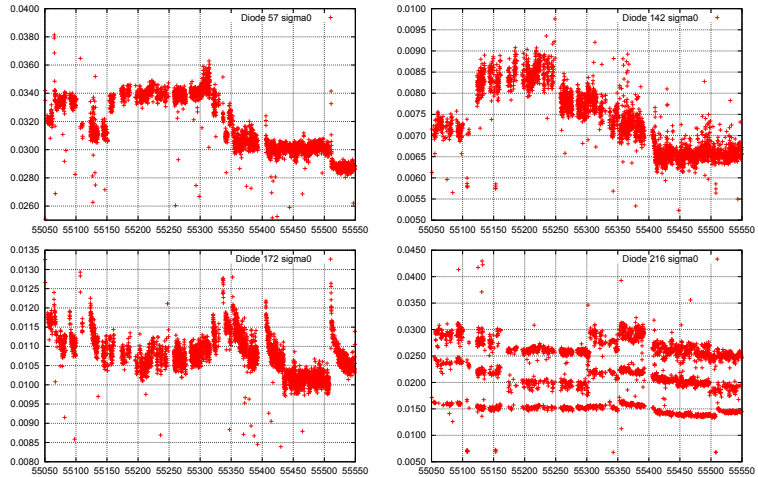


Figure 2.5: The white noise floor σ_0 shows a complicated time dependence with a different pattern for each module, and to a lesser degree for the detector within the module. The four panels show the measured σ_0 values per CES for four different detectors. Large changes in σ_0 happen at discrete events which are common for the detectors. The response to these events differs, but is correlated between the detectors, as illustrated in figure 2.6.




Figure 2.6: The time-dependence of the white noise floor σ_0 is strongly correlated between the detectors. This figure shows the W-band season correlations in σ_0 .




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

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

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

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

where $N_{dd'ff'} \equiv \langle n_{df}n_{d'f'}^\dagger \rangle = \mathcal{F}_{fi}N_{dd'ff'}\mathcal{F}_{ff'}^{-1}$ is the frequency-domain noise covariance matrix between detectors d and d' and frequencies f and f' , and $C_{dd'f}$ is the detector correlations, and where all parameters vary from CES to CES.

2.1.1 Noise estimation

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




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

1/f profile

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

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

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

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

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

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

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

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

The remaining parameters can be determined with a nonlinear search.

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

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

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

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

The performance of the resulting masked $1/f$ estimator at recovering the input parameters for noise only simulations is illustrated in figure 2.10, which shows that the estimator is unbiased. However, this unbiasedness only carries over to real data if no significant signal is left unmasked, and if the real noise actually follows a $1/f$ profile. This is investigated in figure 2.11. It shows the season co-added⁶ detection of a deviation between the data and model for one detector, expressed as the number of standard deviations per 1.25 mHz bin in frequency. This can be calculated by noticing that $\frac{P_f}{\psi(f)}$ follows a scaled chi-squared distribution with 2 degrees of freedom and mean 1. Summing these in bins over the whole season, we arrive at

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

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

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

⁶Using only CESes from CMB patches that pass the data quality cuts.

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

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

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

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

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

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

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

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

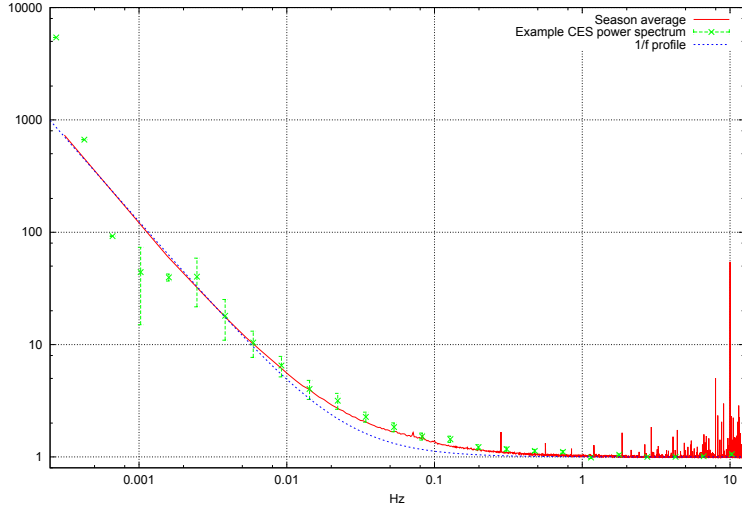


Figure 2.9: The measured noise profile compared to the model. The green points are the measured power spectrum for a given CES and detector, where the white noise floor has been normalized to 1. The red curve is the season average of such curves for the same detector. The blue curve is an asymptotically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




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




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

Correlations

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

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

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

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

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

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

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

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

2.1.2 Filters

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

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

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




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

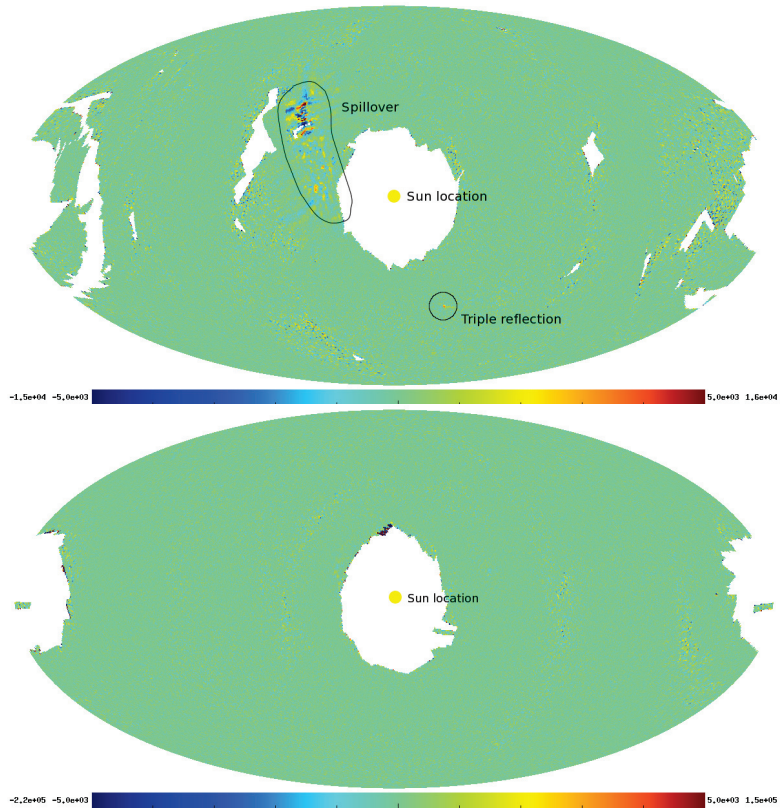


Figure 2.13: Map of the stokes Q component of the sky in Sun-relative coordinates for the detector and deck orientation that shows the strongest contamination from the Sun. The center of the map is the position of the sun. The central $\sim 30^\circ$ are unexposed due to sun avoidance. The first and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

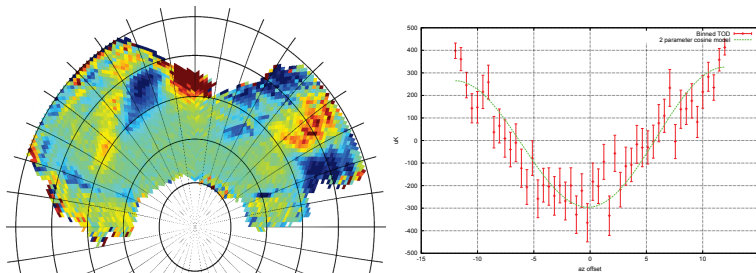


Figure 2.14: Left panel: A Stokes Q map in horizontal coordinates based on the CMB scans of the Q-band observations at a deck orientation of 165° . In these coordinates, any signal from the sky will average down, while any function of az and el , such as the ground signal, will add up. Due to their different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

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

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

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

Combining these cutoffs with 1/f-type apodization⁸, we get a total filter profile of

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

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

Ground pickup is not localized to a well-defined set of frequencies. They are spread out over the same frequencies as the CMB itself, and hence can't be handled with a bandpass filter. It is formed by sidelobes hitting the ground, and because the ground, unlike the sky, is stationary with respect to the telescope, the observed ground pickup will depend only on the telescope azimuth, elevation, deck orientation, the detector's position and orientation in the focal plane, and the ground signal itself. Hence, during a CES, where elevation and deck orientation are constant, and the ground properties presumably do not change noticeably, the ground pickup will be a function of azimuth only. We therefore employ a filter in azimuth to reduce the impact of the ground. Implementation-wise, this is handled as a generalized filter (section 4.5.2) with a basis choice⁹ of

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

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

⁸We also evaluated cosine apodization of the form

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

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

⁹Chebyshev or Legendre polynomials would be other good alternatives.

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

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

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

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

2.1.3 Cuts

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

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

The noise consistency cuts can be divided into two subclasses: Direct cuts based on data that will actually be used in the analysis (i.e. data that survive the filter), and indirect cuts which use the filtered data as a proxy for detecting effects that are likely to also contaminate the unfiltered data. For direct cuts, the goal is to detect deviations from noise, and a cut threshold of 4-5 sigma¹¹ is high enough to avoid significantly truncating the noise distribution through false positives while still being strict enough to detect all but the weakest effects.

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

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

Name	Cut Q (%)		Cut W (%)		Description
	P	T	P	T	
Static	12.3	5.3	9.7	8.9	Broken detectors, missing data or precalibration data.
Gain	7.4	0.0	11.5	50.0	Cuts detectors where no gain model has been established due to missing calibration observations.
Noise fit	7.4	0.0	9.1	8.3	CES-detectors where the 1/f-profile fit failed. This usually happens due to bad weather.
Noise outlier	9.3	40.8	12.3	55.4	CES-detectors with σ_0 or α more than 5 sigma away from that detector's median, or f_{knee} more than 5 sigma above the median.
ADC residuals	4.0	19.3	6.5	37.6	Cuts if evidence for a non-linear ADC response is greater than 5 sigma. In practice, this is also sensitive to bad weather.
Weather	8.2	6.4	9.9	8.5	Cuts if the rms of the 10 s-rms values of the TOD is more than 5 sigma away from the season median, and similar for the 30 s-rms values. Also cuts if the current PWV > 5 mm.
Fourier χ^2	16.2	39.9	18.5	31.7	Validates the χ^2 of frequency ranges by computing the number of standard deviations $\delta(\chi^2) \approx \left \frac{\chi^2 - n}{\sqrt{2n}} \right $, where $n = 2N_{\text{freq}}$ is the number of degrees of freedom. Cuts if $\delta(\chi^2) > 4$, $\delta(\chi^2_{\text{scan} \pm 10 \text{mHz}}) > 10$, $\delta(\chi^2_{0 \text{Hz} \dots 0.2 \text{Hz}}) > 10$, $\delta(\chi^2_{10 \text{Hz} \dots}) > 20$ or $\delta(\chi^2_{\text{filtered}}) > 4$.
TOD χ^2	6.0	24.3	6.4	25.0	Cuts if a χ^2 based on every 10th sample of the filtered TOD is more than 4 sigma away from the expectation value, or if the sample with the highest absolute value in the filtered TOD is higher than 7 times its rms.
Azimuth χ^2	2.7	9.1	3.5	8.5	Cuts if azimuth structure in the TOD is detected at more than 4 sigma.
Spike	5.4	21.1	7.8	24.3	Cuts if the 1 Hz and 1.2 Hz spikes associated with the cryostat, and their first 6 harmonics are detected at more than 20 sigma, or if any 0.1 Hz interval between 0.2 Hz and 4.5 Hz (Q) or 9.5 Hz (W) has a χ^2 excess of more than 7.5 sigma.
Sun	8.6	9.9	2.2	0.0	Rejects CES-detectors if any of their samples expressed in Sun-relative coordinates hit an area with a season-average detection of the sun of more than 5 sigma. See the text for details.
Conservative	18.3	17.0	16.5	15.1	Rejects all detectors for a given CES if more than 40% of the detectors fail the weather, TOD, Fourier, azimuth, spike or sun cuts.
Total	41.7	65.7	41.7	91.6	Less than the sum because of a large degree of overlap.

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

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

The Sun cut

As illustrated in figure 2.13, the telescope sidelobes are quite complex, exhibiting both narrow, intense regions like the triple reflection sidelobe and broader but weaker regions like the spillover sidelobe. Given some region r in sun-relative coordinates¹³, we can determine the strength of detection of the sidelobe in this region by building a map m and its covariance C in this region, and calculating $\chi^2 = m^T C^{-1} m$.

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

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

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

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

¹³See 2.2 for the definition of this coordinate system

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

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

2.2 Pointing

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

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

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

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

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

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

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

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

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

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

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

rotation. The apparent to astrometric rotation R_{ea} is a nonlinear effect taking into account the precession, nutation and aberration etc. due to the Earth's orbit¹⁵.

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

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

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

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

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

Equation (2.24) does not prescribe how to calculate provide ψ from ψ_0 .

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

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

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

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

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

2.2.1 Verifying the pointing

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

¹⁵This is provided by the NOVAS library.

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

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

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

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

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

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

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

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

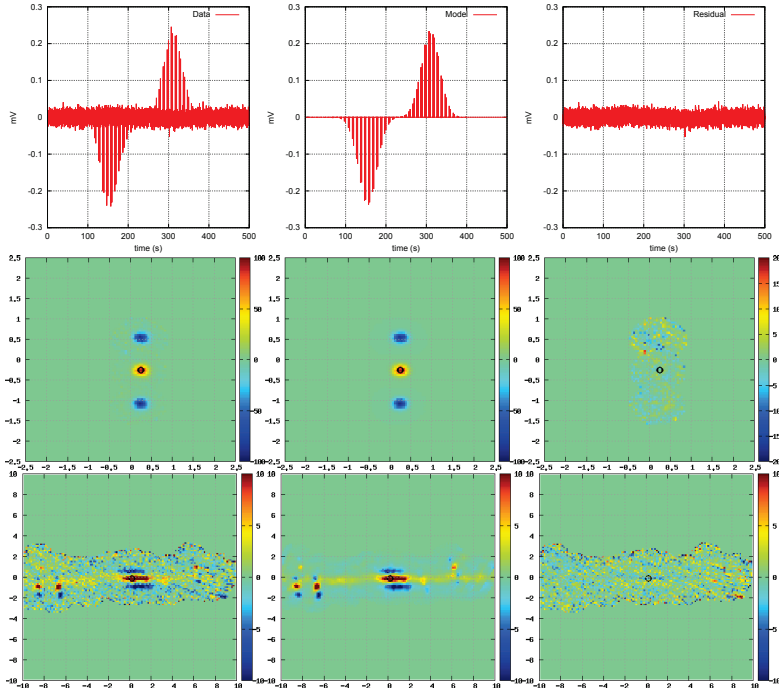


Figure 2.15: Illustrations of the fitting of a source model to observations. Top: time-ordered data for a scan of Jupiter (left), the best fit model (middle), and the residual (right). Middle: The same data naively projected to pixel domain using a fast, simplified map-making procedure which ignores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

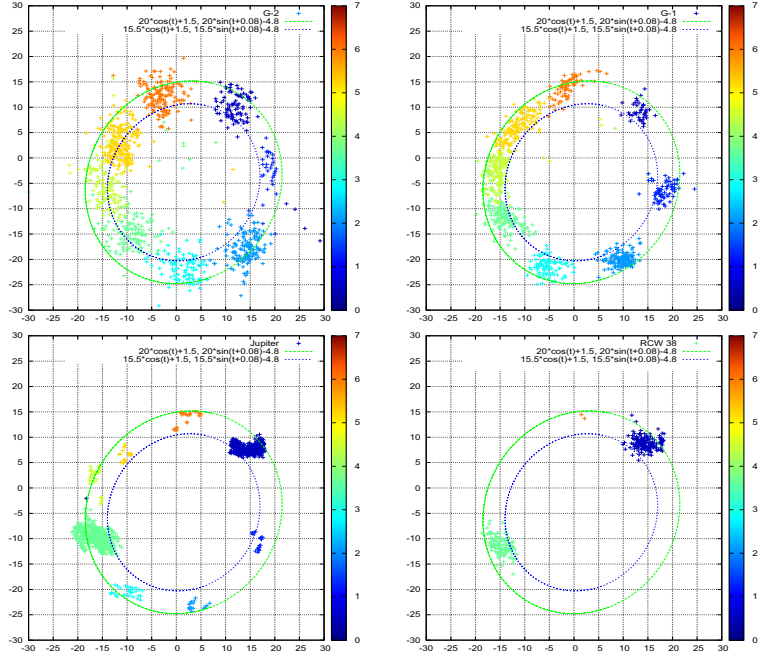


Figure 2.16: The pointing residuals based on single-CES temperature measurements of four objects: the two galactic patches, Jupiter and RCW 38. The horizontal and vertical axes show $(\alpha_{\text{obs}} - \alpha_{\text{true}}) \cos(\epsilon_{\text{true}})$ and $\epsilon_{\text{obs}} - \epsilon_{\text{true}}$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

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

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

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

2.2.2 Mount model

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

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

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

fore be safely ignored.

However, of the remaining effects, none are capable of producing the elliptical pointing residuals in figure 2.16¹⁶, nor can they explain the peculiar elongation of the scatter in the angular direction in the plots or the “U” shape observed in the residuals of Jupiter, the highest signal-to-noise source used in the calibration¹⁷. We therefore adopted a phenomenological mount model which allows for some unanticipated extra degrees of freedom in order to match the observed pointing scatter. The resulting model has 14 parameters, and is summarized in table 2.5. Incorporating these rotations into the full rotation from detector to sky results in

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

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

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

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




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

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

2.2.3 Full-season validation of the mount model

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

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




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

tered into the pointing analysis were based on temperature detectors.

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

We performed three tests to investigate these possibilities:

1. Comparison of an $N_{\text{side}} = 2048$ full-season co-added map of Jupiter¹⁸ based on the secondary total power data stream¹⁹ from the polarization detectors, with the expected size of a point source convolved

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

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

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

2. A similar test for the $N_{\text{side}} = 2048$, full-season co-added map of the point source PMN J0538-4405 in patch CMB-2 (fig. 2.20), which finds a best-fit FWHM of $14.2' \pm 1.0'$. Comparing this to a temperature detector beam of $13.1' \pm 0.2'$, we find a scatter FWHM of $5.5' \pm \frac{2.5'}{5}$.
3. We compared the full W-band season CMB temperature power spectrum and fit it to the model $C_l = AC_l^{\Lambda\text{CDM}} e^{-l(l+1)\sigma_{\text{FWHM}}^2(8\log 2)^{-1}}$ to determine the effective beam, and found a FWHM of $14.1' \pm 0.1'$, resulting in a scatter FWHM of $5.2' \pm 0.7'$.

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

2.3 Gain and beam

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

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

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

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

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

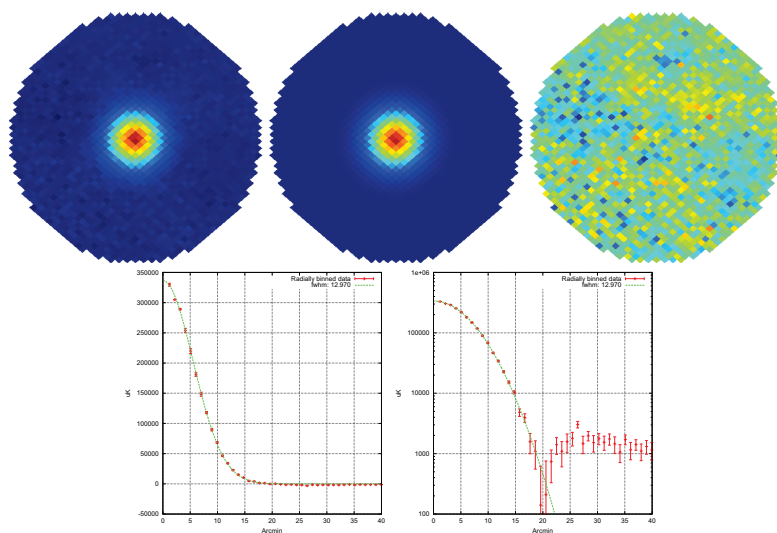


Figure 2.19: Top left: Co-added $N_{\text{side}} = 2048$ temperature map of Jupiter from the QUIET W-band season in object-centered coordinates based on the secondary total power data stream from the polarization detectors, in arbitrary units. Top middle: Best fit Gaussian with a major FWHM of $13.2'$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

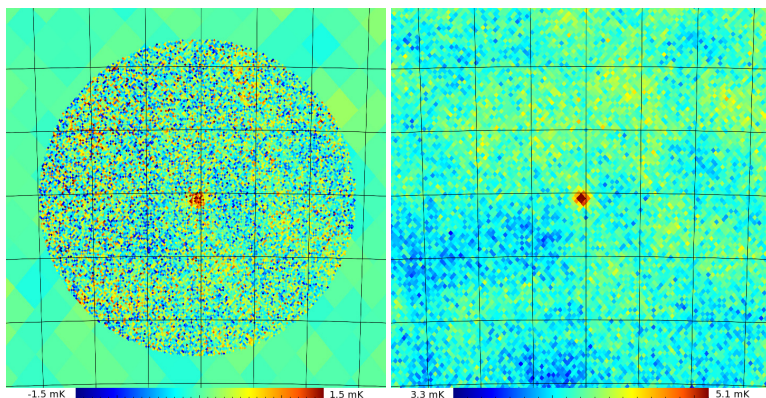


Figure 2.20: Left: A full-season co-added $N_{\text{side}} = 2048$ temperature map of the source PWM J0538-4405 in patch CMB-2, based on the primary data stream from the temperature detectors. Right: WMAP's 7 year W-band $N_{\text{side}} = 1024$ temperature map of the same source. The object can be seen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

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

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

2.3.1 Polarization gains

A thorough analysis of the polarization gains for the Q-band season was performed in [28], where it was found that gain measurements based on the Moon, Tau A, elevation dips and a wire-grid²⁰ were correlated with the measured temperature of the electronics T , resulting in a gain model of the

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

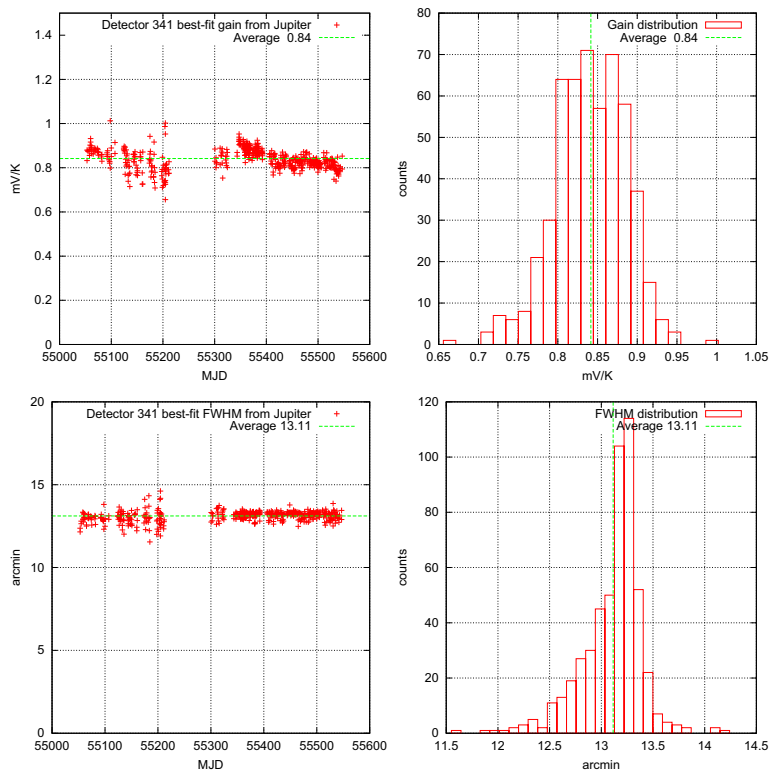


Figure 2.21: An example of gain and beam size estimates produced by fitting a Gaussian profile to Jupiter observations (see section 2.2.1) for a single detector. The top row shows the gain estimates as a function of time (left) and a histogram of these (right). There is some evidence of time-dependence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

form

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

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

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

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

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

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

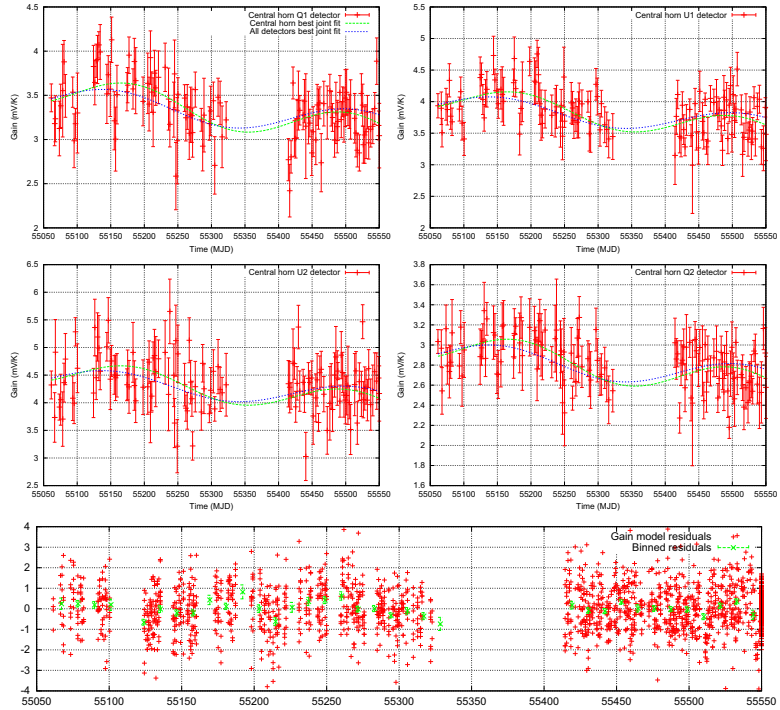


Figure 2.22: Top: Polarization gain measurements from Tau A and the wire-grid for the four detectors in the central module. These are shown as red points with error bars. A roughly seasonal modulation is clearly visible, as well as a slight downwards trend, and the shape of these is the same for the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

Chapter 3

Null tests

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

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

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

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

3.1 Signal-less validation

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

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

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

3.2 Map null-tests

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

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

which should be chisquare-distributed with $N_{\text{dof}} = N_{\text{pix}} N_{\text{stokes}}$ degrees of freedom, with a mean of N_{dof} and a standard deviation of $\sqrt{2N_{\text{dof}}}$. A significant excess in this number indicates that there remains unfiltered junk in the data¹, while a significant shortfall only can be produced by an overestimate of the noise.

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

¹Or that the noise has been under-estimated.

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

3.3 Pseudo- C_l null-tests

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

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

The function f allows us to choose which parts of the maps we wish to test. We are ultimately interested in computing the binned angular power spectra² C_b^{EE} , C_b^{BB} and C_b^{EB} . It therefore makes sense to make f simply be the pseudo- C_l approximation of power in these bins,

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

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

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

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

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

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

An advantage of this variant of null-tests is that by processing the data and simulations the same way, any bias in the map-making procedure is automatically taken into account. Hence, we do not need to use the full, unbiased maximum likelihood map-making equation; much faster, but biased, binned maps³ can be used instead. Together with the lack of need for the covariance matrix, this results in the pseudo- C_l null-tests being 1-2 orders of magnitude faster than the maximum likelihood null-tests, greatly increasing the practical number of null-tests.

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

3.4 Null test suite

Table 3.4 shows a summary of the W-band null test suite together with the final probability to exceed (PTE) for each combination of test and CMB patch⁴. The suite consists of 23 individual data splits, each of which results in one test per multipole bin per EB-combination for each patch, for a total of $N_{\text{test}}N_{\text{bin}}N_{\text{spec}} = 23 \cdot 20 \cdot 3 = 1380$ tests per patch, or 5520 tests in total.

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

³See section 4.4.1.

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

Split	Sensitive to	1	2	3	4
Time (MJD)	Gain drifts	0.24	0.01	0.89	0.77
Focal plane radius	Ground, beam ellipticity	0.15	0.51	0.21	0.61
ADC nonlinearity	Uncorrected nonlinearity, weather	0.68	0.95	0.54	0.02
Power at f_{scan}	Ground, weather	0.25	0.35	0.53	0.76
Power at 10 Hz	Electronic problems	0.41	0.56	0.53	0.48
Elevation	Ground, mount	0.58	0.08	0.39	0.55
Deck angle	Ground, mount	0.54	0.12	0.27	0.77
Azimuth	Ground, mount	0.76	0.95	0.66	0.11
T_{encoder}	Weather, electronics	0.08	0.09	0.24	0.06
$\Delta T_{\text{encoder}}$	Weather, electronics	0.13	0.54	0.09	0.51
T_{cryostat}	Weather, electronics	0.17	0.86	0.31	0.26
$\Delta T_{\text{cryostat}}$	Weather, electronics	0.86	0.51	0.07	0.68
Gain	Gain model	0.55	0.81	0.01	0.34
Water vapour (PWV)	Weather	0.52	0.30	0.79	0.94
Wind	Pointing	0.33	0.02	0.57	0.94
T_{ambient}	Weather	0.62	0.16	0.42	0.92
f_{knee}	Weather, noise model	0.21	0.81	0.49	0.03
σ_0	Weather, noise model	0.25	0.41	0.06	0.52
Module mean ν	Strong foregrounds	0.58	0.27	0.58	0.87
Assembly board	Electronic problems	0.55	0.38	0.60	0.25
TOD rms variability	Weather	0.93	0.18	0.47	0.17
Q vs. U detectors		0.73	0.36	0.99	0.29
$I \rightarrow QU$ leakage	Leakage misestimation	0.16	0.11	0.91	0.43
Total χ^2	Deviation from expected power	0.31	0.08	0.31	0.43
Max χ^2	Individual null-failures	0.63	0.02	0.72	0.80
Mean χ shift	Power excess	0.20	0.27	0.21	0.97

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

Chapter 4

Map making

4.1 Measuring the sky with a scanning telescope

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

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

4.2 From sky to time-ordered data

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

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

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

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

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

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

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

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

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

Here i is the pixel index while α is the signal component within the pixel¹. Our job, then, is to expand $P_{dti\alpha}$ into something concrete enough to be implementable. A very general such expansion is

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

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

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

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

This daunting expression can be dramatically simplified by adding a few reasonable assumptions. The most important of these comes from assuming a constant, detector-independent, circular beam². A circular B will commute with R , and if B additionally is constant, we can apply it to the map once and for all, avoiding the need to convolve the map with the beam at every time step,

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

Here $\tilde{m} = Bm$ is the beam-smoothed map, and I have used the approximation that the g and τ commute³.

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

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

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

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

³The gain is ideally constant, but detectors are imperfect, and the gain may in practice drift gradually. However, it will never be allowed to drift at anywhere near the time scale of the detector time response. Hence, $g_{dt} \tau_{dt'} \approx \tau_{dt'} g_{dt'}$.

a_{dh} .

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

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

The biggest remaining expense in equation (4.6) is the detector time convolution τ . The importance of this factor is strongly dependent on the type of detector used. For amplifier-based radiometers it is practically a delta function, and can be ignored, while for bolometers it takes the form $\tau_{dtt'} = e^{-\frac{t'-t}{z_d}}$ for $t' > t$ and 0 otherwise, with z_d normally being 10 ms or smaller. The effect of τ is similar to the effect of sampling at discrete time steps – it acts as a low-pass filter that limits sensitivity at high frequencies and hence, via the scanning pattern, the smallest scales on the sky. The scanning pattern is typically chosen with this in mind, so that τ has negligible effect on the scales of interest for the experiment, and it is therefore usually safe to ignore this factor even for bolometers.

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

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

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

4.3 The noise

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

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

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

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

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

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

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

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

with n being the number of time samples, so

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

This expression only depends on the time difference $\Delta t = t - t'$, so we only need to store the time correlation function⁵ $N_{dd'}^{\pm 1}(\Delta t) = N_{dd'0\Delta t}^{\pm 1}$.

4.4 From time-ordered data to sky

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

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

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

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

which gives a maximum likelihood estimator of

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

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

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

This is an unbiased estimator

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

with covariance

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

For implementation purposes it will be convenient to introduce the quantity

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

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

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

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

This filtered estimator has covariance

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

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

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

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

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

such that

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

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

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

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

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

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

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

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

so that

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

The observant reader will notice that this form of the filtered estimator is equivalent to an unfiltered map estimator based on the data model $d = Pm + n + g$, where g is an extra Gaussian component with covariance G . That is, an alternative way of implementing filtering is by assuming the presence of extra noise. This leads to the same map estimator, but a simpler form for the covariance: $M = (P^T(N + G)^{-1}P)^{-1}$. However, for this to be valid the extra noise must *actually be present* in the data, otherwise the covariance of the final map will be overestimated⁶. If signal contaminants like bad weather and ground pickup could be accurately modeled as Gaussians, this would be an optimal and efficient way of handling them. This is unlikely to be the case in practice.

4.4.1 Binned maps

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

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

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

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

4.5 Filters

The job of the filters is to remove parts of the data that cannot be incorporated into the noise model, either because it is not Gaussian, not stationary, or too poorly known. Common effects that need to be filtered out include atmospheric disturbances, sidelobe pickup from the ground and the sun, glitches in the electronics, etc. (see section 2.1.2). A good filter should remove as much as possible of these effects while removing as little as possible of the actual signal⁷ and this requires a rough model for where the effects appear, such as a frequency range, certain time intervals, etc., so that these regions can be specifically targeted.

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

4.5.1 Frequency filters

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

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

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

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

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

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




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

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

4.5.2 General filters

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

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

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

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

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

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

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

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




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

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

Note that we are not assuming that the covariance G is a good model for the contamination here – if we had a good model of it we could treat it as a noise component instead⁸. Instead, what is being assumed here is that G is much larger than the variance of the spurious signal, so that to a good approximation the latter no longer matters. A simplifying assumption that guarantees this is to let the nonzero parts of $G \rightarrow \infty$ by setting $C^{-1} = 0$. In addition to removing C from equation (4.32), it also results in $(N_{FF} + G)^{-2}N_{FF} = (N_{FF} + G)^{-1}$, greatly simplifying the application of the Woodbury identity to equation (4.31).

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

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

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

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

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

⁸See section 4.4.

this simplifies somewhat to

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

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

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

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

4.6 Practical implementation of the map-making equations

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

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

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

4.6.1 Implementing $P^T D V$

By using equation (4.7), we see that

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

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

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

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

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

```

W = 0
for all d, t, α, I, h do
  WpdhtαI += ψdhtα[τDV]dtI
end for

```

4.6.2 Implementing $P^T DP$

Starting again from equation (4.7), we find

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

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

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

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

```

W = 0
for all d, d', t, h, h', α, β do
  for -tcorr ≤ Δt ≤ tcorr do
    Wpdhtpd'h'(t+Δt)αβ += ψdhtα[τTDτ]dd'Δt0 ψd'h'(t+Δt)β
  end for
end for

```

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

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

4.6.3 Solving the map-making equation

Equation (4.22) for solving for the final map and its covariance matrix consist of a straightforward dense matrix operation⁹, and can be handled by a linear algebra package like LAPACK¹⁰ [35]. In principle the matrices involved are all symmetric and positive definite, and could be efficiently solved by using Cholesky factorization. Sadly, this is sabotaged by the effect of filtering, which tends to make the covariance matrix poorly conditioned or singular. For example, a high-pass filter which assigns zero weight to low-frequency modes will make the average of the map completely indeterminate. The map covariance matrix will therefore have an infinite eigenmode corresponding to the average of the map. Similar effects apply to a lesser degree to other modes that are damped by filters¹¹. One must therefore use slower, more robust methods such as eigenvalue decomposition.

4.6.4 Implementation in the QUIET pipeline

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

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

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

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

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

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

Building r , A_f^{-1} and A_{ff}^{-1} is done in the program **tod2map**, which consists of about 4000 lines¹² of Fortran 90 code, which processes constant elevation scans in parallel using MPI, and co-adds these into the final r , A_f^{-1} and A_{ff}^{-1} . The computation time for this step scales linearly with the total length of the time-ordered data involved, and required $\sim 10^4$ CPU hours to process all the CMB patches in the Q-band analysis, and $\sim 5 \cdot 10^4$ CPU hours for the W-band analysis due to the larger amount of data involved.

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

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

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

smoothed by

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

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

Memory-wise, the cost is driven by the need to store the two A -matrices, though their symmetry permits us to only store half of each. To avoid artificial loss of resolution due to pixelization, the pixels should be significantly smaller than the beam, which in the Q- and W-bands have a FWHM of 27' and 12' respectively¹⁴. We therefore require a HEALPix N_{side} of at least 256 and 512 corresponding to a pixel diameter of 13.7' and 6.9' for the two frequency bands. For a typical QUIET CMB patch, this results in $\sim 10^4$ pixels in the Q-band and $\sim 4 \cdot 10^4$ in the W-band, which coupled with two polarization components per pixel in the case of a polarization only analysis results in $\sim 2 \cdot 10^4$ and $\sim 8 \cdot 10^4$ degrees of freedom in the map, and the square of these numbers in the matrices. Using single-precision (4 bytes floating point numbers) for the matrices, this amounts to about 1.5 GB (25 GB) of memory to store one full matrix for the Q-band (W-band) analysis.

4.7 Biaslessness and deconvolution

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

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

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

¹⁴See section 2.

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

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

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

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

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

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

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

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

¹⁶While primarily a means of visualization, it is also possible to use this as a cut before further data analysis. This could for example be done if one has evidence that certain eigenmodes are contaminated by ground pickup, atmosphere, or similar¹⁷. In this case, the corresponding modes in the covariance matrix should be assigned infinite variance, reflecting the fact that the value of these eigenmodes is being ignored:

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

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

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

4.8 Results

4.8.1 CMB maps

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

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

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

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

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

4.8.2 Foreground maps

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




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




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




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

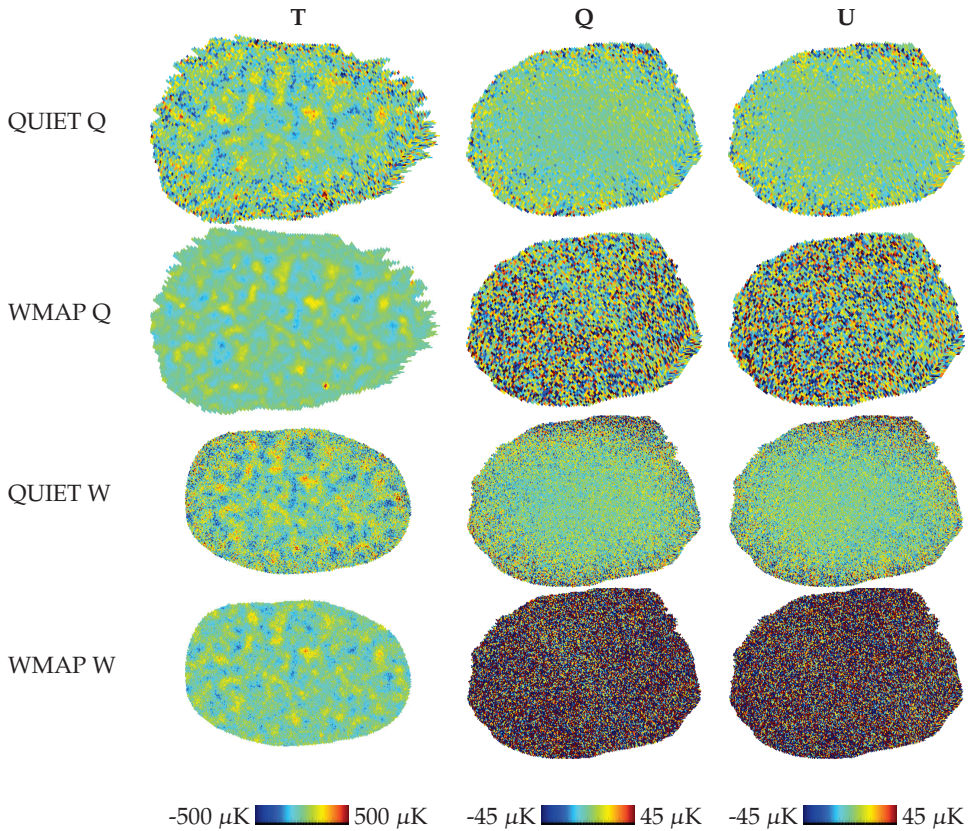


Figure 4.6: Maps for CMB patch 3, centered on galactic $l = 304.6^\circ$, $b = -69.1^\circ$, with a diameter of $\sim 25^\circ$. The columns are temperature (left), Stokes Q polarization (middle) and Stokes U polarization (right). Row 1-2: The QUIET Q-band result compared to the WMAP Q-band map, both at $N_{\text{side}} = 256$. Row 3-4: The QUIET W-band result compared to the WMAP W-band map, both at $N_{\text{side}} = 512$. The S/N is lower here than for patch 1, but it is still possible to make out an E-mode signal in the central, lowest-noise region. The temperature map is consistent with WMAP, but noisier.




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




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

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

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

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

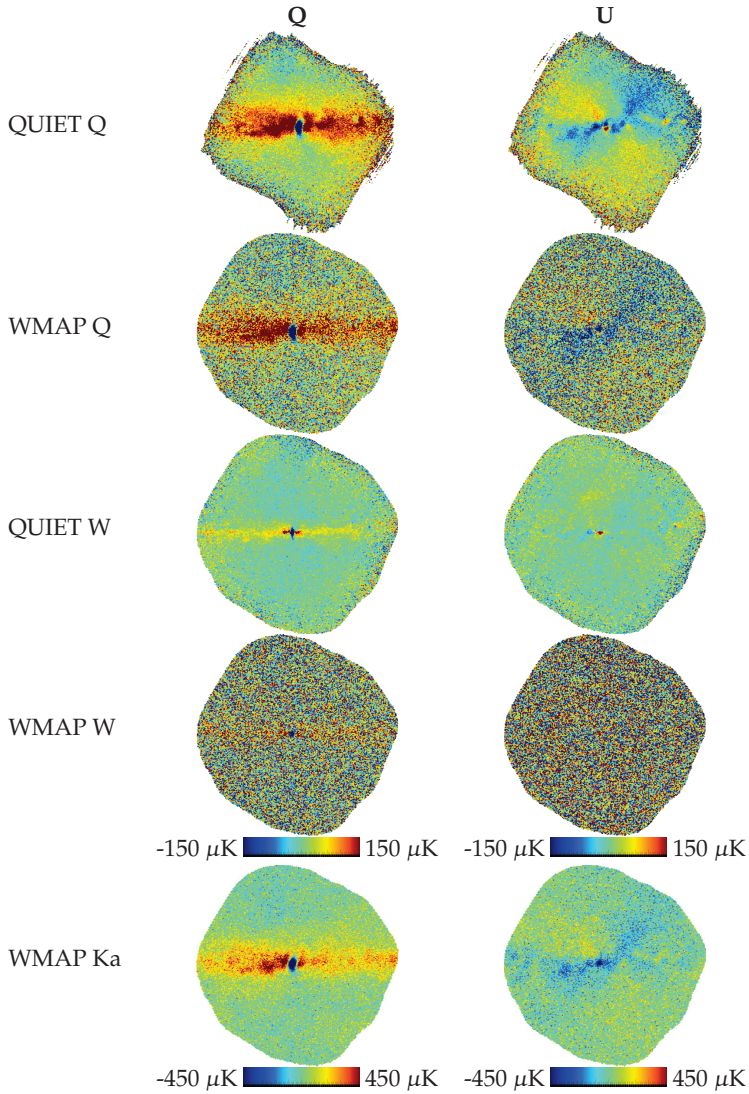


Figure 4.9: Maps of QUIET’s central galactic patch, centered at $l = 0.1^\circ$, $b = -0.1^\circ$, with a diameter of $\sim 20^\circ$. The QUIET Q- and W-band results are consistent with WMAP, but WMAP’s high noise level makes a visual comparison difficult. I therefore provide the high S/N WMAP Ka-band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

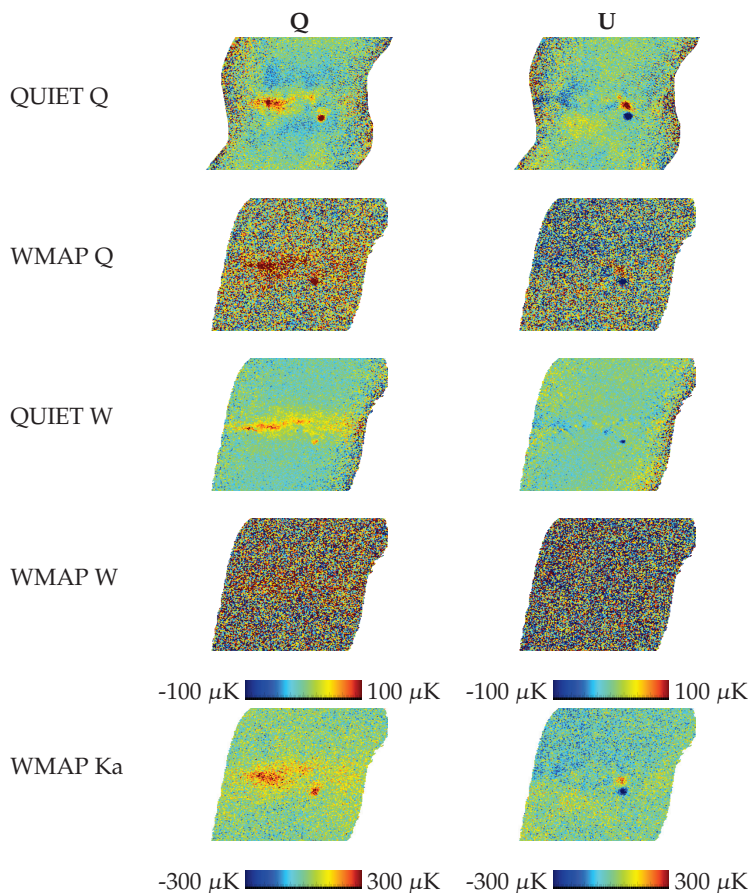


Figure 4.10: Maps of QUIET's other galactic patch, centered at $l = 329^\circ$, $b = 0^\circ$, with a diameter of $\sim 20^\circ$. The columns are the Q and U Stokes parameters respectively. The top two rows show the QUIET and WMAP maps for the Q-band, while the next two rows have the corresponding maps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

Chapter 5

Power spectrum and parameter estimation

Though map-making is a necessary step of the analysis pipeline, and maps are very useful for e.g. foreground analysis, they are not the final scientific goal of CMB experiments. As discussed in chapter 1, the CMB anisotropies are understood to be ultimately sourced by quantum fluctuations set up during inflation, and while the value of these fluctuations at any given position is random, their statistical distribution is governed by the physics from inflation to the surface of last scattering, providing us with a powerful probe of the physics of the early universe. If these fluctuations are Gaussian, homogeneous and isotropic as indeed they seem to be, then they can be completely described by a covariance matrix which is block-diagonal in spherical harmonic basis¹:

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

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

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

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

are given by [11]

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

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

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

For a given cosmological model one can calculate the predicted power spectrum in terms of that model's parameters², and observational bounds on the power spectrum can therefore be used to constrain the cosmological parameters, or exclude the model altogether. Our primary interest in the sky maps is therefore to use them to estimate the power spectrum and ultimately the cosmological parameters.

5.1 Power spectrum estimation

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

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

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

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

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

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

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

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

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

5.1.1 Brute force evaluation

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

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

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

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

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

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

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

Binning

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

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

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

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

That is, each bin is the average of the flattened⁴ spectrum in multipole ranges, with the different spectra and cross-spectra mapping on to different bins (i.e. not being averaged together). Expanding back from bins to multipoles can be done with the operator

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

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

Direct evaluation of cosmological parameters

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

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

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

Practicality

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

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

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

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

5.1.2 Pseudo- C_l estimation

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

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

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

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

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

only Monte Carlo simulations respectively⁵ the unbiased power spectrum can be easily recovered by solving for C in equation (5.16):

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

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

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

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

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

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

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

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

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

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

5.1.3 Gibbs sampling

It is possible to have both the optimality and exactness of the brute force estimator while avoiding its $\mathcal{O}(n^3)$ complexity. As in the PCL approach above, the key to efficiency is to find a way to use the fast full-sky expressions in equations (5.3-5.5 and 5.7), as these benefit from the speed of $\mathcal{O}(N_{\text{pix}}^{\frac{3}{2}})$ full-sky spherical harmonics transformations. The PCL approach to this is to fill in the missing parts of the sky with zeroes⁶, a choice which makes it necessary to bias-correct the spectrum later. Gibbs sampling, on the other hand, makes it possible to use the full-sky expressions by producing full-sky samples of what the noiseless CMB could look like, which are then amenable to efficient spherical harmonics transformations [51, 52, 53].

Mathematically, this is based on the observation that

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

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

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

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

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

through Conjugate Gradients iteration [51, 32]. ω_1 and ω_2 are here vectors of independent samples from the standard normal distribution $N(0, 1)$. Conceptually, this amounts to Wiener-filtering m , and then adding enough fluctuations to compensate for the bias in power introduced by the wiener filter, resulting in a map with covariance S as required.

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

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

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

Parameter estimation

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

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

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

The result is a set of samples $\{\theta, s\}$ that directly provide the probability of the parameters $P(\theta|\text{data})$. As C is a function of θ , the first step is unchanged from above. The second step differs in that we do not have a simple expression for $\theta(C)$, preventing us from simply drawing a C from the distribution in equation (5.7), and translating it into a set of parameters. But we can combine the probability density function of said Wishart distribution with a short MCMC chain in θ to still sample $\theta \leftarrow P(\theta|s)$ at relatively low cost⁷.

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

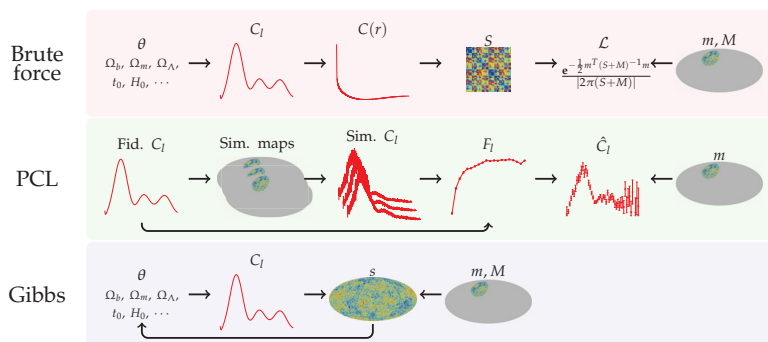


Figure 5.1: A sketch of the three power spectrum estimation methods described in this chapter. *Top*: Brute force maximum likelihood estimation. The likelihood of a set of cosmological parameters is computed by translating them into a signal covariance matrix S via the power spectrum C_l and the two-point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

5.2 Application to QUIET

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

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

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

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

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

where we have assumed for simplicity that tensor- and scalar perturbations have the same primordial spectral behavior. C_l^{scalar} is here the power spectrum in the absence of tensor modes, and C_l^{tensor} is the tensor-only spectrum for the case $r = 1$, as provided by CAMB [40]. If we know C_l^{scalar} , we can calculate C_l for each r and evaluate its likelihood using equation (5.10). But in practice, C_l^{scalar} is among the quantities we are trying to estimate, and must therefore be marginalized over to reflect our ignorance of it. In practice, this would involve an integral in as many dimensions as there are bins in our power spectrum, which is unfeasible with brute force evaluation⁸.

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




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




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

We therefore approximate the marginal with a partial conditional:

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

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

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

5.3 Results

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

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

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

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

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

5.3.1 Temperature spectra

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

5.3.2 Foregrounds

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

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

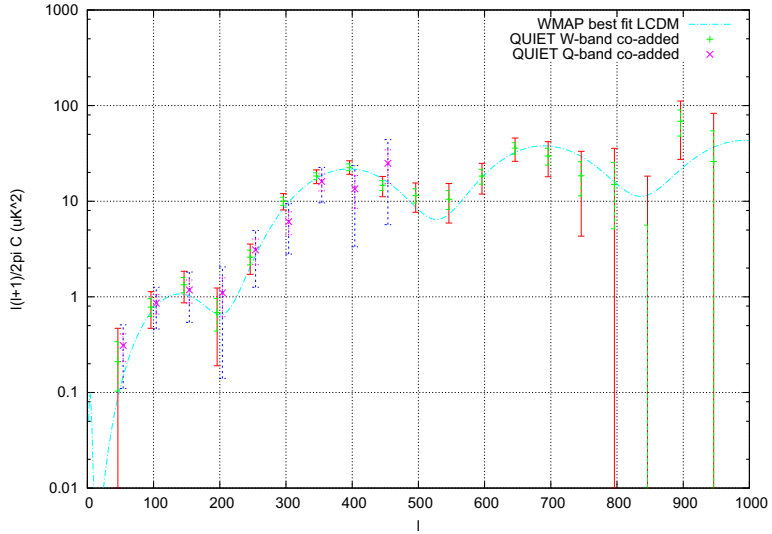


Figure 5.4: The QUIET EE power spectrum based on data from all four CMB patches. The W-band results are indicated using “+”-symbols with 1 sigma (green) and 2 sigma (red) error bars, while the Q-band results use “x”-symbols with 1 sigma (purple) and 2 sigma (blue) statistical error bars. These are compared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




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




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




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




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




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




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




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

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

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

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

5.3.3 Systematic errors

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

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

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

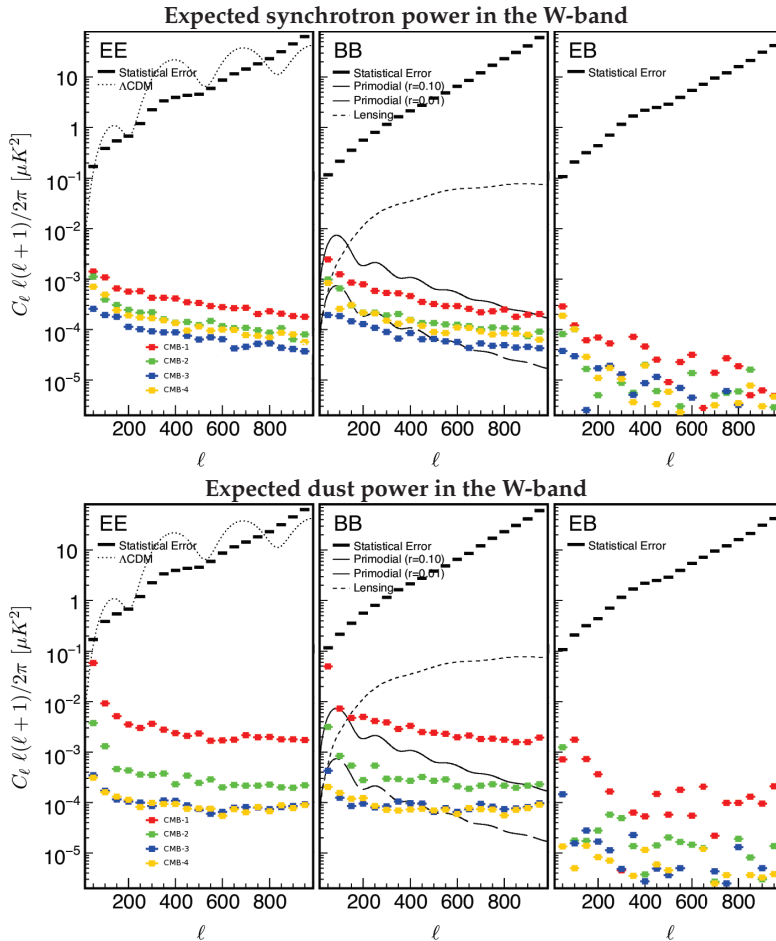


Figure 5.12: The expected power spectrum contamination from synchrotron (top) and dust (bottom) in the W-band, compared to QUIET’s statistical error and the Λ CDM expectation for EE modes, lensing B-modes and primordial B-modes for $r = 0.1$ and $r = 0.01$, based on the Planck foreground model v. 1.7.4.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

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

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

5.3.4 Parameters

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

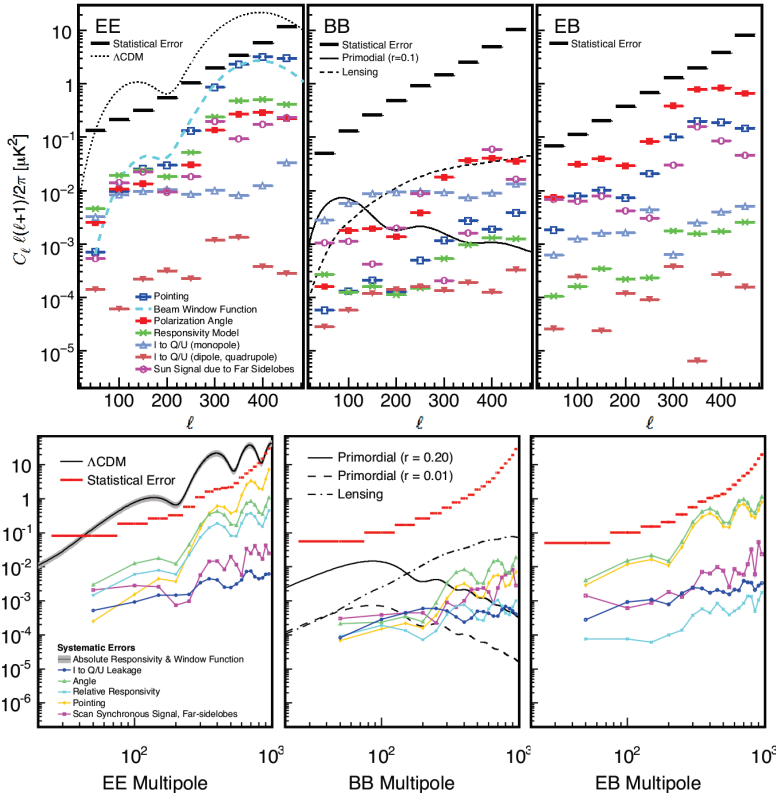


Figure 5.13: The Q-band (top) and W-band (bottom, preliminary) systematic effects compared to the statistical error and the expected signal. With the exception of the absolute gain, the systematic errors are all significantly smaller than the statistical errors. The absolute gain uncertainty is relatively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




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

Chapter 6

Non-gaussianity

The power spectrum and parameter estimation in the previous chapter is all based on the assumption that the CMB is a Gaussian random field. Theoretically, the primordial perturbations derive from quantum fluctuations during inflation, and these are expected to be in the ground state of the harmonic oscillator, which is a Gaussian [58, 1]. The subsequent evolution towards the surface of last scattering preserves this to linear order in perturbation theory, which is still a good approximation, resulting in a CMB which should be very close to Gaussian¹.

However, this conclusion can change based on the model of inflation and the importance of second and higher order effects during the subsequent evolution of the universe, neither of which are settled issues [60, 61]. Testing the Gaussianity of the CMB provides an indirect way of probing these issues, and has become an active field of research [62, 63, 59]. As there are infinitely many ways of being non-Gaussian, a large set of tests have been suggested. These can be divided into targeted tests, which test for a specific kind of non-Gaussianity², and general tests, which can find any kind of non-Gaussianity, at the cost of being less sensitive. Examples of the latter are the Kolmogorov-Smirnov test, the Cramér-von Mises statistic and the Anderson-Darling test³, which all work by comparing the empirical cumulative distribution function with the theoretically expected one.

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

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

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

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




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

random”⁴. This analysis later became the basis of a highly publicised claim for evidence for the Conformal Cyclic Cosmology hypothesis [66], which in turn prompted three independent follow-up studies [67, 68, 69]. Neither of these follow-ups found any significant detection. Subsequent correspondence⁵ made it clear that the reason for the difference in significance was due to Gurzadyan and Penrose [66] using simulations with much lower variance than the standard model due to the claim in Gurzadyan et al. [65] that the CMB is only “20% random”, a claim that had gone unnoticed by most of the community, but which would be revolutionary if true.

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

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

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

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

finitely many samples be distributed according to [64]:

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

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

This relation can then be used to test the hypothesis that a set of samples were drawn from a given distribution. However, the pixels s_i within the 1.5° CMB disks do not fulfil the criterion of being independently identically distributed due to the CMB anisotropies' correlatedness⁶, which are described in pixel space by the covariance $S_{ij} + N_{ij}$ as provided by equation (5.8) and the noise properties of the experiment. But it is possible to *whiten* s and thus make it eligible for the test: The vector $s' \equiv S^{-\frac{1}{2}}s$ will follow the distribution $N(0, 1)$, and thus be independently identically distributed if $s \sim N(0, S + N)$.

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

⁶Ignoring these correlations amounts to assuming a flat angular power spectrum.

Chapter 7

Summary and outlook

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

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

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

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

technology as QUIET-2.

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

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

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

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

¹Other upcoming experiments will also improve our knowledge of foregrounds:

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

Chapter 8

Bibliography

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

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

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

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

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

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

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

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

Part II
Papers

Paper I

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

QUIET Collaboration

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

Paper II

The Q/U Imaging Experiment
QUIET Collaboration
ApJ (to be submitted)

Paper III

Bayesian noise estimation for non-ideal CMB experiments

I. K. Wehus, S. K. Næss, H. K. Eriksen

ApJS, Volume 199, Issue 1, article id. 15 (2012)

Paper IV

Application of the Kolmogorov-Smirnov test to CMB data: Is the universe really weakly random?

S. K. Næss

A&A, Volume 538, pp. A17 (2012)

Application of the Kolmogorov-Smirnov test to CMB data: Is the universe really weakly random?

Sigurd K. Næss

Institute of theoretical Astrophysics, University of Oslo, P.O.Box 1029 Blindern, 0315 Oslo, Norway

Preprint online version: May 26, 2011

Abstract

A recent application of the Kolmogorov-Smirnov test to the WMAP 7 year W-band maps claims evidence that the CMB is “weakly random”, and that only 20% of the signal can be explained as a random Gaussian field. I here repeat this analysis, and in contrast to the original result find no evidence for deviation from the standard Λ CDM model. Instead, the results of the original analysis are consistent with not properly taking into account the correlations of the Λ CDM power spectrum.

Key words. cosmic background radiation

1. Introduction

In astronomical data analysis, it is often useful to be able to test whether a set of data points follows a given distribution or not. For example, many analysis techniques depend on instrument noise being Gaussian, and to avoid bias, one must check that this actually is the case. There are many different ways in which two distributions can differ, and correspondingly many different ways to test them for equality. The simplest ones, such as comparing the means or variances of the distributions, suffer from the problem that there are many ways in which distributions can differ that they cannot detect no matter how many samples are available. For example, samples from a uniform distribution can easily pass as Gaussian if one only considers the mean and variance.

The popular Kolmogorov-Smirnov test (K-S test) resolves this problem by considering the cumulative distribution functions (CDF) instead: Construct the empirical CDF of the data points and find its maximum absolute difference K from the theoretical CDF. Due to the limited number of samples, the empirical CDF will be noisy, and K will therefore be a random variable with its own CDF, which in the limit where the number of samples goes to infinity is given by

$$P(x < K) = F_{KS}(\sqrt{N_{\text{obs}}}K) \quad (1)$$

with

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

In contrast with the simplest tests, this test can detect any deviation in the distributions, but may require a large number of samples to do so, especially in the tails of the distribution.

Recently, a series of papers (Gurzadyan et al. 2011; Gurzadyan & Kocharyan 2008; Gurzadyan et al. 2010)

has applied this test to WMAP’s cosmic microwave background (CMB) maps, resulting in the remarkable claim that the CMB is “weakly random”, with only 20% of the CMB signal behaving as one would expect from a random Gaussian field. This result went on to be used in a much discussed series of papers (Gurzadyan & Penrose 2010a,b, 2011) claiming a strong detection of concentric low-variance circles in the CMB, which was taken as evidence for Conformal Cyclic Cosmology. Other groups failed to significantly detect the circles (Wehus & Eriksen 2010; Moss et al. 2011; Hajian 2010). The difference in significance was due different CMB models: Wehus & Eriksen (2010); Moss et al. (2011); Hajian (2010) used realizations of the best-fit Λ CDM power spectrum, while Gurzadyan & Penrose (2010b) used a “weakly random” CMB model.

Both in order to resolve this issue, and because a weakly random universe would be a strong blow against the Λ CDM model in its own right, it is important to test this result.

2. Method

Before applying the K-S test, one must be aware of its limited area of validity: Equation (1) requires an infinite number of independently identically distributed samples, while CMB maps actually consist of a limited number of correlated samples. However, both the correlations and number of samples can be compensated for, as we shall see.

2.1. Application of the K-S test to correlated data

Though the K-S test is not immediately applicable to a correlated data set, it is possible to perform an equivalent test on a transformed set of samples. The question we are trying to answer with the K-S test is “Do the samples follow the theoretical distribution?”. The truth or falseness of this is preserved if we apply the same transformation to both the samples and the distribution we test them against, and to be able to use the K-S test, the logical transformation to use is a *whitening* transformation, which results in an




Figure 1. Λ CDM two point correlation function after applying the WMAP W-band beam and the HEALPix (Górski et al. 2005) nside 512 pixel window.

independent, identical distribution for the samples. With original samples \mathbf{d} with covariance matrix \mathbf{C} , the whitened (uncorrelated with unit variance) samples \mathbf{r} are given by:

$$\mathbf{r} = \mathbf{C}^{-\frac{1}{2}} \mathbf{d} \quad (3)$$

Thus, to test whether the data points $\mathbf{d} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{C})$, we can test the equivalent hypothesis $\mathbf{r} \leftarrow \mathbf{N}(\mathbf{0}, \mathbf{1})$.

In the case of CMB maps, both the data itself and the noise is expected to be Gaussian, so the obvious theoretical distribution here is $N(0, \mathbf{S} + \mathbf{N})$, where the CMB signal covariance matrix \mathbf{S} is given by the two-point correlation function:

$$S_{ij} = \sum_{l=0}^{\infty} \sqrt{\frac{2l+1}{4\pi}} C_l B_l P_l(\cos(|\mathbf{p}_i - \mathbf{p}_j|)) \quad (4)$$

$P_l(x)$ are the Legendre polynomials normalized to $\frac{1}{2\pi}$, and \mathbf{p}_i and \mathbf{p}_j are the direction vectors for pixel i and j in the disk. C_l is the Λ CDM angular power spectrum, while B_l accounts for the beam and pixel window. \mathbf{N} is instrument dependent, but for the WMAP W-band CMB map we will use here, the noise is nearly diagonal, and given by the corresponding W-band RMS map.

2.2. Application of the K-S test with few samples

The other problem we need to account for is our finite number of samples. In this case equation (1) is only approximate. For most uses of the test, this approximation is good enough, especially when employing analytical expressions for improving the quality of the approximation for low numbers of samples (von Mises 1964). For example, when performing a single test to accept or reject a test distribution, a bias of a few percent in the confidence with which the hypothesis is rejected is not important.

However, when making statistics for a large number of such test results, such a bias may make the results ambiguous. Given a set of experiments with a corresponding set of maximum deviations $\{K_i\}$, the corresponding probabilities $\{p_i = P(x < K_i)\}$ should be uniformly distributed if the




Figure 2. When applying the K-S test to samples known to come from the correct distribution, the resulting values $\{p_i = P(x < K_i)\}$ should be uniformly distributed, but when working with a limited number of samples, the Kolmogorov distribution is only approximate, and the actual CDF of the results, $G(p)$, differs from the ideal $G_{\infty}(p) = p$. This is shown in the upper panel for the case of 540 samples per experiment, where $G(p)$ is the solid line and $G_{\infty}(p)$ is dashed. The lower panel shows the deviation between the two, which is of the order of 1% in this case (but larger with fewer samples).

samples actually follow the theoretical distribution, and a histogram of $\{p_i\}$ should therefore be flat. Deviations from this indicate that the theoretical distribution does not accurately describe the samples. However, the approximate equation (1) also introduces a small non-uniformity in $\{p_i\}$ even if the samples actually do follow the distribution. To avoid the ambiguity this causes, we will instead compute a numerical correction function mapping the approximate p to the true p' .¹

To build up the correction, we simulate a large number² of experiments, each with the same number of samples as the actual data set, but drawn directly from the theoretical distribution. Thus, for these, $\{p_i\}$ should be uniform, with a CDF of $G_{\infty}(p) = p$. However, since equation (1) is inexact, for small numbers of samples, the actual CDF is $G(p) \neq G_{\infty}(p)$. The mapping between the approximate p and true p' is given by $G(p) = G_{\infty}(p') \Rightarrow p' = G_{\infty}^{-1}(G(p)) = G(p)$. Thus, for a limited number of samples

$$P(x < K) = G(F_{\text{KS}}(\sqrt{N_{\text{obs}}}K)). \quad (5)$$

Figure 2 illustrates the correction function for $5 \cdot 10^6$ simulations of 540 each. For this many samples, the correction is only of the order of 1%.

¹ What we do here is essentially replacing the analytical Komolgorov distribution (equation (1)) with a numerical distribution. This could also be done without using the analytical distribution as a basis, at a small cost in clarity.

² The number necessary depends on the level of accuracy desired. The noise in the estimate of $G(p)$ propagates to the final results. To make this a subdominant noise contribution, the number of simulations should be at least as large as the number of actual experiments, preferably much higher.




Figure 3. A randomly selected disk before (left) and after (right) the whitening operation. The samples are strongly correlated and thus unsuitable for the K-S test before the transformation, but afterwards no correlations are visible and the variance is 1. Note that whitening the data does not mean that we are “forcing” the K-S test to pass. The whitened data will only end up matching $N(0, 1)$ after whitening if they followed our theoretical distribution $N(0, C)$ before.

3. Does Λ CDM fail the K-S test?

With this in hand we can finally apply the K-S test on CMB data. Following Gurzadyan et al. (2011), we randomly pick 10 000 disks with a radius of 1.5 degrees from the WMAP 7 year W-band map (Jarosik et al. 2011), with the region within 30 degrees from the galactic equator excluded. Each disk contains on average 540 pixels, which are whitened using equation (3). A typical disk before and after the whitening operation can be seen in Fig. 3. After whitening, the values should follow the distribution $N(0, 1)$ if our model is correct.

The histogram of resulting probabilities $\{p_i = P(x < K_i)\}$ from of applying equation (5) to the hypothesis $\mathbf{r} \leftarrow N(\mathbf{0}, \mathbf{1})$ is shown in Fig. 4, together with the 68% and 95% intervals from 300 simulations. The data and simulations are consistent, and follow a uniform distribution as expected³: The CMB map is fully consistent with Λ CDM + WMAP noise as far as the K-S test is concerned.

This is dramatically different from the curve found by Gurzadyan et al. (2011), which was strongly biased towards *low* values. Low values of $P(x < K)$ would mean that the empirical CDF of the samples matches the theoretical one *too well*, i.e. even better than samples drawn directly from the theoretical distribution.

What could cause Gurzadyan et al. to get results so different from ours? One way biasing $P(x < K)$ low is by basing the parameters of your test distribution on the values themselves. However, even without doing this, it is possible to get low values if the values used in the K-S test are *correlated*. This is also consistent with the presentation given by Gurzadyan et al. (2011) who apparently applied the K-S test directly to the raw samples \mathbf{d} , or equivalently, that they model the pixel values as coming from a 1-dimensional dis-

³ It should be noted that the histogram bins are not completely independent for two reasons: Firstly, some disks are going to overlap, meaning that the same samples enter into several different K-S tests, and secondly, while our transformation has made the samples within each disk independent, the correlation between different disks is still present.




Figure 4. Histogram of results of the K-S test. Each panel compares the results from properly taking the correlations into account (solid line) with those one gets from ignoring them (dashed line), together with 68% and 95% intervals (dotted lines) from simulations. The upper panel corresponds to using samples further than 30 degrees away from the galactic equator, while the lower panel instead uses the WMAP KQ85 analysis mask. In both cases, both the map and the simulations pass the K-S test when taking the correlations into account, while if the are ignored, the K-S test fails in the same way Gurzadyan et al. (2011) reported.

tribution. To check this, I repeated the analysis, this time using the theoretical distribution $d \leftarrow N(\mu, \sigma^2)$, where μ and σ^2 are the measured mean and variance of the samples in the disk. The result is also shown in Fig. 4. This time, the bias towards low values is clearly recreated.

It therefore seems likely that Gurzadyan et al.’s reported “weak randomness” is the result of not properly taking the CMB’s correlations into account. One is, of course, free to use whatever distribution one wants as the theoretical distribution in a K-S test, even a model where the CMB pixels are independently identically distributed, with no correlations at all. The problem lies in the interpretation of the test results. For Gurzadyan et al. (2011), the K-S test results are clearly not uniform, indicating that the chosen theoretical distribution has been disproven. However, Gurzadyan et al. then go on to create a set of simulations (linear combinations of 20% Gaussian and 80% static signal) that fail the test in the same way as the WMAP map does. But having two sets of samples fail the K-S test the

same way does not prove that they have the same properties. It simply means that the chosen test distribution was a poor choice.

4. Kolmogorov maps

While Gurzadyan et al.'s Kolmogorov statistics are biased by not taking the correlations into account, the approach of making sky maps of K-S test results introduced in Gurzadyan et al. (2009) is still an interesting way to search for regions of the sky that do not follow the expected distribution. Making an unbiased Kolmogorov map straightforwardly follows the procedure in Sect. 3, with the main difference being the selection of pixels. Instead of randomly selecting disks, we now systematically go through nside 16 pixels, using the 1024 nside 512 subpixels inside each one as the samples. These are then tested against $N(0, C)$ by whitening them via equation (3) and then comparing the whitened samples to $N(0, 1)$.

The result is the nside 16 map of $P(x < K)$ shown in Fig. 5. Regions that pass the test have a value uniformly distributed between 0 and 1, and we see that this applies to the CMB-dominated areas of the sky, while areas dominated by the galaxy fail the test as expected.

For comparison, Fig. 5 also includes the result of making the same map while ignoring correlations. In this case, the whole sky fails the test: The CMB-dominated areas are biased low, while the galaxy is biased high. This map is similar to the map in Gurzadyan et al. (2009), which is also too low outside the galaxy, and too high inside, which is, again, consistent with Gurzadyan et al. applying the K-S test directly to the raw samples.

5. Summary

The Kolmogorov-Smirnov test is a useful and general way of testing whether a data set follows a given distribution or not. However, it only applies to independently identically distributed samples. The CMB is strongly correlated, and thus not immediately compatible with the test. However, this can be resolved by the application of a whitening transformation, replacing the hypothesis $\mathbf{d} \leftarrow N(\mathbf{0}, C)$ with the equivalent $C^{-\frac{1}{2}}\mathbf{d} \leftarrow N(\mathbf{0}, 1)$. With this, we find that the Λ CDM passes the K-S test. This is incompatible with the original analysis by Gurzadyan et al. (2011), which claimed detection of an unknown non-random component making up 80% of the CMB based on the CMB failing the K-S test there. It turns out that this analysis did not take the CMB correlations into account, which we confirm by producing the same failure of the K-S test when we skip the whitening step. When the correlations are handled properly, there is no need for a weakly random universe.

Acknowledgements. The author would like to thank Hans Kristian Eriksen for useful discussion and comments.

References

- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, arXiv:astro-ph/0409513
 Gurzadyan, V. G., Allahverdyan, A. E., Ghahramanyan, T., et al. 2009, A&A, 497, 343, 0811.2732
 Gurzadyan, V. G., Allahverdyan, A. E., Ghahramanyan, T., et al. 2011, A&A, 525, L7+, 1011.3711




Figure 5. nside 16 map of $P(x < K)$ based on pixels from the WMAP 7 year nside 512 W-band map. This time, the K-S test is performed on the 1024 nside 512 pixels inside each nside 16 pixel instead of on disks. The upper panel uses $N(0, \mathbf{S} + \mathbf{N})$ as the theoretical distribution (by testing the whitened data against $N(0, 1)$), while the lower panel ignores the correlations, instead using $N(\mu, \sigma^2)$, where μ and σ are the measured mean and standard deviation of the samples. The former passes the test outside the galactic plane, while the latter fails everywhere, being biased low outside the galaxy.

- Gurzadyan, V. G., Kashin, A. L., Khachatryan, H. G., et al. 2010, Europhysics Letters, 91, 19001, 1002.2155
 Gurzadyan, V. G. & Kocharyan, A. A. 2008, A&A, 492, L33, 0810.3289
 Gurzadyan, V. G. & Penrose, R. 2010a, ArXiv e-prints, 1011.3706
 Gurzadyan, V. G. & Penrose, R. 2010b, ArXiv e-prints, 1012.1486
 Gurzadyan, V. G. & Penrose, R. 2011, ArXiv e-prints, 1104.5675
 Hajian, A. 2010, ArXiv e-prints, 1012.1656
 Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, ApJS, 192, 14, 1001.4744
 Moss, A., Scott, D., & Zibin, J. P. 2011, JCAP, 4, 33, 1012.1305
 von Mises, R. 1964, Mathematical Theory of Probability and Statistics, ed. von Mises, R.
 Wehus, I. K. & Eriksen, H. K. 2010, ArXiv e-prints, 1012.1268