Searching for B-modes with QUIET: A reanalysis of the first-season Q-band data

Tone Melvær Ruud

Thesis submitted for the degree of
Master of Science in Astronomy

Institute of Theoretical Astrophysics
University of Oslo

June, 2012

Abstract

The Q/U-Imaging ExperimenT (QUIET) is a ground-based coherent polarimeter experiment, primarily targeting the detection of low-ℓ B-mode polarization in the CMB. Such modes are a signature of primordial tensor perturbations, also known as gravitational waves. The existence of tensor perturbations is as yet unproven, but uniquely predicted by theories of cosmological inflation. The B-modes are thus frequently referred to as "the smoking gun of inflation", and efforts towards their detection are many and fierce. The first season of QUIET observations, split between the Q- and W- frequency bands, was carried out in 2008-2010. The analysis of the Q-band data has been completed, and was published in late 2010.

QUIET uses two independent data analysis pipelines. The *Maximum Likelihood* (M-L) pipeline is written and maintained by the QUIET participants at the University of Oslo. QUIET data analysis using the M-L pipeline is the topic for this thesis. I attempt to give a broad and comprehensible account of the process of data analysis for a non-ideal project in radio cosmology, investigating each step in the transformation from raw observational data to final estimates of power spectra and cosmological parameters. Following this narrative, the pipeline is applied in a reanalysis of the QUIET Q-band data set.

Since the completion of the original Q-band analysis, significant changes have been made to the M-L pipeline. A reanalysis of the Q-band data serves as a validation of the new features. We find signs of improved performance in some places, while at other points there is need of further assessment. Most importantly, we find signs of a bias in the polarization null test suite, which is thought to be caused by ground pickup, although this will have to be confirmed.

The reanalysis also has significant educational merit, as achieving a solid understanding of the data analysis process is among my main motives for undertaking this project. The modest size of the Q-band data set makes it well suited for demonstration.

Acknowledgements

So long, and thanks for all the fish.

As this thesis is all but finished, I enjoy feelings of pride and relief, sprinkled with nostalgia and seasoned with the slightest touch of disbelief. These pages, into which a quite ridiculous amount of work has been laid, mark the conclusion of 19 consecutive years of education. The opening words, courtesy of Douglas Adams, are my salute to all those who have helped me through these years.

A select few do, of course, deserve mention. Firstly, I thank my excellent supervisor, Hans Kristian Eriksen, for guidance throughout the year, for showing great faith in my abilities, and for giving me the opportunity of joining the QUIET team. My thanks also to the rest of the local QUIET group, Ingunn K. Wehus and Sigurd K. Næss, for patiently helping me whenever I got stuck - a special thank you to Sigurd for acting as my personal help-desk for all technical matters. I count myself very lucky indeed to have been surrounded by such approachable and enthusiastic people.

My appreciations to that wandering motivational speech, Nicolaas Groeneboom, whose words of praise for Hans Kristian helped me make my decision when I was stuck in choosing between supervisors. I have yet to regret that choice!

I would like to thank all the students at ITA for great companionship. You have all made life in the basement worth remembering.

A most special thank you goes to my three fabulous friends, Thea Marcelia Sletten, Veronica Øverbye and Sarah Thaysen, who have been my faithful companions through six years at the University. It has been long hours, hard work, and much frustration, but you have made it worth the while. Looking back, I see the laughter, the long and winding discussions, the joy of getting to the finish line together, the solving world problems over a cup of coffee. You are such strong and inspiring creatures, and I am proud and lucky to have worked with you. (What *is* the radius of the Universe, anyway?)

Next, I would like to express my gratitude and love for my family. I thank my parents for teaching me to be curious, and to look for answers. I thank my wonderful, loving sister Marit, for her endless support and very welcome distractions. In particular, I thank my father for always supporting my highly impractical interest in astronomy, including accompanying me to hilltops at 5 o'clock in the morning to catch a glimpse of a cloud-ridden partial eclipse. Having inherited certain patterns of sleep, I know that cannot have been easy.

Finally, to the love of my life, Christer Tronsmed: Thank you for making it all so easy. Thank you for making me strong, for believing in me, for keeping me fed, and for making every day a good one. Thank you for making me feel safe, and for being so endlessly patient. I will never cease to be amazed at my luck in finding you.

Tone M. Ruud, Oslo, 1st June 2012

Contents

List of Figures

Part I

- The past is a foreign country -

Introduction

Chapter 1

Cosmology at a glance

> Cosmology, *n.*
>
> 1. *The science or theory of the universe as an ordered whole, and of the general laws which govern it. Also, a particular account or system of the universe and its laws.*
> 2. Philos. *That branch of metaphysics which deals with the idea of the world as a totality of all phenomena in space and time.*
>
> *Origin mid 17th cent.: from French* ***cosmologie*** *or modern Latin* ***cosmologia****, from Greek* ***kosmos*** *'order' or 'world' +* ***-logia*** *'discourse'*
>
> *The Oxford English Dictionary*

In the widest sense of the word, one could probably state with reasonable confidence that cosmology has existed for as long as humanity itself. Looking at the world around us and asking questions about it - what is it made of, how does it work, how did it get there in the first place - is, as far as we can tell, one of the things that distinguishes humans from the host of other creatures that share our Earth. Our ability to gather knowledge about our surroundings, store it, and share it, has helped us to an understanding of the world that far surpasses what a single mind can achieve within the span of a lifetime. We are all standing on shoulders - not of giants, but of a succession of more or less ordinary people who all contributed a small amount to the mountain of human knowledge. The view from the summit, for those who care to look, is quite marvellous.

However, despite thousands of years' worth of astronomers' diligent sky studies, our knowledge of the Universe barely extended beyond our own galaxy by the beginning of the 20th century. Without the aid of modern technology, science had been powerless to answer the questions of cosmology, leaving the field open for the caperings of philosophers and theologians. Such matters are of little concern to us, however valuable they might be. In this thesis, we will restrict ourselves to the science that is modern cosmology.

When working on a master's thesis - or as a professional scientist, for that matter - it is easy to get sucked into the details, delving ever deeper into some small matter, until one is all but engulfed. That is as it should be. The willingness to follow a trail to its end is what keeps scientific progress going, even on areas where we think we have figured out the more important stuff. But once in a while, we should all remember to step back and look at the big picture: Marvel at the beauty of it, and appreciate the great ideas that formed it, remembering those who contributed. In this chapter and

the next, I will sketch the main lines of modern cosmology, to form the backdrop for the more serious business awaiting us. We will begin by studying the physical theories, and move on in chapter 2 to recapitulate the history of this largest of all the branches of science.

1.1 The Standard Model

Today, cosmologists have a fairly good idea of how the Universe works. Formulating a successful cosmological model is no small feat, as it should be able to satisfactorily predict and explain all the large-scale properties of the Universe. The current standard model of cosmology, known, for reasons that will soon become clear, as ΛCDM, consists almost entirely of physics discovered within the last 100 years, and still it has managed to give us a consistent picture of how our Universe works. In the following, I will make a tour of the key ingredients and theoretical foundations of the cosmological Standard Model. For introductory cosmology, there are plenty of good books to choose from, although the pace of the progress makes sure they become outdated rather fast. As the source of this chapter, I have chosen S. Weinberg's *Cosmology* [1] and S. Dodelson's *Modern Cosmology* [2].

1.1.1 General relativity, the expanding Universe and the cosmological principle

Spacetime is the term we apply for the fabric of space and time - the intangible "something" in which everything happens. Albert Einstein introduced this concept in his special theory of relativity, negating the previously dominant notion that space and time are independent and static. To describe spacetime, we use a *metric*, generally in the form of a tensor $g_{\mu\nu}$. The metric tensor uniquely describes any space, since it contains the recipe for measuring distances in it.

We define an invariant quantity, the *line element*, by which we can move between different metrics. It is defined by

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \tag{1.1}$$

where x^μ are generalised coordinates. (Note the use of Einstein's summation convention in this expression - there is an implicit summation over equal indices.) On this most general form, the line element is viable independent of spacetime geometry.

Einsteins *general* theory of relativity, or GR for short, forms the basis upon which many cosmological models are built, including the current standard model. In Einstein's vision, gravity is not a force, but rather a manifestation of the shape of spacetime. When a particle is affected by what Newtonian theory refers to as the force of gravity, Einstein maintains that it is, in fact, falling freely along a *geodesic curve* - a straight line in curved spacetime. Furthermore, he professed that it is the presence of matter that causes spacetime to curve in the first place. In the words of U.S. physicist John Wheeler: "Matter tells space how to curve, and space tells matter how to move". Mathematically, this relationship is conveyed via the *Einstein equations*,

$$E_{\mu\nu} = 8\pi G T_{\mu\nu}. \tag{1.2}$$

This is a tensor equation. The Einstein tensor, $E_{\mu\nu}$, is a function of the metric and its derivatives, thus depending solely on the geometry of spacetime. The energy-momentum tensor tells of the presence and movement of matter and energy in said spacetime. As we habitually restrict ourselves to four-dimensional spacetime, eq. 1.2 should hide a total of 16 differential equations. In practice there are 10, since the tensors in question are symmetric. These equations are our starting point for creating models of the Universe. The recipe is relatively simple: Choose a metric and matter/energy components, insert them into eq. 1.2, and solve.

Our metric of choice is most often the *Friedmann-Robertson-Walker* (FRW) metric, named after Friedmann, who first derived it, and Robertson and Walker, who proved it to be the only metric allowing for exact homogeneous and isotropic solutions of Einstein's equations. We justify our choice with the *cosmological principle*. This is the idea that there should be nothing special about Earth's position in the Universe. In technical terms, this translates to the Universe being homogeneous and isotropic. Homogeneity implies that the Universe should have the same basic properties at all points, while isotropy means that we should see the same basic properties in every direction. Other possibilities may be concieved, but are often cumbersome to work with, and so far, the homogeneous and isotropic solutions have done the trick. However, theorists should always keep in mind that the cosmological principle is an assumption, and could be wrong.

In spherical coordinates, the line element in the FRW metric is

$$ds^2 = -c^2dt^2 + a(t)^2\left(\frac{dr^2}{\sqrt{1-kr^2}} + r^2\left(d\theta^2 + \sin^2\theta d\phi^2\right)\right), \tag{1.3}$$

where the parameter k allows for open, closed or flat solutions. Inserting $k = 0$, we recognise ordinary Euclidean space within the parentheses. The interesting term is the *scale factor*, $a(t)$. This time-dependent quantity gives the metric the possibility of evolving, so that the fabric of spacetime could be stretching or contracting as we speak, causing distances between stationary objects to change with time. As it happens, this is exactly what we observe. By measuring the redshifts of other galaxies we can find their speed relative to ours, and from this we have found that the Universe expands. For small redshifts, it is seen that the velocity at which a galaxy recedes from us depends linearly on distance, as given by the famous Hubble's law,

$$v = H_0 d. \tag{1.4}$$

H_0 is the present-day Hubble constant, the current value of which is approximately $70\,\mathrm{km/s/Mpc}$. It is often given in terms of the reduced Hubble constant h, defined by $H_0 = 100h\,\mathrm{km/s/Mpc}$.

Because of this universal expansion which, according to our best theories, has trundled along since the very beginnings of our Universe, both redshift and the scale factor may be used as measures of time. These quantities both fulfil the requirements of any time variable, which is to be strictly increasing and uniquely determining any point in the past, present or future. Due to their limited numerical range, they are often far more practical than using cosmic time, i.e. "seconds since the Big Bang". The cosmic redshift time variable is used in the sense of "the amount a photon has been redshifted when arriving at Earth after having travelled through expanding spacetime since the

time identified by the scale factor $a(t)$". It is measured by the parameter z,

$$1 + z(t) = \frac{\lambda_{recieved}}{\lambda_{emitted}} = \frac{a_0}{a(t)}. \tag{1.5}$$

Here, and throughout this chapter, subscript 0 means the present-day value of a quantity. The present-day scale factor a_0 is usually set to 1.

1.1.2 A cosmic inventory

When calculating the most basic properties of the Universe, we allow ourselves the simplifying measure of assuming perfect homogeneity and isotropy, overlooking minor details such as clusters of galaxies. In the *really* big picture, even the largest structures we know of become negligible, and we may look at the Universe as a perfect, isotropic fluid. Under this assumption, the Einstein and energy-momentum tensors become diagonal, and the spatial components are interchangeable, so the set of 10 equations reduce to 2. These two are generally known as the *Friedmann equations*.

$$\text{1st Friedmann equation:} \qquad \dot{a}^2 + kc^2 = \frac{8\pi G}{3}\rho a^2. \tag{1.6}$$

$$\text{2nd Friedmann equation:} \qquad \frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right). \tag{1.7}$$

The first is found by solving the time-time component ($\mu = \nu = 0$) of eq. 1.2, while the second can for instance be deduced from the trace of the tensor equation, $E^\mu_\mu = 8\pi G T^\mu_\mu$. The dots signify the derivative with respect to cosmic time.

From these simple equations, we can find the expansion history of the smooth background Universe, depending on its energy content and curvature. The solutions we will find are the simplest cosmological models. Albeit not very realistic, they are a useful starting point, as they demonstrate how the various constituents of a universe affect observables such as its age and expansion velocity. But before we can solve the equation set, we need to adress the fact that we have two equations, but four unknowns, the scale factor $a(t)$, the curvature k, the energy density $\rho(t)$ and the pressure $p(t)$. The two latter are generalised, by which I mean that they apply to the density and pressure of any and all of the energy components present.

The curvature parameter has three distinct cases, positive, negative, or zero. Apart from that, its value is irrelevant, so we may assign it integer values, $k \in [-1, 0, 1]$. The cases may then be solved separately, and we have effectively reduced the number of unknowns by one. Finally, we know from classical thermodynamics that the density and pressure of a fluid are not independent of each other. For any fluid, there is a relation between the two, which we refer to as an *equation of state* (EOS). All the fluids that take part in the simplest cosmological models have state equations of the form

$$p = w\rho c^2.$$

The value of the parameter w distinguishes one fluid from another by giving the specific density-pressure relationship for each type of fluid. When we insert a specified function $p(\rho)$ into the Friedmann equations, we can find the evolution of the scale factor for a universe containing the fluid described by that particular equation of state.

Conservation of energy and momentum in an expanding universe demands that the *covariant* derivative[1] of the energy-momentum tensor must be zero, $T^{\mu}_{\nu;\mu} = 0$. Applied to a universe with perfect homogeneity and isotropy, this implies

$$\dot{\rho} = -3\frac{\dot{a}}{a}\left(\rho + \frac{p}{c^2}\right). \tag{1.8}$$

(We note that the same relation may be found by demanding adiabatic expansion, or by differentiating eq. 1.6 and inserting it into eq. 1.7.) Eq. 1.8 may be integrated using our general equation of state, with the present time as boundary, to achieve the density as a function of the scale factor,

$$\rho = \rho_0 \left(\frac{a_0}{a}\right)^{3(1+w)}. \tag{1.9}$$

This equation demonstrates how the expansion of the Universe depends directly on its energy content.

Helpful quantities

Some concepts beg introduction before we move on. When working with expanding universes, distance becomes somewhat ambiguous. Physical distance is no longer constant, even when considering stationary objects. To alleviate the matter, we introduce comoving coordinates. These may be viewed as a coordinate system, or grid, that follows the universal expansion. Each galaxy, which we imagine for the moment is not moving with respect to spacetime, has a fixed position on this grid. This allows us to uniquely identify any point in the Universe, and follow it through history, as it moves away from us. The r, θ and ϕ from the line element are comoving coordinates. Based on this, we may use the line element to find a definition of physical, or "proper", distance, d_P. By proper distance, we mean the instantaneous length of the straightest possible path between two points. We find it by integrating the line element,

$$d_P(r,t) = \int |ds| = a(t)\int_0^r \frac{dr'}{\sqrt{1-kr'^2}}.$$

This expression clearly illustrates the nature of the scale factor as the ratio between a distance at some specified time, and the same distance, at any other time.

Once we have the proper distance of an object, we may also find its velocity,

$$v(r,t) = \frac{d}{dt}d_P = \dot{a}\int_0^r \frac{dr'}{\sqrt{1-kr'^2}} = \frac{\dot{a}}{a}d_P.$$

Hence we define the Hubble rate as $H(t) = \dot{a}/a$. This allows us to keep the original form of Hubble's law, as we first saw it in eq. 1.4, although it is now valid for all points in time.

We will find it useful to look at the special case of a universe with a spatially flat geometry, $k = 0$. If we apply this condition to eq. 1.6, we find

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho. \tag{1.10}$$

[1] A generalisation of the directional derivative applicable to curved spaces as well as to flat[3].

Into this, we insert the Hubble rate and solve for the density. The present-day value of this expression is what we call the *critical density*,

$$\rho_{cr} = \frac{3H_0^2}{8\pi G}. \tag{1.11}$$

It is the total energy density in the Universe today, if and only if its spatial geometry is, on average, Euclidean.

Assembling the model

Using the pieces we have gathered, we can assemble a simple model of our own Universe. We will need to include more than one type of energy. The Friedmann equations enable the description of multiple-component universes by way of the generalised pressure and density parameters. We simply allow $p = \sum_i p_i$ and $\rho = \sum_i \rho_i$, where each component i has its own equation of state. According to the Standard Model, the energy components present in the Universe are:

Baryonic matter All nuclei and electrons. Technically, electrons are leptons, not baryons, but the term has stuck. Refers to "cold", i.e. nonrelativistic, matter, exerting a negligible pressure on its surroundings. Often referred to as dust, luminous matter, or pressureless matter. EOS parameter $w = 0$.

Cold dark matter Pressureless matter, interacting solely by gravity. Referred to as CDM. EOS parameter $w = 0$.

Radiation Mostly photons, possibly other relativistic, massless particles. Has positive ("outward") pressure. EOS parameter $w = \frac{1}{3}$.

Dark energy Unknown form of energy, characterised by having negative pressure. EOS parameter $w < 0$. Often represented by Einstein's cosmological constant, Λ, which describes dark energy with constant energy density as function of time. In that case, $w = -1$. Forms of dark energy with variable energy densities may be concieved.

In many cases, we prefer to look at the relative densities rather than absolute ones. Consequently, we define for any energy component i, the relative density with respect to the critical, $\Omega_i = \rho_i/\rho_{cr}$. As for the generalised density, we here have $\Omega = \sum_i \Omega_i$. Now, if we divide eq. 1.10 by the critical density and consider the present-day value again, we find another useful form of the flatness condition, $\Omega_0 = \sum_i \Omega_{i\,0} = 1$.

To complete the picture, we conjure up a "curvature density", Ω_k,

$$\Omega_{k0} = -\frac{kc^2}{a_0^2 H_0^2}.$$

This does not make as much physical sense as the other Ω's, but it does enable us to write $\Omega_0 + \Omega_{k0} = 1$, which allows for any curvature[2].

[2]The choice of present day is essentially arbitrary from a non-anthropocentric point of view, and these expressions are equally valid for any other time. Under some circumstances it is convenient to re-define the Ω's and critical density to be time-dependent functions in their own right, but here we prefer to deal with the present-day values.

Best-estimate values of cosmological parameters for the ΛCDM model

Parameter	Value	Description
$\Omega_b h^2$	$0.02258^{+0.00057}_{-0.00056}$	Physical baryon density
$\Omega_c h^2$	0.1109 ± 0.0056	Physical dark matter density
Ω_Λ	0.734 ± 0.029	Dark energy density
n_s	0.963 ± 0.014	Scalar spectral index
$\Delta^2_R\,(k_0 = 0.002\mathrm{Mpc}^{-1})$	$(2.43 \pm 0.11) \times 10^{-9}$	Curvature fluctuation amplitude
τ	0.088 ± 0.015	Reionisation optical depth
Quantities derived from the above:		
t_0	13.75 ± 0.13 Gyr	Age of the Universe
H_0	71.0 ± 2.5 km/s/Mpc	Hubble constant
Ω_b	0.0449 ± 0.0028	Baryon density
Ω_c	0.222 ± 0.026	Dark matter density

Table 1.1: Best-fit estimates of the 6 parameters defining the ΛCDM model, as well as some useful derived quantities, found using the WMAP7 data. The model assumes flat spatial geometry. Courtesy of the WMAP Science Team[4].

By way of equations 1.5, 1.6, 1.9, 1.11, the Hubble rate and the definition of the Ω's, we find our favourite version of the first Friedmann equation for a universe with dark and luminous matter, radiation, possible curvature and dark energy. It is expressed as a function of the redshift,

$$\left(\frac{H(z)}{H_0}\right)^2 = (\Omega_{b0} + \Omega_{c0})(1+z)^3 + \Omega_{r0}(1+z)^4 + \Omega_{k0}(1+z)^2 + \Omega_\Lambda. \tag{1.12}$$

The Standard Model of cosmology, ΛCDM, is not as simple as the above model. It includes deviations from perfect homogeneity. This is a crucial point, needed for the evolution of structure in the Universe. However, eq. 1.12 does well in describing a smooth background universe to function as the "backdrop", something to build on when trying to describe these cosmic inhomogeneities. The full ΛCDM model is reliant on six parameters which have to be fitted to the data. Table 1.1 shows a set of best-fit estimates of these parameters. They are based on the best data set we have to date: Seven years' worth of full-sky observations, performed by the WMAP satellite.

We have now covered the basics of how we describe the Universe as we see it today. But the structures that surround us are not static. As we strain to observe objects ever further off, we see back in time, and find evidence of an evolving Universe. The next step in our brief account must therefore be to describe what we know about its origins.

1.2 The early history of the Universe

In the beginning, the Universe came into being. We do not know how, or why, or what came before, or even if "before" can be said to have any meaning at all. Quite possibly, the beginning of our Universe marks an absolute limit as to how far human knowledge can reach. If anything existed before, the information of it may well have been erased, like a cosmic reboot.

The Big Bang theory states that time and space came into existence some 14 billion years ago. Of the initial minuscule unit of time, called the *Planck era*, which lasted approximately 10^{-43} seconds, we know nothing, although theoretical physicists have many ideas, involving quantum gravity and unification of all the fundamental forces. Then followed - or so we think - the era of *inflation*, lasting a further tiny fraction of a second, in which the Universe is supposed to have undergone exponential expansion. This epoch and its consequences are of great importance for this thesis, and will be treated separately in chapter 3.

Particle physicists are in the process of constructing plausible theories for what happened next; from the process of *baryogenesis*, where the quarks making up all baryonic matter came into existence, to the production and annihilation of antimatter. These theories are still in the early developmental stages, and we will leave them be for now. We move ahead until the epoch of *Big Bang Nucleosynthesis*, BBN, which set in when the temperature had sunk to ~ 1 MeV, and the Universe was about a second old.

1.2.1 New elements in the cosmic stew

The one-second-old Universe was in a hot and dense state, a "primordial soup" of highly interacting particles, maintaining a thermal equilibrium. Generally, a particle species in an expanding universe will remain in thermal equilibrium if the interaction rate of whatever collisions or decay process it is involved in, is larger than the current Hubble rate. As the expansion of space-time causes the temperature to fall, those particle species with the lower interaction rates will fall out of equilibrium, no longer interacting significantly with the rest of the particle fauna. In the case of CDM particles, for instance, this transition, which we call freeze-out, will have happened at a very early time, when the Universe was still less than a second old. Cosmic neutrinos will also have decoupled from the rest of the plasma at temperatures a little above 1 MeV. Remaining in thermal equilibrium are photons, relativistic electrons and positrons, and nonrelativistic baryons.

At this early stage, photons constituted the dominant term of the energy density. We see this by investigating eq. 1.9: The energy density of radiation falls off faster than that of matter (a^{-4} versus a^{-3}), so extrapolating back in time, we find that $\rho_r > \rho_m$ until some time between BBN and photon-baryon decoupling.

As long as the temperature of the plasma is higher than the binding energy of light nuclei such as deuterium (^{2}H) and helium (^{3}He, ^{4}He), any such elements will be torn apart by energetic photons shortly after forming, and they will not make up any significant part of the plasma. These binding energies typically lie in the range of a few MeV, but because of the high photon density, accumulation of nuclei heavier than ^{1}H do not set in until $T \sim 0.1$ MeV. Before this can happen, the neutron-proton equilibrium will break.

Neutrons are unstable when they are not bound to a proton in a nucleus, and will decay into protons and electrons. As long as the temperature of the plasma is high enough, the neutron population is replenished by the reverse reaction, caused by electron-proton collisions. At $T \sim 1$ MeV, this can no longer occur, and the neutron population will decrease steadily until the plasma has cooled enough for the onset of nucleosynthesis. How long this takes depends on the current expansion rate, which in turn depends on the energy density. The same applies to the abundances of the various

light elements, like helium and lithium. When, and for how long, these elements can form, is given by the temperature of the plasma. Hence, the current neutron and light element abundances in the Universe are excellent probes for determining cosmological parameters, making BBN one of the foremost pillars on which the Big Bang theory rests.

1.2.2 *And there was light:* Recombination and the birth of the CMB

The binding energy of neutral hydrogen is 13.6 eV. Due to the still high ratio of photons to baryons, the plasma temperature has to decrease to $\sim$ 1 eV before electrons can remain bound to single protons and heavier nuclei, forming neutral atoms. Up to this point, free electrons, protons and photons are tightly coupled due to Compton and Coulomb scattering, making the optical depth of the plasma very large. In effect, the Universe is completely opaque.

Approximately 300,000 years after the Big Bang, the electrons started binding to protons in a process most frequently known as *recombination*, although (photon-baryon) decoupling is a more correct term. At decoupling, the free electron density dropped to near zero, taking with it the rate of Compton scattering. The Universe became transparent, allowing the photons to free-stream, carrying with them information about the local temperature where they were last scattered. These photons have moved through the expanding Universe ever since, becoming steadily redshifted by the expansion of spacetime. The shape of their frequency spectrum, which is that of a blackbody due to the thermal equilibrium at the time of recombination, does not change. Today, this field of radiation has a mean temperature of $\sim$ 2.73K, peaking in the microwave range of the electromagnetic spectrum. We know it as the *cosmic microwave background radiation*, or CMB, and it has been our prime source of information of the early Universe since its discovery in 1965.

1.3 The Cosmic Microwave Background Radiation

After its discovery, the CMB has been at the centre of cosmological research. Although it is not our only source for probing the cosmos, it outperforms the "competition" in how readily it provides us with new information. One simply needs to know how to decipher it.

So far we have allowed ourselves to look upon the Universe as being completely smooth. This, we know, is not true. At the time of recombination, when the CMB was formed, there were small inhomogeneities throughout the cosmic fluid. Our job, as cosmologists, is to map out these primordial inhomogeneities and compare their properties to those of our theoretical models. Sadly, this is not as simple as all that. The Universe might well be infinite, so making a direct map of it is an impossible task. Then there is the small matter of the time that has passed. After their last scatter at the end of recombination, the CMB photons we detect have been streaming through the evolving Universe for more than 13 billion years, being subjected to gravitational deflection and red- or blueshifting all the while. We need a detailed understanding of their journey in order to read the state of the early Universe from the pattern of the present-day CMB photons. A useful concept here is the *surface of last scattering*, which may be pictured as a spherical shell around us, made up of the points where today's

observed CMB photons were last scattered, at the end of recombination. This surface is our "sample" of the early Universe.

1.3.1 A detour on decomposition

We begin by tackling the size problem. Whether it is infinite or not, the Universe is certainly too large to draw a map of. Rather, we look at its statistical properties. This is a reasonable approach, since we believe the processes responsible for the distribution of matter in the Universe to be stochastic in nature. We would like to find the *distribution* of the inhomogeneities, giving the probability of finding a denser-than-average region of a certain size. Our weapons of choice are harmonic decompositions.

There are several methods of harmonic decomposition. What they all have in common, is that they are based on classes of orthonormal functions spanning a space. They allow us to decompose a field, however complicated, into a series of independent harmonic functions - modes - which may be treated separately, mathematically as well as in physical interpretation. This property is vital for our cause.

When working in flat space, we employ Fourier decomposition. This technique decomposes any function into an infinite sum of sines and cosines, i.e. wave functions defined on $\mathbb{R}^n$. In three-dimensional space it has the form

$$f(\vec{x}) = \int_{-\infty}^{\infty} \mathcal{F}(\vec{k}) e^{i\vec{k}\cdot\vec{x}}\, d^3\vec{k} \quad \text{and} \quad \mathcal{F}(\vec{k}) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} f(\vec{x}) e^{-i\vec{k}\cdot\vec{x}}\, d^3\vec{x}, \tag{1.13}$$

where we have employed the complex exponential form of the wave functions. Using this, we are able to map any 3D scalar field into the function $\mathcal{F}(\vec{k})$, which gives the amplitude of each harmonic component as function of wavenumber (inverse wavelength).

The second of our favourite decompositions is based on *spherical harmonics*. These are, as the name indicates, wavefunctions defined on a spherical surface. They are known collectively as $Y_{\ell m}$, where the two indices determine the characteristics of each "mode", like the wavenumber does for a sine or cosine. The spherical harmonics form an orthogonal basis for the surface of a sphere, and may thus be used to decompose any function defined on this space. The spherical harmonics transformation of a function $f(\theta, \phi)$ takes the form

$$f(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi). \tag{1.14}$$

The coefficients $a_{\ell m}$, like the Fourier coefficients, describe "how much" of a given mode is needed in the decomposition of the function. Basic relations concerning the spherical harmonics are listed in §C.1 of the appendices.

A third useful tool is the Legendre multipole expansion, which is, among other things, used when modelling the perturbed CMB photons. For an introduction to this technique and more on the Fourier and spherical harmonics decompositions, the committed reader is encouraged to seek out any good textbook on mathematical physics, such as Hassani [3].

1.3.2 The power of the power spectra

Armed with the tools just described, we are able to construct a quantity that contains all the important information of an inhomogeneity field in just one function: The power

spectrum.

A power spectrum is a condensation of the decomposition coefficients of a field, giving the distribution of the field's power over structures of different sizes. In the case of a 3D Fourier decomposed field, the power spectrum $P(k)$ is found by

$$\langle \mathcal{F}(\vec{k})\mathcal{F}(\vec{k}')\rangle = (2\pi)^3 P(k)\delta^3(\vec{k}-\vec{k}'), \tag{1.15}$$

where $\delta^3(\vec{k}-\vec{k}')$ denotes the three-dimensional Dirac delta function. We see that this reduces the original field into a one-dimensional function, containing only the information we are interested in, which is how much of the field's total energy that is contained in a mode of wavenumber $|\vec{k}| = k$. We can use this strategy to describe the distributions of all matter and energy components of the Universe.

In a similar fashion, we can define a power spectrum based on the spherical harmonics decomposition of a spherical function - for instance, the CMB temperature field. But first, we should look at how this field is related to the *global* temperature field.

The global temperature field is a way of picturing the distribution of background photons throughout the Universe. It fits our purpose well to express it as a sum of a mean term and a small deviation,

$$T(t,\vec{x},\hat{p}) = T_0(t)(1+\Theta(t,\vec{x},\hat{p})),$$

where $T_0(t)$ is the true mean temperature in the Universe at time t, while the Θ describes the deviation from the mean as a function of time, spatial position $\vec{x}$, and the direction of the momentum of the photons, $\hat{p}$. Because we are tiny humans stuck in a tiny region of space, we cannot measure the whole of $T(t,\vec{x},\hat{p})$. Hence we define the local value to be the *CMB temperature field*, $T(\hat{n})$,

$$T(\hat{n}) \equiv T(t_0,\vec{x}_0,\hat{p}) = T_0(1+\Theta(\hat{n})).$$

The $\hat{p}$ directional coordinate has been substituted by the $\hat{n}$. This unit vector signifies the direction from which we see each photon arriving, since we naturally cannot see photons moving in any direction but ours.

While the mean temperature is its own field of study, it is not of any particular consequence for surveys of the CMB anisotropy. QUIET being of the latter category, we are more interested in the deviation field, which we analyse using the spherical harmonics decomposition,

$$\Theta(\hat{n}) = \sum_{\ell=0}^{\infty}\sum_{m=-\ell}^{\ell} a_{\ell m}Y_{\ell m}(\hat{n}). \tag{1.16}$$

In a move similar to the one used for the Fourier decomposed fields, we construct a power spectrum of the CMB anisotropy by way of the decomposition coefficients,

$$\langle a_{\ell m}a^*_{\ell' m'}\rangle = \delta_{\ell\ell'}\delta_{mm'}C_\ell, \tag{1.17}$$

where the δ_{ij}'s are the Kronecker delta function, and we have assumed that all the $a_{\ell m}$'s for a given ℓ have the same variance. This follows from our assumption of statistical isotropy. If one were to find that the Universe is not isotropic after all, we would have to account for the m's in the power spectrum, too.

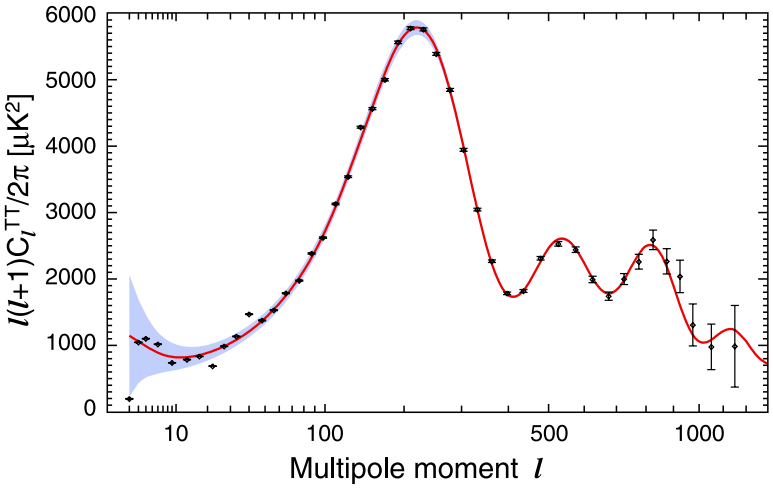

Figure 1.1: The temperature (TT) power spectrum for the seven-year WMAP data set. The solid line is the predicted spectrum for the best-fit ΛCDM model. The error bars on the data points represent measurement errors while the shaded region indicates the uncertainty in the model spectrum arising from cosmic variance. Figure courtesy of the WMAP science team[4].

When talking of *the* power spectrum, it is usually the C_ℓ's one is referring to - the power spectrum of the CMB temperature anisotropy field. The reason this particular power spectrum is singled out, is that it forms one of the central pillars of modern CMB research. It is the prime data product of all observational experiments, essential in cosmological model building, and one of the chief measures by which we compare theory to observation. We have just described how we find the observational power spectrum, and so we turn to the model.

The theoretical power spectrum is based on a conception of the early Universe. In general, this would be expressed in the form of a primordial power spectrum for the spatial distribution of matter and energy, as predicted by the theory one wants to test. Nowadays, a variety of theories of cosmological inflation dominate this field. With the help of known physics, good algorithms, and plenty of CPU time, we can simulate the interaction of all the components of the model universe, and evolve the system into a model of the present-day Universe. Specifically, we extract the function describing the present-day global temperature inhomogeneity field, project it onto a sphere to simulate the CMB temperature field $\Theta(\hat{n})$, and calculate a new power spectrum to compare with the observed one,

$$C_\ell = \frac{1}{(2\pi)^3} \int P(k)\Theta_\ell^2(k)\, dk. \tag{1.18}$$

Here, Θ_ℓ is the coefficients of the Legendre decomposition of $\Theta(t_0, \vec{k})$ (the Fourier decomposition of $\Theta(t_0, \vec{x}, \hat{p})$), while $P(k)$ is the primordial power spectrum for the dark matter perturbation. The mode in the three-dimensional photon distribution associated with each ℓ is often referred to as a *multipole*, and accordingly, the ℓ's are the *multipole moments*.

Figure 1.1 illustrates the comparison of a theoretical and observational power spec-

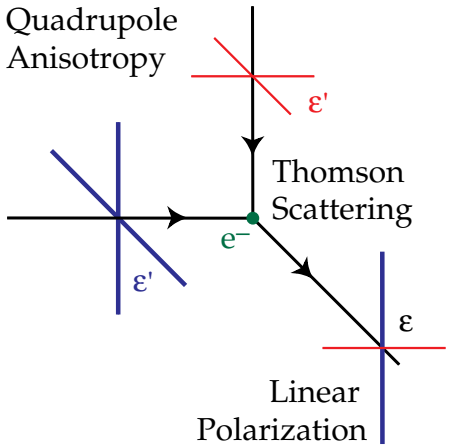

Figure 1.2: Linear polarization arising from Thomson/Compton scattering of incident quadrupolar radiation. Figure courtesy of W. Hu and M. White. [5]

trum, showing the temperature power spectrum for the current standard model alongside the observational spectrum of the current best CMB survey - WMAP[4]. It is easy to see why the model has won its high regard, since the fit is rather extraordinary. In chapter 2, we will look more closely on the slow and gradual process leading up to these marvellous results.

1.3.3 The polarized CMB

So far we have only discussed the power spectrum of the temperature anisotropy. But light carries more information than just temperature. It also has the property of polarization, meaning that the intensity is not evenly distributed among all possible orientations of the electric and magnetic fields, but has some preferred direction. For the unfamiliar reader, an introduction to the basic notation used in the description of polarized light is given in appendix A.

There are only a small set of effects contributing to polarization patterns in the CMB. The first and foremost of these is Compton scattering (or Thomson scattering, the low-energy limit of the former) at the time of recombination.

Compton scattering is an exceedingly effective way of producing polarized radiation: Relative to a given direction of scattering, all transverse polarization will be transmitted, while parallel polarization will be impeded. However, this effect may easily be cancelled if the incident radiation comes from several directions at once. When considering unpolarized incident radiation coming from all directions, we find that the only distribution causing the scattered radiation to be polarized is the quadrupole. In figure 1.2 is given a simple illustration to justify why this is. The red/blue lines indicate possible orientations of the electric fields. Each incident ray has equal intensity in both orientations of the electric field, while the transmitted radiation has more intensity in the vertical direction. In other words, it possesses a net polarization. This would not have happened, had the two incident rays been of equal intensity, as in the case of an isotropic photon distribution.

Formally, one may define a perturbation field Θ_P of the polarized background, just

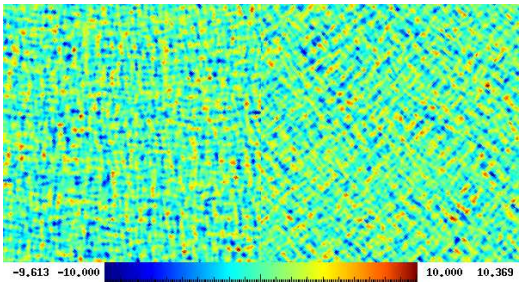

(a) E-mode map. Q(left), U(right). Range $\pm 10\mu$K.

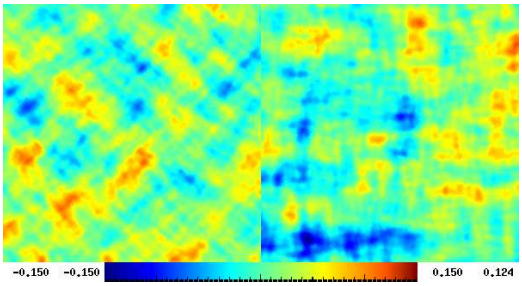

(b) B-mode map. Q(left), U(right). Range $\pm 0.15\mu$K.

Figure 1.3: Simulated maps of Q- and U-polarization patterns on the sky, as predicted by the standard model. Images courtesy of Sigurd K. Næss.

like we did for the photon temperature. By taking into account the full, evolving distributions of photons and electrons, one finds that the polarization caused by Compton scattering in the early Universe follows the relation

$$\Theta_P = \frac{3\tau}{2}(1-\mu^2)\Theta_2, \tag{1.19}$$

where $\mu = \hat{k}\cdot\hat{n}$; $\hat{k}$ is the wave vector of a perturbed mode in the incident radiation, and $\hat{n}$ is the direction into which a photon is scattered. Firstly, we note of course the key role of the quadrupole photonic moment, Θ_2. The presence of the geometrical factor μ indicates that the polarization pattern will vary with the orientation of the original perturbation, as well as with the position of the observer. Finally, the optical depth is a factor because a high density, resulting in rapid scatterings, will randomise photon orientations, and any net polarization will diffuse quickly, along with the quadrupole itself.

When the first stars began producing ionizing radiation, free electrons were reintroduced into the cosmic fauna. We call this incident *reionisation*. The new free electrons cause additional polarization on large scales. Furthermore, gravitational lensing of the CMB photons as they pass by clusters of galaxies may change the pattern of existing polarization. Still, the set of effects we have to take into account when considering the polarized sky is rather smaller than the multitude of things affecting the temperature power spectrum, which is a nice trait from the viewpoint of those who have to interpret it. When building a full universe model, we would include these effects into an augmented form of eq. 1.19.

For a more thorough review of the technicalities of the polarized CMB, including the derivation of eq. 1.19, chapter 10 of Dodelson [2] is highly recommended.

1.3.4 Decomposing the polarized sky

At the moment, we know of no processes that might cause circular polarization of the CMB, so the problem of analysing the polarized background is restricted to the linear case. Observational studies chart the polarization pattern by recording the degree of

Q- and U-polarization at all points on some part of the sky[3]. The resulting field takes the form of a vector field on a spherical surface[4]. To analyse it, we apply a result from vector calculus, stating that a vector field on a spherical surface may always be decomposed into the sum of an irrotational ("curl free") and a solenoidal ("divergence free") vector field. Due to their similarity to the familiar electric and magnetic fields of simple charge distributions, these components are often referred to as E-modes and B-modes, respectively.

As it turns out, this division also carries a physical significance. The scalar (density) perturbations responsible for creating polarization by Compton scattering is only capable of producing the E-mode pattern. To create B-modes, we need *tensor* perturbations. These are predicted to be produced by cosmological inflation, and manifest as *gravitational waves*. B-mode polarization has not yet been observed, and the fates of many theories depend on the outcome of the ongoing search. We will explore this topic further in chapter 3.

Figure 1.3 shows an example of the patterns we would expect to find on the sky, given E-mode or B-mode polarization as predicted by the standard model with a small amount of tensor perturbations. We note that the orientation of the B-mode pattern with respect to the Q- and U-polarization axes is the opposite of what it is for the E-modes. We also observe that the dominating anisotropy mode is on a large scale for the B-modes, and on smaller scales for the E-modes, and that the amplitude of the B-modes is $\sim 10\%$ of the E-modes. This is due to the differing physical origin of the two types: While the E-modes are mainly created by anisotropy modes that are well inside the horizon at the end of recombination, we will find in chapter 3 that the B-modes are due to weaker, super-horizon perturbations.

When we computed the power spectrum of the temperature anisotropy (eqs. 1.16, 1.17), we used a spherical harmonic decomposition. We choose a similar strategy for the polarization, but with one important difference: While temperature is a scalar field, the linear polarization field is a tensor quantity, and dependent upon the orientation of the reference coordinate system. Its properties under rotation identify it as a *spin* ± 2 *field*, and it follows that the *spin-weigthed spherical harmonics* ${}_sY_{\ell m}$, $s = \pm 2$, would be an appropriate basis for their decomposition. These functions have the same basic properties as the "regular" spherical harmonics, the latter being identical to the $s = 0$ case.[5]

The decomposition of the polarization field is given by

$$(Q \pm iU)(\hat{n}) = \sum_{\ell,m} a_{\ell m}^{(\pm 2)}\, {}_{(\pm 2)}Y_{\ell m}(\hat{n}) \;\Leftrightarrow\; a_{\ell m}^{(\pm 2)} = \int (Q \pm iU)(\hat{n})\, {}_{(\pm 2)}Y_{\ell m}(\hat{n})\, d\Omega.$$

From this we construct the E-mode and B-mode coefficients by

$$a_{\ell m}^{E} = -\frac{1}{2}\left(a_{\ell m}^{(+2)} + a_{\ell m}^{(-2)}\right) \quad \text{and} \quad a_{\ell m}^{B} = -\frac{1}{2i}\left(a_{\ell m}^{(+2)} - a_{\ell m}^{(-2)}\right),$$

[3]Preferrably we would map all of the sky. This requires a space-based detector, which is rather expensive, so most projects make do with smaller patches. The Planck satellite will produce full-sky polarization maps.

[4]We keep in mind that polarization "vectors" are not true vectors, since they have no direction.

[5]Consult Seljak & Zaldarriaga[6] for further details on analysing the polarized CMB.

and finally, this gives the polarization angular power spectra,

$$C_\ell^{EE} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\langle a_{\ell m}^E a_{\ell m}^{E*}\rangle \quad \text{and} \quad C_\ell^{BB} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\langle a_{\ell m}^B a_{\ell m}^{B*}\rangle.$$

1.4 Getting to know our Universe

The Universe is a big place. It follows that cosmology is a large field, and summarising it is no easy task. In this chapter, we have touched upon some central elements - for the most part, barely scraping the surface, although it should serve well enough as motivation when we move on to more serious matters.

One of the major problems in cosmology is finding ways to observe and describe the biggest structures known to man without drowning in the sheer size and complexity of it all. As we have seen, statistical methods are of the utmost importance for this endeavour, flanked by appropriate approximations. As cosmology evolved from metaphysical musings into the computationally intensive, thriving science it is today, we have found ways to map and analyse the large-scale properties of the cosmos in a way that allows us to gain concrete knowledge of things unimaginably far away, in space as well as time.

Large-scale surveys are the key to unlocking many a cosmic mystery. Mapping the positions and compositions of stars and galaxies has provided crucial information on the distribution of baryonic and dark matter. By doing spectroscopic investigations revealing the relative abundances of elements, we were able to test the predictions of Big Bang Nucleosynthesis, and provide the first solid piece of evidence in favour of the Big Bang theory. Several large projects have been undertaken with the goal of mapping e.g. galaxy distributions, or the radial velocities of objects as functions of distance. As a result, we steadily improve our understanding of the mechanics doing the directing in the cosmic theater. Such surveys may even enlighten us on the shape of the primordial matter power spectrum. And, of course, there is the CMB, towering above the rest in its remarkable usefulness. Countless scientists have built their careers around its study, a trend that will certainly continue over the foreseeable future.

With the expansion of CMB science into the realm of polarized light, cosmology has moved on to a new level. Including cross-correlations, we now have not one, but six power spectra to play with - six functions to measure with ever-increasing accuracy, compare with theoretical predictions and interpret into factual statements of the evolution of the Universe. The future of cosmological research is indeed promising.

Chapter 2

The history of modern cosmology

> *It does not make any difference how beautiful your guess is. It does not make any difference how smart you are, who made the guess, or what his name is; if it disagrees with experiment it is wrong.*
>
> *Richard P. Feynman*

The history of scientific cosmology is short, but eventful. In this chapter, we will trace the unfolding of this young science, and follow the development of the central ideas as they were formed in the light of a steady drip of new observations. This is as much an effort to see the greater context of the work we are doing, as an homage to the pioneers that wrested the first secrets from the early Universe. For a thorough review of the history of modern cosmology and astrophysics, I would recommend M. Longair's *The Cosmic Century*[7].

2.1 1916 - 1965: Early years

Given our severely limited view of the Universe, the dominating cosmological model for the best part of the last 2000 years were one in which the stars were fixed points of light upon an outer sphere, surrounding the system of the Sun and the planets. Galileo Galilei's substantial contribution to the development of the telescope marked the beginning of the end of this view, as this new tool revealed that the heavens had many more secrets to be discovered. The misty band of the Milky Way turned out to be made up of stars, too close together to be seen by eye. After a meticulous mapping effort involving a large number of telescopes of his own design, british astronomer William Herschel discovered in the early 19th century that the stars were distributed around us in a three-dimensional, disk-like structure. In 1924-25, following two years of observations using the new 2.5 m Hooker telescope at Mt.Wilson observatory, Edwin Hubble presented his results verifying, for the first time, that ours is just one of many separate galaxies[8]. It was suddenly clear that the Universe was an awful lot bigger than we had thought it was.

2.1.1 The birth of a science

In 1916, Albert Einstein published his general theory of relativity[9], which we have already encountered in §1.1.1. Einstein's theory, with its revolutionary new view of the fabric of spacetime, introduced a framework for building mathematical models of the

Universe as a whole, allowing for the first time for the formulation of true scientific theories of cosmology. Consequently, this event is generally considered to be the beginning of scientific cosmology.

2.1.2 Model development

An infinity of solutions to the Einstein equations exist, but only a relatively small subset are physically viable. In 1919, general relativity, as well as Einstein himself, became famous following the theory's success in predicting the deflection of starlight by the sun, as measured during a solar eclipse that year. Within a few years after this breakthrough, two main, competing classes of solutions had emerged. Einstein himself believed strongly that the Universe had to be static, and so chose to include into his theory a cosmological constant, Λ. Without the inclusion of Λ, which was a geometric property of spacetime, Einstein's equations have no static solutions.

Einstein published his preferred cosmological model as early as 1917[10]. Soviet mathematician Alexander Friedmann discovered a class of non-static, homogeneous and isotropic solutions, which were published in 1922 and 1924[11, 12]. His results were obtained independently by Georges Lemaître in 1927, and refined further by Howard Percy Robertson and Arthur Geoffrey Walker in the 1930's, and so this solution is known as the Friedmann-Lemaître-Robertson-Walker metric (or a variant of this, depending somewhat on geography)[13, 14, 15].

Meanwhile, new observations ensured a gradual broadening of our horizon, aided by rapid technological development. In 1929, Edwin Hubble combined his own galaxy surveys with the spectroscopic galaxy redshift measurements of Vesto Slipher, and found that the grand majority of galaxies move away from us at a speed that increases linearly with distance[16]. This confirmed that the Universe did indeed seem to be expanding. At this discovery, Einstein was forced to abandon his static universe, and is famously known to have referred to the introduction of the cosmological constant as his "biggest blunder".

However, the case for the expanding universe was not so easily won. It is a characteristic of the non-static Friedmann cosmologies that the resulting universe have a finite age. If the laws of physics have not changed, the expansion of the Universe today implies that all matter must, some finite time ago, have been packed into a point of infinite temperature and density. This idea, although fascinating, may be seen as somewhat disturbing, and a lot of people fought hard for the view of the Universe as something that should be infinitely old and unchanging.

For the following decades, most cosmologists divided into two camps. The one, led by Fred Hoyle, advocated a "steady state" universe. In 1948, Hoyle formulated a theory that allowed the Universe to be static while expanding via a mechanism of continuous creation of matter[17]. The competing theory, that the Universe started out in a singularity, was first formulated by Lemaître in 1931 and further developed by George Gamow[18, 19][1]. It eventually became famous under the title *Big Bang theory*, a term originally coined, in a scathing manner, by the opponent Hoyle.

[1]Gamow's 1948 article, in which he introduced Big Bang Nucleosynthesis, is known for more than its content. The listed authors are "Alpher, Bethe and Gamow", but Hans Bethe was in fact never involved in the article. Gamow had simply added his name to the list to complete the "$\alpha - \beta - \gamma$" pun.

2.1.3 Fighting for the title

The debate between the Big Bang and Steady State theories raged on. Some people concocted other ideas as well, such as the 'Tired Light' theory of Fritz Zwicky, which was meant to explain Hubble's results without the need of spacetime expansion. However, this never won through, and for a long time, the race was tied between the two main contestants. As more observations were done and the models were refined, the Big Bang theory gained support. It performed better than its opponent at predicting the observed abundances of light elements, as well as being able to explain why such objects as quasars, very young galaxies, are exlusively seen at the very edge of the visible Universe.

The breakthrough came in 1965, when physicists Arno Penzias and Robert Wilson, experimenting with a large and very sensitive radio antenna at Bell Laboratories, found a homogeneous and isotropic background signal of about 3 K[20]. Less than a decade earlier, in 1948, Ralph Alpher and Robert Herman had found that the afterglow of the Big Bang should exist in the form of blackbody radiation homogeneously distributed throughout the Universe[21]. In 1964, Soviet physicists Doroshkevich and Novikov found that this radiation should be a detectable signal - a cosmic microwave background[22]. At the time when Penzias and Wilson made their discovery, a team of astrophysicists at Princeton University, only a few miles away, were preparing a search for the elusive signal, unaware of the two engineers that had just stumbled upon it.

Penzias and Wilson were awarded the Nobel Prize of Physics for their discovery in 1978. And, seeing as the Steady State theory had neither predicted the CMB, nor was able to explain its existence, the detection of the CMB marked the end of the discord in the cosmological community. Although the Big Bang theory was far from perfect, a consensus arose, and attention was turned towards fine-tuning it.

2.2 1965 - 1996: The plot thickens

The decades following World War II saw significant development on the areas of detector technology and rocket science. For the first time in history, one was able to do observations outside the human-visible part of the electromagnetic spectrum. This fueled the progress in astrophysics and cosmology, which was considerable over the 70's and 80's. Large galaxy surveys were conducted, and considerable attention given to the distribution of matter in the Universe.

2.2.1 New ingredients

Several of the central features of what we now know as the Standard Model of Cosmology were introduced into the framework over the years following the discovery of the CMB. In the 70's, observations of the rotational velocities of galaxies indicated that more matter was present than what was visible (e.g. [23]). Some new type of particle, interacting only through gravity, had to be present. This concept, which had been introduced as early as 1933 by Fritz Zwicky, was given the fitting name *dark matter*.

Another, even more significant addition to the theory was inflation. This groundbreaking new idea, proposed by Alan Guth in 1981, said that the Universe might have undergone a period of exponential expansion shortly after the Big Bang[24]. As we

will see in chapter 3, this resolved some of the classical Big Bang theory's more serious issues. Also, inflation would give rise to small inhomogeneities in the primordial matter distribution, which would show up as fluctuations in the CMB power spectrum.

The idea that said fluctuations ought to exist had, in fact, been around for some time before the introduction of inflationary theory. While the simplest cosmological models describe a homogeneous universe, it is painfully obvious that our Universe does not fit this description. The existence of stars, galaxies and humans, for instance, is a severe deviation from this condition. The accumulation of matter into a highly dense region, like a galaxy, cannot occur unless there is some small overdensity to seed a gravitational instability. The evolution of such overdensities into galaxies were studied in detail in the early 60's, and limits for their amplitudes were calculated[25]. In 1967, Rainier Sachs and Art Wolfe calculated what was to be known as the Sachs-Wolfe effect, showing that those necessary density perturbations on the surface of last scattering would give rise to anisotropies in the CMB[26]. They were expected to be tiny - Sachs and Wolfe predicted anisotropies of the order of 1 percent of the background intensity - promising a real challenge for the young science.

2.2.2 Hunting for anisotropies

The convincing case for the anisotropies combined with the fact that one had not yet confirmed that the CMB's thermal spectrum was that of a blackbody, resulted in fierce effort directed towards CMB observational projects through the 70's and 80's. David Wilkinson of Princeton University, who had been on the team that was beaten to the discovery of the CMB by Penzias and Wilson, was involved in some of the earliest attempts.

Due to the opaqueness of the atmosphere in the microwave region, it was quickly realised that the best chances of detecting structure in the CMB lay in airborne experiments, or even better - satellite missions. Since the latter involve very large investments in time as well as money, the earliest results came from experiments mounted on sounding rockets or balloons. Unsuccessful attempts at detecting anisotropies, such as the balloon missions ARGO (1988) and FIRS (1989), resulted in the setting of upper limits to their amplitudes[27, 28]. Also notable was the Soviet satellite mission RELIKT-1, launched in 1983, whose data initially amounted to nothing. A reanalysis carried out several years later showed positive detection of anisotropies[29], but before these results could win the attention they deserved, another project had entered the stage: The *COsmic Background Explorer*, or COBE.

2.2.3 COBE

NASA's COBE satellite was launced in 1989 atop a Delta rocket, and went into a sun-synchronous low-Earth polar orbit. It was concieved as early as 1974, but was delayed several times due first to financial issues, then to the accident of the space shuttle Challenger in 1986. It was to gather data for four years, and carried instruments for three separate experiments, DIRBE, FIRAS, and DMR.

DIRBE (*Diffuse InfraRed Background Experiment*) mapped the full sky in several frequency bands, covering the infrared part of the spectrum, and producing the first full-sky map of the infrared background. FIRAS (*Far-InfraRed Absolute Spectrophotometer*)

aimed to detect the intensity of the CMB at a large range of frequencies around the peak. The goal was to affirm, once and for all, the blackbody nature of the radiation.

The results were staggering. The thermal spectrum measured by FIRAS had no measurable deviation from the blackbody spectrum corresponding to a body with temperature of 2.725 ± 0.001K[30]. John Mather, lead scientist of FIRAS, presented the first results at a meeting of the American Astronomical Society in January 1990, six weeks after the satellite's launch, based on a mere nine minutes of observation. The resulting curve fit was so spectacular, it recieved a standing ovation[2].

The last of the three, the DMR (*Differential Microwave Radiometer*) was designed to detect the CMB anisotropies. Its angular resolution was poor, a humble 7 degrees, but DMR compensated by being a full-sky survey, always adjusted so as to avoid the intense microwave radiation from the Earth and the Sun. It made use of differential radiometers, measuring the difference between the temperature at two separate places, rather than simply measuring the temperature directly. Instrumental noise is generally independent of what one is observing, so this approach ensures cancellation of most of the noise, allowing the instrument a greater chance at detecting the extremely weak signal. In 1992 the DMR maps were good enough to give definite statistical proofs of existing anisotropies of order 10^{-5}[31]. When the final data were released in early 1996, the experimenters could demonstrate a positive detection of true features in the CMB sky map[32].

George Smoot, lead investigator of DMR, and John Mather recieved the Nobel Prize in Physics in 2006 for their work on COBE.

2.3 1996 - present: Precision cosmology

The Nobel comittee acknowledged COBE as marking the beginning of precision cosmology. This piece of solid evidence triggered a cascade of new projects. The shape of the CMB anisotropy power spectrum, now confirmed to exist, depends strongly on the value of various cosmological parameters, and as such it constitutes an unsurpassable method by which cosmologists may test their theories. A fierce collecting of data ensued. As with the hunt for the CMB fluctuations, a satellite mission would be ideal, and planning one began shortly after the final COBE results were released. Meanwhile, a number of smaller projects managed to produce results, and important discoveries were made within few years of COBEs major breakthrough.

2.3.1 Ground- and balloon-based experiments

Ever since Friedmann's 1922 article it had been known that the FRW metric allowed for three distinct geometrical options: Space can be open, closed, or flat, depending on the energy content. Come the 90's all the evidence pointed towards there being far too little matter, dark as well as light, for either of the latter two to be the case. Observation of galaxies and galaxy clusters suggested a total matter density of about a third of the critical density needed for the Universe to be flat. However, the theory of inflation, which did such a good job on explaining things like the origin of structure, also caused flatness, contradicting the obvious conclusion of an open Universe.

[2]For a captivating account of this rare occasion by award-winning science writer Marcus Chown, see "The Nine-Minute Spectrum", available online at http://www.marcuschown.com/afterglowsample.htm

For some time, the debate circled three possibilities, all of which were acceptable with respect to the currently available data: Either the Universe has a matter density of a third of the critical density, and is open. Else it is flat, with the sum of light and dark matter amounting to exactly the critical density. Or one could bring Einstein's old cosmological constant in from the cold, to represent some new, undetected form of energy, allowing for a flat Universe in which the matter density is still only a third of the critical value. The shape of the CMB power spectrum would break the tie.

A multitude of experiments, mounted on telescopes on the ground or flown in balloons, was performed before the end of the millennium(e.g.[33, 34]). The upper panel in figure 2.1 show the status in 1997. Compared to two theoretical models corresponding to flat and open universes, the data seem to favour the former, but are not at all conclusive. This figure also serves as an illustration of the richness of new experiments in this period.

In 1998, new observations of distant supernovae were conducted independently by two projects; the *Supernova Cosmology Project* (SCP) and the *High-Z Supernova Search Team* (HST). Their results bore strong evidence towards the universal space-time expansion actually accelerating, counting in favour of the suddenly popular cosmological constant[35, 36]. The unknown form of energy sourcing the acceleration was dubbed "dark energy", and soon became a standard part of cosmological models. Adam Riess and Brian Schmidt of the HST and Saul Perlmutter of the SCP were awarded the 2011 Nobel Prize in Physics for their results in these projects.

The next breakthrough came with BOOMERanG (*Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics*). This cleverly named experiment made use of a telescope flown in a balloon in a high-altitude orbit ($\sim$ 38 km) around the South Pole, thus avoiding most atmospheric absorption of microwave radiation. The experiment was the first to detect the CMB anisotropies at a such a range of scales that it alone could give enough evidence to discern between the possible theoretical models. After their first flight in 1998, in which the ballon did a full circumvolution of the South Pole, the BOOMERanG team was able to produce the first high-resolution maps of the CMB. Their data, when they were published in 2000, gave clear preference to the flat Universe model[37]. This conclusion was shortly fortified by the results of a similar experiment, MAXIMA (*Millimeter Anisotropy eXperiment IMaging Array*)[38]. As is clear from the second panel of figure 2.1, the evidence for a flat Universe as contrasted to an open one had at this point become irrefutable.

2.3.2 WMAP and the victory of ΛCDM

In February 2003, the first year of data from the WMAP satellite was released. The Wilkinson Microwave Anisotropy Probe was what everyone had been waiting for. Funded by NASA, as COBE had been, this mission was chosen out of several candidates to follow COBE, and the work involved several of the same people. Its original title was MAP, but this was added to in honour of David Wilkinson, by now a renowned cosmologist, who died shortly before the launch in 2001. The satellite was sent into an orbit about the L2 Lagrangian point, 1.5 million kilometers from the Earth, an advantageous position in that the Sun, Earth and Moon stay close together, covering only a small piece of the sky. The goal was to produce full-sky, high-resolution maps of the CMB, using differential detectors in a manner similar to that of COBE. The task was

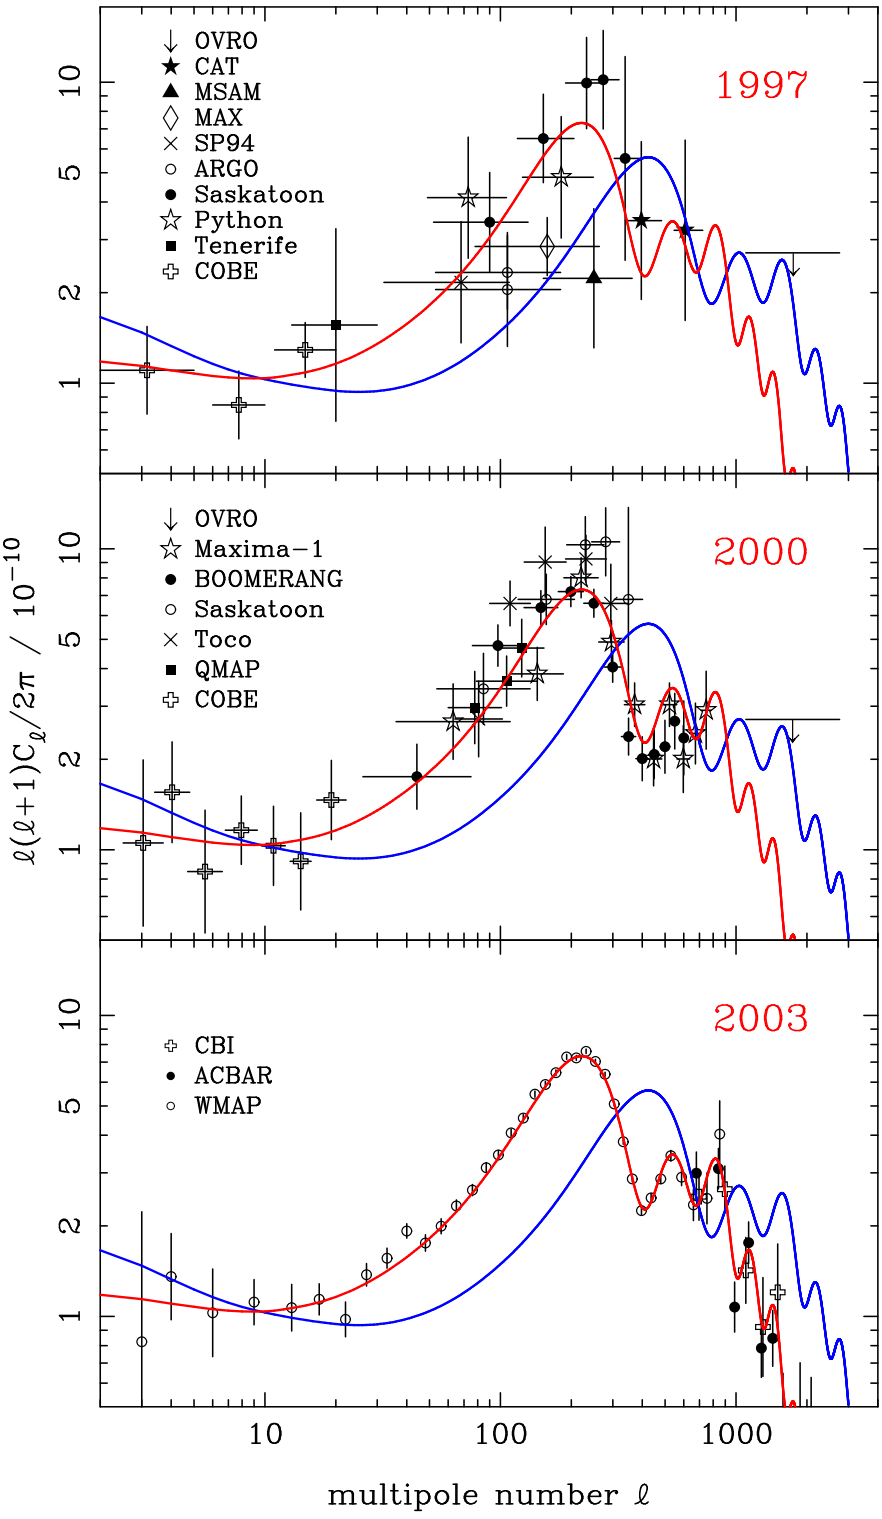

Figure 2.1: Results of various CMB surveys, compared with theoretical spectra. The red curve corresponds to a flat Universe, the blue to an open one, for otherwise identical parameters. These are $[n, \Omega_m, \Omega_b, h] = [1, 0.3, 0.05, 0.65]$. Figure courtesy of J. A. Peacock. [39]

achieved with great success.

Based on the data, the WMAP team produced best-fit estimates of a number of cosmological parameters; the density parameters of the different energy components, the Hubble parameter, spectral index, age of the Universe and so on. The resulting model fits so well to all available observational evidence, it has been accepted by the cosmological community and is widely recognised as the Standard Model of cosmology. The standard model universe contains roughly 70% dark energy, represented by the cosmological constant, and about 25% cold dark matter. Hence it is commonly referred to as ΛCDM[40].

In the final panel of figure 2.1 we get a clear picture of the impact that WMAP had after only a year of operation. It scanned the sky for seven years, and the material it provided still benefits the community today. The next big satellite-based CMB mission is the European Space Agency's Planck satellite, which is currently in operation, and is expected to provide CMB sky maps of unprecedented accuracy.

2.3.3 Polarization experiments and the road ahead

Given the standard model's solid performance since its emergence, substantial changes to its main components does not seem likely. Currently, the focus of those working with observation of the CMB temperature anisotropy is refining the estimates of cosmological parameters, and determining the shape of the power spectrum to ever higher accuracy and smaller angular scales. The biggest pieces of the puzzle seem to be in place, and it is time to work on ever smaller ones. However, as we saw in chapter 1, there is more to the CMB than temperature.

The polarization of the CMB bears promises of new and exiting discoveries. CMB polarization experiments are at the current frontier of cosmological observational research. These projects are far more demanding than detecting the temperature spectrum, since the signals we are looking for are extremely faint: E-mode anisotropies have been detected at amplitudes of the order of 10μK, weaker than the temperature anisotropies by almost an order of magnitude. The as yet undetected B-mode spectrum is expected to be weaker still, by at least one order of magnitude. Undaunted by this, cosmologists all over the world are working eagerly on polarization experiments even as we speak.

Much energy is devoted to detecting the E-mode spectrum to higher accuracy, but more exciting still is the search for the elusive B-modes. Whatever the results of this search will be, the rewards in terms of theory development will certainly be ample. Should we, for instance, successfully probe to accuracies where the leading theories predict B-modes without detecting them, those theories would fall, and we would have to find new ones to replace them. And, of course, one can hardly overlook the promise of glory to those who manage to produce the first positive B-mode detection, if such modes should prove to exist. New discoveries of this magnitude are rare in modern science - after all, all the easy ones have already been taken - so this is certainly a motivation for many.

The *Q-U Imaging ExperimenT*, or QUIET, which will be the subject of this thesis, is one of the many projects focused on B-mode detection. Thus it is only reasonable that we should end this historical review by taking a brief look at the progress so far. The first positive detection of CMB E-mode polarization were reported in 2002 by the

Degree Angular Scale Interferometer(DASI), a ground-based experiment observing from the Amundsen-Scott South Pole Station research facility[41]. Since then, several other experiments have contributed to the gradually improving E-mode data, including the ground-based CAPMAP, CBI, and QUaD, and the balloon-based BOOMERanG and MAXIPOL (the successor of MAXIMA)[42, 43, 44, 45, 46]. WMAP has also produced polarization data, although the quality is not comparable to that of their temperature measurements[47]. For the most part, the experiments have also mapped the TE cross-correlation spectrum, and produced upper limits to the B-mode spectrum on a variety of angular scales. An illustration of the situation as of mid-2009 is given in figure 2.2(a).

The cosmological parameter most intimately connected to the B-modes is the *tensor-to-scalar ratio*, r. This quantity, which will be defined and discussed in §3.3, gives the relative power of tensor perturbations to scalar perturbations, and is expected to be small. More importantly, it is directly connected to the energy scale of inflation, and is highly sensitive to the details of the theory. This makes r an excellent measure for discerning between inflationary theories.

Figure 2.2(b) shows a composition of the expected r's for different inflationary theories, alongside estimated likelihoods of some chosen experiments. To the right, we see the likelihood for r based on data from WMAP, ACT (*Atacama Cosmology Telescope*) and large-scale structures (LSS). These are only quite weak upper limits to r, but the models predicting the larger values of r (such as $V(\phi) \sim \phi^4$) are still ruled out. The likelihood curves in the centre of the plot are forecasts for QUIET-II for different patch configurations. We see that one large patch gives better constraints than four smaller ones, but more importantly, this figure shows that QUIET-II can be expected to rule out or detect values of r down to the *Lyth bound* at $r \lesssim 0.008$. This boundary is important, as it distinguishes between "small field" and "large field" models (see chapter 3). The figure also shows the projected likelihood of a future satellite mission, *CMBPol*, in the case of $r \approx 0.001$.

The current best limit on r based solely on B-mode data is $r < 0.72$ (at 95% confidence). This result was procured by the *Background Imaging of Cosmic Extragalactic Polarization* (BICEP) experiment in 2010[48]. Like DASI and QUaD, BICEP operated from the South Pole. Better limits may be found if one includes the data sets of other surveys in the analysis. The current best limit is $r < 0.24$ with 95% confidence, based on WMAP7 data, measurements of baryonic acoustic oscillations (BAOs) in the galaxy distribution, and improved measurements of the Hubble parameter[49][3].

In chapter 4 we will look into some of the difficulties observational projects face when searching for B-modes, and we will see that QUIET has many advantageous features to deal with them. Hopes are that QUIET will reach the sensitivities needed to boldly go where no detector has gone before.

[3]Slightly better limits have been found by using WMAP data, BAOs and supernova observations ($r < 0.22$ using WMAP5, $r < 0.20$ using WMAP7), however these estimates are not as good in terms of systematics control, so they are not quoted as the current best estimate.

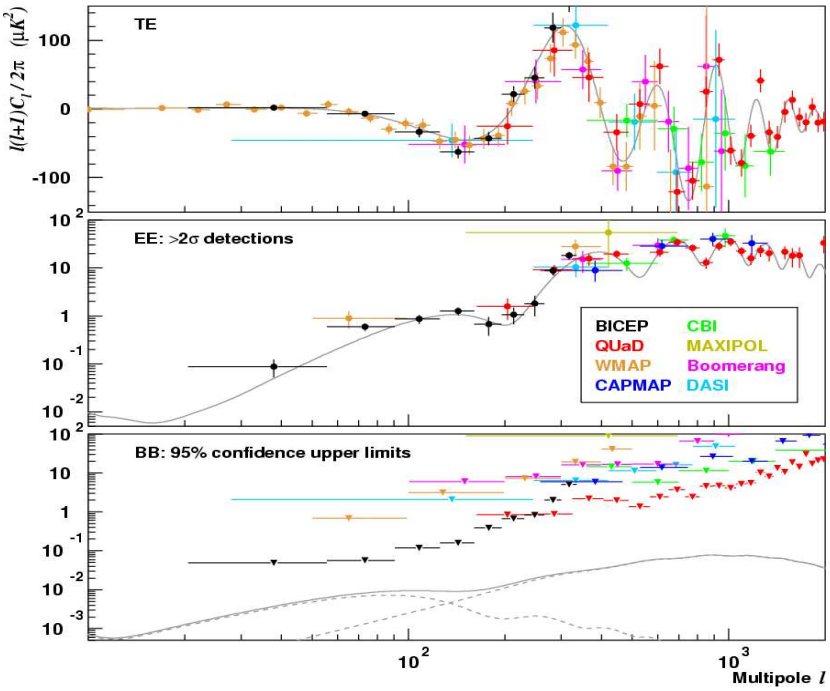

(a) Selected polarization power spectrum measurements as of June 2009. Theoretical B-mode spectra for $r = 0.1$ shown in the bottom panel. Figure courtesy of BICEP team. [48]

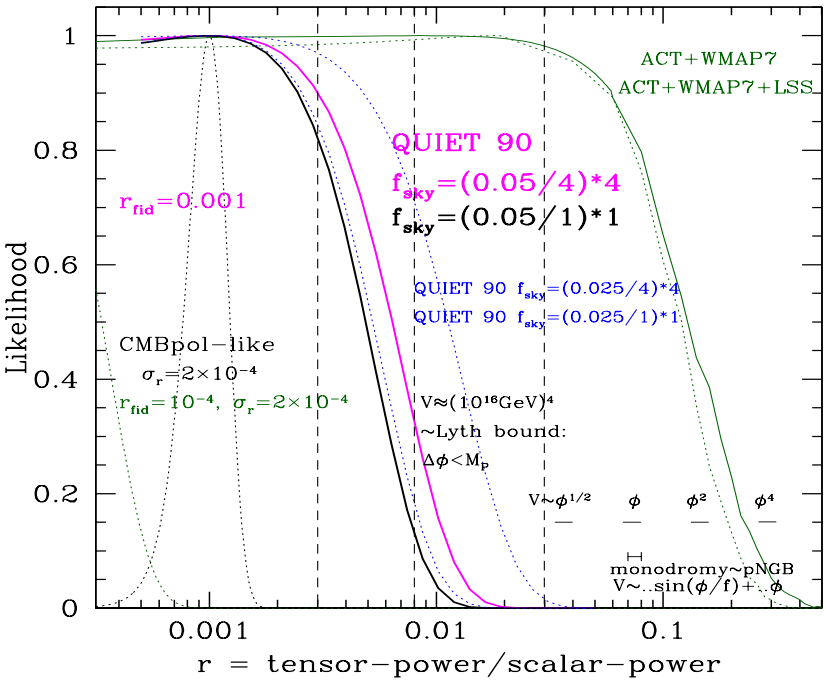

(b) Expected and observed likelihoods for measurements of r, as compared to model predictions. Figure courtesy of J. R. Bond.

Figure 2.2: The current status of the exploration of the polarized CMB and the attempt to measure r

Chapter 3

Cosmological inflation

If you can't explain it simply, you don't understand it well enough.

Albert Einstein

Cosmological inflation is the term we apply to an era of exponential spacetime expansion when the Universe was very young, often stated to occur around $t \sim 10^{-36}$ - $t \sim 10^{-33}$. American cosmologist Alan Guth proposed this mechanism in 1981, as a way of solving certain problems that plagued the then standard hot big bang model[24]. Since its introduction, this concept has grown into a plethora of models, differing in the details, but largely agreeing on the general picture.

Despite having some unresolved issues of its own, such as the conspicuous lack of a physical mechanism for its actuation, inflation is currently unrivalled in the position of go-to theory for cosmologists in need of initial conditions for their models. With good reason: As we will see in the following, the theory performs outstandingly well at providing an environment from which a universe such as our own might develop. However, we are not completely comfortable in the current situation. While the data fit very well with the predictions of an inflationary era, we have yet to find conclusive evidence, and we lack competing theories by which to judge the quality of the fit. As was stated by Scott Dodelson, "...one of the current problems in cosmology is that there is really no viable alternative to inflation."[2, p.139]

Gravitational waves are hoped to remedy this situation. All inflationary theories predict perturbations to the metric in the form of propagating gravitational waves, which would in turn produce large-scale B-mode polarization in the CMB. No other known mechanism could produce these, so finding them would indeed be a breakthrough.

3.1 Introduction to primordial cosmological inflation

A vital part of cosmology is the devising of theories of the Universe's first moments. From these, we derive initial conditions for the equations of structure formation[1], and simulate model universes to compare with our own, as was briefly discussed in §1.3.2. Often, this long and winding road is the only way by which we can test our theories,

[1]We cannot let our simulations run from time $t = 0$, since our equations are not applicable in the limit where $T \to \infty$. We need a theory of quantum gravity to describe these early times, a thing we as yet lack.

since the events we are trying to describe are believed to have happened at energies far surpassing the reach of even our largest particle accelerators.

Cosmological inflation is such a theory (more precisely, it is a class of such theories). A detailed survey of this field would be far outside the scope of an introductory chapter like this one. In the following, I will try to give a brief overview, emphasising some few points that are most relevant to this thesis. Before dealing with the theory itself, however, one should take the time to consider the problems that inspired it.

3.1.1 Dark clouds on the horizon

Physicists mislike fine-tuning. We would rather not have an understanding of the Universe that requires the random aquirement of *just* the right initial conditions from a very large space of possibilities. It is not strictly impossible for our Universe to have evolved from such an origin, but highly improbable. What we want is a mechanism that ensures the development of any random initial configuration into something that we can recognise. Such a mechanism was lacking from the standard hot big bang cosmology as of 1980, meaning that very particular initial conditions was needed for the model to do its job properly. Among the issues calling for fine-tuning were advanced matters arising from theoretical high-energy physics, which predicted detectable frequencies of occurences we have not been able to find evidence of (magnetic monopoles and so-called *topological defects*). While these matters are outside the scope of this chapter, I will briefly discuss the *horizon* and *flatness problems*, which are somewhat more intuitive.

The horizon problem

As nothing can travel faster than light, there is a limit as to how far away two particles can be from each other while maintaining any form of causal contact. Naturally, this limit depends on the amount of time available for transferring information. To quantify this, we define the *particle horizon*, a spherical shell with a radius equal to the maximum proper distance travelled by a photon since time $t = 0$. For a given point in time, all particles separated by less than this distance may be viewed as "knowing of each other", while those further off are causally disconnected.

The proper distance to the particle horizon is found by evaluating the line element for a photon (which we know from special relativity is always zero). We find it to be

$$d_{PH}(t) = a(t) \int_0^t \frac{c\,dt}{a(t')}. \tag{3.1}$$

The standard hot big bang model describes the Universe as being radiation dominated at early times and matter dominated at later times. By solving the first Friedmann equation (1.6) for these cases, we may estimate the size of the particle horizon at the time of last scattering. From this, we may in turn judge the approximate angular size of the areas on the present-day CMB sky that should have correlated properties in a universe correctly described by the standard hot big bang theory. We find it to be of the order of 1° (e.g. 1.6° in [1, p.205]. The precise value depends on what assumptions have been made). But when we look at the sky, we see the same temperature everywhere, implying there has been a thermal equilibrium with only small fluctuations of the order of 10^{-5}, between areas that could never have been in causal contact. For the model to

work, we thus need to demand that the Universe started out at a uniform temperature for no apparent reason. This is the horizon problem.

The flatness problem

By the time of Guth's article, it was generally agreed upon that the Universe must be flat, or at least very nearly so[24]. In §1.1.2 we found that the first Friedmann equation could be expressed as $\Omega + \Omega_k = 1$. For a more or less flat universe, this means

$$|\Omega - 1| = |\Omega_k| = \frac{|k|c^2}{a(t)^2 H(t)^2} \ll 1. \tag{3.2}$$

Applying the results of a matter- or radiation-dominated universe close to flatness, we find that this quantity is in fact increasing with time: $\Omega - 1 \propto t^{2/3}$ for matter, and $\propto t$ for radiation[1, p.202]. In short, flatness is an unstable equilibrium in a standard hot big bang universe. A perfectly flat universe would stay flat, but any deviation from that condition causes the system to diverge towards an open or closed geometry. Again, one might come up with estimates of initial values that would make this problem go away. Guth himself arrived at the conclusion that the initial value of H would have to be fine-tuned to an accuracy of 10^{-55} for the classical model to fit with our Universe[24].

3.1.2 Vacuum energy and the de Sitter universe

The key to solving the above problems - and more, as it would turn out - was found in the *de Sitter* universe model. This classical solution of the Einstein equations, describing an empty universe with a non-zero cosmological constant, was proposed by Willem de Sitter in 1917[50]. Although de Sitter presented his solution as being static, it later became clear that he had used coordinates that were not comoving, and his model was in fact equivalent to an FRW metric with zero curvature and an exponential scale factor[1, p.45].

The first Friedmann equation for a de Sitter universe takes the form

$$\frac{\dot{a}}{a} = \pm\sqrt{\frac{\Lambda}{3}},$$

which is a constant. Since we observe that H_0 is positive today, we choose the positive root, and find $H = \sqrt{\frac{\Lambda}{3}}$ at all times. The solution to the above equation then becomes

$$a(t) = a(t_0)e^{H_0(t-t_0)}. \tag{3.3}$$

As opposed to the radiation- or matter-dominated universe models, this solution has $\ddot{a} > 0$. This accelerating expansion is what solves the flatness and horizon problems, so it will be our requirement for achieving inflation.

In terms of the Friedmann equations, $\ddot{a} > 0$ translates to $p < -\frac{\rho c^2}{3}$. Energy density can not be negative, so this means the pressure must be. As we noted in §1.1.2, the cosmological constant has $p = -\rho c^2$, so it is eligible to be the driving force behind inflation. We often cite *Lorentz invariant vacuum energy* (LIVE) as a candidate for the physical quantity causing the Λ term. This comes from the fact that if one applies a condition

of Lorentz invariance to the energy-momentum tensor for empty space, we get an equation of state with $w = -1$, which is identical to the cosmological constant EOS. Note, however, that one can consider other forms of vacuum energy with more complicated equations of state. Such energies are sometimes referred to as *quintessence*[51].

We consider an epoch between t_i and t_f where the scale factor evolved as in eq. 3.3. The ratio between the scale factor at each end of this epoch is

$$\frac{a(t_f)}{a(t_i)} = e^N \quad \text{where} \quad N \equiv H_i(t_f - t_i).$$

We use this N, the so-called *number of e-foldings*, to quantify how much the universe was inflated. If we assume that the Universe was radiation-dominated before the inflationary epoch, we can use that $H(t) \propto t^{-1}$ to estimate the value of the Hubble parameter during inflation. Using $t_i = 10^{-36}s$ and $t_f = 10^{-34}$ gives $N \sim 100$. This corresponds to an increase of the scale factor of the order of 10^{43}. By comparing theory to observations we find that ≥ 55 e-folds are necessary to solve the horizon and flatness problems[52].

Applying eq. 3.3 to eq. 3.2, we find that $|\Omega(t)-1| = \frac{|k|c^2}{a_i^2 H_i^2} e^{-2H_i(t-t_i)}$ during inflation. In other words, whatever the deviation from flatness was before inflation set in, it will have changed by a factor e^{-2N} at the end of it, effectively eliminating our flatness problem. As for the particle horizon, it will be inflated at the same rate as the scale factor. This follows logically from the nature of the scale factor, but one might easily verify the matter by inserting the appropriate relations for $a(t)$ into eq. 3.1. Given a long enough inflationary epoch, this problem will also be eliminated.

3.1.3 Scalar fields and slow rolling

Although we solved our problems (the monopole and similar problems also disappear, because space gets so heavily inflated the probability of observing these exotic objects become negligible), we are far from done. Apart from a rather vague invocation of "vacuum energy", we have no physical mechanism to drive inflation. Also - rather more critical - there is no way of stopping the exponential expansion. The latter was a standing problem for Guth as well, although his model was not as simple as the one I just sketched. Others improved his initial theory to make it physically viable (e.g. Linde 1982[53]).

In the role of physical mechanism causing inflation, we star the so-called *scalar fields*. In itself, this is a very generic term, implying only that for each point in space we may assign a scalar quantity. The reason for this obscurity is that while the generic scalar field plays its part almost to perfection, we do not have a plausible field theory with which to identify it. In fact, although theoretical scalar fields are treated in quantum field theory, no fundamental scalar field has yet been observed. Guth and Linde ([24, 53]) proposed to use the Higgs field, but today we know that this can not be the case. The constraints on the hypothetical Higgs field are too strict, so it is no longer a viable candidate for driving cosmological inflation[2]. So for the time being, we use this sock-puppet field as a place-holder, until somebody comes up with a better suggestion.

The scalar field is denoted ϕ. We begin by assuming it to be homogeneous, $\phi(\mathbf{x}, t) = \phi(t)$. Since the field must necessarily have energy, it has an energy-momentum tensor, and we might solve the Einstein equations (1.2) for it. For the homogeneous case, we

find for the time-time component[2]

$$\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi), \tag{3.4}$$

and for the spatial components

$$p = \frac{1}{2}\dot{\phi}^2 - V(\phi), \tag{3.5}$$

where the first term is associated with the kinetic energy of the field, and $V(\phi)$ is the potential energy of the field. Invoking energy conservation as given by eq. 1.8, we may use the above relations to find an equation of motion for the field,

$$\ddot{\phi} + 3H\dot{\phi} + V' = 0, \tag{3.6}$$

where $V' = dV(\phi)/d\phi$ and $\dot{\phi} = d\phi/dt$. This is an interesting result. We recognise the form of eq. 3.6 as identical to the equation of motion of a particle moving in a potential, where ϕ is the coordinate of the particle and $3H$ is a friction term. Using this as an analogy for the dynamics of the field is very common in the literature. The scalar field is sometimes referred to as an *inflaton*; this is mostly in the particle-analogy context.

For the field to have negative pressure, we see that it must have more potential than kinetic energy. A possible way of achieving this is to have the field trapped in a false vacuum, with little or no kinetic energy and a potential energy higher than that of the ground state of the system. In this case, the energy density is approximately constant, causing the field to act as a cosmological constant, which we have seen will lead to the exponential space-time expansion. In terms of the particle analogy, this is the situation where the particle is trapped in a local minimum of a potential landscape. The only way for it to escape to the true ground state is by quantum mechanical tunnelling. Guth proposed such a mechanism, but he and others soon realised that the tunnelling method was not effective enough to overcome the rapid expansion of the areas still in the false vacuum state. The solution to this problem is found in the *slow-roll approximation*.

In the slow-roll approximation, we demand that

$$\frac{1}{2}\dot{\phi}^2 \ll V(\phi),$$

which, in addition to assuring we have negative pressure, implies that the kinetic energy of the field is changing very slowly. Again invoking the particle analogy, this translates to the particle slowly rolling down a gently sloping potential towards the true minimum (this is also the origin of the term). Under this approximation we find $\ddot{\phi} \ll V'(\phi)$, which reduces eq. 3.6 to

$$\dot{\phi} \approx -\frac{V'}{3H},$$

indicating that the inflaton rolls at close to terminal velocity. This is often cited as a slow-roll condition. To quantify what we mean by "slow", we define two *slow-roll parameters*, ϵ and η, which constrain the shape of the potential to make sure the inflaton

[2]Here and throughout the rest of the chapter it should be noted that authors' choices of which constant terms to ignore vary a lot. For simplicity, I will aim at avoiding most of the constants.

rolls slowly enough, and that inflation lasts for a sufficient amount of time. The slow-roll parameters are defined as

$$\epsilon \equiv \frac{d}{dt}\left(\frac{1}{H}\right) \propto \left(\frac{V'}{V}\right)^2 \quad \text{and} \quad \eta \equiv \frac{1}{H}\frac{\ddot{\phi}}{\dot{\phi}} \propto \frac{V''}{V}.$$

This definition assures that both parameters vanish in the limit where the kinetic energy of the field is zero. The requirement for maintaining inflation may now be expressed as $\epsilon \ll 1$ and $|\eta| \ll 1$. One will find that $\ddot{a} > 0$ as long as $\epsilon < 1$, so this limit is often given as a definition of the end of inflation.

Once the field has reached the true vacuum state, it will undergo a few damped oscillations before settling at the bottom of the potential well. During this process, the energy released by the scalar field will be transferred to the other fields in the young Universe, in a process known as *reheating*. It is believed that this energy is converted into particles - but that is a different story.

3.2 Quantum fluctuations and the origin of structure

So far, we have looked at the case of a homogeneous scalar field. This is of course not the full picture. One of the great strengths of inflationary theory lies in that it gives us a mechanism for producing primordial inhomogeneities in the matter distribution. These are needed for the growth of structure in the Universe. Observations of said structures may in turn give us evidence as to which, if any, of our inflationary theories are correct. In this section I will endeavour to give a simple account of how inflation causes inhomogeneities. Subsequently we will tackle the issue of how to find evidence of inflation from the CMB.

During inflation, the dominating fields would have been the metric and the scalar field. Matter and radiation were coupled to these, but their effects are negligible as long as the scalar field dominates. Before inflation, the area corresponding to our present-day visible Universe was very small, hot and dense, well within the limits of causal contact, and thus in thermal equilibrium. On this smooth background, there would have been quantum fluctuations. By using first-order perturbation theory, we can model how these fluctuations were inflated by the exponential space-time expansion. As a mode is whisked outside the particle horizon, the corresponding inhomogeneities are "frozen" into the background, and nothing can change them until the regular spacetime expansion brings the mode inside the horizon once more. Due to their origin as quantum fluctuations in an equilibrium environment, these inhomogeneities are expected to be small, statistically isotropic, and to have a Gaussian distribution. Once inflation stopped, the scene was set: The Universe had a pattern of "background" inhomogeneities on various scales, easily described by power spectra, as in eq. 1.15. These power spectra form the initial conditions for our equations of structure formation.

3.2.1 Scalar and tensor fluctuations

Since the metric is a 16-component tensor, its perturbations may take on more than one form. Luckily, the *scalar-vector-tensor decomposition theorem*, first deduced by E. M. Lifshitz in 1946 (see e.g. [54]), states that the perturbations may be decomposed into

three independent forms; scalar, vector, and tensor. These are completely decoupled, and can be treated separately.

Scalar perturbations to the FRW metric may be expressed as

$$\begin{aligned} g_{00} &= -(1 + 2\Psi(\mathbf{x}, t)), \\ g_{0i} &= 0, \\ g_{ij} &= a(t)^2\, \delta_{ij}(1 + 2\Phi(\mathbf{x}, t)), \end{aligned}$$

where δ_{ij} is the Kronecker delta function, and Φ and Ψ are functions giving the first-order scalar perturbation. In contrast to this, *tensor perturbations* have no component in the time-dimension. They take the form

$$g_{ij} = a(t)^2(\delta_{ij} + \mathcal{H}_{ij}),$$

where $\mathcal{H}_{ij}$ is a divergenceless, symmetric, traceless matrix. If we choose our axes such that the perturbations lie in the x-y-plane, it is given by

$$\mathcal{H}_{ij} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{3.7}$$

where the h's are functions of space and time. As for vector perturbations, we will not treat them here, as they are not expected to contribute significantly in main-stream theories.

The scalar perturbations to the various fields, metric, scalar field, matter, and radiation, are all coupled, as one would expect from the decomposition theorem. All are determined by the scalar field perturbation $\delta\phi$. On the other hand, in models with scalar field-driven inflation, the tensor modes do not couple to any other perturbations. The initial power spectra of scalar and tensor modes, as predicted by slow-roll inflation, are

$$P_\Phi(k) \propto \left.\frac{H^2}{k^3\epsilon}\right|_{k=aH} \quad \text{and} \quad P_h(k) \propto \left.\frac{H^2}{k^3}\right|_{k=aH}. \tag{3.8}$$

The expression is evaluated when the mode with wave number k crosses the horizon, which is identified by $k = aH$ (aH is the size of the comoving horizon).

We note the similarity between the scalar and tensor power spectra. The important difference lies in the factor of ϵ in the denominator of the scalar perturbation spectrum, which causes us to expect that the tensor perturbations are by far the weaker contribution.

3.3 Evidence for inflation

We have come to perhaps the most important part of this chapter, which is looking at what observables we can use to test the predictions of inflation. The tensor modes are crucial in this matter, since their existence is a clear and unambiguous prediction of all inflatonary theories. A detection of tensor perturbations would be tantamount to a verification of inflation.

The importance of measuring the energy scale of inflation is second only to making that first detection, since it is chiefly through energy levels that we differ between

inflationary models. By "energy scale", we mean for instance the value of the scalar field $V(\phi)$ at the time where the scales we see in the CMB left the horizon, often denoted by $\phi = \phi_*$. We find that $V \propto H^2$ for a scalar-field dominated slow-roll inflationary universe (through the Friedmann eqns.), so the power spectra in eq. 3.8 carry information on this quantity. While P_Φ depends on both H^2 and ϵ, and thus on both V and its derivative, P_h depends directly on V. In other words, the amplitude of tensor perturbations work as a direct measure of the energy scale of inflation. We also find that the ratio of the two power spectra is very useful, and we call this quantity the *tensor-to-scalar ratio*, r. It is commonly defined as

$$r = \frac{P_h(k_0)}{P_\Phi(k_0)} \propto \epsilon,$$

although some variations occur[3]. The fraction should be evaluated at some specified pivot scale k_0 (a common choice is $k_0 = 0.002\text{Mpc}^{-1}$), chosen so as to minimise the correlations between the amplitude and tilt of the power spectra.

Of course, we have not yet been able to measure neither r nor the amplitude of the tensor modes. However, limits have been set on these quantities based on what observations we do have. As we will see below, there are more quantities that are related to V and its derivatives, and given enough of them we are able to reconstruct, via Taylor expansion, the shape of the potential around the time of CMB mode horizon exit. (See e.g. [55], summary of several sources given in [52].) It has been found that the scalar field potential and r should be related as [52]

$$V^{1/4} = 1.06 \times 10^{16}\text{GeV}(r/0.01)^{1/4}.$$

We should take note of the role of r in this expression. If we should find a value $r > 0.01$, the energy scale of inflation is comparable to that of the particle physicists' Grand Unified Theories (GUTs), which would be an immensely important result.

Another quantity in which we take great interest, is the "field excursion" between the times of CMB scale horizon exit and the end of inflation, $\Delta\phi = \phi_* - \phi_{end}$. In particular, we are interested in the relative size of $\Delta\phi$ and the Planck mass, $m_{Pl} = \sqrt{\frac{\hbar c}{8\pi G}} \approx 2.43 \times 10^{18}\text{GeV}/c^2$. Models that predict $\Delta\phi > m_{Pl}$ are known as *large field* models, and if such a model was confirmed, it would contribute significantly to high energy quantum field theory (perhaps even to the development of a quantum theory of gravity). It has been found that the field excursion should depend on r as [52]

$$\frac{\Delta\phi}{m_{Pl}} \gtrsim 1.06 \times \left(\frac{r}{0.01}\right)^{1/2}.$$

The boundary between large-field and small-field models is found to lie at $r \approx 0.008$. This has become known as the *Lyth bound*, after David Lyth, who first pointed out its significance[56].

[3]It might also be noted that this definition is fairly new. It used to be more common to define r in terms of the angular power spectra, along the lines of $r = C_\ell^{BB}/C_\ell^{EE}$, evaluated at some small ℓ (common choices are 2 and 10). This is not exactly equivalent to the current version, but the change is only marginal, and one may still look to r to get a feeling for the ratio between the B-mode and E-mode power spectra.

Spectral indices and scale-invariance

Considering the power spectra for Φ and h, we see that except for the k-dependence, they are proportional to H, and thus approximately constant. A spectrum where $k^3P(k) \equiv \mathcal{P}(k)$ is constant[4] is known as a *Harrison-Zel'dovich-Peebles* spectrum, or simply *scale-invariant*. We expect scale-invariant primordial power spectra even without inflation, so this is not much of a prediction. However, we remember that in the slow-roll approximation, $\dot{\phi}$ is not exactly zero, and thus H is not perfectly constant. To quantify a possible deviation from this criterium, we often express the power spectra by way of a *spectral index*,

$$P_\Phi(k) = A_S\, k^{-3} \left(\frac{k}{H_0}\right)^{n_S-1} \quad \text{and} \quad P_h(k) = A_T k^{n_T-3},$$

or equivalently,

$$n_S - 1 \equiv \frac{d\ln \mathcal{P}_\Phi}{d\ln k} \quad \text{and} \quad n_T \equiv \frac{d\ln \mathcal{P}_h}{d\ln k}. \tag{3.9}$$

The convention is to define these such that $n_S = 1$ implies scale-invariance for the scalar power spectrum, while the corresponding requirement for the tensor spectrum is $n_T = 0$.

With model-specific expressions for the primordial power spectra, eq. 3.9 can be used to relate the spectral indices to the slow-roll parameters. For a fairly simple single-field model, one treatment finds $n_T = -2\epsilon \propto r$ and $n_S - 1 = 2\eta - 6\epsilon$[52]. The power spectra and the corresponding spectral indices can be translated into predictions on the present-day angular power spectra through transformations like eq. 1.18, and thus we are able to place observational constraints on them. If we should manage to produce constraints on the spectral indices that excludes scale-invariance, it could be reckoned as evidence in favour of inflation. In fact, n_S is one of the six parameters defining the ΛCDM model. As we saw in table 1.1, its current best-fit estimate is $n_S = 0.963 \pm 0.014$, which does exclude the value 1.

Through their connections to the slow-roll parameters, the spectral indices may also be utilised for determining the scalar field potential.

Gravitational waves

Solving the Einstein equations for a metric perturbed like in eq. 3.7 gives solutions for the h's on the form

$$\ddot{h} + 3\frac{\dot{a}}{a}\dot{h} + k^2 h = 0,$$

which we immediately recognise as a damped wave equation. This means that tensorial perturbations to the metric give rise to propagating disturbances, a.k.a. *gravitational waves*. General relativity predict the existence of such waves, but because the effect is very weak, we have yet to observe them directly[5]. The fact that the waves are dampened by the expansion of the Universe hardly helps.

[4] Again, conventions vary. Many authors use an implicit k^3-scaling of the power spectrum. This treatment follows Dodelson[2].

[5] Indirect evidence have been obtained through observations of energy loss from the Hulse-Taylor binary pulsar. Hulse and Taylor were awarded the 1993 Nobel Prize in Physics for this work.

Once again, the CMB comes to our rescue. After inflation, inhomogeneities in the metric were "frozen" until the particle horizon grew large enough for them to be smoothed out. Hence the modes that were super-horizon at the time of recombination, corresponding to $\ell \lesssim 100$, would have left their mark on the CMB. Inflation is the only mechanism we know of that can produce these large-scale tensor modes. However, in the case of the temperature power spectrum, the cosmic variance at small scales makes positive identification of this effect near impossible.

B-modes in the polarized CMB

In §1.3.4, we looked into the decomposition of the polarized CMB into *E-modes* and *B-modes*. The hallmark of an E-mode is that the strength of the polarization changes in a direction parallel or orthogonal to the orientation of the polarization. As it happens, the quadrupole moment of the density perturbation responsible for most of the polarization in the CMB is only capable of producing this type of polarization patterns. Gravitational disturbances are needed to twist the orientation of polarization, forming B-modes. On smaller scales, this can happen by way of gravitational lensing as the polarized CMB photons pass by large concentrations of matter on their way towards us. On large scales, however, only the primordial gravitational waves set up by inflation could possibly cause B-mode patterns in the polarized CMB. Large-scale B-modes are thus deservedly referred to as "the smoking gun of inflation". Hence, any detection of large-scale B-modes, however weak, would count as proof of primordial gravitational waves. Sadly, there are limits as to what we can hope to detect. Estimates have been made suggesting that the smallest observable r is $\sim 10^{-4}$ (e.g. [57]).

In §1.3.2 we introduced the formulae for computing observational and theoretical angular power spectra for temperature, and in 1.3.4 we saw the corresponding observational spectrum for the two polarization components. In light of the inflationary results we have found in this chapter, it is now time to complete the picture by introducing the theoretical polarization power spectra. For polarization component $X \in [E, B]$ where $Y \in [s, t]$ implies either scalar or tensor modes, they are

$$C_{\ell,Y}^{XX} = \int k^2 P_h(k) (\Theta_{P,Y}^{XX})^2 dk,$$

where the full derivation is given by Seljak and Zaldarriaga [6], who also showed that the EB and TB cross-correlation spectra should be identically zero, and that scalar perturbations cannot cause B-modes. The Legendre decompositions of the polarization fields, Θ_P, vary according to polarization component and sourcing perturbation (scalars or tensors). An example, corresponding to $\Theta_{P,s}^{EE}$, was given in eq. 1.19.

Through our theoretical and historical account, we have found justification for why we expect primordial B-modes to exist, and seen the importance of detecting them. The rest of this thesis will focus on the effort towards detecting B-modes with QUIET.

Chapter 4

The Q/U Imaging ExperimenT

My goal is simple. It is a complete understanding of the universe, why it is as it is and why it exists at all.

Stephen Hawking

The Q/U Imaging ExperimenT (QUIET) is, as we saw in §2.3.3, one of numerous experiments in observational cosmology targeting the detection of primordial B-mode polarization. Using groundbreaking new detector technology, the QUIET collaboration hopes to unearth new evidence to bring us closer to a conclusion on the questions of our Universe's very first moments.

QUIET involves research teams from institutions around the globe, with a majority localised in the United States. Several of the collaborators have previous experience from earlier CMB experiments, such as CAPMAP, CBI and QUaD, all of which are projects that have made positive E-mode detections.

The participating group at the University of Oslo is responsible for one of two data analysis pipelines. This pipeline, and the results thereof, are the main focus of this thesis, and will as such be treated thoroughly in due time. In this chapter, I will give a short introduction to the other aspects of QUIET, which will hopefully serve as clarifying background material for the analysis to come.

4.1 Facts and figures

4.1.1 Time and place

QUIET is a ground-based CMB polarization experiment. It is designed as a two-part project, of which the first phase, QUIET-I, was largely intended as a pathfinder for the second phase, QUIET-II[1]. The first phase is now in the process of completing data analysis. Observations were carried out between October 2008 and December 2010, giving more than 10,000 hours of data, divided between two frequency bands (Q and W). The data analysis for the Q-band has been completed, and was published by the end of 2010 [58], while the W-band analysis is still in progress. QUIET-II, to which we will return in chapter 7, is in the planning and design phase, and will hopefully be in operation within 3-5 years, given sufficient funds[59].

[1]There does not seem to be complete agreement as to what the project should be called, as several names are in use. The more popular alternatives are QUIET2, and QUIET90.

The QUIET instrument was deployed at the Llano de Chajnantor Observatory, located on the Chajnantor plateau in the Atacama Desert of the Chilean Andes. Lying at an altitude of more than 5,000 m, this area is reportedly the driest on the planet. The thin atmosphere, dry climate and sparse population makes Atacama one of the prime locations in the world for astronomical observations in the FIR- and microwave frequency range. The observatory is currently operated by the California Institute of Technology (Caltech), although plans for a takeover by the University of Oslo are underway. Llano de Chajnantor has hosted a number of astronomical experiments, including another CMB polarization experiment, the *Cosmic Background Imager* (CBI)[60].

4.1.2 The QUIET instrument

The first issue one meets when looking to study polarized phenomena, is the question of how to extract the polarization from the incident radiation. As described in appendix A, we measure polarization in terms of the Stokes' parameters I, Q, U and V, and these may generally be determined by splitting the radiation and recombining it in a clever fashion.

One of two fundamentally different detection strategies are generally chosen by CMB polarization experiments. One uses coherent[2] amplifier radiometers, which extracts the polarization by first combining the splitted signal, then detecting it. The other employs bolometers, and does the opposite, detecting the total incident power first. Since different methods tend to have different strengths and weaknesses, it is of great importance that both approaches are used in the search for B-modes[59].

To achieve the sensitivity needed to detect the faint B-mode signature, we need strength in numbers: There are limits to the sensitivity a single detector can possibly have, so to reach the levels we need, we have to combine multiple detectors. Radiometers have traditionally been rather large, awkward, and expensive instruments, poorly suited for use in detector arrays. However, a recent breakthrough at the *Jet Propulsion Laboratory* (JPL) has changed the scene. With a footprint of a mere few cm^2 each (5.1cm $\times$ 5.1cm for the Q-band modules, 3.2cm $\times$ 2.9cm for W-band), the QUIET detector modules are easily stacked in large numbers on a single focal plane. Each module combines coherent *High-Electron-Mobility-Transistor* (HEMT) amplifiers, phase switches, and detector diodes. The phase switches and amplifiers are all in the form of tiny *Monolithic Microwave Integrated Circuit* (MMIC) chips, making manufacture relatively fast and easy. Thanks to the revolutionary modules, QUIET is currently the only B-mode experiment using coherent amplifier technology. [59, 61]

The QUIET modules are assembled in two separate reciever arrays, a 19-element Q-band array and a 91-element W-band array. They operate at $\sim$ 44 GHz and $\sim$ 95 GHz, respectively. The light is gathered by a 1.4m crossed Mizuguchi-Dragone dual-reflective telescope, and is then directed into an array of circular, corrugated feedhorns. Each horn leads to a septum polarizer, which splits the radiation into left- and right-handed circularly polarized components. These are then fed into a QUIET module, where the components are amplified and modulated at 4 kHz or 50 Hz[3], before they pass through "hybrid couplers" which combine the "legs" to form four output streams, corresponding

[2] Coherent recievers detect phase as well as intensity.

[3] This corresponds to multiplying the signal by +1 or -1, depending on the state of the phase switch.

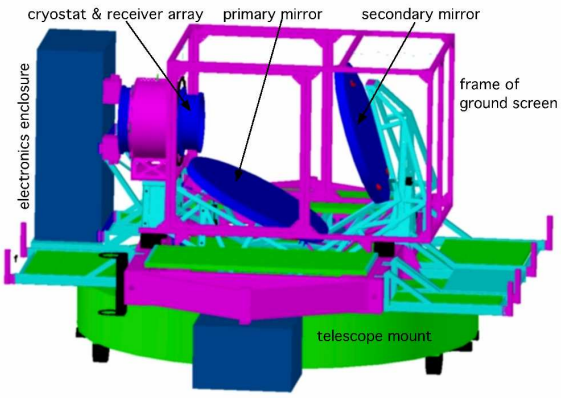

Figure 4.1: A schematic depiction of the QUIET instrument on the CBI mount. Figure from QUIET (2010). [58]

to the Q, U, -U and -Q Stokes' parameters. These are detected by the module's four detector diodes. Measuring Q and U simultaneously is a feature unique for QUIET. [63]

Two of the modules (six for W-band) are dedicated to measuring the temperature anisotropy rather than the polarization. For these TT-modules, the septum polarizers are replaced by so-called *ortho-mode transducers* (OMTs) and "magic tees", splitting the signal from each horn between two modules, and allowing for differential temperature measurements.

To minimise instrumental thermal noise, the feedhorn array, septum polarizers and detector modules are all enclosed in a cryostat, cooling the complex to $\sim$ 20K ($\sim$ 25K for W-band). The telescope, cryostat and an electronics enclosure are installed on the mount previously used by CBI[60]. Figure 4.1 shows the arrangement. Not shown is the ground screen, which encloses the whole system. In the photo of the observatory (fig.4.2), which was taken during the W-band observations, the full ground screen is in place - it is the white box and cylinder hiding the telescope. The upper ground screen (the cylindrical part) was not installed until January 2010, after the full Q-band

Figure 4.2: QUIET observing at Chajnantor Observatory, Atacama Desert, Chile. Image courtesy of the Kavli Institute for Cosmological Physics. [62]

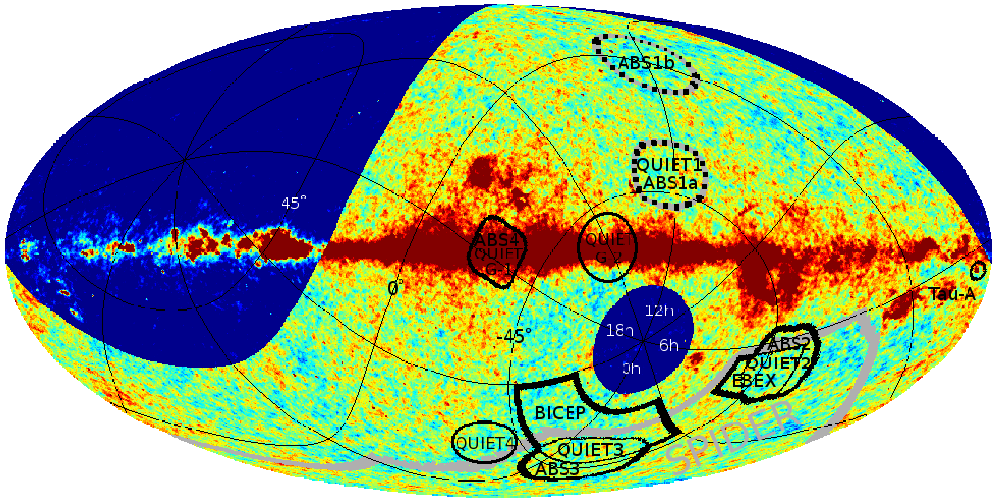

Figure 4.3: Observational patches for QUIET, alongside the fields of a select few future experiments (with the exeption of BICEP, which has already completed its first phase). The part of the sky observable from the Chajnantor site is indicated by the background, which is a WMAP Q-band temperature map. An equatorial coordinate system has also been indicated. Figure courtesy of Ingunn K. Wehus.

season was completed. This means that the Q-band data has sidelobe contamination and ground pickup problems which the W-band data has not.

4.1.3 Patches and scanning strategy

QUIET-I observed a set of four CMB patches, chosen for their low foregrounds and suitable elevation, each covering $\sim 450\,\mathrm{deg}^2$. In addition, two patches in the galactic plane were observed, mainly for purposes of foreground science, and approximately 10% of the observation time of each season was dedicated to observations of various objects, for purposes of calibration. Figure 4.3 shows a montage of QUIET patches alongside those observed by a selection of other experiments, on a background of the WMAP Q-band temperature map. All patches and objects of observation, their positions, and purposes, are listed in table 4.1.

The patches are observed in the following way: The telescope is pointed at the leading edge of a patch, and will scan it by periodic azimuthal motion until the sky rotation causes the patch to leave the telescope's field of view. The mount is capable of scanning at $\sim 6°/\mathrm{s}$ in azimuth, and the scan frequency is typically around 0.1 Hz. We choose to keep the elevation constant during a scan to minimise the noise due to changing atmospheric conditions. By scanning at constant elevation, we restrict our field of view to an atmospheric slice having approximately constant thickness, so any variation across the scanned area will be only the smaller, local variations.

A scan as this lasts approximately 60-90 min. Then the pointing is readjusted in elevation and azimuth, and the process is repeated. The data gathered during one such period is called a *Constant Elevation Scan* (CES). The CES is one of the fundamental units of data used in the analysis. During analysis, we use the term CES to define the

Patch	RA	Dec	Purpose
Patch 2a/QUIET-1	181°	-39°	CMB science
Patch 4a/QUIET-2	78°	-39°	CMB science
Patch 6a/QUIET-3	12°	-48°	CMB science
Patch 7b/QUIET-4	341°	-36°	CMB science
Patch Gc/G-1	266°	-29°	Galactic science, calibration
Patch Gb/G-2	240°	-53°	Galactic science, calibration
Tau A	84°	22°	Calibration (Detector polarization angles, beam, absolute gain)
RCW38	135°	-48°	Calibration of TT modules
The Moon	-	-	Calibration (Detector polarization angles, pointing, relative gain, I→Q/U leakage)
Jupiter	-	-	Calibration of TT modules
Venus	-	-	Calibration of TT modules
Sky dips/CAS	-	-	Calibration (Relative gain, I→Q/U leakage)

Table 4.1: Summary of observational targets for QUIET-I. A total of 3458 (7576) hours of observations were gathered during the Q(W)-band season, of which 77(72)% were dedicated to the CMB patches. [58, 63, 64]

data from a single scan, for all the diodes. These data are in the form of time-streams, and we will often refer to them as *time-ordered data* (TOD). The term TOD is not restricted to the time-streams from single scans, so for the sake of clarity, we often refer to one such single scan, single diode time-stream as a *CES-diode*.

Between each CES, a "sky dip", or *Constant Azimuth Scan* (CAS) is performed by making the telescope "nod" three times, keeping the azimuth orientation constant. This gives us valuable information on the varying conditions of the atmosphere, which is used for gain calibration. Once a week, larger CAS'es are also run, scanning over the full elevation range of the mount (approx. $42° - 80°$).

The CBI mount provides our telescope with three rotational degrees of freedom. Besides the elevation and azimuth axes which we have already mentioned, we have the opportunity of rotating about the optical axis of the primary mirror (the "boresight"), habitually referred to as "deck rotation". A 45° deck rotation was performed once a week troughout the observational season. This is to minimise any possible biasing of the data due to the fixed detector polarization axes which form the reference coordinate system by which the Q and U Stokes' parameters are defined. The combination of deck rotations and the natural sky rotation serves to randomise the detector polarization axes, thus modulating the fixed sky polarization signal. The $I \rightarrow Q/U$ *leakage*, one of the major systematic effects in the Q-band data, was significantly suppressed by this strategy[58, 63].

4.2 Strengths and limitations: The scope of QUIET

The eventual goal of QUIET is to achieve sufficient sensitivity to be able to constrain the tensor-to-scalar ratio r to < 0.01. In §3.3 we saw that this value marks a limit

between two major classes of inflationary models; those with GUT scale energy, and those with less. The Lyth bound could also be within reach. This means that the success of QUIET-II would bring with it either a detection of the fabled B-modes, or a rejection of several significant theoretical models. Both scenarios would be a major contribution to modern cosmology.

QUIET-I do not carry such high ambitions, but is still a project of significant merit, beyond merely paving the way for its successor. The QUIET-I W-band analysis is expected to result in competitive new bounds on the BB power spectrum. The sensitivity of the QUIET modules is such that even the small Q-band array has produced a constraint $r = 0.35^{+1.06}_{-0.87}$, which is the second best constraint in the literature to date[59]. And regardless of the results, QUIET has several unique features that make us a significant contribution to the field of CMB research: The newly developed QUIET modules, the project design giving excellent systematics[4] control, and our choice of frequencies, which are somewhat outside the beaten path.

There are some limits to the precision any project can possibly attain. For Earth-based CMB science, these are mainly dictated by thermal noise, cosmic variance and limited resolution. To quantify the limits to the accuracy with which we can theoretically determine the power spectra, we may use the *Knox formula*,

$$\Delta C_\ell^{XX} = \sqrt{\frac{2}{(2\ell+1)f_{sky}}}[C_\ell^{XX} + w_X^{-1}W_\ell^{-1}], \qquad X \in (T, E, B).$$

This expression, defined analytically by Lloyd Knox[65], tells how the ideal uncertainty of a power spectrum depends on a combination of instrumental effects and cosmic variance: f_{sky} is the fraction of the sky observed, $w_X^{-1/2}$ is the noise per pixel, and W_ℓ is the window function, which tells what area of sky we are observing. The formula implies that the larger the window and the weaker the noise, the smaller the uncertainty, up to a certain limit set by the cosmic variance, which we can never quite escape.

It is the challenge of experimental scientists to design an experiment capable of reaching the theoretical limits, to glean as much information from the object of scrutiny as is humanly possible. The various technical details involved in QUIET noise control and reduction will be summarised in the following. However, procurement of reliable results does not only depend upon the equipment. The method of data analysis is equally significant, and deserves mention in this context (the subject will be treated more extensively in later chapters).

QUIET uses two separate methods of data analysis; both of these have strengths as well as weaknesses, but they complement each other in such a way that the final results are of higher quality than we could have managed with only one. We practice a blind testing policy, requiring the data to pass a number of pre-determined test criteria *before* anyone looks at them. This is to avoid researcher bias, which is an ever-present danger in all parts of science - especially so since we are looking for something we *really* want to find.

[4] By *systematics*, we mean non-random noise that can potentially bias our results. Noise, both systematic and random, is dealt with in more detail in §5.2.1.

4.2.1 Procedural and instrumental factors

Among the foremost of the factors deciding a project's performance is the sensitivity - a measure of how faint an object the detector is able to detect. The QUIET-I arrays were built with the intention of being able to improve upon the existing E-mode measurements and B-mode constraints. In the case of the Q-band array, this meant aiming at sensitivity per horn of 300 $\mu\mathrm{K}s^{1/2}$, 60 $\mu\mathrm{K}s^{1/2}$ for the whole array. Array sensitivity scales as $N^{-1/2}$, where N is the number of detectors, which is why we would like as large an array as we can get. The sensitivity for each detector has to be calculated based on measured quantities, for instance as the ratio of noise to gain. The final values reported for the array sensitivities over the full season were 69 $\mu\mathrm{K}s^{1/2}$ for the Q-band, 85 $\mu\mathrm{K}s^{1/2}$ for the W-band. [66, 58]

The interval of power spectrum multipoles that we can attempt to estimate is limited below by the finite patch size, and above by the finite beam width of the detectors. There are no exact limits in terms of the multipole moment ℓ, but we may use the beam and patch sizes to estimate an interval of reliable detections. For instance, it seems reasonable that an anisotropy mode would have to have a wavelength comparable to or smaller than the angular size of the patch in order for us to be able to detect it with decent accuracy. The QUIET patches are $\sim 20^\circ$ across. This corresponds to $\ell = \frac{360^\circ}{\lambda} \approx 18$, which gives us the general area of our lower limit. The final desicion on what multipoles to include was made empirically, by assessing the magnitude of error bars of power spectra as a function of ℓ.

In the other end of the spectrum, the resolution is given by the size of the field of view of a detector, which is characterised by a *beam function*. The beam, which ideally should be radially symmetrical, gives the spatial distribution of the total power being detected in any one sample. In other words, we cannot say *exactly* where the sampled light came from, because the field of view of each feedhorn is a small, but extended area. We use the FWHM of the beam function as a limit to the field of view, as the amount of light coming from outside that range will be negligible. Thus the beam defines our "pixel size". The QUIET Q-band feedhorn array has a beam FWHM of $27.3'$, which corresponds to $\ell \sim 395$. Similarly, the W-band beam has FWHM of $13.0'$, corresponding to $\ell \sim 830$.

After evaluating error magnitudes for all possible multipoles, the QUIET team has chosen to give power spectrum estimates for the Q-band on the interval $\ell \in [25, 475]$, and for the W-band at $\ell \in [25, 1000]$.

During a complex undertaking like an observational project, there are many pitfalls for the unwary. Any number of things could go wrong, resulting in some form of contamination of the signal. In QUIET, we have to be especially wary of effects that might cause spurious polarization, so an important element of the analysis and validation of our results is a dedicated survey of potential systematics, producing estimates of the severity of the various concieved effects. We quantify signal contamination by way of comparison to theoretical BB power spectra: Given a source of error, what would r have to be for it to have a statistically significant effect on the power spectrum? The tests are performed by simulating contaminated data and feeding them through the data analysis pipeline, thus allowing us to judge the effect. All such potentially dangereous factors are listed in table 4.2.

Besides our careful analysis and cleaning procedures, giving us an excellent degree

Source of systematic error	**Projected effect**
Time-dependence and error of gain	Amp. of all powerspectra, independent of ℓ
Beam shape uncertainty	Amp. of all powerspectra, increasing with ℓ
Telescope pointing uncertainty	Distortion of pol. maps, spurious B-modes
Gain shifts (esp. within CES'es)	Map distortions
Polarization angle uncertainty	E↔B leakage
Instrumental imperfections	I → Q/U leakage (spurious polarization)
Sidelobe contamination	Excess power
Type B glitch (electronics error)	Spurious polarization
Data selection bias	Spurious B-modes

Table 4.2: A summary of potential systematics. The effects are grouped by severity: The first category is confirmed to have an effect on the power spectrum uncertainties. The second category are effects assessed, and found to be significantly below the level of the statistical power spectrum errors. The final group consists of effects that were otherwise evaluated, and found to be unimportant.

of control over potential sources of error, the QUIET project design comprises a number of features contributing to the prospected good results. The compact telescope design, employing Mizuguchi-Dragone optics, has got very good properties with respect to systematics, and a large, approximately circular field of view. HEMT amplifiers are likewise known for their many advantages, such as low noise levels. To assure that instrumental $1/f$ noise[5] do not contaminate the scan frequency and its harmonics, where we expect significant CMB signal contribution, we scan at a high speed, so that the scan frequency is well above the $1/f$ knee frequency. Furthermore, coherent reciever architecture has an inherent advantage with respect to instrumental noise, since this mainly arises during amplification of the signal: We extract the linear polarization by subtracting the two "legs" of the signal, which have been subjected to the same noise-inducing factors. The subtraction effectively cancels most of this instrumental noise, meaning that instruments like ours may be operated at comparably high temperatures.

Finally, QUIET uses a brand new technique called *double demodulation*, in which the signal is modulated twice (by using phase switches, as was mentioned in §4.1.2) before detection, and demodulated after. This approach is very efficient in suppressing spurious instrumental polarization, as well as other instrumental effects and effects of athmospheric fluctuations. It also eliminates the need for additional optical components for modulating the polarized signal. As all components can be sources of noise, it is advisable to keep an instrument as simple as possible. [63]

4.2.2 Observational conditions

While there is a plethora of measures to be taken, variables to control and components to adjust in our instrument and algorithms, there is little we can do to control the heavens. External factors, from local weather to vast, interstellar clouds of gas, are persistently present and annoyingly unpredictable. The best we can do is to avoid what

[5] A very common type of noise, which is inversely proportional to some power of the frequency.

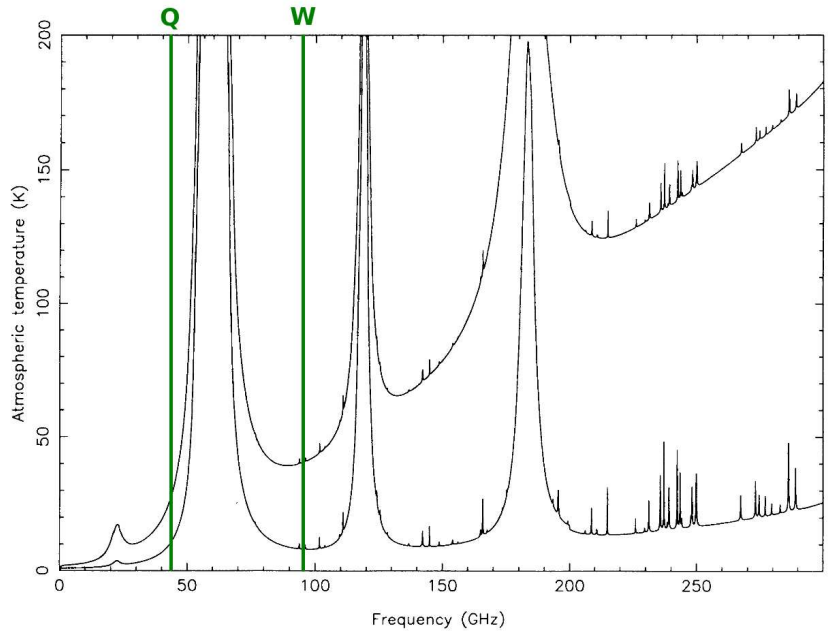

Figure 4.4: Atmospheric brightness temperature at zenith as function of frequency. The lower curve represents a high-altitude observational point, specifically the South Pole, at 3.6 km above sea level, while the upper curve is calculated for a sea level location. The QUIET frequencies are indicated. Figure adapted from Brandt et.al. [67]

we can, and make thorough measurements of the rest, so that we may subtract the effects from our data at a later time.

First of all, we need to make sure the atmosphere is reasonably transparent, by choosing location and frequencies wisely. With the exeption of the "windows" at radio- and visible frequencies, there is generally a non-negligible rate of radiation absorption in the atmosphere, although its value is highly frequency-dependent. To illustrate this, figure 4.4 shows a plot of the atmospheric brightness temperature for the GHz range in which the CMB peaks. The brightness temperature is a measure of how much the atmosphere radiates, and thus how much of an obstruction it will be when we are trying to see what is beyond it. Most prominent here are the absorption lines of oxygen and water vapour at 60GHz, 118 GHz and 182GHz. These frequencies are naturally shunned by ground-based astronomical projects.

By the two curves in the plot, we see the clear advantage of seeking out high-altitude locations. Drier air than what is used in these data ($5\,\mathrm{g/m^3}$) would improve the conditions even more. The Atacama desert has both the low humidity and the altitude, which explains the popularity of the location.

Since the signal we are seeking is a polarized one, we must guard ourselves against non-cosmological polarized sources. To do that, we would preferrably have a detailed map of the polarized sky, giving us the foreground power in all directions. Such maps exist in temperature, but there has not, as yet, been conducted any dedicated full-sky polarization surveys. WMAP produced polarized maps, but these have insufficient sensitivity and cannot meet our needs by themselves. Planck will conduct polarization

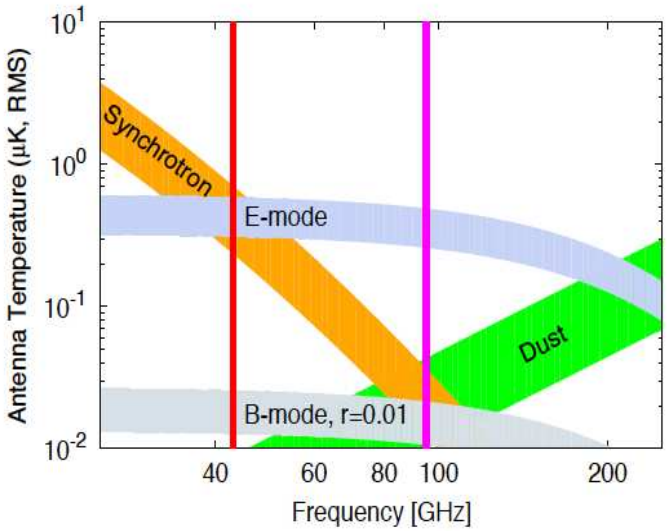

Figure 4.5: Illustration of the frequency dependence of galactic polarized foregrounds, as compared to the CMB E-mode signal and a prediction for the B-modes, as yet undetected. The QUIET frequencies are indicated. Figure courtesy of Yuji Chinone[68].

measurements, but these data are still years away from publication, and we need to make do without them. More subtle approaches are in order.

As far as we know, the dominant sources of polarized foregrounds are dust emission and synchrotron radiation from our own galaxy. We are familiar with the processes responsible for these emissions, and might therefore make predictions about their polarized power based on the available temperature maps. Lacking better data, one will generally choose to look for the CMB at frequencies where the temperature foregrounds indicate low polarized foreground. However, we need to be able to control our results, and the mapping of foregrounds is thus an important secondary goal for many projects (and primary for some). In this process, we have an advantage in that both synchrotron and dust emission have distinctly frequency-dependent intensities, as illustrated in figure 4.5. By probing the foreground power at a frequency where it is dominant, we may use the known spectral index of the foreground to predict its power at the frequency where we are looking for the CMB. Most CMB polarization experiments observe at the dust-dominated frequencies of order 100GHz, since these are frequencies where bolometers are applicable. These are in control of the dust foregrounds, but have no way of estimating the synchrotron radiation. At the present time, the only two projects capable of probing synchrotron power are QUIET and Planck.

As figure 4.5 shows, the E-modes should be detectable without taking polarized foregrounds into account, merely by making good choices of frequency and patch - a fact demonstrated by successful E-mode detections by several experiments. Sadly, this is not the case for the B-modes. Even at the W-band frequency, where the total polarized foreground should be minimal, the foreground power is enough to outshine the expected B-mode power. If we are to have any hope of observing these, we will need excellent foreground control. The QUIET Q-band array is designed for that primary purpose.

Part II

- No time like the present -

QUIET data analysis

Chapter 5

The QUIET data analysis pipeline

> *At the last dim horizon, we search among ghostly errors of observation for landmarks that are scarcely more substantial. The search will continue. The urge is older than history. It is not satisfied and it will not be oppressed.*
>
> *Edwin P. Hubble*

Heavy computational issues are a hallmark of data analysis in modern cosmology. As the engineers come up with increasingly sensitive instruments, the data sets generated go from large to enormous, while at the same time the demands for accuracy are sharpened, as we turn our attention to ever fainter objects. To do the heavy lifting, we design a program complex, which we refer to as a *pipeline*, since - put simplistically - it allows us to insert data in one end, and get results from the other.

In QUIET, we use two independent pipelines, A and B, each with its strengths and weaknesses, in order to ensure the best possible quality of our results. A team at the Kavli Institute for Cosmological Physics (KICP) at the University of Chicago leads pipeline A, which uses a pseudo-C_ℓ framework. Our group, at the University of Oslo, are in charge of pipeline B, which takes a maximum-likelihood approach. Pipeline B has several advantages, for instance allowing for the estimation of cosmological parameters without the use of analytical approximations, but is computationally heavy[58]. Pipeline A is more efficient, and can thus afford to undertake a far larger set of tests. It is pipeline B which is the focus of this thesis, so that when the following treatment refers to "the pipeline", the B is implicit.

In this chapter, I will describe the principal features of the pipeline, with emphasis on justifying the actions taken and explaning the problems met along the way. The flowchart in figure 5.1 gives a schematic overview of the path of the data, indicating the order and input/output of the programs involved. The data analysis relies heavily on some selected statistical tools. For the sake of completeness, relevant definitions and concepts are given in appendix B.

5.1 Data management

The amount of raw data from the two QUIET observing seasons is large. Before doing anything towards analysis, we therefore take some steps to reduce the size of the data, and sort the information in such a way that we can use it effectively.

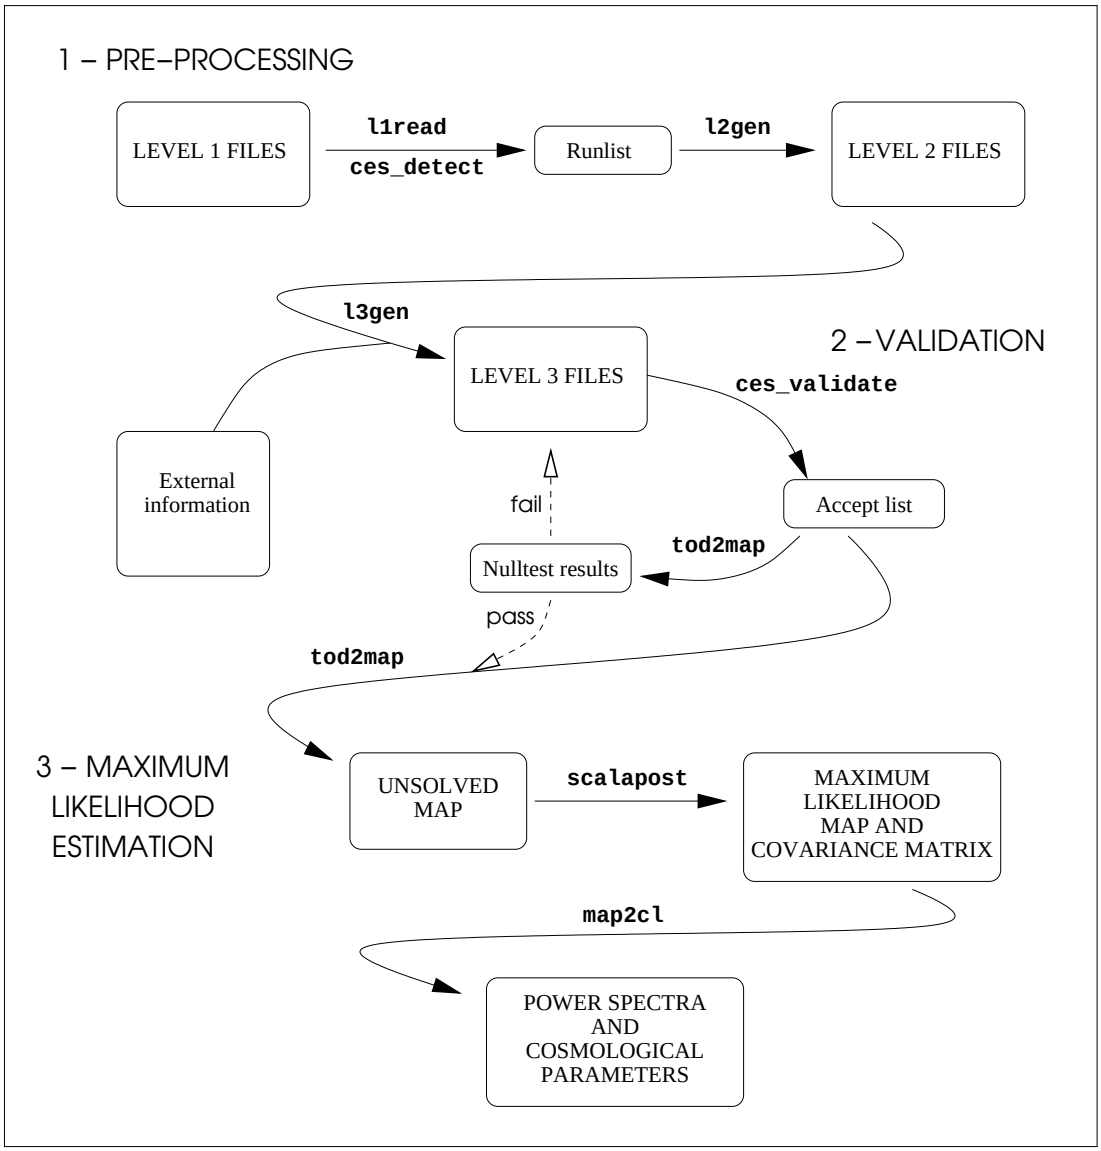

Figure 5.1: Flowchart for pipeline B, starting at the level 1 files. The arrows indicate the running of programs, while the boxes are data containers. Boxes with capitalised labels contain observational data and condensates thereof.

During observation in Chile, all observations and measured quantities, often referred to as *level 0* data, are written to disks (DVD or Blu-Ray) which are mailed to KICP in Chicago. There it is written to *level 1 files*, which are available for downloading. The level 1 files contain 100 Hz data streams (in ADC1 units) and supplementary data. Before preprocessing, the Q-band dataset weighs in at approximately 2-3 Tb, while the W-band dataset is roughly five times larger, ~ 15 Tb.

Converting this to a more manageable format is slow and heavy work, but the software is written so that it should be a one-time endeavour. The data are read (and converted to volts) by the program l1read, and then scanned by ces_detect. As the name indicates, this is a procedure in which we analyse the form of the data to find features which identify the start- and endpoints of scans. For instance, the hallmarks of a CES is regular variations in azimuth, and constant elevation and deck angle. The result of this is a runlist, giving the times for all scans throughout the season, and what level 1 files contain the data (usually several). The runlist is read by the next program, l2gen.

Most important of the tasks performed by l2gen is the demodulation and subsequent downsampling of the data. The former is the final step of the *double demodulation* referred to in §4.2.1. At this point, the form of the samples in the 100 Hz time stream is such that we get the "true" data by transforming it to a 50 Hz stream in which each sample is the difference between two neighbouring samples in the original stream (and then divided by two). This serves to undo the phase switching that was done before detection.

When Fourier transforming a signal sampled at a frequency f, we know that the largest frequency component in the decomposition will be $f/2$. The 50 Hz data stream can then carry information up to 25 Hz. However, due to the size of our beam, we found that we have no sensitivity to such high-frequent modes, so using this high sampling rate would be a waste. Hence the stream is downsampled to 25 Hz, by extracting the average of every two samples.

In addition to this data reduction, l2gen does some smaller tasks, like calibrating the boresight pointing. Finally, the data are stored in a new set of files, the *level 2 files*, one for each scan. These are typically $\sim 100 - 400$ Mb each, well within what is manageable.

5.2 Calibration

During the months of observations, numerous supplementary quantities were recorded, such as weather conditions and matters of telescope pointing. For successful analysis, we need to associate each unit of data with information about the circumstances under which it was procured. This serves to calibrate the data to a common standard, since we will want to add all the pieces together at a later stage.

The recorded data in the level 2 files form the basis from which we will build our knowledge of each scan. Using these, as well as several external sources, we derive all the attributes of each CES that we will need in the subsequent analysis, and finally store the fully calibrated data set in *level 3 files*. The program responsible for this is, of course, called l3gen. Besides computing physical measures, l3gen provides a number of so-called "chi-squares" for quantities that vary over the time span of one scan. The

chi-squares (χ^2) are statistics that are expected to follow the chi-squared distribution. (For a short introduction to this distribution, see appendix B.) Later in this section, we will look more closely at the most important tasks performed by l3gen. But first, here is a list of all the tasks covered by the program:

- Pointing: Individual pointing in galactic coordinates for each feedhorn, and object-centered pointing for the boresight (i.e. pointing relative to particular objects, like the Sun or the Moon).
- Hit pixels: Mapping all pixels on the sky, relative to a Nside=2048 HEALPix map, that the telescope has observed during a scan.
- Scan frequency: Extracts the scan frequency by way of the sampling rate and the boresight azimuth pointing.
- Power spectra: Each CES-diode time-stream is Fourier transformed, and corresponding power spectra are made. These are needed for several of the subsequent items.
- Noise model parameters: A noise model is fit to each CES-diode power spectrum. This is described further in §5.2.1.
- Diode correlations: Computes correlation coefficients between the signals of pairs of diodes, as function of frequency.
- APEX weather stats: Stores relevant weather data from the nearby APEX site, as provided in external files.
- Gain: Diode gain as function of time is modelled. See §5.2.3.
- Diode-specific parameters: Quantities descriptive of a single CES-diode.
  - Type B χ^2: A χ^2-statistic revealing the presence of the type B electronics glitch (also referred to as ADC non-linearity).
  - Weather variables: We compute the standard deviation of the RMS of the TOD, binned in 10 or 30 second intervals. This gives a measure of the variability of the noise on these timescales, which is related to the amount of atmospheric movement (a.k.a. weather) at the time.
  - Single-ces interval χ^2: A χ^2-statistic is computed for the CES-diode power spectra as compared to the aforementioned noise model, in small intervals centered on 3 chosen (spike-prone) frequencies: The scan frequency ($\sim$0.1 Hz), 1.2 Hz and 10 Hz.
  - Sidelobe parameters: Coordinates of the sidelobes, and whether or not they have Sun contamination.
- CES statistics: Single parameters for each CES, valid for the whole array.
  - Time, in MJD (Modified Julian Date) and local time (medial time represents each scan uniquely)
  - Average boresight pointing in horizontal coordinates (az, el, dk)

  - Mean and standard deviation of PWV (Precipitable Water Vapour) and humidity, as provided by APEX
  - Mean of wind speed and ambient temperature, as provided by APEX
  - Mean and standard deviation of temperatures in cryostat and electronics enclosure

- Azimuth filter parameter: Finding the optimal order of the Chebyshev polynomial used in the azimuth filter. See §5.2.2.

- Bandpass filter parameters: Calculates appropriate limits for the highpass- and lowpass-filters for each CES-diode, see again §5.2.1.

5.2.1 Systematics, noise and noise modelling

When we direct our equipment skywards, we want to chart the cosmic microwave background radiation. But where there is a background, there is most often also a foreground. As we saw in §4.2.2, we need to take into account radiating objects such as stars, galaxies, and interstellar matter that will necessarily get in our way. There is also the small problem of the Earth's atmosphere forming a protective blanket between us and the Universe, and the fact that no piece of our complex instrument is perfect. In short, by the time we get our greedy cosmologist's hands on the CMB signal, it has all but disappeared in the pile of garbage coming out of our high-tech and horribly expensive telescope. One could weep.

All scientists doing experiments must deal with the problem of noise. We may define noise loosely as "everything an experiment picks up when looking for something else". Data analysis is all about finding a way through this jungle, hopefully unearthing the treasure of a true signal, while avoiding being eaten by snakes.

We generally group noise effects into two main categories: Systematic effects, and random noise[1]. The systematics are by far the most dangerous kind. These are measurement biases, due, for instance, to faulty calibration, that will keep an estimated property from converging towards the true value. Random effects, being random, are easier to deal with, because of their easily predictable statistical properties. When looking for a signal as weak as the primordial CMB B-modes, one needs a very high degree of systematics control, and excellent methods for cleaning the data.

Some terminology

The data from a scanning telescope will be in the form of discrete time series, of which the vector is a suitable mathematical representation. In the loose definition of noise as "everything that's not the signal" we may decompose the data, $\mathbf{d}$, into a sum of signal, $\mathbf{s}$, and noise, $\mathbf{n}$,

$$\mathbf{d} = \mathbf{s} + \mathbf{n}. \tag{5.1}$$

The vectors may be time-, frequency- or pixel-ordered, as the situation demands. This picture works equally well at representing the data from a single scan, as for a succession if them.

[1] It should be noted that there is a tendency to reserve the term "noise" to the latter category, while the former is referred to as "systematics", or occasionally other, less friendly terms.

In an attempt to quantify the amount of noise in the data, we use the *signal-to-noise ratio*, S/N. As the term indicates, this is the ratio between the signal power and the noise power, and should preferrably be high. ($S/N \sim 5$ could safely be reckoned as high.) The signal-to-noise of a single CES is not, of course, high. Given that a typical single-CES map has a noise amplitude of $\sim 200 - 300\mu$K, while the strongest polarization component, the E-modes, lie at less than a tenth of that, we may safely assume that the signal in each CES-diode is negligible[69]. In the rest of this section, we will therefore focus of the qualities of a single CES-diode, in order to chart the noise effects.

In the time- or pixel-ordered domains, noise, signal and systematics are all in a jumble. Hence we apply a Fourier transform. The discrete Fourier transform of the time-ordered noise vector takes the form

$$n_{t_i} = \sum_{j=0}^{N-1} n_{\nu_j} \exp\left(-\frac{i2\pi\nu_j t_i}{N}\right). \tag{5.2}$$

This is a linear transformation, and may thus also be expressed as a matrix equation:

$$\mathbf{n}_t = \boldsymbol{\mathcal{F}}\, \mathbf{n}_\nu, \quad \text{where} \quad \mathcal{F}_{ij} = \exp\left(-\frac{i2\pi\nu_j t_i}{N}\right) \tag{5.3}$$

As we know, the coefficients n_ν of the transformed noise vector are complex quantities, spanning as many negative frequencies as the physically relevant, positive ones. We would rather look at the power spectrum,

$$\mathcal{P}_n(\nu_i) = n_{\nu_i} n^*_{\nu_i}, \quad \text{where} \quad n^*_{\nu_i} = n_{\nu_{-i}}. \tag{5.4}$$

This power spectrum is the key to sorting out the mess of noise and systematic effects, as the Fourier decomposition splits the effects into independent wavemodes. The power spectrum gives how the noise power is distributed as function of frequency. Hence individual effects will be identifiable by their signatures in the power spectrum, allowing us to remove any dangereous components, and model the rest. Hopefully, this will lead us to a data set that is clean enough that we may extract the elusive CMB signal.

The noise power spectrum

A fairly typical example of a QUIET Q-band CES-diode data stream is shown in figure 5.2. 5.2(a) shows the unfiltered time-ordered data. Most prominent here is the slow oscillatory background trend, to which we will return shortly. In l3gen, all the TODs like this one are Fourier transformed, and a power spectrum is computed according to eq. 5.4. The result is shown in panel 5.2(c).

During first inspection of the power spectrum, certain features stand out. There is a declining trend on the lower frequencies, followed by stabilising at the higher end of the spectrum. This fits a common noise pattern known as the $1/f$ *profile*: The noise power is dominated by a correlated term, decreasing as the inverse of the frequency (hence the term "$1/f$ noise"), and an uncorrelated term, also known as white noise.

$1/f$ noise is caused by slowly varying trends in the signal caused by the electronics, and will thus dominate the low-frequency end of the spectrum. The slow oscillation

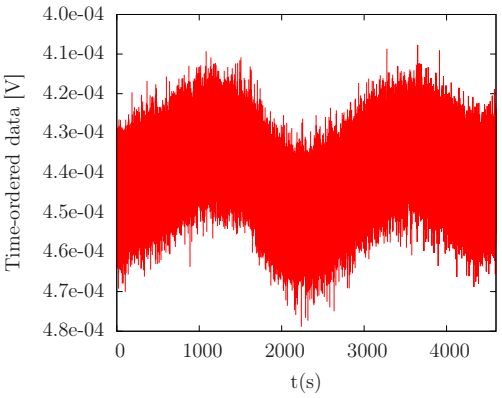

(a) TOD before filtering

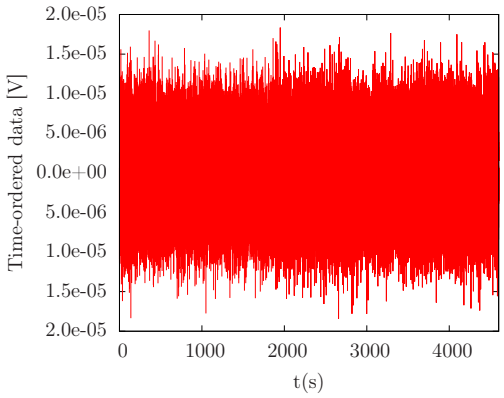

(b) TOD after filtering

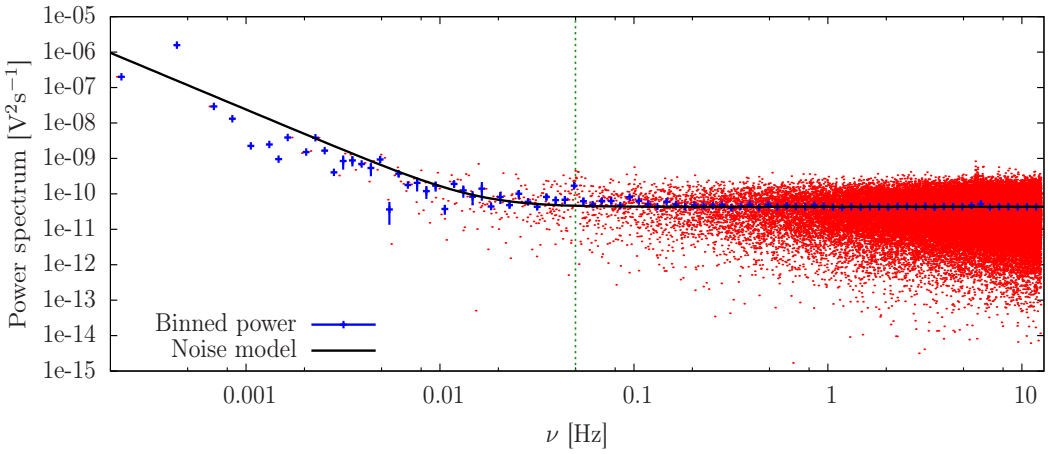

(c) Fourier domain power spectrum, shown with binned power and scan frequency at $\nu = 0.05$ Hz. Noise model parameters: $\sigma_0 = 6.556 \cdot 10^{-6}$, $f_{knee} = 1.582 \cdot 10^{-2}$, $\alpha = -2.289$

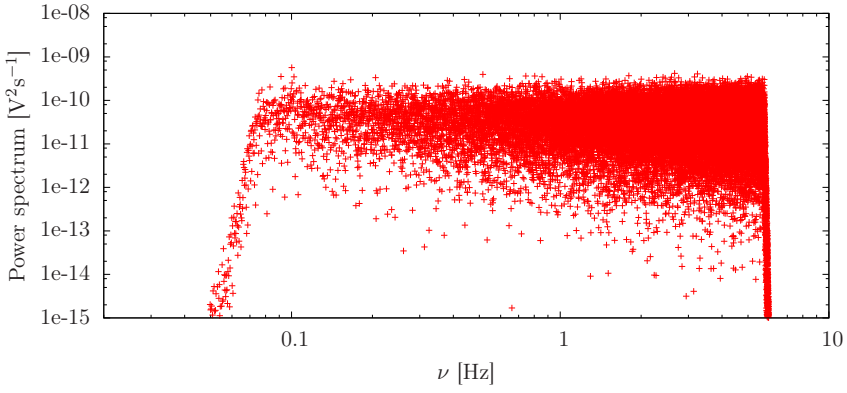

(d) Power spectrum after application of apodized filter. Filter parameters: $\nu_{\text{high}} = 1.5\nu_{\text{scan}}$, $\alpha_{\text{high}} = -30$, $\nu_{\text{low}} = 5.737$ Hz, $\alpha_{\text{low}} = 300$

Figure 5.2: Example of data from the Q-band season: Patch 4a, CES 2160, module 7, diode 0.

we saw in panel 5.2(a) is an example of this[2]. It is far too slow to be the signature of something on the sky, so we may be quite sure we do not want it in our data. Moreover, if we do not remove it, such slow variations will show up as unphysical stripes in our maps. For instance, if we compare the instrument response at, say, $t_1 = 1000\,s$ and $t_2 = 2000\,s$, we find that there is quite a large difference. At these times, the module field of view is sweeping over two different parts of the sky, due to sky rotation. When this TOD is translated into a map, we will see the pixels in the sky slice observed at $t \sim t_1$ as warmer than average, and vice versa for those observed at $t \sim t_2$. Naturally, we want our maps to resemble the true sky as closely as possible, so filtering out the $1/f$ noise is an evident necessity.

The white noise "floor" characterising the higher frequencies in panel 5.2(c) is the signature of random, thermal noise, caused, for instance, by the thermal movements of charge carriers in electrical circuits. The white noise should follow a Gaussian distribution with mean zero (a consequence of the Central Limit Theorem, see §B.3.1). As we noted in §4.1.2, all parts through which the signal must pass before detection are enclosed in an environment cooled to about 20K. This reduces the white noise level, but we can never eliminate it fully.

We fit a model to the power spectrum based on the $1/f$ profile, parameterising the noise in terms of three quantities σ_0, f_{knee} and α. This is done in steps: First we calculate the RMS value of $\mathbf{n}_\nu$ on a limited range of frequencies (normally, on [2 Hz - 4.5 (9.5) Hz] for Q(W)), to find σ_0. Then we mask out certain frequencies known to contain irregularities, and fit the $1/f$ profile with f_{knee} and α. If f_{knee} is found to be large enough to have interfered with the estimation of σ_0, the latter will be adjusted accordingly.

The analytic form of the model is [69]

$$N(\nu) = \sigma_0^2 \left[1 + \left(\frac{\nu}{f_{knee}} \right)^{\alpha} \right] \tag{5.5}$$

Here, σ_0 is the RMS value of the white noise, f_{knee} is the point on the spectrum where the uncorrelated noise starts dominating over the $1/f$ trend, and α is the spectral index of the $1/f$ noise, determining the slope of the curve. A model of this form is plotted atop the sample spectrum in panel 5.2(c). In the ideal case, this function should account for the general behaviour of the noise, all else being random fluctuations about it. The white noise level and knee frequency are detector-specific quantities, remaining fairly constant throughout the season.

Looking more closely, we find that there are some features of the noise power spectrum that disagree with the model. This is most easily seen by comparing the binned power to the noise model. In this example, there is a significant spike around the scan frequency, and there is also some slight spiking tendency at 6 Hz. The scan frequency spike is mainly caused by the telescope picking up signal from stationary objects. Ground-contamination in the sidelobes is significant in the Q-band data, and less so in the W-band, due to the presence of the upper ground screen. Slowly changing weather patterns may also contribute extra power at or near the scan frequency. Harmonics of

[2]Interestingly, in this case we can easily recognise the point in the power spectrum corresponding to this background wave. By visual estimate, we judge the frequency of the mode to be $\sim 4 \cdot 10^{-4}$ Hz. In panel 5.2(c) we find that the point with by far the most power (remembering that the scale is logarithmic) is indeed at around this frequency.

the scan frequency spikes may also appear. Spikes at 6 Hz and above ($>$ 10 Hz for W-band) are common, although there was only a hint of it in this particular scan. The origin of the high-frequency spikes is a little unclear. Electronic factors, such as noise aliasing from the power-line frequencies of 50-60 Hz, are likely to be at least part of the cause[66, chpt.4].

At this point, one should note upon the role of the normal and χ^2-distributions. We are familiar with the fact that the Fourier transform of a normally distributed function is itself a normal distribution. This implies that the Fourier coefficients of our transformed noise data are normally distributed variables, and it follows that the noise power spectrum, consisting of squares of such variables, is χ^2-distributed. Hence we will be using this distribution when testing our hypotheses concerning the distribution of the noise. Since

$$\mathcal{P}_n(\nu_i) = n_{\nu_i} n^*_{\nu_i} = \Re(n_{\nu_i})^2 + \Im(n_{\nu_i})^2,$$

each point in the power spectrum is drawn from a χ^2-distribution with 2 degrees of freedom. Combined with the choice of logarithmic axes, this is the cause of the apparent vertical "widening" of the power spectrum towards the high-frequency end.

The apodized filter function

Later on in the pipeline, we will coadd the data from each CES to produce maps of all the patches. As long as our noise is purely random, this process will diminish the noise power, while the small, but constant contribution from the CMB signal will add up, eventually amounting to a nice, large signal-to-noise ratio. On the other hand, systematic noise contributions will not cancel, and might contaminate our signal and bias our results. To avoid this, we apply filters. The form of the noise power spectrum strongly suggests using frequency filtering, that is, removal of troublesome frequencies before the data are transformed back to the time domain. This raises the question of signal loss: What do we lose by removing parts of the data?

Below the scan frequency, there is little signal. There may be some traces of large-scale modes, but these will be too large for us to constrain properly. The scan frequency itself does have significant signal. As we noted in the previous section, objects that are stationary with respect to the telescope will be detected at the scan frequency and its harmonics, while signal from objects that move slowly will be smeared out around these frequencies. The CMB is of the latter type, due to the rotation of the sky during the scan. So are weather effects, while ground pickup will fall in the former category. Because of these troublesome factors, we have no choice but to tune our highpass filter so as to cut above the scan frequency, even though this causes loss of CMB signal. However, we find that the fraction of signal lost is acceptably small compared to the total signal contained in the higher harmonics[66, chpt.4].

As for the high-frequency end, we found in §4.2.1 that the beam size limits what multipoles we are sensitive to. Our fastest scans have scan frequencies of $\sim$ 0.1 Hz, giving an average velocity of $1.5^\circ \mathrm{s}^{-1}$. A signal detection of 4.5 Hz then implies a mode of wavelength of $\sim 0.33^\circ$, which is too small to be detected with the Q-band array beam size of almost 0.5°. (The argument is identical for the case of the W-band, where the lowpass filter frequency is set to 9.5 Hz.) Consequently, a lowpass filter set to suppress modes above 4.5 Hz will only remove frequencies that we are incapable of observing, and will not cause us any significant loss of signal[70, chpt.6].

Finally, we know that removing data will cause the signal-to-noise ratio in our final maps to be lower than it could ideally have been. This is a price we are willing to pay to get rid of the dangereous systematic effects.

We employ an *apodized* filter, which is a band-pass filter with smooth edges. Its mathematical form is

$$F(\nu) = \left[\left(1 + \left(\frac{\nu}{\nu_1}\right)^{\alpha_1}\right) \left(1 + \left(\frac{\nu}{\nu_2}\right)^{\alpha_2}\right) \right]^{-1}, \tag{5.6}$$

where the two sets of indices signify a highpass component and a lowpass component, respectively. Figure 5.2(d) shows what the power spectrum looks like after this filter has been applied. We note that the edges are slanting rather than vertical. As we know, finite Fourier series are not particularly good at representing step functions, so had we used a "top hat" filter, we would have caused spurious harmonics in our filtered TOD.

We find the filtered TOD by applying the inverse of eq. 5.3 to the filtered power spectrum, and arrive at the data stream shown in panel 5.2(b). The removal of the slow oscillation is evident. Many more modes have been removed from the data, but this cannot be seen by eye. What we do see, is that the filtered TOD is centered on zero. This is because the average of a Fourier transformed function is stored in the $\nu = 0$ Fourier coefficient, and will thus be cut by our highpass filter. This is fortunate. The average instrument response varies a great deal through the season, due to shifting ambient factors like temperature. This will drown out any measurement of the average background, rendering us unable to constrain the average sky temperature, while at the same time causing stripes in the maps by the same mechanism as the $1/f$ noise would.

During the original Q-band analysis, all CES-diodes were filtered with the same function. This had parameters $\nu_1 = 2\nu_{scan}$, $\alpha_1 = -40$ for the highpass filter, and $\nu_2 = 4.5$Hz, $\alpha_2 = 200$ for the lowpass filter. Prior to the W-band analysis, the pipeline was improved on several points, including allowing for individual tuning of the filter parameters to each CES-diode, using the noise model parameters. This way, the filters should be able to remove all systematic features that are detectable in a single CES-diode, while keeping as much data as possible.

5.2.2 Azimuth filter

As mentioned in §4.1.2, the QUIET Q-band observations had significant ground pickup in the sidelobes, because the original ground screen failed to sufficiently inhibit these signals. An additional upper ground screen was ordered as soon as this was realised, but it was not ready to be installed until the Q-band observations were finished. To counter this problem, a template model of the ground signal was accumulated throughout the season. This ground model was subtracted from the data during analysis. As part of the improvements to the pipeline, we have since replaced this strategy with the technique of *azimuth filtering*, which should be a more flexible and robust way of dealing with this type of contamination.

Since the sky moves with respect to the telescope as we scan, the stationary CMB signal will be smeared out in azimuth. Thus we expect the observed, uncontaminated data to be uniform as function of azimuth angle. We exploit this fact when we target sidelobe contamination without making use of a ground model. In l3gen, the TOD is binned in azimuth, and a Chebyshev polynomial is fit to the corresponding function.

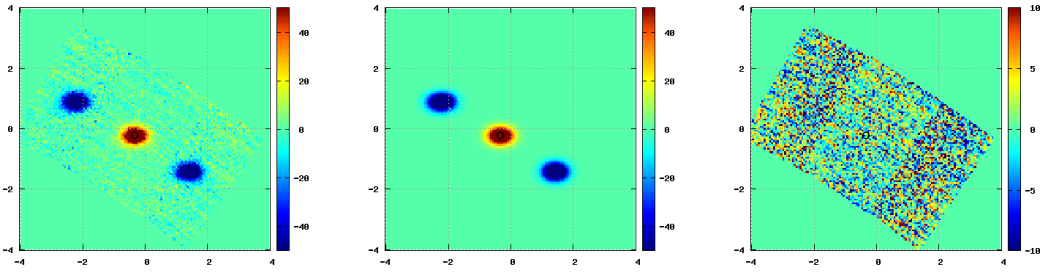

Figure 5.3: Sample map from point_scan. Left panel shows differential-temperature map of Jupiter, middle panel shows simulated model, and right panel shows the difference. Note that the scale on the right panel is four times that of the other plots.

We compute χ^2's to assess what order of polynomial gives the best fit. This function constitutes the azimuth filter. During mapmaking (to which we will turn in §5.5), we remove this function from the data, and make corresponding corrections to the covariance matrix using the Woodbury matrix identity. This identity is also used in the subsequent process of *eigencutting*, which is treated in §5.5.3, hence I will not elaborate on it here.

5.2.3 Gain modelling

The gain of the telescope is a measure of what voltage response we get when looking at a source of a given temperature. In the literature this quantity is often referred to as *responsivity*. Knowing the gain is essential, since we cannot translate between voltages in the instrument and temperatures on the sky without it. Gain determination is thus a vital piece of the calibration puzzle for all experiments in radio astronomy.

Gain is diode-specific and dependent on the temperature in the electronics enclosure. Various data gathered throughout the season are used to study these dependencies, so that we may produce reasonable gain models for the analysis. The absolute gain is measured by making observations of sources with known temperature throughout the season, and comparing that to the instrument response. The targets used for this were listed in table 4.1. For the temperature modules, Jupiter is the main reference, while observations of Venus and RCW38, a dense star cluster, provide additional controls. For polarization, our chief source is Tau A, also known as the Crab Nebula. This famous supernova remnant is also a pulsar wind nebula, and the strongest polarized source on the microwave sky, making it a popular target for calibration[71]. We calibrate the absolute gain of the central module by observing this source. The rest of the array gains are determined relative to this one, using chiefly sky dips, but also observations of the Moon and Tau A. [58]

One of our tools for calibrating gain is the program point_scan, which iteratively fits between a fiducial point source and observations of that source. The program uses as basis an object with known position, temperature and extent. It then proposes possible values of the telescope's corresponding pointing, voltage response and beam shape. A TOD is fabricated for each set of values, to represent what we would have observed if those values had been found to be correct. We find a χ^2-statistic describing

the difference between each model TOD and a real one. This process is repeated until the χ^2 is minimised. A sample map from a temperature calibration run for Jupiter is shown in figure 5.3. The characteristical pattern with three spots is a result of the two differential-temperature horns, where each module registers the positive signal of one horn, and the negative signal of the other. Combining the two causes a pattern such as this one.

In l3gen, diode gain is modelled as function of time, using the calibrated values as basis. A set of diode-specific gain model parameters are provided via input file. This file also specifies which gain model to use. The pipeline currently uses two different ways of modelling the gain, and we choose a method based on the available information. For the Q-band data, we use a gain model format which relies on recorded instrument temperatures. We also have the possibility of modelling the gain as being a piecewise linear function of time, the simplest case of which gives constant gain.

5.3 Data selection

The next step in the pipeline is the selection of data. There are a number of factors that we know can make our observations unreliable in one way or another. Guided by the quantities we procured during the calibration, we evaluate each CES or CES-diode, depending on the nature of the trait under scrutiny. Judgement is passed by way of a set of tuned criteria. This process is performed by the program ces_validate, the product of which is an *accept list*, assigning to every CES-diode a sentence in binary form: 1 means it is accepted for inclusion in the further analysis, 0 means rejection. Below is a summary of the factors by which we make our choices.

Static/input cuts: Cuts dead diodes or missing data files, and any other specific cuts, such as the scans performed during a certain period in July 2009, when there were problems with the hardware calibration. The affected CES'es are known to be void. If issues arise with known scans, these may be cut manually by entering them in an input accept list which the program reads. Scans given in this way will be omitted from the validation process.

Gain cut: Cuts CES-diodes with zero gain.

Elementary cuts: Simple "sanity checks", cutting CES-diodes if their parameters are unphysical or outside the valid range. At present, this amounts to testing that the noise model parameters σ_0 and f_{knee} are positive, and that f_{knee} is less than 100 (this is the default starting value used by the noise model fitting procedure. A value larger than this means the fit has failed).

APEX cuts: Cuts a CES if the registered PWV (*Precipitable Water Vapour*, the total amount of water in a vertical column of the atmosphere, had it fallen as rain), as provided by the APEX weather statistics, is above a certain threshold.

Type B cut: The type-B glitch was an effect of a discontinuity in the electronics responsible for the translation from analogue to digital signal. This caused significant spurious noise in individual diodes[70], and needs to be corrected for. l3gen

computed a χ^2-statistic by way of fitting a linear model to the heavily downsampled TOD. The glitch would cause poor model fit, giving a large χ^2 value for the given CES-diode. Should this test fail, the CES-diode is cut.

Weather cuts: Although operations are shut down during periods of very bad weather, some scans are taken under conditions that may or may not have been too poor. Unruly weather conditions cause more atmospheric movement, resulting in higher levels of noise in the data. To control this, we compute a diode-specific 5σ confidence interval for the distribution of the weather statistics calculated during calibration, for all the CES'es. Those CES-diodes falling outside this interval are cut. For additional security, we check the total share of accepted diodes for each CES. If too many have been cut due to poor weather, we deem the credibility of the rest of the CES to be low, and we reject the whole thing.

Noise model parameters: In a similar fashion to the weather statistics, we compute 5σ confidence intervals for the distributions of the noise model parameters σ_0 and f_{knee}, to which each CES-diode is compared. In the case of the knee frequency, it is also compared to an absolute threshold value of 1 Hz or 0.2 Hz, for temperature and polarization diodes respectively.

Sun cut: Around the edges of the telescope's field of view, radiation from outside the area we wish to detect may sometimes enter the optical path. The regions available for this sort of unwanted detection are referred to as "sidelobes", and they need to be carefully monitored, since powerful sidelobe signals may cause contamination in our data. If the Sun is within the sidelobes, the affected CES-diodes are cut.

FFT cuts: We test our hypotheses about the noise distribution, as described in §5.2.1, by comparing the χ^2-statistics of the CES-diode power spectrum to given thresholds. All of these statistics are computed for a sufficiently large number of samples that the χ^2-distribution is approximately normal. All are scaled to fit a standard normal distribution ($\frac{X-\mu}{\sigma} = Z$), since this makes the testing very simple: The value of the scaled statistic is equal to its deviation from the mean, in units of the standard deviation. The acceptance thresholds are given in the same unit.

Of this group of tests, the most important (and hence the one with the strictest threshold) is the test for the filtered data, since this is the part of each CES-diode we will use in mapmaking. We supplement this with testing other intervals: The unfiltered low and high frequencies (< 0.2 Hz and > 10 Hz), and narrow intervals about the scan frequency and other spike-prone frequencies (and their harmonics), including a sweep of the CMB frequencies for spikes anywhere.

TOD cuts: As we have noted earlier, we expect the TOD itself to be normally distributed, so we compute a χ^2-statistic to test this (using downsampled data so as to avoid any correlation between closely spaced samples). Some TODs have sharp spikes that would not necessarily cause the χ^2-test to fail, so we also test the value of the largest outlier of each. And finally, we compute a χ^2 for the distribution of the TOD as function of azimuth angle, as this should, given observations of random noise and weak CMB signal, be fairly uniform.

Criterium	**Threshold P (T)**	**Danger**
APEX PWV	5 mm	X
Type B χ^2	5 (10) σ	
Weather statistic	5 (10) σ	X
Weather CES limit	40%	
σ_0, f_{knee}	5 σ	
Max. f_{knee}	0.2 (1) Hz	
CMB freq. χ^2	4 (10) σ	X
Low freq. χ^2	10 (100) σ	X
High freq. χ^2	20 (100) σ	X
Scan freq. χ^2	10 σ	X
1 Hz and 1.2 Hz spike χ^2	20 (100) σ	X
General spike χ^2	7.5 (20) σ	X
TOD χ^2	4 (10) σ	X
TOD outlier	7 σ	X
Azimuth χ^2	4 (7) σ	X
Sun cut	all	X
Accept ratio	40%	

Table 5.1: Table of threshold values used for data selection. The rightmost column gives which cut criteria are considered when we evaluate each CES before the accept ratio cut.

Accept ratio cut: After all the other cuts have been performed, we check the ratio of rejected to accepted diodes in each scan. We cut the whole CES if too many diodes have been rejected for "dangereous reasons". These are reasons that can be expected to have affected all diodes in a CES. Table 5.1 indicates which cut criteria are considered "dangereous".

Many of these cuts depend on threshold values which are supplied via input parameter files. The calibration, validation, and null testing steps must be done repeatedly in order to find suitable values for these thresholds. We need them to be just right; not too high, lest we should cut more data than is necessary, and not too low, possibly allowing systematic noise to sneak through our defences. The current standard threshold values for the QUIET data analysis pipeline is listed in table 5.1. The thresholds are the same for Q- and W-band, although some of the intervals on which the various statistics are computed are band-specific. Also, some thresholds are different for the temperature diodes, since these have other noise-related properties.

5.4 Validation by null tests

Before moving on to make maps and power spectra, we need to validate our data, to make sure the calibration and cutting process has worked according to plan. To avoid researcher bias, we use a strategy by which we can screen the data for bias without looking at proper data products. A set of tests and criteria for pass or fail are decided upon, and no maps or power spectra are produced before the full set of tests are passed.

Jackknife	**Code**	**W**		**Q**	
		T	P	T	P
Modified Julian Date	mjd	X	X	X	X
Module position on array (Center or rim)	inout	X	X		X
Type B χ^2	typeb	X	X	X	X
Scan synchronous signal (Scan freq. χ^2)	sss	X	X	X	X
Noise power spectrum χ^2: 10 Hz interval	10hz	X	X		
Boresight pointing: Elevation	el	X	X	X	X
Deck rotation angle (in 45° sectors)	dk	X	X	X	X
Boresight pointing: Azimuth	az	X	X	X	X
Electronics enclosure temperature	tenc	X	X		
Rate of change of enclosure temperature	dtenc	X	X		
Cryostat temperature	cryo	X	X		
Rate of change of cryostat temperature	dcryo	X	X		
Gain	gain	X	X	X	X
Precipitable water vapour	pwv	X	X		
Wind speed	wind	X	X		
Ambient temperature	tamb	X	X		
Noise model knee frequency	fknee	X	X	X	X
Noise model RMS (σ_0)	sigma0	X	X	X	X
Module centre frequency	modfreq	X	X		
Module Attachment Board (MAB)	mab	X	X		
Weather statistic for 10s intervals	weather	X	X	X	X
Stokes parameter Q or U	qu		X		X
I→Q leakage	leak		X		X
Temperature leakage*	tleak	X			

Table 5.2: List of jackknives for the QUIET null test suite used in pipeline B. Less statistics were recorded in the Q-band season, hence the smaller set of tests.

This process is known as *null testing*, and it is a vital piece of the QUIET analysis framework. The main goal for validation is to make sure no significant systematics remain in our data set. There are many factors that can cause bias in our data, and correspondingly many ways biases can cancel each other so we do not see them. The null testing procedure can overcome this obstacle by testing the hypothesis of randomness in many different ways at once.

If a data set is unbiased, it holds only random noise and the CMB signal. Since the CMB is constant in the map, splitting the data set in two and subtracting the one half from the other should result in a map of pure, white noise. One such split with subsequent analysis of the difference is what we refer to as a null test. We list all the factors we would like to check for bias, and for each design a rule of where to do the split. For instance, if we want to know if there are changes over time that may cause bias, we group the data by the time of observation - before or after the median value. Such a rule is called a jackknife. Currently, we have a set of 23 jackknives in our null test suite for W-band, and a subset of 14 used for Q-band. A smaller set still is used for temperature. (The *tleak* test is a new addition which I have not yet been able to

test, as the Q-band data currently do not contain the necessary information to make this split.) The null test procedure is carried out by the program tod2map.

The process goes like this: First, we sort the data into two sets by rule of the current jackknife, and create binned maps of each half. We supply the program with a stack of simulated noise-only data sets, all of which are split the same way as the true data. We subtract all the pairs of half-maps from each other to get what we refer to as a "diffmap", and compute a binned power spectrum of the spherical harmonic transform of each diffmap[3]. Each power spectrum is then normalised according to the mean and standard deviation of each bin, making sure all bins are distributed according to the standard normal distribution. Finally, we compute χ^2-statistics to compare the data to the simulations. For each jackknife, the χ^2 which is the sum of the squared elements of all the pseudo-C_ℓ's of the data is compared to the distribution of the corresponding χ^2's from all the simulations, by way of a *Probability To Exceed* (PTE). The PTE gives the probability of getting a χ^2-value for the data which is larger than what is observed, under the null hypothesis that the diffmap is pure white noise.

Finally, we compute three statistics for the entire set of jackknives, which are also compared to the simulations in the same manner as the individual tests. These are total and maximum χ^2, and the *mean chi shift*. The χ for each jackknife is made by summing all the pseudo-C_ℓ's, without squaring them[4]. In the case of an unbiased diffmap, the χ should, of course, be close to zero. The mean chi shift is, as the name implies, the mean deviation from zero of all the χ's, and gives us a measure of the bias of the data.

For the null test suite to pass, we demand that all the PTE's lie within the interval $[0.04, 0.96]$. In preparation for the W-band analysis, the full null test and validation machinery of pipeline B has been automated.

5.5 Mapmaking

Once our data are sufficiently clean and have passed the null test suite, we are ready for the next big step: Mapmaking. In principle, one could produce power spectra and our final goal, the cosmological parameter estimates, directly from the TODs. Given the size of the data set of every modern CMB experiment (even COBE had $\sim 10^8$ data points), this is completely unfeasible, since the time consumption of this brute-force process scales as $\mathcal{O}(n^3)$. Instead, we rearrange our data from time-ordering to pixel-ordering, so that each element corresponds to the data recieved from a specific point on the sky. We call this vector the *map*, $\mathbf{m}$. Given a good enough method, mapmaking constitutes an effective way of reducing the amount of data without losing any information. It also carries the extra bonus that the sky map is an interesting product in and of itself. [72]

QUIET uses the HEALPix pixelisation scheme for all maps. (HEALPix software is also used to some extent by the pipeline.) [73]

5.5.1 The mapmaking equation

Our starting point for making maps is the TOD. We need a recipe - a mapmaking equation - for moving from TOD to map. (The program doing this in the pipeline is,

[3] We refer to these power spectra as "pseudo-C_ℓ's", since they are computed from a limited patch, not the full sky.

[4] Note that this is not, in fact, a chi-distributed variable.

of course, called tod2map.) The simplest, linear mapmaking strategies take the form

$$\mathbf{m} = \mathbf{W}\mathbf{d}, \tag{5.7}$$

where $\mathbf{W}$ is some method-dependent matrix, and $\mathbf{d}$ is the TOD. We recall from eq. 5.1 that the TOD may be expressed ideally as a sum of signal and gaussian noise. (For the moment, we ignore systematics. We will account for more realistic situations shortly.) We expand this view to accomodate a pixel-valued view of the signal,

$$\mathbf{d}_t = \mathbf{P}\mathbf{s}_p + \mathbf{n}_t. \tag{5.8}$$

$\mathbf{P}$ is the *pointing matrix*, an $(n_t \times n_{pix})$-matrix translating between time- and pixel-valued vectors, which we can find by way of the registered telescope pointing. The simplest type of pointing matrix will consist only of ones and zeroes, with one non-zero element per time step, corresponding to the pixel on which the beam was centered. This is dependent on the instrument beam being symmetric, which is luckily the case for QUIET. For the QUIET differential temperature modules, $\mathbf{P}$ will have two non-zero entries per time step, corresponding to the pixels seen by the two horns. One entry will have the value $+1$ while the other is -1, since our data corresponds to the difference between the signal from the two horns. For the polarization modules, each point in time corresponds to one pixel only, but in addition, the pixel is associated with a detector polarization angle ψ, which gives the orientation of the reference coordinate systems by which Q and U are defined. [74]

While the elements of the pointing matrix are known (at least insofar as we have succeeded in calibrating them), the signal and noise vectors are both unknown, so we lack enough information to solve this system. Hence we define the map-vector $\mathbf{m}$ to be an estimate of $\mathbf{s}$, and use this as a guide when finding a suitable mapmaking strategy. We would like the estimate to be unbiased, i.e. $\langle \mathbf{m} \rangle = \mathbf{s}$, and preferrably to have as small variance as is possible. [72]

There are many possible approaches to mapmaking (e.g. [72, 75]). As we noted in the introduction to this chapter, our pipeline has a maximum likelihood approach to this, as well as to the subsequent power spectrum and parameter estimation. The maximum-likelihood map (or equivalently, the minimum variance map) is the solution to eq. 5.7 that maximises the signal likelihood. Given that our TOD can be expressed like in eq. 5.8, $\mathbf{d} - \mathbf{P}\mathbf{s}$ follows a multivariate normal distribution, $\mathbf{n} \sim \mathcal{N}(0, \mathbf{N})$, where $\mathbf{N} = \langle \mathbf{n}\,\mathbf{n}^\mathsf{T} \rangle$ is the covariance matrix of the time-ordered noise vector. Then the likelihood of the signal is [69]

$$\mathcal{L}(\mathbf{s}) \;=\; \mathcal{P}(\mathbf{d}|\mathbf{s}) \;\propto\; \exp\left(-\frac{1}{2}(\mathbf{d} - \mathbf{P}\mathbf{s})^\mathsf{T}\mathbf{N}^{-1}(\mathbf{d} - \mathbf{P}\mathbf{s})\right).$$

We notice the exponent is actually a chi-squared distributed variable, which may be expressed in terms of the log-likelihood,

$$-2\ln\mathcal{L} \;\propto\; (\mathbf{d} - \mathbf{P}\mathbf{s})^\mathsf{T}\mathbf{N}^{-1}(\mathbf{d} - \mathbf{P}\mathbf{s}) \;\sim\; \chi^2\,.$$

So we see that maximising the likelihood with respect to the signal is equivalent to minimising the χ^2,

$$\frac{\partial\mathcal{L}}{\partial\mathbf{s}} = 0 \quad \Rightarrow \quad \frac{\partial\chi^2}{\partial\mathbf{s}} = 0.$$

Differentiating the χ^2, then, gives us the maximum likelihood-estimate for the signal, meaning that we must also exchange $\mathbf{s}$ for $\mathbf{m}$,

$$-\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}(\mathbf{d}-\mathbf{P}\mathbf{m})-(\mathbf{d}-\mathbf{P}\mathbf{m})^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{P}=0.$$

Knowing that in general, $(\mathbf{AB})^\intercal=\mathbf{B}^\intercal\mathbf{A}^\intercal$, and that since $\mathbf{N}$ is a symmetric matrix, so is $\mathbf{N}^{-1}$, giving $(\mathbf{N}^{-1})^\intercal=\mathbf{N}^{-1}$, we see that the two terms are each other's transpose. Since they are also vectors, this means they must be equal. Hence we have

$$\begin{aligned}
&\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}(\mathbf{d}-\mathbf{P}\mathbf{m})=0,\\
\Rightarrow\quad &\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{P}\,\mathbf{m}\;=\;\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{d},\\
\Rightarrow\quad &\mathbf{m}\;=\;(\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{d},
\end{aligned}\tag{5.9}$$

which is the well-known mapmaking equation for a maximum likelihood map. [75]

This picture can be modified to accomodate for the less-than-ideal circumstances met in most real-life experiments. Using the notation of eq. 5.7, any mapmaking method of the form

$$\mathbf{W}\;=\;(\mathbf{P}^\intercal\mathbf{A}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{A},\tag{5.10}$$

where $\mathbf{A}$ is any non-singular matrix, will produce unbiased maps. To confirm this, we see that

$$\begin{aligned}
\langle\mathbf{m}\rangle &= \langle\,(\mathbf{P}^\intercal\mathbf{A}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{A}(\mathbf{P}\mathbf{s}+\mathbf{n})\,\rangle\\
&= (\mathbf{P}^\intercal\mathbf{A}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{A})\,\langle\,\mathbf{P}\mathbf{s}+\mathbf{n}\,\rangle, \qquad \langle\mathbf{n}\rangle=0\\
&= \mathbf{s}.
\end{aligned}$$

Worth noting in particular is the identity [72]

$$\mathbf{W}\mathbf{P}=(\mathbf{P}^\intercal\mathbf{A}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{A}\mathbf{P}=\mathbf{I}.\tag{5.11}$$

We see that eq. 5.10 gives the maximum likelihood map if $\mathbf{A}=\mathbf{N}^{-\mathbf{1}}$. In our case, this is not enough. We need to account for the filters we discussed in §5.2. Frequency filters, like the apodized filter function given in eq. 5.6, can be expressed as diagonal matrices in the Fourier domain, $\mathbf{F}_\nu = F(\nu')\delta_{\nu\nu'}$. To incorporate the filter into the mapmaking equation, we substitute $\mathbf{N}^{-1}$ with $\mathbf{N}^{-\mathbf{1}}\mathbf{F}$, which is the time-domain noise covariance matrix after the filter has been applied. Setting $\mathbf{A}=\mathbf{N}^{-\mathbf{1}}\mathbf{F}$ gives the mapmaking equation

$$\mathbf{m}\;=\;(\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{F}\mathbf{P})^{-\mathbf{1}}\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{F}\mathbf{d},\tag{5.12}$$

which is the approach used by QUIET. For reasons of efficiency, tod2map does not compute the map itself, but the constituent matrices $\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{F}\mathbf{P}$ and $\mathbf{P}^\intercal\mathbf{N}^{-\mathbf{1}}\mathbf{F}\mathbf{d}$. Filters that cannot be applied in the Fourier domain, such as our azimuth filter, are computed as corrections to these constituent matrices via the Woodbury matrix identity. [58]

5.5.2 The noise covariance matrix

Given that we have filtered the data and passed the null tests, we should now expect to have a Gaussian noise distribution. If that is the case, the map-domain noise covariance matrix constitutes a complete description of the uncertainty in the map[76]. That we can compute unbiased maps with complete covariance matrices is one of the great strengths of pipeline B.

The noise covariance matrix for the map is given by [76]

$$\mathbf{N_m} = \langle (\mathbf{m} - \langle \mathbf{m} \rangle)(\mathbf{m} - \langle \mathbf{m} \rangle)^\mathsf{T} \rangle.$$

On this, we apply that $\langle \mathbf{m} \rangle = \mathbf{s}$, and insert equations 5.7 and 5.8 to find

$$\begin{aligned}
\mathbf{N_m} &= \langle (\mathbf{W}(\mathbf{Ps} + \mathbf{n}) - \mathbf{s})(\mathbf{W}(\mathbf{Ps} + \mathbf{n}) - \mathbf{s})^\mathsf{T} \rangle \\
&= \langle (\mathbf{Wn})(\mathbf{Wn})^\mathsf{T} \rangle \\
&= \mathbf{W} \langle \mathbf{n}\mathbf{n}^\mathsf{T} \rangle \mathbf{W}^\mathsf{T} \\
&= \mathbf{W}\,\mathbf{N}\,\mathbf{W}^\mathsf{T},
\end{aligned}$$

where we have applied the identity of equation 5.11 to cancel the signal. Now we insert our mapmaking equation,

$$\mathbf{N_m} = (\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP})^{-\mathbf{1}}\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{F}\,\mathbf{N}\,\mathbf{F}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{P}(\mathbf{P}^\mathsf{T}\mathbf{F}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{P})^{-\mathbf{1}}).$$

Knowing that $\mathbf{N}^{-\mathbf{1}}\mathbf{F}$ is symmetric helps clean this up a bit, and our final expression becomes

$$\mathbf{N_m} = (\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP})^{-\mathbf{1}}(\mathbf{P}^\mathsf{T}\mathbf{F}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP})\,(\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP})^{-\mathbf{1}}, \tag{5.13}$$

which is the form employed by pipeline B. It is quite common to simplify this expression further by assuming $\mathbf{F} = \mathbf{F}^2$, arriving at the expression $\mathbf{N_m} = (\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP})^{-\mathbf{1}}$. This is strictly correct only if one had used a "top-hat" filter, i.e. a step function. For QUIET, it was found that using this expression biased the χ^2 of a noise-only map $\mathbf{m_n}$, $\chi^2 \equiv \mathbf{m_n}^\mathsf{T}\mathbf{N_m}^{-\mathbf{1}}\mathbf{m_n}$, by approximately 3σ. [58]

Like with the map, only the parts of $\mathbf{N_m}$ are computed by tod2map. $\mathbf{P}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP}$ is needed in the map, so only one more matrix, $\mathbf{P}^\mathsf{T}\mathbf{F}^\mathsf{T}\mathbf{N}^{-\mathbf{1}}\mathbf{FP}$, has to be assembled for us to have all the constituents of the covariance matrix as well as the map.

The time-domain noise covariance matrix

In equation 5.13, $\mathbf{N_m}$ is given in terms of the time-domain noise covariance matrix $\mathbf{N}$. We need to look closer at this quantity, since the pipeline is dependent upon computing and storing it.

In principle, the noise covariance $\mathrm{Cov}\,(n(t), n(t'))$ is defined on $\mathbb{R}^2$, since theoretically, both t and t' span the whole of $\mathbb{R}$. Luckily, this does not mean we have to store an infinite matrix. Firstly, our data are discrete, with a limited sampling rate and total number of samples, which reduces the matrix dimensions to $(n_{samp} \times n_{samp})$. Secondly, we assume that the noise is stationary, meaning the the covariance is only dependent on the spacing of the time coordinates, $N(t_i, t_j) = N(\Delta t_{i,j})$. Thirdly, there will be some maximal time interval above which the correlation is negligible, and we call this the *correlation time*, Δt_{max}. We find that each row in $\mathbf{N}$, corresponding to time t_i, has

a finite number of nonzero elements that fulfil $|t_i - t_j| \leq \Delta t_{max}$. These elements are the same in all the rows, while their horizontal position within the matrix is shifted as i varies. This form is known as a *Toeplitz matrix*, which is a type of matrix that has favourable properties under transformation, and is easily stored, since its information may be summarised in a single vector. [69]

The behaviour of the noise varies from one CES to the next. Weather conditions, for instance, can affect the correlation time of a CES. However, separate CES'es are independent, and so a covariance matrix for more than one CES will be block-diagonal, with each block being the time-domain noise covariance matrix of a single CES.

To tie everything together, we need to relate $\mathbf{N}$ to the Fourier-domain noise model $N(\nu)$ that was given in eq. 5.5. Applying equation 5.3 to the definition of $\mathbf{N}$, we find

$$\begin{aligned}\mathbf{N} &= \langle\, \mathbf{n_t} \mathbf{n_t^*} \rangle \\ &= \langle\, (\mathcal{F}\mathbf{n}_\nu)(\mathcal{F}\mathbf{n}_\nu)^* \rangle \\ &= \mathcal{F} \langle\, \mathbf{n}_\nu \mathbf{n}_\nu^* \rangle\, \mathcal{F}^* \\ &= \mathcal{F} \mathbf{N}_\nu\, \mathcal{F}^*,\end{aligned}$$

where $\mathbf{N}_\nu$ is the diagonal Fourier domain noise covariance matrix. We compute an estimate of this matrix using our noise model, $\mathbf{N}_\nu = N(\nu')\delta_{\nu\nu'}$ [69].

Diode correlations

One small matter remains before we move on. So far, we have only treated single CES-diodes, or uncorrelated ones. While we can state freely that there are no correlations between individual CES'es, this is not the case for diodes within the same CES. Including those, we find that the expression for $\mathbf{N}_\nu$ must be augmented.

The time-ordered data-vector for a full CES is assembled by simple concatenation, making the full $\mathbf{N}_\nu$ block-diagonal as long as no diode correlations are counted. Once those are included, our matrix will become

$$N_{ij}^{\text{total}} = \mathbf{N}_{\nu_i \nu_j}, \quad i, j \in [1, n_{\text{diodes}}],$$

where the diagonal elements are single-diode $\mathbf{N}_\nu$'s, and the off-diagonal elements are given by

$$\mathbf{N}_{\nu_i \nu_j} = \gamma_{ij} \sqrt{\mathbf{N}_{\nu_i} \mathbf{N}_{\nu_j}}.$$

The correlation coefficients γ were computed in l3gen.

5.5.3 Map postprocessing

Some post-processing of the pixel-valued data is needed before we move on with our analysis. Despite our careful filtering, there are still elements present in the map that we will need to remove before continuing with the analysis. The most dangereous of these are modes with low signal-to-noise ratio, since these will be especially sensitive to systematics. For instance, edge pixels that are only intermittently observed will not benefit from the scanning strategy and deck rotations that so effectively suppress systematics in pixels that are present in most (if not all) scans. High-ℓ modes have

wavelengths that are far smaller than the fully covered part of the patch, and will get high S/N, but the low-ℓ modes are at risk.

To take care of the map post-processing, the QUIET pipeline includes the utility program scalapost, which uses the software library ScaLAPACK[5] to effectively perform linear algebra operations on tod2map output. It has a number of specialised routines. Among the foremost of these are the solve and finalmap options, which assembles the final maps and covariance matrices from the constituent matrices discussed in the previous section. Hence these routines are indispensable complements to tod2map. Solve is applied if mapmaking has been run without jackknives and the data are contained in a single set. However, we will often have need of jackknife-split sets. For each jackknife applied to the data, tod2map will output two files, each containing the ingredients to make a map and covariance matrix for one half of the data. The two can be combined in different ways depending on our needs. If we want to look for signs of bias, we compute a difference map by directly subtracting the two halves,

$$\mathbf{m}^{\rm diff} = \frac{1}{2}(\mathbf{m}_1 - \mathbf{m}_2), \quad \text{with} \quad \mathbf{N}_{\mathbf{m}}^{\rm diff} = \frac{1}{4}(\mathbf{N}_{\mathbf{m}1} - \mathbf{N}_{\mathbf{m}2}).$$

To get the map we would have found if we had run the mapmaking without a jackknife, we must perform a weighted sum of the two halves, to account for any differences of the noise properties between the two.

$$\mathbf{m}^{\rm sum} = (\,\mathbf{N}_{\mathbf{m}1}^{-1} + \mathbf{N}_{\mathbf{m}2}^{-1}\,)^{-1}(\,\mathbf{N}_{\mathbf{m}1}^{-1}\,\mathbf{m}_1 + \mathbf{N}_{\mathbf{m}2}^{-1}\,\mathbf{m}_2\,), \quad \text{with} \quad \mathbf{N}_{\mathbf{m}}^{\rm sum} = (\,\mathbf{N}_{\mathbf{m}1}^{-1} + \mathbf{N}_{\mathbf{m}2}^{-1}\,)^{-1}.$$

These operations are handled by finalmap.

In addition to assembling the maps and covariance matrices, these routines apply masks to remove edge pixels and point sources[6], and check for negative eigenvalues. $\mathbf{N_m}$ and $\mathbf{N_m^{-1}}$are both computed and stored on disk, so that they need only be computed once.

After running solve or finalmap, the necessary components for producing power spectra are present, and we may move on should we wish to do so. However, there is still the issue of the possibly contaminated low-S/N modes. Whether there are systematics or not, we will need to remove some modes from those maps we want to look at, because the filters applied during mapmaking have an unfortunate effect on the maps: When we filter, we assign some parts of the map very high variance (equivalent to setting parts of $\mathbf{N}^{-1}$ to something very small). Modes affected by this alteration are hardly constrained at all, and may show up with a very high amplitude in the maximum-likelihood map. Although the high variance means such modes make little contribution in the power spectrum estimation, they are a nuisance in the map, since their amplitude will drown out more interesting, signal-carrying modes. Several scalapost routines are dedicated to resolving this issue.

We have two slightly different takes on the problem. The simplest one, called ℓ-cutting (the routine is lcut), works by removing from the map and covariance matrix all significant contribution from ℓ-modes less than a given threshold. In the official QUIET Q-band analysis, this approach was used, with thresholds of $\ell = 5$ and $\ell = 25$ for polarization and temperature maps, respectively[58]. The other one, which we

[5] www.netlib.org/scalapack

[6] To identify point sources within our patches we use the WMAP 7-year point source catalogue[58].

call *eigencutting*, is numerically similar, but somewhat more sophisticated. In future analyses eigencutting will be preferred, and consequently I will describe the procedure in some detail.

Eigencutting

The scalapost routines eigencut and cutoff targets the low-S/N modes by aiming at highly correlated components in the map-domain noise. At this point, we want no strong modes in the map besides those contributed by actual signal, and a strong noise mode would necessarily be a large contribution to $\mathbf{N_m}$. The eigencutting technique effectively deals with such troublesome features without having to resort to the carpet-bombing tactic of erasing everything below a certain ℓ.

The eigencutting routines are based on the technique of *eigenvalue decomposition*, also known as spectral decomposition or eigendecomposition. (This tool is thoroughly treated in any good textbook on mathematical physics, such as Hassani [3].) Eigenvalue decomposition is the decomposition of any diagonalisable $n \times n$ matrix $\mathbf{A}$ into elements given by its eigenvalues and -vectors,

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}, \tag{5.14}$$

where $\mathbf{\Lambda}$ is a diagonal matrix with non-zero entries equal to the eigenvalues of $\mathbf{A}$, $\Lambda_{ii} = \lambda_i$, and $\mathbf{V}$ is an $n \times n$ matrix with columns given by the corresponding eigenvectors. If $\mathbf{A}$ is also known to be real and symmetric, its eigenvectors may always be chosen so as to form an orthonormal set, simplifying eq. 5.14 to

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{\mathsf{T}}. \tag{5.15}$$

Since the eigenvalue matrix is diagonal, this expands to

$$\mathbf{A} = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^{\mathsf{T}} + \lambda_2 \mathbf{v}_2 \mathbf{v}_2^{\mathsf{T}} + \cdots + \lambda_n \mathbf{v}_n \mathbf{v}_n^{\mathsf{T}}.$$

Note that these are outer products, so that each term is an $n \times n$ matrix, and each eigenvalue denotes how much that "eigenmode" contributed to the total matrix.

We exploit this decomposition to translate the map from pixel-space to the eigenspace of the noise covariance. If there is correlated noise in the map between points of angular separation θ, all corresponding elements in $\mathbf{N_m}$ will have a large value. Hence they must be part of an eigenmode with a large eigenvalue. Modes thus identified should be removed from the map, and to avoid biasing the map, we must also assign them infinite variance. After all, when we have given parts of the map an arbitrary value, we no longer have any knowledge of its original state, so the uncertainty associated with that mode can rightly be said to be infinite.

To identify the modes in the map, we must perform a change of basis to find the map vector that corresponds to the diagonalised covariance matrix. To see how to do this, we exploit the invariance of scalar quantities. Readily available for our purpose is the χ^2 of the map. Combining the definition of the χ^2 with the eigenvalue decomposition of

the noise covariance matrix, we find

$$\begin{aligned}\chi^2 &= \mathbf{m}^\top \mathbf{N_m^{-1}}\, \mathbf{m} \\ &= \mathbf{m}^\top \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top \mathbf{m} \\ &= \mathbf{m}'^\top \mathbf{\Lambda}\, \mathbf{m}', \qquad \mathbf{m}' \equiv \mathbf{V}^\top \mathbf{m} \\ &= \sum m'_i\, \lambda_i^{-1} m'_i. \end{aligned} \tag{5.16}$$

This is the key to the eigencutting procedure. We perform the eigenvalue decomposition and change of basis, then identify and set to infinity (a problem in itself, with which we will deal in a moment) the λ_i's that are too large (using either a specified threshold value, or cutting the largest n modes), and zero out the corresponding elements in the transformed map. Transform back, and voilà - we have a cleaner map.

Numerically, there are some tricks involved in this process. When estimating the power spectra, we will have need of the *inverse* eigencutted noise covariance matrix. Furthermore, since we used the χ^2 to find the change of basis for the map, we need to perform the decomposition on $\mathbf{N_m^{-1}}$ rather than on $\mathbf{N_m}$. Inverting matrices is computationally expensive, so all problems involving them should be tackled as intelligently as possible. Enter the *Woodbury matrix identity*, also known as the Sherman-Morrison-Woodbury formula, or the matrix inversion lemma, which concerns the inverse of a rank k correction to an $n \times n$ matrix. It takes the form

$$(\mathbf{A} + \mathbf{UCV})^{-1} = \mathbf{A^{-1}} - \mathbf{A^{-1}U}(\mathbf{C^{-1}} + \mathbf{VA^{-1}U})^{-1}\mathbf{VA^{-1}} \tag{5.17}$$

where the matrices $\mathbf{A}, \mathbf{U}, \mathbf{C}$ and $\mathbf{V}$ have dimensions $n \times n$, $n \times k$, $k \times k$ and $k \times n$, respectively. Inspecting this expression, we see that while the left-hand side is the inverse of an $n \times n$ matrix, the right-hand side only demands the inversion of the (presumably much smaller) $k \times k$ matrix, besides $\mathbf{A}^{-1}$ itself, which we may assume is already known.

In our case, $\mathbf{A}$ is of course $\mathbf{N_m}$, $\mathbf{C}$ is the rank 1 (i.e. scalar) eigenvalue λ_i corresponding to the mode we would like to eliminate, while $\mathbf{U}$ and $\mathbf{V}$ are the eigenvector $\mathbf{v}_i$ and its transpose, $\mathbf{v}_i^\top$. This means we are dealing with the $k = 1$ special case of eq. 5.17, sometimes referred to as the Sherman-Morrison formula, where the matrix inversion reduces to inverting a scalar. This solves our other problems, too. While setting an eigenvalue to infinity would pose some serious numerical issues, setting its inverse to zero is laughably simple.

The Woodbury identity for the inverse corrected noise covariance matrix takes the form

$$\begin{aligned}\mathbf{N_{m\,(cut)}^{-1}} &= (\mathbf{N_m} + \lambda_i^\infty \mathbf{v}_i \mathbf{v}_i^\top)^{-1} \\ &= \mathbf{N_m^{-1}} - \mathbf{N_m^{-1}} \mathbf{v}_i\, (\mathbf{v}_i^\top\, \mathbf{N_m^{-1}}\, \mathbf{v}_i)^{-1} \mathbf{v}_i\, \mathbf{N_m^{-1}},\end{aligned}$$

where we have used that $(\lambda_i^\infty)^{-1} = 0$. While the original expression involves the inversion of an $(n \times n)$ matrix, we have managed to reduce it to a simple series of matrix multiplications (the expression inside the parentheses is a scalar), involving only the original $\mathbf{N_m^{-1}}$, which we already have.

To illustrate the effect this process has on the map, an example is given in figure 5.4. The problem of large-amplitude, large-variance modes is evident, as we cannot

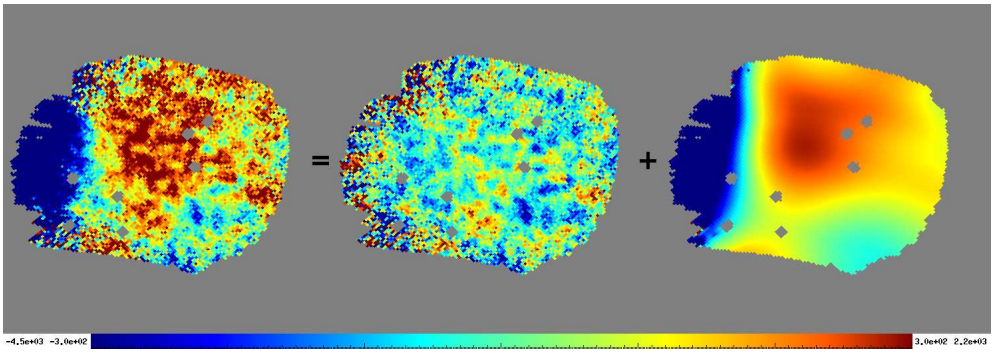

Figure 5.4: QUIET Q-band temperature maps, patch 2a, demonstrating the effect of eigencutting. 38 modes, shown in the right-hand map, has been cut from the left-hand map, resulting in the middle map. Range: $\pm 300\mu$K. The holes in the patch is due to the use of a point-source mask.

read any information about the true temperature of the sky by looking at the left-hand map. After cutting, however, the map is useful. The cut modes on the right contributed nothing of value, and we can throw them away without further ado. We should of course keep in mind that the maps we estimate power spectra by are not cut as much as is done here. Unless we find indication that it is necessary, we might not cut anything at all.

5.6 Power spectrum estimation

We are now nearing the end of the long and winding road of the data analysis, and are ready to estimate power spectra. Including cross-correlations, there is a total of six spectra which we may attempt to measure: TT, EE, BB, TE, TB and EB, where the latter two are expected to be zero, as we noted in §3.3. We briefly discussed the development of theoretical power spectra in chapters 1 and 3. Here I will describe how we produce observational power spectra from our maps, and in chapter 6 we will compare these to the current best-fit model spectra. The processes described in this and the next sections are all performed by the program map2cl.

As has already been mentioned, pipeline B employs maximum likelihood estimation of power spectra and cosmological parameters. After discussing the theoretical side of this method, I will comment on some technicalities of its execution.

5.6.1 Maximum-likelihood power spectrum estimation

Our method of estimating power spectra and cosmological parameters is based on the chief principle of Bayesian analysis, which is that the probability distribution of some outcome, in this case the power spectrum (or, equivalently, a set of cosmological parameters), will change under the influence of any evidence we might have. In this respect we talk of *prior* and *posterior* probabilities, meaning probabilities before and after they have been altered by evidence. In our case, the posterior we would like to map is

$P(C_\ell|\mathbf{m})$. It is given by [77]

$$P(C_\ell|\mathbf{m}) \propto P(\mathbf{m}|C_\ell)P(C_\ell).$$

The prior, $P(C_\ell)$, is usually taken to be uniform, while the term $P(\mathbf{m}|C_\ell)$, the probability of observing the map given a certain theoretical universe described by C_ℓ, is also known as the *likelihood function* of the model,

$$P(\mathbf{m}|C_\ell) = \mathcal{L}(C_\ell).$$

If both the noise and the signal are Gaussian, this function contains all the information in the experiment. If we maximise the likelihood on the space spanned by all possible C_ℓ's, we should find the model that fits optimally with our data. [78]

We find the likelihood by considering the data on hand. We have a map consisting of Gaussian noise and assumed Gaussian signal, $\mathbf{m} = \mathbf{s}_p + \mathbf{n}_p$. These are mutually uncorrelated quantities, since their origins are completely unrelated, meaning that the covariance matrix of the full map is simply the sum of the corresponding matrices for each of these vectors by themselves,

$$\mathbf{C} = \mathbf{S} + \mathbf{N}_\mathbf{m}\,. \tag{5.18}$$

As we know, $\mathbf{N}_\mathbf{m}$ is calculated by the pipeline. The signal covariance matrix has an analytical expression and is dependent on the theoretical C_ℓ, thus forming the link connecting the model Universe to the observed one. Its derivation is rather messy, but is included in appendix C. A crucial point is the fact that the statistical isotropy of the Universe implies that, like with the time-domain noise covariance, the signal covariance is only dependent upon the angular distance of the points in question. This implies the equality of the signal covariance of points of angular separation θ, and the *two-point correlation function* $C(\theta)$[7].

Due to our imperfect equipment and methods of representing the sky, the signal we find from our data is not a perfect rendering of the CMB. During observation, the signal is smoothed by the beam, and it is convolved further by the subsequent pixelisation. These processes may be described in spherical harmonics, taking a convenient form as functions of ℓ. We may multiply C_ℓ with these functions to account for the convolution. The beam function is denoted b_ℓ, and is instrument-specific. We use the N_{side}-dependent HEALPix pixel window function w_ℓ for the pixelisation. HEALPix software provides this function[73]. Including these corrections, the signal covariance matrix for temperature is

$$S_{ij} = C(\theta) = \frac{1}{4\pi}\sum_{\ell=0}^{\infty}(2\ell+1)P_\ell(\cos\theta)b_\ell^2 w_\ell^2 C_\ell, \tag{5.19}$$

where i, j are subject to the constraint $\hat{n}_i \cdot \hat{n}_j = \cos\theta$. The case for polarization is equivalent, only with basis functions that can account for the spin properties of the polarization field. (An example of such functions was briefly discussed in §1.3.4.)

Now we can set up a joint probability distribution for the full map vector. Under the assumption of Gaussianity, it has the form of a multivariate normal distribution,

[7]Despite its somewhat misleading name, $C(\theta)$ is in fact a covariance.

[78]

$$\mathcal{L}(C_\ell) = P(\mathbf{m}|C_\ell) = \frac{1}{(2\pi)^{\frac{N_p}{2}}|\mathbf{C}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}\mathbf{m}^\mathsf{T}\mathbf{C}^{-1}\mathbf{m}\right), \tag{5.20}$$

where N_p is the number of pixels in the map. This is the function we need to maximise over the space of all physically realisable C_ℓ's.

5.6.2 Technical matters

Hardly surprising, we need some tricks up our sleeves in order to accomplish the above task. Straight-forward brute-force is not an option, given the very large space of possible (albeit not necessarily very likely) power spectra we would have to consider.

The likelihood is a multi-dimensional function, where the number of dimensions is given by the number of multipoles we want to estimate. For a perfect, all-sky survey, each ℓ is independent, and could be estimated without any knowledge of the others. This is not the case for real surveys, which are always imperfect in some way. In general, an experiment cannot independently estimate ℓ's that are closer together than $\Delta\ell \lesssim 2\pi/\theta$, where θ is the maximal angular extent of the survey[78]. Hence one will always divide the power spectrum into bins of a certain size, rather than try to determine each ℓ separately. This also serves to reduce the number of free parameters in the calculation. In QUIET, we use bins of $\Delta\ell = 50$.

To find the maximum-likelihood spectrum, we do not need to map the full multi-dimensional likelihood function. Rather, we look for the maximum by way of an iterative procedure, which works as follows: We begin with some initial guess for the power spectrum, for instance $C_\ell = 0$. We hold all the bins C_b constant, except one, which is allowed to vary by some quantity δC_b. We choose the value of δC_b which maximises the likelihood, using Newton-Raphson optimisation. We loop through all the bins in this fashion. Then the binned power spectrum is updated using the set of δC_b's we found in the first round, and the process starts over. We continue until the δC_b's are all sufficiently close to zero.

We are still not quite finished. Computing the likelihood function is very heavy work, so we use a few more simplifying measures. Firstly, rather than computing the actual likelihood function, we choose to work with the log-likelihood, which is numerically easier to handle, and has the same maxima. It is

$$\begin{aligned} -2\ln\mathcal{L} &= \mathbf{m}^\mathsf{T}\mathbf{C}^{-1}\mathbf{m} + 2\ln\left((2\pi)^{N_p/2}|\mathbf{C}|^{1/2}\right) \\ &= \mathbf{m}^\mathsf{T}\mathbf{C}^{-1}\mathbf{m} + \ln|\mathbf{C}| + N_p\ln(2\pi). \end{aligned} \tag{5.21}$$

Following the method described by Bond et. al. (1998), we use a *quadratic estimator* to find the likelihood-maximising δC_b's. By using a truncated Taylor expansion of the log-likelihood, we find a local, quadratic approximation to the function. Under this approximation, we solve for the zeroes of $\partial\ln\mathcal{L}/\partial C_b$, to find the optimal correction δC_b. The quadratic estimator is an approximation to this correction, and has the form

$$\delta C_b = \frac{1}{2}\sum_{b'}\mathbf{F}^{-1}_{bb'}\,\mathrm{Tr}\left[(\mathbf{m}\mathbf{m}^\mathsf{T} - \mathbf{C})(\mathbf{C}^{-1}\mathbf{C}_{,b'}\mathbf{C}^{-1})\right].$$

$\mathbf{F}_{bb'}$ is the Fisher matrix,

$$\mathbf{F}_{bb'} = \frac{1}{2}\,\mathrm{Tr}\left(\mathbf{C}^{-1}\mathbf{C}_{,b}\mathbf{C}^{-1}\mathbf{C}_{,b'}\right).$$

The approximation lies in that we use the Fisher matrix rather than the *curvature matrix* $\mathcal{F}_{bb'} \equiv -\frac{\partial^2 \ln \mathcal{L}(C_\ell)}{\partial C_b \partial C_{b'}}$, which contains the full information of the curvature of the function. The relation between the two is $\mathbf{F}_{bb'} = \langle \mathcal{F}_{bb'} \rangle$. The quadratic estimator has the fortunate property that when applied iteratively, it will find the exact location of the likelihood peak. [78]

There is an additional complication: For eq. 5.20 to be a valid probability distribution, the covariance matrix $\mathbf{C} = \mathbf{S} + \mathbf{N}$ must be positive definite[8]. Hence we have the demand $\mathbf{S} > -\mathbf{N}$. The quadratic estimator does not follow this constraint, so we add a safety mechanism in the form of a step length multiplier α. For each iteration, the change in C_b will be $\alpha\, \delta C_b$, where $0 < \alpha \leq 1$. The multiplier is adjusted for each step, to ensure that we never step into regions where the likelihood diverges.

This strategy is not the most sophisticated approach to power spectrum estimation. A different possibility, allowing for the inclusion of more free parameters, is to use a CMB Gibbs sampler, as described in Eriksen et.al.(2008) [77].

5.7 Estimation of cosmological parameters

Our final step is to make estimates of cosmological parameters based on the power spectra. In principle, this does not need to be a separate operation. One might have applied the maximum-likelihood estimation directly to the desired set of parameters. Numerically, however, this two-step approach is preferrable, since finding a power spectrum acts as data compression, making the subsequent parameter estimation easier. [78]

QUIET does parameter estimation based on the two polarization power spectra, EE and BB. We do a fairly simple brute-force grid evaluation of the pixel-space likelihood in which the binned power spectrum of the data is compared to a test spectrum, which in turn is a function of a fiducial spectrum. As fiducial, we use the WMAP 7-yr best-fit ΛCDM model.

The easiest case is the one where the test spectrum is parameterised by one parameter only, $C_\ell = aC_\ell^{fid}$, where a is the amplitude of the spectrum relative to the fiducial. In the original Q-band analysis, this approach was chosen for both polarization spectra, fitting a parameter q to the EE spectrum, and the tensor-to-scalar ratio r to the BB spectrum. The corresponding 1-dimensional likelihood function is [58]

$$\mathcal{L}(a) \propto \frac{1}{\sqrt{|\mathbf{C}(a)|}} \exp\left(-\frac{1}{2}\mathbf{d}^\intercal \mathbf{C}^{-1}(a)\mathbf{d}\right) \tag{5.22}$$

where $\mathbf{C}$ is the covariance matrix from eq. 5.18. The likelihood is evaluated on a grid in a to find the maximum likelihood.

Numerically, this process is easy enough that we may add another degree of freedom, should we so choose, and the present pipeline supports this. The current standard parameterisations for the EE and BB power spectra, respectively, is

$$C_\ell^{EE} = qC_\ell^{fid}\left(\frac{\ell}{\ell_{pivot}}\right)^n \quad \text{and} \quad C_\ell^{BB} = rC_\ell^{fid}, \tag{5.23}$$

where n is a spectral index allowing for a tilt of the best-fit spectrum with respect to the fiducial.

[8] A real $n \times n$ matrix $\mathbf{A}$ is positive definite if, for all non-zero vectors $\mathbf{z}$ with real entries, $\mathbf{z}^\intercal \mathbf{A}\mathbf{z} > 0$.

Chapter 6

Reanalysis of the QUIET-I Q-band data

> *If you wish to make an apple pie from scratch, you must first invent the universe.*
>
> *Carl Sagan*

As noted in chapter 4, the official QUIET-I Q-band analysis has been completed, and was submitted to the *Astrophysical Journal* by the end of 2010[58]. Since then, significant alterations have been made to the pipeline - most with the intention of improving the efficiency and accuracy with which the program can handle the large W-band data set. Re-running the Q-band analysis serves as a test of this new version of the pipeline: Should the results improve with respect to the official numbers, it would imply that the changes we have made have been good. Conversely, if we find a poorer outcome with the new pipeline, we will have an indication that something might be amiss. In addition, since the Q-band data set is relatively small, it is well suited for fulfilling one of the main objectives of this thesis, which is to present and understand the Oslo branch of the QUIET data analysis framework.

In the following, I will present the results of the reanalysis. Final results will be compared to the current best-fit ΛCDM model, as well as the official QUIET Q-band results. I have chosen to split the analysis into three main sections: Galactic, CMB temperature and CMB polarization. Up to and including the calibration step (l3gen), these groups are treated largely identically, but after this point the analyses diverge.

The interval at which the CES-diode RMS noise level (σ_0) is first estimated spanned frequencies from 2.0 to 4.5 Hz in the first Q-band analysis. For reasons which will be made clear in §6.3.2, we chose to widen this interval. For all the data, frequencies from 0.3 to 4.5 Hz were used for this purpose.

6.1 Galactic observations

As we saw in chapter 4, we have made observations of two areas in the galactic plane. Patch gc covers the centre of the galaxy, while patch gb lies some 30° away. Eventually, these observations will be used for purposes of foreground science, but this is not a priority as long as we are working on the CMB polarization data. In its current form, the pipeline is not fit to analyse these data, and significant adaptations will have to be

Cut efficiencies (% of CES-diodes remaining after each cut)

Cut type	Temperature		Polarization	
	Patch gb	Patch gc	Patch gb	Patch gc
Removing dead diodes	100.00	100.00	92.65	92.65
Static/input cuts	99.66	93.12	92.33	86.27
Gain cut	99.66	93.12	92.33	86.27
Elementary cuts	99.66	93.12	92.33	86.27
APEX cuts	99.66	93.12	92.33	86.27
Type-B cuts	84.24	73.94	88.97	84.31
Weather cuts	82.46	73.94	85.05	83.02
Noise model σ_0 cut	81.19	72.49	84.09	80.92
Noise model f_{knee} cut	64.03	64.15	82.64	79.59
TOD χ^2 cut	46.57	42.33	82.52	79.50
TOD outlier cut	45.38	40.08	81.59	79.01
TOD azimuth cut	38.77	33.40	80.60	78.61
Scan frequency χ^2	23.35	19.44	67.71	75.75
Low frequency χ^2	23.35	19.38	67.41	75.11
All frequencies χ^2	22.46	19.38	67.41	75.11
High frequency χ^2	22.46	19.38	66.74	74.32
1.2 Hz spike	22.46	19.38	66.74	74.32
1.0 Hz spike	22.46	19.38	66.64	74.32
General spike	22.42	19.38	66.56	74.31
Sidelobe sun cut	21.95	19.25	63.11	72.82
Accept ratio cut	21.44	19.25	60.05	72.21
Total no. CES-diodes after cuts	506	291	12047	9281

Table 6.1: Summary of automated cut procedures for all galactic observations. We did a total of 189 CES'es of patch gc and 295 of patch gb during the Q-band season. The temperature data are naturally scarce, as there were only two detectors.

made when we are ready to start this task. For the time being, we content ourselves with making maps.

6.1.1 Cuts and filtering

In sections 5.2 and 5.3 we saw how the pipeline automatically models noise, filters out systematics, and rejects bad scans. All these steps are based on the condition that our CMB patches have low foregrounds, and that little to no signal should be detectable in any one CES-diode. This is not the case for the galactic patches. There are strong point sources in and around the galactic plane, in temperature as well as in polarization. Most notable among these is the presumed supermassive black hole at the centre of the galaxy, Sgr A*, which is a strong source of polarized synchrotron radiation. Should we presume to analyse the galactic patches without taking the necessary precautions, we will most likely end up throwing out most of our data, erroneously interpreting it as contaminated.

Table 6.1 gives the results of the automated cuts for the galactic patches. For

temperature, the final accept rates are very low. We see that large shares of the data are removed because of the TOD χ^2, TOD azimuth and scan frequency χ^2 cuts. This fits well with what we would expect for data containing strong, localised sources. We also note that the noise model f_{knee} fit seems to have failed quite often. The situation is not as bad in polarization, leading us to conclude that the point sources are not as many and as strong here.

Given that the automated cuts are so poorly suited to the galactic observations, we cannot be certain that all bad scans have been successfully cut. To make galaxy maps, some manual editing must thus be expected. For the polarization data, the accepted CES'es are many enough that the auto-generated list may be used as a starting point. When making the maps in the following section, each single-CES map were inspected visually, and scans that were obviously poorly calibrated were removed from the accept list. For the temperature data, new accept lists have been made from scratch, since the automated cuts had rejected practically every scan that covered the plane of the galaxy. The new accept lists gave temperature data accept rates of $\sim 82\%$ and $\sim 66\%$ for patch gb and patch gc, respectively, which is a great improvement.

6.1.2 Maps of galactic patches

In QUIET, we have two different strategies for mapmaking. So far, we have only discussed the *solved* map and corresponding noise covariance matrix, which we found by using eqs. 5.12 and 5.13. Computing these is quite heavy work, but necessary for the subsequent parameter estimation. However, we can also make a naive, *binned* map, which is found using the simpler mapmaking equation, 5.9. The binned map has the underlying assumption that the noise is uncorrelated and standard normally distributed, which gives the simple noise covariance $\mathbf{N}^{-1} = \mathbf{I}$. This makes the mapmaking a fast and easy operation. For the CMB patches, there is often little difference between the binned map and the solved, eigencutted one. The strength of the latter lies in that we have complete knowledge of the noise covariance.

However, our filters do not play well with the binned maps. When applying a filter and including the possibility of $\mathbf{N}^{-1} \neq \mathbf{I}$, our data vector is altered, $\mathbf{d} \rightarrow \mathbf{N^{-1}Fd}$. Since this is not taken into account in the naive mapmaking equation, the binned map will be biased. If the map has areas of high correlation, which is the case for the galactic patches with their strong, localised sources, the consequences of undue filtering are clearly visible. To demonstrate this, the top row of figure 6.1(a) shows a binned map of patch gc with the filtering that is standard for a CMB patch. Below these, for comparison, are included the same maps when all but a modest $1/f$ filter (applying inverse noise-weighting to each Fourier component) have been turned off.

There are two types of stripes in the filtered, binned maps, both of which are completely unphysical. The most glaring of them, going from lower left to upper right in the Stokes' Q map, is due to the azimuth filter. (There are azimuth filter stripes in the U map as well, but they are less conspicuous.) When we bin the TOD in azimuth, we will get a large contribution in one bin each time the galactic centre is observed. The azimuth filter interprets this as contamination, and tries to remove it from the data. However, the filter cannot judge whether the contamination is localised in elevation or not. Hence it will remove power from all pixels within the contaminated azimuth bin. This works well for removing ground pickup, but when the filter is faced with a point

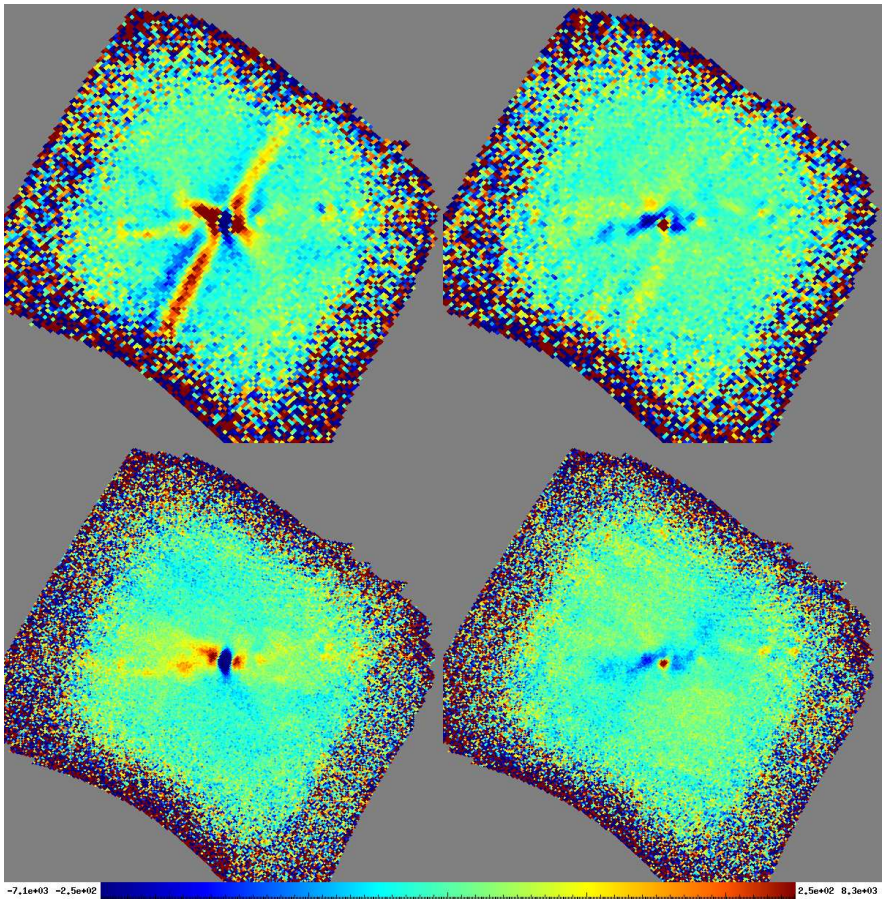

(a) Binned polarization maps of the galactic centre.The upper maps have $N_{side} = 256$ and standard filtering, while the lower have $N_{side} = 512$ and only $1/f$ filtering. Range $\pm 250\mu$K.

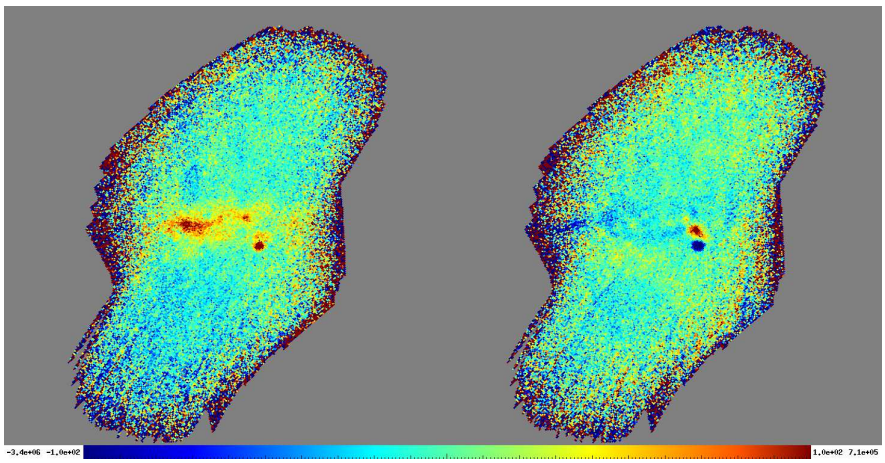

(b) Binned polarization maps for patch gb, $1/f$ filtering only, $N_{side} = 512$. Range $\pm 100\mu$K.

Figure 6.1: Binned maps of Q- (left) and U- (right) polarization for the galactic patches.

source, it makes the type of stripe we see here, because it removes power from parts of the map that did not have excess power to begin with.

The other type of stripe is a little less conspicuous, but we see it clearly when comparing the map with and without filtering: The blue galactic centre has elongated "sidelobes" of the opposite sign, along the scanning direction (perpendicular to the direction of the azimuth filter stripe). This is due to the highpass filter, which removes the lower end of the frequency spectrum. In CMB scans, this works to hinder stripes caused by slow temperature fluctuations in the instrument or atmosphere. However, the galactic centre contributes so much power to the TOD that the average is offset by it. When the average is then removed from the affected CES-diode, the areas surrounding the galactic centre will be wrongly calibrated due to this offset.

The lower maps in figure 6.1(a) clearly show that we are better off with less filtering when making maps of the galactic centre. However, we must remember that the filters are there for a reason. While these maps show little of the filter-induced effects (the galactic centre sidelobes may still be at least partly spurious. We cannot be certain without doing a more thorough analysis, making solved maps and comparing with WMAP), we now see signs of stripes of the kind described in §5.2.1. However, with the high S/N of the galactic patches, abstaining from filtering gives decent results. Figure 6.1(b) shows the equivalent maps for patch gb, in which there is very little, if any, visible contamination.

In the case of the temperature maps, we do not have the benefit of pretty binned maps. Because we use differential temperature detection, bright spots in binned maps will be accompanied by dark shadows, like those that could be seen in the point source map in figure 5.3. When we combine scans from a whole season, during which the deck angle was changed at regular intervals, this gives dark rings around the bright points in the galactic plane. Figure 6.2(a) demonstrates this effect in the binned temperature map of patch gb. Since we want maps with as little of such artefacts as possible, the temperature maps should be made with the proper mapmaking equation. For the sake of appearance, we would like to use $N_{side} = 512$. When computing the full covariance matrix, this requires a lot of space. The unmasked patch gc has ~ 45000 pixels in $N_{side} = 512$, so the corresponding covariance matrix would need more than 7 Gb of memory, even in single presicion. To speed up the process, we apply a conservative mask that removes everything but a band around the galactic plane, where the most interesting structures are. With this mask, the number of pixels was halved, giving significantly shorter computation time.

The resulting maps, after eigencutting, are given in panels (b) and (d) of figure 6.2. A section of the WMAP7 Q-band temperature map corresponding to patch gc, is included for comparison, in panel (c). We notice that there is still some degree of contamination in the solved maps, that of patch gc in particular. A diagonal stripe through the galactic centre and remnants of the dark ring-structures can be seen, the latter possibly caused by uncertainties in the pointing. Still, the QUIET map compare well to WMAP, when we consider the very modest amount of data that went into it.

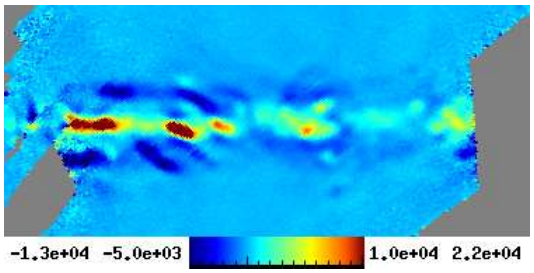

(a) Binned map of patch gb

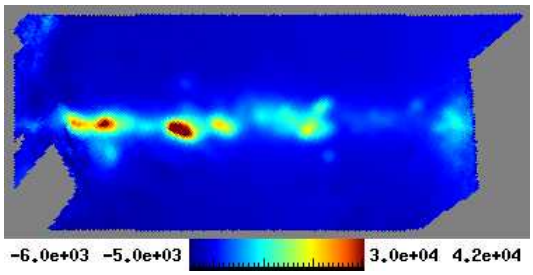

(b) Solved map of patch gb, 95 modes cut

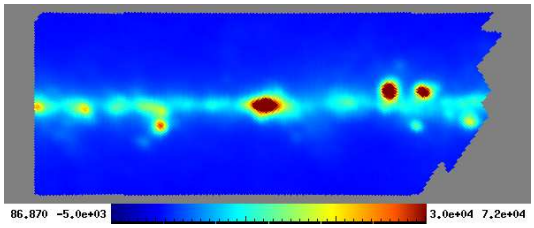

(c) WMAP7 map of patch gc

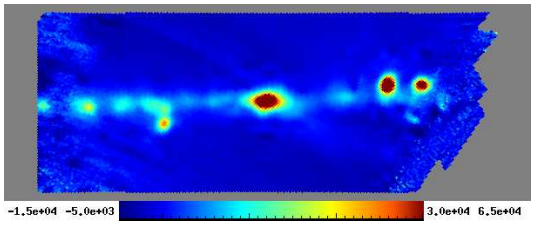

(d) Solved map of patch gc, 174 modes cut

Figure 6.2: Temperature maps for the galactic patches. Temperature range $[-5, 10]$ mK for panel a, $[-5, 30]$ mK for the rest.

6.2 Temperature analysis of CMB patches

Measuring the CMB temperature anisotropy is not in itself an important goal for QUIET, as we can hardly expect much with so few detectors. The prime purpose of the temperature measurements is calibration of pointing, beams and sidelobes. We have no lack of sources with known temperatures, by which we may calibrate our temperature data. Furthermore, the combination of temperature and polarization measurements allows us to compute TE and TB cross-correlation spectra without depending on data from other experiments, which is an advantage. In the following are given the results of the reanalysis of the Q-band CMB temperature data.

6.2.1 Gain estimation

As was discussed in §5.2.3, we need to calibrate the instrument gain so that our maps and power spectra match the true temperatures of the CMB. This also allows us to directly compare our results to those of other experiments. After the conclusion of the original Q-band analysis, something has apparently happened to the temperature gain values, although we have not been able to establish what or why. Whatever the reasons were, upon running a new analysis on temperature data, the amplitudes of the maps and power spectra with respect to WMAP were off by a factor of $\gtrsim 2$ in the maps, ~ 5 in the power spectra. Gain calibration for temperature has therefore been repeated. Two sources, Jupiter and the galactic centre, have been used for this purpose, in two independent calibration runs. As we will shortly see, both targets have strengths as

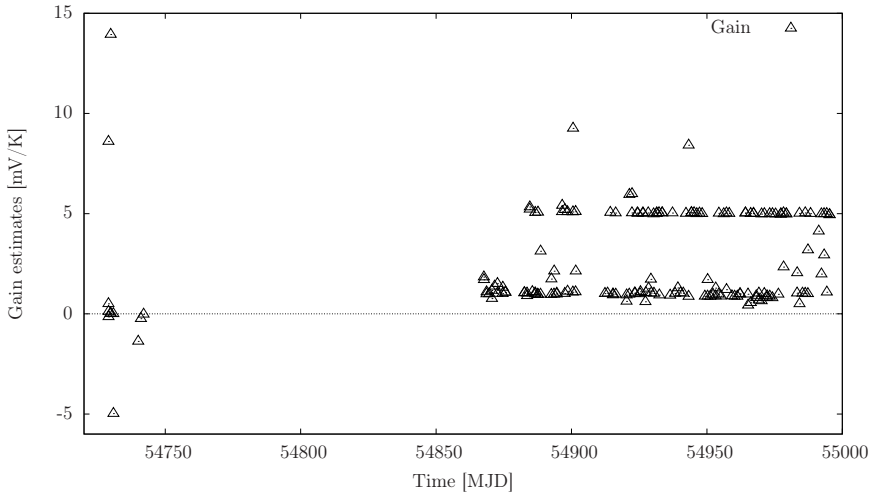

Figure 6.3: Estimated diode gain for module 18, diode 0, based on the Q-band season's observations of patch gc. The lower level corresponds to successful estimates.

well as weaknesses, so we have found it useful to employ both. This allows us to choose the better of the two sets of estimates for our final results.

The antenna temperature response from observing Jupiter is well known, and tabulated as function of frequency. For the galaxy, we use the WMAP Q-band map as the fiducial, but this comes with an additional difficulty: While our Q-band detectors have bandpass centre frequency of ~ 43.5 GHz, the WMAP Q-band data were taken at 41 GHz[79]. In the Q-band, the galactic centre signal is dominated by synchrotron radiation. We know this to have a strong frequency dependence, decreasing as ν^{-3}. The shift in frequency between QUIET and WMAP is enough to cause our gain estimations from the observations of the galaxy to be off by a small, but significant factor.

We find that to correct for the frequency shift, we need to scale our maps by the factor

$$\left(\frac{43.5\,\mathrm{GHz}}{41\,\mathrm{GHz}}\right)^{-3} \frac{a2t(43.5\,\mathrm{GHz})}{a2t(41\,\mathrm{GHz})} \approx 0.84,$$

and correspondingly, the power spectra should be scaled by $0.84^2 \approx 0.71$. The factor $a2t(\nu)$ is the antenna-to-thermodynamic temperature correction factor. We need it because the frequency dependence of the intensity of the galactic synchrotron radiation applies to antenna temperature, i.e. instrument response, while our maps and power spectra are in units of thermodynamic, i.e. "true", temperature. This conversion factor is

$$a2t(\nu) = \frac{(e^x - 1)^2}{x^2 e^x}, \quad \text{where} \quad x = h\nu/k_B T_0,$$

where k_B and h are the Boltzmann and Planck constants, and $T_0 = 2.725$ K is the mean temperature of the sky.

Since Jupiter have no such problems associated with it, it has been our prime source for calibrating temperature gain. However, since it is effectively a point source, the shape of the beam must be calibrated alongside the gain. This additional factor increases the possibility of errors. The galaxy centre, on the other hand, is an extended source as

compared to our beam, meaning that the beam shape will be of little importance in the calibration. The final power spectra from the two sets of gain estimates will be given in §6.2.5. We will find that these results lend credit to the galactic gain estimates, thus these are the ones used for the results presented in the following sections.

Figure 6.3 gives an example of a fairly typical gain estimation outcome. For sufficiently strong sources, the iterative parameter fit is performed for each individual CES-diode. There will always be some unsuccessful fits, for instance if the source was not within the field of view of all the horns, or if there was too much noise and the program did not manage to identify the source. Very strong sources, like Jupiter, will generally have fewer failed fits. The failed fits largely manifest as outliers in plots of estimated diode gain as function of time. We may sometimes find that there are large concentrations of failed fits, like in the example given here. Without any additional information, it is not evident which of the two "levels" correspond to successful gain estimations. However, we may control the estimates using maps of the type we encountered in figure 5.3. Using this strategy we have found that the upper level correspond to failed fits. The successful fits in the lower level indicate that the gain is approximately constant through the season.

The final gain estimates were produced by a two-step procedure: After visual inspection of all temperature diode gain estimates, strict limits for the gain were set to avoid the heavier outliers and the larger concentrations of failed fits. The mean and standard deviation per diode were computed for the remaining values. The subset of values lying within 3σ from the mean was extracted, and a second mean value computed. This value was used as diode gain in the analysis.

6.2.2 Automated data cut results

Table 6.2 gives a summary of the automated cuts on the temperature data. The immediate impression is that the final accept rates are very low, even though this set of cuts is not particularly strict. The accept rates for the temperature data used in the published QUIET Q-band analysis were in fact far lower: 12.4%, 6.9% and 6.8% for patches 2a, 4a, and 6a, respectively (7b was excluded due to too little data altogether)[58]. The main problem with the temperature data is their very high sensitivity to atmospheric disturbances, and so they are more prone to systematics than the polarization data. Since these are Q-band data, we must also expect significant ground pickup. The tabulated values show that among the cut criteria responsible for removing the larger portions of data, we find the scan frequency χ^2 test and the noise model f_{knee} test, while the weather cut only removes a rather small amount. This indicates that ground contamination is the most severe of the contaminations.

The cut thresholds for temperature are adjusted for the poorer data quality (see table 5.1), since the polarization cut criteria would all but eradicate our temperature data. Hence our accepted temperature data are not as clean as the polarization set, but since the signal is 1-2 orders of magnitude stronger, we can live with this.

6.2.3 Null test results

Given that we used milder cut criteria for temperature than for polarization, we have to face the fact that these data are not as clean as we could hope. We cannot expect them to perform as well in null tests as the expectations we have of the polarization

Cut efficiencies (% of CES-diodes remaining after each cut)

Cut type	Patch 2a	Patch 4a	Patch 6a	Patch 7b
Removing dead diodes	100.00	100.00	100.00	100.00
Static/input cuts	98.85	94.97	91.02	87.27
Gain cut	98.85	94.97	91.02	87.27
Elementary cuts	98.85	94.97	91.02	87.27
APEX cuts	98.85	94.97	91.02	87.27
Type-B cuts	89.86	73.40	59.73	68.01
Weather cuts	89.19	72.89	58.29	66.85
Noise model σ_0 cut	87.84	71.84	56.44	65.14
Noise model f_{knee} cut	69.62	50.83	40.24	47.86
TOD χ^2 cut	67.44	50.15	39.56	47.20
TOD outlier cut	64.63	49.17	38.75	46.74
TOD azimuth cut	63.12	48.35	38.06	43.17
Scan frequency χ^2	49.63	37.99	31.06	31.87
Low frequency χ^2	49.61	37.98	31.03	31.87
All frequencies χ^2	49.58	37.98	31.03	31.87
High frequency χ^2	49.58	37.98	31.03	31.87
1.2 Hz spike	49.58	37.98	31.03	31.87
1.0 Hz spike	49.58	37.98	31.03	31.87
General spike	49.56	37.98	30.97	31.87
Sidelobe sun cut	46.25	34.09	29.74	30.59
Accept ratio cut	46.03	33.70	29.64	30.59
Total no. CES-diodes after cuts	3215	1984	1902	788

Table 6.2: Summary of automated cut procedures for all CMB temperature data.

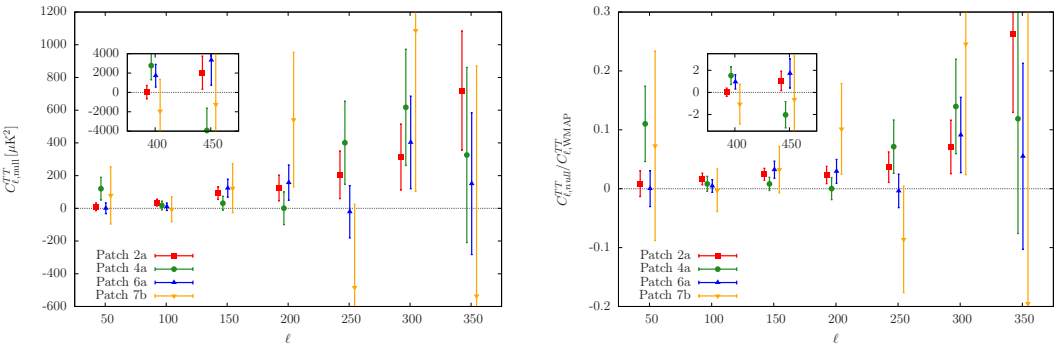

(a) Power spectra for deck angle null test

(b) Ratio of null-spectra to best-fit ΛCDM spectrum

Figure 6.4: Deck angle null test evaluation plots for CMB temperature data. The power spectra are maximum likelihood estimates with $\pm 1\sigma$ error bars.

data. However, since the signal is so strong (the next section will show that it is clearly visible by eye in the maps), we do not need to place such strict demands on the null test outcomes. The important point is to confirm that the statistical errors outweigh the systematic errors.

In terms of the null test parameters PTE and σ, this set of temperature data fails the tests rather spectacularly. There is most certainly systematic contamination left in the temperature maps, but how much? To judge this, we evaluate the power spectra of the difference maps for each null test. We look at the spectra themselves, as well as the ratios of the null-spectra to a fiducial power spectrum. Figure 6.4(a) shows the four patch null-spectra for a nulltest splitting the data by deck angle. There is clear evidence of systematics, since the spectra are not everywhere consistent with zero. At $\ell = 300$ in particular, all patches are consistent with non-zero power. To find how dangereous these deviations are, we computed the ratio between these null-spectra and the current best-fit ΛCDM model, based on WMAP7 data[79]. The result is shown in figure 6.4(b). These results are not bad: For ℓ's of 100-250, the systematic error is within 10% of the expected signal, even within 5% on several points. For the higher multipoles, the situation is worse, but this is to be expected.

6.2.4 Maps

Figure 6.5 show temperature patch maps for the current data set. They have undergone eigenvalue cutting to remove highly correlated noise modes, and have also been scaled by the gc-gain correction factor of 0.84. For comparison, the figure shows the same fields in the WMAP 7-yr Q-band temperature map. The modes that were cut from the QUIET maps have also been removed from the WMAP maps. The figure also gives the difference between WMAP and QUIET for each of the four patches.

While the QUIET maps are more noisy, it is still easy to recognise the structures seen in the WMAP maps. Given the very modest size of our data set, this clearly demonstrates the excellent sensitivity of the QUIET detectors. However, inspecting the column of difference maps, we see features inconsistent with pure noise. The difference maps do in fact appear to have remnants of CMB signal, implying that the amplitude of the QUIET maps is somewhat too high. This indicates that the gain estimates we have used are not quite correct. Given how roughly those estimates were produced, this is hardly surprising.

6.2.5 Power spectra

We have found maximum-likelihood power spectrum estimates for the four patch maps presented in the previous section. In the published Q-band analysis, the maps were subjected to ℓ-cutting up to $\ell = 25$ before likelihood evaluation. To assess how many modes, if any, we should cut, we look at the χ^2 for each eigenmode of the noise covariance matrix of the difference map. This quantity is computed during post-processing. A high χ^2 indicates a mode that deviates significantly from the condition of white noise in the difference map. In the original analysis, it was seen that the first few modes generally had very much higher χ^2's than the rest of the modes, thus justifying the cut. For this data set, however, no such tendency is found, and we may conclude that there is no need to cut modes before making power spectra. Thus it would seem that the changes

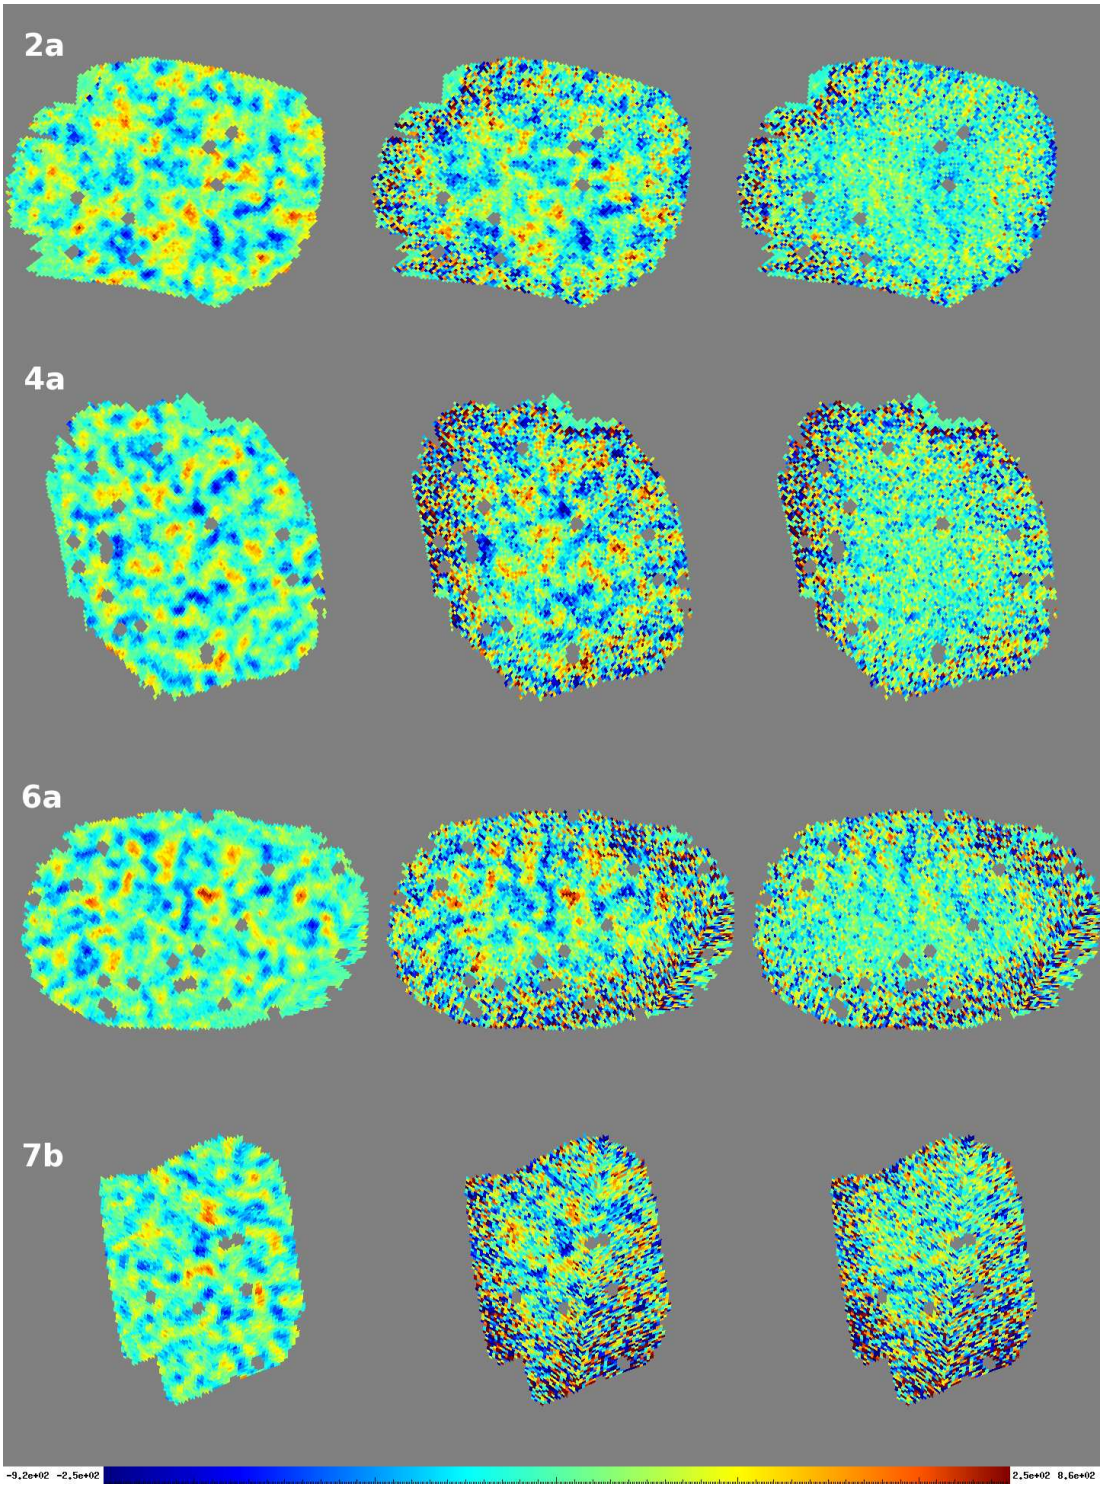

Figure 6.5: Temperature maps for all QUIET CMB patches. Middle column show eigencut QUIET Q-band maps. The left column shows WMAP 7-yr maps, treated to the same cuts as the QUIET maps. The right column gives the difference between QUIET and WMAP. No. of cut eigenmodes: 2a - 324, 4a - 297, 6a - 514, 7b - 63 (Eigenvalue cutoff threshold: 10^{-2}). Temperature range: $\pm 250 \mu$K.

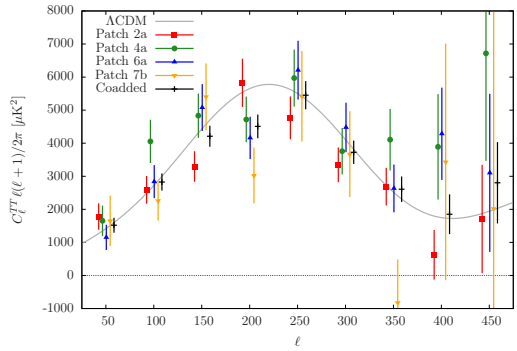

(a) Individual and coadded spectra

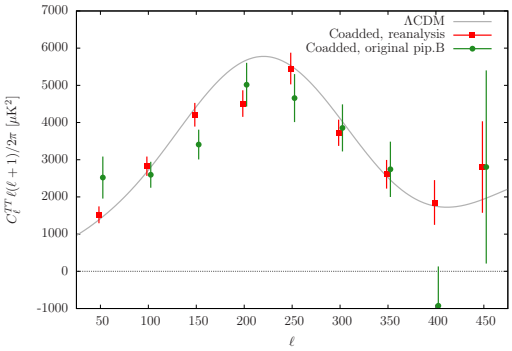

(b) Comparison of coadded spectra from original and new analysis

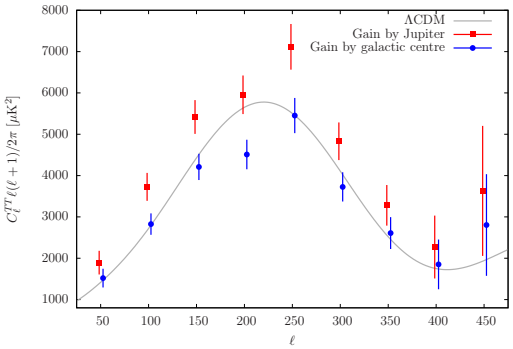

(c) Comparison of coadded spectra from the two gain models

Figure 6.6: Maximum-likelihood TT power spectra for Q-band reanalysis with $\pm 1\sigma$ error bars, as compared to WMAP7 best-fit ΛCDM model[79].

that have been made to the pipeline after the original analysis have, in fact, improved its performance.

From the power spectra of each patch, we can make a combined spectrum by 'co-adding' the components, i.e. computing a weighted average over the patches for each bin:

$$C_b^{tot} = \frac{\sum_i C_b^i/(\sigma_b^i)^2}{\sum_i (1/\sigma_b^i)^2}, \tag{6.1}$$

where the sum runs over all patches. The corresponding variance is

$$\sigma_b^{tot} = \sqrt{\frac{1}{\sum_i (1/\sigma_b^i)^2}}. \tag{6.2}$$

Doing this assumes that for each bin, the four samples are independent and normally distributed with the same mean. This is what we would expect if our patch power spectra were unbiased samples of the true power spectrum. As we found in §6.2.3, we

have not succeeded in cleaning all systematics from our temperature data, so we cannot be certain that there is no bias. A joint analysis of all four maps would have given the total maximum-likelihood spectrum, but for simplicity, we may use the coadded spectrum, which is a fair approximation.

The four patch power spectra are shown in figure 6.6(a), with the coadded spectrum and the WMAP7 best-fit ΛCDM temperature power spectrum[79]. Also shown is a comparison of the coadded spectra for the original and new analyses. This is given in figure 6.6(b). We see a definite improvement. In six of the nine bins, the new value is visibly closer to the ΛCDM curve than the original, and we note that the error bars are smaller. The larger size of the current data set is quite probably the cause of both factors. As this analysis is not as rigid in terms of testing for systematics and bias-removal as is necessary for results that are to be published, we have a data set which is quite a lot larger than what was used for the QUIET Q-band article. Since the CMB temperature anisotropy signal is so stong, we find reasonable results, despite there being contaminations in our data.

Figure 6.6(c) shows the coadded TT power spectra found using the two different gain models discussed in §6.2.1. The spectrum based on galaxy gains has been scaled by the correction factor of 0.71, and we see that it fits quite well to the ΛCDM spectrum. The Jupiter-based spectrum, on the other hand, is too high. This is the reason why the galaxy gains were used as basis for all the results given in this and the previous sections. The poor quality of the Jupiter-based gain model can probably be explained by there being fewer scans of Jupiter than of patch gc, and somewhat more spread in the estimated values. A more thorough gain estimation procedure could be expected to yield better results.

6.3 Polarization analysis

The polarization data analysis is, of course, the most important piece of the QUIET puzzle. In the following are presented the results of a reanalysis of the full set of Q-band polarization data.

6.3.1 Automated data cut results

The results of automated data cuts using standard threshold values as given in table 5.1, are listed in table 6.3. In the original analysis, the accept rates were 84.3%, 70.0%, 61.4% and 74.2%, respectively, for patches 2a-7b. The new accept rates are 10-20 percentage points lower for all patches. This is not ideal, but as there have been some problems with getting the null tests to pass, we have not found room for softening any cut thresholds. In this set, diode 16 has been excluded, as our null test investigations showed it to have poor noise-model fit throughout the season.

Studying the table, we find that the scan frequency χ^2 cut is among the cut criteria that remove the larger portions of data. This is a strong indication of significant ground pickup. As was noted in §5.2.2, the ground model that was used during the original Q-band analysis has since been substituted by azimuth filtering. While this should be a better strategy in the long run, it would seem that the comparably high levels of ground contamination in the Q-band data was removed more completely with the old method.

Cut efficiencies (% of CES-diodes remaining after each cut)

Cut type	Patch 2a	Patch 4a	Patch 6a	Patch 7b
Removing dead diodes	92.65	92.65	92.65	92.65
Static/input cuts	91.59	87.99	84.33	80.85
Gain cut	91.59	87.99	84.33	80.85
Elementary cuts	91.59	87.99	84.33	80.85
APEX cuts	91.59	87.99	84.33	80.85
Type-B cuts	88.95	84.16	77.99	77.52
Weather cuts	86.99	78.35	68.05	73.39
Noise model σ_0 cut	85.96	77.43	67.40	72.52
Noise model f_{knee} cut	84.66	76.39	66.71	71.43
TOD χ^2 cut	84.46	76.32	66.66	71.43
TOD outlier cut	83.30	76.06	66.44	71.37
TOD azimuth cut	82.19	75.12	65.27	70.60
Scan frequency χ^2	77.18	69.60	56.17	59.73
Low frequency χ^2	76.91	69.23	55.85	59.59
All frequencies χ^2	76.91	69.23	55.85	59.59
High frequency χ^2	76.23	68.59	55.28	59.00
1.2 Hz spike	76.23	68.59	55.28	59.00
1.0 Hz spike	76.09	68.48	55.20	58.92
General spike	75.98	68.37	55.13	58.84
Sidelobe sun cut	71.27	62.21	49.08	56.62
Accept ratio cut	69.42	59.97	45.56	54.25
Total no. CES-diodes after cuts	41213	30013	24849	11879

Table 6.3: Summary of automated cut procedures for all CMB polarization data.

6.3.2 Null test results

The final results for the null tests on the polarization data are given in table 6.4. As explained in §5.4, we have tested against a set of 50 simulated, noise-only sky maps. All maps were split according to a set of appropriate jackknives. For each difference map, pseudo-power spectra were computed, with corresponding χ^2-statistics. The χ^2's for the data map are compared to the distribution formed by the χ^2's of the simulations, resulting in the quantities given in the table. PTE (Probability To Exceed) gives the probability of obtaining a χ^2 for the data larger than the one observed. We expect the PTE of the full set of tests to be fairly uniformly distributed between 0 and 1. We supplement the PTE by σ, which tells how far away our statistic is from the mean of the noise-map distribution of χ^2's, in units of the standard deviation.

We find from the tabulated values that this set of tests do not quite pass, although the results are not horrible, either. There are some tests which fail in that they have very high or very low PTE. Here, the corresponding σ's give an insight into the gravity of the situation. For instance, the test splitting the patch 2a data by time of observation (mjd) has PTE=0.0, which is a clear fail. However, χ^2 that gave this value is still only 2.7σ away from the mean of the distribution. (We find that this is, in fact, the test that got the largest σ.) To help us draw a conclusion from all these figures, figure 6.7 shows

Jackknife	**Patch 2a**		**Patch 4a**		**Patch 6a**		**Patch 7b**	
	PTE	σ	PTE	σ	PTE	σ	PTE	σ
mjd	0.000	2.7419	0.200	0.8721	0.740	-0.6876	0.540	-0.1445
inout	0.920	-1.2581	0.340	0.2732	0.200	0.6241	0.780	-0.8092
typeb	0.400	0.2438	0.260	0.5043	0.820	-0.8995	0.020	2.2072
sss	0.100	1.1377	0.120	1.0304	0.100	1.3821	0.560	-0.2225
el	0.120	1.0925	0.660	-0.5614	0.780	-0.8039	0.560	-0.2516
dk{0:45:2}	0.580	-0.3184	0.940	-1.5711	0.540	-0.2930	0.180	0.9207
az{0:180:2}	0.200	0.8061	0.240	0.6298	0.660	-0.5050	0.300	0.5392
gain	0.820	-0.9598	0.380	0.1937	0.760	-0.8515	0.900	-0.9880
fknee	0.440	0.0504	0.000	2.5586	0.640	-0.3902	0.300	0.3594
sigma0	0.660	-0.4979	0.420	0.1147	0.140	0.8778	0.020	2.1596
modfreq	0.620	-0.4709	0.040	2.1051	0.540	-0.1707	0.140	1.1667
weather	0.380	0.2811	0.160	1.1180	0.700	-0.5627	0.160	0.8338
qu	0.180	1.0078	0.400	0.2530	0.680	-0.5540	0.140	0.9282
leak	0.260	0.3587	0.260	0.8704	0.740	-0.4258	0.460	0.0992
Total chisq	0.160	0.8750	0.020	1.7171	0.740	-0.7315	0.060	1.6225
Max chisq	0.260	0.4042	0.400	-0.0419	0.640	-0.4821	0.300	0.2927
Mean chi shift	0.080	1.5601	0.080	1.6296	0.000	2.5905	0.000	2.5572

Table 6.4: Summary of null test suite for QUIET Q-band polarization data.

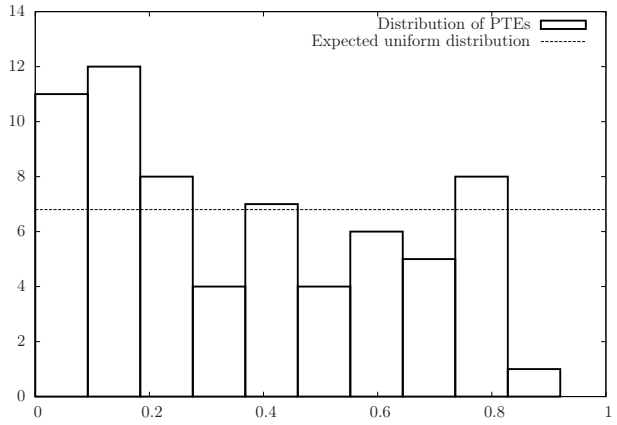

Figure 6.7: Histogram of all PTEs for the CMB patch polarization null suite.

a histogram of all the tabulated PTE's, with an indication of the expected uniform distribution. We find a bias towards low PTE's, i.e. high χ^2's, and this is mirrored by the mean chi shift, which is also showing signs of a positive bias.

Before settling for a configuration, we performed several sets of null tests, while varying mainly the filter parameters. To assess what caused the tests to fail, we studied the distributions of various CES-diode statistics, most notably the distribution of the CMB frequency interval χ^2. This quantity is normalised, and should be distributed as $\sim \mathcal{N}(0,1)$. A high CMB χ^2 indicates that the noise model is set too low, while a low χ^2 means the noise model is too high. We found the full distribution of CMB χ^2's to be too wide, ($\sigma \approx 1.34$). This seemed to be caused by some diodes having significant positive bias, while others had negative bias, thus causing the full distribution to have mean approximately zero, but with too much spread. We also found a trend in which CES-diodes with low f_{knee} had high χ^2's, and vice versa. These facts seemed to indicate a somewhat poor fit of the noise model. To amend the situation, the lower limit of the frequency interval used for the initial noise model estimation was moved from 2 Hz to 0.3 Hz, to get a better basis for fitting the model. This limit was originally high to avoid the σ_0-estimation becoming biased for CES-diodes with high f_{knee}. A correction for this has been added to the pipeline (adjusting σ_0 down if f_{knee} is high), so the wider interval should now be safe. This reduced the spread of the distribution of the χ^2's, but there is still a positive bias in CMB χ^2. This looks to be the main problem with this data set. With ground contamination in mind, this final set of tests, as well as the subsequent analysis, has been run using lower limit values on the highpass and azimuth filters, to be used instead of the individually fit filter parameters if these are below the limits. The values used were $\nu \geq 2.5\nu_{scan}$ for the highpass filter frequency, and polynomial order 10 for the azimuth filter. We do this to strengthen the filters, although it decreases our final signal-to-noise ratio in the maps.

It should be noted that the jackknife "modfreq", splitting the diodes by their bandpass centre frequencies, should not have been part of the Q-band null test suite, and was included by mistake. We have not registered individual module centre frequency variations for the Q-band array - they are all given the same value (44 Hz). Using this test on the Q-band data gives a completely random split, and we should not expect the resulting numbers to give any information. One should keep in mind that this test can have had an influence on the three final parameters, and as such the null test suite should at some point be run again. However, it does not affect the final conclusion, since there are other individual tests failing.

In a strict first-time analysis, we could not have moved on before these tests had all passed, but since the Q-band data has already been analysed once, the "box is open", and we can allow ourselves the luxury of continuing with the analysis despite these somewhat sketchy results.

6.3.3 Maps

Having settled for a set of cuts and filter limits, we move on to producing our final data products. The maps are post-processed differently depending on whether we want to study them or use them for power spectrum estimation. In the latter case, we only cut a few modes, and only if we have reason to believe they might distort our estimates. On the other hand, when we want to study the maps themselves, modes of high variance

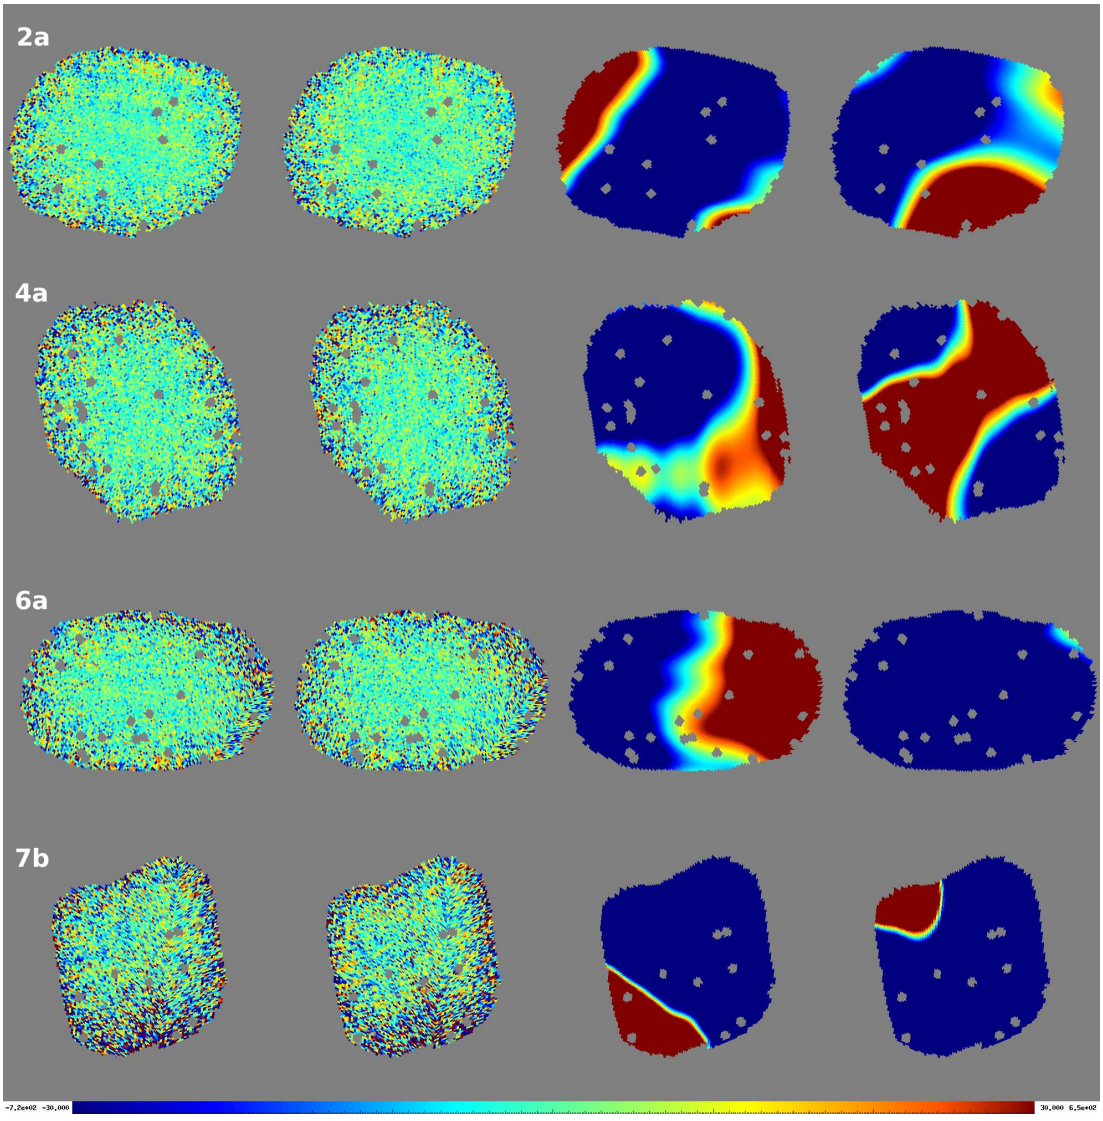

Figure 6.8: Polarization reanalysis maps for Q-band. In the two columns on the left are maps of Stokes' components Q and U, while the columns on the right show the modes that have been cut from these maps. No. of cut eigenmodes: 2a - 46, 4a - 46, 6a - 54, 7b - 83 (eigencut threshold 10^{-2}). Temperature range $\pm 30\mu$K for all maps.

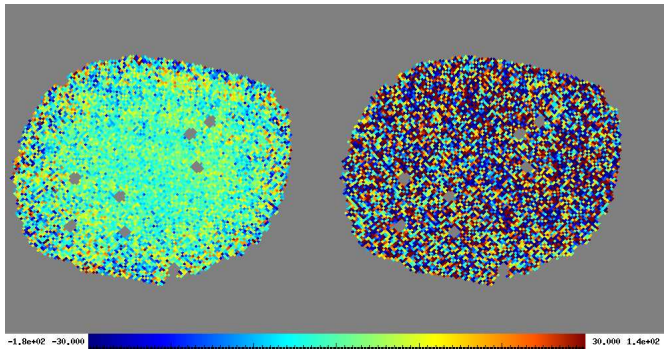

Figure 6.9: Q polarization maps of patch 2a, for QUIET Q-band reanalysis (left) and WMAP 7-yr Q-band (right). The temperature range and no. of cut modes is identical to the case in fig. 6.8. The QUIET modules' superior sensitivity is evident.

and amplitude need to be removed, since they overshadow the interesting modes while contributing no information. By testing a few possibilities, we find a suitable threshold value for the $\mathbf{N_m}$ eigenvalues, such that the resulting maps are sufficiently clean.

The maps for this analysis are given in figure 6.8. In addition to the maps themselves, maps of the eigenmodes that have been cut have also been included, to highlight the large difference in amplitude between the cut modes and those we keep. This is evident in that the temperature range suitable for polarization maps causes the large-mode maps to be almost completely saturated.

When looking at the temperature maps, we had the benefit of being able to compare our maps to WMAP results. But as we know, WMAP had little polarization sensitivity, so this comparison is useless here. To demonstrate, figure 6.9 compares the Q-polarization maps of patch 2a, for QUIET Q-band and WMAP.

In figure 1.3 we saw an example of the polarization patterns predicted by theory. Despite our cleaning efforts, there is no such visible signal in the final maps. Given the relative strengths of the E-modes and the synchrotron foreground in the Q-band, which we saw in figure 4.5, this is as expected. The Q-band array is small and intended to give us control of the foregrounds. Although the Q-band data is sufficient to give positive E-mode detections, as was found in the original analysis, it is not surprising that it is not visible in the map.

6.3.4 Power spectra

As with the temperature data in §6.2.5, we look to the χ^2 per eigenmode of the difference map made from a deck-angle split of the data, for guidance on how many modes must be cut from the maps before making the power spectra. What we found is shown in figure 6.10. For the plot to be legible, only the first 100 modes are shown, but there were no visible differences in the graphs' behaviours though the rest of the modes. Now, we see that the first few modes do not stand out in any way. The very first eigenmode for patch 2a is rather high, but there are comparable peaks further on, and we see no systematic trend of higher χ^2's in the lower end. From this we can draw the conclusion that there is no immediate reason to eigencut these data.

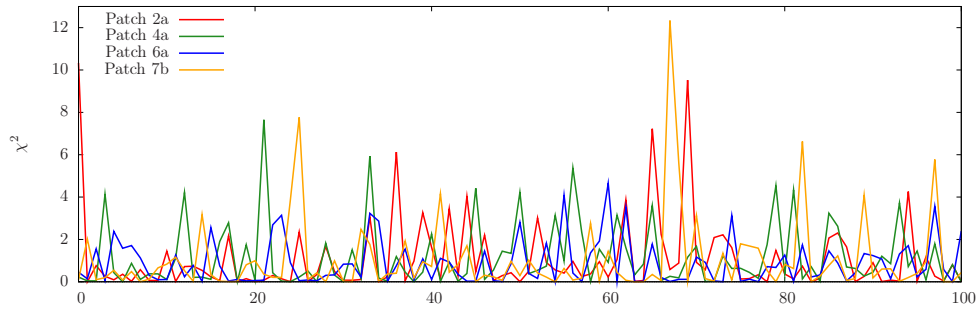

Figure 6.10: χ^2 per eigenmode for the first 100 modes of the deck-angle split CMB polarization data sets.

The set of power spectra eventually chosen is given in figure 6.11, shown with WMAP7 best-fit ΛCDM spectra for comparison[79], and the coadded spectra from the original Q-band analysis. There are some issues with the error bars in the first bin - they seem falsely small for two of the patches. An indication towards this is the relative size of the error bars for the different patches. We expect the statistical error of each data set, quantified by the error bar, to correlate with the amount of data available. The more data, the less the uncertainty. From table 6.3 we find the number of accepted CES-diodes in each patch, and conclude that we should expect patch 2a to have the smaller error bars, followed by 4a, 6a and 7b. For the most part, this seems to fit well with what we see in these plots. But in the first bin, patches 4a and 6a have far smaller error bars than 2a, in both the EE and the EB power spectra.

After looking at the set of four power spectra (and the corresponding coadded spectrum) for three different eigencut levels (0, 10 and 20 modes), I chose to cut 10 modes from the patch 4a and 7b data, while cutting nothing from the other two, since this was the combination giving the lesser number of troublesome errorbars in the first bin. However, this bin is still a little on the dubious side. For instance, the first bin of the coadded EE spectrum shown in figure 6.11(a) seems to be almost completely determined by the patch 6a data point. The reason for these problems are found in the way we compute the point of maximum likelihood, and the corresponding error bars, as we saw in §5.6.2. For most bins, the likelihood peak is found well within the region where $\mathbf{C}$ is positive definite. We use the second derivative at the maximum likelihood point to estimate the error bars, a method that works well as long as the likelihood is reasonably bell-shaped around this value. However, sometimes the likelihood peak is found on the boundary of the permissible area. Then the local second derivative will, as a rule, not be a good way to find the error bar. Typically, the local likelihood will be close to linear, and the second derivative correspondingly small. This gives falsely small error bars, like we see for both patches 4a and 6a in the first bin.

To resolve these issues, we could possibly implement an alternative way of computing the error bar for instances where the likelihood maximum is found on a boundary. A quicker solution is to run a joint analysis on all four patches.

The first bin was dubious for the BB spectra as well, but in this case, we have a neat solution. We did not, of course, detect B-modes in the original Q-band analysis,

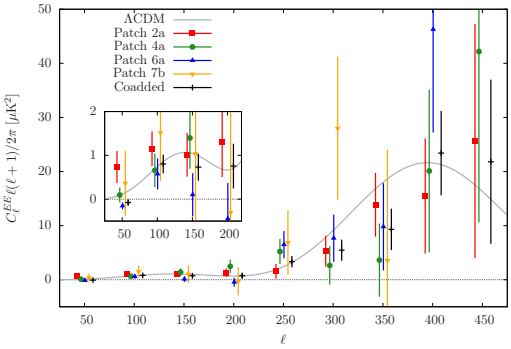

(a) EE, all patches and coadded

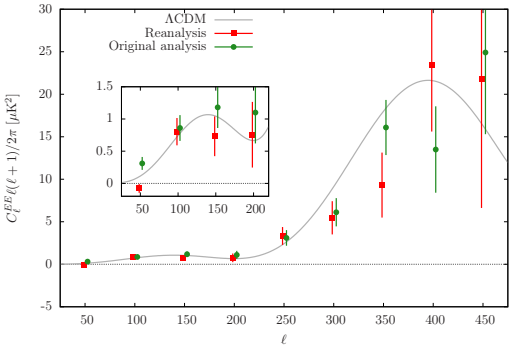

(b) EE, original and reanalysis, coadded

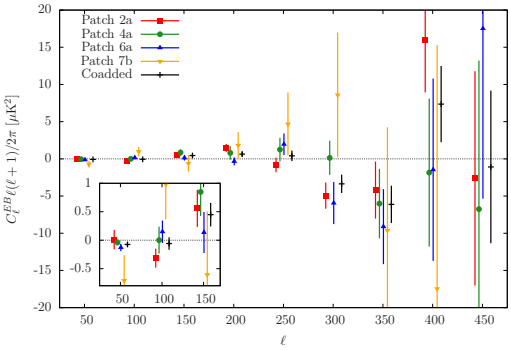

(c) EB, all patches and coadded

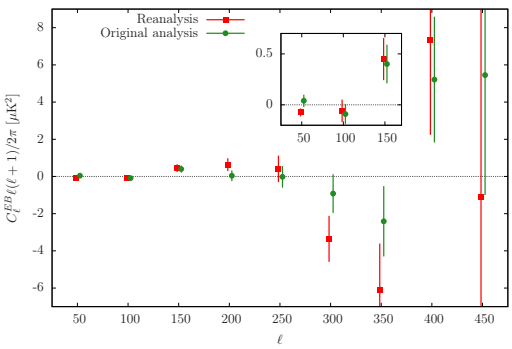

(d) EB, original and reanalysis, coadded

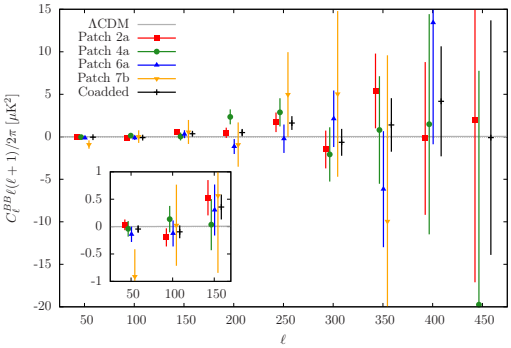

(e) BB, all patches and coadded

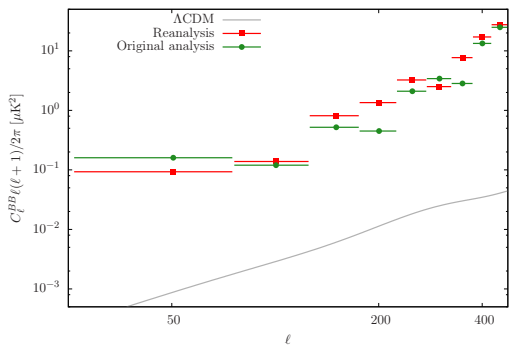

(f) BB, original and reanalysis, 2σ upper limits to the coadded spectra

Figure 6.11: Maximum likelihood polarization power spectra with $\pm 1\sigma$ error bars, as compared to WMAP7 best-fit ΛCDM model and original Q-band analysis results.

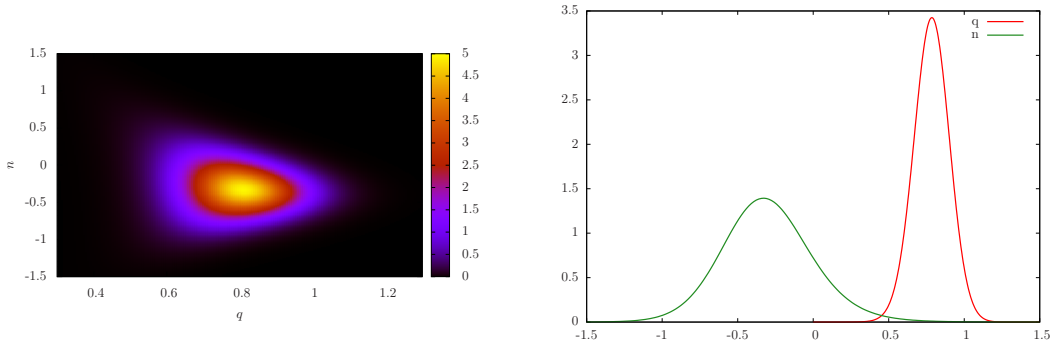

(a) Likelihood function for q and n

(b) Marginal probabilities for q and n

Figure 6.12: Likelihood and probability densities from parameter estimation on coadded EE spectrum

so we expect no such thing now, and the B-mode power spectrum should accordingly be consistent with zero. We have the opportunity of computing the final error bars under the assumption that the best-fit spectrum is zero. We have applied this to the presented BB power spectra, to get more credible error bars.

The original Q-band analysis concluded that the first bin of patch 2a showed signs of foreground contamination, since it had more power than expected from the ΛCDM model, to 3σ significance. Looking at figure 6.11(a) we find this to be consistent with the new spectrum, as the 2a value of the first bin is significantly larger than the other values. As for the rest of the spectra, the results look reasonable enough. The errorbars on the new spectra are slightly larger than for the original analysis, which is a consequence of our smaller accept rates.

6.3.5 Estimates of cosmological parameters

Based on the power spectra we found in the previous section, we can make estimates of cosmological parameters, following the treatment in §5.7. For the sake of lucidity, however, the estimates for the EE spectrum presented here are based on somewhat simpler assumptions.

Following eq. 5.23, we will find the two parameters q and n based on the EE spectrum, utilising the coadded spectrum. To find the point of maximum likelihood, we use simple grid evaluation, like in the full analysis, but rather than using the full likelihood function given in eq. 5.22, we employ a simpler, Gaussian likelihood function

$$\mathcal{L}(C_b) = \prod_b e^{-\frac{1}{2}\left(\frac{C_b - C_b^d}{\sigma_b^d}\right)^2},$$

where the product runs over all bins b, and d identifies the binned power spectra and corresponding σ's we have estimated based on the data. This expression is then evaluated for a set of different test power spectra C_b, parameterised by q and n.

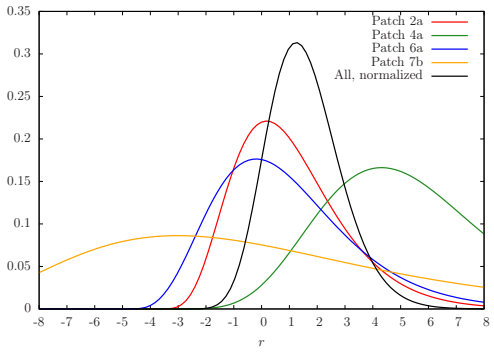

Figure 6.13: Likelihood functions from parameter estimation on BB spectrum.

The likelihood function obtained from the coadded EE spectrum shown in figure 6.11(b) is mapped in 6.12(a). Because of the poor error bars in the first bin for patches 4a and 6a, and the signs of foreground contamination in patch 2a, the first bin of the coadded spectrum has been substituted by the first bin of the patch 7b spectrum for this parameter estimation. (It was found that while removing patches 4a and 6a from the first bin had significant effect on the resulting parameters, removing patch 2a had very little effect.) As fiducial, I used the current WMAP7 best-fit ΛCDM spectrum, including lensing effects and a small tensor contribution ($r < 0.36$)[79]. The marginal probabilities for q and n are plotted in 6.12(b). The maximum likelihood point of the joint distribution is at $q = 0.81, n = -0.33$. For the marginal distributions, we find $\bar{q} = 0.79, \sigma_q = 0.12, \bar{n} = -0.30, \sigma_n = 0.31$, while the peaks are at $q = 0.79, n = -0.33$. (The slight discrepancy between the average and maximum likelihood point for n is due to a slight skewness in the distribution.) These values state that the best-fit EE spectrum should be slightly lower than the ΛCDM spectrum, and tilted so that the high-ℓ end is somewhat decreased. Looking at figure 6.11(b), this seems a reasonable outcome.

For estimating r, we attempted the same strategy as for q and n. However, the outcome was a highly improbable one, and we concluded that the simplified likelihood function was insufficient in this case. Hence we have used the full likelihood for estimating r. The fastest way to find the total likelihood for all four patches is to analyse each patch separately, and multiply the resulting curves. The likelihood curves for individual patches and the total are shown in figure 6.13. The corresponding r estimate is $r = 1.28 \pm 1.29$, i.e. $r < 3.86$ (95% CL). We conclude, for a start, that our data are consistent with no tensor perturbations.

For comparison, the r-estimate from pipeline B in the original Q-band analysis was $r = 0.52^{+0.97}_{-0.81}$. The results of the reanalysis of the polarization data are consistently poorer than what was found in the original analysis. Indications towards this being caused by ground pickup are strong, although further investigations will have to be made to confirm this hypothesis. Since we have also found improvements in some areas, we permit ourselves to be optimistic about the overall performance of the pipeline after its renovation. If ground pickup is found to be the only cause of the current problems, we should be in the clear with regards to future analyses, since we are not about to repeat the mistake of having a too small ground screen.

Part III

- Start the future -

Predictions and conclusions

Chapter 7

QUIET-II

Prediction is very difficult, especially if it's about the future.

Niels Bohr

The first phase of QUIET is nearing completion. The W-band analysis is in its last stages, and will hopefully be ready for publication before long. But as we know from chapter 4, the plans for QUIET includes a second phase, QUIET-II. A proposal has been submitted to the Major Research Instrumentation (MRI) program of the U.S. National Science Foundation (NSF), from which we hope to get funding to build our new, larger, and better instrument.

Despite the project's future being still uncertain, the QUIET-II participants are not idle. There are many matters, large and small, that needs to be dealt with in the planning stage. We need to make certain that everything we have learnt from QUIET-I is taken into account, so that we take full advantage of the successful features, and no mistakes are repeated. Evaluating and, if necessary, making adjustments to the set of observed patches is among the matters that need attention, and in this, I have been involved. I will describe that process in some detail, as an example of the considerations that lie behind the final specifications of a project like ours. But first, I will present the key facts and figures of QUIET-II, and the lessons we have learnt from its predecessor.

7.1 Project specifications

QUIET-II will for the most part be a repetition of QUIET-I, only larger. The prime objective is the detection of the primordial B-mode polarization power spectrum, or if not, then to constrain the tensor-to-scalar ratio to the level of $r < 0.01$. Our secondary goals are to make improved E-mode measurements, detect small-scale lensed B-modes, and constrain the low-frequency galactic synchrotron power in the low-foreground parts of the sky. [59]

The main detector array will be a 499-element W-band (95 GHz) array, with 483 polarization-sensitive modules, and the remaining 16 devoted to temperature measurement. This array will be supplemented by a 13-element Ka-band array. The smaller array will cover the full Ka-band, with 3 detectors at 30 GHz (Ka-low) and 10 at 37 GHz (Ka-high). Two of the Ka-high detectors will be differential temperature monitors, while the rest will be polarization-sensitive modules. The Ka-band modules will have twice the sensitivity of the W-band modules (noise properties of the amplifiers are frequency-dependent). This, combined with the ν^{-3} frequency dependence of the

synchrotron foreground, which we saw in figure 4.5, means that the Ka-band season will achieve twice the synchrotron signal-to-noise ratio of the W-band season. Hence, should the Ka-band detect the synchrotron foreground, we may estimate its amplitude as well as its spectral index, thanks to the two Ka-band frequencies. If we do not detect any synchrotron foreground at Ka-band, we may safely conclude that any signal detected by the W-band array cannot be this foreground. Since using the Q-band for this would demand many more detectors and longer integration time, we will not be taking any more observations in the Q-band. [59]

If all goes according to plan, the instrument will be fully assembled and tested by mid-2015, and will then be fielded at Chajnantor Observatory, where it will gather data for 3.5 years. Like in QUIET-I, we will field the two arrays sequentially. The Ka-band array will be mounted first, and will observe for 6 months, taking an estimated ~ 2000 hours of good data, given accept rates similar to those of QUIET-I. It will then be replaced by the W-band array, which will observe for a full 3 years, from which we expect to gain a total of ~ 12000 usable data-hours. [59]

The QUIET group at the University of Oslo will be the data analysis centre for QUIET-II. The amount of data will necessarily be a lot larger than for QUIET-I. The pipeline has been written with this in mind, meaning that our facilities are fully capable of dealing with this.

7.2 Learning from QUIET-I

As we noted in chapter 4, one of the main goals for QUIET-I was for it to be a pathfinder for its larger successor. Naturally, we need to investigate the path it found, in order to make sure QUIET-II becomes every bit as good as its potential. Here is listed a selection of consequential factors.

Mount and ground screen The plans for QUIET-II includes purchasing a new, modern mount. Uncertainties of pointing was one of the dominating sources of systematic error in the QUIET-I data, brought on by the suboptimal performance of the old CBI mount. For the next phase, we will have a mount designed for fast and frequent pointing changes, which is necessary for our scanning strategy. The new mount will also manage a greater elevation range, allowing us to observe as low as the horizon. Although CMB observations must necessarily be taken well away from the horizon to avoid ground contamination, this will give us the option of on-site calibration by observing artificial sources on the ground. [59]

Including a proper ground-screen is a self-evident point. As we have seen in the previous chapter, having an unsufficient ground screen allowed significant contaminations into our Q-band data, and our experience from the W-band season goes to show that we have a lot to gain from this simple measure.

Telescope and optical design The design of the telescope and main optical components was found highly satisfactory, and will be kept. The corrugated feed horns are easily manufactured and have good performance. For cost efficiency, the design have been scaled up by a factor of 1.7, allowing us to reuse the original 91-element W-band feedhorn array. This will form the central component of the 499-element array, which will be completed by the addition of six 68-element feedhorn arrays

arranged around the central piece like petals. To accomodate the large array, we will need a 2.4 m telescope. [59]

The temperature modules should be distributed evenly around the focal plane, rather than being gathered at one side like they were in QUIET-I. As the prime purpose of these modules is calibration, they should be representative of the whole focal plane. Our pointing calibration would benefit greatly from this.

Hardware calibrators For calibration, QUIET-I has used external sources, such as WMAP, and observations of the sky. This strategy is not ideal, and the analysis has been delayed more than once because of problems with the calibration. In QUIET-II, we would preferrably have specialised equipment for calibration. For instance, one might use artificial sources placed in front of the telescope to calibrate the beam and polarization angles.

Module optimisation Significant progress has been made on the detector front since the fabrication of the QUIET-I instrument. When building the QUIET MMIC modules, we could not test the modules' noise and gain properties at cryogenic temperatures without destroying them. Hence the arrays had rather a lot of spread in these quantities, and some diodes turned out to be altogether useless. Since then, however, new technology has been developed at the California Institute of Technology (Caltech), and we are now able to test the components thoroughy, and pick those that have the better performance. Furthermore, HEMT amplifiers with even lower noise levels have been developed. While the theoretical noise temperature per module was ~ 95 K in the QUIET-I W-band, the new amplifiers will bring this quantity down to a mere 36 K. [59]

An additional design change is that the QUIET-I monolithic modules have been split into three separate parts. While still fitting into the same small footprint, this makes production and testing easier. The type of waveguide (the element that combines the modulated signal into the linearly polarized components) has been changed to a type which has less signal loss. All of these improvements add together, and we now project a W-band module sensitivity of 300 $\mu\mathrm{Ks}^{1/2}$, which is far better than the original modules. With the low sky temperature at the observation site and 3 years of observation time, this is estimated to give us a map sensitivity of $0.09\mu\mathrm{K}/\deg^2$, which should be enough to constrain the tensor-to-scalar ratio almost to the level of the Lyth bound. [59]

7.3 Patch selection

The set of fields observed in the first phase of QUIET, including targets for calibration, were given in table 4.1. These were chosen for their low foreground and good visibility: Since our observation site is at a low latitude, the greater part of the sky rises and sets each day. While experiments in Antartica can observe the same patch around the clock, we have to switch as each patch moves below our field of view. Hence our CMB patches should be spaced so that one is always visible, ensuring optimal usage of observation time.

Other factors must also be considered when choosing patches. For instance, we would prefer it that our patches are not overshadowed by the Sun or Moon for a large

part of the available time. Also, we want to coordinate our patch choices with other projects. QUIET-II will measure or constrain the galactic synchrotron radiation in low-foreground parts of the sky, areas that have yet to be charted to the accuracy needed for cleaning CMB maps. Hence our measurements will be a valueable contribution to other CMB polarization experiments that only observe on higher frequencies, and are unable to assess any synchrotron contamination. On the other hand, we have no way of constraining the dust foregrounds, which dominate at frequencies > 100 GHz. Sharing observation patches with other projects will thus be a win-win situation, where we can share our foreground knowledge for the common good.

Our original patch choice has been put to the test in QUIET-I, and we must now assess whether we think we can do better for QUIET-II. As a starting point, we need a reference map by which we can judge the foreground power. We use the WMAP 7-year K-band map, since the synchrotron foreground is most intense here. From the Q and U polarization maps, we compute an RMS map of the total polarized power, $P = \sqrt{U^2 + Q^2}$. The map is shown in figure 7.1(a). In this map, each pixel gives the RMS value of P in a surrounding disk of radius 10°. This corresponds roughly to the size of one of our CMB patches, meaning that a patch centered at one good pixel in this map should have little foreground throughout. In figure 7.1(a) the CMB and galactic patches of QUIET-I have been indicated, and we see that they are all within the colder areas of the map[1], although patch 2a lies rather close to areas where the galactic radiation is more intense.

Based on our experience with QUIET-I, we have found a few points for which improvements are in order:

- During Q-band analysis, evidence of synchrotron foreground contamination to 3σ significance have been detected in patch 2a[58]. This patch should accordingly be moved further away from the galactic centre.
- Patch 7b lies too close to 6a, and ended up with too little observation time. It has not been of much use and should be moved or dropped.
- Patch gb was not of as much use as we would have liked, and if another area could be found covering more interesting sources, this patch could well be moved.
- Tau A was not included in the set of patches recieving daily observations during QUIET-I. Given its very important role in calibration, it seems appropriate to include it in this set, and assign it daily observation time of ~ 1 hour.
- The QUIET-I patches had little overlap with other projects. Achieving some overlap is considered to be of significant importance for the next phase.

7.3.1 The selection procedure

I have written a small utility program for patch analysis, to provide the information necessary to assess and adjust the set of patches. The program computes the foreground RMS map, and sorts the observable pixels by foreground power. The map and list of

[1] The brighter area near patch 4a is an artefact caused by WMAP's scanning pattern. A similar spot is visible in the upper left quadrant.

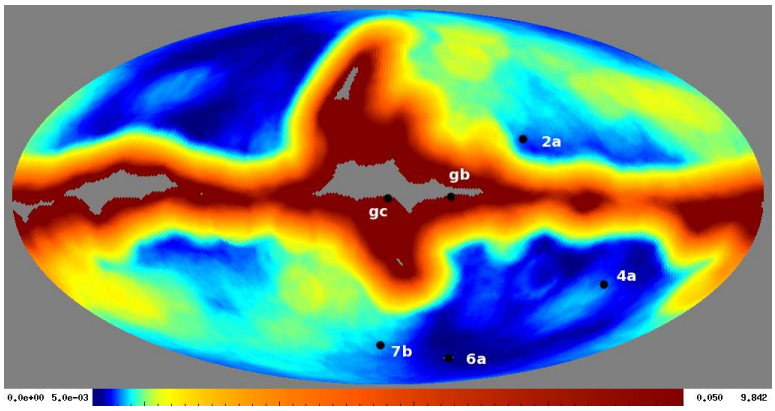

(a) RMS map of polarized K-band sky, with QUIET-I patches indicated. Temperature range 0.005-0.05 mK.

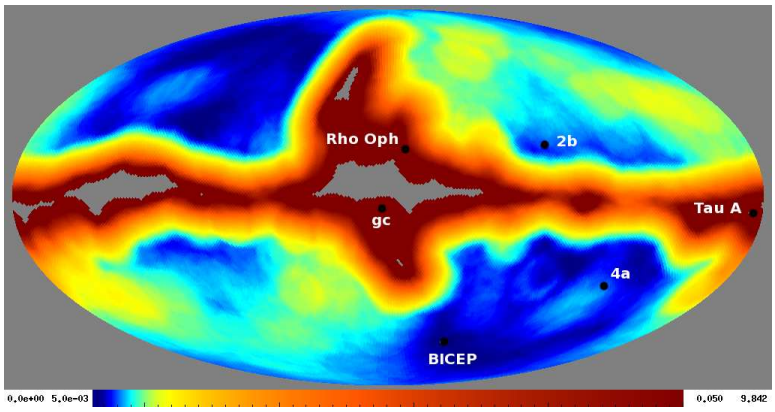

(b) The proposed set of patches for QUIET-II, on the same foreground map.

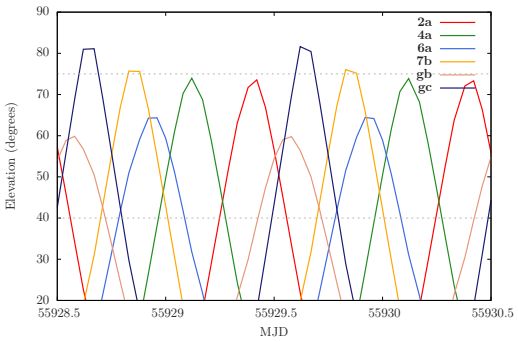

(c) Local patch elevation vs. time (MJD) for QUIET-I patches

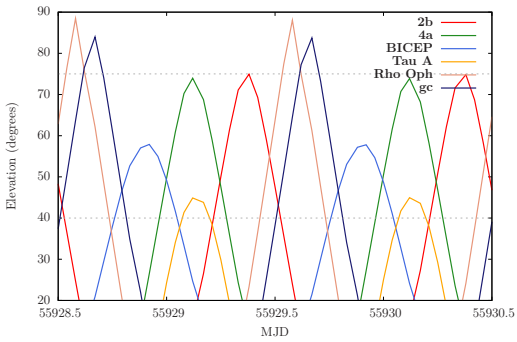

(d) Local patch elevation vs. time (MJD) for proposed QUIET-II patches

Figure 7.1: RMS maps and patch elevation plots for patch selection.

QUIET-I observed patches

Patch	Long	Lat	RMS WMAP7 [mK]	RMS Planck [μK]	Visible hours	Sun	Moon
Gb	329.10	0.00	2.66806E-01	3.273E+01	7.0823	0.0000	0.0007
Gc	0.08	-0.07	8.23478E-01	6.667E+01	5.4146	0.1800	0.0073
2a	292.20	22.80	7.88654E-03	2.485E-01	7.4879	0.1496	0.0048
4a	243.20	-35.30	7.99249E-03	1.354E-01	7.3262	0.1524	0.0048
6a	304.60	-69.10	6.65492E-03	2.200E-01	7.4872	0.1008	0.0013
7b	7.00	-62.00	8.34647E-03	3.512E-01	6.3491	0.1645	0.0059

QUIET-II patch proposal

Patch	Long	Lat	RMS WMAP7 [mK]	RMS Planck [μK]	Visible hours	Sun	Moon
Tau A	184.56	-5.97	1.15656E-01	3.398E+00	2.9401	0.0098	0.0092
Rho Oph	350.96	18.91	6.44300E-02	6.051E+00	5.1449	0.1731	0.0078
Gc	2.70	-4.60	1.39425E-01	1.058E+01	5.3726	0.1779	0.0079
2b	282.20	20.80	7.76493E-03	2.533E-01	7.4904	0.1482	0.0050
4a	243.20	-35.30	7.99249E-03	1.354E-01	7.4872	0.1524	0.0048
BICEP	318.00	-60.60	6.42952E-03	2.637E-01	6.9457	0.0000	0.0006

Table 7.1: Observed and proposed patches, QUIET-I and QUIET-II.

cold pixels are used as a guide when looking for possible new patch locations. When coordinates for test patches are provided, the program will compute the elevation of the patch centre, in local horizontal coordinates, as function of time. It also outputs other useful quantities, as we will see shortly.

Plots of elevation as function of time for a set of patches were among the prime tools for assessing a set and judging where changes should be made. Figure 7.1(c) shows such a plot for the original QUIET-I patches. The mount limits the available range of elevations to $[40°, 75°]$. It is of the utmost importance to have at least one patch within this interval at all times. Hence we use plots like figure 7.1(c) to look for gaps and find where to move the patches to achieve the optimal utilisation of the available time.

After numerous iterations of choosing and evaluating sets of coordinates, I arrived at the proposed set of patches listed in table 7.1. This has been submitted to the QUIET team for further evaluation. The table summarises the important quantities of each patch for both the old and new sets, for comparison. The listed quantities are the galactic coordinates of all patch centra, their RMS polarized foreground levels with respect to WMAP 7-yr K-band skymap as well as the Planck Skymodel[80], the maximum available observation time per day, and the fractional amount of a year during which the Sun or Moon is within a 30° radius of the patch centre. The patch centra are indicated in the foreground RMS map in figure 7.1(b), and the corresponding elevation plot is shown in figure 7.1(d).

The following should be noted about this set of patches:

- A patch covering the Rho Ophiuchi cloud complex has been included as a replacement for patch gb. This object is a dense, dark, star-forming nebula, and a

possible source of a type of spinning dust emission peaking in the W-band, observed in the Magellanic clouds[81]. It is believed that deep polarization maps of such objects may be a valuable contribution to our knowledge of the galaxy and CMB foregrounds.

- The two foreground patches, patch gc and Rho Oph, are not quite centered at the objects of their names. The galactic centre is, by the definition of the coordinate system, at (0.00,0.00) in galactic long/lat coordinates, while Rho Oph is centered at long/lat (353.22, 16.53). The patches were moved a short distance away from these points because the original coordinates caused both patches to be above the upper elevation limit at the same time, creating a gap where no patches were within the visible area. Due to the size of the observed areas, both objects will still be well within their respective patches.

- The patch called BICEP has been as a replacement for QUIET-I patches 6a and 7b. This patch coincides with the observed area of the BICEP experiment. The BICEP patch has already been observed at 100 GHz and 150 GHz, for three seasons (2006-2008)[48], and observations will continue with BICEP2[82]. Overlap with this project would hence be of great advantage to us as well as to the BICEP team and the community at large. The patch is in a suitable area for QUIET, with low foregrounds and little contamination by the Sun or Moon. Ideally, its maximal elevation should have been a little higher. However, with the elevation limits given in this plot, the patch is observable for $\lesssim 7$ hours, so if we observed it through the full available time, that should be sufficient. Furthermore, the elevation limits used here are certainly conservative, as we expect to have a far better mount in QUIET-II, so the elevation of this patch may not be a problem at all.

- Patch 2a has been moved approximately 10° in an attempt to lessen the galactic foreground contamination. To stress that the patch has been changed, I have dubbed the new variant patch 2b.

- Patch 4a performed well during phase 1, and there were no obvious reasons to move it, so it has remained unchanged, as the only patch of the set. This patch has significant overlap with the planned target of the balloon-borne project EBEX, which will be centered at $\sim (250, -38)$ (galactic coordinates), and cover approximately $350\,\mathrm{deg}^2$[83].

Chapter 8

Conclusion

> *Space is big. You just won't believe how vastly, hugely, mind-bogglingly big it is. I mean, you may think it's a long way down the road to the drug store, but that's just peanuts to space.*
>
> *Douglas N. Adams*

We have come to the point of conclusion. Through this thesis, we have investigated the subject of QUIET data analysis, in a process that could deservedly be called a "crash course in observational CMB cosmology". The chief motivation for this endeavour has been to gain an understanding, as complete as possible, of how one finds answers from data, and the obstacles one meets along the way, in this thriving field of research. Now is the time to look forward, but first, we should summarise what has been done and what we have found.

8.1 Thesis summary

The first part of the thesis gave a broad background, with the primary purpose of providing motivation and the necessary background information for the subsequent treatment. We began with covering the basics of our current knowledge of cosmology, including an introduction to the concept of power spectra, the cosmologist's weapon of choice. We then gave a short historical review of the development of modern cosmology from its beginnings in the early 20th century. Moving on to more central matters, a chapter was devoted to the theories of cosmological inflation, which form the theoretical foundations for the ongoing search for the elusive B-modes. In the final introductory chapter, we presented the QUIET experiment.

The main work of the thesis was presented in the second part. The pipeline chapter described the full process of data analysis using the maximum-likelihood pipeline. Starting at where we recieve the data, we moved through the initial tasks of calibration, data selection and null testing, in which the large quantity of raw data was split into appropriate pieces, sorted, scrutinised, filtered, and quality controlled. We then moved on to tackle the problem of mapmaking. The transformation of CMB data into maps has become a self-evident piece of the data analysis machinery for most, if not all, CMB observational projects, as it is an excellent method of data compression while at the same time being the most intuitive way of inspecting the data. We derived the mapmaking equation for the maximum-likelihood map, and discussed the alterations needed to take filtering into account. We looked at the possibilities for post-processing of the

maps, using masks and the cutting of eigenmodes of the noise covariance matrix to remove unwanted features from the maps. We then described how QUIET uses a local quadratic estimator to find the maximum-likelihood power spectra, and finally how we estimate best-fit cosmological parameters using brute-force grid evaluation.

Chapter 6 gave the results of the full Q-band reanalysis. The analysis had three main parts, treating the galactic observations, CMB temperature data and CMB polarization data separately. We made maps of the galactic patches, and looked at the consequences of uncritical filtering. In the analysis of the CMB temperature data, we encountered problems with the gain model, and had to make new gain estimates. Having been less strict in cutting possibly contaminated data, our TT power spectra came from a significantly larger data set than what was used in the original Q-band analysis. Hence we found smaller error bars, and a visibly better fit to the fiducial ΛCDM model spectrum. In the polarization analysis, we encountered problems in the null testing phase. We found a positive bias which we were not able to remove completely, although we strongly suspect the notorious ground contamination to be the culprit. Despite this, we carried out the rest of the analysis, using somewhat less data that what was used in the published results. The final power spectra and the estimate of the tensor-to-scalar ratio r, consequently came out somewhat poorer than what was originally found.

In chapter 7, we looked ahead to QUIET-II, the upscaled successor of the nearly completed QUIET-I. We reviewed the plans for the project, and discussed some points where our experiences from the pathfinder project has demonstrated room for improvement. We rounded off with an account of the evaluation and re-adjustment of the set of observed patches in preparation for QUIET-II.

8.2 Outlook

Although the publication of the QUIET-I W-band analysis should be just around the corner, there are ample possibilities for doing more science on the QUIET data. High on the list is the completion of the work done in this thesis. Perhaps the most important point would be to continue the investigations into the polarization data null test results, to find what caused the tests to fail, and confirm or disprove the hypothesis that ground pickup is the cause. A proper ground screen is a self-evident component in a continuation of QUIET, so if this is what causes the bias in the reanalysis of the Q-band data, we should be in the clear. However, if this is not the cause, we should need to do further investigations, to make quite certain that there are no unknown deficiencies in our pipeline.

The road ahead will depend on whether we recieve the funding we need to build the QUIET-II instrument. If not, we will have to make do with what we have, which could for instance mean delving into the analysis of the foreground patches. There are also plans for a project called FOCUS, which involves renewed observation with the QUIET Q-band array. The pipeline could be stream-lined so as to be applicable to other projects.

The optimal outcome would, of course, be that we are granted the funds for which we have applied and are able to bring QUIET-II to life. I am hopeful that this thesis has succeeded in conveying the potential of the QUIET experiments. The future holds great promise for the field of CMB cosmology, and it is my belief that QUIET deserves to be in the race as we close in on the answer to the question of the B-modes.

Part IV

Appendices

Appendix A

Physical technicalities

A.1 Polarization

Electromagnetic waves are made up of oscillating, mutually perpendicular electric and magnetic fields. We use the term *polarization* to describe the orientation of these fields. Polarization may be either linear, implying that the orientation of the fields is fixed in time, or circular, in which case the field vectors are rotating as the wave propagates.

All photons carry the property of polarization. However, in astrophysics and cosmology, we seldom encounter single photons. Rather, we look at the total radiation from some object, and find whether there is *net* polarization. For most sources, the orientation of each photon's polarization is random, thus giving no net polarization. But some types of objects, for instance ones with strong magnetic fields, have properties causing a degree of alignment of the photons. It is from such sources that we may observe polarized light.

To describe the properties of polarized radiation, we employ *Stokes' parameters*. With respect to fixed, orthonormal coordinate systems (x, y) or (a, b), where the latter is rotated 45° with respect to the former, these are

$$\begin{aligned} I &= |E_x|^2 + |E_y|^2 = |E_a|^2 + |E_b|^2, \\ Q &= |E_x|^2 - |E_y|^2 = -2\,\mathrm{Re}(E_a E_b^*), \\ U &= 2\,\mathrm{Re}(E_x E_y^*) = |E_a|^2 - |E_b|^2, \\ V &= 2\,\mathrm{Im}(E_x E_y^*) = 2\,\mathrm{Im}(E_a E_b^*), \end{aligned}$$

where the E_i are the components of the electric field vector (we do not describe the magnetic field vector, since it is easily found when the electric field is known). I represents the total intensity of the beam. Q gives the degree of linear polarization aligned with the coordinate system (x, y). U is equivalent to Q, but with reference to (a, b). The close relation between Q and U is seen in the way they change roles when one chooses the rotated coordinate system. Finally, V gives the degree of circular polarization, which may be right- or left-handed (with reference to the direction of rotation of the electrical field vector). (Interestingly, if the reference coordinate system had also been rotating, V would be found the same way as Q or U, as the difference between squared amplitudes in left- and right-handed rotating systems.) From these, we may

define a quantity I_P, giving the "total intensity of polarization",

$$Q^2 + U^2 + V^2 = I_P^2 \leq I^2.$$

This makes $\frac{I_P}{I}$ the degree of polarization of the beam. For cosmological purposes, V is seldom of interest, as we know of no cosmological sources of circularly polarized light.

The temperature of the background sky is a scalar field. Not so with polarization. For each point on the sky, we need more than one quantity to describe the polarization properly - it is a *tensor field*. Besides Q and U (and V if needed), we need to know the degree of polarization, and the orientation of the system of reference, since the measured values of Q and U will change if we rotate our detector. Given a rotation through an angle ϕ, the parameters transform as

$$\begin{aligned} Q' &= Q\cos(2\phi) + U\sin(2\phi), \\ U' &= -Q\sin(2\phi) + U\cos(2\phi). \end{aligned}$$

More compactly, this transformation, applied to a field where Q and U depend on the direction $\hat{n}$ in which we're looking, may be expressed as

$$(Q + iU)(\hat{n}) \rightarrow e^{\mp 2i\phi}(Q + iU)(\hat{n}).$$

This transformation allows us to recognise the Stokes' Q and U parameters as *spin* ± 2 *fields*.

Appendix B

Statistical technicalities

All CMB analysis is firmly rooted in statistics, due to the stochastic nature of the physical processes responsible. The normal and χ^2 - distributions are especially central in our work. In the following section are presented some of the basic properties of these functions, since they are used to a large extent in the thesis.

B.1 Basics

A *stochastic*, or random, variable is a quantity whose value is essentially non-deterministic, primarily describing the outcome of some process. To any random variable, be it continuous or discrete, we may assign a *probability distribution*, being a function defined on the space of possible outcomes for the variable, giving the probability for each outcome. The *mean* or *expected value* of a random variable X with probability distribution $f(x)$ is

$$\mu \equiv \langle X \rangle = \int_{-\infty}^{\infty} x f(x)\, dx.$$

By this definition, we easily find that expectation value is a linear operator, meaning that for random variables X, Y and constants a, b we have

$$\langle aX + bY \rangle = a\langle X \rangle + b\langle Y \rangle.$$

Also, if and only if X and Y are independent, $\langle XY \rangle = \langle X \rangle \langle Y \rangle$.

The *variance* of a random variable is a measure of the spread of the variable's probability distribution. It is defined as

$$\mathrm{Var}\, X = \langle (X - \mu)^2 \rangle.$$

The square root of the variance is called the standard deviation, σ.

The *covariance* of a pair of random variables X and Y is a measure of whether they seem to be varying dependently,

$$\mathrm{Cov}\,(X, Y) = \langle (X - \mu_X)(Y - \mu_Y) \rangle.$$

Correlation is a close relative of the covariance: The correlation coefficient of a pair of random variables is their covariance normalised by the standard deviation of each,

$$\rho_{X,Y} = \mathrm{corr}(X, Y) = \frac{\mathrm{Cov}\,(X, Y)}{\sigma_X \sigma_Y}.$$

This normalisation makes sure $\rho_{X,Y} \in [-1, 1]$. The endpoints of this interval signifies perfect positive or negative correlation (i.e. the relationship between the two variables is linear).

A *random vector* is a vector of which each component is a random variable. For a k-dimensional random vector $\mathbf{X}$ we define a mean vector $\boldsymbol{\mu}$,

$$\boldsymbol{\mu} = [\mu_1, \mu_2, \ldots, \mu_k].$$

The pairwise covariance of all the elements of $\mathbf{X}$ is contained within the corresponding *covariance matrix*,

$$\Sigma_{i,j} = \mathrm{Cov}(X_i, X_j) = \langle (X_i - \mu_i)(X_j - \mu_j) \rangle.$$

Equivalently we may give the whole matrix as

$$\boldsymbol{\Sigma} = \langle (\mathbf{X} - \boldsymbol{\mu}_\mathbf{X}) (\mathbf{X} - \boldsymbol{\mu}_\mathbf{X})^\mathsf{T} \rangle.$$

B.2 Samples and populations

When dealing with large data sets, we rarely get the opportunity to find the true mean and standard deviation. More often than not, it's either unfeasible or impossible to check every piece of the puzzle. Hence we introduce the concepts of sample and population: When looking for the average weight of, say, squirrels, we can hardly catch and weigh every squirrel on the planet. (If you're averaging over time as well, you're in even deeper trouble.) Rather, we try to find a group of squirrels - a *sample* - that are fairly representative of the total *population*. If our sample is *unbiased* and sufficiently large, the sample mean will be a good estimate of the true population mean. This way of thinking is central in practical statistics: We are frequently dealing with quantities that certainly exist, but are impossible to measure, so we will have to make do with estimates. The relevant terminology is introduced below.

In a sample of N observations, the sample mean $\bar{x}$ is the average observation, and its expectation value is the population mean,

$$\bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i \,, \qquad \langle \bar{x} \rangle = \mu.$$

The sample variance is similarly an estimate of the population variance,

$$s^2 = \frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{x})^2 \,, \qquad \langle s^2 \rangle = \sigma^2.$$

The sample mean is itself a random variable, with an associated probability distribution. (We may picture doing a series of identical experiments, each resulting in a sample mean, and taking the mean of those to estimate the true mean - the average weight of squirrels, for instance.) The standard deviation of this distribution is called the *standard error of the mean*, $\mathrm{se}_{\bar{x}}$. It is related to the sample standard deviation by

$$\mathrm{se}_{\bar{x}} = \frac{s}{\sqrt{n}},$$

where n is the number of experiments.

B.3 Important distributions

B.3.1 The normal distribution

The normal distribution $\mathcal{N}(\mu, \sigma)$ is a probability distribution of the form

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

This famous function is often referred to as a "Gaussian", and has mean μ and standard deviation σ. $\mathcal{N}(0,1)$ is referred to as the *standard normal distribution*. Standard normal variables are conventionally denoted Z, and are especially useful in calculations. Any Gaussian variable may be mapped into $\mathcal{N}(0,1)$ by normalising via the mean and standard deviation, $Z = \frac{X-\mu}{\sigma}$.

The normal distribution has the fortunate property that it remains normal under Fourier transform, $\mathcal{F}\left\{\exp\left(-\alpha x^2\right)\right\} = \sqrt{\frac{\pi}{\alpha}} \exp\left(-\frac{\pi\nu^2}{\alpha}\right)$.

A k-dimensional vector $\mathbf{x}$ where $x_i \sim \mathcal{N}(\mu_i, \sigma_i)$ has a *multivariate normal distribution*,

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{k/2}|\mathbf{\Sigma}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{\mathsf{T}}\mathbf{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right). \tag{B.1}$$

The normal distribution appears frequently in physical problems. The reason for this lie in the *Central Limit Theorem*. This theorem states that the distribution of the mean for a set of independent random variables will approach a normal distribution if the number of variables is large enough.

B.3.2 The chi-squared distribution

The *chi-squared distribution* with k degrees of freedom, or χ^2_k, is the distribution of a sum of the squares of k independent standard normal variables,

$$Q = \sum_{i=1}^{k} Z_i^2 \quad \Rightarrow \quad Q \sim \chi^2_k.$$

Its mean is

$$\langle Q \rangle = \frac{1}{k}\sum_{i=1}^{k} Q_i = \frac{1}{k}\sum_{i=1}^{k}\left(\sum_{i=1}^{k} Z_i^2\right) = k,$$

since the standard deviation of a standard normal variable is 0. Using this result, as well as the other properties of the normal distribution, we find that

$$\operatorname{Var} Q = \langle (Q - \langle Q \rangle)^2 \rangle = 2k.$$

The shape of the χ^2-distribution depends on the parameter k. Since it describes a sum of independent random variables with finite means and variances, it converges to a normal distribution for sufficiently high k. ("Sufficiently high" frequently amount to $k > 50$.) This is given by the central limit theorem. The definition also implies that the sum of any two or more χ^2-variables is itself distributed according to the χ^2-distribution, with degrees of freedom equal to the sum of the contributions.

A point worth noting is that the exponent of the multivariate normal distribution, when the elements of the normal vector in question are independent, is a chi-squared variable,

$$(\mathbf{x} - \boldsymbol{\mu})^{\intercal} \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \sim \chi^2 \,. \tag{B.2}$$

So for a normally distributed vector with k elements, say, one describing thermal noise, one would expect the product in eq. B.2 to have a value close to k, with a variance close to $2k$.

Appendix C

The signal covariance matrix

In chapter 5, we looked at the likelihood function relating the CMB power spectra to the map. The quantities in this function were all found directly from the data, with the exception of the signal covariance matrix. This quantity holds the key to linking cosmological theory to our observations, and is thus a vital component in our analysis. Its deduction is rather technical, and follows here. We begin by summarising the mathematical tools needed for the job.

C.1 Useful relations

The following, as well as more extensive treatments on the same material, may be found in any decent book on mathematical physics, such as Hassani [3].

Central to our calculations are the *spherical harmonics*, a set of functions defined on a spherical surface, which form the solutions to the angular part of the Laplace equation ($\nabla^2\phi = 0$) in spherical coordinates. They are

$$Y_{\ell m}(\theta,\phi) = (-1)^m \sqrt{\frac{2\ell+1}{4\pi}\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos\theta) e^{im\phi}, \tag{C.1}$$

where $P_\ell^m(\cos\theta)$ is the associated Legendre function. The spherical harmonics obey the orthonormality relation

$$\int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} Y_{\ell m}(\theta,\phi)\, Y^*_{\ell' m'}(\theta,\phi)\, d\Omega = \delta_{\ell\ell'}\,\delta_{mm'}. \tag{C.2}$$

The parameters ℓ and m are integers constrained by $-\ell \leq m \leq \ell$ and $\ell \in [0,\infty\rangle$, while δ_{ij} is the Kronecker delta function. From these properties, we see that the $Y_{\ell m}$'s form a complete set of orthonormal functions. Thus they may be used as a basis for the Hilbert space of square-integrable functions. This makes them very applicable in physics, since physical problems very often require functions to be square-integrable (i.e. $\int_{-\infty}^{\infty} |f(x)|^2\, dx < \infty$).

The spherical harmonics may be used to expand any function defined on the unit sphere. Referring to a point (θ,ϕ) on the sphere by its unit vector $\hat{n}$, defined as the vector that connects the origin and the point on the surface, the spherical harmonics

expansion of our CMB signal function takes the form

$$s(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{n}). \tag{C.3}$$

The $a_{\ell m}$'s are amplitudes, defined so that each denotes "how much" of a given $Y_{\ell m}$ is needed in the decomposition of the function.

Now, if we multiply equation C.3 with $Y^*_{\ell m}(\hat{n})$ and integrate over all directions, we may invoke the orthonormality relation of eq. C.2 and observe that every term of the right hand side vanishes exept the one where $\ell = \ell'$, $m = m'$. This yields the inverse transformation,

$$\begin{aligned} \int s(\hat{n})\, Y^*_{\ell m}(\hat{n})\, d\Omega &= \int \left(\sum_{\ell'=0}^{\infty} \sum_{m'=-\ell'}^{\ell'} a_{\ell' m'} Y_{\ell' m'}(\hat{n}) \right) Y^*_{\ell m}(\hat{n})\, d\Omega, \\ &= a_{\ell m} \delta_{\ell \ell'}\, \delta_{m m'}, \\ &= a_{\ell m}. \end{aligned} \tag{C.4}$$

Inserting this into eq. C.1 we find that the first harmonic is a real constant, $Y_{00} = Y^*_{00} = 1/\sqrt{4\pi}$ (found using $P^0_0 = P_0 = 1$). This we may in turn use on the orthonormality relation (eq. C.2) to find the integral of any $Y_{\ell m}$ or $Y^*_{\ell m}$,

$$\begin{aligned} &\int Y^*_{\ell m} Y_{00}\, d\Omega = \delta_{\ell 0}\delta_{m 0} \quad \text{using} \quad Y_{00} = 1/\sqrt{4\pi} \;\Rightarrow \\ &\int Y^*_{\ell m}\, d\Omega = \sqrt{4\pi}\, \delta_{\ell 0}\delta_{m 0}. \end{aligned} \tag{C.5}$$

The calculation is identical when solving for $\int Y_{\ell m} d\Omega$. One simply enters Y^*_{00} rather than Y_{00}. Hence we find that all spherical harmonics integrate to 0, exept for the $\ell = 0, m = 0$ mode, the integral of which is 4π.

To complete our toolkit, we provide some relations connecting the spherical harmonics to the Dirac delta function and the Legendre polynomials. We begin with the Addition Theorem for Spherical Harmonics,

$$P_\ell(\cos\theta) = \frac{4\pi}{2\ell + 1} \sum_{m=-\ell}^{\ell} Y^*_{\ell m}(\hat{n}_1) Y_{\ell m}(\hat{n}_2), \tag{C.6}$$

where $\cos\theta = \hat{n}_1 \cdot \hat{n}_2$ (and we note that we are now talking of plain Legendre polynomials, not the associated kind).

Furthermore, we have a useful normalisation integral for the Legendre polynomials,

$$\int_{-1}^{1} P_\ell(x) P_m(x)\, dx = \frac{2}{2\ell + 1} \delta_{\ell m}. \tag{C.7}$$

And finally, an expansion relating the Dirac delta function to the Legendre polynomials,

$$\delta(x - x') = \sum_{\ell=0}^{\infty} \frac{2\ell + 1}{2} P_\ell(x) P_\ell(x'). \tag{C.8}$$

Now we are armed and ready, and may move on to the problem at hand.

C.2 Deducing the signal covariance matrix

To recapitulate, we are looking for a formula for the elements of the signal covariance matrix **S**. Of these, we know a few things: Due to our assumption of an isotropic universe, the covariance of any two points of the same angular separation should be the same. The covariance matrix must also be symmetric, since elements $S_{i,j}$ and $S_{j,i}$ necessarily represent the same two pixels. Both these properties are computationally fortunate, since they help reduce the number of calculations needed. So in effect, we may freely swap $S_{i,j}$ for the equivalent quantity $C(\theta)$, also known (somewhat inaccurately, since it's technically a covariance) as the two-point correlation function. $C(\theta)$ yields the covariance between any two points of angular separation θ. The angle and the indices i, j are connected via $\hat{n}_i \cdot \hat{n}_j = \cos\theta$.

The two-point correlation function is defined as

$$C(\theta) = \frac{\int_{\Omega_1}\int_{\Omega_2} s(\hat{n}_1)s(\hat{n}_2)\,\delta(\hat{n}_1\cdot\hat{n}_2-\cos\theta)\,d\Omega_1 d\Omega_2}{\int_{\Omega_1}\int_{\Omega_2}\delta(\hat{n}_1\cdot\hat{n}_2-\cos\theta)\,d\Omega_1 d\Omega_2}. \tag{C.9}$$

Looking at this formula, we see that it represents the mean value of the product of the signal at all possible pairs of points with the given angular separation. We recognise this as the covariance we were after (remembering that the signal has mean 0 by definition). Now for some function juggling:

To simplify eq. C.9, we should like to do something about that ugly denominator (let's call it $\mathcal{D}$). We substitute in eq. C.8,

$$\mathcal{D} = \int_{\Omega_1}\int_{\Omega_2}\left(\sum_{\ell=0}^{\infty}\frac{2\ell+1}{2}P_\ell(\hat{n}_1\cdot\hat{n}_2)P_\ell(\cos\theta)\right)d\Omega_1 d\Omega_2.$$

The sum over ℓ and the integration over angles are independent operations, so we may swap the order in which they are applied. Also, we invoke eq. C.6 to replace the Legendre polynomial within the integral,

$$\mathcal{D} = \sum_{\ell=0}^{\infty}\frac{2\ell+1}{2}P_\ell(\cos\theta)\int_{\Omega_1}\int_{\Omega_2}\left(\frac{4\pi}{2\ell+1}\sum_{m=-\ell}^{\ell}Y^*_{\ell m}(\hat{n}_1)Y_{\ell m}(\hat{n}_2)\right)d\Omega_1 d\Omega_2.$$

The sum over m may also be moved outside the integral. The fractions partly cancel, and we are left with

$$\mathcal{D} = 2\pi\sum_{\ell=0}^{\infty}P_\ell(\cos\theta)\sum_{m=-\ell}^{\ell}\left(\int_{\Omega_1}Y^*_{\ell m}(\hat{n}_1)d\Omega_1\right)\left(\int_{\Omega_2}Y_{\ell m}(\hat{n}_2)d\Omega_2\right),$$

where we have used that $Y^*_{\ell m}(\hat{n}_1)$ may be considered a constant with respect to integration over Ω_2, and vice versa. Since we already know what integrating over a single

$Y_{\ell m}$ gives, we find

$$\begin{aligned}\mathcal{D} &= 2\pi \sum_{\ell=0}^{\infty} P_\ell(\cos\theta) \sum_{m=-\ell}^{\ell} (\sqrt{4\pi}\,\delta_{\ell 0}\delta_{m0})(\sqrt{4\pi}\,\delta_{\ell 0}\delta_{m0}),\\ &= 8\pi^2 \sum_{\ell=0}^{\infty} P_\ell(\cos\theta) \sum_{m=-\ell}^{\ell} \delta_{\ell 0}\delta_{m0},\\ &= 8\pi^2 P_0(\cos\theta),\\ &= 8\pi^2.\end{aligned}$$

How nice! Now we may turn back to eq. C.9,

$$C(\theta) = \frac{1}{8\pi^2}\int_{\Omega_1}\int_{\Omega_2} s(\hat{n}_1)s(\hat{n}_2)\,\delta(\hat{n}_1\cdot\hat{n}_2 - \cos\theta)\,d\Omega_1 d\Omega_2.$$

As with the denominator, we begin by substituting the delta function,

$$C(\theta) = \frac{1}{8\pi^2}\int_{\Omega_1}\int_{\Omega_2} s(\hat{n}_1)s(\hat{n}_2)\left(\sum_{\ell=0}^{\infty}\frac{2\ell+1}{2}P_\ell(\hat{n}_1\cdot\hat{n}_2)P_\ell(\cos\theta)\right)d\Omega_1 d\Omega_2.$$

After rearranging a bit, we invoke the Addition theorem, and split up the integrals using the same reasoning as last time:

$$\begin{aligned}C(\theta) &= \frac{1}{8\pi^2}\sum_{\ell=0}^{\infty}\frac{2\ell+1}{2}P_\ell(\cos\theta)\int_{\Omega_1}\int_{\Omega_2} s(\hat{n}_1)s(\hat{n}_2)\left(\frac{4\pi}{2\ell+1}\sum_{m=-\ell}^{\ell} Y^*_{\ell m}(\hat{n}_1)Y_{\ell m}(\hat{n}_2)\right)d\Omega_1 d\Omega_2,\\ &= \frac{1}{4\pi}\sum_{\ell=0}^{\infty} P_\ell(\cos\theta)\sum_{m=-\ell}^{\ell}\left(\int_{\Omega_1} s(\hat{n}_1)Y^*_{\ell m}(\hat{n}_1)d\Omega_1\right)\left(\int_{\Omega_2} s(\hat{n}_2)Y_{\ell m}(\hat{n}_2)d\Omega_2\right),\\ &= \frac{1}{4\pi}\sum_{\ell=0}^{\infty} P_\ell(\cos\theta)\sum_{m=-\ell}^{\ell} a_{\ell m}a^*_{\ell m},\end{aligned}$$

where the last equality follows from the inverse spherical harmonic transformation shown in eq. C.4. Now we only need to introduce our old friend the power spectrum,

$$C_\ell = \langle |a_{\ell m}|^2\rangle = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} a_{\ell m}a^*_{\ell m}, \tag{C.10}$$

which gives us our final result,

$$C(\theta) = \frac{1}{4\pi}\sum_{\ell=0}^{\infty}(2\ell+1)P_\ell(\cos\theta)C_\ell. \tag{C.11}$$

Note that this expression is slightly idealised, since it does not take into account the convolution of the signal that follows from observing the sky at finite resolution. The signal is convoluted twice, by the instrument beam and by pixelisation. In terms of the power spectrum, this translates to substituting $C_\ell \to b_\ell^2 w_\ell^2 C_\ell$, where b_ℓ is the beam and w_ℓ is the HEALPix pixel window function.

Bibliography

[1] Steven Weinberg. *Cosmology.* Oxford University Press, 2008.

[2] Scott Dodelson. *Modern Cosmology.* Academic Press, 2003.

[3] Sadri Hassani. *Mathematical Physics - A Modern Introduction to its foundations.* Springer, 1999.

[4] N. Jarosik et al. Seven-year *Wilkinson Microwave Anisotropy Probe* (WMAP) observations: Sky maps, systematic errors, and basic results. *The Astrophysical Journal Supplement Series*, 192(2), 2011.

[5] W. Hu and M. White. A CMB polarization primer. *New Astronomy*, 2(4):323–244, October 1997.

[6] M. Zaldarriaga and U. Seljak. All-sky analysis of polarization in the microwave background. *Physical Review D*, 55(4):1830–1840, February 1997.

[7] Malcolm Longair. *The Cosmic Century: A History of Astrophysics and Cosmology.* Cambridge University Press, 2006.

[8] E. P. Hubble. Cepheids in spiral nebulae. *The Observatory*, 48:139–142, 1925.

[9] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 49:769–822, 1916.

[10] A. Einstein. Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie. *Koniglich Preussische Akademie der Wissenschaften, Sitzungsberichte*, pages 142–152, 1917.

[11] A. A. Friedmann. Über die Krümmung des Raumes. *Zeitschrift für Physik*, 10(1):377–386, 1922.

[12] A. A. Friedmann. Über die Möglichkeit einer Welt mit konstanter negativer Krümmung des Raumes. *Zeitschrift für Physik*, 21(1):326–332, 1924.

[13] G. Lemaître. Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Societe Scientifique de Bruxelles*, 47:49–59, 1927.

[14] H. P. Robertson. Kinematics and World-Structure. *Astrophysical Journal*, 82:284–301, November 1935.

[15] A. G. Walker. On Milne's theory of world-structure. *Proceedings of the London Mathematical Society*, s2-42(1):90–127, 1937.

[16] E. P. Hubble. A Relation Between Distance and Radial Velocity Among Extra-Galactic Nebulae. *Proceedings of the National Academy of Sciences*, 15(3):168–173, 1929.

[17] F. Hoyle. A New Model for the Expanding Universe. *Monthly Notices of the Royal Astronomical Society*, 108:372–382, 1948.

[18] G. Lemaître. Contributions to a British Association Discussion on the Evolution of the Universe. *Nature*, 128:704–706, 1931.

[19] R. A. Alpher, H. Bethe, and G. Gamow. The Origin of Chemical Elements. *Physical Review*, 73(7):803–804, April 1948.

[20] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419–421, July 1965.

[21] R. A. Alpher and R. Herman. Evolution of the Universe. *Nature*, 162(4124):774–775, November 1948.

[22] A. G. Doroshkevich and I. D. Novikov. Mean Density of Radiation in the Metagalaxy and Certain Problems in Relativistic Cosmology. *Soviet Physics Doklady*, 9:111–113, 1964.

[23] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R = 4kpc) to UGC 2885 (R = 122 kpc). *Astrophysical Journal*, 238:471–487, 1980.

[24] A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D*, 23(2):347–356, 1981.

[25] I. D. Novikov. On the possibility of the appearance of large scale inhomogeneities in the expanding Universe. *Journal of Experimental and Theoretical Physics*, 46:686–689, 1964.

[26] R. K. Sachs and A. M. Wolfe. Perturbations of a cosmological model and angular variations of the microwave background. *Astrophysical Journal*, 147:73, 1967.

[27] P. de Bernardis et al. A short wavelength measurement of the Cosmic Background Radiation Anisotropy. *The Astrophysical Journal Letters*, 360:L31–L34, 1990.

[28] L. A. Page, E. S. Cheng, and S. S. Meyer. A large-scale Cosmic Microwave Background anisotropy measurement at millimeter and submillimeter wavelengths. *Astrophysical Journal*, 355(1):L1–L4, 1990.

[29] I. Strukov et al. Anisotropy of relic radiation in the RELICT-1 experiment and parameters of Grand Unification. *Physics Letters B*, 315:198–202, 1993.

[30] J. C. Mather et al. A preliminary measurement of the cosmic microwave background spectrum by the Cosmic Background Explorer (COBE) satellite. *Astrophysical Journal, Part 2 - Letters*, 354:L37–L40, 1990.

[31] G. F. Smoot et al. Structure in the COBE differential microwave radiometer first-year maps. *Astrophysical Journal, Part 2 - Letters*, 396(1):L1–L5, 1992.

[32] C. L. Bennett et al. Four-year COBE DMR cosmic microwave background observations: Maps and basic results. *Astrophysical Journal Letters*, 464:L1–L4, 1996.

[33] S. R. Platt et al. Anisotropy in the Microwave Sky at 90 GHz: Results from Python III. *Astrophysical Journal Letters*, 475:L1, 1997.

[34] E. S. Cheng et al. Detection of Cosmic Microwave Background Anisotropy by the Third Flight of MSAM. *ArXiv Astrophysics e-prints*, May 1997.

[35] S. Perlmutter et al. Measurements of Ω and Λ from 42 high-redshift supernovae. *Astrophysical Journal*, 517(2):565–586, June 1999.

[36] A. G. Riess et al. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astronomical Journal*, 116(3):1009, September 1998.

[37] P. de Bernardis et al. A Flat Universe from High-Resolution Maps of the Cosmic Microwave Background Radiation. *Nature*, 404:955, 2000.

[38] S. Hanany et al. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of $10' - 5^\circ$. *Astrophysical Journal*, 545(1):L5–L9, December 2000.

[39] J. A. Peacock. Large-Scale Surveys and Cosmic Structure. ArXiv Astrophysics e-prints: astro-ph/0309240, September 2003.

[40] D. N. Spergel et al. First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. *Astrophysical Journal Supplement Series*, 148(1):175–194, September 2003.

[41] J. M. Kovac et al. Detection of Polarization in the Cosmic Microwave Background using DASI. *Nature*, 420(6917):772–787, December 2002.

[42] D. Barkats et al. First Measurements of the Polarization of the Cosmic Microwave Background Radiation at Small Angular Scales from CAPMAP. *Astrophysical Journal Letters*, 619(2):L127–L130, February 2005.

[43] J. L. Sievers et al. Implications of the Cosmic Background Imager Polarization Data. *Astrophysical Journal*, 660(2):976–987, May 2007.

[44] P. Ade et al. First Season QUaD CMB Temperature and Polarization Power Spectra. *Astrophysical Journal*, 674(1):22–28, February 2008.

[45] T. E. Montroy et al. A Measurement of the CMB $\langle EE \rangle$ Spectrum from the 2003 Flight of BOOMERANG. *Astrophysical Journal*, 647(2):813–822, August 2006.

[46] J. H. P. Wu et al. MAXIPOL: Data analysis and results. *Astrophysical Journal*, 665(1):55–66, August 2007.

[47] D. Larson et al. Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Power Spectra and WMAP-Derived Parameters. *Astrophysical Journal Supplement Series*, 192(2), February 2011.

[48] H. C. Chiang et al. Measurement of Cosmic Microwave Background Polarization Spectra from Two Years of BICEP Data. *Astrophysical Journal*, 711(2):1123–1140, March 2010.

[49] E. Komatsu et al. Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *Astrophysical Journal Supplement Series*, 192(2), 2011.

[50] W. de Sitter. On Einstein's theory of gravitation and its astronomical consequences. *Monthly Notices of the Royal Astronomical Society*, 78:3–28, November 1917.

[51] Øyvind Grøn and Sigbjørn Hervik. *Einstein's General Theory of Relativity With Modern Applications in Cosmology*. Springer, 2007.

[52] D. Baumann et al. CMBPol Mission Concept Study: Probing Inflation with CMB Polarization. *AIP Conferece Proceedings*, 1141:10–120, June 2009.

[53] A. D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393, February 1982.

[54] Edmund Bertschinger. Cosmological Perturbation Theory and Structure Formation. ArXiv Astrophysics e-prints: astro-ph/0101009, December 2001.

[55] S. M. Leach and A. R. Liddle. Microwave Background Constraints on Inflationary Parameters. *Monthly Notices of the Royal Astronomical Society*, 341(4):1151–1156, June 2003.

[56] David H. Lyth. What would we learn by detecting a gravitational wave signal in the Cosmic Microwave Background Anisotropy? *Physical Review Letters*, 78(10):1861–1863, March 1997.

[57] L. Knox and Y. S. Song. Limit on the Detectability of the Energy Scale of Inflation. *Physical Review Letters*, 89(1), July 2002.

[58] The QUIET Collaboration. First season QUIET observations: Measurements of CMB polarization power spectra at 43 GHz in the multipole range $25 \leq l \leq 475$. *Astrophysical Journal*, 741:111–128, 2011.

[59] The QUIET Collaboration. QUIET-II proposal for funding by the NSF Major Research Instrumentation (MRI) Program. Submitted to NSF in January 2012.

[60] S. Padin et al. The Cosmic Background Imager. *Publications of the Astronomical Society of the Pacific*, 114:83–97, January 2002.

[61] Kieran Cleary. Coherent polarimeter modules for the QUIET experiment. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77412H, 2010.

[62] KICP Research Highlights. First QUIET CMB Polarization Results from the Atacama Desert. `http://kicp.uchicago.edu/research/highlights/highlight_2010-12-15.html`. [Online; accessed 07-May-2012].

[63] Immanuel Buder. Q/U Imaging Experiment (QUIET): A Ground-Based Probe of Cosmic Microwave Background Polarization. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77411D, 2010.

[64] Laura Newburgh. Measuring CMB polarization with QUIET: The Q/U Imaging ExperimenT. In *On Recent Developments in Theoretical and Experimental General Relativity, Astrophysics and Relativistic Field Theories*, The Twelfth Marcel Grossman Meeting on General Relativity, page 2145, 2010.

[65] Lloyd Knox. Determination of inflationary observables by cosmic microwave background anisotropy experiments. *Physical Review D*, 52(8):4307–4318, October 1995.

[66] Laura Newburgh. *The Q/U Imaging ExperimenT (QUIET): The Q-band Reciever Array Instrument and Observations*. PhD thesis, Columbia University, 2010.

[67] W. N. Brandt et al. Separation of Foreground Radiation from Cosmic Microwave Background Anisotropy using Multifrequency Measurements. *Astrophysical Journal*, 424(1):1–21, March 1994.

[68] Yuji Chinone. *Measurement of Cosmic Microwave Background Polarization Power Spectra at 43 GHz with Q/U Imaging ExperimenT*. PhD thesis, Tohoku University, 2011.

[69] I. K. Wehus, S. K. Næss, and H. K. Eriksen. Bayesian Noise Estimation for Non-Ideal Cosmic Microwave Background Experiments. *Astrophysical Journal Supplement*, 199(1):15, March 2012.

[70] Colin A. Bischoff. *Observing the Cosmic Microwave Background Polarization Anisotropy at 40 GHz with QUIET*. PhD thesis, The University of Chicago, 2010.

[71] J. Aumont et al. Measurement of the Crab Nebula polarization at 90 GHz as a calibrator for CMB experiments. *Astronomy and Astrophysics*, 514:A70, December 2010.

[72] Max Tegmark. How to make maps from Cosmic Microwave Background data without losing information. *Astrophysical Journal Letters*, 480(2):L87–L90, May 1997.

[73] K. M. Gorski et al. The HEALPix Primer. ArXiv Astrophysics e-prints: astro-ph/9905275, May 1999.

[74] D. Sutton et al. Fast and precise map-making for massively multi-detector CMB experiments. *Monthly Notices of the Royal Astronomical Society*, 407(3):1387–1402, September 2010.

[75] T. Poutanen et al. Comparison of Map-Making Algorithms for CMB Experiments. *Astronomy and Astrophysics*, 449(3):1311–1322, April 2006.

[76] R. Keskitalo et al. Residual noise covariance for *Planck* low-resolution data analysis. *Astronomy and Astrophysics*, 522:A94, November 2010.

[77] H. K. Eriksen et al. Joint Bayesian Component Separation and CMB Power Spectrum Separation. *Astrophysical Journal*, 676(1):10–32, March 2008.

[78] J. R. Bond, A. H. Jaffe, and L. Knox. Estimating the Power Spectrum of the Cosmic Microwave Background. *Physical Review D*, 57(4):2117–2137, February 1998.

[79] NASA. Legacy Archive for Microwave Background Data Analysis (LAMBDA). `http://lambda.gsfc.nasa.gov/`. [Online; accessed 14-may-2012].

[80] The Planck Sky Model Project. `http://www.apc.univ-paris7.fr/~delabrou/PSM/psm.html`. [Online; accessed 17-may-2012].

[81] P. A. R. Ade et al. Planck Early Results XVII. Origin of the submillimetre excess dust emission in the Magellanic clouds. *Astronomy and Astrophysics*, 536:A17, December 2011.

[82] R. W. Ogburn IV et al. The BICEP2 CMB polarization experiment. In *Millimeter, Submillimeter and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77411G, 2010.

[83] P. Oxley et al. The EBEX experiment. In *Infrared Spaceborne Remote Sensing XII*, volume 5543 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, November 2004.