Morphological Studies of the CMB: Non-standard Models and Foregrounds

Theresa R. Jaffe

München 2006

Morphological Studies of the CMB: Non-standard Models and Foregrounds

Theresa R. Jaffe

Dissertation
an der Fakultät für Physik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Theresa R. Jaffe
aus Honolulu, USA

München, den 31. August 2006

Erstgutachter: Prof. Dr. Simon D. M. White

Zweitgutachter: PD Dr. Peter Schücker

Tag der mündlichen Prüfung: 9. Oktober 2006

Contents

List of Figures

List of Tables

Zusammenfassung

Neue Messungen der kosmischen Mikrowellenhintergrundstrahlung (CMB) haben die bisher genaueste Bestimmung der Parameter des Standard-ΛCDM-Modells ermöglicht, aber die Daten enthalten auch verblüffende Anomalien, die mit den Annahmen eines statistisch isotropen und gaußschen Universums nicht konsistent sind. In dieser Arbeit befassen wir uns mit den möglichen Ursachen solcher Anomalien, indem wir die Morphologie der Hintergrundstrahlung untersuchen. Wir finden eine unerwartete Korrelation zwischen den CMB-Anisotropien und dem räumlichen Muster der Temperaturverteilung, wie es in einem Bianchi-Universum des Typs VII$_h$ erzeugt wird, d.h. in einem anisotropen Universum, welches eine universelle Drehung bzw. Vortizität zulässt. Wie wir feststellen ist dieses Modell inkompatibel mit anderen Beobachtungen, jedoch hat die Korrektur der Hintergrundstrahlung um eine solche Komponente den überraschenden Effekt, dass sie viele der Anomalien der von WMAP gemessenen Temperaturverteilung beseitigt. Dieses Ergebnis deutet darauf hin, dass möglicherweise ein alternatives kosmologisches Modell benötigt wird, welches eine solche Morphologie erzeugt. Eine ähnliche Kreuzkorrelationsmethode wenden wir auf die Vordergrundemission im Mikrowellenbereich an, um das spektrale Verhalten der galaktischen Emissionsprozesse zu untersuchen. Die Ergebnisse geben Aufschluss über die überraschend geringe Intensität der Bremsstrahlung, sowie über die Natur der anomalen, mit Staub korrelierten Emission, die bei niedrigen Frequenzen dominiert. In Situationen, in welchen *a priori* keine Information über die räumliche Struktur vorhanden ist, die bei der Suche nach nicht-gaußschen Signalen Anhaltspunkte liefern kann, kommt als komplementäre Methode die Phasenwinkelstatistik zur Anwendung. Wir wenden diese Methode sowohl auf kompakte topologische Modelle als auch auf Überreste der Vordergrundemission an, und eine vorläufige Analyse zeigt, dass sie ein leistungsfähiges Instument für die Untersuchung von nicht-gaußschen Strukturen und Anisotropie sein könnte.

Abstract

Recent measurements of the Cosmic Microwave Background (CMB) have allowed the most accurate determinations yet of the parameters of the standard ΛCDM model, but the data also contain intriguing anomalies that are inconsistent with the assumptions of statistical isotropy and Gaussianity. This work investigates possible sources of such anomalies by studying the morphology of the CMB. An unexpected correlation is found between the CMB anisotropies and a temperature pattern generated in a Bianchi Type VII_h universe, i.e., an anisotropic universe allowing a universal rotation or vorticity. This model is found to be incompatible with other observations of the cosmological parameters, but correcting for such a component can serendipitously remove many of the anomalies from the *WMAP* sky. This result indicates that an alternative cosmological model producing such a morphology may be needed. A similar cross-correlation method applied to the microwave foregrounds studies the variation of the spectral behaviours of the Galactic emission processes across the sky. The results shed light on the unexpectedly low free-free emission amplitude as well as the nature of the anomalous dust-correlated emission that dominates at low frequencies. As a complementary method, phase statistics apply to situations where no *a priori* knowledge of the spatial structure informs the search for a non-Gaussian signal. Such statistics are applied to compact topological models as well as to foreground residuals, and a preliminary analysis shows that these may prove powerful tools in the study of non-Gaussianity and anisotropy.

Chapter 1

Introduction

The cosmic microwave background (CMB) is one of the most important observational pillars on which modern cosmology rests. It was almost discovered in 1940 through the detection of rotational excitation in CN molecules in the ISM (work by McKellar reported by Adams 1941), but the significance of the measured temperature was not recognised, as relic radiation from a hot big bang was only later predicted by Gamow (1946). Penzias & Wilson (1965) detected it by accident, just beating Dicke et al. (1965), who were in the process of building a instrument expressly for that purpose. The measurements of the background temperature and of CMB fluctuations resoundingly confirmed the big bang scenario and partly motivated the next idea to revolutionise cosmology, namely inflation theory.

The CMB originates at the most distant surface that is currently directly observable and provides one of the most important tools for studying both the contents and shape of the universe. The advent of $WMAP$[1] (Bennett et al., 2003a) brings us into the era of "precision cosmology", where we can use the CMB in conjunction with other datasets to constrain cosmological models and their parameters with extraordinary precision. The constraints from the CMB as well as from other observations of large scale structure, supernova distances, etc., have all agreed remarkably well on a "concordance cosmology" that describes the universe as accelerating, flat, composed mostly of dark energy ($\sim 73\%$) and dark matter ($\sim 23\%$), with only $\sim 4\%$ baryonic matter. Despite this concordance, however, there remain unexpected features in the microwave sky that may indicate problems with the standard model of cosmology.

The CMB contains information not only about the very early universe but also about the very local environment. Except in a narrow window in frequency space, the cosmic component is dominated by foreground emission from the Galaxy. From one point of view, this is an annoyance to be removed from the data in order to study the cosmic anisotropies. But it is also a rich source of information about the emission processes in the ISM and the structure of the Galaxy.

This introduction is intended to provide the necessary background for the non-specialist.

---

[1] Wilkinson Microwave Anisotropy Probe

How the CMB informs the standard model of cosmology is briefly described and some of the challenges for interpreting the observations. Of particular interest are anomalies in the data that violate the assumptions of statistical isotropy and Gaussianity. The following chapters will then describe morphological studies of the CMB that not only investigate non-standard models but also apply to the more prosaic but no less interesting problem of Galactic foregrounds.

This thesis represents work done primarily in collaboration with A. J. Banday, K. M. Górski, H. K. Eriksen, F. K. Hansen, S. Hervik, R. D. Davies, C. Dickinson, and R. J. Davis. Chapters 2, 3 and 5, and § 7.3 represent work published with these these coauthors. Chapter 5 in particular was largely written by Davies and Dickinson, though the quantitative results presented are my work. In each case, a note at the beginning explains specific contributions by these collaborators. Otherwise, the work is my own, though the first person plural will be used throughout.

1.1 The CMB and the Physics of the Early Universe

There are several comprehensive reviews of CMB physics including White et al. (1994), Hu & Dodelson (2002), Hu & White (1997), web resources such as E. Wright's[2] and W. Hu's[3], and textbooks such as Partridge (1995), Kolb & Turner (1990), Liddle & Lyth (2000), Peacock (1999), and Dodelson (2003). Here, I will give only a brief overview and refer the reader to these resources and the citations for further details.

1.1.1 The Expanding Universe

In the history of cosmology, one of the most fundamental assumptions has long been the Copernican Principle that there are no privileged observers. This implies that the universe ought to be the same everywhere, or homogeneous[4], and that it should look the same in every direction, i.e., isotropic[5]. Observations have generally confirmed these properties, but it is important to note that some such common-sense ideas have been proved wrong in the past, including most famously the idea that the universe ought to be eternal and unchanging. (In fact, there are worrying indications that we exist at a somewhat special *time* in the evolution of the universe.)

Einstein's formulation of general relativity (see Appendix A.1) showed that a static universe is not possible without a fudge factor, a vacuum energy known as the the cosmological constant. Observations of the Hubble expansion lead to the idea of a dynamic

---

[2]http://www.astro.ucla.edu/~wright/cosmolog.htm

[3]http://background.uchicago.edu/

[4]Misner et al. (1973): "Homogeneity of the universe means, then, that through each event in the universe there passes a space-like 'hypersurface of homogeneity' (physical conditions identical at every event on this hypersurface)."

[5]Misner et al. (1973): "Isotropy of the universe means that, at any event, an observer who is 'moving with the cosmological fluid' cannot distinguish one of his space directions from the others by any local physical measurement."

universe evolving from a hot and dense beginning, or a big bang, as an alternative to the steady-state universe of, e.g., Bondi & Gold (1948). It was only the first observations of the CMB, the relic radiation from this hot birth, that put to rest the idea of the steady-state theory for most cosmologists, and thereafter, the standard model of cosmology has been the big bang scenario. That the cosmological constant has come back to haunt us in the unexpected detection of an accelerated expansion is an indication that there are further surprises in store.

1.1.2 Primordial Perturbations and Inflation

The Problem

Even after the big bang scenario became widely accepted, cosmologists were left with several observational puzzles to explain.

Given that the Universe is not static but expanding from an initial big bang that occurred roughly 14 Gyrs in the past, it is clear that the most distant objects we observe have only relatively recently entered our horizon (see Appendix A.2). Distant regions of the universe cannot have been causally connected in the past, and yet they look qualitatively the same. The CMB in particular has the same temperature in every direction to one part in 10^5 (when corrected for non-cosmological components). This is known as the horizon problem.

The second mystery was the fact that, although the universe is very homogeneous, it is not perfectly so. We observe structure at all scales from stars to galaxy clusters to the cosmic "foam" (voids) or "web" (filaments). Recent observations of the CMB have confirmed corresponding fluctuations in the blackbody temperature. Though the physics of how gravitational collapse forms such objects given an initial density perturbation are relatively well understood, the source of the initial seed was unknown.

Kinematic observations had shown that the geometry of the universe is close to flat. Among the homogeneous solutions to Einstein's equations describing the relation of spatial curvature and density, the flat Friedmann-Robertson-Walker (FRW) metric is a special case (see Appendix A.1). Any small deviation from flatness will grow as expansion proceeds, making the current near-flatness after 14 Gyrs of expansion extremely unlikely. This apparent fine-tuning of the contents of the universe to be at exactly the critical value is known as the flatness problem.

Inflation

Inflation is an idea proposed by Guth (1981) that provides a natural mechanism to explain these observations. In the simplest terms, inflation refers to a period of exponential space-time expansion due to a negative pressure term. This can arise during the slow roll of an "inflaton" field down a nearly flat potential (see Appendix A.3), and the details of this field and potential vary from theory to theory. During the expansion, quantum fluctuations are generated in the inflaton field, which to first order acts like a simple harmonic

oscillator bounded by the horizon. That horizon stays roughly constant during the very brief inflation period, but the perturbations are then expanded out of that horizon. The result is a universe seeded with scale-invariant perturbations that eventually collapse into the observed structures. Regions that were causally connected in the first 10^{-36} s (before inflation) can *now* be larger than the current horizon size due to the effectively superluminal expansion during inflation. The expansion then also effectively erases any observable curvature or anisotropy in the geometry of space-time.

The inflationary mechanism for generating primordial perturbations generally implies that they are adiabatic[6], nearly scale invariant, isotropic, and generally follow a Gaussian distribution. Inflation, it must be noted, currently lacks a physical foundation, but the mechanism has proved extremely useful for interpreting the increasingly precise measurements of both the CMB as well as the matter distribution and expansion history. Before inflation, however, Harrison, Zel'dovich, and Peebles predicted that a scale invariant spectrum of fluctuations was the natural source for structure formation. If the universe is self-similar on all scales, this requires perturbations of the scale of the horizon to always have the same amplitude, and this leads to a scale-invariant primordial spectrum (see Appendix A.5). Thus the measurement of such a spectrum does not by itself prove inflation.

See the abovementioned references for more details.

1.1.3 Physics and the Power Spectrum

COBE-DMR[7] first confirmed that, once the non-cosmological components have been removed (see 1.2), the CMB does indeed have fluctuations of the order 10^{-5}. The challenge is then to extract the details of the physics in the early universe from the nature of those fluctuations. To study a stochastic process, we use a statistical representation of the data, namely the power spectrum.

Definitions

In the harmonic representation, the microwave sky temperature, T, in the direction $\hat{\mathbf{n}}$, is represented as the sum of spherical harmonics:

$$T(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \tag{1.1}$$

where

$$a_{\ell m} = \int d\Omega Y_{\ell m}^{*}(\hat{\mathbf{n}}) \Delta T(\hat{\mathbf{n}}). \tag{1.2}$$

---

[6]I.e. isoentropy, rather than isocurvature; see Appendix A.4. Observations have shown that the perturbations are at least mostly adiabatic, and the discussion will assume such. Though the simplest models predict only adiabatic perturbations, isocurvature perturbations are generated by, for example, two-field inflation models or axions, in which case the fluctuations are also non-Gaussian (see § 1.3).

[7]Cosmic Background Explorer, Differential Microwave Radiometer
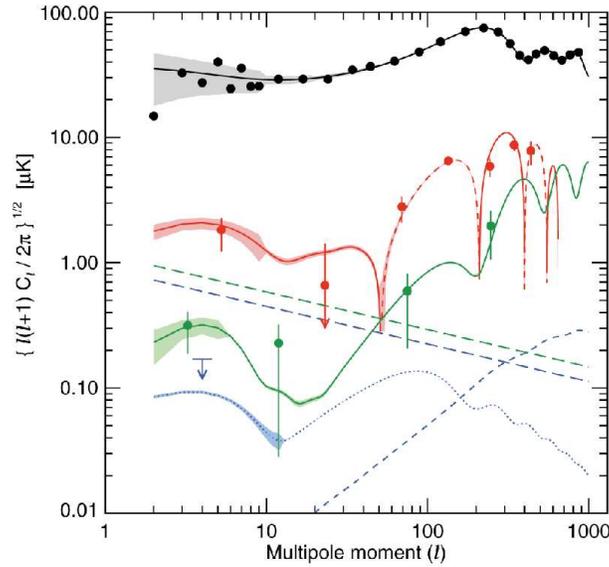
Figure 1.1 Examples of auto- and cross-correlation power spectrum from *WMAP*. Black is the TT, red the TE, green the EE, and blue the BB. Cosmic variance errors are shown as shaded bands. The straight dashed lines give the power-law foreground models, and the rising dashed line the BB lensing signal. Courtesy the *WMAP* Science Team.

The power spectrum characterises the distribution of the coefficients $a_{\ell m}$, i.e.,

$$C_\ell \equiv \langle a_{\ell m}^* a_{\ell m} \rangle \tag{1.3}$$

where the angle brackets represent the ensemble average. Under certain assumptions discussed in § 1.3, it is this distribution that encodes the physics of the early universe. (It is related to the two-point angular correlation function; see § 1.3.1.) Note that what is usually plotted is $\Delta T^2 \equiv \ell(\ell+1)C_\ell/2\pi$.

The above refers to the temperature of the CMB, and a similar formalism is used to study the linear polarisation signal that has the added complexity of an orientation associated with each line of sight in addition to an amplitude. Using the Stokes parameters I, Q, and U, we can decompose the polarisation signal into scalar "electric" (zero curl) E-mode and pseudo-scalar "magnetic" (zero divergence) B-mode polarisation components, with the corresponding power spectra C_ℓ^{EE} and C_ℓ^{BB} and cross-power spectra C_ℓ^{TE}, etc. To express the polarisation with these two components neatly separates the different physical sources of linear polarisation as discussed below.

Examples from *WMAP* are shown in Fig. 1.1, while Fig. 1.2 (a) illustrates schematically the main processes that shape the CMB temperature power spectrum and are summarised in the next section.

Sachs-Wolfe Effect

In green in Fig. 1.2 (a) is the Sachs-Wolfe contribution, the primary effect that turns initial density perturbations into temperature perturbations. If we define the space-space part of the metric fluctuation as $-\Phi \propto \delta\rho$ (see equation A.2), then photons originating from within an overdense region at the last scattering surface are redshifted by Φ from climbing out of the potential well to reach the observer, and likewise blueshifted from an underdense region.

If the perturbations are adiabatic, then such photons are also hotter due to the time-time perturbation, or Newtonian potential, $\Psi \sim -\Phi$, which cancels $1/3$ of the gravitational redshift in the matter dominated case. Likewise, photons from underdense regions are blueshifted. The general expression for the net redshift is

$$\frac{\Delta T}{T} = -\frac{\delta a}{a} = -\frac{2}{3(1 + p/\rho)}\Psi \qquad (1.4)$$

(By contrast, if the initial perturbations are isocurvature, the temperature fluctuation from the time-time perturbation changes sign and adds to the redshift from an overdensity.)

At the largest angular scales (smaller values of ℓ) corresponding to the largest physical scales, fluctuations in the potential remain larger than the horizon up until recombination. The potentials have not been subject to any causal physical processes since imprinted, and these CMB anisotropies therefore reflect the primordial fluctuations. The low-ℓ part of that green curve looks flat, reflecting the nearly scale-invariant nature of the primordial perturbation spectrum. This is known as the Sachs-Wolfe plateau.

Spectral Slope

If the primordial fluctuation spectrum follows a power law, then at large angular scales (via the Sachs-Wolfe effect above and neglecting other effects discussed below) for wave number k, we have a temperature power spectrum $P(k) \equiv <|\delta_k|^2> \propto A_s k^n$, where A_s is the amplitude of scalar perturbations (as distinct from a possible tensor component), and $n \sim 1$ the scalar index. In this case, transforming to the angular power spectrum gives

$$C_\ell \propto A_s \frac{\Gamma^{\frac{(9-n)}{2}}\Gamma(\ell + \frac{n-1}{2})}{\Gamma(\frac{3+n}{2})\Gamma(\ell + \frac{5-n}{2})}, \qquad (1.5)$$

where for $n = 1$, we then have $C_\ell^{-1} \propto \ell(\ell + 1)$, or $\Delta T^2 = $ constant, i.e., a flat spectrum (see Bond & Efstathiou 1987).

This primordial spectrum is then modified by the physical processes that occur between inflation and the observer. To determine the transfer function requires solving the Boltzmann equation for the photon-baryon plasma from inflation to recombination and modelling the secondary anisotropies introduced during the photon's propagation toward the observer through an evolving Universe. Some of these effects will be described below. The large-scale spectrum, the Sachs-Wolfe plateau, tells us that the primordial spectrum
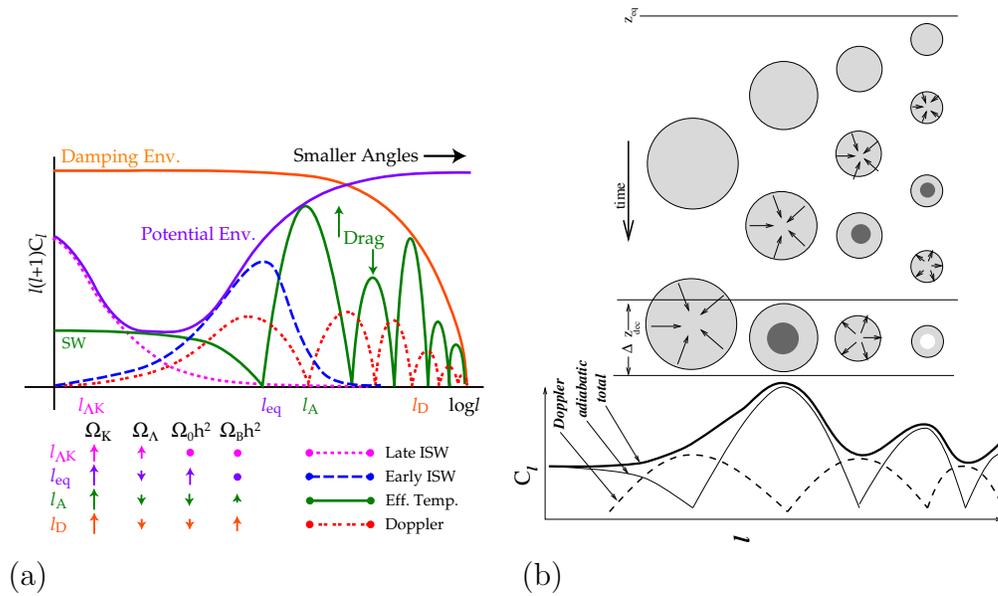
Figure 1.2 (a) Contributions to the power spectrum from various physical processes Hu et al. (1997). (b) Cartoon illustrating the acoustic peak structure from Lineweaver (2003).

was nearly scale invariant, but the processes that modify the spectrum on the small scales must be well understood in order to constrain this slope tightly enough to study inflation itself. Much of the excitement surrounding the 3-year *WMAP* release relates to the claimed detection of $n < 1$ at over 2σ significance (Spergel et al. 2006) and its implications for inflation.

Diffusion Damping

Because the coupling between the photons and baryons is imperfect, photon diffusion will erase some of the small scale perturbations. This effect is known as diffusion or Silk damping. The damping is exponential, as shown by the solid orange curve in Fig. 1.2 (a) and depends on the optical depth, τ, via a visibility function, $\dot{\tau}e^{-\tau}e^{-(k/k_D)^2}$, that is also exponential in the wave number. (The factor k_D gives the characteristic damping scale.)

Acoustic Peaks

Smaller fluctuations enter the horizon before recombination and begin to evolve based on the interaction between the gravitational attraction tending to collapse overdense regions and the photon pressure resisting that compression. The result is acoustic oscillations in the photon-baryon plasma, and these are reflected as peaks and troughs in CMB power. (Note that this depends on all of the wave modes being in phase, which inflation produces naturally.) Oscillations that reach their extrema of compression in dark matter potential wells when recombination freezes in the CMB anisotropies will lead to peaks in the CMB power at the acoustic peak scale, $\ell_A = \pi D/s_*$, depending on the angular diameter distance

to the last scattering surface, D, and the sound horizon at recombination, s_*. Slightly smaller scales (larger ℓ) corresponds to oscillations that have passed such a maximum and are instead caught between compression and rarefaction, leading to Sachs-Wolfe minima, and so on with harmonics at each $\ell \sim n\ell_A$.

The position of the acoustic peaks represents a well understood physical size, the sound horizon at last scattering. It therefore provides an accurate probe of the angular diameter distance to that surface, and thereby measures the curvature of space in between. In a positively curved universe, geodesics converge, and therefore a given length scale will appear larger on the sky, while in a negatively curved universe, geodesics diverge, so a given length will appear smaller. This means that positive curvature shifts the first acoustic peaks to lower ℓ, and vice versa.

Baryon Drag

Baryon drag imposes an asymmetry between modes that reach maximal compression within a potential well (dominated by non-interacting dark matter) and modes that reach maximal rarefaction within such wells, leading to a difference in the power spectrum peak height. The series of peaks and troughs in the CMB therefore reflects the sound speed of the plasma, the horizon size at last scattering, and the baryon fraction. Intuitively, this can be though of as a spring with a weight on one end. Instead of oscillating around its natural minimum, the weight shifts the extrema, where in this case, the maximal compression peaks are amplified and the rarefaction peaks lowered.

Doppler

During the fluid oscillations that generate the acoustic peak structure described above, bulk fluid flows give rise to a Doppler component. The fluid velocity is maximum for perturbations exactly between maximal compression and expansion, so this term is out of phase with the SW component of the acoustic peaks. This is shown by the orange curve of Fig. 1.2 (a) and by the dashed grey curve in (b).

Polarisation

Primordial linear polarisation arises from Thompson scattering of free electrons in a local quadrupole as shown in Fig. 1.3 (a). Before recombination, the optical depth of the plasma is high, so the local quadrupole is weak, and on small scales, diffusion erases it entirely. The primordial linear polarisation power is therefore at most $\sim 1/10$ that of the temperature perturbations.

As mentioned above, the measured I, Q, and U of linear polarisation are decomposed in the E-mode and B-mode polarisation. (Thompson scattering only generates linear polarisation, so the Stokes V parameter describing circular polarisation is not relevant.) One way to express the difference is that in E-modes, the polarisation vector is parallel or perpendicular to the direction of greatest change, while for B-modes, it is at $45°$.

The local $m = 0$ quadrupolar arising from scalar perturbations is relatively easy to envision, as shown in Fig. 1.3 (b) and (c). The plane wave of a scalar perturbation gives rise to fluid flows and thus the local quadrupole, which lines up the resulting polarisation axis parallel or perpendicular to the plane of the wave. In the E- and B-mode decomposition, the result is purely E-mode. The vector and tensor cases are harder to envision. (Vorticity results in velocity differences that give rise to $m = 1$ quadrupolar components, while a passing gravity wave stretches or compresses the photons, creating a local $m = 2$ quadrupole. See, e.g., Hu & White 1997.) In these cases, the polarisation is not simply orientation parallel to or perpendicular to the plane of the perturbation but instead a mix of E- and B-modes.

Since the primordial perturbations are primarily scalar perturbations, the resulting polarisation signal is primarily an E-mode signal. The B-mode signal is much weaker, coming from tensor perturbations, or gravity waves, at early times. (Vorticity is not expected to be relevant, as only scalar density perturbations and gravity waves arise naturally from inflation.) These are shown for the best-fit *WMAP* model in Fig. 1.1.

Figure 1.3 (d) shows how the temperature cross-correlates with the polarisation E-modes and does not correlate with the B-modes. This is why no such C^{TB} spectrum is shown in Fig. 1.1. Instead, the cross-correlation spectra C^{TB} and C^{EB} are useful diagnostics of noise and systematic effects.

Integrated Sachs-Wolfe Effect

The Integrated Sachs-Wolfe Effect describes the change in a photon's energy as it passes through evolving potentials when those potentials are still in the linear regime. (The non-linear ISW is known as the Rees-Sciama effect.) If a photon passes through a static well, the energy gained when it enters is cancelled by that lost when it climbs out, leaving no net effect. But if the well decays before the photon leaves, it retains that extra energy and is blueshifted. If the well deepens, it loses energy and is redshifted. The amount of energy change can be determined by integrating over the change in the potential, or $-\dot{\Phi}$. But the second general relativistic effect must also be taken into account: this can be thought of as the change in the potential causing a stretching or contracting of the space itself, which stretches or contracts the photon. So if the well is decaying, the energy gained by the photon is augmented by the effective contraction of the space, and vice versa, adding a term $\dot{\Psi}$ to the integral.

(Note that the two contributions to each of the SW and ISW case are physically the same, coming from Φ and Ψ in both cases. But the way that they combine is opposite. Adiabatic perturbations mean that photons are hotter inside the primordial potential wells, so the Φ and Ψ partly cancel. But in the case of the ISW effect, an energy loss due to a deepening well is reinforced by an energy loss due to the stretching of the photons.)

The Early ISW, the blue curve in Fig. 1.2 (a), arises when the universe is not fully matter-dominated at last scattering, and the potential fluctuation of modes that are smaller than the sound horizon begin to decay due to radiation pressure. At late times, potentials begin to decay when the matter dominated era ends (and either curvature or vacuum
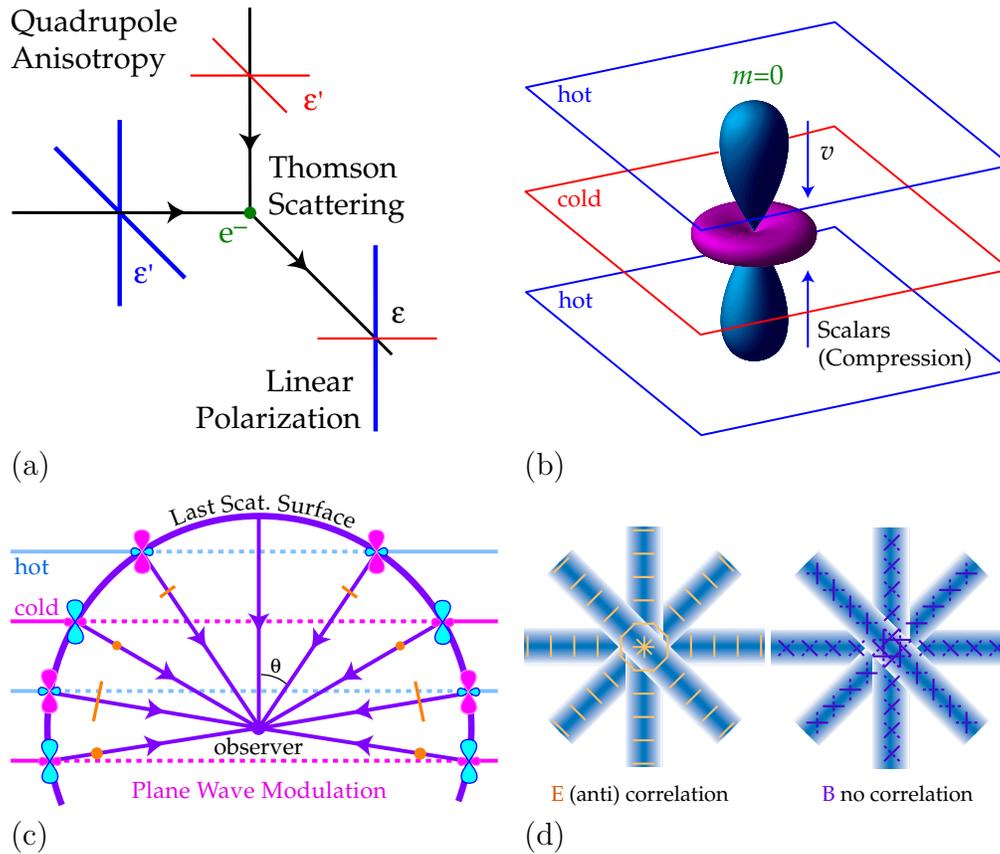
Figure 1.3 Illustration of polarisation signal arising from Thompson scattering of a quadrupolar radiation field (a). A plane wave density field gives rise to fluid flows and a local quadrupole (b). The polarisation pattern observed is shown in (c), and the cross-correlation between the polarisation E and B-modes with the temperature is shown in (d). (All from Hu & White 1997.)

energy dominate), giving rise to the large-scale component shown in pink.

Sunyaev-Zel'dovich Effect

Not shown on Fig. 1.2 is the small-scale contribution from the spectral distortion introduced when the photons pass through the hot ionised IGM of a galaxy cluster, known as the Sunyaev Zel'dovich (SZ) effect. Through inverse Compton scattering, lower energy photons are boosted in energy, so there is a temperature decrement in the Rayleigh-Jeans tail and a corresponding increment in the Wien regime. The electron temperature, T_e, and number density, n_e, determine the magnitude of the spectral distortion via the Compton y-parameter,

$$y \equiv \int \sigma_{\mathrm{T_e}} n_e \frac{kT}{m_e c^2} d\ell, \tag{1.6}$$

where ℓ is here the line element along the line-of-sight, and the temperature distortion is then

$$\frac{\Delta T}{T} = y[x \tanh(x/2) - 4] \tag{1.7}$$

with $x \equiv \hbar\nu/kT$. The Rayleigh-Jeans decrement becomes $\Delta T/T \sim -2y$. The spectral crossover is at 217 GHz for the thermal effect regardless of the temperature of the electrons. A variant of this thermal effect is the kinetic SZ effect where the proper motion of the cluster shifts this distortion spectrum. The SZ effect essentially adds point sources that are cool shadows at low frequencies and hot sources at high frequencies. This effect important for studying the cluster distribution, and when combined with X-ray observations, can provide an independent measure of the Hubble constant.

Other Secondaries

As shown in Fig. 1.7, the dominant component of CMB temperature difference maps is the dipole, a non-cosmological term induced by the observer's motion relative to the CMB. This is a function of the Sun's velocity, $V_{\odot}$, relative to the CMB, which is not known independently with any accuracy. For this reason, the $C_{\ell=2}$ term is then never used in cosmological analysis. *WMAP* measures a dipole of 3.346 ± 0.017 mK in the direction $(l, b) = (263°.85 \pm 0°.1, 48°.25 \pm 0°.04)$ (Bennett et al., 2003b), implying a velocity of ~ 370 km s^{-1}.

At late times, the universe becomes ionised again (due either to UV radiation from the first generation of stars or to AGN activity), and the total optical depth along the photon propagation path dampens fluctuations at smaller angular scales via Thompson scattering. This reionisation also introduces a secondary E-mode polarisation anisotropy signal through the local quadrupole, and the angular size of this signal gives the redshift to reionisation.

As mentioned above, the ISW effect in the non-linear regime of structure formation is known as the Rees-Sciama effect. Another interaction of CMB photons with matter is through gravitational lensing, a closely related phenomenon. This describes an effect which

conserves surface brightness but deflects the photons and thus affects the angular power spectrum, essentially smoothing the acoustic peaks. It also generates B-mode polarisation by changing the spatial distribution of the primordial E-mode signals. Lastly, it generates a small amount of power that is only noticeable at the smallest scales where the primordial signal is damped. See Fig. 1.6.

Cosmic Parameters

All of these (except the non-cosmic dipole) depend on the following basic set of cosmological parameters. The values given are those for the flat ΛCDM model from the *WMAP* 3-year analysis (available on the LAMBDA[8] site) including data from several balloon-born missions, galaxy surveys, supernova, etc.:

- the Hubble expansion rate, H_0, or $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}) = 0.708^{+0.015}_{-0.016}$, which is related to the age of the universe, $t_0 = 13.84 \pm 0.14$ Gyr;

- the fractional matter density, $\Omega_\mathrm{m} = 0.262 \pm 0.016$ (Ω is the density of a given component as a fraction of the critical density at which the geometry is flat, i.e., Ω is the physical density ρ divided by the critical density, $\rho_c \equiv 3H^2/8\pi G$, which is related to $\omega \equiv \Omega h^2 = \rho/(1.88 \times 10^{-26}\,\mathrm{kg/m^3})$);

- either the total density, Ω_tot, or the curvature, $\Omega_k \equiv K$, which for the other values given here is assumed to be zero, but from independent analyses allowing it to vary is $\Omega_k = -0.015^{+0.020}_{-0.016}$;

- the dark energy density, $\Omega_\Lambda = 0.738 \pm 0.016$;

- the baryon density, $\Omega_\mathrm{b} = 0.0437 \pm 0.0014$;

- the dark matter density, Ω_dm (redundant, determined by Ω_tot, Ω_Λ, and Ω_b);

- the the optical depth to last scattering, $\tau = 0.070^{+0.027}_{-0.028}$, which is related to the redshift to reionisation, $z_\mathrm{ion} = 9.4^{+2.6}_{-2.7}$ (depending on the reionisation model);

- the primordial power spectrum index, $n_s = 0.94^{+0.015}_{-0.014}$;

- the amplitude of the primordial spectrum of scalar perturbations at 0.05 Mpc, A_s, or alternatively the galaxy fluctuation amplitude, $\sigma_8 = 0.751^{+0.032}_{-0.031}$ (which also includes dependence on n_s); a related parameter is b, the Galaxy bias factor;

- θ, the ratio of the sound horizon, s_*, to the angular diameter distance, D, which is a function of $\rho_\gamma/\rho_\mathrm{b}$, Ω_Λ, Ω_m, H_0, Ω_tot, and w;

More complicated models may include any of the following optional parameters:

- $w \equiv p/\rho$, the dark energy equation of state;

---

[8] http://lambda.gsfc.nasa.gov/

- $\alpha \equiv dn_s/d\ln k$, the running of the spectral index;

- r, the tensor-to-scalar ratio;

- n_t, the spectral index of the tensor perturbations.

Codes such as CMBFAST[9] or CAMB[10] calculate power spectra for given cosmological parameters. These are then compared to the observed power spectrum to measure the cosmological parameters. For a summary of the cosmological parameters and their relations with each other, see Table 1 of Tegmark et al. 2004. Some useful scaling relations that relate the features of the power spectrum to the cosmological parameters can also be found there as well as in the Appendix of Hu et al. 2001.

Measuring these using the CMB alone runs into the problem of degeneracies where the same power spectrum features may arise from different combinations of parameters. Some of these degeneracies are discussed in Chapter 3, such as that between the scalar perturbation amplitude, A_s and the damping due to the optical depth, τ. It therefore requires complementary datasets to break these degeneracies. For example, there is a degeneracy between the matter and dark energy densities when measured by the acoustic peak scale in the CMB power spectrum. But supernova measurements giving the redshift-luminosity relation have a different degeneracy and the intersection of their confidence regions gives a much more accurate measure than either CMB or supernovae alone. In fact, there are in many cases a variety of different measurements of cosmological parameters, and generally, the agreement is surprisingly good. The "cosmic concordance" for the matter and dark energy densities is shown in Fig. 1.4.

1.2   Observations

1.2.1   Measurement to Cosmology

For a good review of radio astronomy in the context of the CMB, see Partridge (1995). The following is a brief summary of the issues involved.

The microwave sky is difficult to observe from the ground due to absorption by water vapour in the atmosphere, so the first challenge in CMB observations is to place the detectors in a high and dry location. Though this is possible from some observatories on the ground, many measurements are from either balloon-born detectors or satellites.

Measuring the temperature to the accuracy required for cosmological analysis is the next challenge and requires highly sensitive detectors. One option is high electron mobility transistors (HEMTs), which are linear photon amplifiers that are limited by quantum noise that increases with energy. They are thus used at lower radio frequencies such as in *WMAP* and Planck's LFI. (see Chapter 8 for more on Planck.) Bolometers are direct detectors

---

[9]http://www.cmbfast.org/
[10]http://camb.info/
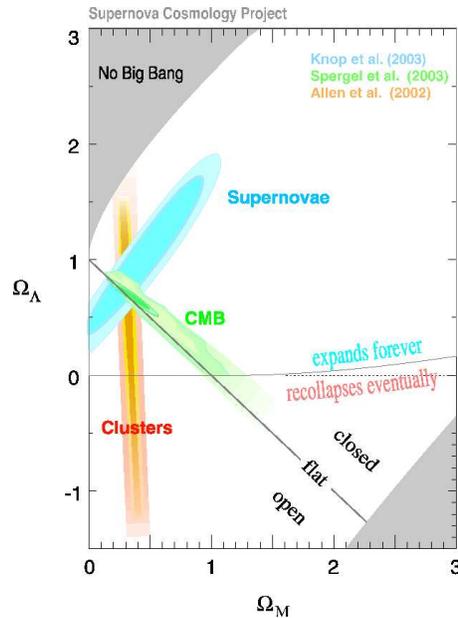
Figure 1.4 Cosmic concordance from the Supernova Cosmology Project (http://supernova.lbl.gov/).

that are more effective but only useful at higher frequencies, as they are dominated by background noise at low frequencies. These are used in Planck's HFI, for example.

Both *COBE*-DMR and *WMAP* are differential instruments that measure not the absolute temperature (like *COBE*-FIRAS) but the difference in temperature from two offset telescopes. This reduces the effects of such systematics as inevitably result from variations in the instrument temperature and gain. Likewise, the observation pattern, or scanning strategy, is also extremely important, as the best way to minimise the effects of such systematics on the resulting sky map is to observe large fractions of the sky in short periods of time repeatedly. The scan pattern also has implications for the efficiency of the map-making process.

Turning a time-ordered series of measurements at different positions with a given imperfect detector into a map of the measured sky temperature is itself a non-trivial task. It involves solving a large set of linear equations, i.e., inversion of a matrix whose size is proportional to the number of pixels. Matrix inversion scales as N_{pix}^3, and the available computing power is not enough to invert directly[11]. This problem is usually solved iteratively, though future missions may be able to perform direct inversion using supercomputers. See, e.g., Tegmark (1997) or Stompor et al. (2002) for discussions of map-making methods.

The next challenge is to measure the power spectrum for comparison with theoretical

---

[11] *COBE*-DMR mapped the sky into 6144 pixels. Maps from the *WMAP* first-year release had over 3 million pixels, and Planck will have even more. Each factor of two increase in ℓ_{max} implies a factor of four increase in the number of pixels.

predictions. In practice, we cannot know the true underlying distribution C_ℓ, or ensemble distribution, but rather can only measure the average over the one sample Universe we are given:

$$\hat{C}_\ell = \frac{1}{(2\ell + 1)} \sum_{m=-\ell}^{\ell} a_{\ell m}^* a_{\ell m} \qquad (1.8)$$

For a statistically isotropic sky (see next section), this measure is an unbiased estimate for the true underlying distribution, C_ℓ. For high values of ℓ, there are many modes in this average, but for low ℓ, there is a larger degree of uncertainty due to the small sample size. This is known as cosmic variance:

$$\langle \Delta \hat{C}_\ell^2 \rangle = \frac{2C_\ell^2}{(2\ell + 1)f} \qquad (1.9)$$

where f is the fractional sky coverage.

In addition to cosmic variance, the power spectrum must be extracted from a sky map that is likely to be incomplete, contaminated with foregrounds, and noisy (see, e.g., § 1.2.3). Even for missions that scan the full sky, the Galactic foreground is too strong to be successfully modelled and removed (see § 1.2.3), so all analyses must mask out this region of the sky. The non-Gaussian correlations that such a sky cut and non-uniform observation pattern introduce between harmonic coefficients complicates the power spectrum estimation. Efstathiou (2004) reviews the more common methods, and his recommended hybrid estimator is used for the 3-year *WMAP* analysis. At large angular scales, the maximum likelihood method that uses the full covariance matrix (taking into account both noise and sky cut) can be applied. But at smaller scales, the matrix is too big and a quadratic estimator is applied using a simpler representation of the coupling matrix and the pseudoC_ℓs, i.e., the distribution of pseudoharmonics calculated by taking the integral in equation (1.2) only over the region of the sky outside the mask. These are also called pseudo-C_ℓ estimators.

Lastly, once we have our estimated power spectrum, we want to use it to measure the cosmological parameters. Again, this is a non-trivial statistical and computational problem to solve due to the large number of possible parameters and degeneracies among them. One of the most successful methods has been the Monte Carlo Markov Chain (MCMC) analysis as implemented in, e.g., CosmoMC[12]. This method uses sampling to map out the multidimensional likelihood space of the data. At each sample point representing a given set of cosmological parameters, the likelihood can be calculated at that point simply by comparing the data to the power spectrum produced by CAMB. After an initial burn in period, the Markov Chain samples the underlying distribution, and therefore the number density of samples in the chain gives the probability distribution over the parameter space without having to calculate it analytically. The results can be used to generate marginalised likelihood contours in parameter space such as those shown in Fig. 1.4 and in Fig. 3.4.

---

[12]http://cosmologist.info/cosmomc/

1.2.2   Telescopes

Pre-*WMAP*

After radar technology profoundly influenced the course of WWII, astronomical radio observations took off and (eventually) revolutionised cosmology. Satellite telecommunications work at Bell Labs in the 60's led to the accidental discovery of the CMB as excess noise of 3.5 ± 1.0 K at 7.35 cm when the antennas were pointed at the sky (Penzias & Wilson, 1965). Once the possibility of pigeon droppings in the antenna was ruled out as the source, Penzias and Wilson called up the road to Princeton, where Dicke's group was able to provide an interpretation of the noise as a cosmic background temperature. That group had been in the process of developing a dedicated detector to measure this expected radiation and were effectively scooped by Penzias and Wilson. Only a few months later, they had their own measurement, at 3 cm, of of 3 ± 0.5 K (Roll & Wilkinson, 1966), and before long, measurements at several wavelengths showed consistency with a blackbody spectrum, while scans over the sky confirmed the isotropy to $\Delta T/T \lesssim 10^{-3}$.

But it took another twenty-five years for the next major advancement in CMB observations: the Cosmic Background Explorer (*COBE*). *COBE* consisted of three instruments: the Diffuse Infrared Background Experiment (DIRBE) in the infrared, a Differential Microwave Radiometer (DMR) to map the temperature anisotropies, and a Far Infrared Absolute Spectrophotometer (FIRAS) explicitely designed to compare the spectrum precisely with a blackbody. It was only with FIRAS that the spectrum was measured with high enough precision to confirm its blackbody nature over a broad range of frequencies, as shown in Fig. 1.5. DMR then for the first time detected the predicted fluctuations at roughly one part in 10^{-5}. (See Bennett et al. 1996 and references therein.)

Though more sensitive than any previous mission, the *COBE*-DMR data are very noisy, and the power spectrum measurements therefore do not reach beyond the Sachs-Wolfe plateau at relatively low ℓ. Balloon-born experiments measuring patches of sky with high precision were the next big step for CMB cosmology. Figure 1.6 (a) shows the data from some of these experiments, including ARCHEOPS (Benoît et al., 2003), MAXIMA (Lee et al., 2001), and BOOMERanG (Ruhl et al., 2003). Each of these measured the power spectrum at smaller angular scales and detected the acoustic peaks. As described above, the position of the first peak confirmed the existence of dark energy as indicated by supernova observations.

WMAP

The design of *WMAP* (see, e.g., Bennett et al. 2003a) explicitly aimed to improve not only the sensitivity but also the reliability of CMB measurements. The nearly back-to-back differential optical system was designed to minimise systematic effects and the radiometers passively cooled in order to provide a very stable thermal environment. The scan strategy covers a large area of the sky frequently and from different relative orientations, minimising systematics due to any beam asymmetries as well as gain or temperature variations. The satellite orbits the the second Sun-Earth Lagrange point and keeps the telescope shielded
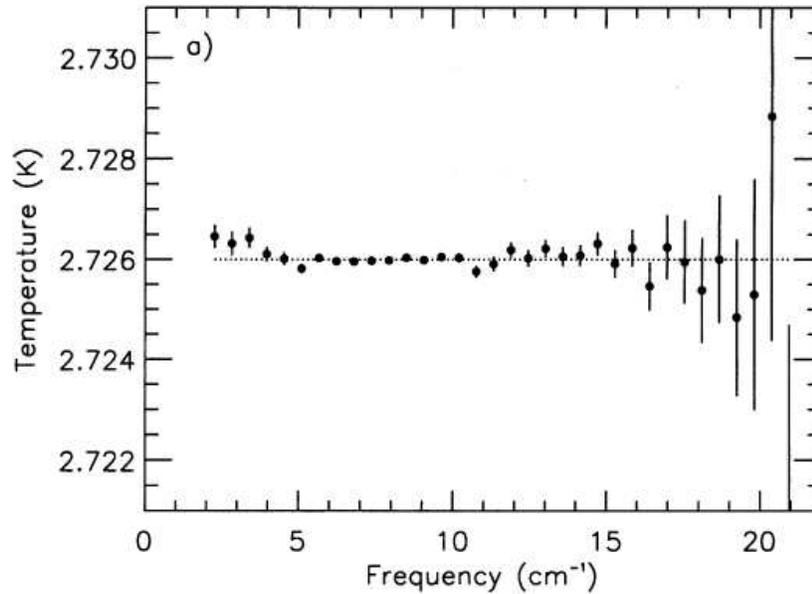
Figure 1.5 *COBE*-FIRAS measurements of the temperature of the CMB. From Mather et al. (1994). Note that a later recalibration by Fixsen et al. (1996) revised the measurement to 2.728 K.

from the influence of the Sun, Earth, and Moon.

WMAP ushered in the era of precision cosmology by providing the most precise measurements yet over the full sky and a set of five frequency bands from 23 to 94 GHz (Bennett et al., 2003b; Spergel et al., 2006). The DMR maps are compared to *WMAP* in Fig. 1.7, and the improvement in angular resolution is visually quite striking. A comparison of the power spectra in Fig. 1.6 likewise shows significant improvement at all but the smallest angular scales. Table II from Wang et al. (2003) summarises the "last stand before *WMAP*" parameter measurements compared to the newer results showing improvement particularly for Ω_Λ, Ω_b, and h due to the better positions of the first two acoustic peaks.

The future of CMB measurements in the post-*WMAP* era will be discussed in the Conclusions.

1.2.3   Foregrounds

Non-cosmological emission, mostly from the Galaxy, complicates the analysis of the CMB from microwave sky maps, but the data are useful for studying the emission processes themselves. Foreground analysis then depends on whether the aim is simply to remove any non-cosmological signal or to understand the physical nature of the different foreground components.

There is a narrow window in frequency space where the cosmic background dominates the foreground components for much of the sky, as shown in Fig. 1.8. Multi-frequency

(a)   (b)
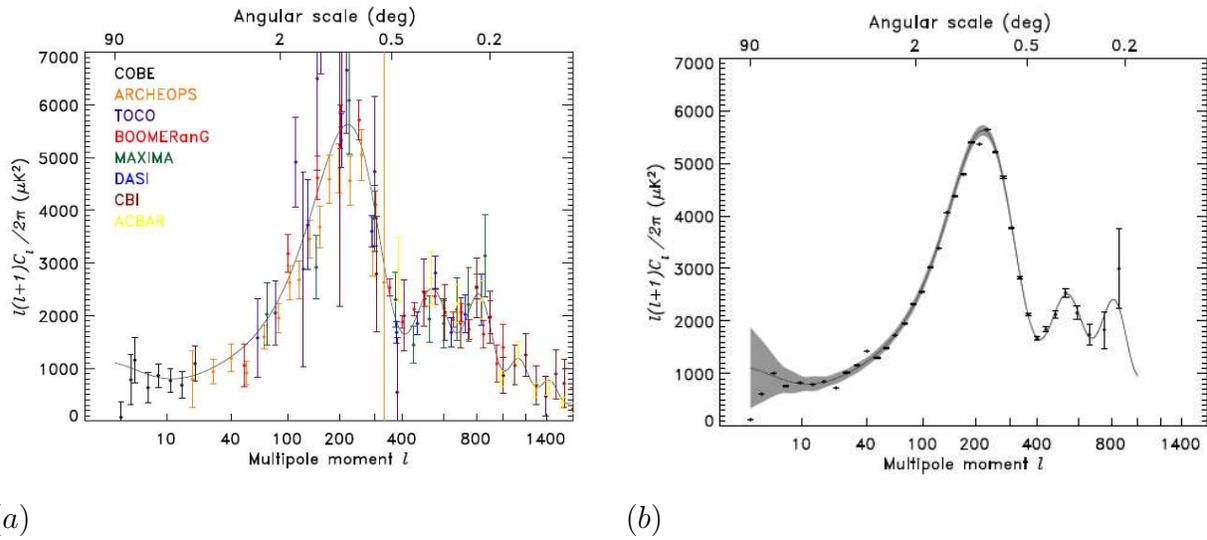
Figure 1.6 Comparison of pre-*WMAP* (a) and *WMAP* first-year (b) power spectra. (For a fair comparison of the state of our knowledge, the different measurements on the left would be combined, but the point of this plot is to show the first detections of the acoustic peaks and to emphasize the improvement *WMAP* gave over all previous experiments.) Figures from Hinshaw et al. (2003a).
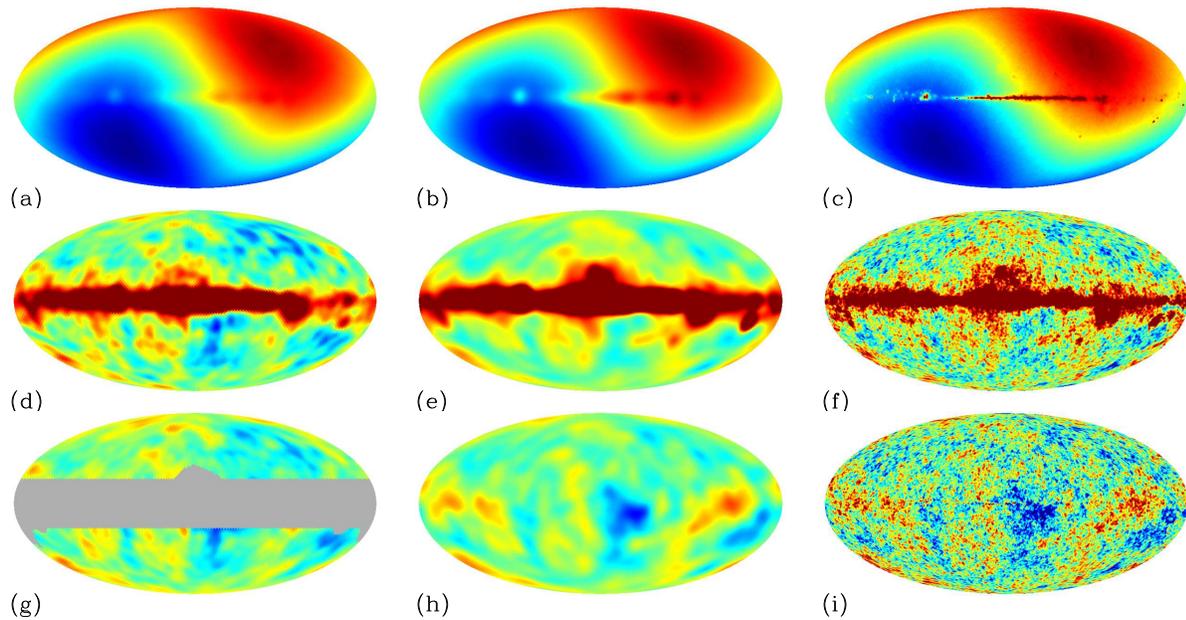
Figure 1.7 Microwave sky maps for *COBE*-DMR (left) and *WMAP* (middle), both smoothed to 10° FWHM, and *WMAP* at 1° FHWM (right). The top row shows the data as dominated by the non-cosmological dipole, the middle row dominated by the foregrounds (53 GHz for DMR, Q-band for *WMAP*), and the bottom row, cleaned (coadded for DMR, ILC for *WMAP*).
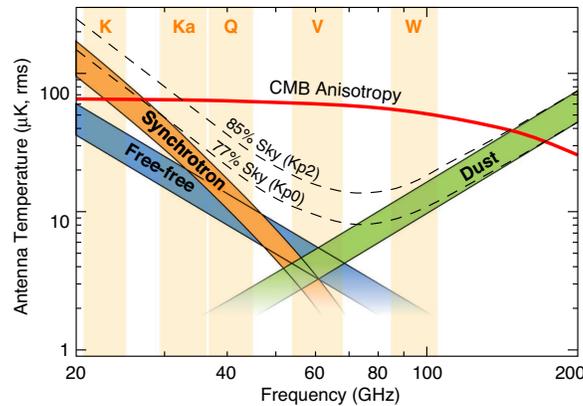
Figure 1.8 Spectral behaviour of microwave foregrounds compared to CMB and the different frequency bands of *WMAP*. From Bennett et al. (2003a).

missions like *WMAP* include not only the frequency where the foregrounds are expected to be minimum (the V-band) but also observations at surrounding frequencies to help understand the spectral behaviours of the different components. To achieve the sensitivity required to constrain cosmological parameters as well as to test for statistical isotropy and Gaussianity (see § 1.3) requires a very good foreground subtraction even in the V-band. This remains one of the primary challenges in CMB analysis.

Components

The three known foreground components in the microwave regime are: synchrotron emission from cosmic-ray electrons gyrating in the galactic magnetic fields; thermal dust emission from dust grains heated by UV radiation that is then re-emitted with a peak in the IR; and free-free (bremsstrahlung) emission from electron-ion interactions in the ionised ISM. There is a fourth component at low frequencies whose nature is still the subject of much debate. Because of its spatial correlation with dust, it is often called anomalous dust, but its spectral index, steeper than free-free and harder than synchrotron, leads the *WMAP* team to consider it a hard synchrotron component associated with dusty star-forming regions. (Such a hard component would be produced by higher energy cosmic ray electrons, which in turn represent a younger and more spatially localised component. Cosmic ray electrons likely come from supernova remnants in dusty star forming regions, thus providing a possible explanation for the spatial correlation with the thermal dust emission seen in the infrared.) But a variety of observations of individual regions have indicated that this component has a spectral turnover around ~ 10 GHz, which synchrotron emission would not show. This is instead the signature of emission due to spinning dust.

   Observations at other frequencies can be used to predict the microwave emission for each of these components but with several caveats. First, such templates themselves may contain noise, systematic errors, calibration uncertainties, etc. Furthermore, the spectral behaviour

of each component is not uniform over the sky and not necessarily well known. This makes it difficult to extrapolate from other observations widely separated in frequency. A template from the radio regime will be dominated by the component with the steeper spectrum, and a harder component that is morphologically different will dominate at microwave bands and therefore not spatially match the low-frequency template. This problem will be discussed in more detail in Chapters 4 and 5.

The middle panels of Fig. 1.7 show the strength of this emission in the plane relative to the CMB for both the DMR 53 GHz and *WMAP* Q-band at 41 GHz. No analysis is currently accurate enough to successfully remove the galactic plane, since the foregrounds there are so much brighter than the CMB, implying that even a very small fractional residual would still dominate the cosmic component. Assuming a plane-parallel model of the galaxy, this emission should roughly follow a cosecant drop-off at higher latitudes. For the *COBE*-DMR analysis, a large galactic cut ($|b| \gtrsim 20°$) was imposed to prevent galactic residuals from contaminating the analysis. With *WMAP*'s smaller beam, more of the sky can be used, but from 15% to 27% of the sky must still be masked, depending on the analysis. (The middle panel of Fig. 1.7 appears free of foregrounds by eye, but a higher resolution version shows residual point sources in the plane.) Though the Galactic foregrounds are strongest in the plane, they are not negligible even at high latitudes.

In addition to Galactic emission, extragalactic point sources are a significant source of foreground contamination at small angular scales. The known sources are excluded in the standard masks used in most of these analyses. The remaining unresolved sources are then modelled as a systematic error component. Again, observations at other wavelengths help to predict the number of such sources in the microwave band, but here as well spectral variation limits the accuracy of these predictions.

Figure 1.9 is a reference showing some of the notable features of the Galaxy as well as bright point sources and the ecliptic.

Analysis methods

Because of the problems mentioned above, foreground analysis is usually about making trade-offs to gain better knowledge of some components by making assumptions about others. The analysis then differs depending not only on the signal-to-noise ratio in the data and the available frequency coverage but also on whether the aim is to isolate the CMB or to study the foregrounds.

- **Template fitting** is perhaps the simplest approach, effectively treating the CMB as a noise term, assuming observations at other wavelengths can be scaled to each microwave band with one amplitude for the full sky, and performing a least-squares fit. In addition to the problem of the spectral variation, there is also the problem that the CMB may happen to correlate or anti-correlate with one of the templates, introducing a systematic error. This not only leads to increased residual foregrounds in the CMB estimate itself but also makes it difficult to interpret the fit values in terms of the emission processes. One way to address this problem is to constrain
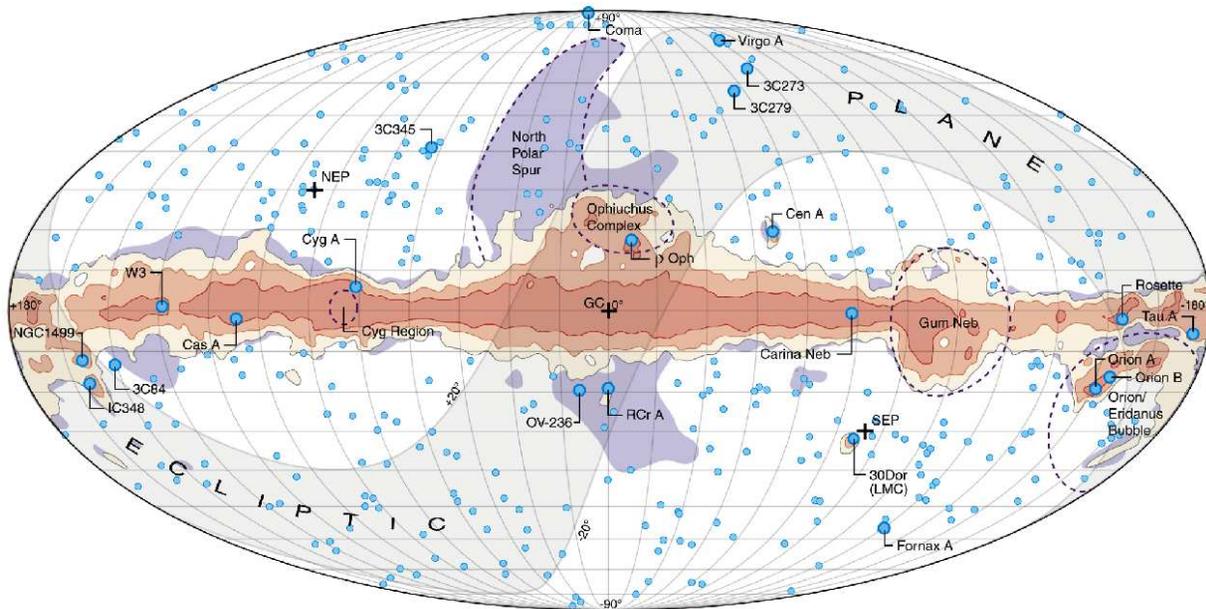
Figure 1.9 Notable features of the Milky Way and a few objects, including significant microwave point sources. Courtesy *WMAP* Science Team.

the fits by forcing the foreground fit amplitudes to follow a fixed spectral behaviour, e.g., a power law, based on knowledge of the spectrum of a given component in other regimes. The trade-off is the loss of information about the true spectral behaviour of the foregrounds in the microwave regime. This work will use template fitting in a variety of contexts and discuss the different implementations and limitations in each case. See sections 2.2.1, 4.2, and 5.4.1.

Template fitting is one of three different foreground analyses performed by the *WMAP* team. For power spectrum estimation at all but the largest scales, they use foreground cleaned maps from a constrained template fit. For the first-year analysis, the map of synchrotron emission at 408 MHz by Haslam et al. (1982) provided the synchrotron template. Some of the possible implications of this are discussed in § 4. For the 3-year results, instead of the Haslam synchrotron template, they form a template from the difference of two *WMAP* bands, K-band minus Ka. This difference map (in thermodynamic temperature units) cancels the CMB leaving only the various foreground residuals and noise. They expect that synchrotron should be the dominant component, and by using this map instead of the 408 MHz map, some of the spectral variation is taken into account. The trade-off is then that it is not easy to interpret the results in the context of understanding the foregrounds themselves. Whatever the dominant component of this K–Ka template, it still contains significant and unknown amounts of both free-free and anomalous emission. Disentangling the mixing requires making assumptions that then limit the conclusions one can draw from the fit values. See § 7.2 for discussion. In § 7.3, we simply want to remove any

foreground in order to study the cosmic component, so here we do use the K–Ka template. But for this reason, the foreground analyses in Chapters 4 and 5 use the Haslam 408 MHz map instead, as there, our aim is to study the foregrounds.

- The internal linear combination (ILC) map is created with a **minimum variance method** that assumes nothing about the foregrounds themselves (and in turn learns little about them). This method simply adds the five *WMAP* frequency bands with coefficients obeying the two constraints that the cosmic component must add to one and the variance in the result must be minimum. To take into account some of the spectral variation, the combination coefficients are determined independently over twelve regions on the sky. The trade-off for this flexibility is then discontinuities in the recombined result, which then requires smoothing, both of which then complicate the noise properties of the final map, making in unsuitable for use in power spectrum estimation at small scales. The end result is that shown in the lower right of Fig. 1.7, and it looks good by eye at larger angular scales. Point sources are still visible along the Galactic plane, however, and foreground residuals can remain even at high latitudes (see § 5.A). The *WMAP* team originally intended this map only for visualisation in the first-year release, but in the 3-year analysis, they claim it is good for $\ell \lesssim 10$ and use it in the cosmological analysis as part of the maximum likelihood power spectrum estimation at large scales.

- To get a better understanding of the foregrounds themselves, they use the **maximum entropy method (MEM)**, which uses the templates and spectral models only as priors. Press et al. (1988) define the entropy, S, as "the logarithm of the number of microscopically distinct configurations that all have the same macroscopic observables (i.e., consistent with the observed macroscopic state)." In this case, the macroscopic observables are the observed sky temperatures, $T(\nu, p)$ in each band ν and pixel p, and the configurations are the parameters describing the different components of the hypothesised signal, i.e., the model, $T_m(\nu, p)$. We then want to maximise the probability of the hypothesis given the data, namely $P(T_m(\nu, p)|T(\nu, p))$. Using Bayes theorem (7.1) as well as the definition of the negentropy, $H \equiv -S = \lambda(p) \sum_c T_c(p) \ln[T_c(p)/P_c(p)]$, we minimise $-\ln[P(T(\nu, p)|T_m(\nu, p))] = 1/2\chi^2(p) + H(p)$Here, $T_c(p)$ is the predicted temperature of component c at pixel p and $P_c(p)$ is the prior probability of that component. The templates from observations at other frequencies and their expected spectral behaviours form these priors, and the parameter λ controls how closely the fit results must remain to those priors. This method is applied iteratively, the prior updated each time using the residuals from the current solution.

    Maps of the three resulting components are shown on the top row of Fig. 1.10. For the priors on the bottom row of that plot, the templates are scaled to the appropriate band following theoretical predictions with a single power-law index across the sky for each component as assumed in the analysis. In general, the visual agreement is good , with the most significant deviations apparent in the "synchrotron" component,

(a) MEM synchrotron (Ka−band) (b) MEM free−free (Ka−band)   (c) MEM dust (W−band)

(d) Haslam (Ka−band)   (e) Finkbeiner H$_\alpha$ (Ka−band)   (f) FDS dust (W−band)

$0\mu K$ ▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰▰ $300\mu K$
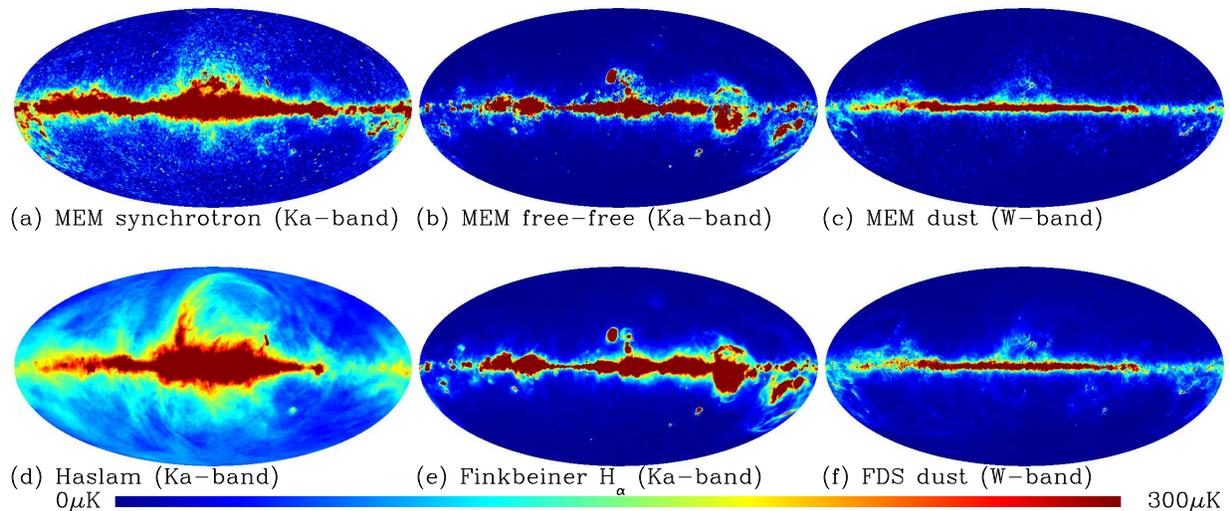
Figure 1.10 Foreground components resulting from *WMAP* MEM analysis (Hinshaw et al. 2006 and references therein) on the top row and the input templates on the bottom row. The latter are scaled to the given bands assuming theoretical predictions: $\beta = -3$ for synchrotron and $T_e = 8000K$ for free-free. (The dust prediction at W-band assumes $\beta = 2$.)

where the unknown anomalous emission is significant. Though this method is useful in separating different foreground components, the results have noise properties that cannot be easily characterised, again making them unsuitable for power spectrum estimation.

- Another interesting method is **independent component analysis (ICA)**, a method of separating statistically independent components using an equivalent (or greater) number of independent measurements, e.g., frequencies, that are linear mixes of the components and noise. Unlike MEM or template-based methods, this method can be used for *blind* source separation, requiring no assumptions about the spatial or spectral distribution of the foregrounds. The assumptions relate to the statistical properties of the components, namely that they be independent and that all but at most one be non-Gaussian. This method does assume that the frequency dependence does not vary across the sky. But when the different frequencies are fairly close, the result of such an analysis is more accurate than using templates that are from different regimes entirely. Spectral matching ICA (SMICA) is a variant based on covariance matching in Fourier space.

Tegmark et al. (2003) use two alternative methods to isolate the CMB component without assumptions about the foregrounds. The first, which we shall refer to as the TOH map, is related to the ILC method though combined in in harmonic space and using higher resolution than the ILC at 1°. This map resembles the ILC closely, though its power spectrum appears to have fewer foreground residuals at smaller scales. The second map

is a straightforward Wiener filter, which reduces the very small scale residuals along the plane but is not an unbiased estimate of the CMB. We will use the former in Chapter 2 as an alternative to the *WMAP* ILC map. Eriksen et al. (2004a) improve on the ILC method itself using Lagrange multipliers, and we will also use this "LILC" map in Chapters 2 and 6. (The *WMAP* ILC will then be referred to as the WILC.)

1.3 Statistical Isotropy and Gaussianity

As described above, the physics of the early universe is studied through the CMB power spectrum, C_ℓ, where the estimate of the underlying distribution comes from averaging the modes over the only sky we are given. In using this measured quantity, it is generally assumed that the anisotropies arise from an isotropic Gaussian stochastic process defined on the celestial sphere. (This is the simplest assumption for such quantum fluctuations during inflation, as described above.) In brief, this means that the statistical properties of the fluctuations can be characterised by a Gaussian distribution whose properties are the same in all directions. Specifically, this implies that

$$\langle a^*_{\ell'm'}a_{\ell m}\rangle = \delta_{\ell'\ell}\delta_{m'm}C_\ell \tag{1.10}$$

i.e., that all of the statistical information is contained in the power spectrum with no correlations between different modes and no information in higher order moments.

At large angular scales, it is generally simply assumed that the cosmic anisotropies are statistically isotropic (SI) and Gaussian. But it is clear from the top rows of Fig. 1.7 that the observed sky is highly non-Gaussian, and any uncertainty in the removal of non-cosmological foregrounds will leave non-Gaussian residuals in the data.

Gaussianity and isotropy arise naturally in most inflationary scenarios, though at very small angular scales, different inflation models can be studied by the different non-Gaussian signatures arising from, e.g., multi-field models, curvatons, etc. Such signatures are often studied by simply parameterising the primordial perturbation, Φ, as linear in some Gaussian random field ϕ_G but with an additional nonlinear term, e.g., a quadratic term, such that $\Phi = \phi_G + f_{\mathrm{NL}}\phi_G^2$. Then the level of non-Gaussianity expected is characterised by the value of the nonlinear coefficient, f_{NL}.

A variety of studies have investigated this issue. The *WMAP* team have performed several analyses that have confirmed these assumptions. Minkowski functionals describe the properties of the hot and cold regions at different threshold levels. The other method they use is the bispectrum, essentially the Fourier analogue of the three-point correlation function. Each of these tests returns nothing unexpected according to Komatsu et al. (2003) and Spergel et al. (2006), though it should be noted that these tests place limits on the value of f_{NL} as parameterised above rather than a generic non-Gaussian signal. Hajian et al. (2005) define a test called the bipolar power spectrum designed explicitly to look for statistical anisotropy, and they also find the sky to be as expected.

Though these tests and others confirm the standard picture, there have also been some surprising detections of violations of SI and Gaussianity.

1.3.1   Queerness of the Quadrupole and the Oddness of the Octopole

The low quadrupole amplitude had been a source of speculation for some time before *WMAP* confirmed the apparently low DMR measurement. (Hinshaw et al. 2006 give $\Delta T_2 \sim 200\mu K^2$ compared to the prediction from the best-fit model of $\sim 1200\mu K^2$.) In the first year data release, the octopole was also somewhat low, though not outside the cosmic variance uncertainty. These have been dubbed "queerness of the quadrupole" and the "oddness of the octopole" by Copi et al. (2006b).

The dearth of 2-point power at the largest angular scales can also be seen in another way. Figure 1.11 shows the angular correlation function:

$$C(\theta) = \overline{T(\hat{e}_1)T(\hat{e}_2)}_\theta \qquad (1.11)$$

i.e., the average correlation over all pixels separated by a given angular distance θ. Under the assumption of statistical isotropy, this can be related to the power spectrum by

$$C(\theta) = \frac{1}{4\pi} \sum_{\ell=0}^{\infty} (2\ell + 1) C_\ell P_\ell(cos(\theta)) \qquad (1.12)$$

Apparent from the figure is that above $\ell \sim 60$, there is almost no signal in the angular correlation function. This is closely related to the low quadrupole and octopole seen in Fig. 1.6.

There are a few things to note in comparing these plots. There is a difference between the curves representing the angular correlation in the ILC over the full sky (in solid black) and that over the sky where the Galactic plane region is masked (dashed black). This evidence for anisotropy is likely due to foreground residuals in the galactic plane. In the first-year data release, the *WMAP* team emphasised that the ILC was for visualisation purposes only, not for cosmological analysis. In the 3-year release, however, they state that the map is valid for cosmological purposes at large angular scales ($\ell \lesssim 10$) and they use it in their maximum likelihood power spectrum estimation for low ℓ (Hinshaw et al., 2006). That the angular correlation function at large scales changes with the sky cut indicates that there are still problems with this map at large scales and that the increased octopole in the 3-year release may be due to foreground contamination.

The quadrupole amplitude, despite these possible foreground effects, is quite low and remains so whatever the analysis. What is less clear is how significant it is. Initial analyses such as de Oliveira-Costa et al. (2004b) give a probability of such a low quadrupole as 1 in 20, but note that this is an *a posteriori* statistic: it applies to the one data point chosen for being furthest from the expectation. The *WMAP* team consider the quadrupole consistent with the best-fit ΛCDM model; it is not unexpected to see a few outlying points in the power spectrum, and the quadrupole happens to be one such. Other analyses such as Land & Magueijo (2005) confirm that the amplitude of the quadrupole itself is not statistically anomalous. But Copi et al. (2006b) examine the significance of the power deficit at large scales using the angular correlation function outside the Galactic plane region and find that it is anomalous at the 99.97% level.
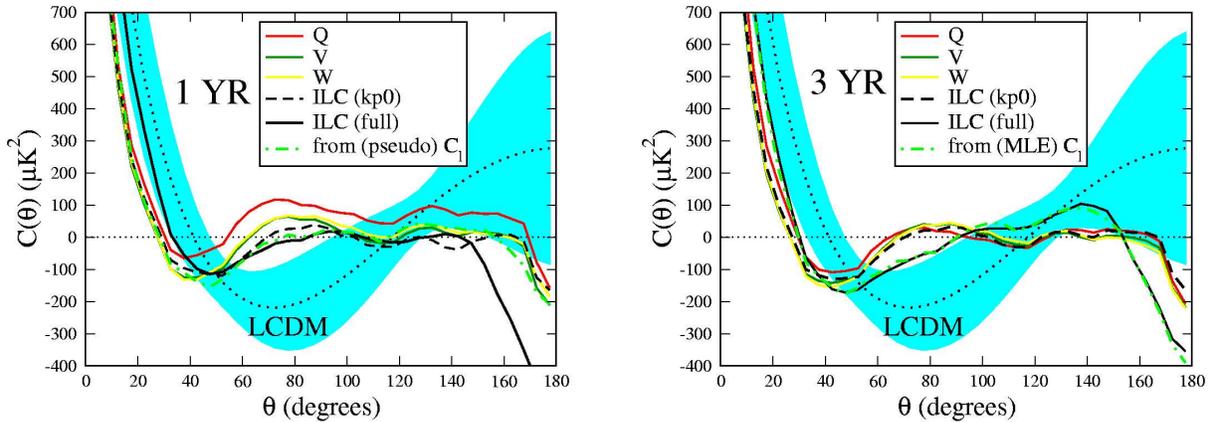
Figure 1.11 The angular correlation function of the *WMAP* data from Copi et al. (2006b). The expectation in the best-fit theoretical model is the dotted black line with cosmic variance uncertainty shown as the blue shaded region.

1.3.2   Axis of Evil

Land & Magueijo (2005) and Copi et al. (2006b) in independent analyses find that, while the quadrupole and octopole individually are consistent with isotropy and Gaussianity, the alignment of these two modes with each other (and perhaps extending up to $\ell = 5$) rules out isotropy at greater than the 99% level. Land & Magueijo (2005) dubbed this alignment the "the axis of evil".

The alignment is related to the planarity of these two modes. Figure 1.12 shows clearly that the hot and cold spots in the quadrupole and octopole lie along a plane in both cases and that the two planes are aligned with each other. This can be seen in the fact that the multipole vectors lie along similar planes and the normal vectors all lie in the same region in the southwest.

Though the symmetry axis points roughly toward Virgo, this is also suspiciously near the equinox as well as the dipole direction, which suggests that the anomaly may be related to systematic effects. In particular, the quadrupole is affected by a recalibration of the gain variation between the first- and 3-year *WMAP* releases. The change is small, but Copi et al. (2006b) point out that the effect of the correction is roughly perpendicular to the axis of evil. This leads them to speculate that there remains an unidentified systematic explanation with a similar morphology to explain the anomaly.

1.3.3   Cold Spot

In searching for non-Gaussianity or violations of statistical isotropy using the spherical harmonic representation, it is difficult to assess the morphology of any such signals because the functions have structure over the full sky. Wavelets are a powerful tool for analysing the morphology of the CMB, since in contrast, they are by definition functions over limited

(a)

(b)
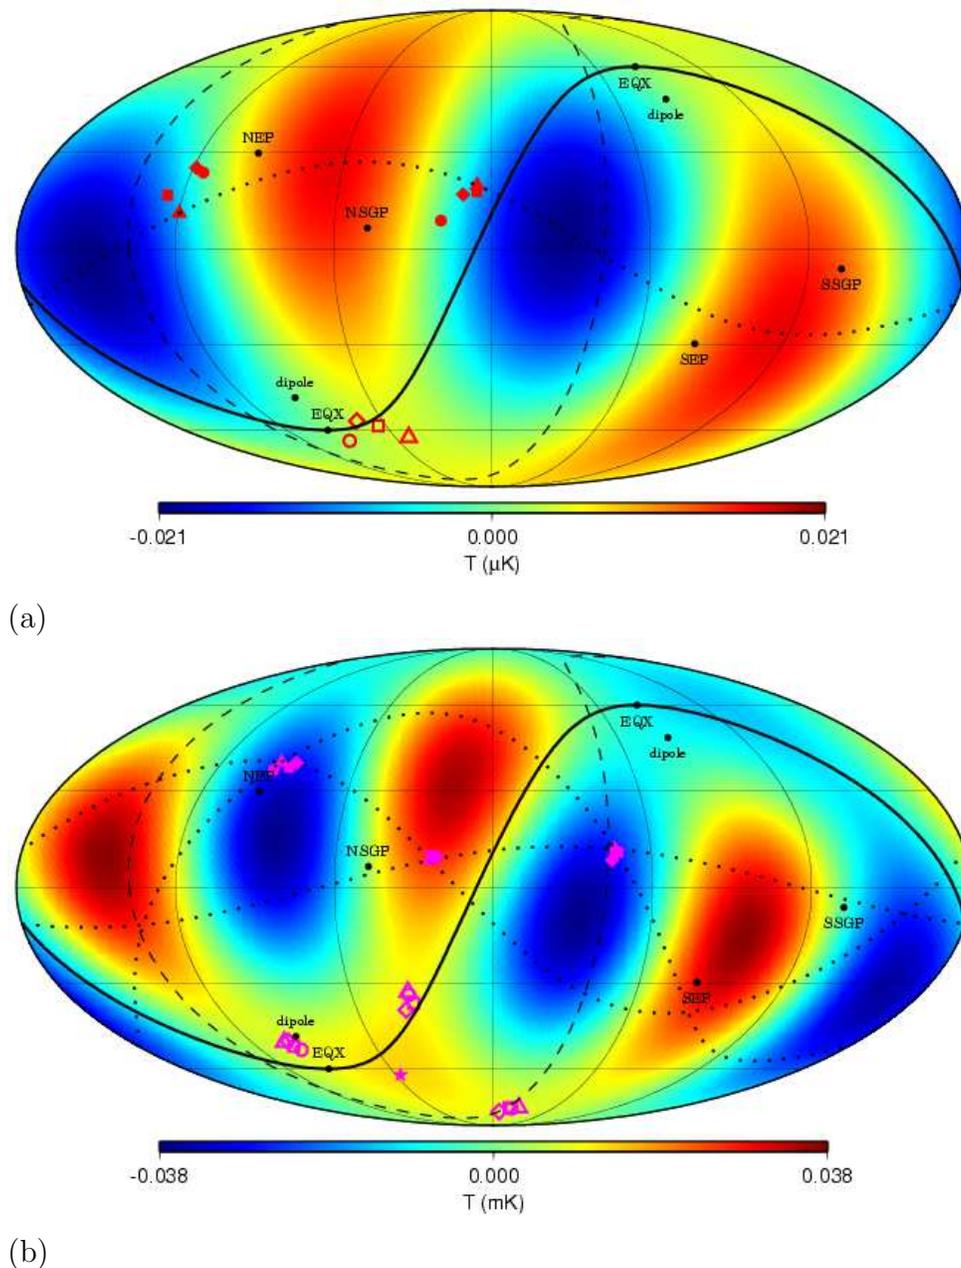
Figure 1.12 The quadrupole (a) and octopole (b) from Copi et al. (2006b) Fig. 1. Overplotted are the ecliptic plane (solid line) and the supergalactic plane (dashed line), with labels for the equinoxes (EQX), ecliptic poles (NEP and SEP), and supergalactic poles (NSGP and SSGP). The solid symbols show the multipole vectors (for different versions of the full sky, foreground cleaned data), and the open symbols the normal vectors. The dotted lines are the great circles connecting the pairs of multipole vectors for the 3-year ILC map. The solid magenta start in (b) shows the direction of the maximum angular momentum dispersion axis in the 3-year ILC octopole.

areas. As in the case of spherical harmonics, SI and Gaussianity ought to be reflected in the distributions of wavelet coefficients. Several groups have therefore used them to test these hypotheses.

Cruz et al. (2005) find a significant deviation from Gaussianity in the coefficients for spherical Mexican hat wavelets (SMHW). The source of this non-Gaussian signal is an unusually cold spot in the southeast. They quantify this spot as anomalous a the 99.8% level and show that excluding this region leaves a Gaussian sky. A more sophisticated technique using the higher criticism statistic was applied by Cayón et al. (2005) and confirms both the detection and its significance.

1.3.4   Power Asymmetry

Hansen et al. (2004b) and Eriksen et al. (2004b) study the CMB power spectrum in opposite hemispheres and find an asymmetry that is anomalous at the 2 to 3σ level. Figure 1.13 shows the asymmetry measured for several maps. The asymmetry appears in the *COBE*-DMR data, so it is not related to any *WMAP* systematic. The stability of the asymmetry over a range of frequencies and sky cuts as well the different orientation of the asymmetry seen in Galactic foreground templates makes it unlikely to be due to foreground residuals.

As discussed above, parameter estimation using the power spectrum depends on the assumption that the sky is statistically isotropic and Gaussian. This power asymmetry then throws doubt on the validity of some of those measurements. Hansen et al. (2004a) examine the effect of such an asymmetry and find that the optical depth, τ, in particular is affected due to its sensitivity to the ratio of large to small-scale power. The first-year *WMAP* estimate for τ was unexpectedly high at 0.17. Hansen et al. find that, under certain prior assumptions, the optical depth measured in the northern hemisphere is effectively zero while in the southern hemisphere, it is high at 0.24. Unsurprisingly, the *WMAP* value lies between the two. (In the 3-year *WMAP* results, the measured optical depth drops to 0.09, still roughly between the two hemisphere values. The latter measurement results partly from the polarisation signal. The asymmetry analysis with the 3-year data is ongoing.) It is difficult then to credit the detection of reionisation from such an anisotropic signal. This example emphasises the importance of testing the statistical isotropy and Gaussianity of the data.

1.3.5   Other

Here is a brief summary of some other interesting results, but by no means an exhaustive list of all analyses.

- Hansen et al. (2004c) found an asymmetry in the properties of "hills, lakes, and saddles" in the temperature distribution of thresholded maps. The axis defining the maximum asymmetry also lies near the ecliptic poles.

- Eriksen et al. (2004b, 2005) analyse N-point statistics in addition to the power spectrum and also find asymmetry at larger angular scales, again aligned with the ecliptic.
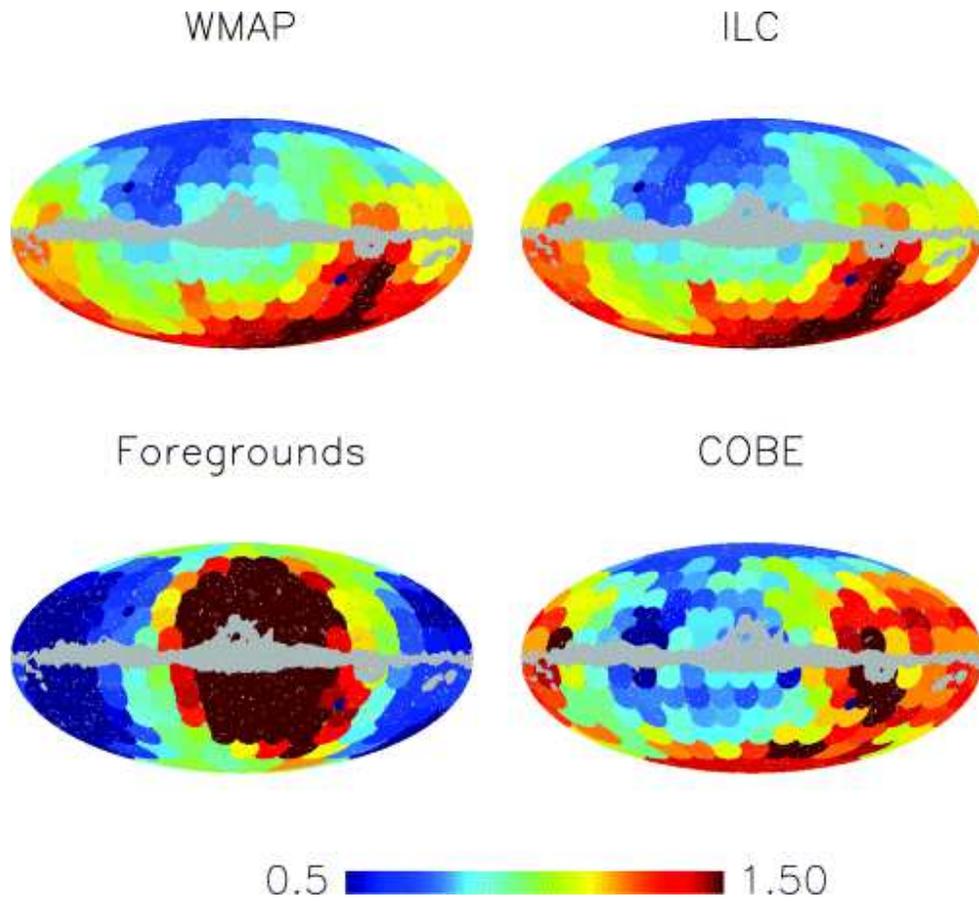
Figure 1.13 Asymmetric distribution of CMB power measured by Hansen et al. (2004b). Each coloured disc shows the ratio of power between a hemisphere defined about the axis through the centre of the disc and the opposite hemisphere. The upper left is the result using the *WMAP* V+W Coadded map, the upper right the ILC, lower left the foreground templates, and the lower right the *COBE*-DMR map. The small blue discs marks the ecliptic poles.

- Larson & Wandelt (2004) examine the temperature distribution using extrema statistics and again find a similar asymmetry in the variance between the north and south ecliptic hemispheres.

- Using Minkowski functionals and the skeleton (the zero-contour line of a map), Eriksen et al. (2004c) also find an asymmetry between Galactic hemispheres.

- An independent bispectrum analysis by Medeiros & Contaldi (2006) revealed indications of a non-Gaussian signal due to foreground contamination.

- Phase statistics have been applied by several groups (Chiang et al., 2003; Coles et al., 2004; Chiang et al., 2006). They report detections of non-Gaussianity that they attribute to foreground residuals. These methods and results will be discussed in more detail in Chapter 6.

- McEwen et al. (2005, 2006) examine the skew in the coefficients of two directional wavelets (i.e., wavelets that, unlike SMHW, are not azimuthally symmetric). They confirm the cold spot and find an additional non-Gaussian signal that is morphologically different from the above anomalies, lying largely in the northwest.

None of these are 5σ detections that would rule out statistical isotropy or Gaussianity with confidence, so it is unclear just how worried cosmologists should be. But the fact that there is a small collection of $\sim 3\sigma$ violations is certainly suspicious. If we should not necessarily be worried, at the very least, we should be curious.

Many of the above analyses have ruled out contamination by any known foreground based on both morphology and spectral behaviour. Neither is there a known systematic that could produce such anomalies. We shall consider, then, the possibility that the anomalies are truly cosmic in nature and represent problems with the standard cosmological model.

1.4   Non-standard Cosmological models

It must be remembered that Einstein's theory of general relativity is essentially a theory of geometry and does not predict the geometry, homogeneity, isotropy or any global topology[13]. The anomalies above have raised speculation that the universe may not have a trivial topology and may not be isotropic. As mentioned above, the flat FRW model is one solution for the special case of an isotropic and homogeneous universe. It is the simplest and as such favoured in the absence of evidence to the contrary. (There is still debate about what would constitute the "simplest" or most "natural" topology, with arguments made for both compact and infinite possibilities.) But it is possible that some of the anomalies described above may be the first such evidence.

---

[13]See Appendix A.6 for a definition of terms.

Below are introduced a few examples of homogeneous but non-standard models that are particularly interesting in light of *WMAP* data and will be discussed further in the following chapters.

1.4.1   Compact Topologies

Many, including Einstein, have argued that a compact universe is simpler and more "natural" than an infinite universe (Levin et al. 1998 and references therein, or Misner et al. 1973, p. 704).

De Oliveira-Costa et al. (1995) among others considered "small universe" models as a possible explanation for the observed homogeneity of the universe. In a compact, multiply connected universe, if the size of the fundamental cell is much smaller than the current horizon, then we observe many copies of the same universe, which would then appear homogeneous at the largest scales. This idea was put forward as an alternative to the inflationary paradigm. It is interesting to note that these models were considered uninteresting based on the *COBE* observation of too *much* power at large scales to allow a small enough cell size to explain the homogeneity. At the time, the quadrupole was not considered anomalous, in part due to the loose constraints on the CMB power spectrum.

But as the power spectrum was measured with increasing accuracy, the lack of power on large scales and the axis of evil revived speculation that the universe may not be infinite. If instead it is compact, then fluctuations would not be larger than the size of the fundamental cell. They would show a drop-off in power at scales approaching this size and non-Gaussian correlations between such modes. In the multi-connected case such as the torus, an asymmetry in the fundamental cell size in different directions could lead to the alignments of the largest modes. If the fundamental size is smaller than the horizon, an observer would see multiple copies of an individual object in different directions. For the CMB, this would correspond to identical temperature fluctuations along "circles in the sky" at the self-intersections of the surface of last scattering. (See, for example, Fig. 7 of Riazuelo et al. 2004b.)

De Oliveira-Costa et al. (2004) speculated that this axis of evil might be due to a compact universe model with one smaller dimension in the direction of Virgo. They apply two tests for compact models, the S-statistics (discussed in § 6.2.4) and matched circles, but find no evidence that the data are inconsistent with an infinite universe. An independent analysis by Cornish et al. (2004) searching for back-to-back circles also turned up no evidence of a compact universe. Note that these searches must assume a certain topology and range of sizes to test, since computational resources do not allow all possible pairs of circles to be tested in a feasible amount of time. Kunz et al. (2006) apply a method using the full covariance matrix to test the sky for a given topology, and again find no evidence for a finite universe for the cases they examine.

Another interesting topology is the Poincaré Dodecahedron, which can be imagined in its 2-dimensional analogue of a sphere tiled with 12 pentagons. Luminet et al. (2003) show how this model is consistent with the large-scale power measured by *WMAP*. Roukema et al. (2004) find a possible confirmation with a matched-circles test, but the statistical

significance of the possible detection is marginal.

Yet these models remain of interest. Linde (2004) discusses the preference of certain inflation theories for such compact models. Barrow et al. (2001), by contrast, discusses how compact topologies may explain the observed universe without the need for inflation.

1.4.2 Bianchi Models

In the context of Riemannian geometry, before general relativity was invented, the Bianchi classification system defined the isometry classes of Riemannian 3-manifolds. In other words, the Bianchi types I through IX describe the possible geometries of the Universe. Those which admit the FRW model are types I, V, VII$_0$, VII$_h$, and IX. (Note that in this work, in the context of Bianchi models, the topology is assumed to be non-compact. As mentioned above, the topology can limit the geometry, and not all Bianchi models are allowed in compact topologies. See Barrow et al. 2001, Table 1.)

Types VII$_0$, VII$_h$ and IX are the most general solutions to the Einstein equations for homogeneous, anisotropic universes in the flat, open, and closed cases respectively. (Types I and V are special cases of VII$_0$ and VII$_h$.) Each contains the corresponding FRW model as a special case. The generic anisotropic cases allow universal rotation or vorticity about a given axis and differential expansion or shear along that axis. Vorticity means that at every point in the universe, the rest of the universe appears to be rotating about an axis running through that point. Shear means that objects in the direction along that axis appear to be receding at different velocities from objects away from it.

This anisotropic geometry imprints an anisotropic pattern on the CMB due the rotation of the geodesics themselves. Photons propagating from the surface of last scattering to the observer will be red- or blueshifted differently depending on the path they take. Examples are shown in Fig. 2.1. Note that a universe with vorticity must then have a handedness. Another important aspect of the type VII$_h$ models is the resulting asymmetry; in the case of hyperbolic, negatively curved geometry, the geodesics are focused in one direction along the symmetry axis. Flat type VII$_0$ models have no such asymmetric focusing, but retain the spiral structure and are consistent with type VII$_h$ in the limit as the density approaches critical from below. By contrast, closed type IX models have *no* spiral structure, and therefore, as Barrow et al. point out, a detection of any such spiral would unambiguously rule out a closed universe, however large the uncertainty about the value of Ω_{tot}.

Note that these models are not consistent with the inflationary paradigm. The early exponential expansion is expected to wipe out any anisotropy just as it would any curvature. The CMB has been used to place ever lower limits on the amount of vorticity and shear, starting with Collins & Hawking (1973), who used the fact that *no* anisotropies had yet been detected in the CMB to place fairly loose limits on vorticity. Kogut et al. (1997) and Bunn et al. (1996) tightened these limit significantly with *COBE*-DMR data. Their limits came from the lack of detection of any such spiral pattern. But as will be shown in Chapter 2, the Bianchi type VII$_h$ are particularly interesting in light of the anomalies in the *WMAP* sky.

1.5   Aims

The aim of this work is to examine the morphology of the observed microwave sky, both the cosmic and galactic components.

In light of the anomalies described above, the question of the topology and geometry of the universe deserves renewed scrutiny. In the case that an anisotropic geometry results in a deterministic pattern, we use template-fitting methods. Chapter 2 describes a method to search for a correlation between any such template anisotropy pattern and the data. The challenge in this case is the fact that we have no a priori knowledge of the relative orientation of the model on the sky. The total convolver algorithm provides an efficient way to search over the three-dimensional space of Euler rotation angles. We tested this method using Bianchi type VII_h models largely for the pragmatic reason that the derivation of the CMB anisotropy pattern had already been done by Barrow et al. (1985). Our intention was to use first-year *WMAP* data to tighten the limits on vorticity placed by Kogut et al. (1997) and Bunn et al. (1996) with *COBE*-DMR data. The results proved surprisingly interesting, and in Chapter 3, we examine more thoroughly the viability of these models. Section 7.3 provides an update using the *WMAP* 3-year data release and additional analysis ruling out the influence of foregrounds and noise on the results.

The galactic foregrounds are, from a technical point of view, easier to study with template fitting, because the orientation is fixed to the galactic reference frame. But unlike the cosmic anisotropy component, the spectral behaviour of each foreground can vary across the sky, which then complicates such analysis. Chapter 4 looks at this issue by examining the implications of neglecting such spectral variation. This then motivates Chapter 5, where we analyse the sky in small enough patches that the spectral index can reasonably be assumed constant over the region. Here, the difficulty is then that the noise in the data as well as the uncertainty in the morphology of the templates both become important. But we can still measure some of the underlying spectral variation, which is particularly interesting for the anomalous dust component.

It is not always the case, however, that we wish to study something with a known morphology. We may be interested in anisotropic topological models or simply foreground residuals representing failures in our knowledge of the spatial and spectral distribution of the galactic emission. In these cases, we need tools that study the morphology without templates, and for this, we find that phase statistics are a promising tool. Chapter 6 describes how CMB phases contain information about the issues of statistical isotropy and Gaussianity and how to construct tests for violations due either to compact topologies or foregrounds.

Chapter 2

Fast and Efficient Template Fitting of Deterministic Anisotropic Cosmological Models Applied to *WMAP* Data

Note: this chapter was published as Jaffe et al. (2006a). § 2.6 includes work done by my collaborators, including analysis of the quadrupole, power spectrum asymmetry, and wavelet kurtosis. Section 2.8 has been added.

2.1   Introduction

The widely accepted model in cosmology, the so-called Concordance Model, is that of an isotropic and homogeneous universe with small anisotropies generated by primordial fluctuations in the inflationary field. These anisotropies are present in the cosmic microwave background (CMB), which should then be statistically isotropic and Gaussian. Many CMB studies therefore examine the CMB from a statistical point of view with the intention of testing for violations of these properties. Alternative cosmological models have not, however, been completely ruled out, and there are several anomalies in the *WMAP* data that indicate that such models merit further investigation by alternate means.

We investigate methods for testing any deterministic anisotropic cosmological model. The predicted anisotropy template can be compared to the data using fitting techniques in both pixel and harmonic space to search for correlations. We present a description of these methods and apply a fast and efficient algorithm for searching the full sky for the best orientation of a template relative to the data. We test these methods with both full- and incomplete-sky datasets, and use simulations to characterise the significance of the results.

Motivated by the morphology of several detected violations of Gaussianity and/or isotropy in the *WMAP* data (de Oliveira-Costa et al., 2004b; Eriksen et al., 2004b; Hansen

et al., 2004b; Vielva et al., 2004), we test our methods using Bianchi type VII$_h$ models and the *WMAP* first-year data. A preliminary analysis was published in Jaffe et al. (2005), where we reported on a surprisingly significant detection of a Bianchi model at the 99.7% significance level compared to simulations. Here we present an improved search of the model space, confirm the basic result, and discuss in detail issues such as foreground contamination and chance alignments.

2.2   Methods

2.2.1   Template Fitting

Given any anisotropy pattern that contributes to the data as an additional component of the observed microwave sky (whether topological in origin, as in the case of Bianchi models, or foreground), we perform a fit of the template to the *WMAP* data as has been done in the past by, e.g., Górski et al. (1996) and Banday et al. (1996) for foreground analysis. The best-fit amplitude α for a template vector $\mathbf{t}$ compared to a data vector $\mathbf{d}$ can be measured by minimising

$$\chi^2 = (\mathbf{d} - \alpha\mathbf{t})^T \mathbf{M}_{\mathrm{SN}}^{-1}(\mathbf{d} - \alpha\mathbf{t}) = \tilde{\mathbf{d}}^T \mathbf{M}_{\mathrm{SN}}^{-1}\tilde{\mathbf{d}}, \tag{2.1}$$

where $\mathbf{M}_{\mathrm{SN}}$ is the covariance matrix including both signal and noise for the template-corrected data vector $\tilde{\mathbf{d}} \equiv \mathbf{d} - \alpha\mathbf{t}$. Solving for α then becomes

$$\alpha = \frac{\mathbf{t}^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}}{\mathbf{t}^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{t}}. \tag{2.2}$$

To compare multiple template components to a given datasets, e.g., different foregrounds, the problem becomes a matrix equation. In the case where we have N different foreground components, we define

$$\tilde{\mathbf{d}} = \mathbf{d} - \sum_{k=1}^{N} \alpha_k \mathbf{t}_k \tag{2.3}$$

and

$$\mathbf{M}_{\mathrm{SN}} = \left\langle \tilde{\mathbf{d}}\tilde{\mathbf{d}}^T \right\rangle = \mathbf{M}_{\mathrm{S}} + \mathbf{M}_{\mathrm{N}}. \tag{2.4}$$

In this case, minimising $\tilde{\mathbf{d}}^T \mathbf{M}_{\mathrm{SN}}^{-1}\tilde{\mathbf{d}}$ leads to the following set of equations,

$$\sum_{j=1}^{N} \mathbf{t}_k^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{t}_j \alpha_j = \mathbf{t}_k^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}. \tag{2.5}$$

This is the simple system of linear equations $\mathbf{A}x = \mathbf{b}$, where

$$A_{kj} = \mathbf{t}_k^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{t}_j,$$
$$b_k = \mathbf{t}_k^T \mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}$$
$$x_k = \alpha_k.$$

Where only one template is present, this reduces to equation (2.2) above.

The errors $\delta\alpha_\nu^k$ are the square root of the diagonal of $\mathbf{A}^{-1}$. The matrix $\mathbf{A}$ gives information about the cross-correlation between the templates themselves.

Note that the above is equally valid in pixel space or harmonic space. In the former, it is very easy to account for incomplete sky coverage or to remove, e.g., the Galactic Plane, by simply including in the data vectors only the relevant pixels, and likewise by including only the corresponding rows and columns of the covariance matrix. The noise in pixel space is usually well represented by a diagonal matrix representing uncorrelated pixel noise. But the signal covariance matrix in pixel space is large and not sparse, which makes harmonic space more convenient when possible. In harmonic space and under the assumption of Gaussianity, the signal covariance is diagonal, and with the approximation of uncorrelated noise that is uniform over the sky, the noise covariance can be made to be so as well. The difficulty in harmonic space is the sky coverage. As discussed by Mortlock et al. (2002), the coupling matrix to cut out the Galactic Plane using a cut the size of $|b| > 20°$ becomes numerically singular for resolutions of $\ell_{max} > 50$. Cuts such as the conservative Kp0 mask defined by the *WMAP* team remove more of the sky and, due to their structure, the coupling matrix is more difficult to compute.

We define a method that applies harmonic space fitting to the full-sky cases using highly processed maps discussed in §2.3.2. This allows us to increase the computational efficiency using the algorithm described in §2.2.2. Here, we use a uniform mean noise approximation that has a diagonal covariance matrix. We use pixel space fitting for each band separately in the cut-sky analysis where the Galactic Plane region is masked out. (At the *WMAP* signal to noise level, little would be gained by simultaneously fitting the different bands, and the memory and CPU requirements to invert the matrix would become onerous.) Again, we use a diagonal noise approximation that this time takes into account the observation pattern but not the effects of smoothing. Comparisons of fits with fully correct noise to those using these approximations show that the results do not vary significantly (at most a few percent, or a small fraction of the error bar.) All codes have been cross-checked with identical inputs to confirm identical outputs.

Note that the cosmic monopole and dipole are not known, and though a best-fit dipole is subtracted from the data in the map making process, small residual monopole and dipole terms remain in the data. For this reason, we cannot include this component in the fit. In harmonic space, any monopole and dipole terms can simply be excluded or ignored by setting, e.g., $C_1 = C_2 = 10^8\mu K^2$. In pixel space, we fit the monopole and dipole simultaneously as independent components. See §2.4.3 for discussion.

The above method determines the best-fit amplitude for a given template at a fixed orientation relative to the sky. For foreground fitting, that is generally all that is required, but in the search for an anisotropic cosmological component, there are two additional complexities. Firstly, we have no *a priori* guess for where the symmetry axis may be pointing and must thus search the entire sky. §2.2.2 describes the algorithm we use to do this quickly and efficiently. Secondly, we may have an infinite number of possible templates (for example, parameterised as described in §2.3.1) among which we want to find the "best", so in addition to determining the best-fit amplitude for each template, we need a way to

compare how well different templates fit the data and to select the most interesting. §2.2.3 discusses how we address this.

2.2.2 Total Convolver

The search for the best orientation of a template compared to the data requires that we evaluate the statistic α described in the previous section at every possible relative orientation of the template and data. Working in harmonic space allows us to use an algorithm based on Fourier transforms to speed up this search significantly.

In the case of full-sky analysis, the location or orientation of the template does not affect the error, i.e., $\delta\alpha$ is invariant. Then the maximum of α is found at the maximum of the numerator in equation (2.2) above, $\mathbf{t}^T\mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}$. Neglecting for the moment the covariance matrix, the quantity to be maximised is simply the convolution of the data with the template. We seek the maximum over all possible locations and orientations, and this can be found efficiently using the total convolver algorithm described in Wandelt & Górski (2001), which was originally developed for map making using instrument beams.

This algorithm decomposes the Euler angles into what amounts to a scan pattern and then takes advantage of the form the convolution takes in harmonic space to simplify the calculation. The rotation operator $D(\Phi_2, \Theta, \Phi_1)$ can be factored into $D(\phi_{\mathrm{E}}, \theta_{\mathrm{E}}, 0)D(\phi, \theta, \omega)$, where a predefined scan pattern determines θ_{E} and θ, which in the case of full-sky coverage are both $\pi/2$ so that the set of angles $(\phi_{\mathrm{E}}, \phi, \omega)$ covers the full sky at all possible orientations. See Wandelt & Górski (2001) Figure 1. (In *only* this context of total convolution on the full sky, ϕ corresponds to the polar angle and ϕ_{E} to the azimuthal angle. Elsewhere in this paper, these are represented by the more common θ and ϕ.) Defining $T(\phi_{\mathrm{E}}, \phi, \omega) \equiv \mathbf{t}^T\mathbf{d}$ as the quantity to be maximised, $b_{\ell m}$ as the spherical harmonic coefficients of the template $\mathbf{t}$, and $a_{\ell m}$ as that for the data $\mathbf{d}$, the convolution is then (Wandelt & Górski, 2001, eqs. 9 & 8)

$$T_{mm'm''} = \sum_l a_{\ell m} d^\ell_{mm'}(\theta_{\mathrm{E}}) d^\ell_{m'm''}(\theta) b^*_{\ell m''} \tag{2.6}$$

$$T(\phi_{\mathrm{E}}, \phi, \omega) = \sum_{m,m',m''} T_{mm'm''} e^{im\phi_{\mathrm{E}}+im'\phi+im''\omega}, \tag{2.7}$$

where $d^\ell_{mm'}(\theta)$ is the real function such that $D^\ell_{mm'}(\phi_2, \theta, \phi_1) = e^{-im\phi_2} d^\ell_{mm'}(\theta) e^{-im'\phi_1}$. The problem has then become simply to calculate $T_{mm'm''}$ and Fourier transform to $T(\phi_{\mathrm{E}}, \phi, \omega)$ to find the maximum. To take into account the signal and noise covariance, we simply use a "whitened" data vector, $\mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}$.

The total convolver can find the best-fit position with an accuracy limited only by the resolution of the inputs. The positional accuracy is $\pi/\ell_{\max}$, which for our analysis is $2.8°$. Note that this is larger than the size of a pixel at the usual HEALPix resolution of $N_{\mathrm{side}} = \ell_{\max}/2$.

It should also be noted that searching the full sky will *not* return an unbiased estimate for the amplitude. Simulations with a known input value for a particular template at a

known position will, on average, have slightly higher amplitudes returned by the search. If the correct template location is simply fit to an ensemble of simulations with additional CMB and noise, the returned amplitudes will have a Gaussian distribution with the correct mean and variance, but the same is not true when one is searching for the best location and orientation as well. This is because the search is seeking the maximum, and the resulting distribution is a form of extreme value distribution[1], which introduces a small positive bias in the results. For realistic situations with CMB and noise in addition to the component we are fitting, the total convolver is likely to find a maximum amplitude a small distance away from the true position. How different the amplitudes and positions are on average depends on the particular case in question, since it is a function of how dominant the template is compared to the CMB and noise, how much the template structure changes over angular distance, etc. This is quantified for the particular case in question in §2.4.2 using simulations.

This method is approximately two orders of magnitude faster than performing the fit in harmonic space over a grid of individual rotations one at a time. The disadvantage is the storage requirement for the matrix T, which increases with the third power of the resolution and becomes over 2GB for a HEALPix resolution of $N_{\text{side}} = 128$, or angular resolution of $42'$.

2.2.3 "Best" fit model and significance

As mentioned above, when it is not one unique template for which we are testing but rather a set of possibilities, we need not only to find the best fit of each to the data but also to find the best fit among the possible models. Depending on how the model space is parameterised, there can be an infinite number of possibilities. Previous studies seeking upper limits on shear and vorticity Kogut et al. (1997); Bunn et al. (1996) used two different statistics to determine the "best"-fit model.

Given a model, Kogut et al. define the best-fit position and amplitude in terms of $\Gamma = \alpha/\delta\alpha$. They used *COBE* data, for which no full-sky analysis was possible. In the case of incomplete sky analysis, the amount of template structure that is masked changes the significance of the fit. A large amplitude where most of the structure is masked by the Galactic Plane cut is not as interesting as a lower amplitude fit where the structure is included. By finding the maximum not of α but of Γ, they attempt to find the most significant fit rather than simply the maximum amplitude.

Bunn et al. use a different statistic to accomplish the same effective selection. They define $\eta_1 \equiv (\chi_0^2 - \chi_1^2)/\chi_0^2$, where χ_1^2 is as in equation (2.1), and χ_0^2 the corresponding statistic for the data by itself, uncorrected for any anisotropic component. The difference is then an indication of how much better the data fit the (statistically isotropic CMB) theory after correction for the anisotropic model.

Finding the maximum of Γ is equivalent to finding the maximum of η_1 (though Bunn et al. use a different statistical method.) So for a given model, either statistic can be used

---

[1]See, e.g., `http://mathworld.wolfram.com/ExtremeValueDistribution.html`

to find the best-fit amplitude and position. But it becomes more complicated to compare one model to another in order to determine which model fits the data better.

The problem with the simple approach used by Kogut et al. (1997) as well as in our preliminary analysis Jaffe et al. (2005) of using Γ or η_1 to find the best model is that the distribution of these values for chance alignments is not the same for all models. Though they are generally quite similar, differences in the tails of the distributions mean that a given value of Γ has a slightly different significance for different models. This means that finding the maximum of Γ might have missed other models that are significant, but where the tail of the distribution doesn't reach as high in Γ. In other words, the significance of the fit found in our original result is not incorrect, but it is possible that such an analysis fails to detect another significant model.

For this more complete analysis, we analyse a set of LILC simulations Eriksen et al. (2004a); Eriksen et al. (2005), using the above formalism to characterise the distributions of α values for a given model. In this analysis, for a given model, we compare the α value (equivalently Γ, since $\delta\alpha$ does not change for a given model on the full sky) for the *WMAP* data against the ensemble of simulations. We can then quantify the significance of a given model fit to the data based on the percentage of LILC simulations where the model fits with a lower amplitude. This gives clearer indication of which are the most interesting models than a simple Γ or χ^2 statistic. Comparison of the results using α or η_1 in this way show there is little difference between the two in terms of how significant a given fit to the data is against the simulations. In the following analysis, we use the numbers for α only.

2.2.4 Visualisation: cross-correlation signal maps

It is helpful to be able to visualise what parts of the sky are driving a particular fit. To do this, we simply note that the numerator of equation (2.2), $\mathbf{t}^T\mathbf{M}_{\mathrm{SN}}^{-1}\mathbf{d}$, can be rewritten in pixel space as $\sum_p[\mathbf{L}^{-1}\mathbf{t}]_p[\mathbf{L}^{-1}\mathbf{d}]_p$, where $\mathbf{L}$ is the "square root" of the covariance matrix $\mathbf{M}_{\mathrm{SN}}$, or its lower triangular decomposition found from, e.g., Cholesky decomposition. A simple visualisation is to turn this into a map, where each pixel contains the product of $[\mathbf{L}_{\mathrm{SN}}^{-1}\mathbf{t}]$ and $[\mathbf{L}_{\mathrm{SN}}^{-1}\mathbf{d}]$ at that pixel. This map shows exactly what regions on the sky drive the fit at a given orientation. This is particularly important when certain regions of the sky are known to be contaminated; these plots show whether or how much those regions affect the fit. Examples will be shown in §2.5.2.

2.3 Data and simulations

Here we describe the particular class of models we investigate and the datasets used in the analysis.

2.3.1   Bianchi Models

Bianchi type VII_h refers to the class of spatially homogeneous generalisations of Friedmann universes that include small vorticity (universal rotation) and shear (differential expansion) components. (Type VII_0 includes the flat Friedmann-Robertson-Walker model, and VII_h includes that with negative spatial curvature as special cases.) Barrow et al. (1985) solve the geodesic equations to derive the induced CMB anisotropy by linearising the anisotropic perturbations about the Friedmann models. Their solution does not include any dark energy component, which is a significant shortcoming considering the preponderance of evidence that now points to $\Omega_\Lambda \sim 0.7$. But we examine them first as a test of our template fitting methods and secondly because of the intriguing possibility that they may explain several anomalies in the data.

Following the prescription in Barrow et al. (1985), we construct a template for the anisotropy induced by vorticity (ω) and shear (σ). Bianchi type VII_h models are parameterised by the current total energy density Ω_0 and a parameter x (Collins & Hawking, 1973):

$$x = \sqrt{\frac{h}{1 - \Omega_0}},\qquad (2.8)$$

where h is related to the canonical structure constants and is that to which the type VII_h refers; see Kogut et al. (1997),Bunn et al. (1996), and Barrow et al. (1985). This parameter can be understood as the ratio of the scale on which the basis vectors change orientation to the Hubble radius (present values.) The resulting temperature anisotropy pattern is then described by (Barrow et al., 1985, eq. 4.11)

$$\frac{\Delta T}{T} = \left(\frac{\sigma}{H}\right)_0 \{[B(\theta_R) + A(\theta_R)]\sin(\phi_R) \pm [B(\theta_R) - A(\theta_R)]\cos(\phi_R)\},$$

where A and B are also functions of x and Ω_0 and include integrals over conformal time that trace the geodesic from the surface of last scattering to observation. The angles θ_R and ϕ_R are not the observing angles; those are rather $\theta_{\rm ob} = \pi - \theta_R$ and $\phi_{\rm ob} = \pi + \phi_R$. The sign on the $\cos(\phi_R)$ term (or alternatively, the ϕ_R to $\phi_{\rm obs}$ transformation) determines the handedness. Then σ determines the amplitude of the fluctuation, and x the pitch angle of the spiral. The vorticity is then

$$\left(\frac{\omega}{H}\right)_0 = \frac{\sqrt{2}(1 + h)^{1/2}(1 + 9h)^{1/2}}{6x^2\Omega_0}\left(\frac{\sigma}{H}\right)_0.\qquad (2.9)$$

Note that the shear and vorticity values in our original paper, Jaffe et al. (2005), contain an error in amplitude, though the basic conclusions are not affected.

Equation (2.9) can be rewritten as

$$\frac{\Delta T}{T} \propto \cos(\phi_R \pm \tilde{\phi}).\qquad (2.10)$$

In other words, for a given θ_R, the temperature variation follows a $\cos(\phi_R)$ dependence. The phase shift $\tilde{\phi}$ is ultimately a function of θ_R and the two physical parameters, x and

Ω_0. The result is a spiral pattern with approximately $N = 2/\pi x$ twists. The smaller the x, the smaller the scale at which the basis vectors change their orientations and the tighter the resulting spiral. In the case of $\Omega_0 < 1$ models, geodesic focusing leads to an asymmetry wherein the spiral structure appears compressed in one direction along the rotation axis.

The template is calculated as $\frac{\Delta T}{(\sigma/H)_0}$, i.e. the contents of the curly brackets in equation (2.9) times the average CMB temperature, so that the shear $(\sigma/H)_0$ is the amplitude of the template to be found by fitting it against the CMB anisotropies, ΔT. Examples are shown in Figure 2.1, where the template is plotted without the normalisation by the shear. In generating all of our Bianchi templates, we have taken the redshift to the surface of last scattering, or recombination, as $z_{\mathrm{rec}} = 1100$. (Changing to, e.g., $z_{\mathrm{rec}} = 1000$ lowers the amplitude of the anisotropy by $\sim 15\%$, implying a corresponding increase in the value of the shear $(\sigma/H)_0$ for a given ΔT.)

We make the simple and pragmatic assumption that the anisotropy induced by the geometry simply adds to the statistically isotropic and Gaussian component.

We examine a grid of such models over $0.1 \leq \Omega_0 \leq 1.0$ in increments of 0.05 and over $0.1 \leq x \leq 10.0$ in increments of 0.05 in the interval $0.1 \leq x \leq 1.0$ and then logarithmically sampled up to $x = 10$. A finer grid was also examined surrounding the best-fit model, $0.52 \leq x \leq 0.68$ and $0.42 \leq \Omega_0 \leq 0.58$ in increments of 0.02. For the largest values of x, the spiral has almost disappeared (because the scale on which the basis vectors change orientations becomes larger than the horizon size), and so models of higher x are self-similar. Smaller values of x start to become physically unrealistic. Collins & Hawking (1973) point out that for $x \sim 0.05$, the characteristic length scale over which basis vectors change orientation becomes comparable to the size of large scale structure, which means that lower values are ruled out by observations of large scale homogeneity. Furthermore, as discussed in §2.5.3, small values of x require higher precision analysis than is feasible.

2.3.2   Data

For this work, we are interested only in large scale structure. In all of the following, unless otherwise noted, we use maps in HEALPix[2] format Górski et al. (2005) at a resolution of $N_{\mathrm{side}} = 32$ and smoothed to an effective beam of FWHM 5.5°, with harmonics up to $\ell_{\mathrm{max}} = 64$.

The following full-sky maps are used in this analysis (where all *WMAP* data products are from the first year data release):

- The full-sky *WMAP* Internal Linear Combination (WILC) map released by the *WMAP* team. See Bennett et al. (2003c). This map is formed by taking linear combinations of the different bands such that the foregrounds, each of which has a different spectral dependence from the CMB, are removed leaving only CMB. The different weights of the linear combination are determined solely by the data, via minimum variance, rather than any prior assumptions about the foreground behaviour.

---

[2]http://healpix.jpl.nasa.gov/
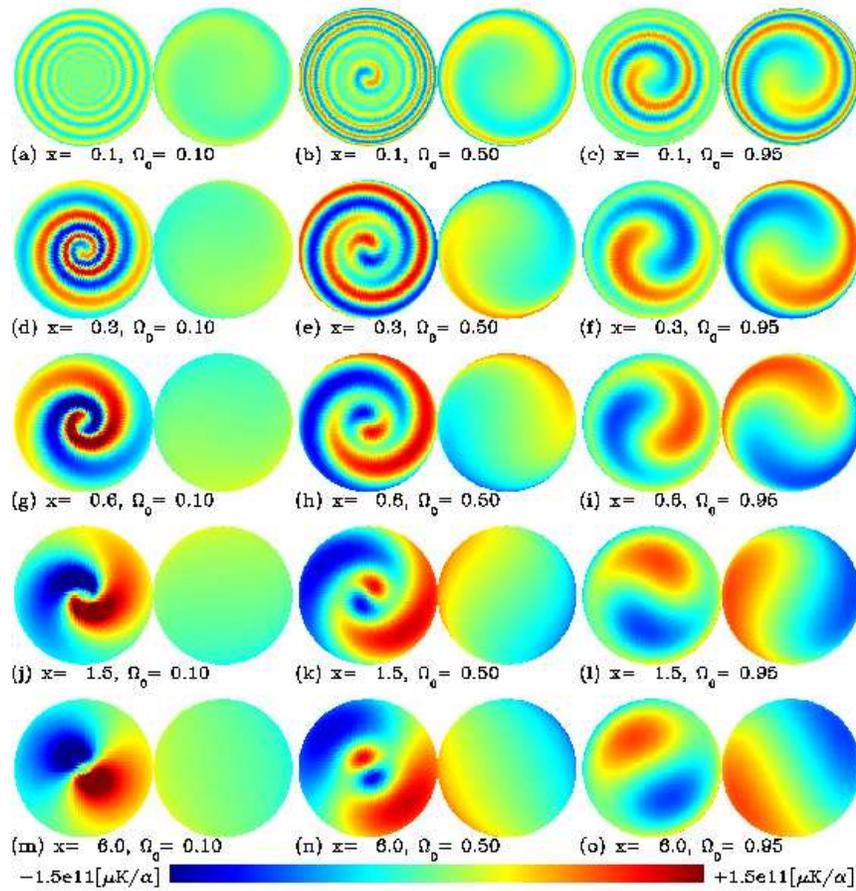
Figure 2.1 Examples of left-handed Bianchi anisotropy templates in orthographic projection, all on a common colour scale to show the relative amplitudes. These must be multiplied by a factor of $\alpha = (\sigma/H)_0$, i.e., the shear (realistically of order $< 10^{-9}$), in order to give the amplitude of the observed anisotropy in μK. Note that these have been rotated by $\beta = -90°$ to move the centre of the structure from the $-\hat{z}$ pole (as defined in equation 2.9) to the Galactic Centre.

- The Lagrange Internal Linear Combination (LILC) map of Eriksen et al. (2004a); Eriksen et al. (2005). The weights used to form the WILC map are slightly suboptimal with respect to the minimum-variance criterion (Eriksen et al., 2005), and this is corrected in the LILC map, which uses Lagrange multipliers to compute the ILC weights.

- The foreground-cleaned map of Tegmark et al. (2003), hereafter TOH. This map is also generated by a linear combination of bands, where in this case, the weights are determined in harmonic space.

All of these maps contain residual foreground emission, some of which is visible by eye along the Galactic Plane and some of which extends to high latitudes. It should be noted that none of these maps is intended for high-precision CMB analysis, but we nevertheless use them in the following to locate the best-fit Bianchi template by full-sky convolution. Simulations show that these fits are affected by two opposing biases (see §2.2.2 and §2.5.3) that are larger than the effects of the foreground residuals (see §2.4.2), thus justifying our use of these maps despite their known disadvantages. In general, we use the full-sky maps initially to locate best-fit axis for each Bianchi model (see §2.4.2), and then verify the amplitude using partial-sky algorithms on the following additional data:

- *WMAP* uncorrected maps for each of the five frequency bands, coadded from each differencing assembly using noise weighting; see Bennett et al. (2003b). Also noise-weighted, coadded combinations of bands Q+V, V+W, Q+V+W, Q-V, V-W, Q-W.

- Kp0 intensity mask, excluding 23.2% of the pixels where the K-band intensity is high and also 0.6° around known point sources, downgraded to $N_{\mathrm{side}} = 32$.

and observations at other wavelengths as foreground templates:

- the Finkbeiner et al. (1999) model for thermal dust emission (hereafter FDS);

- the Schlegel et al. (1998) 100μm intensity dust template (hereafter SFD), which is used an alternative to the FDS model; see discussion in §2.5.2;

- the Finkbeiner (2003) Hα template, with dust correction $f_{\mathrm{d}} = 0.5$;

- the Dickinson et al. (2003) Hα template with no dust correction, which is used as an alternative to the Finkbeiner template;

- the Haslam et al. (1982) 408MHz map of synchrotron emission processed by Davies et al. (1996).

These foreground components are fit simultaneously to each band over the incomplete sky using the Kp0 mask, which reduces the effects of foreground contamination on the fit amplitude (see §2.4.3.) Note that though we are simultaneously fitting the foreground components, these templates are not accurate enough in the Galactic Plane region for full sky fits to be reliable.

2.3.3   Gibbs samples

In addition to the *WMAP* data products, we also analyze a set of Gibbs sampled maps that were generated by the method described by Jewell et al. (2004),Wandelt et al. (2004), and Eriksen et al. (2004d). Effectively, this method samples the space of CMB signal maps that are consistent with the data, taking into account both noise characteristics and limited sky coverage. Thus, each single Gibbs sample represents a full-sky, noiseless CMB signal consistent with the data assuming Gaussianity, and the distribution of such maps describes the full CMB signal posterior distribution.

Such sampled maps may thus be analyzed very efficiently using the total convolver method described above, since neither sky cut nor non-uniform noise complicate the analysis. These allow us to avoid the problem of foreground residuals in the Galactic Plane, since this region of the Gibbs samples contains only CMB signal that is either consistent with the structure outside the plane, in the case of large enough scales, or entirely Gaussian random, in the case of smaller scales. The ensemble of fit results then reflects how well the template fits the CMB signal posterior distribution. In the following, we analyze ensembles of 1000 samples corresponding to each of the three cosmologically important *WMAP* Q, V, and W bands.

2.3.4   Assumed signal covariance

Given that we are searching for evidence of anisotropy, the description of the expected signal covariance is not trivial. Bianchi models in particular are not compatible with inflation theory and do not make any prediction for fluctuations at the surface of last scattering. Clearly, a self-consistent theory is required to explain the observed anisotropies in addition to the Bianchi component, and in particular, that theory must be consistent with the acoustic peaks now detected at smaller scales. No such theory currently exists, but we note that the Harrison-Zel'dovich power-law spectrum prediction predates inflation theory. Because it has been shown to match the data very well on small scales, we use the inflationary prediction as a starting point.

The signal covariance expected after subtraction of any Bianchi component is then assumed to be that of Gaussian, isotropic CMB fluctuations fully characterized by the power spectrum. We use the best-fit *WMAP* theoretical power-law spectrum to perform our fit. One could then refine the input spectrum based on the result (i.e., do a new parameter estimation using the corrected sky) and iterate. In the present analysis, however, we do not aim to improve the power spectrum estimation. Template fitting proves to be insensitive to the assumed power spectrum. (The fit result changes by less than 3% when using a flat, $Q = 18\mu\mathrm{K}$ power spectrum instead.) So for the purposes of this analysis, the best-fit *WMAP* theoretical power-law spectrum is sufficient.

2.4   Performance, bias, and accuracy

In order to interpret the results of the analysis using real data, we need first to quantify the effects described above. The model selection accuracy, the bias due to the maximisation over rotations, any bias due to foreground residuals, and the distribution of chance alignments are all effects that we can quantify using simulations.

These are generated by the LILC simulation pipeline of Eriksen et al. (2004a); Eriksen et al. (2005). The simulations start with a Gaussian CMB signal generated from an assumed power spectrum and are then smoothed to the beam width of each *WMAP* differencing assembly. Pixel noise is added, uncorrelated and following the instrument properties and observation pattern described in Bennett et al. (2003b). Lastly, the three foreground components above are added to create simulated raw data for each of the ten differencing assemblies. The LILC algorithm is then used to reconstruct the corresponding processed, foreground-cleaned sky. Though these are known to underestimate somewhat the amount of residual emission along the Galactic plane, they provide a vital indication of the morphology and approximate amount of such residuals that may be present in the WILC or LILC maps.

We apply the fitting methods outlined above to the ensemble of LILC simulations, with and without an additional known anisotropic signal, to characterise how well the methods perform. In most of the below, a set of 1000 simulations were used in the full grid searches and cut-sky pixel space fitting. An expanded ensemble of 10 000 LILC reconstructions was used to refine the significance measures for the two best-fit models found as described in §2.5.

2.4.1   Model selection accuracy

First, we add a known Bianchi component (the particular template and amplitude found in our initial analysis, Jaffe et al. (2005)) to a set of LILC simulations and perform the full sky search over all rotations (using the total convolver) and over the grid of models. We find that the most significant model returned is close ($\pm \sim 0.1$ in x and Ω_0) to the correct model in $\sim 50\%$ of cases. Among the other $\sim 50\%$, a qualitatively different model was found to be the best-fit, but the correct model was still found to be over 99% significant in most cases. In other words, only in $\sim 23\%$ of realizations was the correct model not detected.

We must then see if we can distinguish the correct model from a false detection by other means such as incomplete sky fits with simultaneous foreground template fitting. These give an idea how much the full-sky fit is affected by residuals in the Galactic Plane. Furthermore, models that appear far apart in the model space may in fact be fitting to the same CMB structure. We therefore select the several most significant models to examine in more detail. Then we look at what structures are driving the fits and how they behave when the galactic plane is excluded and foreground templates simultaneously fit. These tools give an additional qualitative way to compare different model fits.

2.4.2   Full-sky fitting accuracy

Next, we consider a known Bianchi component added to the input noiseless, pure CMB realization (as opposed to the LILC reconstruction) and see how well its position and amplitude are recovered by the full-sky fit. For 1000 simulations, a Bianchi component (at the same position and amplitude as our best-fit against the real data) is added to the input CMB sky and then fit using the total convolver method described above. In $\sim 80\%$ of realizations, the returned fit is within $5°$ (approximately the beam width) of the correct location. (In the orientation angle, it is less accurate due to the self-similarity of the spiral structure under such rotations. The returned orientation is within $10°$ in 52% of the simulations.) The amplitudes average $\sim 7\%$ *higher* than the input value (as noted in §2.2.2), with an RMS error of about 80% the calculated error. Neither of these facts is unexpected, since these values are the selected maxima, and their distribution is not Gaussian. The results are quantitatively the same for the LILC reconstructed skies, indicating that the foreground residuals do not introduce a significant additional bias in the case where a real Bianchi component is being fitted. Note that simulations where the input Bianchi model has an amplitude a factor of ~ 3 higher show a much smaller *relative* bias ($\sim 1\%$), as one would expect.

2.4.3   Cut-sky fitting accuracy

The cut-sky fits are performed with the Bianchi model at the fixed location found as the best-fit using the full-sky total convolver method. As described in §2.2.2, there is a bias introduced by the selection of the maximum amplitude position. This bias will also be reflected in the cut-sky fits, though masking out the Galactic Plane should remove some of the bias due to residual foreground emission.

For fits to the raw data outside the Kp0 mask, eight template components are fit simultaneously to each band: the three foreground templates described in §2.3.2, a monopole term, the three spherical harmonics representing the real-valued dipole terms, and the Bianchi template.

For simulations with no additional Bianchi component, the results show amplitudes on average 6% lower than those from the LILC fits. This is further indication that chance alignments are affected by residuals in the plane, since the exclusion of that region tends to lower the fit amplitude.

Simulations with an additional Bianchi component at a known position and amplitude were run through the same pipeline, i.e., first the full-sky LILC reconstruction was used to find the best-fit location, then that location used to fit the template to the cut sky in pixel space. As described above, the total convolver will return a position that is very close to the true position but where the fit amplitude happens to be highest due to CMB and noise contributions. These will also affect the cut-sky fits, which also show a bias of $\sim 3\%$. This is lower than the bias in the full-sky fits, showing that a few percent of the full-sky bias is due to residuals in the galactic plane region. The relative drop in amplitude between the full- and cut-sky fits for true detections is on average half the drop in the case of chance
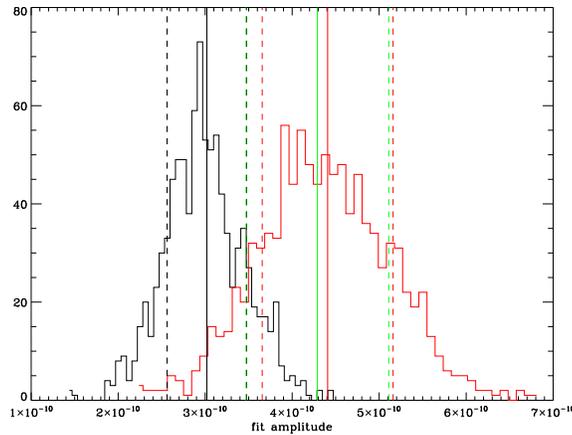
Figure 2.2 Distributions of fit results for the Bianchi component for 1000 simulations without (black) and with (red) a Bianchi component added. Vertical solid lines show the means, and vertical dashed the actual errors. The vertical green line shows the true value and expected errors.

alignment detections.

Figure 2.2 shows what these distributions look like for the fit to 1000 simulations in the V band, both in the case where a Bianchi component is added (red) and where it is not (black). Also plotted as vertical lines are the mean and RMS errors on the distributions, as well as in green the true value and expected errors. The small bias in the value of the Bianchi fit is seen in the distance between the vertical red and green lines.

Note that in all these cases, the bias in the fits affects the absolute amplitude (i.e., shear) estimate, but not the significance of the fit, since the ensemble of simulations used to estimate the significance is also affected by such a bias. The expected bias in the amplitude is also much smaller than the error bar. Therefore this does not affect our basic results, namely the particular best-fit model, its location, its approximate amplitude, and its approximate significance relative to chance alignments.

2.4.4 Chance alignments

For a given sky realization, we find the best model as described in Section 2.2.3 and then simply compare the amplitude of that fit against the ensemble of amplitudes for that model relative to Gaussian simulations to estimate the significance. Visual inspection of the *WMAP* sky maps shows no obvious Bianchi component, so any such signal must remain at or below the level of the stochastic component. Chance alignments may therefore either *cancel* a Bianchi-induced signal or give a false positive . The former effect was quantified in section 2.4.1 at $\sim 23\%$, but the latter is more difficult to quantify.

The Family-wise Error Rate (FWER), the expected number of false detections when testing m hypotheses, is $\sim mp$ when p is the probability of one false detection. If 3σ is the detection threshold (implying $p \sim 0.003$) and one tests 100 different hypotheses

(or models), the FWER is then 0.3, meaning one gets a false detection somewhere in the model space one third of the time. Over our grid of Bianchi parameters, the models are not independent (since models close in (x, Ω_0) space will resemble each other closely), so we cannot determine *a priori* what the true frequency of false detections would be, but we can get this from the ensemble of Gaussian simulations.

We perform the full-sky search using the total convolver over the grid of Bianchi models and find the best-fit model for each realization. We find that a false detection due to a chance alignment that has a significance of 99.7% occurs in $\sim 17\%$ of the cases. A better comparison might be to use the χ^2 representing the goodness of the fit. We then compare the statistic (defined above in §2.2.3) $\eta_1 \equiv (\chi_0^2 - \chi_1^2)/\chi_0^2$, namely the relative improvement in the χ^2 when the Bianchi model is subtracted. We find that by this measure, $\sim 10\%$ of the best chance alignments are as good fits as our best-fit model to the *WMAP* data (see §2.5.) Note, however, that these statistics are dependent on the assumed amount of large scale power.

The above numbers simply imply that a detection of a Bianchi model with an amplitude higher than in 99.7% of simulations is more than four times as likely to be real as it is to be a chance alignment, in the absence of all other information.

2.5 Application to the first-year *WMAP* data

Armed with the information gained from the analyses of simulations, we can now examine the fits to the real data.

2.5.1 Fits over model space grid

Using the total convolver to find the best orientation, we fit the grid of Bianchi models to each of the WILC, LILC, and TOH full-sky processed maps. Figure 2.3 shows filled contours over this grid for the LILC. (The results for the WILC and TOH look very similar.) For each point on the grid corresponding to a model of the given (x, Ω_0), the template is fit to the LILC map and the colour indicates the significance estimate of the resulting amplitude, i.e., the fractional number of Gaussian LILC simulations (out of 1000) with lower amplitude. As discussed in §2.2.3, we use a finer grid and better method for determining the best-fit model and thereby select a slightly different model then the analysis in Jaffe et al. (2005). But it is apparent from the right hand column of Figure 2.3 that the significance as a function of the Bianchi parameters x and Ω_0 is flat in the region ± 0.1 in both x and Ω_0 about the maximum.

We find that the most significant fit is found with a right-handed Bianchi template of $x = 0.62$ and $\Omega_0 = 0.5$ when that template is rotated to a position and orientation given by Euler angles (following the total convolver's "zyz" convention about fixed axes) $(\Phi_2, \Theta, \Phi_1) = (42°, 28°, -51°)$. As defined in §2.3.1, the spiral structure of the unrotated model is centred on the south pole (or $-\hat{z}$ axis), so this rotation places the centre of that structure at galactic longitude and latitude of $(l, b) = (222°, -62°)$ and changes it's
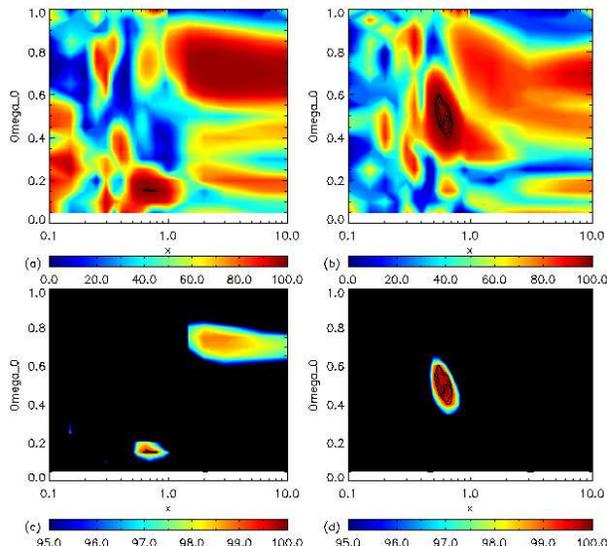
Figure 2.3 Significance as percentage of LILC simulations whose best-fit chance alignment amplitude is lower. Left (right) column: left (right)-handed models. Over plotted contours are at 99.3, 99.5, 99.7, and 99.9%. Two colour scales are used to show the global structure (top) as well as that near the peaks (bottom).

orientation about that location by $\Phi_1 = -51°$. This model fits at an amplitude of $\left(\frac{\sigma}{H}\right)_0 = 4.29 \times 10^{-10}$, which is higher than 99.7% of the 10 000 simulations. This model and the best-fit from previous work Jaffe et al. (2005), at $x = 0.55$, are almost identical.

All models near this best-fit (x, Ω_0) return the same location for the centre of the spiral within 3° but vary the orientation (Euler angle Φ_1) up to 36°. The broad spiral in all of these models is very self-similar under these rotations, so the change is driven largely by the precise locations of the paired hot and cold spots.

Looking at Figure 2.3, one can see that more than one model appears "significant" in the sense of fitting with an amplitude above 99% of the amplitudes found fitting that same model to Gaussian simulations. As discussed above in §2.4.1, this is unsurprising, and we must examine each of these models in more detail.

The full resolution ILC map is shown along with the best-fit Bianchi model on the same scale and the corrected ILC map in Figure 2.4. A summary of all fit results is shown in Tables 2.1 and 2.2. The expected bias in these results is discussed in §2.4.3. The following sections describe the two most interesting models in more detail.

2.5.2   Two best fits

Left-handed model $(x, \Omega_0) = (0.62, 0.15)$

The most significant left-handed model, at 99.4%, is at $(x, \Omega_0) = (0.62, 0.15)$. This model was not found in our earlier work Jaffe et al. (2005), because it is only in a fairly small region of the model space that this fits with any significance, and our previous, coarser grid

Table 2.1. Fitted template amplitudes: right-handed $(x, \Omega_0) = (0.62, 0.5)$.

Map	$\left(\frac{\sigma}{H}\right)_0$ $(\times 10^{-10})$	$\left(\frac{\omega}{H}\right)_0$ $(\times 10^{-10})$	$P\left(|\alpha_{\rm sim}| < |\alpha_{\rm obs}|\right)$ %
WILC	4.33 ± 0.82	9.58	99.8
LILC	4.29 ± 0.82	9.49	99.7
TOH	4.03 ± 0.82	8.92	98.6
K[a]	$2.59(4.13) \pm 0.83$	5.72	16.7(99.1)
KA[a]	$3.50(4.09) \pm 0.83$	7.74	86.9(99.0)
Q[a]	$3.76(4.11) \pm 0.83$	8.31	95.6(99.1)
V[a]	$3.99(4.19) \pm 0.83$	8.82	98.1(99.5)
W[a]	$4.08(4.35) \pm 0.82$	9.03	99.1(99.8)
QVW[a]	$3.84(4.15) \pm 0.83$	8.49	96.8(99.2)
VW[a]	$3.99(4.22) \pm 0.83$	8.84	98.2(99.6)
Q-V[a]	$0.06(0.11) \pm 0.02$	0.13	99.0(100.0)
V-W[a]	$-0.05(-0.08) \pm 0.02$	0.11	93.8(99.0)
Q-W[a]	$0.01(0.04) \pm 0.02$	0.02	25.0(83.2)
Q[b]	4.09 ± 0.10^c	9.04	-
V[b]	4.11 ± 0.10^c	9.09	-
W[b]	4.12 ± 0.11^c	9.12	-

Note. — Amplitudes of the best-fit model derived from various combinations of data and various methods as described in the text. The full sky was used in the analysis of the WILC, LILC, TOH, and Gibbs samples, while the Kp0 mask was imposed for the remaining maps.

[a]Simultaneous fits with foreground components. In parentheses are the values using the SFD dust template instead of the FDS, and the Dickinson et al. (2003) Hα instead of Finkbeiner (2003).

[b]Average over 1000 Gibbs samples.

[c]Errors are RMS variation over Gibbs samples.

Table 2.2.   Fitted template amplitudes: left-handed $(x, \Omega_0) = (0.62, 0.15)$.

Map	$\left(\frac{\sigma}{H}\right)_0$ $(\times 10^{-10})$	$\left(\frac{\omega}{H}\right)_0$ $(\times 10^{-10})$	$P\left(|\alpha_{\mathrm{sim}}| < |\alpha_{\mathrm{obs}}|\right)$ %
WILC	2.39 ± 0.47	22.31	97.8
LILC	2.49 ± 0.47	23.29	99.4
TOH	2.45 ± 0.47	22.94	99.0
K[a]	$2.33(3.31) \pm 0.50$	21.76	96.3(99.9)
KA[a]	$2.24(2.63) \pm 0.50$	20.93	94.8(99.3)
Q[a]	$2.29(2.50) \pm 0.50$	21.42	96.0(98.9)
V[a]	$2.33(2.44) \pm 0.50$	21.81	96.7(98.6)
W[a]	$2.32(2.46) \pm 0.49$	21.69	96.3(98.4)
QVW[a]	$2.30(2.48) \pm 0.50$	21.46	96.1(98.8)
VW[a]	$2.34(2.44) \pm 0.50$	21.85	96.7(98.6)
Q-V[a]	$0.03(0.02) \pm 0.02$	0.27	78.6(63.1)
V-W[a]	$-0.06(-0.10) \pm 0.02$	0.55	96.4(99.9)
Q-W[a]	$-0.03(-0.08) \pm 0.02$	0.29	73.1(99.3)
Q[b]	2.10 ± 0.11^c	19.67	-
V[b]	2.08 ± 0.11^c	19.46	-
W[b]	2.09 ± 0.09^c	19.53	-

Note. —   See Table 2.1 for notes.
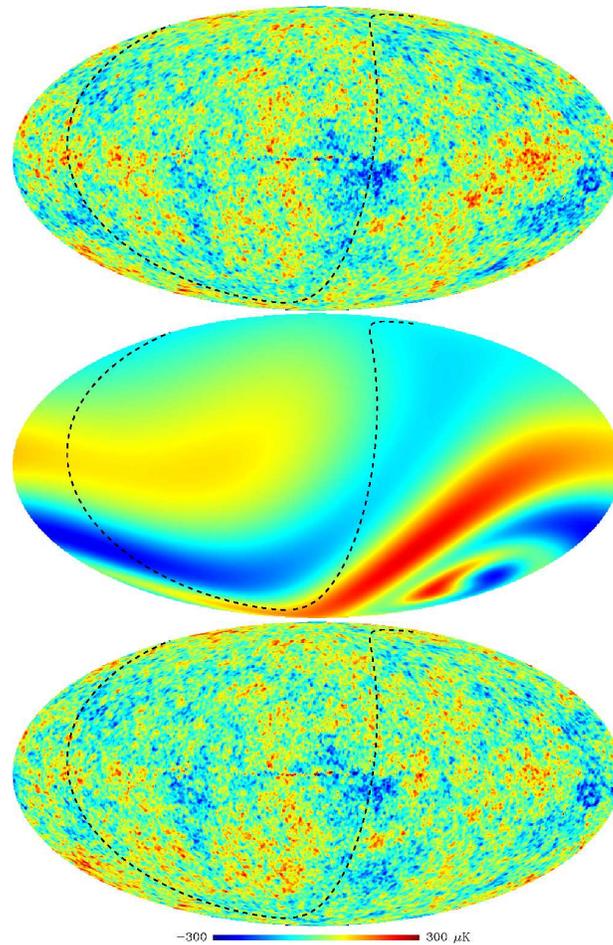
Figure 2.4 (*top*) *WMAP* Internal Linear Combination map. (*middle*) Best-fit Bianchi VII$_h$ template (enhanced by a factor of four to bring out structure). (*bottom*) Difference between WILC and best-fit Bianchi template; the "Bianchi-corrected" ILC map. Overplotted on each as a dotted line is the equator in the reference frame that maximises the power asymmetry as described in section 2.6.3.
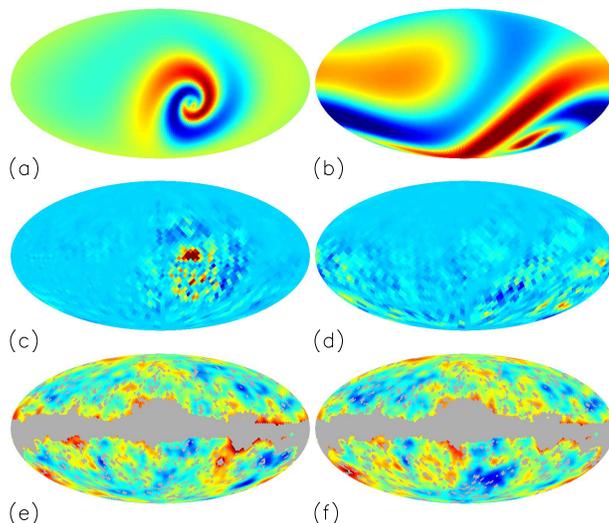
Figure 2.5 Two significant models: the left column shows the best-fit left handed model with $(x, \Omega_0) = (0.62, 0.15)$, while the right column is the best-fit overall model, right-handed with $(x, \Omega_0) = (0.62, 0.5)$. The top row shows the template amplified by a factor of three to bring out the structure. The middle row shows the corresponding cross correlation map (see §2.2.4) scaled from -1% to 2%. The bottom row shows the "corrected" *WMAP* Q+V+W map scaled from -150 to $150\mu K$. The grey region is the excluded region of the Kp0 mask.

effectively straddled the peak in Ω_0. The best-fit location for this model puts the centre of the structure at $(l, b) = (320, -20)$, which is closer to the Galactic Centre region than the best-fit right-handed model, raising the question of how much it is driven by foreground residuals.

Cut-sky fits give fit amplitudes for this component that are 8% lower and significances of $\sim 96\%$ in most cases. Furthermore, the Galactic Centre region tends to draw the template in simulations; the best-fit location among the simulated LILC maps for this model is twice as likely to be found in the area around $(0°, -20°)$ as should be expected from a uniform distribution. The only thing that all of the LILC simulations have in common is foregrounds, so this is an indication that there is some residual there that is a weak attractor. One possibility is the "free-free haze" described by Finkbeiner (2004a) (see also Patanchon et al. 2004), though this haze does not match up well with the template structure, the two show little cross-correlation, and inclusion of Finkbeiner's haze template in the simultaneous fitting does not alter the fit amplitude of the Bianchi model.

In Figure 2.5, it looks like the fit should be largely driven by the cold region below the Galactic Centre. The cross-correlation maps described in §2.2.4 do show correlation there but also indicate that the fit is largely driven by a very strong signal in the Galactic Plane. Figure 2.5 shows these maps for both this model and the best-fit right handed model. Where the right handed model shows relatively uniform correlation over the hemisphere about the best-fit axis, this model shows a rather concentrated region including a very strong driver on the Galactic Plane.

The Gibbs samples throw further doubt on this model. Among the 1000 Gibbs samples in each of Q, V, and W bands, this model fits at the same approximate location as for the LILC map less than half of the time. Where the location was the same, the amplitude of the best-fit is significantly lower for the ensemble of Gibbs sampled maps, which drop over 15% in amplitude to a mean of 2.1×10^{-10}, indicating that some of the structure in the data that drives the fits is not consistent with the posterior CMB distribution as determined by the Gibbs sampling technique. Furthermore, this model is almost as likely to fit near the location of the best-fit right-handed model instead of near the Galactic Centre. This is largely driven by the cold spot.

In summary, this model is quantitatively less significant than the best-fit right-handed model based on the cut sky and Gibbs sample fit values. Furthermore, the morphology indicates that foreground residuals drive the full-sky fit.

Right-handed model $(x, \Omega_0) = (0.62, 0.5)$

Figure 2.3 shows that the best-fit model is this right-handed model.

The amplitude of the best-fit Bianchi component varies somewhat across the different frequencies, in all cases lower than the full-sky amplitude fit with the LILC. As discussed in §2.4.3, this is likely due to small foreground residuals, but does not mean that the detection is a false positive; the same effect occurs in simulations that include a Bianchi component. The amplitude in the W band, where the least foreground residuals are expected, is still higher than $\sim 99\%$ of simulations. The K and Ka band fits are significantly lower when the FDS dust and Finkbeiner Hα templates are used, but are consistent with the other bands when the SFD dust and Dickinson Hα templates are used instead. It is known that foreground subtraction is a problem even at high latitudes in the K and Ka bands, and these residuals are clearly affecting the low frequency fits. Looking at the residuals of the two fits shows that the difference may be driven by a small region around $(l, b) = (300°, -15°)$ where the dust templates differ strongly. The higher frequency fits, however, are more consistent. The difference maps, e.g., Q-V, should contain no CMB component but only foreground residuals and noise. The fact that the Bianchi component amplitude found from these maps is less than 2% of the coadded map amplitude is an indication that such residuals are not contributing significantly to the fit.

The results of fitting the Gibbs sampled maps show that for this model, the amplitude is quite stable over the ensemble of Gibbs samples, with, e.g., a mean of $4.12 \pm 0.1 \times 10^{-10}$ in the W band compared to 4.08×10^{-10} for the cut-sky fit to the raw data. Since the Gibbs samples represent the posterior CMB distribution, taking into account foregrounds and iterating over the power spectrum, these results are a strong indication that the fit is due primarily to CMB signal.

Figure 2.5 shows the cross-correlation map as described in §2.2.4, which give a visual indication of what regions drive the fit. Unlike the left handed model (top), which shows one concentrated region in the Galactic Plane to be driving the fit, this model correlates over more than half the sky at moderate levels. One can see that the cold spot does partly drive the fit, but not particular region can be said to dominate. Fits to the combined QVW

and VW data where the cold spot is excluded (in a $10°$ radius around $(l, b) = (209, -57)$) have comparable amplitudes (6% lower) than fits where the region is included. Further Gibbs samples were also computed while masking this region. Full sky searches using these samples show that fewer than 20% return positions more than $10°$ from the original location, and amplitudes which are on average 15% lower (which is within the calculated error bar.) These results confirm that the cold spot does affect but does not exclusively drive the fit amplitude.

2.5.3　Location and orientation accuracy

As mentioned above, where a Bianchi component was added to simulations at a known location, the full-sky search with the total convolver returned the correct position within $5°$ in $\sim 80\%$ of realizations. The uncertainty in the location is due to the CMB fluctuations, which are quite comparable to the Bianchi component at the amplitude detected.

To determine how the amplitude changes with the position and orientation of the template compared to the data, we take the best-fit Bianchi model and fit it to the LILC on a grid of fixed positions within $20°$ of the best-fit position. Results are shown in Figure 2.6. The orientation is not very sensitive in this model, whose spiral structure is self similar under rotations about its symmetry axis; only the precise positions of the hot and cold spots affect the variation with orientation angle. The amplitude drops by 1% when the orientation is $4°$ off. The location of the symmetry axis is a bit more sensitive, where the amplitude drops by 3% at $2°$. The fact that the total convolver at this resolution uses steps of $2.8°$ means that its best-fit amplitude can be several percent off of the actual maximum. All the fits to the simulations as well as the data are subject to this same uncertainty. If we assume the worst, that the LILC amplitude was found at its true maximum (i.e., the true axis of symmetry happened to lie exactly on the centre of one of the total convolver's bins) and the simulations all at $1.4°$ away from their true maxima (i.e., the axis exactly between bins) and have true values correspondingly higher, the comparative significance could then be overestimated by only 0.5%. The likely effect is of course much smaller.

Using the LILC map at higher resolution, $\ell_{max} = 128$, gives an accuracy in the total convolver of $\pi/\ell_{max} = 1.4°$. The position returned is identical, with only the orientation one step of $1.4°$ different.

The above applies to the best-fit model at $(x, \Omega_0) = (0.62, 0.5)$, but other models have structure at different angular scales. In particular, for the region of small x and Ω_0, where a tightly wound spiral is even more tightly focused in one hemisphere, the fit amplitudes are far more dependent on the exact position. Because the total convolver resolution is π/ℓ_{max}, our analysis is not as sensitive for this region of model space as it would be for a higher resolution analysis. In these cases, the difference of a few degrees can mean a large difference in amplitude. Simulations show that, though the location returned is the closest bin to the true location, the amplitude of a model $(x, \Omega_0) = (0.1, 0.1)$ is underestimated by $\sim 20\%$ on average due to the limited resolution. Increasing the resolution of the analysis to HEALPix $N_{side} = 64$ increases the mean and brings it closer to the correct value, but it is still underestimated. (Higher resolution analysis with the total convolver is not feasible
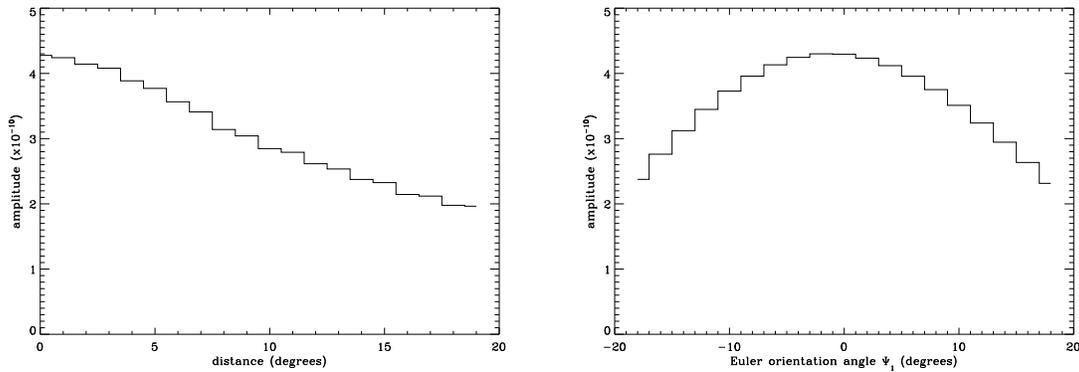
Figure 2.6 *(left)* Average fit amplitude as the location of the template varies from the best-fit position. For these fits, the orientation of the template is unchanged. *(right)* Fit amplitude as the orientation of the template, i.e., the Euler angle γ, is changed. For these fits, the location in longitude and latitude is unchanged. Note that this grid finds a preferred orientation $2°$ from that found by the total convolver (due to the slightly different grids used.)

due to the memory and CPU requirements.) In the region of model space where $x > 0.25$ and $\Omega_0 > 0.25$, this effect drops to less than a few percent.

A more detailed look at these models at increased resolution ($N_{\text{side}} = 64$) shows no evidence that the lower resolution analysis missed a significant detection. But the limits placed on shear and rotation are less stringent than they would be were a higher resolution analysis feasible.

2.5.4   DMR fit

Our best-fit amplitude is below the upper limit DMR could place on the shear. Using this model, a fit to the DMR data gives $\left(\frac{\sigma}{H}\right)_0 = 3.38 \pm .98 \times 10^{-10}$, which is within our best-fit error bar for the *WMAP* data, but which is not distinguishable from a chance alignment for DMR. Kogut et al. (1997) report a distribution of Γ values for chance alignments up to 4.5. Our fit value and error give $\Gamma = 3.4$, and though this value comes from different methods and assumptions, it is roughly comparable.

2.5.5   Sensitivity to assumed power spectrum

As mentioned in §2.3.4, assumptions about the cosmological parameters go into this analysis from the beginning with the choice of the signal covariance matrix. In effect, we are assuming that the CMB signal consists of an anisotropic Bianchi-induced component plus a statistically isotropic, Gaussian random field described completely by its power spectrum, which is taken to be the *WMAP* best-fit theoretical power-law spectrum. As we are searching for evidence of a model that affects the power spectrum at large scales and that is inconsistent with inflation, this approach obviously lacks consistency.
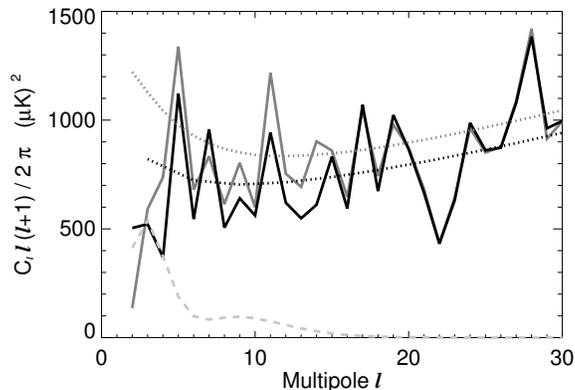
Figure 2.7 Comparison of power spectra. The gray and black *solid* lines show the power spectrum estimated from the coadded V+W map before and after correcting for the Bianchi template, respectively. The *dotted* gray and black lines shows the theoretical best-fit power spectra from the *WMAP*-team analysis and Hansen et al. (2004a) respectively. The latter is a fit to the northern hemisphere data alone. The *dashed* grey line is the power in the Bianchi template alone.

We have verified, however, that changing the assumed parameters and using, e.g., a flat $Q = 18\,\mu\text{K}$ power-law spectrum, or a completely implausible spectrum, has little effect (less than 3%) on the resulting best-fit amplitude and position for the Bianchi component. In fact, the power spectrum affects only the estimated significance of the result, as that significance is dependent on the expected level of large scale CMB structure that drives chance alignments. As shown in Figure 2.7, correction for this Bianchi model lowers the large scale power. Our significance estimates are based on simulations generated assuming a higher level of large scale power, so the significance of the detection would increase when compared to an ensemble consistent with the corrected power spectrum.

2.6   Implications

There are several interesting results based on *WMAP* first year data that are inconsistent with the assumptions of isotropy and Gaussianity and that are immediately relevant to this study. De Oliveira-Costa et al. (2004), Land & Magueijo (2005), and Copi et al. (2006a) (and sources therein) examine the low-ℓ multipoles of the foreground-cleaned data and find that, in addition to the anomalously low quadrupole amplitude, the preferred axes of the quadrupole and octopole are anomalously well aligned in the direction of $(l, b) = (-110°, 60°)$. Eriksen et al. (2004b) and Hansen et al. (2004b) show that there is a significant difference in large scale power between two hemispheres (roughly aligned with the ecliptic) at the 98-99% level, with significantly more power in the South. Vielva et al. (2004) and Cruz et al. (2005) detect non-Gaussianity in the *WMAP* combined Q-V-W map using spherical wavelets; they find significant kurtosis in the wavelet coefficients at a scale of 10° and identify a cold spot at $(l, b) = (209°, -57°)$ as the probable source. Our

choice of models was partly motivated by the morphology of these anomalies, and indeed, subtracting for our best-fit Bianchi template corrects them.

2.6.1   Quadrupole amplitude

The quadrupole amplitude has been considered anomalously low since *COBE*; see de Oliveira-Costa et al. (2004b) and references therein. As pointed out by Jaffe et al. (2005), the correction for this Bianchi component raises the low quadrupole amplitude to a value more consistent with the theoretical power spectrum. This result is unchanged with the best-fit model of this work, since the models are almost the same. Should this be considered "fine tuning"?

We can simulate the situation by taking as the the primordial quadrupole the *WMAP* quadrupole (as derived by Bielewicz et al. (2004)) minus the quadrupole of our best-fit Bianchi model. If we then add the Bianchi quadrupole at random orientations, we can see how likely it is that the resulting total quadrupole be as low as the observed *WMAP* quadrupole. We find that the likelihood is $\sim 5\%$. This implies that the level of "fine tuning" required to end up with the low observed quadrupole is not exceptional.

We further take a set of 1000 simulated Gaussian CMB skies, with and without a Bianchi component, and fit our best-fit Bianchi model to them. The "corrected" quadrupole is on average $\sim 5\%$ lower than the original, which is to be expected considering that the fit is a least squares solution. In contrast, using the real LILC data, the correction has the effect of raising the quadrupole. This happens in over $\sim 20\%$ of the simulations, so while this is not the average behaviour, it is not extraordinary.

2.6.2   Low-ℓ alignment and planarity

De Oliveira-Costa et al. (2004), Land & Magueijo (2005), and Copi et al. (2006a) discuss the statistically anomalous alignment of the quadrupole and octopole in the *WMAP* data. The preferred axes of the $\ell = 2$ and $\ell = 3$ multipoles are only $7°$ apart (roughly in the direction of $(l, b) = (-110°, 60°)$), which is anomalous at the 99.3% level compared to simulations. After subtracting the best-fit Bianchi template, these axes lie $74°$ apart, consistent (at 27%) with the statistically isotropic simulations (see Figure 2.8).

The planarity of the low-ℓ multipoles has also been considered somewhat anomalous (see de Oliveira-Costa et al. (2004b) and Land & Magueijo (2005) for a discussion.) The t-statistic defined by de Oliveira-Costa et al. (2004b) provides a measure of this planarity. Again, subtracting the Bianchi template lowers the significance of the the low-ℓ multipoles. The planarity of the octopole in particular drops from a significance of $\sim 90\%$ (depending on whether the WILC or LILC is used) to $\sim 50\%$. Figure 2.8 (b) and (d) shows how the planarity of the octopole is disrupted. This will also impact the results of multipole vector analyses such as that of Copi et al. (2006a).

(a) ILC l=2   (b) ILC l=3

(c) ILC−B l=2   (d) ILC−B l=3
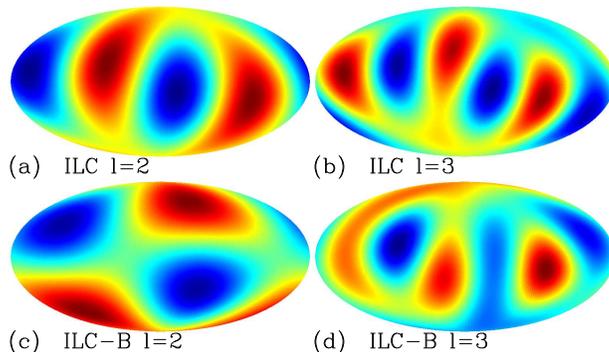
Figure 2.8 Low-ℓ multipoles of the WILC corrected for the Bianchi component (*bottom row*) or uncorrected (*top row.*)

2.6.3   Large-scale power asymmetry

Eriksen et al. (2004b) and Hansen et al. (2004b) reported that the large-scale power ($\ell \lesssim 40$) in the *WMAP* data is anisotropically distributed over two opposing hemispheres (in the reference frame where the z-axis points toward $(l, b) = (57°, 10°)$; see Figure 2.9), with a significance of 3σ compared with simulations. Repeating the analysis and adopting the Kp2 sky coverage, we compare the corrected V+W *WMAP* map with 2048 simulations. We find that $\sim 14\%$ of the simulations have a larger maximum power asymmetry ratio than the Bianchi-corrected map, whereas only 0.7% have a larger ratio than the uncorrected data (see Figure 2.9). It is apparent that the maximum power ratio between any two hemispheres is significantly suppressed after subtracting the Bianchi template – no asymmetry axis is found at any statistically significant level. It is apparent from that figure, however, that some residual power asymmetry remains. This comes largely from the range $20 < l < 40$ where the Bianchi template has little power, indicating that a model with more small scale structure may be needed.

2.6.4   Wavelet kurtosis

Vielva et al. (2004) and Cruz et al. (2005) used a wavelet technique to detect an excess of kurtosis in the wavelet coefficients and isolate an unusually cold spot ($\sim 3\sigma$ significance relative to Gaussian simulations) at Galactic coordinates $(l, b) = (209°, -57°)$. Referring again to Figure 2.4, we see that a cold spot is indeed present at the right location, in the form of the centre of the spiral.

    We therefore also repeat the analysis of Vielva et al. (2004), and compute the kurtosis of the wavelet coefficients as a function of scale from both the WILC and the corresponding Bianchi-subtracted map. A $|b| < 20°$ galactic cut is imposed in this case, for computational convenience.

    The results from this exercise are reported in Figure 2.10 After subtracting the Bianchi template, the significance of the southern hemisphere anomaly is greatly reduced, and no new non-Gaussian features have been introduced. See also Addendum.
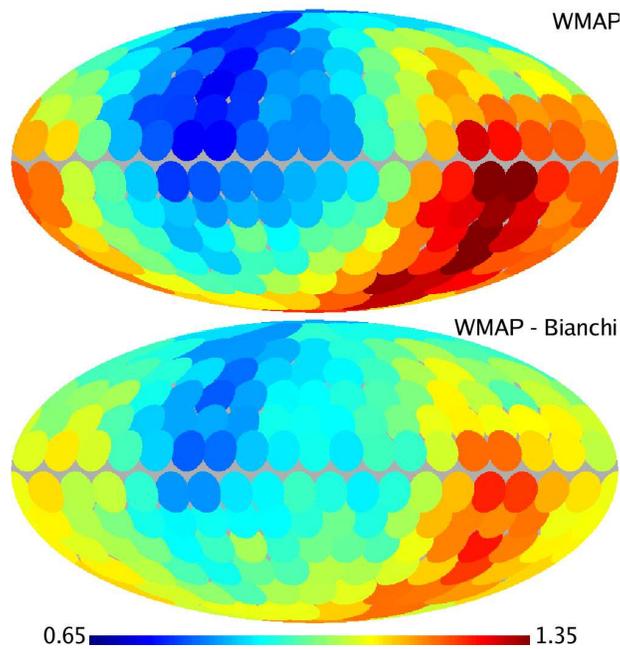
Figure 2.9 Power ratio between hemispheres in *WMAP* ILC (*top*) and in *WMAP* ILC corrected for best-fit Bianchi component (*bottom.*)
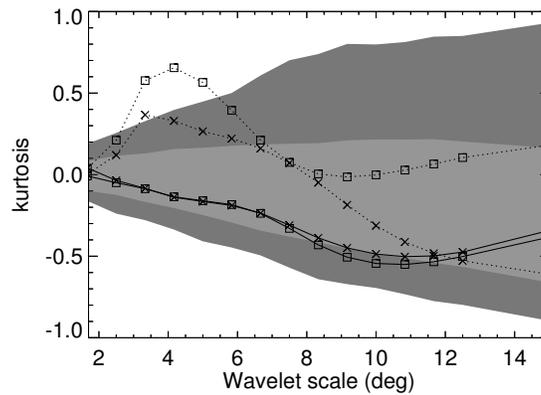
Figure 2.10 Kurtosis in wavelet coefficients. The *boxes* and *crosses* show the kurtosis before and after subtracting the Bianchi template, respectively, computed from the southern (*dotted* line) and northern (*solid* line) Galactic hemisphere.

2.7   Discussion and Conclusions

We have considered a fast and efficient method for fitting a template to the full sky in harmonic space and finding the best-fit location and orientation. The total convolver algorithm evaluates the correlation between the sky and the template at every possible relative orientation using fast Fourier transforms. With this algorithm, the search for the best-fit becomes two orders of magnitude faster than the corresponding search performed one rotation at a time. This method, along with pixel-space simultaneous foreground fitting, provides a powerful tool for testing any deterministic model for anisotropy in the CMB. Simulations generated by the LILC pipeline allow us to quantify the bias, to investigate the effects of foreground contaminants, and to show how well each of these methods detects a known input.

We have applied this method to the first year *WMAP* data to search for evidence of shear and vorticity using templates derived for Bianchi type VII$_h$ universes. We find a surprisingly significant correlation between the *WMAP* data and a right-handed Bianchi model with $x = 0.62$, $\Omega_0 = 0.5$, and shear of $\left(\frac{\sigma}{H}\right)_0 = 4.3^{\pm.8} \times 10^{-10}$, implying a vorticity of $\left(\frac{\omega}{H}\right)_0 = 9.5 \times 10^{-10}$. The centre of the spiral structure lies at approximately $(l, b) = (222°, -62°)$. Simulations show that this amplitude is likely to be biased by $\sim 7\%$, implying a true amplitude closer to 4.0×10^{-10}. Incomplete sky fits, simultaneous foreground fitting, and fits to a set of Gibbs samples are all consistent with this amplitude and indicate that confusion with Galactic emission is unlikely to contribute significantly to this detection.

Correcting the *WMAP* data for the effect of the best-fit model solves several problems seen in the data. The corrected maps show significantly reduced power asymmetry between any two hemispheres. The correction also eliminates the non-Gaussian kurtosis in the wavelet coefficients detected by Vielva et al. (2004) and Cruz et al. (2005), raises the low measured quadrupole by a factor of two, and disrupts the planarity of the octopole and its anomalous alignment with the quadrupole. In short, the data appear far more Gaussian and isotropic after correction.

The original analyses by Kogut et al. (1997) and Bunn et al. (1996) were limited by the signal to noise level in the DMR instrument. Our best fit result is just under their upper limit but still significant due to *WMAP*'s greatly improved signal to noise. Furthermore, the Kogut analysis searched a coarse ($\sim 10°$) grid of possible locations and orientations, while with the total convolver, we can efficiently search a finer grid.

How likely is it that our best-fit model is a true detection rather than a chance alignment? Considering the best-fit model by itself and comparing its fit amplitude to simulations, it is higher than 99.7% of simulations. But the simulations also show that $10 - 20\%$ of Gaussian, statistically isotropic skies will have one of the Bianchi models appear as significant. Considering the fact that the sky is approximately Gaussian and isotropic, one would not expect to find a more definitive detection based on template fitting alone. But the distribution of chance alignments in the simulations is sensitive to the amount of large scale power assumed, and that is lowered by the Bianchi correction to the *WMAP* data. Furthermore, the cumulative probability that a chance alignment not only fits at the level of our best-fit model but also has the effect of resolving the several anomalies in the data

must also be considered in any qualitative judgement of the significance of this result.

Further improvement to the data will not refine these measures significantly, because at the *WMAP* sensitivity level, the analysis is already very close to the expected distribution of chance alignments in the absence of noise. Improved foreground subtraction will, however, remove some of the possible confusion and bias. But neither higher resolution nor higher signal to noise data should change this result nor be able to provide additional information concerning the question of whether the fit is a real detection of vorticity and shear. To answer that question will require additional verifiable predictions for the effects of vorticity and shear on other observables.

But in the context of the anomalies that this hypothesis can explain, the possible detection is certainly provocative. The most important result of this analysis is that a model with vorticity and shear can explain the observed asymmetry in the CMB anisotropies and the non-Gaussian cold spot. Note that this asymmetry exists only in the $\Omega_0 < 1$ versions of these Bianchi models. Significant evidence currently indicates that Ω_0 is very close to one, so our best fit model cannot be considered physically realistic. However, as mentioned in §2.3.1, Barrow et al. did not include any dark energy component. Furthermore, the Bianchi model does not include a mechanism to generate structure at the surface of last scattering. A self-consistent theory is required that can explain the small scale fluctuations, and in particular the acoustic peaks, in the context of a Universe with shear and vorticity. But from a pragmatic point of view, one can conclude that, regardless of the viability of the particular Bianchi model, this result gives a measure of the significant deviation from isotropy in the data.

We consider this result to be further motivation for considering ideas outside of the so-called Concordance Model of cosmology. There are anomalies in the data that are inconsistent with the theory of a Gaussian, statistically isotropic universe, and Bianchi models are only one such anisotropic model that merits investigation. We have demonstrated a method of template fitting that can be applied to test any model that makes a deterministic prediction for an anisotropy pattern in the CMB. The best-fit Bianchi model provides a template temperature pattern that can explain the observed anomalies in the data and that describes the morphology theorists may need to reproduce in considering alternatives to the standard cosmological model.

2.8   Addendum

A more sophisticated statistical analysis of the wavelet coefficients has also been done by Cayón et al. (2006). They apply a Higher Criticism (HC) statistic using a significance test in the context of determining whether or not the data reject the null hypothesis. This method is applied to the sky both corrected and uncorrected for the best-fit Bianchi template described above. The results indicate that correcting for the template at its best-fit location and amplitude reduces the non-Gaussian signal to insignificance. Figure 2.11 shows the HC statistic for each wavelet scale compared to the expected distribution for Gaussian simulations. Where the uncorrected sky (stars) shows a strong non-Gaussian
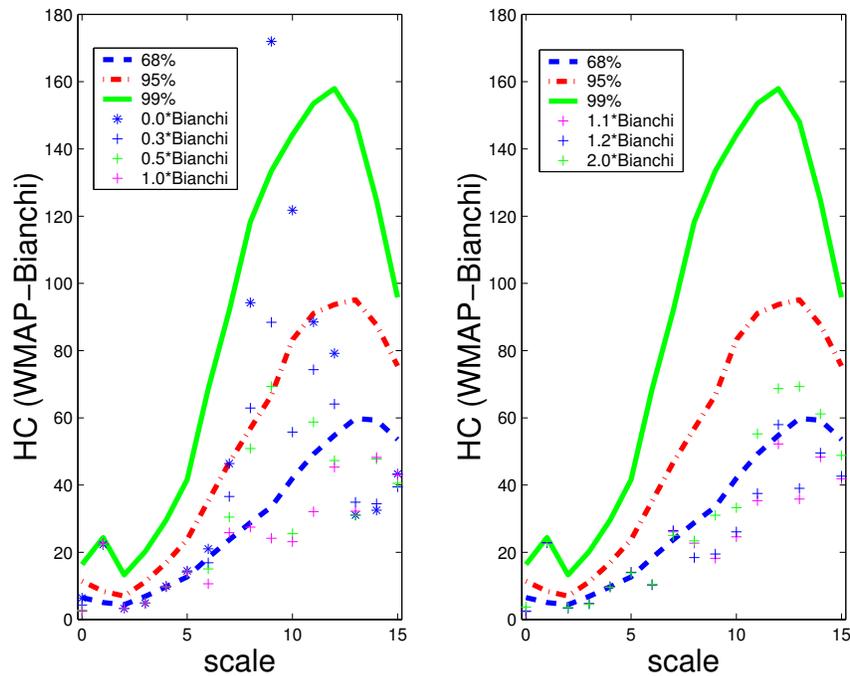
Figure 2.11 Higher Criticism statistic for the corrected sky. Different colors show the results for the Bianchi template at different amplitudes relative to the best-fit amplitude (i.e., 0 is uncorrected, 1 is corrected at the best-fit amplitude). The lines show the significances from comparision with Gaussian simulations. From Cayón et al. (2006).

signal, this signal is not present in the sky that has been corrected for the Bianchi template. Various amplitudes of that template, relative to the best-fit amplitude, are also shown for comparison.

A map of the HC statistic for each pixel on the sky shows that the signal is still visible, though not statistically significant, when the template is subtracted at its best-fit amplitude. Cayón et al. find that a 20% increase in the template amplitude completely erases it. This is within the error bar of the least-squares fit, but it is also possibly an indication that a slightly different model, with more power on smaller scales, may better match the data.

Chapter 3

On the Viability of Bianchi Type VII$_h$ Models with Dark Energy

Note: this chapter was published as Jaffe et al. (2006b). The derivation in § 3.2.1 was done by S. Hervik, who also wrote the text of that section.

3.1 Introduction

While cosmology appears to be converging on a "concordance model" that describes the universe as inflationary and isotropic, there remain unexplained anomalies in the CMB data, and other models are not yet ruled out. The *WMAP* data provide some of the most accurate measurements yet of the cosmic microwave background and contribute to high accuracy determinations of cosmological parameters (Bennett et al., 2003b; Spergel et al., 2003). However, there are several studies that show that at large scales, the CMB is not statistically isotropic and Gaussian, as predicted by inflation theory.

De Oliveira-Costa et al. (2004), Land & Magueijo (2005), and Copi et al. (2006a) and sources therein examine the low-ℓ multipoles of the foreground-cleaned data. These studies find that, in addition to the anomalously low quadrupole amplitude, the preferred axes of the quadrupole and octopole are anomalously well aligned in the direction of $(l, b) = (-110°, 60°)$ and that both of these modes are unusually planar. Furthermore, Eriksen et al. (2004b) and Hansen et al. (2004b) show that there is a significant difference in large scale power between the two hemispheres roughly aligned with the ecliptic, almost perpendicular to the quadrupole and octopole axes. This asymmetry is significant at the 98-99% level, with significantly more power in the South. And lastly, Vielva et al. (2004) and Cruz et al. (2005) detect non-Gaussianity in the *WMAP* combined Q-V-W map using spherical wavelets. They find significant kurtosis in the wavelet coefficients at a scale of $10°$ and identify a cold spot at $(l, b) = (209°, -57°)$ as the source (see also Cayón et al. 2005.)

In Jaffe et al. (2005, 2006a), we examined a particular set of anisotropic cosmological models: the Bianchi type VII$_h$ class of spatially homogeneous generalisations of Friedmann

universes that include small vorticity (universal rotation) and shear (differential expansion) components. These models were derived by Barrow et al. (1985) and include no cosmological constant component in the total energy density. Surprisingly, we found evidence that one of these models correlates with the CMB sky and that correcting for it leaves a sky which is statistically isotropic by each of the above measures. The best-fit model found by that study required $\Omega_{\mathrm{m}} = 0.5$, implying, $\Omega_k = 0.5$, i.e., a significantly negatively curved universe. Though the large-scale power overall is lowered after correction for this component, the quadrupole is more than a factor of two larger than uncorrected and now within the cosmic variance errors about the WMAP best fit theoretical spectrum. The quadrupole and octopole, originally $7°$ apart with a likelihood of only 0.3% (compared to simulations), are now $74°$ apart, with a likelihood of 27%. The planarity of the low-ℓ modes, anomalous at the 90% to 99% level, become average. The large-scale power spectra in the two hemispheres are in agreement after correction, the significance of the asymmetry down from 99% to 86% after correction. Likewise, the significance of the wavelet kurtosis drops from over 99% to below 68%. See Jaffe et al. (2005, 2006a) and Cayón et al. (2006) for details, and the former specifically for the power spectra and maps after the best-fit Bianchi structure has been subtracted.

Land & Magueijo (2006) subsequently considered flat Bianchi models and sought explicitly to resolve the problems of the low quadrupole and the low-ℓ alignments using a statistic constructed to achieve that purpose. Their analysis does indeed find a model that fixes these anomalies, but it remains statistically insignificant as a detection. Our original result has the benefit of being a detection that is based entirely on a simple least-squares fit of the Bianchi template to the data, yet we find that it serendipitously resolves several anomalies without that requirement having been built-in to the search algorithm. The importance of our result, therefore, lies in the fact that it resolves the problems of the low quadrupole, low-ℓ alignments, power asymmetry, and non-Gaussian cold spot.

However, both analyses neglect the fact that the existing Bianchi solutions include no dark energy component. Furthermore, the best-fit results require an energy content that is consistent neither with other astronomical observations nor with the CMB itself. In this work, we present a modification of the Barrow et al. Bianchi type VII$_h$ solution so that it includes a cosmological constant term in the evolution. We discuss the impact of that term on the structure of the resulting CMB anisotropy pattern, particularly the degeneracy that is introduced in the model space by the addition of Ω_Λ. We then discuss the viability of Λ models that are morphologically identical to our best-fit $\Omega_\Lambda = 0$ models by considering the constraints imposed by different measurements of the cosmological parameters.

3.2   Bianchi Models with $\Lambda > 0$

3.2.1   Solution

First we shall generalise the equations in Barrow et al. (1985) to include a cosmological constant Λ. The basic assumption is that the type VII$_h$ universe is close to FRW.

We start with the equations of motion using expansion-normalised variables (see, e.g., Wainwright & Ellis, 1997). The equations describing the evolution of Bianchi type VII_h universes with a tilting perfect fluid can be found in Coley & Hervik (2005) and Hervik et al. (2005a). In our case, we consider a tilting perfect fluid with $\gamma = 1$ (dust) and a cosmological constant. The equations in Coley & Hervik (2005) can easily be generalised to include a cosmological constant by adding a matter term $(\Omega_\Lambda, \gamma_\Lambda = 0)$ wherever there is a matter term (Ω, γ).

Furthermore, we assume that the universe is close to a FRW at all times and that the tilt velocity of the dust fluid is small; hence we assume:

$$\Sigma^2 \ll 1, \ |\mathbf{N}_\times| \ll 1, \ v_1 \ll 1, \ |\mathbf{v}| \ll 1. \tag{3.1}$$

Here, $\Sigma^2 = \sigma^2/(3H^2)$ is the expansion-normalised shear scalar, $\mathbf{N}_\times$ is a complex curvature variable, and v_1, $\mathbf{v} = v_2 + iv_3$ are the tilt components (see Coley & Hervik 2005 for details). The curvature variables A and $\bar{N}$ need not be small, and they are related to the group parameter h as follows:

$$A^2 \approx 3h\bar{N}^2.$$

We also assume that the parameter h is not too small: $h \geq \mathcal{O}(1)$. (As $h \to 0$, the assumptions made for this derivation break down and a qualitatively different solution for type VII_0 is needed. See Hervik et al. 2005b.)

Given the above assumptions, the deceleration parameter is

$$q \approx \frac{1}{2}\Omega_\text{m} - \Omega_\Lambda, \tag{3.2}$$

and the equation of motion for the Hubble scalar is

$$H' = -(q+1)H, \tag{3.3}$$

where the prime indicates the derivative with respect to the dimensionless time T defined by $dt/dT = 1/H$, and t is the cosmological time.

At the lowest order, the tilt components that are related to the vorticity are the components $v_2 + iv_3 \equiv \mathbf{v}$ (in the notation of Coley & Hervik, 2005). These induce non-zero shear components $\Sigma_{12} + i\Sigma_{13} \equiv \mathbf{\Sigma}_1$ via the linear constraint equation, which in the FRW limit reduces to:

$$\mathbf{\Sigma}_1 \bar{N}(i - 3\sqrt{h}) + \Omega_\text{m}\mathbf{v} = 0. \tag{3.4}$$

Moreover, close to FRW the curvature $K \equiv A^2$ and matter densities Ω_m and Ω_Λ are related via the Friedmann equation

$$1 = \Omega_\Lambda + K + \Omega_\text{m}. \tag{3.5}$$

We also define x to be

$$x \equiv \left(\frac{h}{K_0}\right)^{\frac{1}{2}} = \left(\frac{h}{1 - \Omega_{\Lambda 0} - \Omega_{\text{m}0}}\right)^{\frac{1}{2}}. \tag{3.6}$$

Eq. (3.4) now reduces to Barrow et al.'s eqs.(4.6) and (4.7) by dropping the subscript 0 and replacing Ω with Ω_m. The equations of motion can now be solved.

By choosing $T_0 = 0$, $x = e^{\alpha_0} H_0$, we can relate z to T and the time variable τ in Barrow et al. (1985) to the redshift z:

$$1 + z = e^{-T}, \tag{3.7}$$

$$\frac{dT}{d\tau} = H e^T e^{\alpha_0}. \tag{3.8}$$

The Hubble scalar, K, Ω_m, and Ω_Λ can then be written

$$H(z) = H_0 \mathcal{H}(z), \tag{3.9}$$

$$\mathcal{H}(z) = \left[\Omega_{\Lambda 0} + K_0 (1+z)^2 + \Omega_{m0}(1+z)^3 \right]^{\frac{1}{2}}, \tag{3.10}$$

$$K = \frac{K_0 (1+z)^2}{\mathcal{H}^2(z)}, \tag{3.11}$$

$$\Omega_m = \frac{\Omega_{m0}(1+z)^3}{\mathcal{H}^2(z)}, \tag{3.12}$$

$$\Omega_\Lambda = \frac{\Omega_{\Lambda 0}}{\mathcal{H}^2(z)}, \tag{3.13}$$

where $\Omega_{\Lambda 0}, K_0, \Omega_{m0}$ are the present values of the expansion-normalised Λ, curvature and ordinary matter (dust) and obey the constraint:

$$1 = \Omega_{\Lambda 0} + K_0 + \Omega_{m0}.$$

We also note that $K = \Omega_k > 0$ implies negative curvature. The shear components Σ_{12} and Σ_{13}, and the remaining curvature variable are given by

$$\Sigma_{12} = \frac{\Sigma_{12,0}(1+z)^3}{\mathcal{H}(z)}, \tag{3.14}$$

$$\Sigma_{13} = \frac{\Sigma_{13,0}(1+z)^3}{\mathcal{H}(z)}, \tag{3.15}$$

$$\bar{N} = \frac{\bar{N}_0(1+z)}{\mathcal{H}(z)}. \tag{3.16}$$

Conformal time becomes

$$\tau - \tau_0 = -\frac{1}{x} \int_0^z \frac{dz}{\mathcal{H}(z)}, \tag{3.17}$$

which can be used in the Barrow expressions for s and ψ. The integrals of Barrow et al. equations (4.12) and (4.13) become

$$\int_0^{z_E} \frac{s(1-s^2)\sin\psi(1+z)^2 dz}{(1+s^2)^2 \mathcal{H}(z)} \tag{3.18}$$

$$\int_0^{z_E} \frac{s(1-s^2)\cos\psi(1+z)^2 dz}{(1+s^2)^2 \mathcal{H}(z)} \tag{3.19}$$

with $C_3 = 4$, where z_E is the redshift at photon emission. In calculating the constant C_1, the density Ω_0 should now be Ω_{m0}, i.e. the matter density only.

The following assumptions should be kept in mind regarding this solution:

- The tilt velocity is small. In a dust- or Λ-dominated Bianchi type VII_h universe, the tilt asymptotically tends to zero at late times. Using the above relations, it can also be shown that for the case of interest, it remains small back to the last scattering surface as well.

- The universe is close to FRW at all times. This assumption is required to make the problem tractable. In a Λ-dominated universe, the type VII_h models become close to FRW. For the cases in question, plugging the numbers into the above relations shows that this assumption still holds at the surface of last scattering. (At very early times, however, this is not the case.) It is likely that models for which this assumption does not hold at last scattering would then have greater shear effects, but the details would require a more complicated derivation.

- Barrow et al. (1985) only consider the contribution from the vorticity that directly involves the shear variables Σ_{12} and Σ_{13}. The (expansion-normalised) shear scalar, $\Sigma^2 = \sigma^2/(3H^2)$, is given by

$$\Sigma^2 = \Sigma_+^2 + |\boldsymbol{\Sigma}_\times|^2 + \Sigma_{12}^2 + \Sigma_{13}^2 \tag{3.20}$$

  and involves therefore all of the shear components. In Barrow et al. (1985) and here, the other shear components have simply been ignored in the calculation of $\Delta T/T$. Considering additional shear degrees of freedom is in principle possible but in practice difficult. The effect of such additions is one of the largest unknowns in this analysis.

3.2.2   Properties

In the presence of the cosmological constant, a universe that is very nearly flat now and that was very nearly flat at the time of last scattering can become significantly negatively curved at intermediate redshifts. We can see this from equation (3.11) above.

Figure 3.1 shows how $K(z)$ evolves for different values of $\Omega_{\Lambda 0}$ and Ω_{m0}. At high redshift, the matter dominates and the curvature vanishes, and at low redshift, Ω_Λ can dominate, depending on its exact value. At intermediate redshifts, however, the universe may become more negatively curved. For a currently almost flat universe with parameters close to the observed values of $\Omega_{\Lambda 0} \sim 0.7$ and $\Omega_{m0} \sim 0.3$, the curvature is at most only a few percent (black curve). Only for very small matter densities and very large dark energy densities does the curvature become large at intermediate redshifts. To reproduce the observed asymmetry with these $\Lambda > 0$ models requires either a large current negative curvature (pink and blue dot-dashed curves) or a very small matter density (green dashed curve).
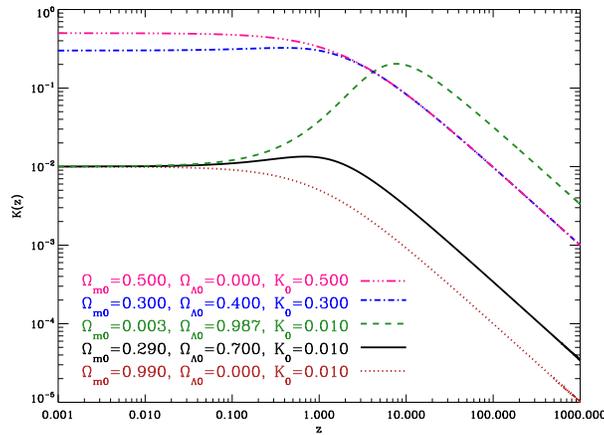
Figure 3.1 Examples of the evolution of $K(z)$. See text.

The addition of a Λ component adds a third parameter to the Bianchi parameter space, but there is a degeneracy to the resulting morphology of the Bianchi-induced pattern. The same structure in the induced anisotropy pattern can be reproduced by different parameter combinations. The addition of dark energy results in a tighter spiral for a given value of x, because the redshift to the surface of last scattering corresponds to a larger time difference. The geodesics have therefore completed more rotations since recombination for a given value of x when $\Omega_\Lambda > 0$. Secondly, the focusing of the spiral depends on the curvature between the observer and the last scattering surface. Significant focusing can arise from a universe that is very open now, or from a universe where Λ is large enough to give negative curvature at intermediate redshifts (as shown in Fig. 3.1.)

Three combinations of parameters, for example, can reproduce the same apparent structure on the sky :

1. $(x, \Omega_{\Lambda 0}, \Omega_{m0}) = (0.62, 0, 0.5)$ – i.e., no Λ and a matter content half critical, giving a large current negative curvature (the model found in Jaffe et al., 2006a);

2. $(x, \Omega_{\Lambda 0}, \Omega_{m0}) = (0.8, 0.4, 0.3)$ – i.e., approximately the observed matter content (baryon plus dark) , with some Λ but still a large current negative curvature;

3. $(x, \Omega_{\Lambda 0}, \Omega_{m0}) = (4.0, 0.987, 0.003)$ – i.e., a much smaller than observed matter content, but a nearly flat current curvature ($\Omega_k = 0.01$); the large Λ causes a large negative curvature at intermediate redshifts.

The real difference between these models is the amplitude: (3) has an amplitude $\sim 80\%$ of (1), and (2) has an amplitude of ~ 8 *times* that of (1). The shear and vorticity values obtained from fitting these templates to the data would change accordingly.

Effectively, this implies that a given structure on the sky characterised by the amount of geodesic focusing and the number of spiral turns is the same for all models along a line in the three dimensional parameter space of $(\Omega_{k0}, \Omega_{\Lambda 0}, x)$. The degeneracy in the templates is

only broken by the amplitude of the variation. But $(\sigma/H)_0$ is what we measure by fitting a template to the sky, so we cannot distinguish among the degenerate models without an independent measurement of the shear.

It is not straightforward to calculate where such lines of degeneracy lie. Instead, we determine this empirically by generating a grid of models and comparing them to the previously determined best-fit model.

3.3 Viability of a Bianchi Type VII$_h$ Model

We now address the question of whether the degeneracy of the Bianchi type VII$_h$ models with $\Omega_{\Lambda 0} > 0$ described in §3.2 gives us enough freedom to define a Bianchi model that is morphologically identical to the best-fit model in Jaffe et al. (2005, 2006a) and that is also consistent with estimates of the cosmological parameters.

3.3.1 Matter Density and 'Geometric' Degeneracy

Efstathiou & Bond (1999) describe the limitations of constraining cosmological parameters with CMB data alone. In particular, there is a 'geometric' degeneracy in the curvature and dark energy density (or equivalently, the matter density.) Degenerate models have: 1) the same values of the physical baryon density , $\omega_{\mathrm{b}} = \Omega_{\mathrm{b}}h^2$, and cold dark matter density, $\omega_{\mathrm{c}} = \Omega_{\mathrm{c}}h^2$; 2) the same primordial fluctuation spectrum; and 3) the same value of the parameter

$$\mathcal{R} = \sqrt{\frac{\omega_{\mathrm{m}}}{\omega_k}} \begin{cases} \sinh[\sqrt{\omega_k}y] & \text{if} \omega_{\mathrm{k}} > 0, \\ \sqrt{\omega_k}y & \text{if} \omega_{\mathrm{k}} = 0, \\ \sin[\sqrt{\omega_k}y] & \text{if} \omega_{\mathrm{k}} < 0, \end{cases} \tag{3.21}$$

where

$$y = \int_{a_{\mathrm{rec}}}^{1} \frac{da}{\sqrt{\omega_{\mathrm{m}}a + \omega_k a^2 + \omega_\Lambda a^4}} \tag{3.22}$$

(where a_{rec} is the scale factor at recombination). These models will have almost the same power spectra. Contours of constant $\mathcal{R}$ in model space Ω_k versus Ω_Λ are shown in Figure 3.2, with red highlighting the curve intersecting the best-fit, flat, *WMAP*-only parameters (Spergel et al., 2003, Table 1). Figure 3.3 shows power spectra for those degenerate models (calculated using CMBFAST.[1]) When normalised, the spectra are identical at the acoustic peaks, and the degeneracy only fails at large angular scales.

In the three dimensional parameter space of $(\Omega_{k0}, \Omega_{\Lambda 0}, x)$, the geometric degeneracy forms a surface, while the Bianchi models with identical structure form a line. The two need not intersect, but in the case of the observed *WMAP* CMB-only power spectrum of Spergel et al. (2003) and best-fit Bianchi model of Jaffe et al. (2005, 2006a), they do. There is a Bianchi model that has the identical structure to the best-fit $\Omega_\Lambda = 0$ model and

---

[1]http://www.cmbfast.org/
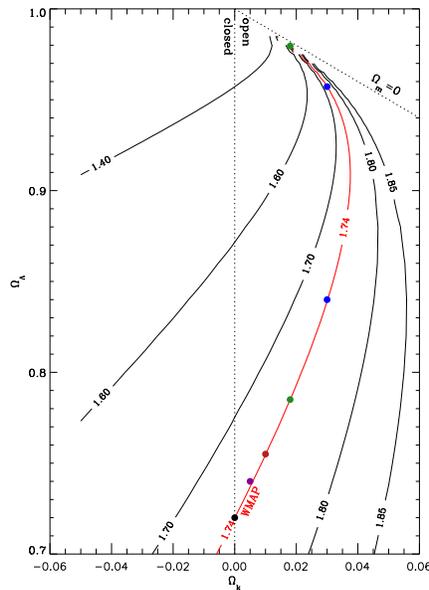
Figure 3.2 Degeneracy contours on the $\Omega_\Lambda - \Omega_k$ plane for $\omega_b = 0.023$ and $\omega_m = 0.13$. Along the red line denoting $\mathcal{R}$ for the *WMAP* best-fit parameters, the circles mark example power spectra that are shown in Figure 3.3.

lies on the *WMAP* degeneracy curve. It has parameters $(\Omega_{k0}, \Omega_{\Lambda 0}, x) = (0.028, 0.96, 2.5)$. But how viable is this region of parameter space?

As shown in Figure 3.3, the geometric degeneracy is broken at large angular scales. Models with high $\Omega_{\Lambda 0}$ and low Ω_{m0} predict too much large scale power. *WMAP* data alone place relatively loose constraints on Ω_{m0}, but even these rule out such low values as required for the Bianchi models. (The power spectrum of the sky corrected for the Bianchi component has less large-scale power overall (Jaffe et al., 2005, 2006a) than the uncorrected power shown in that figure, so the problem will only become worse.)

To quantify these limits, we examine the posterior likelihood from the *WMAP* data alone using the COSMOMC[2] code of Lewis & Bridle (2002), which implements a Monte Carlo Markov Chain method. For a simple look to compare the *WMAP* constraints to the Bianchi degeneracy, we explore the five-parameter space of: θ (the ratio of the sound horizon to the angular diameter distance), τ (optical depth), Ω_k (curvature), n_s (spectral index of scalar perturbations), and A_s (amplitude of scalar perturbations), with fixed parameters $\omega_b = 0.023$ (the physical baryon density), $\omega_{dm} = 0.107$ (the physical dark matter density), $f_\nu = 0$ (neutrino fraction of dark matter density), $w = -1$ (dark energy equation of state), $r \equiv A_t/A_s = 0$ (ratio of tensor to scalar fluctuations). Other parameters such as H_0, $\Omega_{\Lambda 0}$, and Ω_{m0} are derived from this set. (The free parameters have only such priors as defined by ranges that are much broader than any realistic uncertainty.) The resulting constraints in the $\Omega_m - \Omega_\Lambda$ plane (marginalising over the other parameters) are

---

[2]http://cosmologist.info/cosmomc/
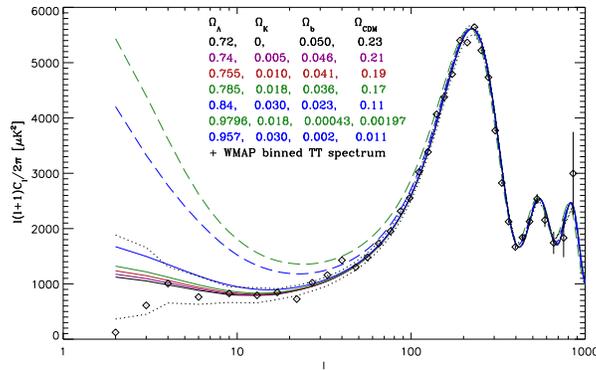
Figure 3.3 Power spectra for several models degenerate with the *WMAP* data (marked with circles in figure 3.2.) The solid versus dashed green and blue curves are where there are two degenerate models at the same value of Ω_k; the dashed is that with the higher Ω_Λ. The spectra are normalised so that the amplitude at the first acoustic peak is always the same. The dotted black line shows the cosmic variance uncertainty for each data point around the *WMAP* model.

shown in Figure 3.4 in red. All Bianchi models with the structure of the best-fit model lie well outside the 95% confidence region from the *WMAP* data alone. Allowing the physical matter densities to vary as well but including other datasets, the constraints (in magenta) are even tighter.

Tegmark et al. (2004) give $\Omega_{\rm m0} = 0.57^{+0.45}_{-0.33}$ from *WMAP* data alone, which tightens to $\Omega_{\rm m0} = 0.40^{+0.10}_{-0.09}$ adding data from the Sloan Digital Sky Survey (SDSS). See also Sánchez et al. (2006) for CMB+2dFGRS results.

3.3.2 Optical Depth

Figure 3.3 shows that the geometric degeneracy breaks at large angular scales, and the models with high Ω_Λ have too much large scale power. But there is an additional degeneracy if we allow the optical depth, τ, to vary. Adjusting τ and $A_{\rm s}$ such that $A_{\rm p} \equiv A_{\rm s}e^{-2\tau}$ remains constant has exactly the effect we need of modifying only the large scale power while leaving the peak heights unchanged (Tegmark et al., 2004). Figure 3.5 shows how this works, but that the effect is not large enough. Even a τ of zero does not bring the power at large angular scales within range of the data, and that model is inconsistent with the large-scale peak in the TE spectrum.

One can also ask if the addition of the Bianchi component can affect the variance of the large scale TE cross-power, resulting in an incorrect estimate of τ due to a chance alignment of the polarisation signal with the Bianchi structure. Simulations with no reionisation and with an added Bianchi temperature pattern (the same location and shear amplitude as our best-fit model) show that this is not the case. The cross-power remains flat at low ℓ with the expected variance.
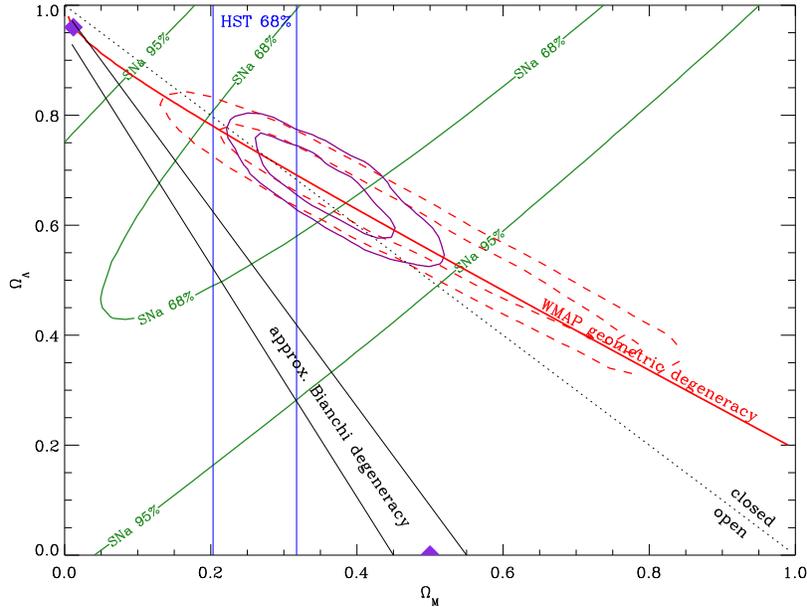
Figure 3.4 Parameter space by Ω_{m}. In green are the supernovae constraints (Knop et al., 2003). The solid red line is the same geometric degeneracy curve shown in Figure 3.2. In blue are the HST Key Project (Freedman et al., 2001) constraints from $\Omega_{\mathrm{m0}} = \omega_{\mathrm{m}}/h^2$ (assuming fixed ω_{m} and ω_{c}). The black solid lines show an approximate representation of where the Bianchi degeneracy lies, from the original model at $(x, \Omega_{\Lambda 0}, \Omega_{\mathrm{m0}}) = (0.62, 0, 0.5)$ to the one that lies on the *WMAP* geometric degeneracy curve at $(0.028, 0.96, 2.5)$ (each shown with a violet diamond.) The likelihood contours from *WMAP* data alone (computed using CosmoMC; see text) are shown with the red dashed lines. In solid magenta are the contours from *WMAP*, supernovae, HST, and SDSS data combined, where ω_{m} and ω_{c} also vary.
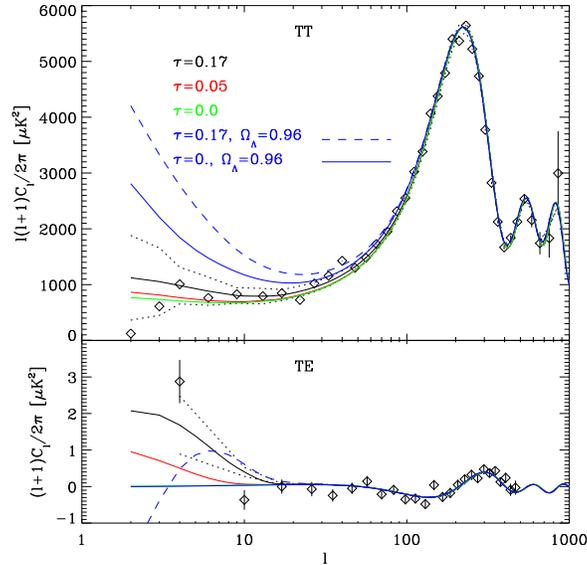
Figure 3.5 Power spectra as a function of optical depth, τ, each normalised such that the first peak heights are all the same (effectively changing A_s.) The top panel is the TT power, and the bottom the TE cross-power. Lowering τ lowers the large scale power, but not enough to bring the high Ω_Λ models down to the range of the data.

Note that Hansen et al. (2004a) find different values of τ derived from fitting temperature data in the Northern and Southern hemispheres as defined in the frame of reference that maximises the power asymmetry (see Hansen et al. 2004b.) The additional large-scale structure in the South that the Bianchi template reproduces could contribute to this effect; Jaffe et al. (2005, Fig. 2) shows a comparison of the corrected power spectrum with the theoretical spectra fit to the uncorrected full sky and to the Northern hemisphere. If a polarisation signal were also produced in a Bianchi topology and correlated with the temperature anisotropy, it would influence the measured optical depth from the low-ℓ TE peak as well.

3.3.3   Supernovae, H_0, etc.

The degeneracy contours in Figure 3.2 are plotted for constant values of $\omega_m = 0.13$ and $\omega_b = 0.023$ (from Spergel et al., 2003). For different points in the $(\Omega_{k0}, \Omega_{\Lambda0})$ space, Ω_{m0} changes and therefore the value of $h \equiv H_0/(100 \text{ kms}^{-1}\text{Mpc}^{-1})$ changes, increasing with increasing $\Omega_{\Lambda0}$ along the degeneracy curve. For models of high $\Omega_{\Lambda0}$ that lie on the geometric degeneracy curve, H_0 reaches values over 300. This is ruled out at high significance by the *WMAP* data itself; Tegmark et al. (2004) give $H_0 = 48^{+27}_{-12}\text{kms}^{-1}\text{Mpc}^{-1}$. Independent determinations of Hubble's constant also rule out these values, e.g., the HST Key Project (Freedman et al., 2001) value of $72\pm8\text{kms}^{-1}\text{Mpc}^{-1}$. See, e.g., Spergel et al. (2003, Table 3)

for other estimates of the Hubble constant.

High-redshift supernovae observations can break the geometric degeneracy by independently placing constraints that are nearly perpendicular to the *WMAP* constraints in $\Omega_{\mathrm{m}} - \Omega_\Lambda$ space (see Figure 3.4.) The degeneracy curve for the best-fit Bianchi model compared to the *WMAP* data lies outside the 98% confidence contour determined by the Supernova Cosmology Project (Knop et al., 2003, fit no. 6).

It is worth pointing out that it is not clear whether type Ia supernovae are truly standard candles out to high redshift. The uncertainty, however, is not enough to accommodate such low values of $\Omega_{\mathrm{m}0}$.

The addition of cosmic vorticity and shear would, of course, influence the determinations of such parameters. The sky coverage of supernova, for example, is fairly extensive outside the Galactic Plane region, but the sample size is small. Studies such as Kolatt & Lahav (2001) have ruled out significant dipole or quadrupole asymmetry in the expansion, but a test for a more complicated anisotropy structure induced by the vorticity would require far more supernova in the sample and a good sky coverage. Furthermore, as the current value of the shear expansion is very small, distance measurements to relatively low redshift (compared to the CMB) objects may not be sensitive enough to detect it. This also implies that such a small current shear would not significantly alter measurements of other cosmological parameters such as the Hubble constant that are dependent on relatively low-redshift observations.

3.3.4   Other Bianchi Models

One can also ask, then, how much one can vary the parameters of the Bianchi model and still have a statistically significant fit. The most interesting region on Figure 3.4 is near $(\Omega_{\mathrm{m}}, \Omega_\Lambda) = (0.15, 0.75)$, where the *WMAP* and supernovae constraints approach the region of the best-fit Bianchi model. This model resembles the best-fit model, but since it is closer to flat than the Bianchi degeneracy curve, there is less geodesic focusing. As a result, the model is not so good a fit to the data, and has a significance, compared to Gaussian realizations, of only 85%, compared to over 99% for the model at the same Ω_Λ but on the Bianchi degeneracy curve.

Considering that the *WMAP* data somewhat favour a closed universe, one might ask about closed models with vorticity and shear. These are Bianchi type IX models, also discussed in Barrow et al. (1985). Unlike the open type VII$_h$ models, however, closed models exhibit neither geodesic focusing nor the spiral pattern, even in the presence of vorticity. Barrow et al. use them to place limits on vorticity and shear simply using the quadrupole. Such closed models do not reproduce the morphology needed to explain the power asymmetry or the cold spot, and it would be impossible to claim any detection with only the quadrupole as an observable.

3.3.5 Other Dark Energy Models

The cosmological constant is the simplest form of dark energy, and all observations so far remain consistent with it. Other models are not ruled out, however, and in some cases are favoured. Alternatives come in many varieties, some with physical motivation, others constructed to give a particular result. (See, e.g., Padmanabhan 2003.) Here, we address the question of whether an alternative dark energy model can bring our Bianchi pattern any closer to the constraints imposed by the data.

Dark energy models are characterised by their equation of state, $p = w\rho$, where a cosmological constant Λ corresponds to a model with a constant $w = -1$. Alternative theories allow w to vary with time, as in "quintessence" and "k-essence" models. A rather *ad hoc* parameterisation is often used of the form $w = w_0 + w_1 z$, which allows comparison of generic dark energy models with supernova data (Wang et al., 2004; Dicus & Repko, 2004).

The derivation in §3.2 can be generalised for any dark energy model by the appropriate substitution into equations (3.10) and (3.13) of Ω_X instead of Ω_Λ. We then have

$$\mathcal{H}(z)^2 = \Omega_{X0} f_X(z) + K_0(1+z)^2 + \Omega_{m0}(1+z)^3 \tag{3.23}$$

where the function f_X is derived from the equation of state as

$$f_X(z) = \exp\left\{ 3 \int_0^z \frac{dz'}{1+z'} \left(1 + w(z')\right) \right\} \tag{3.24}$$

Examples for extreme values of w_0 and w_1 are shown in Figure 3.6. As was shown in Figure 3.1, we can look at the evolution of the curvature as a function of time for these models and see if any display the required negative curvature at intermediate redshift. Using this parameterisation, or alternatively the parameterisation $w = w_0 + w_1\frac{z}{z+1}$, does not give much additional freedom, however, to create a viable model. Ultimately, the amount of curvature and therefore focusing is still largely a function of the relative densities of matter and dark energy and of the current curvature. Manipulating the evolution of the dark energy equation of state simply changes the time scales on which the transitions between the different components of the density occur.

3.3.6 Small-scale Structure

We have so far ignored an additional issue with these models related to the stochastic component of the CMB. The analysis in Jaffe et al. (2005, 2006a) assumes that the observed CMB anisotropies consist of two independent signals: the predicted Bianchi pattern, and a stochastic, statistically isotropic component. The latter may be generated via inflation or another mechanism, but if it is statistically isotropic at the surface of last scattering, it might no longer be statistically isotropic when observed after travelling through a Bianchi universe.
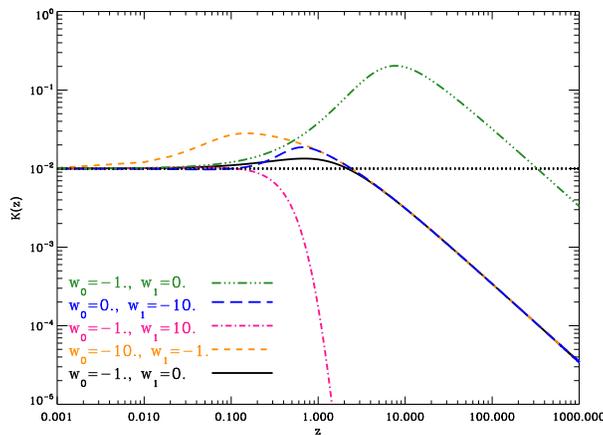
Figure 3.6 Examples of the evolution of $K(z)$ for different dark energy models where the current curvature is constrained to be small ($K_0 = 0.01$). $\Omega_{m0} = 0.29$ for all except the green triple-dot-dashed curve, where $\Omega_{m0} = 0.003$. This is the curve shown in Fig. 3.1 with a small matter density, high Ω_Λ, and approximately the amount of curvature needed to create the best-fit Bianchi model.

The power spectrum at small scales might deviate from predictions due to additional structure (dependent on orientation relative to the preferred axis) introduced by propagation in a Bianchi metric. We note that the asymmetry observed in Eriksen et al. (2004b) and Hansen et al. (2004b) extends to $\ell \sim 40$, while the Bianchi model has structure only up to $\ell \sim 15$. Geodesic focusing might cause such a power asymmetry, though it might also require that the asymmetry continue to the smallest angular scales, which is not observed. Hansen et al. (2004b) also find that some unexplained outliers in the *WMAP* power spectrum are associated with different hemispheres.

Furthermore, the generation of the fluctuations at the last scattering surface might also be affected by the anisotropy, though in the case in question, that anisotropy remains very small ($\lesssim 10^{-4}$) at $z \sim 1000$. Essentially, we expect that the Bianchi models could also be constrained by the lack of deviations in the power spectrum from the best-fit *WMAP*/concordance model, although detailed predictions for such deviations are needed.

The consistency of the acoustic peak scale with the flat concordance model would be very difficult for our best-fit Bianchi model to explain.

3.3.7   Early Universe

Considering the success of Big Bang Nucleosynthesis (BBN) in explaining light element abundances, any model that significantly changes the physical processes at those early times must be ruled out. As discussed in §3.3.1, to reproduce the morphology of the best-fit model requires a very small total matter density, namely $\Omega_m = 0.012$. If there were no dark matter, this value is still smaller than current BBN predictions for the total baryon density, e.g., the Burles et al. 2001 prediction of $\Omega_b = 0.02 \pm 0.002$. An additional concern

is the impact of the anisotropic expansion itself at the time of BBN.

Several studies have thus attempted to place limits on shear expansion and rotation by examining their effect on the relative Helium abundance, Y. Barrow (1976) originally showed that such analysis can place strong constraints on shear. The interesting exception, however, was type VII$_h$ models, where Barrow found that the CMB remained a stronger constraint. Later studies showed that including more complicated effects can reverse the trend of Y with the shear (see, e.g., Juszkiewicz et al. 1983). Barrow (1984) revisited the issue and showed that in some cases, even extremely anisotropic models may still have the observed Helium abundance. Bianchi type VII$_h$ models, however, have not been treated in detail since the result of Barrow (1976). After demonstrating for several cases, not including VII$_h$, that Y increases for anisotropic models, Matzner et al. (1986) *conjecture* that the same trend applies for *all* anisotropic models, but they admit that this is difficult to prove.

In short, even ignoring the problem of the baryon density, it is possible that such Bianchi models are strongly ruled out by the effect of the anisotropy on BBN and the observed Helium fraction. But this has not been definitively proved.

3.4 Discussion and Conclusions

We have presented solutions for Bianchi type VII$_h$ type universes that include a dark energy term and examined how their morphological properties change over the parameter space. The addition of dark energy adds a degeneracy such that different combinations of the three parameters $(\Omega_{\mathrm{m}}, \Omega_\Lambda, x)$ can lead to the same observed structure as in the best-fit model of Jaffe et al. (2005, 2006a). A template can be constructed that has the identical structure of that best-fit model and also falls on the geometric degeneracy curve for the parameters as measured by *WMAP* data.

This model, however, lies well outside the 95% confidence region in $\Omega_{\mathrm{m}} - \Omega_\Lambda$ space for *WMAP* data, ruled out by the overprediction of large scale power. It also lies outside the 95% likelihood contours from the Supernova Cosmology Project, and is further inconsistent with the measurement of the Hubble constant from the Hubble Key Project. Bianchi models that are more consistent with these other measurements are no longer good fits to the *WMAP* large scale structure.

One of the most difficult problems for these models is to account for the acoustic peak structure. The anisotropy at early times might influence the nature of the fluctuations at last scattering, and the geometry could affect the power spectrum on small angular scales due to the geodesic focusing between last scattering and the observer. Detailed predictions for these effects would be needed, but it is difficult to envision such an anisotropic scenario that happened to reproduce the observed acoustic peak structure, mimicking the concordance cosmology so well.

There is currently no prediction for the CMB polarisation anisotropy in a Bianchi universe, but such a topology-induced signal could provide an additional test of these models. If the preferred direction indicated in the temperature data is also reflected in

the full sky polarisation data (expected from further *WMAP* data releases and, eventually, from Planck), there will be even more motivation to consider non-standard models.

We have shown that our best-fit Bianchi type VII$_h$ model is not compatible with measured cosmological parameters, despite the additional freedom from adding dark energy. It is worth reiterating, however, that the serendipitous discovery of a theoretically derived template that correlates well with the data also happens to resolve several anomalies that cannot be explained in the standard picture. These particular models may not be viable, but lacking any plausible scenario for systematics or foregrounds to be the source of the anomalies, non-standard models that reproduce a similar morphology merit continued interest.

Chapter 4

Cross-Correlation in WMAP Foreground Analysis

In the previous chapters, the microwave foregrounds have been treated as an obstacle to CMB analysis. Template fitting using foreground predictions and the data in all *WMAP* bands was used simply to rule out the possibility that the cross-correlation signal between the data and the Bianchi template was due to foreground contamination. This chapter and the following treat the fitting of the foreground templates themselves in detail, first from the point of view of removing them and second from the point of view of understanding them.

The purpose of this chapter is to discuss the template-fitting method in the context of foreground subtraction, particularly its limitations, and to motivate the analysis that follows. The approach in Chapter 5 is to study the foreground emission processes themselves, and the different sources of Galactic foreground emission are discussed in more detail there and in references therein.

4.1   Galactic Foreground Templates

There are four morphologically and spectroscopically distinct components to the foreground emission from the Galaxy in the microwave regime. Thermal dust, synchrotron, and free-free are the three that are best understood. The thermal dust component has a rising spectrum in the microwave and is the dominant foreground at the higher frequencies such as *WMAP* W-band. The non-thermal synchrotron and free-free components contribute most at the low frequencies, the former having a steeper power-law index and therefore increasingly dominant at longer wavelengths. Both *COBE*-DMR and *WMAP* were designed to observe the microwave sky where the cosmic signal dominates in a window between these components. The DMR data revealed the presence of a fourth component often referred to as anomalous dust for its apparent correlation to the thermal dust morphology but very different spectral behaviour. For each of these components, we use observations at other wavelengths to predict the emission in the microwave bands.

(a) SFD dust

(b) FDS dust

(c) SFD dust T

(e) DDD H$_\alpha$, $f_d = 0$

(f) DDD H$_\alpha$, $f_d = 0.5$

(g) Finkb. H$_\alpha$ - DDD, $f_d = 0$

(h) Haslam 408 MHz

(i) *WMAP* K-Ka

(j) K-Ka-Haslam
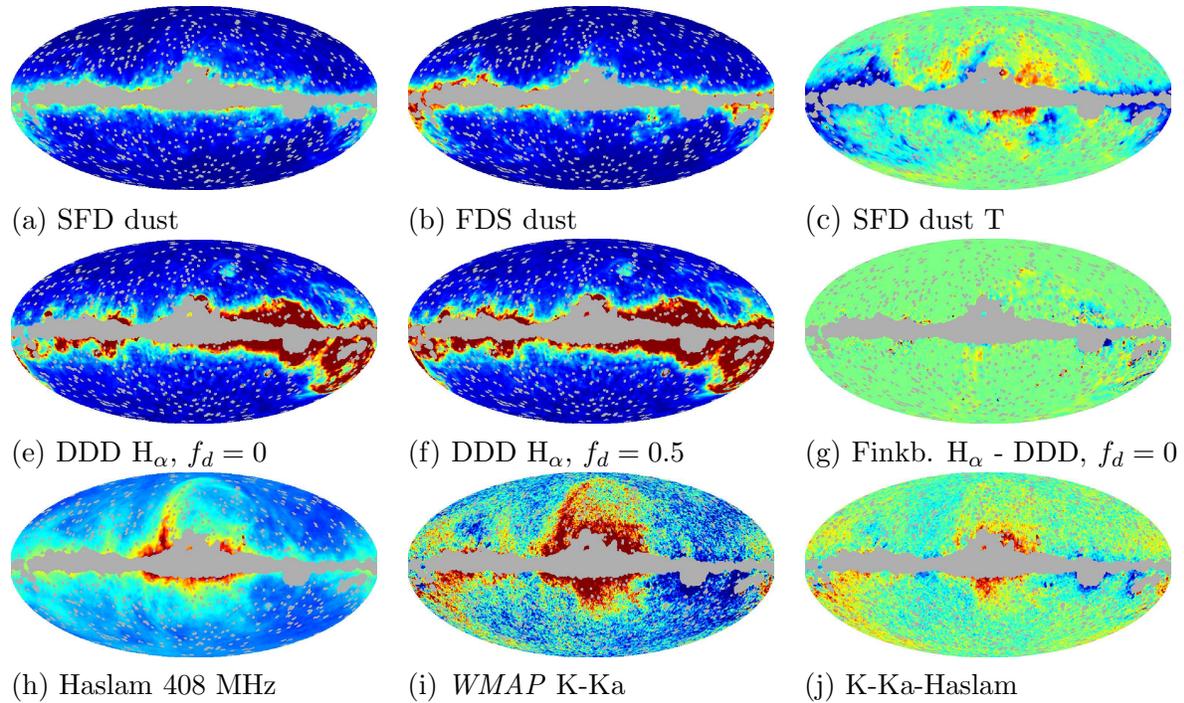
Figure 4.1 Foreground template comparisons. H$_\alpha$ is shown for K-band, synchrotron for Q, and dust for W-band. Each is shown at the amplitude it fits to the data. The K-Ka map is corrected for the dust and H$_\alpha$ to leave only synchrotron. See text. (Emission templates scale from $0 - 100\mu K$, except for synchrotron to $50\mu K$; dust color temperature, T, from $18 - 20K$; difference maps $\pm 50\mu K$.)

4.1.1 Dust

Thermal dust is strong in the infrared, which gives us two templates often used to model the *thermal* dust emission expected at microwave bands:

- Schlegel, Finkbeiner, & Davis (1998; hereafter SFD) map of dust emission at 100 μm from DIRBE and IRAS. This template is limited in its ability to predict the emission at other wavelengths by the variation in dust temperature across the sky.

- Finkbeiner, Davis, & Schlegel (1999; hereafter FDS) model 8 prediction for dust emission at 94GHz (available on the LAMBDA web site). This template is based on the SFD 100μm map but characterises the dust emissivity using a two component model and the 100/240μm flux ratio (from *COBE*-FIRAS). This attempts to correct the emissivity for the variation in dust colour temperature. The two component model fit gives a thermal spectrum following a power law with $\beta = 1.7$, though the *WMAP* team assume $\beta = 2$ in their analysis.

These templates are described in more detail in § 5.2.3. The top row of Fig. 4.1 shows these two templates as well as the SFD colour temperature .

It is now clear, however, that there is an additional foreground component that correlates well with the dust distribution but that rises at lower frequencies. Finkbeiner (2004a) proposes a rather *ad hoc* model for this "anomalous dust" component by multiplying the FDS dust template by $T_{\mathrm{d}ust}^2$, and we also examine this template fit to the low frequency bands in § 4.4.4. Chapter 5.2.3 will discuss this component in more detail.

4.1.2 Free-Free

Free-free emission is not the dominant component at any wavelength, and therefore cannot be observed directly. Hα is a good tracer of ionised hydrogen, because it arises from the recombination of the same free electrons and ions. We have several predictions for free-free emission based on Hα surveys:

- Dickinson, Davies, & Davis (2003; hereafter DDD) composite of WHAM and SHASSA data, where the latter is used only where no WHAM coverage is available. These data are optionally corrected for dust absorption with fraction $f_d = 0.5$ (i.e., assuming uniform mixing of the dust and gas) using the abovementioned SFD dust template. This map has a nominal resolution of $\sim 1°$ FWHM.

- Finkbeiner (2003) composite including SHASSA, VTSS, and WHAM data, in that order, where available. This template, with no dust correction, is available at the LAMBDA website. This map has a nominal resolution of 6'.

These templates are discussed in more detail in § 5.2.2. The middle row of Fig. 4.1 shows the DDD template uncorrected for dust absorption, with a correction assuming $f_d = 0.5$, and the difference between the DDD and Finkbeiner (both uncorrected) templates.

Note that in addition to baseline calibration differences (see § 5.2.2), these two maps use different methods to reconcile the different beams of the WHAM and SHASSA surveys, resulting in maps that are in fact at different resolutions where WHAM is used. In effect, to compare the two, the Finkbeiner map must be smoothed with a beam of $(60^2 - 6^2)^{1/2}$ arcmin FWHM so that it has a nominal resolution of $1°$ as well. But the DDD map requires additional smoothing in the WHAM regions, as no smoothing had been applied there, and a $1°$ *top-hat* window function is closer to a Gaussian of FWHM $37'$. For full sky analysis at low resolution (e.g., in the case of the Bianchi analysis of Chapter 2), this difference has no effect. But for the analysis of small regions described in Chapter 5, the appropriate smoothing is taken into account, leading to very comparable results between the two templates. But care should be taken in using the DDD template at $1°$, since in fact, the beam properties of the template are slightly different in the WHAM and SHASSA regions.

Figure 4.2 shows the coverage of the different Hα surveys used in the two templates. Also shown are the individual regions that will be examined in Chapter 5.

The main problem with either of these templates is the variation across the sky of the electron temperature that is not easily taken into account.

4.1.3   Synchrotron

For synchrotron emission, we use the Haslam et al. (1982) map at 408 MHz, where it is least contaminated by the free-free emission. Again, the spectral variation across the sky is the main problem when a template based on observations at one wavelength (and one far from the microwave band) is used. See 5.2.4 for more details.

An alternative adopted by the *WMAP* team in their 3-year analysis is to use the K–Ka difference map, which they expect to be dominated by synchrotron emission, that takes into account some of this spectral variation. This choice certainly helps with removing the foregrounds from the higher frequency bands, but it is complicated to separate the Galactic components using a template that must be some combination of synchrotron, free-free, and anomalous dust emission. Furthermore, it appears that the anomalous dust-correlated component in fact dominates at K–Ka. Chapter 7 will discuss this in more detail, in particular the physical interpretation of the emission in K–Ka.

The bottom row of Fig. 4.1 shows the Haslam 408 MHz map on the left and the synchrotron component of the K minus Ka map for comparison in the middle, with the difference on the right. K minus Ka will also have free-free and dust emission, so in order to plot only the K–Ka synchrotron component, the others are subtracted. (We fit Haslam, DDD Hα , and FDS dust to K and to Ka separately. We subtract the fitted DDD and FDS components from each of K and Ka, then construct the resulting K–Ka map, which is now largely synchrotron and noise.) The templates are shown at the amplitudes they fit to the Q-band. In the difference map on the right is what Finkbeiner (2004a) dubs a "free-free haze"; see § 4.4.4.

The K–Ka map will be discussed further in Chapter 7. For the rest of this chapter and Chapter 5, the Haslam template will be used following the *WMAP* first-year analysis.
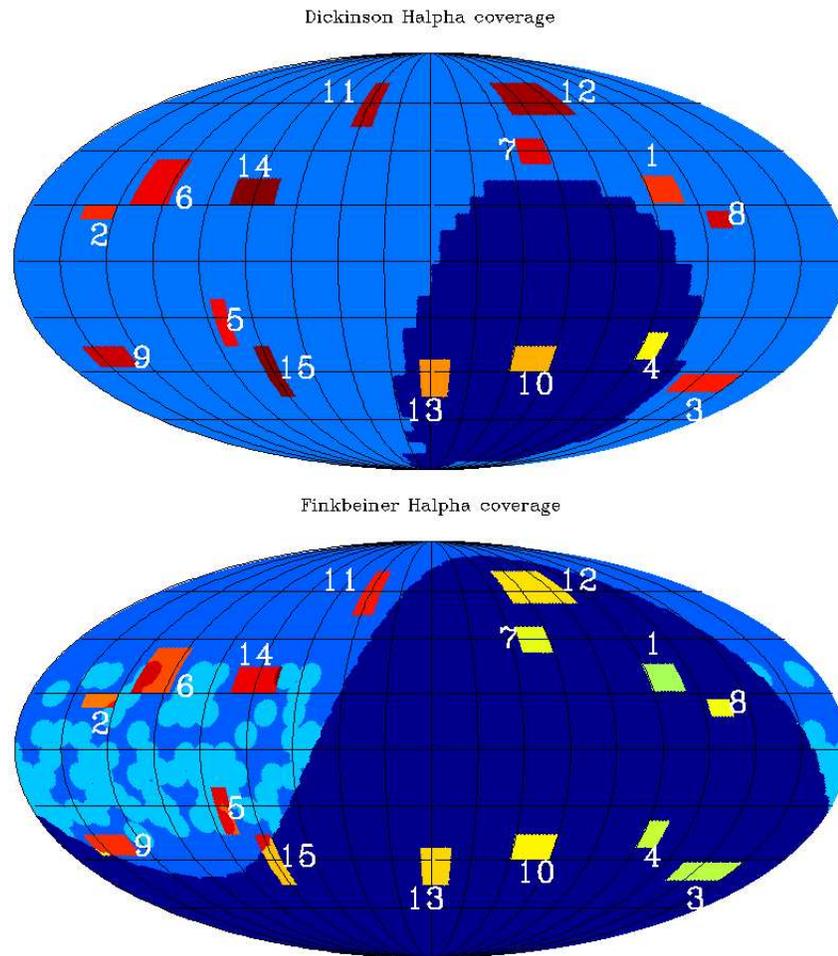
Figure 4.2 Dickinson Hα coverage (top): medium blue shows WHAM coverage, and dark blue SHASSA. Finkbeiner Hα coverage (bottom): pale blue shows VTSS coverage, medium blue shows WHAM coverage, dark blue SHASSA. Shades of red and yellow show regions defined in §5.

dataset	band	beam (FWHM °)	noise σ_0	Reference
WMAP K	22.8 GHz	0.82	1424μK	Bennett et al. (2003b)
WMAP Ka	33.0 GHz	0.62	1449μK	Ibid.
WMAP Q	40.7	0.49	2211μK	Ibid.
WMAP V	60.8	0.33	3112μK	Ibid.
WMAP W	93.5	0.21	6498μK	Ibid.
Haslam	408 MHz	0.85	-	Haslam et al. (1982)
SFD dust	100μm	0.102 (6.1')	-	Schlegel et al. (1998)
FDS dust	-	6.1'	-	Finkbeiner et al. (1999)
Finkbeiner Hα	-	nominally 6'	-	Finkbeiner (2003)
DDD Hα	-	nominally ~ 1	-	Dickinson et al. (2003)

Table 4.1 Properties of the maps used in the analysis.

4.2 Method

In §2.2.1, the basic method is described, and the following gives the more general formalism including fitting multiple bands with spectral constraints.

4.2.1 Simultaneous fitting of multiple templates to multiple bands

To compare multiple template components to multiple datasets, e.g., different foregrounds and different frequencies, the problem has one more level of complexity. The data vector $\mathbf{d}$ becomes the joint vector of all the frequencies, and the covariance matrix is composed of blocks, each of which is the covariance between the two relevant frequencies. In the case where we have the sky maps $\mathbf{d}$ in M frequencies and N different foreground components, and following the notation of Górski et al. (1996), we define:

$$\tilde{\mathbf{d}}_{\uplus} = \left(\begin{array}{c} \mathbf{d}_{\nu 1} - \sum_{k=1}^{k=N} \alpha_{\nu 1}^k \mathbf{t}_k \\ ... \\ \mathbf{d}_{\nu M} - \sum_{k=1}^{k=N} \alpha_{\nu M}^k \mathbf{t}_k \end{array} \right) \tag{4.1}$$

and

$$\mathbf{M}_{SN\uplus} = <\tilde{\mathbf{d}} \cdot \tilde{\mathbf{d}}^T> = \left(\begin{array}{ccc} \mathbf{M}_{SN\nu 1} & ... & \mathbf{M}_S \\ ... & ... & ... \\ \mathbf{M}_S & ... & \mathbf{M}_{SN\nu M} \end{array} \right) \tag{4.2}$$

The off-diagonal blocks include only the signal covariance, since the noise is assumed to be uncorrelated between bands. In this case, minimising $\tilde{\mathbf{d}}^T \cdot \mathbf{M}_{SN}^{-1} \cdot \tilde{\mathbf{d}}$ leads to the set of equations:

$$\sum_{j,\nu'} \mathbf{t}_k^T \cdot (\mathbf{M}_{SN\uplus}^{-1})_{\nu\nu'} \cdot \mathbf{t}_j \alpha_{\nu'}^j = \sum_{\nu''} \mathbf{t}_k^T \cdot (\mathbf{M}_{SN\uplus}^{-1})_{\nu\nu''} \cdot \mathbf{d}_{\nu''} \tag{4.3}$$

which is the simple system of linear equations $\mathbf{Ax} = \mathbf{b}$, where

$$A_{\nu\nu'}^{kj} = \mathbf{t}_k^T \cdot (\mathbf{M}_{SN\oplus})_{\nu\nu'}^{-1} \cdot \mathbf{t}_j$$

$$b_\nu^k = \sum_{\nu''} \mathbf{t}_k^T \cdot (\mathbf{M}_{SN\oplus})_{\nu\nu''}^{-1} \cdot \mathbf{d}_{\nu''}$$

$$x_\nu^k = \alpha_\nu^k \tag{4.4}$$

(The pair k and ν define one element, e.g., the amplitude α_ν^k of the k^{th} template in the ν^{th} band. The matrix $(\mathbf{M}_{SN\oplus})_{\nu\nu''}^{-1}$ is the $\nu\nu''$ block of the full inverse covariance matrix.) Where only one band and one template are present, this reduces to equation 2.2 above. Again, for incomplete sky coverage (e.g. using a Galactic mask or small regions as in Chapter 4), the data vector simply includes only those pixels of interest, and the covariance matrices only the corresponding rows and columns.

The errors $\delta\alpha_\nu^k$ are the square root of the diagonal of $\mathbf{A}^{-1}$.

4.2.2 Constrained fits

Constrained fits are useful where, in the simultaneous fitting of several bands, we can help reduce degeneracies by assuming that some or all templates scale at different frequencies with a known relationship, e.g., following a power law, leaving only one free parameter for the given template. Here, our system of linear equations $\mathbf{Ax} = \mathbf{b}$ is then of reduced dimension. Instead of $N \times M$ unknowns x_ν^j, we have M unknowns for each unconstrained template and only one for each constrained. The system to solve becomes:

$$A_{ij} = \sum_{\nu\nu'} \mathbf{c}_{\nu i}^T \cdot (\mathbf{M}_{SN\oplus})_{\nu\nu'}^{-1} \cdot \mathbf{c}_{\nu' j}$$

$$b_i = \sum_{\nu\nu'} \mathbf{c}_{\nu i}^T \cdot (\mathbf{M}_{SN\oplus})_{\nu\nu'}^{-1} \cdot \mathbf{d}_{\nu'}$$

$$\tag{4.5}$$

where the $\mathbf{c}_{\nu i}$ represent the constrained templates. For example, in the case where we have two frequencies and three templates, and the first template is unconstrained but the others are constrained, we have

$$\mathbf{c} = \begin{pmatrix} 1 * \mathbf{t}_0 & 0 * \mathbf{t}_0 & 1 * \mathbf{t}_1 & 0 * \mathbf{t}_1 & 1 * \mathbf{t}_2 & 0 * \mathbf{t}_2 \\ 0 * \mathbf{t}_0 & 1 * \mathbf{t}_0 & 0.2 * \mathbf{t}_1 & 0 * \mathbf{t}_1 & 0.3 * \mathbf{t}_2 & 0 * \mathbf{t}_2 \end{pmatrix} \tag{4.6}$$

representing the constraints that the amplitude of template $\mathbf{t}_1$ in the second band (or row) be 0.2 times its amplitude in the first band, and that the amplitude of template $\mathbf{t}_2$ in the second band be 0.3 times its amplitude in the first band. In this case, the unknowns correspond to the columns of $\mathbf{c}$ that have non-zero elements: $\mathbf{x} = (\alpha_{\nu_0}^0, \alpha_{\nu_1}^0, \alpha^1, \alpha^2)$.

4.2.3   Application

The application of the above methods used in this Chapter includes the following:

- Monopole and dipole

  Both the cosmic signal and galactic foregrounds all have monopole and dipole components, but in the data, the monopole and dipole are complicated by the differential nature of the *WMAP* measurements, zero-point uncertainties in the templates, and the fact that the dominant dipole component is the non-cosmic term induced by the Earth's motion. Rather than trying to disentangle them, these multipoles can be simply removed from the analysis in either of two ways: by setting the uncertainty effectively to infinity (in practice, a large number) in the assumed power spectrum that determines the signal covariance; or by fitting monopole and dipole template simultaneously. See § 4.3.2 for discussion. For the analysis in this chapter, we use the latter option.

- Resolution and covariance

  The analysis here is performed at a resolution of $N_{\mathrm{side}} = 128$, for a total of 196,608 pixels. The raw *WMAP* data maps as well as all templates are smoothed to a common resolution of $1°$ and converted to μK in *antenna* temperature, the natural units for discussing Galactic emission processes. (Note that the maps at this resolution and smoothing are *undersampled*. A higher resolution pixelisation would make the matrix inversion much more onerous. The undersampling does not cause a problem for this pixel-space analysis, as we verify with simulations.)

  The anisotropy pattern in the best-fit Bianchi model was a large scale structure, having little or no power above $\ell \sim 20$, so for that analysis, low resolution, highly smoothed maps could be used and the full covariance matrix could be inverted using roughly 500 GB of memory and 12 hours clock time on a desktop machine. The memory usage scales as N_{pix}^2, and the inversion scales at least as N^3, so it is clear that the full covariance matrix cannot be used for a full sky analysis at this resolution. For this reason, we use a diagonal approximation to this matrix. The simulations described in § 4.3 show that this method is unbiased, simply losing accuracy due to the approximation. In the next chapter, we use maps at the same resolution but over small patches of sky where the matrix inversion becomes tractable.

- Mask

  The Bianchi analysis used the most conservative Galaxy cut—the Kp0 mask—in order to minimise the effect of possible foreground residuals on the fits to the Bianchi templates. For analysis of the foregrounds themselves, we use the standard Kp2 cut, making use then of some emission closer to the plane. (These masks are based on the emission in the K-band, where Kp2 uses a higher threshold than Kp0. In all cases, the point sources are also masked. See Bennett et al. 2003b.) The choice of mask does influence the analysis, as is discussed in § 4.4.2.

4.3   Simulations and testing

We use simulations with foreground components added at known amplitudes to test how well different fitting methods recover the input. A set of 1000 simulations are created at the resolution of the *WMAP* data products, i.e., at HEALPix of $N_{side} = 512$, using the beam width and noise properties of the respective channels. These maps are then smoothed (via convolution in harmonic space) to a common resolution of $1°$ and downgraded to $N_{side} = 128$, as are the data. Different combinations of foreground residuals are added using the templates described in §4.1.

   We can also use these simulations to create the corresponding ILC map for each realisation by using the same weights found for the *WMAP* ILC. By using the same weights, we do not simulate the ILC process itself (though that is done in Eriksen et al. 2004a), but instead we ensure that similar amounts of foreground residuals are found in the resulting simulated ILC as are expected in the real ILC.

   Most of the tests discussed in the following were performed on both the full sky as well as in the small regions studied in Chapter 5. The two cases show the same behaviour.

4.3.1   Cross-talk

These simulations can be used to test how much cross-talk—correlation among the different templates themselves—affects the fits, particularly in the small regions chosen in an attempt to minimise such effects. For example, does the existence of small amounts of additional dust in the small, Hα-dominated regions affect the individual fits using the Hα template alone? Does fitting simultaneously three templates to a region with only Hα emission result in the Hα power aliased into the other two templates? Several test cases were examined:

- First, the foregrounds (FDS dust, DDD Hα, and Haslam synchrotron) were added at the amplitudes matching the full sky fitting done by the *WMAP* team for the dust (listed in Table 3 of Bennett et al. (2003c)) and at theoretical values for Hα (see Table 1 of same) and synchrotron (with a spectral index of $\beta = -3$.) The three templates were then fit simultaneously to each individual band. The results show that the means are in agreement with the true values and the calculated errors represent the variance in the results. Different regions may be slightly over- or underestimated systematically. Table 5.8 column "3to3[†]" shows how well on average the inputs are recovered for each region.

- Next, to the same simulations, we fit only the Hα template, to see how much of the power of the other foreground components are aliased into the Hα amplitude for the "individual" fits. Table 5.8 column "1to3[†]" shows that the Hα estimates *are* overestimated due to the other emission, even in regions dominated by Hα emission.

- Then we fit the three template components to simulations where only Hα was included to determine how much that component would be aliased into the other tem-

plates. This is essentially no different from fitting three components to simulations with three foregrounds; in this case, the true amplitudes of two of them are zero. In both cases, all the estimates remain correct on average with the expected variance.

4.3.2   Monopole and Dipole

As mentioned above, the fitting must be made to be insensitive to the monopole and dipole terms in the data. The most straightforward way to accomplish this is to fit monopole and dipole terms simultaneously. An alternative approach is to make the fit insensitive to these by calculating the expected covariance from a power spectrum with $C_0 = C_1 = 10^8 \mu \mathrm{K}^2 \sim \infty$. But the former approach is more intuitively obvious and allows us to see what monopole and dipole terms are fit. Simulations with an additional non-cosmological monopole and dipole show that either method is effective at removing any effect of such signals on the foreground fits.

Note that in the case of small regions, fitting out the monopole and dipole have the effect of removing the fit sensitivity to a gradient across the region being fit. The fit results end up with huge dipole terms, offset with the monopole term, which place the dipole such that it fits to such gradients. Simulations have shown, however, that the fit results do not change significantly when the dipole is fit out as well, or only the monopole. All of these methods give very similar results for the small regions analysed in Chapter 5, varying only small fractions ($\sim 10\%$ on average) of the calculated errors.

4.3.3   Stability

Tests where the simulations match exactly the templates being fit are not, of course, an accurate representation of the situation in reality. The templates we are using are imperfect matches to the data, and therefore those simulations give only the best-case expected results.

But we can learn how template uncertainties might affect our fits by using the different available predictions for dust and free-free emission. For example, simulations show that the fit results are fairly stable against changes in the way that the signal and noise covariance is approximated or in how the monopole is fit out. But fitting slightly different templates to the same simulations give less stable results. Due to the slight mismatch between the template and the data, the fit results are more affected by the stochastic noise and cosmic components. In the small region fits described in Chapter 5, some regions are more stable than others, depending on the differences in the templates in the particular area of the sky. This should not affect the results significantly, but it does imply that the real uncertainties are slightly larger than the calculated error bars.

4.3.4   Constrained fits

Simulations were created where the dust was added at the WMAP fit amplitude but the free-free and synchrotron were added following theoretical power-law predictions, with

Band	Dust	H_α	Haslam
Input values			
K	6.30	11.40	5.73
Ka	2.40	5.24	1.89
Q	1.50	3.38	1.01
V	0.90	1.45	0.30
W	1.20	0.59	0.08
Unconstrained			
K	6.30 ± 0.35	11.41 ± 0.88	5.74 ± 0.79
Ka	2.40 ± 0.34	5.25 ± 0.87	1.90 ± 0.78
Q	1.50 ± 0.34	3.39 ± 0.86	1.02 ± 0.77
V	0.90 ± 0.32	1.46 ± 0.82	0.31 ± 0.73
W	1.20 ± 0.29	0.59 ± 0.72	0.09 ± 0.65
Constrained			
K	6.30 ± 0.38	11.41 ± 1.27	5.75 ± 1.08
Ka	2.40 ± 0.35	5.25 ± 0.58	1.89 ± 0.36
Q	1.50 ± 0.37	3.38 ± 0.38	1.01 ± 0.19
V	0.90 ± 0.38	1.45 ± 0.16	0.30 ± 0.06
W	1.20 ± 0.34	0.59 ± 0.07	0.08 ± 0.02

Table 4.2 Fit results from constrained fitting to 1000 simulations. The mean fit value is shown with its RMS variance. For the Ka- through W-bands of free-free and synchrotron, which are constrained to follow power laws with $\beta = -2.1$ and -3.0 respectively, the entry is simply the appropriate multiple of the K-band fit value.

indices of -2.1 and -3.0 respectively using the DDD and Haslam templates. Table 4.2 shows the results, which are in good agreement with the input. For these simulations, where the inputs match the fit templates, the constrained fits do not show any improvement over the unconstrained fits. But in the real-world context of imperfect templates and an additional unknown anomalous component, constraining the components that we expect to follow a theoretical power-law index can be very useful.

4.4 Fit Results for **WMAP**

Both the *WMAP* team (in Bennett et al. 2003c and Hinshaw et al. 2006) and Finkbeiner (2004a) have used template fitting to study and subtract the Galactic foregrounds from the *WMAP* data. Previously, a detailed analysis of foregrounds with *COBE*-DMR data was done by Banday et al. (2003) with a slightly different template-fitting formalism. While the *WMAP* and DMR data have been shown to be consistent, the results of the foreground analyses are not entirely so. In particular, the DMR analysis found the amplitude of free-free emission to be consistent with electron temperatures of 5000-9000 K, while in the case of *WMAP*, the fit results show unexpectedly low values for free-free emission.

The following subsections discuss the fits at different resolutions and sky cuts, using different templates, and comparisons with DMR.

4.4.1   Smoothing and Sky Cut

In order to compare the *WMAP* fits to the DMR fits, we analyse the *WMAP* data at several smoothing scales including using a beam of FWHM 7°, roughly DMR's beam. The results of several *WMAP* fits with different inputs to investigate the effects of the smoothing are summarised in Table 4.3 (*WMAP* first year data only). Since a diagonal approximation to the covariance is used, the uncertainties are determined by simulations and shown only for the "default" column. They will be roughly the same when alternate templates are used, though they will go up accordingly with a larger sky cut.

(Note that when the *WMAP* data are smoothed with a large beam, the masks used in the usual analysis are no longer sufficient to completely exclude emission from the Galactic Plane region. We therefore expand the mask by smoothing it with the same beam. The mask is originally zero near the plane and one off the plan, and the smoothed mask has a gradient. We rather arbitrarily chose to keep only those pixels where the smoothed mask value remains above 0.95, which corresponds to cutting out almost one additional FWHM. This does not apply to the DMR mask, which already takes into account the large beam of that instrument.)

Table 4.3 shows small variations in the fit values as a function of the smoothing, but more striking are the variations depending on the cut used to mask the Galactic plane. This may be an indication that the Galactic gradients in the templates do not match the sky or that the fits are driven by serious mismatches in isolated regions near the plane. Further investigation is needed to explain these differences. Fig. 7.5 (g) shows the residuals from fitting the foreground templates to the MEM foreground solutions, and Fig. 4.1 shows residuals in the synchrotron component of K–Ka compared to the Haslam template. Both clearly show areas of mismatch near the galactic plane and something of a gradient, both of which will affect the fits with a strong dependence on both the cut and the degree of smoothing. See also Fig. 4.4.

To compare the SHASSA and WHAM data, we use the Finkbeiner Hα template where it uses SHASSA data and compare to using the DDD template where it uses WHAM (since the former uses SHASSA in the overlap regions and the latter uses WHAM; see Fig. 4.2). These results are fairly consistent with each other as seen in Table 4.3, though the WHAM fit results are systematically somewhat lower, increasingly so the smaller the smoothing beam applied. This is likely due to the different beams of the instruments, particularly the fact that the WHAM data in the DDD map have a 1° *top hat* beam, which is effectively smaller than a 1° Gaussian beam. In the analysis described in Chap. 5, we use an additional smoothing for the DDD template where WHAM data is used in order to better compare the template with the *WMAP* data. There, we find a very good agreement between the two templates.

Band	FWHM	WMAP	Default	ff05	SHASSA	WHAM	kp0	constr	minusHalpha
					dust				
	1	6.30	6.68($\pm$0.37)	6.25	6.62	6.64	6.84	6.85	6.22
k	3	-	6.70($\pm$0.48)	6.33	6.63	6.64	7.11	6.88	6.10
	7	-	7.05($\pm$0.77)	6.74	6.98	6.97	7.27	7.25	6.22
	1	2.40	2.44($\pm$0.37)	2.29	2.24	2.24	2.61	2.36	2.17
ka	3	-	2.45($\pm$0.47)	2.30	2.20	2.20	2.65	2.37	2.13
	7	-	2.55($\pm$0.76)	2.39	2.26	2.26	2.73	2.50	2.15
	1	1.50	1.26($\pm$0.36)	1.16	1.01	1.02	1.42	1.22	1.07
q	3	-	1.27($\pm$0.47)	1.18	0.98	0.98	1.45	1.21	1.04
	7	-	1.36($\pm$0.75)	1.26	1.04	1.03	1.51	1.23	1.08
	1	0.90	0.73($\pm$0.35)	0.70	0.48	0.48	0.92	0.24	0.62
v	3	-	0.77($\pm$0.45)	0.74	0.47	0.48	0.98	0.21	0.62
	7	-	0.86($\pm$0.72)	0.83	0.54	0.54	1.04	0.18	0.67
	1	1.20	1.02($\pm$0.31)	1.01	0.76	0.76	1.21	0.98	0.94
w	3	-	1.06($\pm$0.40)	1.05	0.76	0.77	1.26	0.97	0.96
	7	-	1.14($\pm$0.64)	1.13	0.82	0.82	1.32	0.95	1.02
					free-free				
	1	4.60	7.61($\pm$0.96)	5.94	8.06	7.80	5.03	7.32	11.40
k	3	-	7.22($\pm$1.22)	5.77	7.68	7.53	4.14	6.86	11.40
	7	-	6.87($\pm$1.67)	5.77	7.20	7.11	2.86	6.41	11.40
	1	2.10	3.30($\pm$0.95)	2.83	3.82	3.67	0.55	3.37	5.20
ka	3	-	3.20($\pm$1.20)	2.75	3.71	3.62	0.08	3.15	5.20
	7	-	3.07($\pm$1.64)	2.61	3.31	3.28	0.11	2.95	5.20
	1	1.30	1.96($\pm$0.93)	1.71	2.38	2.29	-0.63	2.17	3.30
q	3	-	1.85($\pm$1.19)	1.62	2.27	2.22	-1.09	2.03	3.30
	7	-	1.76($\pm$1.63)	1.54	1.91	1.90	-1.01	1.90	3.30
	1	0.50	0.60($\pm$0.89)	0.55	0.83	0.79	-1.58	2.17	1.40
v	3	-	0.49($\pm$1.13)	0.45	0.68	0.66	-2.03	2.03	1.40
	7	-	0.38($\pm$1.55)	0.38	0.30	0.31	-2.02	1.90	1.40
	1	0.10	0.04($\pm$0.79)	0.07	0.20	0.18	-1.80	0.38	0.60
w	3	-	-0.03($\pm$1.00)	0.00	0.11	0.09	-2.18	0.35	0.60
	7	-	-0.08($\pm$1.38)	-0.01	-0.18	-0.17	-2.14	0.33	0.60
					synchrotron				
	1	5.60	5.60($\pm$0.82)	5.83	5.75	5.77	5.17	5.26	5.20
k	3	-	5.75($\pm$0.94)	5.93	5.94	5.95	5.26	5.47	5.41
	7	-	5.79($\pm$1.17)	5.91	5.92	5.93	5.17	5.55	5.63
	1	1.50	1.78($\pm$0.81)	1.83	2.24	2.25	1.11	1.73	1.66
ka	3	-	1.76($\pm$0.92)	1.81	2.29	2.30	1.12	1.80	1.65
	7	-	1.68($\pm$1.15)	1.73	2.35	2.35	1.15	1.83	1.60
	1	0.50	0.85($\pm$0.80)	0.88	1.39	1.40	0.18	0.92	0.77
q	3	-	0.84($\pm$0.91)	0.87	1.46	1.47	0.20	0.96	0.77
	7	-	0.75($\pm$1.13)	0.78	1.54	1.54	0.23	0.98	0.70
	1	-0.20	0.23($\pm$0.76)	0.24	0.84	0.84	-0.30	0.92	0.18
v	3	-	0.22($\pm$0.87)	0.22	0.90	0.91	-0.29	0.96	0.17
	7	-	0.12($\pm$1.08)	0.12	0.97	0.97	-0.28	0.98	0.08
	1	-0.30	-0.02($\pm$0.67)	-0.02	0.55	0.55	-0.51	0.08	-0.05
w	3	-	-0.04($\pm$0.77)	-0.04	0.61	0.61	-0.50	0.08	-0.07
	7	-	-0.13($\pm$0.95)	-0.14	0.66	0.66	-0.50	0.08	-0.16

Table 4.3 Comparison of fits at different resolutions. **WMAP:** from Table 3 of Bennett et al. (2003), using a Finkbeiner Hα template *corrected for dust*. **Default:** diagonal covariance approximation, FDS dust, DDD ($f_d = 0.0$) Hα, Kp2 sky coverage (extended by smoothing); uncertainties from simulations. **ff05:** fits with DDD corrected for dust using $f_d = 0.5$. **SHASSA:** using the Finkbeiner Hα only over SHASSA coverage; **WHAM:** using DDD ($f_d = 0.0$) Hα only over WHAM coverage. **kp0:** using the Kp0 mask (extended by smoothing) and otherwise default. **constr:** constrained fits in five bands with power law Hα using $\beta = -2.1$ and synchrotron $\beta = -3.0$, otherwise default. **minusHalpha:** fit of dust and synchrotron with Hα subtracted at theoretical levels.

4.4.2   DMR comparison

A detailed foreground analysis using the *COBE*-DMR data is presented in Banday et al. (2003). The method used is a similar template fitting to that described above but performed in harmonic space and using a formalism that takes into account the sky cut as described in Górski et al. (1996). Table 4.4 shows comparisons with DMR fits described in Banday et al. (2003). Using the DMR standard cut on the *WMAP* data smoothed to 7° gives comparable results in the Ka and W bands (33 and 93.5 GHz, respectively) compared to DMR's 31.5 and 90 GHz bands. But again, there are significant differences among the fit values using different sky cuts.

The full sky free-free fits done by the *WMAP* team as well as the region fits described in section 5 are all lower than predicted (see, e.g., Dickinson et al. 2003). The *WMAP* first-year fit value of 4.6 $\mu K/R$ at 23 GHz is less than half the expected value of 11.4 assuming an electron temperature of 8000K, though the *WMAP3* results give values near 6.5 $\mu K/R$. Note that they use an Hα template corrected for dust absorption, which means that these values cannot be compared directly with the DMR numbers or those in Chapter 5, where no such correction is applied. The difference can be seen in the "Default" *versus* "ff05" columns of Table 4.3.

Fits using DMR data in Banday et al. (2003) were roughly consistent with electron temperatures of 5000-9000 K, and the *WMAP* data have been compared to the DMR data and found to be generally consistent. So it is not obvious where such a difference is coming from, but as discussed above, the fit results are sensitive to both the smoothing and the sky cut, likely due to mismatches between the templates and the microwave sky near the Galactic plane. As mentioned in § 4.3.3, such mismatches affect the stability of the results so that a small change, e.g. in the mask or the noise properties, may lead to a large change in the resulting fit values.

| Band | FWHM | DMR | | | *WMAP* | | |
		15°	DMR	30°	15°	DMR	30°
				dust			
	1	-	-	-	6.0	3.2	2.8
ka	3	-	-	-	3.7	3.1	2.8
	7	6.5(±0.7)	5.0(±1.7)	4.0(±2.0)	3.2	3.0	2.9
	1	-	-	-	1.3	1.2	1.0
w	3	-	-	-	1.1	1.3	1.1
	7	2.0(±0.5)	2.5(±1.0)	3.0(±1.3)	1.0	1.2	1.2
				free-free			
	1	-	-	-	5.2	4.2	6.7
ka	3	-	-	-	6.4	2.7	4.4
	7	5.0(±0.9)	3.5(±2.0)	6.5(±3.5)	6.0	2.3	1.8
	1	-	-	-	0.6	-0.2	0.2
w	3	-	-	-	0.7	-0.5	-0.5
	7	2.0(±0.4)	1.0(±1.0)	-2.0(±2.5)	0.7	-0.6	-1.9
				synchrotron			
	1	-	-	-	-0.6	1.1	1.0
ka	3	-	-	-	0.8	1.2	1.1
	7	1.0(±0.8)	1.0(±1.0)	1.0(±1.5)	1.1	1.3	1.3
	1	-	-	-	-0.5	-0.4	-0.3
w	3	-	-	-	-0.4	-0.4	-0.3
	7	-0.5(±0.5)	-0.5(±0.5)	-1.0(±0.7)	-0.4	-0.4	-0.1

Table 4.4 Comparison of *WMAP* and DMR fits. For each, three masks are used to exclude the Galactic plane: a symmetric 15° cut, the DMR mask, or a symmetric 30° cut. Approximate errors for *WMAP* fits can be seen in the default column of Table 4.3, though they will be larger or smaller for these cuts.

4.4.3   *WMAP*-team analysis

The *WMAP* team use similar template fitting techniques to create foreground subtracted maps from which power spectra are extracted for cosmological parameter estimation. As described in §6 of Bennett et al. (2003c), the full sky outside of the Kp2 mask is fit to the three templates plus a constant offset component. (The analysis in Hinshaw et al. (2006) is similar.) For the power spectrum analysis (in contrast to the numbers given in Table 3 of the Bennett et al. 2003c), the fit is performed for Q, V and W bands with power-law constraints applied to the free-free and synchrotron templates. We perform similar fits, both constrained and unconstrained, and obtain comparable results as seen in Table 4.3.

Given the spectral variation expected in both free-free and synchrotron emission, the template approach is far from ideal, but it is a practical necessity. In § 4.5, we examine the effect of such spectral variations on the power spectrum after the template-based foreground subtraction.

For the 3-year results, the *WMAP* team improve the foreground subtraction by using the K–Ka map instead of the Haslam synchrotron template. This map traces any low-frequency foreground not traced by the dust and Hα templates, regardless of its source. Using K–Ka gives an improved cosmic component, with less residual foreground emission, but the results are difficult to interpret in terms of studying the foregrounds themselves. The K–Ka template is then an unknown combination of synchrotron, free-free, and possibly anomalous dust emission, and it is difficult to disentangle these components.

Bennett et al. (2003c) attribute the anomalous, low-frequency, dust-correlated emission to synchrotron with a hard spectral index. But other analyses support the hypothesis that this emission is due to the electric dipole emission of spinning dust grains. (See Draine & Lazarian 1999 and references therein.)

The analysis described in Chapter 5 will show that the anomalous component has a spectral index between the K and Q-bands of $\beta \approx -2.85$, though this hardens (within the uncertainties) to -2.6 using the *WMAP* 3-year data (§ 7.1). If indeed this anomalous emission is due to hard synchrotron coming from dusty regions of the ISM, then we can extrapolate it to the 408 MHz regime and subtract it from the Haslam map at that frequency. (There is uncertainty in the zero-points, so this comparison is only approximate.) Figure 4.3 shows that with an index as steep as -2.7, this hypothesis is ruled out by the overprediction at 408 MHz (indicated by the blue regions of negative flux in the residuals in Fig. 4.3 (c)). With an index of -2.6, the overprediction is not as obvious, but it remains a problem. This lends support to the spinning dust theory that predicts a turnover in the spectrum around 20 GHz (see Draine & Lazarian 1998a).

4.4.4   Finkbeiner analysis

Finkbeiner (2004a) examined the foreground components using a template approach that used a slightly different χ^2 minimisation technique. This analysis assumes that the free-free, thermal dust, and synchrotron components are well understood and constrains them in order to focus on the anomalous component and such residuals. We reproduce this result

(a) 408 MHz

(b) FDS 94 GHz

(c) resid, $\beta = -2.7$

(d) resid, $\beta = -2.6$

$-10.$K ▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬ $50.$K
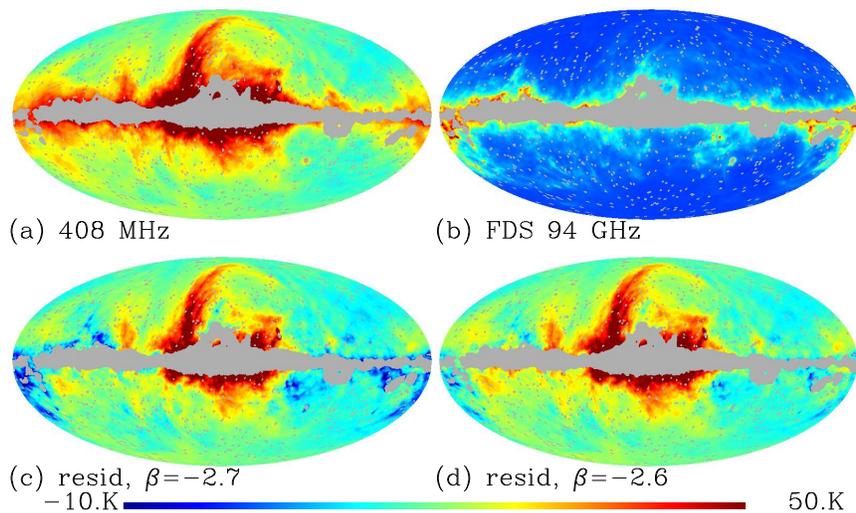
Figure 4.3 Extrapolation of a hypothetical hard synchrotron component (neglecting zero-point uncertainties); (a) the 408 MHz map of Haslam et al. (1982); (b) FDS dust prediction at 94 GHz extrapolated to 408 MHz assuming a spectral index of $\beta = -2.7$; (c) the residual soft synchrotron component, i.e. (a) minus (b); (d) as (c) but assuming $\beta = -2.6$.

with the fitting methods described above in Fig. 4.4. One key difference is that Finkbeiner subtracts the ILC estimate of the CMB and then expects only noise in the residuals. But the ILC contains residual foreground emission as described in § 5.A. For example, if the free-free emission does follow a power law with $\beta = -2.1$, then subtracting the ILC from the maps means that you end up with only 99% of that emission at K-band and only 84% of it at W-band. This will introduce errors in the fitting.

The top row shows the data in each band, and in the last column, the ILC estimate for the CMB component. In each subsequent row is shown the map above minus the component in the last column above. First the CMB is subtracted, then the free-free, synchrotron, thermal dust, and anomalous dust. The last row shows the residuals, which would be only noise if the templates were correct but which show some structure remains at low frequencies. The K and Ka-band residuals are averaged to form what Finkbeiner (2004a) called the "free-free haze".

The free-free emission is constrained to follow a power-law index of $\beta = -2.1$, and the amplitude fit to the data. Likewise, the synchrotron is constrained by $\beta = -3.0$. The thermal dust component as well as the ILC estimate for the CMB are both subtracted before the fit, the former using the FDS dust template and an index of $\beta = 1.7$. This leaves the anomalous component, which Finkbeiner fits using $FDS \times T_{dust}^2$, a template he describes as "empirically" motivated and with no clear theoretical motivation. (We discuss the factor T^2 further in Chapter 5.) The residuals on the bottom row of Fig. 4.4 show that there remains emission concentrated around the galactic centre region (and some oversubtraction elsewhere along the plane.) Due to the spectral behaviour of this "haze", Finkbeiner originally attributed it to free-free emission and creates a template by simply

averaging the K and Ka residuals. (Finkbeiner 2004b offers an alternative explanation for this haze as synchrotron emission from highly relativistic e^+e^- pairs produced by dark matter particle annihilation.)

This figure is particularly useful in showing the morphology of each component and of the data in each band for comparison. What it tells us about the foreground components themselves is less clear. Finkbeiner simply attributes to a mysterious haze any failure in the templates to accurately model the data, though obviously this cannot explain regions where the fit residuals are negative. It is these residuals, particularly visible in the K-band, that illustrate why the template fits will be sensitive to both smoothing and sky cuts.

As to the nature of these residuals, further study is certainly needed. It is not clear how much they may represent a new component spatially distinct from those traced by any of the templates (e.g., hard synchrotron) or whether they are simply residuals due to spectral variation or relatively isolated inaccuracies in those templates that drives the fits to be under- or overestimated.
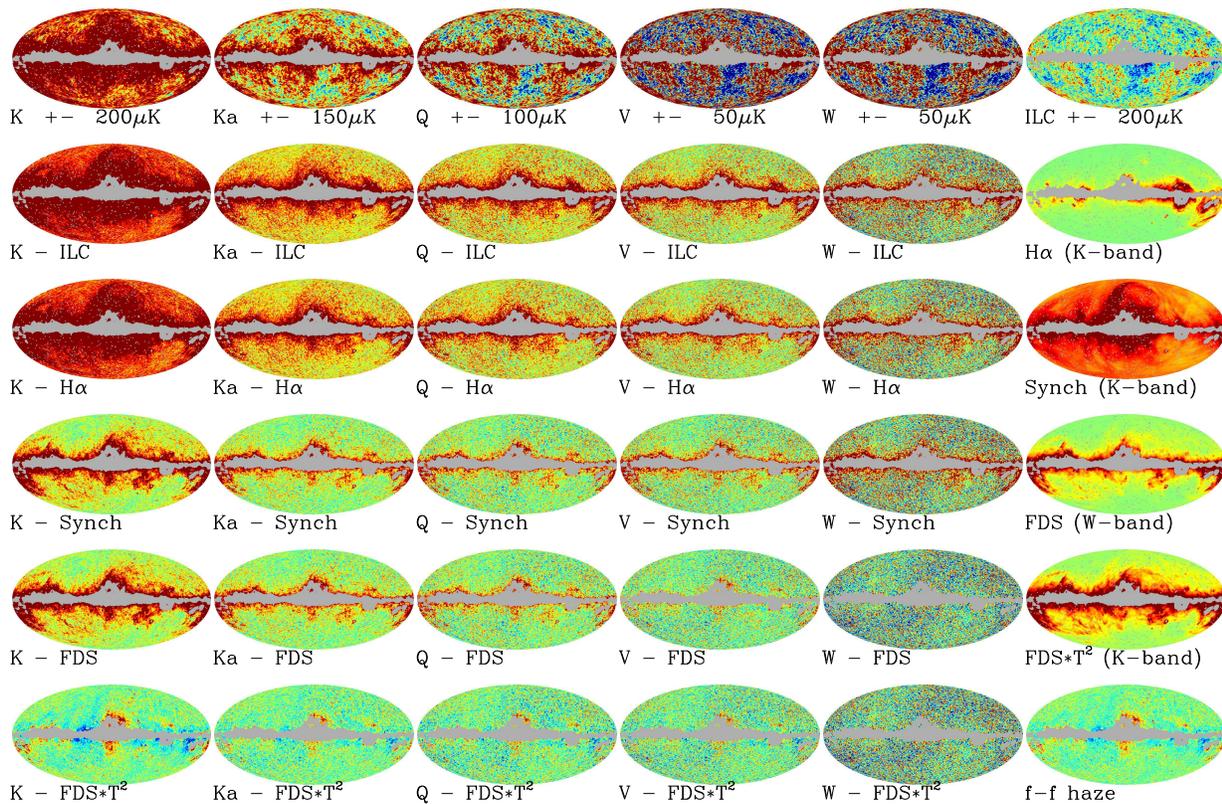
Figure 4.4 Peeling away foreground components one at a time from the *WMAP* data using template fitting. See text.

4.5 Spectral variations in synchrotron

4.5.1 Maps

Figure 4.5 shows further difficulties of the template approach. Each map plotted is the difference between the solution found using the full sky and that found using only the quadrant shown. In other words, the templates are fit to the data on that quadrant and the residuals, which should be purely CMB, compared to those found where the fits are found over the full sky (always outside the Galactic mask). Positive signal (red) represents where either the full-sky fit oversubtracts or the quadrant fit undersubtracts, and negative (blue) the reverse. These residuals can then be compared to the map of microwave sky features in Fig. 1.9.

Compare, for example, fits to the northern and southern hemispheres in the galactic reference frame, plots (a) and (d). Clearly, there is a different spectral behaviour in the two hemispheres, likely dominated by the north polar spur. The full-sky fit is effectively finding an average, while the hemisphere fits are able to reduce the residuals. The same effect is seen wherever the spur is included in the quadrant being fit. Region (c) shows a problem in the Galactic gradient. In region (k), the residuals look like mismatches near the Galactic centre resembling the "haze" described in § 4.4.4 and in the free-free near the Gum nebula, while in (l), there are strong residuals around the Orion complex. All of these residuals are due to the fact that each foreground component has a different spectral behaviour in these different clouds.

4.5.2 Pseudo-C_ℓs

We would like to estimate the level of foreground contamination that remains in the resulting power spectrum due to an analysis that uses such templates therefore fails to take into account the spectral variation of each across the sky. We start as follows:

- Following the *WMAP* team analysis in the first year data release, we fit the FDS dust, Finkbeiner Hα, and Haslam templates to the Q-, V- and W-band maps outside the Kp2 cut. (Unlike the *WMAP* team analysis, we do this in pixel space using a diagonal approximation to the signal-plus-noise covariance matrix. We also use data and templates that have all been smoothed to a common resolution of 1° FWHM and downgraded to HEALPix $N_{side} = 128$. This should not affect the point of this analysis, since both analyses use the same templates.) The Hα and synchrotron fits are constrained to follow power laws with indices of -2.15 and -3, respectively.

- We then divide the sky into hemispheres and quadrants in each of three reference systems: that defined by the Galactic plane, that which maximises the power asymmetry (Eriksen et al., 2004b; Hansen et al., 2004b), and that defined by the Ecliptic plane. In each of these sectors, we perform the same fit using Q, V, and W and the four foreground components, with the same spectral constraints.
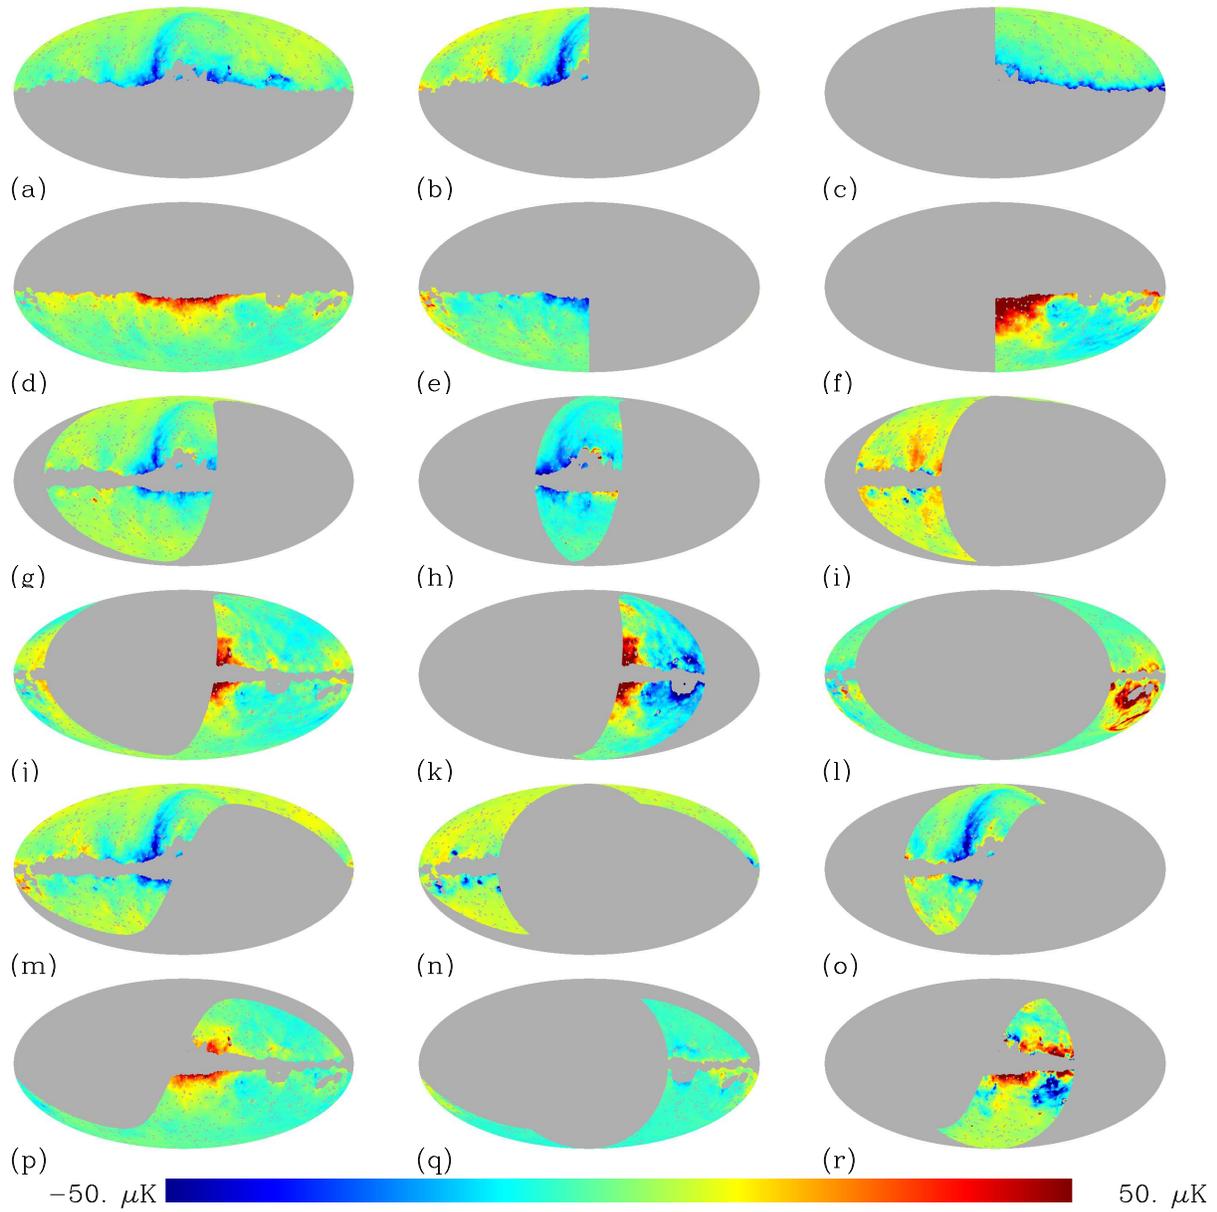
Figure 4.5 Difference maps showing the Q-band residuals from fits to the given sector minus the residuals from fits to the full sky.

- We then calculate the so-called 'pseudo-C_ℓ' power spectra—i.e., that calculated over the incomplete sky—for the fit residuals in each of these sectors and compare them to the residuals in that sector from the full sky fits.

The resulting differences in the residuals have two sources. On a smaller region of sky, a least-squares fit can find a different solution with fewer residuals in that region than there would be if the fit were over the full sky. This can be quantified using a set of simulations analysed the same way. The second effect is what we are looking for, the changes in the fits due to spatially varying behaviour of the foregrounds that is not taken into account in a template fit that determines one correlation coefficient for the sky.

The differences between the pseudo-C_ℓs are shown in Figure 4.6 as solid lines colour-coded for each band. (We plot $\Delta C_\ell = C_\ell^{\text{quad}} - C_\ell^{\text{full}} = < a_{\ell m}^{\text{quad}*} a_{\ell m}^{\text{quad}} > - < a_{\ell m}^{\text{full}*} a_{\ell m}^{\text{full}} >$, not $< (a_{\ell m}^{\text{quad}} - a_{\ell m}^{\text{full}})^* (a_{\ell m}^{\text{quad}} - a_{\ell m}^{\text{full}}) >$. This is therefore not the power spectrum of the *differences* between the fits, but the difference in the power spectra. Rather than taking into account the phases of the differences, then, this is meant to give an idea of how the resulting average power changes.) As mentioned, it is expected that in smaller regions, least-squares fits will give a different solution, often with fewer residuals. Therefore, we use the same procedure with 1000 simulations and use the average and variance to characterise this effect, shown as the pale coloured band for each frequency. In that figure, if, for example, the solid red curve representing the Q-band residuals were significantly outside the pink region, this would be an indication that the data contain spectral variations that do change the fits in different quadrants more than expected.

The results show that the fits in different sectors of the sky in general do not differ from the fits over the full sky more than one would expect, though there are a few outlying points in the Q-band. In particular, the biggest contamination appears to come from the synchrotron to the south and east of the Galactic centre or the Gum nebula, plots (f), (j), (k), (l), and (r). The corresponding maps of the residuals in the Q-band are shown in Fig. 4.5.

Any such effects remain much smaller than cosmic variance itself, but it is still unclear from this simple analysis what effect they could have on the parameter estimation. The *WMAP* team use only the V and W bands for parameter fitting, because as seen in these figures, the Q band (in red) shows the highest level of foreground contamination. For the 3-year analysis, they improve this situation by using the K–Ka difference map as a synchrotron template. This then accounts for much of the spectral variation in the synchrotron component and leads to much improved fits.

4.5.3 Simulations

We also create simulations where the synchrotron template is added with a different spectral index in each pixel. The first-year analysis by the *WMAP* team concluded that the anomalous component was essentially dust-correlated hard synchrotron. We approximate this variation by computing the index at each pixel derived from comparison of the *WMAP* K-band map to the Haslam 408 MHz map. The former has been corrected for the CMB
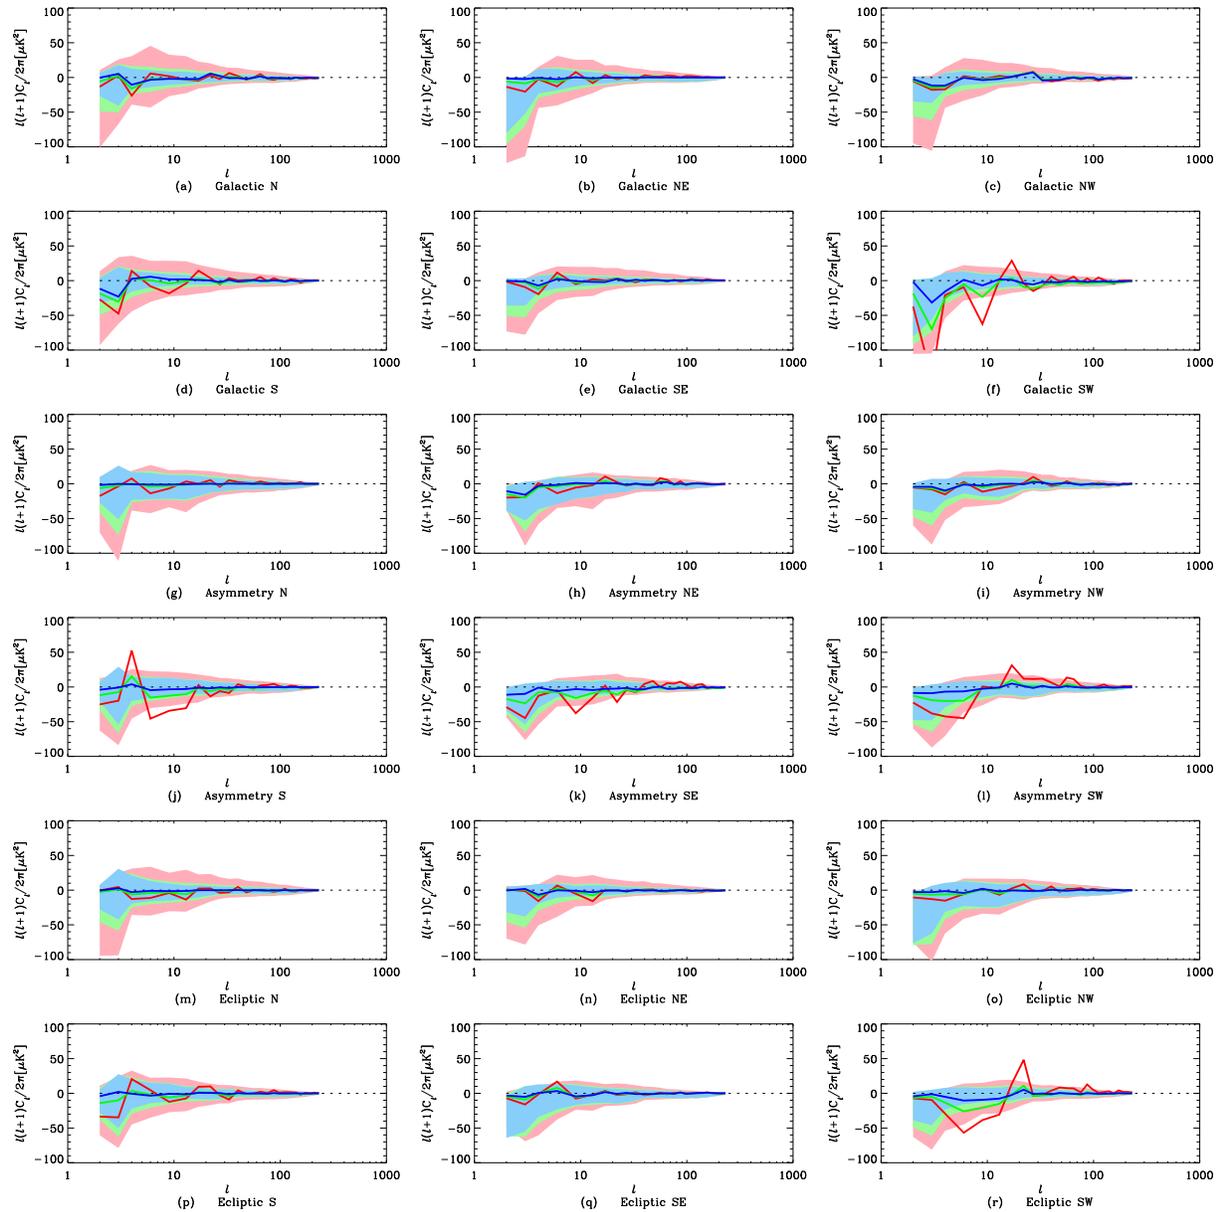
Figure 4.6 Difference in pseudo-C_ℓ power between residuals where templates were fit over the given sector and where fit over the full sky. The corresponding maps are shown in Fig. 4.5. Subtitles indicate the reference coordinate system: Galactic, Asymmetry (i.e. oriented such that the power asymmetry described in Eriksen et al. 2004b; Hansen et al. 2004b is maximised between the northern and southern hemispheres), and Ecliptic. Hemispheres or quadrants for the North, South, Northeast, etc. are indicated by "N", "S", "NE", etc. *WMAP* Q-band is shown in red, V-band in green, and W-band in blue. In all cases, the Kp2 mask is also applied. For each colour, the solid line shows the difference in power from the data, while the shaded regions show roughly the expected differences determined from a set of simulations. Note that for the range covered by the data, the cosmic variance uncertainties are larger than the y-axis range. The data are binned following the *WMAP* binning.

ILC estimate (ignoring the foreground aliasing described in § 5.A) and the free-free (DDD) emission and assuming the thermal dust in K-band is negligible. (Zero-point uncertainties are ignored. Changing the zero-point will change only the average spectral index in the simulations and then simply shift the fit results, and what we are interested in is not their absolute value but the residuals do to the variation.) This is meant only to give an approximate idea of the amount and morphology of spectral variation expected if indeed the K-band is dominated by synchrotron. (In fact, as discussed above and further in Chapter 5, the dominant component in the K-band is more likely to be due to spinning dust rather than hard synchrotron, but for this exercise, we are interested in assessing the effect of such a possible spectral variation.) Figure 4.7 (a) shows the result. We then use this index map (smoothed to remove noise and any ill-defined regions where the correlation is negative) to generate simulations with such spectrally varying synchrotron emission added.

We then look at the effect this has on the resulting power spectrum derived from constrained fits following the *WMAP* first-year analysis. First we do the constrained fit over the full sky outside the Kp2 cut, where free-free and synchrotron are constrained to follow the usual power laws and dust is free to fit in each band. From the fit residuals, we compare the pseudo-C_ℓ power spectra obtained from simulations with a fixed synchrotron index to those from the varying index model. The resulting average difference in the power spectra is shown in Fig. 4.7 (b), where even in the worst band, Q, the differences are a small fraction of the cosmic variance errors.

We conclude, then, that at angular scales above $1°$, systematic errors due to foreground residuals at reasonable expected levels in the high frequency bands remain quite small in comparison with the statistical errors due to cosmic variance. But again, further work is needed to determine what impact this would have on parameter estimation. The differences may be small, but they are systematic in ℓ and could bias parameters that depend on large-scale structure. Eriksen et al. (2006b) reanalyse the *WMAP* 3-year data and find a similar problem at large-scales related to the mask that has just such an effect on the low-ℓ power and affects the estimate of n_s.

4.6   Conclusions

This chapter has described several of the limitations of template fitting in foreground analysis. In particular, the spectral variations across the sky of the emission components make it difficult to construct an accurate template from observations at other frequencies. Template fits are then inaccurate and will vary depending on the resolution and on what regions of the sky are included in the fit. The residuals that may remain in the power spectra may be small relative to the cosmic variance errors, but it is not clear that the results of parameter estimation remain unbiased (see Eriksen et al. 2006b for example). It is likewise difficult to predict the influence of such residuals on calculations of other statistics such as the multipole vectors, the low-ℓ alignments, etc. But one can always simply examine the variation of any measurement with frequency to characterise the possible influence of foregrounds. In a likelihood analysis, one can marginalise over uncertainties in the template

fits, which effectively carries them through to be included in the calculated errors for the resulting power spectrum.

An improved understanding of the foregrounds is still needed for temperature analysis, and the problem only gets worse for polarisation analysis, where the cosmic signal no longer dominates the foregrounds even at high latitudes. The following chapter presents a study of the foregrounds in small regions of the sky in order to examine such spectral variations. This analysis represents only the first step toward a better understanding of the foregrounds that would eventually allow a cleaner separation of them from the cosmic component. The focus, however, will be on the foregrounds themselves, which are of interest in their own right for the insight they give into the emission processes of the ISM.
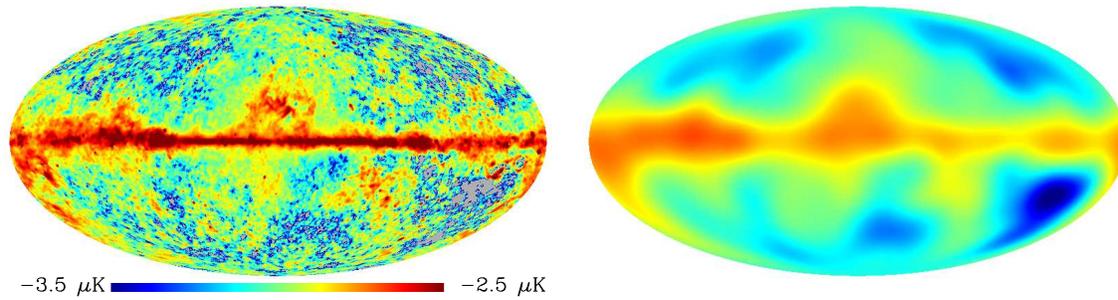
(a)
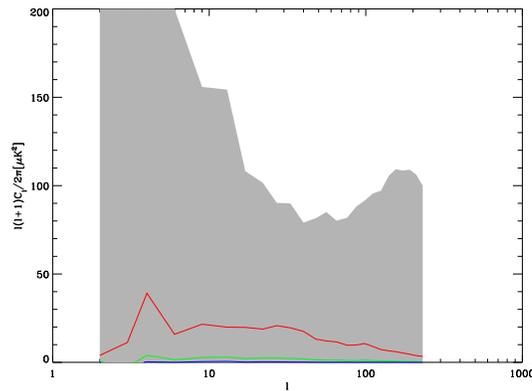
(b)

Figure 4.7 Map (top left) showing an estimate of the synchrotron spectral index deduced from comparing the Haslam 408MHz map to the *WMAP* K-band estimated synchrotron component from subtracting the ILC CMB estimate and the predicted free-free emission (if $T_e = 8000K$). (This attributes any anomalous emission then to synchrotron.) The grey regions show where the K-band synchrotron estimate shows negative emission. The top right map is the same but smoothed with beam of 20° FWHM. Beneath are difference spectra for simulations with varying spectral index compared to with constant index. Red, green, and blue show Q-, V-, and W-band, respectively, while the grey band shows the cosmic variance uncertainties.

Chapter 5

A Determination of the Spectra of Galactic Components Observed by *WMAP*

Note: the following chapter was published as Davies et al. (2006). T. R. Jaffe performed all the cross-correlation analysis that is used after demonstrating that the TT fitting technique does not work well even in regions dominated by one foreground component. The text of Sections 5.4.1 and 5.B, Tables 5.3, 5.4, 5.5, 5.7, 5.8, and 5.9, and Figures 5.4 and 5.5 were done by T. R. Jaffe. Updated Tables and Figure using WMAP 3-year *data are given in* § *7.1.*

5.1   Introduction

The all-sky observations by the Wilkinson Microwave Anisotropy Probe (*WMAP*: Bennett et al. 2003b) provide unprecedented data on Galactic emission components in the frequency range 23 to 94 GHz, with a high precision estimate of the CMB power spectrum. As CMB studies move to higher precision it becomes necessary to determine the various components of Galactic foreground emission to higher and higher accuracy. As an example of this requirement, the question of the glitch in the power spectrum at multipole $\ell = 40$ (Hinshaw et al., 2003b)[1] is debated and various sources have been proposed including a Galactic origin. In an analysis of structures in the *WMAP* CMB map derived after removing Galactic foregrounds, Hansen et al. (2004a) find a range of asymmetric structures on scales of tens of degrees. References to the many analyses which have detected asymmetries or non-Gaussian structures may be found in this paper.

The role of a Galactic component in the above scenarios is unclear, but cannot be ruled out. Of particular relevance to this discussion is the fact that each of the foreground components has a spectral index that varies from one line of sight to another so using a single spectral index can lead to significant uncertainties in the corrections required.

---

[1]The glitch at $\ell = 40$ is still present in the new 3-year *WMAP* data Hinshaw et al. (2006).

It is obvious that the foregrounds that can be studied with *WMAP* data are of interest in their own right. In comparing the maps at the 5 frequencies of *WMAP* (23, 33, 41, 61 and 94 GHz) with the free-free, synchrotron dust templates it is possible to clarify important properties of the emission. For the free-free one can derive the electron temperature distribution in the brighter regions of the Galaxy (near the Galactic plane and in the Gould Belt system). In the case of the synchrotron emission significant information on the spectral index variations across the sky can be established. Even more important, data are available to help clarify the FIR-correlated emission.

The greatest insight is likely to come from *WMAP* for the dust-correlated emission. The present situation is far from clear. The dust-correlated emission was observed in the *COBE*-DMR data (Kogut et al., 1996) but was originally thought to be free-free. Leitch et al. (1997) suggested the excess emission could be hot (10^6 K) free-free emission. Draine & Lazarian (1998a,b) moved attention to the dust itself as the emission source through dipole emission from spinning grains, referred to as "spinning dust". They also considered an enhancement to the thermal emissivity produced by thermal fluctuations in the grain magnetisation (Draine & Lazarian, 1999), but this explanation is less favoured by the data. The lower frequency Tenerife results show that it was incompatible with free-free (Jones et al., 2001) while the Tenerife data at 10, 15 and 30 GHz (de Oliveira-Costa et al., 1999, 2000, 2002) provided strong evidence for dust at intermediate Galactic latitudes emitting a spectrum of the form expected by spinning dust. A reanalysis of the intermediate and high Galactic latitude data taken by *COBE* and supplemented by 19 GHz observations (Banday et al., 2003) led to similar conclusions. Finkbeiner et al. (2004) used 8.35 and 14.35 GHz data in combination with *WMAP* data and found a similar spectrum for a different environment in the Galactic ridge (effectively $|b| < 4 \deg$) in the inner Galaxy ($l \sim 15 \deg$ to $45 \deg$); the effect was not so clear-cut in the central regions of the Galaxy ($|l| < 7.\deg 5$). The first targeted search was carried out by Finkbeiner et al. (2002) where they found a rising spectrum over the $5 - 10$ GHz range for 2 diffuse clouds, which was interpreted as tentative evidence for spinning dust. Finkbeiner (2004a) also considers intermediate latitudes in the *WMAP* data in which he finds an anomalous component compatible with spinning dust or a hot gas (10^6 K) component, but inconsistent with a traditional free-free spectral index. New results from the Cosmosomas experiment (Watson et al., 2005) also detect strong anomalous emission from the Perseus molecular cloud with a rising spectrum in the range $11 - 17$ GHz, suggestive of spinning dust. In contrast the *WMAP* team (Bennett et al., 2003c) gave a radically different interpretation of the dust-correlated emission, considering it to be synchrotron emission supposedly from star-forming regions associated with the dust; again this analysis was at intermediate and high Galactic latitudes.

Our present approach is to identify regions away from the Galactic plane which are expected to be dominant in one of the three foreground components, free-free, synchrotron or dust and to derive the spectrum for each component. Five regions covering angular scales of $3 \deg$ to $20 \deg$ were chosen for each component, based on foreground template maps, making 15 in all. Such a selection is intended to minimise the potential cross-talk between the various physical components. Moreover, by considering regions which are

largely dominated by well-known objects selected at a given frequency, it is likely that the spectral behaviour is uniform over the region in question thus supporting the use of a template based comparison. We are also interested in evaluating spectral variations over the sky, and intend that any region to region scatter should reflect this. Two complementary analyses of each region were considered.

The classical T-T plot approach can provide a detailed look at the distribution of the data. In order to minimise cross-talk with the CMB background, each of the five bands must be corrected for this component as described by Bennett et al. (2003c) employing an internal linear combination (ILC) of the *WMAP* sky maps. For the high-latitude sky considered here, this corresponds to a single set of linear coefficients for each of the 5 frequencies. The ILC map at high latitudes is therefore simply given by 0.109K −0.684Ka −0.096Q +1.921V −0.25W. Unfortunately, subtracting the ILC CMB map changes the relative levels of foreground emissions at each frequency depending on the spectral characteristics of a given component. This "aliasing" effect (see appendix5.A), together with other potential cross-talk between the foreground components, renders the method useful only for visualisation and qualitative analysis. Instead, all quantitative results in this paper are derived using a cross-correlation (C-C) method, similar to the approach taken by Banday et al. (2003). The C-C analysis does not rely on a given CMB map. Instead, the CMB is taken into account internally by including a CMB component into the covariance matrix (see section 5.4.1), and the the various correlations are solved for simultaneously.

Section 5.2 describes the foreground templates used in this analysis while Section 5.3 gives the considerations for selecting the 15 regions for investigation. The cross-correlation analysis and results are presented in Section 5.4. The spectrum of each component is discussed further in Section 5.5. A comparison with the new 3-year *WMAP* data is described in Section 5.6 and overall conclusions given in Section 5.7.

5.2 Templates used in the analysis

The present analysis of the *WMAP* data seeks to quantify the Galactic foreground components of free-free, synchrotron and dust emission using appropriate templates of each component. The approach is similar to that of Banday et al. (2003) in their analysis of the *COBE*-DMR data but with the difference that the current work relates to selected areas rather than the full sky with the Galactic plane removed. Our analysis is made on an angular scale of 1 deg, the smallest that is feasible with the templates available, namely, 1 deg in the Hα free-free map, $0.^\circ85$ in the 408 MHz synchrotron map and $0.^\circ82$ in the K-band of *WMAP*. The basic properties of the main maps used in the current analysis are summarised in Table 5.1.

5.2.1   *WMAP* data

We use the 1st-year *WMAP* data (Bennett et al., 2003b) provided in the HEALPix[2] pix-elisation scheme, with a resolution parameter of $N_{side} = 512$, which can be obtained from the LAMBDA website[3]. The data consist of 5 full-sky maps covering the frequency range 23 GHz (K-band) up to 94 GHz (W-band); see Table 5.1. For our lower resolution analysis and to compare to the foreground templates, these maps (as well as all templates) are smoothed to a common resolution of 1° and converted to μK of antenna temperature. The smoothed maps are then downgraded to a HEALPix resolution of $N_{side} = 128$, with a total of 196608 pixels.

We convert from thermodynamic temperature to brightness temperature units i.e. the Rayleigh-Jeans convention[4]. This corresponds to a correction of 1.4 per cent at 23 GHz increasing to 25 per cent at 94 GHz. Bright point sources are masked using the mask templates provided by the *WMAP* team. They are based on various catalogues covering a wide range of wavelength domains masking almost 700 sources in total (see Bennett et al. 2003b for details); pixels within 0°.6 radius of a source are blanked. This operation typically removes ~ 10 per cent of the pixels in each region. Fainter sources not included in the mask are not expected to make a significant change in the results presented here. Any bright sources still remaining would be easily identified in the maps.

The effective centre frequency of each band depends on the continuum spectrum of the foreground being considered. We adopt the values given by Jarosik et al. (2003) which apply to the CMB blackbody spectrum. Reference to Page et al. (2003) shows that the effective frequencies for synchrotron and free-free respectively are 1.0 and 0.7 per cent lower while those of the thermal (vibrational) dust are 0.7 per cent higher. Thus using the frequencies appropriate for the CMB will not have a significant affect on our estimates of spectral index for the various foregrounds.

5.2.2   The Hα free-free template

The only effective free-free template at the intermediate and high Galactic latitudes used in the present study comes from Hα emission. In this analysis we use the all-sky Hα template described in Dickinson et al. (2003, hereafter DDD) which is a composite of WHAM Fabry-Perot survey of the northern sky (Haffner et al., 2003) which gives a good separation of the geocoronal Hα emission and of the SHASSA filter survey of the southern sky (Gaustad et al., 2001). Baseline effects may be significant in the SHASSA data where information is lost on scales $> 10\,\mathrm{deg}$ due to geocoronal emission. To correct for the Galactic gradient with latitude, a baseline correction was applied assuming a cosecant law for declinations further south ($-30\,\mathrm{deg}$) where WHAM data are not present. On the angular scales of the WHAM data ($1\,\mathrm{deg}$) the sensitivity of both surveys are comparable

---

[2]http://www.eso.org/science/healpix/.

[3]http://lambda.gsfc.nasa.gov/.

[4]The conversion factor from thermodynamic units to brightness (antenna) units (the "Planck correction") is given by the derivative of the Planck function: $x^2 e^x/(e^x - 1)^2$ where $x = h\nu/k_b T_{CMB}$.

Dataset	Frequency/ Wavelength	Beamwidth (FWHM °)	Reference
Haslam	408 MHz	0.85	[1]
WMAP K	22.8 GHz	0.82	[2]
WMAP Ka	33.0 GHz	0.62	[2]
WMAP Q	40.7 GHz	0.49	[2]
WMAP V	60.8 GHz	0.33	[2]
WMAP W	93.5 GHz	0.21	[2]
FDS8 dust	94 GHz	0.10	[3]
Hα	656.2 nm	~ 1	[4]

Table 5.1 Properties of the main maps used in the analysis. References are [1]: Haslam et al. (1982); [2]: Bennett et al. (2003b); [3]: FDS; [4]: Dickinson et al. (2003).

at ~ 0.1 Rayleigh (R). Recently Finkbeiner (2003, hereafter F03) has produced an all-sky Hα map by including data from the VTSS filter survey (Dennison et al., 1998). This map contains structure down to 6 arcmin in scale, but has effectively variable resolution due to the different resolutions of the WHAM and SHASSA surveys. The differences between the Dickinson et al. and the Finkbeiner maps are determined to be less than 15 per cent in R over the common power spectrum range ($\ell = 2 - 200$). The largest discrepancies are apparent near the "cross-over" region of the datasets where baseline levels have been determined in a different way. In these regions, baseline uncertainties are typically ~ 1 R. For the majority of the sky, the baseline levels are tied to the WHAM data which contains baseline uncertainties of $\lesssim 0.1$ R (Haffner et al., 2003). The Hα solutions do not change appreciably when the Finkbeiner Hα map is used showing the similarity between the 2 templates.

When using the Hα map as a template for the free-free emission it is necessary to correct for the foreground dust absorption. Dickinson et al. (2003) used the 100 μm map given by Schlegel et al. (1998, hereafter SFD98) to estimate an absorption correction in magnitudes at the Hα wavelength of $A(\text{H}\alpha) = (0.0462 \pm 0.0035)D^T f_d$ where D^T is the SFD temperature-corrected 100 μm intensity in MJy sr^{-1} and f_d is the fraction of dust in front of the Hα in the line of sight. A value of $f_d \sim 0.5$ expected under the assumption that the ionised gas and dust are coextensive along the line of sight (i.e. uniformly mixed). Dickinson et al. (2003) find $f_d \sim 0.3$ and accordingly $A(\text{H}\alpha)$ is < 0.2 mag over most of the intermediate and high latitude sky where $D^T < 5$ MJy sr^{-1} ; at latitudes below ~ 5 deg the absorption is too high to make a reasonable estimate of the true Hα intensity. Banday et al. (2003) use *COBE*-DMR and 19-GHz data to place a upper limit of $f_d \lesssim 0.35$ assuming $T_e = 7000$ K. This confirms that zero correction is required for high Galactic latitudes ($|b| \gtrsim 20°$). It is worth noting that for the *WMAP* 1-year analysis (Bennett et al., 2003c), which uses the Finkbeiner Hα map, $f_d = 0.5$ was adopted for the entire sky. At high Galactic latitudes, the dust column density is small enough for this to have almost negligible effect ($A(\text{H}\alpha) \lesssim 0.1$); the variance of the Hα map corrected by this value is 20-30 per cent larger than for an uncorrected map, depending on the galactic mask employed.

We therefore adopt the uncorrected template in the following analysis.

The conversion of dust-corrected Hα intensities to emission measure (*EM* in units of cm^{-6} pc) and then to free-free emission is well-understood. The brightness temperature T_b can be related to EM using $T_b \propto T_e^{-0.35} \nu^{-2.1} \times EM$ and thus requires a knowledge of the electron temperature T_e of the ionised gas. For the *WMAP* bands K, Ka, Q, V and W this corresponds to 11.4, 5.2, 3.3, 1.4 and 0.6 μK R^{-1} respectively at $T_e = 8000$ K; see Dickinson et al. (2003) for details.

A number of estimates are available for T_e in regions of the Galaxy relevant to the present intermediate and high latitude study, namely at galactocentric distances $R \sim R_0$. Shaver et al. (1983) used RRLs from Galactic H$\textsc{ii}$ regions to establish a clear correlation of T_e with R; their result was

$$T_e\,(R) = (3150 \pm 110) + (433 \pm 40)\,R \tag{5.1}$$

The following similar relationship was found by Paladini et al. (2004) from a larger sample which contained many weaker sources

$$T_e\,(R) = (4170 \pm 120) + (314 \pm 20)\,R \tag{5.2}$$

At $R \sim R_0$, in the local region, these expressions indicate that $T_e = (7200 \pm 1200)$ K. It is possible that the T_e of diffuse H$\textsc{ii}$ emission at a given galactocentric distance may be different from that of the higher density H$\textsc{ii}$ regions on the Galactic plane. There are strong indications from observation and theory that the diffuse ionised gas will have a higher electron temperature than in the density bounded H$\textsc{ii}$ regions which contain the ionising stars (Wood & Mathis, 2004). RRLs give another route to identifying the free-free component of the Galactic foreground and may be useful at low Galactic latitudes when the Hα signal is heavily absorbed by foreground dust.

5.2.3   The dust template

Dust has two broadband emission components in the frequency range 10 to 1000 GHz. The anomalous emission is dominant at the lower end while the thermal (vibrational) component is responsible for the higher end. We will see that the anomalous component is the strongest for the *WMAP* frequencies of 23, 33 and 41 GHz while the vibrational component dominates at 94 GHz. We clearly need templates for both components if we are to quantify accurately the anomalous emission at 61 and 94 GHz.

COBE-DIRBE full-sky maps at 100, 140 or 240 μm with 0°.7 resolution have commonly been used as tracers of the thermal dust component (Kogut et al., 1996). However, the most sensitive full-sky map of dust emission is the 100 μm data at 6 arcmin resolution from *IRAS*. These data have been recalibrated using *COBE*-DIRBE data and reanalysed to give reduced artifacts due to zodiacal emission and to remove discrete sources (SFD98). In a preliminary analysis, we utilised the latter to help define specific dust fields of interest, and then to examine the dust emissivity of the 15 selected regions (see section 5.3). Fig. 5.1 shows T-T plots for one of the regions in K, Ka and Q-bands of *WMAP* against the
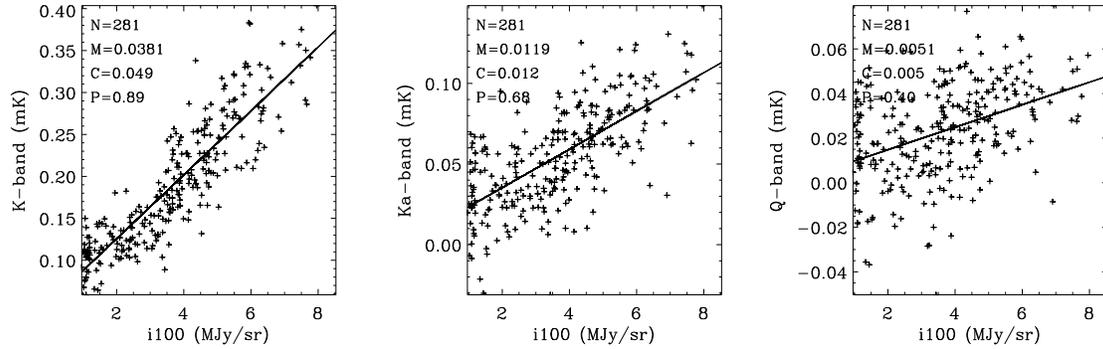
Figure 5.1 T-T plots for region 6 (middle row of Fig. 5.3) clearly showing the *WMAP* dust-correlated emission at K, Ka and Q bands against the 100 μm SFD98 map. The best-fitting line is plotted along with the number of pixels N, the y-intercept C, slope M and Pearson correlation coefficient P.

100 μm map. The dust-correlation is striking, particularly at K- and Ka-bands. Moreover, the emissivity (μK/(MJy sr^{-1})) was determined to vary by a factor of ~ 2.5 from cloud to cloud. However, this scatter was found, at least in part, to be driven by variations in the dust temperature. Finkbeiner et al. (1999, hereafter FDS) recognised the importance of this for predictions of the dust contribution at microwave wavelengths, and developed a series of models based on the 100 and 240 μm maps tied to *COBE*-FIRAS spectral data in the range 0.14 to 3.0 mm. The preferred model 8 (hereafter FDS8) has a spectral index $\beta \approx +1.7$ over the *WMAP* frequencies. For the work undertaken in this paper, we adopt the FDS8 predicted emission at 94 GHz as our reference template for dust emission. Note that in previous work, correlations were often referenced to the SFD98 100 μm template, in units of μK/(MJy sr^{-1}). To convert these values to correlations relative to FDS8, they should be divided by ~ 3.3.

5.2.4   The synchrotron template

The synchrotron emission of the Galaxy is best studied at low frequencies (< 1 GHz) where it is least contaminated by other emission (principally free-free emission from the ISM). Studies at these frequencies show that the temperature spectral index ($T_b \propto \nu^\beta$) has typical values of $\beta = -2.55$ and -2.8 (Lawson et al., 1987) at 38 and 800 MHz respectively. Reich & Reich (1988) demonstrated a range of spectral index values $\beta = 2.3 - 3.0$ between 408 and 1420 MHz, with a typical dispersion $\Delta\beta = \pm 0.15$. At higher frequencies β is expected to increase by ~ 0.5 due to radiation losses in the relativistic CR electrons responsible for the synchrotron emission.

We use the 408 MHz map by Haslam et al. (1981, 1982) which is the only all-sky map with adequate resolution (51 arcmin) at a sufficiently low frequency. It has a brightness temperature scale which is calibrated with the 404 MHz $8°\!.5 \times 6°\!.5$ survey of Pauliny-Toth
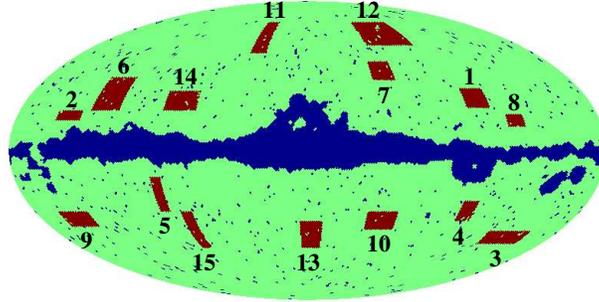
Figure 5.2 Mollweide full-sky map showing the 15 selected regions in red with the masked regions in blue, based on the Kp2 Galactic plane mask and ∼ 700 bright radio sources. The Galactic centre is in the middle and longitude increases to the left.

& Shakeshaft (1962). The 1.4 GHz northern sky map with a resolution of 35 arcmin made by Reich & Reich (1986) and the 2.3 GHz map at a resolution of 20 arcmin from Jonas et al. (1998) are employed to provide frequency coverage at GHz frequencies when assessing the spectral index of emission regions selected in the present study.

We note that spurious baseline effects have been identified in these surveys (Davies et al., 1996) which can affect β determinations for weaker features. In the current study we select stronger emission regions for comparison with the *WMAP* data; such strong regions are essential when extending the spectra to the highest map frequencies (94 GHz) where $\beta \sim -3.0$.

5.3   Field selection

The fields selected for study were chosen on the basis that one of the 3 foregrounds (free-free, dust or synchrotron emission) was dominant in each field. This criterion can be satisfied at intermediate Galactic latitudes, well away from the Galactic plane where the foregrounds are inevitably confused. As a result, our study will sample conditions in the Local Arm or adjacent spiral arms.

The angular scale of this study was determined by the largest beamwidths of critical elements in the data sets. The beamwidth of the 408 MHz survey is 51 arcmin, that of the WHAM Hα survey is 1 deg and that of the *WMAP* lowest frequency, 22.8 GHz, is 49 arcmin (Table 5.1). All the other data sets had a higher resolution. Accordingly a beamwidth of 1.deg 0 was chosen as the appropriate resolution and the analysis of the data sets was undertaken by smoothing to this resolution.

The choice of a 1.deg 0 resolution determined the size of structures in the *WMAP* maps which could be studied to best effect. The best signal-to-noise ratio would be achieved in structures on a scale of several resolution elements which therefore contain a number of

Table 5.2 Summary of 15 selected regions of sky.

Field No.	Dominant Emission	Longitude Range	Latitude Range	Description.
1	Free-free	$245\,\mathrm{deg}-260\,\mathrm{deg}$	$+21\,\mathrm{deg}-+31\,\mathrm{deg}$	Northern edge of Gum Nebula.
2	Free-free	$140\,\mathrm{deg}-155\,\mathrm{deg}$	$+15\,\mathrm{deg}-+20\,\mathrm{deg}$	Disc-like structure above Galactic plane.
3	Free-free	$200\,\mathrm{deg}-230\,\mathrm{deg}$	$-41\,\mathrm{deg}--48\,\mathrm{deg}$	Eridanus complex - within southern Gould Belt system.
4	Free-free	$250\,\mathrm{deg}-260\,\mathrm{deg}$	$-25\,\mathrm{deg}--35\,\mathrm{deg}$	Southern edge of Gum Nebula
5	Free-free	$90\,\mathrm{deg}-97\,\mathrm{deg}$	$-13\,\mathrm{deg}--30\,\mathrm{deg}$	Below plane in northern sky.
6	Dust	$118\,\mathrm{deg}-135\,\mathrm{deg}$	$+20\,\mathrm{deg}-+37\,\mathrm{deg}$	$l=125\,\mathrm{deg}$ dust spur, NCP region (the "duck").
7	Dust	$300\,\mathrm{deg}-315\,\mathrm{deg}$	$+35\,\mathrm{deg}-+45\,\mathrm{deg}$	Outer edge of northern Gould Belt system.
8	Dust	$227\,\mathrm{deg}-237\,\mathrm{deg}$	$+12\,\mathrm{deg}-+18\,\mathrm{deg}$	Above plane in southern sky.
9	Dust	$145\,\mathrm{deg}-165\,\mathrm{deg}$	$-30\,\mathrm{deg}--38\,\mathrm{deg}$	Orion region in southern Gould Belt.
10	Dust	$300\,\mathrm{deg}-320\,\mathrm{deg}$	$-30\,\mathrm{deg}--40\,\mathrm{deg}$	Below plane southern sky.
11	Synchrotron	$33\,\mathrm{deg}-45\,\mathrm{deg}$	$+50\,\mathrm{deg}-+70\,\mathrm{deg}$	Middle section of North Polar Spur.
12	Synchrotron	$270\,\mathrm{deg}-310\,\mathrm{deg}$	$+55\,\mathrm{deg}-+70\,\mathrm{deg}$	Outermost section of North Polar Spur.
13	Synchrotron	$350\,\mathrm{deg}-5\,\mathrm{deg}$	$-35\,\mathrm{deg}--50\,\mathrm{deg}$	Southern bulge in synchrotron sky.
14	Synchrotron	$70\,\mathrm{deg}-90\,\mathrm{deg}$	$+20\,\mathrm{deg}-+30\,\mathrm{deg}$	A "weak" northern spur.
15	Synchrotron	$76\,\mathrm{deg}-84\,\mathrm{deg}$	$-30\,\mathrm{deg}--50\,\mathrm{deg}$	A southern spur.

independent data points. Many of the features would be elongated or contain structure. Another selection requirement was that the field should contain a smooth background covering approximately half the area. This was essential in identifying the feature and separating its emission from underlying emission. The features studied typically had structure on scales of 3 deg to 10 deg.

5 fields were chosen in which each of the three Galactic foregrounds were dominant. The synchrotron fields were selected from the Haslam et al. (1982) 408 MHz map, the free-free fields from the Dickinson et al. (2003) Hα map and the dust fields from the SFD98 100 μm map. Table 5.2 lists the 15 fields of the present study. For each field the dominant emission and the Galactic coordinates are given along with a short description of the field.

Fig. 5.2 shows the position of the 15 selected regions overlaid on the Kp2 intensity mask and source mask ($\sim$ 700 sources in total) used by the *WMAP* team (Bennett et al., 2003c). Fig. 5.3 shows 3 regions (regions 4,6 and 11) with an ILC-subtracted K-band, Hα, 100 μm and 408 MHz data. The dominant foreground in each region (see Table 5.2) is clearly seen along with the correlated emission at K-band.

5.4   Cross-correlation analysis

The cross-correlation (C-C) method used here is a least-squares fit of one map to one or more templates. In the absence of any covariance information on the residuals, and assuming no net offsets (i.e., monopoles), this method is equivalent to the classical T-T method when only one template map ($T_1 \equiv \mathbf{t}$) is compared to the data ($T_2 \equiv \mathbf{d}$). The advantage of the C-C method is that we can fit several components simultaneously and that we can include information about the CMB thorough its signal covariance rather than having to correct for it. The issues of CMB subtraction and correlated components are discussed further in Appendices A & B.

5.4.1   Method

The cross-correlation measure, α, between a data vector, $\mathbf{d}$ and a template vector $\mathbf{t}$ can be measured by minimising:

$$\chi^2 = (\mathbf{d} - \alpha\mathbf{t})^T \cdot \mathbf{M}_{SN}^{-1} \cdot (\mathbf{d} - \alpha\mathbf{t}) = \tilde{\mathbf{d}}^T \cdot \mathbf{M}_{SN}^{-1} \cdot \tilde{\mathbf{d}} \tag{5.3}$$

where $\mathbf{M}_{SN}$ is the covariance matrix including both signal and noise for the template-corrected data vector $\tilde{\mathbf{d}} \equiv \mathbf{d} - \alpha\mathbf{t}$. Solving for α then becomes:

$$\alpha = \frac{\mathbf{t}^T \cdot \mathbf{M}_{SN}^{-1} \cdot \mathbf{d}}{\mathbf{t}^T \cdot \mathbf{M}_{SN}^{-1} \cdot \mathbf{t}} \tag{5.4}$$

To compare multiple template components $\mathbf{t}_j$, *e.g.*, different foregrounds, to a given dataset, the problem becomes a matrix equation. Górski et al. (1996) describe the method in harmonic space, which is fundamentally no different from pixel space. In the case where
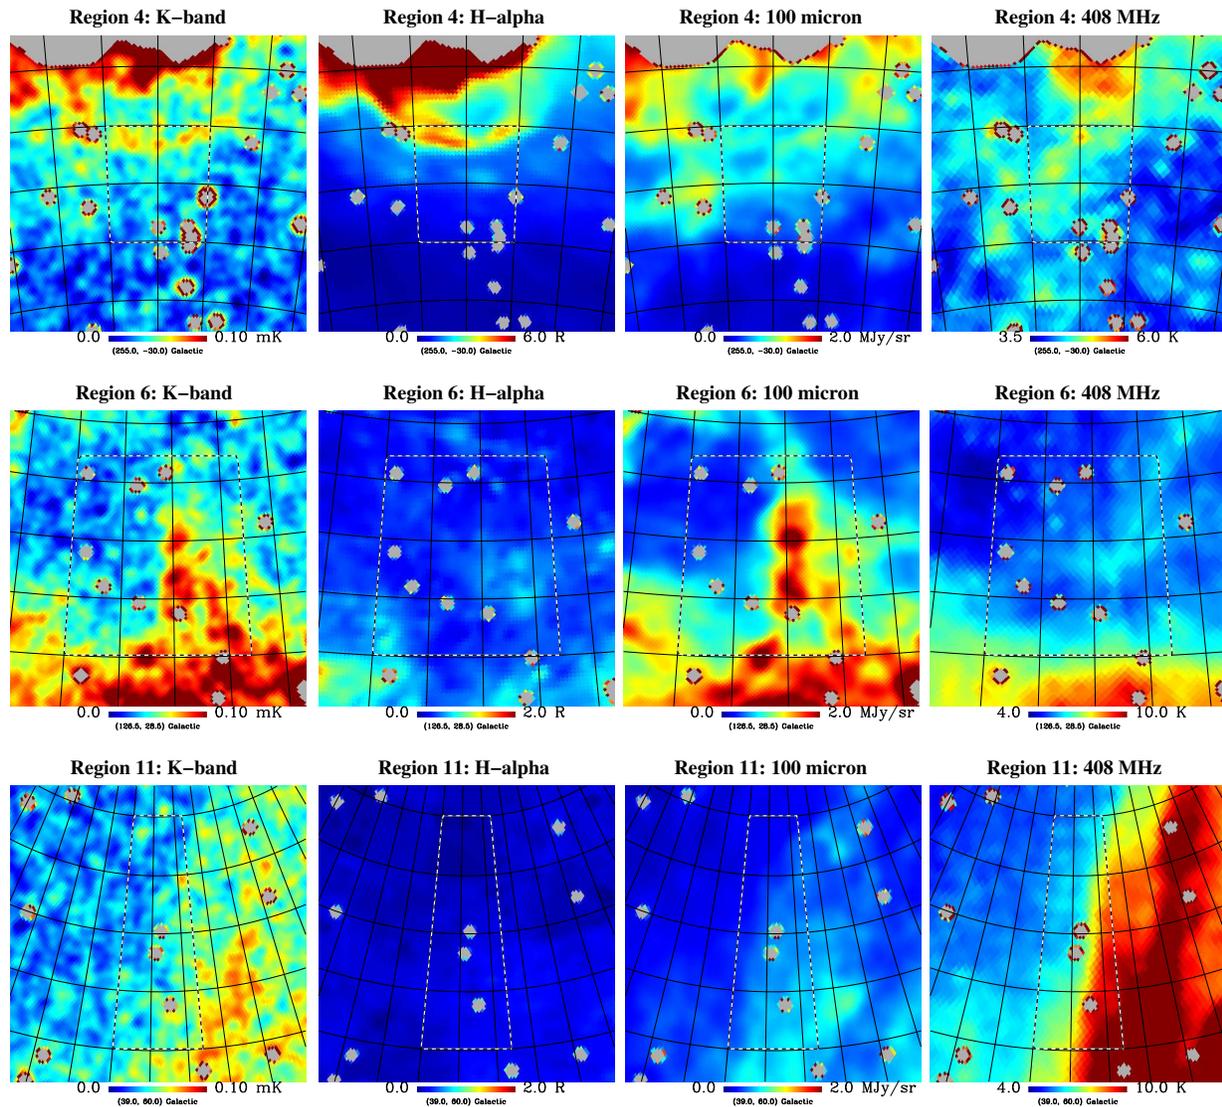
Figure 5.3 Maps of region 4 (Hα (free-free) dominated; *top* row), region 6 (dust dominated; *middle* row), and region 11 (synchrotron dominated; *bottom* row). From left to right are maps at *WMAP* K-band, Hα, SFD98 100 μm dust intensity and 408 MHz. Galactic coordinates are shown. Each map, with a pixel resolution $N_{\mathrm{side}} = 256$, covers a 25 deg $\times$25 deg area with 1 deg resolution. The dotted black/white line delineates the actual areas used for the T-T plots and cross-correlation analyses. Grey areas are the standard *WMAP* Kp2 mask and extragalactic sources mask.

we have N different foreground components, we end up with the simple system of linear equations $\mathbf{A}\mathbf{x} = \mathbf{b}$, where

$$A_{kj} = \mathbf{t}_k^T \cdot \mathbf{M}_{\mathrm{SN}}^{-1} \cdot \mathbf{t}_j,$$
$$b_k = \mathbf{t}_k^T \cdot \mathbf{M}_{\mathrm{SN}}^{-1} \cdot \mathbf{d},$$
$$x_k = \alpha_k. \tag{5.5}$$

When only one template is present, this reduces to equation (5.4) above.

The signal covariance is that for theoretical CMB anisotropies, $M_{ij}^S = \frac{1}{4\pi} \sum_{\ell=0}^{\infty} (2\ell + 1) C_\ell B_\ell^2 P_\ell(\hat{n}_i \cdot \hat{n}_j)$, where B_ℓ is the Gaussian beam of 1° FWHM. The power spectrum, C_ℓ, is taken from the *WMAP* best fit ΛCDM power-law spectrum (Bennett et al., 2003b). The noise covariance is determined from the uncorrelated pixel noise as specified for each pixel in the *WMAP* data, and subsequently convolved as described above.

For each of the fifteen regions, the data vector includes only those pixels of interest, and the covariance matrices are only the corresponding rows and columns. These regions vary in size from 230 to 1140 pixels at this resolution.

The errors $\delta\alpha_k$ are the square root of the diagonal of $\mathbf{A}^{-1}$. The simultaneous fitting of multiple template components allows us to deal with the fact that, though the regions are chosen to be dominant in one given component, they are not entirely free of the other components. The simultaneous fitting of multiple foreground components allows such cross-talk to be quantified.

5.4.2   Results of the cross-correlation analysis

At each *WMAP* band the emissivity of the 3 foreground components (free-free, dust and synchrotron) has been estimated as a ratio of template brightness; μK R^{-1}, μK/μK$_{\mathrm{FDS8}}$ and μK K^{-1} respectively. The analysis was a joint solution derived for all 3 components simultaneously. For each of the components we also made solutions for the all-sky (Kp2 cut) *WMAP* data. For completeness, the solutions for all 3 components and all 15 regions are given in Appendix 5.C (Table 5.9).

5.4.3   Free-free emission

The 5 isolated regions with strong Hα features $(1-5)$ are listed in Table 5.2. Free-free emission should in principle be detectable at radio frequencies from 408 MHz up to and including the *WMAP* frequencies. Accordingly we have used the additional radio surveys at 1.4 GHz (Reich & Reich, 1988) and 2.3 GHz (Jonas et al., 1998) to confirm that the regions exhibit a general flattening of spectral index between 408 MHz and 2.3 GHz suggesting that there is considerable free-free emission, at least relative to any synchrotron component at these frequencies. The peak Hα intensities in the 5 maps are in the range $10-40$ R. The results of the analysis for K- and Ka-bands, where free-free emission will be strongest, are given in Table 5.3. For comparison we list the fits for the two Hα templates (DDD and F03) and find overall, there is good agreement between them. The results indicate a

lower electron temperature of roughly $T_e = 4000 - 5000$ K rather than the often-assumed $T_e = 8000$ K. However, it is important to note that there is variation in this ratio by a factor of ~ 2 from region to region. The average of the 5 fields is consistent with the high Galactic latitude solution (Kp2 cut).

5.4.4   Anomalous dust emission

The dust emissivity for all fields for the 5 *WMAP* frequency bands is given in Table 5.4. The emissivity is that relative to the FDS8 prediction for the W-band and is given separately for the raw *WMAP* data and for Hα corrected data. Also shown is the range of dust temperature in each field taken from SFD98. The five fields $(6 - 10)$ that were selected on the basis that they exhibited dust emission only weakly confused by synchrotron or Hα (free-free) emission produced the most significant correlations amongst the 15 fields studied. In addition field numbers 2,5,11,14 and 15 also show significant dust correlations; this is because the anomalous dust emission is the dominant foreground at the lower *WMAP* frequencies $(20 - 40$ GHz). However, in these latter fields some confusion from synchrotron and free-free might have been expected. The possible effect of a contribution from free-free emission has been tested by subtracting the Hα template converted to free-free brightness temperature with an electron temperature $T_e = 4000$ K, consistent with our previous findings. The free-free correction is small relative to the dust and the dust results remain remarkably robust.

It can be seen from Table 5.4 that there is a spread of a factor of ≈ 2 in the emissivity of dust clouds. The range of dust emission spectral index as determined for example by the ratio of K to Ka-band emissivity of individual clouds is less than this (~ 1.4 in the dust-dominated regions) and may also be a significant result.

The spectrum of the anomalous dust emissivity is best given by the average of the clouds listed in Table 5.4. The average values are seen to be slightly higher than that of the full-sky (Kp2 cut). Note that the traditional vibrational dust component is negligible in the K, Ka and Q-bands and is only dominant at W-band.

The anomalous dust emissivity, when corrected for vibrational emission, shows an average spectral index β of -2.85 (varying from ≈ -2.8 to -3.8 in the dust-dominated regions) and is discussed further in section 5.5.2.

5.4.5   Synchrotron emission

The 408 MHz all-sky map is used as the basic synchrotron template for comparison with the *WMAP* data. The fits between the 408 MHz and the *WMAP* maps (bands K, Ka and Q) are given in the upper part of Table 5.5. The spectral index of synchrotron emission between the *WMAP* bands can be derived from the 408 MHz-correlated signal at each *WMAP* frequency and is given in the bottom part of Table 5.5. Note that the implied spectral index can be misleading where the synchrotron is not detected at an amplitude higher than its error bar (upper part of Table 5.5). It can be seen that, except of region 11 at K-band, the results for individual regions are not significant at the 2σ level. Nevertheless, these

Field No.	Template	$\frac{T_{\rm K}}{T_{{\rm H}_\alpha}}$ μK R^{-1}	$\frac{T_{\rm Ka}}{T_{{\rm H}_\alpha}}$ μK R^{-1}	H$_\alpha$ intensity range R
1	F03	$11.3^{\pm4.8}$	$9.8^{\pm4.4}$	1-10
	DDD	$9.8^{\pm4.4}$	$6.2^{\pm4.1}$	
2	F03	$4.7^{\pm3.0}$	$0.7^{\pm2.8}$	3-15
	DDD	$5.2^{\pm2.9}$	$1.4^{\pm2.7}$	
3	F03	$5.5^{\pm3.0}$	$3.1^{\pm2.6}$	2-14
	DDD	$2.3^{\pm3.5}$	$-1.7^{\pm3.2}$	
4	F03	$7.4^{\pm2.6}$	$4.4^{\pm2.3}$	3-16
	DDD	$7.2^{\pm2.1}$	$1.2^{\pm1.1}$	
5	F03	$9.6^{\pm1.1}$	$4.8^{\pm1.0}$	3-40
	DDD	$10.1^{\pm1.2}$	$5.1^{\pm1.1}$	
Avg.	F03	$8.6^{\pm0.9}$	$4.4^{\pm0.8}$	
	DDD	$8.5^{\pm0.9}$	$3.0^{\pm0.7}$	
Kp2	F03	$7.7^{\pm0.9}$	$3.7^{\pm0.9}$	
	DDD	$7.5^{\pm0.9}$	$3.6^{\pm0.9}$	
	T$_e$=4000K	8.0	3.6	
	T$_e$=5000K	8.9	4.1	
	T$_e$=6000K	9.8	4.5	
	T$_e$=7000K	10.6	4.9	
	T$_e$=8000K	11.4	5.2	

Table 5.3 Free-free emission as determined by CC method for K-and Ka-band *WMAP* data. Units are brightness temperature T (μK) relative to unit Hα intensity I_{H_α} (R). The expected values are shown at the bottom from the predictions given by Dickinson et al. (2003) assuming different electron temperatures. The range of Hα intensities in each region is also given in the last column. For each region, two templates are used: the Finkbeiner (2003) template, F03, or the Dickinson et al. (2003), DDD. Full sky fits, outside the Kp2 cut, are also shown; here, only a diagonal approximation to the covaraince matrix is used, and the uncertainties are determined using simulations (see Appendix 5.B) and the DDD template only.

Field No.	FDS range (μK)	Dust T range (K)	Dust emissivity relative to FDS				
			K	Ka	Q	V	W
1	4.4-10.3	17.6-18.5	$8.4^{\pm5.9}$ / $8.9^{\pm5.7}$	$-0.1^{\pm5.1}$ / $0.5^{\pm5.0}$	$2.4^{\pm4.3}$ / $3.4^{\pm4.2}$	$-1.0^{\pm3.3}$ / $-0.3^{\pm3.2}$	$-2.3^{\pm2.6}$ / $-1.6^{\pm2.5}$
2	14.2-41.3	16.5-17.5	$6.6^{\pm1.6}$ / $6.2^{\pm1.6}$	$1.8^{\pm1.3}$ / $0.9^{\pm1.3}$	$2.0^{\pm1.2}$ / $0.5^{\pm1.2}$	$-0.8^{\pm1.1}$ / $-2.1^{\pm1.0}$	$0.3^{\pm0.9}$ / $-0.3^{\pm0.9}$
3	1.6-10.1	17.6-18.5	$12.1^{\pm5.2}$ / $7.7^{\pm4.7}$	$5.5^{\pm4.5}$ / $1.6^{\pm4.1}$	$7.2^{\pm3.9}$ / $3.8^{\pm3.6}$	$4.8^{\pm3.5}$ / $2.7^{\pm3.3}$	$-0.7^{\pm3.1}$ / $-0.8^{\pm2.9}$
4	2.5-10.4	17.8-18.3	$3.9^{\pm6.1}$ / $2.9^{\pm5.6}$	$7.3^{\pm5.1}$ / $4.8^{\pm4.9}$	$6.4^{\pm4.3}$ / $5.0^{\pm4.3}$	$2.5^{\pm3.6}$ / $2.0^{\pm3.5}$	$-2.0^{\pm3.0}$ / $-2.3^{\pm3.0}$
5	6.9-30.2	17.2-18.6	$12.6^{\pm2.1}$ / $13.1^{\pm2.0}$	$2.6^{\pm1.8}$ / $2.5^{\pm1.8}$	$2.6^{\pm1.6}$ / $1.7^{\pm1.6}$	$-0.3^{\pm1.4}$ / $-1.6^{\pm1.4}$	$-0.8^{\pm1.2}$ / $-1.6^{\pm1.2}$
6	**2.8-51.1**	**15.7-18.1**	$\mathbf{6.7^{\pm0.7}}$ / $\mathbf{6.5^{\pm0.7}}$	$\mathbf{2.9^{\pm0.6}}$ / $\mathbf{2.6^{\pm0.6}}$	$\mathbf{2.1^{\pm0.6}}$ / $\mathbf{1.9^{\pm0.6}}$	$\mathbf{0.8^{\pm0.5}}$ / $\mathbf{0.8^{\pm0.5}}$	$\mathbf{2.3^{\pm0.4}}$ / $\mathbf{2.2^{\pm0.4}}$
7	**6.2-17.0**	**17.9-18.7**	$\mathbf{8.1^{\pm3.9}}$ / $\mathbf{7.8^{\pm3.8}}$	$\mathbf{2.8^{\pm3.5}}$ / $\mathbf{2.8^{\pm3.5}}$	$\mathbf{0.2^{\pm3.2}}$ / $\mathbf{-0.3^{\pm3.2}}$	$\mathbf{0.6^{\pm2.9}}$ / $\mathbf{0.4^{\pm2.9}}$	$\mathbf{0.5^{\pm2.6}}$ / $\mathbf{0.2^{\pm2.6}}$
8	**6.4-25.8**	**17.1-18.0**	$\mathbf{8.4^{\pm2.5}}$ / $\mathbf{8.2^{\pm2.5}}$	$\mathbf{2.6^{\pm2.3}}$ / $\mathbf{2.6^{\pm2.3}}$	$\mathbf{2.4^{\pm2.2}}$ / $\mathbf{2.5^{\pm2.2}}$	$\mathbf{2.4^{\pm2.0}}$ / $\mathbf{2.5^{\pm2.0}}$	$\mathbf{3.1^{\pm1.8}}$ / $\mathbf{3.1^{\pm1.8}}$
9	**9.6-97.5**	**15.4-17.8**	$\mathbf{7.3^{\pm0.6}}$ / $\mathbf{7.3^{\pm0.6}}$	$\mathbf{2.9^{\pm0.6}}$ / $\mathbf{2.9^{\pm0.6}}$	$\mathbf{1.8^{\pm0.5}}$ / $\mathbf{1.8^{\pm0.5}}$	$\mathbf{1.4^{\pm0.5}}$ / $\mathbf{1.4^{\pm0.5}}$	$\mathbf{1.9^{\pm0.4}}$ / $\mathbf{1.9^{\pm0.4}}$
10	**3.5-31.7**	**16.8-18.3**	$\mathbf{12.0^{\pm1.4}}$ / $\mathbf{12.2^{\pm1.3}}$	$\mathbf{4.7^{\pm1.2}}$ / $\mathbf{5.0^{\pm1.2}}$	$\mathbf{1.8^{\pm1.1}}$ / $\mathbf{1.8^{\pm1.1}}$	$\mathbf{0.2^{\pm1.0}}$ / $\mathbf{0.3^{\pm1.0}}$	$\mathbf{1.3^{\pm0.9}}$ / $\mathbf{1.3^{\pm0.9}}$
11	1.9-7.6	17.6-18.3	$19.0^{\pm7.9}$ / $17.8^{\pm7.8}$	$9.9^{\pm6.8}$ / $8.3^{\pm6.7}$	$6.6^{\pm6.1}$ / $7.0^{\pm6.1}$	$3.1^{\pm5.4}$ / $3.5^{\pm5.4}$	$1.4^{\pm4.7}$ / $1.3^{\pm4.7}$
12	2.2-5.9	17.6-18.2	$15.6^{\pm7.5}$ / $15.4^{\pm7.5}$	$5.3^{\pm6.6}$ / $5.1^{\pm6.6}$	$4.6^{\pm5.8}$ / $4.4^{\pm5.8}$	$-7.1^{\pm5.0}$ / $-7.3^{\pm5.0}$	$-2.1^{\pm4.2}$ / $-2.2^{\pm4.2}$
13	1.9-10.5	17.5-18.5	$3.9^{\pm6.4}$ / $3.6^{\pm6.3}$	$-3.7^{\pm5.7}$ / $-3.0^{\pm5.7}$	$1.6^{\pm4.9}$ / $1.5^{\pm4.9}$	$-2.2^{\pm3.8}$ / $-2.2^{\pm3.8}$	$-1.0^{\pm3.1}$ / $-1.0^{\pm3.1}$
14	3.7-9.8	17.5-18.5	$8.3^{\pm4.8}$ / $7.1^{\pm4.7}$	$7.6^{\pm3.9}$ / $7.1^{\pm3.9}$	$1.8^{\pm3.0}$ / $1.1^{\pm3.0}$	$3.0^{\pm2.3}$ / $2.2^{\pm2.3}$	$-0.6^{\pm1.8}$ / $-0.7^{\pm1.8}$
15	6.1-16.3	17.2-18.3	$8.1^{\pm3.3}$ / $8.1^{\pm3.3}$	$-1.0^{\pm3.0}$ / $-0.8^{\pm3.0}$	$-1.0^{\pm2.7}$ / $-0.9^{\pm2.7}$	$-1.3^{\pm2.5}$ / $-1.4^{\pm2.5}$	$-0.5^{\pm2.2}$ / $-0.4^{\pm2.2}$
Avg.			$7.8^{\pm0.4}$ / $7.7^{\pm0.4}$	$3.0^{\pm0.4}$ / $2.8^{\pm0.4}$	$2.0^{\pm0.3}$ / $1.8^{\pm0.3}$	$0.8^{\pm0.3}$ / $0.6^{\pm0.3}$	$1.6^{\pm0.3}$ / $1.4^{\pm0.3}$
Kp2			$6.6^{\pm0.3}$ / $6.6^{\pm0.3}$	$2.4^{\pm0.3}$ / $2.4^{\pm0.3}$	$1.4^{\pm0.3}$ / $1.4^{\pm0.3}$	$0.9^{\pm0.3}$ / $0.8^{\pm0.3}$	$1.2^{\pm0.3}$ / $1.1^{\pm0.3}$

Table 5.4 Dust-correlated emissivity in the 5 *WMAP* bands from a CC analysis using the FDS dust prediction. For each region, the first row gives the result from a simultaneous fit of the three foreground components, while the second row gives that for a fit of two foregrounds to data with the Hα (free-free) subtracted assuming T$_e$=4000K. Also shown are the FDS8 predicted dust emission intensity range for each region as well as the SFD dust temperature. The five dust-dominated regions are highlighted in bold face. Full sky fits, outside the Kp2 cut, are also shown; here, only a diagonal approximation to the covaraince matrix is used, and the uncertainties are determined using simulations.

Field	Synchrotron fit amplitudes		
	$\frac{K}{408}$ $(\times 10^6)$	$\frac{Ka}{408}$ $(\times 10^6)$	$\frac{Q}{408}$ $(\times 10^6)$
11	$4.28^{\pm 1.76}$	< 2.20	< 0.90
12	< 6.22	< 1.82	< 0.64
13	< 6.81	< 2.38	< 1.31
14	< 4.67	< 0.90	< 0.12
15	< 8.52	< 2.20	< 1.66
Avg.	$2.82^{\pm 0.82}$	< 0.87	< 0.30
Kp2	5.45 ± 0.79	< 2.93	< 2.06
Field	Synchrotron spectral index β		
	$\frac{K}{408}$	$\frac{Ka}{408}$	$\frac{Q}{408}$
11	$-3.07^{+0.09}_{-0.13}$	< -2.97	< -3.03
12	< -2.98	< -3.01	< -3.10
13	< -2.96	< -2.95	< -2.94
14	< -3.05	< -3.17	< -3.47
15	< -2.90	< -2.97	< -2.89
Avg.	$-3.18^{+0.06}_{-0.09}$	< -3.18	< -3.27
Kp2	$-3.01^{+0.03}_{-0.04}$	< -2.90	< -2.84

Table 5.5 Synchrotron spectral index β ($T_b \propto \nu^\beta$) between 408MHz and the *WMAP* K, Ka, and Q bands. Where the fit amplitude is less than a 2σ detection, the 2σ upper limit is shown instead ($< 2\sigma$).

spectral index values are somewhat steeper than those calculated between GHz frequencies and the *WMAP* frequencies. This is likely to be due to the effect of spectral ageing of the CR electrons which produce the synchrotron emission in the Galactic magnetic field. The average of the five regions is significant at K-band with $\beta = -3.18$ and the full-sky value (outside of the Kp2 cut) is -3.01. The uncertainties shown for the Kp2 fits are based on simulations and are larger than might be expected since they use a diagonal approximation to the full covariance matrix M in equation (4) (see Appendix 5.B for details).

5.5   Discussion

5.5.1   Free-free emission

Free-free emission is the weakest foreground component at *WMAP* frequencies for intermediate and high Galactic latitudes. By selecting five Hα -dominated regions, we have been able to quantify the Hα -correlated free-free emission in these regions as tabulated in section 5.4.3 (Table 5.3).

We note that the current analysis is for intermediate and high latitude H$\textsc{ii}$ regions

which are therefore associated with the local spiral arm such as the Gould Belt system and the Gum Nebula; they most likely lie at $|z| \lesssim 200$ pc from the Galactic plane. This class of HII regions is different from the more compact regions confined to the Galactic plane with a width of $z \approx 60$ pc (Paladini et al., 2003, 2004). As mentioned in section 5.2.2, these HII regions have a mean $T_e = 7200 \pm 1200$ K compared to the value of around $4000 - 5000$ K found in the present study. Note that for the brightest region (5) where T_e is most accurately determined, $T_e \sim 6500$ K. This discrepancy has an unclear origin, especially given that an identical result is determined for the entire high-latitude sky as defined by the *WMAP* Kp2 mask. It may be indicative of problems associated with the Hα template itself, or in the conversion of Hα flux to the free-free brightness temperature. However, variations by a factor of ~ 2 are also seen from region to region.

It is of interest to note that the electron temperature derived from radio recombination line studies of extended HII regions such as the Gum Nebula have an average value of 7000 K (Woermann et al., 2000). A study of diffuse foregrounds in the *COBE*-DMR at 7 deg angular scales (Banday et al., 2003) found that Hα correlations with the DDD Hα template were more or less consistent with $T_e \sim 7000$ K for $|b| > 15$ deg. For $|b| > 30$ deg, lower values were preferred but with larger error bars. Banday et al. (2003) also analysed 19 GHz data with a 3 deg beam, which favoured lower values even for $|b| > 30$ deg. These results suggest that the discrepancy may be scale-dependent and therefore might be related to the different beam shapes of the WHAM and SHASSA Hα surveys for angular scales comparable to the beam size (~ 1 deg).

5.5.2 Dust emission

The dust-correlated emission is the dominant foreground component in the *WMAP* bands and its spectral properties can be derived for the individual clouds included in the present study. The spectra of all fifteen regions are shown in Fig. 5.4. It is immediately seen that the spectral slopes of each of the clouds over the range from K- to V-band are quite similar. Also, all the clouds show a turn-up in emissivity at W-band where thermal emission becomes dominant.

A further significant result is that the emissivity relative to the FDS8 prediction varies by a factor of 2 from cloud to cloud.

The average spectral emissivity for the clouds is shown by black filled circles in Fig. 5.4. The average spectral index from K- to Q-band, is -2.4, shown as a dashed black line in Fig. 5.4. The average spectral index in the range K–Ka-band and Ka–Q-band are -2.6 and -1.9. An estimate of β from K-band to higher frequencies depends sensitively upon the vibrating dust contribution in these bands and requires a knowledge of its spectral index. Assuming $\beta = +1.7$, we use a simple χ^2 test to find the best fit to the data over a grid of values for the anomalous dust spectral index, the anomalous dust amplitude, and the vibrating dust amplitude. The result gives a spectral index for the anomalous component of -2.85 and shows that the FDS8 prediction at W-band is underestimated by 30 per cent. Assuming a steeper thermal index of 2.0 or 2.2 results in an anomalous index of -2.75. The data minus the vibrating dust fit are shown in red in Fig. 5.4 along with the best-fit

anomalous power law.

One proposed explanation for the anomalous dust-correlated emission in the low fre-
quency *WMAP* bands, motivated by its spectral behaviour, is that it represents a hard
synchrotron component, morphologically different in the *WMAP* bands from the soft syn-
chrotron component traced by the 408 MHz emission. This hard synchrotron emission
would correlate with dust in regions of active star formation. We find that this anomalous
component has a spectral index from the K- to Q-band of $\beta = -2.85$ when the FDS8 ther-
mal dust prediction is subtracted, assuming $\beta = +1.7$ (see Fig. 5.4). We then extrapolate
this component to 408 MHz to see how much of this hard emission would be seen at that
low frequency. If there is no spectral hardening between 408 MHz and *WMAP*, then a
spectral index of $\beta = -2.85$ would imply more emission at 408 MHz than is observed by
a factor of $\gtrsim 2$ in many regions. If the thermal dust spectrum is steeper, e.g., $\beta = +2.2$,
then (as discussed above) the anomalous index flattens slightly to -2.75, but that still
overpredicts the emission at 408 MHz.

We now consider the relevance of the dust temperature to our results, using as a proxy
the SFD98 colour temperature based on the ratio of the DIRBE 100- and 240- μm data at
a resolution of 1.deg 1.

Comparing the dust temperature to the emissivity in these regions (as well as over the
sky outside the Kp2 cut), it appears that in general, the strongest anomalous emission
(relative to the FDS8 prediction) comes from the coldest regions. This is particularly
striking in the two dust regions which have the smallest error bars, 6 and 9, which can be
seen to dominate the averages in Table 5.4 and Fig. 5.4. Fig. 5.5 shows the dust temperature
in region 6, which can be compared to the 100 μm SFD98 map shown in Fig. 5.3. The fit
amplitudes in these regions show the lowest cross-correlation with the low-frequency data,
where the emission is not thermal but comes from the anomalous component. Table 5.4
shows that the K-band emission (where strongly detected) is lowest relative to the FDS8
template in regions 2, 6, and 9, which have the lowest dust temperatures as well. (Not
apparent from the table, which shows only the dust temperature *range*, is that the emission
comes from the coldest parts of the region).

Finkbeiner (2004a) has examined the foreground residuals after subtracting the FDS8
prediction and proposed that an anomalous dust template could be better constructed
using FDS8$\times$T^2. By comparing the χ^2 values for template fits for the Kp2 mask using
FDS8$\times$T^n for different values of n, we find that formally the best value for n in the K-band
is 1.6 and that it drops to zero at the higher frequencies, as one would expect. The same
exercise repeated on the fifteen regions, however, gives a large scatter in the preferred value
of n, ranging from 0 to over 5 (the limit of the range tested).

The average dust emissivity among all of the regions is slightly higher at low frequencies
(particularly K-band) than that over the full sky (outside the Kp2 cut). Can this be
explained by the full sky fits being driven by the even colder emission near the Galactic
Plane at the anti-centre? Fits of the FDS8 template to the hemisphere around the Galactic
centre and around the Galactic anti-centre do indeed show that the fit values around the
anti-centre are from 10 to 30 per cent lower than fits around the Galactic centre, depending
on the band. These differences appear to confirm the indication in the smaller regions that
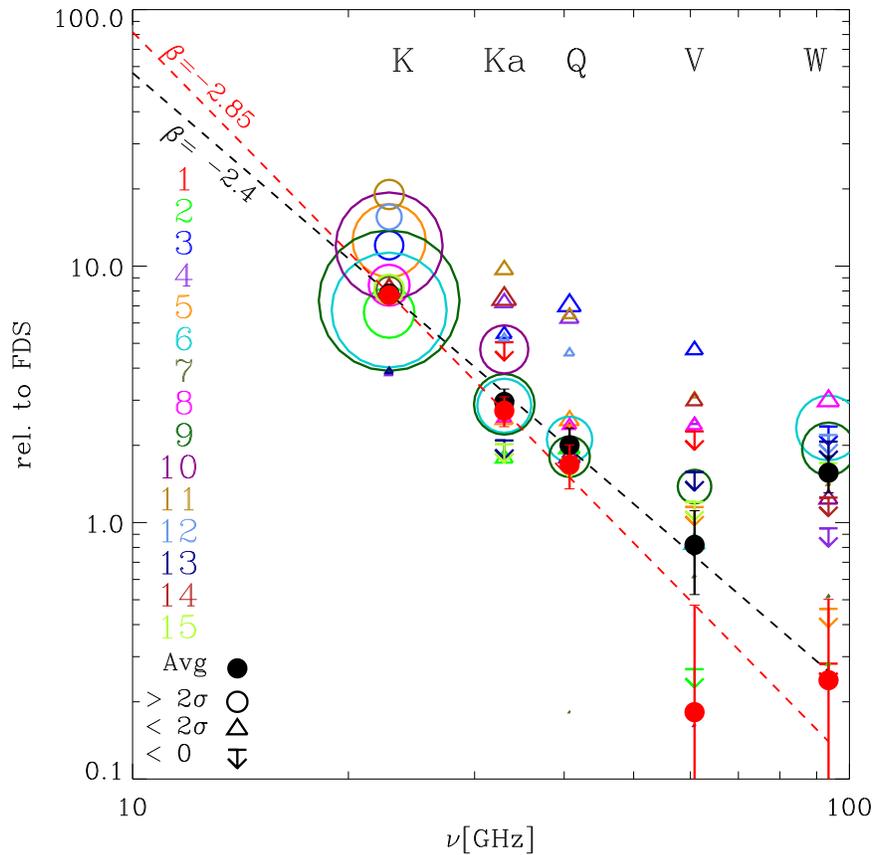
Figure 5.4 Summary of dust emissivities (antenna temperature units, relative to FDS8 at 94 GHz) from C-C analysis. Each region emissivity is plotted with a different colour open symbol whose size is proportional to the signal-to-noise ratio. The symbol is a circle if the signal-to-noise is greater than two or a triangle if it is positive but less than two. Where the fit value is negative, the 1σ upper limit is shown as an arrow downward, and values more than one sigma negative are not shown (though all values are included in the weighted average). The filled black circle shows the weighted average surrounded by its error bar. The dashed line represents the best-fit spectral index over K through Q bands. Red filled circles and the red dashed line show the average and fit when corrected for the best-fit thermal dust component assuming $\beta = +1.7$; see text in §5.5.2.
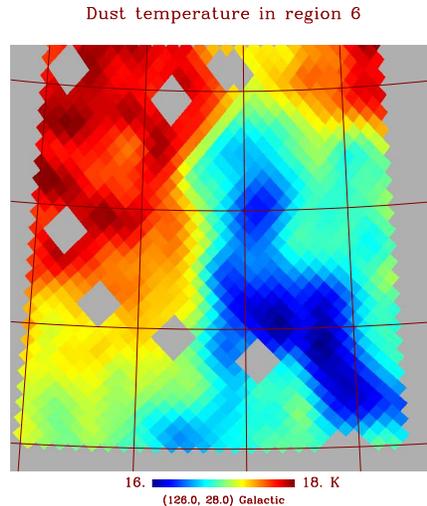
Figure 5.5 SFD dust temperature for region 6, showing that the emissivity, as seen in Fig. 5.3, comes from the coldest regions.

the emissivity of the anomalous component is lower from colder dust.

5.5.3   Synchrotron emission

As in the case of free-free emission the synchrotron emission, defined as the 408 MHz-correlated emission, is weaker than the dust-correlated emission at *WMAP* frequencies $(20 - 90$ GHz) over most of the sky. Except for the strongest synchrotron feature, the North Polar Spur (region 11), the 408-MHz correlations are marginal, even at K-band. Nevertheless the average spectral index of the 5 synchrotron-selected fields (Table 5.5) indicates an increasing slope relative to 408 MHz of $\beta = -3.18$ to < -3.27 from K to Q-band. The full-sky fits (Kp2) indicates $\beta = -3.01$ from 408 MHz to 23 GHz. At higher WMAP frequencies, the fits are not significant. The errors are probably too conservative due to the diagonal covariance assumption used in the Kp2 fit (see Appendix 5.B). Nevertheless, the results are in good agreement with the *WMAP* values reported in Bennett et al. (2003c).

Using low frequency maps at GHz frequencies (408 MHz, 1420 MHz, 2326 MHz from Haslam et al. (1982), Reich & Reich (1986) and Jonas et al. (1998), respectively) we find that the average spectral index for the 5 maps in the GHz range is -2.91. This steepening with frequency is indicative of ageing of the relativistic electrons in these fields. This can be compared with the results of the Cosmosomas experiment which found $\beta \ -3.16$ from 408 MHz to 13 GHz for $|b| > 30\,\mathrm{deg}$ (Fernández-Cerezo, 2006). Our limited data does not show evidence for strong variation of the synchrotron spectral index at *WMAP* frequencies from field-to-field. The discrepancy between the average of the regions and the Kp2 cut is likely to be due to the dominance of the North Polar Spur which is known to have a steeper index relative to the sky average.

5.6 Comparison with *WMAP3*

Whilst this paper was being completed, the *WMAP* team released their 3-year results to the community. For this new analysis, several semi-independent studies of the foreground contamination of the data were undertaken (Hinshaw et al., 2006). For the purposes of gaining physical insight into the nature of the Galactic foregrounds, a maximum-entropy (MEM) technique was applied. However, for the purposes of cleaning the data for cosmological studies, a template subtraction method was adopted. As with the first year analysis, the F03 Hα template was employed as a tracer of free-free emission and the FDS8 model normalised at 94 GHz used for thermal dust emission. For synchrotron emission an internally generated template, comprising the difference of the K- and Ka-bands, was constructed. There are several aspects of these foreground results that merit comment in the present paper.

Hinshaw et al. (2006) have determined a free-free to Hα ratio of $\sim$6.5 μK R^{-1} based on fits to the F03 template. This is completely consistent with the mean values derived from the 5 free-free regions in this paper, particularly after adjusting the *WMAP* coefficient 20-30 per cent upwards to compensate for the increased amplitude of their $f_d = 0.5$ dust corrected Hα template. The MEM analysis finds a slightly higher ratio $\sim$8 μK R^{-1}, but also considerable variation (by a factor of $\sim$2) depending on location, as we have also found.

They adopted the difference between the observed K- and Ka-band emission as a tracer of synchrotron emission was intended to compensate for the problem with assuming a fixed full-sky spectral index in order to extrapolate the Haslam 408 MHz sky map to *WMAP* frequencies. Such a procedure is clearly in contradiction with the spectral index studies of Reich & Reich (1988) between 408 and 1420 MHz, which showed large variations of spectral index across the sky. Of course, utilising the K–Ka map as a foreground correction template is, to some extent, independent of whether the dominant foreground contribution is due to synchrotron, anomalous dust, or a combination thereof. By utilising what Hansen et al. (2006) have referred to as an internal template, it is likely that the synchrotron morphology is well traced over the frequencies of interest. Moreover, fitting this template to the remaining sky maps with a global scale factor per frequency is likely to be quite accurate, even in the presence of modest departures from a single spectral index. This treatment does not contradict our own studies, since our intention is to study the variations in spectral behaviour over the sky relative to the 408 MHz survey (Section 5.5.3).

Maintaining a consistent approach to their treatment of the first year data, the *WMAP* team have not attempted to directly address the issue of the anomalous dust correlated component. Rather, the MEM solutions were allowed only to produce what may be interpreted as a combined synchrotron/anomalous dust solution at each frequency, with no attempt made to disentangle the two components. Hinshaw et al. (2006) comment that it is not possible, using the *WMAP* data alone, to distinguish between anomalous dust emission and flatter spectrum synchrotron emission that is well correlated with dusty star-forming regions. This is particularly true given that a putative spinning dust component can exhibit a similar spectral shape to synchrotron emission over the 20 - 40 GHz frequency range.

However, Page (2006) in their foreground modelling efforts for the polarisation analysis, determine a high-polarisation-fraction component of the synchrotron emission that is well correlated with the Haslam template, and a low-polarisation component with a dust-like morphology.

Lagache (2003) has noted that the anomalous dust component is predominantly associated with regions of low HI column density, and proposes that this indicates that small dust grains are responsible for the emission. Such grains will emit radiation via rotational degrees of excitation, but since the grains will not align well with magnetic field lines, then a low level of polarisation is expected. Page (2006) also find a steep spectrum for the polarised synchrotron component of -3.2 between K- and Ka-bands, although this unexpectedly flattens at higher frequencies. Nevertheless, the overall picture seems to be quite consistent with what is found in this work – a steep synchrotron component correlated with the Haslam template plus a spinning dust contribution that is the dominant foreground over the 20-60 GHz range for temperature anisotropy, but is subdominant to synchrotron in polarisation due to the low associated polarisation fraction.

Finally, we note that the *WMAP* team impose a constraint on their thermal dust template fits that the derived dust coefficients must have a spectral index of ~ 2, rather than the value of 2.2 used in the first year analysis, or the value of 1.7 predicted by the FDS8 model. This makes only a very minor difference to the spinning dust spectrum because thermal dust emission is negligible at in the lowest *WMAP* bands.

5.7   Conclusions

In our study of the free-free, dust and synchrotron foreground components in the *WMAP* data we have chosen a selection of fields which are intended to have minimal cross-contamination from other components. Each of the 3 components has been quantified in terms of a mean value of the emissivity in each of the 5 *WMAP* bands.

Fig. 5.6 shows the emission, in thermodynamic units, as expected using the values of emissivity we have determined in conjunction with the Hα , 100 μm dust and the 408 MHz templates outside the Kp2 mask. These are calculated using r.m.s. values of 5.9 K, 2.6 R and 6.8 μK for 408 MHz, Hα , and FDS8 model at 94 GHz, respectively. The data points are the Kp2 solutions from the C-C analysis. The curves are foreground models; synchrotron with $\beta = -3.1$ normalised to K-band, free-free with $\beta = -2.14$ for $T_e = 4000$ K and vibrational dust emission for $\beta = +1.7$ normalised to W-band. The magenta curve is a spinning dust model from Draine & Lazarian (1998a,b)[5], scaled to fit the data points from K- to Q-band. The curves plotted in Fig. 5.6 are therefore not strictly "best-fits" to these data points and are plotted to depict the approximate amplitude and spectral dependencies of the 4 Galactic components at high latitudes (outside the Kp2 cut). The dominance of the dust emission is evident. Also the similarity of the dust and the synchrotron spectrum at *WMAP* frequencies is superficially evident; these may be separated by lower frequency

---

[5]Spinning   dust   models   were   downloaded   from:
http://www.astro.princeton.edu/~draine/dust/dust.html/.
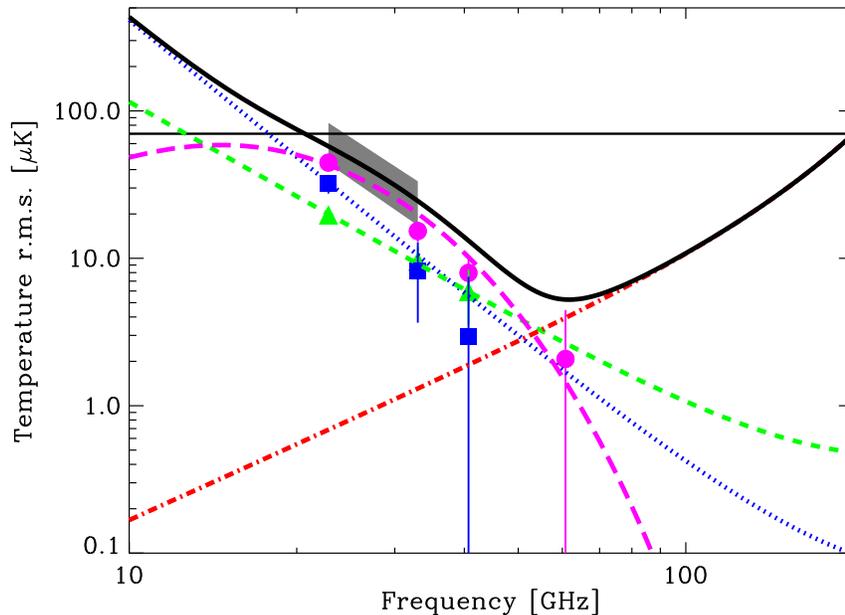
Figure 5.6 R.m.s. fluctuations spectrum (thermodynamic temperature units) of foreground components at 1 deg resolution. The symbols are the Kp2 solutions from the C-C analysis for synchrotron (*blue squares*), free-free (*green triangles*) and anomalous dust-correlated emission (*magenta circles*). The curves represent synchrotron for $\beta = -3.1$ (*blue dotted line*), free-free for $\beta = -2.14$ (*green short-dashed line*), vibrational dust with $\beta = +1.7$ (*red dot-dashed line*), Draine & Lazarian spinning dust model (*magenta long-dashed line*) and all the foreground models combined in quadrature (*thick black line*). The CMB fluctuations are shown at 70 μK r.m.s (*thin black line*). The grey region shows the variation of dust emissivity between K- and Ka-bands in regions where significant ($> 2\sigma$) dust-correlated emission was detected.

(5 − 15 GHz) data as shown for example by de Oliveira-Costa et al. (2004a) and Watson et al. (2005). We note that dust (anomalous and thermal) is the dominant foreground over the *WMAP* and Planck bands. The thick black curve in Fig. 5.6 is the total of the 4 model curves, combined in quadrature, corresponding to the approximate total foreground r.m.s. level; the minimum foreground contamination of the CMB is at ~ 70 GHz for total intensity. The integrated foreground spectrum is relatively simple when sampled sparsely in frequency (e.g. *WMAP*). This is why the *WMAP* team find that the spectrum from 408 MHz to V-band is well fitted by a simple power law, although it says little in itself about the underlying foreground components.

One of the most interesting results of the present study is the variation in the emissivity of the dust-correlated emission in the *WMAP* bands, typically by a factor of ~ 2. The grey shaded region of Fig. 5.6 indicates the variation of the dust-correlated component between

K- and Ka-bands from the C-C analysis of 15 regions where a significant ($> 2\sigma$) detection was found. There is clearly considerable variation in dust emissivity at K- and Ka-bands. Fig. 5.6 also shows that the average emissivity in the regions is higher than typical value seen outside the Kp2 cut (see Table 5.4). The effect on producing an anomalous dust template is profound because it is the dominant foreground between 20 and 60 GHz. If the anomalous dust emission is polarised at ~ 10 per cent this could approach the synchrotron emission as an important polarised foreground in the $20 - 60$ GHz range. These considerations are important for cleaning CMB maps from missions such as *WMAP* and Planck as well and high sensitivity ground-based experiments such as CLOVER[6] and QUIET[7].

The *WMAP* team have instigated some debate over the origin of the anomalous dust correlated component, and have preferred an interpretation in terms of a hard synchrotron contribution from star-forming regions that are strongly associated with dust. The 5 synchrotron regions selected in this paper are dominated by well-known structures on the sky away from star-forming regions. As such, it might be expected that the derived synchrotron indices would be steep, and this is indeed the case. It is also expected that cross-talk with other physical components should be minimised. An important result is that anomalous emission was detected in 11 out of the 15 regions of studied here. We consider this to be strong evidence against a synchrotron origin for the anomalous component, although the exact nature of the dust emission mechanism requires lower frequency measurements to elucidate the detailed spectral behaviour. Certainly, the anomalous emission constitutes the dominant foreground component over the 20-60 GHz frequency range.

Further work is required to understand the origin of this variation in dust emissivity by using other physical properties of dust such as its size and temperature. New data in the critical radio frequency range $\sim 5 - 15$ GHz will be vital for a clearer definition of the anomalous dust spectrum. Polarisation data will be particularly important for understanding the physical mechanism that produces the anomalous emission which is expected to be polarised at different levels (e.g. Draine & Lazarian 1999). For example, spinning dust emission is expected to be only weakly (few per cent) polarised (Lazarian & Draine, 2000), whereas the synchrotron emission is known to be highly polarised.

5.A   Aliasing of foregrounds due to CMB subtraction

Some studies of the *WMAP* data remove the ILC map before performing the correlation analysis in order to minimise the impact of the CMB structure on that analysis. However, this introduces a new problem. The ILC is constructed as a linear combination of the five *WMAP* frequency bands in such a way that i) the CMB signal is conserved; ii) the variance of the final map is minimised. The latter condition does *not* guarantee the absence of residual foregrounds from the ILC. Therefore, when one subtracts the ILC from the individual *WMAP* frequency maps the foreground contribution is altered by an amount

---

[6]http://www-astro.physics.ox.ac.uk/research/expcosmology/groupclover.html

[7]http://quiet.uchicago.edu/

depending on the residuals in the ILC map. Since we do not know a priori the actual properties of the foregrounds present in the data, it is not possible to absolutely correct for the effect of this 'aliasing' of foregrounds from the ILC into the individual frequency maps. However, as shown in Eriksen et al. (2004a), the likely residual level can be predicted under various assumptions about the foreground spectral behaviour. This at least gives some insight into the impact of the ILC subtraction on derived foreground properties.

For the high latitude fields considered in this paper, the 5 coefficients that define the ILC map for each frequency band are constant.[8] This allows us to calculate the effect of the ILC subtraction, for a given spectral index, both in terms of relative amplitude of the aliased signal as well as the effect on the derived spectral index.

The numbers in Table 5.6 represent the fraction of foreground signal present in the corrected frequency map in antenna temperature units, assuming a specific spectral behaviour for that foreground. The CMB correction due to the ILC is given by

$$0.109K - 0.684Ka - 0.096Q + 1.921V - 0.250W.$$

For these coefficients, it turns out that the effect can be relatively small for some of the foregrounds and frequencies of the *WMAP* data. For example, a synchrotron component, with $\beta = -3.1$ is increased at the 3 per cent level at K-band or 10 per cent at Ka-band. For synchrotron, the strong aliasing at high frequencies is not important since the actual foreground level is well below that of the other components. For Hα , with a flatter spectral index, the effect is $\sim$1 per cent lower at K-band and $\sim$15 per cent at W-band. It appears that the K-, Ka- and Q-band fits are reliable tracers of the synchrotron and free-free spectral indices up to a modest correction factor. The anomalous emission appears to have a relatively steep spectral index similar to that of synchrotron thus we might expect aliasing at similar levels to the synchrotron emission for $\beta = -2.85$. This may not be the case if the spinning dust models of Draine & Lazarian are correct models of the emission mechanism.

The aliasing effect is most strongly seen in W-band for the vibrational dust component. For $\beta = +2.0$, the dust component will be *reduced* by as much as 40 per cent. This is enough to change the effective spectral index of this component by a significant amount; a true spectral index of +1.7 increases to +1.9 after ILC subtraction. Furthermore, the dust correlation coefficients at low frequency have oversubtracted thermal dust contributions.

It is our contention that making an ILC subtraction before foreground analysis introduces difficulties of interpretation that may invalidate conclusions unless the effect of foreground aliasing is handled correctly.

5.B The reliability of T–T versus C–C fits

The TT method is a straightforward linear fit of two datasets and has two basic drawbacks: firstly, only one template can be compared to the data, and secondly, only independent

---

[8]Bennett et al. (2003c) also derive coefficients in 12 regions of the Galactic plane.

β	K	Ka	Q	V	W
Free-free					
-2.0	0.98	0.96	0.93	0.86	0.71
-2.1	0.99	0.98	0.97	0.93	0.84
-2.2	1.00	1.00	1.00	0.99	0.99
Thermal dust					
1.7	-5.19	-2.25	-1.24	-0.08	0.54
2.0	-6.97	-2.75	-1.43	-0.03	0.62
2.2	-8.34	-3.08	-1.54	0.00	0.66
Anomalous dust					
-2.5	1.02	1.05	1.08	1.21	1.53
CNM	1.21	1.37	1.66	4.96	101.54*
WNM	1.11	1.34	1.88	13.5	2921.2*
WIM	1.15	1.37	1.83	8.96	574.35*
DL 98	1.17	1.36	1.71	6.01	146.88*
Synchrotron					
-2.7	1.03	1.07	1.12	1.35	1.98
-2.9	1.03	1.09	1.16	1.49	2.51
-3.1	1.03	1.10	1.19	1.62	3.09

Table 5.6 The fraction of foreground signal present in the corrected frequency map in antenna temperature units, assuming a specific spectral behaviour for that foreground. For example,. the Q-band contains 119 per cent of the synchrotron contribution expected if the synchrotron emission has a spectral index of $\beta = -3.1$. Therefore, in the absence of cross-talk with other foreground components, a synchrotron template fit to the Q-band should be downweighted by a factor of 1.19. For anomalous dust, we also tabulate the expected aliasing signal for the Cold Neutral Medium (CNM), Warm Neutral Medium (WNM), and Warm Ionised Medium (WIM) from models due to Draine & Lazarian (1998a,b). *The W-band aliasing factor is large, but since the WNM spectrum falls off very rapidly this factor results only in physical amplitudes closer to the other channels.

pixel noise (effectively, a diagonal covariance matrix) can be taken into account. The fifteen regions were chosen such that one component is dominant in an effort to get around the first problem. The ILC can be subtracted from the data in order to deal with the second. The CC method is more complicated but allows simultaneous fitting of multiple templates with a full signal covariance matrix.

We use simulations with foreground components added at known amplitudes to test how well these different fitting methods recover the input. A set of 1000 simulations are created at the resolution of the *WMAP* data products, i.e. at HEALPix of $N_{side} = 512$, using the beam width and noise properties of the respective channels. The foregrounds are added at the theoretical levels for free-free (assuming an electron temperature of T_e=8000K) and for synchrotron (assuming a spectral index of $\beta = -3.0$), while the dust is added at the levels found by Table 3 of Bennett et al. (2003b), approximating both thermal and anomalous contributions. These maps are then smoothed (via convolution in harmonic space) to a common resolution of 1° and downgraded to $N_{side} = 128$, as are the data.

These simulations can be used to test how much cross-talk still affects fits using only one template in the small regions chosen in an attempt to minimise such problems. In other words, does the existence of small amounts of additional dust and synchrotron emission in the Hα-dominated regions affect the individual fits using the Hα template alone? Looking at the elements of the fit matrix A (defined in §5.4.1) gives an indication of how the different templates correlate with eachother. This is shown for a few regions in Table 5.7. But analysis of simulations is needed to quantify the effects on the fit results.

First, we compare fits to data where the ILC estimate of the CMB component is subtracted, which is necessary in the TT method but which introduces the small aliasing effects discussed in Appendix 5.A. Table 5.8 columns labelled "1to3" give results where one template is fit to simulations with three foregrounds; these show that the Hα estimates *are* overestimated due to the other emission, even in regions dominated by Hα emission. The amount depends on the region, but some regions are overestimated by as much as 30%. This affects both methods when only one template is used. But where the CC method is given the three templates to fit simultaneously, the results are then roughly consistent with the input, as seen in columns labelled "3to3". The CC and TT methods also both give roughly correct results when only one foreground component is present in the data, as seen in the "1to1" columns.

The bias due to the ILC subtraction is clear in Table 5.8 in the "1to1" columns or the first "3to3" column. As expected for an index of $\beta = -2.1$, the free-free emission is underestimated by $\sim 1\%$. Likewise, the synchrotron is overestimated by approximately 3%. The dust is underestimated by $\sim 6\%$, as expected for the mixture of thermal and anomalous dust represented by the input amplitudes, taken from the *WMAP* best fits (Bennett et al. 2003b, Table 3.)

We conclude from the above and from the information in Table 5.8 that the only un-biased estimate of the foreground fit amplitudes comes from the full CC method using the three template simultaneously and the CMB signal plus noise covariance matrix against the raw data. The bias introduced by the ILC-subtraction is relatively small (if the foregrounds follow roughly the expected spectral behaviour), but the cross-talk among templates pre-

	FDS dust	DDD H$_\alpha$	Haslam	Const.
Region 5				
FDS dust	1.00	0.23	0.59	0.62
DDD H$_\alpha$	0.23	1.00	0.20	0.14
Haslam	0.59	0.20	1.00	0.82
Const.	0.62	0.14	0.82	1.00
Region 9				
FDS dust	1.00	0.26	0.33	0.30
DDD H$_\alpha$	0.26	1.00	0.42	0.44
Haslam	0.33	0.42	1.00	0.91
Const.	0.30	0.44	0.91	1.00
Region 14				
FDS dust	1.00	0.34	0.51	0.61
DDD H$_\alpha$	0.34	1.00	0.29	0.33
Haslam	0.51	0.29	1.00	0.73
Const.	0.61	0.33	0.73	1.00

Table 5.7 Fit matrix showing effectively the amount of cross-correlation among the different templates. Each value is $t_i^T M^{-1} t_j / \sqrt{(t_i^T M^{-1} t_i)(t_j^T M^{-1} t_j)}$. The regions were chosen as those where the dominant component had the smallest error bar.

vents the TT method from being accurate enough in most cases. Only the CC method with the full covariance, shown in the "3to3†" column of Table 5.8 gives unbiased results.

5.C   Full fit results

Table 5.9 lists the solutions for all 15 regions, for the 3 components of synchrotron, free-free and dust. Bold face numbers represent the 5 fields that were chosen to be dominated by each foreground component.

	TT...		CC...				
	1to3	1to1	1to3	1to1	3to3	1to3†	3to3†
				free-free (11.4)			
1	1.02±0.11	0.99±0.11	1.03±0.10(0.10)	0.99±0.10(0.10)	0.99±0.10(0.11)	1.12±0.33(0.34)	1.00±0.34(0.35)
2	1.08±0.09	0.99±0.09	1.18±0.07(0.08)	0.99±0.07(0.08)	0.99±0.08(0.09)	1.13±0.13(0.16)	1.00±0.14(0.16)
3	1.21±0.07	0.99±0.07	1.21±0.07(0.07)	0.99±0.07(0.07)	0.99±0.09(0.09)	1.09±0.16(0.17)	0.99±0.17(0.18)
4	1.14±0.06	0.99±0.06	1.14±0.05(0.06)	0.99±0.05(0.06)	0.99±0.07(0.08)	1.08±0.16(0.17)	1.01±0.17(0.18)
5	1.32±0.03	0.99±0.03	1.25±0.03(0.03)	0.99±0.03(0.03)	0.99±0.03(0.03)	1.04±0.06(0.06)	1.00±0.06(0.06)
				dust (6.3)			
6	1.06±0.03	-	1.07±0.03(0.03)	-	0.94±0.03(0.03)	1.05±0.11(0.11)	1.00±0.11(0.11)
7	1.05±0.22	-	1.06±0.22(0.22)	-	0.94±0.23(0.25)	1.13±0.59(0.60)	1.00±0.59(0.61)
8	0.87±0.15	-	0.89±0.15(0.16)	-	0.93±0.15(0.16)	0.98±0.36(0.37)	1.01±0.36(0.37)
9	0.98±0.05	-	0.98±0.05(0.05)	-	0.94±0.05(0.06)	1.03±0.10(0.10)	0.99±0.10(0.10)
10	1.03±0.06	-	1.03±0.06(0.06)	-	0.94±0.06(0.06)	1.10±0.20(0.21)	1.00±0.21(0.22)
				synchrotron (5.73)			
11	1.32±0.08	-	1.31±0.08(0.08)	-	1.04±0.11(0.13)	1.11±0.27(0.28)	1.01±0.28(0.29)
12	1.15±0.09	-	1.15±0.09(0.09)	-	1.03±0.10(0.10)	1.05±0.24(0.26)	1.01±0.24(0.26)
13	1.24±0.13	-	1.24±0.12(0.13)	-	1.03±0.15(0.16)	1.08±0.32(0.35)	1.02±0.32(0.35)
14	1.22±0.07	-	1.22±0.06(0.06)	-	1.04±0.08(0.09)	1.08±0.23(0.24)	1.00±0.23(0.24)
15	0.65±0.18	-	0.66±0.17(0.19)	-	1.01±0.23(0.25)	0.90±0.44(0.45)	1.01±0.44(0.45)

Average fit amplitude as fraction of true

Table 5.8 Comparison of resulting fit amplitudes as fraction of true. The expected foreground residuals after ILC subtraction in the K band are: for free-free, assuming $\beta = -2.1$, 0.99% ; for dust, added at the *WMAP* best-fit values, approximating a combination of thermal and anomalous dust, 0.94%; for synchrotron, assuming $\beta = -3.0$, 1.03%.

Field	K	Ka	Q	V	W
			free-free		
1	**9.8**$^{\pm4.4}$	**6.2**$^{\pm4.1}$	**6.0**$^{\pm3.9}$	**3.4**$^{\pm3.6}$	**3.7**$^{\pm3.2}$
2	**5.2**$^{\pm2.9}$	**1.4**$^{\pm2.7}$	**0.3**$^{\pm2.6}$	**-0.3**$^{\pm2.5}$	**0.7**$^{\pm2.2}$
3	**2.3**$^{\pm3.5}$	**-1.7**$^{\pm3.2}$	**-2.2**$^{\pm3.0}$	**-2.5**$^{\pm2.8}$	**0.8**$^{\pm2.5}$
4	**7.2**$^{\pm2.1}$	**1.2**$^{\pm1.1}$	**0.9**$^{\pm0.6}$	**0.3**$^{\pm0.4}$	**-0.2**$^{\pm0.3}$
5	**10.1**$^{\pm1.2}$	**5.1**$^{\pm1.1}$	**3.3**$^{\pm1.0}$	**1.6**$^{\pm0.9}$	**0.7**$^{\pm0.8}$
6	-4.7$^{\pm9.9}$	-12.3$^{\pm7.4}$	-12.1$^{\pm6.3}$	-5.1$^{\pm5.8}$	-10.9$^{\pm5.2}$
7	3.5$^{\pm11.9}$	6.8$^{\pm9.0}$	-5.7$^{\pm7.8}$	-0.5$^{\pm7.2}$	-5.2$^{\pm6.5}$
8	-22.9$^{\pm16.4}$	0.7$^{\pm13.5}$	-5.4$^{\pm12.3}$	2.5$^{\pm11.3}$	-13.1$^{\pm10.2}$
9	9.2$^{\pm15.5}$	-0.3$^{\pm11.0}$	11.4$^{\pm9.5}$	-1.2$^{\pm8.7}$	0.1$^{\pm7.9}$
10	15.3$^{\pm15.2}$	13.6$^{\pm7.2}$	2.7$^{\pm3.9}$	5.7$^{\pm2.3}$	1.1$^{\pm1.7}$
11	-16.6$^{\pm19.1}$	-29.4$^{\pm13.2}$	12.9$^{\pm11.5}$	9.6$^{\pm10.4}$	-2.9$^{\pm9.3}$
12	9.0$^{\pm12.9}$	1.6$^{\pm9.8}$	-9.6$^{\pm8.4}$	-7.3$^{\pm7.7}$	-5.2$^{\pm7.0}$
13	3.7$^{\pm14.0}$	15.1$^{\pm7.2}$	-3.5$^{\pm3.7}$	-1.5$^{\pm2.2}$	-0.6$^{\pm1.7}$
14	-4.4$^{\pm8.0}$	-2.2$^{\pm5.8}$	-1.2$^{\pm4.9}$	-4.9$^{\pm4.5}$	5.8$^{\pm4.0}$
15	4.7$^{\pm17.5}$	19.5$^{\pm12.5}$	9.5$^{\pm10.6}$	-9.5$^{\pm9.7}$	2.6$^{\pm8.7}$
			dust		
1	8.4$^{\pm5.9}$	-0.1$^{\pm5.1}$	2.4$^{\pm4.3}$	-1.0$^{\pm3.3}$	-2.3$^{\pm2.6}$
2	6.6$^{\pm1.6}$	1.8$^{\pm1.3}$	2.0$^{\pm1.2}$	-0.8$^{\pm1.1}$	0.3$^{\pm0.9}$
3	12.1$^{\pm5.2}$	5.5$^{\pm4.5}$	7.2$^{\pm3.9}$	4.8$^{\pm3.5}$	-0.7$^{\pm3.1}$
4	3.9$^{\pm6.1}$	7.3$^{\pm5.1}$	6.4$^{\pm4.3}$	2.5$^{\pm3.6}$	-2.0$^{\pm3.0}$
5	12.6$^{\pm2.1}$	2.6$^{\pm1.8}$	2.6$^{\pm1.6}$	-0.3$^{\pm1.4}$	-0.8$^{\pm1.2}$
6	**6.7**$^{\pm0.7}$	**2.9**$^{\pm0.6}$	**2.1**$^{\pm0.6}$	**0.8**$^{\pm0.5}$	**2.3**$^{\pm0.4}$
7	**8.1**$^{\pm3.9}$	**2.8**$^{\pm3.5}$	**0.2**$^{\pm3.2}$	**0.6**$^{\pm2.9}$	**0.5**$^{\pm2.6}$
8	**8.4**$^{\pm2.5}$	**2.6**$^{\pm2.3}$	**2.4**$^{\pm2.2}$	**2.4**$^{\pm2.0}$	**3.1**$^{\pm1.8}$
9	**7.3**$^{\pm0.6}$	**2.9**$^{\pm0.6}$	**1.8**$^{\pm0.5}$	**1.4**$^{\pm0.5}$	**1.9**$^{\pm0.4}$
10	**12.0**$^{\pm1.4}$	**4.7**$^{\pm1.2}$	**1.8**$^{\pm1.1}$	**0.2**$^{\pm1.0}$	**1.3**$^{\pm0.9}$
11	19.0$^{\pm7.9}$	9.9$^{\pm6.8}$	6.6$^{\pm6.1}$	3.1$^{\pm5.4}$	1.4$^{\pm4.7}$
12	15.6$^{\pm7.5}$	5.3$^{\pm6.6}$	4.6$^{\pm5.8}$	-7.1$^{\pm5.0}$	-2.1$^{\pm4.2}$
13	3.9$^{\pm6.4}$	-3.7$^{\pm5.7}$	1.6$^{\pm4.9}$	-2.2$^{\pm3.8}$	-1.0$^{\pm3.1}$
14	8.3$^{\pm4.8}$	7.6$^{\pm3.9}$	1.8$^{\pm3.0}$	3.0$^{\pm2.3}$	-0.6$^{\pm1.8}$
15	8.1$^{\pm3.3}$	-1.0$^{\pm3.0}$	-1.0$^{\pm2.7}$	-1.3$^{\pm2.5}$	-0.5$^{\pm2.2}$
			synchrotron		
1	3.3$^{\pm2.4}$	0.7$^{\pm1.0}$	-0.7$^{\pm0.5}$	0.0$^{\pm0.3}$	-0.1$^{\pm0.2}$
2	8.2$^{\pm3.5}$	0.4$^{\pm1.5}$	-0.9$^{\pm0.8}$	0.8$^{\pm0.5}$	0.7$^{\pm0.4}$
3	3.7$^{\pm2.2}$	1.3$^{\pm0.8}$	0.4$^{\pm0.4}$	-0.0$^{\pm0.2}$	-0.1$^{\pm0.2}$
4	4.2$^{\pm3.0}$	1.2$^{\pm1.0}$	1.0$^{\pm0.5}$	0.5$^{\pm0.3}$	0.1$^{\pm0.2}$
5	4.3$^{\pm2.2}$	1.7$^{\pm1.1}$	-0.2$^{\pm0.7}$	-0.7$^{\pm0.4}$	0.1$^{\pm0.3}$
6	2.8$^{\pm1.3}$	-0.3$^{\pm0.4}$	0.1$^{\pm0.2}$	-0.2$^{\pm0.1}$	-0.2$^{\pm0.1}$
7	3.8$^{\pm2.6}$	0.6$^{\pm1.0}$	-0.5$^{\pm0.5}$	-0.1$^{\pm0.4}$	0.3$^{\pm0.3}$
8	7.2$^{\pm4.2}$	0.6$^{\pm2.2}$	0.0$^{\pm1.5}$	-0.2$^{\pm1.0}$	0.4$^{\pm0.8}$
9	0.2$^{\pm3.3}$	0.1$^{\pm1.6}$	2.1$^{\pm1.0}$	-0.5$^{\pm0.6}$	-0.5$^{\pm0.5}$
10	1.2$^{\pm2.1}$	0.8$^{\pm0.9}$	-0.4$^{\pm0.5}$	0.6$^{\pm0.3}$	0.1$^{\pm0.2}$
11	**4.3**$^{\pm1.8}$	**-0.0**$^{\pm1.1}$	**-0.6**$^{\pm0.8}$	**0.1**$^{\pm0.6}$	**0.6**$^{\pm0.5}$
12	**2.9**$^{\pm1.6}$	**0.6**$^{\pm0.6}$	**0.1**$^{\pm0.3}$	**0.1**$^{\pm0.1}$	**0.1**$^{\pm0.1}$
13	**2.8**$^{\pm2.0}$	**0.7**$^{\pm0.8}$	**0.5**$^{\pm0.4}$	**0.0**$^{\pm0.2}$	**0.0**$^{\pm0.2}$
14	**1.6**$^{\pm1.5}$	**-0.4**$^{\pm0.6}$	**-0.5**$^{\pm0.3}$	**-0.0**$^{\pm0.2}$	**-0.0**$^{\pm0.1}$
15	**3.0**$^{\pm2.8}$	**-0.8**$^{\pm1.5}$	**-0.2**$^{\pm0.9}$	**-0.4**$^{\pm0.6}$	**-0.7**$^{\pm0.5}$

Table 5.9 Full fit results for all components in all regions. The three templates (DDD H_α, FDS8 dust, and Haslam 408MHz, plus a constant offset) are fit simultaneously to each band individually.

Chapter 6

Phase analysis

6.1 Introduction

Statistical isotropy (SI) and Gaussianity in the CMB fluctuations are simple to define; as described in 1.3, they imply that the power spectrum contains a full description of the information in the data. But since Gaussianity and SI are effectively defined by the *absence* of any other information, to characterise *violations* (i.e., $\langle a_{\ell m}^* a_{\ell' m'} \rangle \neq \delta_{\ell \ell'} \delta_{mm'} C_\ell$) is much more difficult. Such violations imply higher order statistical moments that encode correlations between different scales, preferred directions, "m-preference", non-random phases, etc.

A variety of statistical measures have been defined and applied to the *WMAP* data to test for such violations, among which phase statistics are particularly interesting. In the spherical harmonic expansion of the sky given in equation (1.2), the morphological information is contained in the phases of the harmonic coefficients $\phi_{\ell m}$, where

$$a_{\ell m} = |a_{\ell m}| e^{i \phi_{\ell m}}. \tag{6.1}$$

If the sky is a realisation of a statistically isotropic stochastic Gaussian process on the sphere (more briefly, a Gaussian random field, or GRF hereafter), then the real and imaginary parts of the harmonic coefficients are independent Gaussian random variables whose variance is C_ℓ. This implies that the phases $\phi_{\ell m}$ are randomly and uniformly distributed on the unit circle. Any coherent structure on the sky or correlation between modes will manifest itself as correlation in these phases. (Any sky cut will also induce correlations between modes, as will be discussed in § 6.4. For the following analysis, the full sky is used.)

Figure 6.1 demonstrates the importance of the phases. The CMB, which is largely Gaussian and isotropic, has largely random phases, while on the map of the Earth, it is easy to pick out by eye the non-Gaussian and anisotropic structures. By using the phases from the Earth map, we can make a sky that has the same power spectrum as the CMB and yet is clearly not representative of an isotropic stochastic process. Looking at the power spectrum alone gives only a very limited picture of the sky.

(a) $a_{\ell m} = |a_{\ell m}^{CMB}| \exp(i\phi_{\ell m}^{CMB})$ (b) $a_{\ell m} = |a_{\ell m}^{\oplus}| \exp(i\phi_{\ell m}^{\oplus})$

(c) $a_{\ell m} = |a_{\ell m}^{CMB}| \exp(i\phi_{\ell m}^{\oplus})$ (d) $a_{\ell m} = |a_{\ell m}^{\oplus}| \exp(i\phi_{\ell m}^{CMB})$
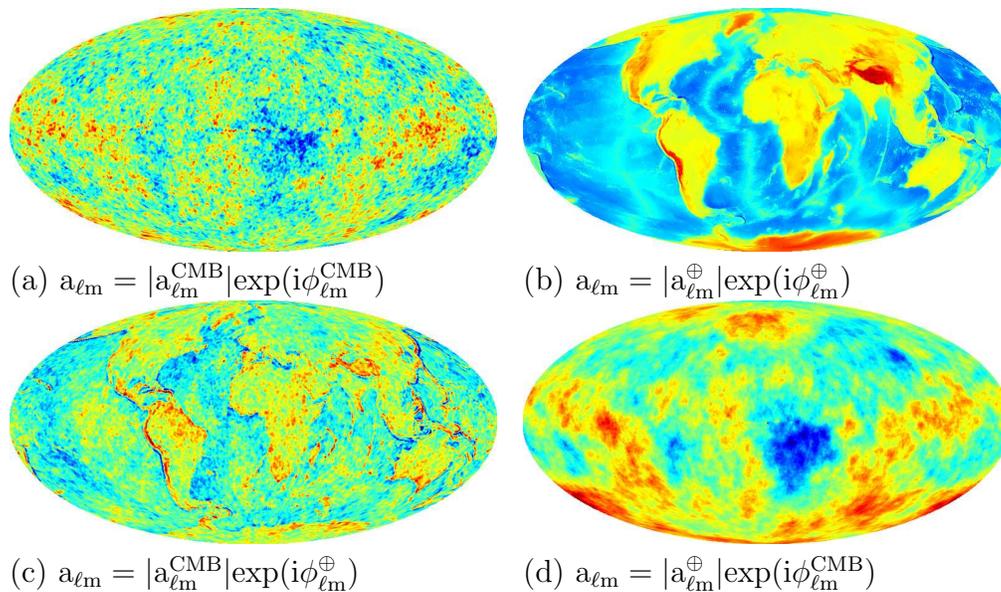
Figure 6.1 Example of the importance of phase to morphology. On the top row are maps of the CMB and of the Earth, and on the bottom row are the same maps but with phases exchanged. The map with the harmonic amplitudes from the CMB but with the phases from the Earth still looks like the Earth, and vice versa. (Figure adapted from talk by P. Coles.)
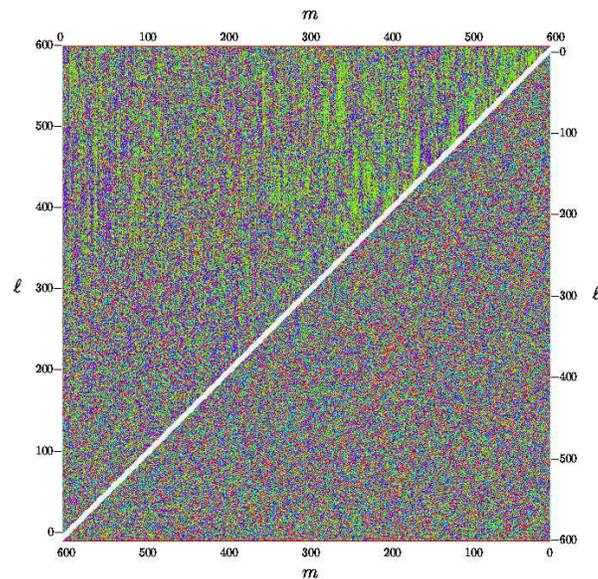
Figure 6.2 Colour-coded phase gradient D_ℓ (see § 6.2.1) for two foreground cleaned maps from *WMAP* first-year data (from Chiang et al. 2003). The eye easily picks out the non-Gaussian signal for the top map, while the bottom looks completely random. The top is the "foreground cleaned" map from Tegmark et al. (2003), while the bottom is their Wiener filtered map. The former has structure along the galactic plane that is obvious by eye but that has been effectively filtered out of the Wiener map.

Chiang et al. (2003) plot the phases (or phase gradients; see § 6.2) colour-coded from blue to red for phases from zero to 2π, which allows the eye to easily spot non-Gaussianity. Figure 6.2 shows this plot for two maps based on *WMAP* first-year data. These maps have had different foreground cleaning methods applied, and the cleaning appears more successful in one case than the other, where there are clear non-Gaussian signatures remaining. Figures 6.3 and 6.4 show maps and similar phase gradient plots for some toy models to illustrate the effects of some likely foreground residuals. The diffuse plane induces correlations at constant m for small ℓ (large scales), while a point source in the plane induces such striping at high ℓ (small scales). A point source off the plane makes a more complicated signal for $m \sim \ell$, and a collection of point sources shows a variety of structures bye eye. These indicate that the signal seen in the top half of Fig. 6.2, the striping at constant m, is likely to be due to a residual point sources along the galactic plane.

Chiang et al. (2003) applied a "return mapping" technique and Coles et al. (2004) use Kuiper's statistic, both to measure how much phases "return" or clump in a particular region of $[0, 2\pi)$. Naselsky et al. (2005) continue this work, relating it to Pearson's random walk statistics (see Stannard & Coles 2005.)

It should be noted that the spherical harmonic coefficients, and in particular their phases and phases statistics, are not rotationally invariant. The return mapping analyses of Chiang et al. (2003) and Chiang et al. (2006) were applied in one orientation, while the Kuiper's statistic analysis of Coles et al. (2004) used a set of random rotations and examined the distribution. Naselsky et al. (2005) discuss the rotational properties of their phase statistic, and point out that the fact that the non-invariance of phases could be used to test for foreground residuals, but they only mention the possibility.

The following analysis extends this previous work to look at the rotational properties of the statistics for both Gaussian and non-Gaussian datasets. We show that the properties of the statistics at different orientations contain information about the morphology of the non-Gaussian correlations in the CMB. Knowing the morphology of the non-Gaussian signal we are looking for helps to construct a more sensitive statistical test.

6.2   Methods

Detecting a non-random distribution of phases is not a trivial proposition. The non-randomness may manifest in an infinite number of possible ways. The methods below look for clustering in the phase values or in the phase gradients, $D \equiv \phi_{\ell'm'} - \phi_{\ell m}$. The clustering can be measured in either the one-dimensional distribution on the unit circle or in a two-dimensional plot of one set of phases versus another.

6.2.1   Return mapping statistic

Chiang et al. (2003) apply their "phase mapping" technique (see references therein) to the WMAP first year results. The basic idea is that for a given pair $(\Delta l, \Delta m)$, a return map is the plot of $\phi_{l+\Delta l, m+\Delta m}$ versus $\phi_{l,m}$. If the phases are randomly distributed, this will be
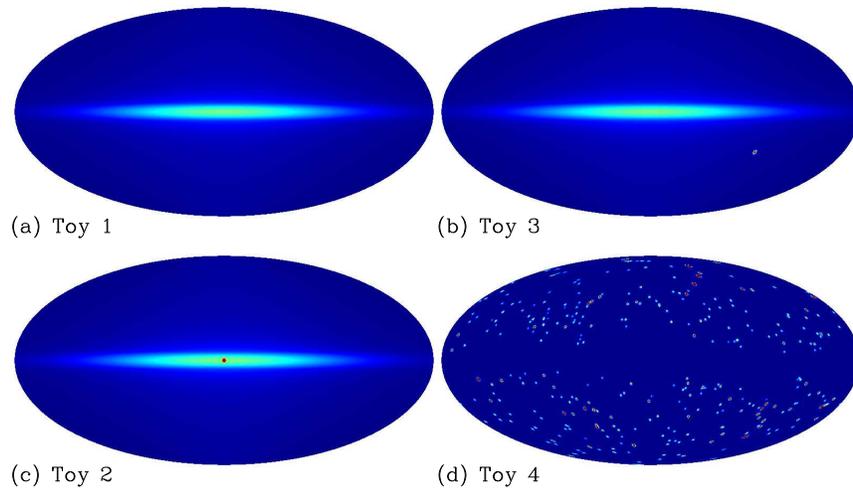
Figure 6.3 Example maps whose phases are shown in Fig. 6.4. Plots (a), (b), and (c) show toy models with a diffuse plane, an added point source at the galactic centre (c) or in the southeastern quadrant (b). Toy 4 (d) is a map made from the *WMAP* point source catalogue.
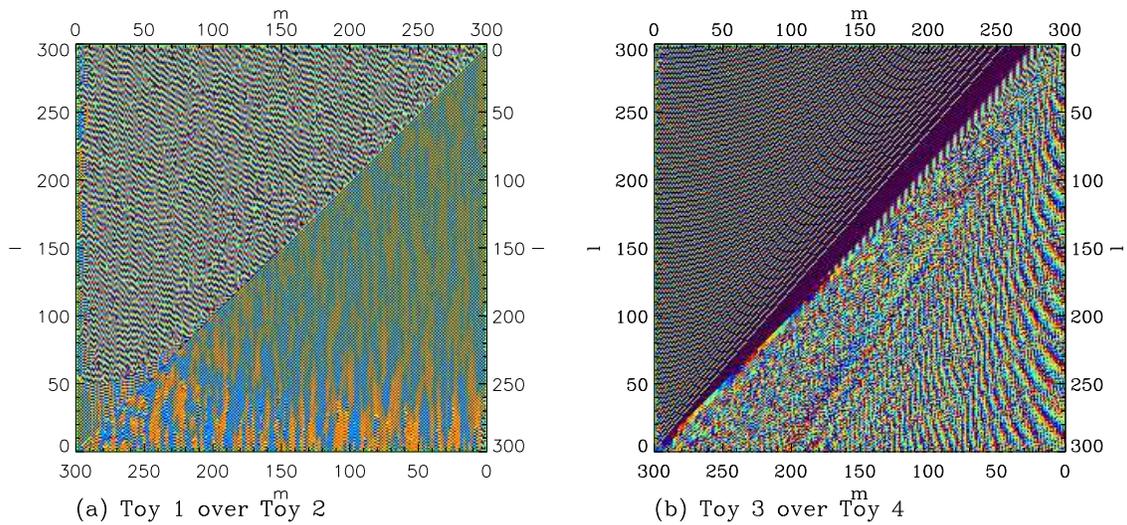
Figure 6.4 Phase gradients for the examples shown in Fig. 6.3.

a scatter plot filling in uniformly the square of $[0, 2\pi)^2$. Measuring how evenly distributed these pairs are gives a measure of Gaussianity.

So first, for a given gradient $(\Delta l, \Delta m)$, they collect all pairs $(\phi_{l+\Delta l, m+\Delta m}, \phi_{l,m})$, and bin them onto a discrete grid, where at each bin (i, j), the value $p(i, j)$ represents how many pairs fall in that bin. Then they smooth that map (see Fig. 6.5 left column for examples) and calculate the statistic:

$$\bar{\chi}_R^2 = \frac{1}{M} \sum_{i,j} \frac{[p(i,j) - \bar{p}]^2}{\bar{p}} \tag{6.2}$$

where M is the number of pixels in the binned return map, $p(i, j)$ is the number of points falling in that bin on the smoothed map, and $\bar{p}$ is the average value. In the following, this $\bar{\chi}_R^2$ will be referred to as the "R-statistic". Their several papers describe the expected distribution of these $\bar{\chi}_R^2$ values, which is approximately Gaussian for large samples, as can be verified by simulations.

They collect an ensemble of such return maps, which is, for a range of interest $l_{min} \leq l \leq l_{max}$, the set of return maps for all possible $(\Delta l, \Delta m)$ in that range. They then examine the distribution of the $\bar{\chi}_R^2$ values.

6.2.2 Kuiper's statistic

Coles et al. (2004) outline a method of measuring the non-Gaussianity of each multipole in the CMB using Kuiper's Statistics. Kuiper's statistic (often referred to as V but here called the "K-statistic") reflects the degree of "clumping" in a set of values. Since the randomness of the phases is a weak test of Gaussianity, Kuiper's statistic can be applied to the set of phases for a particular map; e.g., for each mode l, Kuipers can be calculated over the the set of phase gradients $D_m(l) = \phi_{l,m+1} - \phi_{l,m}$, or the contrary. To account for the effect of the choice of coordinate system (it is not rotationally invariant), Coles et al. calculate the statistic for a set of randomly rotated maps. They then compare the binned *distribution* of values against those for a set of Gaussian simulations that provide the null hypothesis. A simple χ_K^2 statistic can measure the deviation of the distribution for a given CMB map from the expected distribution from the simulations. A rough guess for the significance of the "non-Gaussianity" of that mode is then the percentage of Gaussian simulations that had a lower χ_K^2 for that mode.

Coles et al. (2004) applies this statistic to specific slices of the data. The statistic for $D_m(l)$ refers to the Kuiper's statistic for the collection of phase gradients $\phi_{l,m+1} - \phi_{l,m}$ with $0 < m \leq l$, evaluated at all l, while $D_\ell(m)$ is that for the collection of $\phi_{l+1,m} - \phi_{l,m}$ with $m \leq l \leq l_{max}$ evaluated for all m. This gives a separate measure of Gaussianity for each multipole or "m-pole". One can obviously slice or group the data in other ways, since Gaussianity implies that all such sets of phases or gradients be randomly distributed. One could in theory tailor the analysis to use the slices that pick out a particular non-Gaussian signal.
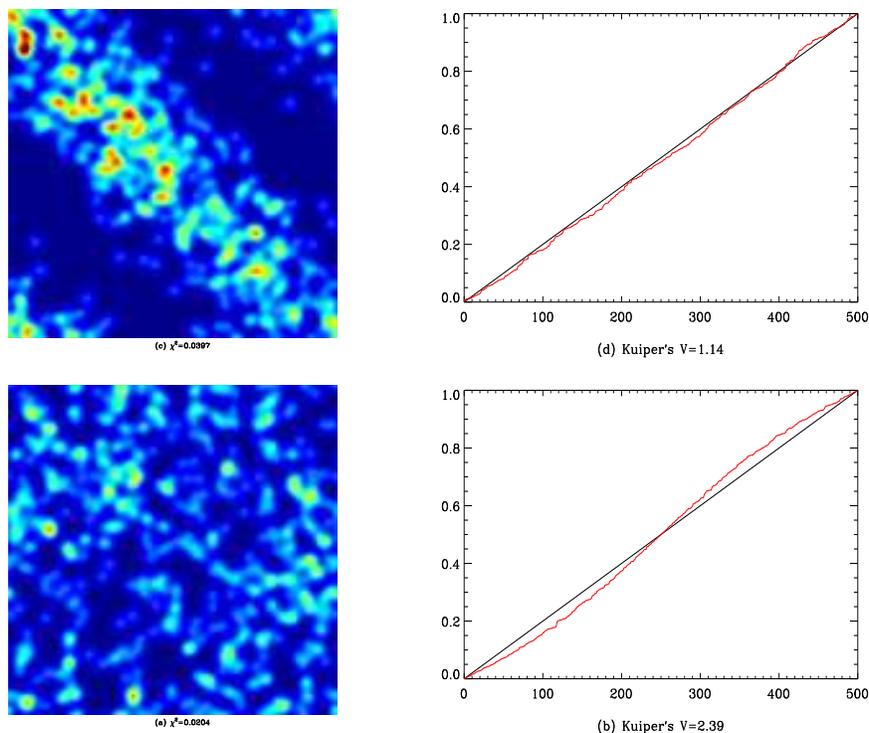
Figure 6.5 Return mapping statistic compared to Kuiper's. The left column shows the smoothed return maps from which the R-statistic is calculated, while the right column shows the sorted gradients over which Kuiper's statistic is calculated. The top row shows a case where the R-statistic detects something when the K-statistic does not, and on the bottom row, vice versa.

Methods using Kuiper's statistic (K-statistic) and the return mapping statistic (R-statistic) are complementary. Intuitively, it is clear that a clumping of values in $\phi_{\ell'm'} - \phi_{\ell m}$ will also form clumps on a return map, so often where one statistic is high, the other is as well. But there are situations where one statistic detects a non-randomness that the other does not. Figure 6.5 shows an example for a set of phase pairs constructed to illustrate the point. (This situation does not correspond directly to any particular morphology on the sky, as that set of phase pairs could arise in any number of ways via different slices of the data, different gradients, etc. Both of those examples were constructed as illustrations.)

6.2.3   Statistical mapping

As described above, phase statistics are not rotationally invariant. The analyses of Chiang et al. (2003); Chiang et al. (2006) are limited to the Galactic reference frame. The authors argue that this is the "natural" frame to discuss non-Gaussian signatures which are presumed to be related to Galactic foreground emission. The Coles et al. (2004) analysis is somewhat more sophisticated in that it uses an ensemble of random rotations to generate

a *distribution* of phase statistics, and it is this distribution that they study. But neither analysis takes advantage of the fact that the variations of these statistics over different rotations contains additional information about the morphology of the correlations.

The goal of this Chapter is to show how this rotational information can be used to study the non-Gaussian signals in different models and to construct tests tailored to detect them. Rather than choosing one coordinate system or averaging over random rotations, we examine the statistic at a set of orientations and look at the morphology of the result. So we construct a statistical map for a given sky, where each pixel contains the value of a non-Gaussian statistic such as those defined above as calculated in the reference frame where the z-axis points through that pixel. Figure 6.6 shows an example of the R-statistic applied to a cubic torus universe (see §6.3.1 for a description of these). (In that case, no rotation about the original z-axis is performed, i.e., the third of the Euler angles. Note that the Kuiper's statistic test is invariant under such a rotation, but the return mapping statistic is not. See below.) Figs. 6.6 (g) and (h) clearly show non-Gaussian signatures along great circles perpendicular to the axes of symmetry of this model, while at other orientations, that signature disappears. Using this information about where the non-Gaussian signal is expected to be strongest makes a more sensitive test. If it is not apparent *a priori*, simulations of different non-Gaussian skies (e.g., different topologies) can show how to best optimise the test.

Unlike the test using Kuiper's statistic, the return mapping statistic is not invariant under rotations about the z-axis. This means that simply taking a randomly oriented sky and applying only the two Euler rotations will not necessarily result in the strongest non-Gaussian signal, since the third orientation angle remains random. In the examples in Fig. 6.6, the simulations were generated with the symmetry axes aligned with the coordinate axes, but reality will not be so convenient. To set up the test to be most effective, we need to determine the symmetry axes of the sky and start with it in the correct orientation.

6.2.4 S-statistic

De Oliveira-Costa et al. (2004a) references therein describe several statistics related to preferred axes and symmetry properties of sky maps. In particular, their S-statistic is helpful for cubic torus models. This statistic measures the degree of symmetry between two hemispheres and is defined by:

$$S(\hat{\mathbf{n}}_i) \equiv \frac{1}{N_{pix}} \left[\delta T(\hat{\mathbf{n}}_j) - \delta T(\hat{\mathbf{n}}_{ij}) \right]^2$$

where $\hat{\mathbf{n}}_{ij}$ denotes "the reflection of $\hat{\mathbf{n}}_j$ in the plane whose normal is $\hat{\mathbf{n}}_i$", i.e.,

$$\hat{\mathbf{n}}_{ij} = \hat{\mathbf{n}}_{\mathbf{j}} - 2(\hat{\mathbf{n}}_i \dot{\hat{n}}_j) \hat{\mathbf{n}}_i$$

The smaller the value of $S(\hat{\mathbf{n}})$, the more symmetry exists. (Note as well that the statistic is identical at $\hat{n}$ and at $-\hat{n}$.) If the S-statistic is very low at the North pole for example, that means that there is a high degree of symmetry between the temperature at

(a) GRF CMB   (b) T444 CMB   (c) T333 CMB   (d) T222 CMB

(e) GRF R−map   (f) T444 R−map   (g) T333 R−map   (h) T222 R−map

(i) GRF K−map   (j) T444 K−map   (k) T333 K−map   (l) T222 K−map
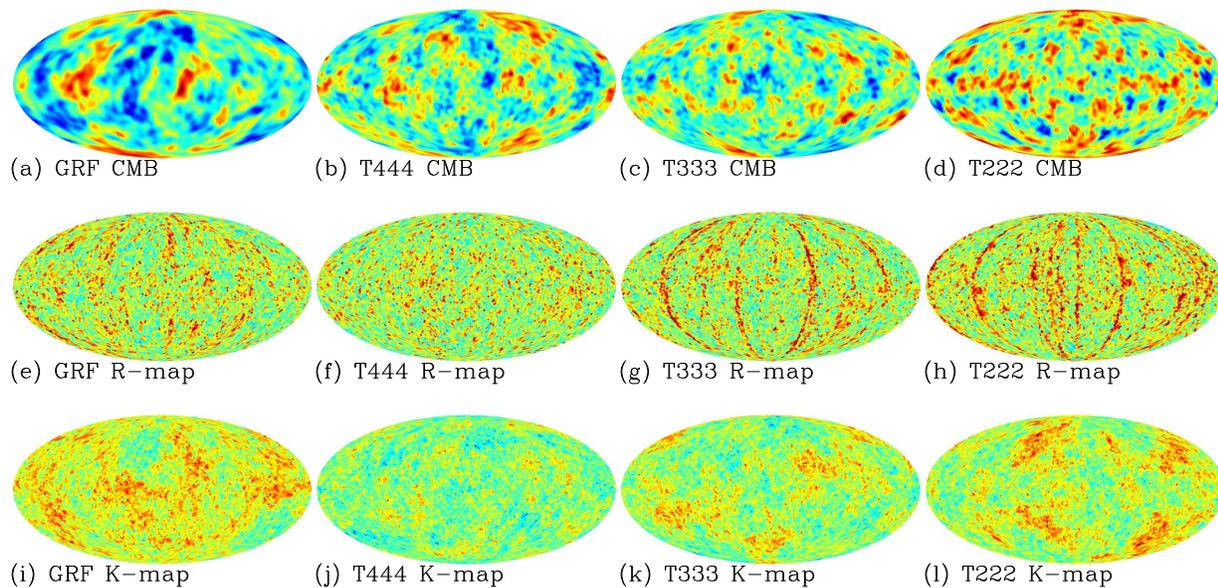
Figure 6.6 Example of return mapping and Kuiper's statistic analyses for cubic torus universes of different cell sizes. The input CMB (a single randomly chosen realisation) is shown on the top row, each rotated to where it's S-statistic minima lie on the x and z coordinate axes. The resulting R-statistic and K-statistic maps are shown in the middle and bottom rows, respectively. Left to right are realisations for Gaussian random or torus models T444, T333, and T222 (see § 6.3.1). The non-Gaussian signal measured by the R-statistic is clear for the torus models with small cell sizes. The Kuiper's statistic is not effective for this model, however.
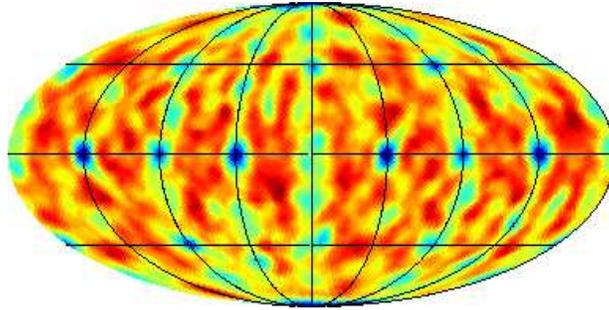
Figure 6.7 Example of S-statistic map for the cubic torus simulation with $L = 2 * R_H$ (see 6.3.1) shown in Fig. 6.6 (d).

a point above the galactic plane and its counterpart at the opposite latitude and the same longitude. An example is shown in figure 6.7.

De Oliveira-Costa et al. (1996) show how toroidal models can be detected using this statistic. By itself, it is not a very sensitive test. (Using COBE, they found a limit of $L \geq 3h^{-1}$Gpc for the smallest dimension.) But finding the minima of this statistic and orienting the map such that they lie on the coordinate axes allows us to then test the return mapping statistic at each rotation from this starting point. We can then narrow the test to only where we expect the non-Gaussian signature to be strongest and maximise the signal to noise.

Note that different topological models may be more effectively detected using different statistical tests. De Oliveira-Costa et al. (2004) define several statistics, as do Land & Magueijo (2005). We find the S-statistic most useful for the toroidal models we will consider.

6.3 Application

6.3.1 Cubic torus simulations

Compact topologies can introduce non-Gaussian behaviour in the CMB fluctuations, and the tools developed by Coles et al. and Chiang et al. may be able to detect them. Dineen et al. (2005) apply the Kuiper's statistical test to compact models of Rocha et al. (2004) and find that very small sized Universes are detectable with their method. As described above, however, their analysis did not take advantage of the additional morphological information in the rotations.

We apply our method to the similar models of Riazuelo et al. (2004a,b). These models are based on the WMAP best fit ΛCDM parameters (first year results), with $R_{LSS} = 6.23 * R_H$ and have scale lengths $L = 2 * R_H$, $L = 3 * R_H$, and $L = 4 * R_H$ (which will be referred to hereafter as T222, T333, and T444). (R_{LSS} and R_H are the radius of the last scattering surface and the Hubble radius, respectively.)

Figure 6.8 show the averaged "super maps" (Chiang et al., 2003) for GRF and T222
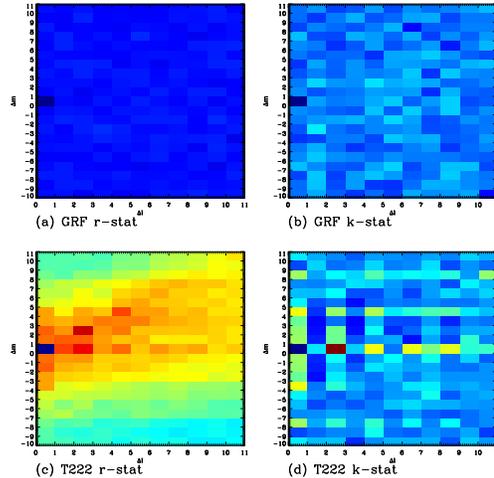
Figure 6.8 "Super maps" for GRF (*top*) and T222 (*bottom*) simulations, averaged over 1000 realisations. The colour indicates the degree of non-Gaussianity, and the scale is set to bring out the structure of the non-Gaussianity in the T222 cases. The *left* (*right*) column shows the results using the R-statistic (K-statistic).

simulations, both using the R- and K-statistics. For the K-statistic applied to the T222, there is increased non-Gaussian signal at specific gradients $\Delta_\ell = \{0, 2, 4, 6, 8\}$ and $\Delta_m = \{0, \pm 4, \pm 8\}$. For the R-statistic, the non-Gaussian signal is seen at many other gradients as well. In any case, these "super maps" can serve as a guide to extract the non-Gaussian signal with the highest signal-to-noise. With simulations, we can find out how sensitive a test it is possible to construct.

Calculating the S-statistic for cubic torus models shows symmetry both aligned with the axes of the torus as well as with the axes pointing toward the next nearest copies. See Figure 6.7 for an example. Finding the minima in the S-statistic and looking for the corresponding minima expected at multiples of $\pi/4$ intervals around gives an indication of this topology. Additionally, if these axes can be found with the S-statistic, the map can then be rotated such that the return mapping statistic will also show the cubic geometry when plotted as a map, as shown in figure 6.6. (Simulations show that the Kuiper's statistic is not as sensitive as the return mapping statistic for these models.)

Our method is then as follows:

- For each map, calculate the S-statistic in all directions and rotate the map such that the minima of that statistic are at the z- and x-axes.

- In this orientation, calculate the return mapping statistic at orientations along great circles of constant longitudes $n\pi/4$ (including points within $\pm 3°$ to account for some of the uncertainty). At each orientation, we calculate $\bar{\chi}_R^2$ for gradients $0 \leq \Delta\ell \leq 10$ and $-10 \leq \Delta m \leq 10$, using phases for harmonic coefficients $2 \leq \ell \leq 32$. (We use a smoothing scale of $R = 2$ for a binning of $M = 128^2$ pixels; see Chiang et al. 2003.)
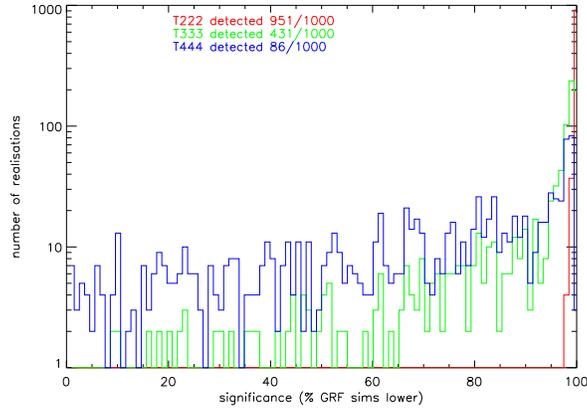
Figure 6.9 Sensitivity of return mapping analysis to different sized cubic torus models: red, $L = 2R_{\mathrm{H}}$; green, $L = 3R_{\mathrm{H}}$; blue, $L = 4R_{\mathrm{H}}$. The histogram counts the number of realisations that has the given statistical significance compared to Gaussian simulations.

- Compare the set of return mapping statistics for the map to an ensemble of simulations analysed identically. For each of 1000 simulations, we perform the above analysis, and then at each point, we calculate the expected mean and RMS variation for the return mapping statistic, $\bar{\chi}^2_{\mathrm{R}}$. Thus for each realisation, we have a set of statistics and can calculate how much it varies from this expectation (using, e.g., a χ^2 to compare.) A given realisation can then be assigned a significance based on where it falls in the distribution of such statistics for the ensemble of Gaussian random simulations.

We test this method using realisations for cubic torus models with cell lengths of two, three, and four times the Hubble radius. The sensitivity of the test is shown in figure 6.9. For the smallest cubic torus models, T222, the return mapping statistic almost always detects non-Gaussianity at the 99% level. For larger cell sizes, the sensitivity of the method drops to detecting fewer than half of the T333 models and 8% of the T444 models. This may be improved by fine tuning the test to use only those gradients and harmonics which show the strongest correlations for the given topology, but this work remains to be done.

It is difficult to compare the effectiveness of this method with that of Dineen et al. (2005). Their "Model 1" is the cubic torus case, and with $j \equiv L/R_{\mathrm{h}} = 0.5$ should be comparable to our T333.[1] But their analysis does not test many different skies generated for each topology, but at most five and only for the $j = 0.5$ hypertorus case. (Confusingly, what they refer to as a "realisation" is a new set of random rotations.) They simply state that the results are inconclusive due to the variation from one simulation to another.

---

[1] R_{h} is the radius of the horizon, while R_{H} is the Hubble radius.

(a) T441 R−map   (b) T441 K−map
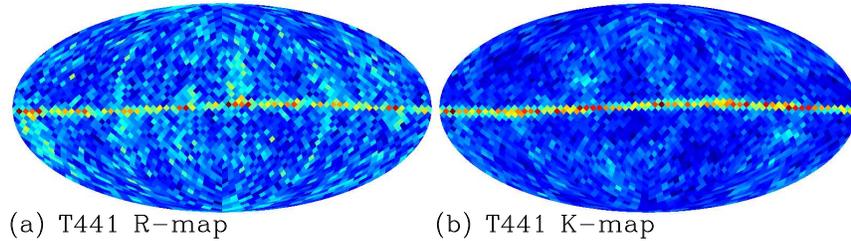
Figure 6.10 Example statistical maps for asymmetric torus model where the torus has $L_x = L_y = 4 * R_\mathrm{H}$ and $L_z = R_\mathrm{H}$. Perpendicular to this axis (i.e., horizontally), both the Kuiper's statistic (left) and the return mapping statistic (right) show a strong non-Gaussian signal.

6.3.2   Asymmetric torus simulations

In cubic torus universes, the length of the fundamental cell size is the same in each of the three directions, so these models have three equivalent symmetry axes. The phase correlations for asymmetric torus models, where one axis is much smaller than the others, manifest themselves in a slightly different way from the cubic torus case. Instead of three circles perpendicular to each axis, there is unsurprisingly strong non-Gaussian signal along one circle perpendicular to the shortest axis and faint signals corresponding to the other axes. Furthermore, the K-statistic here turns out to be at least as sensitive as the R-statistic for that shortest axis. See Figure 6.10.

Clearly, then, such mapping techniques would apply to any such compact models where $L_x \neq L_y \neq L_z$, where the R-statistic would increase at orientations perpendicular to each axis with a strength dependent on the cell size along that axis.

6.3.3   Bianchi simulations

Just as simulations show that the return mapping statistic is most effective for the cubic torus models, likewise we find that the Kuiper's statistic is more effective for some Bianchi models. As shown in Fig. 2.1, however, the morphology of these models varies with the parameters (x, Ω_0), so a statistical test which works well on, e.g., a highly asymmetric model might not work well on other cases. Furthermore, these models are morphologically quite different from the torus cases and do not display the same symmetries. It is therefore more difficult to construct a test which would be sensitive to vorticity in general with a random symmetry axis .

Tests with a few highly asymmetric models show no clear morphology to the non-Gaussian signal that we can use to maximise the signal to noise. Figure 6.11 shows a symmetric example where there is a strong non-Gaussian signal perpendicular to the rotation axis. The plot shows that the K-statistic is more effective for this case than the R-statistic, but that it requires a dominant Bianchi component in order for this signal to rise significantly above the noise.

These tests can perhaps be refined with further work to isolate the best possible combi-

(a) Map with $\alpha=2\mathrm{x}10^{-11}$ (b) R–map   (c) K–map

(d) Map with $\alpha=1\mathrm{x}10^{-10}$ (e) R–map   (f) K–map
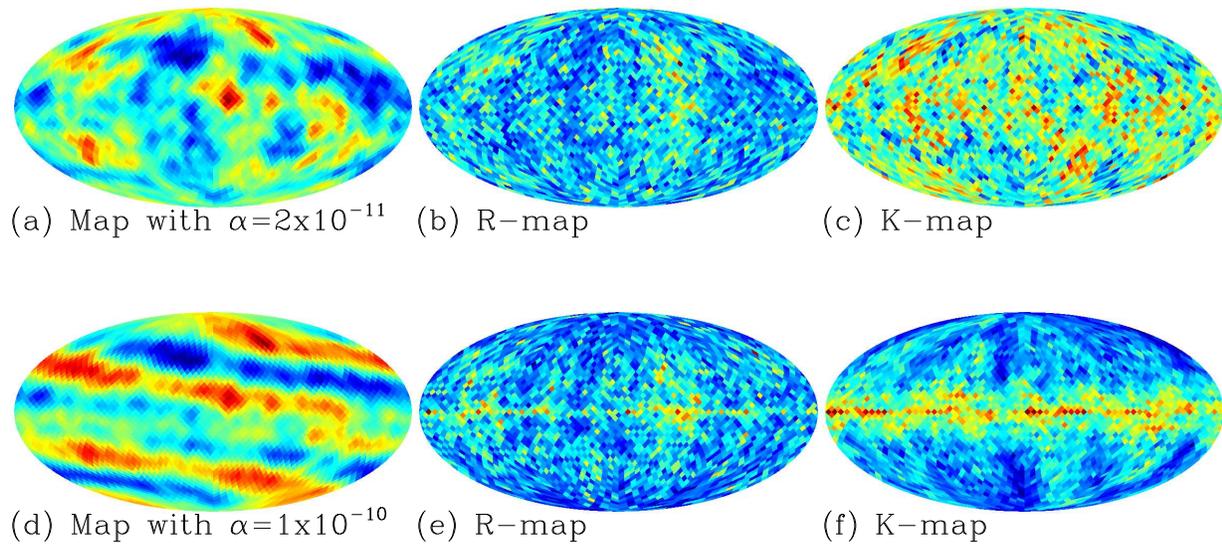
Figure 6.11 Statistical maps for simulated skies with CMB and a Bianchi component, $(x, \Omega) = (0.2, 0.99)$, at the given amplitude, α (see § 2.3.1).

nations of multipoles, gradients, and rotations. But the preliminary analysis suggests that that these methods are not easily applicable to detecting the signature of vorticity.

6.3.4   LILC simulations

Chiang et al. (2003); Chiang et al. (2006) and Coles et al. (2004) all attribute the non-Gaussianity they detect to residual foreground emission in the cleaned maps. Eriksen et al. (2005) describes simulated LILC data which characterise the foreground residuals one can expect in the real ILCs. This is naturally limited by our knowledge of these foregrounds, and we expect the real ILC map to have more residuals than the simulations, particularly along the plane, but they do indicate at least an approximate level and morphology of the residuals.

We can apply our phase analysis to these simulations to determine both the average behaviour and the likelihood of detecting these residuals on a given sky. As in the case of non-trivial topological models, we can see if they can guide us to the most sensitive test for the predicted morphology. They may also be able to inform our judgement of whether the non-Gaussian signals detected in the data are consistent with what we expect from foreground residuals.

Figure 6.12 shows the *average* super maps for 1000 simulations where the LILC is either in the Galactic reference frame or rotated such that the Galactic plane residuals are perpendicular to the equator or likewise diagonal. (These choices are motivated by the assumption that the Galactic plane itself would be the dominant structure in any such residuals, which may or may not be correct.) These maps show strong phase correlations for

(a) GRF K−stat   (b) LILC K−stat   (c) Perp. LILC K−stat   (d) Diag. LILC K−stat
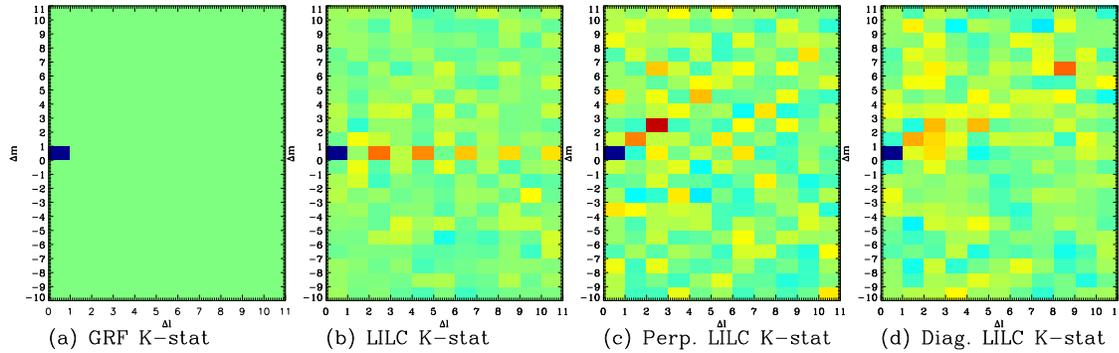
Figure 6.12 "Super" maps showing average significance of Kuiper's statistic for GRF simulations compared to LILC simulations at several orientations. (a) gives those for Gaussian random field (GRF) simulations, while the subsequent plots give those for LILC simulations (b) in the Galactic reference frame and in reference frames where the Galactic plane runs (c) vertically (perpendicular to the coordinate equator) and (d) diagonally.

$\Delta\ell$ even and $\Delta m = 0$ as expected when the planar residuals are aligned with the coordinate equator. The correlations then mix into other modes depending on the orientation of the residuals. At all three orientations, there is some non-Gaussian signal at all gradients, though this is not detectable for an individual realisation as it remains well under the variance. The strongest correlation on average appears not to be in the Galactic frame but the perpendicular frame (c). Note that the return mapping statistic, not shown, does not appear to detect these residuals at any orientation or gradient.

We have generated such super maps for several different ranges of scales, namely $2 \leq \ell \leq 64$, $2 \leq \ell \leq 16$, $16 \leq \ell \leq 32$, and $32 \leq \ell \leq 64$, and find that the phase correlations come from the smaller scales. In particular, the super maps for $\ell \leq 16$ and $16 \leq \ell \leq 32$ do not show a non-Gaussian signal on average. The signal is therefore likely due to residual point sources along the plane. This makes it unclear, then, whether the non-Gaussianities detected by Coles et al. (2004), for example (and which we confirm), at large angular scales, can be attributed to foreground residuals. Obviously, there remain residuals not mapped by the LILC simulations, but based on our current knowledge of the morphology of the Galactic emission, they are unlikely to explain that detected non-Gaussianity.

What is not clear from those averaged super maps is that the variance among realisations is quite large, and these signals cannot necessarily be seen for an individual sky. If we look at the pair which gives the strongest signal in the orientation which is likewise strongest, which is $\Delta m = 2$ and $\Delta\ell = 2$ in the perpendicular orientation, according to Figure 6.12, then we find that even this test only gives a strong ($> 99\%$) detection in ~ 30 out of 1000 simulations. An individual ILC simulation is usually not distinguishable from an individual GRF simulation. Though if we average enough of them, we start to see the non-Gaussian signal, it is simply too weak to expect any individual case to show a clear signal.
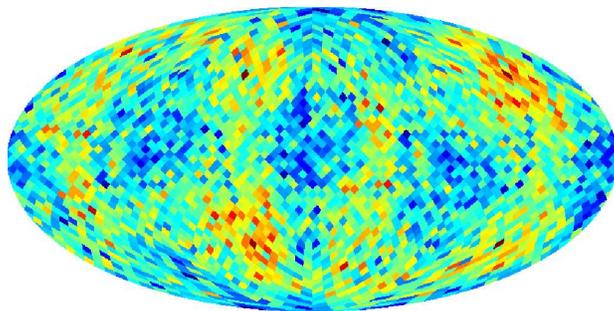
Figure 6.13 Average K-statistic map for LILC simulations, arbitrary units.

But with the "super" maps as our guide, we can construct statistical maps as described in §6.2.3 to see if the additional morphological information will help. We use only the ranges $-4 \leq \Delta m \leq 4$ and $0 \leq \Delta \ell \leq 4$ and the full range of harmonics up to $\ell \leq 64$. The results are not as informative for the LILCs as for the compact models. Figure 6.13 shows that there is no clear morphology to the non-Gaussian signal. In this case, we cannot improve on the technique used in Coles et al. (2004) of examining the distribution over all orientations and comparing to the distribution for the null hypothesis, using for example a simple χ^2 misfit statistic. But again, the variation among different realisations is great, and of the 1000 simulations tested, only ~ 80 of the LILC simulations are anomalous at the $> 95\%$ level by this measure, and only 35 at the $> 99\%$ level.

We must conclude from the above that this method is therefore not an effective tool for analysis of foreground residuals. The simulations indicate that the strongest non-Gaussian signal arising from known foregrounds comes from smaller scales, confirming that what Chiang et al. (2003) detect is likely such residuals but that the detection of Coles et al. (2004) at larger angular scales may not be.

6.3.5   WMAP data

Armed with the above information on the non-Gaussian signatures due to compact topologies or foreground residuals, we turn to the *WMAP* data itself in the form of the LILC map (Eriksen et al., 2004a). As mentioned, incomplete sky coverage introduces correlations between modes which will confuse any true non-Gaussian signals. The internal linear combination maps are the best full sky, foreground cleaned maps of the CMB, despite the known foreground residuals that remain. We shall perform analyses like those described above to see if the phase statistics can tell us about the foreground residuals in the data or the possibility that the universe is finite.

First, we can look for the likely foreground residuals using the perpendicular orientation found in § 6.3.4. The left side of Fig. 6.14 shows the "super" maps for different ranges of ℓ. But as noted, the signal that appears when averaging over an ensemble of simulations is not visible in a single realisation, so it is unsurprising that there is nothing visible in these figures.

K-statistic   R-statistic
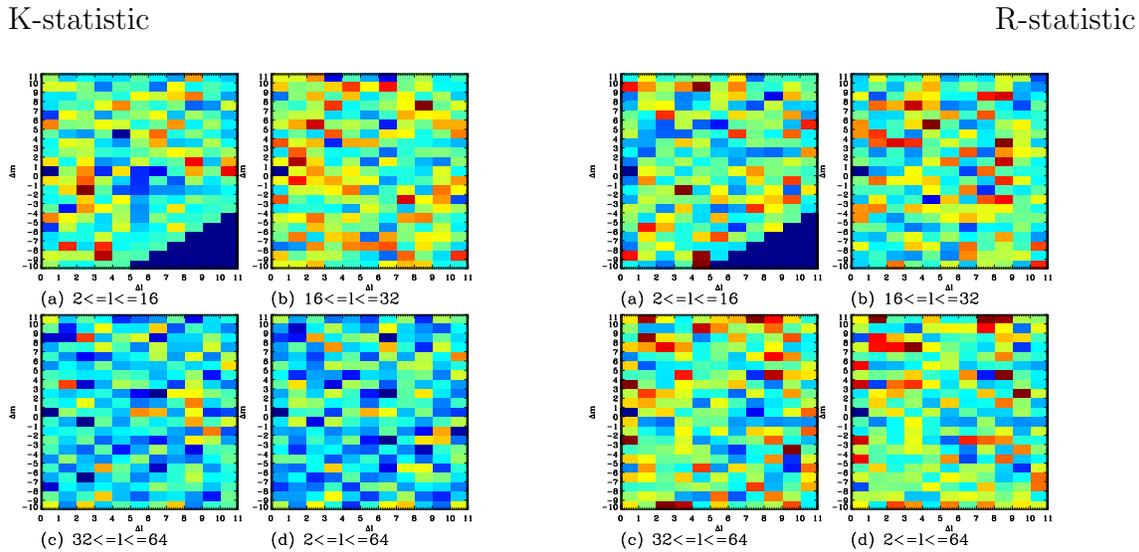
Figure 6.14 "Super" maps from the LILC map for different ranges of ℓ as shown, on the left the R-statistic and on the right the K-statistic (perpendicular orientation).

Nor do we see any indication of the compact topology signature shown in the R-statistic shown at right in Fig. 6.14. The statistical maps are shown in Fig. 6.15, and again, there is little obvious structure though we have attempted to isolate the possible signals following the analysis above of cubic torus models and the LILC simulations. As noted in § 6.3.1, however, a universe with a cell size of only three Hubble radii is only detectable in roughly half of realisations by this method. Therefore, the implied lower limit this method can place on such a compact topology is closer to two Hubble radii. Lastly, for completeness, we look directly at two phase gradients in Fig. 6.16 and see no indication of non-Gaussianity.

Further work is needed to characterise the various phase statistic distributions and to refine these tests to detect the known galactic contamination in these maps. Indeed, point sources are visible by eye along the plane, though this analysis has been at larger angular scales. And though these preliminary results do not show much of interest, as described above, there are many different ways to slice the data in order to search for the correlations.

6.4   Discussion and conclusions

Kunz et al. (2006) have developed a method designed to detect topological models such as hypertorus or slab spaces. Their method uses the full theoretical covariance matrix of the spherical harmonic coefficients, and their results show that it may be far more sensitive than the phase method described here. Both the Kunz et al. analysis and the phase analysis with cubic tori presented here were carried out using simple simulations lacking contributions from late-time ISW, noise, or foreground residuals. It is unclear how each method will be affected by these problems for real-world data or how the computational

(a) R−map tot unrtd

(b) K−map tot unrtd

(c) R−map (Δl,Δm)=(2,2) rtd

(d) K−map (Δl,Δm)=(2,2) rtd

(e) R−map (Δl,Δm)=(2,0) rtd

(f) K−map (Δl,Δm)=(2,0) rtd
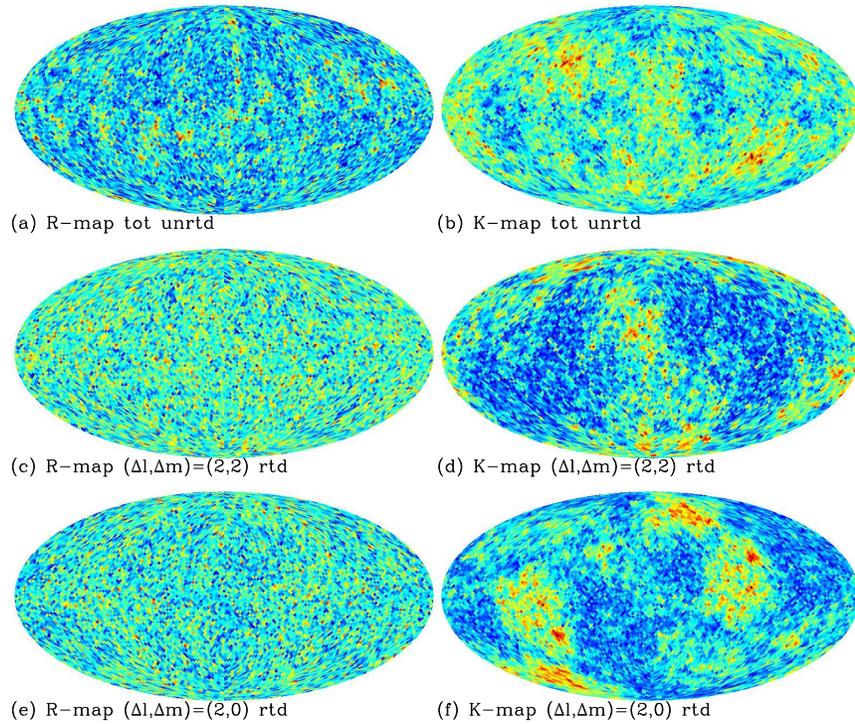
Figure 6.15 Statistical maps for the LILC. The top row shows the total for $0 \leq \Delta \ell \leq 10$ and $-10 \leq \Delta m \leq 10$ with the map in its galactic orientation. The next rows are for the particular values of $\Delta \ell$ and Δm which appear to maximise the signals according to the simulations in § 6.3.1 and 6.3.4. For these rows, the map has been rotated to the orientation that places the axis of maximum symmetry (defined by the S-statistic in § 6.2.4) at the north pole.
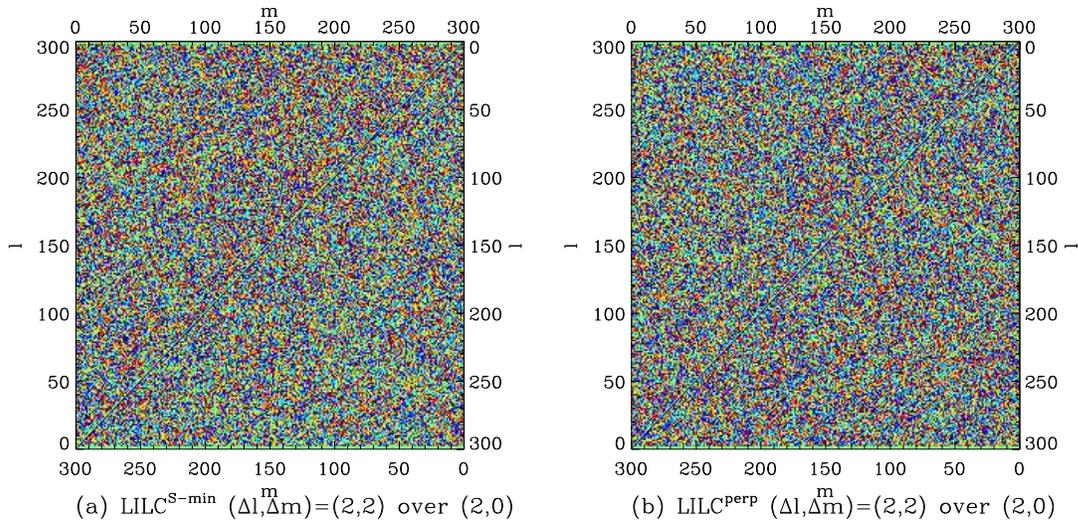
Figure 6.16 Phase gradients for the LILC in the Galactic orientation (a) and the perpendicular orientation (b).

requirements scale. The phase analysis may also be improved upon as described above by further customising the test to include only the gradients and modes where the signal to noise is highest. Furthermore, the method of Kunz et al. requires calculating the covariance matrix for every possible compact model they wish to test in order to look for any signal. Phase statistics are more sensitive when directed by such simulations, as described above, but they can be applied in the absence of any *a priori* model. In any case, phase analysis remains an interesting and promising method for detecting non-Gaussianity that can be adapted to different contexts.

One of the limitations of the methods described above is the fact that they require full sky coverage and are therefore currently only applicable to highly processed, foreground-cleaned maps like the LILC, where foreground residuals that remain may affect the result. The next step toward developing these tools is to define a phase analysis on the cut sky. One could simply compute the pseudoharmonics, i.e. just the harmonic transform of the masked sky, and examine these phases. The sky cut will, however, have introduced correlations between different harmonic modes. For the full sky, the spherical harmonics are orthogonal:

$$K_{\ell\ell'mm'} \equiv \sum_p Y_{\ell m}^*(p) Y_{\ell'm'}(p)\Omega_p = \delta_{\ell\ell'}\delta_{mm'}, \qquad (6.3)$$

(where Ω_p is the area of a pixel). But summing over the incomplete sky, the second inequality does not hold. Fig. 6.17 shows some examples. For an azimuthally symmetric cut (upper left triangles), the coupling introduced is only for $m = m'$, but for a more complicated mask, many modes are coupled.
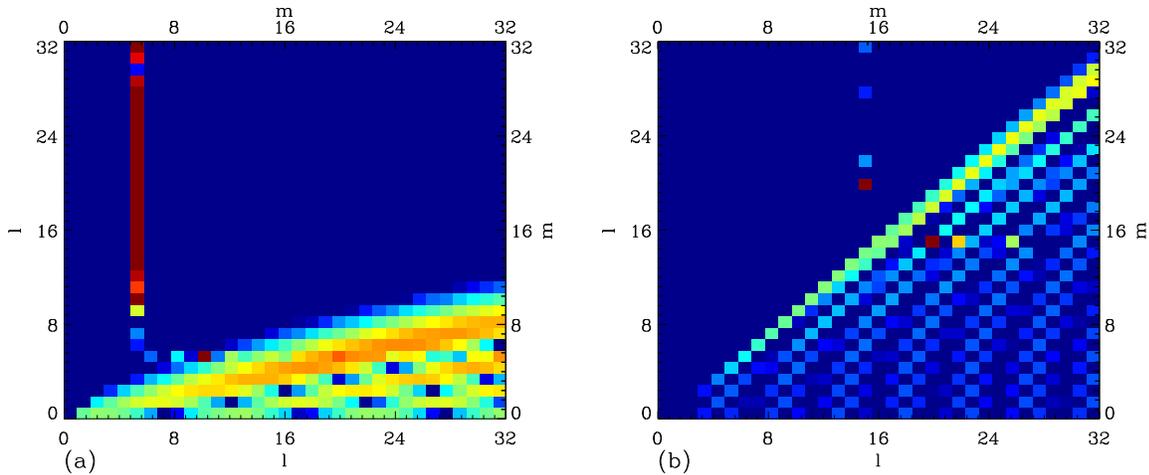
Figure 6.17 Spherical harmonic coupling due to incomplete sky coverage (arbitrary units). Each plot shows a slice of the matrix, $K_{\ell\ell'mm'}$, where on the left, we set $(\ell', m') = (10, 5)$ and on the right $(\ell', m') = (20, 10)$. In each plot, the lower right triangle is for the Kp0 mask while the upper left is for a symmetric cut of $|b| > 20°$.

That coupling between harmonics will then give a spurious non-Gaussian signal. With Monte Carlo simulations to characterise the correlations induced by the sky cut, it is possible that a test could still be defined using the pseudoharmonic phases. This approach may prove to lack sufficient sensitivity for any realistic application. An alternative is to use functions defined to be orthonormal on the cut sky. The latter approach is limited by computational tractability to large scales (e.g., $\ell \lesssim 50$ for reasonable sky cuts), but these are the scales that are interesting for studying possible large-scale anisotropy. At very small scales, we can study small patches of the sky that can be approximated as flat, and therefore Fourier phase statistics may be applicable to studying the primordial non-Gaussianity arising in some inflationary models or perhaps the effects of correlated noise.

Phase statistics have also been used to study foregrounds by, e.g., Naselsky et al. (2006) and references therein. Cross-correlation in harmonic space with phases is similar to that in pixel space, as is defining filters based again on assumptions that the residuals be Gaussian and statistically isotropic. The phases of harmonic functions defined on the full sky are less intuitive for studying angular structures or the spectral variation of foregrounds, but they may nevertheless be a useful tool for testing for the presence of residuals.

Chapter 7

Updates with the *WMAP* 3-year Data

The *WMAP* 3-year data release came after the work in the previous chapters was either published or at least submitted for publication. Rather than merging new material into previously published results, we here give brief updates. The first subsection gives the results of the 15 regions analysis on the 3-year data. The second subsection discusses a change in the *WMAP* analysis from the first to the 3-year data release, the use of the K–Ka template to map the dominant non-thermal emission at low frequencies. This template is then also used in the Bianchi 3-year update (third subsection).

7.1 Foreground Analysis on Regions

The analysis in Chapter 5 has been repeated with the *WMAP* 3-year data. The results are shown in Tables 7.1 through 7.3. The anomalous dust spectrum is shown in Fig. 7.1. These tables are similar to those in Chapter 5 but with additional columns giving the change between the 3-year and first-year values as a fraction of the uncertainty.

The results in general are very comparable, with only a few regions (as expected) showing changes larger than the error bars. There is some increase in the dust emissivity in the V bands, which is due to an unusually large (2σ) increase the one region that has a small error bar and thus strongly affects the average. The fitted anomalous dust index shown in Fig. 7.1 is thus a bit flatter for the combined 3-year data than for the first-year, now -2.6 instead of -2.85. As shown in Fig. 4.3, even this value overpredicts the anomalous component in dust regions where the 408 MHz map shows no such emission, if (as assumed by the *WMAP* team) it is a hard synchrotron component and if there is no spectral flattening.

We have also compared the fits from the first-year data release given in Chapter 5 with those to the first-year maps from the 3-year data release. The difference here is primarily the beam calibration. (There is a large-scale gain calibration difference as well, but for these small regions, that should have no effect.) This does introduce some small differences

Field No.	Template	$\frac{T_{\mathrm{K}}}{T_{\mathrm{H}_\alpha}}$ μK R^{-1}	WMAP3-1 $\Delta_{\mathrm{K}}/\sigma$	$\frac{T_{\mathrm{Ka}}}{T_{\mathrm{H}_\alpha}}$ μK R^{-1}	WMAP3-1 $\Delta_{\mathrm{Ka}}/\sigma$
1	F03	$11.3^{\pm4.6}$	0.0	$10.0^{\pm4.1}$	0.0
	DDD	$9.7^{\pm4.2}$	-0.0	$6.1^{\pm3.9}$	-0.0
2	F03	$6.0^{\pm2.9}$	0.4	$2.0^{\pm2.7}$	0.4
	DDD	$7.0^{\pm2.8}$	0.6	$2.2^{\pm2.6}$	0.3
3	F03	$8.3^{\pm2.8}$	0.9	$5.8^{\pm2.4}$	1.0
	DDD	$3.2^{\pm3.4}$	0.3	$0.3^{\pm3.0}$	0.6
4	F03	$8.2^{\pm2.4}$	0.3	$2.6^{\pm2.0}$	-0.8
	DDD	$7.1^{\pm1.7}$	-0.0	$0.3^{\pm0.8}$	-0.8
5	F03	$9.0^{\pm1.0}$	-0.5	$4.2^{\pm0.9}$	-0.6
	DDD	$9.8^{\pm1.2}$	-0.2	$4.9^{\pm1.1}$	-0.2
Avg.	F03	$8.6^{\pm0.8}$	0.1	$4.1^{\pm0.7}$	-0.3
	DDD	$8.4^{\pm0.9}$	-0.0	$2.0^{\pm0.6}$	-1.4
Kp2	F03	$7.6^{\pm0.9}$	-0.1	$3.7^{\pm0.9}$	0.0
	DDD	$7.5^{\pm0.9}$	-0.1	$3.6^{\pm0.9}$	0.0
	T_e=4000K	8.0		3.6	
	T_e=5000K	8.9		4.1	
	T_e=6000K	9.8		4.5	
	T_e=7000K	10.6		4.9	
	T_e=8000K	11.4		5.2	

Table 7.1 Free-free emission as determined by CC method for K-and Ka-band *WMAP* data. Also shown is the change, Δ, from the first-year results in Table 5.3 as a fraction of the error bar.

Field No.	K	$\frac{\Delta_K}{\sigma}$	Ka	$\frac{\Delta_{Ka}}{\sigma}$	Q	$\frac{\Delta_Q}{\sigma}$	V	$\frac{\Delta_V}{\sigma}$	W	$\frac{\Delta_W}{\sigma}$
				Dust emissivity relative to FDS						
1	$7.8^{\pm 5.4}$	-0.1	$0.5^{\pm 4.4}$	0.1	$3.0^{\pm 3.3}$	0.1	$2.7^{\pm 2.2}$	1.1	$-0.9^{\pm 1.7}$	0.5
	$8.1^{\pm 5.3}$	-0.1	$0.9^{\pm 4.3}$	0.1	$3.7^{\pm 3.3}$	0.1	$3.3^{\pm 2.2}$	1.1	$-0.5^{\pm 1.7}$	0.5
2	$6.6^{\pm 1.4}$	0.0	$1.1^{\pm 1.2}$	-0.5	$1.3^{\pm 1.0}$	-0.6	$0.1^{\pm 0.9}$	0.8	$0.2^{\pm 0.8}$	-0.1
	$6.4^{\pm 1.4}$	0.1	$-0.2^{\pm 1.2}$	-0.9	$-0.9^{\pm 1.0}$	-1.2	$-1.8^{\pm 0.9}$	0.3	$-0.8^{\pm 0.8}$	-0.5
3	$11.5^{\pm 4.8}$	-0.1	$3.5^{\pm 4.0}$	-0.5	$9.2^{\pm 3.2}$	0.5	$1.8^{\pm 2.8}$	-0.9	$-0.6^{\pm 2.4}$	0.0
	$7.6^{\pm 4.4}$	-0.0	$0.1^{\pm 3.7}$	-0.4	$5.2^{\pm 3.1}$	0.4	$-0.5^{\pm 2.7}$	-1.0	$-1.0^{\pm 2.3}$	-0.1
4	$8.3^{\pm 5.7}$	0.7	$6.6^{\pm 4.6}$	-0.1	$9.4^{\pm 3.7}$	0.7	$3.4^{\pm 3.0}$	0.3	$0.4^{\pm 2.5}$	0.8
	$7.3^{\pm 5.4}$	0.8	$2.8^{\pm 4.5}$	-0.4	$6.5^{\pm 3.7}$	0.3	$2.3^{\pm 2.9}$	0.1	$-0.4^{\pm 2.4}$	0.6
5	$11.7^{\pm 2.0}$	-0.4	$3.6^{\pm 1.6}$	0.6	$4.3^{\pm 1.4}$	1.0	$0.2^{\pm 1.2}$	0.4	$1.9^{\pm 1.0}$	2.2
	$12.2^{\pm 1.9}$	-0.4	$2.8^{\pm 1.6}$	0.1	$2.0^{\pm 1.4}$	0.2	$-2.2^{\pm 1.2}$	-0.4	$0.5^{\pm 1.0}$	1.7
6	$\mathbf{6.7^{\pm 0.7}}$	**-0.0**	$\mathbf{2.8^{\pm 0.6}}$	**-0.0**	$\mathbf{2.3^{\pm 0.5}}$	**0.4**	$\mathbf{1.9^{\pm 0.4}}$	**2.2**	$\mathbf{2.3^{\pm 0.3}}$	**-0.1**
	$\mathbf{6.6^{\pm 0.7}}$	**0.1**	$\mathbf{2.6^{\pm 0.6}}$	**0.0**	$\mathbf{2.2^{\pm 0.4}}$	**0.5**	$\mathbf{1.9^{\pm 0.4}}$	**2.2**	$\mathbf{2.2^{\pm 0.3}}$	**-0.0**
7	$\mathbf{8.7^{\pm 3.7}}$	**0.1**	$\mathbf{2.0^{\pm 3.2}}$	**-0.2**	$\mathbf{0.9^{\pm 2.8}}$	**0.2**	$\mathbf{0.9^{\pm 2.4}}$	**0.1**	$\mathbf{-1.5^{\pm 2.1}}$	**-0.8**
	$\mathbf{8.0^{\pm 3.6}}$	**0.0**	$\mathbf{2.1^{\pm 3.2}}$	**-0.2**	$\mathbf{0.6^{\pm 2.8}}$	**0.3**	$\mathbf{0.7^{\pm 2.4}}$	**0.1**	$\mathbf{-1.5^{\pm 2.1}}$	**-0.7**
8	$\mathbf{9.0^{\pm 2.4}}$	**0.2**	$\mathbf{3.4^{\pm 2.2}}$	**0.4**	$\mathbf{3.0^{\pm 2.0}}$	**0.3**	$\mathbf{0.9^{\pm 1.8}}$	**-0.8**	$\mathbf{2.6^{\pm 1.6}}$	**-0.3**
	$\mathbf{8.9^{\pm 2.4}}$	**0.3**	$\mathbf{3.4^{\pm 2.2}}$	**0.3**	$\mathbf{3.0^{\pm 2.0}}$	**0.2**	$\mathbf{1.0^{\pm 1.8}}$	**-0.8**	$\mathbf{2.7^{\pm 1.6}}$	**-0.3**
9	$\mathbf{7.5^{\pm 0.6}}$	**0.3**	$\mathbf{2.9^{\pm 0.5}}$	**-0.0**	$\mathbf{1.8^{\pm 0.5}}$	**0.0**	$\mathbf{1.2^{\pm 0.4}}$	**-0.3**	$\mathbf{1.6^{\pm 0.4}}$	**-0.7**
	$\mathbf{7.5^{\pm 0.6}}$	**0.3**	$\mathbf{2.8^{\pm 0.5}}$	**-0.1**	$\mathbf{1.8^{\pm 0.5}}$	**-0.1**	$\mathbf{1.3^{\pm 0.4}}$	**-0.2**	$\mathbf{1.6^{\pm 0.4}}$	**-0.7**
10	$\mathbf{11.4^{\pm 1.3}}$	**-0.4**	$\mathbf{4.3^{\pm 1.1}}$	**-0.4**	$\mathbf{1.1^{\pm 0.9}}$	**-0.7**	$\mathbf{-1.1^{\pm 0.8}}$	**-1.3**	$\mathbf{0.7^{\pm 0.7}}$	**-0.6**
	$\mathbf{11.6^{\pm 1.3}}$	**-0.4**	$\mathbf{4.7^{\pm 1.1}}$	**-0.2**	$\mathbf{1.2^{\pm 0.9}}$	**-0.6**	$\mathbf{-1.0^{\pm 0.8}}$	**-1.4**	$\mathbf{0.8^{\pm 0.7}}$	**-0.6**
11	$17.0^{\pm 7.4}$	-0.3	$9.9^{\pm 6.1}$	0.0	$6.6^{\pm 5.3}$	0.0	$5.1^{\pm 4.5}$	0.4	$6.0^{\pm 3.9}$	1.0
	$16.7^{\pm 7.4}$	-0.1	$10.0^{\pm 6.1}$	0.3	$6.7^{\pm 5.3}$	-0.1	$5.4^{\pm 4.5}$	0.3	$5.8^{\pm 3.8}$	1.0
12	$13.5^{\pm 7.1}$	-0.3	$2.5^{\pm 5.8}$	-0.4	$3.6^{\pm 4.7}$	-0.2	$-0.3^{\pm 3.8}$	1.4	$-2.8^{\pm 3.2}$	-0.2
	$13.2^{\pm 7.1}$	-0.3	$2.3^{\pm 5.8}$	-0.4	$3.4^{\pm 4.7}$	-0.2	$-0.5^{\pm 3.8}$	1.4	$-2.9^{\pm 3.2}$	-0.2
13	$5.6^{\pm 6.1}$	0.3	$-0.9^{\pm 5.2}$	0.5	$5.4^{\pm 4.0}$	0.8	$1.9^{\pm 2.8}$	1.1	$-0.6^{\pm 2.2}$	0.1
	$3.9^{\pm 6.0}$	0.1	$-0.7^{\pm 5.2}$	0.4	$5.4^{\pm 4.0}$	0.8	$2.0^{\pm 2.8}$	1.1	$-0.6^{\pm 2.2}$	0.1
14	$9.1^{\pm 4.5}$	0.2	$9.4^{\pm 3.4}$	0.4	$5.3^{\pm 2.3}$	1.2	$4.5^{\pm 1.6}$	0.6	$2.7^{\pm 1.2}$	1.8
	$8.0^{\pm 4.5}$	0.2	$9.0^{\pm 3.4}$	0.5	$4.5^{\pm 2.3}$	1.1	$3.6^{\pm 1.6}$	0.6	$2.3^{\pm 1.2}$	1.7
15	$5.8^{\pm 3.1}$	-0.7	$-0.7^{\pm 2.7}$	0.1	$2.0^{\pm 2.4}$	1.1	$-0.4^{\pm 2.1}$	0.3	$2.4^{\pm 1.8}$	1.3
	$5.8^{\pm 3.1}$	-0.7	$-0.6^{\pm 2.7}$	0.1	$2.0^{\pm 2.4}$	1.1	$-0.4^{\pm 2.1}$	0.4	$2.5^{\pm 1.8}$	1.3
Avg.	$7.8^{\pm 0.4}$	-0.0	$2.9^{\pm 0.3}$	-0.1	$2.2^{\pm 0.3}$	0.6	$1.3^{\pm 0.2}$	1.6	$1.6^{\pm 0.2}$	0.3
	$7.7^{\pm 0.4}$	0.0	$2.7^{\pm 0.3}$	-0.3	$1.9^{\pm 0.3}$	0.3	$1.0^{\pm 0.2}$	1.3	$1.5^{\pm 0.2}$	0.2
Kp2	$6.6^{\pm 0.3}$	-0.3	$2.4^{\pm 0.3}$	-0.1	$1.4^{\pm 0.3}$	0.1	$0.8^{\pm 0.3}$	-0.0	$1.1^{\pm 0.3}$	-0.0
	$6.5^{\pm 0.3}$	-0.3	$2.4^{\pm 0.3}$	-0.1	$1.4^{\pm 0.3}$	0.1	$0.8^{\pm 0.3}$	-0.0	$1.1^{\pm 0.3}$	-0.0

Table 7.2 Dust-correlated emissivity in the 5 *WMAP3* bands. Also shown is the change, Δ, from first-year results in Table 5.4.

Field	Synchrotron fit amplitudes					
	$\frac{K}{408}$ $(\times 10^6)$	$\Delta_{\mathrm{K}}/\sigma$ $(\times 10^6)$	$\frac{Ka}{408}$ $(\times 10^6)$	$\Delta_{\mathrm{Ka}}/\sigma$	$\frac{Q}{408}$	$\Delta_{\mathrm{Q}}/\sigma$
11	$3.15^{\pm 1.48}$	-0.6	< 1.78	0.2	< 0.88	0.7
12	$3.24^{\pm 1.28}$	0.2	< 0.57	-1.2	< 0.28	-0.6
13	< 5.51	-0.1	$1.14^{\pm 0.51}$	0.5	< 0.53	-1.1
14	< 3.62	-0.1	< 0.57	0.2	< 0.34	1.5
15	< 5.43	-0.6	< 0.51	-0.4	< 0.20	-0.9
Avg.	$2.35^{\pm 0.63}$	-0.6	< 0.48	-0.4	< 0.16	-0.0
Kp2	5.46 ± 0.79	-0.1	< 2.95	-0.1	< 2.09	0.1

Field	Synchrotron spectral index β		
	$\frac{K}{408}$	$\frac{Ka}{408}$	$\frac{Q}{408}$
11	$-3.15^{+0.10}_{-0.16}$	< -3.01	< -3.03
12	$-3.14^{+0.08}_{-0.12}$	< -3.27	< -3.28
13	< -3.01	$-3.12^{+0.08}_{-0.14}$	< -3.14
14	< -3.11	< -3.27	< -3.24
15	< -3.01	< -3.30	< -3.35
Avg.	$-3.22^{+0.06}_{-0.08}$	< -3.31	< -3.41
Kp2	$-3.01^{+0.03}_{-0.04}$	< -2.90	< -2.84

Table 7.3 Synchrotron fits between 408MHz and the *WMAP3* K, Ka, and Q bands. Also shown is are the changes, Δ, from the first-year results in Table 5.5.
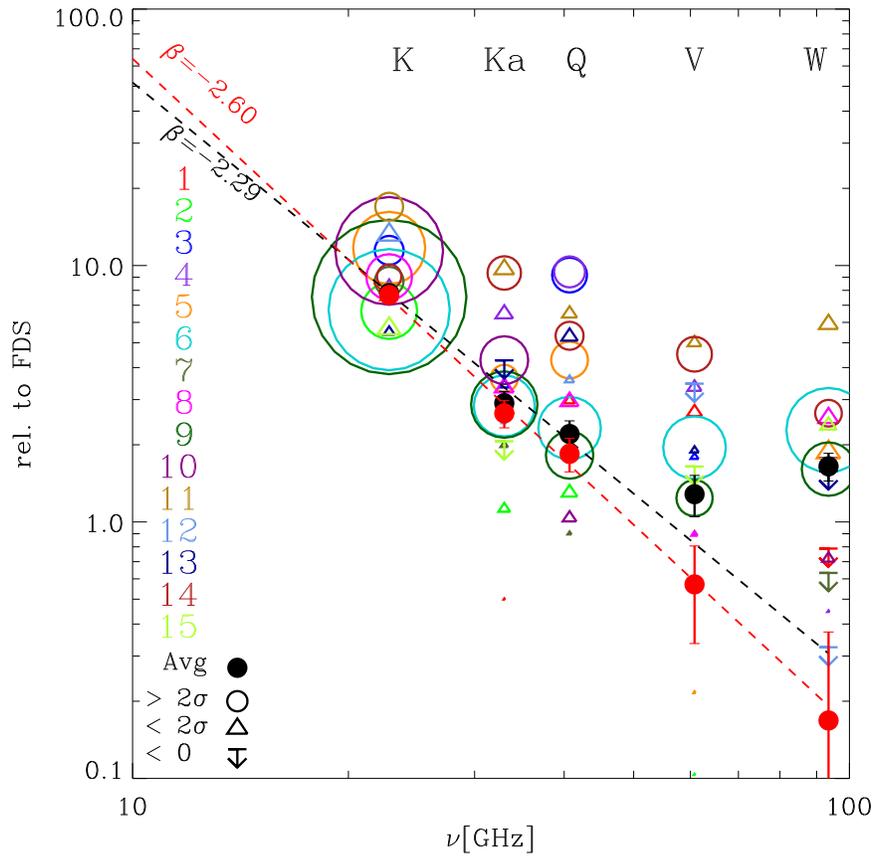
Figure 7.1 Summary of dust emissivities (antenna temperature units, relative to FDS8 at 94 GHz) from C–C analysis for *WMAP 3-year* analysis. See Fig. 5.4 and text in §5.5.2.

(a) K–Ka−Synch−FF   (b) FF (Finkb.)

(c) Dust (FDS)   (d) Synch. (Haslam)

0. μK   300. μK
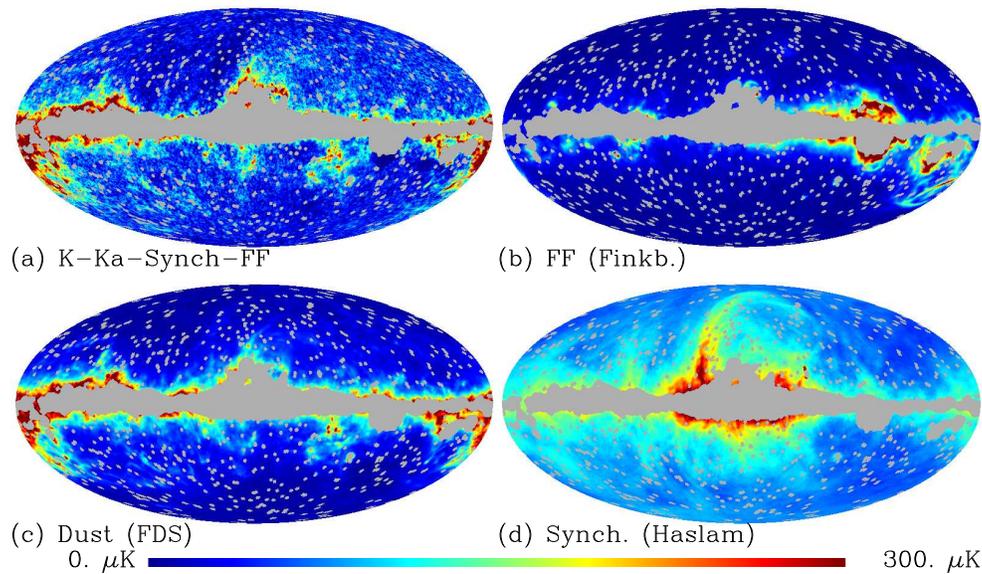
Figure 7.2 Comparison of the anomalous emission in the *WMAP* K–Ka map with the template maps from other frequencies. In (a) is the K–Ka map, where the synchrotron and free-free emission based on fitting the Haslam and Finkbeiner templates has been subtracted, leaving only the anomalous component (and a very small predicted thermal dust component.)

in the fit results, though again, they are within the error bars. Further differences from year-to-year also show that the noise does have an effect on these fits, again within the error bars. The differences from the analyses of the first-year data release and of the 3-year combined data are then due to both of these effects. The flattening of the spectrum of the anomalous dust-correlated component is due to such difference in the few regions that dominate the Q- and V-band averages. Because they remain within the uncertainties, these differences do not affect our conclusions.

7.2   K–Ka Internal Template

The *WMAP* team treat the anomalous emission as a hard synchrotron component. But as shown in Fig. 4.3, the spectral behaviour of this component implies too much emission in many regions of the sky at 408 MHz if there is no spectral flattening. Furthermore, it is striking how the morphology of the anomalous emission so closely resembles that of the dust. Figure 7.2 shows the map of K–Ka, where the best-fit free-free (from Finkbeiner Hα) and synchrotron (from Haslam) components subtracted. The residuals should primarily be the anomalous component, and they follow the dust very closely.

The *WMAP* 3-year analysis uses the K–Ka map itself as an internal template. They assume that it is dominated by synchrotron emission, with a free-free contribution, but in contrast to the first-year analysis, they do not constrain the fit to force this component to follow a particular power-law spectrum. (Such a constraint would be unrealistic given the

mix of physical components included.) They instead constrain the dust template to follow the thermal spectrum with $\beta = +2.0$. As before, the free-free is constrained by $\beta = -2.14$. The K–Ka template then fits to whatever emission does not resemble free-free or thermal dust. There remain spectral variations in the different components in the K–Ka template, but this template remains far more effective (i.e., leaves less visible residual emission) at removing the non-thermal foregrounds, whatever their source, than using Haslam. The interpretation of the resulting fit amplitudes then requires consideration about how the components are mixed in K–Ka. They assume that the K–Ka not mapped by Hα or thermal dust is synchrotron and that it has a spectral index near -2.9. (Completely aside from the issue of the anomalous emission, the north polar spur can be seen to cause problems as shown in Fig. 4.5.)

Figure 7.3 (a) shows the spectra (in RMS antenna temperature over the masked sky) of *unconstrained* template fits to *WMAP* 3-year data. The lines show the expectation assuming the given spectral index and: synchrotron scaled from 408 MHz with $\beta = -3.0$; free-free and anomalous dust scaled to K-band fits; thermal dust at predicted amplitude. The circles show the fit values for each colour-coded component where the Haslam template was used for synchrotron, and the agreement is generally good. It is important to note that the fits also show that even in the K-band, the dominant component is the anomalous dust-correlated emission and not the synchrotron emission correlated with the 408 MHz map.

The squares show what one would expect to get for template fits if the spectra follow the expectation and if instead of Haslam, the K–Ka template is used for synchrotron. These are calculated assuming that there are three morphologically distinct components and that the dust and Hα templates are accurate tracers of two of them. Then any morphologically different component, assumed to be synchrotron, must be fit by the K–Ka template. So the K–Ka amplitude will be determined solely by the synchrotron, and the other components by the residuals. Since K–Ka includes both dust and free-free emission, the corresponding RMS values for those templates will then be lower and the "synchrotron" higher. It is important to remember that these values are now the RMS associated with the template, not with the individual component, since now they are mixed.

The triangles show what actually comes out from such fits with the K–Ka template, and in general, they agree with the expectation (squares). The dust values (red triangles) show a slightly flatter spectrum than expected, though the squares are within their errors. The squares are the predictions assuming that synchrotron emission drives the fit of the K–Ka template to the data and that the dust template then picks up whatever is left of the dust-correlated emission. Given the behaviour of the fits with Haslam, this seems unlikely. The K–Ka fits will be driven by the anomalous dust-correlated emission, which it presumably traces more accurately than the FDS dust prediction. This simply illustrates the difficulty in separating these components. Though it undoubtedly leads to reduced residuals to use this template instead of Haslam when the dust and Hα templates are free to fit without spectral constraints, it is less clear how well it works when those constraints are applied. Figure 7.3 (b) is the same as (a) except that the triangles now represent fits where the dust is constrained to follow a thermal spectrum ($\beta = 1.7$). The results
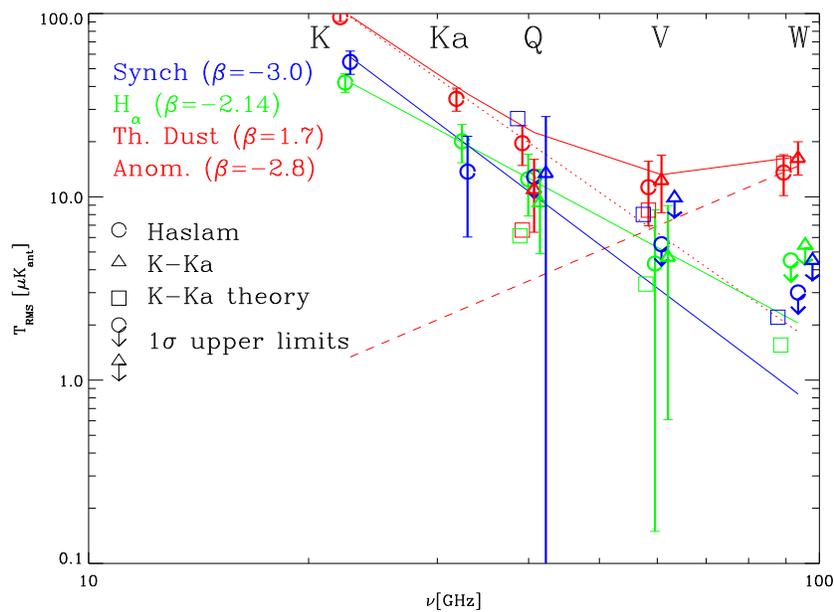
(a)
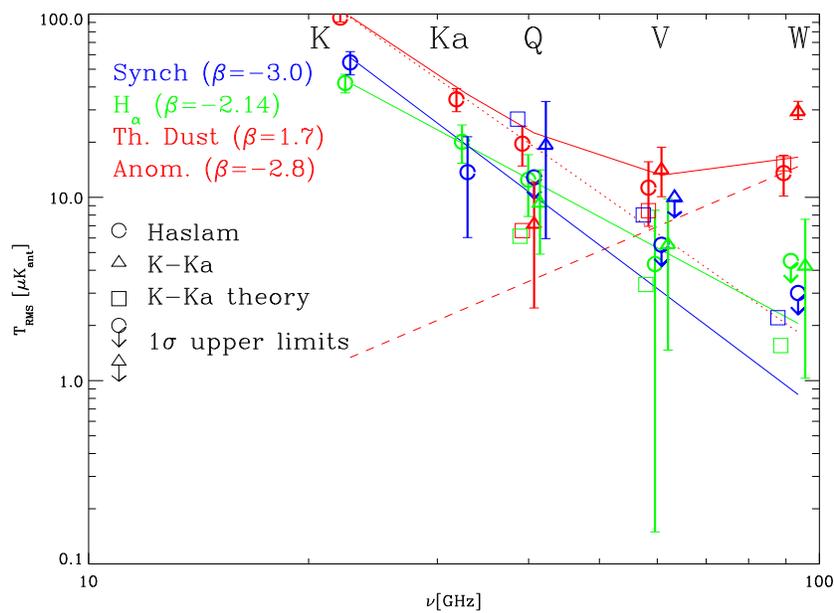
(b)

Figure 7.3 Spectra of foreground template fits using Haslam versus using K–Ka, using an unconstrained (*top*) or constrained (*bottom*) thermal dust spectrum. See text.

are very similar except in the W-band, where the K–Ka fit amplitude (not shown) is now 1σ negative and the thermal dust value is high. (The results are very similar when the constraint is $\beta = 2.0$.)

But the fact that the unconstrained fit values do not differ very significantly from what one would expect from anomalous dust indicates that using K–Ka works reasonably well even though it may not be dominated by synchrotron. This is what the *WMAP* team use to create cleaned maps for cosmological analysis (though they impose several more constraints on the fits.) Interpreting the fit results is complicated, but from the point of view of simply removing the foregrounds, using the internal template is an improvement. We then also use the K–Ka map instead of Haslam in our reassessment of the Bianchi template with the *WMAP* 3-year data.

7.3   Bianchi Type VII$_h$ Models

Note: this section was published as Jaffe et al. (2006b). Redundant text has been removed, the discussion of the Bridges et al. (2006) analysis expanded, and references changed to previous chapters.

Chapters 2 and 3 discuss the unexpected detection of a correlation between the CMB sky in the *WMAP* first-year data release and a Bianchi Type VII$_h$ template. Here, we examine the newly released 3-year data described in Hinshaw et al. (2006) and show that the signal remains constant both in the combined 3-year data and in each individual year, despite differences in the data processing pipelines and calibration and an additional debiasing of the internal linear combination map.

7.3.1   Methods

The methods used are described in detail in Chapter 2. In the 3-year release, an additional "debiasing" is applied to the internal linear combination map (WILC3) that attempts to correct for the fact that the minimum variance method effectively seeks to anti-correlate the CMB and foregrounds, which particularly affects the galactic plane region. (See Hinshaw et al. 2006 for details.)

Following the *WMAP* 3-year analysis, we use the *WMAP* K–Ka map as a tracer of synchrotron instead of the 408 MHz of Haslam et al. (1982), which will account for some variation in the synchrotron spectral index as well as *any* additional foreground component at low frequencies that is not traced by the dust or Hα templates, including the anomalous emission. See § 7.2 for more discussion of this map.

In addition to fits to the maps for the individual bands from each year, we also fit year-to-year difference maps to characterise the effects of noise and the MEM foreground maps to estimate the effect of foreground residuals not traced by the templates.
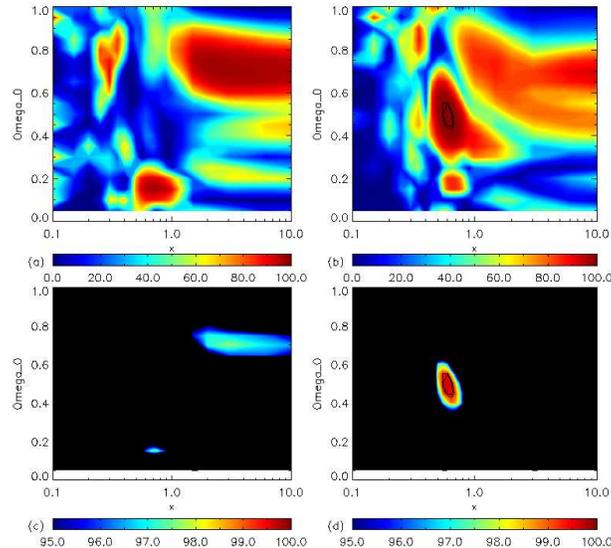
Figure 7.4 Significance contours as percentage of LILC simulations whose best-fit chance alignment amplitude is lower. The left (right) column shows the left-handed (right-handed) models. Over plotted contours are at 99.3 and 99.5%. Two colour scales are used to show the global structure (top) as well as that near the peaks (bottom).

7.3.2   Results

The results of the model-space search using the WILC3 map over the full sky are presented in Fig. 7.4. As in Jaffe et al. (2006a), there is a region of high significance in the parameter space for right-handed models around $(x, \Omega_0) = (0.6, 0.5)$ seen in the right-hand column of Fig. 7.4. In the previous work, there was also a possible detection of a left-handed model at $(x, \Omega_0) = (0.6, 0.15)$, which we examined in detail and concluded was due to foreground contamination in the galactic plane region. The left-hand column of Fig. 7.4 shows that this model is not a significant fit to the 3-year data. This change is likely due to the differences introduced by the debiasing applied to the WILC3 in the 3-year release. This analysis repeats the search for the best-fit orientation of the template relative to the data and finds the same position as the previous best-fit to within one bin of $2.8°$ at this resolution. Using the full-sky WILC3 map, the best-fit model amplitude drops a few percent from 4.33×10^{-10} to 4.19×10^{-10}, which corresponds to a change in significance from 99.8% to 99.5%.

In Fig. 7.5 are difference maps that show the amplitude and morphology of noise, calibration, and map-making effects which might influence these fits. (The monopole and dipole terms are removed.) Fig. 7.5 (a) shows the Ka-band difference between the 3-year and 1-year map, where the time variability in a few point sources (shown here highly smoothed) is the dominant component. Fig. 7.5 (b) shows the map-making and calibration differences for the 1$^{\text{st}}$ year data in Ka-band between the first- and third-year releases, where again, the structure is partly due to the smoothing of differences around point sources

(a) Ka 3yrs−yr1

(b) Ka yr1−oyr1

(c) W 3yrs−yr1

(d) W yr1−oyr1

(e) W yr3−yr2

(f) W yr2−yr1

(g) MEM Q resid ii

(h) Bianchi*0.25

−10.00 μK   10.00 μK

Monopole and dipoles removed
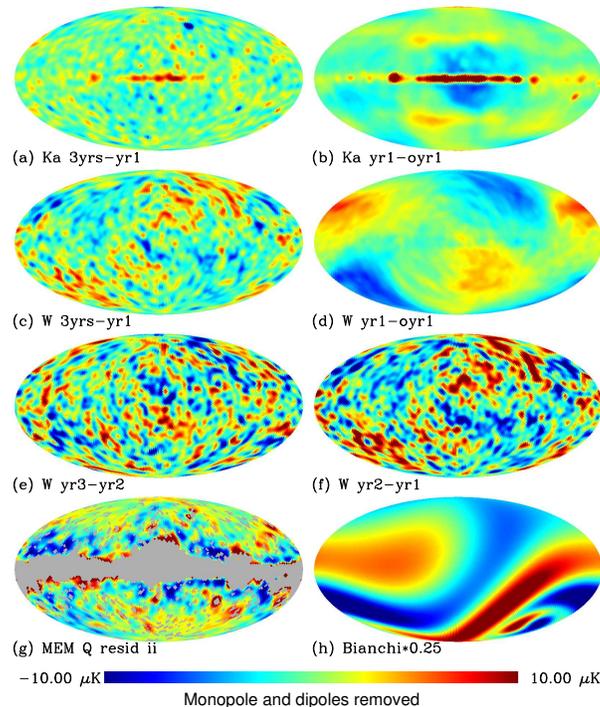
Figure 7.5 Difference maps showing noise and calibration effects. See text.

along the plane as well as the time-varying gain calibration. See Jarosik et al. (2006) for a discussion of these differences. Figs. 7.5 (c) and (d) are the same for the W-band, where the point sources are no longer a significant contaminant, but the gain calibration produces some differences. Figs. 7.5 (e) and (f) show two examples of the noise in the W-band between pairs of individual years. Fig. 7.5 (g) gives the residuals from fitting the foreground templates to the MEM Q-band foregrounds (see below), where those residuals are expected to be largest. Lastly, Fig. 7.5 (h) shows the Bianchi template at one quarter of its best-fit amplitude for comparison. The Bianchi anisotropy amplitude is clearly much higher than any of these effects.

In Fig. 7.6, we plot the fits to year-to-year difference maps, i.e. pure noise, for each differencing assembly. The scatter clearly increases at high frequency, as expected due to the increased noise. W4 in particular shows significantly higher scatter than the other W-band assemblies, which is unsurprising considering that it has the highest level of $1/f$ noise (see Jarosik et al. 2003). We have also performed fits to maps constructed from the MEM foreground solutions to test for the effect of residuals not traced by the templates. We find that the maximum Bianchi amplitude fitting to these maps is less than 3% of the WILC3 fit amplitude, highest for the Q-band fits as expected for foregrounds that increasingly dominate at lower frequencies. The fits are plotted in Fig. 7.7, and the grey bands reflect the uncertainty in the fit results due to the noise and foreground residuals. The error bars are the statistical uncertainties almost entirely due to the CMB.

Table 7.4 shows that the results of the Bianchi fitting to the *WMAP* data are very
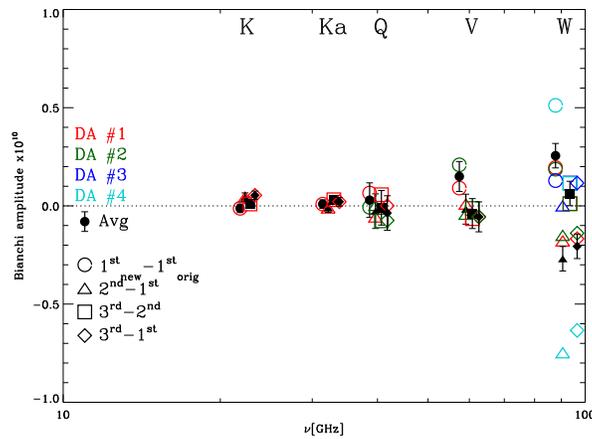
Figure 7.6 Fits to year-to-year difference maps by assembly. The different symbols show different pairs of years, while the colour indicates the differencing assembly. The filled black symbols are the averages over the different assemblies.

stable when the K–Ka map is used as a synchrotron tracer. Although they are slightly lower than the WILC3 fit amplitude, they remain near or above the 99% significance level. Using the combination of templates FDS dust, Finkbeiner Hα, and Haslam for synchrotron as in our original analysis, there remains some low-frequency residual which anti-correlates with the CMB and therefore lowers the fit amplitudes. But the K–Ka map is a better template as it traces whatever low-frequency foreground emission is not traced by the dust and free-free templates, and these results are largely unvarying with frequency. Using the DDD instead of the Finkbeiner Hα template and the SFD instead of the FDS thermal dust template again gives very similar results. There is a small difference at W-band where the SFD template is perhaps not accounting for the thermal dust as well as the FDS template does. (Not shown are fits where only one of the two is changed at a time; changing only the Hα template yields identical results.) All of these fits demonstrate that the result is robust when the best available templates are used.

The difference maps, which have no cosmic signal, give fit amplitudes of only 3% or less of the signal, consistent with the estimate given by the MEM fits. (The significances of these difference fits reflect the fact that there are more residuals in the data than in simulations where the templates are exactly correct.) The 3-year fit values are also plotted in Fig. 7.7 with fits to noise and foreground maps as well for comparison. (The difference between the 3-year fits and the individual years is too small to see on this plot.) The grey band around the fit value obtained with the WILC3 shows that, though they are slightly lower, the cut-sky fits do not vary in excess of what is expected due to such small foreground residuals and noise. Even with these effects, the fits are consistently at the 99% significance level compared with simulations (shown by the dashed line).

Map	Amplitudes $(\sigma/H)_0 \times 10^{10}$												$P(|\alpha_{\mathrm{sim}}| < |\alpha_{\mathrm{obs}}|)$		
	1$^{\mathrm{st}}$-*year*			2$^{\mathrm{nd}}$-*year*			3$^{\mathrm{rd}}$-*year*			3-year			3-year %		
WILC										4.19 ± 0.82			99.50		
	i	*ii*	*iii*	*i*	*ii*	*iii*	*i*	*ii*	*iii*	*i*	*ii*	*iii*	*i*	*ii*	*iii*
K	2.92	4.08	4.03	2.96	4.12	4.08	2.95	4.10	4.08	2.92	4.08	4.04	41.6	98.9	98.8
KA	3.68	4.08	4.03	3.70	4.09	4.05	3.70	4.08	4.04	3.69	4.08	4.04	93.9	99.0	98.8
Q	3.87	4.07	4.03	3.88	4.07	4.04	3.91	4.10	4.07	3.89	4.08	4.05	97.3	99.1	98.9
V	4.09	4.13	4.17	4.06	4.10	4.14	4.03	4.07	4.11	4.06	4.10	4.14	98.9	99.2	99.3
W	4.15	4.15	4.35	4.06	4.05	4.26	4.13	4.13	4.33	4.11	4.11	4.31	99.1	99.1	99.8
QVW	3.95	4.10	4.10	3.94	4.08	4.09	3.95	4.09	4.10	3.95	4.09	4.09	98.0	99.1	99.1
VW	4.10	4.14	4.20	4.06	4.09	4.17	4.05	4.08	4.15	4.07	4.10	4.17	98.9	99.2	99.3
Q-V	-0.06	0.08	0.05	-0.07	0.07	0.04	-0.05	0.10	0.07	-0.06	0.08	0.05	99.1	99.9	97.6
V-W	-0.13	-0.09	-0.16	0.11	0.15	0.08	0.03	0.06	-0.01	0.00	0.04	-0.03	6.5	86.8	71.1
Q-W	-0.20	-0.01	-0.11	0.04	0.22	0.12	-0.02	0.16	0.06	-0.06	0.12	0.02	98.1	100.0	62.0

Table 7.4 Amplitudes (i.e. the shear value $(\sigma/H)_0 \times 10^{10}$) of the best-fit model derived from various combinations of data and various methods as described in the text of Jaffe et al. (2006a). Uncertainties are shown for the ILC only, and they are roughly the same for all, since the noise is negligible for this low-resolution analysis and the structure outside the Galactic cut. For the WILC maps, the total convolver method is used on the full sky. For the individual bands, the Kp0 mask was imposed and foreground templates fit simultaneously for the remaining maps: in columns labeled *(i)*, we used the FDS dust, Finkbeiner H$_\alpha$, and Haslam for synchrotron; in columns labeled *(ii)*, we used *WMAP* K-Ka for synchrotron; in columns labeled *(iii)*, we use alternate templates SFD for dust and DDD for H$_\alpha$, with again *WMAP* K-Ka for synchrotron. The last columns show the significance measure for the 3-year results, i.e. the percentage of simulations with lower amplitude.
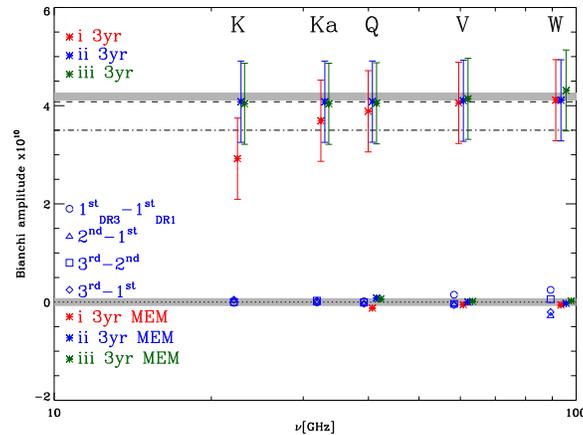
Figure 7.7 *WMAP* 3-year fit amplitudes for the three different sets of templates labelled columns *i*, *ii*, and *iii* in Table 7.4. Also plotted are fit results from difference maps between years and from the MEM foreground maps. These indicate the degree to which noise and foreground residuals may be contaminating the fits, shown as the grey band. This band is also shown around the best-fit WILC3 amplitude for comparison. The dashed and dot-dashed lines represent the approximate amplitudes corresponding to 99% and 95% significance, respectively.

7.3.3　Bayesian Evidence

One of the difficulties interpreting this result has been to judge its significance, in both the qualitative and quantitative senses of the word. Our analysis is a simple frequentist approach, and though there is a surprising detection, we do not claim that the result is a statistically strong detection. Its interest is primarily in the serendipitous result that the sky, when corrected for the template, is statistically isotropic and Gaussian, lacking the anomalies described in § 1.3.

But it is an interesting question whether or not the data justify the extra parameters needed to specify the Bianchi template. Follow-up work on this issue has been done by Bridges et al. (2006), where the authors perform a more sophisticated statistical analysis to attempt to determine whether or not the data require the template. Bayes theorem,

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)}, \tag{7.1}$$

describes the posterior probability of parameter θ given a model or hypothesis, M, and the data, D. $P(D|\theta, M)$ is the likelihood of the data given parameter θ of model, M. $P(\theta|M)$ is the prior probability of the parameter θ given the model, M. The denominator, $P(D|M)$ is the evidence of the model, M. Usually, we are interested in finding, for example, the best value of θ in a known model M, so the denominator is an arbitrary normalisation. But when we are not certain of the model itself, the Bayesian evidence is used to compare different models.

The Bayesian evidence, or marginal likelihood, is defined as

$$E \equiv P(D|M) = \int d\theta P(D|\theta, M) P(\theta|M) \qquad (7.2)$$

and to compare two models, M_0 and M_1, we examine the Bayes factor:

$$B_{10} \equiv \frac{E(M_1)}{E(M_0)}. \qquad (7.3)$$

From the definition of the evidence, it is clear that models where a large region of the parameter space has high likelihood given the data will have a larger value of E, and vice versa. This effectively penalises models with too many unnecessary parameters, naturally implementing Occam's Razor[1]. The Bayes factor B_{10} is usually used to compare a simpler model or "null" hypothesis, M_0, with a more complicated model, M_1. If the latter model is not well constrained by the data, and there is an overly-large parameter space within that model, then its Bayesian evidence will be small and the Bayes factor small. It is the Bayes factor that, if large enough, implies that the more complicated model is justified by the data. Note that such Bayesian methods are sensitive to the choice of priors, $P(\theta|M)$. In the case of the Bridges et al. analysis, the choice was made to investigate only right-handed models, and the parameter space of (Ω, x) must be restricted. The implications of their choices are not clear.

Bridges et al. (2006) compare the Bayesian evidence for a cosmological model with added Bianchi parameters to the standard cosmological model. They confirm our result from Chapter 3 that there is no evidence for a Bianchi template with matter and dark energy densities consistent either with other observations or with the small-scale CMB structure. They then treat the densities needed to create the Bianchi morphology as additional parameters unrelated to the background cosmology. It becomes an arbitrarily parameterised template divorced from its physical origins as an anisotropic cosmology. (This is in effect what we do in Chapter 2, where we use a power spectrum consistent with the concordance model and independently vary Ω_0 for the Bianchi template.)

They find that the data *do* appear to require this template, even with the six additional parameters. The Bayes factor comparing the compound model to the null hypothesis is $\ln B_{10} \sim 2$, or "substantial" on the Jeffreys' scale (Jeffreys, 1961). This is not "strong", much less "decisive", but it is an indication that the detection is not simply a chance alignment or the effect of fitting with extra parameters.

It is also important to consider this result in the broader context of the anomalies. To judge the importance of the template, one would need to quantify the likelihood of the data including all of the $\sim 3\sigma$ anomalies, rather than the standard likelihood, $\mathcal{L} \propto \exp(-\chi^2)$, which assumes a Gaussian, statistically isotropic distribution. Then one could quantify the improvement after correcting for the template not only by the improved χ^2 but also

---

[1]W. Shakespeare: *"There are more things in heaven and earth, Horatio, Than are dreamt of in your philosophy."* W. V. O. Quine: *"Possibly, but my concern is that there not be more things in my philosophy than are in heaven and earth."*

by the fact that the violations of Gaussianity and statistical isotropy have disappeared as well.

We conclude, therefore, that in the absence of an unknown systematic effect which could explain both the anomalies and the correlation, the *WMAP* data require an addition to the standard cosmological model that resembles the Bianchi morphology.

Chapter 8

Conclusions

The preceding chapters have described several morphological analyses of the microwave sky. Template-fitting methods have been demonstrated to apply to both the question of the isotropy of the cosmic component and to the task of foreground separation. A preliminary analysis with harmonic phases also shows promise in addressing non-Gaussianity from whatever source. This chapter will outline outstanding questions and work to be done, as well as prospects from future experiments.

8.1 Outstanding Questions

The work presented here represents only a fraction of the work yet to be done. Typically, we are left not with fewer but simply different questions than we started with.

- **Can a viable anisotropic model be found to explain the anomalies described in § 1.3?** Chapter 2 showed that the morphology of an open universe with vorticity (i.e., a Bianchi type VII_h model) not only matches the data but, when subtracted, leaves a statistically isotropic and Gaussian sky. But in Chapter 3, we find that this model cannot be reconciled with either the larger-scale CMB fluctuations or other observations that require a universe close to flat. We now need an alterative theory to explain an anisotropic signal of a similar morphology.

- **Do the polarisation data display similar violations of statistical isotropy and Gaussianity?** This will be challenging to answer due to the dominance of foregrounds and in any case will probably require Planck (see § 8.2.4) sensitivity. Secondary anisotropies at large angular scales due to reionisation may also make this question difficult to answer. But the data will provide crucial additional information about the isotropy and Gaussianity of the universe.

- **What is the nature of the residual foreground emission at low frequencies, that which does not correlate with any of the templates derived from other wavelengths?** Is the residual structure seen in Fig. 4.4 indeed a "free-free haze" or

a hard synchrotron component (from highly relativistic electrons resulting from dark matter annihilation) or other? Spectral studies of separate regions including these residuals could shed light on this issue.

- **Is the anomalous dust-correlated emission due to spinning dust?** The spectral index of $\beta \sim -2.8$ found in Chapter 5 likely implies an overprediction at 408 MHz if indeed the emission is a hard synchrotron component and has no spectral break. Prospects for answering this question will be discussed in § 8.2.3.

- **Why does the anomalous dust-correlated emission appear to have lower emissivity in regions where the dust is colder?** In the coldest regions studied in Chapter 5, the cross-correlation values are among the lowest, indicating the least amount of anomalous emission. The sample size is small, so it is not clear if this is coincidence or something important about the nature of the anomalous component (see, e.g., Lagache 2003).

- **Why is the electron temperature inferred from the free-free to Hα ratio much lower than expected?** Both the individual Hα regions studied in Chapter 5 as well as the average over the high-latitude sky appear to have less free-free emission than expected (though the variations among the regions can be as much as a factor of two). Is this due to a problem in the emissivity conversion, the Hα templates, or real variations in the electron temperature? This result could have important implications for studies of the physics of the ISM.

- **What limits can phase statistics place on compact universe cell sizes?** The phase analysis methods described in Chapter 6 may prove a useful tool for studying topology. The variation over different realisations of a given topology is fairly large, which makes it difficult to draw conclusions from a single observed sky, and more work is needed to optimise these methods and test their effectiveness.

- **Can a rotationally invariant phase statistic be defined?** An alternative approach to using the rotational information is to define a statistic that is insensitive to the sky orientation. This may be less sensitive than isolating a non-Gaussian signal by using the information in the rotational nature of the phases, but it has the advantage of requiring no *a priori* knowledge about the orientation.

- **Can phase statistics be applied to primordial non-Gaussianities at small angular scales to study inflation?** This may depend on the inflation model as well as on how well secondary anisotropies on small scales such as the SZ effect can be removed.

8.2 Prospects

8.2.1 Alternative Explanations for the Anomalies

Many ideas have been proposed to explain one or more of the anomalies described in § 1.3. Here is a brief summary of some of the more interesting or plausible:

- Rakić et al. (2006) tested whether the Rees-Sciama effect from the Local Supercluster could explain the alignments. The RS effect is essentially the non-linear ISW effect, namely the distortion induced by the local supercluster affecting the CMB at large angular scales. But they find that the effect from a spherically symmetric model does not match the observed axis of evil.

- A related idea is the "Local Pancake", i.e., weak lensing of the CMB dipole by local a structure; see Vale (2005) and Cooray & Seto (2005). A pancake-shaped mass concentration such as the Great Attractor and the Shapley supercluster, moving relative to both the Earth and the CMB would induce a lensing that would distort the motion-induced dipole, mixing it into other low-ℓ modes. The uncertainty here is whether the estimated mass and velocity needed to cause the "axis of evil" is consistent with observations of the local universe. (Note that this weak lensing is another way of looking at the Rees-Sciama effect. Here, the changing potential is not due to the gravitational collapse but relative motion.)

- Inoue & Silk (2006) suggest that suitably placed large-scale voids could account for several anomalies. Again, the RS effect is the physical process, but in this case due to expanding voids. An appropriately placed pair of voids could explain the quadrupole-octopole alignment and planarity, and a third the non-Gaussian cold spot. One could also imagine an asymmetric distribution of such voids to explain the power asymmetry. In all cases, one would want confirmation of the existence of such voids, which would require a deeper full-sky galaxy redshift survey than currently available, or at very least justification for the possibly non-Gaussian size, shape, and number needed to explain the anomalies.

- As mentioned in § 1.4.1, an anisotropic compact topology such as an asymmetric torus could explain the "axis of evil". Another interesting possibility is the Picard horn topology. This is an infinite but multiply connected space of finite volume and hyperbolic geometry that could explain the power deficit at large angular scales. Yet another idea inspired by the low quadrupole is the "ellipsoidal universe" (referring to the ellipsoidal shape of the surface of last scattering in the presence of a different expansion rate along one axis) discussed by Campanelli et al. (2006). Though the low quadrupole itself may not be significant, such ideas will remain interesting until the "axis of evil", which is significant, is explained.

- Moffat (2005) proposes an inhomogeneous universe to explain not only the alignments but also the power asymmetry and acceleration without dark energy. This seems to

simply redefine the problem, where now it is the inhomogeneity that requires an explanation.

- Abramo et al. (2006) shows that the thermal SZ effect from a simple geometrical model of the local supercluster (LSC) could explain the low quadrupole and alignments. Dolag et al. (2005), by contrast, use a constrained hydrodynamic simulation of the local structure and find that the resulting SZ effect cannot explain the low-ℓ anomalies. With only observations in the Rayleigh-Jeans regime, these results cannot easily be verified. But Planck's HFI observations could definitively confirm or refute this possibility.

- Frisch (2005) discusses a possible source of foreground contamination from interstellar dust grains in the solar heliosheath. This idea is particularly interesting considering the coincidence of the axis of evil with the ecliptic. This signal would have a time variability that would help to detect it, but that variability is on the scale of the 22-year solar cycle and would therefore likely require combining *COBE*, *WMAP*, and Planck data.

None of these ideas has yet proved sufficiently convincing, and it's not clear that we will ever have a satisfying explanation. It is possible that these anomalies are simply chance. But many of these possibilities could in theory be verified or falsified by other observations. This work is intended to emphasise the fact that the anomalies should not be dismissed as statistical flukes without further investigation into alternatives, whether systematics, foregrounds, or non-standard models.

8.2.2 Non-standard Models

The difficulty in interpreting the result of the Bianchi analysis in Chapter 2 is the possibility that the detection may be a chance alignment of the template with the microwave background. It is statistically more likely to be a true detection than a false positive, and the Bayesian evidence supports the additional parameters, but not by a strong enough degree to be decisive. Combined with the fact that the Bianchi model is not consistent with the concordance cosmology (Chapter 3), an independent confirmation would be necessary to revive the idea of vorticity. If the template represents a truly anisotropic signal in the background, either due to vorticity (somehow to be reconciled to the measured energy densities) or due to an as-yet unknown process that results in a signal of very similar morphology, there ought to be other observables.

Polarisation data is a promising possibility. It will be difficult, even with Planck's sensitivity, because we still do not yet understand the polarised foregrounds that dominate the cosmic signal; see below. But the fact that most of the Bianchi structure is well away from the Galactic plane will help. The effect of the anisotropic geometry on polarisation has not yet been studied. But it is quite possible that a corresponding polarisation signal would be introduced that might be distinguishable from that associated with the stochastic fluctuations themselves. The secondary anisotropies generated during reionisation may

make this more difficult to detect. In any case, it will be very interesting to see whether or not the polarisation data reflect the anomalies in the temperature fluctuations at large scales. Either answer would help to understand their origin.

Studying compact topology, where there is in general no deterministic pattern, is more challenging. The matched-circles tests could in theory provide a definitive answer, but this is complicated by the late-time ISW effect. This means that though we may be looking at the same region of the surface of last scattering in different directions, the photons propagate to us through differently evolving potentials. This is effectively an additional correlated noise term that makes the analysis that much more challenging. Sensitive, full-sky surveys of the galaxy distribution out to high redshift may allow us to map the ISW effect accurately enough to take this into account. But a further challenge arises from that fact that, even though the topology may be non-trivial, if its scale length is too large, there are no matched circles to detect. In that case, it may still be possible to detect the topology using phase statistics as described in Chapter 6 or with other methods of studying the non-Gaussian correlations such as that of Kunz et al. (2006).

In the absence of any connection to a fundamental physical theory of the origin of the universe itself, the issue of topology could remain more an abstract philosophical question than science. But investigating it ought to remain a priority for cosmologists until and unless another convincing explanation for the anisotropies in the data is found. The statistical isotropy of the sky is never only an abstract philosophical question, as it is a vital assumption in parameter estimation. Our so-called "cosmic concordance" and "precision cosmology" both hang on this assumption, and again, until and unless inflation is proved (or another physical cause of isotropy), it is an assumption we must continue to test.

8.2.3   Galactic Emission

If Stephen Weinberg is right[1], then the study of microwave foregrounds is certainly one of the places where the action is. Despite it being in front of our noses in cosmological terms, there is still a great deal we don't know about the ISM. Among the surprises that came out of the first sensitive observations of the microwave sky was the anomalous dust-correlated emission. Less surprising is the fact that the different components vary in their spectral behaviour across the sky. This was always going to be a problem, since though extensive surveys of the sky have been done in the radio, the microwave regime is not easily accessible from the ground and is thus even more challenging.

Both of these facts make foreground analysis somewhat messy, as was discussed in detail in Chapters 4, 5, and 7. The work in Chapter 5 might be extended to attempt to map out the spectral variations in each component by performing such fits in small regions that tile the full sky. More complicated methods have also been explored such as fitting a set of spectral models to individual pixels using MCMC methods as in Eriksen et al. (2006a). But in addition to more sophisticated analysis tools, we need a better physical understanding of each source of foreground emission.

---

[1] "My advice is to go for the messes – that's where the action is." (Weinberg, 2003)

Anomalous Dust-correlated Emission

The jury is still out on the origin of the anomalous dust-correlated emission. The hypothesis favoured by the *WMAP* team is that this component is hard synchrotron radiation (Bennett et al., 2003c). The highest-energy cosmic ray electrons lose energy more quickly as they diffuse away from where they were accelerated. Therefore, the spatial distribution of such high-energy electrons should remain fairly concentrated around the supernova remnants, which in turn lie in dusty star forming regions. Therefore a hard synchrotron component might well be spatially correlated with the thermal dust emission. To confirm this possibility requires a better knowledge of the spectrum of cosmic-ray electrons at their acceleration site as well as of the processes by which they lose energy as they diffuse away. Observations of the diffuse γ-ray emission in the Galaxy may shed light on the issue, as the cosmic-ray electrons also produce γ-ray emission through inverse Compton scattering.

One obvious criticism of this view is that, considering the strikingly tight spatial correlation of these components (Fig. 7.2), it makes sense to assume that a dust-correlated emission comes from the dust itself. A plausible mechanism of dipole radiation from spinning dust grains has been proposed (Draine & Lazarian, 1998a,b) that allows us to distinguish it from a synchrotron component by its spectral turn-over near 10-15 GHz. This component has been detected via this turn-over in different regions of the sky (de Oliveira-Costa et al., 1999, 2000, 2002; Finkbeiner et al., 2002, 2004; Finkbeiner, 2004a; Watson et al., 2005), though a full sky survey at this frequency is not currently available. One further distinction between these hypotheses is in the different predicted polarisation signals, and as mentioned in § 5.6, the steep polarisation spectrum at K and Ka is an indication that the anomalous emission is not hard synchrotron. (It is interesting to note that the polarised spectrum appears to harden toward V and W, which is consistent with neither hypothesis, though this is perhaps due to confusion with other components.)

Chapter 5 studies the spectral behaviour of this anomalous component in different regions and finds that it overpredicts the dust-correlated emission that should be present in the Haslam map at 408 MHz if there is no spectral flattening between. This work also finds that the emissivity depends on the dust temperature, such that the coldest regions show the lowest inferred emissivity from the anomalous component. This emission is shown to be the dominant foreground at the lower frequency *WMAP* bands and that its emissivity appears to vary roughly by a factor of 2 from region to region. These results may prove important in confirming any hypothesis for the origin of this mysterious foreground.

Polarised Foregrounds

While galactic foregrounds dominate the cosmic temperature signal only near the plane or outside the frequency window around 60 GHz, the polarised emission from the Galaxy dominates the cosmic signal over the full sky and at all frequencies (except for some angular scales). The polarisation components of the Galactic emission are also significantly less well understood than their temperatures, and due to Faraday rotation, cannot easily be extrapolated from observations in the radio.

Synchrotron emission is strongly polarised perpendicular to the magnetic field and can have a polarisation fraction up to $\sim 75\%$. But along a given line of sight, and observed with a given beam, that field is not uniform and the signal is thus reduced significantly by mixing, though it still dominates the cosmic component. A detailed 3-dimensional knowledge of the galactic magnetic field structure is needed in order to accurately model and remove the polarised synchrotron.

At higher frequencies, polarised dust becomes the dominant foreground (the foreground minimum moving from 60 GHz for temperature to ~ 80 GHz for polarisation). This is simply due to the alignment of asymmetric dust grains perpendicular to the magnetic field. The degree of polarisation depends on the details of the grain sizes and composition, how well they align with the field, and how uniform that field is along the line of sight. The polarisation fraction may be anywhere from 1% to 10%.

The magnetic field can also be studied through the Faraday rotation of extragalactic radio sources as well as galactic pulsars and through the polarisation of starlight due to absorption by aligned dust grains.

Future

Part of what we need to progress on these issues will soon be available: more frequency coverage, such as the nine frequency bands that will be provided by Planck. Those nine bands will make it easier to separate components with different spectral dependencies, particularly when combined with other data.

To study the anomalous dust, observations at 10-15 GHz will be needed in order to see if the spectrum turns over as expected in the spinning dust models. But polarisation, which Planck should measure sensitively over the full sky, will also help answer that question, as different models for the anomalous component predict different polarisation signatures. The full-sky survey by the Akari[2] (a.k.a. Astro-F or IRIS - InfraRed Imaging Surveyor) infrared satellite will be particularly useful for studying the spatial distributions of different dust components. This will shed light on whether the distribution of small spinning dust grains is consistent with this explanation for the anomalous foreground.

Another interesting project in the more distant future is the SKA[3] survey of Faraday rotation to map out the geometry of the Milky Way's magnetic fields. This will help our understanding of the Galactic synchrotron emission, though SKA is not expected to be on-line until 2019. In the more immediate future is the GALFACTS[4] continuum survey at ~ 1.4 MHz from Aricebo that will also provide important data for studying the magnetic fields through the synchrotron emission in the radio.

---

[2]http://www.ir.isas.ac.jp/ASTRO-F/
[3]http://www.skatelescope.org/
[4]http://www.ras.ucalgary.ca/GALFACTS/

8.2.4   Future CMB Observations

Polarisation is the next big thing in CMB science. As mentioned above, polarisation data will be important not only for in studying the Galactic foregrounds but may help as well to determine if cosmic component is statistically isotropic and Gaussian on large scales. The detection of B-mode polarisation, which could confirm the existence of primordial gravitational waves, remains one of the primary goals of future missions including the Planck satellite[5] and a variety of ground-based telescopes. The gravity wave signal may dominate only at the larger angular scales ($\ell \lesssim 150$), as smaller angular scales become rapidly dominated by the secondary polarisation signal from weak lensing of the much larger E-modes by intervening matter. To unambiguously detect gravity waves then requires measuring both the large and small scales of the power spectrum in order to distinguish the two components of the B-modes. The reward for such a detection would be one of the first direct probes of the expansion rate *during* inflation and thereby the inflaton field and potential (see Eqtn. A.22).

The Planck satellite (The Planck Collaboration, 2006), scheduled for launch in 2008, will provide as much of an improvement over *WMAP* as that mission did over *COBE* in providing sensitive full-sky measurements of both temperature and linear polarisation over a broad frequency range from 30 to 850 GHz (to 353 GHz for polarisation). The Low Frequency Instrument (LFI) uses HEMTs at 30, 44, and 70 GHz, while the High Frequency Instrument (HFI) uses bolometers at 100, 143, 217, 353, 545, and 857 GHz. The angular resolution ranges from 5 arcmin at the higher frequencies to 33 at the lowest. The two instruments will provide independent measurements that will help study and minimise systematics, while the frequency coverage will be invaluable in studying and removing foregrounds. From an orbit about the second Sun-Earth Lagrange point (L_2), Planck will observe 85° away from an anti-Sun spin axis about which it rotates every minute, thus sweeping out circles of about one degree per day.

Planck will do more than simply refine previous measurements. Measuring the power spectrum up to $\ell \sim 3000$ will allow not only the most precise determination of the spectral index of scalar perturbations, n_s, but also the first possible detection of its running, $dn_s/d\ln k$, and thereby distinguish different models of inflation. The spectral coverage over the SZ cross-over will allow precise measurements of both the thermal and kinetic SZ effects, measuring the clusters' peculiar velocities and hot intracluster gas and providing additional insight into structure formation.

Planck will also provide far more sensitive full-sky polarisation measurements than *WMAP*. Measurement of C_ℓ^{BB} from the largest scales down to $\ell \sim 1000$ may be able to distinguish gravitational waves against the stronger signal from lensing, depending on the tensor-to-scalar ratio. But there are also a variety of ground-based or balloon-born missions designed explicitly to detect B-mode polarisation more accurately on small angular scales, such as Clover[6] (at 97, 150, and 225 GHz) and Quiet[7] (at 40 and 90 GHz) in the Atacama

---

[5]http://planck.esa.int/

[6]http://www-astro.physics.ox.ac.uk/research/expcosmology/groupclover.html

[7]http://quiet.uchicago.edu/index.php

desert, and QUaD[8] (a.k.a. QUEST, at 100 and 150 GHz) at the South Pole, to name only a few. Verde et al. (2006) provide a summary of these future missions and how well they can expect to be able to constrain inflation.

8.3 Last Word

There are some who look at the concordance model in this era of precision cosmology and suppose that we are approaching the "end game", that we are reaching the limit of what science can tell us about the universe (Horgan, 1998). Personally, I suspect that we are no more nearing the end of cosmology than the "end of history" (Fukuyama, 1989). The number of outstanding questions remains long, and those questions no less fundamental than those answered in the 20^{th} century. Among the mysteries currently confronting cosmologists are: dark matter, dark energy, inflation, local geometry, global topology, the fate of the universe, the formation of the first objects, galaxy formation and evolution, the values of the fundamental constants, quantum gravity, etc.

A more valid question than that of whether or not we are nearing the end game is whether or not the current standard model of cosmology is a modern version of epicycles (see, e.g., Narlikar 2005). Three of the most fundamental components of that model, dark matter, dark energy, and inflation, are completely unknown. The terms themselves represent phenomena that have been postulated to explain observations that do not fit in with the rest of what we think we know.

It is here that the CMB becomes so very important. As mentioned in the introduction, the CMB is not a good probe of dark energy, but it is vital to study both inflation and dark matter. One can argue that dark matter was invented to explain the failure of our theory of gravity (in the context of galaxy rotation curves) and likewise inflation our theory of geometry (in the context of spatial flatness and isotropy). But the acoustic peak structure in the CMB power spectrum would be very difficult to explain without both dark matter and inflation. There are other ideas, however, such as the Quasi-Steady-State Cosmology (Narlikar et al., 2003, QSSC), or modified Newtonian dynamics (MOND) as an alternative to dark matter, but these have so far failed to explain the acoustic peak structure (see, e.g., Slosar et al. 2005). Inflation itself has long remained beyond the reach of any experiment to test, but in the near future observations of the primordial power spectrum as well the large-scale B-mode polarisation signal may allow us to begin to reduce the number of possible models.

The CMB is one reason that the situation today is quite a bit different than even twenty years ago. It is only recently that the amount of data has even been comparable to the number of theories. The sensitive dependence of the statistical properties of the CMB on both the physics of the early universe and its evolution provides one of the most stringent tests any theory will have to pass in order to be accepted.

---

[8]`http://www.astro.cf.ac.uk/groups/instrumentation/projects/quad/`

Appendix A

Definitions

The following derivations and relations are taken from Kolb & Turner (1990), Misner et al. (1973), Peacock (1999), and Liddle & Lyth (2000).

A.1 Cosmological Equations

- **The Friedmann-Robertson-Walker (FRW) metric**[1] describes a homogeneous and isotropic space:

$$ds^2 \equiv g_{\mu\nu}dx^\mu dx^\nu = a^2(t)\left\{-d\eta^2 + \frac{dr^2}{1 - kr^2} + r^2 d\theta^2 + r^2 sin^2\theta d\phi^2\right\} \qquad (\text{A.1})$$

  where the sign of k determines that of $\Omega_{\text{tot}} - 1$ and thus the geometry as described in § A.6. The cosmic scale factor is a, here defined such that the current value is one.

  A linear perturbation about this metric is given (in the conformal Newtonian gauge) by

$$ds^2 = a^2(t)\left\{-(1 + 2\Psi)d\eta^2 + (1 - 2\Phi)\frac{dr^2}{1 - kr^2} + r^2 d\theta^2 + r^2 sin^2\theta d\phi^2\right\}, \qquad (\text{A.2})$$

  where Φ is the space-space curvature perturbation and Ψ is the time-time, or Newtonian, potential, $\Psi = dt/t \sim -\Phi$.

- **Einstein equation:**

$$R_{\mu\nu} - \frac{1}{2}\mathcal{R}g_{\mu\nu} \equiv G_{\mu\nu} = 8\pi G T_{\mu\nu} + \Lambda g_{\mu\nu} \qquad (\text{A.3})$$

  where $R_{\mu\nu}$ is the Ricci tensor, $\mathcal{R}$ the Ricci scalar, $g_{\mu\nu}$ the metric, $G_{\mu\nu}$ the Einstein tensor, $T_{\mu\nu}$ the stress-energy tensor, and Λ the cosmological constant, or vacuum

---

[1]In spherical polar coordinates and units where $\hbar = c = 1$. Conformal time is sometimes referred to as τ, but we shall reserve that for optical depth and instead use η.

energy density. The stress-energy tensor for a perfect fluid is $T_{\mu\nu} = \mathrm{diag}(\rho, \mathrm{p}, \mathrm{p}, \mathrm{p})$, where ρ is the density and p the pressure, and the equation of state is given by $w \equiv p/\rho$.

- **The Friedman equation** comes from the 0-0 component of the Einstein equation with the FRW metric:

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{\rho}{3M_{\mathrm{Pl}}^2} - \frac{k}{a^2} + \frac{\Lambda}{3} \tag{A.4}$$

where $M_{Pl} \equiv (8\pi G)^{-1/2}$. Alternatively, the curvature and vacuum energy are expressed as density components:

$$H^2(a) \equiv H_0^2 \mathcal{H}^2(a) = H_0^2 \left(\Omega_{\mathrm{m0}} a^{-3} + K_0 a^{-2} + \Omega_{\Lambda 0}\right) \tag{A.5}$$

where we have used $\Omega \equiv \rho/\rho_{\mathrm{c}}$ and $\rho_{\mathrm{c}} \equiv 3H/8\pi G = 3H M_{\mathrm{Pl}}^2$ is the critical density[2]. (We use K instead of Ω_{k}.) The subscript zero refers to the current value.

- Out of the Einstein equation and the conservation of the stress-energy tensor also comes the first law of thermodynamics, or **the continuity or fluid equation:**

$$\dot{\rho} = -3\frac{\dot{a}}{a}(\rho + p). \tag{A.6}$$

(This can also be seen as the energy conservation law for adiabatic expansion, $dE = -pdV$, and where $E = V\rho$.)

- Putting these together, we have **the acceleration equation:**

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}, \tag{A.7}$$

With these equations, we can solve for the evolution of the universe in each phase dominated by one component with equation of state $w \equiv p/\rho$.

- Density, ρ:

$$\rho \propto a^{-3(1+w)} = \begin{cases} a^{-4} & \text{radiation-dominated} \\ a^{-3} & \text{matter-dominated} \\ \text{const.} & \text{vacuum-dominated} \end{cases} \tag{A.8}$$

- Scale factor, a:

$$a \propto t^{2/3(1+w)} = \begin{cases} t^{1/2} & \text{radiation-dominated} \\ t^{2/3} & \text{matter-dominated} \\ \exp(Ht) & \text{lambda-dominated}(w = -1) \\ t & \text{curvature-dominated}(w = -1/3) \end{cases} \tag{A.9}$$

---

[2]We use the *reduced* Planck mass, M_{Pl}, differing by a factor of $\sqrt{8\pi}$ from the usual Planck mass, m_{Pl}

- Temperature, T:

$$T \propto \begin{cases} a^{-1} & \text{decoupled-massless} \\ a^{-2} & \text{decoupled-massive} \end{cases} \tag{A.10}$$

- Time, t, or age of the universe, t_0:

$$t \equiv \int_0^{a(t)} \frac{da'}{\dot{a}'}, \tag{A.11}$$

which gives the age of a *flat* universe of

$$t_0 = \frac{2}{3} H_0^{-1} \Omega_\Lambda^{-1/2} \ln \left[\frac{1 + \Omega_\Lambda^{1/2}}{(1 - \Omega_\Lambda)^{1/2}} \right]. \tag{A.12}$$

To study the evolution of fluid perturbations, we also need:

- **the Euler equations** for a perfect fluid are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \qquad \text{a.k.a. the continuity equation} \tag{A.13}$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} = -\rho \nabla \phi - \nabla \mathbf{p}, \tag{A.14}$$

$$\nabla^2 \phi = 4\pi G \rho, \qquad \text{a.k.a. the Poisson equation.} \tag{A.15}$$

A.2   Distances

- The **particle horizon**, R_h, is the comoving distance a particle may have travelled, or the maximum distance between any two causally connected events. This is equivalent to the **conformal time**, η,

$$R_\mathrm{h} = \eta = \int_0^{t_0} \frac{dt}{a} = \int_0^1 \frac{da}{a^2 H} = \int_0^\infty \frac{dz}{H(z)}, \tag{A.16}$$

(up to the constant speed of light, c, taken to be equal to one throughout). In a flat, matter dominated universe, this distance is roughly 13 Gpc. The size of the horizon at last scattering, when projected onto the sky at the redshift of the CMB, corresponds to about $1°$.

- The **event horizon** is the maximum comoving distance to any event in the future that may be causally connected to the present, t_0:

$$R_\mathrm{e} = \int_{t_0}^\infty \frac{dt}{a}. \tag{A.17}$$

Whether or not this value is finite depends on the cosmological model. In the big rip scenario, the event horizon will be finite, as eventually, the rate of expansion will overtake the increase in the particle horizon. In a flat model without dark energy, the expansion slows asymptotically to zero, and the event horizon will be infinite.

- The **Hubble radius** is the comoving length scale of the universe and is defined by

$$R_{\rm H} = \frac{1}{aH}, \tag{A.18}$$

now $R_{\rm H} \approx 4.3$ Gpc. Note that it is very different from the horizon. In effect, it gives the distance beyond which the recession velocity exceeds the speed of light, and as the expansion decelerates, this radius increases. Inflation refers to the idea of a period in the early universe where this radius *decreased*, which is also the case at late times in vacuum-dominated, accelerated expansion (see § A.3). The Hubble radius does not mean that we can never observe beyond this distance, of course, since a photon emitted at that source will eventually cross into our Hubble volume (assuming that the Hubble radius is not decreasing.)

See Davis & Lineweaver (2004) for a discussion of common misconceptions or misstatements regarding horizons, the Hubble radius, etc.

A.3   Slow-roll Inflation

Inflation, by definition, is an exponential expansion, $\ddot{a} > 0$. Using the Friedmann equation (A.4) and the acceleration equation (A.7), this can also be stated as $d(H^{-1}/a)/dt < 0$ or $p < -\rho/3$. The former statement means that the Hubble length measured in comoving coordinates decreases, i.e., the characteristic length scale of the universe is actually getting smaller relative to the expansion. The latter describes the possible equations of state, including that of the cosmological constant, where $p = -\rho$.

If you define the inflaton field, ϕ, moving in a potential, $V(\phi)$, such that

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V(\phi) \tag{A.19}$$

and

$$p_\phi = \frac{1}{2}\dot{\phi}^2 - V(\phi), \tag{A.20}$$

then using the continuity equation (A.6), then you obtain the scalar wave equation:

$$\ddot{\phi} + 3H\dot{\phi} = -\frac{dV}{d\phi}, \tag{A.21}$$

also that of a ball rolling down a hill with friction. Substituting into the Friedman equation (A.4), we get

$$H^2 = \frac{1}{3M_{\rm Pl}^2}\left[\frac{1}{2}\dot{\phi}^2 + V(\phi)\right]. \tag{A.22}$$

Using the definition of inflation, $\ddot{a} > 0$, and substituting the pressure and density above into the acceleration equation (A.7), we get $\dot{\phi}^2 < V(\phi)$. In other words, inflation only

happens when the potential energy dominates, and one way to ensure this is to have a flat potential down which the inflaton field rolls slowly.

The number of e-foldings from the start t_s to the end t_e of inflation is given by

$$N \equiv \ln \frac{a(t_e)}{a(t_s)} = \int_{t_s}^{t_e} H dt = \frac{1}{M_{Pl}^2} \int_{\phi_e}^{\phi_s} \frac{V}{V'} d\phi \qquad (A.23)$$

Note: Guth's original inflation started with the universe in a false vacuum state, i.e., a local minimum of the potential. But this model suffers from the "graceful exit problem", where the reheating mechanism—nucleation of bubbles of true vacuum—fails to produce the particles needed for a hot big bang. This is now referred to as "old inflation", and the slow-roll scenario as "new inflation".

A.4   Adiabatic versus Isocurvature Fluctuations

An "adiabatic", "curvature", or "isoentropy" perturbation is a perturbation in the total energy density, $\delta\rho \equiv (\rho - \langle\rho\rangle)/\langle\rho\rangle \neq 0$, and thus a perturbation in the curvature, where the entropy, s, is unperturbed. From this, it follows that:

- the *specific entropy*, or matter-to-radiation ratio, remains constant for all matter species, i, such that

$$\delta \left(\frac{n_\gamma}{n_i} \right) \propto \delta\rho_i - \frac{3}{4}\delta\rho_\gamma = 0; \qquad (A.24)$$

- therefore, all components are perturbed equally in *number* density;

- from $\rho_\gamma \propto T^4$, then the temperature *within* the perturbation is

$$\frac{\Delta T}{T} = \frac{1}{4}\delta\rho_\gamma = \frac{1}{3}\delta\rho_b \qquad (A.25)$$

(not to be confused with the effective temperature observed from outside the perturbation; the Sachs-Wolfe effect is described in § 1.1.3).

By contrast, in "isocurvature" perturbations, the total energy density and thus the curvature are not perturbed, $\delta\rho = \delta\rho_\gamma + \delta\rho_m = 0$. Instead the fluctuation corresponds to a change in the equation of state, since it changes not only the density but the composition.

A.5   Harrison-Zel'dovich-Peebles Spectrum

Expressing the density perturbation, $\delta\rho$, in Fourier space, we have

$$\delta\rho = (2\pi)^{-3} \int \delta_k \exp(-i\mathbf{k} \cdot \mathbf{x}) \mathbf{d^3k}. \qquad (A.26)$$

The density power spectrum is

$$P(k) \equiv \langle |\delta_k|^2 \rangle. \tag{A.27}$$

Assuming that there is no preferred length scale implies that the power spectrum follows a power law:

$$\langle |\delta_k|^2 \rangle \propto k^n \tag{A.28}$$

To see how a density perturbation evolves when it is still outside the horizon, define the total density in a perturbation about the mean density, ρ_0, as $\rho \equiv (\rho_0 + \rho_1)$, and apply the Friedmann equation (A.4) for both the unperturbed universe, $H^2 + k_0/a^2 \propto \rho_0$, and the perturbed, $H^2 + k_1/a^2 \propto (\rho_0 + \rho_1)$. Combining, we find that a relative perturbation evolves as $\rho_1/\rho_0 \propto (\rho_0 a^2)^{-1}$. Substituting the evolution of ρ_0 from equation (A.8), we have

$$\delta \equiv \frac{\rho_1}{\rho_0} \propto \begin{cases} a^2 & t < t_{\text{eq}} \\ a & t > t_{\text{eq}} \end{cases} \tag{A.29}$$

Using the definition of conformal time, η, from equation (A.16), and the expansion history from equation (A.9), we then have

$$\delta \propto \eta^2 \tag{A.30}$$

The argument for a scale-invariant power spectrum begins with the assumption that the universe is self-similar. Choosing a given length scale, x, and averaging over all modes in a box of that size effectively filters out all modes smaller than x, or $k > 1/x$, leaving

$$\langle \delta^2 \rangle \propto \int_0^{1/x} \langle |\delta_k|^2 \rangle d(k^3) = \int_0^{1/x} k^n k^2 dk = x^{-(n+3)}, \tag{A.31}$$

so that the total perturbation in the box is, on average,

$$\delta \propto x^{-(n+3)/2}. \tag{A.32}$$

But what is a sensible choice for a length scale? The only natural length scale in an expanding universe is the size of the horizon, $R_{\text{h}} = \eta$. So if at a given time, we have $\delta \propto R_{\text{h}}^{-(n+3)/2}$, and the evolution of a given perturbation goes as $\eta^2 = R_{\text{h}}^2$, then we have

$$\delta(R_{\text{h}}) \propto R_{\text{h}}^{-(n-1)/2} \tag{A.33}$$

which is constant for $n = 1$. This means that the grown of perturbations balances the scale-dependence. Thus the idea of a fractal universe, one that is self-similar at all scales, leads to a scale-invariant power spectrum of perturbations, or $P(k) \propto k^n$ where $n = 1$.

This argument has variously been made by Harrison, Zel'dovich, and Peebles and thus may be referred to by any combination of these names.

A.6 Geometry and Topology

There is often a confusion of terms regarding the geometry, topology, and dynamics of the Universe. Here, we will define these terms in the context of cosmology. For more details, see, e.g., Misner et al. (1973).

- Local Geometry

  Geometry, as commonly discussed, is essentially a local quantity describing the intrinsic curvature of spacetime, sometimes called the 3-geometry. In the FRW metric (the simplest solution to the Einstein field equations), the geometry or curvature is determined by the sign of $\Omega_k \equiv \Omega_{\rm tot} - 1$.

  - $\boldsymbol{\Omega_k = 0}$: A Euclidean or flat geometry is that with which we are most familiar, where parallel lines remain parallel at all points, and the angles of a triangle add up to 180°. This is often called "critical".

  - $\boldsymbol{\Omega_k > 0}$: A spherical, or positively curved, geometry is where parallel lines will converge in both directions and the angles of a triangle will sum to more than 180°, as in the case of the familiar two-sphere. This is is often called "closed".

  - $\boldsymbol{\Omega_k < 0}$: A hyperbolic, or negatively curved, geometry means that parallel lines will diverge in both directions and the angles of a triangle will make less than 180°. The two-dimensional analogy to the "pseudosphere" is of a sort of saddle, and this geometry is usually called "open".

  Note that the two-dimensional analogies are only to help visualise the geometry in a familiar three dimensions. One can discuss them as embedded in a higher dimensional space, but note that this has nothing to do with space-time or the topology of the universe.

- Global Geometry, or Topology

  Topology refers to the shape and extent of the manifold itself, i.e., whether it is simply or multiply connected, infinite or compact, etc. The simplest to imagine are the three possible simply connected cases. These can be infinite (an infinite plane in the flat case or an infinite pseudosphere in the hyperbolic, negatively curved case) or compact (a three-sphere in the positively curved case).

  One of the infinite number of possible examples of multi-connected spaces is the three-torus. Instead of an infinite plane, one can imagine a finite square tiling the infinite plane such that a path that "leaves" the square on one side re-enters it from the opposite direction. Alternatively, one can envision identifying the opposite edges of the square, rolling it up both ways to form a donut-like torus. The shortcoming of the latter analogy is that the local geometry of the three-torus can be flat or even hyperbolic, which the two-dimensional analogy cannot represent. (This is related to the difference between what Misner et al. 1973, p. 336 refers to as intrinsic

versus extrinsic curvature. In this work, the term curvature refers only to intrinsic curvature: "'intrinsic' in the sense that it is defined by, and depends exclusively on, measurements of distance made within the hypersurface.") The torus is particularly interesting in that it is the only compact, orientable flat space that preserves homogeneity (see, e.g., Levin et al. 1998).

So far, there is no physical theory that *predicts* the topology of the universe. General relativity deals only with its geometry.

Note that some confusion arises from the term "closed" having a different meaning for geometry as for topology. In the latter context, spatially or topologically closed could refer to finite or compact topology. In this work, only the term compact will be used for such to avoid confusion with the idea of positive curvature. A universe may have a compact topology and yet be negatively curved.

- Dynamics

  There is some additional confusion regarding the relation between the curvature and the eventual fate of the universe. One often sees the word "closed" refer either to the positively curved geometry or to the idea that the universe will recollapse in a "big crunch", and likewise "open" for negative curvature leading to a "big rip". It is also said that the curvature governs the expansion evolution. This confusion may have arisen when dark energy was not considered viable, and in that case, the correspondence between curvature and the expansion evolution was correct.

  As explained by Misner et al. (1973, p. 747), it is not so simple when dark energy is involved and instead also depends on the value of Λ. A large positive value of Λ may imply that the total energy density is larger than critical, and therefore the curvature is positive, but the dark energy will then continue to drive the expansion forever. Likewise, a negative value could imply that, despite a hyperbolic curvature due to the small total energy density, the universe would recollapse. This can be seen in Fig. 1.4 in the regions marked "expands forever" versus "recollapses eventually".

  (Note that the "No Big Bang" region of that figure represents the situation where, in the expansion evolution, the curvature radius reaches a minimum in the past rather than becoming arbitrarily small. At that minimum, the Hubble parameter goes to zero. It was to allow a static univers that the cosmological constant was inserted into the Friedmann equation. But such a static universe is unstable: any perturbation to the density will drive H away from zero.)

  For this work, "open", "closed", and "critical" or "flat" will refer to the intrinsic geometry we can measure, not to the eventual fate of the universe.

Note that the topology and geometry are generally independent, but the former can limit the possibilities for the latter. A simply connected topology can only be compact if it is positively curved. But a multiply connected topology can be compact and yet be flat or open. (Imagine rolling up a piece of paper with a triangle drawn on it. This illustrates how

the flat intrinsic geometry on the piece of paper is not changed by the cylindrical topology, unlike, for example, a sphere, which cannot be made by such an operation. A three-torus is a similar situation in higher dimension, though the two-dimensional analogue we envision embedded in three dimensions cannot represent it.) In the case of hyperbolic geometry, the curvature scale is also related to the topology scale.

The CMB is one of the most powerful tools for studying the geometry and topology, since it comes from the region of the universe that is the most distant we can currently directly observe.[3] Note that it is not in fact much good in studying dark energy models, because most of the information it provides concerns only one point in time, giving only one measurement of the angular diameter distance to a single redshift. As discussed above, Λ changes the dynamics such that the expansion history is not governed purely by the total energy density or geometry that we can determine using the CMB. This is why supernovae are so useful as a complementary measurement, since we can then measure the luminosity-redshift relation over a large range of redshifts.

---

[3]Using electromagnetic radiation, we cannot see beyond recombination, as the universe was then opaque. But primordial gravity waves or neutrinos, for example, if detected, would go back to even further.

References

Abramo, L. R., Sodre, L. J., & Alexandre Wuensche, C. 2006, (astro-ph/0605269)

Adams, W. S. 1941, ApJ, 93, 11

Banday, A. J., Dickinson, C., Davies, R. D., Davis, R. J., & Górski, K. M. 2003, MNRAS, 345, 897

Banday, A. J., Górski, K. M., Bennett, C. L., Hinshaw, G., Kogut, A., & Smoot, G. F. 1996, ApJL, 468, L85+

Barrow, J. 1976, MNRAS, 175, 359

Barrow, J. D. 1984, MNRAS, 211, 221

Barrow, J. D., Juszkiewicz, R., & Sonoda, D. H. 1985, MNRAS, 213, 917

Barrow, J. D., Kodama, H., & Ahluwalia, D. V. 2001, International Journal of Modern Physics D, 10, 785

Bennett, C. L., Banday, A. J., Górski, K. M., Hinshaw, G., Jackson, P., Keegstra, P., Kogut, A., Smoot, G. F., Wilkinson, D. T., & Wright, E. L. 1996, ApJL, 464, L1+

Bennett, C. L., Bay, M., Halpern, M., Hinshaw, G., Jackson, C., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wilkinson, D. T., Wollack, E., & Wright, E. L. 2003a, ApJ, 583, 1

Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., Wright, E. L., Barnes, C., Greason, M. R., Hill, R. S., Komatsu, E., Nolta, M. R., Odegard, N., Peiris, H. V., Verde, L., & Weiland, J. L. 2003b, ApJS, 148, 1

Bennett, C. L., Hill, R. S., Hinshaw, G., Nolta, M. R., Odegard, N., Page, L., Spergel, D. N., Weiland, J. L., Wright, E. L., Halpern, M., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., & Wollack, E. 2003c, ApJS, 148, 97

Benoît, A., Ade, P., Amblard, A., Ansari, R., Aubourg, É., Bargot, S., Bartlett, J. G., Bernard, J.-P., Bhatia, R. S., Blanchard, A., Bock, J. J., Boscaleri, A., Bouchet, F. R.,

Bourrachot, A., Camus, P., Couchot, F., de Bernardis, P., Delabrouille, J., Désert, F.-X., Doré, O., Douspis, M., Dumoulin, L., Dupac, X., Filliatre, P., Fosalba, P., Ganga, K., Gannaway, F., Gautier, B., Giard, M., Giraud-Héraud, Y., Gispert, R., Guglielmi, L., Hamilton, J.-C., Hanany, S., Henrot-Versillé, S., Kaplan, J., Lagache, G., Lamarre, J.-M., Lange, A. E., Macías-Pérez, J. F., Madet, K., Maffei, B., Magneville, C., Marrone, D. P., Masi, S., Mayet, F., Murphy, A., Naraghi, F., Nati, F., Patanchon, G., Perrin, G., Piat, M., Ponthieu, N., Prunet, S., Puget, J.-L., Renault, C., Rosset, C., Santos, D., Starobinsky, A., Strukov, I., Sudiwala, R. V., Teyssier, R., Tristram, M., Tucker, C., Vanel, J.-C., Vibert, D., Wakui, E., & Yvon, D. 2003, AAP, 399, L19

Bielewicz, P., Górski, K. M., & Banday, A. J. 2004, MNRAS, 355, 1283

Bond, J. R. & Efstathiou, G. 1987, MNRAS, 226, 655

Bondi, H. & Gold, T. 1948, MNRAS, 108, 252

Bridges, M., McEwen, J. D., Lasenby, A. N., & Hobson, M. P. 2006, (astro-ph/0605325)

Bunn, E. F., Ferreira, P. G., & Silk, J. 1996, PhRvL, 77, 2883

Burles, S., Nollett, K. M., & Turner, M. S. 2001, ApJL, 552, L1

Campanelli, L., Cea, P., & Tedesco, L. 2006, (astro-ph/0606266)

Cayón, L., Banday, A. J., Jaffe, T., Eriksen, H. K., Hansen, F. K., Gorski, K. M., & Jin, J. 2006, MNRAS, 369, 598

Cayón, L., Jin, J., & Treaster, A. 2005, MNRAS, 362, 826

Chiang, L.-Y., Naselsky, P. D., & Coles, P. 2006, (astro-ph/0603662)

Chiang, L.-Y., Naselsky, P. D., Verkhodanov, O. V., & Way, M. J. 2003, ApJL, 590, L65

Coles, P., Dineen, P., Earl, J., & Wright, D. 2004, MNRAS, 350, 989

Coley, A. & Hervik, S. 2005, CQGrav, 22, 579

Collins, C. B. & Hawking, S. W. 1973, MNRAS, 162, 307

Cooray, A. & Seto, N. 2005, JCAP, 12, 4

Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2006a, MNRAS, 367, 79

—. 2006b, (astro-ph/0605135)

Cornish, N. J., Spergel, D. N., Starkman, G. D., & Komatsu, E. 2004, PhRvL, 92, 201302

Cruz, M., Martínez-González, E., Vielva, P., & Cayón, L. 2005, MNRAS, 356, 29

Davies, R. D., Dickinson, C., Banday, A. J., Jaffe, T. R., Górski, K. M., & Davis, R. J. 2006, MNRAS, 758

Davies, R. D., Watson, R. A., & Gutierrez, C. M. 1996, MNRAS, 278, 925

Davis, T. M. & Lineweaver, C. H. 2004, PASA, 21, 97

de Oliveira-Costa, A., Tegmark, M., Davies, R. D., Gutiérrez, C. M., Lasenby, A. N., Rebolo, R., & Watson, R. A. 2004a, ApJL, 606, L89

de Oliveira-Costa, A., Tegmark, M., Devlin, M. J., Haffner, L. M., Herbig, T., Miller, A. D., Page, L. A., Reynolds, R. J., & Tufte, S. L. 2000, ApJL, 542, L5

de Oliveira-Costa, A., Tegmark, M., Finkbeiner, D. P., Davies, R. D., Gutierrez, C. M., Haffner, L. M., Jones, A. W., Lasenby, A. N., Rebolo, R., Reynolds, R. J., Tufte, S. L., & Watson, R. A. 2002, ApJ, 567, 363

de Oliveira-Costa, A., Tegmark, M., Gutierrez, C. M., Jones, A. W., Davies, R. D., Lasenby, A. N., Rebolo, R., & Watson, R. A. 1999, ApJL, 527, L9

de Oliveira-Costa, A., Tegmark, M., Zaldarriaga, M., & Hamilton, A. 2004b, PRD, 69, 063516

Dennison, B., Simonetti, J. H., & Topasna, G. A. 1998, PASA, 15, 147

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, ApJ, 142, 414

Dickinson, C., Davies, R. D., & Davis, R. J. 2003, MNRAS, 341, 369

Dicus, D. A. & Repko, W. W. 2004, PRD, 70, 083527

Dineen, P., Rocha, G., & Coles, P. 2005, MNRAS, 358, 1285

Dodelson, S. 2003, Modern Cosmology (Academic Press)

Dolag, K., Hansen, F. K., Roncarelli, M., & Moscardini, L. 2005, MNRAS, 363, 29

Draine, B. T. & Lazarian, A. 1998a, ApJL, 494, L19+

—. 1998b, ApJ, 508, 157

—. 1999, ApJ, 512, 740

Efstathiou, G. 2004, MNRAS, 349, 603

Efstathiou, G. & Bond, J. R. 1999, MNRAS, 304, 75

Eriksen, H., Banday, A. J., Górski, K., & Lilje, P. B. 2005, [astro-ph/0508196]

Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004a, ApJ, 612, 633

—. 2005, ApJ, 622, 58

Eriksen, H. K., Dickinson, C., Lawrence, C. R., Baccigalupi, C., Banday, A. J., Górski, K. M., Hansen, F. K., Lilje, P. B., Pierpaoli, E., Seiffert, M. D., Smith, K. M., & Vanderlinde, K. 2006a, ApJ, 641, 665

Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004b, ApJ, 605, 14

Eriksen, H. K., Huey, G., Saha, R., Hansen, F. K., Dick, J., Banday, A. J., Górski, K. M., Jain, P., Jewell, J. B., Knox, L., Larson, D. L., O'Dwyer, I. J., Souradeep, T., & Wandelt, B. D. 2006b, (astro-ph/0606088)

Eriksen, H. K., Novikov, D. I., Lilje, P. B., Banday, A. J., & Górski, K. M. 2004c, ApJ, 612, 64

Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., Wandelt, B. D., Larson, D. L., Górski, K. M., Levin, S., Banday, A. J., & Lilje, P. B. 2004d, ApJS, 155, 227

Fernández-Cerezo, S. e. a. 2006, MNRAS, submitted (astro-ph/0601203)

Finkbeiner, D. P. 2003, ApJS, 146, 407

—. 2004a, ApJ, 614, 186

—. 2004b, (astro-ph/0409027)

Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, ApJ, 524, 867

Finkbeiner, D. P., Langston, G. I., & Minter, A. H. 2004, ApJ, 617, 350

Finkbeiner, D. P., Schlegel, D. J., Frank, C., & Heiles, C. 2002, ApJ, 566, 898

Fixsen, D. J., Cheng, E. S., Gales, J. M., Mather, J. C., Shafer, R. A., & Wright, E. L. 1996, ApJ, 473, 576

Freedman, W. L., Madore, B. F., Gibson, B. K., Ferrarese, L., Kelson, D. D., Sakai, S., Mould, J. R., Kennicutt, R. C., Ford, H. C., Graham, J. A., Huchra, J. P., Hughes, S. M. G., Illingworth, G. D., Macri, L. M., & Stetson, P. B. 2001, ApJ, 553, 47

Frisch, P. C. 2005, ApJL, 632, L143

Fukuyama, F. 1989, The National Interest, Summer

Gamow, G. 1946, PhRv, 70, 572

Gaustad, J. E., McCullough, P. R., Rosing, W., & Van Buren, D. 2001, PASP, 113, 1326

Górski, K. M., Banday, A. J., Bennett, C. L., Hinshaw, G., Kogut, A., Smoot, G. F., & Wright, E. L. 1996, ApJL, 464, L11+

Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759

Guth, A. H. 1981, PRD, 23, 347

Haffner, L. M., Reynolds, R. J., Tufte, S. L., Madsen, G. J., Jaehnig, K. P., & Percival, J. W. 2003, ApJS, 149, 405

Hajian, A., Souradeep, T., & Cornish, N. 2005, ApJL, 618, L63

Hansen, F., Banday, A., Eriksen, H., Górski, K., & Lilje, P. 2006, (astro-ph/0603308)

Hansen, F. K., Balbi, A., Banday, A. J., & Górski, K. M. 2004a, MNRAS, 354, 905

Hansen, F. K., Banday, A. J., & Górski, K. M. 2004b, MNRAS, 354, 641

Hansen, F. K., Cabella, P., Marinucci, D., & Vittorio, N. 2004c, ApJL, 607, L67

Haslam, C. G. T., Klein, U., Salter, C. J., Stoffel, H., Wilson, W. E., Cleary, M. N., Cooke, D. J., & Thomasson, P. 1981, AAP, 100, 209

Haslam, C. G. T., Stoffel, H., Salter, C. J., & Wilson, W. E. 1982, A&AS, 47, 1

Hervik, S., van den Hoogen, R., & Coley, A. 2005a, CQGrav, 22, 607

Hervik, S., van den Hoogen, R. Lim, W. C., & Coley, A. 2005b, CQGrav, in press

Hinshaw, G., Nolta, M. R., Bennett, C. L., Bean, R., Dore', O., Greason, M. R., Halpern, M., Hill, R. S., Jarosik, N., Kogut, A., Komatsu, E., Limon, M., Odegard, N., Meyer, S. S., Page, L., Peiris, H. V., Spergel, D. N., Tucker, G. S., Verde, L., Weiland, J. L., Wollack, E., & Wright, E. L. 2006, (astro-ph/0603451)

Hinshaw, G., Spergel, D. N., Verde, L., Hill, R. S., Meyer, S. S., Barnes, C., Bennett, C. L., Halpern, M., Jarosik, N., Kogut, A., Komatsu, E., Limon, M., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003a, ApJS, 148, 135

—. 2003b, ApJS, 148, 135

Horgan, J. 1998, The End of Science (Abacus)

Hu, W. & Dodelson, S. 2002, ARA&A, 40, 171

Hu, W., Fukugita, M., Zaldarriaga, M., & Tegmark, M. 2001, ApJ, 549, 669

Hu, W., Sugiyama, N., & Silk, J. 1997, Nature, 386, 37

Hu, W. & White, M. 1997, NewA, 2, 323

Inoue, K. T. & Silk, J. 2006, (astro-ph/0602478)

Jaffe, T. R., Banday, A. J., Eriksen, H. K., Górski, K. M., & Hansen, F. K. 2005, ApJL, 629, L1

—. 2006a, ApJ, 643, 616

—. 2006b, (astro-ph/0606046)

Jaffe, T. R., Hervik, S., Banday, A. J., & Górski, K. M. 2006b, ApJ, 644, 701

Jarosik, N., Barnes, C., Bennett, C. L., Halpern, M., Hinshaw, G., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003, ApJS, 148, 29

Jarosik, N., Barnes, C., Greason, M. R., Hill, R. S., Nolta, M. R., Odegard, N., Weiland, J. L., Bean, R., Bennett, C. L., Dore', O., Halpern, M., Hinshaw, G., Kogut, A., Komatsu, E., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., & Wright, E. L. 2006, (astro-ph/0603452)

Jeffreys, H. 1961, Theory of Probability (Oxford University Press)

Jewell, J., Levin, S., & Anderson, C. H. 2004, ApJ, 609, 1

Jonas, J. L., Baart, E. E., & Nicolson, G. D. 1998, MNRAS, 297, 977

Jones, A. W., Davis, R. J., Wilkinson, A., Giardino, G., Melhuish, S. J., Asareh, H., Davies, R. D., & Lasenby, A. N. 2001, MNRAS, 327, 545

Juszkiewicz, R., Bajtlik, S., & Górski, K. 1983, MNRAS, 204, 63P

Knop, R. A., Aldering, G., Amanullah, R., Astier, P., Blanc, G., Burns, M. S., Conley, A., Deustua, S. E., Doi, M., Ellis, R., Fabbro, S., Folatelli, G., Fruchter, A. S., Garavini, G., Garmond, S., Garton, K., Gibbons, R., Goldhaber, G., Goobar, A., Groom, D. E., Hardin, D., Hook, I., Howell, D. A., Kim, A. G., Lee, B. C., Lidman, C., Mendez, J., Nobili, S., Nugent, P. E., Pain, R., Panagia, N., Pennypacker, C. R., Perlmutter, S., Quimby, R., Raux, J., Regnault, N., Ruiz-Lapuente, P., Sainton, G., Schaefer, B., Schahmaneche, K., Smith, E., Spadafora, A. L., Stanishev, V., Sullivan, M., Walton, N. A., Wang, L., Wood-Vasey, W. M., & Yasuda, N. 2003, ApJ, 598, 102

Kogut, A., Banday, A. J., Bennett, C. L., Górski, K. M., Hinshaw, G., Smoot, G. F., & Wright, E. I. 1996, ApJL, 464, L5+

Kogut, A., Hinshaw, G., & Banday, A. J. 1997, PRD, 55, 1901

Kolatt, T. S. & Lahav, O. 2001, MNRAS, 323, 859

Kolb, E. W. & Turner, M. S. 1990, The Early Universe (Westview Press)

Komatsu, E., Kogut, A., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Verde, L., Wollack, E., & Wright, E. L. 2003, ApJS, 148, 119

Kunz, M., Aghanim, N., Cayón, L., Forni, O., Riazuelo, A., & Uzan, J. P. 2006, PRD, 73, 023511

Lagache, G. 2003, AAP, 405, 813

Land, K. & Magueijo, J. 2005, PhRvL, 95, 071301, astro-ph title: "The Axis of Evil"

—. 2006, MNRAS, 312

Larson, D. L. & Wandelt, B. D. 2004, ApJL, 613, L85

Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, MNRAS, 225, 307

Lazarian, A. & Draine, B. T. 2000, ApJL, 536, L15

Lee, A. T., Ade, P., Balbi, A., Bock, J., Borrill, J., Boscaleri, A., de Bernardis, P., Ferreira, P. G., Hanany, S., Hristov, V. V., Jaffe, A. H., Mauskopf, P. D., Netterfield, C. B., Pascale, E., Rabii, B., Richards, P. L., Smoot, G. F., Stompor, R., Winant, C. D., & Wu, J. H. P. 2001, ApJL, 561, L1

Leitch, E. M., Readhead, A. C. S., Pearson, T. J., & Myers, S. T. 1997, ApJL, 486, L23+

Levin, J., Scannapieco, E., & Silk, J. 1998, CQGrav, 15, 2689

Lewis, A. & Bridle, S. 2002, PRD, 66, 103511

Liddle, A. R. & Lyth, D. H. 2000, Cosmological Inflation and Large-scale Structure (Cambridge University Press)

Linde, A. 2004, JCAP, 0410, 004

Lineweaver, C. H. 2003, in Proceedings of the New Cosmology Summer School, Vol. (astro-ph/0305179)

Luminet, J.-P., Weeks, J. R., Riazuelo, A., Lehoucq, R., & Uzan, J.-P. 2003, Nature, 425, 593

Mather, J. C., Cheng, E. S., Cottingham, D. A., Eplee, Jr., R. E., Fixsen, D. J., Hewagama, T., Isaacman, R. B., Jensen, K. A., Meyer, S. S., Noerdlinger, P. D., Read, S. M., Rosen, L. P., Shafer, R. A., Wright, E. L., Bennett, C. L., Boggess, N. W., Hauser, M. G., Kelsall, T., Moseley, Jr., S. H., Silverberg, R. F., Smoot, G. F., Weiss, R., & Wilkinson, D. T. 1994, ApJ, 420, 439

Matzner, R., Rothman, T., & Ellis, G. F. R. 1986, PhRvD, 34, 2926

McEwen, J. D., Hobson, M. P., Lasenby, A. N., & Mortlock, D. J. 2005, MNRAS, 359, 1583

—. 2006, (astro-ph/0605305

Medeiros, J. & Contaldi, C. R. 2006, MNRAS, 367, 39

Misner, C. W., Thorne, K. S., & Wheeler, J. A. 1973, Gravitation (W. H. Freeman and Company)

Moffat, J. W. 2005, JCAP, 10, 12

Mortlock, D. J., Challinor, A. D., & Hobson, M. P. 2002, MNRAS, 330, 405

Narlikar, J. V. 2005, in International Cosmic Ray Conference, 37–+

Narlikar, J. V., Vishwakarma, R. G., Hajian, A., Souradeep, T., Burbidge, G., & Hoyle, F. 2003, ApJ, 585, 1

Naselsky, P., Chiang, L.-Y., Olesen, P., & Novikov, I. 2005, PRD, 72, 063512

Naselsky, P. D., Novikov, I. D., & Chiang, L.-Y. 2006, ApJ, 642, 617

Padmanabhan, T. 2003, Phys. Rep., 380, 235

Page, L., Barnes, C., Hinshaw, G., Spergel, D. N., Weiland, J. L., Wollack, E., Bennett, C. L., Halpern, M., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., & Wright, E. L. 2003, ApJS, 148, 39

Page, L. e. a. 2006, (astro-ph/0603450)

Paladini, R., Burigana, C., Davies, R. D., Maino, D., Bersanelli, M., Cappellini, B., Platania, P., & Smoot, G. 2003, AAP, 397, 213

Paladini, R., Davies, R. D., & DeZotti, G. 2004, MNRAS, 347, 237

Partridge, B. 1995, 3K: The Cosmic Microwave Background Radiation (Cambridge University Press)

Pauliny-Toth, I. K. & Shakeshaft, J. R. 1962, MNRAS, 124, 61

Peacock, J. A. 1999, Cosmological Physics (Cambridge University Press)

Penzias, A. A. & Wilson, R. W. 1965, ApJ, 142, 419

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1988, Numerical Recipes in C++ (Cambridge University Press)

Rakić, A., Räsänen, S., & Schwarz, D. J. 2006, MNRAS, 369, L27

Reich, P. & Reich, W. 1986, A&AS, 63, 205

—. 1988, AAP, 196, 211

Riazuelo, A., Uzan, J.-P., Lehoucq, R., & Weeks, J. 2004a, PRD, 69, 103514

Riazuelo, A., Weeks, J., Uzan, J.-P., Lehoucq, R., & Luminet, J.-P. 2004b, PRD, 69, 103518

Rocha, G., Cayón, L., Bowen, R., Canavezes, A., Silk, J., Banday, A. J., & Górski, K. M. 2004, MNRAS, 351, 769

Roll, P. G. & Wilkinson, D. T. 1966, PhRvL, 16, 405

Roukema, B. F., Lew, B., Cechowska, M., Marecki, A., & Bajtlik, S. 2004, AAP, 423, 821

Ruhl, J. E., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Contaldi, C. R., Crill, B. P., de Bernardis, P., De Troia, G., Ganga, K., Giacometti, M., Hivon, E., Hristov, V. V., Iacoangeli, A., Jaffe, A. H., Jones, W. C., Lange, A. E., Masi, S., Mason, P., Mauskopf, P. D., Melchiorri, A., Montroy, T., Netterfield, C. B., Pascale, E., Piacentini, F., Pogosyan, D., Polenta, G., Prunet, S., & Romeo, G. 2003, ApJ, 599, 786

Sánchez, A. G., Baugh, C. M., Percival, W. J., Peacock, J. A., Padilla, N. D., Cole, S., Frenk, C. S., & Norberg, P. 2006, MNRAS, 366, 189

Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525

Shaver, P. A., McGee, R. X., Newton, L. M., Danks, A. C., & Pottasch, S. R. 1983, MNRAS, 204, 53

Slosar, A., Melchiorri, A., & Silk, J. I. 2005, PRD, 72, 101301

Spergel, D. N., Bean, R., Dore', O., Nolta, M. R., Bennett, C. L., Hinshaw, G., Jarosik, N., Komatsu, E., Page, L., Peiris, H. V., Verde, L., Barnes, C., Halpern, M., Hill, R. S., Kogut, A., Limon, M., Meyer, S. S., Odegard, N., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2006, (astro-ph/0603499)

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003, ApJS, 148, 175

Stannard, A. & Coles, P. 2005, MNRAS, 364, 929

Stompor, R., Balbi, A., Borrill, J. D., Ferreira, P. G., Hanany, S., Jaffe, A. H., Lee, A. T., Oh, S., Rabii, B., Richards, P. L., Smoot, G. F., Winant, C. D., & Wu, J.-H. P. 2002, PRD, 65, 022003

Tegmark, M. 1997, ApJL, 480, L87+

Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, PRD, 68, 123523

Tegmark, M., Strauss, M. A., Blanton, M. R., Abazajian, K., Dodelson, S., Sandvik, H., Wang, X., Weinberg, D. H., Zehavi, I., Bahcall, N. A., Hoyle, F., Schlegel, D., Scoccimarro, R., Vogeley, M. S., Berlind, A., Budavari, T., Connolly, A., Eisenstein, D. J., Finkbeiner, D., Frieman, J. A., Gunn, J. E., Hui, L., Jain, B., Johnston, D., Kent, S., Lin, H., Nakajima, R., Nichol, R. C., Ostriker, J. P., Pope, A., Scranton, R., Seljak, U., Sheth, R. K., Stebbins, A., Szalay, A. S., Szapudi, I., Xu, Y., Annis, J., Brinkmann, J., Burles, S., Castander, F. J., Csabai, I., Loveday, J., Doi, M., Fukugita, M., Gillespie, B., Hennessy, G., Hogg, D. W., Ivezić, Ž., Knapp, G. R., Lamb, D. Q., Lee, B. C., Lupton, R. H., McKay, T. A., Kunszt, P., Munn, J. A., O'Connell, L., Peoples, J., Pier, J. R., Richmond, M., Rockosi, C., Schneider, D. P., Stoughton, C., Tucker, D. L., vanden Berk, D. E., Yanny, B., & York, D. G. 2004, PRD, 69, 103501

The Planck Collaboration. 2006, (astro-ph/0604069)

Vale, C. 2005, (astro-ph/0509039)

Verde, L., Peiris, H. V., & Jimenez, R. 2006, JCAP, 1, 19

Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., & Cayón, L. 2004, ApJ, 609, 22

Wainwright, J. & Ellis, G. F. R., eds. 1997, Dynamical Systems in Cosmology (Cambridge University Press)

Wandelt, B. D. & Górski, K. M. 2001, PRD, 63, 123002

Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, PRD, 70, 083511

Wang, X., Tegmark, M., Jain, B., & Zaldarriaga, M. 2003, PRD, 68, 123001

Wang, Y., Kostov, V., Freese, K., Frieman, J. A., & Gondolo, P. 2004, JCAP, 12, 3

Watson, R. A., Rebolo, R., Rubiño-Martín, J. A., Hildebrandt, S., Gutiérrez, C. M., Fernández-Cerezo, S., Hoyland, R. J., & Battistelli, E. S. 2005, ApJL, 624, L89

Weinberg, S. 2003, Nature, 426, 389

White, M., Scott, D., & Silk, J. 1994, ARA&A, 32, 319

Woermann, B., Gaylard, M. J., & Otrupcek, R. 2000, MNRAS, 315, 241

Wood, K. & Mathis, J. S. 2004, MNRAS, 353, 1126

Acknowledgements

"You're very clever, young man, very clever," said the old lady. "But it's turtles all the way down." – S. Hawking, *A Brief History of Time*