

Testing the cosmological principle of isotropy: local power-spectrum estimates of the *WMAP* data

F. K. Hansen,^{1★} A. J. Banday^{2★} and K. M. Górski^{3,4★}

¹*Dipartimento di Fisica, Università di Roma ‘Tor Vergata’, Via della Ricerca Scientifica 1, I-00133 Roma, Italy*

²*Max-Planck-Institut für Astrophysik, Karl-Scharzschild-Strasse 1, Postfach 1317, D-85741 Garching bei München, Germany*

³*JPL, M/S 169/327, 4800 Oak Grove Drive, Pasadena, CA 91109, USA*

⁴*Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland*

Accepted 2004 July 13. Received 2004 June 9; in original form 2004 April 20

ABSTRACT

We apply the Gabor transform methodology proposed by Hansen et al. to the *WMAP* data in order to test the statistical properties of the cosmic microwave background (CMB) fluctuation field and specifically to evaluate the fundamental assumption of cosmological isotropy. In particular, we apply the transform with several apodization scales, thus allowing the determination of the positional dependence of the angular power spectrum with either high spatial localization or high angular resolution (i.e. narrow bins in multipole space). Practically, this implies that we estimate the angular power spectrum locally in discs of various sizes positioned in different directions: small discs allow the greatest sensitivity to positional dependence, whereas larger discs allow greater sensitivity to variations over different angular scales. In addition, we determine whether the spatial position of a few outliers in the angular power spectrum could suggest the presence of residual foregrounds or systematic effects. For multipoles close to the first peak, the most deviant local estimates from the best-fitting *WMAP* model are associated with a few particular areas close to the Galactic plane. Such deviations also include the ‘dent’ in the spectrum just shortward of the first peak which was remarked upon by the *WMAP* team. Estimating the angular power spectrum excluding these areas gives a slightly higher first Doppler peak amplitude.

Finally, we probe the isotropy of the largest angular scales by estimating the power spectrum on hemispheres and reconfirm strong indications of a north–south asymmetry previously reported by other authors. Indeed, there is a remarkable lack of power in a region associated with the North ecliptic Pole. With the greater fidelity in ℓ -space allowed by this larger sky coverage, we find tentative evidence for residual foregrounds in the range $\ell = 2$ –4, which could be associated with the low measured quadrupole amplitudes and other anomalies on these angular scales (e.g. planarity and alignment). However, over the range $\ell = 5$ –40 the observed asymmetry is much harder to explain in terms of residual foregrounds and known systematic effects.

By reorienting the coordinate axes, we partition the sky into different hemispheres and search for the reference frame which maximizes the asymmetric distribution of power. The North Pole for this coordinate frame is found to intersect the sphere at $(80^\circ, 57^\circ)$ in Galactic colatitude and longitude over almost the entire multipole range $\ell = 5$ –40. Furthermore, the strong negative outlier at $\ell = 21$ and the strong positive outlier at $\ell = 39$, as determined from the global power spectrum by the *WMAP* team, are found to be associated with the Northern and Southern hemispheres, respectively (in this frame of maximum asymmetry). Thus, these two outliers follow the general tendency of the multipoles $\ell = 5$ –40 to be of systematically lower amplitude in the north and higher in the south. Such asymmetric distributions of power on the sky provide a serious test for the cosmological principle of isotropy.

★E-mail: frodekh@roma2.infn.it (FKH); banday@MPA-Garching.MPG.DE (AJB); Krzysztof.M.Gorski@jpl.nasa.gov (KMG)

Key words: methods: data analysis – methods: statistical – techniques: image processing – cosmological parameters – cosmology: observations.

1 INTRODUCTION

One of the fundamental assumptions of cosmology is that the Universe is isotropic. Such an assumption necessarily implies therefore that the statistical properties of the cosmic microwave background (CMB) should be the same in all directions on the sky. The FIRAS experiment on the *COBE* satellite (Smoot et al. 1992; Bennett et al. 1996) demonstrated that the mean temperature of the CMB is isotropic to high precision, but it is only now with the superior sensitivity of the *Wilkinson Microwave Anisotropy Probe* (*WMAP*; Bennett et al. 2003) data that the isotropic nature of the angular fluctuations in the CMB can be tested.

In this paper we will consider the variation in the angular power spectrum of the CMB determined from patches sampling different locations on the sky. In an isotropic Universe, there should be no preferred direction and these locally determined angular power spectra, after allowing for some sample variance, should demonstrate the same expectation value over all positions.

We adopt the Gabor transform method proposed by Hansen et al. (2002) and Hansen & Górski (2003) and estimate the power spectrum in discs with 9.5° and 19° radius and then on hemispheres, thus allowing either high spatial localization or high angular resolution, both of which are used to probe the directional dependence of the anisotropy. The discs on which we estimate the power spectrum are positioned such that they cover as much area as possible without overlapping so much as to introduce strong correlations between the locally estimated spectra. With the smaller discs, we can test the isotropy of the spectrum at a high spatial resolution, but small patches limit the resolution in multipole space. By increasing the size of the discs we test the isotropy of the spectrum at a refined resolution in multipole space but at the cost of a low spatial resolution. We apply different tests of consistency between the spectra in different multipole bins, comparing with a Monte Carlo ensemble generated with the best-fitting model spectrum obtained by the *WMAP* team and using the noise and beam properties of *WMAP*. We have chosen to focus on the co-added W+V map as these are the channels in which the foreground contamination is expected to be lowest. In all cases, only data outside of the *WMAP* Kp2 sky cut are included in the analysis.

With this approach we can also test for the presence of contaminating residual foregrounds or systematic effects. Indeed, several tests of Gaussianity of the *WMAP* data (Chiang et al. 2003; Coles et al. 2004; Copi, Huterer & Starkman 2004; Komatsu et al. 2003; Magueijo & Medeiros 2004; Naselsky, Doroshkevich & Verkhodanov 2003a; Naselsky et al. 2003b; Park 2004; Vielva et al. 2004, Cabella et al. 2004; Cruz et al. 2004; Hansen et al. 2004; Eriksen et al. 2004b,a; Land & Magueijo 2004; Larson & Wandelt 2004; Naselsky, Doroshkevich & Verkhodanov 2004a; Naselsky et al. 2004b) have reported the presence of non-Gaussian signals at different scales, and attributed such results to residual foreground contamination. Note that such contamination could give rise to an anisotropic distribution of the power spectrum. For example, incorrectly subtracted Galactic foregrounds would most likely manifest asymmetry in the Galactic frame of reference, with an in-

creasingly likely level of residual contamination towards the plane. Similarly, systematic effects would plausibly align with the ecliptic frame of reference reflecting the scanning strategy of the satellite. In addition, contamination of the data due to Solar System emission would be confined to the ecliptic plane. For these reasons, we compare the statistical distribution of the power-spectrum estimates in both the Galactic and ecliptic Northern and Southern hemispheres, as well as the polar and equatorial regions for these reference frames.

We will in particular check the positional dependence of some features in the power spectrum observed by *WMAP* (Hinshaw et al. 2003): two at large scales corresponding to bins centred around $\ell = 21$ and $\ell = 39$, one close to the first Doppler peak and one close to the first trough (the latter was not seen in the analysis by *WMAP*). Moreover we investigate the recent claims of an asymmetry of the large-scale structure between the Northern and Southern Galactic hemispheres (Copi et al. 2004; Park 2004; Vielva et al. 2004; Cruz et al. 2004; Eriksen et al. 2004a,b; Hansen et al. 2004; Land & Magueijo 2004; Larson & Wandelt 2004). As the asymmetry is reported to be seen at the largest scales, we use the results of the power spectra estimated on hemispheres where the multipole resolution is large enough to resolve different features in the range $\ell = 2-63$.

The outline of the paper is as follows. In Section 2 we outline the method described in Hansen et al. (2002) and Hansen & Górski (2003) to estimate the power spectrum on a patch on the sky using the so called Gabor transforms. The details of the data set used in the paper is outlined in Section 3. In Section 4 we show the results of the local *WMAP* spectra estimated on discs of radius 9.5° . We check for consistency between the locally estimated spectra in different multipole bins and different directions. To check the results at a higher multipole resolution, the analysis is repeated in Section 5 with discs of radius 19° . In Section 6 we check for an asymmetry in the large-scale structure and report results of the *WMAP* spectra estimated on hemispheres. In particular, we test the robustness of the results using other frequency channels and a variety of sky cuts. Finally, Section 7 summarizes our results.

2 THE GABOR TRANSFORMATION ON THE SPHERE

In (Hansen et al. 2002, from now on HGH) we introduced the method of Gabor transforms on the sphere and showed how its application could be used to analyse the CMB power spectrum, using the pseudo-spectral information from a given patch on the sky. The pseudo-spectrum is used as input data to a likelihood analysis for an axisymmetric patch on the sky. For such a sky partition, the correlation properties of the sky signal can be calculated analytically. In this paper, we extend the method to an arbitrary asymmetric patch on the sky, by using Monte Carlo pre-calculations of the correlation matrix. In this section we will outline the method as described in HGH. The details of the extension to arbitrary patches are deferred to the Appendix; here we will focus on the results.

We define the pseudo-power-spectrum \tilde{C}_ℓ as the power spectrum obtained using the spherical harmonic transform on a CMB temperature field $T(\hat{n})$ multiplied by a window function $G(\hat{n})$. Letting $a_{\ell m}$

be the spherical harmonic coefficients of the CMB field and C_ℓ the corresponding power spectrum, the corresponding quantity on the windowed sky is indicated by a tilde and can be written as

$$\tilde{C}_\ell = \sum_m \frac{\tilde{a}_{\ell m}^* \tilde{a}_{\ell m}}{2\ell + 1}, \quad (1)$$

where

$$\tilde{a}_{\ell m} = \int d\hat{n} T(\hat{n}) G(\hat{n}) Y_{\ell m}^*(\hat{n}). \quad (2)$$

This pseudo-spectrum can be written in terms of the full-sky spectrum by means of the following formulae:

$$\langle \tilde{C}_\ell \rangle = \frac{1}{2\ell + 1} \sum_{\ell'} C_{\ell'} K(\ell, \ell') \quad (3)$$

where the kernel $K(\ell, \ell')$ can be calculated analytically by the formulae given in HGH, or by Monte Carlo methods as described here in the Appendix.

The method proposed in HGH was to sample a disc-like region of the observed sky $T(\hat{n})$ using a top-hat or a Gaussian window $G(\hat{n})$. The pseudo-power-spectrum of this disc was then used as the input vector to a likelihood analysis, assuming the distribution of the \tilde{C}_ℓ to be Gaussian. It was demonstrated that this is a good approximation when the multipole ℓ is sufficiently high, which for a Gaussian window with 15° full width at half maximum (FWHM) corresponded to $\ell \sim 100$. HGH employed the following likelihood Ansatz:

$$\mathcal{L} = \frac{\exp(-\frac{1}{2} \mathbf{d}^T \mathbf{M}^{-1} \mathbf{d})}{\sqrt{2\pi \det \mathbf{M}}}, \quad (4)$$

where the elements of the data vector are

$$d_i = \tilde{C}_{\ell_i} - \langle \tilde{C}_{\ell_i}^S \rangle - \langle \tilde{C}_{\ell_i}^N \rangle. \quad (5)$$

Here the first term is the observed pseudo-power-spectrum including noise and the second term is the expected signal pseudo-spectrum obtained from equation (3). The last term contains the expected noise pseudo-spectrum for the patch. Similarly, the matrix expressing the correlation between the pseudo-power-spectrum coefficients can be written as

$$M_{ij} = \langle \tilde{C}_{\ell_i} \tilde{C}_{\ell_j} \rangle - \langle \tilde{C}_{\ell_i} \rangle \langle \tilde{C}_{\ell_j} \rangle = M_{ij}^S + M_{ij}^N + M_{ij}^X, \quad (6)$$

where there is one contribution from the signal (S), one from the noise (N) and one from signal–noise mixing (X). The form of these matrices and how they can be computed are discussed in HGH for an axisymmetric region of the sky. Here, in the Appendix, we show how to extend the calculation of these matrices to an arbitrarily shaped patch using a Monte Carlo approach. Note that this likelihood expression is a function of the underlying full-sky power spectrum C_ℓ which we want to estimate. The best estimate is the power spectrum which maximizes this likelihood.

Note that one cannot use all \tilde{C}_ℓ in the analysis, as the correlation matrix will then be overdetermined. Following HGH, the optimal spacing in ℓ -space is given by the width of the kernel and can be expressed approximately as $\Delta\ell = 115/\theta$ where θ is the radius of the window in degrees. We will call the number of samples (\tilde{C}_ℓ) to be used, N_{sam} . Also, one cannot estimate all the full-sky power-spectrum coefficients C_ℓ as there is not enough information in a disc to obtain an estimate of every single multipole. For this reason we use a limited number N_{bin} of flat power-spectrum bins. As discussed in HGH this number should be smaller than N_{sam} to allow for at least three or four observations (\tilde{C}_ℓ) per bin. A more complete discussion

on the number of input multipoles and full-sky bins to use is given in HGH.

In HGH it was also demonstrated that one can simultaneously include the pseudo-power-spectra from several discs in order to analyse a full-sky observation. In this case, the log-likelihood can be written as

$$L = \sum_{I=1}^{N_{\text{disc}}} \mathbf{d}_I^T \mathbf{M}_I^{-1} \mathbf{d}_I + \sum_{I=1}^{N_{\text{disc}}} \ln \det \mathbf{M}_I, \quad (7)$$

where d_I is the full data vector of disc number I , \mathbf{M}_I is similarly the correlation matrix of disc I and N_{disc} is the number of discs considered. This form of the likelihood assumes the \tilde{C}_ℓ between different discs to be uncorrelated. For minimally overlapping discs this can remain a good approximation. It will be demonstrated later that for the set of discs used here, only a tiny correlation between the estimated power spectra of overlapping discs is introduced.

The purpose of this paper is to use these locally estimated spectra to investigate the isotropy of the CMB in the WMAP data. We will assume the best-fitting WMAP running-index spectrum, make a Monte Carlo ensemble of estimated power spectra on discs of different sizes and compare these to the spectra estimated on the WMAP data set.

3 PREPARING THE DATA

In this paper we analyse the publicly available WMAP data¹ after correction for Galactic foregrounds. We consider data from the V and W frequency bands only since these are less likely to be contaminated by residual foregrounds. The maps from the different channels are co-added using inverse noise-weighting:

$$T = \frac{\sum_c T_c / \sigma_c^2}{\sum_c 1 / \sigma_c^2},$$

where c runs over the frequency channels V1, V2, W1, W2, W3, W4. In the simulations we have been generating maps smoothing with a beam and adding noise equal to the effective beam and noise of the co-added W+V channel (also available on the LAMBDA Web site). The input power spectrum in the simulation was the best-fitting WMAP running-index spectrum. Most of the analysis in this paper has been performed using the WMAP Kp2 sky cut, including the point-source mask. In those cases where, in order to test the robustness of the results, we have used either an alternative power spectrum, different frequency channels or different masks, the differences have been clearly stated.

It should be noted that the WMAP team have utilized the cross-spectra between different channels for their estimation of the power spectrum, which is therefore less dependent on the exact noise properties of the experiment than in our case. Here we have adopted the approximate white-noise model provided by the WMAP team, and uncertainties in this model could result in erroneous estimates on smaller angular scales where the multipoles are noise dominated. In fact, it has been determined that the noise level estimated using the tail of the pseudo-power-spectrum on a given disc is slightly higher than that obtained using the white-noise model. In principle, this could be accounted for using the 110 noise realizations per channel, which include all of the effects of the full data processing pipeline, which have been made available by the WMAP team on the LAMBDA Web site. In practice, this number is not sufficient to

¹ These can be obtained at the LAMBDA Web site: <http://lambda.gsfc.nasa.gov/>




Figure 1. The position and numbering of the 9.5° discs on which the local power spectra have been calculated.

Table 1. The power-spectrum bins for the 9.5° disc estimates.

bin	1	2	3	4	5	6	7
range	$\ell = 2-63$	$\ell = 64-113$	$\ell = 114-163$	$\ell = 164-213$	$\ell = 214-263$	$\ell = 264-313$	$\ell = 314-363$
centre	$\ell = 33$	$\ell = 89$	$\ell = 139$	$\ell = 189$	$\ell = 239$	$\ell = 289$	$\ell = 339$
bin	8	9	10	11	12		
range	$\ell = 364-413$	$\ell = 414-463$	$\ell = 464-513$	$\ell = 514-563$	$\ell = 564-613$		
centre	$\ell = 389$	$\ell = 439$	$\ell = 489$	$\ell = 539$	$\ell = 589$		

converge to an accurate estimate of the statistical properties of the noise for individual discs.

4 RESULTS OF THE SPECTRA ESTIMATED ON SMALL DISCS

In this section, we test the isotropy of the CMB with high spatial sensitivity by estimating the power spectrum on discs of radius 9.5° . As a consequence of the small disc size, the multipoles are strongly coupled and wide power-spectrum bins must be employed. Using a set of simulated maps we have created a Monte Carlo ensemble of estimates of the power spectrum on individual small discs from the disc set, and additionally estimates of the power spectrum corresponding to larger sample regions by combining several discs. These are then compared to the estimates derived from the *WMAP* data.

4.1 The analysis

With the method described above, we have estimated the angular power spectrum for 164 discs with an angular radius of 9.5° uniformly distributed on the part of the sphere outside of the *WMAP* Kp2 sky cut. The disc position and disc size was chosen in order to cover

as much as possible of the observed sky with as many samples (discs) as possible without losing too much multipole resolution (i.e. to maintain enough power-spectrum bins to resolve the shape of the power spectrum). The overlap of the discs was also kept to a minimum in order to reduce correlations between the locally estimated spectra. We sampled the pseudo-spectrum in $N_{\text{sam}} = 84$ points which we used as input data to the likelihood estimation procedure. We then estimated $N_{\text{bin}} = 12$ power-spectrum bins in the multipole range $\ell = 2-850$. As the correlation matrix in the likelihood can become singular if the observed area is too small, we have excluded all discs where less than 90 per cent of the disc area remains after application of the Kp2 mask. In Fig. 1 we show the remaining 130 discs with the numbering which will be used in the following treatment. The distance in the latitudinal θ direction between the disc centres going from the Galactic poles to the equator is roughly 14° and in the longitudinal ϕ direction the centres are all at an equal distance for a given Galactic latitude. The 12 power-spectrum bins b are summarized in Table 1. As described in HGH, the C_ℓ coupling kernel has a similar width at all scales, giving a similar resolution over all multipoles. Therefore we have adopted bins of a fixed size $\Delta\ell$ in the whole multipole range.

Fig. 2 shows the disc-disc correlation matrix for the bins 1, 4 and 7 (centred at $\ell = 33$, $\ell = 189$ and $\ell = 339$). In Fig. 3 some slices




Figure 2. The disc–disc correlation matrix for the multipole bins 1, 4 and 7 centred at $\ell = 33$, $\ell = 189$ and $\ell = 339$, respectively. The disc size used here was 10° . The shaded areas show the off-diagonal elements of the correlation matrix which are above 5 per cent of the diagonal. All the off-diagonal elements are below 20 per cent of the diagonal.

of the correlation matrix are shown for disc 0, which has significant overlap with several other discs (1–5), for disc 6, which has overlap with two other discs (1 and 16) and for disc 60, which shows little intersection with other regions. The other discs plotted were selected because of the properties of certain of their multipole bins, as will become evident later in this section. For the lowest bin one can see some correlation between several discs, whereas for the higher bins the correlations are limited to the neighbouring discs. Taking into account the extension of the structures at the largest scales, a stronger correlation could be expected and in the figure one can even see correlation between non-overlapping discs for the lowest multipole bin. Note that for higher-multipole bins, the off-diagonal elements of the disc–disc correlation matrix are at a maximum 10–20 per cent of the diagonal. Moreover the disc–disc correlations for these multipoles are only significant between strongly overlapping discs. Finally Fig. 4 shows the bin-to-bin correlation matrix for a single disc. One can see that, apart from a weak anticorrelation between neighbouring bins, the disc estimates are essentially uncorrelated.

4.2 The results of the analysis applied to the WMAP data on individual discs

As a first test of isotropy, we will look at the distribution of the power-spectrum estimates over the sphere. In Fig. 5 we have plotted the best-fitting WMAP running-index spectrum (solid line), the same spectrum binned as in our estimation procedure (histogram) and the power-spectrum results of the WMAP team with this binning (crosses). The shaded areas indicate the spread of our power-spectrum estimates on the WMAP data over the different discs. The bands represent the areas containing 67 and 95 per cent of the disc estimates for a given bin. Note that the mean of the disc estimates close to the first peak (bin 4) lies below the best-fitting prediction. This corresponds to the ‘dip’ in the power spectrum pointed out by the WMAP team. Moreover, the distribution of estimates for the spectrum does not contain strong outliers due to some particular location on the sky.

The aim of the disc estimation procedure is to check the isotropy of the CMB fluctuation field. But as explained in the introduction,

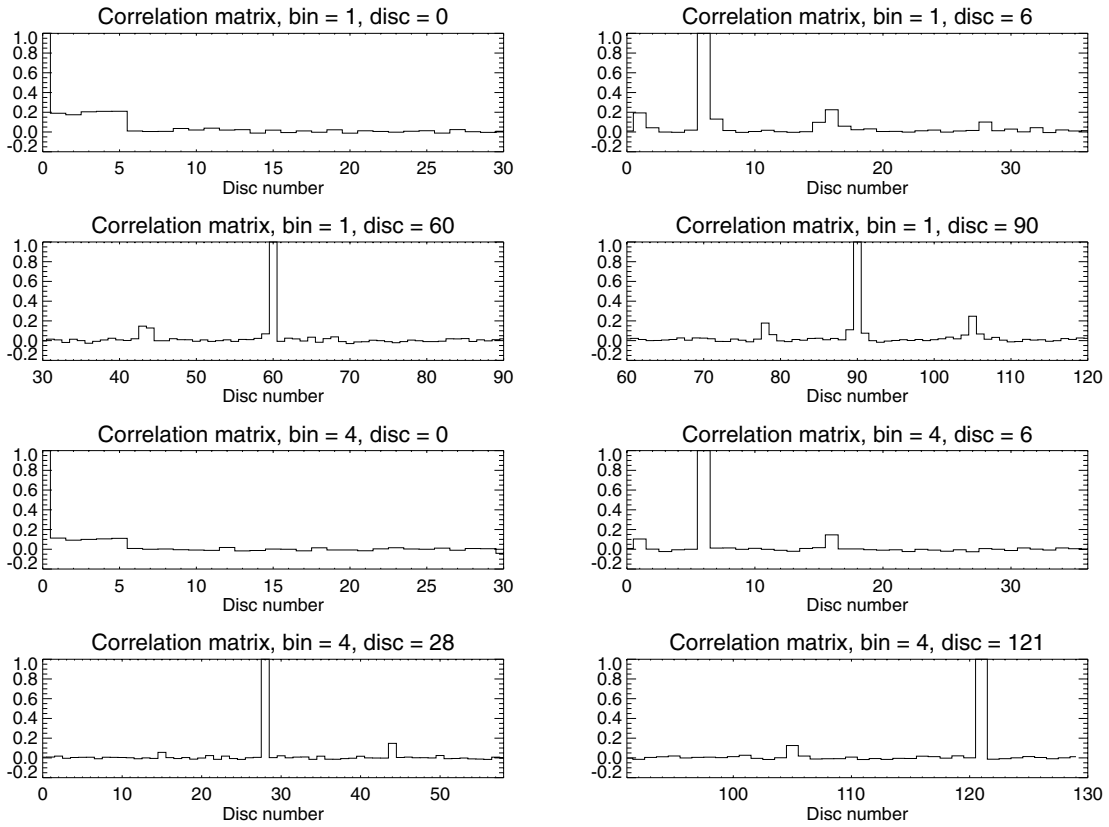


Figure 3. Slice of the disc correlation matrix for some selected discs at the lowest multipole bin $\ell = 33$ and a bin centred close to the first power-spectrum peak $\ell = 189$. We have chosen the discs 0, 6 and 60 in order to study the correlations for discs which have many (disc 0) and few (disc 6 and 60) overlapping neighbours. The other discs have been chosen as particular bin values of these discs have been found in the *WMAP* data. We have chosen particular discs close to the ecliptic poles as some particular features are found in this area (see the text and Fig. 9).

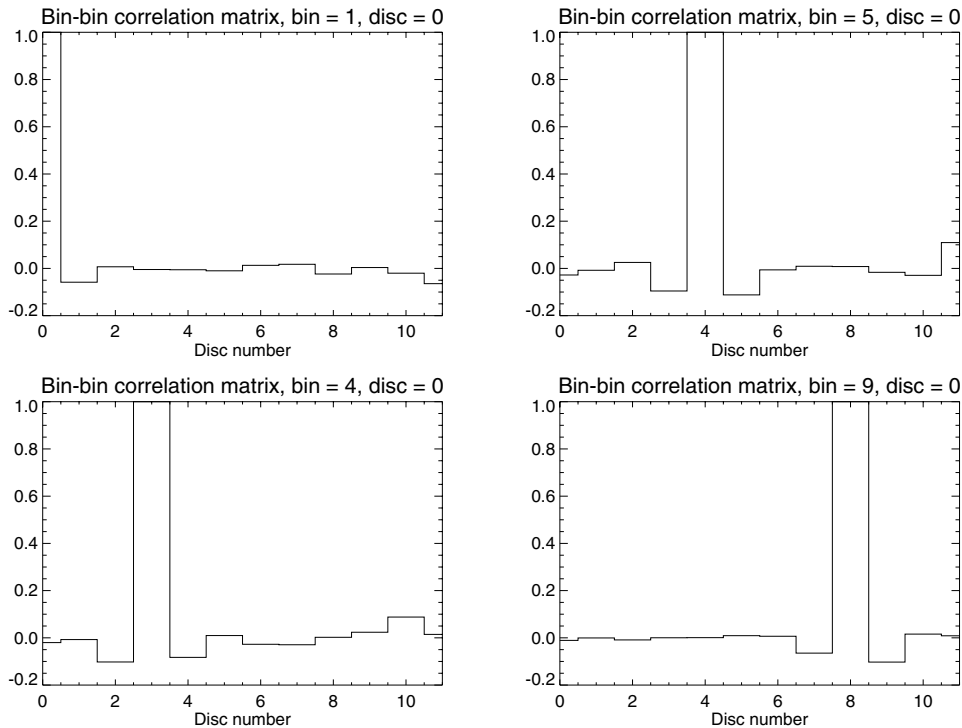


Figure 4. Slice of the bin-to-bin correlation matrix for disc 0 taken at bins 1, 4, 5 and 9.




Figure 5. The best-fitting *WMAP* running-index power spectrum (solid line). The histogram shows the spectrum binned using the same bins as for the disc estimates. The crosses show the binned full-sky *WMAP* spectrum obtained by the *WMAP* team and the shaded areas show the spread of the *WMAP* spectrum over 130 discs of radius 9.5° , found from the Gabor analysis. The two shaded areas indicate where 67 and 95 per cent of the 130 *WMAP* disc estimates are contained. Note that the shaded areas are *not* error bars for the *WMAP* estimates, but the spread found by estimating the spectrum in different positions.




Figure 6. The best-fitting *WMAP* running-index power spectrum (solid line). The histogram shows the spectrum binned in the same way as the disc estimates. The two shaded areas indicate where 67 and 95 per cent of the disc estimates in the simulations are contained. The error bars at each bin show the same distribution for the *WMAP* data, the left being the 1σ level and the right being the 2σ level.

if residual foregrounds or systematic effects are present in the data set, these could show up as asymmetries in the Galactic or ecliptic frames of reference. Therefore, we evaluate the distribution of the power-spectrum estimates in the Galactic and ecliptic Northern and Southern hemisphere as well as the polar and equatorial regions in these reference frames. As a first check, we looked at the distribution of all the 130 disc estimates compared to the same distribution in the 6144 simulations. In Fig. 6 the error bars at each power-spectrum bin show where 67 (left bar) and 95 (right bar) per cent of the disc estimates of *WMAP* data fall (same as the shaded zones in Fig. 5). The shaded zones show the same distribution taken from

the simulations. We see that the width of the distributions for *WMAP* is in good agreement with the simulations.

In Fig. 7 we show the disc distributions of the power spectrum on the Northern and Southern Galactic hemispheres (NGH and SGH, upper-left plot), Galactic polar and equatorial region (GPR and GER, upper-right plot), the Northern and Southern ecliptic hemispheres (NEH and SEH, lower-left plot) and the ecliptic polar and equatorial regions (EPR and EER, lower-right plot). The error bars indicate the 1σ spread of the power spectrum over the discs in the given regions for *WMAP*, whereas the shaded zones show the corresponding 1σ and 2σ spreads for the simulated maps. We find the distributions of




Figure 7. The best-fitting *WMAP* running-index power spectrum (solid line). The histogram shows the spectrum binned in the same way as the estimates on 9:5 discs. The two shaded areas indicate where 67 and 95 per cent of the disc estimates in the *WMAP* map are contained. The error bars at each bin show the same distribution (67 per cent level) for a part of the *WMAP* data. Upper left: Northern Galactic hemisphere (left bar)/Southern Galactic hemisphere (right bar); upper right: Galactic polar region (left bar), Galactic equatorial region (right bar); lower left: Northern ecliptic hemisphere (left bar)/Southern ecliptic hemisphere (right bar); lower right: ecliptic polar region (left)/ecliptic equatorial region (right).

the *WMAP* spectra consistent with the simulations, the only exception being bin 1 for the Northern/Southern (Galactic and ecliptic) hemispheres shown in the plots in the left column. For the Northern hemisphere, this bin value is below the average value for all discs within the 1σ band (indicated by the left error bar). For the Southern hemisphere, however (the right error bar), the distribution seems consistent with simulations.

In Fig. 8 we investigate further the difference between the distribution of the power in bin 1 between the Northern and Southern hemisphere. The figure shows the power in this bin as a function of disc number (given in Fig. 1). The discs on the left side (1–65) belong to the Northern Galactic hemisphere and the discs on the right side (66–129) to the Southern hemisphere. The figure confirms the observation that the large-scale power estimated in the discs on the Northern hemisphere is systematically lower than in the Southern. Note that for the Northern hemisphere only one-fourth of the discs

have values above the average whereas for the Southern hemisphere the values are distributed equally above and below the average. Below the lower 1σ band there are only half as many discs in the Southern hemisphere as in the Northern. The significance of this north–south asymmetry will be studied further in Section 6 using hemispheres for which the multipole resolution at large scales is higher.

Fig. 9 attempts to summarize the spatial distribution of the power spectra derived from the 9:5 discs. We have plotted disc results for the largest angular scales (bin 1 and 2, upper row), the first peak (bin 4 and 6, middle row) and the first trough (bin 8 and 9, lower row). The colour coding indicates where these bins are above (light red: below 1σ , red: between 1σ and 2σ , dark red: above 2σ) and below (light blue: above 1σ , blue: between 1σ and 2σ , dark blue: below 2σ) the average established by simulations. The bin 1 results seem to suggest that the putative north–south asymmetry is actually




Figure 8. The disc estimates of the bin $\ell = 2-63$. The disc numbers correspond to the disc numbers in Fig. 1. The shaded zones indicate the 1σ , 2σ and 3σ spread of the estimated bin for the given disc as found from Monte Carlo simulations.

associated more with the direction of the ecliptic poles. The exact direction of maximum asymmetry will be investigated further in Section 6.

Bin 4 corresponds to the dent in the full-sky spectrum as estimated by the *WMAP* team. There is some evidence that low amplitudes are concentrated close to the ecliptic poles. It also appears that, for scales corresponding to the first trough in the power spectrum (bins 8 and 9), amplitudes are relatively high close to the Southern ecliptic Pole. These features will be discussed further using the joint likelihood estimates in the next section. Apart from this, we cannot find other deviations of isotropy in the distribution of power looking at this figure. Studying the darkest colours (the 2σ outliers), we do not find that they appear in some particular position on the sky. Also, for 164 discs we would expect about 7–8 discs at 2σ . None of the maps have more 2σ outliers. We will test the isotropy more systematically in the next subsection. Finally, note that the disc–disc correlation matrix for one of the discs with an abnormally low value and one of the discs with an abnormally high value are shown in Fig. 3. We see that there is nothing anomalous with the correlations for these discs.

4.3 Results of the analysis combining several discs

We now proceed to a joint likelihood analysis of the locally estimated spectra as described in HGH and outlined here in equation (7). Fig. 10 shows the results from a likelihood estimation including all 130 discs. The grey bands indicate the 1σ and 2σ levels derived from simulations based on the *WMAP* best-fitting spectrum, and the crosses show our result for *WMAP*. Relative to the simulations, bin 4 (centred at $\ell = 189$) is too low at the 1σ level, consistent with the *WMAP* team results (see Fig. 5). Bin 9 (centred at $\ell = 439$) is a strong outlier, but aside from this our full-sky estimate at this multipole resolution is in good agreement with the best-fitting model.

Nevertheless, in the previous section we found tentative evidence for some asymmetry in the large-scale power distribution. In order to confirm this result, and as a test of residual Galactic foregrounds or systematic effects, we also performed a joint likelihood analysis

of several partitions of discs selected over larger areas of the sky to probe the variation of the power spectrum in the Galactic and ecliptic reference frames. In the Galactic reference frame we have elected to test the effect on the spectrum when excluding discs 45–65 and 111–129 (see Fig. 1) which are the two rows closest to the Galactic plane. The surviving discs on the Northern and Southern hemispheres were analysed both separately and together. The spectrum for the three rows of discs closest to the Galactic plane were also computed, for the Northern and Southern hemisphere separately and together. The results are collected in Fig. 11. A corresponding analysis was then performed in the ecliptic reference frame, excluding those discs with nominal centres further than 50° from the ecliptic poles. The results are shown in Fig. 12.

The most striking feature in these plots is the north–south asymmetry visible both in the Galactic and ecliptic frame of reference on large angular scales. Moreover one can see that the low value of bin 4 (centred at $\ell = 189$) seems to be strongest (below the 1σ band) in the ecliptic polar regions (both north and south) whereas it is weaker (inside of 1σ) in the equatorial region (both north and south). This seems to confirm our observation in Fig. 9 and a similar observation made by the *WMAP* team (Hinshaw et al. 2003). However, comparing the ratio of this bin value from the ecliptic polar and ecliptic plane regions to simulations, we find that it is not a significant difference in power. In fact most of the simulations had a higher ratio between the ecliptic polar regions and ecliptic plane for this bin. At the multipole resolution obtained by 9.5 discs we find bin 4 to be consistent with the best-fitting model and isotropically distributed on the sky.

Bin 9 (centred at $\ell = 439$) appears to be of particularly high amplitude both in Galactic polar regions, where one would expect foregrounds to be less dominant, and in the Southern ecliptic polar regions. There are no obvious systematic effects or foregrounds which could cause higher power at this particular scale over such diverse regions of the sky. One plausible solution, however, is connected to the fact that the white-noise model adopted is an approximation, i.e. if the estimated noise level is slightly too low, this would explain the outlier in bin 9 and also the fact that the multipole bins above $\ell > 400$ are systematically above the best-fitting model.

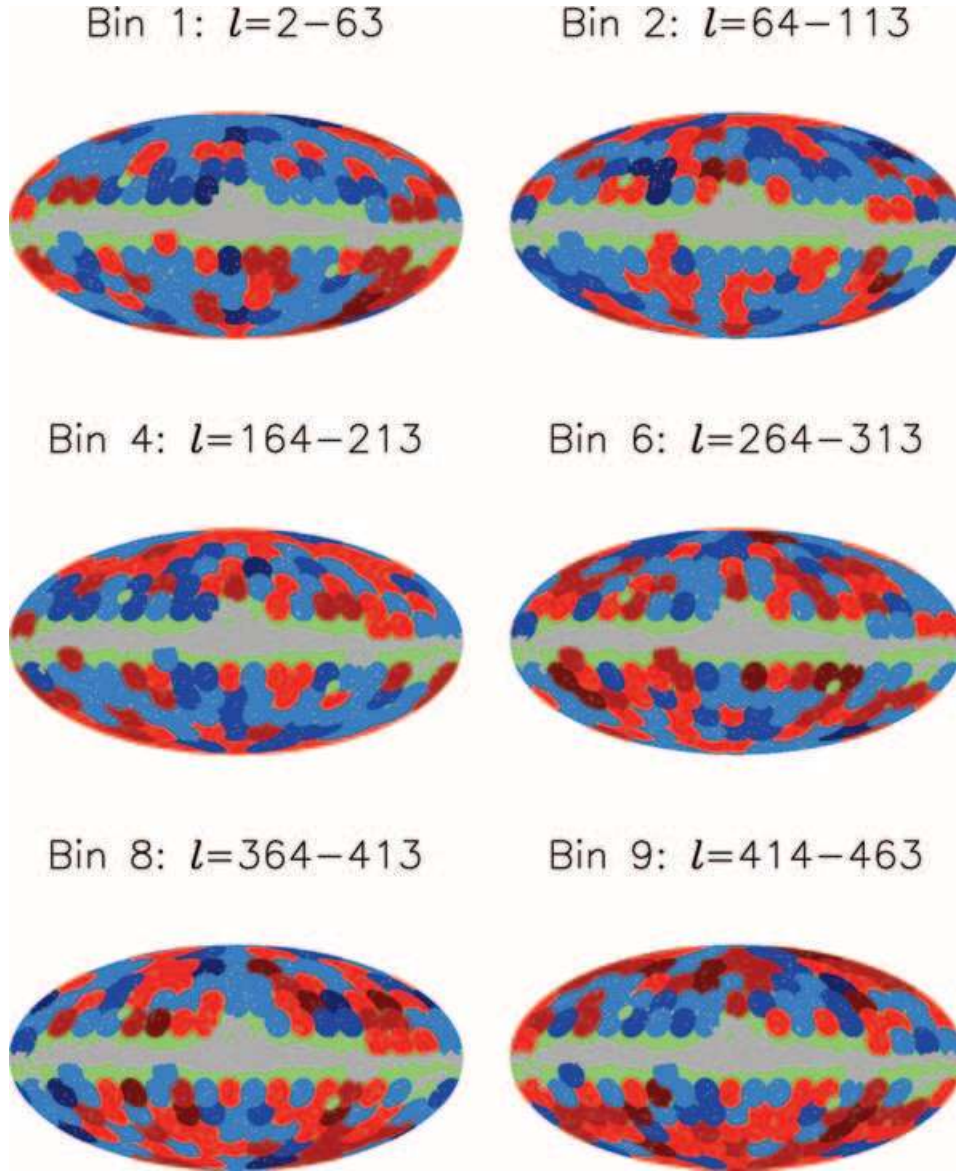


Figure 9. The local power in the bin $\ell = 2-63$ (upper-left plot), $\ell = 64-113$ (upper-right plot), $\ell = 164-213$ (middle-left plot), $\ell = 264-313$ (middle-right plot), $\ell = 364-413$ (lower-left plot) and the bin $\ell = 414-463$ (lower-right plot) estimated on $9^{\circ}.5$ discs. The discs with power above the average set by simulations have yellow (below 1σ), red (between 1σ and 2σ) and dark red (above 2σ) colour. The discs with power below the average have blue colour: light blue (above 1σ), blue (between 1σ and 2σ) and dark blue (below 2σ). The two green discs show the positions of the ecliptic poles.

Furthermore, if this is the case then it is necessary for the *WMAP* team to provide a better model for the noise on these angular scales for future work. Alternatively, a similar enhancement in power could result if the unresolved foreground due to point sources has not been removed correctly. For source removal, we have adopted the procedure described in Hinshaw et al. (2003) which is intended for full-sky measurements. If the distribution of unresolved sources is not completely uniform, applying this procedure to the local disc spectra could introduce errors at high multipoles. However, this is unlikely in practice as the point source distribution is essentially Poissonian in nature.

The conclusion to be drawn from the small-disc analysis is that there is little evidence for departures from isotropy on small angular scales. However, we found a possible north–south asymmetry on the largest scales, which we will return to later in Section 6.

5 RESULTS ON MEDIUM-SIZED DISCS

In this section we aim to test the isotropy of the power distribution on the sky with greater spectral resolution by using discs of radius 19° . In particular, we take advantage of the higher resolution in multipole space to place more bins close to the first Doppler peak in order to test whether its location in ℓ -space exhibits any spatial dependence as traced by the disc positions.

5.1 The analysis

Using the same criteria as for the small discs, we have positioned 34 discs of radius 19° on the sphere as shown in Fig. 13. The distance between the discs and Galactic poles is 35° for the first ring of discs and 65° for the second ring. Again, the distance between




Figure 10. The joint likelihood estimation of 130 discs of radius 9.5° . The histogram shows the mean and the shaded areas the 1σ and 2σ levels from 1536 simulated maps with the same input spectrum. The crosses show the result of the same estimation procedure on the *WMAP* data.

Table 2. The power-spectrum bins for the 19° disc estimates.

bin	1	2	3	4	5
range	$\ell = 2-63$	$\ell = 64-113$	$\ell = 114-163$	$\ell = 164-188$	$\ell = 189-213$
centre	$\ell = 33$	$\ell = 89$	$\ell = 139$	$\ell = 176$	$\ell = 201$
bin	6	7	8	9	10
range	$\ell = 214-238$	$\ell = 239-263$	$\ell = 264-288$	$\ell = 289-313$	$\ell = 314-363$
centre	$\ell = 226$	$\ell = 251$	$\ell = 276$	$\ell = 301$	$\ell = 339$
bin	11	12	13	14	15
range	$\ell = 364-413$	$\ell = 414-463$	$\ell = 464-513$	$\ell = 514-563$	$\ell = 564-613$
centre	$\ell = 389$	$\ell = 439$	$\ell = 489$	$\ell = 539$	$\ell = 589$

discs in the ϕ direction is the same for all discs at a given Galactic latitude. The number of input samples of the pseudo-spectrum was chosen to be $N_{\text{sam}} = 133$ and the number of estimated bins $N_{\text{bin}} = 16$. We have positioned smaller and more bins around the first peak in order to constrain its location in ℓ -space more accurately. The bins are defined in Table 2. We do not show the disc–disc and bin–bin correlation matrices since these are very similar those for the small discs. Even for the small bins around the first peak, the correlations are limited to a small anticorrelation between the neighbouring bins.

5.2 Results on the *WMAP* data

Fig. 14 shows the result of the joint likelihood estimation of the power spectrum by combining all 34 discs. We find that the result is consistent with the small-disc result (Fig. 10). Specifically, there is a dip at the 1σ level at bin 4 centred at $\ell = 176$ and the spectrum is slightly too high after the first peak. There is also a strong outlier at bin 12 centred on $\ell = 439$, but otherwise the inferred full-sky spectrum is consistent with the *WMAP* best-fitting model at this higher multipole resolution. Nevertheless, we proceed to assess the

significance of the spatial distribution of the bin values and the position of outliers as a test of isotropy.

In Fig. 15 we show the positional dependence of the first two bins, the ‘dip’ at bin 4 ($\ell = 164-188$), bin 6 ($\ell = 214-238$) and bin 8 ($\ell = 264-288$) which are at or just follow the first peak, and bin 12 ($\ell = 414-463$) in the first trough. It is clear by comparison with Fig. 9 that the distribution of these bins is largely consistent with the small-disc results. Note, however, that bin 4 has three negative 2σ outliers close to the Galactic plane. This is also the case for bin 8 which has four 2σ outliers close to the Galactic plane. In our simulations, we find three 2σ outliers in the Galactic region for bin 4 in 10 per cent of the maps, so this feature should not be regarded as highly significant. However, four 2σ outliers in this region for bin 8 are found only in 2 per cent of the maps.

Assuming that these deviations might have other than a cosmological origin, it is difficult to speculate what their origin might be. Certainly, errors in Galactic foreground modelling could be a plausible source of the outliers, either due to under- or overprediction of the template-based contribution in these regions, resulting in either over- or undersubtraction of power. Whatever the cause of the features, we examine their impact on the power-spectrum estimation

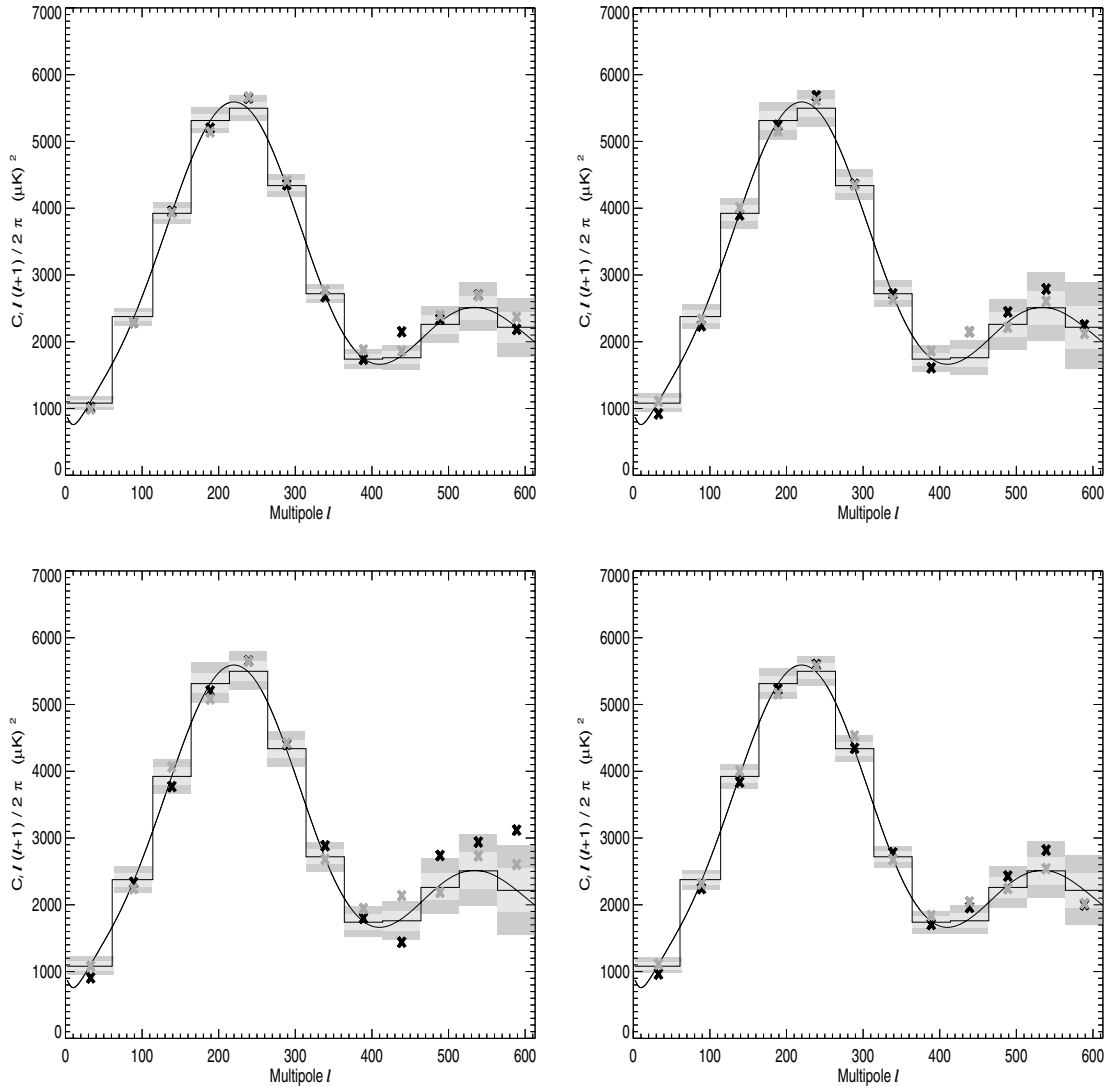


Figure 11. Results of the power-spectrum analysis in different parts of the sky using a joint analysis of $9^{\circ}5$ discs in the Galactic reference frame (same as Fig. 10 on a limited number of discs). The shaded zones indicate the 1σ and 2σ error bars calculated in each plot for one of the spectra shown. The error bars of the other spectrum are very similar and are not shown. Upper-left plot: black crosses: polar regions (discs 0–44, 66–110), grey crosses: equatorial regions (discs 29–65, 95–129). Upper-right plot: black crosses: polar north (discs 0–44), grey crosses: polar south (discs 66–110). Lower-left plot: black crosses: equatorial region north (discs 29–65), grey crosses: equatorial region south (discs 95–129). Lower-right plot: black crosses: Northern hemisphere (discs 0–65), grey crosses: Southern hemisphere (discs 66–129).

by excluding the discs associated with the 2σ outliers in bins 4 and 8, and repeating the joint estimate of the power spectrum using the remaining 19° discs. In Fig. 16 we see the resulting power spectrum. Interestingly, the dent before the first Doppler peak has now disappeared and a slightly higher peak amplitude is preferred.

Finally we exploit the higher multipole space resolution to estimate the position and amplitude of the first peak in different directions. Following Page et al. (2003), we approximate the first peak as a Gaussian. Using the estimated power spectra we applied a maximum-likelihood procedure to find the best-fitting peak position and amplitude. Fig. 17 shows the result of the peak position estimates on 512 simulated maps. The shaded zones indicate the 1σ and 2σ spread of the peak position. The circles show the mean 1σ and 2σ contours for the individual disc estimates. Note that any specific disc will have significance levels different from the mean levels shown by the circles. The crosses show the peak position on

the individual 19° discs using the *WMAP* data, and the bold dark cross shows the joint result of the 34 discs of the *WMAP* data. We find the estimated peak position, amplitude and the distribution of the position and amplitude in individual discs consistent with the best-fitting model. The white bold cross shows the position of the first peak when excluding those discs with possible residual Galactic contamination. In this case, the agreement with the model spectrum is improved. In Fig. 18 we show the discs where the peak position and amplitude are above (red) or below (blue) the mean, and where they are within the 1σ level (bright colour), between 1σ and 2σ (medium) and outside of 2σ (dark colour). We find no particular positional dependence of the first peak position and amplitude.

In conclusion, the power on small scales appears to be isotropically distributed, although there is some evidence that contamination of several discs close to the Galactic plane may be associated with the dent in the power spectrum close to the first peak. When




Figure 12. Results of the power-spectrum analysis in different parts of the sky using a joint analysis of 9:5 discs in the ecliptic reference frame (same as Fig. 10 on a limited number of discs). The shaded zones indicate the 1σ and 2σ error bars calculated in each plot for one of the spectra shown. The error bars of the other spectrum are very similar and are not shown (this is not the case for the ecliptic polar regions versus ecliptic plane, but still the qualitative results are independent of the set of error bars chosen). Upper-left plot: black crosses: polar regions (distance to the poles less than 50°), grey crosses: equatorial regions (distance from the poles larger than 50°). Upper-right plot: black crosses: polar north (distance to the North Pole closer than 50°), grey crosses: polar south. Lower-left plot: black crosses: equatorial region north (distance from the poles larger than 50° but inside the Northern hemisphere), grey crosses: equatorial region south. Lower-right plot: black crosses: Northern hemisphere, grey crosses: Southern hemisphere.

these discs are excluded from the analysis, the peak amplitude is somewhat higher.

6 RESULTS OF THE SPECTRA ESTIMATED ON HEMISPHERES

In previous sections we have presented evidence for a possible asymmetry in the large angular scale power present in the *WMAP* data. In order to pursue this effect further, and to allow greater spectral resolution in the low-order multipole bins, we have computed the power-spectrum amplitudes determined for the Northern and Southern hemispheres as defined in a particular coordinate system. Using the small- and medium-sized discs, the multipole resolution was limited and the asymmetry could only be studied using a relatively large multipole binning range. By estimating the power spectrum

on hemispheres the positional dependence of smaller multipole intervals can be investigated in order to check if the asymmetry is related to specific angular scales. In addition, we will study the two strong outliers at $\ell = 21$ and $\ell = 39$ found in the results of the global analysis obtained by the *WMAP* team.

6.1 The analysis

In order to check the positional dependence of the lowest multipoles we have computed the power spectra on hemispheres defined with respect to those coordinate frames for which the North Pole pierces the centre of the discs in Fig. 1, plus a further set of 34 discs which were previously rejected in the small-disc analysis since they were sufficiently compromised by the Kp2 mask. Clearly, all 164 disc positions can now be used since they are merely used to define




Figure 13. The position and numbering of the 19° discs on which the local power spectra have been calculated.




Figure 14. The full-sky *WMAP* power spectrum from a joint likelihood estimation of 34 discs of radius 19° . The histogram shows the binned *WMAP* best-fitting running-index spectrum, the shaded areas the 1σ and 2σ levels from 512 simulated maps with the same input spectrum. The crosses show the result of the joint disc estimation procedure on the *WMAP* data.

reference directions. For each hemisphere, bandpower estimates were computed for the first 64 multipoles in bins of width 3, defined by $\ell_b \in \{\ell_{\min}, \ell_{\max}\}$ where $\ell_{\min} = 3b - 1$ and $\ell_{\max} = 3b + 1$. The spectra were then compared to a Monte Carlo ensemble of hemisphere spectra estimated on simulated maps.

6.2 Results on the Northern and Southern hemispheres

In Fig. 19 we have plotted the spectra estimated for the Northern and Southern hemispheres in the Galactic and ecliptic reference frames. We have also included the two hemispheres for which the large-

scale asymmetry is found to be largest (which will be defined in the last part of this section). In each plot, the histogram shows the binned best-fitting *WMAP* spectrum (the full spectrum shown as a solid line) used in the simulations (2048 for each hemisphere), the shaded zones show the 1σ and 2σ error bars from the simulations on the given hemisphere, the grey thick crosses show the *WMAP* global estimate of the given bin and the black crosses show our estimate on the given hemisphere of the *WMAP* map. Note that the error bars shown as shaded zones are valid only for the hemisphere results. The full-sky *WMAP* results have error bars smaller by a factor of roughly $\sqrt{2}$.

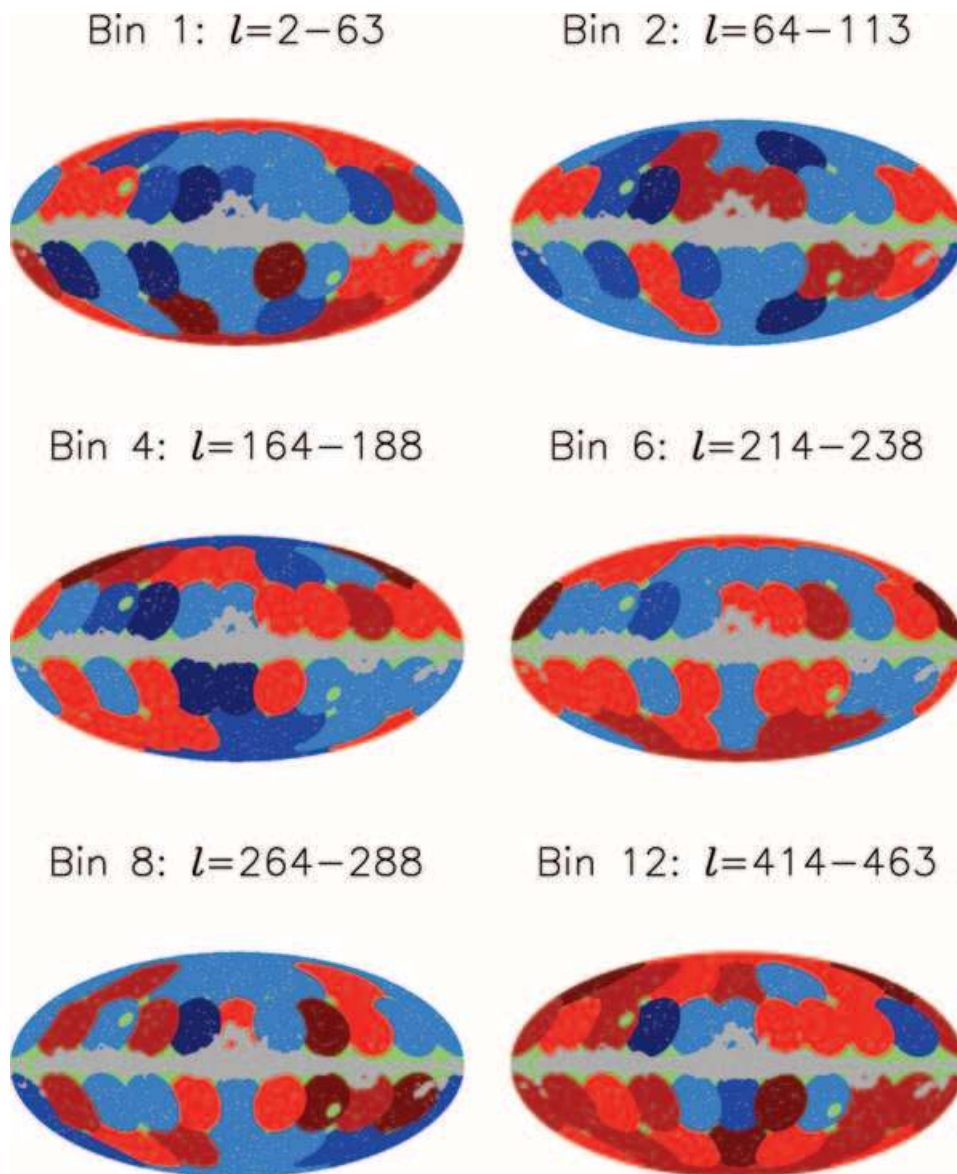


Figure 15. The local power in the bin $\ell = 2-63$ (upper-left plot), $\ell = 64-113$ (upper-right plot), $\ell = 164-188$ (middle-left plot), $\ell = 214-238$ (middle-right plot), $\ell = 264-288$ (lower-left plot) and the bin $\ell = 414-463$ (lower-right plot) estimated on 19° discs. The discs with power above the average set by simulations are red and the discs with power below the average are blue. The intensity of the colour indicates whether the power is within 1σ (light), between 1σ and 2σ (medium) and outside of 2σ (dark). The two green discs show the positions of the ecliptic poles.

On the largest angular scales, note that bin 1 ($\ell = 2-4$) which includes the quadrupole is low in the Northern Galactic hemisphere, exceptionally so in the ecliptic reference frame, yet takes values close to the best-fitting average prediction in the corresponding Southern hemispheres. Moreover, in the reference frame of maximum asymmetry the amplitude no longer seems to be anomalous.

Next we consider those bins centred at $\ell = 21$ and $\ell = 39$ which were determined to be outliers in the analysis of Hinshaw et al. (2003). It is striking that the negative outlier at $\ell = 21$ is particularly strong for the Northern hemispheres of the ecliptic and maximum asymmetry reference frames, but absent in the corresponding Southern hemispheres. In the case of the $\ell = 39$ bin, the opposite trend is observed with the outlier being associated with only the Southern ecliptic hemisphere. We suspect that these features may be related to the asymmetry suggested by the small-disc results since both fea-

tures are consistent with the observed trend of high power in the Southern hemisphere and low power in the Northern hemisphere (relative to the canonical best-fitting model).

A systematic effect could preferentially produce artefacts in the ecliptic reference frame (as a consequence of the scan pattern of the WMAP satellite), but it is difficult to envisage how an asymmetric signal could manifest itself over a range of angular scales. Indeed, the suppression of power in the Northern hemisphere would itself require a convenient alignment of structure generated by the parasitic signal with the underlying CMB anisotropy. An astrophysical origin might also be aligned with the ecliptic reference frame, but such signals would be predominantly confined to the ecliptic plane region.

Fig. 20 shows the positional dependence of the three bins described above. The discs delineate the location of the centres of the




Figure 16. The full-sky *WMAP* power spectrum from a joint likelihood estimation of 19 discs of radius 19° . The histogram shows the binned *WMAP* best-fitting running-index spectrum, the shaded areas the 1σ and 2σ levels from 512 simulated maps with the same input spectrum. The crosses show the result of the joint disc estimation procedure on the *WMAP* data.




Figure 17. The position and amplitude of the first peak. The solid line shows the best-fitting *WMAP* power spectrum. The shaded zones show the result of the full-sky power-spectrum estimation using the joint likelihood estimation procedure of 34 discs with radius 19° . The zones indicate the 1σ and 2σ spread of the peak position over 512 simulations, and the dark bold cross shows the result of the same procedure applied to the *WMAP* data. The white bold cross shows the peak position when some discs showing indications of Galactic foreground contamination are excluded. The crosses show the peak positions estimated on individual discs using the *WMAP* data, and the two circles show the corresponding 1σ and 2σ spread over 34 discs in 512 simulations. Note that individual discs may have larger 1σ and 2σ contours and are therefore not automatically outside of their individual contours even if they are outside of the circles in the plot.

hemispheres (i.e. where the North or South Pole direction vectors pierce the sphere), red indicate a bin value above the average, blue below. The values within 1σ , between 1σ and 2σ and outside of 2σ are indicated by light, medium and dark colours, respectively. From these figures one can easily see how the distribution of bins 7 ($\ell = 21$) and 13 ($\ell = 39$) gives rise to differences in the Northern and Southern ecliptic hemispheres as well as the differences between the polar and equatorial ecliptic regions. We observe that although bin 1 is lower than average for almost all hemispheres, the distribution is still asymmetric as there are particularly low values around the

Northern ecliptic hemisphere. The opposite is the case for bin 13 which is generally high, but particularly around the Southern ecliptic hemisphere. For bin 7 the bin values are distributed both above and below the average value.

Finally, we utilize the hemisphere estimates to check the positional dependence of the large-scale power over larger multipole intervals which include the three interesting ranges above. Using the above-defined 3-bins, we constructed new larger bins. Initially, we cross-check the bin $\ell = 2-63$ for which we detected a north-south asymmetry using the smaller 9.5° and 19° discs. Fig. 21 indicates the




Figure 18. The peak position (upper plot) and peak amplitude (lower plot) estimated on 19° discs. Red colour indicates that for the given disc, the estimate on *WMAP* was above the mean of simulations and blue colour shows that the *WMAP* estimate was below the mean. The significance is indicated by the intensity of the colour: light means within 1σ , medium means between 1σ and 2σ and dark means outside of 2σ . The two green dots show the position of the ecliptic poles.

ratio of the bandpower amplitude between the hemisphere centred at the point where the disc is plotted and the value in the opposite hemisphere. A similar asymmetric structure is present in this map as found for the small and medium discs (Figs 9 and 15). By considering the bins $\ell = 2-4$, $\ell = 5-16$, $\ell = 17-28$, $\ell = 29-40$ and $\ell = 41-63$ it appears that the asymmetry seen for $\ell = 2-63$ seems to be dominated by the $\ell = 5-16$ and $\ell = 29-40$ multipole ranges. For these two ranges the asymmetry pattern is almost identical. For multipoles beyond $\ell = 40$ the asymmetry seems to be gone. There is also a multipole range $\ell = 17-28$ which does not have exactly the asymmetry pattern of the other multipoles. Nevertheless, we conclude that separate multipole ranges for $\ell = 5-40$ show similar asymmetric patterns.

6.3 The significance of the maximum asymmetry

In order to avoid the introduction of bias in our interpretation of the significance of the asymmetry results, in particular due to the consideration of hemispheres defined only with respect to ‘special’ coordinate systems, we have adopted the following approach.

(i) The power spectrum for each hemisphere centred on the 164 possible orientations on the sky was computed for 2048 simulated maps.

(ii) The results were binned over the desired interval in ℓ and the ratio of bandpower r computed between a specific (Northern) hemisphere and its opposite (Southern) hemisphere.

(iii) For each simulation, the axis with the maximum ratio, r_{\max} , was determined (NB: this may or may not be the same as the axis of maximum asymmetry for the *WMAP* data themselves).

Table 3 summarizes the number of simulations with a higher maximum ratio r_{\max} than the maximum ratio in the *WMAP* data for different multipole ranges. Table 4 shows the position of the pole of maximum asymmetry for the same ℓ -intervals. We note that the position of maximum asymmetry for those bin ranges starting with $\ell = 2$ are all at the Galactic North Pole. Later, we will speculate that the lowest multipole band, $\ell = 2-4$, may contain residual Galactic contamination, although it is not clear how this would lead to the North Galactic Pole being the preferred asymmetry direction. For the remaining multipole combinations, there are intervals where only 0–5 per cent of the simulated maps show a similar asymmetry to the data, and nearly all of these share the same axis of asymmetry, specifically that for which the North Pole of the reference frame pierces ($80^\circ, 57^\circ$) in galactic coordinates.²

Fig. 22 serves to emphasize that this asymmetry cannot be associated with one particular region on the sky, but is an extended feature seen by the locally estimated spectra in most directions. In fact, this is simply Fig. 8 – showing the local large power estimates on small discs – replotted in our preferred reference frame. We see that within 60° of this North Pole, the bin $\ell = 2-63$ is particularly low (only four discs above the average, but several discs below the lower 1σ and 2σ levels) whereas around this South Pole, the bin values are particularly high (very few discs below the lower 1σ level but several above the upper 1σ and 2σ level).

Fig. 23 shows the percentage of simulated maps with a higher maximum asymmetry for each single bin. Note that there is a dip for the bins close to $\ell = 21$ and a peak close to $\ell = 39$ where there are particular features in the full-sky power spectrum as detailed above. At the first feature, the asymmetry is particularly large, whereas for the second feature the spatial distribution is particularly uniform. In Fig. 24 we show the position of the hemispheres with the 10 highest (black discs) and lowest (white discs) powers for the bins $\ell = 2-40$ (large discs), $\ell = 8-40$ (second-largest discs), $\ell = 5-16$ (second-smallest discs) and $\ell = 29-40$ (smallest discs). We see that the maximum/minimum positions are in good agreement for all these multipole ranges. We also see that the 10 maximum/minimum positions are all relatively close to each other. We conclude that the asymmetry, apart from being similarly orientated for several multipole ranges, is also highly significant for separate ranges as compared to 2048 simulations and is not associated with one particular spot on the sky.

6.4 The asymmetry in other channels and with other sky cuts

In order to constrain possible foreground residual contributions to our analysis, we have analysed the data with more aggressive Galactic cuts – the *WMAP* Kp0 mask and a $30^\circ \times 2$ (30° north of equator and 30° south of equator) Galactic cut – and also considered data at a lower frequency – the combined Q1+Q2 channels evaluated with the Kp0 cut. To save CPU time, we have only computed pseudo-power-spectrum estimates, with simple corrections for the sky coverage, beam response and noise contribution, rather than employing the full maximum-likelihood machinery.³ Tables 5 and 6

² Here, θ and ϕ are measured in the HEALPIX convention thus corresponding to *co*-latitude and longitude.

³ We first verified that this approach when applied to the co-added W+V channels rendered results consistent with the maximum-likelihood power-spectrum estimation.

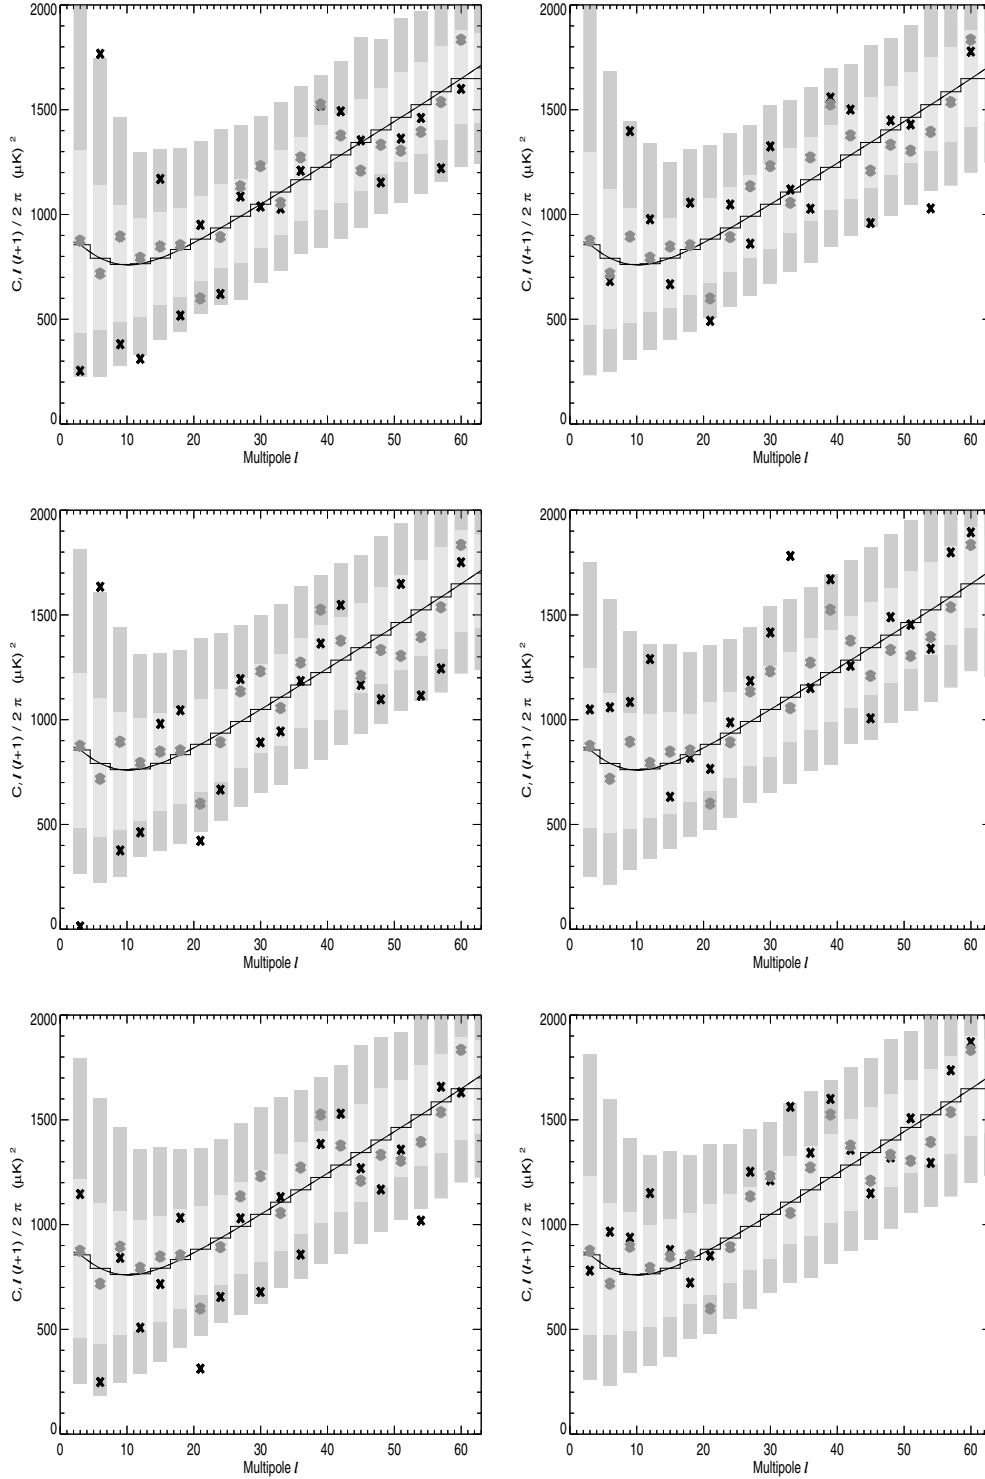


Figure 19. Results of power-spectrum estimation on hemispheres: the left column shows the northern spectra, the right column shows the southern spectra. The first row is taken in the Galactic reference frame, the second row in the ecliptic reference frame and the last row in the reference frame of maximum asymmetry, where the North Pole points at $(80^\circ, 57^\circ)$. The solid line shows the best-fitting *WMAP* spectrum, the histogram is the same spectrum binned, the shaded zones show the 1σ and 2σ error bars from simulations of the given hemisphere, grey crosses show the *WMAP* estimate on the full sky (with the Kp2 mask) and the black crosses show our estimate on the given hemisphere. Please note that the shaded zones show the error bars for a hemisphere; the error bars for the *WMAP* estimates (grey crosses) are smaller by a factor of roughly $\sqrt{2}$.

summarize the results. We have considered the Galactic reference frame and the reference frame of maximum asymmetry and found the fraction of simulations (from a total of 512 realizations) with a higher asymmetry ratio than for the *WMAP* data over specific multi-

pole ranges. We see that when excluding the lowest multipoles, the asymmetry persists for any sky cut and channel in either frame of reference. Furthermore, when considering the Kp0 cut, the strength of the asymmetry for the W+V and *Q*-band maps is identical. If




Figure 20. The local power in the bin $\ell = 2-4$ (upper plot), $\ell = 20-22$ (middle plot) and $\ell = 38-40$ (lower plot) estimated on hemispheres centred on the positions indicated by the discs. The hemispheres with power above the average set by simulations are indicated by red discs and the hemispheres with power below the average with blue discs. The significance is indicated by the intensity of the colour: light means within 1σ , medium means between 1σ and 2σ and dark means outside of 2σ . The two green discs show the positions of the ecliptic poles.

Galactic foregrounds had played an important role, one would expect some clear differences between these two frequency channels. The fact that the asymmetry is present at a significant level in several frequency channels and with different sky cuts excludes a simple explanation based on systematic effects and foregrounds.

As a useful cross-check which should help to mitigate against systematic effects as the cause of the observed structure, the same hemisphere exercise has been performed for the co-added 53+90 GHz *COBE*-DMR map, considering multipoles in the range $\ell = 5-20$ where the signal is dominant. A comparison of the results with the *WMAP* data is shown in Fig. 25. We see immediately the same asymmetric structure in both data sets even if the *COBE*-DMR map is

somewhat noisier, the significance of the asymmetry being at the $\sim 1\sigma$ level. The fact that the *COBE*-DMR data are susceptible to different parasitic signals compared to *WMAP* argues against an explanation for the results in terms of a systematic effect.

Finally, in the same figure, we have plotted the positional dependence of the power ratio for the Internal Linear Combination (ILC) map and the effective total foreground template contribution to the *WMAP* data. There is no clear correlation between the structures seen in the foregrounds and the direction of maximum asymmetry, which again argues against Galactic foregrounds as the cause of the asymmetry. Although evidence has been found for the presence of residual foregrounds in the ILC map (Naselsky et al. 2003a), the excellent agreement between maps corrected in two completely different and independent ways – either by the direct subtraction of foreground templates or by the ILC method – does not support the hypothesis that such residuals are responsible for the asymmetric structure seen in the data. Even more importantly, the lack of frequency dependence of the asymmetry seen both here and in Eriksen et al. (2004a), Hansen et al. (2004) and Vielva et al. (2004) is unlikely to be a consequence of residuals from the Galactic synchrotron, free-free or dust emission which have very strong frequency spectrum dependencies. Over the range $\ell = 2-4$ the distribution of power in the foregrounds template shows the same morphology as for the interval $\ell = 5-20$ seen in the plot. If we then compare the distribution of power for the multipoles $\ell = 2-4$ determined from the *WMAP* data (Fig. 21) with the distribution of the foregrounds template, there is a striking similarity. This strengthens the hypothesis that this bin is contaminated by foregrounds as we noted above looking at the direction of maximum asymmetry. Thus, Galactic contamination in this bin could contribute to the low value of the quadrupole, but cannot easily explain the asymmetry seen in the other multipole ranges.

We conclude that the remarkable lack of large-scale structure in the area around the ecliptic North Pole [more precisely around the point $(80^\circ, 57^\circ)$] over the multipole range from $\ell = 5$ to $\ell = 40$ cannot easily be explained by either Galactic foregrounds or systematic effects.

7 CONCLUSIONS

In this paper, we have applied a method of local power-spectrum estimation to small regions of the CMB sky as measured by the *WMAP* satellite in order to search for evidence of spatial dependence in the computed spectra beyond that due to sampling variance alone. In this way, we are able to confront one of the central tenets of modern cosmology, namely that of cosmological isotropy. As an important consequence of this analysis, we are also able to mitigate against foreground or systematic artefacts in the data. Specifically, we have estimated the angular power spectrum in different spatial directions within discs of radius 9.5° and 19° and on hemispheres. Since by increasing the disc size we increase the resolution of the power-spectrum estimation in multipole space but lower the sensitivity to the spatial dependence of the estimate, the combination of these scales allows us to achieve high fidelity in both real and ℓ -space.

On small angular scales and specifically for two power-spectrum bins, $\ell = 164-188$ and $\ell = 264-288$, close to the first Doppler peak, we found that the highest deviations, relative to the best-fitting model preferred by the *WMAP* team, was seen in the power-spectrum estimates for those medium-sized discs close to the Galactic plane. This was also the case for an interval in ℓ -space spanning the ‘dent’ in the power spectrum as estimated by the global analysis of the *WMAP* team. By excluding those discs close to the Galactic plane

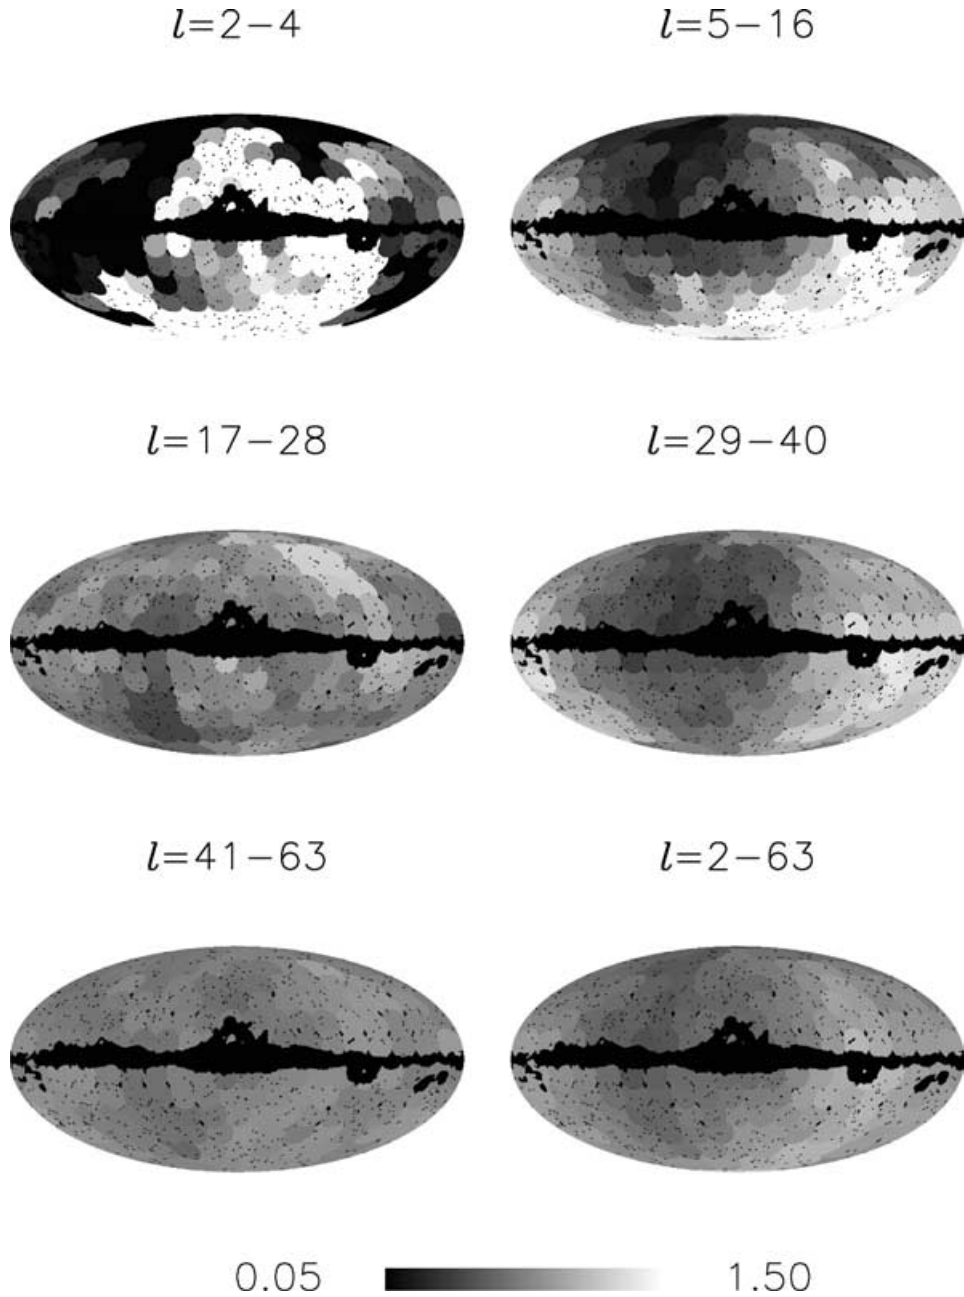


Figure 21. The discs show the centres of the hemispheres where the power spectrum has been estimated. The colours show the ratio of the power-spectrum bin $\ell = 2-4$ (upper left), $\ell = 5-16$ (upper right), $\ell = 17-28$ (middle left), $\ell = 29-40$ (middle left), $\ell = 41-63$ (lower left) and $\ell = 2-63$ (lower right) between the opposite hemispheres.

with a local spectrum deviating by more than 2σ from the best-fitting model, we find a resulting power spectrum with no dent and a slightly higher Doppler peak. We consider this to be a tentative indication of possible residual Galactic foreground contamination.

When studying the distribution of large-scale power using the small- and medium-sized discs, we found that the Northern Galactic and ecliptic hemispheres seemed to show a remarkable lack of large-scale power. In order to investigate further this apparent lack of isotropy, we estimated the power spectrum for the first 60 multipoles on hemispheres centred on various directions on the sky. The multipole range $\ell = 2-4$ was found to be strongly asymmetric with the Galactic north–south axis as the axis of maximum asymmetry.

Furthermore we found that the spatial dependence of the variations in amplitude for this multipole range was similar to the morphology seen in the integrated foreground template contribution, suggesting that the low-order multipoles may remain compromised by residual Galactic emission (this was also noted by de Oliveira-Costa et al. 2004; Eriksen et al. 2004c; Schwarz et al. 2004). The other multipoles between $\ell = 5$ and $\ell = 40$, however, are asymmetric with the axis pointing in the direction $(80^\circ, 57^\circ)$ in Galactic colatitude and longitude, close to the ecliptic axis. In the Northern hemisphere of the reference frame of maximum asymmetry, almost the full multipole range between $\ell = 5$ and $\ell = 40$ is below the average amplitude. The strong negative outlier at $\ell = 21$ found in the *WMAP* global

Table 3. The asymmetry ratio as computed in a coordinate frame selected such that the asymmetry is maximized. The numbers indicate the percentage of simulations with a higher maximum ratio than that found in the *WMAP* data for the given multipole range. The top row shows the first multipole and the left column shows the last multipole in the multipole range. Note that in this case the data value is compared against values derived from the simulations which may or may not have the same preferred axis as the data.

bin	$\ell = 2$	$\ell = 5$	$\ell = 8$	$\ell = 14$	$\ell = 20$	$\ell = 26$	$\ell = 32$	$\ell = 38$	$\ell = 44$
$\ell = 19$	0.2	7.6	20.3	23.3					
$\ell = 25$	0.1	12.5	24.5	74.9	24.9				
$\ell = 31$	0.1	5.5	16.8	24.8	7.0	15.4			
$\ell = 37$	0.2	0.7	4.7	2.2	0.5	1.8	1.7		
$\ell = 40$	0.3	0.7	4.7	9.0	0.9	13.7	35.1	90.9	
$\ell = 43$	0.5	2.4	10.5	23.9	6.2	23.4	41.9	52.4	
$\ell = 49$	0.5	3.4	5.6	14.5	2.5	8.5	20.7	36.7	90.3
$\ell = 55$	1.9	1.4	5.9	5.9	2.8	2.3	4.7	14.0	16.8
$\ell = 61$	0.3	1.1	5.5	13.4	4.6	23.6	40.7	77.7	75.2

Table 4. Position of the axis of maximum asymmetry for the multipole ranges given in Table 3. The top row shows the first multipole and the left column shows the last multipole in the multipole range. The given position is the centre of the hemisphere with most power of the pair of opposite hemispheres. The angles (θ, ϕ) are given in degrees in Galactic coordinates, using colatitude as in the HEALPIX software package, i.e. θ goes from 0 to 180° going from the North to the South Pole.

bin	$\ell = 2$	$\ell = 5$	$\ell = 8$	$\ell = 14$	$\ell = 20$	$\ell = 26$	$\ell = 32$	$\ell = 38$	$\ell = 44$
$\ell = 19$	(0,0)	(14,72)	(28,180)	(100,331.6)					
$\ell = 25$	(0,0)	(14,72)	(54,180)	(54,23)	(80,57)				
$\ell = 31$	(0,0)	(80,57)	(80,76)	(80,38)	(80,57)	(100,28)			
$\ell = 37$	(0,0)	(80,57)	(80,57)	(80,38)	(80,57)	(80,38)	(67,6)		
$\ell = 40$	(0,0)	(80,57)	(80,57)	(80,57)	(80,57)	(80,57)	(67,346)	(126,203)	
$\ell = 43$	(0,0)	(80,57)	(80,57)	(80,57)	(80,57)	(100,85)	(152,72)	(113,86)	
$\ell = 49$	(0,0)	(80,57)	(100,85)	(100,85)	(100,85)	(100,85)	(152,72)	(113,86)	(126,113)
$\ell = 55$	(0,0)	(80,57)	(80,57)	(139,47)	(152,72)	(139,47)	(152,72)	(152,72)	(126,46)
$\ell = 61$	(0,0)	(80,57)	(80,57)	(80,57)	(80,57)	(80,57)	(152,72)	(152,72)	(54,0)

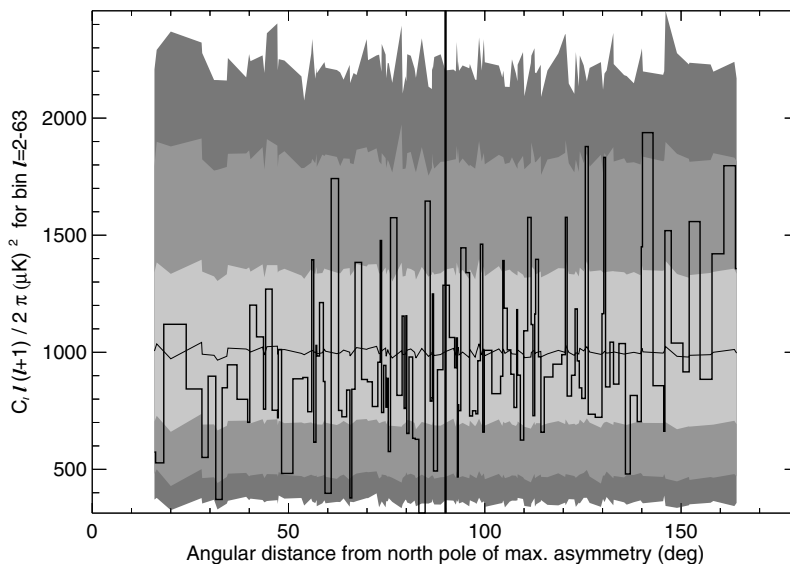


Figure 22. The disc estimates of the bin $\ell = 2-63$. The discs are ordered such that the values to the left are close to the North Pole of the axis ($80^\circ, 57^\circ$) of maximum asymmetry and the values to the right are close to the corresponding South Pole. The angular distance from this North Pole is shown. The shaded zones indicate the 1σ , 2σ and 3σ spread of the estimated bin for the given disc as found from Monte Carlo simulations.

estimate of the power spectrum is associated specifically with this hemisphere. Conversely, in the Southern hemisphere of maximum asymmetry, almost all multipoles in the range $\ell = 5-40$ are above the average amplitude, and a positive outlier at $\ell = 39$, as seen

in the *WMAP* global spectrum, is found. The asymmetry does not appear to be associated with a particular anomalous region on the sky, but extends over a large area, as evidenced by the small- and medium-disc results.




Figure 23. The percentage of maps with a higher asymmetry than in the *WMAP* data for a single multipole bin.

We have checked the possibility that incorrectly removed foregrounds could cause the observed large-scale asymmetry by testing the dependency on Galactic cut and frequency channel. The observed asymmetry is remarkably stable with respect to both frequency and sky coverage, thus arguing against this possibility. As a further check, we have verified that the ILC map, which is constructed without the use of foreground templates, demonstrates the same level of asymmetry. Although Naselsky et al. (2003a) have found evidence for residual foregrounds at large scales in the ILC map, it would require a particular conspiracy of the Galactic synchrotron, free-free and dust residual emission to generate identical morphologies in both this and the observations over several frequencies. Of course, one cannot completely rule out an unknown foreground component with a flat CMB-like frequency spectrum as the cause of the asymmetry, but no convincing candidates have been proposed.

In terms of a systematic error origin for the asymmetry, we make the following observations. Since the scanning strategy of the experiment is related to the ecliptic frame of reference, one might propose that a systematic effect could manifest itself in that frame, which is close to that for maximum asymmetry. Nevertheless, a more detailed consideration indicates that this will rather depend on the nature of the parasitic signal – its time period, how it couples to each detector,

Table 5. The percentage of simulations with a higher ratio of the pseudo-spectra between the Northern and Southern Galactic hemisphere for the given multipole, frequency channel and sky cut.

channel:	W+V Kp2	W+V Kp0	W+V $30^\circ \times 2$	Q Kp0
$\ell = 2-19$	0.7	1.4	3.1	1.4
$\ell = 2-40$	0.9	1.4	3.5	1.4
$\ell = 2-61$	1.0	1.4	3.9	1.4
$\ell = 5-40$	0.0	0.4	2.1	0.4
$\ell = 8-19$	0.0	0.0	0.4	0.0
$\ell = 8-40$	0.0	0.0	0.4	0.0
$\ell = 8-61$	0.0	0.0	0.4	0.0
$\ell = 20-40$	0.0	0.4	4.5	0.4
$\ell = 20-61$	0.0	0.2	6.4	0.2

Table 6. The percentage of simulations with a higher ratio of the pseudo-spectra between the maximally asymmetric hemispheres [with the axis piercing ($80^\circ, 57^\circ$)] for the given multipole, frequency channel and sky cut.

channel:	W+V Kp2	W+V Kp0	W+V $30^\circ \times 2$	Q Kp0
$\ell = 2-19$	0.6	1.2	1.9	1.4
$\ell = 2-40$	0.7	1.4	0.0	1.4
$\ell = 2-61$	0.8	1.4	3.1	1.4
$\ell = 5-40$	0.0	0.0	0.0	0.0
$\ell = 8-19$	0.0	0.0	0.0	0.0
$\ell = 8-40$	0.0	0.0	0.0	0.0
$\ell = 8-61$	0.0	0.0	0.0	0.0
$\ell = 20-39$	0.0	0.0	3.3	0.0
$\ell = 20-61$	0.0	0.0	5.1	0.0

to the motion of the satellite and, indeed, to the differential nature of the detectors. Even an apparently quasi-periodic signal in the time-ordered observations may be smeared over the celestial sphere in many configurations in the final reconstituted sky map. In particular, it is difficult to imagine a systematic signal which produces a different effect in the Northern or Southern hemispheres, rather than a difference between the (ecliptic) poles and the equator, and which also demonstrates no frequency dependence. Moreover, the results seem unlikely to be compromised by systematics since we find supporting evidence (albeit at reduced statistical significance) for similar features in the *COBE*-DMR maps, which are susceptible to *different* parasitic signals.




Figure 24. The discs show the positions of the hemispheres with the 10 highest (black discs) and 10 lowest (white discs) bin values. The power-spectrum bins considered were $\ell = 2-40$ (large discs), $\ell = 8-40$ (second-largest discs), $\ell = 5-16$ (second-smallest discs) and $\ell = 29-40$ (smallest discs).

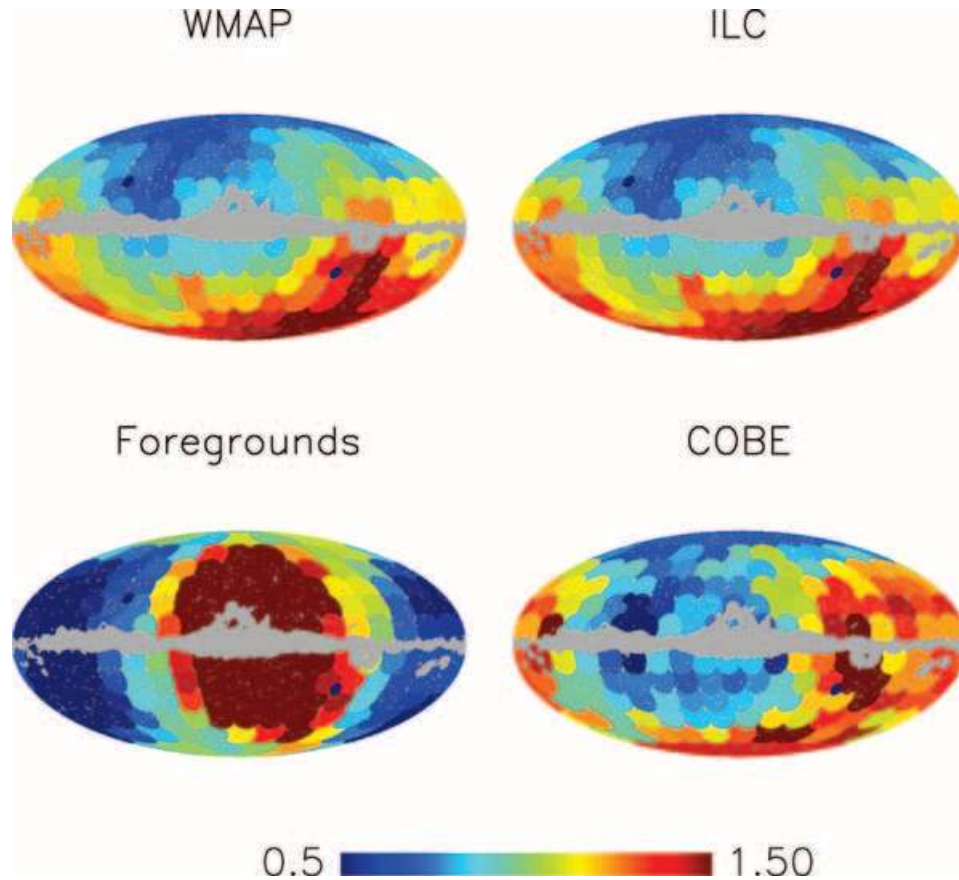


Figure 25. The discs show the centres of the hemisphere on which the power spectrum has been estimated. The colour in each disc represents the ratio between the power in the given hemisphere and the opposite hemisphere in the multipole range $\ell = 5-20$. The upper row shows the *WMAP* W+V channel and the ILC map, the lower row shows the *COBE* data and the foreground template used for *WMAP*.

Similar asymmetric structures have been determined by a number of other groups (Copi et al. 2004; Park 2004; Vielva et al. 2004; Cruz et al. 2004; Eriksen et al. 2004a,b; Hansen et al. 2004; Land & Magueijo 2004; Larson & Wandelt 2004) using several alternative analysis techniques. Besides finding asymmetry, non-Gaussian features have also been detected in the Northern and/or Southern hemispheres. Since some of the non-Gaussian estimators used are power-spectrum dependent, the detection of non-Gaussianity on opposite hemispheres may be related to the uneven distribution of large-scale power studied in this paper. Given the large number of detections with different methods on different sky cuts and frequency channels, the possibility that the *WMAP* data do indeed contain unexpected properties on large scales cannot be excluded. Although there remains evidence for Galactic residuals in the *WMAP* data (Naselsky et al. 2003a,b), there are no clear indications that such residuals are the originators of the asymmetry. Having eliminated systematic errors as the source of the signal, the intriguing possibility is raised that the cosmological principle of isotropy is violated and/or that fundamentally new physics on large scales in the Universe is required. Further clarification of this scenario awaits further observations from *WMAP*, and ultimately the forthcoming *Planck* satellite mission.

ACKNOWLEDGMENTS

We are grateful to Gary Hinshaw, Hans Kristian Eriksen and Domenico Marinucci for useful discussions. FKH acknowledges

financial support from the CMBNET Research Training Network. We acknowledge use of the HEALPIX (Górski, Hivon & Wandelt 1998) software and analysis package for deriving the results in this paper. This research used resources of the National Energy Research Scientific Computing Centre, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC03-76SF00098. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

REFERENCES

- Bennett C. L. et al., 1996, *ApJ*, 464, L1
- Bennett C. L. et al., 2003, *ApJS*, 148, 1
- Bond J. R., Jaffe A. H., Knox L., 2000, *ApJ*, 533, 19
- Cabella P., Hansen F. K., Marinucci D., Pagano D., Vittorio N., 2004, *Phys. Rev. D*, 69, 063007
- Chiang L., Naselsky P. D., Verkhodanov O. V., Way M. J., 2003, *ApJ*, 590, L65
- Coles P., Dineen P., Earl J., Wright D., 2004, *MNRAS*, 350, 983
- Copi C. J., Huterer D., Starkman G. D., 2004, *Phys. Rev. D*, 70, 043515
- Cruz M., Martinez-Gonzalez E., Vielva P., Cayon L., 2004, preprint (astro-ph/0405341)
- de Oliveria-Costa A., Tegmark M., Zaldariagga M., Hamilton A., 2004, *Phys. Rev. D*, 69, 063516
- Eriksen H. K., Hansen F. K., Banday A. J., Górski K. M., Lilje P. B., 2004a, *ApJ*, 605, 14

Eriksen H. K., Novikov D. I., Lilje P. B., Banday A. J., Górski K. M., 2004b, *ApJ*, 612, 64
 Eriksen H. K., Banday A. J., Górski K. M., Lilje P. B., 2004c, *ApJ*, 612, 633
 Górski K. M., Hivon E., Wandelt B. D., 1998, in Banday A. J., Sheth R. K., Da Costa L., eds, *Analysis Issues for Large CMB Data Sets*. ESO, Print-partners Ipskamp, NL, p. 37 (astro-ph/9812350); HEALPIX HOMEPAGE: <http://www.eso.org/science/healpix/>
 Hansen F. K., Górski K. M., 2003, *MNRAS*, 343, 559
 Hansen F. K., Górski K. M., Hivon E., 2002, *MNRAS*, 336, 1304
 Hansen F. K., Cabella P., Marinucci D., Vittorio N., 2004, *ApJ*, 607, L67
 Hinshaw G. et al., 2003, *ApJS*, 148, 135
 Komatsu E. et al., 2003, *ApJS*, 148, 119
 Land K., Magueijo J., 2004, preprint (astro-ph/0405519)
 Larson D. L., Wandelt B. D., 2004, preprint (astro-ph/0404037)
 Magueijo J., Medeiros J., 2004, *MNRAS*, 351, L1

Naselsky P. D., Doroshkevich A. G., Verkhodanov O. V., 2003a, *ApJ*, 599, L53
 Naselsky P. D., Verkhodanov O. V., Chiang L. Y., Novikov I. D., 2003b, preprint (astro-ph/0310235)
 Naselsky P. D., Doroshkevich A. G., Verkhodanov O. V., 2004a, *MNRAS*, 349, 695
 Naselsky P. D., Verkhodanov O. V., Chiang L. Y., Novikov I. D., 2004b, preprint (astro-ph/0405523)
 Page L. et al., 2003, *ApJS*, 148, 233
 Park C. G., 2004, *MNRAS*, 349, 313
 Schwarz D. J., Starkmann G. D., Huterer D., Copi C. J., 2004, CERN-PH-TH/2004-052
 Smoot G. F. et al., 1992, *ApJ*, 396, L1
 Vielva P., Martinez-Gonzalez E., Barreiro R. B., Sanz J. L., Cayon L., 2004, *ApJ*, 609, 22

APPENDIX A: MONTE CARLO PRE-CALCULATIONS

For a general Gabor window G_i as a function of pixel number i , one can write the pseudo-spherical harmonic coefficients as

$$\tilde{a}_{\ell m} = \sum_{\ell' m'} a_{\ell' m'} h(\ell, \ell', m, m') \quad (\text{A1})$$

where

$$h(\ell, \ell', m, m') = \sum_i G_i Y_{\ell m}^i Y_{\ell' m'}^i. \quad (\text{A2})$$

To use the likelihood Ansatz on the pseudo-spectrum coefficients, we need to sample the observed spectrum at N^{in} multipoles ℓ_i . How dense the sampling is, depends on the size of the window (Hansen et al. 2002). Also, we want to estimate the underlying full-sky spectrum in N^{bin} flat bins b .

For calculating the likelihood, one needs the theoretical expectation value $C_i \equiv \langle \tilde{C}_{\ell_i} \rangle$ and the correlation matrix $C_{ij} \equiv \langle \tilde{C}_{\ell_i} \tilde{C}_{\ell_j} \rangle - \langle \tilde{C}_{\ell_i} \rangle \langle \tilde{C}_{\ell_j} \rangle$. These can be expressed, using the above formulae as

$$C_i = \sum_b C_b \frac{1}{2\ell_i + 1} \underbrace{\sum_{\ell' \in b} \frac{B_{\ell'}}{\ell'(\ell' + 1)} \sum_{mm'} h^2(\ell_i, \ell', m, m')}_{K(i,b)} \quad (\text{A3})$$

and

$$C_{ij} = \sum_{bb'} C_b C_{b'} \frac{2}{(2\ell_i + 1)(2\ell_j + 1)} \underbrace{\sum_{\ell' \in b} \sum_{L' \in b'} \sum_{mM} \sum_{m'M'} h(\ell_i, \ell', m, m') h(\ell_j, \ell', M, m') h(\ell_i, L', m, M') h(\ell_j, L', m', M')}_{K'(i,j,b,b')} \quad (\text{A4})$$

For a general Gabor window, the factors $K(i, b)$ and $K(i, j, b, b')$ are very heavy-going to calculate using analytic formulae. We propose to calculate these factors using a Monte Carlo approach.

The Monte Carlo approach to the calculation of these factors can be done as follows;

(i) Generate a set of independent random Gaussian variables $A_{\ell m}$ with variance

$$\langle A_{\ell m} A_{\ell' m'} \rangle = \delta_{\ell \ell'} \delta_{m m'} \frac{B_{\ell}}{\ell(\ell + 1)}. \quad (\text{A5})$$

(ii) For each bin b , calculate a set of skies T_i^b as a function of pixel number i defined as

$$T_i^b = \sum_{\ell' \in b} \sum_{m'} A_{\ell' m'} Y_{\ell' m'}^i. \quad (\text{A6})$$

(iii) For each sky, calculate the pseudo-spectrum $\tilde{a}_{\ell m}^b$,

$$\tilde{a}_{\ell m}^b = \sum_i T_i^b G_i Y_{\ell m}^i \quad (\text{A7})$$

(iv) Form the pseudo-cross-bin spectra $C_i^{bb'}$ defined as

$$C_i^{bb'} = \sum_m \frac{\tilde{a}_{\ell_i m}^b \tilde{a}_{\ell_i m}^{b'}}{2\ell_i + 1} \quad (\text{A8})$$

and the covariance matrices $C_{ij}^{bb'}$

$$C_{ij}^{bb'} = (C_i^{bb'} - \langle C_i^{bb'} \rangle)(C_j^{bb'} - \langle C_j^{bb'} \rangle). \quad (\text{A9})$$

Using the formulae above, one sees that $\langle C_i^{bb} \rangle = K(i, b)$ and $2\langle C_{ij}^{bb'} \rangle = K'(b, b', i, j)$ for $b \neq b'$ and $\langle C_{ij}^{bb} \rangle = K'(b, b, i, j)$. Using about 1000 simulations, we get an accuracy of the C_ℓ estimate that is within 1 per cent of the estimate found using the analytical expression.

When noise is present, the correlation matrix C_{ij} can be written as

$$C_{ij} = C_{ij}^S + C_{ij}^X + C_{ij}^N, \quad (\text{A10})$$

where C_{ij}^S is the signal correlation matrix calculated above, $C_{ij}^N = \langle \tilde{C}_{\ell_i}^N \tilde{C}_{\ell_j}^N \rangle - \langle \tilde{C}_{\ell_i}^N \rangle \langle \tilde{C}_{\ell_j}^N \rangle$ which can be calculated analytically for white noise and by Monte Carlo of noise realizations for correlated noise. Finally there is a mixed signal–noise term which can be written as

$$C_{ij}^X = \sum_b C_b \underbrace{\frac{4}{(2\ell_i + 1)(2\ell_j + 1)} \sum_{\ell' \in b} \sum_{mM} \sum_{m'} h(\ell_i, \ell', m, m') h(\ell_j, \ell', M, m') h'(\ell_i, \ell_j, m, M)}_{K''(i, j, b)}, \quad (\text{A11})$$

where $h'(\ell, \ell', m, m') = \sum_i \sigma_i^2 G_i^2 Y_{\ell m}^i Y_{\ell' m'}^i$ in the case of white noise for a given noise pixel variance σ_i^2 . This factor is also heavy-going to calculate for a general window and for general noise. It can, however, be calculated using a similar Monte Carlo approach:

(i) Generate a set of independent random Gaussian variables $A_{\ell m}$ with variance

$$\langle A_{\ell m} A_{\ell' m'} \rangle = \delta_{\ell \ell'} \delta_{m m'} \frac{B_\ell}{\ell(\ell + 1)}. \quad (\text{A12})$$

(ii) For each bin b , calculate a set of skies T_i^b as a function of pixel number i defined as

$$T_i^b = \sum_{\ell' \in b} \sum_{m'} A_{\ell' m'} Y_{\ell' m'}^i. \quad (\text{A13})$$

(iii) Generate a noise realization T_i^N .

(iv) For each sky, calculate the signal pseudo-spectrum $\tilde{a}_{\ell m}^b$,

$$\tilde{a}_{\ell m}^b = \sum_i T_i^b G_i Y_{\ell m}^i \quad (\text{A14})$$

and the noise pseudo-spectrum

$$\tilde{a}_{\ell m}^N = \sum_i T_i^N G_i Y_{\ell m}^i. \quad (\text{A15})$$

(v) Form the signal–noise pseudo-cross-bin spectra C_i^{NSb} defined as

$$C_i^{NSb} = \sum_m \frac{\tilde{a}_{\ell_i m}^b \tilde{a}_{\ell_i m}^N}{2\ell_i + 1} \quad (\text{A16})$$

and the covariance matrices

$$C_{ij}^{NSb} = C_i^{NSb} C_j^{NSb}. \quad (\text{A17})$$

Using the formulae above, one can show that $4\langle C_{ij}^{NSb} \rangle = K''(i, j, b)$. The noise part converges more slowly than the signal part, and for that reason we suggest calculating five noise realizations for each signal realization so that a total of 5000 noise realizations are used.

Finally, we note that the estimates of the lowest multipoles are biased (Bond, Jaffe & Knox 2000). This comes from the fact that we are estimating the variance of the likelihood. This normally introduces a bias of the form $1/\sqrt{\sigma^2 + 1}$ which in our case is $2\ell/(2\ell + 1)$. Correcting with this factor, the bias vanishes.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.