

Poincaré dodecahedral space parameter estimates

B. F. Roukema¹, Z. Buliński¹, and N. E. Gaudin²

¹ Toruń Centre for Astronomy, Nicolaus Copernicus University, ul. Gagarina 11, 87-100 Toruń, Poland
e-mail: boud@astro.uni.torun.pl

² École nationale supérieure de physique de Strasbourg, Université Louis Pasteur, Bd. Sébastien Brant, BP 10413,
67412 Illkirch Cedex, France

Received 26 July 2008 / Accepted 25 October 2008

ABSTRACT

Context. Several studies have proposed that the preferred model of the comoving spatial 3-hypersurface of the Universe may be a Poincaré dodecahedral space (PDS) rather than a simply connected, infinite, flat space.

Aims. Here, we aim to improve the surface of last scattering (SLS) optimal cross-correlation method and apply this to observational data and simulations.

Methods. For a given “generalised” PDS orientation, we analytically derive the formulae required to exclude points on the sky that cannot be members of close SLS-SLS cross-pairs. These enable more efficient pair selection without sacrificing the uniformity of the underlying selection process. For a sufficiently small matched circle size α and a fixed number of randomly placed points selected for a cross-correlation estimate, the calculation time is decreased and the number of pairs per separation bin is increased. Using this faster method, and including the smallest separation bin when testing correlations, (i) we recalculate Monte Carlo Markov Chains (MCMC) on the five-year Wilkinson Microwave Anisotropy Probe (WMAP) data; and (ii) we seek PDS solutions in a small number of Gaussian random fluctuation (GRF) simulations in order to further explore the statistical significance of the PDS hypothesis.

Results. For $5^\circ < \alpha < 60^\circ$, a calculation speed-up of 3–10 is obtained. (i) The best estimates of the PDS parameters for the five-year WMAP data are similar to those for the three-year data; (ii) comparison of the optimal solutions found by the MCMC chains in the observational map to those found in the simulated maps yields a slightly stronger rejection of the simply connected model using α rather than the twist angle ϕ . The best estimate of α implies that, given a large-scale auto-correlation as weak as that observed, the PDS-like cross-correlation signal in the WMAP data is expected with a probability of less than about 10%. The expected distribution of ϕ from the GRF simulations is not uniform on $[-\pi, \pi]$.

Conclusions. Using this faster algorithm, we find that the previous PDS parameter estimates are stable to the update to five-year WMAP data. Moreover, for an infinite, flat, cosmic concordance model with Gaussian random fluctuations, the chance of finding both (a) a large-scale auto-correlation as weak as observed; and (b) a PDS-like signal similar to what is observed is less than about 0.015% to 1.25%.

Key words. cosmology: cosmological parameters – cosmology: observations – cosmology: theory

1. Introduction

During the past half-decade, attention has focussed on the Poincaré Dodecahedral Space (PDS) as a potentially better model of comoving space than the infinite flat model (Luminet et al. 2003; Roukema et al. 2004; Aurich et al. 2005a,b; Gundermann 2005; Key et al. 2007; Niarchou & Jaffe 2007; Caillerie et al. 2007; Lew & Roukema 2008; Roukema et al. 2008). If the hypothesis that the comoving spatial section of the Universe is a PDS is correct, then it should be possible to estimate the astronomical coordinates of the fundamental domain, as has been tried by Roukema et al. (2008, hereafter, RBSG08). Moreover, successive improvements in data and analysis methods should yield successively more precise and more accurate estimates of these astronomical coordinates.

Here, we aim to improve on the method of optimising the cross-correlation ξ_C (Eq. (1), RBSG08) of cosmic microwave background (CMB) temperature fluctuations between copies of the surface of last scattering (SLS) presented in Roukema et al. (2008) and also applied by Aurich (2008). We investigate whether or not this can lead to improved parameter estimates. We also test a small number of simulations to see whether the infinite flat model can reproduce the observational signal.

In RBSG08, the method was introduced by pointing out that algebraically, it is an extension of using the identified circles principle first published by Cornish et al. (1998). Here, we make use of the fact that the cross-correlation method is not only an algebraic extension of the identified circles principle, but its relation to the identified circles principle can also be interpreted in terms of identified *annuli*. This leads to algebraic/trigonometric relations that enable a faster calculation of cross-correlations, without sacrificing the uniformity of the underlying random selection of points on the sky. Moreover, use of these relations also increases the numbers of pairs per separation bin, leading to less relative Poisson noise in individual cross-correlation estimates. The annulus interpretation and the associated relations are presented in Sect. 3.1.

In Sect. 3.2, use of this faster method to recalculate Monte Carlo Markov Chains (MCMC) for the five-year release of the WMAP data is described. Since the new method increases the numbers of pairs per separation bin, the smallest separation bin available when estimating correlations should not be as noisy as without this new method, so we include this smallest separation bin.

Use of this faster method also makes it practical to perform a consistent analysis of both the WMAP map and a small number

of Gaussian random fluctuation (GRF) WMAP simulated maps, in order to estimate the expectation of a PDS-like signal under the assumption of a simply connected model. In order to model this, we need to consider the following.

One of the key motivations from WMAP data for studying multiply connected models has been the lack of structure on angular scales larger than about 60° , as has been noted by several authors (e.g. Spergel et al. 2003; Luminet et al. 2003; Aurich et al. 2005b). As was noted in RBSG08, in particular in Eq. (10), if the PDS model is incorrect, then cross-correlations of temperature fluctuations at pairs of points that are implied (by the PDS hypothesis) to be spatially close should, on average, be much smaller than the auto-correlations at the same small spatial separation scale. This is because if the PDS model is wrong, then these pairs are in reality widely separated rather than closely separated, so their correlations should, on average, be small.

However, this is only a statistical statement: even for the case that the PDS model is physically wrong, it is possible that some orientations of a PDS model, by chance, happen to give cross-correlations that violate Eq. (10) of RBSG08. How probable is it for chance large-scale correlations in particular directions to mimic a PDS-like signal? This clearly depends on the amplitude of the large length-scale auto-correlation. If this amplitude is low or high, then the chance of finding PDS-like signals using the optimal cross-correlation method should be low or high respectively. Since the range of matched circle sizes α studied in RBSG08 included circles with $\alpha \leq 60^\circ$, most of the cross-correlations are from pairs of points separated observationally by angles $\geq 60^\circ \sim 1$ rad.

The latter is the minimum scale above which Spergel et al. (2003) quantified the lack of temperature-temperature auto-correlations at large scales in the WMAP data, with their parameter S , defined

$$S = \int_{-1}^{1/2} [C(\theta)]^2 d\cos\theta, \quad (1)$$

where $C(\theta)$ is the two-point auto-correlation function at an angular separation of θ on the SLS (Sect. 7, Eq. (9), Spergel et al. 2003). Hereafter, we call this S_ξ , in order to distinguish it from the S parameter used for matched circle analyses. The two parameters have very different meanings.

The chance of observing S_ξ to be as small as that observed was estimated by Spergel et al. (2003) as 0.15% for an infinite, flat, cosmic concordance model with a fixed spectral index of density perturbations. Efstathiou (2004), using what should be a more accurate method if we assume a simply connected model, estimated the chances to be much higher, from 3% to 12.5% depending on which sky map is analysed and which (if any) galactic cut mask is used (Table 5, Efstathiou 2004).

We can quantify the relation between large length-scale auto-correlations and the chance of a PDS signal in the case that the PDS is wrong by writing that *the lower the observed value of S_ξ , the less likely it is that a non-PDS model will give a PDS-like signal*. Hence, since we expect that the two properties, a low S_ξ value and a low chance of a PDS-like cross-correlation signal occurring in a flat, infinite, cosmic concordance model, are related, it would not be useful to estimate their probabilities independently. Instead, what is of interest to investigate is the chance that a PDS-like cross-correlation signal occurs, *given that S_ξ is as low as that which is observed*. Conservatively, we can calculate an upper limit to this conditional probability if we use simulated sky maps with S_ξ values a little higher than the observed estimate of S_ξ . Moreover, here we use the kp2 galactic cut mask,

for which Efstathiou (2004) finds the highest probability for the observed S_ξ value to occur in an infinite, flat, cosmic concordance model, so we should obtain a conservative upper limit to the probability of both (a) an S_ξ value as low as that observed; and (b) a PDS-like cross-correlation signal similar to that observed occurring. This is described in more detail in Sect. 3.3.

Results are presented in Sect. 4, discussed in Sect. 5 and conclusions are given in Sect. 6. For general background on spherical, multiply connected spaces, see Weeks (2001), Gausmann et al. (2001), Lehoucq et al. (2002) and Riazuelo et al. (2004). See the references cited above for background on cosmic topology. Comoving coordinates are used when discussing distances (i.e. “proper distances” at the present epoch, Weinberg 1972, equivalent to “conformal time” if $c = 1$) and the Hubble constant is written $H_0 \equiv 100 \text{ h km s}^{-1} \text{ Mpc}^{-1}$.

2. Observations and simulations

2.1. Observations

The analysis described in Sect. 3.2 uses the Internal Linear Combination (ILC)¹ all-sky map of the five-year WMAP data (Hinshaw et al. 2009) and the foreground cleaned, Wiener filtered version of the same five-year data published by the Tegmark et al. (2003) group (TOH)² The “kp2” mask to eliminate the Galactic Plane and other likely contaminating regions, covering about 15% of the sky, is used throughout this paper unless otherwise noted³. We do not smooth these maps.

2.2. Simulations and associated observational map

The analysis described in Sect. 3.3 uses a version of the WMAP map and sky simulations prepared consistently from three-year observational and simulated maps in the Q , V and W frequency bands, weighted by inverse noise as given in Hinshaw et al. (2003), and smoothed by a Gaussian of $FWHM$ 1° , as described in Sect. 2 of Lew & Roukema (2008)⁴. This version of the observational map is referred to hereafter as “INC3”⁵.

As mentioned above, the expected value of S_ξ from the infinite, flat, cosmic concordance model is higher than that observed. However, there is a (small) chance in the model that such a low S_ξ value occurs in any single realisation of the model, such as the one in which we live. Since we wish to test the chance that a PDS-like cross-correlation signal occurs *given that S_ξ is as low as that which is observed*, we generate simulations using the observational estimates of the spherical harmonic amplitudes C_l of the temperature fluctuations as estimated in Hinshaw et al. (2007), rather than using the mean values implied by an infinite flat model. The phases of the spherical harmonics are randomised. Gaussian noise is simulated according to the properties and scanning strategy of each differencing assembly and added to each simulated map.

¹ http://lambda.gsfc.nasa.gov/data/map/dr3/dfp/wmap_ilmc_5yr_v3.fits

² http://space.mit.edu/home/tegmark/wmap/wiener5yr_map.fits

³ Data file: http://lambda.gsfc.nasa.gov/data/map/dr2/ancillary/wmap_kp2_r9_mask_3yr_v2.fits; map projection: http://lambda.gsfc.nasa.gov/product/map/current/map_images/f02_int_mask_b.png

⁴ The maps can be downloaded from <http://cosmo.torun.pl/GPLdownload/MCMC/sims-from-LR08-project/>.

⁵ Inverse Noise Coadded 3-year.

Since a lot of the large-scale power in what are considered as the best estimates of the cosmological signal in the WMAP data lie close to the Galactic Plane (e.g. Tegmark et al. 2003), simulations based on spherical harmonics with the same C_l 's as this signal but different phases will most of the time have high large-scale power which does not happen to lie inside the kp2 cut. Hence, simulations made with the Hinshaw et al. (2007) C_l estimates and masked with the kp2 mask will, in general, have S_ξ estimates *larger* than that which is actually observed for the part of the signal outside of the kp2 cut rather than approximately equal to it.

This is not due to an error in the Hinshaw et al. (2007) estimates (assuming a simply connected model⁶). Even though these estimates for $l \leq 30$ are made using nearly the same kp2 cut that we use here (Hinshaw et al. 2007, use a kp2 cut degraded in resolution), the spherical harmonics are nevertheless functions of the full sky. Hinshaw et al. (2007) use a maximum likelihood method which uses the information from the cut sky to infer the best estimate of the C_l values for these functions covering the full sky. Unsurprisingly, it implicitly extrapolates from the cut sky to the full sky, and recovers some of the information masked in the Galactic Plane.

Since we wish to test the chance that a PDS-like cross-correlation signal occurs given that S_ξ is as low as that which is observed, we have two obvious approaches to choose from in order to have simulations which are statistically comparable with the observations. One approach would be to analyse the full sky for both observations and simulations. In this case, the S_ξ estimates should be approximately equal. However, a large part of the signal would be that from inside the Galactic Plane. The risk of implicitly testing the properties of foreground contaminants rather than of the cosmological signal would be high.

The other approach is to analyse the cut sky (with the kp2 mask) for both observations and simulations. This decreases the risk that our analysis will be affected by contamination from the Galactic Plane. However, for the cut sky, many of the simulations will have S_ξ higher than S_ξ of the observations. In this case, it will be necessary to select those simulations with the lowest values of S_ξ not greater than S_ξ of the observational map.

This will not give an exact statistical match between simulations and observations. However, since greater S_ξ should increase the chance of cross-correlations occurring, this should give an upper limit for estimating the probability of cross-correlations occurring in a simply connected model, i.e. a conservative estimate. This is the approach we adopt here.

We estimate S_ξ for each of the observational and simulation using

$$S_\xi \equiv \int_{r_{\text{SLS}}}^{2r_{\text{SLS}}} [\xi_A(r)]^2 d\frac{r}{r_{\text{SLS}}}, \quad (2)$$

where r_{SLS} is the radius of the SLS, $\xi_A(r)$ is the (comoving) spatial auto-correlation function as given in Eq. (4) of RBSG08, and r is the comoving spatial separation of a pair of points on the SLS. This integral differs from that in Eq. (9) of Spergel et al. (2003), since here we focus on spatial separation, while Spergel et al. (2003) use an orthogonal projection of the angle between a pair of sky positions. The minimum angular separation used by the latter is $60^\circ \approx 1$ rad, so here we use r_{SLS} for the minimum separation. Provided that we estimate S_ξ in the same way for both observational data and simulations, the difference between

spatial and projected angular definitions of S_ξ should not affect the questions of interest here.

3. Method

The method of using Markov Chain Monte Carlo simulations to optimise the cross-correlation ξ_C of temperature fluctuations between copies of the surface of last scattering (SLS) in the covering space S^3 , modelled as $S^3 \subset \mathbb{R}^4$, is described in Sect. 3 of RBSG08.

Since it is important that the choice of pairs of points is as uniform as the nature of the observations allows, the selection of points on a copy of the SLS in RBSG08 was chosen uniformly on the SLS, i.e. on the 2-sphere. The binning into bins of pair separations was calculated only after the pair separations have been calculated by applying a holonomy transformation to one of the members of the pair.

From a computational point of view, applying the holonomy is the most intensive step in the calculation, requiring the multiplication of a 4×4 matrix by a 4-vector, as a rotation in \mathbb{R}^4 . This matrix calculation is carried out 12 times for a given pair of points, since all of the 12 holonomy transformations mapping a point to one of its neighbouring copies of the SLS must be examined.

The set of 20 holonomy transformations that maps the fundamental domain to the next layer in the direction of the hyperspherical equator could, in principle, be used too. However, for these to give matched circles as small as 5° , this would require the total matter-energy density to be $\Omega_{\text{tot}} \gtrsim 1.03$, which is uncomfortably high given present observational estimates of Ω_{tot} . To see this, replace $\pi/10$ in Eq. (15) of RBSG08 by $\pi/6$, which is half the geodesic length of any member of this set of holonomy transformations (Clifford translations), and use $R_C = (c/H_0)(\Omega_{\text{tot}} - 1)^{-1/2}$, where R_C is the radius of curvature, i.e. the radius of S^3 modelled as a subset of \mathbb{R}^4 . Moreover, it is not clear how a ‘‘generalised’’ twist parameter could be used if both sets (or yet further sets) of holonomy transformations were to be used in a single estimate of the cross-correlation function. The twist for this set of holonomy transformations is $\pm\pi/3$, not $\pm\pi/5$. For these reasons, we consider just the 12 immediately neighbouring copies of the fundamental domain, as in RBSG08.

In this paper, we argue that after uniformly selecting points on the 2-sphere, a filtering of these points is possible in a way that excludes only those points that are certainly not members of any close SLS-SLS pairs, for a given maximum pair separation r_2 . This should enable shortening the time to find a given number of close pairs for a constant number of points N_p , as well as finding a larger number of close pairs for the same value of N_p , since points that are certainly not members of any close pair are rejected before the iteration over pairs of points is started. The latter improvement to the algorithm should enable practical use of the method for shorter pair separations.

3.1. Preselection of potential members of close pairs: α_1, α_2

Figures 1 and 2 show how a close pair used in a cross-correlation estimate relates to the two copies of the SLS and the angle α_\pm that separates it, on a copy of the SLS, from a dodecahedral face centre. Using the spherical cosine formula, the lower triangle in the two figures has the relation

$$\cos \frac{x}{R_C} = \cos \frac{r_{\text{SLS}}}{R_C} \cos \frac{\pi}{5} + \sin \frac{r_{\text{SLS}}}{R_C} \sin \frac{\pi}{5} \cos \alpha_\pm \quad (3)$$

where α_+ and α_- are for Figs. 1 and 2 respectively.

⁶ For a multiply connected model, the distributions of the various a_{lm} 's are to some degree dependent on one other, so a method which assumes they are independent can at best give an approximate result.

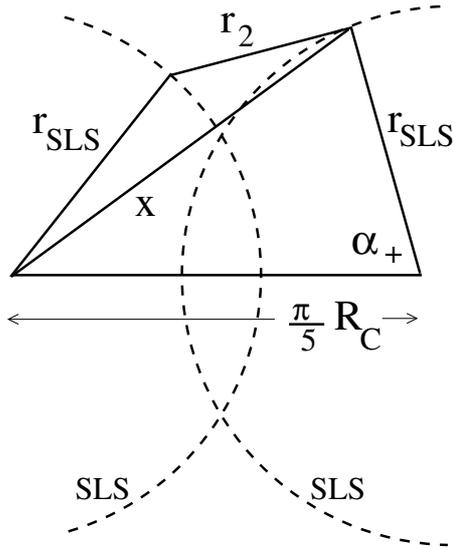


Fig. 1. Relation of a spatial geodesic of length r_2 joining a “close” pair of points in space to two copies of the SLS, and the angle α_+ separating the member of the pair on the right-hand SLS from the dodecahedral face centre, i.e. from the spatial geodesic joining the two SLS copies. This figure shows a spatial geodesic “external” to the matched circles (intersection between the two copies of the SLS). The centres of the two copies of the SLS (2-spheres) are separated by $(\pi/5)R_C$. See Sect. 3.1.

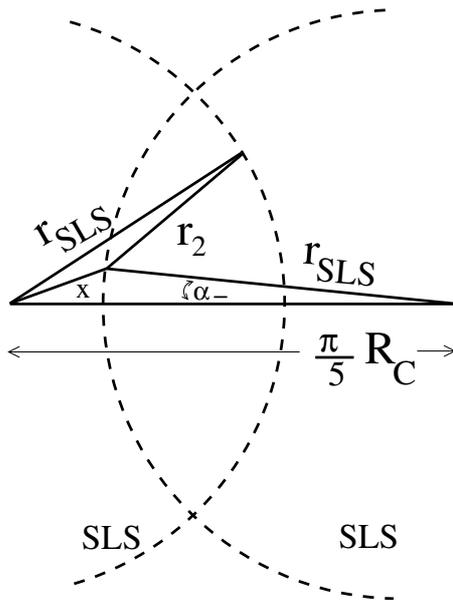


Fig. 2. As for Fig. 1, showing a spatial geodesic, of length r_2 , “internal” to the matched circles, and angle α_- separating the member of the pair on the right-hand SLS from the dodecahedral face centre.

Clearly, α_+ is maximised when x is maximised. For fixed r_2 and r_{SLS} , x is maximised when $x = r_{\text{SLS}} + r_2$, i.e. when the upper triangle degenerates into a single line segment. Hence, the angle α_+ is maximised when $x = r_{\text{SLS}} + r_2$.

Similarly, α_- is minimised when x is minimised, provided that $r_2 \leq r_{\text{SLS}}$, which is the situation most interesting for the SLS-SLS cross-correlation method, since close pairs are the most useful. For fixed r_2 and r_{SLS} , x is minimised when $x = r_{\text{SLS}} - r_2$, i.e. when the upper triangle degenerates into a single line segment. Hence, the angle α_- is minimised when $x = r_{\text{SLS}} - r_2$.

Hence, by symmetry, for a given separation $r_2 \leq r_{\text{SLS}}$, the minimum and maximum boundaries for defining an annulus

$$\alpha_- \leq \alpha \leq \alpha_+ \quad (4)$$

on the SLS around a dodecahedron face centre, in order to include all points that can potentially be members of a close pair separated by a spatial geodesic distance of at most r_2 for the corresponding pair of matched faces of the fundamental domain, are

$$\alpha_{\pm} = \cos^{-1} \left[\frac{\cos \frac{r_{\text{SLS}} \pm r_2}{R_C} - \cos \frac{r_{\text{SLS}}}{R_C} \cos \frac{\pi}{5}}{\sin \frac{r_{\text{SLS}}}{R_C} \sin \frac{\pi}{5}} \right]. \quad (5)$$

When r_2 is larger than the separation of dodecahedral face centres (Eq. (32) in RBSG08), i.e. when

$$r_2 > 2R_C \left(\frac{r_{\text{SLS}}}{R_C} - \frac{\pi}{10} \right), \quad (6)$$

the derivation leading to the expression for α_- in Eq. (5) is no longer valid. Instead, the lower limit α_- should be set to zero.

As mentioned above, the preselection enabled by Eq. (5) can be applied in two ways:

- (i) a point that fails to be a member of a “close” pair for a given maximum separation r_2 in all of the 12 directions of holonomy transformations to adjacent copies of the point in the covering space can be removed from the list of uniformly selected points; and
- (ii) when iterating through pairs of points and holonomy transformations g_i ($i = 1, \dots, 12$), a pair of points for which at least one of the two points does not satisfy the condition $\alpha \leq \alpha_+$ can be rejected without calculating the spatial separation of the pair.

When r_2 is large, α_+ will also be large, and the six pairs of annuli may be sufficiently wide that together they cover the whole sky. In this case, effect (i) will not occur. For a small enough separation r_2 , the effect should occur. In that case, more points can be uniformly selected from S^2 according to a uniform distribution, and again tested, until the required number of points is obtained. In that case, a higher fraction (though not 100%) of pairs defined by this set of points will be useful for the cross-correlation function calculation. This should increase the number of pairs per bin, especially for the smallest bins, which have the fewest numbers of pairs. This would be useful for high resolution calculations. Both effects (i) and (ii) should increase the calculation speed for a given number of points N_p , since they avoid having to carry out unnecessary matrix multiplications.

3.2. Use of preselection on the five-year WMAP data

The MCMC analysis is performed as in RBSG08, using the ILC and TOH five-year WMAP data (Sect. 2.1), but including the smallest separation bin, i.e. using the full range of separations

$$d/r_{\text{SLS}} < 40/90, \quad (7)$$

i.e.

$$d \lesssim 4.4 h^{-1} \text{ Gpc} \quad (8)$$

for $\Omega_{\text{tot}} = 1.01$, matter density $\Omega_m = 0.28$ and SLS redshift $z_{\text{SLS}} = 1100$, corresponding approximately to angles on the SLS $0 \leq \theta_d \leq 25^\circ$. All other parameters are kept as in Sect. 3.6 of RBSG08. In particular, this includes the five parameters for orientation of the fundamental dodecahedron (galactic longitude

and latitude of one face centre (l, b) and a rotation parameter θ around the axis defined by (l, b)), the matched circle size α and the “generalised” twist phase ϕ when matching opposite faces. The GPL (GNU General Public Licence) program CIRCLES⁷ is used.

3.3. Analysis of simulations

Although the method presented in Sect. 3.1 should speed up the calculation, analysing a large number of simulations is still time-consuming. Here, we estimate S_ξ for 50 simulations and the WMAP observational map described in Sect. 2.2 and select those 20 simulations whose S_ξ estimates are smallest, provided that S_ξ is not less than S_ξ of the observational map. In each estimate, 10 000 points selected randomly from a uniform distribution on the sky outside of the kp2 mask were used. Using each of these 20 simulations and the observational map, four MCMC chains with random starting points in the parameter space described in Eq. (28) of RBSG08 are carried out.

4. Results

4.1. Benchmarking

An example set of calculation times and number of pairs in the smallest separation bin are shown in Table 1, for one calculation of the auto-correlation and cross-correlation functions at an arbitrary PDS orientation and twist, for $\alpha = 5^\circ$ and $\alpha = 60^\circ$, which determine the ratio r_{SLS}/R_C . The speed-up factors range from about 3–10, depending on both α and r_2 .

For $d < r_2 = 4.4 h^{-1}$ Gpc, the “annulus outer radii” are $\alpha_+ = 45.0^\circ, 90.9^\circ$ for $\alpha = 5^\circ, 60^\circ$ respectively. This is clearly too large to allow any removal of points from the list of potentially useful pairs, i.e. effect (i) in Sect. 3.1 does not occur: N_A and N_C are negligibly affected by the pair preselection mechanism. However, the labelling of points to record which annuli they can potentially be pairs of does yield a speed-up through effect (ii), i.e. by factors of about 7 and 3 for $\alpha = 5^\circ$ and 60° respectively.

For a much smaller maximum pair separation, i.e. $d < r_2 = 0.4 h^{-1}$ Gpc, even though the annulus outer radii are still quite large, i.e. $\alpha_+ = 13.6^\circ, 62.9^\circ$ for $\alpha = 5^\circ, 60^\circ$ respectively, both effects (i) and (ii) occur. That is, not only is there a speed-up by a factor of about 7–10, but there is also an increase in the number of pairs in the smallest bin for calculating the cross-correlation function, by factors of about 30 and 7 for $\alpha = 5^\circ$ and 60° respectively.

In practice, use of a small maximum pair separation bin to get cosmologically significant results will be complicated by the relatively larger contributions from the Doppler and ISW effects, from residual foreground contamination, from the differences between various versions of the all-sky CMB map at these resolutions (Aurich et al. 2006), and from the absence of the large scale signal. Nevertheless, the concern expressed in Sect. 5.5 of RBSG08 that improvements to the algorithm would be geometrically quite complex, at the risk of introducing biases to the method, appears to have been overcome, reducing one obstacle to small scale work using this type of method.

⁷ Version CIRCLES-0.3.2.1 was used for calculating the MCMC chains for the five-year WMAP maps, and version CIRCLES-0.3.8 was used for calculating the chains for the INC3 observational and simulation maps. Various versions of CIRCLES are downloadable from <http://adjani.astro.umk.pl/GPLdownload/dodec/>. These and earlier versions of the software require medium to advanced GNU/LINUX, FORTRAN77 and C experience for a scientific user.

Table 1. Example of benchmarking on a 3 GHz processor ^a.

α °	r_2 $\approx h^{-1}$ Gpc	α_+^b °	N_p	t^c s	N_A^d	N_C^d
5	0.4	...	1000	5	1317	52
5	0.4	...	2000	22	5268	206
5	0.4	...	8000	375	83 476	3900
5	4.4	...	1000	5	1317	52
5	4.4	...	2000	23	5268	206
5	4.4	...	8000	377	83 476	3900
5	0.4	13.6	1000	0	6994	1911
5	0.4	13.6	2000	2	28 085	7860
5	0.4	13.6	8000	41	449 308	126 395
5	4.4	45.0	1000	0	1317	52
5	4.4	45.0	2000	3	5268	204
5	4.4	45.0	8000	53	83 476	3880
60	0.4	...	1000	5	1408	192
60	0.4	...	2000	23	5650	736
60	0.4	...	8000	354	90 242	12 070
60	4.4	...	1000	5	1408	192
60	4.4	...	2000	24	5650	736
60	4.4	...	8000	369	90 242	12 070
60	0.4	62.9	1000	0	3532	1518
60	0.4	62.9	2000	3	14 290	5986
60	0.4	62.9	8000	48	228 426	94 887
60	4.4	90.9	1000	2	1408	183
60	4.4	90.9	2000	8	5650	706
60	4.4	90.9	8000	130	90 242	11 536

^a The pseudo-random number generator has the same initial seed for each calculation; ^b an estimate of α_+ from Eq. (5) is shown in the cases where pair preselection as described in Sect. 3.1 is used; ^c calculation time; ^d numbers of pairs in the smallest bin N_A, N_C for ξ_A, ξ_C respectively.

In a full MCMC chain, the matched circle size α will vary between the limits illustrated in Table 1. The actual speed-up factor (and increase in numbers of pairs per bin if this occurs) will depend on the particular path of the MCMC chain through parameter space. In the calculations leading to the results described below, the range $5^\circ < \alpha < 60^\circ$ was retained.

4.2. Parameter estimates from the five-year WMAP data

For both maps of the five-year data as described in Sect. 2.1 (ILC and TOH), $N_{\text{chain}} = 16$ MCMC chains were run, each starting with different random seeds, using the kp2 mask. Each run had 12 000 steps and the first 2000 steps of each were discarded⁸. Figure 3 shows the sky positions (l, b) _{$i=1,12$} implied by the (l, b, θ) triples in the MCMC chains for which $P > 0.5$ (Eq. (25), RBSG08), for the ILC and TOH versions of the five-year WMAP observational data. The “probability” function P used for optimisation by the MCMC procedure is that defined in Eq. (25) in RBSG08. This is not a true probability function.

Similarly to what was done in RBSG08, the 16 chains are grouped together into four groups, each of four chains. For a given group, steps 2001 to 12 000 from each of the four chains are concatenated.

We make the convergence requirements a little more stringent than was described in Sect. 4.1 of RBSG08. That is, we start from an initial angular radius of $\beta_1 = 30^\circ$ (covering most of the sphere), decrease by 1° for the next 20 iterations, and then

⁸ The MCMC chains used in this paper can be downloaded for independent analysis from the file http://adjani.astro.umk.pl/GPLdownload/MCMC/mcmc_RBG08.tbz

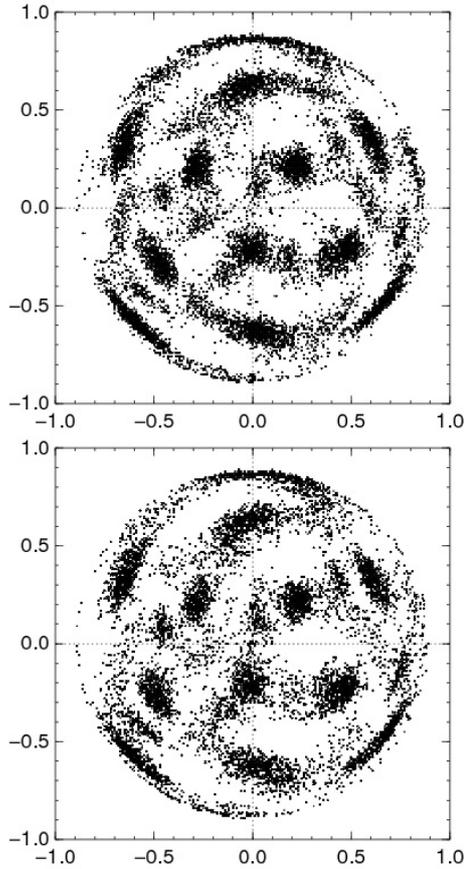


Fig. 3. Full sky map [Lambert azimuthal equal area projection (Lambert 1772), centred on the North Galactic Pole (NGP), with the 0° meridian as the positive vertical axis and galactic longitude increasing clockwise] showing the optimal orientation of dodecahedral face centres based on 160 000 steps in 16 MCMC chains, using the five-year ILC map (*upper panel*) and the five-year TOH map (*lower panel*), and the kp2 mask, showing face centres for which $P > 0.5$.

remain constant at $\beta_{j \geq 21} = 10^\circ$ until the iteration for that face number converges or until a total of 40 iterations has been reached. The analyses of the four concatenated groups of chains give what are considered to be four independent estimates in order to get an estimate of the uncertainties due to our MCMC estimation method.

The resulting numerical estimates are listed in Table 2. The columns show minimum “probability” P_{\min} , face number i , number n of MCMC steps contributing to the estimate obtained from the final iteration, galactic longitude l and latitude b , and the standard error in the mean between the four estimates of different sets of MCMC chains, in great circle degrees, $\sigma_{\langle(l,b)\rangle}$. These values do not differ significantly from those in Tables 1 and 4 (for the kp2 mask) of RBSG08.

The MCMC states for α and ϕ in the final radii of convergence, and their means and standard errors in the mean, are shown in Fig. 4 and Table 3. Histograms of the distribution of ϕ are shown in Fig. 5. Again, these values do not differ significantly from those in Table 2 of RBSG08. The TOH map shows a small offset between the best estimate of ϕ found (for $P > 0.5$) using our convergence algorithm and that at which the histogram of ϕ states peaks. The former is $\phi = 30.4^\circ$, i.e. a little below 36° , while the latter (lower panel of Fig. 5) is a few degrees above 36° . The difference can reasonably be attributed to the fact that the former uses a convergence algorithm in

Table 2. Sky positions of the best estimate of the six dodecahedral face centres for the five-year ILC map with the kp2 mask.

P_{\min}	i^a	n	l	b	$\sigma_{\langle(l,b)\rangle}$
0.4	7	5141	182.8	62.3	1.1
0.4	12	5239	305.5	44.4	1.2
0.4	3	7097	45.0	49.5	0.7
0.4	5	5812	115.9	19.5	1.2
0.4	8	5056	174.6	-3.0	2.9
0.4	10	4707	239.8	13.9	2.1
0.5 ^b	7	2736	181.1	62.2	1.4
0.5	12	2838	305.7	44.5	1.3
0.5	3	3712	45.4	49.3	0.7
0.5	5	3005	114.6	18.7	1.6
0.5	8	2732	178.1	-0.5	1.6
0.5	10	2570	239.9	13.8	1.6
0.6	7	1429	179.3	62.6	1.1
0.6	12	1487	306.0	44.3	1.2
0.6	3	1919	45.5	48.8	0.8
0.6	5	1544	113.6	18.4	2.0
0.6	8	1387	175.6	-1.6	3.0
0.6	10	1419	237.7	14.5	1.6

^a The face centres are ordered according to the ordering in Table 1 of RBSG08. The other 6 faces are directly opposite with identical errors; ^b the estimate for $P_{\min} = 0.5$ corresponds to the points shown in the upper panel of Fig. 3, based on 160 000 steps of MCMC chains.

multi-dimensional parameter space, whereas the latter is the peak of a projected distribution using a single parameter.

As was discussed in Sect. 4.2 of RBSG08, when α is not too large, small changes in the comoving separation between pairs lead to relatively large changes in α . Does the use of the smallest bin in the present analysis, and/or the use of the slightly improved quality between the three-year and five-year WMAP data help overcome the large uncertainty in α ? Visual inspection of Fig. 4 suggests that the MCMC chains on the ILC five-year map favour $22^\circ \lesssim \alpha \lesssim 35^\circ$, while those on the TOH map favour $15^\circ \lesssim \alpha \lesssim 35^\circ$. The problem of approximate degeneracy in α , which is presumably sensitive to moderate levels of systematic error, remains.

4.3. Simulations

4.3.1. Estimates of S_ξ

The value of S_ξ (Eq. (2)) in the INC3 observational map (using the kp2 mask) is

$$S_\xi^{\text{INC3}} = 963(\mu K)^4. \quad (9)$$

The 50 simulations have S_ξ in the range $1170(\mu K)^4 < S_\xi < 8645(\mu K)^4$, i.e. up to about an order of magnitude higher S_ξ than in the observations. As mentioned in Sect. 2.2, this is because the WMAP cosmological signal has a lot of power close to the Galactic Plane and we use the kp2 cut sky. In order to use the simulations best able to test the hypothesis that a simply connected universe with an observationally valid large scale auto-correlation can give a PDS-like signal, we select the 20 of these with the lowest S_ξ values, i.e. in the range

$$1170(\mu K)^4 < S_\xi < 3782(\mu K)^4, \quad (10)$$

i.e. with up to about 3.9 times larger S_ξ than in the observations. Since at large scales (outside of the kp2 cut), these simulations are more correlated than the observations, we should obtain an

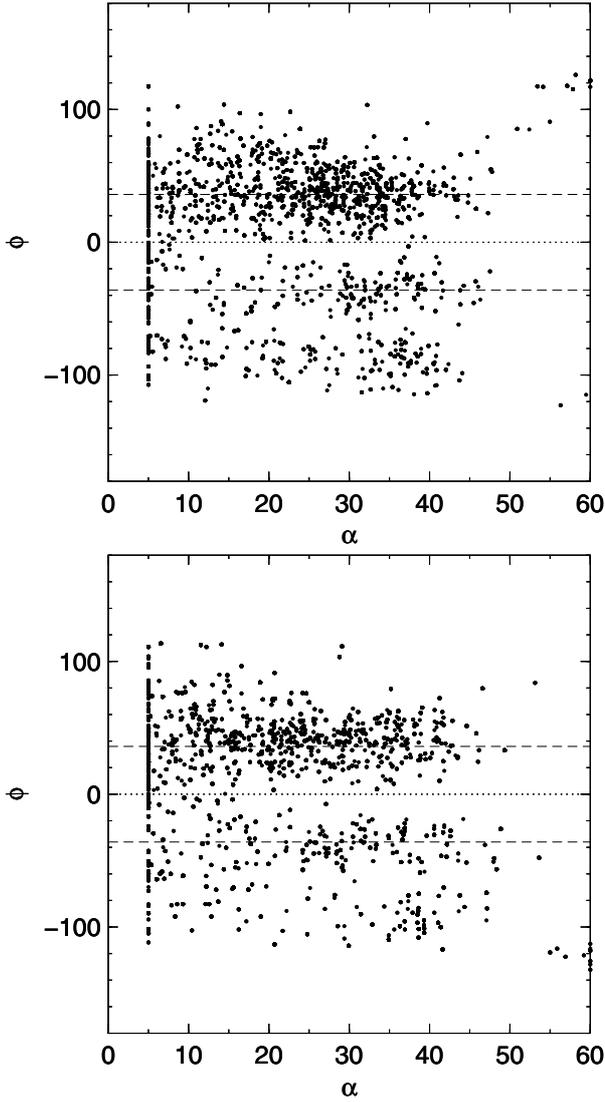


Fig. 4. Distribution of α and ϕ states where $P > 0.5$ in the MCMC chains of the dodecahedral solution used in Table 2, for the ILC (*upper panel*) and TOH (*lower panel*) maps. In this and similar figures, lines indicating $\pm 36^\circ$ are shown. These are not fit to the data.

upper limit to the the frequency of detecting PDS-like signals, assuming that the PDS hypothesis is wrong.

As a check on the amount of large scale power present in the Galactic Plane in different versions of the map of cosmological signal, we estimate S_ξ for the ILC and TOH five-year maps with and without the kp2 mask. Table 4 shows these estimates. Both maps have much more power without the cut than with the cut, as expected. Moreover, while the difference in estimates of S_ξ for the two maps differs by nearly a factor of two for the full sky, it differs by only 10% for the cut sky. This confirms the advantage in analysing the cut sky rather than the full sky: there is approximate consensus between these two different methods of generating the map. The INC3 estimate of $S_\xi = 963(\mu K)^4$ is also close to these two estimates.

4.3.2. MCMC chains

While the speed improvement introduced in this paper makes it possible to analyse a set of simulations rather than just an observational map, carrying out large numbers of MCMC chains still

Table 3. Estimates of matched circle radius α and twist phase ϕ from the Integrated Linear Combination (ILC) and Tegmark et al. (2003) (TOH) versions of the five-year WMAP data.

Map	P_{\min}	n^a	α	$\sigma_{\langle\alpha\rangle}$	ϕ	$\sigma_{\langle\phi\rangle}$
ILC	0.4	5508.67	20.3	0.7	38.7	1.6
ILC	0.5	2932.17	20.5	1	37.8	1.2
ILC	0.6	1530.83	20.1	0.7	39.8	1.1
TOH	0.4	4802.08	21.2	1.4	30.6	4.8
TOH	0.5	2955.58	20.4	0.8	32.1	2.6
TOH	0.6	1655.92	20.3	0.7	27.5	5.5

^a The number of steps n can be a non-integer since for a given MCMC step, it is possible that some of the face centres fall within the convergence radius of the final iteration as described in Sect. 4.2, but other face centres do not.

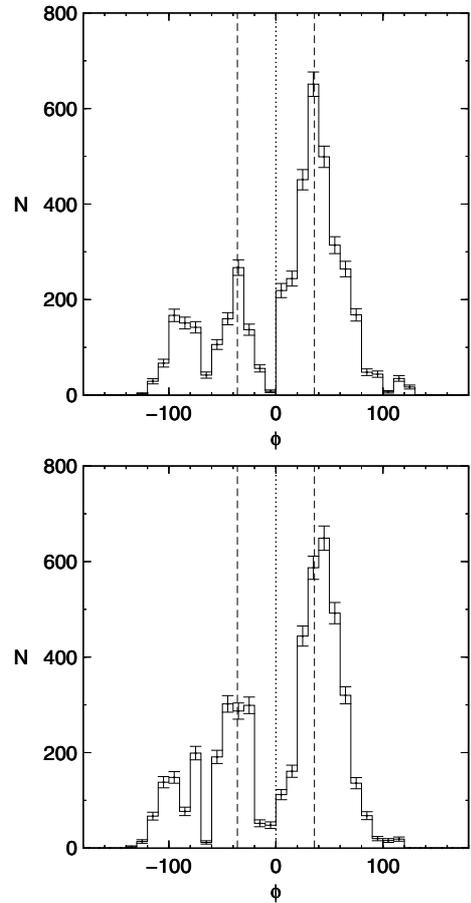


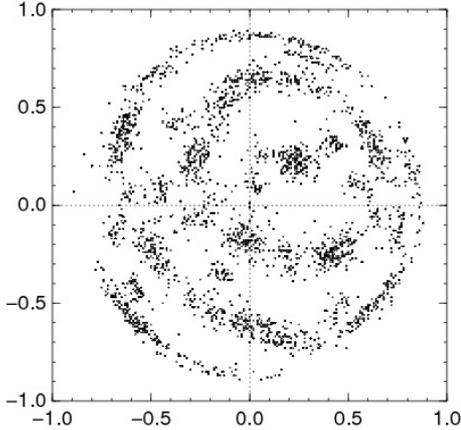
Fig. 5. Histogram of the twist angle ϕ shown in Fig. 4, with \sqrt{N} error bars, for the ILC (*upper panel*) and TOH (*lower panel*), excluding states with $\alpha \leq 15^\circ$.

remains computationally prohibitive. For this reason, we carry out only a small number of MCMC chains (four) on each data set (INC3 observational data set or simulation), and do not attempt to estimate uncertainties on the individual optimal parameters for a given data set.

We first concatenate together steps 2001 to 12 000 of each chain in a group, i.e. we ignore 2000 burn-in steps. This concatenated chain is considered to be a single chain for the iterative procedure of estimating parameter values. We modify the method of choosing an initial rough estimate of the optimal dodecahedral face centre positions that is used to start the iterations

Table 4. S_ξ estimates with and without the kp2 mask for the ILC and TOH five-year maps.

Map	kp2 (μK) ⁴	No mask (μK) ⁴
ILC	1012	4851
TOH	1136	2749

**Fig. 6.** Full sky map showing the optimal orientation of dodecahedral face centres based on 40 000 steps in 4 MCMC chains, using the three-year INC map and the kp2 mask, showing face centres for which $P > 0.5$.

towards a more precise estimate by randomly selecting a set of dodecahedral face centres $[(l, b)$ chosen from a uniform distribution on S^2 , θ chosen from a uniform distribution on $(0, 2\pi/5)$.] This risks causing the convergence to be less accurate, but since this is applied in the same way for both simulations and observations, this should not introduce any statistical bias for comparison of observations to simulations.

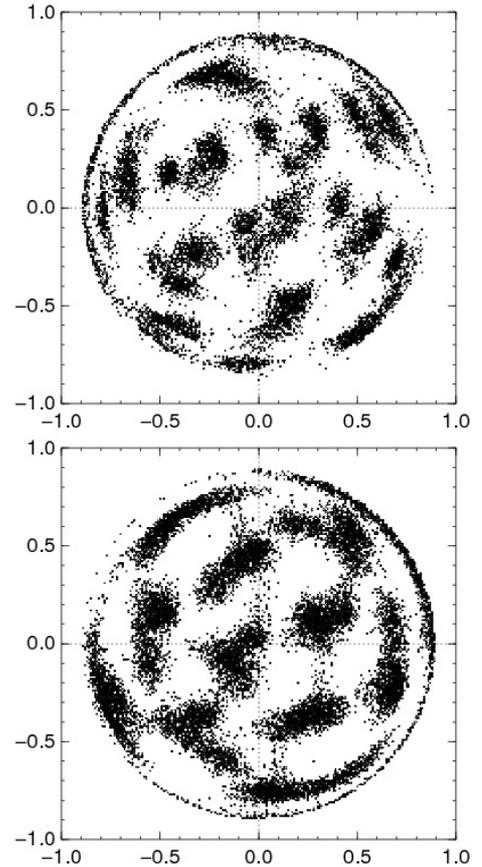
We also make our iteration a little more stringent than that described in Sect. 4.2. In the first iteration of parameter estimation, we estimate the dodecahedral face centres starting from a random initial set as just described, and in the following iterations, we converge on both dodecahedral faces centres and α , ϕ simultaneously. The values of β_j (see Sect. 4.2) are unchanged. The set of points within the convergence radius $\beta_{j \geq 21} = 10^\circ$ is used for the final estimates of α and ϕ .

4.3.3. Optimal dodecahedron orientation: (l, b, θ) space

Figure 6 shows the optimal set of face centres resulting from the four MCMC chains for the INC3 observational map. While the sharpness of the optimal signal does not appear as strongly as in the results using larger numbers of chains in RBSG08 and in Fig. 3 for the five-year data here, it is clearly consistent in position. Since we use many fewer chains here, it is unsurprising that the signal appears weaker. This should not be a problem for the present purpose, since the uncertainties from using a small number of chains should be statistically equivalent for both INC3 observational data and simulations.

On the other hand, the relative weakness of the signal and the smaller numbers of chains require using slightly lower minimum probability thresholds for estimating the best solution. Below, we cite results for $P_{\min} = 0.3, 0.4, 0.5$ rather than the earlier thresholds of $P_{\min} = 0.4, 0.5, 0.6$.

Examples of optimal dodecahedral orientations, for the two simulations whose S_ξ estimates are lowest (i.e. are closest to that

**Fig. 7.** Full sky map showing the optimal orientation of dodecahedral face centres based on 40 000 steps in 4 MCMC chains, using the two simulated three-year INC maps with lowest S_ξ estimates (simulations 92 and 90, in the upper and lower panels respectively), and the kp2 mask, showing face centres for which $P > 0.5$.

of the observations), are shown in Fig. 7. In one case, there is clearly a problem in converging on a single set of dodecahedral face centres, while in the other, there does appear to be more or less convergence to a single solution.

4.3.4. Circle size α and twist phase ϕ

Figure 8 shows that convergence in α and ϕ occurs for the INC3 observational map in a similar way to that of the five-year observational maps. In contrast, Fig. 9 shows that the two simulations whose S_ξ estimates best match that of the data, i.e. those whose best dodecahedral face centres are shown in Fig. 7, both “escape” towards the lower limit $\alpha = 5^\circ$. The median circle size is $\alpha = 5^\circ$ in both cases. This effect can be expected due to the increased relative Poisson error when comparing fewer numbers of pairs of pixels. Aurich (2008) found a similar effect, describing it as “drifting” towards “large L ”, which corresponds here to small α . Several previous authors have found that independently of whether or not a cosmic topology signal is present, relatively higher cross-correlations for zero separation pairs, i.e. for pairs on exactly matched circles, have been found to occur as the circle size α approaches zero. For example, see Fig. 2 in Cornish et al. (2004), Figs. 4–6 in Roukema et al. (2004), or Fig. 3 in Lew & Roukema (2008), where various definitions of a normalised correlation statistic S are shown to increase as α decreases towards zero. Hence, it can be expected that MCMC chains will be drawn towards the lower α limit.

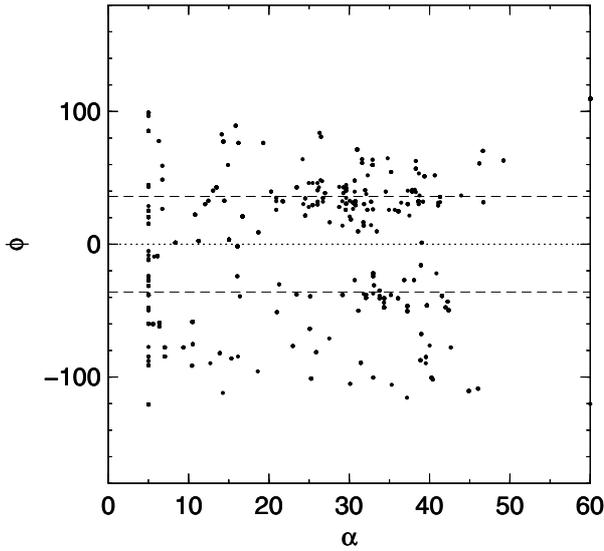


Fig. 8. Distribution of α and ϕ states where $P > 0.5$ in the MCMC chains for the INC3 observational map.

In cases where this occurs, chains spend a large amount of time at this limit, but cannot go below it. For this reason, the median of the α estimates (above a minimum “probability” used in the MCMC chains, e.g. $P > P_{\min} = 0.5$), should more accurately represent the optimal region favoured by the MCMC chains than the mean. Thus, here we use the median of α .

Figure 10 shows that nearly all the simulations escape to the lower α limit. Most of the simulations have best estimate α values less than 10° , and two have α estimates just a few degrees higher. Only two of the simulations have an optimum α estimate anywhere near the circle size α of the INC3 observational best estimate given in Table 5: $\alpha^{\text{INC3}} = 30.8^\circ$ (for $P_{\min} = 0.5$).

Do the latter two simulations, which are similar to the observational map in the sense that an optimal solution away from low α limit is found, have sufficiently similar characteristics to the observational map such that the observational map can be considered a chance realisation statistically similar to these two simulations? Figure 11 does not support this. The two simulations are among the simulations that have the highest S_ξ estimates, well above three times that present in the observational map. This is consistent with what is expected: the stronger the auto-correlation on large scales, the higher the chance is of getting high cross-correlations between apparently distant parts of the sky over a large number of pixel pairs, rather than escaping to the low α limit where there are relatively few pixel pairs. This can be quantified as follows.

Since the α distributions are constrained from below, they are unlikely to be Gaussian. So, estimating the significance of the correlation between α and S_ξ among the simulations is best done using a non-parametric statistic. Spearman’s rank correlation ρ and Kendall’s rank correlation τ one-sided tests with a positive correlation as the alternative hypothesis give probabilities that α and S_ξ are unrelated of 2.5% and 2.4% respectively. Even if we arbitrarily remove the two simulations with high α from the data set, the same two two rank correlation tests give probabilities that α and S_ξ are unrelated of 12.6% and 12.5% respectively.

This supports the visual impression from Fig. 11. For low S_ξ , an optimal cross-correlation in GRF simulations should escape to the low α limit. The point representing the observational simulation appears to be quite exceptional.

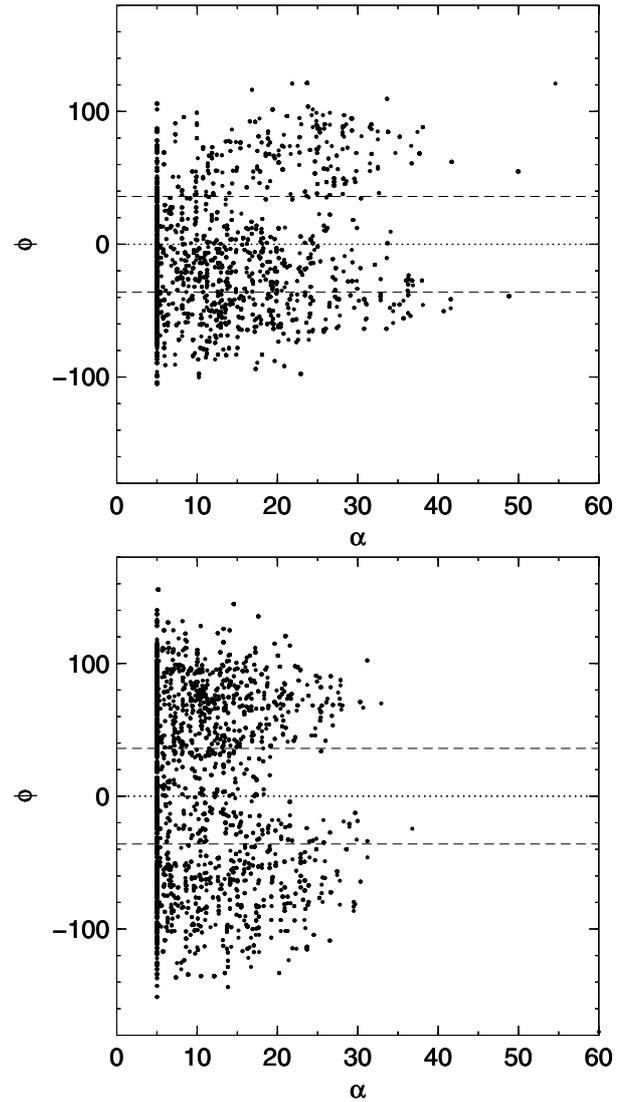


Fig. 9. Distribution of α and ϕ states where $P > 0.5$ in the MCMC chains for the two simulated three-year INC maps with lowest S_ξ estimates (simulations 92 and 90, in the upper and lower panels respectively).

Table 5. Estimates of matched circle radius α and twist phase ϕ for the INC3 observational map.

P_{\min}^a	n^b	α	ϕ
0.3	294.3	32.7	27.6
0.4	197.2	30.8	31.2
0.5	108.8	30.8	32.5

^a Minimum probability; ^b number of MCMC steps contributing to the estimate.

Do the two simulations with high α estimates (numbered 58 and 80, with α estimates of 24.5° and 31.9° respectively) have convergent estimates of dodecahedral face centres and ϕ ? Figure 12 indicates a poor convergence of dodecahedral face centres for simulation 58 and what looks like the superimposition of a strong primary and a weak secondary convergence in the case of simulation 80. Moreover, Fig. 13 shows that both of these two simulations have quite strongly bimodal distributions in ϕ rather than favouring any individual optimal value of the

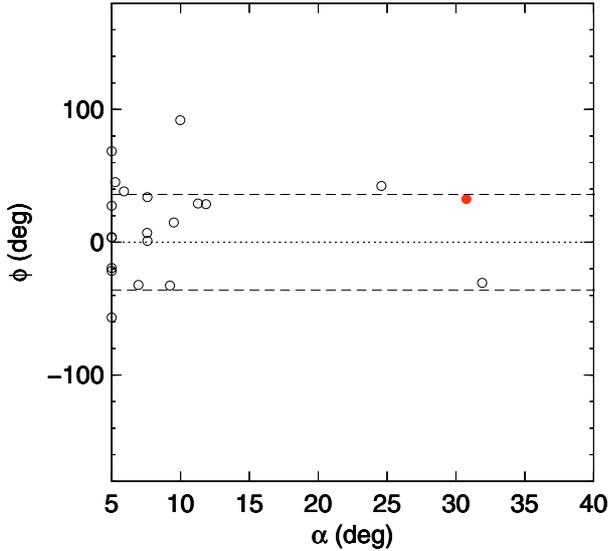


Fig. 10. Median circle sizes α (see Sect. 4.3.4) and twist angles ϕ (mean) for each of 20 simulations (empty circles) and the WMAP observational map (solid circle, values given in Table 5), analysed using the steps with $P > 0.5$ (see Eqs. (25), (26) of RBSG08). In each case, the 10 000 steps following 2000 burn-in steps of each of 4 MCMC chains started at random points in parameter space are concatenated for this analysis.

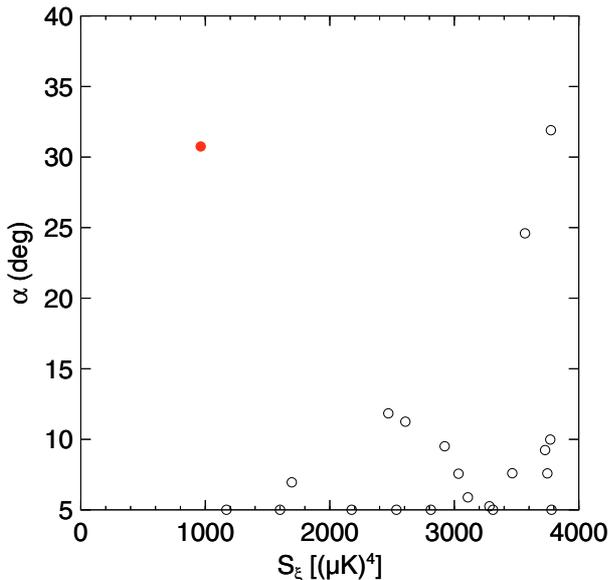


Fig. 11. Median circle sizes α as a function of S_ξ for each of the 20 simulations (empty circles) and the WMAP observational map (solid circle), analysed using the steps with $P > 0.5$.

twist ϕ . This is quite different behaviour to that in Fig. 8 for the observational map.

However, in order to be conservative, let us suppose that simulations 58 and 80 converge well enough in comparison to the observational map that we can consider them to have convergent MCMC solutions with $\alpha \gg \alpha^{\text{limit}} = 5^\circ$. This gives us an estimate

$$P(\alpha \gg \alpha^{\text{limit}} | S_\xi \leq 3.9S_\xi^{\text{INC3}}) \lesssim 10\%, \quad (11)$$

where $\alpha \gg \alpha^{\text{limit}}$ represents the event of getting a non-Poisson-noise signal at least 10° (the MCMC step size) away from the lower limit of α^{limit} .

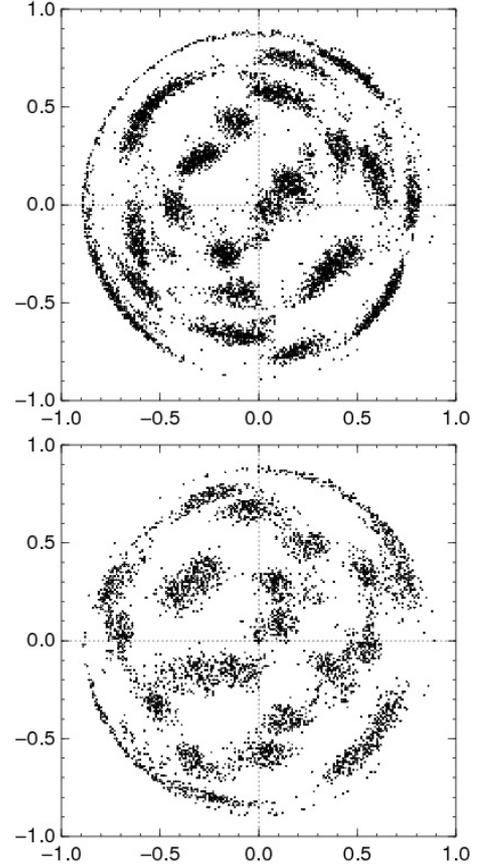


Fig. 12. Full sky map showing the optimal orientation of dodecahedral face centres based on 40 000 steps in 4 MCMC chains, using the two simulated three-year INC maps with the highest α estimates (simulations 58 and 80, in the upper and lower panels respectively), and the kp2 mask, showing face centres for which $P > 0.5$.

4.3.5. Distribution of the optimal twist phase ϕ

It is clear in Fig. 10 that the simulations give a distribution of best estimate twist angles ϕ different from the uniform distribution on $[-\pi, \pi]$ described in Eqs. (12) and (13) of RBSG08 and discussed in Sect. 5.4 of that paper. Figure 14 shows the distribution as a histogram. A two-sided Kolomogorov-Smirnov test between the distribution of the best estimates of ϕ from the 20 simulations and a uniform distribution on $[-\pi, \pi]$ rejects equality with $P = 0.01$. The values of ϕ for the simulations, shown in Figs. 10 and 14, all lie in the range $(-100^\circ, +100^\circ)$, and mostly seem to cluster even closer to $\phi = 0$.

A likely explanation is that this is a consequence of the anti-correlation in the auto-correlation function measured in WMAP sky maps at nearly antipodal scales by different authors using different methods (Fig. 16, Spergel et al. 2003; Fig. 1, RBSG08). The Hinshaw et al. (2007) C_l values should implicitly include the information that there is an antipodal anti-correlation. Estimates of the auto-correlations of the simulations confirm that an anti-correlation of about $-100(\mu\text{K})^2$ at $\phi = \pm\pi$ is present, so that the simulations correctly reproduce this characteristic of the observational data. This anti-correlation implies that MCMC chains should disfavour $\phi = \pm\pi$ when correlating pairs on a matched circle pair, and, hence, generally disfavour $\phi = \pm\pi$ for pairs on “matched annuli”. This is discussed further in Sect. 5.3.2.

Since a uniform distribution on $[-\pi, \pi]$ is clearly wrong, a reasonable hypothesis must be made regarding the intrinsic,

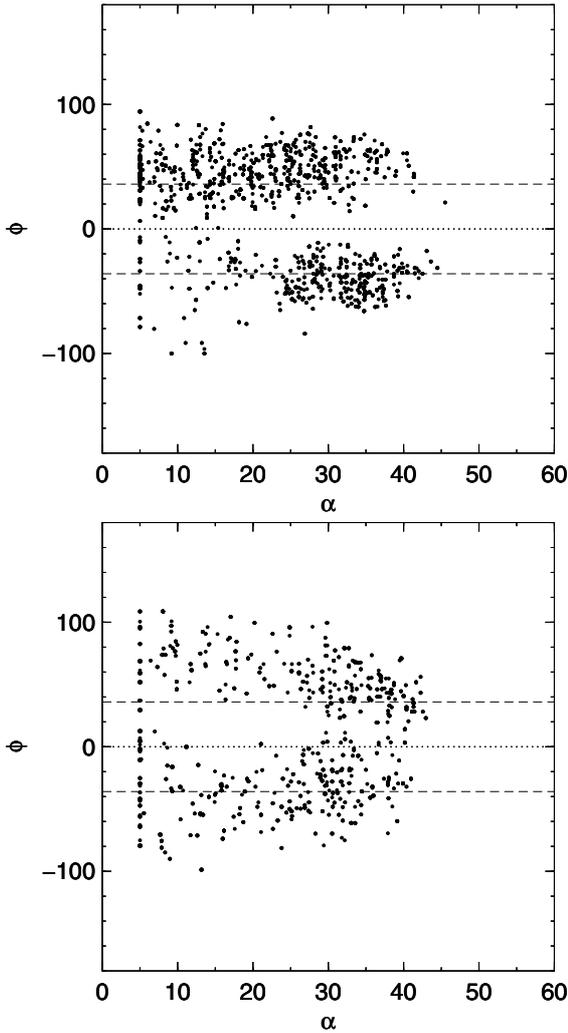


Fig. 13. Distribution of α and ϕ states where $P > 0.5$ in the MCMC chains for the two simulated three-year INC maps with highest α estimates (simulations 58 and 80, in the upper and lower panels respectively).

expected distribution of best estimates of ϕ . Given the numerical results from the simulational analyses and the presence of the $\phi = \pm\pi$ anticorrelation, which should favour ϕ away from $\pm\pi$ and towards zero, we assume a Gaussian distribution, $f(\phi)$, centred on $\phi = 0$ with width σ_ϕ estimated by the rms of ϕ in the simulations.

The parameters of this distribution for different P_{\min} values, and the Kolmogorov-Smirnov probabilities that the simulational values are consistent with a Gaussian distribution, are given in Table 6. For the lower thresholds, $P_{\min} = 0.3$ and $P_{\min} = 0.4$, the Gaussian distribution hypothesis is mildly rejected by the Kolmogorov-Smirnov test. If we set the mean and standard deviation of the Gaussian distribution to be the mean and standard deviation of the simulational ϕ estimates, then the three Kolmogorov-Smirnov probabilities for $P_{\min} = 0.3, 0.4, 0.5$ are $P_{\text{KS}} = 54\%, 84\%, 97\%$ respectively. Thus, it is clear that the distributions are consistent with Gaussianity if we use the estimated means rather than force symmetry about $\phi = 0$.

Could there be a reasonable justification for using a non-zero mean ϕ ? The only possible source of systematic asymmetry is that the noise simulations follow noise patterns in the observational data, which are not perfectly symmetrical. It is not obvious

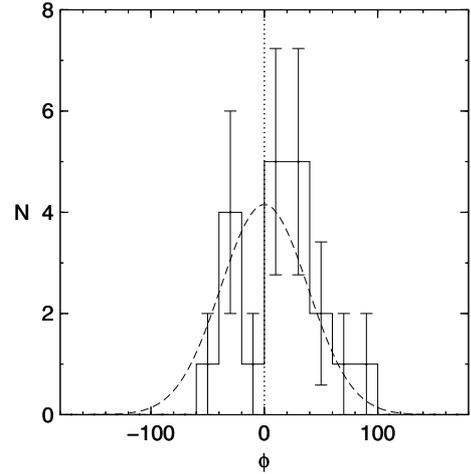


Fig. 14. Histogram of the optimal twist angle ϕ (shown in Fig. 10), with \sqrt{N} error bars, together with a Gaussian distribution of width 38.4° , centred at zero (see Table 6).

Table 6. Properties of the distribution of best estimates of ϕ from the 20 simulations.

P_{\min}	ϕ_{rms}^a	P_{KS}^b
0.3	32.9	5.9%
0.4	38.5	13.1%
0.5	38.4	32.8%

^a Root mean square width of the distribution; ^b two-sided Kolmogorov-Smirnov probability that the simulational values are consistent with a Gaussian distribution centred on zero, of width $\sigma_\phi = \phi_{\text{rms}}$.

that this asymmetry is sufficient to justify assuming an asymmetry in the expected distribution of ϕ . Moreover, the means and standard errors in the mean for the three P_{\min} thresholds are $12.8 \pm 7.1^\circ$, $14.4 \pm 8.4^\circ$, $12.1 \pm 8.6^\circ$ respectively, showing no statistically significant difference from zero.

Given a Gaussian distribution in ϕ centred on zero with width $\sigma_\phi = \phi_{\text{rms}}$ as listed in Table 6, the probability that ϕ is as close to $\pm\pi/5$ as the observational value (Table 5) is close to $+\pi/5$ is 22%, 12.8%, 9.4% for $P_{\min} = 0.3, 0.4, 0.5$ respectively, i.e.

$$P\left(\min(|\phi \pm \pi/5|) < |\phi^{\text{INC3}} - \pi/5| \mid S_\xi \leq 3.9S_\xi^{\text{INC3}}\right) \lesssim 22\%. \quad (12)$$

Since the values ϕ_{rms} are themselves not too far from $\pi/5$, it is clearly not so improbable that $\phi \approx \pm\pi/5$, compared to what would be expected from a uniform distribution on $[-\pi, \pi]$. In other words, this suggests that using the present method, an estimate of ϕ is not as good a discriminator between a chance PDS-like signal and an intrinsic, physical signal as it would be if the expected distribution were uniform on $[-\pi, \pi]$, i.e. there is another *topological degeneracy* (cf. Aurich et al. 2005a) in CMB all-sky maps.

On the other hand, is it just a coincidence that $\phi_{\text{rms}} \sim \pi/5$? In Sect. 5.3.2 below we discuss this question. It is possible that the empirical C_l spherical harmonic spectrum, even with randomised phases, may encode more cosmic topology information than might naively be expected.

4.3.6. Probability of rejecting the simply connected, infinite, flat model

The analyses of these simulations indicate that the requirement for the MCMC chains to avoid “escaping” to the lower limit in circle size α gives a stronger constraint against the simply connected, infinite, flat model than the requirement that could potentially exclude the PDS model, i.e. the requirement that $\phi \approx \pm\pi/5$.

Let us write the PDS-like characteristics of the WMAP observational data which we have tried to reproduce by simulations as follows. The data

- (i) have a large scale cutoff in structure statistics;
- (ii) yield a solution with $\alpha \gg \alpha^{\text{limit}}$ and $\phi \approx +\pi/5$ when using the MCMC method for optimising the cross-correlation ξ_C for the “generalised” PDS.

Rewrite these as

- (i) $S_\xi \lesssim S_\xi^{\text{INC3}}$;
- (ii) ξ_C yields $\alpha \gg \alpha^{\text{limit}}$ and $\min(|\phi \pm \pi/5|) < |\phi^{\text{INC3}} - \pi/5| \leq 8.4^\circ$,

where we use the “worst” estimate of ϕ from Table 5, i.e. $\phi = 27.6^\circ$, for $P_{\text{min}} = 0.3$. For convenience, we write these even more compactly as

- (i) $S_\xi \lesssim S_\xi^{\text{INC3}}$;
- (ii) ξ_C^{WMAP} .

Since both $S_\xi \approx S_\xi^{\text{INC3}}$ and ξ_C^{WMAP} are present in the WMAP data, the probability of these two characteristics both occurring in the simply connected, infinite, flat model can be written $P[(S_\xi \approx S_\xi^{\text{INC3}}) \cap \xi_C^{\text{WMAP}}]$. From the results above, in particular from Eqs. (11) and (12), we can write

$$\begin{aligned}
 P(\xi_C^{\text{WMAP}} | S_\xi \lesssim S_\xi^{\text{INC3}}) &= P[\alpha \gg \alpha^{\text{limit}} \text{ and} \\
 &\quad \min(|\phi \pm \pi/5|) < |\phi^{\text{INC3}} - \pi/5| \\
 &\quad \left| S_\xi \lesssim S_\xi^{\text{INC3}} \right] \\
 &\leq P(\alpha \gg \alpha^{\text{limit}} \mid S_\xi \lesssim S_\xi^{\text{INC3}}), \\
 &\leq P(\alpha \gg \alpha^{\text{limit}} \mid S_\xi \lesssim 3.9 S_\xi^{\text{INC3}}), \\
 &\lesssim 10\%. \tag{13}
 \end{aligned}$$

Are the probabilities that $\alpha \gg \alpha^{\text{limit}}$ and $\phi \approx \pm\pi/5$ independent of one another, so that instead of the Eq. (13), we can write

$$\begin{aligned}
 &P[\alpha \gg \alpha^{\text{limit}} \text{ and} \\
 &\quad \min(|\phi \pm \pi/5|) < |\phi^{\text{INC3}} - \pi/5| \mid S_\xi \lesssim S_\xi^{\text{INC3}}] \\
 &= P(\alpha \gg \alpha^{\text{limit}} \mid S_\xi \lesssim S_\xi^{\text{INC3}}) \times \\
 &\quad P[\min(|\phi \pm \pi/5|) < |\phi^{\text{INC3}} - \pi/5| \mid S_\xi \lesssim S_\xi^{\text{INC3}}] \approx 2.2\%, \tag{14}
 \end{aligned}$$

again using Eqs. (11) and (12)? Spearman’s rank correlation ρ and Kendall’s rank correlation τ one-sided tests where $\min(|\phi \pm \pi/5|)$ decreases as α increases as the alternative hypothesis both give probabilities that these two parameters are unrelated of 7.9%. A two-sided test gives 16% in both cases. While neither of these rejections of the hypothesis that the two parameters are unrelated is highly significant, they are strong enough rejections that it would be premature to assume that the two parameters are

independent. Hence, the probability estimate in Eq. (14) cannot (yet) be assumed to be valid.

As mentioned above, Spergel et al. (2003) estimated that $P(S_\xi \lesssim S_\xi^{\text{WMAP}}) \sim 0.15\%$ for an infinite, flat, cosmic concordance model, with a fixed spectral index of density perturbations. Efstathiou (2004) estimated $P(S_\xi \lesssim S_\xi^{\text{WMAP}}) \sim 3\text{--}12.5\%$. The latter calculation reconstructs unobserved structure hidden behind the galactic mask by assuming a simply connected model. For the purposes of testing the simply connected model hypothesis, this is internally consistent.

Using these two estimates to write two respective estimates of $P(S_\xi \lesssim S_\xi^{\text{INC3}}) \approx P(S_\xi \lesssim S_\xi^{\text{WMAP}})$, we have

$$\begin{aligned}
 P[(S_\xi \lesssim S_\xi^{\text{INC3}}) \cap \xi_C^{\text{WMAP}}] &= \\
 &P(S_\xi \lesssim S_\xi^{\text{INC3}}) P(\xi_C^{\text{WMAP}} | S_\xi \lesssim S_\xi^{\text{INC3}}) \tag{15}
 \end{aligned}$$

by the definition of conditional probability. Hence, from Eq. (13), we have

$$\begin{aligned}
 P[(S_\xi \lesssim S_\xi^{\text{INC3}}) \cap \xi_C^{\text{WMAP}}] &< 0.015\%, \\
 P[(S_\xi \lesssim S_\xi^{\text{INC3}}) \cap \xi_C^{\text{WMAP}}] &< 1.25\% \tag{16}
 \end{aligned}$$

for the lower and high estimates of $P(S_\xi \lesssim S_\xi^{\text{INC3}})$ respectively.

In other words, the simultaneous existence of both of these two properties of the WMAP data, one a generic characteristic of small universe models and the other highly specific to the PDS, is unlikely with a probability of about 99.99% or 99% depending on whether the Spergel et al. (2003) or Efstathiou (2004) estimates of $P(S_\xi \lesssim S_\xi^{\text{WMAP}})$ are used. This result only requires a frequentist approach to probabilities. Bayesian modelling of prior beliefs is not invoked.

5. Discussion

5.1. Matched annuli and calculational speedup

The preselection method described in Sect. 3.1 leads to faster calculation times by a factor of about 3–10 (Sect. 4.1). This is by eliminating most calculations of pair separations for pairs which are not useful for the cross-correlation calculation. Is it possible to improve this algorithm even further? Given that we have Eq. (5), the number of calculations required for a given total number of pairs could, in principle, be reduced by another small factor as follows.

Randomly select a point p_i from a uniform distribution on S^2 , considered to be the left-hand copy of the SLS in either Figs. 1 or 2. For each of the 12 holonomy transformations g_j to adjacent copies of the fundamental domain, check if the angle α_{ij} on the SLS from p_i to the dodecahedral face centre for g_j satisfies $\alpha_- \leq \alpha_{ij} \leq \alpha_+$. If this constraint is satisfied, then choose a second point p'_i randomly from a uniform distribution on the circle defined by the intersection of the right-hand copy of the SLS in Figs. 1 or 2 and the 2-sphere centred at p_i , having radius r_2 .

By construction, all of the pairs (p_i, p'_i) selected in this way are at separation r_2 , and should be statistically equivalent to generating a full set of pairs of which both members are uniformly selected on S^2 and then selecting those whose separations are close to r_2 . A loop over values of r_2 will give cross-correlations over the desired range of separations. This construction would bring the calculation method closer to the

identified circles method itself, with the difference that instead of correlating points lying precisely along the circles, points “near” the circles are correlated.

This raises the question of differences between using the present method and using the identified circles method with “thickened” circles. The most likely differences between these two methods would depend on how “thickening” is defined and how overlapping thickened circles are dealt with.

“Thickening” would require a somewhat arbitrary choice of an averaging procedure. For example, a circle could be divided into equal angular intervals and thickened so that individual bands of the annulus are internally averaged to obtain the temperature fluctuation at that angular position around the circle/annulus. Alternatively, a Gaussian smoothing could be used. However, neither of these methods would take into account the fact that the comoving spatial separation of a pair of points “bordering” opposite regions of the annulus does not change linearly as the width of the annulus increases. In contrast, the present method calculates and uses the spatial separations directly.

The problem of overlapping thickened circles follows from the fact that twelve pairs of annuli on the 2-sphere intersect each other many times. The greater the thickness of the annulus, the greater the number of randomly selected points in these intersections. If a correlation statistic is calculated over each circle separately and then averaged, then points which are members of these intersections contribute several times (typically two or three) more to the final statistic than points which are members of only one annulus. If chance fluctuations at some pairs of these intersections happen to have high correlations or anti-correlations, then the final statistic might be biased by these pairs, since it would implicitly assume that the pairs are independent of one another, even though this is false in some cases.

The use of numerical simulations should make this more of a problem of excessive noise rather than systematic bias. It is difficult to see any simple way in which a method directly based on the identified circles principle could avoid this problem, whether or not its final effect is statistical bias or rather an extra source of noise. For the SLS optimal cross-correlation method, this problem does not exist, since both points are selected uniformly on S^2 and selected afterwards based on their spatial separations (the method as presented in RBSG08), or else preselected in a way that is equivalent to this (the present paper).

Hence, results from comparing a series of thickened circle pairs for a range of thicknesses to use of the SLS optimal cross-correlation method should be expected to be different, for both geometrical and statistical reasons.

5.2. Five-year WMAP data

The five-year WMAP data (Sect. 4.2) give best estimates of the PDS model similar to those for the three-year data. Since most of the improvement in the WMAP data is at small angular scales, this is unsurprising.

5.3. Simulations

Given the infinite, flat model with GRF as a null hypothesis, the probability that the observed WMAP data could be a random realisation of this model, i.e. that both a large-scale cutoff in power and a specifically PDS-like signal appear in the WMAP data, is estimated in Eq. (16) as 0.015% or 1.25% depending respectively on whether we use the Spergel et al. (2003) or Efstathiou (2004) estimates of the former. This appears to be

a significant rejection of the infinite flat model. Is there any way to avoid this interpretation? Other properties of the cosmological component of the WMAP data, unlikely in the infinite flat model, have been noted by many authors. It is unlikely that these different properties are statistically fully independent of one another.

On the other hand, in this case we have a physically motivated model, motivated from the most fundamental level: the spatial section of the Universe is a 3-manifold, i.e. it must have a “shape”. As Schwarzschild (1900, 1998) stated a little over a century ago, that shape may well be multiply connected rather than simply connected.

More recently, the generic predictions of a cutoff in power at large scales were made when only the COBE data were available (Starobinsky 1993; Stevens et al. 1993). The identified circles principle (Cornish et al. 1998), of which the present method is an extension, and the matching of the Poincaré Dodecahedral Space hypothesis to the large scale lack of power and the estimates of $\Omega_{\text{tot}} \sim 1.01\text{--}1.02$ (Luminet et al. 2003), were published without the knowledge that the present version of the method, allowing for a “generalised” PDS of arbitrary twist angle, would be applied to the WMAP data. It seems physically unreasonable, therefore, not to combine the probabilities of the different signs of cosmic topology.

In this case, to argue that the probability in Eq. (16) has underestimated, either the Spergel et al. (2003) or Efstathiou (2004) probability estimates of a low S_ξ have to be increased, or the estimate of the probability of a PDS-like signal, given a low S_ξ , (Eq. (13)) has to be increased. We have conservatively taken the maximum probability estimate in Table 5 of Efstathiou (2004) for the former.

For the latter, we have conservatively assumed that the two simulations with $\alpha \gg \alpha^{\text{limit}}$ have a signal given by the MCMC method similar enough to the signal in the observational map that we can set this probability at $P(\xi_C^{\text{WMAP}} | S_\xi \lesssim S_\xi^{\text{INC3}}) \leq 10\%$. To increase this probability significantly, it would be necessary to argue that *many or most* of the simulations have convergence characteristics, best estimates of $\alpha \gg \alpha^{\text{limit}}$ and best estimates of $\phi \approx \pm\pi/5$ similar to those for the observational map. Given the results presented in Sect. 4.3, a systematic error of this sort able to satisfy this seems unlikely.

On the contrary, Fig. 11 and the rank correlation statistics (Sect. 4.3.4) favouring a positive correlation between α and S_ξ suggest that $P(\xi_C^{\text{WMAP}} | S_\xi \lesssim S_\xi^{\text{INC3}})$ is *smaller* than what we have been able to estimate with a small number of simulations. This is consistent with what was argued in RBSG08, i.e. that the lower the amplitude of the large-scale auto-correlations, the less chance there should be of cross-correlations occurring in the absence of a PDS-like signal or for a wrong orientation of the PDS model.

This suggests two possible alternative approaches to that used here. Either we could generate simulations so that *most* of the set of simulations have S_ξ about as low as that observed – with the kp2 cut – or we could generate simulations for the full uncut sky. However, both of these approaches have problems. The problems in both cases arise from the fact that a large amount of the large scale power is estimated to lie close to the Galactic Plane, but the estimates of how much this power is and of how precisely we can measure the detailed fluctuations close to the Galactic Plane are highly uncertain. For example, Table 4 shows that the estimates for S_ξ for the five-year ILC and TOH maps differ by only 10% when the kp2 cut is used, but differ by nearly a factor of 2 for the full sky.

Suppose that we generate simulations so that most of the set of simulations have S_ξ about as low as that observed, using the kp2 cut. In this case, we implicitly assume that both the extrapolations from outside of the kp2 cut to inside it, as well as the direct estimates for the full sky, vastly overestimate the power inside the kp2 cut. For testing a multiply connected model hypothesis, this could have some validity given that the spherical harmonics are not statistically independent from one another. However, this is not the null hypothesis that is to be tested with the simulations. The aim is to test the null hypothesis of a simply connected, infinite, flat model with Gaussian random fluctuations. This model implies a relation between structures inside and outside of the kp2 cut, which needs to be included in the simulations if we wish to correctly test the hypothesis. In other words, the statistics of fluctuations generated by simulations which are designed to mostly have small S_ξ outside of the kp2 cut are unlikely to be statistically equivalent to those of fluctuations generated in the way performed here. In the present work, the full set of simulations was generated by the Hinshaw et al. (2007) estimates of the C_l values, and a subset of these was selected with the criterion that S_ξ outside of the kp2 cut must be as close as possible to S_ξ of the observations outside of the kp2 cut. These two different methods of generating simulations are distinct.

The second alternative to the present method would be to generate simulations for the full uncut sky, and to run MCMC chains on both these and the observational map. However, because of the intrinsic difficulty in correcting for emission in the Galactic Plane and as is indicated in Table 4 for two different versions of the all-sky map of the cosmological signal, there would then be a much greater systematic uncertainty in the results. Hence, both of these two alternative approaches have disadvantages relative to the method used here.

In addition to our main results of probability estimates, it has become clear that the properties of the expected distribution of ϕ (for the null hypothesis of an infinite, flat model) are not simple. This can be understood generically by realising that the mathematical procedure we are using is pattern matching. A cross-correlation for the correct mapping between two genuinely correlated copies of a single pattern will necessarily yield a high value. However, for a pattern which is sufficiently complex, both cross-correlations for incorrect mappings between two genuinely correlated copies of a single pattern, and also cross-correlations for arbitrary mappings between two uncorrelated patterns may in some cases yield high values due to the complexities of the patterns and chance correlations. From first principles, modelling this is unlikely to be simple. For this reason, simulations provide an algorithmic shortcut to estimating the likely distribution of ϕ , given a certain family of patterns. Gaussian random fluctuations from a given C_l spectrum are one such family of patterns. For a different C_l spectrum or different statistical properties of the fluctuations, different characteristics of the expected distribution of ϕ may occur. Here we discuss some characteristics of interest.

5.3.1. Tendency of an MCMC chain to favour regions of ϕ approximately symmetric around zero

Figure 8 in RSBG08 and Figs. 4 and 8 each indicate that in addition to the main cross-correlation signal, there is also an additional, weak, secondary signal with an approximately equal but opposite value of ϕ . This is clear in Table 3 of RSBG08 and Fig. 5 for the five-year ILC and TOH maps. Since only four MCMC chains were carried out for each of the INC3 real and

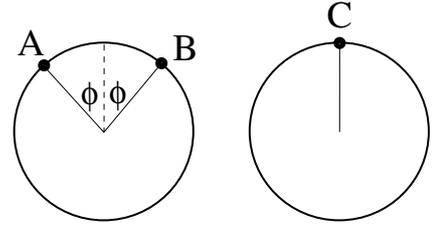


Fig. 15. Schematic diagram showing two circles on the SLS on which cross-correlation is to be optimised, for twist $\phi = 0$. The fluctuations are zero everywhere on the two circles except at the points A and B on the first circle and C on the second circle. At all three points the fluctuation amplitude is +1.

simulational maps for the main INC3 analysis, we carried out 12 additional MCMC chains on the real INC3 map in order to create a histogram equivalent to those in Fig. 5. The resulting histogram of the ϕ distribution from the 16 chains is shown in Fig. 16. This is clearly consistent with those in Fig. 5, in the sense that the $+\pi/5$ peak strongly dominates over the $-\pi/5$ peak.

Is it physically reasonable that a PDS model would give both $-\pi/5$ and $+\pi/5$ as valid twist angles, even where one twist gives a much weaker signal than the other? At most one of these can indicate the correct 3-manifold of comoving space. A possible explanation for both twist angles to be favoured in the MCMC chains might be that the density perturbations went through some sort of resonance process during an early epoch. In that case, it is conceivable that a harmonic created at that early epoch would still remain present. The optimal cross-correlations in the simulations, in particular those shown in Figs. 9 and 13, also have bimodal distributions of preferred values of ϕ . Moreover, they appear to give nearly equal weight to solutions with positive and negative values of ϕ of about the same absolute value. This suggests that some type of pattern in the fluctuations can tend to cause an MCMC chain to favour regions of ϕ approximately symmetric around zero.

The following schematic diagram shows that at least one specific pattern of fluctuations on the sky can lead to some degree of symmetry of optimal ϕ values, to the extent to which the real (or simulated) pattern mimics the idealised pattern. Figure 15 shows two circles on opposite sides of the sky, seen by someone external to the SLS, looking approximately but not nearly along the axis joining the two circles. If we approximate this region of the covering space S^3 as an approximately flat region for intuitive simplicity, then the mapping corresponding to a “generalised” holonomy transformation from one circle to the other is a translation followed by a twist. If the twist is either $+\phi$ or $-\phi$, then the summed cross-correlation is 1. If the twist is anything other than $\pm\phi$, including zero, then these two idealised circles will have a zero cross-correlation.

The original identified circles principle, i.e. the present method in the limit of zero pair separations, would imply that both of these patterns are equally optimal. By extension, our present method will find the same result for this idealised pattern.

This schematic situation is clearly highly simplified. Apart from the fact that the fluctuation patterns are unlikely to be as simple as in Fig. 15, here we have assumed that the MCMC chain is fixed at a particular orientation on the sky and matched circle size. In reality, the chains are free to change both of these. Changing orientation and/or circle size would weaken the symmetry provided by the pattern in Fig. 15. Hence, to calculate the relevance of this schematic pattern realistically from first principles would be quite complex. However, given that we have

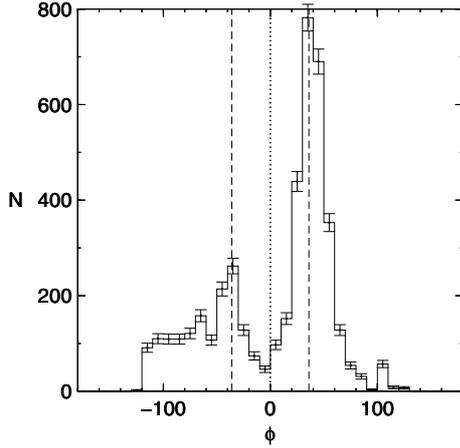


Fig. 16. Distribution of the twist angle ϕ in 16 MCMC chains run on the INC3 observational map, for states where $P > 0.5$ (the four chains shown in Fig. 8 plus 12 additional MCMC chains), with \sqrt{N} error bars, excluding states with $\alpha \leq 15^\circ$.

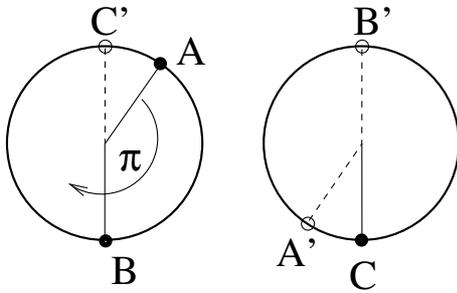


Fig. 17. Schematic diagram of two circles on the SLS for twist $\phi = 0$, illustrating that the observed anti-correlation at antipodal points on the SLS corresponds to an anti-correlation at twist $\phi = \pi$ on the circles. Solid circles at points A–C indicate positive fluctuations, and hollow circles at the antipodal points A'–C' indicate negative fluctuations.

realistic numerical simulations (for the infinite, flat model), we do have an illustration here of at least one basic pattern which could lead to best solutions which are approximately symmetric in ϕ around two values $\pm\phi$, even though only one of these two values is the “true” value (physical or simulated) of the map.

5.3.2. Properties of the distribution of optimal ϕ values

Statistical properties of the distribution of optimal ϕ values found in the analyses of the simulations (Fig. 14) are different from what was intuitively expected. Not only is the distribution of ϕ *not* uniform on $[-\pi, \pi]$, but it is moderately well fit by a Gaussian centred on $\phi = 0^\circ$ of width $\sigma_\phi \equiv \sqrt{\langle \phi^2 \rangle} \approx \pi/5$ (Table 6). Why is the distribution of optimal ϕ values not uniform on $[-\pi, \pi]$? Why should σ_ϕ be approximately $\pi/5$?

Consider the twist angle $\phi = \pi$. On matched circles, this represents pairs of antipodal points, independently of the matched circle size α . Figure 17 illustrates this schematically. In general, a pair of points on a pair of would-be matching circles are separated by a twist $\phi \ll \gamma$, where γ is the angle on the SLS separating the pair of points. The exception is that $\phi = \pi$ when $\gamma = \pi$. This separation corresponds to a spatial geodesic separation between the pair of points in the covering space S^3 of $2r_{\text{SLS}}$.

For antipodal points, ξ_A is slightly negative (Fig. 16, Spergel et al. 2003; Fig. 1, RBSG08). In the simulations analysed here,

this anti-correlation at $\phi = \pi$ is approximately $-100(\mu\text{K})^2$ in most simulations.

The MCMC chains are optimised to find positive cross-correlations, not anti-correlations. Hence, it should, in fact, be expected⁹ that optimal cross-correlations should cluster towards $\phi = 0$ and away from $\phi = \pi$.

We test this with a toy model as follows. The angular scale above which the auto-correlation becomes negative in Fig. 16 of Spergel et al. (2003) is $\arccos(-0.9) = 154^\circ$. To first order, we can approximately think of this as antipodal pairs of 25° discs on the sky being anti-correlated. So, we start with one of the simulated maps and randomly select 50 antipodal pairs from a uniform distribution on the sky. For each antipodal pair, we calculate the mean temperature fluctuation within 10° of each pole. We choose this radius to be smaller than 25° so that there is a fair chance that the disc we use will mostly cover a single “anti-correlated” disc, of full size 25° . If we had chosen the full disc size, then frequently there would be too little overlap for an anti-correlation to be measured. If the product of the mean temperature fluctuations in the two discs of the pair is negative, then we multiply the fluctuations in one member of the pair by -1 . This is done for all 50 antipodal disc pairs. If there were no overlap between these discs, then this would cover about 80% of the full sky.

While this procedure is *not* likely to give a valid statistical set of fluctuation patterns for testing the infinite, flat universe hypothesis, it should be sufficient for qualitatively testing the hypothesis that the anti-correlation is responsible for the concentration of the expected distribution of ϕ away from π and towards zero. Moreover, since the anti-correlation at the antipodes provides some contribution to S_ξ , this procedure decreases the minimum, mean, and median values of S_ξ for the 20 simulations from 1170, 2928, 3071 $(\mu\text{K})^4$ to 467, 1494, 1169 $(\mu\text{K})^4$ respectively, i.e. by about a factor of two. The maximum value of S_ξ increases slightly, from 3782 to 4012 $(\mu\text{K})^4$.

Figure 18 shows the result of searching for optimal PDS solutions using four MCMC chains per modified simulation, and analysing the chains as before. The distribution of best ϕ estimates is shown in Fig. 19. Several effects of the antipodal anti-correlation reversal are visible in the two figures. Firstly, several optimal correlations occur close to the antipodes, i.e. at $\phi \sim \pi$. Secondly, the chains which escape to α^{limit} have optimal ϕ values spread over a wider range than without the modification. These two effects are qualitatively consistent with the hypothesis that the antipodal anti-correlation plays an important role in concentrating the expected ϕ distribution for an infinite, flat model towards zero.

A third effect can also be seen in Fig. 18; a fairly large number of chains no longer escape to α^{limit} . Since we reverse the fluctuation sign in one element of each of 50 antipodal pairs, it is clear that this does not only affect antipodal correlations, it can also create correlations at several different scales, which did not exist in the original simulation. This can create highly artificial attractors for the MCMC chains.

The combination of these effects does not yield a distribution of ϕ which looks uniform by inspection of Figs. 18 and 19. A Kolmogorov-Smirnov two-sided test comparing the ϕ distribution to a uniform distribution on $[-\pi, \pi]$ yields a probability of 16%, i.e. the distribution is only marginally consistent with a uniform distribution. Clearly, although the non-uniform distribution of ϕ is clearly influenced by the antipodal anti-correlation, the latter alone is insufficient to explain it.

⁹ This fact was missed in RBSG08.

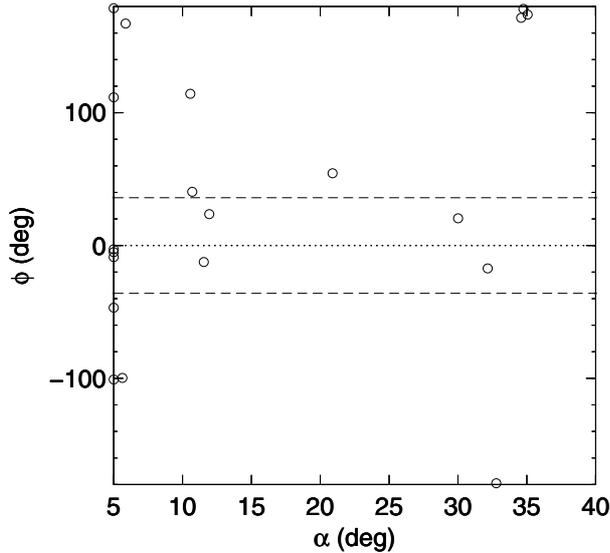


Fig. 18. Median circle sizes α and twist angles ϕ (mean) for each of the 20 simulations, analysed using the steps with $P > 0.5$, modified as described in Sect. 5.3.2 by inverting the anti-correlation in 50 randomly selected pairs of 10° radius antipodal patches on the sky.

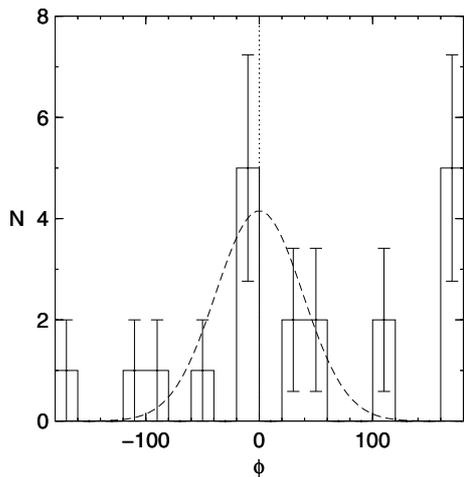


Fig. 19. Histogram of ϕ values shown in Fig. 18. The Gaussian distribution shown in Fig. 14 is reproduced here.

Given that the observational map does include an antipodal anti-correlation, and possibly other features contributing to the non-uniform expected distribution of ϕ , to what degree should the optimal cross-correlations favour ϕ close to zero, and what shape should the distribution of ϕ take? For example, if σ_ϕ as defined above is estimated on the interval (without periodicity) $[-\pi, \pi)$, then how small should it be?

The simulations are generated from spherical harmonics using the spherical harmonic mean coefficients C_l of Hinshaw et al. (2007), obtained from the WMAP data, but with random phases (and simulated noise). We choose those of the simulations with the lowest S_ξ estimates. These simulations contain much of the same information that is in the observational data set. After all, this is the point of simulations.

We know that the MCMC chains in the data with the correct phases favour $\phi = +\pi/5$, and have an anti-correlation at $2r_{\text{SLS}}$. Randomising the phases of the spherical harmonics while retaining the C_l values should yield many statistical properties that are similar to that of the map with the correct phases. So,

the simulated maps must statistically allow ϕ to be at least as high as $\pi/5$, but do not necessarily have to allow it to be much higher. Since the random phases have no way of favouring right-handed twists over left-handed twists, a distribution centred at $\phi = 0$ and including values roughly up to $\pm\pi/5$ is consistent with the information that the simulations should statistically contain. However, the reason why the distribution does not extend to optimal twists with $|\phi| \gg \pi/5$ remains an open question for future work.

6. Conclusion

By use of some spherical trigonometry, it is possible to speed up the Markov Chain Monte Carlo cross-correlation method of testing a cosmic topology hypothesis described in RBSG08 by a factor of about 3–10. This is shown in Fig. 1 and Eq. (5). This could, in principle, make it practical to make calculations at higher resolution than before. However, the physical interpretation of the calculations would be ambiguous because of several systematic effects listed above.

Moreover, for low matched circle sizes and low maximum pair separations, use of Eq. (5) can increase the numbers of pairs per separation bin, thereby decreasing the noise. A further improvement in the method and the relation between this method as a test of “matched annuli” and tests of matched circles are discussed above in Sect. 5.1.

Applying the faster method to the ILC and TOH versions of the WMAP five-year data, we find little significant change in the best estimate parameters for a Poincaré dodecahedral space model of the Universe, compared to those given in RBSG08. Depending on the minimum pseudo-probability level used in the MCMC chains that is used for estimating the twist ϕ , the optimal value of ϕ is a few degrees higher and a few degrees lower than $+36.0^\circ$ in the ILC and TOH maps respectively (Table 3).

We also applied the faster method to a small number of simulated skies. The WMAP observations confirmed the generic cosmic topology prediction (Starobinsky 1993; Stevens et al. 1993) of a cutoff in structure statistics such as the temperature-temperature fluctuation auto-correlation on large angular or spatial scales. Here, we estimated the weakness of the large scale auto-correlations using a statistic S_ξ (Eq. (2)) similar to that of Spergel et al. (2003). The low observed value of S_ξ implies that cross-correlations on these scales should usually be weak, so that MCMC chains as used here should have difficulty finding a region of parameter space with high optimal cross-correlations. Hence, to conservatively test the infinite, flat, concordance model hypothesis, on the assumption that the real observations have a low value of S_ξ due to being a single realisation of a random process, we analysed simulated skies which use the Hinshaw et al. (2007) C_l estimates and randomised phases of the spherical harmonics. The observational map was calculated from the three least contaminated frequency WMAP bands, Q , V and W , and the Gaussian random fluctuation simulations were created using an equivalent analysis pipeline (Lew & Roukema 2008).

For simulated and observed skies with the kp2 Galactic cut, these simulations using the Hinshaw et al. (2007) C_l estimates generally yield overestimates of S_ξ for the cut sky. Since we prefer to analyse the cut sky in order to minimise galactic contamination effects, it was necessary to select those simulations with the lowest values of S_ξ , so that the simulations were statistically as similar as possible to the observations outside of the kp2 cut.

The results of running MCMC chains on the 20 simulations with the lowest S_ξ estimates from among 50 simulations were that only two simulations gave optimal cross-correlation solutions which did not escape to the lower limit in circle size α , where small number statistics favour fortuitous cross-correlations. This indicates a conditional probability of finding a non-noise solution of about 10% (Eq. (11)).

The distribution of optimal twists ϕ from the simulations showed that despite the low value of S_ξ , the expected distribution of ϕ in the simulations was not uniform. Instead, it is consistent with a Gaussian distribution centred near $\phi = 0$, of width $\sim 33\text{--}38^\circ$. Possible reasons for the nature of this distribution are discussed above, in Sects. 5.3.1 and 5.3.2.

Assuming this numerical, Gaussian fit to the expected (for an infinite, flat model) distribution of ϕ , and using the estimate of ϕ obtained from the observational map (Table 5) that is most discrepant from $\pi/5$, we found that ϕ could be expected to be as close to $\pm\pi/5$ as the observational value is close to $+\pi/5$ with a probability of about 22% (Eq. (12)).

Both of these probabilities are conditional on S_ξ being low, which itself is unusual for an infinite, flat model (e.g. Spergel et al. 2003; Efstathiou 2004). Hence, for an infinite, flat, cosmic concordance model with Gaussian random fluctuations, we find that comparing an observational map with simulational maps gives an estimate of the chance of finding both (a) a large scale autocorrelation as weak as that observed, and (b) a PDS-like, optimal cross-correlation signal similar to that observed to be about 0.015% or 1.25% for the Spergel et al. (2003) or Efstathiou (2004) estimates of the probability of low S_ξ respectively (Eq. (16)).

Acknowledgements. Thank you to Bartosz Lew for numerous helpful and insightful comments, and to the anonymous referee who provided constructive and thoughtful recommendations. Usage of the Nicolaus Copernicus Astronomical Center (Toruń) computer cluster is gratefully acknowledged. Use was made of the WMAP data (<http://lambda.gsfc.nasa.gov/product/>), of the Centre de Données astronomiques de Strasbourg (<http://cdsads.u-strasbg.fr>), of GNU Octave command-line, high-level numerical computation software (<http://www.gnu.org/software/octave>), the GNU project R environment for statistical computing and graphics (<http://www.r-project.org/>) and the GNU PLOTUTILS plotting package.

References

- Aurich, R. 2008, *class. Quant. Gra.*, 25, 225017
Aurich, R., Lustig, S., & Steiner, F. 2005a, *Class. Quant. Gra.*, 22, 3443
Aurich, R., Lustig, S., & Steiner, F. 2005b, *Class. Quant. Gra.*, 22, 2061
Aurich, R., Lustig, S., & Steiner, F. 2006, *MNRAS*, 369, 240
Caillerie, S., Lachièze-Rey, M., Luminet, J.-P., et al. 2007, *A&A*, 476, 691
Cornish, N. J., Spergel, D. N., & Starkman, G. D. 1998, *Class. Quant. Gra.*, 15, 2657
Cornish, N. J., Spergel, D. N., Starkman, G. D., & Komatsu, E. 2004, *Phys. Rev. Lett.*, 92, 201302
Efstathiou, G. 2004, *MNRAS*, 348, 885
Gausmann, E., Lehoucq, R., Luminet, J.-P., Uzan, J.-P., & Weeks, J. 2001, *Class. Quant. Gra.*, 18, 5155
Gundermann, J. 2005, arXiv preprints [arXiv:astro-ph/0503014]
Hinshaw, G., Spergel, D. N., Verde, L., et al. 2003, *ApJS*, 148, 135
Hinshaw, G., Nolta, M. R., Bennett, C. L., et al. 2007, *ApJS*, 170, 288
Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, *ApJS*, in press [arXiv:0803.0732]
Key, J. S., Cornish, N. J., Spergel, D. N., & Starkman, G. D. 2007, *Phys. Rev. D*, 75, 084034
Lambert, J. 1772, *Anmerkungen und Zusätze zur Entwerfung der Land und Himmelscharten*. In *Beiträge zum Gebrauche der Mathematik und deren Anwendung*, pt. 3, sec. 6. (English translation: Notes and Comments on the Composition of Terrestrial and Celestial Maps, Ann Arbor, University of Michigan 1972)
Lehoucq, R., Weeks, J., Uzan, J.-P., Gausmann, E., & Luminet, J.-P. 2002, *Class. Quant. Gra.*, 19, 4683
Lew, B., & Roukema, B. F. 2008, *A&A*, 482, 747
Luminet, J., Weeks, J. R., Riazuelo, A., Lehoucq, R., & Uzan, J. 2003, *Nature*, 425, 593
Niarchou, A., & Jaffe, A. 2007, *Phys. Rev. Lett.*, 99, 081302
Riazuelo, A., Weeks, J., Uzan, J., Lehoucq, R., & Luminet, J. 2004, *Phys. Rev. D*, 69, 103518
Roukema, B. F., Lew, B., Cechowska, M., Marecki, A., & Bajtlik, S. 2004, *A&A*, 423, 821
Roukema, B. F., Buliński, Z., Szaniewska, A., & Gaudin, N. E. 2008, *A&A*, 486, 55
Schwarzschild, K. 1900, *Vier. d. Astr. Gess*, 35, 337
Schwarzschild, K. 1998, *Class. Quant. Gra.*, 15, 2539
Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, *ApJS*, 148, 175
Starobinsky, A. A. 1993, *J. Exper. Theor. Phys. Lett.*, 57, 622
Stevens, D., Scott, D., & Silk, J. 1993, *Phys. Rev. Lett.*, 71, 20
Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. 2003, *Phys. Rev. D*, 68, 123523
Weeks, J. 2001, *The Shape of Space* (2nd Ed.) (Manhattan: Marcel Dekker)
Weinberg, S. 1972, *Gravitation and cosmology: Principles and applications of the general theory of relativity* (New York: Wiley)