

HEMISPHERICAL POWER ASYMMETRY IN THE THIRD-YEAR *WILKINSON MICROWAVE ANISOTROPY PROBE* SKY MAPS

H. K. ERIKSEN,^{1,2,3,4} A. J. BANDAY,⁵ K. M. GÓRSKI,^{3,4,6} F. K. HANSEN,^{1,2} AND P. B. LILJE^{1,2}

Received 2007 January 4; accepted 2007 March 14; published 2007 April 10

ABSTRACT

We consider the issue of hemispherical power asymmetry in the third-year *WMAP* data, adopting a previously introduced modulation framework. Computing both frequentist probabilities and Bayesian evidences, we find that the model consisting of an isotropic CMB sky modulated by a dipole field gives a substantially better fit to the observations than the purely isotropic model, even when accounting for the larger prior volume. For the ILC map, the Bayesian log-evidence difference is ~ 1.8 in favor of the modulated model, and the raw improvement in maximum log likelihood is 6.1. The best-fit modulation dipole axis points toward $(l, b) = (225^\circ, -27^\circ)$, and the modulation amplitude is 0.114, in excellent agreement with the results from the first-year analyses. The frequentist probability of obtaining such a high modulation amplitude in an isotropic universe is $\sim 1\%$. These results are not sensitive to data set or sky cut. Thus, the statistical evidence for a power asymmetry anomaly is both substantial and robust, although not decisive, for the currently available data. Increased sky coverage through better foreground handling and full-sky and high-sensitivity polarization maps may shed further light on this issue.

Subject headings: cosmic microwave background — cosmology: observations — methods: statistical

1. INTRODUCTION

While the first-year results from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) experiment (Bennett et al. 2003) overall clearly supported the currently popular inflationary cosmological model, describing a flat, isotropic, and homogeneous universe seeded by Gaussian and adiabatic fluctuations, a disturbing number of unexpected anomalies on large scales were reported shortly after the public data release. Perhaps the three most important ones were: (1) alignments and symmetry features among low- l multipoles (de Oliveira-Costa et al. 2004; Eriksen et al. 2004a), (2) an apparent asymmetry in the distribution of fluctuation power in two opposing hemispheres (Eriksen et al. 2004b; Hansen et al. 2004), and (3) a peculiar cold spot in the southern hemisphere (Vielva et al. 2004; Cruz et al. 2005). All of these features were subsequently studied extensively by independent groups, and all remain unresolved to the present day.

In 2006 March, the third-year *WMAP* results were released, prompting researchers to revisit the anomalies detected in the first-year data (Bridges et al. 2006; Copi et al. 2007; Jaffe et al. 2006; Land & Magueijo 2007; Martínez-González et al. 2006). Of course, considering that already the first-year data were strongly signal-dominated on the scales of interest, it should come as no surprise that most of these analyses concluded with similar results as for the previous data, although different foreground handling could affect some results.

The *WMAP* team paid particular attention to the question of large-scale power asymmetry in their analyses (Hinshaw et al. 2006; Spergel et al. 2007a). Specifically, in an early version of their paper, Spergel et al. (2007a) approached the problem

from a semi-Bayesian point of view, by defining a parametric model consisting of an isotropic and Gaussian cosmic microwave background (CMB) field modulated by a large-scale function. The power asymmetry anomaly was then addressed by a dipolar modulation field, and the low- l alignment anomalies were studied with a quadrupole modulation field. However, due to several issues with this early analysis, several of which were first addressed by the present paper, the authors decided to remove the corresponding section from the final version of their paper (Spergel et al. 2007b). One example is simple marginalization over noncosmological monopole and dipole components, which was first done by Gordon (2007) in an otherwise identical analysis. A second example was the limited harmonic range considered by Spergel et al. (2007a). Thus, we present in this Letter the first complete modulation analysis that covers the full range of angular scales presented by Eriksen et al. (2004b) and that takes into account all known sources of systematics, such as monopole/dipole and foreground marginalization. We also present the first proper computation of the Bayesian evidence for the modulated model.

Following the first report of the power asymmetry, much effort has been spent by theorists on providing possible physical explanations. Examples range from those questioning the very fundamentals of physics and cosmology (e.g., introducing intrinsically inhomogeneous cosmologies [Moffat 2005; Jaffe et al. 2005], violation of Lorenz invariance [Kanno & Soda 2006], or violation of rotational invariance in the very early universe [Ackermann et al. 2007]) to those essentially considering special cases of established physics (e.g., second-order gravitational effects from local inhomogeneities [Tomita 2005], the presence of local voids [Inoue & Silk 2006], or spontaneous isotropy breaking from nonlinear response to long-wavelength density fluctuations [Gordon et al. 2005]).

2. ALGORITHMS

We now outline the methods used for the analyses presented in the following sections.

¹ Institute of Theoretical Astrophysics, University of Oslo, N-0315 Oslo, Norway.

² Centre of Mathematics for Applications, University of Oslo, N-0316 Oslo, Norway.

³ Jet Propulsion Laboratory, Pasadena, CA 91109.

⁴ California Institute of Technology, Pasadena, CA 91125.

⁵ Max-Planck-Institut für Astrophysik, D-85741 Garching bei München, Germany.

⁶ Warsaw University Observatory, 00-478 Warsaw, Poland.

2.1. Data Model and Likelihood

We model the CMB temperature sky maps as

$$\mathbf{d}(\hat{\mathbf{n}}) = s(\hat{\mathbf{n}})[1 + f(\hat{\mathbf{n}})] + \mathbf{n}(\hat{\mathbf{n}}), \quad (1)$$

where $s(\hat{\mathbf{n}})$ is a statistically isotropic and Gaussian random field with power spectrum C_l , $f(\hat{\mathbf{n}})$ is a dipole modulation field with amplitude less than unity, and $\mathbf{n}(\hat{\mathbf{n}})$ is instrumental noise. Thus, the modulated signal component is an anisotropic, but still Gaussian, random field and therefore has a covariance matrix given by

$$\tilde{s}(\hat{\mathbf{n}}, \hat{\mathbf{m}}) = [1 + f(\hat{\mathbf{n}})]s(\hat{\mathbf{n}}, \hat{\mathbf{m}})[1 + f(\hat{\mathbf{m}})], \quad (2)$$

where

$$s(\hat{\mathbf{n}}, \hat{\mathbf{m}}) = \frac{1}{4\pi} \sum_l (2l+1) C_l P_l(\hat{\mathbf{n}} \cdot \hat{\mathbf{m}}). \quad (3)$$

Taking into account instrumental noise and possible foreground contamination, the full covariance matrix is

$$\mathbf{C}(\hat{\mathbf{n}}, \hat{\mathbf{m}}) = \tilde{s}(\hat{\mathbf{n}}, \hat{\mathbf{m}}) + \mathbf{n} + \mathbf{F}. \quad (4)$$

The noise and foreground covariance matrices depend on the data processing and are described in greater detail in § 3. With these definitions ready at hand, the log likelihood is given by

$$-2 \log \mathcal{L} = \mathbf{d}^T \mathbf{C}^{-1} \mathbf{d} + \log |\mathbf{C}|, \quad (5)$$

up to an irrelevant constant.

2.2. Posterior Distributions and Choice of Parameters

The posterior distribution $P(\theta|\mathbf{d})$ is a primary goal of any Bayesian analysis, θ being the set of all free parameters in the model. For the model defined above, the free parameters can be divided into two groups, namely, those describing the isotropic CMB covariance matrix or C_l , and those describing the modulation field. Both may be parameterized in a number of different ways, and these choices may affect the outcome of the analysis through different prior definitions.

First, for the isotropic CMB component, we choose to parameterize the power spectrum in terms of a simple two-parameter model with free amplitude q and tilt n ,

$$C_l = q \left(\frac{l}{l_0}\right)^n C_l^{\text{fid}}. \quad (6)$$

Here l_0 is a pivot multipole and C_l^{fid} is a fiducial model, in the following chosen to be the best-fit power-law spectrum of Hinshaw et al. (2006). Second, the modulation field $f(\hat{\mathbf{n}})$ is parameterized in terms of a direction $\hat{\mathbf{p}}$ and an overall amplitude A ,

$$f(\hat{\mathbf{n}}) = A \hat{\mathbf{n}} \cdot \hat{\mathbf{p}}. \quad (7)$$

We use flat priors on all parameters in this Letter; the modulation axis is uniform over the sphere, and the amplitude is restricted to $A \leq 0.3$. The power spectrum parameters are restricted to $0.5 \leq q \leq 1.5$ and $-0.5 \leq n \leq 0.5$. These choices are sufficiently generous to include all nonzero parts of the likelihood.

The posterior distribution,

$$P(q, n, A, \hat{\mathbf{p}}|\mathbf{d}) \propto \mathcal{L}(q, n, A, \hat{\mathbf{p}})P(q, n, A, \hat{\mathbf{p}}), \quad (8)$$

is then mapped out using a standard Markov chain Monte Carlo technique. We use a Gaussian proposal density for q , n , and A , and an Euler-matrix-based, uniform proposal density for $\hat{\mathbf{p}}$.

2.3. Bayesian Evidence and Nested Sampling

In a Bayesian analysis, one is not only interested in the set of best-fit parameter values but also in the relative probability of competing models. The most direct way of measuring this is through the Bayesian evidence,

$$E \equiv P(\mathbf{d}|H) = \int P(\mathbf{d}|\theta, H)P(\theta|H)d\theta, \quad (9)$$

which is simply the average likelihood over the prior volume. Typically, one computes this quantity for two competing models, H_0 and H_1 , and considers the difference $\Delta \log E = \log E_1 - \log E_0$. If $\Delta \log E > 1$, the evidence for H_1 is considered substantial; if $\Delta \log E > 2.5$, it is considered strong.

Traditionally, computation of evidences has been a computational challenge. However, Mukherjee et al. (2006) introduced a method called “nested sampling,” proposed by Skilling (2004) to the cosmological community, that allows for accurate estimation of the evidence through Monte Carlo sampling. We implemented this for the priors and likelihood described above and found that it works very well for the problem under consideration.

2.4. Maximum Likelihood Analysis

We also perform a standard frequentist maximum likelihood analysis by computing the maximum likelihood modulation parameters for isotropic Monte Carlo simulations. For these computations, we use a modified version of the evidence code, which we find to be considerably more robust than a simple nonlinear search; while the nonlinear search algorithms often get trapped in local minima, the nested sampling algorithm always finds the correct solution, but of course, at a considerably higher computational expense.

3. DATA

We analyze two versions of the third-year *WMAP* sky maps in the following: the template-corrected Q-, V-, and W-band maps and the “foreground-cleaned” internal linear combination (ILC) map (Hinshaw et al. 2006). All maps are processed as described by Eriksen et al. (2007): They are first downgraded to HEALPix⁷ resolution $N_{\text{side}} = 16$, by additional smoothing to a 9° FWHM Gaussian beam and appropriate pixel window. Second, uniform Gaussian noise of $\sigma_n = 1 \mu\text{K}$ is added to each pixel in order to regularize the pixel-pixel covariance matrix. This combination of smoothing and noise level results in a signal-to-noise ratio of unity at $l = 40$ and strong noise domination at the Nyquist multipole of $l = 47$.

We use two different sky cuts for our analyses. First, given that the Galactic plane is clearly visible in the single-frequency data, our first mask is conservatively defined. This cut is created by expanding the Kp2 mask (Hinshaw et al. 2006) by 9° in all

⁷ See <http://healpix.jpl.nasa.gov>.

TABLE 1
MODULATION MODEL RESULTS

Data (1)	$(l_{\text{bf}}, b_{\text{bf}})$ (deg) (2)	A_{bf} (3)	$\Delta \log \mathcal{L}$ (4)	$\Delta \log E$ (5)	P (6)
ILC ^a	(225, -27)	0.114	6.1	1.8 ± 0.2	0.991
ILC ^b	(208, -27)	0.125	6.0	1.8 ± 0.2	0.991
Q band ^b	(222, -35)	0.124	5.5	1.5 ± 0.2	0.987
V band ^b	(205, -19)	0.127	5.6	1.5 ± 0.2	0.990
W band ^b	(204, -31)	0.121	5.2	1.3 ± 0.2	0.985

NOTES.—Cols. (2) and (3): The marginal best-fit dipole axis and amplitude. Col. (4): The change in likelihood at the posterior maximum, $\Delta \log \mathcal{L} = \log \mathcal{L}_{\text{mod}} - \log \mathcal{L}_{\text{iso}}$, between the modulated and the isotropic model. Col. (5): The Bayesian evidence difference, $\Delta \log E = \log E_{\text{mod}} - \log E_{\text{iso}}$. Col. (6): The frequentist probability for obtaining a lower maximum likelihood modulation amplitude than the observed one, computed from isotropic simulations.

^a Liberal 12.8% sky cut imposed.

^b Conservative 36.3% sky cut imposed.

directions and then manually removing all near-Galactic pixels for which any difference map between two channels are clearly larger than noise. In total, 36.3% of all pixels are rejected by this cut (see Fig. 2). Second, we also adopt the directly downgraded Kp2 cut used by the *WMAP* team that removes 12.8% of all pixels. We use this mask for the ILC map only.

The noise covariance matrix is given by the uniform noise only, $N_{ij} = \sigma_n^2 \delta_{ij}$. For completeness, we have also computed the noise covariance from the smoothed instrumental noise for the V-band data, but we find that this has no effect on the final results, since its amplitude is far below the CMB signal. It is therefore omitted in the following.

As an additional hedge against foreground contamination, we marginalize over a set of fixed spatial templates, t_i , through the covariance matrix $F_i = \alpha_i t_i t_i^T$, $\alpha_i \gtrsim 10^3$. Monopole and dipole terms are always included, and one or more foreground templates. For the V-band and ILC maps, we follow Hinshaw et al. (2006) and adopt V - ILC as our foreground template. For the Q-band data, we marginalize over a synchrotron (Haslam et al. 1982), a free-free (Finkbeiner 2003), and a dust (Finkbeiner et al. 1999) template individually. Finally, for the W-band data, we use the W - ILC difference map. However, we have tried various combinations for all maps, and there is virtually no sensitivity to the particular choice, or indeed, to the template at all, due to the conservative sky cut used.

4. RESULTS

The results from the analysis outlined above are summarized in Table 1. For each map, we report the best-fit dipole axis and amplitude as well as the maximum log likelihood difference and Bayesian evidence difference for the modulated versus the isotropic model. The errors on the evidence are estimated by performing eight independent analyses for each case and computing the standard deviation (Mukherjee et al. 2006). We also compute the probability of obtaining a smaller modulation amplitude than the observed one by analyzing 1000 isotropic Monte Carlo simulations.

Starting with the first case in Table 1, the ILC map cut by a 12.8% mask, we see that the best-fit modulation axis points toward $(l, b) = (225^\circ, -27^\circ)$, and the corresponding modulation amplitude is 0.114. The raw likelihood improvement is $\Delta \log \mathcal{L} = 6.1$. The probability of finding such a high modulation amplitude in intrinsically isotropic simulations is $\sim 1\%$, and, finally, the improvement in Bayesian evidence is $\Delta \log E = 1.8$.

Furthermore, these results are not sensitive to data set or sky

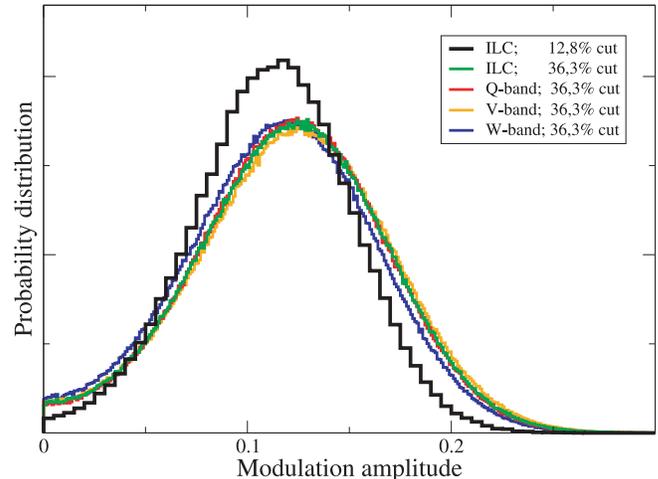


FIG. 1.—Posterior distributions for the dipole modulation amplitude, marginalized over direction and CMB power spectrum.

coverage: Even the Q-band map, which presumably is the least reliable with respect to residual foregrounds, yields a modulation amplitude that is high at the 98.7% (frequentist) confidence level and a Bayesian log-evidence improvement of 1.5. This frequency independence is further illustrated in Figure 1, where we show the marginalized posterior distributions for the modulation amplitudes for each data set. The agreement among data sets is very good.

In Figure 2, we show the dipole axis posterior distribution for the ILC map and 36.3% sky cut. Superimposed on this, we have also marked the first-year asymmetry axis reported by Eriksen et al. (2004b) $[(l, b) = (237^\circ, -10^\circ)]$ in white and also the other axes listed in Table 1. All agree well within 2σ , and this is another testimony to the excellent stability of the effect with respect to statistical method, data set, and overall procedure.

Finally, we note that this model may also partially explain the anomalous cold spot reported by Vielva et al. (2004) and Cruz et al. (2005): by demodulation, the spot would increase its temperature by about 10%, and although still very cold, it would be significantly less extreme. Similar arguments could possibly also be made for the Bianchi VII_b correlation found by Jaffe et al. (2005). These issues will be considered further in future work.

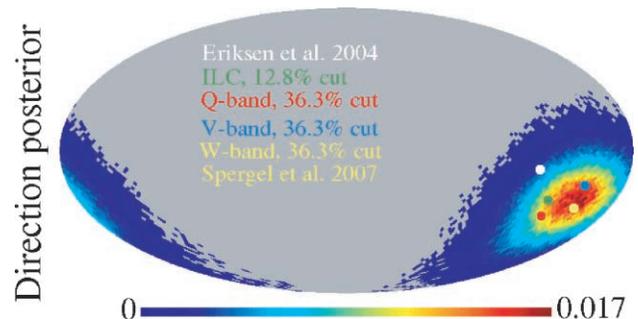


FIG. 2.—Posterior distribution for the dipole modulation axis, shown for the ILC map and 36.3% sky cut, and marginalized over power spectrum and amplitude parameters. Gray sky pixels indicate pixels outside the 2σ confidence region. The circles indicate the axis (1) reported by Eriksen et al. (2004b) in white; (2) for the ILC map with a 12.8% sky cut in green; (3) for the Q, V, and W bands in red, blue, and yellow, respectively. The axis reported by Spergel et al. (2007a) coincides with the W-band axis.

5. CONCLUSIONS

A notable power asymmetry between two opposing hemispheres in the first-year *WMAP* sky maps was reported by Eriksen et al. (2004b). This feature may be observed as strong fluctuations in the southern ecliptic hemisphere but virtually no large-scale structure in the northern ecliptic hemisphere (e.g., Hinshaw et al. 2006).

In this Letter, we have revisited this issue in the third-year *WMAP* data, adopting the statistical framework introduced and applied by Spergel et al. (2007a). With these tools, we find that the evidence for power asymmetry in the *WMAP* data is very consistent with that initially reported for the first-year maps by Eriksen et al. (2004b), and the *WMAP* data clearly suggest a dipolar distribution of power on the sky: the best-fit modulation amplitude is roughly 12% in real space, or about 20% in terms of power spectra. The corresponding dipole direction is $(l, b) \sim (225^\circ, -27^\circ)$. All results are independent of data set choices, i.e., frequency channel or sky cut.

However, the statistical evidence for this effect is still only tentative. In frequentist language, the significance is about 99%,

while in Bayesian terms, the log-evidence difference is ~ 1.5 – 1.8 , corresponding to odds of one to five or six. This is quite comparable to the evidence for $n_s \neq 1$ after the third-year *WMAP* data release, for which the odds are about one to eight in the highest case (Parkinson et al. 2006). Thus, there is still a chance that the effect may be a fluke, and most likely, this will remain the situation until *Planck* provides new data in some 5 years. With additional frequency coverage, a better job can be done on foreground treatment, and more sky coverage can be reliably included in the analysis. Second, full-sky and high-sensitivity polarization data should provide valuable insights on the origin of the effect.

H. K. E. acknowledges financial support from the Research Council of Norway. Some of the results in this Letter have been derived using the HEALPix (Górski et al. 2005) software and analysis package. We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

REFERENCES

- Ackerman, L., Carroll, S. M., & Wise, M. B. 2007, preprint (astro-ph/0701357)
 Bennett, C. L., et al. 2003, *ApJS*, 148, 1
 Bridges, M., McEwen, J. D., Lasenby, A. N., & Hobson, M. P. 2006, preprint (astro-ph/0605325)
 Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2007, *Phys. Rev. D*, 75, 023507
 Cruz, M., Martínez-González, E., Vielva, P., & Cayón, L. 2005, *MNRAS*, 356, 29
 de Oliveira-Costa, A., Tegmark, M., Zaldarriaga, M., & Hamilton, A. 2004, *Phys. Rev. D*, 69, 063516
 Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004a, *ApJ*, 612, 633
 Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004b, *ApJ*, 605, 14
 Eriksen, H. K., et al. 2007, *ApJ*, 656, 641
 Finkbeiner, D. P. 2003, *ApJS*, 146, 407
 Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
 Gordon, C. 2007, *ApJ*, 656, 636
 Gordon, C., Hu, W., Huterer, D., & Crawford, T. 2005, *Phys. Rev. D*, 72, 103002
 Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
 Hansen, F. K., Banday, A. J., & Górski, K. M. 2004, *MNRAS*, 354, 641
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. 1982, *A&AS*, 47, 1
 Hinshaw, G., et al. 2006, preprint (astro-ph/0603451)
 Inoue, K. T., & Silk, J. 2006, *ApJ*, 648, 23
 Jaffe, T. R., Banday, A. J., Eriksen, H. K., Górski, K. M., & Hansen, F. K. 2005, *ApJ*, 629, L1
 ———. 2006, *A&A*, 460, 393
 Kanno, S., & Soda, J. 2006, *Phys. Rev. D*, 74, 063505
 Land, K., & Maguiejo, J. 2007, *MNRAS*, in press (astro-ph/0611518)
 Martínez-González, E., Cruz, M., Cayón, L., & Vielva, P. 2006, *NewA Rev.*, 50, 875
 Moffat, J. W. 2005, *J. Cosmol. Astropart. Phys.*, 10, 12
 Mukherjee, P., Parkinson, D., & Liddle, A. R. 2006, *ApJ*, 638, L51
 Parkinson, D., Mukherjee, P., & Liddle, A. R. 2006, *Phys. Rev. D*, 73, 123523
 Skilling, J. 2004, in *AIP Conf. Proc.* 735, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. von Toussaint (New York: AIP), 395
 Spergel, D. N., et al. 2007a, preprint (astro-ph/0603449)
 ———. 2007b, *ApJ*, in press
 Tomita, K. 2005, *Phys. Rev. D*, 72, 043526
 Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., & Cayón, L. 2004, *ApJ*, 609, 22