Large-scale alignments from *WMAP* and *Planck*

Craig J. Copi,[1][*] Dragan Huterer,[2] Dominik J. Schwarz[3] and Glenn D. Starkman[1,4]

[1]*CERCA/Department of Physics/ISO, Case Western Reserve University, Cleveland, OH 44106-7079, USA*
[2]*Department of Physics, University of Michigan, 450 Church St, Ann Arbor, MI 48109-1040, USA*
[3]*Fakultät für Physik, Universität Bielefeld, Postfach 100131, D-33501 Bielefeld, Germany*
[4]*Physics Department, Theory Unit, CERN, CH-1211 Genève 23, Switzerland*

ABSTRACT

We revisit the alignments of the largest structures observed in the cosmic microwave background (CMB) using the seven and nine-year *Wilkinson Microwave Anisotropy Probe* (*WMAP*) and first-year *Planck* data releases. The observed alignments – the quadrupole with the octopole and their joint alignment with the direction of our motion with respect to the CMB (the dipole direction) and the geometry of the Solar system (defined by the Ecliptic plane) – are generally in good agreement with results from the previous *WMAP* data releases. However, a closer look at full-sky data on the largest scales reveals discrepancies between the earlier *WMAP* data releases (three to seven-year) and the final, nine-year release. There are also discrepancies between all the *WMAP* data releases and the first-year *Planck* release. Nevertheless, both the *WMAP* and *Planck* data confirm the alignments of the largest observable CMB modes in the Universe. In particular, the *p*-values for the mutual alignment between the quadrupole and octopole, and the alignment of the plane defined by the two with the dipole direction, are both at the greater than 3-sigma level for all three *Planck* maps studied. We also calculate conditional statistics on the various alignments and find that it is currently difficult to unambiguously identify a leading anomaly that causes the others or even to distinguish correlation from causation.

Key words: cosmic background radiation – large-scale structure of Universe.

1 INTRODUCTION

Cosmic microwave background (CMB) maps produced by the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) and, most recently, the *Planck* collaborations provide an unprecedented view into the physics of the early Universe. These data help constrain the parameters of the standard cosmological model, Λ cold dark matter (ΛCDM), to the per cent level accuracy. They also point out some intriguing anomalies, particularly on the largest angular scales or at the lowest multipole moments. One anomaly, and the one of interest for this work, is alignments. The quadrupole and octopole are found to be mutually aligned and they define axes that are unusually perpendicular to the Ecliptic pole and parallel to the direction of our motion with respect to the rest frame of the CMB (the dipole direction). Another anomaly, the lack of correlations on large angular scales, is the subject of a companion work (Copi et al. 2013). For reviews of the large-angle anomalies *before* the final *WMAP* and the first cosmological *Planck* data releases, see Bennett et al. (2011) and Copi et al. (2010).

The study of alignments requires precise measurements of the full-sky CMB. Conventional wisdom tells us that the largest scales in the Universe should be well measured. The basic argument is straightforward: measurements by *WMAP* and *Planck* are signal dominated on small angular scales. Large numbers of these measurements are averaged to determine the large-scale structure of the CMB which further reduces the noise leading to very precise determinations. Thus, even though cosmic variance is large and our Universe could have been drawn from a broadly distributed ensemble, our particular realization can be very well measured. Both *WMAP* and *Planck* are measuring the same CMB sky so should agree very well on the largest scales.

Unfortunately, the real-world is more complicated. Though the statistical noise is small, residual contamination due to foregrounds, even after cleaning the maps, remains a source of significant uncertainty. These issues are of paramount importance in the study of alignments. In the nine-year data release, the *WMAP* team states: 'We conclude that our ability to remove foregrounds is the limiting factor in measurement of the cosmological quadrupole+octopole alignment' (Bennett et al. 2013). They then argue that the statistical significance of the alignment is weakened due to this. In a similar context the *Planck* team states: 'Residual foregrounds (mostly on the Galactic plane) present in the four *Planck* CMB estimates could

* E-mail: cjc5@cwru.edu

influence the reconstruction of the low-order multipoles' (Planck Collaboration XXIII 2014). (Note that their four maps include the `Commander-Ruler` map in addition to the three we discuss below.) They proceed to Wiener filter their maps to further reduce contamination which has a small effect on alignments. This observation is not new: for example, Chiang, Naselsky & Coles (2007) found signs of residual foregrounds in the *WMAP* three-year data release on these large scales. Such residual contaminations persist in the data to the present time.

The goal of this paper is to study these large-angle alignments with the *Planck* one-year data and to compare them to the *WMAP* seven and nine-year data. The *Planck* collaboration included a brief discussion of alignments in their extensive study of isotropy in the CMB (Planck Collaboration XXIII 2014). In their presentation, they did not include results from the principal mathematical tool employed here – the multipole vectors; moreover, we use a somewhat different approach to generate the full-sky maps, and use a much larger set of Monte Carlo simulations to obtain the statistical inferences. This work therefore complements the existing detailed study of other tests of isotropy, such as the hemispherical power asymmetries and moments of the temperature field (Planck Collaboration XXIII 2014).

To study the large-angle alignments, it is necessary to define statistics and assign significance to the data based on them. In recent years, the trend has been towards applying a Bayesian analysis to all statistical questions. Broadly speaking, the Bayesian approach is designed to compare models and for parameter estimation within a model. Even the problem of null hypothesis testing in the Bayesian approach is reduced to a model comparison; that between the full model and a subset of the full model with a restricted parameter set, some parameters fixed, etc. When there is a model with no serious competitors, such as in cosmology with ΛCDM, Bayesian statistics struggles to even ask the question of the consistency between the model and the data (though see Starkman, Trotta & Vaudrevange 2008 for a possible, if computationally challenging, approach). At the present time in cosmology, *there are no compelling alternative models* that can account for the anomalies. Clever ideas have been proposed to explain some of the anomalies (e.g. Frisch 2005; Gordon et al. 2005; Alnes & Amarzguioui 2006; Inoue & Silk 2006; Rakić, Räsänen & Schwarz 2006; Pullen & Kamionkowski 2007; Dikarev et al. 2008; Ramirez & Schwarz 2009; Peiris & Smith 2010), but no model that explains all, or even most, of them exists. It is not even clear whether the origin of the anomalies is cosmological, astrophysical foregrounds, systematic (instrumental, map making, etc.), or simply statistical, although it could be argued that since the *Planck* satellite and data reduction is very different from that provided by the *WMAP* satellite, systematic effects are unlikely to explain the existence of shared anomalies. Due to these issues, we adopt the frequentist approach consistent with that used in previous work (see Copi et al. 2006, for example). The frequentist approach is well suited for this specific problem – to address the question of tension between the model and the data, and, if there is one, where this tension lies. This allows the data to point the way towards the source of any potential discrepancy independent of finding a better model to describe it.

This paper is organized as follows. In Section 2, we take a first look at the large-angle *Planck* data and compare the angular power spectrum at the lowest multipoles between *Planck* and *WMAP* releases. In Section 3, we describe the methodology of how we arrive at an ensemble of full-sky maps using harmonic inpainting, and how we correct for the effect of our motion through the CMB rest frame on the quadrupole. In Section 4, we describe the statistics that

Table 1. The power spectrum coefficients, D_ℓ, in units of μK^2 as reported by *WMAP* and *Planck*. All values are based on a maximum likelihood estimator. Since the one-year *WMAP* reported values were based on pseudo-C_ℓ estimators, they have been excluded from this table. The $S_{1/2}$ values have been computed for $\ell_{max} = 100$ unless otherwise stated.

Data release	D_2	D_3	D_4	D_5	$S_{1/2}$ (μK^4)
WMAP 3yr	211	1041	731	1521	8330
WMAP 5yr	213	1039	674	1527	8915
WMAP 7yr	201	1051	694	1517	8938
WMAP 9yr	151	902	730	1468	5797
Planck R1	299	1007	646	1284	8035\rlap[a]

Note. [a]This $S_{1/2}$ has been calculated for $\ell_{max} = 49$ since *Planck* only provides binned values for $\ell \geq 50$.

we use (and that we developed previously for our *WMAP* analyses), and in Section 5, we carry out a frequentist analysis to quantify the various alignments. We conclude in Section 6.

2 STATE OF LARGE-SCALE DATA

Since full-sky data are required for a study of alignments, it is important to understand its current state. A high-level view of the data can be obtained through the power spectrum coefficients

$$D_\ell \equiv \frac{\ell(\ell+1)}{2\pi} C_\ell. \tag{1}$$

The results reported by *WMAP* and *Planck* are summarized in Table 1.[1] In all cases, these values are given in μK^2 and are based on a maximum likelihood estimator. Since the one-year *WMAP* analysis employed a pseudo-C_ℓ-based estimator, it has not been included in the table.

The $S_{1/2}$ statistic defined in the one-year *WMAP* data release (Spergel et al. 2003) is a convenient and discriminating tool to quantify the lack of correlations on large angular scales. The statistic can be calculated as

$$S_{1/2} \equiv \int_{-1}^{1/2} [C(\theta)]^2 d(\cos\theta) = \sum_{\ell=2}^{\ell_{max}} C_\ell I_{\ell\ell'} C_{\ell'}, \tag{2}$$

where the $I_{\ell\ell'}$ are components of an easily calculated matrix. For a more thorough discussion of the $S_{1/2}$ from the *Planck* data see Copi et al. (2013). In Table 1, the $S_{1/2}$ has been calculated for $\ell_{max} = 100$ except for *Planck* which only provides binned C_ℓ for $\ell \geq 50$.

The noise only contribution to the uncertainty in the values in Table 1 is estimated in *WMAP* from the Fisher matrix to be $\sigma_{D_\ell} \sim 10 \, \mu K^2$ in all cases; typically slightly larger in the earlier data releases and slightly smaller in the later ones. This sets the scale for the expected statistical scatter in the data. From the table, we see that the three-year through seven-year *WMAP* data releases are in good agreement and provide a consistent picture of the large-scale Universe. The differences are most likely due to systematic analysis improvements such as beam-shape determination, point-source identification, and masking. Quite surprisingly the nine-year *WMAP* data release provides a markedly different view of the large-scale Universe and even more surprisingly the first *Planck* data

Table 2. The power spectrum coefficients, D_ℓ, in units of μK^2 extracted from cleaned, full-sky maps provided by *WMAP* and *Planck*.

Map	D_2	D_3	D_4	D_5	$S_{1/2}$ (μK^4)
WMAP ILC 1yr	195	1053	834	1667	8190
WMAP ILC 3yr	248	1051	756	1588	8476
WMAP ILC 5yr	243	1052	730	1591	8642
WMAP ILC 7yr	240	1048	731	1593	8528
WMAP ILC 9yr	243	1013	709	1612	8156
Planck NILC	209	863	704	1379	4816
Planck SEVEM	205	798	736	1207	3766
Planck SMICA	239	925	713	1494	6309

release provides yet another different view. In particular, comparing the quadrupole, D_2, we see that the nine-year *WMAP* and first-year *Planck* results differ by a factor of 2, or roughly 15 times the approximate noise error!

An alternative view of the large-scale Universe comes from the foreground cleaned, full-sky maps provided in the data releases (we restrict our analysis to the maps officially released by the *WMAP* and *Planck* collaborations). The D_ℓ extracted from these maps are shown in Table 2. For *WMAP*, the maps are cleaned using the Internal Linear Combination (ILC) method (Eriksen et al. 2004; Hinshaw et al. 2007). For *Planck* multiple cleaning procedures were used and three, full resolution maps based on them, the SMICA, NILC, and SEVEM maps, have been provided (Planck Collaboration XII 2014). Furthermore, for the SMICA and NILC maps a small *inpainting mask* with $f_{sky} \approx 0.97$ has been defined by the *Planck* team. This small region has been inpainted using a constrained realization and it is these inpainted maps that were analysed for the values provided in the table. From these results, we see that the *WMAP* full-sky maps from all data releases are in reasonably good mutual agreement. The one-year map is somewhat discrepant and the nine-year map is also somewhat different in its $S_{1/2}$ value, but otherwise the cleaning procedure employed throughout the years has produced relatively stable results. For the *Planck* maps this is not the case. Though the cleaning procedures do produce quadrupoles more in line with the one from *WMAP*, there is still a large discrepancy between the large-scale structure internally among the *Planck* maps and in comparison with *WMAP*. Moreover, these three maps were not the only cleaned maps produced, just the best three (four including also the Commander-Ruler method), selected by the *Planck* team, based on unpublished criteria. The larger spread of results from the *Planck* maps when compared to the *WMAP* ILC maps of different years does not come as a surprise. While the *WMAP* collaboration published similar incarnations of the same ILC method over the years, the three analysed *Planck* maps are based on (radically) different assumptions about the foreground.

Naturally many details have been ignored in these comparisons. For example, a visual examination of the SEVEM map shows clear signs of small-scale contaminations near the Galactic plane. This region is inpainted in the SMICA and NILC maps and suggests that we should be careful interpreting results based on the SEVEM map. Also, for all of the maps larger masks, called *validity masks*, have been defined by the *Planck* team to define a more conservative (i.e. smaller) set of pixels that are believed to be clean. These masks have sky fractions of $f_{sky} = 0.88$ for the SMICA map, $f_{sky} = 0.93$ for the NILC map, and $f_{sky} = 0.76$ for the SEVEM map. As explained below, we will mostly use the inpainting masks, except for the SEVEM map for which we will use the validity mask.

3 FULL-SKY MAPS

Perhaps contrary to our conventional wisdom, the largest scales in the CMB have been measured precisely but not (yet) necessarily accurately; this is immediately clear by inspecting the angular power as represented in Tables 1 and 2. Therefore, care must be taken when analysing full-sky data and when interpreting the results. In this section, we describe our methodology for arriving at an ensemble of maps that represent the *Planck* full-sky. The two particular techniques that we will apply are inpainting and Doppler correction for the quadrupole.

For the purpose of studying alignments in ΛCDM cosmologies, all that matters is the phase structure of the $a_{\ell m}$ – i.e. the relative values of $a_{\ell m}$ of equal ℓ and all m. The statistics used to study alignments are constructed so as to be independent of the power spectrum. Thus, any power spectrum could be used to produce realizations for studying the distributions of these statistics, provided the $a_{\ell m}$ have random phases, and the underlying distribution is statistically isotropic. Despite this, some of the discussions below will refer to the magnitudes of the $a_{\ell m}$. To be consistent with our companion analysis of the two-point angular correlation function (Copi et al. 2013), we use the *Planck* base best-fitting power spectrum. This power spectrum is freely available along with all the *Planck* data and is based on the parameters listed in the first column of table 2 in Planck Collaboration XVI (2014).

3.1 Inpainting

One approach to handling particularly contaminated regions of a map is to excise those regions and to replace them with synthetic data constrained to agree with the uncontaminated regions of the map. In general, masking the skies introduces coupling among the modes so this must be used with some care. The relationship between the uncontaminated partial map and a representation of the full-sky information (either as a map or spherical harmonic coefficients, $a_{\ell m}$) is an underdetermined linear system. Due to this there is not a unique solution and extra information is required to re-fill the contaminated regions. One set of techniques to achieve this re-filling, that is the extra assumptions imposed and the balance between these assumptions and the type and quality of fit to the uncontaminated regions, is known as inpainting. Inpainting has been studied in numerous contexts and many approaches have been explored for CMB maps (see Inoue, Cabella & Komatsu 2008 and especially Starck, Fadili & Rassat 2013 for a discussion of techniques in the context of low-ℓ reconstruction). The *Planck* analysis has implemented a particular form of inpainting (Benoit-Lévy et al. 2013; Planck Collaboration XII 2014) used in the SMICA and NILC maps discussed above. For other related work on the topic, see Abrial et al. (2008); Dupé et al. (2011); Bucher & Louis (2012); Nishizawa & Inoue (2013).

Inpainting is not a simple cure for contamination. The resulting full-sky information will depend on the extra assumptions (type of inpainting performed) and also on the mask employed. Choosing a large mask will remove the most contamination but will also lead to a large variance, poor determination, of the full-sky information. In this work, our goal is to study the large-scale anomalies present in the full-sky data independent of their origin. If we were interested in only a cosmological origin for alignments, we would restrict the analysis to the cleanest portions of the sky by employing a large mask. A large mask washes out the results making definitive statements about the alignments impossible. On the other hand, employing a small mask allows for definitive statements about the

presence of alignments but reduces the ability to determine their origin. Since our focus is on the existence of alignments in the data as provided, our main results will be based on inpainting using the aforementioned smaller inpainting masks with $f_{\text{sky}} \approx 0.97$.

3.2 Harmonic inpainting

Harmonic inpainting as applied to the CMB (Inoue et al. 2008; Kim, Naselsky & Mandolesi 2012) is an application of constrained Gaussian realizations for reconstructing the $a_{\ell m}$ from a masked sky. It *assumes* the usual statistical properties, Gaussianity and isotropy, but also produces probability distributions which can be used in subsequent analyses. Here we mainly follow Kim et al. (2012), but have changed the notation slightly. The technique is intimately related to pseudo-$a_{\ell m}$ reconstruction as will be evident in what follows.

3.2.1 Algorithm

Consider a temperature map, $T(\hat{\boldsymbol{e}})$, and a mask, $W(\hat{\boldsymbol{e}})$. Here, and in what follows, $\hat{\boldsymbol{e}}$ represents the usual radial unit vector that in Cartesian coordinates may be written as

$$\hat{\boldsymbol{e}} = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta). \tag{3}$$

These quantities may be expanded in spherical harmonics as

$$T(\hat{\boldsymbol{e}}) = \sum_{\ell'm'} a_{\ell'm'} Y_{\ell'm'}(\hat{\boldsymbol{e}}), \quad W(\hat{\boldsymbol{e}}) = \sum_{\tilde{\ell}\tilde{m}} w_{\tilde{\ell}\tilde{m}} Y_{\tilde{\ell}\tilde{m}}(\hat{\boldsymbol{e}}). \tag{4}$$

We may also expand the masked sky in terms of the pseudo-$a_{\ell m}$, denoted by $\tilde{a}_{\ell m}$, and relate this to the previous expansions by

$$\begin{aligned} W(\hat{\boldsymbol{e}})T(\hat{\boldsymbol{e}}) &= \sum_{\ell m} \tilde{a}_{\ell m} Y_{\ell m}(\hat{\boldsymbol{e}}) \\ &= \sum_{\tilde{\ell}\tilde{m}} w_{\tilde{\ell}\tilde{m}} Y_{\tilde{\ell}\tilde{m}}(\hat{\boldsymbol{e}}) \sum_{\ell'm'} a_{\ell'm'} Y_{\ell'm'}(\hat{\boldsymbol{e}}). \end{aligned} \tag{5}$$

This is a standard procedure: we solve for the pseudo-$a_{\ell m}$ by integrating and using the properties of the spherical harmonics. This leads to an integral over three spherical harmonics which may be evaluated in terms of the Gaunt coefficients and most easily represented by the Wigner $3J$ symbols. The end result is

$$\tilde{a}_{\ell m} = \sum_{\ell'm'} F_{\ell m;\ell'm'} a_{\ell'm'}, \tag{6}$$

where

$$\begin{aligned} F_{\ell m;\ell'm'} &\equiv (-1)^m \sum_{\tilde{\ell}\tilde{m}} w_{\tilde{\ell}\tilde{m}} \sqrt{\frac{(2\ell+1)(2\ell'+1)(2\tilde{\ell}+1)}{4\pi}} \\ &\quad \times \begin{pmatrix} \ell' & \tilde{\ell} & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell' & \tilde{\ell} & \ell \\ m' & \tilde{m} & -m \end{pmatrix}. \end{aligned} \tag{7}$$

More compactly we write this in matrix notation as

$$\tilde{\boldsymbol{a}} = \mathbf{F}\boldsymbol{a}. \tag{8}$$

Moreover, we define the following matrices which are simplified using the statistical isotropy of the underlying $a_{\ell m}$,

$$\mathbf{b} \equiv \langle \boldsymbol{a}\tilde{\boldsymbol{a}}^\dagger \rangle = \mathbf{C}_a \mathbf{F}^\dagger, \quad \mathbf{C} \equiv \langle \tilde{\boldsymbol{a}}\tilde{\boldsymbol{a}}^\dagger \rangle = \mathbf{F}\mathbf{C}_a \mathbf{F}^\dagger, \tag{9}$$

where the angle brackets represent an ensemble average and $\mathbf{C}_a$ represents the diagonal matrix of the C_ℓ^{th} from the best-fitting ΛCDM model, $\mathbf{C}_a \equiv \langle \boldsymbol{a}\boldsymbol{a}^\dagger \rangle$.

For a given map and mask, we calculate $\mathbf{F}$ and $\tilde{\boldsymbol{a}}$. From these, harmonic inpainting proceeds by generating realizations of ΛCDM as $a_{\ell m}^{\text{mc}}$, that is, the $a_{\ell m}^{\text{mc}}$ are Gaussian random variables with $|a_{\ell m}^{\text{mc}}| \in N(0, C_\ell^{\text{th}})$ and random phases. These are then corrected giving us constrained realizations as

$$\boldsymbol{a}^{\text{inp}} = \boldsymbol{a}^{\text{mc}} + \mathbf{b}\mathbf{C}^{-1}(\tilde{\boldsymbol{a}} - \mathbf{F}\boldsymbol{a}^{\text{mc}}). \tag{10}$$

We now explain how to efficiently calculate the matrix $\mathbf{C}^{-1}$.

3.2.2 Inversion and marginalization

Calculating the matrix $\mathbf{C}^{-1}$ is complicated by the fact that the underlying map which we are trying to inpaint may contain a monopole and dipole (on the full- or cut-sky). We wish to remove these contributions. This can be done in a number of ways; we will marginalize over them using a method based on that of Slosar & Seljak (2004) with some modifications.

To begin we construct the rectangular coupling matrix, $\mathbf{F}$, as above using the best-fitting ΛCDM C_ℓ^{th} with $C_0^{\text{th}} = C_1^{\text{th}} = 0$. Let $\mathbf{f}$ represent the columns of $\mathbf{F}$ corresponding to the monopole and dipole (typically the first four columns) and $\hat{\mathbf{F}}$ be the remainder of the columns. With this we may marginalize over the monopole and dipole in the correlation matrix by first writing it as

$$\mathbf{C} = \hat{\mathbf{F}}\mathbf{C}_a \hat{\mathbf{F}}^\dagger + \lambda \mathbf{f}\mathbf{f}^\dagger \equiv \hat{\mathbf{C}} + \lambda \mathbf{f}\mathbf{f}^\dagger, \tag{11}$$

where λ is a large parameter, preferably $\lambda \to \infty$. This can be accomplished analytically when calculating the inverse using a result from Rybicki & Press (1992),

$$\begin{aligned} \mathbf{C}^{-1} &= \lim_{\lambda \to \infty} (\hat{\mathbf{C}} + \lambda \mathbf{f}\mathbf{f}^\dagger)^{-1} \\ &= \lim_{\lambda \to \infty} [\hat{\mathbf{C}}^{-1} - \hat{\mathbf{C}}^{-1}\mathbf{f}(\lambda^{-1} + \mathbf{f}^\dagger \hat{\mathbf{C}}^{-1}\mathbf{f})^{-1}\mathbf{f}^\dagger \hat{\mathbf{C}}^{-1}] \\ &= \hat{\mathbf{C}}^{-1} - \hat{\mathbf{C}}^{-1}\mathbf{f}(\mathbf{f}^\dagger \hat{\mathbf{C}}^{-1}\mathbf{f})^{-1}\mathbf{f}^\dagger \hat{\mathbf{C}}^{-1}. \end{aligned} \tag{12}$$

This last line is now easy to calculate. Since $\hat{\mathbf{C}}$ is a correlation matrix it is Hermitian and positive definite so may be efficiently inverted using a Cholesky decomposition. We have an extra matrix we need to invert, the one in parenthesis in the last line of the previous equation. However, this is a 4×4 correlation matrix so is also easy to handle.

Initially, adoption of harmonic inpainting seems risky since it imposes the assumptions of Gaussianity and statistical isotropy on the full-sky $a_{\ell m}$. However, this is outweighed by clear advantages: the inpainting generates statistical distributions of the full-sky $a_{\ell m}$ thus allowing accurate statements for any test that rests on the full CMB sky, and it is straightforward to implement. Crucially, since we have restricted our main analysis to masks with $f_{\text{sky}} \approx 0.97$, the results will be insensitive to the inpainting assumptions.

3.2.3 Results

Two scales appear in harmonic inpainting: we denote by L_{inp} the maximum multipole reconstructed by the inpainting and by L_{max} the maximum multipole used in the reconstruction (i.e. the largest coupling multipole). In other words, the coupling matrix $\mathbf{F}$ defined in equation (7) is a rectangular matrix of size $(L_{\text{inp}} + 1)^2 \times (L_{\text{max}} + 1)^2$. For the work performed here we are only interested in the quadrupole and octopole and (mainly) in the SMICA and NILC maps with their inpainting masks with $f_{\text{sky}} \approx 0.97$. We study the sensitivity of harmonic inpainting with these constraints in mind.

Figure 1. Histograms of the a_{2m} in μK from 5000 harmonic inpaintings of the SMICA map using its inpainting mask for various choices of scales $L_{inp} \times L_{max}$. The histograms have been shifted by the average from the 50 $\times$ 100 case so are centred roughly around zero (and precisely on zero for the black curve). These histograms show little sensitivity in the choice of $L_{inp} \times L_{max}$ and that the a_{2m} are well reconstructed, typically to within about 1 μK.

To study the sensitivity to choices of L_{inp} and L_{max}, we begin with the SMICA map. Figs 1 and 2 show histograms from 5000 reconstructions for various choices of $L_{inp} \times L_{max}$. We have chosen the $L_{inp} \times L_{max} = 50 \times 100$ case as the standard. The histograms in these figures are shifted by the average value from this (50×100) case. The x-axis in the histograms is in units of μK and shows that the $a_{\ell m}$ are reconstructed to within about 1 μK in all cases. This is not surprising: given the small fraction of the sky inpainted, it is expected that **F** is diagonally dominant; the resulting inpainting should not be sensitive to the choice of scales (provided they are large enough) and the $a_{\ell m}$ should be well determined.

With the choice $L_{inp} \times L_{max} = 50 \times 100$, we next compare inpainting of the various *Planck* maps. These histograms based on 5000 reconstructions are shown in Figs 3 and 4. For the SMICA and NILC maps their inpainting masks with $f_{sky} \approx 0.97$ have been used, while for the SEVEM map an inpainting mask was not available, so we used its validity mask with $f_{sky} = 0.76$ instead. For this reason, the distributions for the SEVEM map are significantly broader than the other two. For the most part, the three maps are in very good agreement though there are notable exceptions. In particular, a_{21}^{im} and a_{31}^{re} show significant shifts as compared to their widths between the SMICA and NILC maps. In even more cases, the SMICA and NILC $a_{\ell m}$ are in the tails of the SEVEM distribution. These results provide a graphical view of the levels and properties of residual contaminations in the maps.

3.3 Doppler quadrupole

The study of the CMB temperature anisotropies typically begins with the quadrupole and proceeds to higher multipoles, smaller
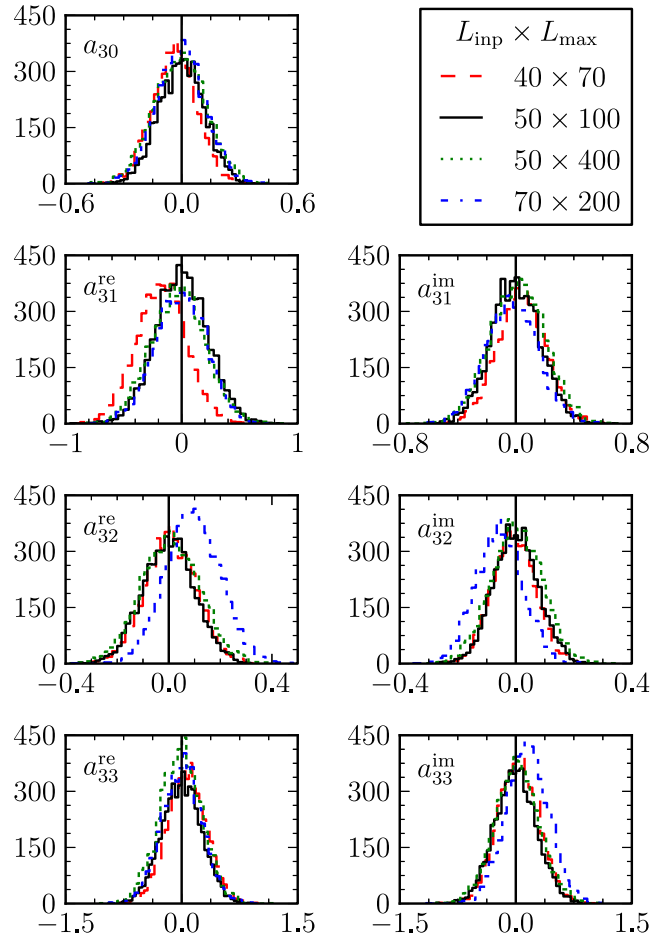

Figure 2. Same as Fig. 1, now for the a_{3m}.

scales. The monopole is not included since its magnitude is not predicted by the theory. The dipole is subtracted since our motion through the Universe with respect to the CMB is at a speed $\beta = v/c \sim 10^{-3}$, whereas the fluctuations are $\Delta T/T \sim 10^{-5}$ so the Doppler dipole is about two orders of magnitude larger than the expected cosmological CMB dipole. A Doppler quadrupole (DQ) – effect of the Sun's proper motion on the quadrupole – is also induced and expected to have a magnitude $\mathcal{O}(\beta^2) \sim 10^{-6}$. Though this is small, it is not negligible, especially in our Universe which has a small cosmological quadrupole and when properties of the maps are being studied.

The CMB temperature (monopole) has been determined from the nearly perfect blackbody to be $T_0 = 2.7255 \pm 0.0006$ K (Fixsen 2009). Our direction through the Universe in Galactic coordinates is $(l, b) = (263°.99 \pm 0°.14, 48°.26 \pm 0°.03)$ with a speed $\beta = (1.231 \pm 0.003) \times 10^{-3}$ (Hinshaw et al. 2009). These values can be compared to those in table 3 of Planck Collaboration V (2014). Note that there is an additional contribution to the speed due to the velocity of the satellite with respect to the Sun. This introduces up to about a 10 per cent variation in β at a non-constant direction with respect to the CMB. This contribution has not been included in the following discussion.

With these values, the Doppler contribution to the quadrupole from the CMB monopole may be calculated as (Peebles & Wilkinson 1968; Kamionkowski & Knox 2003)

$$a_{2m}^{DQ} = T_0 \int [\gamma(1 - \boldsymbol{\beta} \cdot \hat{\boldsymbol{e}})]^{-1} Y_{2m}^*(\hat{\boldsymbol{e}}) \, d\hat{\boldsymbol{e}}, \qquad (13)$$

Figure 3. Histograms of the a_{2m} in μK from 5000 harmonic inpainting of the *Planck* maps for the scales $L_{inp} \times L_{max} = 50 \times 100$. For the SMICA and NILC maps, the areas covered by their inpainting masks with $f_{sky} = 0.97$ have been inpainted. For the SEVEM map, its validity mask with $f_{sky} = 0.76$ has been used. From these histograms, we see that the a_{2m} are mostly consistent among the maps, see the text for details.

Figure 4. Same as Fig. 3, now for the a_{3m}.

where $\gamma \equiv (1 - |\boldsymbol{\beta}|^2)^{-1/2}$. The numerical results of this calculation along with the values from the SMICA map are given in Table 3. We see that while the DQ correction is generally small it is not negligible and in the case of the imaginary part of a_{21} the correction is comparable to the uncorrected value! This has important consequences for any analysis involving the a_{2m}, such as the alignments considered here.

Even higher multipole moments are also induced due to our motion but these are further suppressed by powers of β and need not be considered. For example, the Doppler octopole should have a magnitude $\mathcal{O}(\beta^3) \sim 10^{-9}$. Direct calculation shows that for all components the Doppler octopole correction has a magnitude less than about 10^{-4} of the observed octopole.

In principle, the DQ can be corrected during data reduction. This is naturally accomplished when calibrating off the anisotropies induced by our motion with respect to the rest frame of the CMB (typically referred to as calibrating off the Doppler dipole). For a blackbody with temperature T_0 and motion in the direction $\hat{e}$ with time-dependent velocity $\boldsymbol{\beta}(t)$, the induced anisotropies are

$$\Delta T(\hat{e}, t) = \left[\frac{1}{\gamma(t)(1 - \boldsymbol{\beta}(t) \cdot \hat{e})} - 1\right] T_0. \quad (14)$$

For *Planck*, this expression *was* used to calibrate the low-frequency instrument (Planck Collaboration V 2014). Surprisingly, for the high-frequency instrument the classical approximation for the dipole anisotropy

$$\Delta T(\hat{e}) \approx (\boldsymbol{\beta} \cdot \hat{e}) T_0 \quad (15)$$

was employed (Planck Collaboration VIII 2014). *WMAP* also used this classical approximation (Hinshaw et al. 2003). This means that the *WMAP* ILC maps require the DQ correction and that it is less

clear how to handle the *Planck* full-sky maps. On the one hand, most of the individual frequency band data combined to produce the cleaned full-sky maps have not been DQ corrected (the high frequency bands). On the other hand, one of the cleanest frequency bands, 70 GHz, has been DQ corrected. Given the current state, we have applied the DQ correction to all the cleaned full-sky *Planck* maps.

Subsequent to our original analysis, Planck Collaboration XXIII (2014) was updated to include estimates of DQ correction factors for each of their released combined maps. Unfortunately neither a complete description of how these correction factors were calculated nor all the data required to calculate them were made publicly available. Given the importance of the DQ correction to alignment results we have recalculated some of the statistics as discussed in Section 5.2. We find that the *Planck* DQ correction factors *strengthen* the alignments making the data *less consistent* with ΛCDM. This further shows the importance of applying the DQ correction and having an

Table 3. The spherical harmonic coefficients for the quadrupole from the *Planck* inpainted SMICA map and the Doppler quadrupole (DQ). The a_{2m} values are given in μK. The corrected a_{2m} are calculated by subtracting the first two rows.

Source	a_{20}	a_{21}	a_{22}
Planck SMICA	13.09	$-1.53 + 2.50i$	$-15.50 - 17.09i$
DQ correction	1.46	$0.28 - 2.64i$	$-11.16 - 10.25i$
SMICA corrected	11.63	$-1.81 + 5.14i$	$-14.34 - 16.84i$

accurate estimate of it. Due to the uncertainty in the calculation of the factors we will continue to apply our simple DQ correction as discussed above and again affirm that the significance of the anomalous alignments is sensitive to details such as the DQ correction.

3.4 Large-scale data summary

The comparisons presented here provide a broad view of the state of the existing CMB temperature data on the largest scales. There are many other comparisons that can be made (see Frejsel, Hansen & Liu 2013; Kovács, Carron & Szapudi 2013, for example). On the largest scales, we see that there exist significant discrepancies both between the *WMAP* and *Planck* maps and even internally among the maps produced by the *WMAP* team in different releases and among the maps produced using the various cleaning methods employed by *Planck*. This suggests that one must be careful about assigning a cosmological origin to any large-scale result based on full skies.

Despite the caution that must be exercised when interpreting large-scale, full-sky results, it is important to pursue such work. The discovery of alignments, or any anomaly, that persist through multiple full-sky maps is striking given the different instruments, systematics, cleaning procedures, etc. It strongly suggests that there is at least some fundamental origin to them. The extent to which this origin is cosmological and the statistical significance of such an identification is difficult to determine. Even if the origin of the alignments is ultimately determined to be an as-yet unidentified systematic or an unsubtracted (or mis-subtracted) foreground, it is important to characterize the properties of the anomalies. Only in this way can their cause be isolated and (if it is not cosmological) the data processing pipeline thereby improved, resulting in higher quality large-scale data.

4 ALIGNMENT STATISTICS

The planarity of the octopole and alignment of the quadrupole and octopole in the CMB as observed by *WMAP* was first studied by de Oliveira-Costa et al. (2004) through the maximum angular momentum dispersion. This statistic has subsequently been applied to the *Planck* data (Planck Collaboration XXIII 2014). A more complete picture of CMB alignments is obtained through the use of the multipole vectors Copi, Huterer & Starkman (2004). Here, we study both of these approaches.

4.1 Maximum angular momentum dispersion

The angular momentum dispersion about an axis $\hat{\boldsymbol{n}}$ is defined by de Oliveira-Costa et al. (2004) as

$$[\Delta L(\hat{\boldsymbol{n}})]_\ell^2 \equiv \sum_m m^2 |a_{\ell m}(\hat{\boldsymbol{n}})|^2. \tag{16}$$

Heuristically, $[\Delta L(\hat{\boldsymbol{n}})]_\ell^2$ measures the amount of planarity of structure in the multipole ℓ around an axis $\hat{\boldsymbol{n}}$. Notice that this is not a rotationally invariant quantity and its value is maximized around some axis, $\hat{\boldsymbol{n}}_\ell$. Having found such axes for the quadrupole and octopole, we quantify their alignment through their dot product, $|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$.

To determine the axes $\hat{\boldsymbol{n}}_\ell$ without applying brute force numerical rotations of maps at some resolution, we follow Copi et al. (2006) and use the known rotational properties of the $a_{\ell m}$. Here we briefly review the procedure. Under rotations the spherical harmonic coefficients transform as $\boldsymbol{a}'_\ell = \mathbf{D}^\dagger \boldsymbol{a}_\ell$. It is most convenient to represent

the rotation in terms of the Wigner rotation matrices which, for Euler angles α, β, and γ, may be written as (Edmonds 1960)

$$D_{m'm}^{(\ell)}(\alpha\,\beta\,\gamma) = e^{im'\gamma} d_{m'm}^{(\ell)}(\beta) e^{im\alpha}. \tag{17}$$

Here the $d_{m'm}^{(\ell)}(\beta)$ are the reduced Wigner matrices. From explicit calculation, we find that

$$
\begin{aligned}
(\Delta L)_\ell^2 &= \sum_{m'm''} a_{\ell m'}^* a_{\ell m''} e^{i(m'-m'')\gamma} \\
&\quad \times \sum_m m^2 d_{m'm}^{(\ell)}(\beta) d_{m''m}^{(\ell)}(\beta) \\
&\equiv \sum_{m'm''} H_{m'm''}^{(\ell)}(\gamma) G_{m''m'}^{(\ell)}(\beta) \\
&= \mathrm{Tr}\left[\mathbf{H}^{(\ell)}(\gamma) \mathbf{G}^{(\ell)}(\beta)\right].
\end{aligned}
\tag{18}
$$

To find the extrema of $(\Delta L)_\ell^2$, we take derivatives with respect to the angles reducing the problem to the solution of two coupled non-linear equations

$$\mathrm{Tr}\left[\partial_\gamma \mathbf{H}^{(\ell)}(\gamma)\mathbf{G}^{(\ell)}(\beta)\right] = 0, \quad \mathrm{Tr}\left[\mathbf{H}^{(\ell)}(\gamma)\partial_\beta \mathbf{G}^{(\ell)}(\beta)\right] = 0. \tag{19}$$

It is easy to see that

$$\partial_\gamma H_{m'm''}^{(\ell)}(\gamma) = i(m'-m'') H_{m'm''}^{(\ell)}(\gamma). \tag{20}$$

Both $G_{m'm'}^{(\ell)}(\beta)$ and $\partial_\beta G_{m'm'}^{(\ell)}(\beta)$ may be evaluated quickly and efficiently using known properties of the $d_{m'm}^{(\ell)}(\beta)$ (Edmonds 1960).

We note that for the quadrupole the axis $\hat{\boldsymbol{n}}_2$ may instead be calculated directly from the multipole vectors as $\hat{\boldsymbol{n}}_2 = \hat{\boldsymbol{w}}^{(2;1,2)}$ for $\boldsymbol{w}^{(2;1,2)}$ defined below, equation (23). (For the derivation of this correspondence see Copi et al. 2006.)

4.2 Multipole vectors

The multipole vectors form an irreducible representation of the rotation group SO(3), so provide a basis for expanding any scalar function on the sphere. These vectors contain all the information about the function though, in a form different from the usual spherical harmonic coefficients, $a_{\ell m}$. Various properties of the scalar function are more easily described or detected in one basis than the other, and for considering the alignments the multipole vectors have much to offer.

Consider a scalar function, $f_\ell(\hat{\boldsymbol{e}})$, of pure multipole ℓ. In the usual spherical harmonic decomposition this function is expanded as

$$f_\ell(\hat{\boldsymbol{e}}) = \sum_m a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{e}}). \tag{21}$$

In the multipole vector representation, it is instead written in terms of a scalar, $A^{(\ell)}$, and ℓ unit vectors, $\{\hat{\boldsymbol{v}}^{(\ell;j)} \mid j = 1, \ldots, \ell\}$, as

$$f_\ell(\hat{\boldsymbol{e}}) = A^{(\ell)} \left[\prod_{j=1}^\ell \left(\hat{\boldsymbol{v}}^{(\ell;j)} \cdot \hat{\boldsymbol{e}}\right) - \mathcal{T}_\ell\right]. \tag{22}$$

Here $\mathcal{T}_\ell$ is the sum over all possible traces; it removes the lower order multipoles from the preceding product. Notice that when written in this form the function is manifestly rotationally invariant. Also notice that these vectors are only determined up to a sign so that they really define axes.

This representation is not new: James Clerk Maxwell in his study of electromagnetism discussed the properties of spherical harmonics in a very similar manner (Maxwell 1891). It has been shown that Maxwell's representation is identical to the one given above (Weeks 2004), suggesting the alternative name of Maxwell's multipole vectors. With their introduction into CMB studies an algorithm

for converting from the usual spherical harmonic coefficients, $a_{\ell m}$, to the multipole vectors was provided (Copi et al. 2004) greatly facilitating their computation.

There are many ways in which the multipole vectors can be employed to study the CMB. The most useful way has proven to be through the area vectors (Copi et al. 2004, 2006)

$$\boldsymbol{w}^{(\ell;i,j)} \equiv \hat{\boldsymbol{v}}^{(\ell;i)} \times \hat{\boldsymbol{v}}^{(\ell;j)}. \tag{23}$$

These area vectors define sets of planes. Alignment of these planes with a direction $\hat{\boldsymbol{n}}$ is checked by again using the dot product to define a set of values $\{A_j | j = 1, \dots, n\}$ by

$$A_j \equiv |\boldsymbol{w}_j \cdot \hat{\boldsymbol{n}}|. \tag{24}$$

Here $\boldsymbol{w}_j$ represents any area vector. Different sets of area vectors can be used in different cases. For the work discussed here, we restrict the analysis to the quadrupole, $\ell = 2$, and octopole, $\ell = 3$, vectors. Finally we define two statistics based on these values

$$S \equiv \frac{1}{n} \sum_{j=1}^{n} A_j, \quad T \equiv 1 - \frac{1}{n} \sum_{j=1}^{n} (1 - A_j)^2. \tag{25}$$

See Copi et al. (2006) for more general versions of these statistics and Weeks (2004), Schwarz et al. (2004), Slosar & Seljak (2004), Katz & Weeks (2004), Land & Magueijo (2005a), Land & Magueijo (2005b), Bielewicz et al. (2005), Abramo et al. (2006a), Abramo, Sodré & Wuensche (2006b), Weeks & Gundermann (2007), Gruppuso & Burigana (2009), Gruppuso & Górski (2010) for other tests with multipole vectors.

5 RESULTS

The main results in this work are based on DQ-corrected maps as discussed in Section 3.3. Also, for the *Planck* SMICA and NILC maps the results are based on 5×10^5 harmonic inpaintings using the method discussed in Section 3.2.

5.1 Maximum angular momentum dispersion

The alignment of the quadrupole and octopole as quantified by the maximum angular momentum dispersion axes has been studied in recent data releases (Bennett et al. 2011, 2013; Planck Collaboration XXIII 2014). Planck Collaboration XXIII (2014) used Wiener filtered maps to quantify the alignment through the dot product, $|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$, and determined the fraction of realizations with at least this level of alignment in ΛCDM from 1000 realizations. The results reported by *Planck* are summarized in Table 4; we have roughly estimated the uncertainty in these p-values for the *Planck* analysis based

Table 4. Maximum angular momentum dispersion direction alignments reported by *Planck*. The *Planck* results and p-values are from table 17 of Planck Collaboration XXIII (2014) and are based on Wiener filtered maps. The uncertainties in these p-values are calculated based on their use of 1000 realizations. The values from this work are based on 5×10^5 realizations. See the text for details.

| Map | $|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$ | p-value (per cent) | |
		Planck	This work
Planck NILC	0.974	3.3,0.6	2.59
Planck SEVEM	0.988	1.6,0.4	1.17
Planck SMICA	0.977	3.2,0.6	2.28

Table 5. Maximum angular momentum dispersion direction alignments for maps with and without DQ correction. See the text for details.

| Map | Uncorrected | | DQ corrected | |
	$|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$	p-value (per cent)	$|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$	p-value (per cent)
WMAP ILC 7yr	0.9999	0.006	0.9966	0.327
WMAP ILC 9yr	0.9985	0.150	0.9948	0.511
Planck NILC	0.9902	0.955	0.9988	0.118
Planck SEVEM	0.9915	0.825	0.9995	0.055
Planck SMICA	0.9809	1.883	0.9965	0.338

on the simplifying assumption of Poisson statistics. For comparison, using the $|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3|$ values provided by *Planck* we have recalculated the p-values based on 5×10^5 realizations of ΛCDM. These values are also included in Table 4 and are consistently a little more than one-sigma lower than those provided by *Planck*. While the number of simulations used by the *Planck* collaboration is significantly smaller than our sample, their simulations also include instrumental effects and ours do not. The small difference in p-values suggests that instrumental effects are not dominant and add at most a small correction of the order of the statistical uncertainty of the *Planck* simulation itself.

To study the effect of the extra cleaning provided by the Wiener filtering we have calculated the maximum angular momentum dispersion axes from the full-sky maps provided by *Planck*. For the SMICA and NILC maps the *Planck* inpainted maps were analysed. The results are given in Table 5 as the 'Uncorrected' values. These p-values should be compared to our p-values from Table 4 (last column). We see that the provided maps exhibit somewhat more alignment than the Wiener filtered maps. So something *has* been removed by the Wiener filtering. Whether it is noise or CMB signal is unclear.

Also included in Table 5 are the results from the *WMAP* data releases. The 'Uncorrected' values are consistent with discussions in the *WMAP* seven-year (Bennett et al. 2011) and nine-year (Bennett et al. 2013) analyses. The alignment in the seven-year data is quite remarkable for being almost perfect ($|\hat{\boldsymbol{n}}_2 \cdot \hat{\boldsymbol{n}}_3| \simeq 1$). The change in alignment in the nine-year data is largely attributed to improvements in the asymmetric beam deconvolution (Bennett et al. 2013) and is one example of how analysis improvements affect the alignments. Even so, the *WMAP* maps show more alignment than the *Planck* maps.

The effect of the DQ correction as discussed in Section 3.3 is also included in Table 5. Interestingly, the DQ correction has the opposite effect on the *WMAP* and *Planck* alignments. Since the *WMAP* alignments are so precise, this correction lessens the significance as we would expect; however, for *Planck* we find the alignments become more significant. More importantly, *WMAP* and *Planck* are found to be in *better* agreement with each other after the DQ correction has been applied.

5.2 Multipole vectors

The multipole vectors for the DQ-corrected SMICA map are listed in Table 6 and shown in Fig. 5 plotted in Galactic coordinates. When compared to fig. 3 from Copi et al. (2006), this figure shows that the general features have not changed significantly since the first-year *WMAP* data release. It also provides a visual summary of many of the large-angle anomalies. In particular,

Table 6. Average DQ-corrected multipole vectors from 5×10^5 harmonic inpaintings of the SMICA map. The vector directions are given in Galactic coordinates, (l, b), and their Cartesian equivalents, (x, y, z). These vectors are plotted in Fig. 5.

Vector	l (deg)	b (deg)	x	y	z	Magnitude
$\hat{v}^{(2,1)}$	3.5	14.4	0.967	0.059	0.249	–
$\hat{v}^{(2,2)}$	126.5	13.3	−0.579	0.783	0.229	–
$w^{(2;1,2)}$	63.6	−62.7	0.182	0.366	−0.791	0.890
$\hat{v}^{(3,1)}$	90.5	42.0	−0.007	0.744	0.669	–
$\hat{v}^{(3,2)}$	22.6	9.2	0.911	0.380	0.159	–
$\hat{v}^{(3,3)}$	−47.1	11.8	0.667	−0.716	0.205	–
$w^{(3;1,2)}$	102.5	−47.4	−0.136	0.610	−0.680	0.924
$w^{(3;2,3)}$	−22.8	−77.0	0.192	−0.081	−0.906	0.930
$w^{(3;3,1)}$	35.3	−32.4	0.632	0.447	−0.491	0.916

(i) the Ecliptic plane is seen to carefully thread itself between a hot and cold spot and there is a clear power asymmetry across the Ecliptic plane;

(ii) the planarity of the octopole and the alignment of the quadrupole and octopole planes is clearly visible – note the remarkable near-overlap of the quadrupole and octopole maximum angular momentum dispersion axes;

(iii) the area vectors lie near each other, near the Ecliptic plane, and also near the dipole direction.

To quantify the alignments, we consider the mutual alignment of the quadrupole and octopole area vectors as well as alignments with the Ecliptic plane, north Galactic pole (NGP), and the direction of our motion with respect to the CMB (dipole). This is a subset of the directions considered in Copi et al. (2006). The results based on the *S* and *T* statistics from equation (25) are shown in Table 7, and indicate that alignments persist at the 95 to 99.9 per cent level, with the strongest alignment occurring with the dipole direction ($\geq$99.6 per cent). As we argued above, the spread of values between the different maps gives an idea of the effect of the residual systematic errors due to foregrounds.

As noted above, subsequent to our original analysis Planck Collaboration XXIII (2014) was updated to include estimates for DQ correction factors based on the calibration techniques and weights applied to the individual frequency band maps combined to create the cleaned, full-sky maps. The correction factors were found to be 1.7, 1.5, and 1.7 for the *Planck* NILC, SEVEM, and SMICA maps, respectively. To determine the effect of using these correction factors, we have reanalysed the alignments of the multipole vectors for these three maps. For the *Planck* SMICA map, the direction of the oriented area vector for the quadrupole based on our simple DQ correction ($\hat{n}_2$ and labelled 'Qa' in Fig. 5) moves about $3°2$ when calculated using the *Planck* DQ correction. This is a significantly smaller change than that from the vector without any DQ correction applied (the 'no DQ' vector in Fig. 5). Of more importance, the magnitude of the oriented area vector, $w^{(2;1,2)}$, has changed from 0.890 (see Table 6) to 0.948. In other words, with the *Planck* DQ


Figure 5. Quadrupole and octopole multipole vectors for the DQ corrected SMICA map in Galactic coordinates. The background shows the quadrupole+octopole pattern from the DQ corrected SMICA map. The multipole vectors are shown as circles, in red and labelled 'Qv' for the quadrupole and in black and labelled 'Ov' for the octopole. The direction of the area vectors defined in equation (23), $\hat{w}^{(\ell;i,j)}$, is shown as squares. Again the quadrupole area vector is in red and labelled 'Qa' and the octopole area vectors are in black and labelled 'Oa'. Since the multipole vectors are only determined up to a sign each vector appears twice in the figure. The area vectors have only been plotted in the Southern hemisphere to avoid cluttering the plot. The maximum angular momentum dispersion direction for the octopole, $\hat{n}_3$, is shown as the black star. Since $\hat{n}_2 = \hat{w}^{(2;1,2)}$ it is also represented by the red square. The direction of $\hat{n}_2$ without the DQ correction is shown as the red diamond. For reference also shown in the figure is the Ecliptic plane (black line), the locations of the north (NEP) and south (SEP) Ecliptic poles, and the direction of our motion with respect to the CMB (dipole). The coordinates of the vectors are listed in Table 6.

Table 7. The S and T alignment statistics from equation (25) for various directions. Listed are the p-values in per cent of ΛCDM producing a value larger than that found in the given map based on 10^6 realizations of ΛCDM. The directions tested are the quadrupole+octopole alignment (Q+O), the Ecliptic plane, the north Galactic pole (NGP), and the direction of our motion with respect to the CMB (dipole). The results for the SMICA and NILC maps are based on the average of the S statistic from 5×10^5 harmonic inpaintings of these maps. The results below the '*Planck* DQ Correction' line were calculated using the *Planck* DQ correction factors as discussed in the text.

	Q+O		Ecliptic Plane		NGP		dipole	
Map	S	T	S	T	S	T	S	T
WMAP ILC 7yr	0.22	0.10	2.66	2.70	0.82	0.90	0.18	0.20
WMAP ILC 9yr	0.18	0.08	1.96	1.82	0.79	0.76	0.14	0.15
Planck NILC	1.85	1.05	2.80	3.04	1.41	1.26	0.32	0.19
Planck SEVEM	0.41	0.22	2.52	2.94	0.79	0.92	0.09	0.05
Planck SMICA	1.62	0.93	3.74	4.16	1.56	1.52	0.37	0.30
			Planck DQ Correction					
Planck NILC	0.54	0.27	2.55	2.64	1.14	1.10	0.18	0.14
Planck SEVEM	0.16	0.08	2.31	2.58	0.73	0.89	0.06	0.05
Planck SMICA	0.54	0.28	3.37	3.60	1.32	1.39	0.23	0.23

correction of the *Planck* SMICA map the quadrupole multipole vectors are even more nearly perpendicular. The effect this has on the S and T statistics for the alignments we have studied is given in Table 7 under the '*Planck* DQ Correction' line. In all cases, we see that the *Planck* DQ correction leads to alignments *more unlikely* in ΛCDM.

To make the anomalous nature of the alignments more clear, the results for the S and T statistics from equation (25) of the multipole vectors with the direction of our motion with respect to the CMB (the dipole direction) are shown in Fig. 6. In this figure, the histograms are the expected values from ΛCDM and the vertical lines the range of values found for the various maps. In all cases, the observed alignments reside far in the tail of the expected distributions.

5.3 Interdependence of alignments

The alignments shown in Table 7 are peculiar, particularly the mutual alignment of the quadrupole and octopole area vectors as well


Figure 6. The S and T statistics from equation (25) for the alignment of the multipole vectors with the direction of our motion with respect to the CMB (the dipole direction). The histograms represent the distribution of the S (solid, black line) and T (dashed, red line) statistics from 10^6 realizations of ΛCDM. The shaded regions between the vertical lines represent the range of values found for the CMB maps studied in this work. See Table 7 for the full results.

as that with the dipole direction. A remaining question is the independence of these alignments.

As a first method of addressing this, we follow Copi et al. (2006) and calculate the significance of alignments with various directions *given* the observed relative orientation of the quadrupole and octopole area vectors. That is, given the observed quadrupole and octopole structure, the overall orientation of this pattern on the sky is arbitrary. With this in mind we can address how likely it is for a random orientation of this structure to have the quadrupole and octopole area vectors at least as well aligned with a fixed direction in the sky as observed. Though we think of this as rotating the map, it is equivalent to comparing the alignments to random directions. The result of such a comparison based on 10^6 random directions using the S and T statistics (25) is given in Table 8. We find that the residual p-value for alignment with the dipole and Ecliptic Plane directions, given the mutual quadrupole+octopole alignment, is at the 2 to 6 per cent level, while the alignment with the Galactic pole is not significant. These results are in agreement with results in Copi et al. (2006) (see table 4 in that paper), and indicate that even given the relative location of the quadrupole and octopole area vectors (i.e. their mutual alignment), the Ecliptic plane and dipole alignments are unlikely at the 95 per cent level.

A related question is to ask in realizations of ΛCDM with *at least* the observed alignment of the quadrupole and octopole, what is the p-value for alignment with other directions? The difference between this and the previous question is that the orientation of the quadrupole and octopole is no longer a rigid structure. From Table 7 we see that realizations of ΛCDM with the quadrupole and octopole aligned as closely as in the data are rare. The most frequent alignment occurs for the *Planck* NILC map with p-values of

Table 8. The alignment p-value based on the S and T statistics from equation (25) *given the observed quadrupole+octopole alignment*. The values are quoted in per cent and are based on 10^6 randomly rotated directions. See the text for details.

	Ecliptic Plane		NGP		dipole	
Map	S	T	S	T	S	T
WMAP ILC 7yr	3.27	3.33	12.26	12.74	6.47	6.55
WMAP ILC 9yr	1.30	1.38	12.56	12.35	6.02	6.30
Planck NILC	4.39	5.14	12.26	12.74	5.50	5.08
Planck SEVEM	4.93	6.37	11.75	12.74	3.98	3.17
Planck SMICA	5.17	6.10	12.66	13.15	5.94	6.05

Table 9. The alignment *p*-values based on the *S* and *T* statistics from equation (25) *given at least the observed quadrupole+octopole alignment*. The values are quoted in per cent and based on 10^6 realizations of ΛCDM with at least the quadrupole+octopole alignment observed in the *Planck* `NILC` map. See the text for details.

Map	Ecliptic Plane		NGP		dipole	
	S	*T*	*S*	*T*	*S*	*T*
WMAP ILC 7yr	20.9	18.9	14.3	17.4	8.7	10.8
WMAP ILC 9yr	18.6	16.5	14.3	16.8	7.9	10.0
Planck `NILC`	14.0	14.0	13.6	14.9	6.9	6.2
Planck `SEVEM`	18.8	18.2	13.5	16.6	5.9	5.5
Planck `SMICA`	17.4	17.1	14.4	16.3	7.8	8.1

1.85 and 1.05 per cent for the *S* and *T* statistics, respectively. Using these values as cutoffs, we generate 10^6 realizations of ΛCDM with at least this much alignment, that is, realizations with either a *S* or *T* quadrupole+octopole statistic at least the value from the *Planck* `NILC` map. Note that this means that the number of realizations vary for each map. The ratio of the number of realizations between each pair of maps is given by the ratio of *p*-values from the Q+O columns in Table 7. For each map we then find the conditional probability of alignment with some direction, such as the dipole, *given* at least the observed alignment of the quadrupole+octopole. The resulting *p*-values are given in Table 9. In comparison with Table 8 we see by relaxing the rigid structure of the quadrupole and octopole orientations that the *p*-values have increased. Overall the *p*-values are now roughly 5 to 21 per cent with the more unlikely alignment being with the dipole direction, meaning that the quadrupole+octopole and dipole alignments are the least correlated.

The results from Tables 7 and 9 can be combined using Bayes' theorem to find the conditional probability for any two alignments. Let d_1 and d_2 represent two different alignments, then, $P(d_1|d_2)$ is the conditional probability of alignment with d_1 given the alignment with d_2 and similarly for $P(d_2|d_1)$. These alignments are related as

$$P(d_1|d_2)P(d_2) = P(d_2|d_1)P(d_1). \qquad (26)$$

From the tables provided, we may immediately calculate the conditional probability for the alignment of the quadrupole+octopole given the alignment of any other direction. To do so let $d_2 \equiv$ qo represent the quadrupole+octopole alignment, then

$$P(\text{qo}|\text{d}_1) = P(\text{d}_1|\text{qo})\frac{P(\text{qo})}{P(\text{d}_1)}. \qquad (27)$$

Here both $P(\text{qo})$ and $P(d_1)$ may be read from Table 7 and $P(d_1|\text{qo})$ from Table 9. These conditional probabilities are shown in Table 10. Conditional probabilities for other pairs of alignments can be calculated from repeated application of Bayes' theorem.

Table 10. The quadrupole+octopole alignment *p*-values based on the *S* and *T* statistics from equation (25) *given at least the observed dipole, Ecliptic or Galactic alignment*. The values are quoted in per cent and are obtained using Bayes' theorem and the data from Tables 7 and 9.

Map	Ecliptic Plane		NGP		dipole	
	S	*T*	*S*	*T*	*S*	*T*
WMAP ILC 7yr	1.7	0.7	3.8	1.9	10.6	5.4
WMAP ILC 9yr	1.7	0.7	3.3	1.8	10.2	5.3
Planck `NILC`	9.3	4.8	17.8	12.4	39.9	34.3
Planck `SEVEM`	3.1	1.4	7.0	4.0	26.9	24.2
Planck `SMICA`	7.5	3.8	15.0	10.0	34.2	25.1

If we assume the observed levels of alignment with the Ecliptic plane or the Galactic poles, the quadrupole+octopole alignment remains anomalous at the 1 to 4 per cent level for *WMAP*. For *Planck*, the conclusion is not clear, however all *p*-values remain below 20 per cent. The situation is different for the dipole. We see that assuming the observed level of dipole alignment, the quadrupole+octopole alignment seems to be quite plausible. This suggests that the dipole alignment (which is also the most significant and robust alignment in Table 7) could be the reason for the other observed alignments.

6 CONCLUSIONS

The largest structures in the microwave sky, the quadrupole and octopole, are aligned with one another and with physical directions or planes – the dipole direction and the Ecliptic plane. These alignments, first observed and discussed in the one-year *WMAP* data, have persisted throughout *WMAP*'s subsequent data releases, and are now confirmed in the one-year *Planck* data. On the one hand, this is to be expected: the largest scales are precisely measured and the same CMB sky is observed by both satellites. On the other hand, this is surprising: cleaned, full-sky maps are required to see these alignments, and the removal of foregrounds, along with other systematic effects, makes it challenging to accurately produce full-sky maps on large angular scales.

In this work, we have studied the ILC maps from the seven and nine-year *WMAP* data releases and the `NILC`, `SEVEM`, and `SMICA` cleaned maps from the first-year *Planck* data release. Qualitatively, the main anomalies detected in earlier *WMAP* releases remain: the quadrupole and octopole are aligned with each other; the normal to their average plane is aligned with the dipole – the direction of our motion through the Universe; that normal is also close to the Ecliptic plane, so that the average plane of the quadrupole and octopole is nearly perpendicular to the Ecliptic plane. Finally, as can be seen in Fig. 5 which shows the sum of the quadrupole and octopole from the `SMICA` map, the Ecliptic plane cleanly cuts between a hot and cold spot, thereby separating weaker quadrupole+octopole power in the north Ecliptic hemisphere from the stronger power in the south Ecliptic hemisphere.

Quantitatively, statistics from the maximum angular momentum dispersion (Table 5) and the multipole vectors (Table 7 and Fig. 6) both show strong evidence for the mutual alignment of the quadrupole and the octopole. The *p*-values for at least as much alignment as observed occurring in realizations of ΛCDM are typically less than 0.5 per cent, once one has Doppler-corrected the quadrupole. The exceptions are the `NILC` and `SMICA` multipole vector statistics, *S* and *T*, which have *p*-values of 1 to 2 per cent. These results are strengthened when the *Planck* DQ corrections are applied instead of our simple one.

The alignment of the quadrupole and octopole with the dipole (Table 7) appears at first sight even more robust than their mutual alignments, with *p*-values of less than 0.4 per cent (and as low as 0.05 per cent) in all maps and with both *S* and *T* statistics. The interpretation however is not clear. The dipole includes contributions from several sources, but is almost certainly dominated by the Doppler effect from the Sun's motion through the Galaxy, the Galaxy's motion through the Local Group, and the Local Group's motion through the more distant large-scale structure; all giving comparable contributions. (A dominant or even significant contribution from a cosmological dipole seems remote.) Since these contributions originate from gravity gradients on very different scales, it is difficult to envision physics that would connect the dipole, quadrupole, and octopole. A systematic error in the measurement

or the analysis pipeline could connect them all, but the robustness of the alignment across the two satellites argues against that explanation. (A remote possibility does still remain since *Planck* calibrated off the *WMAP* dipole in their first-year data release.)

As an attempt to disentangle correlation from causation between and among the alignments, we have studied their interdependence by calculating the conditional probability of alignment with a fixed direction given the observed mutual quadrupole+octopole alignment and *vice versa*. Through application of Bayes' theorem, the conditional probability for any two alignments can be deduced from these calculations. For example, the conditional *p*-values for quadrupole+octopole+dipole alignment given either the observed quadrupole+octopole alignment (Table 8) or at least the observed quadrupole+octopole alignment (Table 9) are 3 to 10 per cent. These are consistent with the 4 to 6 per cent found in Copi et al. (2006) for the conditional *p*-value of the quadrupole+octopole+dipole alignment given the observed relative directions of the quadrupole and octopole area vectors. In other words, the alignments with the dipole may well be a distraction – a statistical accident – barring an unknown common *WMAP–Planck* systematic which somehow causes the dipole+quadrupole+octopole alignment.

A priori, less scepticism could be attached to a possible physical explanation for the correlation between the quadrupole+octopole and the Ecliptic plane. If the underlying cosmological quadrupole and octopole were unexpectedly absent, then we could well imagine a Solar system (or even nearby Solar-neighbourhood) source for the quadrupole and octopole correlated with the plane of the Solar system. Nevertheless, there are no proposed viable physical models that correctly reproduce the observed arrangement of quadrupole and octopole extrema lying on a plane perpendicular to the Ecliptic and well separated by it. The statistics offer only weak, and even confusing, guidance as to what is correlation and what, if anything, is causation. The statistical situation is no clearer for correlation with the Galactic pole.

Unfortunately, the fact that the dipole direction simply happens to be just off the Ecliptic plane, which passes about 30° from the Galactic poles, makes establishing the priority of one correlation over another difficult just on the basis of statistics of CMB temperature data. Some, or all, of these correlations are presumably accidental. Solving this puzzle will require data other than just CMB temperature maps, and probably a model that can be tested against such data.

The lack of correlation on large angular scales from cut-sky maps has been presented in a companion work (Copi et al. 2013). A natural next step would be to also study the interdependence of the alignments of the full-sky low multipoles and the lack of correlations. For the *WMAP* releases up to year seven, these were discussed in Rakić & Schwarz (2007) and Sarkar et al. (2011). The corresponding analysis for the final *WMAP* and first-year *Planck* releases will be presented elsewhere, but we expect that the conclusion from Sarkar et al. (2011), namely that the lack of angular correlation and the alignments are uncorrelated anomalies, will remain valid.

In summary, the quadrupole and octopole alignments noted in early *WMAP* full-sky maps persist in the *WMAP* seven-year and final (nine-year) maps, and in the *Planck* first-year full-sky maps. The correlation of the quadrupole and octopole with one another, and their correlations with other physical directions or planes – the dipole, the Ecliptic, the Galaxy – remain broadly unchanged across all of these maps. Consequently, it is not sufficient to argue that they are less significant than they appear merely by appealing to the uncertainties in the full-sky maps – such uncertainties are presumably captured in the range of foreground removal schemes that went into the map making. It similarly seems contrived that the primordial CMB at the last scattering surface is correlated with the local structures imprinted via the Integrated Sachs–Wolfe (ISW) effect in just such a way to generate the observed alignments, as proposed elsewhere (Rakić, Räsänen & Schwarz 2006; Francis & Peacock 2010; Dupé et al. 2011; Rassat & Starck 2013; Rassat, Starck & Dupé 2013), even taking for granted the reliability of the procedure to subtract the ISW signal from the map.

While it may be tempting to explain away the observed large-angle alignments in the CMB by postulating additional, unspecified corrections to the maps, such explanations so far have not been compelling. Numerous corrections have been applied in the data analysis pipelines, and they have also evolved between the initial *WMAP* data releases and the *Planck* first-year release, yet the alignments remain. Furthermore, almost anything done to the maps will lessen the significance of the observed alignments, so, just because a new correction *could* affect the observed alignments, this does not mean that an otherwise unspecified new correction *must* exist. We think it is preferable to acknowledge that the existence of anomalies seen in the *WMAP* and *Planck* maps at large angular scales *may* point to residual contamination in the data or to interesting new fundamental physics.

REFERENCES

Abramo L. R., Bernui A., Ferreira I. S., Villela T., Wuensche C. A., 2006a, Phys. Rev. D, 74, 063506
Abramo L. R., Sodré L., Jr, Wuensche C. A., 2006b, Phys. Rev. D, 74, 083515
Abrial P., Moudden Y., Starck J. L., Fadili J., Delabrouille J., Nguyen M. K., 2008, Stat. Methodol., 5, 289
Alnes H., Amarzguioui M., 2006, Phys. Rev. D, 74, 103520
Bennett C. L. et al., 2011, ApJS, 192, 17
Bennett C. L. et al., 2013, ApJS, 208, 20
Benoit-Lévy A., Déchelette T., Benabed K., Cardoso J. F., Hanson D., Prunet S., 2013, A&A, 555, A37
Bielewicz P., Eriksen H. K., Banday A. J., Górski K. M., Lilje P. B., 2005, ApJ, 635, 750
Bucher M., Louis T., 2012, MNRAS, 424, 1694
Chiang L. Y., Naselsky P. D., Coles P., 2007, ApJ, 664, 8
Copi C. J., Huterer D., Starkman G. D., 2004, Phys. Rev. D, 70, 043515

Copi C. J., Huterer D., Schwarz D. J., Starkman G. D., 2006, MNRAS, 367, 79
Copi C. J., Huterer D., Schwarz D. J., Starkman G. D., 2010, Adv. Astron., 2010, 78
Copi C. J., Huterer D., Schwarz D. J., Starkman G. D., 2013, preprint (arXiv:1310.3831)
de Oliveira-Costa A., Tegmark M., Zaldarriaga M., Hamilton A., 2004, Phys. Rev. D, 69, 063516
Dikarev V., Preuß O., Solanki S., Krüger H., Krivov A., 2008, Earth Moon and Planets, 102, 555
Dupé F. X., Rassat A., Starck J. L., Fadili M. J., 2011, A&A, 534, A51
Edmonds A. R., 1960, Angular Momentum in Quantum Mechanics. Princeton Univ. Press, Princeton, NJ
Eriksen H. K., Banday A. J., Górski K. M., Lilje P. B., 2004, ApJ, 612, 633
Fixsen D. J., 2009, ApJ, 707, 916
Francis C. L., Peacock J. A., 2010, MNRAS, 406, 14
Frejsel A., Hansen M., Liu H., 2013, J. Cosmol. Astropart. Phys., 6, 005
Frisch P. C., 2005, ApJ, 632, L143
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
Gruppuso A., Burigana C., 2009, J. Cosmol. Astropart. Phys., 0908, 004
Gruppuso A., Górski K. M., 2010, J. Cosmol. Astropart. Phys., 3, 019
Hinshaw G. et al., 2003, ApJS, 148, 63
Hinshaw G. et al., 2007, ApJS, 170, 288
Hinshaw G. et al., 2009, ApJS, 180, 225
Inoue K. T., Silk J., 2006, ApJ, 648, 23
Inoue K. T., Cabella P., Komatsu E., 2008, Phys. Rev. D, 77, 123539
Kamionkowski M., Knox L., 2003, Phys. Rev. D, 67, 063001
Katz G., Weeks J., 2004, Phys. Rev. D, 70, 063527
Kim J., Naselsky P., Mandolesi N., 2012, ApJ, 750, L9
Kovács A., Carron J., Szapudi I., 2013, MNRAS
Land K., Magueijo J., 2005a, Phys. Rev. Lett., 95, 071301
Land K., Magueijo J., 2005b, MNRAS, 362, 838
Maxwell J. C., 1891, A Treatise on Electricity and Magnetism, Vol. I, 3rd edn. Clarendon Press, London
Nishizawa A. J., Inoue K. T., 2013, preprint (arXiv:1305.0116)
Peebles P. J. E., Wilkinson D. T., 1968, Phys. Rev., 174, 2168
Peiris H. V., Smith T. L., 2010, Phys. Rev. D, 81, 123517
Planck Collaboration V, 2014, A&A, 571, A5
Planck Collaboration VIII, 2014, A&A, 571, A8
Planck Collaboration XII, 2014, A&A, 571, A12
Planck Collaboration XVI, 2014, A&A, 571, A16
Planck Collaboration XXIII, 2014, A&A, 571, A23
Pullen A. R., Kamionkowski M., 2007, Phys. Rev. D, 76, 103529
Rakić A., Schwarz D. J., 2007, Phys. Rev. D, 75, 103002
Rakić A., Räsänen S., Schwarz D. J., 2006, MNRAS, 369, L27
Ramirez E., Schwarz D. J., 2009, Phys. Rev. D, 80, 023525
Rassat A., Starck J. L., 2013, A&A, 557, L1
Rassat A., Starck J. L., Dupé F. X., 2013, A&A, 557, A32
Rybicki G. B., Press W. H., 1992, ApJ, 398, 169
Sarkar D., Huterer D., Copi C. J., Starkman G. D., Schwarz D. J., 2011, Astropart. Phys., 34, 591
Schwarz D. J., Starkman G. D., Huterer D., Copi C. J., 2004, Phys. Rev. Lett., 93, 221301
Slosar A., Seljak U., 2004, Phys. Rev. D, 70, 083002
Spergel D. N. et al., 2003, ApJS, 148, 175
Starck J. L., Fadili M. J., Rassat A., 2013, A&A, 550, A15
Starkman G. D., Trotta R., Vaudrevange P. M., 2008, preprint (arXiv:0811.2415)
Weeks J. R., 2004, preprint (arXiv:astro-ph/0412231)
Weeks J., Gundermann J., 2007, Class. Quantum Gravity, 24, 1863

This paper has been typeset from a TEX/LATEX file prepared by the author.