

A Monte Carlo comparison between template-based and Wiener-filter CMB dipole estimators

H. Thommesen, K. J. Andersen, R. Aurlen, R. Banerji, M. Brilenkov, H. K. Eriksen, U. Fuskeland, M. Galloway, L. M. Mocanu, T. L. Svalheim, and I. K. Wehus

Institute of Theoretical Astrophysics, University of Oslo, PO Box 1029, Blindern 0315, Oslo, Norway
e-mail: harald.thommesen@astro.uio.no

Received 13 July 2020 / Accepted 8 September 2020

ABSTRACT

We review and compare two different cosmic microwave background (CMB) dipole estimators discussed in the literature and assess their performances through Monte Carlo simulations. The first method amounts to simple template regression with partial sky data, while the second method is an optimal Wiener filter (or Gibbs sampling) implementation. The main difference between the two methods is that the latter approach takes into account correlations with higher-order CMB temperature fluctuations that arise from nonorthogonal spherical harmonics on an incomplete sky, which for recent CMB data sets (such as *Planck*) is the dominant source of uncertainty. For an accepted sky fraction of 81% and an angular CMB power spectrum corresponding to the best-fit *Planck* 2018 Λ CDM model, we find that the uncertainty on the recovered dipole amplitude is about six times smaller for the Wiener filter approach than for the template approach, corresponding to 0.5 and 3 μ K, respectively. Similar relative differences are found for the corresponding directional parameters and other sky fractions. We note that the Wiener filter algorithm is generally applicable to any dipole estimation problem on an incomplete sky, as long as a statistical and computationally tractable model is available for the unmasked higher-order fluctuations. The methodology described in this paper forms the numerical basis for the most recent determination of the CMB solar dipole from *Planck*, as summarized by [Planck Collaboration Int. LVII \(2020\)](#).

Key words. cosmic background radiation – cosmology: observations – diffuse radiation

1. Introduction

The cosmic microwave background (CMB) radiation was discovered in 1965 by [Penzias & Wilson \(1965\)](#) and has been the primary target for several dozens of CMB experiments ever since. The main scientific target for most of these studies has been small variations in intensity and polarization that correspond to cosmic density variations some 380 000 years after the Big Bang. These variations contain a wealth of information about the early history and evolution of the Universe; for a recent analysis, see, for example, [Planck Collaboration VI \(2020\)](#).

The CMB sky features three main physical components. The first is simply a constant blackbody term with a temperature of 2.7255 K ([Fixsen 2009](#)), corresponding to the average temperature of the CMB photons populating the Universe today. This component is often denoted as the CMB monopole, which acknowledges its correspondence to the lowest multipole moment in spherical harmonics space.

The second component is the CMB dipole, which has an amplitude of about 3 mK ([Lineweaver 1997](#)). The CMB dipole is the result of Doppler boosting caused by the motion of the measuring instrument with respect to the CMB rest frame. It may be decomposed into two components, namely the solar dipole, which is caused by the movement of the Solar System around the Milky Way's center, and the orbital dipole, which is generated by the movement of the Earth and the instrument around the Sun. Both components play an important role in CMB experiments as they represent the best available astrophysical calibration source for most experiments. Specifically, the orbital dipole serves as an

invaluable tool for absolute calibration since the Earth–Sun distance and the orbital period are known with very high precision. Likewise, the solar dipole provides an excellent relative calibration target since it is brighter than most other signals, it has a perfectly known frequency spectrum, and it is visible across the entire sky.

The third component is the CMB density fluctuations with typical variations of about 100 μ K. These correspond very closely to a statistically isotropic and Gaussian random field with an angular power spectrum, which can be described by a Λ CDM power spectrum ([Planck Collaboration VI 2020](#); [Planck Collaboration VII 2020](#)). In addition to these CMB sources, real-world microwave observations also contain contributions from astrophysical foregrounds, most notably in the form of synchrotron, free-free, spinning and thermal dust, and CO emission from the Milky Way (e.g., [Planck Collaboration IV 2020](#)).

This paper discusses how to optimally estimate the amplitude and direction of the CMB dipole with data that contain both astrophysical foregrounds and small-scale CMB fluctuations. Obtaining robust estimates for these parameters is important for several reasons. First, since the CMB dipole is used as a calibration source for most experiments, a potential bias in the CMB dipole amplitude translates directly into a corresponding bias in the overall normalization of the angular CMB power spectrum. Second, because the CMB dipole is about four orders of magnitude brighter than the cosmological variations in the large-scale polarization field, it is necessary to estimate the relative gains between detectors within a single frequency channel




Fig. 1. Basic templates used for CMB dipole estimation in the traditional template fitting approach, shown in Galactic coordinates. These correspond to the monopole (*upper left*) and the *x*- (*upper right*), *y*- (*lower left*), and *z*-dipoles (*lower right*).

prior to mapmaking with a precision better than $\mathcal{O}(10^{-4})$ in order to avoid significant bias on the optical depth of reionization, τ (Planck Collaboration VI 2020).

Uncertainties on the dipole parameters result mainly from four different contributors. First, statistical instrumental noise defines a fundamental floor for the overall sensitivity that can be achieved. Second, many systematic effects due to nonidealities in the instrument itself can induce spurious dipoles, including sidelobe pickup, time-variable gain, or ADC corrections. Third, as already mentioned, foreground emission from the Milky Way obscures our view of the CMB, and also carries a dipole moment. To mitigate this effect it is in practice necessary to mask out parts of the Galactic regions. However, working with incomplete sky coverage has the unwanted effect of making the spherical harmonic base functions, $Y_{\ell m}(\theta, \phi)$, lose their orthogonality. This leads to the fourth and last contaminant, which is confusion from the higher-order CMB temperature fluctuations when analyzing partial-sky observations.

The traditional way of CMB dipole parameter estimation with real data has typically followed a fairly simple approach. First, the orbital dipole contribution is removed from the time-ordered data of a given experiment, which are subsequently co-added into pixelized sky maps. Then foreground contamination is suppressed either through some component separation technique or by simple template regression with respect to known foreground tracers. Next, some part of the sky is removed by masking, before finally the CMB dipole is estimated through template fitting with partial-sky and foreground-cleaned data, typically adopting the templates shown in Fig. 1. It is important to note that the monopole is usually included in the fit for marginalization purposes only, to avoid potential inaccuracies in the zero-level determination from biasing the dipole fit. For one specific example of such an implementation, see, for example, Planck Collaboration II (2016).

As instrumental sensitivity has improved through the years, the specific details of each step in this procedure have become more important. For instance, for COBE-FIRAS (Fixsen et al. 1994) the uncertainty on the dipole amplitude due to statistical noise was $6 \mu\text{K}$, while the foreground-induced uncertainty was about $14 \mu\text{K}$. For comparison, Lineweaver (1997) estimated that the uncertainty due to confusion from higher-order CMB fluctuations was $3 \mu\text{K}$, and this particular term was therefore irrelevant for COBE.

The same did not hold true for the Wilkinson Microwave Anisotropy Probe (WMAP) experiment, for which both the raw sensitivity and foreground rejection capabilities improved

massively with respect to COBE. For this reason, the WMAP team replaced the simple dipole template-fitting procedure discussed above with a more sophisticated and optimal Wiener filter method (Hinshaw et al. 2009) originally pioneered by Jewell et al. (2004), Wandelt et al. (2004), Eriksen et al. (2004). The main advantage of this approach is the fact that the higher-order CMB fluctuations are estimated jointly with the dipole parameters, and by assuming that these correspond to a statistically isotropic and Gaussian random field, it is possible to partially reconstruct their properties even inside the Galactic mask.

The main goal of the current paper is to quantify the relative performance of the template fitting and Wiener filter methods. This has recently become a particularly important topic in the context of the *Planck* experiment, for which the sensitivity and control of systematic effects is so high that the total error budget has now become dominated by the higher-order CMB contribution. Specifically, the total instrumental uncertainty on the dipole amplitude is about $1 \mu\text{K}$ (Planck Collaboration I 2020), whereas the CMB confusion term arising from the naive template approach is, as we see later, typically between 1 and $3 \mu\text{K}$, depending on sky fraction. Minimizing this term is therefore critically important. The results we obtain when applying this analysis framework to the latest *Planck* observations are summarized in Planck Collaboration Int. LVII (2020).

The rest of this paper is organized as follows: in Sect. 2 we give a short theoretical introduction on how the dipole parameters are obtained from sky maps, what effect partial sky coverage has on the uncertainties of these parameters, and further present a Wiener filter method for the estimation of these uncertainties. In Sect. 3 we describe the implementation of the uncertainty estimation technique. We present our results in Sect. 4, and we give our conclusions in Sect. 5.

2. Notation and methods

We start our discussion with a review of the traditional approach for estimating the dipole parameters (amplitude A , Galactic longitude l and latitude b) from a CMB map, and a discussion of how partial sky coverage complicates this procedure.

In this paper, we take as a starting point for dipole parameter estimation a cleaned CMB map in which as much foreground emission as possible has been removed; for details on how to perform foreground cleaning, we refer the interested reader to, for example, Planck Collaboration IV (2020) and references therein. However, no component separation technique allows for perfect foreground removal, and there will always be some residual emission, especially in the Galactic plane. One therefore must typically apply a mask to eliminate heavily contaminated regions on the sky. How to optimally estimate the dipole parameters from such a clean but incomplete CMB map is the main topic of this section.

2.1. Traditional template fitting

As a starting point for our analysis, we assume that the clean sky map may be written in the form

$$\mathbf{d} = \mathbf{T}\mathbf{a} + \mathbf{n}, \quad (1)$$

where \mathbf{T} is a matrix containing the monopole and dipole templates in its columns, shown in Fig. 1; \mathbf{a} is a corresponding vector of template amplitudes; and \mathbf{n} is noise. We note that the latter may or may not include higher-order CMB fluctuations. In addition, we define a noise covariance matrix $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^t \rangle$, and a diagonal mask matrix \mathbf{M} that is zero for masked pixels and unity for unmasked pixels.




Fig. 2. Spherical harmonics coupling kernel K for $\ell \leq 60$ for an unmasked (*top*) and a masked (*bottom*) sky. Multipole moments are listed in ℓ -major ordering with element numbering given by $i = \ell^2 + \ell + m + 1$. The sky fraction for the masked case is 36%. The corresponding mask is plotted in Fig. 8.

Assuming that the noise is Gaussian, the maximum-likelihood solution for \mathbf{a} is then given by the so-called “normal equations”,

$$\mathbf{a} = (\mathbf{T}'\mathbf{M}\mathbf{N}^{-1}\mathbf{T})^{-1}\mathbf{T}'\mathbf{M}\mathbf{N}^{-1}\mathbf{d}. \quad (2)$$

The trigonometric relations between these amplitudes and the dipole parameters are

$$A = \sqrt{a_x^2 + a_y^2 + a_z^2} \quad (3)$$

$$l = \arctan(a_y/a_x) \quad (4)$$

$$b = 90 - \arccos(a_z/A). \quad (5)$$

Here, a_x , a_y and a_z are the components of the coefficient vector \mathbf{a} and A , l and b are the dipole amplitude, longitude and latitude respectively. We note that many commonly used implementations of this approach do not implement full inverse variance noise weighting, as described by Eq. (2), but simply adopt $\mathbf{N} = \text{const.}$, and thereby in effect assign equal weight to all pixels. We do the same in the following.

2.2. Complications from partial sky coverage

The dipole parameters in Eq. (3)–(5) are subject to confusion from small-scale CMB fluctuations whenever a mask is applied. To see this, we expand the CMB fluctuation field into spherical harmonics as follows,

$$T(\hat{n}) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{n}), \quad (6)$$

where T is the CMB fluctuation map; $Y_{\ell m}$ are the spherical harmonic base functions; and the $a_{\ell m}$ are the associated weights. With access to the full celestial sphere, these coefficients may be computed as

$$a_{\ell m} = \int_{4\pi} T(\hat{n}) Y_{\ell m}^*(\hat{n}) d\Omega. \quad (7)$$

However, when masking parts of the sky with a mask \mathbf{M} the spherical harmonic base functions are no longer orthogonal, and the new so-called pseudo-harmonic coefficients read (Hivon et al. 2002)

$$\begin{aligned} \tilde{a}_{\ell m} &= \int_{4\pi} \mathbf{M} T(\hat{n}) Y_{\ell m}^*(\hat{n}) d\Omega \\ &= \sum_{\ell' m'} a_{\ell' m'} K_{\ell m, \ell' m'}[\mathbf{M}]. \end{aligned} \quad (8)$$

Here, K is called the coupling kernel, and quantifies the mutual dependence between any two modes $Y_{\ell m}$ and $Y_{\ell' m'}$. The explicit expression for the coupling kernel reads

$$K_{\ell m, \ell' m'} = \int_{4\pi} \mathbf{M}(\hat{n}) Y_{\ell m}(\hat{n}) Y_{\ell' m'}^*(\hat{n}) d\Omega. \quad (9)$$

Figure 2 shows K for multipoles up to $\ell_{\max} = 60$ for two different cases. In the no-mask case, shown in the top panel, all modes are orthogonal and therefore independent. When we apply a mask, shown in the bottom panel, nonzero off-diagonal elements appear. In other words, any higher-order mode will induce a spurious dipole contribution unless properly accounted for. The magnitude of this effect is dependent on the size of the mask, which is demonstrated in the following. Figure 3 highlights the $\ell = 1$ part of the coupling kernel and quantifies the coupling strength between the dipole and higher-order harmonic modes. The top and bottom panel correspond to the Y_{10} and Y_{11} modes respectively. The first peak of each line is the coupling strength with itself and therefore is unity. With access to only 36% of the sky (blue line), the mode coupling strength is significantly higher compared to the case where 92% of the sky is available (orange line). Another interesting observation is that the Y_{11} mode is more strongly coupled to higher modes than the Y_{10} mode. This is true because the former is oriented in the x -direction while the latter is oriented in the z -direction (see Fig. 1), therefore a mask applied along the equator will affect the coupling kernel in different ways.

Figure 4 provides an intuitive illustration of how a mask induces a spurious dipole. The top panel shows an ideal CMB realization drawn from a Λ CDM power spectrum with an identically vanishing dipole moment, $A = 0$. However, one of the largest hotspots on the sky happens, by chance, to align closely with the western half of the Galactic plane. After applying the Galactic mask, which gives this hotspot zero weight in the dipole fit, the result is a net dipole that points in the opposite direction. In this particular case, the result is a spurious dipole of $7.6 \mu\text{K}$ pointing toward the Eastern Hemisphere.




Fig. 3. Spherical harmonics coupling kernel K for dipole modes Y_{10} (top) and Y_{11} (bottom) for two differently masked skies. The blue and orange lines correspond to skies with an unmasked sky fraction of 36% and 92% respectively. The corresponding masks are shown in Fig. 8.

2.3. Wiener filtering and Gibbs sampling

Within the simple template fitting approach described above, the higher-order CMB fluctuations are treated as a random fluctuation term. Since the CMB fluctuations are Gaussian and isotropic, this term does not lead to any bias in the central estimates, but it does increase the variance. An alternative approach is to exploit the assumptions of isotropy and Gaussianity to estimate the CMB signal jointly with the dipole parameters, adopting the following data model,

$$\mathbf{d} = \mathbf{s} + \mathbf{n}, \quad (10)$$

where

$$\mathbf{s} = \mathbf{s}(\hat{n}) = \sum_{\ell,m} a_{\ell m} Y_{\ell m}(\hat{n}) \quad (11)$$

now is an isotropic and Gaussian random field with some angular power spectrum C_ℓ . We note that the previously defined dipole is contained within \mathbf{s} in the form of $\mathbf{T}\mathbf{a} = \sum_{\ell=0}^1 \sum_m a_{\ell m} Y_{\ell m}(\hat{n})$.

To estimate \mathbf{s} and C_ℓ jointly, we employ the Gibbs sampling algorithm described by Eriksen et al. (2004). This algorithm draws samples from the probability density $P(\mathbf{s}, C_\ell | \mathbf{d})$. Since it is difficult to sample directly from this joint distribution, we instead employ Gibbs sampling and perform consecutive sampling from each conditional density, $P(\mathbf{s} | C_\ell, \mathbf{d})$ and $P(C_\ell | \mathbf{s}, \mathbf{d})$, which according to Gibbs sampling theory will converge to being samples from the joint density $P(\mathbf{s}, C_\ell | \mathbf{d})$. Thus, the two Gibbs sampling steps are

$$\mathbf{s}^{i+1} \leftarrow P(\mathbf{s} | C_\ell^i, \mathbf{d}), \quad (12)$$

$$C_\ell^{i+1} \leftarrow P(C_\ell | \mathbf{s}^{i+1}). \quad (13)$$

For more information on the sampling process for the CMB power spectrum in Eq. (13), we again refer to Eriksen et al.




Fig. 4. Illustration of spurious dipole excited from higher-order moments through mask coupling. *Top panel:* an ideal CMB sky with identically vanishing dipole moment. *Bottom panel:* same realization, but with the Galactic plane masked. By accident, one large extended hotspot happens to lie on the Western Hemisphere inside the mask. Once this is removed by the mask, the net unmasked result is a dipole pointing toward the Eastern Hemisphere with an amplitude of $7.6 \mu\text{K}$.

(2004). However, we note that the inverse-gamma sampler described in that paper is only employed for multipoles $\ell \geq 2$ in the current analysis. For the first two elements, we manually set C_ℓ to a numerical large value of $10^{12} \mu\text{K}^2$, which effectively corresponds to imposing no informative priors on the monopole and dipole moments.

In our context we are mostly interested in the map sampling process in Eq. (12). In effect, the sky sample \mathbf{s}^i uses phase information in the data \mathbf{d} outside the mask to extrapolate into the missing pixels. The result is a constrained realization with the assumed power spectrum C_ℓ , such that the full map is a sample from the desired target distribution.

The map sampling process in Eq. (12) is performed in two steps. First we compute the so-called mean field map by solving the Wiener filter mapmaking equation for $\hat{\mathbf{x}}$,

$$(\mathbb{S}^{-1} + \mathbf{Y}'\mathbf{N}^{-1}\mathbf{Y})\hat{\mathbf{x}} = \mathbf{Y}'\mathbf{N}^{-1}\mathbf{d}. \quad (14)$$

Here, \mathbb{S} is a diagonal prior matrix that contains the assumed power spectrum, \mathbf{Y} denotes spherical harmonic transforms, and \mathbf{N} is the noise covariance matrix. The equation above is a Wiener filter and the resulting map $\hat{\mathbf{x}}$ is as such biased. To unbiased the full sample, we have to add a fluctuation term. We obtain the corresponding fluctuation map by solving the following expression for $\hat{\mathbf{y}}$,

$$(\mathbb{S}^{-1} + \mathbf{Y}'\mathbf{N}^{-1}\mathbf{Y})\hat{\mathbf{y}} = \mathbb{S}^{-1/2}\omega_1 + \mathbf{Y}'\mathbf{N}^{-1/2}\omega_2. \quad (15)$$

Here, ω_1 and ω_2 are two independent Gaussian white noise maps with zero mean and unit variance. The full sample \mathbf{s} is then the sum of $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$.




Fig. 5. Illustration of the Gibbs sampling procedure. The full sample s (bottom panel) is the sum of the mean field map \hat{x} (top panel) and a fluctuation map \hat{y} (middle panel).

The key advantage of sampling from the full distribution is that one simultaneously takes into account all multipole scales and all elements of the coupling kernel shown in Fig. 2. To illustrate this visually, Fig. 5 shows the different maps involved in the sampling procedure for some typical mask. The top panel shows the Wiener filter component. We note that this map contains small-scale structures in the unmasked regions, but only smooth structures in the masked regions. However, critically, it is not zero inside the mask. On the contrary, because of the assumptions of statistical isotropy and Gaussianity, the field inside the mask must show some degree of phase correlation with the unmasked regions, and this is precisely the information that allows partial reconstruction inside the mask.

Of course, this extrapolation is only supported for large angular scales. The fluctuation term, shown in the middle panel, therefore compensates for the fluctuation power that is lost due to the mask and noise, such that the sum of the two components, shown in the bottom panel, is a single full-sky map that is consistent with the original data. We note, however, that this map contains a significant stochastic component, and a full ensemble of such Gibbs samples is therefore required to adequately describe both the mean and covariance of the true underlying signal.

3. Simulations

We have now established two different methods for estimating dipole parameters from a foreground-cleaned CMB map with partial sky coverage. In order to assess the relative performance of these two methods, we perform Monte Carlo simulations for both algorithms, and compare the resulting uncertainties.

3.1. Monte Carlo procedure

For the standard template fitting approach, the procedure is defined as follows:

1. Generate N simulated CMB skies $d_{[1,\dots,N]}$ based on a Λ CDM power spectrum.
2. Apply mask \mathbf{M} to each sample.
3. Compute best-fit dipole parameters A , l and b for each sample with Eqs. (2)–(5).
4. Report the standard deviation for each parameter, σ_A , σ_l and σ_b .

For the Wiener filter approach, the procedure is similar, but additionally involves an intermediate sampling loop for each realization:

1. Generate N simulated CMB skies $d_{[1,\dots,N]}$ based on a Λ CDM power spectrum.
2. Apply mask \mathbf{M} to each sample.
3. For each masked CMB map:
 - (a) Draw n full-sky Wiener filter samples $s_{[1,\dots,n]}$.
 - (b) Compute dipole parameters A , l and b for each sample.
 - (c) Compute single-realization standard deviations σ_A^i , σ_l^i and σ_b^i .
4. Report the mean of σ_A^i , σ_l^i and σ_b^i .

The template fitting analysis is implemented using the HEALPix `fit_dipole` routine (Górski et al. 2005), while the Wiener filter analysis is performed using the Commander code Eriksen et al. (2004).

3.2. CMB simulations

For our simulated sky maps, we adopt the same model as in Eqs. (10)–(11), but this time we explicitly include support for an instrumental beam with Legendre expansion b_ℓ ,

$$\mathbf{d} = \mathbf{s} + \mathbf{n} = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} b_\ell a_{\ell m} Y_{\ell m} + \mathbf{n}. \quad (16)$$

We generate multipoles with $\ell \geq 2$ using Healpy’s `synfast`¹ routine, drawn from the *Planck* 2018 best-fit Λ CDM power spectrum (Planck Collaboration VI 2020). For the dipole component, we adopt $A = 3364.0 \mu\text{K}$, $l = 264.1 \text{ deg}$ and $b = 48.3 \text{ deg}$. No monopole is added.

For both the template and Wiener filter approaches, we establish ensembles of 100 realizations. To limit the computational speed involved in the Wiener filter stage, which requires iterative sampling, we choose to perform the analysis at a HEALPix resolution of $N_{\text{side}} = 32$, corresponding to a pixel size of about 1.8 deg. This is sufficient to capture all features that are relevant for dipole estimation.

For high-sensitivity experiments such as *Planck*, the direct template fitting approach is largely insensitive to specific details of the instrumental noise, as the dominant noise contributor is the CMB fluctuations, and \mathbf{n} may be safely disregarded in this framework. However, for the Wiener filter approach some care is

¹ <http://healpix.jpl.nasa.gov>




Fig. 6. Overview of signal and noise power spectra adopted for the Monte Carlo simulations. The solid line shows the best-fit *Planck* Λ CDM power spectrum, and the dashed line shows the same, but convolved with a 6° FWHM Gaussian beam. The dotted line shows the noise spectrum for white uncorrelated noise with $\sigma_0 = 1 \mu\text{K}$ per pixel at a HEALPix resolution of $N_{\text{side}} = 32$. The vertical dotted line shows the harmonic space truncation limit of $\ell_{\text{max}} = 95$.

warranted also for this term. In particular, the details of \mathbf{n} determine how aggressively the estimator is able to extrapolate into the masked region. For the method to be accurate and unbiased, it is important that the data model in Eq. (16) actually is a good representation of the observations in question.

First, since we perform our analysis at $N_{\text{side}} = 32$, the highest resolvable multipole moment is given roughly by $\ell_{\text{max}} \approx 3N_{\text{side}} = 96$. In order to suppress the signal above this ℓ_{max} , which is not supported by Eq. (16), we smooth the simulated CMB realizations with a 6 deg FWHM Gaussian beam with a Legendre expansion given by (Tegmark 1997)

$$b_\ell = \exp \left[-\frac{1}{2} \ell(\ell+1) \left(\frac{\text{FWHM} \cdot \pi}{180} \cdot \frac{1}{\sqrt{8 \ln 2}} \right)^2 \right]. \quad (17)$$

Next, to avoid ringing artifacts from the truncation limit around the mask edge, we have to ensure that the effective signal-to-noise ratio is negligible at ℓ_{max} . We do this by adding regularization noise with a standard deviation of $\sigma_0 = 1 \mu\text{K}$ per pixel. In harmonic space, this corresponds to a flat noise spectrum with an amplitude given by (Tegmark 1997)

$$N_\ell = \sigma_0^2 \cdot \frac{4\pi}{N_{\text{pix}}} \cdot \frac{\ell(\ell+1)}{2\pi}. \quad (18)$$

Figure 6 summarizes the simulated data in terms of signal and noise power spectra. We note that with these choices of parameters, the effective signal-to-noise ratio at $\ell = 95$ is smaller than 0.01. The individual components involved in the simulated map are illustrated in Fig. 7.

Since the elements of the coupling kernel depend on the sky fraction and on the shape of the mask, we repeat the Monte Carlo analysis for a variety of different masks with sky fractions ranging from 20 to 95%, shown in Fig. 8. These masks were already used for a similar purpose in Planck Collaboration III (2020).

4. Results

We are now ready to present the main result of this paper, which is a quantitative comparison of the template fitting and Wiener




Fig. 7. Components used to construct each Monte Carlo realization. These are the higher-order CMB fluctuations (*top panel*), the CMB dipole (*second panel*), and instrumental noise (*third panel*). *Bottom panel*: sum of the three components with a mask super-imposed.

filter approaches to dipole parameter estimation. For each mask shown in Fig. 8, we analyze 100 independent Monte Carlo realizations with both methods.

The main result is summarized in Fig. 9, where we show the mean dipole parameters and their corresponding statistical uncertainties as a function of the sky fraction. From top to bottom, the panels show the dipole amplitude, the longitude, and the latitude. The thick red lines show the mean of the derived solar dipole parameters using the Wiener filter technique, and




Fig. 8. Analysis masks used for the Monte Carlo comparison. The colored areas mark the regions of the sky that are to be ignored in the analysis. Analysis masks are typically characterized by their unmasked sky fraction, which is the fraction of the sky that remains after applying the mask. The unmasked sky fractions of the various masks are 92% (brown), 84% (yellow), 76% (purple), 68% (lilac), 60% (bright orange), 52% (red), 44% (bright red), 36% (green), 28% (cyan) and 20% (light blue).




Fig. 9. CMB solar dipole parameter uncertainties as a function of sky fraction; The thick red and black lines show the posterior means derived by the Wiener filter method and the traditional method respectively; The shaded bands are the corresponding $\pm 1\sigma$ confidence intervals; The horizontal dashed lines mark the true dipole parameters that were used as input for the simulations.

the regions shaded in red show the corresponding $\pm 1\sigma$ confidence intervals. The equivalent results derived using the traditional method are shown in black and gray. We mark the true dipole parameters as horizontal dashed lines.

We find that the uncertainties derived by the Wiener filter technique are significantly reduced compared to the traditional method. This effect is strongest for large sky fractions, for which it is easier to extrapolate into the masked regions. In contrast, for small sky fractions the extrapolation is very unreliable, and the two methods therefore give very similar results. It is worth to mention that the uncertainties derived with the traditional method decrease less smoothly with sky fraction than those of the Wiener filter method, which is a result of the morphologies of the various masks.

In some cases in Fig. 9 it may appear as if the traditional method yields smaller uncertainties than the Wiener filter at small sky fractions, especially for the longitude and latitude




Fig. 10. Absolute uncertainties of solar dipole parameters estimated with Wiener filter method (solid lines) and traditional method (dashed lines) as a function of sky fraction.




Fig. 11. Ratio of uncertainties derived by traditional method and Wiener filter method as a function of sky fraction.

where the upper confidence interval boundary is in fact below that of the Wiener filter method. However, this is not actually the case, since the gray band is also shifted to lower values due to statistical fluctuations of the derived means. To make this more explicit, we plot the absolute uncertainties of the two methods in Fig. 10, and their ratios in Fig. 11. We see that for large sky fractions (above $\approx 85\%$) the Wiener filter uncertainties are reduced by a factor of 10 or more for all parameters. At a sky fraction of about 50%, the uncertainties drop to roughly half of those of the traditional method.

5. Conclusions

In this paper we have quantitatively compared two numerical techniques for estimating the dipole parameters from a CMB map. The first method is basic template fitting regression with partial sky data, while the second method relies on Wiener filtering. The main difference between the two methods lies in their treatment of partial sky observations. Specifically, masking parts of the sky introduces couplings between small-scale CMB fluctuations and the dipole. The traditional template fitting

Table 1. Comparison of Solar dipole measurements from COBE, WMAP, and *Planck*.

Experiment	Amplitude [μK_{CMB}]	Galactic coordinates		Reference
		l [deg]	b [deg]	
COBE ^(a,b)	3358 \pm 23	264.31 \pm 0.16	48.05 \pm 0.09	Lineweaver (1997)
WMAP ^(c)	3355 \pm 8	263.99 \pm 0.14	48.26 \pm 0.03	Hinshaw et al. (2009)
LFI 2015 ^(b)	3365.5 \pm 3.0	264.01 \pm 0.05	48.26 \pm 0.02	Planck Collaboration II (2016)
HFI 2015 ^(d)	3364.29 \pm 1.1	263.914 \pm 0.013	48.265 \pm 0.002	Planck Collaboration VIII (2016)
LFI 2018 ^(b)	3364.4 \pm 3.1	263.998 \pm 0.051	48.265 \pm 0.015	Planck Collaboration II (2020)
HFI 2018 ^(d)	3362.08 \pm 0.99	264.021 \pm 0.011	48.253 \pm 0.005	Planck Collaboration III (2020)
NPIPE ^(a,c)	3366.6 \pm 2.7	263.986 \pm 0.035	48.247 \pm 0.023	Planck Collaboration Int. LVII (2020)

Notes. ^(a)Statistical and systematic uncertainty estimates are added in quadrature. ^(b)Computed with naive dipole estimator that does not account for higher-order CMB fluctuations. ^(c)Computed with Wiener filter estimator that estimates, and marginalizes over, higher-order CMB fluctuations jointly with the dipole. ^(d)Higher-order CMB fluctuations are accounted for by subtracting a dipole-adjusted CMB map from frequency maps prior to dipole estimation.

procedure disregards this coupling effect and simply treats the CMB fluctuations as a noise term. In contrast, the Wiener filter approach exploits the fact that these fluctuations represent a statistically isotropic and Gaussian random field to partially reconstruct the field inside the mask, and thereby reduce the overall uncertainties.

We apply both methods to an ensemble of 100 Monte Carlo realizations for sky fractions ranging from 20 to 95%, and derive uncertainties as a function of sky fraction. We find that the Wiener filter approach leads to significantly reduced uncertainties for typical sky fractions used in this type of analyses. For example, at $f_{\text{sky}} \approx 60\%$ the uncertainties are reduced by a factor of ≈ 3 , while at $f_{\text{sky}} \approx 85\%$ they are reduced by a factor of ≈ 8 .

Table 1 shows a comparison of measurements of the CMB solar dipole made by COBE, WMAP and (various generations of) *Planck*, and is a direct reproduction of Table 10 from [Planck Collaboration Int. LVII \(2020\)](#). Most of these analyses employed sky fractions around 80%. For this sky fraction, we see from Fig. 10 that the uncertainty on the amplitude due to small-scale CMB fluctuations is about $2.5 \mu\text{K}$ using the template fitting approach. In contrast, the total uncertainty for COBE was $23 \mu\text{K}$, and for WMAP it was $8 \mu\text{K}$. As such, the contribution from CMB confusion was subdominant for both these experiments. Nevertheless, it is important to note that WMAP was indeed the first experiment to implement this method for this particular purpose ([Hinshaw et al. 2009](#)), even though it may not have been critically important.

For *Planck*, the situation is fundamentally different. For this experiment, the raw uncertainty from instrumental noise and systematics is smaller than $1 \mu\text{K}$, and the CMB confusion has therefore become a dominant factor. In the low frequency instrument (LFI) processing, this contribution was simply included in the error budget, leading to a final uncertainty of $3 \mu\text{K}$. For the high frequency instrument (HFI), however, a different approach was taken, in that an estimate of the CMB fluctuations was removed from the raw data prior to template fitting. At first sight, this approach appears to eliminate the CMB confusion term entirely, evading the topic discussed in this paper. However, it is important to note that for this approach to be unbiased, the CMB template that is being subtracted must itself have a vanishing dipole moment. Determining the dipole moment of this map is therefore equivalent to the problem described in this paper. For the HFI analyses summarized in Table 1, this determination was performed with a very small mask, which in effect

assumes that the component separation method of choice (see [Planck Collaboration IV 2020](#) for details) is able to remove foregrounds accurately even in the central Galactic plane.

The last row in Table 1 lists results for the most recent *Planck* analysis, which is informally referred to as NPIPE. NPIPE represents the first joint analysis of the *Planck* LFI and HFI data sets, using a common machinery to reduce the raw time-ordered data into final sky maps. One important difference between these maps and earlier versions of the *Planck* data is that the NPIPE maps retain the solar dipole for both LFI and HFI. It is therefore, for the first time, possible to compute a single coherent all-*Planck* dipole with these sky maps. The results from this analysis are presented in Sect. 8 of [Planck Collaboration Int. LVII \(2020\)](#) and employ the Wiener filter methodology described in this paper. The values reported in Table 1 correspond to a sky fraction of 81%, which represents a compromise between maximizing available data and minimizing foreground-induced systematic effects. We believe that the values reported for NPIPE in Table 1 represent the most conservative and statistically robust estimate of the CMB solar dipole published to date.

Acknowledgements. We thank Jeff Jewell and Reijo Keskitalo for useful discussions. This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement numbers 776282, 772253 and 819478. Some of the results in this paper have been derived using the HEALPix package.

References

- Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, *ApJS*, 155, 227
 Fixsen, D. J. 2009, *ApJ*, 707, 916
 Fixsen, D. J., Cheng, E. S., Cottingham, D. A., et al. 1994, *ApJ*, 420, 445
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, *ApJS*, 180, 225
 Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, 567, 2
 Jewell, J., Levin, S., & Anderson, C. H. 2004, *ApJ*, 609, 1
 Lineweaver, C. H. 1997, *Microwave Background Anisotropies*, 16, 69
 Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, 142, 419
 Planck Collaboration II. 2016, *A&A*, 594, A2
 Planck Collaboration VIII. 2016, *A&A*, 594, A8
 Planck Collaboration I. 2020, *A&A*, 641, A1
 Planck Collaboration II. 2020, *A&A*, 641, A2
 Planck Collaboration III. 2020, *A&A*, 641, A3
 Planck Collaboration IV. 2020, *A&A*, 641, A4
 Planck Collaboration VI. 2020, *A&A*, 641, A6
 Planck Collaboration VII. 2020, *A&A*, 641, A7
 Planck Collaboration Int. LVII. 2020, *A&A*, 643, A42
 Tegmark, M. 1997, *Phys. Rev. D*, 56, 4514
 Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511