

# SILC: a new *Planck* Internal Linear Combination CMB temperature map using directional wavelets

Keir K. Rogers,<sup>1\*</sup> Hiranya V. Peiris,<sup>1</sup> Boris Leistedt,<sup>2,1</sup> Jason D. McEwen<sup>3</sup>  
and Andrew Pontzen<sup>1</sup>

<sup>1</sup>*Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

<sup>2</sup>*Department of Physics, New York University, 4 Washington Place, New York, NY 10003, USA*

<sup>3</sup>*Mullard Space Science Laboratory, University College London, Surrey RH5 6NT, UK*

23 August 2016

## ABSTRACT

We present new clean maps of the CMB temperature anisotropies (as measured by *Planck*) constructed with a novel internal linear combination (ILC) algorithm using directional, scale-discretised wavelets — Scale-discretised, directional wavelet ILC or SILC. Directional wavelets, when convolved with signals on the sphere, can separate the anisotropic filamentary structures which are characteristic of both the CMB and foregrounds. Extending previous component separation methods, which use the frequency, spatial and harmonic signatures of foregrounds to separate them from the cosmological background signal, SILC can additionally use morphological information in the foregrounds and CMB to better localise the cleaning algorithm. We test the method on *Planck* data and simulations, demonstrating consistency with existing component separation algorithms, and discuss how to optimise the use of morphological information by varying the number of directional wavelets as a function of spatial scale. We find that combining the use of directional and axisymmetric wavelets depending on scale could yield higher quality CMB temperature maps. Our results set the stage for the application of SILC to polarisation anisotropies through an extension to spin wavelets.

**Key words:** cosmic background radiation – methods: data analysis

## 1 INTRODUCTION

Accurate measurements of the cosmic microwave background (CMB) arguably form the bedrock of modern precision cosmology. In particular, the full-sky multifrequency CMB maps provided by three generations of satellite experiments — *COBE* (Mather et al. 1990; Bogess et al. 1992), *WMAP* (Bennett et al. 2003b) and *Planck* (Planck Collaboration et al. 2011) — represent milestones in our understanding of the cosmological model. However, to obtain a full-sky map of the CMB requires removing instrumental noise and signals due to astrophysical foregrounds (primarily in the Milky Way). Full-sky foreground-cleaned CMB maps are used for a wide variety of scientific purposes (see e. g., Planck Collaboration et al. 2015c,d).

There are numerous methods to perform foreground component separation. They broadly fall into two categories: blind methods which make minimal physical assumptions about the contributing signals and the so-called mixing matrix (which quantifies the strength of different components at different frequencies) and non-blind methods which are based on a physical modelling of the sky components. Examples of non-blind methods include

the Maximum Entropy Method (MEM) (Hobson et al. 1998) and the parametric Bayesian CMB Gibbs sampler Commander (Eriksen et al. 2006, 2008). Correlated Component Analysis (CCA) (Bedini et al. 2005) is a semi-blind method that estimates the mixing matrix based on second-order statistics. Spectral Estimation via Expectation Maximisation (SEVEM) (Martínez-González et al. 2003; Leach et al. 2008; Fernández-Cobos et al. 2012) is a template fitting technique. Examples of so-called blind source separation include the sparsity-based method Local-Generalized Morphological Component Analysis (L-GMCA) (Bobin et al. 2008, 2013) and the Spectral Matching Independent Component Analysis (SMICA) (Cardoso et al. 2008), although the latter work does discuss how the choice of component model affects the blindness of this method. Of particular interest to this work is another blind method, the Internal Linear Combination (ILC), most recently implemented by the Needlet ILC (NILC) (Delabrouille et al. 2009). In its component separation analysis, the *Planck* Collaboration used Commander, NILC, SEVEM and SMICA (Planck Collaboration et al. 2015a). See, e. g., Delabrouille et al. (2009); Bobin et al. (2013) for reviews of CMB component separation methods.


The ILC computes a weighted sum of CMB maps as measured at multiple frequencies. These weights are constrained to sum to unity at each point in the map, ensuring that the CMB signal

\* E-mail: keir.rogers.14@ucl.ac.uk

is conserved, assuming that it is equal at each frequency. Under this constraint, the weights are calculated by minimising the empirical variance of the ILC map, which in turn minimises the variance of the error in CMB reconstruction (assuming the CMB and foregrounds and the CMB and noise are respectively uncorrelated). The variances we minimise are empirical in that they are calculated using the data themselves. In this way, the weights are calculated to remove foreground and noise, revealing the underlying primordial CMB anisotropies. The ILC method was originally used by the *WMAP* Collaboration (Bennett et al. 2003a) and then extended by Eriksen et al. (2004) through an analytical calculation of the weights. One limitation of the original ILC approach is the extent of localisation of the weights. The initial versions calculated different weights in separate parts of the sky (*e.g.*, Bennett et al. 2003a split the Galactic region into 11 parts). In order to further remove local contamination, the weights can be allowed to vary across the sky and also at different spatial scales. Tegmark et al. (2003) made an ILC map allowing the weights to vary at each multipole, as well as within different regions of the sky. A direct extension of this work is to make use of both spatial and frequency information simultaneously using wavelets. The weights are then defined across wavelet scales and within wavelet coefficient maps on the sky.

Wavelets are functions that are localised in both real and frequency space. To analyse full-sky CMB maps, wavelets defined on the surface of a sphere are required. A number of wavelet frameworks on the sphere have been developed recently (Antoine & Vandergheynst 1999, 1998; Baldi et al. 2009; Barreiro et al. 2000; Geller et al. 2008; Geller & Marinucci 2010, 2011; Leistedt et al. 2013, 2015; Narcowich et al. 2006; McEwen et al. 2006; McEwen & Scaife 2008; McEwen et al. 2011; McEwen et al. 2013, 2014; McEwen et al. 2015a,b; Marinucci et al. 2008; Sanz et al. 2006; Starck et al. 2006, 2009; Wiaux et al. 2005; Wiaux et al. 2008). In particular, needlets (Narcowich et al. 2006; Marinucci et al. 2008; Baldi et al. 2009) have been used in the latest generation of ILC methods. Needlets are a set of axisymmetric kernels defined on the surface of a sphere. Each member of the set has compact support in harmonic space over different multipole ranges. When each needlet is convolved with a signal defined on the sphere, the resulting signal (*i.e.*, needlet coefficients) also has compact harmonic support. NILC (Delabrouille et al. 2009) computes its weights by considering needlet scales separately (harmonic localisation) and then different parts of each needlet coefficient map separately (spatial localisation). The needlets are constructed in such a way that the original signal can be recovered from its needlet coefficients with no loss of information (in practice, small losses can be introduced by approximate spherical harmonic transforms). NILC has been very successful at forming clean full-sky CMB maps, which contain very little residual foreground and noise contamination.

In this work, we introduce the Scale-discretised, directional wavelet ILC or SILC, which extends the wavelet ILC framework by localising the calculation of ILC weights in an additional domain. We use wavelets which are not only harmonically-localised but also directional (Wiaux et al. 2008; McEwen et al. 2013; McEwen et al. 2015b). Unlike needlets, which are axisymmetric on the sphere, directional wavelets are non-axisymmetric, *i.e.*, the kernels are “squeezed.” This means that for one wavelet scale, one axisymmetric kernel is replaced by a number of complementary directional kernels, each with a different orientation. When these directional wavelets are convolved with a signal on the sphere, (within each scale) different orientations of signal structure are separated. This directional localisation allows the ILC weights to be additionally fine-tuned to better remove foreground and noise, in particu-



**Figure 1.** The spatial localisation on the sphere of directional, scale-discretised wavelets. Each sub-plot shows a representation of a directional wavelet kernel at different scales, where red, raised parts show positive wavelet response and blue, depressed parts show negative wavelet response. *From left to right, top to bottom:* wavelet scale index  $j$  decreases. The number of directions per wavelet scale  $N = 3$ . Therefore, for complete reconstruction at each scale, the above wavelets would be complemented by two more wavelets of the same size but of a different orientation on the sphere. This figure is adapted from McEwen et al. (2013).


lar for signals with filamentary structure. Furthermore, directional wavelets exhibit exact reconstruction, allowing them to be embedded in an ILC such that no signal is lost.

SILC is being developed with the goal of analysing CMB polarisation components through an extension to spin, directional wavelets (McEwen et al. 2015b; McEwen et al. 2014; Leistedt et al. 2015), which are expected to be well-suited to localising the complex filamentary morphologies of polarised foregrounds. As a precursor step, in this work we test SILC on the scalar temperature field in order to demonstrate the quantitative consistency of its foreground cleaning performance compared with existing component separation methods, and to identify possible optimisations for the extension to spin fields.

Directional wavelets are explained briefly in § 2. In § 3, the SILC algorithm is explained in detail. Various sources of error in the method are considered in § 4. In § 5, we compare our method to previous component-separation methods. The application to *Planck* simulations (§ 6) is followed by application to *Planck* data (§ 7). We discuss the results and error estimation based on the data in § 8 and conclude in § 9.

## 2 DIRECTIONAL WAVELETS

Directional, scale-discretised wavelets on the sphere that support exact reconstruction have been developed in Wiaux et al. (2008), McEwen et al. (2013) and McEwen et al. (2015b), while their localisation properties have been studied in McEwen et al. (2015a). Figure 1 shows an example of the spatial localisation of directional wavelets. Larger wavelet scales have larger kernels, and when these are convolved with signals defined on the sphere (such the CMB and astrophysical foregrounds), signal structure with the same scale and orientation as the wavelet is isolated. The kernels in Fig. 1 are shown for a single direction and (for complete reconstruction) would be complemented by two more sets of kernels of the same sizes but rotated to different orientations. Figure 2 shows an example of the harmonic localisation of directional wavelets (for the



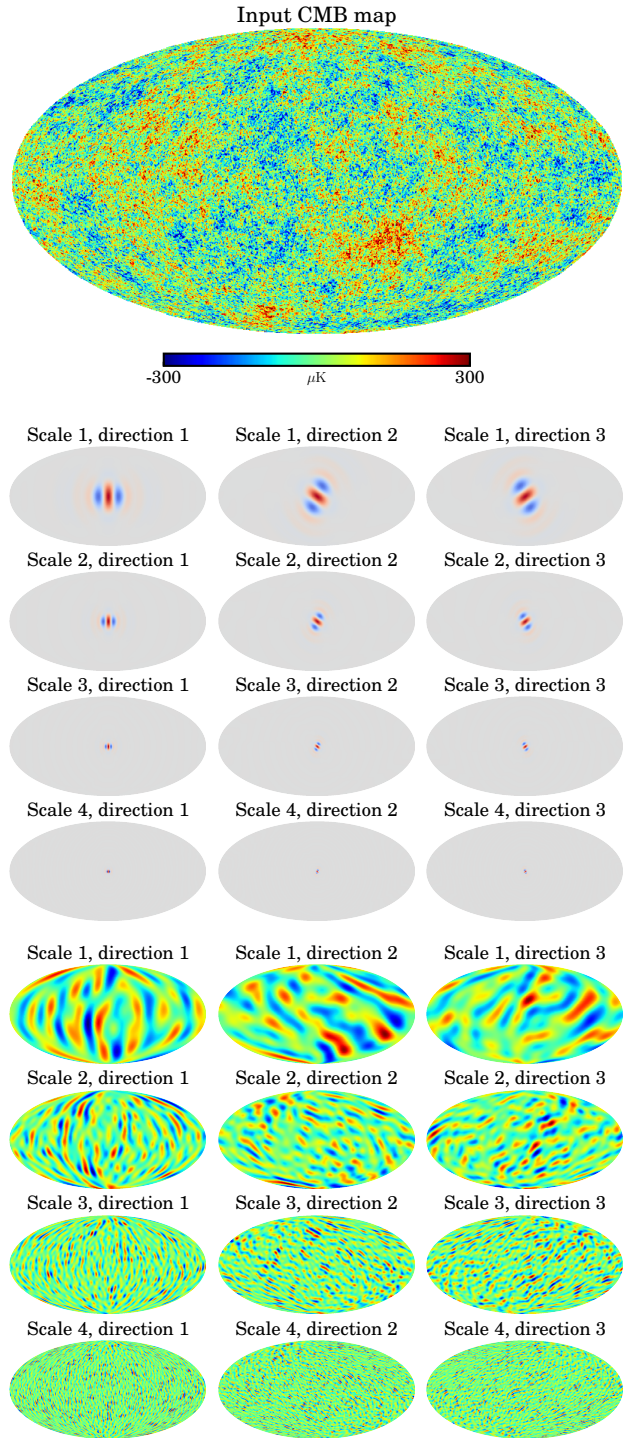
**Figure 2.** The harmonic response of the directional wavelets used in this work, where  $j$  specifies the wavelet scale. Increasing  $j$  corresponds to a smaller wavelet kernel and so a multipole range on smaller scales (*i. e.*, larger multipoles  $\ell$ ). The largest wavelet scale (Scal.) is the scaling function (§ 3.4). The two smallest wavelets are harmonically truncated at  $\ell = 3600$  but are smoothly tapered to zero from  $\ell = 3400$  to  $\ell = 3600$  (the two dotted lines) by the beam tapering discussed in § 3.3. The band-limits of the above wavelets are given in Table 1.

**Table 1.** The harmonic band-limits [ $\ell_{\min}^j, \ell_{\max}^j$ ] of the directional wavelets used in this work.  $\ell_{\text{peak}}^j$  is the multipole at which each wavelet has its maximum response. The final column shows the number of equiangular samples per wavelet coefficient map  $N_{\text{samp}}^j$ .

Wavelet scale $j$	$\ell_{\min}^j$	$\ell_{\text{peak}}^j$	$\ell_{\max}^j$	$N_{\text{samp}}^j$
Scal.	0	64	64	8,385
0	32	64	128	33,153
1	64	128	256	131,841
2	128	256	512	525,825
3	256	512	706	998,991
4	542	705	918	1,688,203
5	705	917	1193	2,850,078
6	917	1192	1551	4,815,856
7	1192	1550	2015	8,126,496
8	1550	2015	2540	12,910,821
9	2116	2539	3048	18,589,753
10	2539	3047	3600	25,930,801
11	3047	3600	3600	25,930,801

wavelets used in this work). The harmonic supports of the wavelets overlap, with each wavelet covering a finite set of multipoles. Figure 3 shows an example of directional wavelet decomposition as applied to the CMB. Although the CMB anisotropies are statistically Gaussian, the CMB spots on the sky demonstrate anisotropy as a function of scale (Bond & Efstathiou 1987). When the CMB is convolved with directional wavelets, structure of different orientations is separated. This further supports the use of directional wavelets in CMB analysis: both filamentary foreground structure and the CMB itself are better localised. This particularly applies in the case of polarisation, as will be discussed in § 9 when we consider extensions to our method. For a mathematical description of directional wavelets, see § 3.4.

In the spherical harmonic transforms used in the computation of directional wavelet coefficient maps, we adopt the sam-



**Figure 3.** The CMB (*top map*) decomposed into directional wavelet coefficient maps (*bottom section*). The wavelet kernels are shown (*middle section*), where red indicates positive response and blue indicates negative. In the full analysis, we also include smaller wavelets than we show above.

pling scheme on the sphere of McEwen & Wiaux (2011) (hereafter MW sampling), rather than, *e. g.*, HEALPix sampling (Górski et al. 2005), although in principle HEALPix could be used if desired. The corresponding sampling theorem of McEwen & Wiaux (2011) shows that the MW sampling scheme requires fewer samples for a band-limited signal than any other sampling theorem.



Additionally, the use of a separation of variables and fast Fourier transforms (FFTs) yields a numerically efficient algorithm. In particular, our spherical harmonic transforms are theoretically exact, unlike HEALPix. This allows one to manipulate signals with the minimal number of samples and to perform the numerous spherical harmonic transforms involved in the ILC algorithm without any loss of information (other than that due to the finite representation of floating point numbers). Our final map products, however, are provided in HEALPix format. Finally, MW sampling of spin signals on the sphere requires no additional computational complexity and this will be vital in the extension of our method to polarisation  $E$  and  $B$  modes (§ 9). Further details on MW sampling are given in § 3.4.

### 3 METHOD

We start by outlining the SILC algorithm. The steps are explained in more detail in the subsequent subsections (§ 3.1 to 3.6). We discuss our numerical implementation in § 3.7.

(1) The raw input data are multifrequency full-sky maps of CMB temperature fluctuations. These maps use the HEALPix format. (See § 3.1.) The model we employ for the raw data is explained in § 3.2.

(2) The maps are “pre-processed” by inpainting in a small point source mask (see § 3.6).

(3) The input maps are converted to thermodynamic (CMB) temperature (if necessary). For *Planck* temperature data, the 545 GHz and 857 GHz maps are converted from spectral flux density per unit solid angle (MJy Sr<sup>-1</sup>) to CMB temperature (K<sub>CMB</sub>) by the unit conversions given in the *Planck* 2015 Release Explanatory Supplement<sup>1</sup>.

(4) The maps are each convolved to have the same effective beam (see § 3.3).

(5) Each input map is converted into a set of wavelet coefficient maps. This separates both the scale and orientation of structure within each map. These wavelet coefficient maps use MW sampling (McEwen & Wiaux 2011). (See § 3.4.)

(6) The ILC method is then applied separately to each wavelet scale and orientation. For each scale and orientation, the multifrequency wavelet coefficient maps are weighted and added to form a single wavelet coefficient map that contains mainly CMB signal, as well as some residual foreground and noise. These weights are allowed to vary from wavelet coefficient to wavelet coefficient. The calculation of these weights is explained in § 3.5.

(7) The final ILC wavelet coefficient maps are synthesised to form the final product: a full-sky map of CMB temperature fluctuations (with some residual foreground and noise). The final map uses the HEALPix format. (See § 3.4.)

(8) The final map is inpainted in a small point source mask (see § 3.6).

<sup>1</sup> *Planck* 2015 Release Explanatory Supplement: UC CC tables ([http://wiki.cosmos.esa.int/planckpla2015/index.php/UC\\_CC\\_Tables](http://wiki.cosmos.esa.int/planckpla2015/index.php/UC_CC_Tables)). For the 545 GHz map, the unit conversion is (58.0356 ± 0.0278) MJy Sr<sup>-1</sup> K<sub>CMB</sub><sup>-1</sup> and for the 857 GHz map, the unit conversion is (2.2681 ± 0.0270) MJy Sr<sup>-1</sup> K<sub>CMB</sub><sup>-1</sup>.

### 3.1 Input data

Our main CMB temperature map products use full-mission 2015 release *Planck* temperature maps as their input<sup>2</sup>. All nine frequency channels are used. At 70 GHz, we use the higher-resolution version at  $N_{\text{side}} = 2048$ . We also use the full-mission Full Focal Plane 8 (FFP8) simulations (Planck Collaboration et al. 2015b) without bandpass mismatch<sup>3</sup>.

### 3.2 Data model

Each full-sky temperature map can be modelled (*e. g.*, Basak & Delabrouille 2012) as

$$T^{\text{OBS},c}(\hat{n}) = \int_{\hat{n}'} d\hat{n}' B^c(\hat{n}, \hat{n}') T^{\text{SIG},c}(\hat{n}') + T^{\text{N},c}(\hat{n}), \quad (1)$$

where the signal component can further be decomposed as

$$T^{\text{SIG},c}(\hat{n}) = a^c T^{\text{CMB}}(\hat{n}) + T^{\text{FG},c}(\hat{n}). \quad (2)$$

$T^{\text{CMB}}(\hat{n})$  is the CMB component at a point on the sky  $\hat{n}$ .  $T^{\text{FG},c}(\hat{n})$  and  $T^{\text{N},c}(\hat{n})$  are respectively the foreground and detector noise components for frequency channel  $c$ .  $a^c$  is the calibration coefficient for the CMB for each channel. The overall signal component is smoothed by a beam function  $B^c(\hat{n}, \hat{n}')$  due to the finite resolution of the observations. However, the noise component is not smoothed by the beam. Here we assume the beam to be circularly symmetric. Therefore, the beam can be represented as a sum over Legendre polynomials,

$$B^c(\hat{n}, \hat{n}') = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} B_{\ell}^c P_{\ell}(\hat{n}, \hat{n}'). \quad (3)$$

We can recast Eq. (1) in the spherical harmonic representation as

$$a_{\ell m}^{\text{OBS},c} = a^c B_{\ell}^c a_{\ell m}^{\text{CMB}} + B_{\ell}^c a_{\ell m}^{\text{FG},c} + a_{\ell m}^{\text{N},c} \quad (4)$$

where  $a_{\ell m}$  are the coefficients of spherical harmonics  $Y_{\ell m}(\hat{n})$ .

### 3.3 Beam convolution

Equation (4) shows that each frequency channel  $c$  has a different beam transfer function  $B_{\ell}^c$ . To replace each beam with a channel-independent resolution, we perform a deconvolution/convolution procedure to give spherical harmonic coefficients

$$a_{\ell m}^c = \frac{B_{\ell}^{\text{EFF}}}{B_{\ell}^c} a_{\ell m}^{\text{OBS},c}, \quad (5)$$

where  $B_{\ell}^{\text{EFF}}$  is the final (effective) beam transfer function of our map products. For *Planck* data, we use a Gaussian beam with a FWHM of  $5'$  as our input beam. We taper this beam to zero from  $\ell = 3400$  to  $\ell = 3600$  using a Fermi function. This suppresses any small-scale power aliasing due to having harmonically-truncated wavelets in this multipole range. Convolution with beam transfer functions ignores the non-axisymmetric component of the beams; these will remain in the input maps but are assumed to be small.

This deconvolution/convolution procedure does not correctly

<sup>2</sup> <http://pla.esac.esa.int/pla>

<sup>3</sup> FFP8 simulations are also available with bandpass mismatch, accounting for differences in the bandpasses of detectors nominally at the same frequency, leading to spurious signals in the frequency maps.



handle the noise component of our input maps. Equation (5) can be expanded (using Eq. (4)) as

$$a_{\ell m}^c = B_{\ell}^{\text{EFF}}(a_{\ell m}^c \text{CMB} + a_{\ell m}^{\text{FG},c}) + \frac{B_{\ell}^{\text{EFF}}}{B_{\ell}^c} a_{\ell m}^{\text{N},c}, \quad (6)$$

where  $a^c$  are the CMB calibration coefficients (not to be confused with the inverse spherical harmonic transform of harmonic coefficients  $a_{\ell m}^c$ ). The final resolution of an ILC map is usually chosen to match the best resolution of the input maps. Therefore, for all but the highest resolution channel and for all  $\ell$ ,  $B_{\ell}^{\text{EFF}} > B_{\ell}^c$ . This has the effect of increasing the noise contribution of the input maps, particularly at high  $\ell$  and for low-resolution maps, where  $B_{\ell}^{\text{EFF}} \gg B_{\ell}^c$ . We use the *Planck* beam transfer functions as provided in the Reduced Instrument Model (RIMO)<sup>4</sup>. For the LFI beams, we use Gaussian approximations with FWHM 32.33', 27.01' and 13.25' for 30, 44 and 70 GHz respectively. Following [Planck Collaboration et al. \(2014b\)](#), the deconvolved beams are thresholded such that the  $B_{\ell}^c$  is set to the value given in the RIMO or 0.001, whichever is larger. This prevents the last term in Eq. (6) from becoming so large that numerical errors occur. Although we lose accuracy in the deconvolution process, the contribution of the channels in the multipole ranges affected is highly attenuated in the ILC weights in any case.

### 3.4 Wavelet analysis and synthesis

The wavelet ILC method requires the decomposition of each band-limited temperature map  $T^c(\hat{n})$  into a set of wavelet coefficient maps  $W^{\Psi^j}$ : in our case, directional, scale-discretised wavelets ([Wiaux et al. 2008](#); [McEwen et al. 2013](#); [McEwen et al. 2015b](#)). A general introduction was provided in § 2 — here we provide some technical details of the implementation. We drop the  $c$  superscript on  $T$  for the rest of this subsection since each map is analysed using the same wavelets. The wavelet coefficients are defined as the directional convolution of  $T$  with wavelets defined on the sphere  $\Psi^j \in L^2(\mathbb{S}^2)$  (specifically those shown in Fig. 2) where index  $j$  denotes the wavelet scale. Importantly, directional wavelets yield coefficients  $W^{\Psi^j}(\hat{\rho})$  that live on the space of three-dimensional rotations, *i. e.*, the rotation group  $\text{SO}(3)$ :

$$W^{\Psi^j}(\hat{\rho}) \equiv \langle T, \mathcal{R}_{\hat{\rho}} \Psi^j \rangle = \int_{\mathbb{S}^2} d\hat{n} T(\hat{n}) (\mathcal{R}_{\hat{\rho}} \Psi^j)^*(\hat{n}), \quad (7)$$

where  $d\hat{n}$  is the usual invariant measure on the sphere,  $*$  denotes complex conjugation and the rotation operator is defined by

$$(\mathcal{R}_{\hat{\rho}} \Psi^j)(\hat{n}) \equiv \Psi^j(\mathbf{R}_{\hat{\rho}}^{-1} \hat{n}), \quad (8)$$

where  $\mathbf{R}_{\hat{\rho}}$  is the three-dimensional rotation matrix corresponding to  $\mathcal{R}_{\hat{\rho}}$ . In these equations,  $\hat{\rho} = (\theta, \phi, \chi) \in \text{SO}(3)$  denotes the Euler angles (in the  $zyz$  convention) with colatitude  $\theta \in [0, \pi]$ , longitude  $\phi \in [0, 2\pi)$  and direction  $\chi \in [0, 2\pi)$ <sup>5</sup>. In other words, the wavelet coefficients probe directional structure in  $T$  with  $\chi$  corresponding to the orientation about each point  $(\theta, \phi)$  on the sphere.

Following the directional construction of scale-discretised wavelets ([Wiaux et al. 2008](#); [McEwen et al. 2013](#); [McEwen et al.](#)

[2015b](#)), wavelets are defined by their spherical harmonic coefficients in factorised form:

$$\Psi_{\ell n}^j \equiv \kappa^j(\ell) s_{\ell n}, \quad (9)$$

where  $\kappa^j(\ell)$  sets the harmonic localisation (Fig. 2) and  $s_{\ell n}$  sets the directional localisation.

In the original definition of scale-discretised wavelets, the size of all harmonic kernels (setting the harmonic localisation of the wavelets) is parameterised by a unique wavelet dilation parameter  $\lambda \in \mathbb{R}_+^*$ ,  $\lambda > 1$ . Similarly, the number of directions is set by a unique azimuthal band-limit  $N$ . These two parameters respectively characterise  $\kappa^j(\ell)$  and  $s_{\ell n}$  for all  $j$ . In this work, we vary  $\lambda$  as a function of multipole in order to allow more flexible harmonic localisation. We achieve this by defining different values of  $\lambda$  in different multipole regions and then stitching together harmonically-truncated wavelets at the region boundaries. We use the values  $\lambda = 2, 1.3, 1.2$  with transitions at the multipoles  $\ell = 512, 2015$ . If at a transition multipole the harmonic peak of the larger wavelet doesn't equal the peak of the smaller wavelet, then a small amount of unit response is used so that the two wavelets can be continuously combined. Wavelets constructed in this manner satisfy the standard admissibility criterion required for exact reconstruction. The harmonic tiling of the resulting wavelets is shown in Fig. 2. The technical details of the construction of each kernel is described in [McEwen et al. \(2015a\)](#). Finally, we use a single parameter  $N$  for all scales, *i. e.*, each wavelet is divided into the same number of directions. However, a possible extension of this work is to vary  $N$  as a function of scale  $j$ , *e. g.*, by using curvelet kernels ([Chan et al. 2015](#)) or other directional optimisations.

In the case of a single parameter  $\lambda$ , the limits of the wavelet harmonic window for scale  $j$  are simply  $(\ell_{\min}^j, \ell_{\max}^j) = (\lambda^{j-1}, \lambda^{j+1})$ , with their peak response at  $\lambda^j$ . In our hybrid scheme, this property remains but  $j$  and  $\lambda$  must be adjusted in each harmonic region. The full details of our tiling are given in Table 1. When the limits of the harmonic windows of the maximum wavelet scales extend beyond the overall band-limit  $\ell_{\max}$ , the windows are truncated at  $\ell_{\max}$ . Finally, note that a scaling function  $W^{\Phi}$  is needed to capture the very low frequency content of the signal. It is axisymmetric and the corresponding scaling coefficients therefore live on the sphere. Here we do not give the full details of the construction of the scaling function or the factors  $\kappa^j(\ell)$  and  $s_{\ell n}$  since these can be straightforwardly reproduced by following previous approaches ([Wiaux et al. 2008](#); [McEwen et al. 2013](#); [McEwen et al. 2015b](#)) and using Table 1.

To apply the ILC algorithm, the above continuous wavelet coefficients must be discretised. Since they live on the rotation group  $\text{SO}(3)$ , we represent them using the sampling scheme of [McEwen et al. \(2015\)](#), which is a generalisation of the MW sampling scheme ([McEwen & Wiaux 2011](#)). Because our wavelets have well-defined band-limits, this approach allows a multiresolution scheme where each scale is pixellated with a minimal number of samples. In practice, the  $j$ -th wavelet scale has a band-limit  $\ell_{\max}^j$  and is only evaluated at locations  $(\theta_t^j, \phi_p^j, \chi_n)$  with  $t \in \{0, 1, \dots, \ell_{\max}^j\}$ ,  $p \in \{0, 1, \dots, 2\ell_{\max}^j\}$  and  $n \in \{0, 1, \dots, N-1\}$ . Although wavelet coefficients are evaluated at discrete samples only, for a band-limited signal they capture the total information content of the underlying continuous wavelet coefficient representation, probed up to harmonic band-limit  $\ell_{\max}^j$  and azimuthal band-limit  $N$ . This is thanks to the sampling theory on the rotation group  $\text{SO}(3)$  of [McEwen et al. \(2015\)](#). In the full ensemble of realisations, the ILC (see § 3.5 for details) has no sensitivity to the choice of coordinate convention for directions  $\chi$ . In a single realisation, there will be some marginal

<sup>4</sup> *Planck* 2015 Release Explanatory Supplement: The 2015 instrument model ([http://wiki.cosmos.esa.int/planckpla2015/index.php/The\\_RIMO](http://wiki.cosmos.esa.int/planckpla2015/index.php/The_RIMO)).

<sup>5</sup> We adopt the  $zyz$  Euler convention corresponding to the rotation of a physical body in a fixed coordinate system about the  $z$ ,  $y$  and  $z$  axes by  $\chi$ ,  $\theta$  and  $\phi$ , respectively.

sensitivity to this choice manifesting in the localisation of the empirical covariances we use. However, this effect is sub-dominant to the choice of  $N$ , on which we concentrate our analysis.

After the ILC method (see § 3.5) has been applied to the sets of wavelet coefficient maps, there is one final map  $W^{j,\text{ILC}}(\hat{\rho})$ , for each wavelet scale  $j$ , living on  $\text{SO}(3)$  and including the multiple orientations  $\chi_0, \dots, \chi_{N-1}$ . The additional axisymmetric scaling coefficients  $W^{\text{ILC}}(\hat{n})$  live on the sphere. The final temperature map  $T^{\text{ILC}}(\hat{n})$  is synthesised by

$$T^{\text{ILC}}(\hat{n}) = \int_{\mathbb{S}^2} d\hat{n}' W^{\text{ILC}}(\hat{n}') (\mathcal{R}_{\hat{n}'} \Phi)(\hat{n}) + \sum_{j=j_{\min}}^{j_{\max}} \int_{\text{SO}(3)} d\hat{\rho} W^{j,\text{ILC}}(\hat{\rho}) (\mathcal{R}_{\hat{\rho}} \Psi^j)(\hat{n}), \quad (10)$$

where  $d\hat{\rho}$  is the usual invariant measure on the rotation group. This final ILC temperature map is pixellated using the HEALPix format from its spherical harmonic coefficients  $T_{\ell m}^{\text{ILC}}$ .

The wavelet analysis and synthesis are performed using the latest version of the S2LET<sup>6</sup> code (Leistedt et al. 2013; McEwen et al. 2015), which in turn relies on the SSHT<sup>7</sup> (McEwen & Wiaux 2011) and S03<sup>8</sup> (McEwen et al. 2015) codes to compute spin spherical harmonics and Wigner transforms exactly and efficiently using the MW sampling scheme. Thanks to the sampling theorem, the wavelet coefficients can be transformed using Wigner transforms without any loss of information (McEwen et al. 2015).

### 3.5 ILC method

Following the wavelet analysis of the input maps (see § 3.4), there is a wavelet coefficient map  $W_{jnk}^c$  for each channel  $c$ , scale  $j$  and orientation  $n$  with a pixel index  $k$ . Using this more compact notation, we conflate the scaling coefficient map with the wavelet coefficient maps as the ILC method applies in exactly the same fashion. The ILC estimate of the CMB signal at each wavelet scale and orientation is defined as a weighted sum of the wavelet coefficient maps at that scale and orientation

$$W_{jnk}^{\text{ILC}} \equiv \sum_{c=1}^{N_c} \omega_{jnk}^c W_{jnk}^c, \quad (11)$$

where  $\omega_{jnk}^c$  are the weights (which are allowed to vary across the scale and orientation of the signal as well as pixel space) and  $N_c$  is the number of input channels.

We impose a constraint on the weights (to ensure that the CMB signal is preserved) such that

$$\sum_{c=1}^{N_c} \omega_{jnk}^c = 1. \quad (12)$$

Assuming that the CMB and foregrounds and the CMB and noise are respectively uncorrelated, the variance of the error in the result is minimised when the variance of the ILC map itself is minimised. The resulting weights are given by

$$\omega_{jnk}^c = \frac{\sum_{c'=1}^{N_c} (R_{jnk}^{-1})^{cc'} a^{c'}}{\sum_{c=1}^{N_c} \sum_{c'=1}^{N_c} a^c (R_{jnk}^{-1})^{cc'} a^{c'}}, \quad (13)$$

where the true covariance matrices at scale  $j$ , orientation  $n$  and pixel

$k$ ,  $(R_{jnk})^{cc'} = \langle W_{jnk}^c W_{jnk}^{c'} \rangle$  (where the angled brackets indicate an ensemble average). For a derivation of Eq. (13), see Tegmark et al. (2003); Eriksen et al. (2004).

In this work, we estimate covariance matrices empirically by the following procedure (as used in Basak & Delabrouille 2012; Planck Collaboration et al. 2015a). We start by calculating at each pixel  $k$ :

$$(R_{jnk}^{\text{approx}})^{cc'} = W_{jnk}^c W_{jnk}^{c'}. \quad (14)$$

We then smooth each element of the above matrix by a Gaussian beam  $w_j(k, k')$  in pixel space to form the empirical estimates of covariance matrices

$$(\hat{R}_{jnk})^{cc'} = \sum_{k'=1}^{N_{\text{samp}}^j} w_j(k, k') (R_{jnk'}^{\text{approx}})^{cc'} \quad (15)$$

where  $N_{\text{samp}}^j$  is the total number of pixels in a given map at scale  $j$ . For computational efficiency, we perform this smoothing in harmonic space:

$$(\hat{R}_{jnk})^{cc'} = \sum_{\ell=0}^{2\ell_{\max}^j} \sum_{m=-\ell}^{\ell} w_j^\ell(r_{jn}^{\ell m})^{cc'} Y_k^{\ell m}, \quad (16)$$

where  $(r_{jn}^{\ell m})^{cc'}$  are the harmonic coefficients of the maps formed by the elements of matrices  $(R_{jnk}^{\text{approx}})^{cc'}$ ,  $w_j^\ell$  is a Gaussian beam transfer function and  $Y_k^{\ell m}$  are the spherical harmonics evaluated at pixel  $k$ .

The size of the Gaussian kernel used to smooth the covariance matrices is chosen to be proportional to the size of the wavelet used to form a particular set of wavelet coefficient maps<sup>9</sup>. In general the estimation of covariance matrices in ILC methods could be further optimised. It may be preferable to dynamically adapt the smoothing kernel used based on local data. Delabrouille et al. (2009) suggested using a larger kernel at high Galactic latitudes where Galactic emission does not vary so much and a smaller kernel towards the Galactic equator where emission is more complex. It could involve masking equatorial regions when estimating the covariance at higher latitudes (somewhat akin to Planck Collaboration et al. (2014b)). It could involve convolving the maps of elements of covariance matrices with the same directional wavelet in order to pick out how the local covariance follows the directional structure of the underlying signal. As mentioned above, in this work, we use a similar method as in previous work for ease of comparison.

It is also worth discussing the upper limit on the summation over  $\ell$  in Eq. (16). We first note the general rule that for the product of two spherical harmonics (Driscoll & Healy 1994)

$$Y_{\ell_1, m_1}(\hat{n}) Y_{\ell_2, m_2}(\hat{n}) = \sum_{L=|\ell_1-\ell_2|}^{\ell_1+\ell_2} a_{L, m_1+m_2} Y_{L, m_1+m_2}(\hat{n}) \quad (17)$$

where  $Y_{L, m_1+m_2}(\hat{n})$  is defined to be zero if  $|m_1 + m_2| > L$ . It follows that for the product of two band-limited maps  $M(\hat{n}) = \sum_{\ell=\ell_1}^{\ell_2} \sum_{m=-\ell}^{\ell} m_{\ell m} Y_{\ell m}(\hat{n})$  and  $N(\hat{n}) = \sum_{\ell=\ell_3}^{\ell_4} \sum_{m=-\ell}^{\ell} n_{\ell m} Y_{\ell m}(\hat{n})$  (where, without loss of generality,  $\ell_1 \leq \ell_4$ ):

$$M(\hat{n})N(\hat{n}) = \sum_{L=0}^{\ell_2+\ell_4} \sum_{M=-L}^L p_{LM} Y_{LM}(\hat{n}) \quad (18)$$

for  $\ell_3 < \ell_2$ , *i. e.*, the limits on  $\ell$  in the two maps overlap (the  $p_{LM}$

<sup>6</sup> <http://www.s2let.org>

<sup>7</sup> <http://www.spinsht.org>

<sup>8</sup> <http://www.sothree.org>

<sup>9</sup>  $\text{FWHM}^j = 50 \sqrt{\frac{1200}{N_{\text{samp}}^j}}$ . This value is the same as used in the NILC implementation on *Planck* data.

are the new harmonic coefficients). (If the limits do not overlap, the lower limit on  $L$  in Eq. (18) becomes  $\ell_3 - \ell_2$ .) The limits on  $\ell$  in the wavelet coefficient maps  $W_{jnk}^c$  are  $(\ell_{\min}^j, \ell_{\max}^j)$  (see § 3.4). Therefore, by Eq. (18) and Eq. (15), the limits on  $\ell$  in the covariance matrix element maps  $(\hat{R}_{jnk}^c)^{cc'}$  are  $(0, 2\ell_{\max}^j)$ ; hence the limits on  $\ell$  in Eq. (16).

Having established the main equations governing the ILC method, we now present the main steps in the ILC algorithm that we use:

- (i) Form the  $(R_{jnk}^{\text{approx}})^{cc'}$  by Eq. (14).
- (ii) Smooth the  $(R_{jnk}^{\text{approx}})^{cc'}$  in harmonic space by Eq. (16).
- (iii) Take the inverse of each covariance matrix at each pixel to form  $(\hat{R}_{jnk}^{-1})^{cc'}$ .
- (iv) Calculate the ILC weights  $\omega_{jnk}^c$  by Eq. (13), where we assume that  $a^c = 1$  for all  $c$  and we substitute the empirical estimates for the inverse covariance matrices.
- (v) Finally, calculate the ILC estimate wavelet coefficient maps  $W_{jnk}^{\text{ILC}}$  by applying Eq. (11).

### 3.6 Point source masking

The input frequency maps are diffusively inpainted in a small point source mask following the method employed by Planck Collaboration et al. (2015d). This recognises that the ILC fails when the CMB is obscured by bright extragalactic sources or complex emission near the Galactic equator. The inpainting removes these sources and attempts to replace them with an extrapolation of the surrounding signal. The mask supplied is taken from the NILC section of Planck Collaboration et al. (2014a) and is constructed from the Planck Catalogue of Compact Sources (PCCS) (Planck Collaboration et al. 2014c, 2015e)<sup>10</sup>. It masks about 2.2% of the whole sky, predominantly along the Galactic equator towards the Galactic centre.

Because of this inpainting, the final ILC map is inpainted within the point source mask. For the purposes of this inpainting, we have split the mask into two, based on the size of its constituent individual contiguous holes<sup>11</sup>. For holes consisting of less than or equal to 800 pixels, we inpaint with a constrained Gaussian realisation following the method of Benoit-Lévy et al. (2013), itself an approximate implementation of the Hoffman-Ribak algorithm (Hoffman & Ribak 1991). For holes consisting of more than 800 pixels (the largest 131 out of 10031), we inpaint with a standard diffusive algorithm (in particular, following the method employed by Planck Collaboration et al. (2015d)). The result is that the ILC map is 1.3% Gaussian inpainted and 0.9% diffusively inpainted. This follows Benoit-Lévy et al. (2013), who do not recommend using their Gaussian inpainting for large holes near the Galactic equator.

### 3.7 Numerical implementation

SILC is implemented in Python and is parallelised. At full Planck resolution ( $N_{\text{side}} = 2048$ ,  $\ell_{\text{max}} = 3600$ ), when run on a 60-core symmetric multiprocessor (SMP) with 1.5 TB RAM and a 24-core

cluster node with 256 GB RAM<sup>12</sup>, the pipeline takes approximately 12 hours per direction. As shown by Eq. (16), we perform spherical harmonic transforms to  $2\ell_{\text{max}} = 7200$ . For a given number of directions  $N$ , the full pipeline takes approximately  $N$  times as long as the axisymmetric limit of our method (when  $N = 1$ ). In our infrastructure, the code was usually memory-limited (due to the very high resolution of the covariance matrix maps  $(R_{jnk}^c)^{cc'}$  at harmonic band-limit  $2\ell_{\text{max}}^j$ ); the amount of parallelisation sometimes had to be reduced to prevent memory overloads on a single node. As mentioned in § 3.4, the wavelet transforms employ the latest version of S2LET (Leistedt et al. 2013; McEwen et al. 2015), written in C with Python wrappers, itself employing SSHT (McEwen & Wiaux 2011) and S03 (McEwen et al. 2015). Despite the use of MW sampling and FFTs, spherical harmonic transforms are the most time-consuming part of the pipeline, again due to the very high resolution of the  $(R_{jnk}^c)^{cc'}$  (for the smallest wavelets, these covariance maps are band-limited at  $\ell = 7200$ ). There is scope to further optimise the implementation. Our wavelet analysis and synthesis functions do not respectively output and take as input wavelet coefficient maps at double-resolution (*i. e.*, a map band-limited at  $\ell_j^{\text{max}}$  sampled at  $2\ell_j^{\text{max}}$ ), requiring additional spherical harmonic transforms to double the resolution. Also, our spherical harmonic transform function does not calculate harmonic coefficients to a multipole less than the band-limit of the input map (*i. e.*, to only calculate  $a_{\ell m}$  for  $\ell < L$  where  $L < \ell_{\text{max}}^j$ ), resulting in excess computation at certain steps in the algorithm. These optimisations are left as further work.

## 4 SOURCES OF ERROR IN THE ILC

By the linearity of the wavelet transform in Eq. (7), the data model in Eqs. (1) and (2) can be recast in wavelet space as

$$W_{jnk}^c = a^c W_{jnk}^{\text{CMB}} + W_{jnk}^{\text{FG},c} + W_{jnk}^{\text{N},c}, \quad (19)$$

where  $W_{jnk}^{\text{CMB}}$ ,  $W_{jnk}^{\text{FG},c}$  and  $W_{jnk}^{\text{N},c}$  are respectively the CMB, foreground and instrumental noise contributions to each wavelet coefficient map. The beams within each component have been absorbed into the component wavelet coefficient maps. Substituting Eq. (19) into Eq. (11) gives

$$\begin{aligned} W_{jnk}^{\text{ILC}} &= \sum_{c=1}^{N_c} a^c \omega_{jnk}^c W_{jnk}^{\text{CMB}} + \sum_{c=1}^{N_c} \omega_{jnk}^c (W_{jnk}^{\text{FG},c} + W_{jnk}^{\text{N},c}) \\ &= W_{jnk}^{\text{CMB}} + \frac{\sum_{c,c'=1}^{N_c} (W_{jnk}^{\text{FG},c} + W_{jnk}^{\text{N},c}) (R_{jnk}^{-1})^{cc'} a^{c'}}{\sum_{c,c'=1}^{N_c} a^c (R_{jnk}^{-1})^{cc'} a^{c'}}, \end{aligned} \quad (20)$$


where the second equality follows by applying the constraint given in Eq. (12) and expanding the weights as given in Eq. (13). Even when the calibration  $a^c$  and the covariance matrices  $(R_{jnk}^c)^{cc'}$  are correct, there is always residual signal in the final ILC wavelet coefficient maps, given by the second term on the RHS of Eq. (20). Due to the linearity of the inverse wavelet transforms, this residual signal will propagate linearly into the final ILC temperature map as calculated by Eq. (10). As explained in § 3.5, this error term is reduced by minimising the empirical variance of the ILC map assuming that the CMB and foregrounds and the CMB and noise are respectively uncorrelated.

<sup>10</sup> The details of its construction are given in Planck Collaboration et al. (2014a). It can be downloaded from <http://pla.esac.esa.int/pla> and is labelled I\_MASK in the NILC data products.

<sup>11</sup> Note that many holes can be large and irregularly-shaped due to the overlapping of smaller circular holes.

<sup>12</sup> The exact specification for our infrastructure is an Intel Xeon E7-4890 2.8 GHz SMP with 4 × 15-core CPUs with 25.6 GB RAM per core, and an Intel Xeon E5-2697 2.7 GHz node with 2 × 12-core CPUs with 10.7 GB RAM per core.



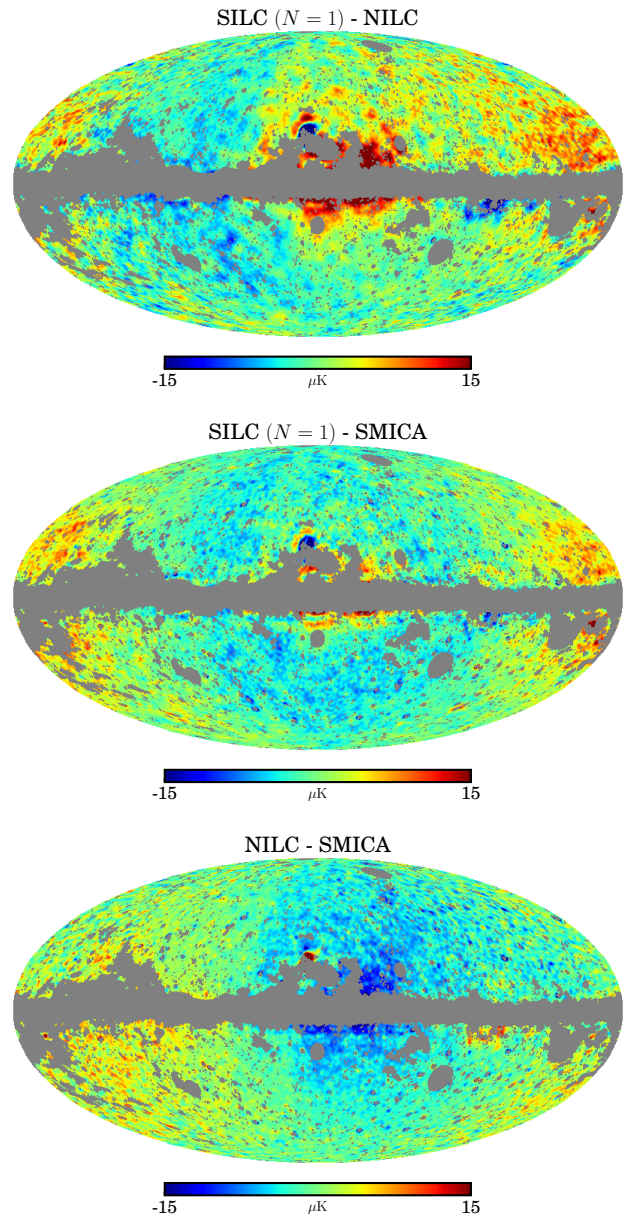


**Figure 4.** *Planck* data. The CMB temperature anisotropies reconstructed using SILC in the axisymmetric limit ( $N = 1$ , FWHM =  $5'$ ,  $N_{\text{side}} = 2048$ ). The grey pixels are the point source mask.

There are additional sources of error in the ILC method. The first is due to inaccuracy in the calculation of covariance matrices  $(R_{jnk})^{cc'}$ , *i. e.*, deviations in the empirical estimate  $(\hat{R}_{jnk})^{cc'}$  from the true covariance  $(R_{jnk})^{cc'}$ . Delabrouille et al. (2009) estimated the first order expansion of the reconstruction error in the ILC map estimate due to this covariance error. They showed that the covariance of the ILC error with the CMB is inversely proportional to the number of “effective modes” used in the ILC calculation. This covariance bias is negative. In our directional wavelet decomposition, our “effective modes” are spherical harmonic coefficients weighted to take account of the fact that the harmonic responses of wavelets overlap in both scale and direction. As  $N$ , the number of orientations probed, increases and so does the number of wavelets, each wavelet coefficient map contains fewer “effective modes.” We therefore conclude that our directional wavelet ILC may be susceptible to this negative ILC bias by increasing  $N$ . Delabrouille et al. (2009) also showed that due to chance correlations between the CMB and foregrounds, the variance minimisation leads to the unintentional cancellation of  $N^c - 1$  CMB modes. For *Planck*,  $N^c = 9$ , whereas for *WMAP*,  $N^c = 5$ . We therefore expect the magnitude of this negative bias to double simply by using more input frequency channels. Also, since this covariance bias is due to the cancellation of CMB modes, Delabrouille et al. (2009) showed that the absolute value is proportional to the CMB power. Therefore, the absolute value of the bias is greatest on large scales where CMB power is concentrated. In general, these biases are best estimated through suites of Monte Carlo simulations.

Another source of error is due to inaccuracy in the calibration  $a^c$  of the CMB. Dick et al. (2010) calculated the consequence of a first order error in  $a^c$  on a multiplicative correction to the CMB term in Eq. (20). They showed that even a small error in calibration can lead to a significant negative multiplicative bias in the CMB term, when the signal-to-noise ratio is large. (Here, the noise in this ratio also includes foreground signal.) They consider the implications for using an ILC on *Planck* data, where the signal-to-noise ratio is larger than for *WMAP* data. They estimate that a 1% error in  $a^c$  can cancel about a third of the CMB signal, while even a 0.1% error in  $a^c$  can remove about 1% of the CMB. Since our main map products use *Planck* data as input, they will be susceptible to this additional negative calibration bias. In this work, we assume that the CMB is calibrated to have unit response for all frequency channels, *i. e.*,  $a^c = 1$  for all  $c$ .

As mentioned in § 3.2 and 3.3, we assume all beams to be




**Figure 5.** *Planck* data. Differences between the axisymmetric limit ( $N = 1$ ) of SILC, NILC and SMICA. The maps have been smoothed to FWHM =  $80'$  and downgraded to  $N_{\text{side}} = 128$ . The grey pixels are the UT78 confidence mask from Planck Collaboration et al. (2015a), which masks the regions of the NILC and SMICA maps not recommended for cosmological analysis. The differences are (from top to bottom) (a) SILC ( $N = 1$ ) - NILC, (b) SILC ( $N = 1$ ) - SMICA and (c) NILC - SMICA.

circularly symmetric. Therefore, non-axisymmetric beam components will propagate into the ILC calculation but are assumed sufficiently small to be ignored.

## 5 COMPARISON TO PREVIOUS WORK


We now consider how SILC compares with existing component separation methods, particularly those adopted for the *Planck* analysis. We applied the axisymmetric limit (when  $N = 1$ ) of SILC to full-mission *Planck* data and compared the results to existing



**Figure 6.** *Planck* data.  $TT$  angular power spectra comparing the axisymmetric limit ( $N = 1$ ) of SILC to NILC and SMICA. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the best-fit  $\Lambda$ CDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ( $\ell < 1500$ ).

*Planck* analyses using the NILC and SMICA methods: the former because it is the closest in spirit to SILC, and the latter because it is the baseline method adopted by the *Planck* Collaboration for high-resolution analyses. Figure 4 shows the CMB reconstructed by the axisymmetric limit of SILC, while Fig. 5 shows the differences between this map and the NILC and SMICA (full-mission 2015 release) CMB maps and the difference between NILC and SMICA. The differences between the three maps are small in magnitude and mostly concentrated at the edges of the Galactic mask towards the Galactic centre, where foreground emission is most intense and complex. Quantitatively, we can compare the mean values and standard deviations of the masked difference maps. The mean values of Figs. 5 (a), (b) and (c) are respectively 0.44,  $-0.63$ , and  $-1.07 \mu\text{K}$ , while the standard deviations are respectively 4.24, 3.38 and  $3.43 \mu\text{K}^2$ . These values are small and similar, suggesting a strong consistency between the three methods. These difference maps have been formed from maps which have been smoothed and downgraded in resolution and so visually highlight differences at the lowest multipoles.

Figure 6 compares point source masked  $TT$  angular power spectra ( $D_\ell = \ell(\ell + 1)C_\ell/2\pi$ ) at the full multipole range of the




**Figure 7.** *Planck* data.  $TT$  angular power spectra comparing the axisymmetric limit ( $N = 1$ ) of SILC to NILC and SMICA. The top panel (a) shows full-sky spectra of inpainted maps. The middle panel (b) shows residuals after subtracting the best-fit  $\Lambda$ CDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ( $\ell < 1500$ ).

three maps (up to  $\ell = 3400$ )<sup>13</sup> with a CMB spectrum derived from the *Planck* 2015  $TT$  and low  $TEB$  likelihood<sup>14</sup>. The SILC spectrum is remarkably similar to that of NILC. This is unsurprising since the axisymmetric limit of SILC (when  $N = 1$ ) is very similar to the NILC method. Nonetheless, there are a number of pipeline differences. In particular, we use a different set of wavelets than the needlets employed in NILC (as discussed in § 3.4), even in the axisymmetric limit, with different harmonic responses. Figure 2 shows the harmonic response of the wavelets used in this work and Table 1 lists their harmonic band-limits  $\ell_{\min}^j$  and  $\ell_{\max}^j$ . The SMICA

<sup>13</sup> In order to estimate full-sky spectra from a masked map, we correct the  $C_\ell$  by dividing by  $f_{\text{sky}} = 0.978$ , a good approximation for a small mask. We elect to use point source masked spectra in order to concentrate our analysis on foreground and noise removal, rather than how maps are inpainted; all three maps are inpainted (at least) within the mask used.

<sup>14</sup> The parameters come from the `base.plikHM.TT.lowTEB` likelihood. The values are available in the *Planck* 2015 Release Explanatory Supplement: 2015 Cosmological parameters and MC chains ([http://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline\\_params\\_table\\_2015\\_limit68.pdf](http://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf)).



**Figure 8.** Planck simulations. Difference between output ILC and input CMB temperature maps from FFP8 simulations. The maps have been smoothed to  $\text{FWHM} = 80'$  and downgraded to  $N_{\text{side}} = 128$ . The grey pixels are the UTA76 confidence mask from [Planck Collaboration et al. \(2015a\)](#), which masks the Galactic region in FFP8 simulations where foreground emission is strongest.


spectrum has lower residuals at higher multipoles than both the axisymmetric limit of SILC and NILC.

Figure 7 compares full-sky angular power spectra of the three maps, including the inpainted point source regions. The spectra are similar to those in Fig. 6. The main difference is the lower noise tail in the SILC map at high multipoles above  $\ell = 1500$  (where all component separation CMB maps are dominated by residual instrumental noise). This is because, unlike NILC and SMICA, we do not Gaussian inpaint the very largest point source holes, but rather use diffusive inpainting (as discussed in § 3.6). The Gaussian inpainting of large irregular holes is poorly constrained and adds residual noise relative to diffusive inpainting.

We have shown that the axisymmetric limit of SILC gives comparable performance to NILC and SMICA. In § 6 and § 7, we “turn on” the directionality of the wavelets and consider the impact on CMB reconstruction from simulated and real data respectively.

## 6 APPLICATION TO PLANCK SIMULATIONS

We now apply SILC to the fiducial full-mission *Planck* FFP8 simulated sky maps, focusing on the impact on CMB reconstruction by increasing directionality as a function of scale. Figure 8 shows the difference between our reconstructed CMB (using  $N = 1$ ) and the input simulated CMB. There are small-magnitude differences particularly at the edge of the Galactic mask where the strength and complexity of foreground emission is greatest. As in Fig. 5, this difference map is at low resolution and so highlights residuals at the lowest multipoles. Figure 9 compares point source masked  $TT$  angular power spectra (up to  $\ell = 3400$ ) of CMB maps reconstructed using values of  $N$  from 1 to 5. It can be seen that the introduction of directionality has the greatest effect at multipoles around  $\ell = 800$ ; the residuals are beginning to converge for  $\ell \gtrsim 2000$ . Figure 10 shows the differences between simulated CMB maps reconstructed using  $N = [2, 3, 4, 5]$  minus the input CMB. The four maps and the axisymmetric difference map in Fig. 8 are almost identical with small magnitude residuals. This is because these low-resolution difference maps again highlight residuals on the very largest scales. However, as discussed in § 3.4, the wavelets we use are constructed to have an axisymmetric scaling function at the very lowest multi-



**Figure 9.** Planck simulations.  $TT$  angular power spectra comparing output ILC using different values of  $N$  and input CMB from FFP8 simulations. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the input CMB spectrum. The bottom panel (c) shows the same residuals at low multipoles only ( $\ell < 1500$ ).

poles. The scaling function we use (as detailed in Table 1) means that no directionality is applied for  $\ell < 32$ .

Figure 11 shows equivalent difference maps as in Fig. 10 but for the simulated CMB reconstructed using directional wavelets at all scales<sup>15</sup>, including for  $\ell < 32$ . It can be seen that the reconstruction errors are significantly larger in magnitude and cover almost the entire sky. The errors are also dominated by the largest scales, in particular a large error in the quadrupole increasing with magnitude as the amount of directionality  $N$  increases. We attribute this effect most probably to the ILC “biases” discussed in § 4, in particular the cancellation of CMB modes due to chance correlations with foregrounds in the ILC variance minimisation. [Delabrouille et al. \(2009\)](#) showed that the absolute value of this effect is largest on large scales where CMB power is concentrated, since the cancelled CMB modes on large scales have the greatest magnitude. Further, as discussed in § 4 and shown in Fig. 11, these errors are expected to increase in magnitude as a function of  $N$ . This is because as  $N$  increases, each directional wavelet coefficient map (the space in which our ILC operates) contains fewer “effective modes”

<sup>15</sup> In particular, the scaling function and  $j = 0$  wavelet are replaced by two directional wavelets with harmonic band-limits  $[1, 60]$  and  $[1, 128]$ .



of the input data and so the error in our empirical covariance estimation is expected to increase. This error propagates to the final maps.

These map reconstruction errors due to the implementation of directionality on the very largest scales are accompanied by increasingly negative power spectrum residuals as  $N$  increases, in particular in the first multipole bin from  $\ell = 2$  to  $\ell = 11$ . This is also indicative of the negative ILC bias due to empirical CMB cancellation, as discussed in § 4 and Delabrouille et al. (2009). The results in Fig. 11 thus motivate the use of an axisymmetric scaling function, which ensures that no directionality is used for  $\ell < 32$  and so reduces the errors in CMB reconstruction. In principle these biases can be estimated and corrected through large suites of simulations, which is beyond the scope of this work.

## 7 APPLICATION TO PLANCK DATA

We now study the application of SILC with increasing directionality to the full-mission *Planck* sky maps. The left column of Fig. 12 shows the full-resolution reconstructed CMB maps as calculated with different values of  $N$  from 2 to 5, which visually appear very similar. The right column of Fig. 12 shows the differences between the CMB reconstructed using  $N = [2, 3, 4, 5]$  minus the axisymmetric limit (when  $N = 1$ ), highlighting the differences at the lowest multipoles. The differences are of largest magnitude towards the Galactic plane where foreground emission is concentrated. This shows how the different wavelet kernels are localising the ILC weights differently in response to the directional structure of the foregrounds and CMB. The differences are small, reflecting the implementation of an axisymmetric scaling function, meaning that no directionality is applied at  $\ell < 32$ . Figure 13 compares point source masked  $TT$  angular power spectra of the CMB reconstructed using values of  $N$  from 1 to 5. The power spectrum residuals from a *Planck* best-fit  $\Lambda$ CDM model remain small for most scales until the reconstructed spectra reach a characteristic noise spectrum for  $\ell \gtrsim 1500$  where the different values of  $N$  converge. At these high multipoles, the ILC solution is dominated by residual instrumental noise. We see the biggest impact from directionality at intermediate multipoles (from  $\ell = 400$  to  $\ell = 1500$ ). For comparison, we plot the SMICA power spectrum. In further support to the discussion in § 5, SILC matches the performance of SMICA. We note that, as with the simulations in § 6, directionality changes the reconstructed CMB power spectrum most significantly at intermediate multipoles around  $\ell = 800$ .

## 8 DISCUSSION

The comparisons in § 5 demonstrate that SILC matches the performance of two previous methods, NILC and SMICA, in both maps and power spectra, with particular similarity between the axisymmetric limit of SILC and NILC, as expected. Both map residuals and power spectra in § 6 and § 7 show that switching on directionality changes CMB reconstruction most significantly at intermediate multipoles  $\ell = 400 - 1500$ . There appears to be little benefit in localising the ILC with directional wavelets at the very smallest scales, where the ILC result is noise-limited. We also adopt an axisymmetric scaling function on the very largest scales, meaning that there is no directionality at  $\ell < 32$ . In Fig. 11, we show the large CMB reconstruction errors arising from using directionality on the largest scales. This motivates the use of an axisymmetric scaling

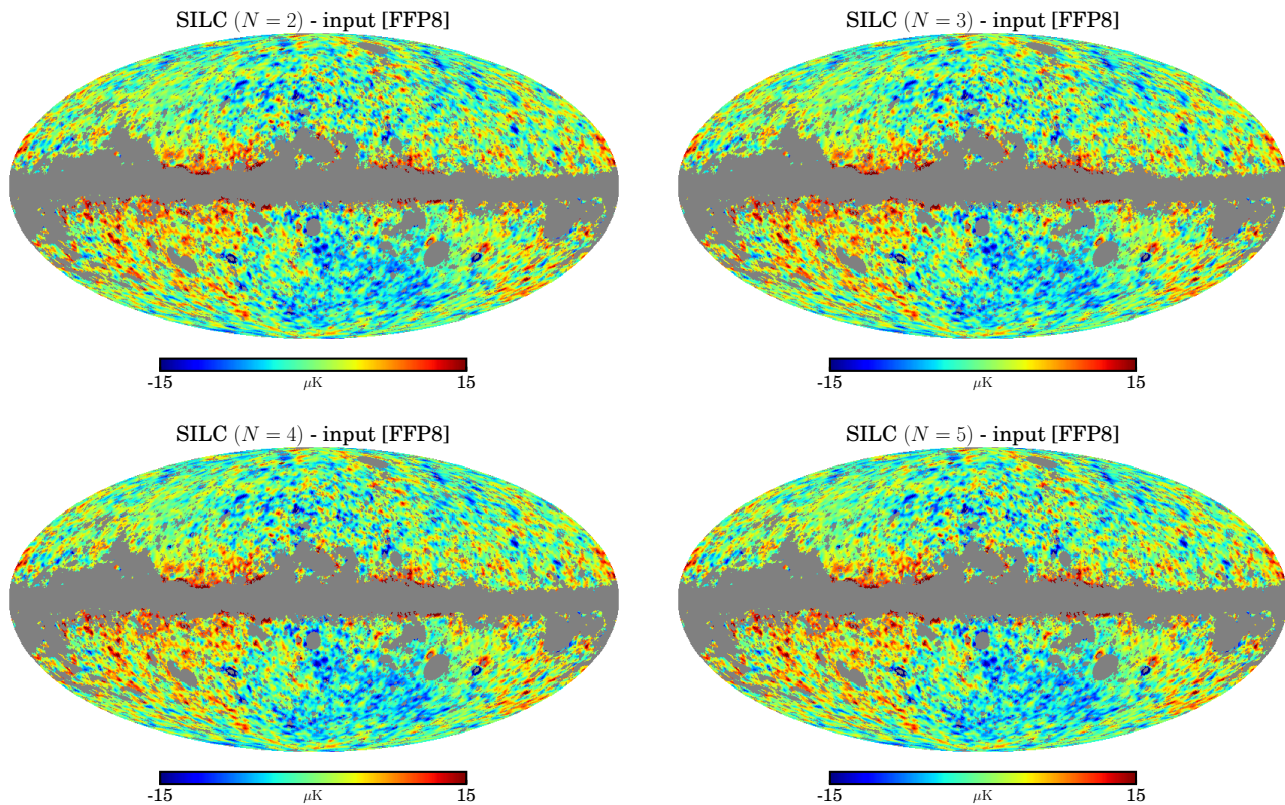
function, which significantly reduces the errors as seen in Fig. 10. In § 6, we sketched out an argument that attributes these errors to empirical CMB cancellation (§ 4). However, the precise source and exact magnitude of any ILC errors are best estimated through suites of simulations.

We have presented this analysis by producing CMB maps (in § 6 and § 7) each with a different single value of  $N$  at all wavelet scales. Our method can be simply extended to allow different values of  $N$  at each wavelet scale. In the same way that each wavelet scale has different harmonic band-limits, they can also have different azimuthal band-limits, optimised as identified above to reduce foreground and noise residuals.

The negative ILC power spectrum biases discussed in § 4 must be quantified in parallel to this directionality optimisation if using the resulting map for power spectrum analyses. It is possible to estimate variance biases in the data through suites of realistic simulations. However, we can also calculate this using the data itself and a fiducial CMB spectrum. In wavelet space, the variance estimator at each wavelet coefficient is  $\langle W^{\text{ILC}} W^{\text{ILC}\dagger} \rangle = \omega^\dagger \langle W W^\dagger \rangle \omega = \omega^\dagger \mathbf{R} \omega = (\mathbf{a}^\dagger \mathbf{R}^{-1} \mathbf{a})^{-1} = (\sum_{c,c'}^{N_c} (R^{-1})^{cc'})^{-1}$ , where each equality follows by applying in turn Eqs. (11), (14) and (13) (from § 3.5) and then expanding the vector notation (the vectors span the space defined by the number of input channels; explicitly, we assume unit CMB calibration  $\mathbf{a} = (1, 1, 1, 1, 1, 1, 1, 1)$ ). In order to calculate the variance bias in the ILC, we can subtract the expected CMB variance  $\sum_{\ell m} C_\ell |\Psi_{\ell m}^j|^2$ . If analysis was done in wavelet space, the above would define the variance bias. If only considering the diagonal terms in wavelet space, it is possible to straightforwardly transform this estimate to real space through an inverse wavelet transform as in Eq. (10), substituting  $\Phi(\hat{\mathbf{n}})$  and  $\Psi^j(\hat{\boldsymbol{\rho}})$  respectively for  $|\Phi(\hat{\mathbf{n}})|^2$  and  $|\Psi^j(\hat{\boldsymbol{\rho}})|^2$ . However, for a full treatment of the variance bias including off-diagonal terms, full wavelet space covariances need to be calculated. Although many off-diagonal terms would decay, this would still be computationally demanding and will be the focus of future research. However, we reiterate that if analysis is carried out in wavelet space, then variance biases can be straightforwardly calculated from the information already contained in the results.

## 9 CONCLUSIONS

We have presented SILC, a new form of internal linear combination that uses directional, scale-discretised wavelets to localise the ILC weighting according to the frequency, spatial, harmonic and, for the first time, morphological information in the CMB and its foregrounds. This is motivated by the anisotropic or filamentary morphology of both the CMB and astrophysical foregrounds in the microwave sky. We have tested SILC on 2015-release *Planck* data and simulations, demonstrating comparable performance to two existing component separation algorithms, NILC and SMICA, and investigated how to optimise the use of morphological information through directionality. We have explored increasing the amount of directionality in the algorithm, showing that on the largest and the smallest scales, the axisymmetric limit of the ILC works well, while at intermediate multipoles (from  $\ell = 400 - 1500$ ), increasing  $N$  (the number of directions per scale) leads to lower residuals. At high multipoles ( $\ell \gtrsim 1500$ ), the input data are already noise-limited, as is the ILC reconstruction, and directionality does not reduce the reconstruction error, as instrumental noise has no directional structure. We adopt an axisymmetric scaling function to analyse the largest scales, meaning that we use no directionality for  $\ell < 32$ .



**Figure 10.** Planck simulations. Differences between output ILC reconstructed using different values of  $N$  and input CMB temperature maps from FFP8 simulations. The maps have been smoothed to  $\text{FWHM} = 80'$  and downgraded to  $N_{\text{side}} = 128$ . The grey pixels are the UTA76 confidence mask. The differences are (from left to right, top to bottom) (a)  $N = 2$ , (b)  $N = 3$ , (c)  $N = 4$ , (d)  $N = 5$  minus the input CMB.

This is motivated by the observation that increasing directionality on large scales gives increased reconstruction errors over the axisymmetric limit. We argue that these errors are due to empirical CMB cancellation in the ILC calculation, though the exact source must be estimated through large suites of realistic simulations. Allowing  $N$  to vary with wavelet scale is analogous to the choice of different harmonic band-limits at different scales.

We conclude that the introduction of directional wavelets allows greater flexibility in the ILC to make use of morphological information at targeted scales. Our multiprocessing implementation takes advantage of the wavelet scales to allow large-scale results to be analysed while small scales are still being processed. Moreover, our wavelet transforms are quick and exact, using MW sampling and fast Fourier transforms (see § 2). As discussed in § 4, the ILC is prone to several sources of error and variance bias, including empirical CMB cancellation. This bias can be estimated through suites of Monte Carlo simulations, but we have also outlined (in § 8) the ability to estimate biases directly from the data, most straightforwardly in wavelet space. We make our map products available at <http://www.silc-cmb.org><sup>16</sup>.

This work on scalar signals (*i. e.*, the temperature  $I$  component of the CMB) can be extended to spin signals (*i. e.*, the polarisation  $Q$  and  $U$  components of the CMB, or, equivalently, the  $E$  and  $B$  modes), by using spin wavelets (McEwen et al. 2015b; McEwen et al. 2014; Leistedt et al. 2015). These are an extension of directional, scale-discretised wavelets to represent spin signals, such as

CMB polarisation, a spin  $\pm 2$  signal. We expect that the directionality will be particularly suited to the anisotropic, filamentary nature of polarised foregrounds when observed on the sky, and in future work will present the application of SILC to CMB polarisation data.

#### ACKNOWLEDGEMENTS

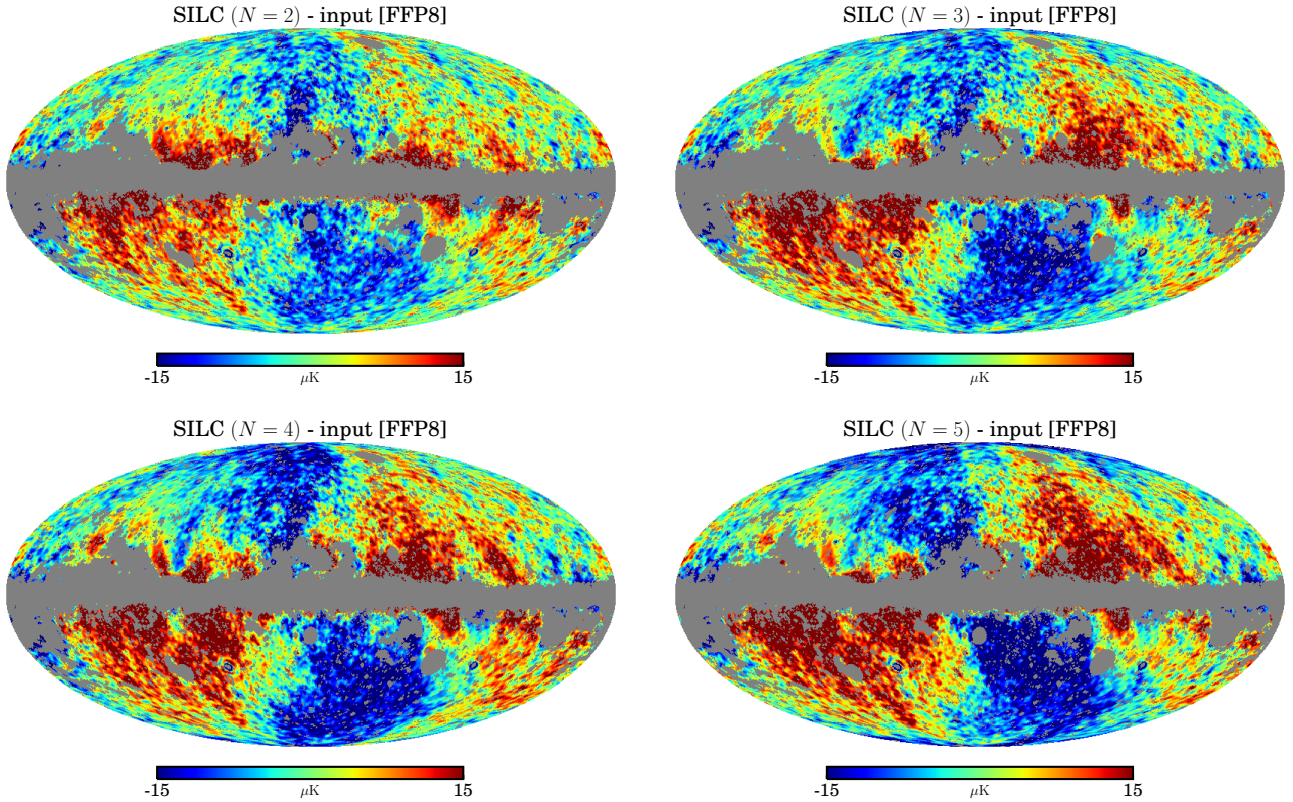
KKR thanks Franz Elsner for valuable discussions. KKR was supported by the Science and Technology Facilities Council. HVP and BL were partially supported by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 306478-CosmicDawn. JDM was partially supported by the Engineering and Physical Sciences Research Council (grant number EP/M011852/1). AP was supported by the Royal Society. Based on observations obtained with *Planck* (<http://www.esa.int/planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA and Canada.

#### REFERENCES

- Antoine J.P., Vanderghaynst P., 1998, *J. Math. Phys.*, 39, 8, 3987  
 Antoine J.P., Vanderghaynst P., 1999, *ACHA*, 7, 1  
 Baldi P., Kerkycharian G., Marinucci D., Picard D., 2009, *Ann. Stat.*, 37 No.3, 1150, [arXiv:math/0606599](https://arxiv.org/abs/math/0606599)

<sup>16</sup> The DOI for our data release is 10.5281/zenodo.44373.



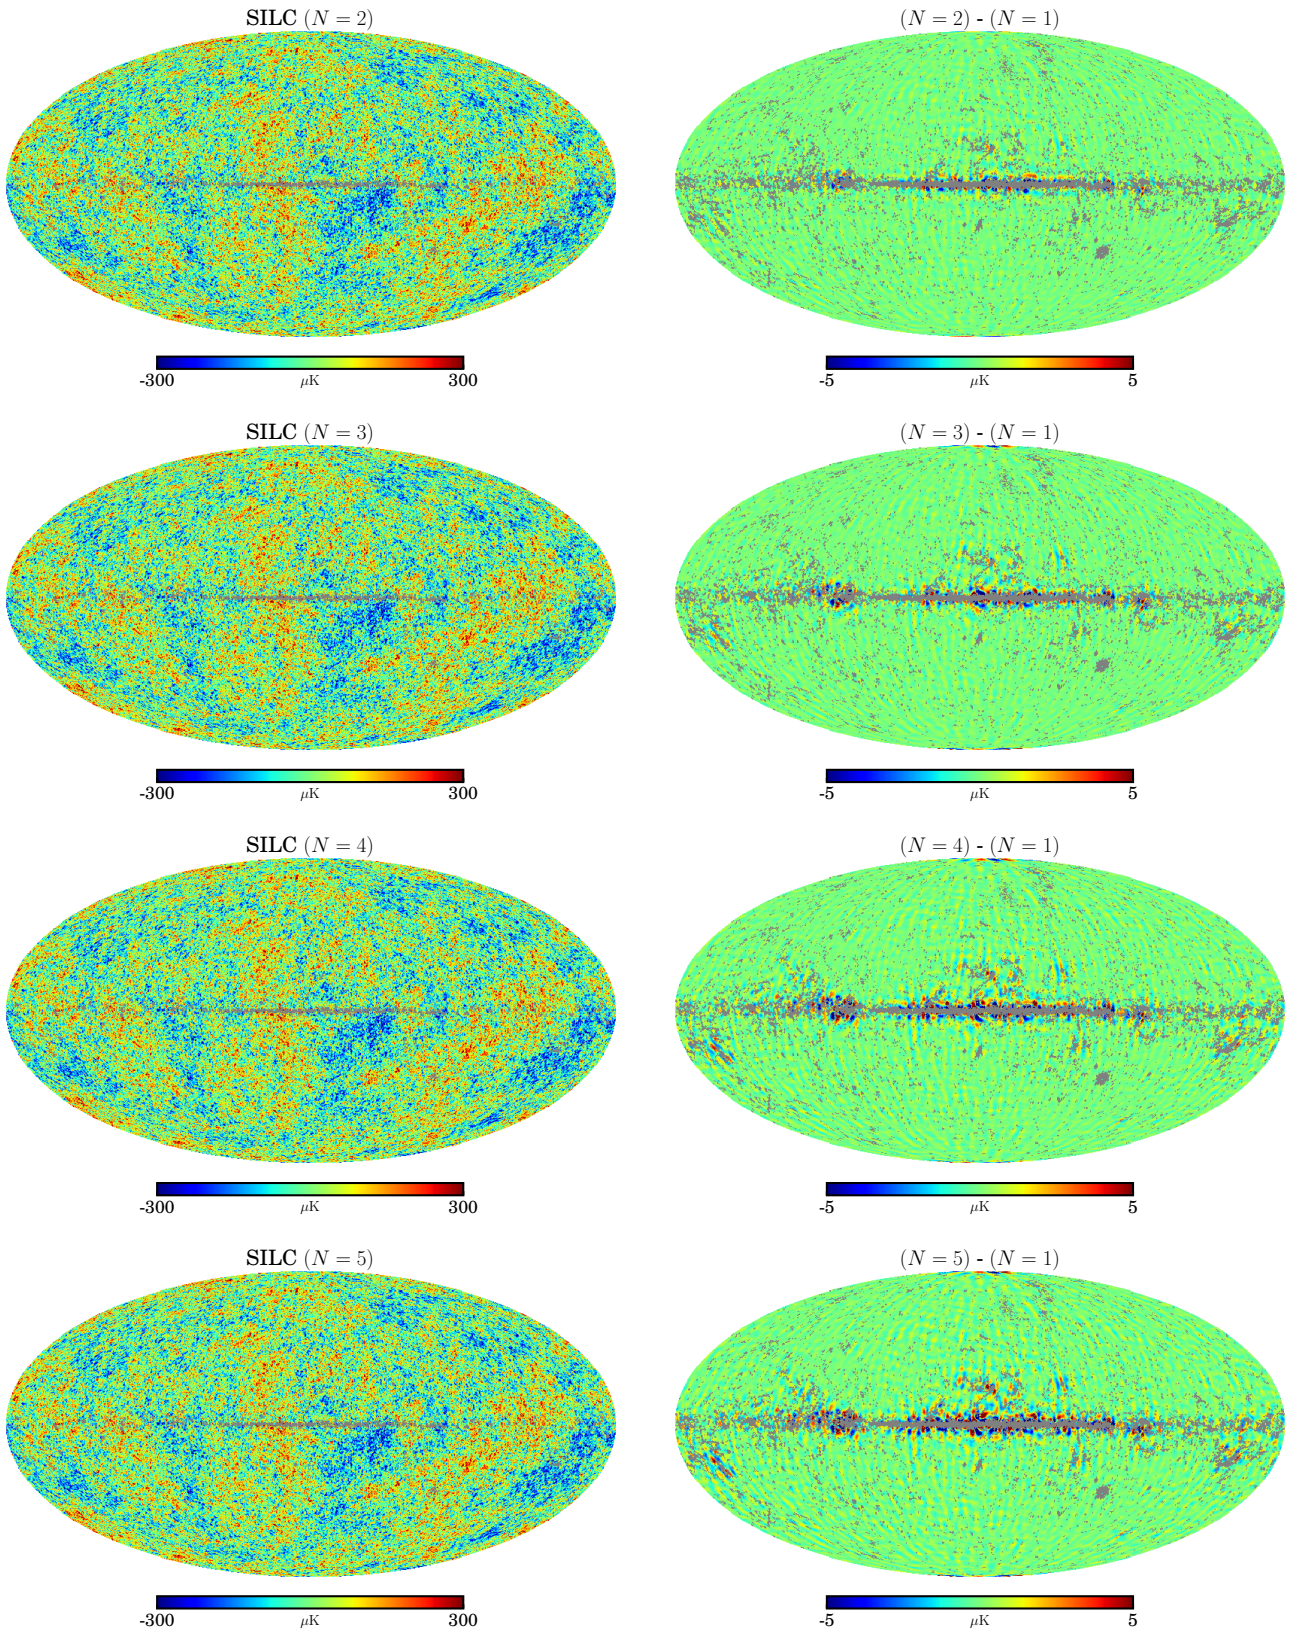


**Figure 11.** Planck *simulations*. Same as Fig. 10 (which uses the recommended wavelets) but here using directional wavelets on large scales ( $\ell < 32$ ), which is not recommended as it leads to increased CMB reconstruction errors as seen above.


Barreiro R.B., Hobson M.P., Lasenby A.N., Banday A.J., Górski K.M., Hinshaw G., 2000, MNRAS, 318, 475, [astro-ph/0004202](#)  
 Basak S., Delabrouille J., 2012, MNRAS, 419, 1163, [1106.5383](#)  
 Bedini L., Herranz D., Salerno E., Baccigalupi C., Kuruoglu E.E., Tonazzini A., 2005, EURASIP Journal on Applied Signal Processing, 2005, 2400, [astro-ph/0407108](#)  
 Bennett C.L., et al., 2003a, ApJS, 148, 97, [astro-ph/0302208](#)  
 Bennett C.L., et al., 2003b, ApJS, 148, 1, [astro-ph/0302207](#)  
 Benoit-Lévy A., Déchelette T., Benabed K., Cardoso J.F., Hanson D., Prunet S., 2013, A&A, 555, A37, [1301.4145](#)  
 Bobin J., Moudden Y., Starck J.L., Fadili J., Aghanim N., 2008, Statistical Methodology, 5, 307, [0712.0588](#)  
 Bobin J., Starck J.L., Sureau F., Basak S., 2013, A&A, 550, A73, [1206.1773](#)  
 Boggess N.W., et al., 1992, ApJ, 397, 420  
 Bond J.R., Efstathiou G., 1987, MNRAS, 226, 655  
 Cardoso J.F., Martin M., Delabrouille J., Betoule M., Patanchon G., 2008, ArXiv e-prints, [0803.1814](#)  
 Chan J.Y.H., Leistedt B., Kitching T.D., McEwen J.D., 2015, IEEE TSP, submitted, [arXiv:1511.05578](#)  
 Delabrouille J., Cardoso J.F., Le Jeune M., Betoule M., Fay G., Guilloux F., 2009, A&A, 493, 835, [0807.0773](#)  
 Dick J., Remazeilles M., Delabrouille J., 2010, MNRAS, 401, 1602, [0907.3105](#)  
 Driscoll J., Healy D., 1994, Advances in Applied Mathematics, 15, 2, 202, ISSN 0196-8858  
 Eriksen H.K., Banday A.J., Górski K.M., Lilje P.B., 2004, ApJ, 612, 633, [astro-ph/0403098](#)

Eriksen H.K., Jewell J.B., Dickinson C., Banday A.J., Górski K.M., Lawrence C.R., 2008, ApJ, 676, 10, [0709.1058](#)  
 Eriksen H.K., et al., 2006, ApJ, 641, 665, [astro-ph/0508268](#)  
 Fernández-Cobos R., Vielva P., Barreiro R.B., Martínez-González E., 2012, MNRAS, 420, 2162, [1106.2016](#)  
 Geller D., Hansen F.K., Marinucci D., Kerkycharian G., Picard D., 2008, PRD, 78, 12, 123533, [arXiv:0811.2881](#)  
 Geller D., Marinucci D., 2010, JFAA, 16, 6, 840, [0811.2935](#)  
 Geller D., Marinucci D., 2011, J. Math. Anal. Appl., 375, 2, 610, [arXiv:1006.3835](#)  
 Górski K.M., Hivon E., Banday A.J., Wandelt B.D., Hansen F.K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759, [astro-ph/0409513](#)  
 Hobson M.P., Jones A.W., Lasenby A.N., Bouchet F.R., 1998, MNRAS, 300, 1, [astro-ph/9806387](#)  
 Hoffman Y., Ribak E., 1991, ApJ, 380, L5  
 Leach S.M., et al., 2008, A&A, 491, 597, [0805.0269](#)  
 Leistedt B., McEwen J.D., Büttner M., Peiris H.V., Vanderghaynst P., Wiaux Y., 2015, ArXiv e-prints, [1502.03120](#)  
 Leistedt B., McEwen J.D., Vanderghaynst P., Wiaux Y., 2013, A&A, 558, A128, [1211.1680](#)  
 Marinucci D., et al., 2008, MNRAS, 383, 539, [arXiv:0707.0844](#)  
 Martínez-González E., Diego J.M., Vielva P., Silk J., 2003, MNRAS, 345, 1101, [astro-ph/0302094](#)  
 Mather J.C., et al., 1990, ApJ, 354, L37  
 McEwen J.D., Büttner M., Leistedt B., Peiris H.V., Vanderghaynst P., Wiaux Y., 2014, in *IAU Symposium*, volume 306 of *IAU Symposium*, 64–67, [1412.1340](#)





**Figure 12.** Planck data. *Left:* CMB temperature anisotropies reconstructed using SILC with different values of  $N$  ( $\text{FWHM} = 5'$ ,  $N_{\text{side}} = 2048$ ). *Right:* differences between CMB temperature maps reconstructed using different values of  $N$  minus the axisymmetric limit  $N = 1$ . The maps have been smoothed to  $\text{FWHM} = 80'$  and downgraded to  $N_{\text{side}} = 128$ . In both columns: the grey pixels are the point source mask (downgraded in resolution as appropriate). From top to bottom: (a)  $N = 2$ , (b)  $N = 3$ , (c)  $N = 4$ , (d)  $N = 5$ .



**Figure 13.** Planck data.  $TT$  angular power spectra comparing different values of  $N$  from 1 to 5 and SMICA. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the best-fit  $\Lambda$ CDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ( $\ell < 1500$ ).

- Planck Collaboration, et al., 2014b, *A&A*, 571, A12, [1303.5072](#)  
 Planck Collaboration, et al., 2014c, *A&A*, 571, A28, [1303.5088](#)  
 Planck Collaboration, et al., 2015a, ArXiv e-prints, [1502.05956](#)  
 Planck Collaboration, et al., 2015b, ArXiv e-prints, [1509.06348](#)  
 Planck Collaboration, et al., 2015c, ArXiv e-prints, [1506.07135](#)  
 Planck Collaboration, et al., 2015d, ArXiv e-prints, [1502.01592](#)  
 Planck Collaboration, et al., 2015e, ArXiv e-prints, [1507.02058](#)  
 Sanz J.L., Herranz D., López-Cañiego M., Argüeso F., 2006, in *EUSIPCO*, [astro-ph/0609351](#)  
 Starck J., Moudden Y., Bobin J., 2009, *A&A*, 497, 931, [arXiv:0902.0574](#)  
 Starck J.L., Moudden Y., Abrial P., Nguyen M., 2006, *A&A*, 446, 1191, [astro-ph/0509883](#)  
 Tegmark M., de Oliveira-Costa A., Hamilton A.J., 2003, *Phys. Rev. D*, 68, 12, 123523, [astro-ph/0302496](#)  
 Wiaux Y., Jacques L., Vanderghenst P., 2005, *ApJ*, 632, 15, [astro-ph/0502486](#)  
 Wiaux Y., McEwen J.D., Vanderghenst P., Blanc O., 2008, *MNRAS*, 388, 770, [0712.3519](#)

- McEwen J.D., Buttner M., Leistedt B., Peiris H.V., Wiaux Y., 2015, *IEEE Signal Processing Letters*, 22, 2425, [1508.03101](#)  
 McEwen J.D., Durastanti C., Wiaux Y., 2015a, ACHA, submitted, [arXiv:1509.06767](#)  
 McEwen J.D., Hobson M.P., Lasenby A.N., 2006, ArXiv, [astro-ph/0609159](#)  
 McEwen J.D., Leistedt B., Büttner M., Peiris H.V., Wiaux Y., 2015b, *IEEE TSP*, submitted, [arXiv:1509.06749](#)  
 McEwen J.D., Scaife A.M.M., 2008, *MNRAS*, 389, 3, 1163, [arXiv:0803.2165](#)  
 McEwen J.D., Vanderghenst P., Wiaux Y.t.c., 2013, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8858 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 0, [1308.5706](#)  
 McEwen J.D., Wiaux Y., 2011, *IEEE Transactions on Signal Processing*, 59, 5876, [1110.6298](#)  
 McEwen J.D., Wiaux Y., Evers D.M., 2011, *A&A*, 531, A98, [arXiv:1108.3900](#)  
 Narcowich F., Petrushev P., Ward J., 2006, *SIAM J. Math. Anal.*, 38, 574, [astro-ph/0403098](#)  
 Planck Collaboration, et al., 2011, *A&A*, 536, A1, [1101.2022](#)  
 Planck Collaboration, et al., 2014a, *A&A*, 571, A1, [1303.5062](#)