

# BEYONDPLANCK VIII. Efficient Sidelobe convolution and corrections through spin harmonics

M. Galloway<sup>11\*</sup>, M. Reinecke<sup>13</sup>, K. J. Andersen<sup>11</sup>, R. Aurlen<sup>11</sup>, R. Banerji<sup>11</sup>, M. Bersanelli<sup>4,9,10</sup>, S. Bertocco<sup>8</sup>, M. Brilenkov<sup>11</sup>, M. Carbone<sup>14</sup>, L. P. L. Colombo<sup>4</sup>, H. K. Eriksen<sup>11</sup>, J. R. Eskilt<sup>11</sup>, M. K. Foss<sup>11</sup>, C. Franceschet<sup>4,10</sup>, U. Fuskeland<sup>11</sup>, S. Galeotta<sup>8</sup>, M. Galloway<sup>11</sup>, S. Gerakakis<sup>14</sup>, E. Gjerløw<sup>11</sup>, B. Hensley<sup>2</sup>, D. Herman<sup>11</sup>, M. Iacobellis<sup>14</sup>, M. Ieronymaki<sup>14</sup>, H. T. Ihle<sup>11</sup>, J. B. Jewell<sup>12</sup>, A. Karakci<sup>11</sup>, E. Keihänen<sup>3,7</sup>, R. Keskitalo<sup>1</sup>, G. Maggio<sup>8</sup>, D. Maino<sup>4,9,10</sup>, M. Maris<sup>8</sup>, A. Mennella<sup>4,9,10</sup>, S. Paradiso<sup>4,9</sup>, B. Partridge<sup>6</sup>, M. Reinecke<sup>13</sup>, M. San<sup>11</sup>, A.-S. Suur-Uski<sup>3,7</sup>, T. L. Svalheim<sup>11</sup>, D. Tavagnacco<sup>8,5</sup>, H. Thommesen<sup>11</sup>, D. J. Watts<sup>11</sup>, I. K. Wehus<sup>11</sup>, and A. Zacchei<sup>8</sup>

<sup>1</sup> Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, U.S.A.

<sup>2</sup> Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, U.S.A.

<sup>3</sup> Department of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland

<sup>4</sup> Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria, 16, Milano, Italy

<sup>5</sup> Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, Trieste, Italy

<sup>6</sup> Haverford College Astronomy Department, 370 Lancaster Avenue, Haverford, Pennsylvania, U.S.A.

<sup>7</sup> Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland

<sup>8</sup> INAF - Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, Trieste, Italy

<sup>9</sup> INAF-IASF Milano, Via E. Bassini 15, Milano, Italy

<sup>10</sup> INFN, Sezione di Milano, Via Celoria 16, Milano, Italy

<sup>11</sup> Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway

<sup>12</sup> Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California, U.S.A.

<sup>13</sup> Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

<sup>14</sup> Planetek Hellas, Leoforos Kifisias 44, Marousi 151 25, Greece

March 22, 2021

## ABSTRACT

We present the treatment of sidelobe power within the BEYONDPLANCK end-to-end Gibbs sampling framework. Sidelobe power is estimated in the time domain using a new implementation of the `libconvqt` convolution algorithm. This procedure is derived and then applied to the *Planck* LFI instrument to estimate the sidelobe signal at each Gibbs sample within the BEYONDPLANCK sampling loop. We use the same sidelobe response functions as the 2018 LFI DPC analysis, which are stored as  $a_{lm}s$  with  $m_{\max} = 100$  and  $l_{\max} = 512$ . The estimated time domain sidelobe signals are subtracted out of **MR: out of  $\Rightarrow$  from?** the sky maps shown in the other BEYONDPLANCK papers, but are presented here for the purposes of completeness and verification. We compare our implementation to the previous Levels implementation and show good agreement between the two codes. We then show mean sidelobe estimates, which match well with previous efforts on this topic, as well as sidelobe RMS maps that characterize the sensitivity of the sidelobe estimates to the sky model provided. We also demonstrate the importance of accurate sidelobe measurements, and discuss problems of calibration and gain estimation that arise when sidelobes are subtracted or modeled incorrectly.

**Key words.** ISM: general – Cosmology: observations, polarization, cosmic microwave background, instrument characterization – Galaxy: general

## 1. Introduction

One of the important systematic effects that must be accounted for in CMB instruments is the telescope stray light, or sidelobes. This is the non-zero response of the detector to areas of the sky outside the main beam, however that is defined. This response can be caused by many different physical effects, such as spurious optical reflections or manufacturing irregularities in the detectors or optical elements. These signal contributions can have far reaching consequences on the observed signal, as they do not behave in the same manner as the main beam signal.

Sidelobe signals can produce many types of errors in CMB analysis pipelines, and they represent a potent source of systematic contamination. As the sidelobe signal originates from locations other than the primary telescope pointing, it can contaminate

inate signals like the CMB and solar dipoles that are used for calibration. In some experiments, the spurious signal can originate from a source not on the sky, such as ground pick up or RF noise. In all cases, sidelobe signal is detrimental to the quality of the final sky maps and parameter estimates, and requires a dedicated removal effort.

Characterizing and correcting these spurious signals is therefore an important part of optimal CMB mapmaking, and requires optimized algorithms to characterize them efficiently. One of the most important of these is the convolution operation, convolving a beam or sidelobe response function with a sky map or model to generate a re-observed map. Convolution of a sky measurement with a beam response function is vital for determining how signal will appear to a telescope, and has been the subject of much active research in the last 20 years.

Full sky convolution on the sphere is a problem that has been important in the CMB field since the earliest satellite measure-

\* Corresponding author: M. Galloway; [mathew.galloway@astro.uio.no](mailto:mathew.galloway@astro.uio.no)

ments. Early experiments like *COBE* and *WMAP* used simple pixel-based convolution approaches, which even for their low resolution required spherical beam approximations (Wu et al. 2001) or limited the applications to large scales (Burigana et al. 2001).

Wandelt & Górski (2001) presented the first harmonic space convolution algorithm, which achieved a large performance gain ( $O(\sqrt{N_{\text{pix}}})$ ) over the pixel-based methods ( $O(N_{\text{pix}})$ ). This initial breakthrough allowed the calculation of these convolutions easily enough that they could be applied to each simulation, instead of requiring a dedicated study requiring months of runtime.

Several years later, Prézeau & Reinecke (2010) developed the Conviqt approach, which was used in the official *Planck* analysis pipelines (Planck Collaboration III 2016, Planck Collaboration Int. XLVI 2016, Planck Collaboration Int. LVII 2020) and in the *Planck* FFP simulations (Planck Collaboration XII 2016). This approach was an improvement over the state of the art, speeding up the computation of the Wigner recursion relationships used in their harmonic space algorithm, as well as providing a standardized, user friendly library, `libconviqt`, that was incorporated into numerous pipelines.

BEYONDPLANCK makes use of a new convolution formalism to correct for the spurious signal from sidelobe pickup, one which is based on Spherical Harmonic Transforms (SHTs). This approach uses the same fundamental algebra as the older Conviqt approach, but instead of directly computing the Wigner Matrix elements, we express the convolution as a Spherical Harmonic Transform and are thus able to leverage the highly optimized `libsharp` SHT library to perform the bulk of the calculations (Reinecke & Seljebotn 2013). Although this new approach was not developed specifically for BEYONDPLANCK, this paper is, to our knowledge, the first to explicitly derive and discuss them.

## 2. Sidelobes, `libconviqt` and `libsharp`

### 2.1. Total Convolution through Spin Harmonics

Given a sky map,  $s(\hat{n})$ , and beam,  $b(\hat{n})$ , our task is to compute a quantity  $c(\vartheta, \varphi, \psi) \in \mathbb{R}$  that represents the convolution of these two fields, with the beam oriented in polar coordinates  $(\vartheta, \varphi)$ , and rotated around its own central axis by  $\psi$ ,

$$c(\vartheta, \varphi, \psi) \equiv \int_{4\pi} s(\hat{n})b(\hat{n}'(\vartheta, \varphi, \psi) - \hat{n}) d\Omega_{\hat{n}}. \quad (1)$$

Here,  $\vartheta$  and  $\varphi$  are the co-latitude and longitude of a location on the sphere, i.e. they have the same meaning as in the HEALPix context.  $s_{l,m_s}$  is the  $a_{l,m}$  representation of the sky signal, and  $b_{l,m_b}$  is the beam in the same representation. Care has been taken to distinguish between  $m_s$  ( $m_{\text{sky}}$ ) and  $m_b$  ( $m_{\text{beam}}$ ), as the two indices will be treated separately in the following derivation.

The original derivation for this problem has been given by Prézeau & Reinecke (2010). Their algorithm involves computing Wigner  $d$  matrix elements via recurrences that are no longer used in the new implementation, so we provide another derivation here, which results in expressions which match the currently employed algorithm more closely.

As shown by Wandelt & Górski (2001), Eq. (1) can be evaluated efficiently in harmonic space as

$$c(\vartheta, \varphi, \psi) = \sum_{l,m_s,m_b} s_{l,m_s} b_{l,m_b}^* [D_{m_s,m_b}^l(\varphi, \vartheta, \psi)]^*, \quad (2)$$

where  $s_{l,m_s}$  and  $b_{l,m_b}$  are the spherical harmonics coefficients of the signal and beam, respectively, and  $D_{m_s,m_b}^l$  is the Wigner

$D$ -matrix, for which efficient recursion relations are available (Prézeau & Reinecke 2010).

In the following, we will show that Eq. (1) can alternatively be expressed in terms of spin-harmonics. The resulting algebra is in principle identical to the recursion relations used by Prézeau & Reinecke (2010), but the implementation is simply repackaged in a format that is significantly easier to implement in practical computer code, since it may use existing and highly optimized spherical harmonics libraries, such as Reinecke & Seljebotn (2013), to perform the computationally expensive parts.

As shown by Goldberg et al. (1967), it is possible to express the Wigner  $D$  matrix as

$$D_{-m_s}^l(\varphi, \vartheta, -\psi) = (-1)^m \sqrt{\frac{4\pi}{2l+1}} {}_s Y_{lm}(\vartheta, \varphi) e^{is\psi}, \quad (3)$$

where  ${}_s Y_{lm}(\vartheta, \varphi)$  is called a spin-weighted spherical harmonic. Inserting this expression into Eq. (2) yields

$$c(\vartheta, \varphi, \psi) = \sum_{l,m_s,m_b} \sqrt{\frac{4\pi}{2l+1}} s_{l,m_s} b_{l,-m_b} {}_{-m_b} Y_{lm_s}(\vartheta, \varphi) e^{im_b\psi}, \quad (4)$$

where we have assumed that the beam is real-valued in position space (implying  $b_{l,-m_b} = (-1)^{m_b} b_{l,m_b}^*$ ), and used the well-known symmetry relations

$$D_{-m_s,-m_b}^l(\vartheta, \varphi) = (-1)^{m_s+m_b} [D_{m_s,m_b}^l(\vartheta, \varphi)]^* \quad (5)$$

$$b_{l,m_b}^* (-1)^{m_b} = b_{l,-m_b}. \quad (6)$$

Separating the sum yields

$$c(\vartheta, \varphi, \psi) = \sum_{m_b} e^{im_b\psi} \sum_{l,m_s} \sqrt{\frac{4\pi}{2l+1}} s_{l,m_s} b_{l,-m_b} {}_{-m_b} Y_{lm_s}(\vartheta, \varphi). \quad (7)$$

**Bruce: The text between Eqn. 7 and Eqn. 9 could provide more explanation and justification. Here is a place where a little more "teaching" would help the average reader.**

In this expression, the term for  $m_b = 0$  can be identified as a spin-0 spherical harmonic transform of the quantity  $\sqrt{4\pi/(2l+1)} s_{l,m_s} b_{l,0}$ , which can be easily computed by a library like `libsharp` (Reinecke & Seljebotn 2013).

For  $m_b \neq 0$ , we consider the contribution of each pair of terms with the same  $|m_b|$ . Since  $c(\vartheta, \varphi, \psi) \in \mathbb{R}$ , these must be conjugate complex with respect to each other, and their combined contribution is therefore

$$e^{im_b\psi} {}_{m_b} S_{l,m_s} + e^{-im_b\psi} {}_{m_b} S_{l,m_s}^* = 2(\cos(m_b\psi) \text{Re}({}_{m_b} S_{l,m_s}) + \sin(m_b\psi) \text{Im}({}_{m_b} S_{l,m_s})), \quad (8)$$

where we have defined

$${}_{m_b} S_{l,m_s} \equiv \sum_{l,m_s} \sqrt{\frac{4\pi}{2l+1}} s_{l,m_s} b_{l,-m_b} {}_{-m_b} Y_{lm_s}(\vartheta, \varphi). \quad (9)$$

This is a spherical harmonic transform of a quantity with spin  $m_b$ , which can also be computed efficiently by `libsharp`. **DJW: I spent a bit of time trying to understand what the purpose was of this paragraph was, but as I understand it now, the idea is once we can get stuff into a format that `libsharp` does well, i.e., a spin- $s$  SHT, we have "solved the problem." It might be instructive to just say early on in this section, "a spin- $s$  SHT is bla", and then later on you can say "This is a spherical harmonic transform of a quantity with spin  $m_b$ ,**

i.e., Equation (8), which can also be computed efficiently by `libsharp`”.

In practice, the transforms in Eq. (9) are implemented by separating  $S$  into its gradient and curl (or  $E$  and  $B$ ) coefficients,  $a_{lm}$  (Lewis 2005),

$$m_b S_{l,m_s} = -(m_b E_{l,m_s} + i m_b B_{l,m_s}), \quad (10)$$

using the symmetry relations  $m_b E_{l,-m_s} = (-1)^{m_s} m_b E_{l,m_s}^*$  and  $m_b B_{l,-m_s} = (-1)^{m_s} m_b B_{l,m_s}^*$ , and the overall minus sign is a convention. Again making use of the symmetry relation in Eq. 6, this results in

$$m_b E_{l,m_s} = -s_{lm_s} \text{Re}(b_{l,m_b}) \quad (11)$$

$$m_b B_{l,m_s} = -s_{lm_s} \text{Im}(b_{l,m_b}). \quad (12)$$

To summarize, efficient evaluation of the convolution integral in Eq. (1) may be done through the following steps:

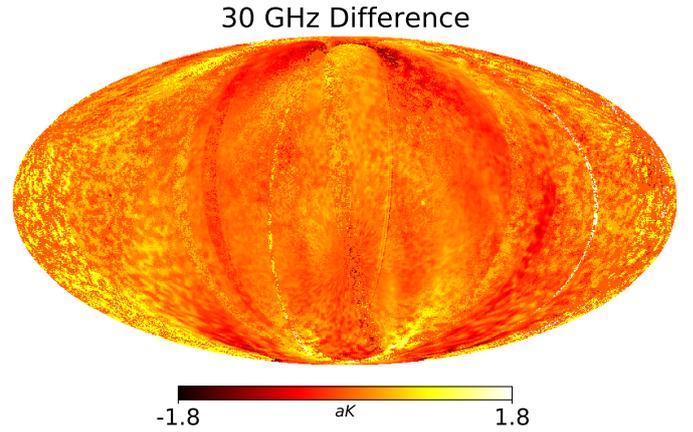
1. For each  $m = 0 \dots m_b$ , pre-compute the spin spherical harmonic coefficients in Eqs. 11–12, and compute the corresponding spin- $m_b$  spherical harmonics transform with an external library such as `libsharp`; this results in a three-dimensional data cube of the form  $c(\vartheta, \varphi, m_b)$ .
2. For each position on the sky,  $(\vartheta, \varphi)$ , perform a Fourier transform to convert these coefficients to  $c(\vartheta, \varphi, \psi)$ , as given by Eq. 7. **DJW: I have to admit, the  $m_b \rightarrow \psi$  stuff was not clear to me until I read Dag’s thesis. It might be helpful to just say explicitly somewhere that the whole point of doing this is getting into the ring format so we can do the practically instantaneous FFT**

In practice, the resulting  $c(\vartheta, \varphi, \psi)$  data object is evaluated at a finite pixel resolution typically set to match the beam bandlimit. To obtain smooth estimates within this data object, a wide range of interpolation schemes may be employed, trading off computational efficiency against accuracy. However, this issue is identical to previous approaches (Wandelt & Górski 2001; Prézeau & Reinecke 2010), and we refer the interested reader to those papers for further details.

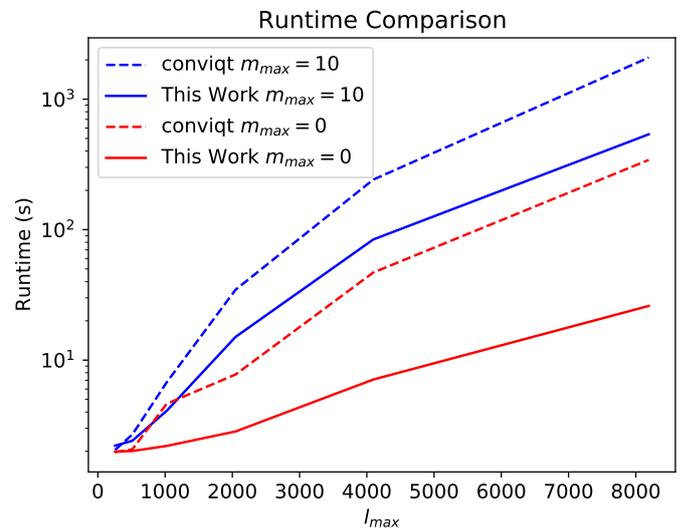
## 2.2. Comparison with `libconvirt`

To compare the results of this new total convolution approach and the older `libconvirt` approach of Prézeau & Reinecke (2010), we performed some simple tests. We used the two algorithms to compute the convolution between the LFI 28M beam and a Commander 30 GHz sky model. The resulting convolution cubes were then observed using LFI’s scanning strategy for the first 365 days of the *Planck* flight, and the map differences are shown in Fig. 1. The convolution cubes were also directly compared for accuracy, and they had an integrated difference at the  $10^{-8}$  level, which indicates only machine-level differences.

The primary advantage of this new approach, in addition to being mathematically nicer, is of course speed. Figure 2 shows runtime comparisons between the two approaches for a test convolution of an elliptical beam with the same set of random sky  $a_{l,m}$ s. For both cases shown,  $m_{\max} = 0$  and  $m_{\max} = 10$ , this work outperforms the old approach at all but the lowest  $l_{\max}$ , where the data read time is the dominant effect. Additionally, for compatibility with the old `libconvirt` approach, this test was performed with an older version of `libsharp`, so we expect that the new algorithm scales more favourably than this with the latest implementation.



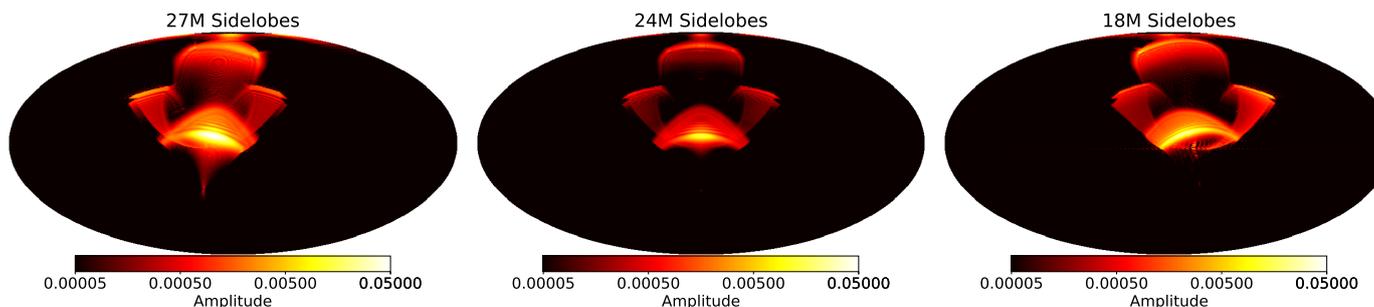
**Fig. 1.** Map level difference of the new SHT convolution algorithm compared to the old `Convirt` approach, observed using the identical pointing of the first year of the *Planck* mission. The differences are at the level of machine precision, showing good agreement between the two algorithms. **Diverging colormap?**



**Fig. 2.** Runtime comparison between the `libconvirt` approach and the new spin-SHT approach for the convolution of an elliptical Gaussian with a set of random sky  $a_{l,m}$ s. This work ties or outperforms the previous approach for all values of  $l_{\max}$  from 256 to 8192 for both  $m_{\max}$  values shown.

## 3. Sidelobe Models

Figure 3 shows characteristic sidelobe response functions evaluated at a fixed frequency on the sky for a detector in each band. The sidelobe response for each detector looks visually quite similar, so only these representative ones are shown here. Each is stored on disk as a set of  $a_{l,m}$ s with  $l_{\max} = 512$  and  $m_{\max} = 100$ .  $a_{l,m}$  coefficients with indexes above these maximums are omitted, which is equivalent to setting them to 0, which is justified as they contain only noise. Heuristically,  $l_{\max}$  defines a scale factor, below which we do not encode sidelobe information, and  $m_{\max}$  is somewhat equivalent a maximum angular resolution of the sidelobe representation. Figure 3 makes it clear that the sidelobe signal is an important contribution to the detector signal at each frequency, and that exact removal of the sidelobe signal is required to avoid biasing other output products.



**Fig. 3.** Maps of the sidelobe response on the sky from the first detector at (left to right) 30 GHz, 44 GHz and 70 GHz. The beam orientation is such that the main beam is pointed directly at the north pole in these maps. The intensities are normalized such that the main beams have unit power at  $l = 0$ .

### 3.1. Main Beam Treatment

In the BEYONDPLANCK analysis, the sidelobe and main beam components of the sky response are separated, and the sidelobes are treated as a nuisance signal similar to the orbital dipole and monopole correction terms, as can be seen in the global parametric model of the data given by Eq. (13).

$$d_{j,t} = g_{j,t} \left[ P_{lp,j} B_{pp',j} \sum_c M_{c,j}(\beta_{p'}, \Delta b p^j) a_{p'}^c + s_{j,t}^{\text{orb}} + s_{j,t}^{\text{fsl}} + s_j^{\text{mono}} \right] + n_{j,t}^{\text{corr}} + n_{j,t}^{\text{w}} \quad (13)$$

The other terms in this equation are discussed in detail in [BeyondPlanck Collaboration \(2021\)](#), but here the main beam effect is denoted as  $B_{pp',j}$  and the sidelobe signal is extracted from the signal contribution and expressed as  $s_{j,t}^{\text{fsl}}$ . This distinction allows the sidelobes to be treated separately from the main beam in all respects. Treating the main beam using the Conviqt formalism of this paper would be possible, but the additional precision required to model it accurately would require much higher  $l_{\text{max}}$ , and therefore greatly increased computational time and memory requirements.

In the BEYONDPLANCK analysis, the main beam is used (in conjunction with the sidelobes) to compute the full  $4\pi$  dipole response, as detailed in Section 3.3. Additionally, a Gaussian main beam approximation is used during component separation to smooth the sky model to the appropriate beam resolution for each channel. During mapmaking, the pixel size for each channel is selected so that a single pixel can comfortably contain the main beam, and the beam is assumed to be pointed at the center of each pixel. In practice, we find this to be a sufficient approximation to get good convergence.

### 3.2. Sidelobe Normalization

The normalization of the sidelobes differs slightly from the normalization used within the Planck LFI collaboration. The official LFI beam products leave a small percentage of the power within the system unassigned due to uncertainties about where it should be. BEYONDPLANCK has chosen to renormalize the beam transfer function such that this power is distributed proportionally at each  $l$  (that is, rescale the beam transfer function  $B_l$  such that its full sky integral  $B_0 = 1$ ). This re-scaling is equivalent to assigning the unknown beam power uniformly over the entire beam, which is likely incorrect, but as we believe this is still preferred to the alternative of not assigning it. This removes the requirement that

the resultant maps be re-normalized by a transfer function after they are generated, and allows the LFI data to be compared to other experiments in a more natural way.

### 3.3. Orbital Dipole and Quadrupole Sidelobe Response

The treatment of the sidelobes is also important while generating orbital dipole and quadrupole estimates. Because Planck is calibrated primarily from the dipole measurements, as discussed in [Svalheim et al. \(2021\)](#), the sidelobe's contribution to the dipole can directly result in an absolute calibration error if not handled appropriately. While the CMB dipole can easily be handled using the the Conviqt approach described in Sec. 2.1, the orbital dipole is not sky-stationary and thus must be handled separately.

BEYONDPLANCK generates orbital dipole and quadrupole estimates directly from the Planck pointing information, using the satellite velocity data which has been stored at low resolution (one measurement per pointing period). With this information, it is possible to estimate the orbital dipole and quadrupole amplitude for each timestep, allowing the time-domain removal of the signal before it contaminates the final products with non-sky-stationary signal artifacts. Additionally, once this signal has been isolated from the raw data, it can be used as an aid in the calibration routines because of its highly predictable structure.

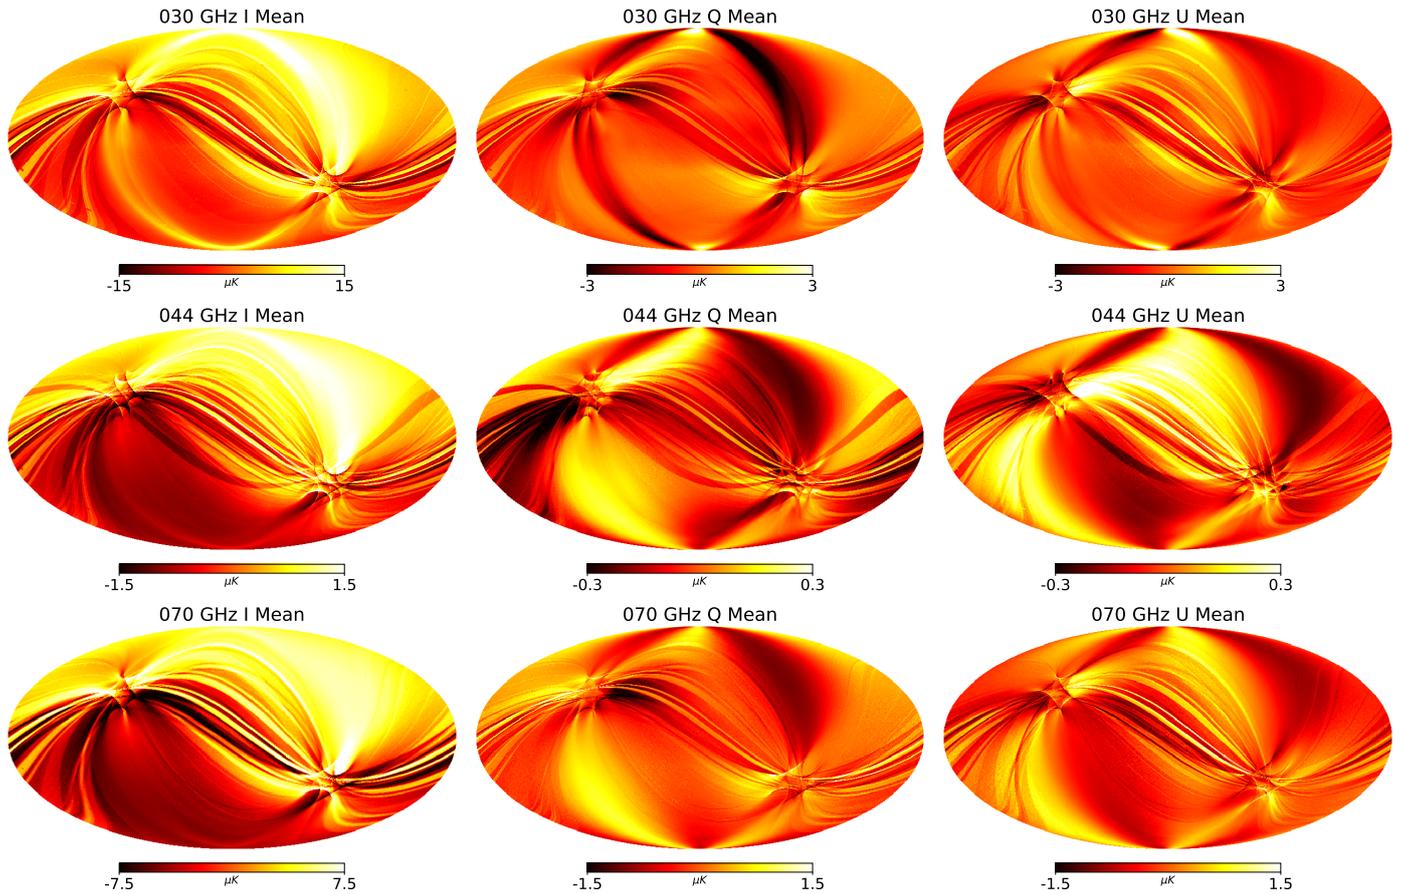
To generate the orbital dipole and quadrupole estimate, BEYONDPLANCK borrows a technique from the Planck NPIPE pipeline ([Planck Collaboration Int. LVII 2020](#), see appendix C). Following this reference, we express the signal  $\tilde{D}$  seen by a detector observing a fixed direction  $\hat{n}_0$  as the convolution of the dipole and quadrupole signal on the sky  $D(\hat{n})$  with the full  $4\pi$  beam response  $B(\hat{n}, \hat{n}_0)$ .

To generate the orbital dipole and quadrupole estimate, BEYONDPLANCK borrows a technique from the *Planck* NPIPE pipeline ([Planck Collaboration Int. LVII 2020](#), see appendix C). Following this reference, we express the signal  $\tilde{D}$  seen by a detector observing a direction  $\hat{n}_0$  as the convolution of the dipole and quadrupole signal on the sky  $D(\hat{n})$  with the full  $4\pi$  beam response  $B(\hat{n}, \hat{n}_0)$ ,

$$\tilde{D}(\hat{n}_0) = \int d\Omega B(\hat{n}, \hat{n}_0) D(\hat{n}). \quad (14)$$

Here, it is useful to break the dipole signal up into three orthogonal components in the standard Cartesian coordinates  $(x, y, z)$ . The orbital dipole and quadrupole in the CMB can be expressed as a Doppler shift in each direction,

$$D(\hat{n}) = T_0 [\beta \cdot \hat{n} (1 + q\beta \cdot \hat{n})] \quad (15)$$



**Fig. 4.** Sidelobe maps at each of the three LFI frequencies. From top to bottom: 30 GHz, 44 GHz and 70 GHz. The left column is the unpolarized sky signal, the central column is the  $Q$  polarization and the right column is  $U$ . Note the difference in the colour scales required to see the same level of detail in all three channels. **DJW:** would be nice to have these using the planck colormap, or some other diverging colormap

where  $\beta$  is the satellite velocity divided by the speed of light  $\beta = \frac{v}{c}$ ,  $T_0$  is the CMB temperature and  $q$  is quadrupole factor dependent on the frequency  $\nu$ , defined by

$$q = \frac{x e^x + 1}{2 e^x - 1}, \quad \text{where } x = \frac{h\nu}{k_B T_0}. \quad (16)$$

Because each of the three coordinates are orthogonal, Eq. (15) can be expanded in Cartesian coordinates as:

$$\begin{aligned} \tilde{D} = T_0 \int d\Omega B(\hat{n}, \hat{n}_0) & \left[ x \beta_x + y \beta_y + z \beta_z + \right. \\ & q \left( x^2 \beta_x^2 + y^2 \beta_y^2 + z^2 \beta_z^2 + \right. \\ & \left. \left. 2xy \beta_x \beta_y + 2xz \beta_x \beta_z + 2yz \beta_y \beta_z \right) \right]. \end{aligned} \quad (17)$$

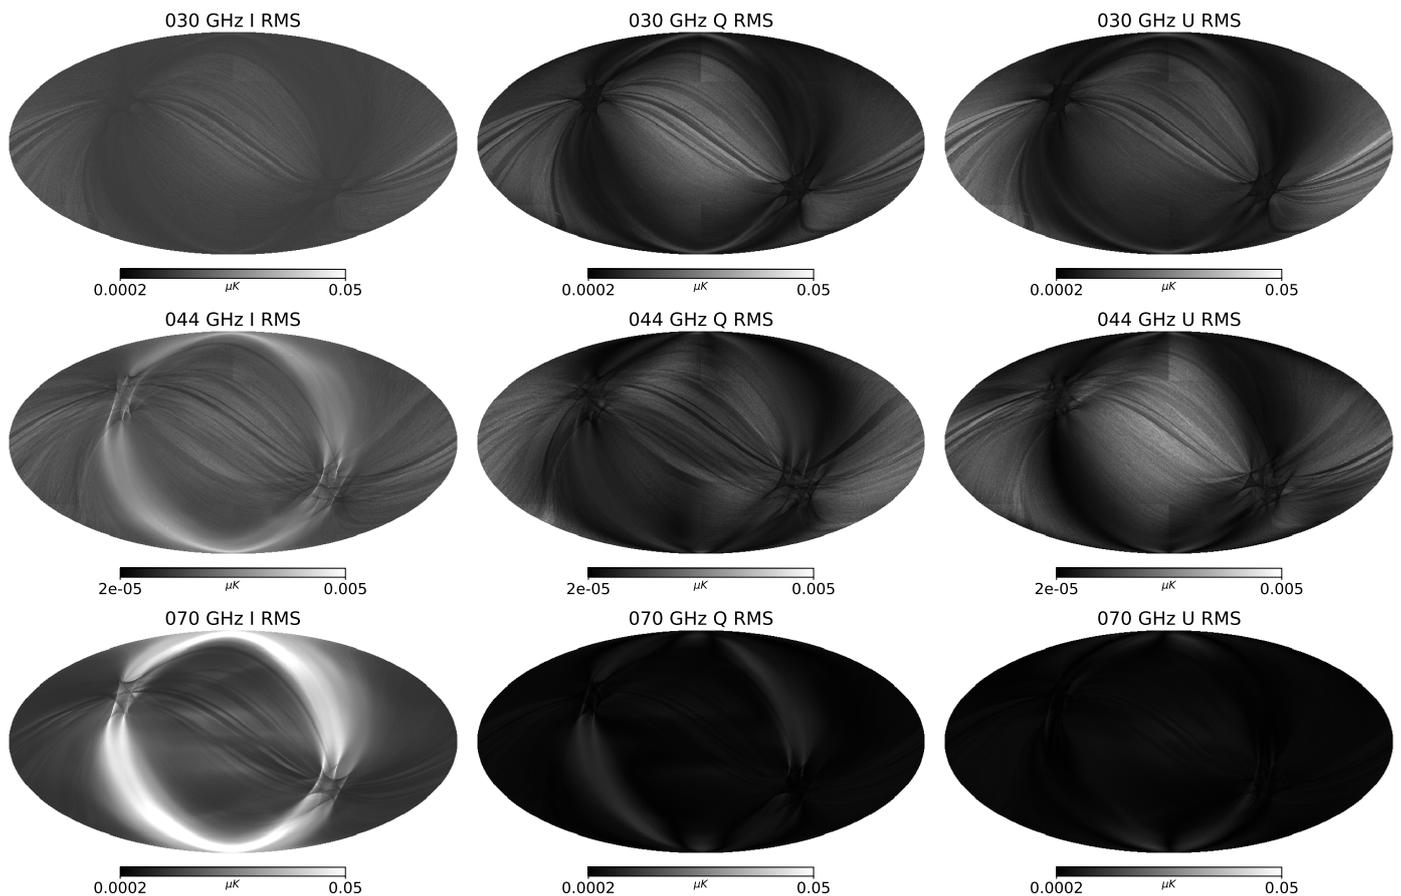
Here,  $\hat{n}$  is a unit direction vector which is the integration variable, and  $\hat{n}_0$  is the fixed direction of the satellite pointing for this timestep. Noting that this integral is independent of the  $\beta$  terms, we can precompute this full sky beam/dipole integral for a fixed beam orientation (we selected the default one of the beam pointing to the "north" pole of the HEALPix map). This gives 9

pre-computed  $S$  factors that can then be applied to the velocity as:

$$\begin{aligned} \tilde{D} = T_0 & \left[ S_x \beta_x + S_y \beta_y + S_z \beta_z + \right. \\ & q \left( S_{xx} \beta_x^2 + S_{yy} \beta_y^2 + S_{zz} \beta_z^2 + \right. \\ & \left. \left. 2S_{xy} \beta_x \beta_y + 2S_{xz} \beta_x \beta_z + 2S_{yz} \beta_y \beta_z \right) \right]. \end{aligned} \quad (18)$$

Once we have precomputed this integral for a fixed orientation, all we must do to compute  $\tilde{D}$  is rotate the pointing vector to the same orientation in which we pre-computed these  $S$  factors, and apply this same rotation to the satellite velocity vector. This means that instead of a full sky integral per timestep, we can instead simply compute a single rotation and then a summation operation, greatly increasing the speeds at which the dipole and quadrupole estimates can be generated in the time domain.

BEYONDPLANCK further accelerates this operation by computing this rotation for only one point in twenty, and using a spline to interpolate between them. This saves the costly operation of calculating a new rotation matrix at each step, and instead relies on the smoothness of the signal to ensure continuity. The algorithm treats the final few points of each pointing period that do not divide evenly into the subsampling factor separately. This allows the use of regular bin widths, which greatly speeds up the splining routines, while the final few points are calculated using the slower rotation matrix technique.



**Fig. 5.** Sidelobe rms maps at each of the three LFI frequencies. From top to bottom: 30 GHz, 44 GHz and 70 GHz. The left column is the unpolarized sky signal, the central column is the  $Q$  polarization and the right column is  $U$ .

## 4. Sidelobe Estimates

### 4.1. Maps

Figure 4 shows the mean sidelobe estimates at each of the three LFI frequencies. These maps show the sidelobe signal for the entire flight co-added across each frequency and then projected into sky coordinates, identically to the way the true sky signal is treated. They differ from the traditional *Planck* method of producing a template map and fitting it out because these signals are produced in the time domain during TOD processing, and thus are always appropriate for the selection of data used by the pipeline. These templates are exactly correct for the maps produced by this pipeline runs, but will not match precisely with analyses that use different data cuts, flagging or channel selection.

In each case, the 50 sample initial burn in period is discarded before computing the means. The sidelobe estimates are not output every sample because they are only used for diagnostic purposes and have little cosmological interest. Instead, they were written every 10 samples, a rate which is adjustable in the parameter file. Consequently, the maps of Fig. 4 are generated using only 90 total samples of the possible 900.

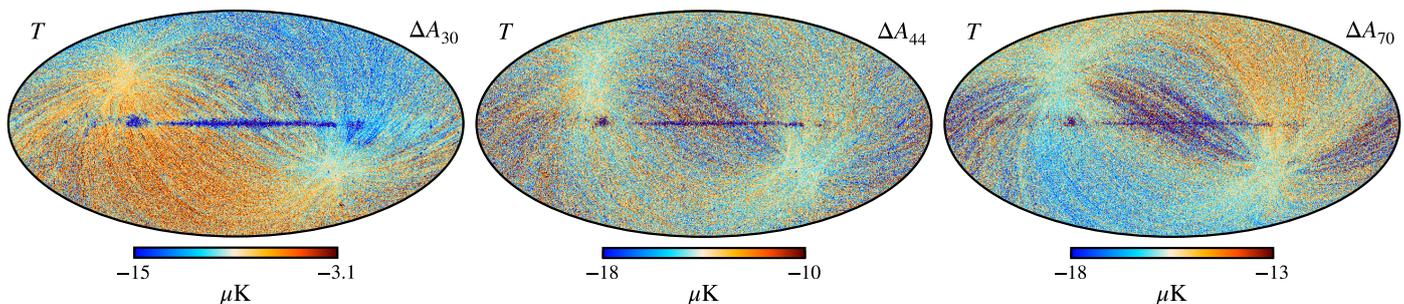
These results look similar to those previously presented in Fig. 7 of [Planck Collaboration et al. \(2016\)](#) (the far sidelobe maps were not shown in the 2018 release as they were very similar to 2015). The main difference is that the results here also consider the sidelobe response to the CMB dipole as part of the spurious signal at this stage, whereas the LFI 2015 DPC analysis showed the sidelobe pickup of dipole-subtracted maps. Both

pipelines see the most significant sidelobe signal at 30 GHz, and the morphology of the ring structures is markedly similar. These rings are caused by the pickup of emission from the galactic plane in the sidelobes as the satellite scans. The galactic foregrounds are the brightest part of the sky at these frequencies and so it is understandable that they would have the greatest effect on the sidelobe pickups.

### 4.2. Stability

In addition to the sidelobe mean maps, the BEYONDPLANCK pipeline outputs can also be used to look at sidelobe stability and statistical variation. Figure 5 shows the rms maps generated from the same sample of sidelobe estimates as was used in Fig. 4. Clear evidence of the scanning pattern can be seen, which is expected. The sharp vertical lines visible in polarization (clearest in 30 GHz Q and U at the top, and 44 GHz U at the top and bottom) have been previously examined by the Planck team, and are caused by a chance alignment between the non-dense Planck scanning strategy and the shape of the HEALPix pixels. For an example of this effect, see Fig. 15 of [Planck Collaboration et al. \(2014\)](#).

These RMS maps cannot be considered true sidelobe error estimates, however, as they do not characterize the response of the system to different models of the physical sidelobe pickup. Instead, they simply show the change in the estimated sidelobe signal as the sky estimates change through component separation. They can be interpreted as regions in the sky where the sidelobe signals are particularly sensitive to changes in sky sig-



**Fig. 6.** Frequency map difference plots at (left to right) 30, 44 and 70 GHz, comparing the mean Gibbs maps from the BEYONDPLANCK release with a maximum-likelihood execution that has no sidelobe correction.

nal, however, so they are still somewhat informative about how small errors in sidelobe modeling might appear in the sky maps.

To have a true propagation of sidelobe error estimates would require sampling over the physical parameters that determine the response on the sky. While sampling the full set of GRASP model parameters is likely to be infeasible due to size and computational time, finding a minimal parameter set that allows for some reasonable variation in the sidelobe response functions would allow the BEYONDPLANCK pipeline to sample over this range of parameters and give physically motivated errors to the sidelobe models. This approach will be developed for future applications such as the LiteBIRD mission (Hazumi et al. 2019).

## 5. Results

To demonstrate the importance of sidelobe corrections on a variety of other products, we performed a simple maximum likelihood run of the Commander code, starting from the same input data as the main pipeline run. In the absence of changes, this run should approximate the mean maps from the sampling pipeline, although obviously without any information about the full sample distribution and errors. To explore the effects of sidelobes, we removed the far sidelobe correction from this secondary pipeline execution, and we compare the results to the main pipeline mean maps.

Figure 6 shows the differences in the frequency maps between the two cases in temperature, where the effects are the most obvious. The most glaring features are the large dipole differences that can clearly be seen at 30 and 70 GHz. These are directly caused by the dipolar component seen in Fig. 4, as this contribution to the total sky signal that was in the sidelobe term is now unaccounted for. This explains the lack of much of a dipole signal at 44 GHz, as here the sidelobe contributions were at a much lower amplitude than the other frequencies. In previous analyses, these dipole contributions were handled through specific modeling of exactly these effects, but this makes it explicitly clear that correct dipole measurements require accurate knowledge of the sidelobe pickup.

The second obvious effect visible in Fig. 6 is a monopole offset of 10–15  $\mu\text{K}$  at all frequencies (note that the colour scales are not centered at 0). This monopole offset follows directly from the dipole problems discussed above. The internal gain of the instrument is calibrated by fitting the observed dipole signal to a model of the dipole components based on relative velocities and geometries (for more details see Gjerløw et al. (2021)). This gain coefficient is therefore directly affected by the spurious dipole signal injected by the sidelobes, which in turn causes an incorrect overall calibration of the maps. This bad calibration feeds directly into the overall amplitude of the CMB monopole at each

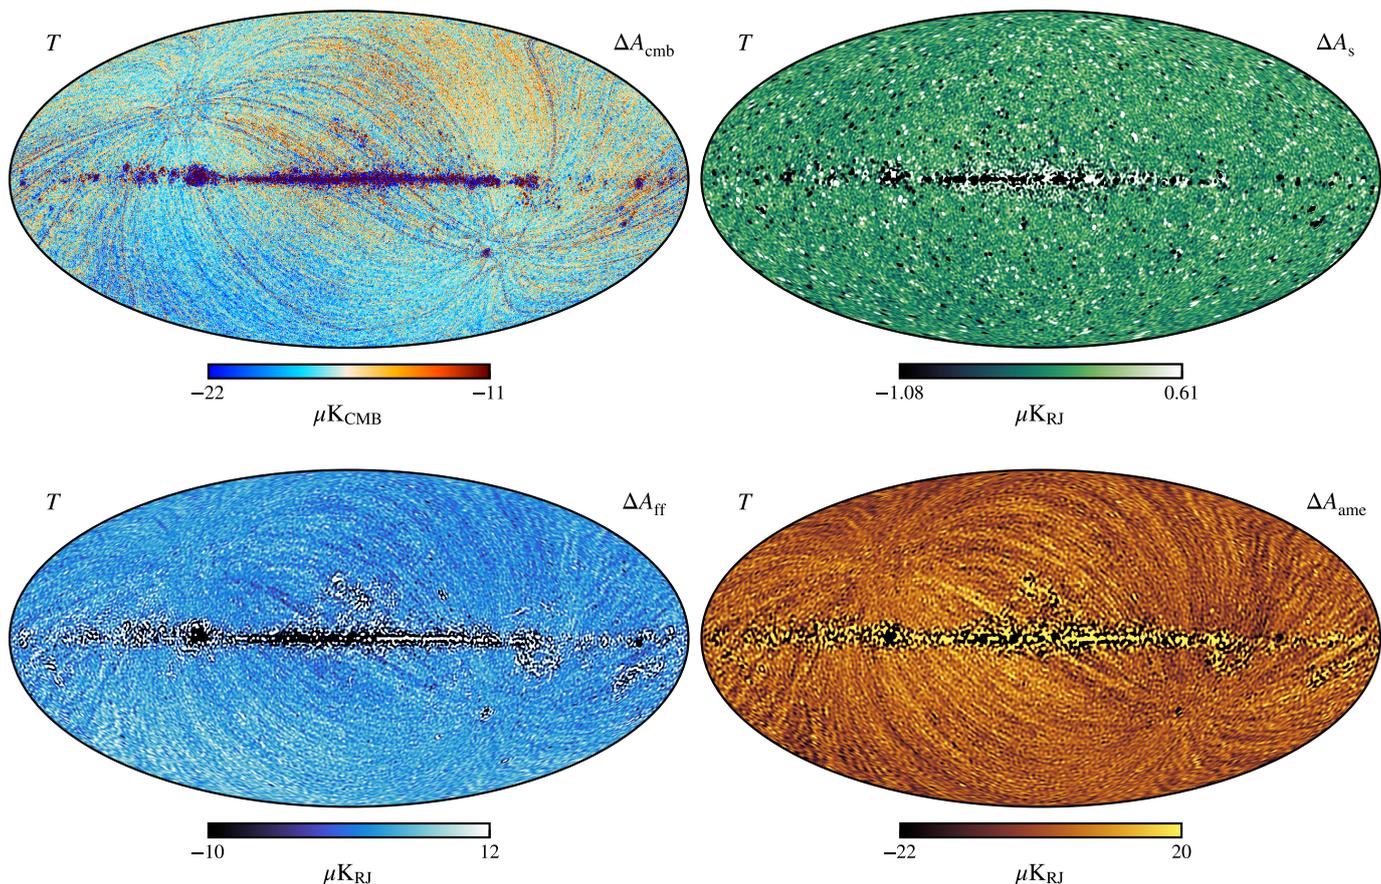
frequency, causing this offset that we see in the difference maps. Once again, this shows the importance of modeling all contributions appropriately, and demonstrates ways in which poor sidelobe knowledge could contribute directly to bad astrophysical sky parameters like the CMB temperature.

Finally, we see two more features in the difference maps that are more localized. The first of these are the ring structures that match the actual sidelobe map structures quite closely, these are of course the same rings from Fig. 4, which are not corrected for in the second pipeline run (where the sidelobe corrections are omitted). Additionally, there are some uniform residuals that are visible in the galactic plane regions of the difference maps. These are caused by calibration mismatch between the detectors at a single frequency. As each of the detectors now sees a slightly different dipole signal on the sky, depending on its specific sidelobe response, their calibrations do not agree with one another, which causes signal residuals which are most visible in the plane where the signal amplitude is highest.

Figure 7 shows the differences in component maps from this same comparison, also in temperature. The CMB as well as the three low-frequency foreground components are estimated using the standard Commander3 technique described in Andersen et al. (2021). The CMB component sees similar issues to the ones seen by the frequency maps above. The mono- and dipoles are incorrect, there are sidelobe-esque stripes and the galactic plane shows a strong residual, all of which are effects that have been seen directly in the frequency maps. This dipole difference seen here is precisely the one that contributes to the difference in calibration between the two different pipeline executions.

The other three low-frequency foreground components (synchrotron, AME and free-free) show less structural difference when compared. They have absorbed some of the monopole offsets present in the maps, as well as some additional sidelobe-like ring structures, but the primary difference can be seen most clearly in the galactic plane. Here, we notice a large residual caused by the inaccurate model of the galactic emission being altered slightly by the gain and calibration differences between the two runs. As the galactic emission is significantly brighter than the rest of the sky, small changes in calibration produce large errors like the ones seen here.

These residual errors of Figs. 6 and 7 must also be present in the BEYONDPLANCK analysis, albeit at much lower levels. We know that our knowledge and modeling of the sidelobes are imperfect, as they are based on limited measurements of the physical LFI sidelobes, and some of the power is unaccounted for. Future applications of the pipeline that aim for a robust  $r \leq 0.01$  measurement will be required to marginalize over the sidelobe uncertainties in some manner, either by directly Gibbs sampling a subset of the instrument parameters or by parameterizing and



**Fig. 7.** Component map difference plots for (clockwise from top left) cmb, synchrotron, AME and freefree emission, comparing the mean Gibbs maps from the BEYONDPLANCK release with a maximum-likelihood execution that has no sidelobe correction.

fitting sidelobe error estimates. We do not believe that the sidelobe contribution causes significant errors in the LFI sample sets produced by BEYONDPLANCK, as it is unlikely to be more than a 10 % error on the sidelobe estimates of Fig. 4. At 30 GHz, this corresponds to at most a 0.05 % error in our temperature maps and a 1 % error in polarization. We do expect however, that as instrumental sensitivities improve, especially in polarization, this sidelobe term will need to be modeled more accurately than has been done up until now.

## 6. Summary and Conclusions

The Conviqt algorithm is the fastest full-sky convolution technique currently available. This paper presents the implementation and verification of Conviqt within the Commander framework as well as characterization of the outputs. We showed estimates of sidelobe signals and RMS maps produced by the BEYONDPLANCK pipeline for the three LFI frequency bands. We also discussed the issues that could arise in an analysis that had poor sidelobe modeling, and showed that this issue would be much more important for planned high-sensitivity CMB projects, which would require better instrumental characterization and novel algorithms for marginalizing over sidelobe errors.

The full-sky sidelobe treatment techniques presented here are easily generalizable to other experiments, and can be tuned to match the required spatial characteristics of other instruments simply by adjusting the  $l_{max}$  and  $m_{max}$  of the sidelobe description. The only requirement for using the code with a new instrument is a HEALPIX-compatible description of the sidelobe response

function per detector. The more accurate this characterization of the instrument is, the better the sidelobe estimate will approximate the true sidelobe contamination in the timestream. We intend to use this formalism to correct sidelobes in the forthcoming Commander3 integration of WMAP, LiteBIRD and HFI, as well as on other as of yet unplanned projects.

The approach presented here is less useful for ground or balloon based experiments where the sidelobe pickup contains radiation from an environmental source. This pickup is not sky-synchronous, and thus cannot be modeled purely as a beam-sky convolution, but must include additional contributions from, for example, telescope baffles, ground pickup or clouds. For these types of experiments, other techniques such as aggressive baffling are likely better suited.

Future CMB experiments such as *LiteBIRD* that are targeting low  $B$ -mode limits may need to consider more complex ways of handling sidelobes and beams. One proposal that has been discussed is the full sky deconvolution of the full beam at each timestep, which could be achieved using a similar framework to the approach discussed here. This would remove the sidelobes as a nuisance signal from the data model of Eq. (13) and instead incorporate them directly into the beam term,  $B_{pp'j}$ . This approach would be feasible for a relatively low-resolution experiment like *LiteBIRD* and will be investigated going forward.

*Acknowledgements.* We thank Prof. Pedro Ferreira for useful suggestions, comments and discussions, and Dr. Diana Mjaschkova-Pascual for administrative support. We also thank the entire *Planck* and *WMAP* teams for invaluable support and discussions, and for their dedicated efforts through several decades without which this work would not be possible. The current work has received funding from the European Union's Horizon 2020 research and innovation programme

under grant agreement numbers 776282 (COMPET-4; BEYONDPLANCK), 772253 (ERC; BRTS2COSMOLOGY), and 819478 (ERC; COSMOGLOBE). In addition, the collaboration acknowledges support from ESA; ASI and INAF (Italy); NASA and DoE (USA); Tekes, Academy of Finland (grant no. 295113), CSC, and Magnus Ehrnrooth foundation (Finland); RCN (Norway; grant nos. 263011, 274990); and PRACE (EU).

## References

- Andersen et al. 2021, A&A, in preparation [arXiv:201x.xxxxx]
- BeyondPlanck Collaboration. 2021, A&A, in preparation [arXiv:2011.05609]
- Burigana, C., Maino, D., Górski, K. M., et al. 2001, A&A, 373, 345
- Gjerløw et al. 2021, A&A, in preparation [arXiv:2011.08082]
- Goldberg, J. N., Macfarlane, A. J., Newman, E. T., Rohrlich, F., & Sudarshan, E. C. G. 1967, Journal of Mathematical Physics, 8, 2155
- Hazumi, M., Ade, P., Akiba, Y., et al. 2019, Journal of Low Temperature Physics, 194
- Lewis, A. 2005, Physical Review D, 71
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, A&A, 571, A7
- Planck Collaboration, Ade, P. A. R., Aumont, J., et al. 2016, A&A, 594, A3
- Planck Collaboration III. 2016, A&A, 594, A3
- Planck Collaboration XII. 2016, A&A, 594, A12
- Planck Collaboration Int. XLVI. 2016, A&A, 596, A107
- Planck Collaboration Int. LVII. 2020, A&A, 643, A42
- Prézeau, G. & Reinecke, M. 2010, ApJS, 190, 267
- Reinecke, M. & Seljebotn, D. S. 2013, A&A, 554, A112
- Svalheim et al. 2021, A&A, in preparation [arXiv:201x.xxxxx]
- Wandelt, B. D. & Górski, K. M. 2001, Physical Review D, 63
- Wu, J. H. P., Balbi, A., Borrill, J., et al. 2001, The Astrophysical Journal Supplement Series, 132, 1–17