

Lawrence Berkeley National Laboratory

LBL Publications

Title

BEYONDPLANCK

Permalink

<https://escholarship.org/uc/item/5dz9h9r2>

Authors

Galloway, M

Reinecke, M

Andersen, KJ

et al.

Publication Date

2023-07-01

DOI

10.1051/0004-6361/202243138

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

BEYONDPLANCK VIII. Efficient Sidelobe Convolution and Corrections through Spin Harmonics

M. Galloway^{11*}, M. Reinecke¹³, K. J. Andersen¹¹, R. Aurlien¹¹, R. Banerji¹¹, M. Bersanelli^{4,9,10}, S. Bertocco⁸, M. Brilenkov¹¹, M. Carbone¹⁴, L. P. L. Colombo⁴, H. K. Eriksen¹¹, M. K. Foss¹¹, C. Franceschet^{4,10}, U. Fuskeland¹¹, S. Galeotta⁸, S. Gerakakis¹⁴, E. Gjerløw¹¹, B. Hensley², D. Herman¹¹, M. Iacobellis¹⁴, M. Ieronymaki¹⁴, H. T. Ihle¹¹, J. B. Jewell¹², A. Karakci¹¹, E. Keihänen^{3,7}, R. Keskitalo¹, G. Maggio⁸, D. Maino^{4,9,10}, M. Maris⁸, S. Paradiso^{4,9}, B. Partridge⁶, A.-S. Suur-Uski^{3,7}, T. L. Svalheim¹¹, D. Tavagnacco^{8,5}, H. Thommesen¹¹, D. J. Watts¹¹, I. K. Wehus¹¹, and A. Zacchei⁸

¹ Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, U.S.A.

² Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, U.S.A.

³ Department of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland

⁴ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria, 16, Milano, Italy

⁵ Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, Trieste, Italy

⁶ Haverford College Astronomy Department, 370 Lancaster Avenue, Haverford, Pennsylvania, U.S.A.

⁷ Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland

⁸ INAF - Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, Trieste, Italy

⁹ INAF-IASF Milano, Via E. Bassini 15, Milano, Italy

¹⁰ INFN, Sezione di Milano, Via Celoria 16, Milano, Italy

¹¹ Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway

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

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

¹⁴ Planetek Hellas, Leoforos Kifisias 44, Marousi 151 25, Greece

January 11, 2022

ABSTRACT

We introduce a new formulation of the `Conviqt` convolution algorithm in terms of spin harmonics, and apply this to the problem of sidelobe correction for `BEYONDPLANCK`, the first end-to-end Bayesian Gibbs sampling framework for CMB analysis. We compare our implementation to the previous `Planck` `LevelS` implementation, and find good agreement between the two codes in terms of accuracy, but with a speed-up reaching a factor of 3–10, depending on the frequency bandlimits, l_{\max} and m_{\max} . The new algorithm is significantly simpler to implement and maintain, since all low-level calculations are handled through an external spherical harmonic transform library. We find that our mean sidelobe estimates for `Planck` LFI agree well with previous efforts. Additionally, we present novel sidelobe rms maps that quantify the uncertainty in the sidelobe corrections due to variations in the sky model.

Key words. Cosmology: observations, polarization, cosmic microwave background — Methods: data analysis, statistical

1. Introduction

One of the important systematic effects that must be accounted for in CMB instruments is telescope stray light and sidelobes (e.g., Barnes et al. 2003; Planck Collaboration III 2014). This is the non-zero response of the detector to areas of the sky outside the main beam, however that is defined. Because microwave telescopes typically work near their diffraction limit, some level of sidelobes is inevitable. Furthermore, their structure can be complicated 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, particularly at large angular scales, as they do not behave in the same sky-stationary 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 (e.g., Planck Collaboration III 2016; Watts et al. 2022). In particular, as the sidelobe response functions often

are broadly distributed, this contamination can confuse important signals like the CMB solar and orbital dipoles that are used for calibration. Sidelobe uncertainties couple directly with foreground emission from diffuse galactic components, producing an important source of contamination. In some experiments a further contaminating signal can originate from a source not on the sky, such as ground pickup or radio-frequency (RF) noise. In all cases, sidelobe signal is detrimental to the quality of the final sky maps and parameter estimates, and requires a dedicated effort to remove it. 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 to convolve a beam or sidelobe response function with a sky map or model to generate a re-observed map.

Full sky convolution on the sphere is a problem that has been important in the CMB field since the earliest satellite measurements. Early experiments like *COBE* (Torralba et al. 1989) either did not model sidelobes at all, or used simple pixel-based convolution approaches which even for their low resolution required

* Corresponding author: M. Galloway; mathew.galloway@astro.uio.no

radially symmetric 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, often referred to as “total convolution”, which achieved a large performance gain over pixel-based methods, by as much as a factor of $O(\sqrt{N_{\text{pix}}})$. This breakthrough allowed the calculation of these convolutions easily enough that they could be applied to each simulation, instead of requiring a dedicated study necessitating months of runtime.

Next, Prézeau & Reinecke (2010) developed the `Conviqt` approach, which was used both by several official *Planck* analysis pipelines (Planck Collaboration III 2016, Planck Collaboration Int. XLVI 2016, Planck Collaboration Int. LVII 2020) and to generate the *Planck* Full Focal Plane (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 the original harmonic space algorithm, as well as providing a standardized, user friendly library, `libconviqt`, that was incorporated into numerous pipelines.

In this paper, we introduce a new formulation of the `Conviqt` algorithm that is based on Spherical Harmonic Transforms (SHTs), rather than directly computing the Wigner matrix elements. We 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 the first to explicitly derive, discuss, and benchmark the method.

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 (ϑ, φ) ,¹ and rotated around its own central axis by ψ ,

$$\begin{aligned} c(\vartheta, \varphi, \psi) &:= s_{l m_s} * b_{l m_b} \\ &\equiv \int_{4\pi} s(\hat{n}) b(\hat{n}'(\vartheta, \varphi) - \hat{n}, \psi) d\Omega_{\hat{n}}. \end{aligned} \quad (1)$$

Here, $s_{l m_s}$ denotes the spherical harmonic coefficients of the sky signal, and $b_{l m_b}$ is the beam in the same representation. In this expression, care must be taken to distinguish between the sky and beam bandlimits, m_s and m_b , as the two indices will be treated separately in the following derivation.

A computationally efficient solution for this problem was derived by Prézeau & Reinecke (2010), who exploited fast recurrence relations for Wigner d matrix elements to evaluate Eq. (1) in harmonic space. In the following, we will show that this equation 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 equations are simply repackaged in a format that is significantly easier to implement in practical computer code, since it may use existing and highly optimized spherical harmonic libraries, such as `libsharp` (Reinecke & Seljebotn 2013), to perform the computationally expensive parts.

¹ In this paper, ϑ and φ are the co-latitude and longitude of a location on the sphere, i.e., they have the same meaning as in the HEALPIX context (Górski et al. 2005).

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 harmonic coefficients of the signal and beam, respectively, and D_{m_s, m_b}^l is the Wigner D -matrix.

This may be expressed as (Goldberg et al. 1967)

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

where ${}_s Y_{lm}(\vartheta, \varphi)$ is the spin-weighted spherical harmonic and the placement of the negative signs are an arbitrary historical convention. Inserting this expression into Eq. (2) yields

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

where we have assumed that the beam is real-valued in position space, and we have used the symmetry relations,

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

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

Pulling the summation over m_b in front of the other sums yields

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

The terms in this outer sum can be arranged in the form $m_b = 0, \pm 1, \pm 2, \dots$. The contribution from $m_b = 0$ can be interpreted 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).

Since $c(\vartheta, \varphi, \psi) \in \mathbb{R}$, we know that the contributions from the pairs $m_b = \pm 1, \pm 2, \dots$ must be complex conjugate with respect to each other, and their combined contribution is therefore

$$\begin{aligned} e^{i m_b \psi} {}_{m_b} S(\vartheta, \varphi) + e^{-i m_b \psi} {}_{m_b} S^*(\vartheta, \varphi) = \\ 2 [\cos(m_b \psi) \text{Re}({}_{m_b} S(\vartheta, \varphi)) - \sin(m_b \psi) \text{Im}({}_{m_b} S(\vartheta, \varphi))], \end{aligned} \quad (8)$$

where we have defined

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

This is a spherical harmonic transform of a quantity with spin m_b , which can also be computed efficiently by `libsharp`.

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

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

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

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




Fig. 1. Map level difference in temperature of the new SHT convolution algorithm compared to the old Conviqt approach, observed using the identical pointing of the first year of the 30 GHz *Planck* detectors. The differences are at the level of machine precision, indicating full agreement between the two algorithms.

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 SHT 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, (ϑ, φ) , perform a Fourier transform to convert these coefficients to $c(\vartheta, \varphi, \psi)$, as given by Eq. (7).

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. 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 `libconviqt`

To compare the results of this new total convolution approach with the older `libconviqt` approach of Prézeau & Reinecke (2010), we evaluate the convolution between the beam for one of the LFI 30 GHz receivers (28M) and a Commander 30 GHz sky model (Andersen et al. 2022; Svalheim et al. 2022) using both methods. The resulting convolution cubes are then observed using LFI’s scanning strategy for the first year of the *Planck* flight. The resulting map differences are shown in Fig. 1. The convolution cubes are also directly compared for accuracy, and found to agree with an integrated difference at the 10^{-8} level, indicative of differences at the level of numerical precision.

Figure 2 compares the runtime between the two approaches for a test configuration with an elliptical beam and a fixed sky model, and with $m_{\max} = 0$ (using only the radially symmetric part of the beam) and $m_{\max} = 10$, respectively. In both cases, the new implementation outperforms the old approach at all but the lowest l_{\max} , where the data read time dominates. Additionally, for compatibility with the old `libconviqt` approach, this test was performed with an older version of `libsharp`, so we expect that the new algorithm scales even more favourably than this with the latest implementation. We note that this is a significant real-life advantage of the new approach: any improvement in SHT libraries, which typically are subject to intensive al-




Fig. 2. Runtime comparison between the `libconviqt` 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. Note the log scale on the y-axis.

gorithm development and code maintenance, translates directly into a computational improvement for the convolution algorithm.

3. Sidelobe Models

Figure 3 shows characteristic sidelobe response functions evaluated at a fixed frequency on the sky for a detector in each *Planck* LFI band. The sidelobe response for each detector within a single *Planck* band look visually quite similar, so only these representative ones are shown here. Each is stored on disk as a set of a_{lm} ’s with $l_{\max} = 512$ and $m_{\max} = 100$.

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 correction term, as can be seen in the BEYONDPLANCK global parametric model of the data:

$$d_{j,t} = g_{j,t} \left[\mathbf{P}_{tp,j} \mathbf{B}_{pp',j} \sum_c \mathbf{M}_{cj}(\beta_{p'}, \Delta b p^j) a_{p'}^c + s_{j,t}^{\text{orb}} + s_{j,t}^{\text{fsl}} \right] + s_{j,t}^{\text{1hz}} + n_{j,t}^{\text{corr}} + n_{j,t}^{\text{w}}. \quad (13)$$

The other terms in this equation are discussed in detail in BeyondPlanck (2022), but here the main beam signal 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 needed to model it accurately would require much higher l_{\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 4π dipole response, as detailed in Sect. 3.3. Additionally, a Gaussian main beam approximation is used during component separation to




Fig. 3. Maps of the sidelobe response on the sky from a representative 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 $l = 0$.

smooth the sky model to the appropriate beam resolution for each channel. During mapmaking, beam effects are ignored and the beam is assumed to be pointed at the center of each pixel.

3.2. Sidelobe Normalization

We adopt a normalization of the sidelobes that differs slightly from the normalization used within the *Planck* LFI collaboration. The *Planck* 2018 LFI beam products leave a small portion (around 1%) of known missing power within the system unassigned due to uncertainties about to which component it should be assigned (Planck Collaboration IV 2016). In the current analysis, we rather adopt the same approximation as for *Planck* DR4 (Planck Collaboration Int. LVII 2020), and renormalize the beam transfer function such that this power is distributed proportionally at each l ; that is, we rescale the beam transfer function B_l such that its full sky integral is $B_0 = 1$. This re-scaling is equivalent to assigning the unknown beam power uniformly over the entire beam. We note that this normalization is in either case always done before any higher-level analysis for both *Planck* 2018 and DR4; the only difference is whether the renormalization must be performed by external users through deconvolution of a non-unity normalized main beam transfer function or not.

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 (Planck Collaboration I 2020; Planck Collaboration Int. LVII 2020; Gjerløw et al. 2022), the sidelobe’s contribution to the dipole can directly result in an absolute calibration error if not handled appropriately. While the CMB Solar dipole can easily be handled using the Conviqt approach described in Sec. 2.1, the orbital dipole is not sky-stationary and thus must be handled separately.

BEYONDPANCK 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.

BEYONDPANCK uses the same technique as *Planck* DR4 (Planck Collaboration Int. LVII 2020, see appendix C) to generate the orbital dipole and quadrupole estimate. That is, we ex-

press 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 4π 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) , and we adopt the convention that the main beam points toward the north pole in our coordinate system.

The orbital CMB dipole and quadrupole can be expressed as a Doppler shift in each direction (Notari & Quartin 2015),

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

where β 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 ν , defined by

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

Inserting these expressions into Eq. (15), one obtains

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

where $\hat{n} = (x, y, z)$ is a unit direction vector that is also the integration variable, \hat{n}_0 is the fixed direction of the satellite pointing for this timestep. Noting that the geometric factors in this expression may be precomputed as

$$S_x = \int x B(\hat{n}, \hat{n}_0) d\Omega_{\hat{n}}, \quad (18)$$

we see that Eq. (17) may be written in the following form,

$$\begin{aligned} \tilde{D} = T_0 \left[S_x\beta_x + S_y\beta_y + S_z\beta_z + \right. \\ \left. q \left(S_{xx}\beta_x^2 + S_{yy}\beta_y^2 + S_{zz}\beta_z^2 + \right. \right. \\ \left. \left. 2S_{xy}\beta_x\beta_y + 2S_{xz}\beta_x\beta_z + 2S_{yz}\beta_y\beta_z \right) \right]. \end{aligned} \quad (19)$$

To compute \tilde{D} for an arbitrary beam orientation, one simply needs to rotate the satellite pointing and velocity vectors into the coordinate system used to define S , and then one can evaluate Eq. (19) very quickly.




Fig. 4. Maps of the sidelobes convolved with the sky 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.

BEYONDPLANCK further accelerates this operation by computing this rotation for only one point in twenty (chosen so as to still fully sample the dipole), and using a spline to interpolate between them. This saves the cost 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.

4. Sidelobe Estimates

4.1. Posterior mean corrections

Figure 4 shows the mean sidelobe signal estimates at each of the three LFI frequencies for the entire mission, co-added across each frequency and projected into sky coordinates, identically to the way the true sky signal is treated. Each map is averaged over 90 Gibbs samples produced in the main BEYONDPLANCK analysis (BeyondPlanck 2022), after discarding burn-in and thinning the remaining chain by a factor of ten.

We note that these maps follow the traditional *Planck* LFI method of sidelobe correction by producing these signals in the time domain during TOD processing. These templates are therefore exactly correct for the maps produced by these pipeline runs, but will not match precisely with analyses that use different data cuts, flagging or channel selection.

These results look visually similar to the corresponding *Planck* DPC results presented in Fig. 7 of Planck Collaboration III (2016). The main difference is that the current results also include the CMB dipole, whereas the LFI 2015 DPC analysis showed the sidelobe pickup of dipole-subtracted maps. We see that the sidelobe signal is strongest at 30 GHz, and that the dominant features in the co-added sky maps consist of a series of rings created by the interplay between the sidelobe pickup and bright Galactic plane features.

Figure 4 also clearly indicates that the overall level of sidelobe pickup at 44 GHz is significantly lower than for the 30 and 70 GHz channels. This is due to the particular location in the focal plane of two of the three 44 GHz feedhorns, which results in a significant under-illumination of both the primary and secondary reflectors of the *Planck* telescope for those two horns (see Fig. 4 of Sandri et al. (2010)).

4.2. Error Propagation

In addition to the posterior mean sidelobe maps, the BEYONDPLANCK pipeline outputs also provide an estimate of the sidelobe stability and statistical variation. Figure 5 shows the rms maps generated from the same sample of sidelobe signal 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 be-




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. Again note the different colour scales.

tween 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 VII \(2014\)](#).

These posterior rms maps cannot be considered true sidelobe error estimates, however, as they do not account for uncertainties in the sidelobe response itself. Rather, they only show the change in the estimated sidelobe signal due to sky model variations from component separation. Full sidelobe error propagation would require sampling over the physical parameters that determine the detectors' sidelobe response on the sky. While sampling the full set of optical model parameters is likely to be infeasible due to excessive computational time, identifying a minimal parameter set that may account for the main potential variations in the sidelobe response functions, and precomputing response functions over a grid of such parameters, would result in physically motivated uncertainties for the sidelobe models. This approach will be developed for future applications such as the *LiteBIRD* mission ([Hazumi et al. 2019](#)).

5. Impact on Frequency and Component Maps

To assess the importance of sidelobe corrections on frequency and component maps, we perform two runs of the `Commander` code, starting from the same input data as the main pipeline run, and with identical random seeds. As a comparison, we remove the far sidelobe correction from one of these secondary pipeline executions, and we difference the results between these two pipelines.

Figure 6 shows the differences in the frequency maps between the two cases in temperature, where the effects are the most obvious. The only large-scale features that can clearly be seen are the small dipole differences (most clearly visible at 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. In previous analyses, these dipole contributions were handled through specific modeling of exactly these effects, but this test makes it explicitly clear that correct dipole measurements require accurate knowledge of the sidelobe pickup.

Next, 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 accounted for in the second pipeline run without sidelobe corrections. 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, again 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. \(2022\)](#). The AME component sees similar issues to the ones seen




Fig. 6. Frequency map difference plots smoothed to one degree at (top to bottom) 30, 44 and 70 GHz, comparing two pipeline executions with the same seed, one of which has no sidelobe correction.

by the frequency maps above. The dipole is slightly 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. The dipole difference seen here is precisely the one that contributes to the difference in calibration between the two different pipeline executions.

The other three components (synchrotron, CMB and free-free) show relatively less structural difference. They have absorbed some of the 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.

Finally, in Fig. 8, we show the correlated noise map difference at 30 GHz between the two runs. Here, we see that some of

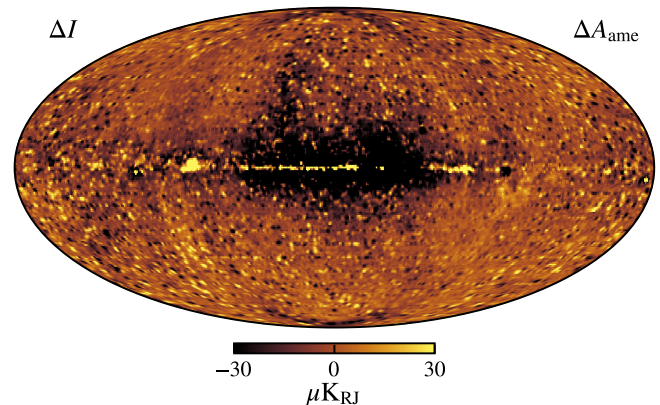


Fig. 7. Component map difference plots for (top to bottom) CMB, synchrotron, AME and free-free emission, comparing two pipeline executions with the same seed, one of which has no sidelobe correction.

the missing sidelobe signal has been accommodated by the correlated noise component. These structures mirror the strongest sidelobe-like signals in the 30 GHz difference map in the top




Fig. 8. Difference in correlated noise, projected into the map domain at 30 GHz comparing two pipeline executions with the same seed, one of which has no sidelobe correction.

panel of Fig. 6. The fact that n_{corr} can accommodate some sidelobe residuals is in fact helpful, as it allows for some leeway in the final sidelobe model, in the sense that small uncertainties and artifacts that are inconsistent between different frequency channels will be mostly absorbed into the correlated noise component, rather than in the sky maps. The differences that we see in the maps, however, indicate that this process is not perfect, as some of the spurious signal still makes it to the final maps without a perfect sidelobe model. For a real-world example of these issues, we refer the interested reader to the on-going BEYONDPLANCK re-analysis of the *WMAP* data, for which far sidelobe contamination appears to be a dominant problem (Watts et al. 2022).

The residual errors seen in Figs. 6 and 7 are also 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 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 very accurately, and the corresponding uncertainties must be propagated properly, for instance using methods similar to those presented in this paper.

6. Summary and Conclusions

This paper presents a formulation of the *Conviqt* algorithm in terms of spin-weighted spherical harmonics. This algorithm is already implemented in the latest versions of the *libconviqt* library, and it has now also been re-implemented directly into *Commander*, where it is used for sidelobe corrections for the BEYONDPLANCK analysis framework. Based on the Monte Carlo samples produced in that analysis, we have presented novel posterior mean and standard deviation maps for each of the three *Planck* LFI frequency bands.

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 spherical harmonic bandpass parameters, 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 note that the approach presented here is less useful for ground or balloon based experiments where the dominant 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.

We also stress that the current implementation only supports sidelobe error propagation for sky model uncertainties, not uncertainties in the actual sidelobe response function itself, and these are very likely to dominate the total sidelobe error budget. Future work should therefore aim to introduce parametric models for the sidelobe response itself, and sample (or, at least, marginalize) over the corresponding free parameters, as these are typically one of the most important unknowns for many experiments.

Finally, we note that future CMB experiments such as *LiteBIRD*, which are targeting low B -mode limits, may need to consider more complex ways of handling sidelobes and beams. The ultimate solution in this respect is 4π beam convolution for every single 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 should be feasible for a relatively low-resolution experiment such as *LiteBIRD*, and will be investigated going forward.

Acknowledgements. We thank Prof. Pedro Ferreira and Dr. Charles Lawrence for useful suggestions, comments and discussions. 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; BITS2COSMOLOGY), and 819478 (ERC; COSMOGLOBE). In addition, the collaboration acknowledges support from ESA; ASI and INAF (Italy); NASA and DoE (USA); Teles, 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. 2022, A&A, in preparation [arXiv:201x.xxxxx]
- Barnes, C., Hill, R. S., Hinshaw, G., et al. 2003, ApJS, 148, 51
- BeyondPlanck. 2022, A&A, in preparation [arXiv:2011.05609]
- Burigana, C., Maino, D., Górski, K. M., et al. 2001, A&A, 373, 345
- Gjerløw et al. 2022, A&A, in preparation [arXiv:2011.08082]
- Goldberg, J. N., Macfarlane, A. J., Newman, E. T., Rohrich, F., & Sudarshan, E. C. G. 1967, Journal of Mathematical Physics, 8, 2155
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
- Hazumi, M., Ade, P., Akiba, Y., et al. 2019, Journal of Low Temperature Physics, 194, 443
- Lewis, A. 2005, Phys. Rev. D, 71, 083008
- Notari, A. & Quartin, M. 2015, Journal of Cosmology and Astroparticle Physics, 2015, 047–047
- Planck Collaboration III. 2014, A&A, 571, A3
- Planck Collaboration VII. 2014, A&A, 571, A7

- Planck Collaboration III. 2016, A&A, 594, A3
Planck Collaboration IV. 2016, A&A, 594, A4
Planck Collaboration XII. 2016, A&A, 594, A12
Planck Collaboration I. 2020, A&A, 641, A1
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
Sandri, M., Villa, F., Bersanelli, M., et al. 2010, A&A, 520, A7
Svalheim et al. 2022, A&A, in preparation [arXiv:201x.xxxxx]
Torral, M., Ratliff, R., Lecha, M., et al. 1989, Antennas and Propagation, IEEE Transactions on, AP-37, 171
Wandelt, B. D. & Górski, K. M. 2001, Phys. Rev. D, 63, 123002
Watts et al. 2022, A&A, in preparation [arXiv:201x.xxxxx]
Wu, J. H. P., Balbi, A., Borrill, J., et al. 2001, The Astrophysical Journal Supplement Series, 132, 1–17