

BEYONDPLANCK XIV. Intensity foreground sampling, degeneracies and priors

K. J. Andersen^{11*}, R. Aurlen¹¹, 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, M. Galloway¹¹, 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. Paradiso⁴, B. Partridge⁶, M. Reinecke¹³, A.-S. Suur-Uski^{3,7}, T. L. Svalheim¹¹, D. Tavagnacco^{8,5}, H. Thommesen¹¹, M. Tomasi^{4,9}, 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 Hållströmin 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 Hållströmin 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 25, 2022

ABSTRACT

We present the intensity foreground algorithms and model employed within the BEYONDPLANCK analysis framework. The BEYONDPLANCK analysis includes a limited set of frequency channels, and it is therefore particularly sensitive to parameter degeneracies. We discuss various priors that are introduced to break these degeneracies, and we improve the previous *Planck*-based *Commander* component separation implementation in four specific ways that are all designed to improve stability and computational efficiency for weakly constrained posterior distributions. These are 1) joint foreground spectral parameter and amplitude sampling, building on ideas from MIRAMARE; 2) component-based monopole determination; 3) joint spectral parameter and monopole sampling; and 4) application of informative spatial priors for component amplitude maps. We find that the only spectral parameter with a significant signal-to-noise ratio using the current BEYONDPLANCK data set is the peak frequency of the anomalous microwave emission component, for which we find $\nu_p = 25.3 \pm 0.5$ GHz; all others must be constrained through external priors. Future work will aim at integrating many more data sets into this analysis, both map and time-ordered based, and thereby gradually eliminating the currently observed degeneracies in a controlled manner with respect to both instrumental systematic effects and astrophysical degeneracies. This work will be organized within the Open Science-based COSMOGLOBE community effort.

Key words. ISM: general – Cosmology: observations, cosmic microwave background, diffuse radiation – Galaxy: general

Contents

1 Introduction	2	3.3 Joint spectral parameter and frequency-band monopole sampling	10
2 Overview of the BEYONDPLANCK sampling framework	3	3.4 Breaking small-scale degeneracies through spatial priors	10
2.1 Data model and Gibbs chain	3	3.4.1 Algorithmic smoothing priors	11
2.2 Data survey	4	3.4.2 Informative Gaussian spatial priors	12
2.3 Astrophysical sky model	6	4 BEYONDPLANCK analysis and posterior distributions	14
3 Commander extensions for efficient intensity-based component separation	6	4.1 Spectral parameter prior tuning	14
3.1 Joint amplitude and spectral parameter sampling	6	4.2 Markov chain trace plots and correlations	17
3.2 Component-based monopole determination	8	4.3 Goodness-of-fit statistics	17
		4.4 Signal component posterior distributions	20
		5 Summary and outlook	25

* Corresponding author: K. J. Andersen; k.j.andersen@astro.uio.no

1. Introduction

The cosmic microwave background (CMB) represents one of our best sources for knowledge of the early universe. The intensity of this radiation peaks at about 161 GHz (Mather et al. 1994), in the microwave frequency range, and a long line of experiments have therefore targeted these wavelengths ever since the CMB was first discovered by Penzias & Wilson (1965). However, there are many sources of radiation in the microwave sky that obscure our view of the CMB, both from our own galaxy and from distant sources (see, e.g., Delabrouille et al. 2013; Planck Collaboration IV 2018, and references therein). All these sources must be modelled to high accuracy in order to establish a clean estimate of the CMB sky.

The current paper presents the low-frequency foreground sky model and component separation algorithms that are used for intensity analysis within the BEYONDPLANCK framework, as applied to the *Planck* (Planck Collaboration I 2020) Low Frequency Instrument (LFI; Planck Collaboration II 2020), *Planck* HFI DR4 857 GHz (Planck Collaboration Int. LVII 2020), *WMAP* (Bennett et al. 2013) and Haslam 408 MHz (Haslam et al. 1982) data. This is a novel end-to-end Bayesian CMB analysis framework that builds on decades of experience gained within the *Planck* collaboration, and its main defining feature is that instrument characterization and calibration is performed jointly with map-making and component separation, resulting in a single statistically consistent model for the full data set. As such, the foreground sky model plays a key role in the process, feeding directly into many aspects of the analysis, from gain and correlated noise estimation via leakage corrections and mapmaking to final component maps, CMB estimates and cosmological parameters. An overview of the full process is provided in BeyondPlanck (2022) and its companion papers.

The main target of the current BEYONDPLANCK project is the *Planck* LFI data set for three main reasons. First, the BEYONDPLANCK analysis considers time-ordered data. Since the LFI data volume is limited, requiring only 1.5 TB of RAM (Galloway et al. 2022a), this data set allows a rapid algorithm development and debug cycle. Second, the inherent LFI systematics are relatively benign due to a highly stable instrumental design (Planck Collaboration III 2016). Third, the BEYONDPLANCK team already has extensive expertise with the LFI data sets following decades of *Planck* analysis work. As such, the LFI data set represents a very useful test case for development purposes.

At the same time, the restricted frequency range of the LFI data, covering 30 to 70 GHz, enforces strict limitations on what foreground parameters are possible to constrain. Furthermore, a conscious choice has been made within the BEYONDPLANCK project that the LFI data should play the main statistical role in the current analysis, as the main goal of this work is to demonstrate the methodology itself, and not to establish the most sensitive component maps possible. Providing the most sensitive sky model is one of the goals of the community-driven and Open Science COSMOGLOBE¹ project, which will employ the computational framework derived here. For this reason, only a minimal set of external data is included in the current analysis, as is required to establish a minimal viable foreground model, but not so much that the statistical power of LFI is dominated by the same external data.

To be specific, the most important astrophysical signal components for *Planck* LFI are CMB, synchrotron, free-free, anomalous microwave emission (AME), and thermal dust emission (Planck Collaboration X 2016; Planck Collaboration IV 2018).

Clearly, constraining a model with five independent components with only three frequency bands represents a mathematically degenerate problem. To break these degeneracies, we add the Haslam 408 MHz survey (Haslam et al. 1982) to constrain synchrotron emission; the *Planck* 857 GHz channel to constrain thermal dust (Planck Collaboration I 2020); and the *WMAP* *Ka*-*V* bands (covering 33 to 61 GHz; Bennett et al. 2013) to constrain AME and free-free emission. Notably, we do not include the intermediate *Planck* HFI frequencies, because these would completely dominate the derived CMB products, and they could in theory also introduce systematic uncertainties that are not possible to control or mitigate within the current framework. The same argument applies to the *WMAP* *K*-band, which has similar or higher statistical constraining power with respect to the low-frequency foregrounds than the LFI 30 GHz channel, and therefore we omit the *WMAP* *K*-band in this work.

With a total of eight frequency maps and five component maps, the model is formally mathematically non-degenerate. However, it is by no means tightly constrained, primarily because the *Planck* LFI and *WMAP* frequency ranges overlap significantly, and therefore probe very similar astrophysical effects. The main theme of the current paper is therefore not the properties of the astrophysical foregrounds themselves, since these have a higher signal-to-noise and less degeneracies (though possibly more systematics) when including all available data sets, as for instance was done in Planck Collaboration X (2016). Instead, its main purpose is to explore the degeneracies that exist within the current model, and show how these may be regularized through the use of informative and physically meaningful priors. This analysis thus demonstrates the role of priors in a classic Bayesian setting, and paves a robust road toward future analyses that will include additional data sets.

The algorithms used in this paper derives most closely from a similar Commander-based (Eriksen et al. 2004, 2008) Bayesian analysis of the *Planck*, *WMAP*, and Haslam et al. data presented by Planck Collaboration IX (2016). The main differences between the two analyses are the following. First, in the previous analysis there was limited feedback between the gain estimation and component separation, as only a single overall absolute calibration factor was fitted during the component separation process. In the current analysis, a full time-dependent time-ordered data model is propagated throughout the analysis. Second, the previous analysis provided only a single maximum likelihood solution for each component together with marginal per-pixel uncertainties. In the present analysis, we provide a full Monte Carlo ensemble of samples drawn from the full posterior distribution, which supports full end-to-end propagation of all uncertainties. Third, the computational framework used in the *Planck* 2015 analysis required uniform resolution across all frequency bands, thus limiting the resolution to that of the instrument with the poorest resolution. This was improved upon in the *Planck* 2018 analysis (Planck Collaboration IV 2018), where a new computational framework that allowed for different resolution and smoothing scales was developed. However, the corresponding analysis only included a single joint low-frequency component in intensity, and did not provide new estimates of the individual low-frequency components. In the current analysis, we provide full-resolution parameter maps for all components listed above, limited only by the signal-to-noise ratio of the data in question.

In this paper, we also implement four important new algorithmic features into the Commander framework that are all designed to improve sampling efficiency and stability for weakly constrained posterior distributions. The first of these is

¹ <http://cosmoglobe.uio.no>

a joint spectral parameter and amplitude sampler that employs the marginal spectral parameter posterior distribution to move quickly through the multi-dimensional parameter space. This algorithm was first introduced in the CMB literature by [Stompor et al. \(2009\)](#), and implemented in the MIRAMARE component separation code by [Stivoli et al. \(2010\)](#). The main advantage of this approach is a significantly reduced Monte Carlo Markov Chain (MCMC) correlation length, and thereby overall lower computational costs. In this paper, we demonstrate this sampler on the peak frequency of the AME emission spectral energy density (SED). However, the importance of this new step will increase significantly when additional spectral parameters are explored in the future, and will for instance be critically important to probe the thermal dust spectral index and temperature efficiently with *Planck* HFI.

The second algorithmic improvement is component-based monopole determination. As discussed at length by, e.g., [Planck Collaboration X \(2016\)](#) and [Wehus et al. \(2017\)](#), one of the most important challenges regarding intensity-based spectral parameter estimation is accurate determination of the zero-levels at each frequency band; any error in this will translate directly into a bias in spectral parameters, which typically manifests itself as a spatial correlation between the spectral parameter map and the corresponding component amplitude map. At the same time, it is important to note that few, if any, CMB experiments actually have sensitivity to the true sky monopole, and they have therefore typically instead resorted to morphology-based algorithms to determine meaningful zero-levels. In this paper, we point out that this problem may be entirely circumvented by instead focusing on determining the zero-levels of the astrophysical *component* maps, rather than individual *frequency* maps. The resulting global sky model may then be used to determine the frequency map offsets. This approach is significantly more transparent from a physical point-of-view (e.g., the CMB temperature perturbation component may be assumed to have an identically vanishing monopole), it automatically guarantees consistency between the zero-levels at different frequency channels, and it gives zero statistical weight to the frequency monopoles in the actual fitting procedure.

Nevertheless, there is a formal degeneracy between the spectral parameters and the frequency monopoles at each step in the algorithm, and to eliminate the Markov chain correlation length increase from this degeneracy, we additionally implement a new joint spectral parameter and monopole sampler as our third algorithmic improvement.

Fourth and finally, we generalize the concept of informative spatial component map priors that was introduced by [Planck Collaboration IV \(2018\)](#) and [Planck Collaboration Int. LVII \(2020\)](#), and use this to apply informative physical Gaussian spatial priors. This can be leveraged to significantly reduce correlations between various sky components on small angular scales, and in particular degeneracies between AME, free-free and CMB ([Colombo et al. 2022](#)) may be alleviated in this manner. Of course, the ideal approach to resolve such degeneracies is not through the use of informative priors, but rather by integrating additional data. The current algorithms allow for a gradual and controlled introduction of such data sets, without introducing pathological artefacts along the way.

The rest of the paper is organized as follows. Section 2 gives a short overview of the BEYONDPLANCK framework, data selection, and sky signal model, while Sect. 3 describes the main sampling algorithms for intensity foreground parameters. The main results are summarized in Sect. 4, before we conclude in

Sect. 5. We note that polarization-based component separation results are discussed separately by [Svalheim et al. \(2022b\)](#).

2. Overview of the BEYONDPLANCK sampling framework

The BEYONDPLANCK project has implemented an integrated end-to-end data analysis pipeline for CMB experiments ([BeyondPlanck 2022](#)), connecting all steps going from raw time-ordered data to cosmological analysis in a self-consistent Bayesian framework, finally realizing ideas originally proposed almost 20 years ago by [Jewell et al. \(2004\)](#) and [Wandelt et al. \(2004\)](#). This methodology allows us to characterize degeneracies between instrumental and astrophysical parameters in a statistically well-defined framework, with uncertainties propagating consistently through all stages of the pipeline. It also seamlessly connects low-level instrumental quantities like gain ([Gjerl w et al. 2022](#)) and correlated noise ([Ihle et al. 2022](#)), bandpasses ([Svalheim et al. 2022a](#)), and far sidelobes ([Galloway et al. 2022b](#)) via Galactic parameters such as the synchrotron amplitude and spectral index (current paper and [Svalheim et al. 2022b](#)), to the angular CMB power spectrum and cosmological parameters ([Colombo et al. 2022](#); [Paradiso et al. 2022](#)). In this section, we provide a brief review of the BEYONDPLANCK sky model, data selection, and sampling scheme, and we refer the interested reader to the various companion papers for full details.

2.1. Data model and Gibbs chain

In BEYONDPLANCK the most basic data sets are raw un-calibrated time-ordered data (TOD), which are modelled as follows,

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

Here j represents a radiometer label, t indicates a single time sample, p denotes a single pixel on the sky, and c represents one single astrophysical signal component. Further, g denotes the instrumental gain; \mathbf{P} denotes the pointing matrix; \mathbf{B}^{symm} and $\mathbf{B}^{\text{asymm}}$ denote the symmetric and asymmetric beam matrix, respectively; \mathbf{a} are the astrophysical signal amplitudes; β are the corresponding spectral parameters; Δ_{bp} are the bandpass corrections; \mathbf{M}_{cj} denotes the bandpass-dependent component mixing matrix; s^{orb} is the orbital dipole; s^{fsl} are the far sidelobe corrections; s^{1hz} represents electronic 1 Hz spike corrections; n^{corr} is the correlated noise; and n^w is the white noise. For a more detailed explanation of the parameters, we refer the reader to [BeyondPlanck \(2022\)](#) and companion papers.

Data sets may also be included in the form of preprocessed pixelized sky maps, \mathbf{m}_v , in which case the above data model simplifies to

$$\mathbf{m}_{v,p} = g_v \mathbf{B}_{pp',v}^{\text{symm}} \sum_c \mathbf{M}_{cv}(\beta_{p'}, \Delta_{bp}^j) a_{p'}^c + n_{v,p}^w. \quad (2)$$

When producing \mathbf{m}_v , all time-dependent quantities (i.e., the far sidelobe, orbital dipole, 1 Hz spike, and correlated noise contributions) in Eq. (1) are explicitly subtracted from the TOD prior to mapmaking, leaving only sky stationary contributions in the final pixelized map. However, at this level only a very limited set of instrumental parameters may be accounted for per frequency,

Table 1. Frequency band summary for the BEYONDPLANCK intensity analysis.

SURVEY	DETECTOR LABEL	CENTRAL FREQUENCY [GHz]	BANDWIDTH [GHz]	BEAM SIZE (FWHM) [arcmin]	HEALPix RESOLUTION (N_{side})	AVERAGE RMS ^a [μK_{CMB} arcmin]	REFERENCE(S)
<i>Planck</i> LFI	30	28.4	5.7	32.4	512	179.3	Planck Collaboration IV (2016)
	44	44.1	8.8	27.1	512	213.8	
	70	70.1	14.0	13.3	1024	189.0	
<i>WMAP</i>	<i>Ka</i>	33.0	7.0	40	512	290.2	Bennett et al. (2013)
	<i>Q1</i>	40.6	8.3	31	512	400.6	
	<i>Q2</i>	40.6	8.3	31	512	380.0	
	<i>V1</i>	60.8	14.0	21	512	517.2	
	<i>V2</i>	60.8	14.0	21	512	446.6	
<i>Planck</i> HFI . . .	857	857	249	10.0 ^b	1024	5984	Planck Collaboration Int. LVII (2020)
Haslam	0.408	0 ^c	60 ^d	512	7.886 ^e	Haslam et al. (1982)

Notes.^(a) Average white noise rms without regularization noise.^(b) The native resolution of 857 GHz is 4.64' FWHM (Planck Collaboration VII 2016), but it is smoothed to 10' FWHM in this analysis.^(c) 408 MHz bandpass profile is assumed to be a δ function.^(d) The native resolution of 408 MHz is 56' FWHM (Haslam et al. 1982), but it is smoothed to 60' FWHM in this analysis.^(e) Unit is K_{CMB} arcmin.

namely an overall absolute calibration factor, g_v , an azimuthally symmetric beam, \mathbf{B}^{symm} , and white noise, $n_{v,p}^w$. This latter expression may be written on the following compact form,

$$\mathbf{m}_v = \mathbf{A}_v(\beta)\mathbf{a} + n_v^w, \quad (3)$$

where $\mathbf{A}_v(\beta)$ is an effective mixing matrix that takes into account both the frequency scaling of each component and beam convolution.

The goal of the Bayesian approach is to sample from the joint posterior distribution,

$$P(\mathbf{g}, \mathbf{n}_{\text{corr}}, \xi_n, \Delta_{\text{bp}}, \mathbf{a}, \beta, C_\ell \mid \mathbf{d}). \quad (4)$$

This is a large and complicated distribution, with many degeneracies. However, exploiting the Gibbs sampling algorithm (Geman & Geman 1984) we may factorize the sampling process into a finite set of simpler sampling steps. In this algorithm, samples from a multi-dimensional distribution are generated by sampling from all corresponding conditional distributions. The BEYONDPLANCK Gibbs chain may be written schematically as follows (BeyondPlanck 2022),

$$\mathbf{g} \leftarrow P(\mathbf{g} \mid \mathbf{d}, \xi_n, \Delta_{\text{bp}}, \mathbf{a}, \beta, C_\ell) \quad (5)$$

$$\mathbf{n}_{\text{corr}} \leftarrow P(\mathbf{n}_{\text{corr}} \mid \mathbf{d}, \mathbf{g}, \xi_n, \Delta_{\text{bp}}, \mathbf{a}, \beta, C_\ell) \quad (6)$$

$$\xi_n \leftarrow P(\xi_n \mid \mathbf{d}, \mathbf{g}, \mathbf{n}_{\text{corr}}, \Delta_{\text{bp}}, \mathbf{a}, \beta, C_\ell) \quad (7)$$

$$\Delta_{\text{bp}} \leftarrow P(\Delta_{\text{bp}} \mid \mathbf{d}, \mathbf{g}, \mathbf{n}_{\text{corr}}, \xi_n, \mathbf{a}, \beta, C_\ell) \quad (8)$$

$$\beta \leftarrow P(\beta \mid \mathbf{d}, \mathbf{g}, \mathbf{n}_{\text{corr}}, \xi_n, \Delta_{\text{bp}}, C_\ell) \quad (9)$$

$$\mathbf{a} \leftarrow P(\mathbf{a} \mid \mathbf{d}, \mathbf{g}, \mathbf{n}_{\text{corr}}, \xi_n, \Delta_{\text{bp}}, \beta, C_\ell) \quad (10)$$

$$C_\ell \leftarrow P(C_\ell \mid \mathbf{d}, \mathbf{g}, \mathbf{n}_{\text{corr}}, \xi_n, \Delta_{\text{bp}}, \mathbf{a}, \beta), \quad (11)$$

where \leftarrow indicates sampling from the distribution on the right hand side.

We note that not all of these steps follow the strict Gibbs approach of conditioning on all other parameters. Most notably for

us, this is the case in Eq. (9) for the spectral parameters sampler, $P(\beta \mid \mathbf{d}, \dots)$, which conditions on all parameters except for \mathbf{a} . Instead, we effectively sample \mathbf{a} and β jointly by exploiting the definition of a conditional distribution, as detailed in Sect. 3.1. The advantage of a joint sampling step is a significantly shorter Markov correlation length as compared to standard Gibbs sampling.

A very convenient property of Gibbs sampling is its modular nature, as the various parameters are sampled within each conditional distribution, but joint dependencies are explored through the iterative scheme. In this paper we are therefore only concerned with the sampling of two of the above steps, namely Eqs. (9) and (10). For all other sampling steps, we refer the interested reader to BeyondPlanck (2022) and references therein.

2.2. Data survey

Table 1 provides an overview of all frequency maps included in the intensity component of the BEYONDPLANCK analysis. For further details on and motivation for using this data set, we refer the interested reader to Sect. 5 in BeyondPlanck (2022), but note again that the main motivation underlying the BEYONDPLANCK analysis is not to derive a novel state-of-the-art intensity sky model, but rather to develop and demonstrate the Bayesian end-to-end analysis framework using *Planck* LFI as a worked case. Accordingly, to ensure that the main results are dominated by *Planck* LFI, all CMB-dominated *Planck* HFI bands and the *WMAP* *K*-band channel are excluded from the analysis; the data summarized in Table 1 represent a minimum set that is able to algebraically resolve all main foreground components relevant for *Planck* LFI. All sky maps, both LFI and others, are discretized using the HEALPix² (Górski et al. 2005) pixelization.

² <http://healpix.jpl.nasa.gov>

Table 2. Summary of main parametric signal models for the temperature analysis. The symbol “ \sim ” implies that the respective parameter has a prior as given by the right-hand side distribution; Uni denotes a uniform distribution within the indicated limits; $N(\mu, \sigma^2)$ denotes a (normal) Gaussian distribution with the indicated mean and variance; and a_i denotes the component amplitude of component i at the given reference frequency $\nu_{0,i}$, and s_i is spectral energy density, i.e., the observed signal at a given frequency ν .

COMPONENT	FREE PARAMETERS AND PRIORS	SPECTRAL ENERGY DENSITY, s_ν [μK_{RJ}]	ADDITIONAL INFORMATION
CMB	$a_{\text{CMB}} \sim \text{Uni}(-\infty, \infty)$	$x = \frac{h\nu}{k_{\text{B}} T_{\text{CMB}}}$ $g(\nu) = (\exp(x) - 1)^2 / (x^2 \exp(x))$ $s_{\text{CMB}} = a_{\text{CMB}} / g(\nu)$	$T_{\text{CMB}} = 2.7255 \text{ K}$
Relativistic CMB quadrupole	$a_{\text{quad}} = T_{\text{CMB}} \beta_{\text{sun}}^2 z^2$	$x = \frac{h\nu}{k_{\text{B}} T_{\text{CMB}}}$ $g(\nu) = (\exp(x) - 1)^2 / (x^2 \exp(x))$ $Q(\nu) = (x/2)(\exp(x) + 1) / (\exp(x) - 1)$ $s_{\text{quad}} = a_{\text{quad}} Q(\nu) / g(\nu)$	$T_{\text{CMB}} = 2.7255 \text{ K}$ $\beta_{\text{sun}} = 1.2343 \cdot 10^{-3}$ $\hat{\beta}_{\text{sun}} = (264.00^\circ, 48.24^\circ)$ $z = \hat{n} \cdot \hat{\beta}_{\text{sun}}$
Synchrotron	$a_{\text{s}} \sim \text{Uni}(-\infty, \infty)$ $\beta_{\text{s}} \sim N(-3.3 \pm 0.1),$ fullsky	$s_{\text{s}} = a_{\text{s}} \left(\frac{\nu}{\nu_{0,\text{s}}} \right)^{\beta_{\text{s}} + C \ln \nu / \nu_{0,\text{s}}}$	$\nu_{0,\text{s}} = 30 \text{ GHz}$ $C = 0, \text{ low signal-to-noise}$
Free-free	$a_{\text{ff}} \sim \text{Uni}(-\infty, \infty)$ $T_{\text{e}} = 7000 \text{ K},$ fullsky	$g_{\text{ff}} = \log \left\{ \exp \left[5.960 - \sqrt{3}/\pi \log \left(\nu_9 T_4^{-3/2} \right) \right] + e \right\}$ $s_{\text{ff}} = a_{\text{ff}} \left(\frac{\nu_{0,\text{ff}}}{\nu} \right)^2 \frac{g_{\text{ff}}(\nu)}{g_{\text{ff}}(\nu_{0,\text{ff}})}$	$\nu_{0,\text{ff}} = 40.0 \text{ GHz}$ $T_4 = T_{\text{e}} / 10^4$ $\nu_9 = \nu / (10^9 \text{ Hz})$ $e = \text{Euler's number}$
AME/ spinning dust	$a_{\text{ame}} \sim \text{Uni}(-\infty, \infty)$ $\nu_{\text{p}} \sim N(22 \pm 3 \text{ GHz}),$ fullsky	$s_{\text{ame}} = a_{\text{ame}} \left(\frac{\nu_{0,\text{ame}}}{\nu} \right)^2 \frac{f_{\text{ame}}(\nu \cdot \nu_{\text{p}0} / \nu_{\text{p}})}{f_{\text{ame}}(\nu_{0,\text{ame}} \cdot \nu_{\text{p}0} / \nu_{\text{p}})}$	$\nu_{0,\text{ame}} = 22.0 \text{ GHz}$ $\nu_{\text{p}0} = 30.0 \text{ GHz}$ $f_{\text{ame}}(\nu) = \text{External template}$
Thermal dust	$a_{\text{d}} \sim \text{Uni}(-\infty, \infty)$ $\beta_{\text{d}} \sim N(1.56 \pm 0.03),$ fullsky	$\gamma = \frac{h}{k_{\text{B}} T_{\text{d}}}$ $s_{\text{d}} = a_{\text{d}} \left(\frac{\nu}{\nu_{0,\text{d}}} \right)^{\beta_{\text{d}}+1} \frac{\exp(\gamma \nu_{0,\text{d}}) - 1}{\exp(\gamma \nu) - 1}$	$T_{\text{d}} = \text{NPIPE template}$ $\nu_{0,\text{d}} = 545 \text{ GHz}$
Radio sources	$a_{\text{src}} > 0$ $\alpha_{\text{src}} \sim N(-0.1 \pm 0.3)$	$s_{\text{src}} = U_{\text{mJy}}(\nu_{0,\text{src}}) a_{\text{src}} \left(\frac{\nu}{\nu_{0,\text{src}}} \right)^{\alpha_{\text{src}}-2}$	$\nu_{0,\text{src}} = 30 \text{ GHz}$ $U_{\text{mJy}}(\nu_{0,\text{src}}) = \text{Unit conversion factor}$

For all non-LFI bands, we adopt nominal bandpass profiles as recommended by the respective references. However, as an exception we adopt the simplified (and commonly used) delta function approximation for the Haslam 408 MHz channel. For LFI 30 GHz, we allow both an absolute bandpass shift for the full frequency band, and relative differences between individual detectors, while for the 44 and 70 GHz channels, we only allow relative detector shifts, but with no overall absolute shifts; see Svalheim et al. (2022a) for further details.

The noise is assumed to be uncorrelated and Gaussian for all channels except *Planck* LFI, with a spatially varying rms as defined by the number of hits per pixel. Again, the only exception is Haslam 408 MHz, which is nominally strongly signal-dominated

per pixel and dominated by systematic uncertainties, not statistical. In this case, we instead adopt a noise rms model that is the sum of an isotropic 0.8 K term (representing statistical uncertainties) and 1 % of the actual map itself, representing multiplicative uncertainties; this is the same approach as taken by Planck Collaboration X (2016). For LFI, correlated noise is accounted for on all angular scales through explicit time-domain sampling, as discussed by Ihle et al. (2022).

In the data model given in Eq. (1), only the orbital CMB dipole and the far sidelobes are modelled with the full asymmetric beams. For astrophysical component modelling, all beams are assumed to be azimuthally symmetric, with window functions, b_ℓ , provided individually by each experiment. No uncertainties

on these are propagated in the current analysis, but support for this will be added in future work. The Haslam 408 MHz and *Planck* DR4 (NPIPE) 857 GHz maps are smoothed from their native resolutions to 60' and 10', respectively. The latter is additionally re-pixelized to a HEALPix resolution of $N_{\text{side}} = 1024$ to reduce CPU and memory requirements; note that this channel still has a higher angular resolution than the 70 GHz LFI channel, which is the highest resolution channel of the main BEYONDPLANCK analysis.

An important and novel aspect of the current analysis is component-based monopole (or “zero-levels” or “offsets”) determination, as discussed in Sect. 3.2. Rather than attempting to set the monopoles for each channel before component separation, we impose physical priors on the monopole for each astrophysical component. This astrophysical model is then used to determine deterministically the zero-level for each frequency map. Reasonable priors may be defined for all components except synchrotron emission, and for this component we instead adopt explicit literature values. Specifically, we adopt a monopole value of $8.9 \pm 1.3 \text{ K}_{\text{CMB}}$ for synchrotron emission at 408 MHz, as estimated by Wehus et al. (2017), and we thereby neglect possible contributions from free-free emission to Haslam 408 MHz outside the very conservative Galactic mask employed by Wehus et al. (2017). We also apply the dipole corrections to the Haslam map derived by the same analysis.

Any additional pre-processing applied to the various maps is kept at a minimal level. Specifically, for *WMAP* we add the *WMAP* solar CMB dipole of $(d, l, b) = (3355 \mu\text{K}, 263.99^\circ, 48.26^\circ)$ to each map (Hinshaw et al. 2009), while for *Planck* 857 GHz, we apply a zodiacal light correction, following Planck Collaboration Int. LVII (2020).

No calibration corrections are applied to any non-LFI data sets, and we thus rely on the calibration of the original analyses for these channels. This is particularly important with respect to the *WMAP* channels, which have a non-negligible impact on the solar CMB dipole; consequently, the final BEYONDPLANCK solar dipole estimate represents a noise-weighted average between the *WMAP* and BEYONDPLANCK-based LFI estimates.

2.3. Astrophysical sky model

The dominant astrophysical foreground components in the *Planck* LFI frequencies are synchrotron, AME, free-free, thermal dust emission and compact radio sources. The modelling of each of these components is detailed in BeyondPlanck (2022) and here we will only review the relevant details for this paper. In Table 2 we summarize the models for each component in terms of free parameters, priors and SEDs. In addition, each diffuse component is modelled in terms of an amplitude sky map, \mathbf{a} , at a given reference frequency ν_0 in brightness temperature units. Scaling to arbitrary frequencies is performed through the SEDs, such that the actual observed signal at a given frequency ν may in general be written

$$s_{\text{RJ}}^i(\nu) = a_i \cdot f_i(\nu, \nu_{0,i}, \beta_i), \quad (12)$$

where i denote the specific component, $\nu_{0,i}$ is the reference frequency of the given component, β_i is a set of component-specific spectral parameters, and f_i is the SED. For diffuse components, a_i is defined in terms of spherical harmonic space with a maximum multipole ℓ_{max} defined for each component, depending on the signal-to-noise ratio and angular resolution of the data sets supporting that component. For instance, synchrotron and AME have lower ℓ_{max} 's than thermal dust and CMB. In addition, as

discussed in Sect. 3.4, we regularize the high- ℓ multipoles of each component with some smoothing prior, either derived from the known physical behaviour of the respective component (e.g., Planck Collaboration X 2016) or by a Gaussian smoothing operator. For compact sources, a_i represents simply the flux density in mJy, with the spectral index α also defined in mJy, and an explicit unit conversion factor, U_{mJy} , converts from flux density to brightness temperature units.

With this notation, the astrophysical sky model used for the current BEYONDPLANCK analysis may be written as follows,

$$s_{\text{RJ}} = (a_{\text{CMB}} + a_{\text{quad}}(\nu)) \frac{x^2 e^x}{(e^x - 1)^2} + \quad (13)$$

$$+ a_s \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s} + \quad (14)$$

$$+ a_{\text{ff}} \left(\frac{\nu_{0,\text{ff}}}{\nu} \right)^2 \frac{g_{\text{ff}}(\nu; T_e)}{g_{\text{ff}}(\nu_{0,\text{ff}}; T_e)} + \quad (15)$$

$$+ a_{\text{ame}} \left(\frac{\nu_{0,\text{ame}}}{\nu} \right)^2 \frac{f_{\text{ame}}(\nu \cdot \frac{\nu_p}{30.0 \text{ GHz}})}{f_{\text{ame}}(\nu_{0,\text{ame}} \cdot \frac{\nu_p}{30.0 \text{ GHz}})} + \quad (16)$$

$$+ a_d \left(\frac{\nu}{\nu_{0,d}} \right)^{\beta_d+1} \frac{e^{h\nu_{0,d}/k_B T_d} - 1}{e^{h\nu/k_B T_d} - 1} + \quad (17)$$

$$+ U_{\text{mJy}} \sum_{j=1}^{N_{\text{src}}} a_{j,\text{src}} \left(\frac{\nu}{\nu_{0,\text{src}}} \right)^{\alpha_{j,\text{src}}-2}, \quad (18)$$

where a_{CMB} and a_{quad} are given in thermodynamic temperature units (K_{CMB}), $a_{j,\text{src}}$ in flux density units (mJy), and all other amplitudes a_i are given in brightness temperature (K_{RJ}). The amplitude of component i is equal to that observed at a monochromatic frequency $\nu_{0,i}$. The sum in Eq. (18) runs over all compact sources brighter than some flux threshold as defined by an external source catalogue. In particular, we adopt the same catalogue as Planck Collaboration IV (2018), which is a hybrid of the AT20G (Murphy et al. 2010), GB6 (Gregory et al. 1996), NVSS (Condon et al. 1998) and PCCS2 (Planck Collaboration XXVI 2016) catalogs comprising a total of 12 192 individual sources.

When comparing the results from the above model with previous work, it is important to note that we fit a straight power-law for the synchrotron SED. This means that the effect of any potential negative curvature between 408 MHz and 30 GHz, as for instance assumed by Planck Collaboration X (2016), will instead be interpreted as a slightly steeper spectral index in the current analysis. This will be explicitly demonstrated in Sect. 4.

For further information regarding this model, and a brief discussion of each individual component, we refer the interested reader to BeyondPlanck (2022) and references therein. The main goal of the present paper is to establish efficient sampling algorithms for the amplitudes and spectral parameters in Eq. (13)–(18).

3. Commander extensions for efficient intensity-based component separation

The specific BEYONDPLANCK computer code implementation is called Commander3 (Galloway et al. 2022a), and this is a direct generalization of the code first introduced for CMB power spectrum estimation purposes by Eriksen et al. (2004) and later generalized to also account for astrophysical component separation by Eriksen et al. (2008); Seljebotn et al. (2014, 2019). It was one of four main component separation algorithms adopted by

the *Planck* collaboration (Planck Collaboration XII 2014; Planck Collaboration X 2016; Planck Collaboration IV 2018; Planck Collaboration Int. LVII 2020). Unless useful for context, we do not distinguish between the different code versions, but refer to all versions simply as *Commander*.

Here we describe the four algorithmic improvements we have made to *Commander*, as introduced in Sect. 1.

3.1. Joint amplitude and spectral parameter sampling

As summarized in Eqs. (5)–(11), the BEYONDPLANCK pipeline implements a Gibbs sampling chain iterating over all free parameters in the data model. While Gibbs sampling in general is a very powerful method for exploring complicated distributions, its main weakness is the inability to probe degenerate distributions. This problem is illustrated for a toy example in the left panel of Fig. 1: The black contours represents the 68 and 95 % confidence limits of a two-dimensional Gaussian distribution with a Pearson’s correlation coefficient of $\rho = 0.99$. The purple point indicates the starting position of a Markov chain, while the dashed, grey horizontal line indicates the baseline of the corresponding conditional distribution $P(\beta | a_{\text{init}})$, which itself is shown as an orange curve. Since the Gibbs sampling algorithm works by moving according to conditional distributions alone, the allowed step size in each iteration is very narrow compared to the full marginal distribution, shown as a blue distribution.

To illustrate the step size effect in the Gibbs sampler, let us consider a Gibbs move in Fig. 1, starting with β from the purple point. The relevant conditional distribution for β is shown as an orange curve along the horizontal dashed gray line passing through the purple point, and we now draw a random value from this distribution. This could for instance be the value indicated by the right-most vertical dashed gray line. According to the Gibbs sampling algorithm, we are now required to draw a sample from the corresponding conditional distribution for a , which is indicated by the orange distribution aligned with the vertical dashed line. One possible outcome from this, after completing one full Gibbs iteration, is the orange point. Now, because each conditional distribution is much narrower than the corresponding marginal distributions, the relative Gibbs step size is very short, and it takes a very long time to move from one side of the joint distribution to the opposite. The result is poor Markov chain mixing and a very long correlation length. As a real-world illustration of this, the orange points in the right panel of Fig. 1 show the 100 first steps of an actual Gibbs chain with this precise target distribution. We see that less than half of the distribution is actually explored, and many thousands of samples will be required in order to probe the full distribution with this algorithm.

This problem is directly relevant for modern Bayesian intensity-based CMB component separation. For experiments like *Planck* and *WMAP*, which have very high signal-to-noise ratios, there are strong degeneracies between the foreground amplitudes, the foreground spectral parameters, and map-level monopoles. Explicitly, if one assumes that all spectral parameters and monopoles are known, then the conditional amplitude uncertainty is very small. Conversely, if one assumes the amplitude and monopoles to be known, then the conditional spectral parameter uncertainties are small. However, when all parameters are unknown, the full uncertainties are significant.

For this reason, Gibbs sampling should usually be considered a last resort to handle an otherwise intractable distribution. If direct joint sampling methods are available, then those are usually more efficient. Fortunately, the Gibbs sampling method can be interleaved by any combination of conditional and joint

steps while still maintaining the requirement of detailed balance (Geman & Geman 1984), and the more steps that can be handled jointly, the more efficient the overall Gibbs chain will be. For the purposes of intensity-based CMB component separation, we therefore introduce a new special-purpose joint amplitude–spectral parameter step by exploiting the definition of a conditional distribution as follows,

$$P(\mathbf{a}, \beta | \mathbf{m}_\nu) = P(\beta | \mathbf{m}_\nu) P(\mathbf{a} | \mathbf{m}_\nu, \beta). \quad (19)$$

The first distribution on the right hand side is the marginal distribution of β with respect to \mathbf{a} , and the second distribution is the conditional distribution of \mathbf{a} with respect to β . This equation therefore implies that we may generate a joint sample by first drawing β from the its *marginal* distribution, and then sample \mathbf{a} from the corresponding *conditional* distribution,

$$\begin{aligned} \beta &\leftarrow P(\beta | \mathbf{m}_\nu) \\ \mathbf{a} &\leftarrow P(\mathbf{a} | \mathbf{m}_\nu, \beta); \end{aligned}$$

note the absence of \mathbf{a} in the first distribution.

We now need to derive sampling procedures for each of these two distributions. According to Bayes’ theorem,

$$P(\mathbf{a}, \beta | \mathbf{m}_\nu) = \frac{P(\mathbf{m}_\nu | \mathbf{a}, \beta) P(\mathbf{a}, \beta)}{P(\mathbf{m}_\nu)}, \quad (20)$$

where $P(\mathbf{m}_\nu)$ is just a normalization factor (often called the “evidence”), $P(\mathbf{a}, \beta)$ denotes optional priors, and the final factor is the likelihood function, $P(\mathbf{m}_\nu | \mathbf{a}, \beta) \equiv \mathcal{L}(\mathbf{a}, \beta)$. From the compact data model in Eq. (3), we note that

$$\mathbf{m}_\nu - \mathbf{A}_\nu(\beta) \mathbf{a} = \mathbf{n}_\nu^w, \quad (21)$$

and since \mathbf{n}_ν^w is assumed to be zero-mean and Gaussian distributed with known variance, we can immediately write down the following expression for the likelihood function,

$$\ln \mathcal{L}(\mathbf{a}, \beta) \propto -\frac{1}{2} \sum_\nu (\mathbf{m}_\nu - \mathbf{A}_\nu(\beta) \mathbf{a})^T \mathbf{N}_\nu^{-1} (\mathbf{m}_\nu - \mathbf{A}_\nu(\beta) \mathbf{a}), \quad (22)$$

where \mathbf{N}_ν is the noise covariance matrix of band ν , which is diagonal in the case of pure white noise. The priors are less well-defined, and are left to the user to determine. In the following, we will adopt Gaussian priors for spectral parameters, and for notational convenience assume no spatial amplitude priors, $\mathbf{S}^{-1} = 0$.

Let us first consider the marginal spectral parameter distribution, $P(\beta | \mathbf{m}_\nu)$. This is derived by integrating Eq. (22) with respect to \mathbf{a} , and this was done by Stompor et al. (2009) and Stivoli et al. (2010) as part of developing the MIRAMARE component separation code. The result reads

$$\begin{aligned} \ln \mathcal{L}_{\text{marg}}(\beta) &= \ln \int d\mathbf{a} \exp \left[-\frac{1}{2} (\mathbf{m} - \mathbf{A}\mathbf{a})^T \mathbf{N}^{-1} (\mathbf{m} - \mathbf{A}\mathbf{a}) \right] \\ &= \text{const} - \frac{1}{2} (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{m})^T (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{m}) \\ &\quad + \frac{1}{2} \ln |\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}|, \end{aligned} \quad (23)$$

where all terms should be interpreted as sums over frequencies. The same authors also introduce a so-called “ridge likelihood”, in which one does not marginalize over \mathbf{a} , but rather sets \mathbf{a} equal to its maximum likelihood value for a given value of β . This may also be analytically evaluated, and is in fact identical to the above expression, except without the last determinant term. We implement support for both options in our codes. Note that this

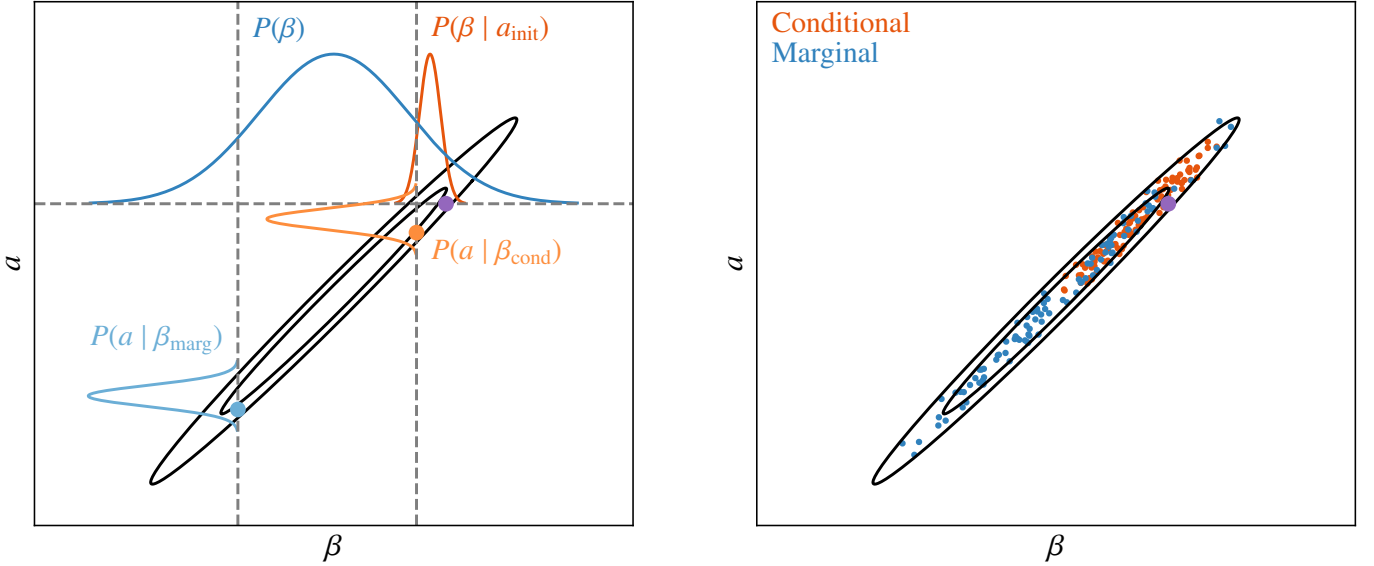


Fig. 1. Illustration of conditional (pure Gibbs; orange) and marginal sampling (blue) algorithms for a highly correlated (Pearson’s correlation coefficient of $\rho = 0.99$) two-dimensional Gaussian distribution (black contours). The initial position $(\beta_{\text{init}}, a_{\text{init}})$ is indicated by a purple dot. *Left:* Comparison of un-normalized conditional $P(\beta | a_{\text{init}})$ distribution evaluated at the initial position and the corresponding marginal $P(\beta)$ distribution; note that the latter is much wider than the former. Assuming β samples drawn at the 10th percentile, the graphs along the vertical lines represent un-normalized conditional distributions $P(a | \beta)$ evaluated at the β values drawn with the conditional (orange) and marginal (blue) distributions of β ; note that the marginal sampling case results in a much longer step length between the initial and final sample values. *Right:* Samples of a standard Gibbs sampling chain (orange) using conditional sampling for both a and β , and a sampling chain (blue) using marginal sampling for β and conditional sampling for a . Both cases show the first 100 samples initialized from the purple point.

expression requires all data to be defined at the same angular resolution, and so all data must be smoothed to a common resolution before evaluating Eq. (23).

The second required distribution is $P(a | m_v, \beta)$. This is a simple multi-variate Gaussian distribution in a , for which there are efficient samplers readily available; see, e.g., Appendix A of [BeyondPlanck \(2022\)](#) for details. One particularly efficient sampling equation reads

$$\left(\sum_v A_v^t N_v^{-1} A_v \right) a = \sum_v A_v^t N_v^{-1} m_v + \sum_v A_v^t N_v^{-1/2} \eta, \quad (24)$$

where η is a vector of random Gaussian $N(0, 1)$ variates. This equation may be solved efficiently using preconditioned Conjugate Gradient methods ([Shewchuk 1994](#)), as discussed by [Seljebotn et al. \(2019\)](#).

These equations are integrated into the main BEYONDPLANCK Gibbs sampling loop as follows: First, we run a short (typically a few hundred steps) standard Metropolis sampler (see Appendix A of [BeyondPlanck 2022](#)) for each spectral parameter, using the product of Eq. (23) and any desired priors to define the accept rate, i.e., the relative number of Metropolis proposals being accepted (which should preferably stay between 0.3 and 0.7 for an efficient sampler). All data are smoothed to a common angular and pixel resolution before evaluating the expression. Immediately following the last Metropolis step, we draw one sample from $P(a | m_v, \beta)$ using Eq. (24); it is critically important that no other parameters are updated between β and a , as the previous value of a is completely inconsistent with the new β value, which is drawn marginally with respect to a .

Returning to Fig. 1, the improvement achieved by this joint two-step sampler is illustrated as blue distributions and points. Starting with the left panel, the fundamental difference between the joint and Gibbs samplers is that the first step in the β -direction is drawn from the full marginal distribution (horizontal

blue distribution) instead of the conditional distribution (horizontal orange distribution). This is much wider, and covers by construction the full width of the underlying target distribution. One single proposal may therefore move from one side of the distribution to the other, and there is no memory of the previous parameter state. To obtain a valid sample, it is however critically important to draw a corresponding sample from the appropriate conditional amplitude distribution (vertical blue distribution) immediately after the marginal move. Correspondingly, the blue points in the right panel shows 100 samples drawn with the joint sampler. In this case, they cover the full distribution very efficiently.

Figure 2 shows a similar comparison for AME ν_p for a test chain that only explores the AME parameters in the BEYONDPLANCK data model with both methods. Also in this case, we see that the marginal sampler explores the full range much more efficiently than the conditional sampler.

3.2. Component-based monopole determination

Next, we consider the problem of monopole determination for CMB experiments, which has long been one of the main challenges for parametric component separation methods (see, e.g., [Planck Collaboration X 2016](#); [Wehus et al. 2017](#), and references therein). The problem stems from the following challenge: For traditional CMB experiments and maximum likelihood map-making methods, there are no data-driven constraints on the monopoles in the derived frequency sky maps. For example, *WMAP* is explicitly differential in nature, measuring only differences between pairs of points, and therefore cannot by construction constrain the zero-level. For *Planck*, the high level of $1/f$ noise prohibits any useful constraints on the zero-levels. An important exception to this is *COBE/FIRAS* ([Mather et al. 1994](#)), which is absolutely calibrated; but its mK-level uncertainties are




Fig. 2. Comparison of AME ν_p chains derived using the conditional (orange; Eq. 22) and marginal (blue; Eq. 23) samplers discussed in the text. For this illustration, all other parameters than a_{AME} and ν_p are fixed.

still orders of magnitude too large to be useful for modern CMB component separation purposes.

For this reason, several indirect methods have been established to determine the frequency map monopoles based on the morphology of the maps themselves. Four examples are mean subtraction in a small region (Planck Collaboration V 2014); fitting a plane-parallel co-secant model (Bennett et al. 2003; Planck Collaboration II 2016); imposing foreground SED consistency between neighboring frequencies (Wehus et al. 2017); and cross-correlation with external data sets with known zero-levels (Planck Collaboration VIII 2016). However, all of these methods have in common the fact that they operate on *frequency maps*, and aim to determine the zero-level at a given frequency channel, before feeding these into traditional component separation algorithms. In this paper, we make the observation that it is in fact much easier to determine the monopoles of the *component amplitude maps*, and then use these to deterministically set the frequency map monopoles through the resulting sky model. The frequency map zero-levels have thus no independent impact on any higher-order analyses (most notably the spectral parameters), but simply adjust themselves to whatever the model dictates at any given moment.

An immediate question is then; how do we in fact determine the component monopoles? This must be done on a case-by-case basis, applying the most natural prior for each component. We note that any true monopole signal in the components that do not agree with the chosen priors will end up in the frequency map monopoles. Starting with the CMB case, this can either be set to zero or 2.7255 ± 0.0006 K (Fixsen 2009), depending on whether one wants sky maps without or with the CMB monopole. In practice, we additionally account for sub-optimal foreground modelling by applying a mask. For the current analysis, we derive the CMB monopole mask from a set of smoothed component amplitude maps, namely by thresholding the sum of synchrotron, AME, free-free and thermal dust emission, all smoothed to 10° FWHM. In addition, we mask out radio sources and any pixel with a reduced normalized χ^2 higher than 5σ . The resulting mask is shown in Fig. 3, and has an accepted sky fraction of $f_{\text{sky}} = 0.64$. The monopole of the CMB component map is set to zero (or 2.7255 K) outside this mask, while simultaneously fitting for (but not modifying) the dipole component.




Fig. 3. Masks used for signal amplitude zero-level and band monopole sampling: (top) Frequency-band monopole and CMB amplitude zero-level, and (bottom) free-free amplitude zero-level, with a free-free amplitude sample at an angular resolution of $30'$ FWHM plotted underneath.

Regarding the free-free component, Planck Collaboration X (2016) found that the measured emission is strongly noise-dominated over large areas of the sky, with no detectable amplitude. Also in this case, we therefore set the monopole to have zero mean outside a conservative mask. In this case, the mask is derived from the free-free amplitude map itself, evaluated at 30 GHz and smoothed to 10° FWHM, and truncated at $5\mu\text{K}_{\text{RJ}}$. In addition, we exclude areas where the other foreground signals are high to account for signal-leakage into the free-free component, similar to the CMB mask, but thresholding the sum of all other components evaluated at 44 GHz. The resulting mask is shown in the bottom panel of Fig. 3, and accepts $f_{\text{sky}} = 0.50$. We do note that any true unmasked free-free signal is by definition positive, and this can bias the monopole also in the noise-dominated regime. Future work should aim to correct for this bias by directly estimating the residual free-free monopole in the unmasked region. We do note, however, that this bias will decrease as more high sensitivity data become available, as more and more of the free-free signal may be masked directly.

In contrast, synchrotron emission as observed by the Haslam 408 MHz map is highly diffuse on the sky, and there are no regions on the sky that can be assumed to be approximately clean of synchrotron emission, i.e., no synchrotron signal. In this case, the best estimates of the synchrotron amplitude are those already estimated for the Haslam 408 MHz map itself. In this paper, we adopt the zero-level correction of $8.9 \pm 1.3 \text{ K}_{\text{CMB}}$ derived by Wehus et al. (2017). We marginalize over the uncertainty by drawing a random offset correction in every Gibbs iteration, as defined by a Gaussian distribution with the quoted mean and standard deviation.




Fig. 4. Cross-correlations between the H I column density N_{HI} and (top) the thermal dust amplitude and (bottom) the AME amplitude. The thermal dust cross-correlation is evaluated at HEALPix resolution $N_{\text{side}} = 64$ and a common angular resolution of $60'$ FWHM, while the AME cross-correlation is evaluated at HEALPix resolution $N_{\text{side}} = 16$ and a common angular resolution of $10'$ FWHM. The lines represent best-fit lines for pixels with H I column densities less than $2 \cdot 10^{20} \text{ cm}^{-2}$ (blue) or $4 \cdot 10^{20} \text{ cm}^{-2}$ (orange). The green curve is the best-fit second degree polynomial to pixels with H I column densities less than $4 \cdot 10^{20} \text{ cm}^{-2}$. The contour lines are plotted at 0.001, 0.01, 0.05, and $0.1 N_{\text{pix}} / (\mu\text{K}_{\text{RJ}} 10^{20} \text{ cm}^{-2} N_{\text{side}}^2)$; where only the lower three contour line values are plotted for AME. The contours have been smoothed for visualization.

For thermal dust emission, we adopt essentially the same approach as the *Planck* HFI DPC (e.g., see [Planck Collaboration III 2020](#)), setting the zero-level through cross-correlation with H I column density observations (e.g., see [Lenz et al. 2017](#)), although with a few minor variations. Firstly, we apply this method to the thermal dust component map, as opposed to individual frequency maps. Secondly, for *Planck* DR4 ([Planck Collaboration Int. LVII 2020](#)) the HFI 545 GHz zero-level was set through a linear fit for pixels with $N_{\text{HI}} < 4 \cdot 10^{20} \text{ cm}^{-2}$. However, as the 545- N_{HI} scatter plot appeared to be non-linear around values of $N_{\text{HI}} = 1.5 \cdot 10^{20} \text{ cm}^{-2}$, they also performed a second degree fit. We show in Fig. 4 a similar scatter plot between N_{HI} and the thermal dust amplitude from one of our Gibbs samples, where we perform both a first and second degree polynomial fit to the plot at $N_{\text{HI}} < 4 \cdot 10^{20} \text{ cm}^{-2}$. Furthermore, we also perform a

linear fit at $N_{\text{HI}} < 2 \cdot 10^{20} \text{ cm}^{-2}$, and we see that the intersection of the linear fit with a lower threshold is close to the intersection of the second degree fit. This raises the question of the uncertainty of the threshold value for the linear fit, as [Planck Collaboration XXIV \(2011\)](#) found a good correlation up to at least $N_{\text{HI}} = 2 \cdot 10^{20} \text{ cm}^{-2}$. We therefore implement this cross-correlation method as a prior on the thermal dust zero-level using a range of thresholds, and for each Gibbs iteration we perform linear fits of N_{HI} and the thermal dust amplitude with N_{HI} threshold values ranging from 1.5 to 4 [10^{20} cm^{-2}] with increments of 0.5. Then we draw the intersection value from a Gaussian distribution given the mean and variance of the linear fits, which is subtracted from the dust amplitude. This way, the uncertainty of the thermal dust amplitude zero-level also propagates through the pipeline.

The zero-level of the AME component is determined using the same procedure, noting that [Planck Collaboration X \(2016\)](#) demonstrated a very tight spatial correlation between AME and thermal dust emission on large angular scales. The only difference with respect to the thermal dust procedure is that we smooth all maps to a common angular resolution of $10'$ FWHM and a HEALPix resolution of $N_{\text{side}} = 16$, and adopt thresholds of 2 to 4 [10^{20} cm^{-2}], with increments of 0.5; both the smoothing and the lower resolution are imposed to reduce the impact of instrumental noise. A scatter plot between AME and N_{HI} is shown in Fig. 4 for one arbitrary Gibbs sample.

With the introduction of component-based monopole priors, all frequency-band monopoles become free parameters, and can be deterministically fitted. Explicitly, for each frequency channel we first subtract the predicted sky model as defined by Eq. (2), and then fit and subtract the residual monopole outside some mask. The mask should be defined such that it excludes areas on the sky prone to foreground mis-modelling, hence we adopt the same mask as we do for the CMB monopole prior, shown in the top panel of Fig. 3. We note that this monopole adjustment needs to be done immediately after any change in any of the component maps, in order to not break the Gibbs chain, fully analogously to the immediate amplitude update that must follow any marginal spectral parameter move discussed in the previous section.

3.3. Joint spectral parameter and frequency-band monopole sampling

Returning to the AME-H I cross-correlation plot in Fig. 4, we notice that the zero-level is associated with a large statistical uncertainty. When sampling the AME peak frequency, ν_p , this uncertain monopole is also directly affected by the resulting SED changes, and corresponding monopole offsets are induced at all frequencies. If ν_p is sampled conditionally with respect to the band monopoles, these will therefore tend to pull ν_p towards the old value, and thereby increasing the overall Markov chain correlation length.

This inefficiency may be alleviated by exploiting the new component-based monopole sampler described in Sect. 3.2. Since all frequency-band monopoles are now deterministically defined by the sky model, these can be adjusted jointly whenever that is modified. We therefore implement internal estimation of frequency band monopoles during the spectral parameter sampling algorithm, such that for each proposed spectral parameter value, we estimate a new band monopole value conditioned on the input amplitude and the proposed parameter value. The frequency band monopoles are updated together with the final parameter at the end of the sampling.




Fig. 5. Comparison of AME ν_p χ^2 distributions with (green) and without (orange) monopole marginalization. These distributions are evaluated using the marginal spectral parameter likelihood given in Eq. (23), but the same qualitative behaviour holds irrespective of which spectral parameter distribution is used (conditional, ridge, or marginal): The distribution becomes significantly wider when marginalizing over band monopoles.

To illustrate the usefulness of this combined sampling step, we generate an idealized simulation that includes only AME signal and noise at each frequency channel, with an angular resolution of 10° FWHM and a HEALPix resolution of $N_{\text{side}} = 16$. We choose an input peak frequency of $\nu_p = 26$ GHz, and adopt a zero-level from HI column density cross-correlation to $2 \mu\text{K}_{\text{RJ}}$. Figure 5 shows the resulting log-likelihood (or χ^2) distributions as evaluated from the marginal definition in Eq. (23), for cases both with (green curve) and without (orange curve) marginalizing over the frequency band monopoles. We see that by marginalizing over the band monopoles the log-likelihood function widens by a factor of 1.5–2. This translates into correspondingly longer Metropolis step sizes in the spectral index sampling steps, and thereby faster exploration of the full posterior distribution. The higher conditional signal-to-noise ratio a given component has, the more important this effect will be.

3.4. Breaking small-scale degeneracies through spatial priors

The final algorithmic improvement presented in this paper is the introduction of informative spatial priors for foreground components, either in the form of purely algorithmic smoothing power spectrum priors or as actual informative Gaussian priors with a non-zero mean. The first of these has already been used in the latest *Planck* analyses (Planck Collaboration IV 2018; Planck Collaboration Int. LVII 2020), but in the following we generalize the approach to non-zero cases, and we show systematically how different choices affect the final results.

Mathematically speaking, the only difference between an informative prior and a smoothing prior is whether a pre-existing mean map is assumed for the astrophysical component in question (in which case the prior is called “informative”), or whether the prior mean is assumed to be zero. Practically speaking, however, there is also an important difference between the prior variances in the two cases, since for informative priors the variance quantifies the allowed level of fluctuations *around the mean map*,

while for smoothing priors it quantifies the allowed level of fluctuations *around zero*. Thus, for informative priors a prior variance of zero is fully acceptable, in which case the output component map will be identical to the prior mean, while for a smoothing prior the variance should be larger than the actual component fluctuations in order to avoid over-smoothing.

Let us start by revisiting the sampling equation for the component amplitude maps, as defined by Eq. (24). This equation provides a sample from a posterior defined only by the likelihood itself. If we additionally want to impose a Gaussian prior on the amplitudes, as defined by a multi-variate Gaussian distribution, $N(\mu, S)$, then this is generalized to

$$(S^{-1} + A^T N^{-1} A) \mathbf{a} = A^T N^{-1} \mathbf{m}_\nu + S^{-1} \mu + A^T N^{-1/2} \boldsymbol{\eta}_1 + S^{-1/2} \boldsymbol{\eta}_2; \quad (25)$$

see Appendix A in BeyondPlanck (2022) for an explicit derivation. In this expression, μ has the same dimension as \mathbf{a} , and represents the prior mean for \mathbf{a} , while S is an associated prior covariance matrix that defines the “strength” of the prior, fully analogous to the usual standard deviation, σ , of a Gaussian univariate prior. Thus, if $S = 0$, the final solution for \mathbf{a} will be identical to μ , while if $S \rightarrow \infty$ (or, equivalently, $S^{-1} = 0$), then the prior term vanishes, and one is left with the original likelihood expression in Eq. (24). For reference, we note that the previous *Planck* analyses (Planck Collaboration IV 2018; Planck Collaboration Int. LVII 2020) set $\mu = 0$ in this equation, and only used S to impose smoothness on \mathbf{a} .

Computationally speaking, introducing informative priors with non-zero means in Eq. (25) represents no additional algorithmic complications compared to the prior-free case: The equation is in both cases solved using the same preconditioned Conjugate Gradient implementation. If anything, the equations are actually a bit easier to solve with informative priors, as they reduce degeneracies between different parameters, and thereby reduce the condition number of the coefficient matrix on the left-hand side of the equation. As pointed out by Seljebotn et al. (2019), from an algorithmic point-of-view a prior may simply be considered to be a new independent data set with sensitivity only to the component in question, and it therefore provides orthogonal information with respect to the likelihood term contributions. Rather, the main challenge regarding priors is how to define them in a useful and controlled manner that does not significantly bias or contaminate the final posterior distribution, and this must be assessed on a case-by-case basis.

In the current BEYONDPLANCK analysis, we follow Planck Collaboration IV (2018), and define S in harmonic space, giving different prior weights to different angular scales. Explicitly, each component map is defined in terms spherical harmonics,

$$\mathbf{a} = \sum_{\ell, m} a_{\ell m} Y_{\ell m}, \quad (26)$$

and we define the prior covariance matrix as

$$S = S_{\ell m, \ell' m'} \equiv \langle a_{\ell m} a_{\ell' m'}^* \rangle = P_\ell \delta_{\ell \ell'} \delta_{m m'}, \quad (27)$$

where P_ℓ is an angular prior power spectrum for \mathbf{a} , which, again, is fully analogous to the standard deviation of a Gaussian prior, but now defined per angular multipole.

3.4.1. Algorithmic smoothing priors

Starting with the algorithmic smoothing priors adopted by Planck 2018, we note that if we define P_ℓ^a to be an estimate of




Fig. 6. Comparison of AME (*left*) and free-free (*right*) amplitude maps derived without any spatial priors in a $20 \times 20^\circ$ field centered on the Galactic South Pole, $(l, b) = (0^\circ, -90^\circ)$, at an angular resolution of $60'$ FWHM. Note the striking anti-correlation between the two maps.

the true angular power spectrum of \mathbf{a} , then the following power spectrum prior,

$$P_\ell = P_\ell^a e^{-\ell(\ell+1)\sigma^2}, \quad (28)$$

simply represents a Gaussian smoothing prior of \mathbf{a} with a smoothing kernel width equal to σ , which often is defined in terms of $\theta_{\text{FWHM}} = \sqrt{8 \ln 2} \sigma$. We call this an algorithmic smoothing prior, as it explicitly pushes the solution to be smooth on small angular scales. We also note that P_ℓ^a does not have to be an accurate estimate of the true component power spectrum, but it should in general be *larger* than the true spectrum, in order to prevent the prior from being overly constraining. This is again fully analogous to choosing a prior width that is wider than the expected target distribution in standard uni-variate analysis problems. Thus, simply setting P_ℓ^a to a constant that is a few times larger than the expected spectrum is usually a perfectly good choice.

It is important to note that Eq. (28) is indeed just a prior, and not a deterministic post-processing smoothing operator. This has both advantages and drawbacks that are important to be aware of when using the products from the analysis. The main advantage is that signal-dominated localized objects (for instance point sources) are not excessively smoothed when applying the smoothing as a prior. The main disadvantage is that the effective angular resolution of the amplitude map becomes spatially varying, and depending on the local signal-to-noise ratio in a given pixel; if the signal-to-noise ratio is high, the angular resolution will be determined by the resolution of the data, while if the signal-to-noise ratio is low, it is given by θ_{FWHM} . This is similar to the GNILC method (Remazeilles et al. 2011), which also implements signal-to-noise dependent angular resolution. The main difference between the two methods is that while GNILC requires regions of different resolutions to be pre-defined, the current approach automatically and dynamically adopts the resolution while solving Eq. (25).

In the current BEYONDPLANCK analysis, for which the main scientific target is CMB power spectrum and cosmological parameter estimation, we do not impose any priors on the CMB component amplitudes, but we do impose a Gaussian smoothing prior for thermal dust emission in intensity with $\theta = 5'$ FWHM and $P_\ell^a = 10^7 \mu\text{K}_{\text{RJ}}^2$ at 545 GHz. For users who are primarily interested in the angular power spectrum of thermal dust emission, it would be more useful to instead impose a ΛCDM spectrum




Fig. 7. Difference maps between the derived amplitude and prior maps for AME for three different 857 GHz scaling factors. From top to bottom, the three panels show scaling factors of $\alpha = 2 \cdot 10^{-5}$, $\alpha = 3 \cdot 10^{-5}$ and $\alpha = 4 \cdot 10^{-5}$.

on the CMB amplitude map, and no priors on the thermal dust spectrum. Then the resulting thermal dust power spectrum would be an unbiased estimator; with the current analysis, the thermal dust spectrum will be biased low on small angular scales due to the smoothing prior. We also impose a Gaussian smoothing prior on synchrotron emission in intensity, with $\theta = 60'$ FWHM and $P_\ell^a = 3 \cdot 10^{14} \mu\text{K}_{\text{RJ}}^2$ at 408 MHz.

3.4.2. Informative Gaussian spatial priors

For AME and free-free emission, we adopt informative priors with $\mu \neq 0$. The reason for this is simply that the limited data combination considered in the current BEYONDPLANCK analysis (see Table 1) is inadequate for constraining all of AME, free-free, and CMB separately without additional information; when future observations from, e.g., C-BASS (King et al. 2010) and QUIJOTE (Génova-Santos et al. 2015) become publicly available and integrated in the analysis, this will hopefully no longer be necessary. As an illustration of the problem, Fig. 6 compares the AME and free-free amplitude maps derived without any pri-




Fig. 8. (*Top:*) AME amplitude prior map, derived by scaling the *Planck* DR4 857 GHz by $\alpha = 3 \cdot 10^{-5}$ and smoothing to $10'$ FWHM. (*Bottom:*) Derived AME amplitude maps for four different spatial prior combinations, $\hat{D}_{\text{AME}}(\ell) = q (\ell/\ell_0)^\beta$. Rows show results for $q = 10^2 \mu\text{K}_{\text{RJ}}^2$ and $10^4 \mu\text{K}_{\text{RJ}}^2$, respectively, while columns show results for $\beta = -1$ and 1 .

ors near the Galactic South Pole; even at a visual level, these two maps are nearly perfectly anti-correlated, with no true constraining power on their own. This also makes other components (most importantly the CMB) susceptible to small systematic residual mismatches between the AME and free-free components, and it significantly increases the CMB noise.

Starting with the AME case, we first note that the prior-free *Planck* 2015 analysis (Planck Collaboration X 2016) found a very strong spatial correlation between their AME and thermal dust component maps at an angular resolution of 1° FWHM. In general, a high degree of correlation between these components is expected from current theoretical AME models (e.g., Erickson 1957; Draine & Lazarian 1998; Ali-Haïmoud 2010; Silsbee et al. 2011; Hensley & Draine 2020), although the specific correlation coefficient depends on model details. Based on these observa-

tions, we adopt the *Planck* DR4 857 GHz map³ (Planck Collaboration Int. LVII 2020) as a spatial mean template for AME, i.e., μ in Eq. (25). However, before it can be inserted into Eq. (28), it must be adjusted in amplitude to account for the mean SED difference between thermal dust emission at 857 GHz and AME at 22 GHz. To do this, we solve Eq. (28) using the 857 GHz map scaled by a range of values between $\alpha = 2 \cdot 10^{-5}$ and $4 \cdot 10^{-5}$ as the AME prior. We then take the difference between the derived amplitude map and the input prior, and adopt the scaling factor for which the difference is smallest as our default prior. Example difference maps are shown in Fig. 7, and we adopt $\alpha = 3 \cdot 10^{-5}$ as our default prior.

³ The *Planck* DR4 857 GHz map is corrected for both zodiacal light emission and a zero-level of $-0.657 K_{\text{CMB}}$, and smoothed to $10'$ FWHM, before adopted as an AME prior.

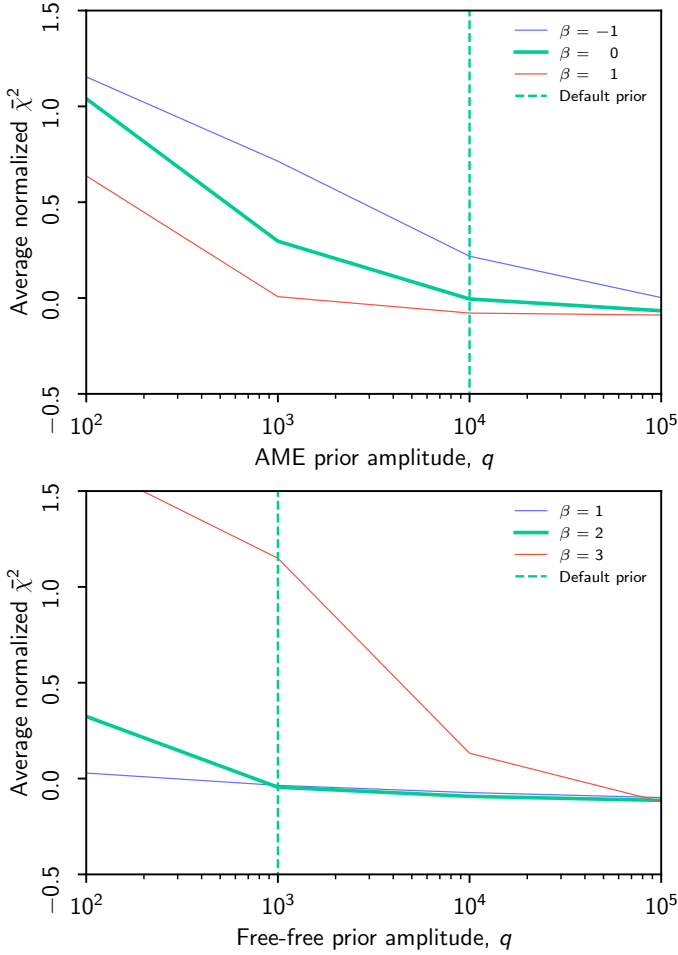


Fig. 9. Average $\bar{\chi}^2$ per $N_{\text{side}} = 16$ pixel as a function of AME (*top panel*) and free-free (*bottom panel*) prior amplitude, q , where $\bar{\chi}^2 \equiv (\chi^2 - n_{\text{dof}})/\sqrt{2n_{\text{dof}}}$ and $n_{\text{dof}} = 15400$. Coloured solid lines show results for different tilt parameters, β , and the intersection between the thick green curve and dashed line indicates the prior combination adopted for the main BEYONDPLANCK analysis.

The final step is to define the strength of this prior, as given by P_ℓ . Ideally, we want the prior to be stronger (i.e., S to be smaller) for the noise-dominated small angular scales, and looser for the signal-dominated large angular scales. To quantify these considerations, we define a power-law prior power spectrum for AME,

$$P_\ell^{\text{AME}} = q \left(\frac{\ell}{\ell_0} \right)^\beta, \quad (29)$$

where q is an overall amplitude at a pivot multipole, ℓ_0 , and β is a tilt parameter. A negative (positive) β results in a stronger (weaker) prior at high multipoles, and a smaller (higher) amplitude q gives a stronger (weaker) prior on all angular scales.

Figure 8 compares the resulting AME amplitude maps for various choices of q and β (bottom panels) with the prior mean map (top panel). Here we see that a high value of $q = 10^4 \mu\text{K}_{\text{RJ}}^2$ leads to very similar solutions for $\beta = -1$ and $\beta = 1$, indicating that the prior is largely irrelevant, and the solution is data dominated. We also note that there is substantial instrumental noise at high latitudes. For $q = 10^2 \mu\text{K}_{\text{RJ}}^2$, the maps are notably smoother at high latitudes, but we also see clear smoothing artefacts near the Galactic plane, corresponding to harmonic space

ringing from the Galactic plane. A value of $\beta = -1$ leads to sharper edges than $\beta = 1$.

To actually determine the parameters used for the final prior, we solve Eq. (28) over a grid in q and β , and evaluate

$$\chi^2(q, \beta) \equiv \sum_{\nu, p} \left(\frac{m_{\nu, p} - s_{\nu, p}(q, \beta)}{\sigma_{\nu, p}} \right)^2 \quad (30)$$

for each configuration, where s is the derived sky model in each case. The results from this evaluation are shown in the top panel of Fig. 9 in terms of the normalized reduced $\bar{\chi}^2 \equiv (\chi^2 - n_{\text{dof}})/(2n_{\text{dof}})$ where n_{dof} is the number of degrees of freedom; for a perfect model fit and $n_{\text{dof}} \gg 1$, this quantity should be distributed approximately as a Gaussian distribution with zero mean and unit standard deviation. For $q = 10^2 \mu\text{K}_{\text{RJ}}^2$ and $\beta = 0$, we see that $\bar{\chi}^2 \approx 1$, which indicates a clear excess residual. However, for $q = 10^4 \mu\text{K}_{\text{RJ}}^2$ and $\beta = 0$, this excess is greatly diminished, while there still is some effect from the prior. In the following, we adopt this latter combination as our default AME prior.

For free-free emission, there are no corresponding full-sky independent spatial templates available in the literature. Observations of $\text{H}\alpha$ (Finkbeiner 2003) or radio recombination lines (RRL; Alves et al. 2015) might serve useful roles, but both are associated with significant short-comings for the purposes of the current analysis: The $\text{H}\alpha$ observations lack most of the Galactic plane signal due to dust absorption, while the RRL observations only cover a part of the Galactic plane. For now, the best available full-sky free-free tracer is in fact the *Planck* 2015 free-free map (Planck Collaboration X 2016), which was based on the same data set as studied in the current paper—but additionally, and critically, also the *Planck* HFI observations. We therefore adopt this map as a spatial prior in the current analysis, while recognizing that this is strictly speaking not admissible in the Bayesian framework; some data (i.e., LFI, *WMAP*, and Haslam) are used twice to constrain free-free emission, and the resulting uncertainties will therefore be under-estimated. In practice, this solution is a way of integrating HFI observations into the analysis without directly affecting the CMB component. A critical goal for near-future work is to integrate HFI observations directly into the analysis in the form of frequency maps, and at that point this informative free-free prior will be removed.

We adopt the same parametric function for free-free emission as for AME, defined by Eq. (29), and adjust the free parameters in the same way. The results from this optimization are shown in the bottom panel of Fig. 9, and we adopt $q = 10^3 \mu\text{K}_{\text{RJ}}^2$ and $\beta = 2$ in this case. A comparison of different prior choices with the actual input prior map is shown in Fig. 10.

The only algorithmic difference with respect to AME is that we additionally impose a Gaussian smoothing for free-free emission, as per Eq. (28) with $\theta = 30'$ FWHM. This is done to account for the fact that the distribution of free-free emission is highly localized on the sky, and therefore requires a high maximum multipole moment to capture all significant structures; the additional Gaussian smoothing ensures that no ringing emerges from the high- ℓ truncation.

4. BEYONDPLANCK analysis and posterior distributions

We now turn our attention to the actual BEYONDPLANCK analysis and intensity component posterior distributions derived from the data combination discussed in Sect. 2.2. We once again note that the goal of the BEYONDPLANCK project is not to derive a new CMB




Fig. 10. (Top:) Free-free amplitude prior map, adopted from the *Planck* 2015 analysis which includes HFI observations. (Bottom:) Derived free-free amplitude map for four different spatial prior combinations, $\hat{D}_{\text{ff}}(\ell) = q (\ell/\ell_0)^\beta$. Rows show results for $q = 10^2 \mu\text{K}_{\text{RJ}}^2$ and $10^4 \mu\text{K}_{\text{RJ}}^2$, respectively, while columns show results for $\beta = 1$ and 3.

state-of-the-art astrophysical component model (given that critical data sets such as *Planck* HFI and *WMAP* *K*-band are not included), but rather to lay the algorithmic groundwork for a statistically robust community-wide sky model, as will be implemented through the Open Science COSMOGLOBE⁴ community effort. As far as Bayesian intensity sky modelling is concerned, *Planck* Collaboration X (2016) still represents the state-of-the-art.

4.1. Spectral parameter prior tuning

Before presenting the BEYONDPLANCK Markov chains and posterior distributions, there is still one task that must be completed before the algorithm described in Sect. 2 is complete, namely

⁴ <http://cosmoglobe.uio.no>

finalizing the informative spectral parameter priors. With the reduced number of data sets included in this work, we have reduced constraining power when sampling spectral parameters, and strong priors are required for most spectral SED parameters. With this in mind, we will assume that all free parameters can only be fitted with a single constant value over the full sky, at least for now. Already at this stage, we fix the thermal dust temperature, T_{d} , at the sky map derived by *Planck* DR4 (*Planck* Collaboration Int. LVII 2020), noting that LFI has no constraining power for this particular parameter.

Even though LFI should have some constraining power of the thermal dust spectral index, β_{d} , the thermal dust and AME components are found to be highly degenerate with the limited data set used in this work. A joint fit of the AME ν_{p} and β_{d} would therefore lead to unphysical results, with preliminary analyses showing that the β_{d} diverged to values $\beta_{\text{d}} > 2$ pulling ν_{p} to much




Fig. 11. AME (left) and free-free (right) amplitude maps derived for two different synchrotron spectral indices. The top row shows results for $\beta_s = -2.8$, while the bottom row shows results for $\beta_s = -3.3$. In the top row we see a negative synchrotron-like (see Fig. 16) imprint that is not present in the bottom row.

higher frequencies. The uncertainty in the β_d value is important for error propagation, and we therefore simply marginalize over the adopted prior, instead of trying to constrain β_d with the current data set.

For each free spectral parameter we create a dedicated sampling mask, in which we exclude regions on the sky where the other components are strong in order to reduce potential modelling mismatch errors to propagate between the various components. These masks are created from the amplitude maps of the modelled components, evaluated at 44 GHz and smoothed to 10° FWHM. In addition, we mask radio sources by thresholding the *Planck* 30 GHz compact source map at three different angular resolutions, namely at native resolution, and at 1° and 10° FWHM. For β_s and ν_p , we exclude regions of the sky where any other component signals is greater than $40 \mu\text{K}_{\text{RJ}}$, while for T_e we exclude the areas where the other components are greater than $50 \mu\text{K}_{\text{RJ}}$. For all parameters we exclude any pixel for which the smoothed radio sources are stronger than $30 \mu\text{K}_{\text{RJ}}$. Finally, we mask out regions of the sky contributing to the largest 15 % of a 10° FWHM χ^2 map to exclude regions with large known modelling errors. The accepted sky fractions of the resulting masks are $f_{\text{synch}} = 0.66$, $f_{\text{ff}} = 0.74$, and $f_{\text{AME}} = 0.66$, respectively.

For each free parameter, we adopt a Gaussian informative prior with some mean and standard deviation, $P(\beta) = N(\mu_\beta, \sigma_\beta^2)$. The prior parameters are informed by literature results, and listed in Table 2. For synchrotron emission, we note that few intensity-based constraints are available for frequencies higher than 30 GHz in the literature, and we therefore adopt $\beta_s = -3.3 \pm 0.1$ as derived from polarization measurements in [Planck Collaboration V \(2020\)](#). We also note that we have attempted to use flatter mean values of $\beta_s = -3.1$ and -3.0 , as suggested from low-frequency surveys, but these result in obvious artefacts in

the current analysis in the form of a significantly over-estimated synchrotron amplitude at 30 GHz. As an example, Fig. 11 compares the AME and free-free amplitude maps derived for two different values of β_s , namely $\beta_s = -2.8$ and -3.3 . While neither of these solutions produce an excess χ^2 , and therefore a free likelihood-driven fit is unable to distinguish between them, it is obvious from visual inspection that the former spectral index leads to clear synchrotron leakage into both the AME and free-free components. With a prior of $\beta_s = -3.3 \pm 0.1$, the nonphysical flat-index solutions are largely excluded, while some parameter space is still allowed toward the steeper end to explore degeneracies. Another approach of reducing the predicted amplitude at CMB frequencies is by introducing a negative curvature in the synchrotron SED, as discussed in Sect. ??, and when comparing the results from this paper with previous results, it is important to ensure that the models are compatible.

Given the listed priors, we perform a coarse χ^2 grid evaluation for each free parameter, conditioning on all other spectral parameter means, but allowing for amplitudes and band monopoles to adjust to the given spectral parameter. The resulting χ^2 's (evaluated at a HEALPix resolution of $N_{\text{side}} = 16$) are shown in Fig. 12.

Starting with AME ν_p , shown in the top panel, we see that the χ^2 is well-defined with a typical best-fit value around 22 GHz. The actual χ^2 values show rapid increases at both lower and higher values with variations ranging in the thousands, and correspondingly, the prior (which is of order unity) is therefore largely irrelevant. It is clear that the current data combination has significant constraining power for ν_p .

The second panel shows similar results for β_s . In this case, we see a rapid χ^2 increase for $\beta \lesssim -2.7$, but otherwise it is slowly increasing for smaller values of β_s in the region $\beta_s < -3.2$, be-




Fig. 12. χ^2 distributions (blue curves) from coarse grid evaluation of each free spectral parameters. From top to bottom, the panels show 1) AME peak frequency, ν_p ; 2) synchrotron spectral index, β_s ; and 3) free-free electron temperature, T_e . The minimum χ^2 value for each parameter has been subtracted in each case. Orange curves show the priors adopted for the given component; see Table 2. The prior values for β_s and ν_p have been scaled by a factor of 100 to fit in the plot with the derived χ^2 values.

coming almost flat. Additionally, we already know from Fig. 11 that spectral indices flatter than $\beta_s \lesssim -2.8$ lead to clearly contaminated AME component maps, even if the χ^2 is not able to identify this. At the same time, the actual χ^2 variations are indeed larger than the prior, and this typically indicates that the algorithm prefers to use this unconstrained degree of freedom to fit other degrees of freedom, for instance modelling errors in the thermal dust model. To avoid pathological solutions, we therefore instead choose to marginalize explicitly over the prior, and disable the likelihood term entirely when sampling this parameter. In other words, we simply marginalize over the adopted prior, but do not attempt to constrain β_s with the current data set.

The same considerations hold to an even greater extent for the last parameter. For the electron temperature, T_e , the χ^2 vari-

ations are entirely spurious, and we therefore disable the likelihood term for T_e .

In summary, the only spectral parameter that the current data set is able to robustly constrain in intensity is the AME peak frequency, ν_p . All others are either drawn from their corresponding priors in the current analysis or frozen, and we note that introducing additional data sets to constrain these parameters is a critically important next step for future work.

With β_s and β_d drawn from priors and T_e frozen, we find that the mask used in the coarse grid sampling of ν_p is too conservative when sampling ν_p , masking out too much of the galactic plane, leading to large dust-like residual signals in the LFI 30 GHz and the *WMAP* *Ka* channels. Additionally, more dust like signal was found to be leaking into the free-free component amplitude. In order to limit these effects, a less conservative mask had to be used. The mask used to sample ν_p in the final BEYONDPLANCK production is generated by excluding all pixels with values above $200 \mu\text{K}_{\text{RJ}}$ of the free-free amplitude at both $30'$ and 2° FWHM angular resolution evaluated at 40 GHz; all pixels with values above $100 \mu\text{K}_{\text{RJ}}$ of the point source amplitudes smoothed with a 1° FWHM beam evaluated at the LFI 30 GHz band frequency; and the regions of the sky contributing to the largest 2.5 % of the χ^2 . The accepted sky fraction of the resulting mask is $f_{\text{AME}} = 0.91$.

4.2. Markov chain trace plots and correlations

With data, sampling algorithms, and priors in place, we are now ready to consider the actual Markov chain results. As described by BeyondPlanck (2022), for the final analysis we produce two independent chains, each with 750 samples.

Figure 13 shows the first 250 samples for a selection of relevant parameters. Several points are worth noting in this figure. First, we note that the burn-in period for most of the foreground-induced parameters is very short, while it is slightly longer for some of the global gain and bandpass instrumental parameters. We remove the first 10 samples for burn-in for the following analysis. This short burn-in is a combination of the novel sampling algorithms introduced in Sect. 3 and the prior choices discussed above. At the same time, we do observe a weak shift in the average of ν_p , where the chains split away from each other around sample 120, which appears to trace some of the more slowly varying instrumental parameters, primarily the total bandpass correction of LFI 30 GHz; this is quite typical, as many global instrumental parameters tend to have long Markov correlation lengths, and these directly affect foreground residuals.

In Fig. 14 we plot Pearson's correlation coefficients between the various parameters. For this particular plot, we have subtracted a running average with a length of ten samples (five samples prior and succeeding) from each function before evaluating the correlation coefficient, in order to highlight sample-by-sample correlations; two parameters may appear to be spuriously correlated if there are long-term gradients, irrespective of their origins.

Several interesting observations may be made from this plot. First, we note that there is a very high correlation between β_d and both the AME peak frequency ν_p and the CMB dipole amplitude $A_{\text{dp}}^{\text{CMB}}$. This is not unexpected, given the critically important role of thermal dust emission at all CMB frequencies. At the same time, this also serves as a useful reminder that several important BEYONDPLANCK results depend directly on official *Planck* results through the use of the HFI 857 GHz frequency band, and the




Fig. 13. Trace plots of a set of sampled component separation parameters and selected instrument parameters, including absolute calibration (g), bandpass shift (Δ_{bp}), and the Solar CMB dipole (A_{dp}^{cmb}) from a naive mono- and dipole estimate using the band monopole mask as described in Fig. 3. The parameter a_{ame} is the AME amplitude evaluated at 10° FWHM and $N_{side} = 16$ of the pixel at Galactic coordinates $(l, b) = (31^\circ, 5^\circ)$. Parameters from chain 1 is plotted in blue and chain 2 in orange.




Fig. 14. Correlation coefficient plot of local deviation, from a running mean of 5 prior and 5 succeeding samples, of the sampled spectral parameters, component cross-correlation intersections, frequency band monopoles, absolute gain calibration, LFI 30 GHz bandpass shift, and CMB dipole amplitude, all as described in Fig. 13.

assumed thermal dust spectral index prior of $\beta_d = 1.56 \pm 0.03$ (Planck Collaboration X 2016; Planck Collaboration IV 2018; Planck Collaboration Int. LVII 2020), and the systematic errors in these are not propagated properly in the current analysis. Integrating HFI observations into the BEYONDPLANCK framework at the TOD level is clearly an important goal of near-term work.

Next, we note that all microwave frequency monopoles are internally strongly correlated, and notably anti-correlated with the component monopoles. Both of these observations make intuitive sense, as if one increases a given component monopole, the frequency monopoles have to decrease to result in a net zero change to the overall model. In addition, all frequency monopoles have to change in the same direction to a given change in the component monopoles.

Accurately accounting for all of these correlations in terms of MCMC samples is perhaps the most important advantage of the Bayesian end-to-end BEYONDPLANCK framework.

4.3. Goodness-of-fit statistics

Before presenting the actual posterior distributions, we consider the goodness-of-fit of the derived model. First, Fig. 15 shows

posterior mean residual maps on the form $r_v = d_v - s_v$ for each of the ten frequency bands included in the analysis, where the average is evaluated over all samples in the Markov chain.

Starting with the LFI channels shown in the top section of the figure, we see that the residuals at high Galactic latitudes are largely consistent with instrumental noise, except for scattered point source residuals at 30 GHz, while at low Galactic latitudes there is an obvious dust morphology residual coming from either AME or thermal dust. At 30 GHz it is also possible to see some bright negative free-free residuals. Overall, though, the fits are performing quite well, with typical residuals smaller than $3 \mu K$ over most of the sky, and the remaining artefacts are relatively easy to mask through Galactic and point source thresholding.

The middle section shows the WMAP channels, plotted on a color range of $\pm 5 \mu K$. The most striking residual feature in this case is a strong dust residual in the *Ka* (33 GHz) channel. This, combined with the weaker dust residuals observed in the LFI channels, strongly suggests that the current single-component AME model adopted for the BEYONDPLANCK analysis is not a statistically adequate model for the actual AME sky. In fact, this shortcoming was already pointed out by Planck Collaboration X (2016), who introduced a second AME component to fit the

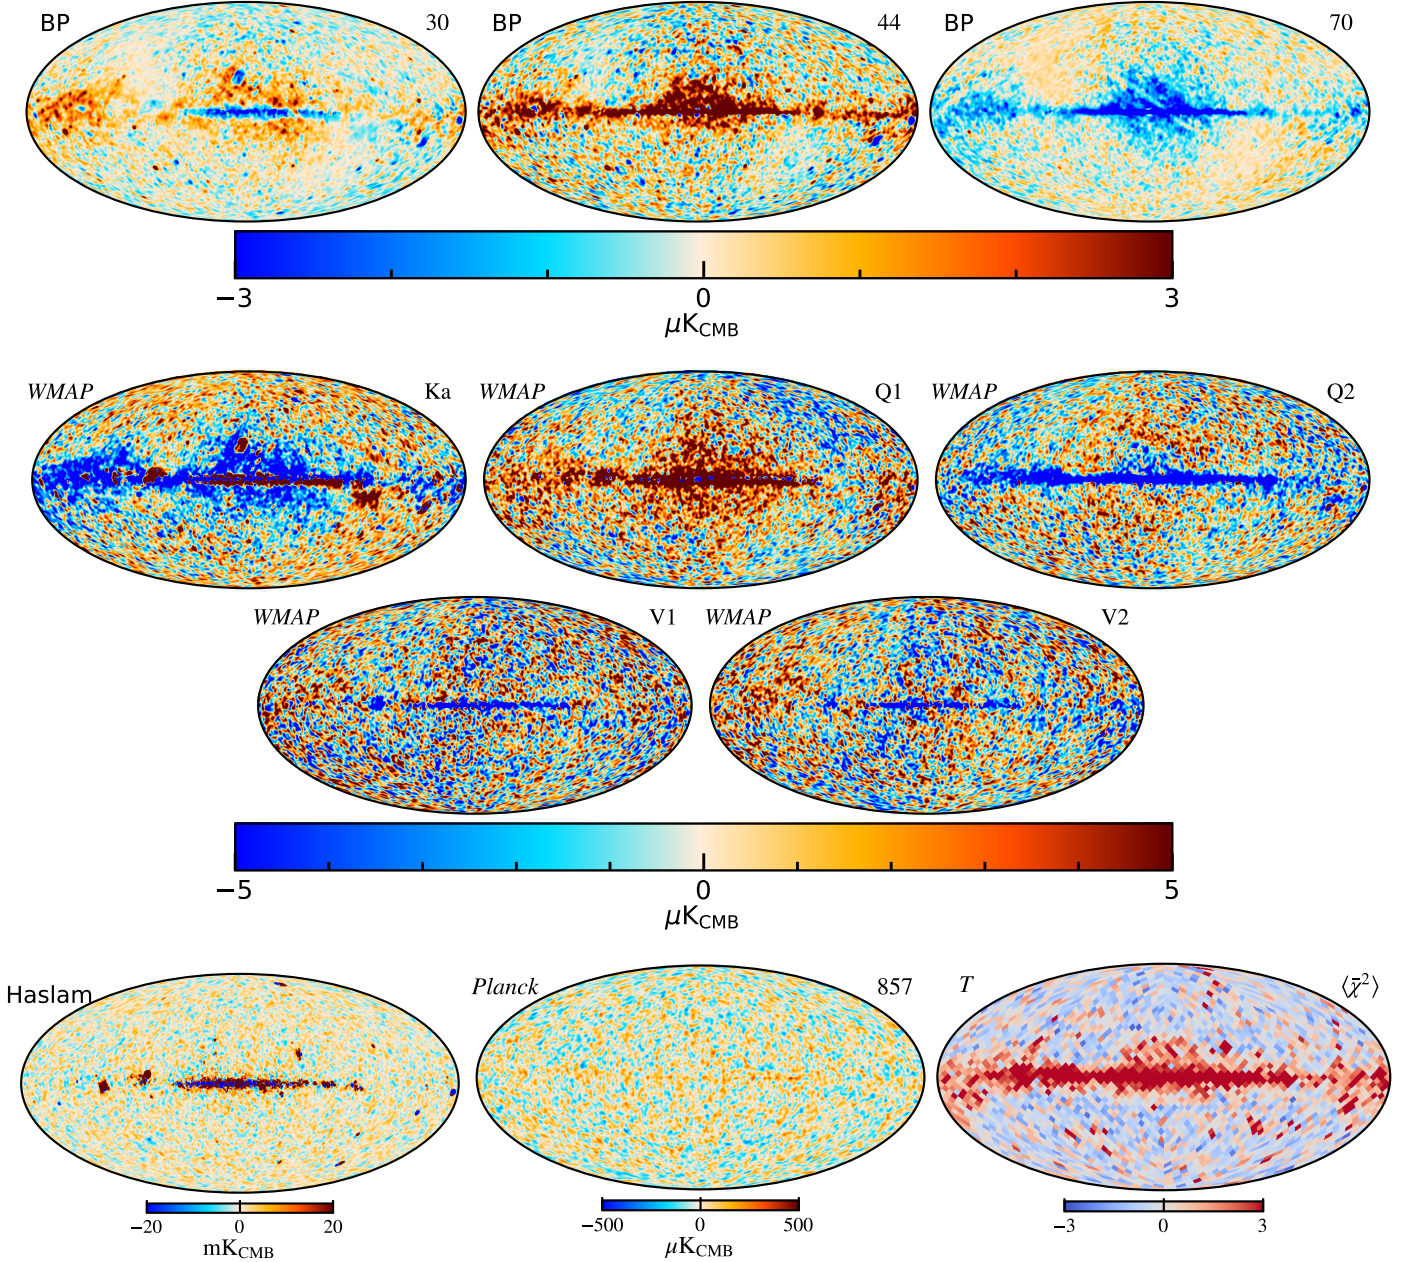


Fig. 15. Mean residual maps for the different frequency channels included in the BEYONDPLANCK analysis. All maps have been smoothed to a common angular resolution of 2° FWHM. The *Planck* LFI residuals are plotted with a range of $3 \mu\text{K}_{\text{CMB}}$, and the *WMAP* residuals are plotted with a range of $5 \mu\text{K}_{\text{CMB}}$. The bottom right panel shows the mean reduced chi-squared $\bar{\chi}^2$ per $N_{\text{side}} = 16$ pixel of the BEYONDPLANCK Gibbs chain.

full contribution. Doing the same with the current data selection would lead to a massively increased noise level for all derived components, and we instead accept the foreground mis-modelling here, and we instead simply make sure to mask out the contaminated regions of the sky in higher-order analyses.

Other notable features in the *WMAP* residuals are large regions of low-level residuals at high Galactic latitudes that do not obviously trace known Galactic components. As discussed by Barnes et al. (2003), an important challenge regarding this data set on large angular scales is sidelobe modelling, and this may also be relevant for the residuals we see in Fig. 15. A re-analysis of the time-ordered *WMAP* data within the BEYONDPLANCK framework is already on-going (Watts et al. 2022).

The last two frequency channels, Haslam 408 MHz and HFI 857 GHz, show very uniform residuals. This is simply due to

the fact that their unique signal-to-noise ratios massively dominate the synchrotron and thermal dust components, respectively, and any potential mis-modelling therefore leak directly into the component amplitude maps. A flat residual should thus not be interpreted as the absence of systematic errors, but rather simply as an indication that these two channels have no significant cross-check in terms of their effect on the signal model by any other channel. (This is not entirely true for the Haslam 408 MHz map, which does have competitors in terms of free-free emission near the Galactic center in *Planck* and *WMAP*, and corresponding residuals may be seen here.)

Table 3. Comparison of frequency channel monopole constraints from BEYONDPLANCK and *Planck*. For *Planck*, all numbers correspond to the *Planck* 2015 analysis (Planck Collaboration X 2016), except for the HFI 857 GHz, which is taken from *Planck* DR4 (Planck Collaboration Int. LVII 2020). Note that the *Planck* 2015 analysis did not fit for the *WMAP* *Ka*-band monopole, but fixed it at the given value; this parameter therefore has no associated uncertainty. Lastly, note that while the *WMAP* and HFI maps are the same in both columns, the LFI maps are not, and they have different zero-levels coming from the TOD processing.

Channel	BEYONDPLANCK	<i>Planck</i>	Unit
LFI 30 GHz	0 ± 6	-17 ± 1	μK_{CMB}
LFI 44 GHz	0 ± 2	11 ± 1	μK_{CMB}
LFI 70 GHz	0 ± 1	16 ± 1	μK_{CMB}
<i>WMAP Ka</i>	16 ± 3	3	μK_{CMB}
<i>WMAP Q1</i>	12 ± 2	2 ± 1	μK_{CMB}
<i>WMAP Q2</i>	11 ± 2	2 ± 1	μK_{CMB}
<i>WMAP V1</i>	6 ± 1	1 ± 1	μK_{CMB}
<i>WMAP V2</i>	6 ± 1	1 ± 1	μK_{CMB}
HFI 857 GHz	-0.65 ± 0.03	-0.72 ± 0.01	K_{CMB}

The bottom right panel of Figure 15 shows the reduced and normalized χ^2 per $N_{\text{side}} = 16$ pixel, as defined by

$$\bar{\chi}^2(p) = \frac{\sum_{v,p' \in p} \left[\frac{d_v(p') - s_v(p')}{\sigma_v(p')} \right]^2 - n_{\text{dof}}}{\sqrt{2n_{\text{dof}}}}, \quad (31)$$

where the sum runs over all pixels within a given low-resolution pixel, and $n_{\text{dof}} = 15400$ is an estimate of the total number of degrees of freedom for each low-resolution pixel. Since n_{dof} is large, this quantity is expected to be $N(0, 1)$ distributed in the ideal case, and Fig. 15 thus quantifies the agreement between the data and the model in units of standard deviations per pixel. Overall, we see that the distribution agrees with the expectation to about 1σ at high Galactic latitudes, except for some compact sources, while at low Galactic latitudes there is a strong residual with a clear dust-like morphology. This χ^2 map serves as an important input for producing masks for higher-order analyses.

4.4. Signal component posterior distributions

Next we will consider the signal component posterior distributions, and we start with the frequency monopoles, as summarized in terms of posterior mean and standard deviations in the second column of Table 3. For comparison purposes, the third column shows corresponding results derived by the *Planck* team; all results except 857 GHz are reproduced from the *Planck* 2015 analysis (Planck Collaboration X 2016), while the 857 GHz result is taken from *Planck* DR4 (Planck Collaboration Int. LVII 2020).

Several important differences between the two sets of results can be identified. First, we note that the BEYONDPLANCK LFI mean monopoles are all zero; this happens by construction during the mapmaking process, as the frequency band monopoles are determined directly from the sky model, and any deviation from this is assigned to the correlated noise component (Ihle et al. 2022). On the other hand, we do see that the uncertainties of the LFI zero-levels are larger at both 30 and 44 GHz, reflecting the difficulty of uniquely determining the AME offset, as discussed in Sect. 3.2. In addition, the uncertainty of the Haslam, and thereby the synchrotron monopole propagates down to the lower frequency *Planck* and *WMAP* channels and contributes to increased

uncertainties of these monopoles. We argue that the uncertainties determined through component-based monopole determination are more realistic than those obtained through morphologically based frequency monopole determination.

For *WMAP*, we note that the our monopole corrections are considerably larger compared to those determined in the *Planck* 2015 estimates. This is caused by two main differences. First, the 2015 analysis adopted the *WMAP Ka*-band explicitly as a fixed “anchor channel” (Planck Collaboration X 2016) that was not allowed to vary in the analysis. In the current analysis, for which we instead impose constraints directly on the component monopoles, this channel is instead associated with a $16\mu\text{K}_{\text{CMB}}$ correction, which is then accounted for by an offset of $-17\mu\text{K}_{\text{CMB}}$ at LFI 30 GHz in the 2015 analysis. Second, we have changed the average synchrotron index from $\beta_s = -3.1$ in the *Planck* 2015 analysis (Planck Collaboration X 2016) to $\beta_s = -3.3$ in the current analysis. Since the synchrotron monopole is set to $8.9 \text{ K}_{\text{RJ}}$ at 408 MHz (see Sect. 2.2), the predicted monopole at *WMAP Ka*-band is $11\mu\text{K}_{\text{CMB}}$ with $\beta_s = -3.1$ and $4\mu\text{K}_{\text{CMB}}$ with $\beta_s = -3.3$, a difference of $7\mu\text{K}_{\text{CMB}}$. The frequency band monopole has to adjust for this difference. Similarly, differences in the zero-level of the other components, especially the AME, will contribute to frequency-band monopole differences in equal fashion. Regarding the offsets for the 857 GHz channel, we see that these agree within 2σ as estimated by BEYONDPLANCK, and the uncertainties are significantly larger (and, we believe, more realistic) in the new approach, which is a reflection of the fact that we are now propagating uncertainties in the thermal dust zero-level as discussed in Sect. 3.2.

Regarding spectral parameters, the AME peak frequency is the only spectral parameter that is fully sampled from the data in this component separation work. For this, we find a posterior mean and standard deviation of $\nu_p = 25.3 \pm 0.5 \text{ GHz}$. From the high correlation coefficient between ν_p and the thermal dust β_d seen in Fig. 14, it is clear that this result is highly dependent on the assumptions made regarding thermal dust modelling in this paper, and future work that includes *Planck* HFI observations will be critically important to make the AME model more robust.

Posterior mean and standard deviation maps for each of the four component maps (synchrotron, free-free, AME, and thermal dust emission) are shown in Figs. 16 and 17, each at their own angular resolution (120, 30, 120, and $10'$) and reference frequency (408 MHz, 40 GHz, 22 GHz, and 857 GHz), and all shown in terms of brightness temperature. In Figure 17, we notice the dark stripes following the *Planck* scanning pattern at high galactic latitudes in the standard deviation map of the free-free component. These stripes are barely visible in the AME standard deviation map while being completely absent from the synchrotron and thermal dust maps. This can be explained by the sampling of the amplitude zero-levels, where the free-free zero-level is the only one not marginalized over in the BEYONDPLANCK Gibbs chain. The base value at higher latitudes of the amplitude RMS is equivalent to the RMS of the sampled zero-levels.

Figure 18 shows mean and standard distribution plots for the compact source component for a $20^\circ \times 20^\circ$ field of the sky in the Northern Galactic hemisphere for each of the three LFI frequency channels.

The four component posterior mean maps may be compared to similar products from the *Planck* 2015 analysis (Planck Collaboration X 2016), and corresponding difference maps are shown in Fig. 19. All maps are smoothed to a common angular resolution of 2° FWHM before differencing, and a free offset has been fitted and subtracted using the frequency-band monopole mask discussed in Sect. 3.3. In addition, for the components that




Fig. 16. Posterior mean amplitude maps for each of the four fitted foreground component; synchrotron (*top left*), free-free (*top right*), AME (*bottom left*), and thermal dust emission (*bottom right*). The angular resolutions of the four maps are 120, 30, 120, and 10' FWHM, respectively.




Fig. 17. Posterior standard deviations for the same maps as shown in Fig. 16. Note that synchrotron and dust emission are potted with linear scaling.




Fig. 18. Partial sky plots of the mean (*top*) and standard deviation (*bottom*) amplitude of the fitted compact sources (point sources) as seen by the three *Planck* LFI detector bands; 30 GHz (*left*), 44 GHz (*middle*) and 70 GHz (*right*). The plots show gnomonic projections of a 20×20 degree patch of the sky centered on 90° longitude and 70° latitude, with the north galactic pole located towards the top-center of the plots. All plots are at native angular resolution and pixelization, see Table 1 for details.

have different reference frequencies in the two analyses (i.e., free-free, AME, and thermal dust emission), a single multiplicative factor has been fitted to take into account SED scaling differences.

Starting with the synchrotron case, we note that typical high-latitude differences are small compared to the overall amplitude of the Haslam 408 MHz map, typically less than $1 K_R$. This is of course entirely expected, since the two analyses are both dominated by the same map. However, we do see differences at both low and high Galactic latitudes; for high latitudes we note that the current analysis explicitly models individual point sources, while in the *Planck* 2015 analysis there was no such component, and compact sources were therefore part of the diffuse components. At low latitudes, the differences are dominated by differences in the free-free, AME and thermal dust emission models.

For free-free emission, we see clear negative imprints of the free-free amplitude map itself at low Galactic latitudes. In this case, we note that while the *Planck* 2015 analysis fitted the elec-

tron temperature pixel-by-pixel, we adopt a single constant value of $T_e = 7000$ K in the current analysis. Second, we see clear positive imprints of AME or dust around the negative free-free imprint at low Galactic latitudes, resulting from degeneracies between the component signals. Lastly, we also note that the current analysis suffers significantly with respect to free-free emission from the absence of the *Planck* HFI channels, which provide both angular resolution and sensitivity to this component.

The AME component residual map is largely dominated by a dipole component aligned with the direction of the Galactic center. In this case, we note that the BEYONDPLANCK analysis estimates the absolute calibration of each *Planck* LFI frequency channel through a joint fit with all frequencies, and this provides more robust estimates of the CMB dipole. We also see a negative free-free imprint at low Galactic latitudes similar to what we see in the free-free difference map. One key feature is the almost circular blob just above and slightly to the left of the Galactic center, which is also clearly seen as a dark blob in the mean

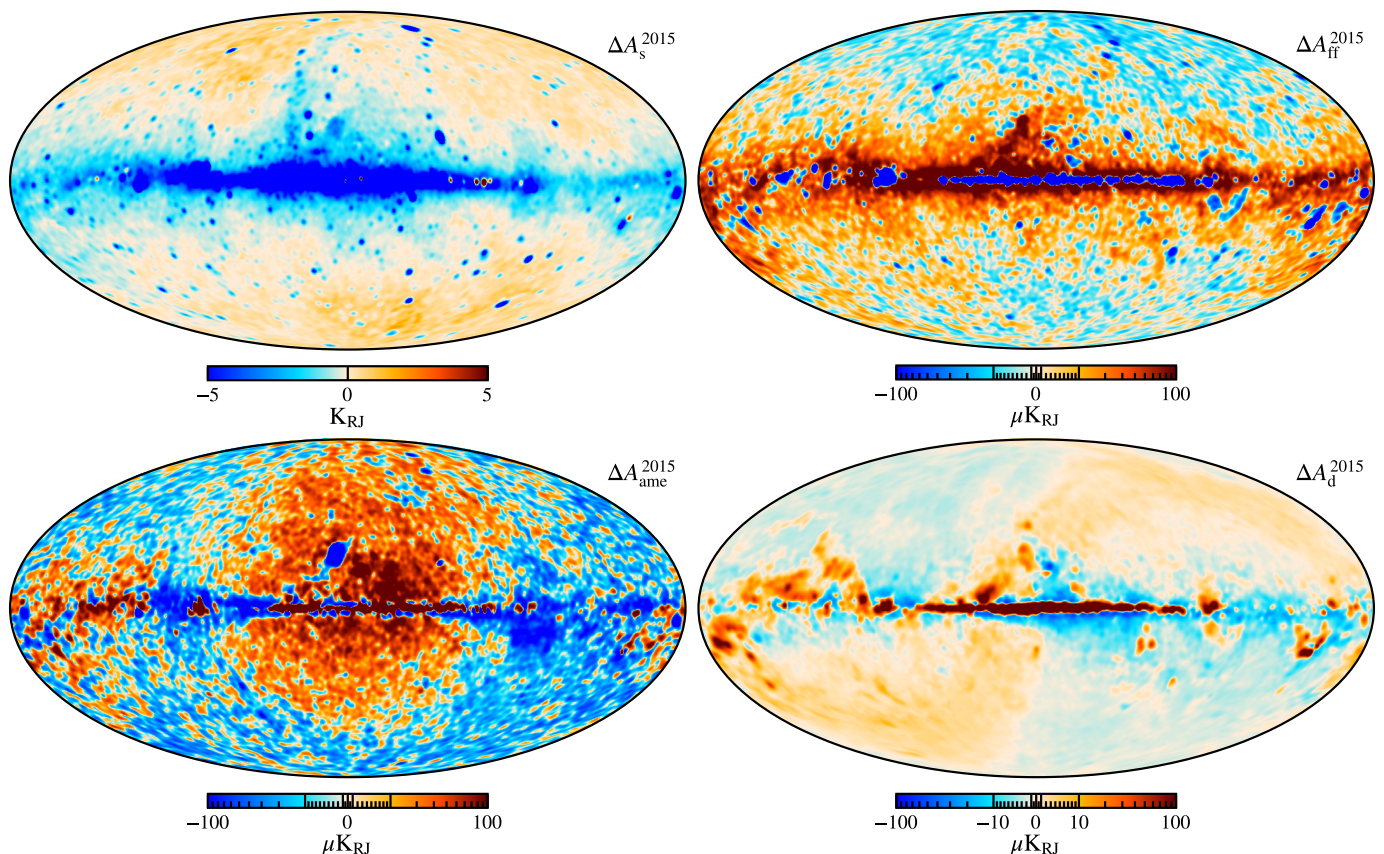


Fig. 19. Difference maps between the component amplitude maps derived in BEYONDPLANCK and the *Planck* 2015 analysis, for synchrotron (top left), free-free (top right), AME (bottom left), and thermal dust emission (bottom right), respectively. All maps have been smoothed to a common angular resolution of 2° FWHM, a relative offset has been fitted and subtracted using the frequency-band monopole mask discussed in Sect. 3.3, and differences in reference frequencies (where relevant) have been accounted for by a single multiplicative scaling factor.

map in Fig. 16. This coincides with a similarly bright circular blob in the free-free amplitude. A similar signal can be seen in the *Planck* 2015 free-free signal (Planck Collaboration X 2016), though it did not affect the AME amplitude in the same way. If we look at the residuals in Fig. 15, we notice the same blob in the LFI 30 GHz and the *WMAP* *Ka* maps, indicating a degeneracy between the modelled AME and free-free components. To break this degeneracy, either more data sets need to be introduced or the free-free T_e and the AME ν_p need to be sampled with spatial variance, or both. This is however not possible with the limited data set in this analysis and therefore needs to be left for future analyses.

Finally, for the thermal dust emission amplitude map, which is essentially defined by the *Planck* 857 GHz frequency map, the differences are explained very closely by the morphology of the thermal dust spectral index map $\beta_d(p)$ presented by Planck Collaboration X (2016). This is also entirely expected, given the fact that we only model β_d in terms of a single spatial constant over the full sky.

5. Summary and outlook

The main goal of the current paper is to establish an efficient Monte Carlo sampling scheme for intensity foregrounds within an end-to-end Bayesian CMB analysis pipeline, as implemented in the BEYONDPLANCK framework. This sampling scheme must be able to operate in both low and high signal-to-noise regimes, and it must be able to incorporate both algorithmic and informative priors in a controlled and transparent manner. Degenera-

cies between different parameters must be explored properly, and it must be possible to propagate corresponding uncertainties to higher-level products.

By far most of the algorithmic elements in the machinery used in this paper were developed and applied within the context of the official *Planck* project, as described in, e.g., Planck Collaboration X (2016), Planck Collaboration IV (2018), and Planck Collaboration Int. LVII (2020). In this paper, we have added four new algorithmic components to this machinery, namely 1) joint amplitude and spectral index sampling, borrowing heavily from ideas already introduced by Stompor et al. (2009) and Stivoli et al. (2010); 2) component-based monopole determination; 3) joint spectral index and monopole sampling; and 4) the application of informative spatial priors. Each of these steps significantly improves the computational efficiency and robustness of the Gibbs sampling based Commander approach.

We stress that the current BEYONDPLANCK results are not intended to define a new state-of-the-art model of the astrophysical sky. Rather, the current framework and analysis constitute a “skeleton” to which additional data sets, both legacy and future, may be added in a controlled fashion, keeping track of both instrumental and astrophysical modelling errors. As more and more data sets are added, the dependency on external priors may be gradually lifted, until, hopefully, all key parameters of the model become data driven. Obviously, the single most important step towards realizing this goal is the introduction of *Planck* HFI TOD observations, which still define the state-of-the-art in terms of full-sky CMB sensitivity to date. On a longer term, we argue that all key data sets in the community should be inte-




Fig. 20. Brightness temperature rms as a function of frequency and astrophysical component for temperature. Each component is smoothed to an angular resolution of 1° FWHM, and the lower and upper edges of each band are defined by masks covering 27 and 88 % of the sky, respectively. We note that foreground rms values decrease nearly monotonically with sky fraction, whereas the CMB rms is independent of sky fraction, up to random variations. The vertical bands represent the frequency range of detector data, where the top panel shows the data employed in this paper and the bottom panel shows some data available for future analysis.

grated into the model, allowing one set to break the degeneracies of others. This is the goal of the Open Source COSMOGLOBE project, and the transition from BEYONDPLANCK to COSMOGLOBE may be illustrated in Fig. 20: The top panel of this figure shows the standard deviation brightness temperature as a function of frequency for each primary intensity CMB foreground as colored bands. The BEYONDPLANCK frequency bands are indicated

by vertical bars. Looking at this figure, noting the large unexplored frequency ranges, it is strikingly obvious why the current data model is significantly prior dominated: More data are desperately needed. The bottom panel of Fig. 20 shows an alternative scenario, in which almost the entire frequency range is covered by including data from several past and planned sky surveys. Providing a computationally and organizationally efficient

platform to make this happen is the goal of COSMOGLOBE, and the BEYONDPLANCK project represents an important step towards realizing this promise.

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); Tekes, Academy of Finland (grant no. 295113), CSC, and Magnus Ehrnrooth foundation (Finland); RCN (Norway; grant nos. 263011, 274990); and PRACE (EU).

References

- Ali-Haïmoud, Y. 2010, *Astrophysics Source Code Library* [ascl:1010.016]
 Alves, M. I. R., Calabretta, M., Davies, R. D., et al. 2015, *MNRAS*, 450, 2025
 Barnes, C., Hill, R. S., Hinshaw, G., et al. 2003, *ApJS*, 148, 51
 Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2003, *ApJS*, 148, 97
 Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20
 BeyondPlanck. 2022, A&A, in preparation [arXiv:2011.05609]
 Colombo et al. 2022, A&A, in preparation [arXiv:201x.xxxxx]
 Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, 115, 1693
 Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, A&A, 553, A96
 Draine, B. T. & Lazarian, A. 1998, *ApJ*, 494, L19
 Erickson, W. C. 1957, *ApJ*, 126, 480
 Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, 676, 10
 Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, *ApJS*, 155, 227
 Finkbeiner, D. P. 2003, *ApJS*, 146, 407
 Fixsen, D. J. 2009, *ApJ*, 707, 916
 Galloway et al. 2022a, A&A, submitted [arXiv:2201.03509]
 Galloway et al. 2022b, A&A, submitted [arXiv:2201.03478]
 Geman, S. & Geman, D. 1984, *IEEE Trans. Pattern Anal. Mach. Intell.*, 6, 721
 Génova-Santos, R., Rubiño-Martín, J. A., Rebolo, R., et al. 2015, *MNRAS*, 452, 4169
 Gjerløw et al. 2022, A&A, in preparation [arXiv:2011.08082]
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Gregory, P. C., Scott, W. K., Douglas, K., & Condon, J. J. 1996, *ApJS*, 103, 427
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
 Hensley, B. S. & Draine, B. T. 2020, *ApJ*, 895, 38
 Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, *The Astrophysical Journal Supplement Series*, 180, 225–245
 Ihle et al. 2022, A&A, in preparation [arXiv:2011.06650]
 Jewell, J., Levin, S., & Anderson, C. H. 2004, *ApJ*, 609, 1
 King, O. G., Copley, C., Davies, R., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7741, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1
 Lenz, D., Hensley, B. S., & Doré, O. 2017, *ApJ*, 846, 38
 Mather, J. C., Cheng, E. S., Cottingham, D. A., et al. 1994, *ApJ*, 420, 439
 Murphy, T., Sadler, E. M., Ekers, R. D., et al. 2010, *MNRAS*, 402, 2403
 Paradiso et al. 2022, A&A, in preparation [arXiv:201x.xxxxx]
 Penzias, A. A. & Wilson, R. W. 1965, *ApJ*, 142, 419
 Planck Collaboration XXIV. 2011, A&A, 536, A24
 Planck Collaboration V. 2014, A&A, 571, A5
 Planck Collaboration XII. 2014, A&A, 571, A12
 Planck Collaboration II. 2016, A&A, 594, A2
 Planck Collaboration III. 2016, A&A, 594, A3
 Planck Collaboration IV. 2016, A&A, 594, A4
 Planck Collaboration VII. 2016, A&A, 594, A7
 Planck Collaboration VIII. 2016, A&A, 594, A8
 Planck Collaboration IX. 2016, A&A, 594, A9
 Planck Collaboration X. 2016, A&A, 594, A10
 Planck Collaboration XXVI. 2016, A&A, 594, A26
 Planck Collaboration I. 2020, A&A, 641, A1
 Planck Collaboration II. 2020, A&A, 641, A2
 Planck Collaboration III. 2020, A&A, 641, A3
 Planck Collaboration IV. 2018, A&A, 641, A4
 Planck Collaboration V. 2020, A&A, 641, A5
 Planck Collaboration Int. LVII. 2020, A&A, 643, A42
 Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, *MNRAS*, 418, 467
 Seljebotn, D. S., Bærlund, T., Eriksen, H. K., Mardal, K. A., & Wehus, I. K. 2019, A&A, 627, A98
 Seljebotn, D. S., Mardal, K.-A., Jewell, J. B., Eriksen, H. K., & Bull, P. 2014, *Astrophys. J. Suppl.*, 210, 24
 Shewchuk, J. R. 1994, *An Introduction to the Conjugate Gradient Method Without the Agonizing Pain*, Edition 1 $\frac{1}{4}$, <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf>
 Silsbee, K., Ali-Haïmoud, Y., & Hirata, C. M. 2011, *MNRAS*, 411, 2750
 Stivoli, F., Grain, J., Leach, S. M., et al. 2010, *MNRAS*, 408, 2319
 Stompor, R., Leach, S., Stivoli, F., & Baccigalupi, C. 2009, *MNRAS*, 392, 216
 Svalheim et al. 2022a, A&A, submitted [arXiv:2201.03417]
 Svalheim et al. 2022b, A&A, submitted [arXiv:2011.08503]
 Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511
 Watts et al. 2022, A&A, in preparation [arXiv:201x.xxxxx]
 Wehus, I. K., Fuskeland, U., Eriksen, H. K., et al. 2017, A&A, 597, A131