

Tensor-to-scalar ratio forecasts for extended LiteBIRD frequency configurations

U. Fuskeland^{1,*}, J. Aumont², R. Aurlien¹, C. Baccigalupi^{3,4,5}, A. J. Banday², H. K. Eriksen¹, J. Errard⁶, R. T. Génova-Santos^{7,8}, T. Hasebe⁹, J. Hubmayr¹⁰, H. Imada¹¹, N. Krachmalnicoff^{3,4,5}, L. Lamagna^{12,13}, G. Pisano¹², D. Poletti^{14,15}, M. Remazeilles^{16,17}, K. L. Thompson^{18,19}, L. Vacher², I. K. Wehus¹, S. Azzoni^{20,9}, M. Ballardini^{21,22,23}, R. B. Barreiro¹⁶, N. Bartolo^{24,25,26}, A. Basyrov¹, D. Beck¹⁹, M. Bersanelli^{27,28}, M. Bortolami²¹, M. Brilenkov¹, E. Calabrese²⁹, A. Carones^{30,31}, F. J. Casas¹⁶, K. Cheung^{32,33,34}, J. Chluba¹⁷, S. E. Clark^{19,35}, L. Clermont³⁶, F. Columbro^{12,13}, A. Coppolecchia^{12,13}, G. D'Alessandro^{12,13}, P. de Bernardis^{12,13}, T. de Haan^{37,38}, E. de la Hoz^{16,39}, M. De Petris^{12,13}, S. Della Torre¹⁵, P. Diego-Palazuelos^{16,39}, F. Finelli^{23,40}, C. Franceschet^{27,28}, G. Galloni³⁰, M. Galloway¹, M. Gerbino²², M. Gervasi^{14,15}, T. Ghigna³⁷, S. Giardiello²⁹, E. Gjerløw¹, A. Gruppuso^{23,40}, P. Hargrave²⁹, M. Hattori⁴¹, M. Hazumi^{38,37,42,9,43}, L. T. Hergt⁴⁴, D. Herman¹, D. Herranz¹⁶, E. Hivon⁴⁵, T. D. Hoang⁹, K. Kohri³⁷, M. Lattanzi²², A. T. Lee^{46,32,38}, C. Leloup⁶, F. Levrier⁴⁷, A. I. Lonappan³, G. Luzzi⁴⁸, B. Maffei⁴⁹, E. Martínez-González¹⁶, S. Masi^{12,13}, S. Matarrese^{24,25,26,50}, T. Matsumura⁹, M. Migliaccio^{30,31}, L. Montier², G. Morgante²³, B. Mot², L. Mousset², R. Nagata⁴², T. Namikawa⁹, F. Nati^{14,15}, P. Natoli^{21,22}, S. Nerval⁵¹, A. Novelli¹², L. Pagano^{21,22,49}, A. Paiella^{12,13}, D. Paoletti^{23,40}, G. Pascual-Cisneros¹⁶, G. Patanchon⁶, V. Pelgrims^{52,53}, F. Piacentini^{12,13}, G. Piccirilli³⁰, G. Polenta⁴⁸, G. Puglisi^{30,34}, N. Raffuzzi²¹, A. Ritacco^{54,49,47}, J. A. Rubino-Martin^{7,8}, G. Savini⁵⁵, D. Scott⁴⁴, Y. Sekimoto^{42,56,37}, M. Shiraishi⁵⁷, G. Signorelli⁵⁸, S. L. Stever^{59,9}, N. Stutzer¹, R. M. Sullivan⁴⁴, H. Takakura^{56,42}, L. Terenzi²³, H. Thommesen¹, M. Tristram⁶⁰, M. Tsuji⁶¹, P. Vielva¹⁶, J. Weller⁶², B. Westbrook³², G. Weymann-Despres⁶⁰, E. J. Wollack⁶³, and M. Zannoni^{14,15}
(LiteBIRD Collaboration)

(Affiliations can be found after the references)

Received 15 February 2023 / Accepted 10 May 2023

ABSTRACT

LiteBIRD is a planned JAXA-led cosmic microwave background (CMB) B -mode satellite experiment aiming for launch in the late 2020s, with a primary goal of detecting the imprint of primordial inflationary gravitational waves. Its current baseline focal-plane configuration includes 15 frequency bands between 40 and 402 GHz, fulfilling the mission requirements to detect the amplitude of gravitational waves with the total uncertainty on the tensor-to-scalar ratio, δr , down to $\delta r < 0.001$. A key aspect of this performance is accurate astrophysical component separation, and the ability to remove polarized thermal dust emission is particularly important. In this paper we note that the CMB frequency spectrum falls off nearly exponentially above 300 GHz relative to the thermal dust spectral energy distribution, and a relatively minor high frequency extension can therefore result in even lower uncertainties and better model reconstructions. Specifically, we compared the baseline design with five extended configurations, while varying the underlying dust modeling, in each of which the High-Frequency Telescope (HFT) frequency range was shifted logarithmically toward higher frequencies, with an upper cutoff ranging between 400 and 600 GHz. In each case, we measured the tensor-to-scalar ratio r uncertainty and bias using both parametric and minimum-variance component-separation algorithms. When the thermal dust sky model includes a spatially varying spectral index and temperature, we find that the statistical uncertainty on r after foreground cleaning may be reduced by as much as 30–50% by extending the upper limit of the frequency range from 400 to 600 GHz, with most of the improvement already gained at 500 GHz. We also note that a broader frequency range leads to higher residuals when fitting an incorrect dust model, but also it is easier to discriminate between models through higher χ^2 sensitivity. Even in the case in which the fitting procedure does not correspond to the underlying dust model in the sky, and when the highest frequency data cannot be modeled with sufficient fidelity and must be excluded from the analysis, the uncertainty on r increases by only about 5% for a 500 GHz configuration compared to the baseline.

Key words. ISM: general – cosmology: observations – cosmic background radiation – polarization – cosmological parameters – Galaxy: general

1. Introduction

One of the key predictions of the current cosmological inflationary paradigm is the existence of a stochastic background of primordial gravitational waves created shortly after the Big Bang (Starobinsky 1980; Sato 1981; Guth 1981; Albrecht & Steinhardt 1982; Linde 1982, 1983). If such waves

do exist, they should induce a particular and unique polarization signature in the cosmic microwave background (CMB) on large angular scales, corresponding to so-called B -mode or divergence-free polarization (Kamionkowski et al. 1997). Detecting this imprint in the CMB ranks among the top observational priorities in modern cosmology, and huge efforts are currently being made to develop the necessary instrumentation, data analysis techniques, and theoretical modeling required for this task.

* Corresponding author: U. Fuskeland,
e-mail: unnif@astro.uio.no

Among the leading international *B*-mode efforts is LiteBIRD (the Lite (Light) satellite for the study of *B*-mode polarization and Inflation from cosmic background Radiation Detection), a satellite concept chosen by JAXA as one of their top priorities for the coming decade, aiming for a launch from 2028 to 2029 (Hazumi et al. 2020; LiteBIRD Collaboration 2022). The development of LiteBIRD has already been ongoing for more than 13 years, evolving gradually from a relatively simple and light-weight concept that originally included only six frequency channels between 60 and 280 GHz (Hazumi et al. 2012; Matsumura et al. 2014) into the current baseline configuration spanning 15 frequency bands between 40 and 402 GHz (Hazumi et al. 2020). This increased sophistication came as a direct response to observational insights gained regarding the astrophysical sky as measured by CMB experiments such as *Planck* (Planck Collaboration I 2020) and BICEP2/Keck (Ade et al. 2016). These observations pointed toward a complex reality in which both polarized synchrotron and thermal dust emission must be modeled and subtracted with exquisite precision, while simultaneously accounting for unpolarized complications such as carbon monoxide and time-dependent zodiacal light emission.

In the following, we revisit the upper limit of the LiteBIRD frequency range, and examine its implications for the instrument’s ability to distinguish between CMB and astrophysical foregrounds. While LiteBIRD Collaboration (2022) has already demonstrated that the baseline configuration fully achieves the primary mission goal – namely to constrain the tensor to scalar ratio with the total uncertainty, δr , down to $\delta r < 0.001$ – it is still desirable to optimize the uncertainties from component separation, which corresponds to one-third of the total budget error on r , in order to increase the overall mission margins. Exploring that is the main motivation for the current work. At the same time, we emphasize that this paper only considers the raw sensitivity and component separation aspects of the instrumental design. We need to perform a comprehensive trade study before the collaboration can decide on a baseline change. The scope of the current paper is simply to determine whether sufficiently significant gains are available to warrant such a study.

To this end, we compared the current baseline configuration with five alternative configurations in which the upper limit was varied between 400 and 600 GHz in steps of 40 GHz. The absolute upper center frequency limit of 600 GHz was dictated by the current bolometer design, which has a hard cutoff at 680 GHz (Novotny & Meincke 1975). To additionally minimize the number of changes required for the overall satellite design, we chose to only modify one of the three LiteBIRD telescopes, namely the High-Frequency Telescope (HFT); the Low- and Mid-Frequency Telescopes (LFT and MFT) have been left unchanged. Frequencies up to 600 GHz can be considered only because LiteBIRD is a space mission, while ground-based experiments are restricted to frequencies below 300 GHz due to the atmospheric windows (Liebe 1981; Pardo et al. 2001).

The main question considered in this paper is the quantitative relationship between the frequency range and the tensor-to-scalar ratio uncertainty marginalized over polarized thermal dust emission. Previous analyses (e.g., Remazeilles et al. 2016) have suggested that a low cutoff on the high frequency side only supports very limited ability to constrain the spectral energy distribution (SED) of thermal dust emission. At the same time, it is well known that the CMB SED falls nearly exponentially above 300 GHz relative to the thermal dust SED (see e.g., Fig. 35 of Planck Collaboration IV 2020), and therefore even relatively small changes in the cutoff frequency may have a

dramatic impact on the confusion between CMB and thermal dust. Quantifying the importance of including frequencies higher than where the CMB SED drops off in a realistic setting is the main goal of the current paper. We note that while we subsequently address LiteBIRD in particular, the main arguments are instrument-agnostic, and the primary conclusions are therefore generally applicable to any future space mission or balloon-borne experiment. Indeed, similar analyses have recently been published for both the PICO (Aurlien et al. 2023) and ECHO (Sen et al. 2022) CMB satellite concepts with consistent results.

This challenge of optimizing the frequency range may be separated into two important and complementary aspects, namely optimizing our sensitivity to the tensor-to-scalar ratio overall and maximizing our ability to reject wrong sky models. The following analysis is organized accordingly. To address the first question, we analyze ideal foreground simulations in which the fitting model matches the simulated sky, and estimate the tensor-to-scalar ratio uncertainty as a function of frequency coverage. In this case, we adopt a standard one-component modified blackbody (MBB) thermal dust model with spatially varying spectral parameters, following closely in the footsteps of LiteBIRD Collaboration (2022). In order to ensure robustness in terms of analysis-dependent details, we employed five independent component-separation methods for this task, three of which implement different variations of parametric fitting (Commander – Eriksen et al. 2008, FGBuster – Errard & Stompor 2019, and Moment Expansion – Vacher et al. 2022), while the last two implement (semi)blind internal linear combination (ILC) fitting (Needlet ILC, NILC and constrained moment ILC, cMILC – Remazeilles et al. 2011, 2021).

To assess the capability of rejecting incorrect data models, we analyzed simulations based on an extended thermal dust model, namely the physically motivated silicate-iron-carbon model (SiFeC; Hensley & Draine 2017). We evaluated the residual χ^2 , and compared this with the derived tensor-to-scalar estimates. For this study, we note that only Commander currently provides χ^2 -based goodness-of-fit estimates, and we therefore only reported results for that code in this part of the paper. However, we note that work is currently ongoing on implementing similar statistics for other methods, and those results will be reported elsewhere.

It is also important to note that the foreground sky in this paper was approximated as a “single layer model,” and no attempt was made to take into account the full 3D structure of the Milky Way. Full line-of-sight integration of 3D effects would introduce additional frequency-decorrelation effects (Tassis & Pavlidou 2015; Chluba et al. 2017; Pelgrims et al. 2021; Ritacco et al. 2023; Vacher et al. 2023a), and the following results therefore represent an optimistic view of the true sky in terms of complexity. In reality, there will be additional structures, both spatially and spectrally, and fully resolving these would in principle require even more data than considered here, for instance in terms of additional frequency bands, wider frequency range, or dedicated 3D constraints on the Galactic magnetic field. The following optimization results should therefore be considered as a lower bound on what is actually needed for a future production analysis.

The paper is organized as follows. In Sect. 2 we provide a brief overview of the LiteBIRD concept and current baseline, and we define a set of alternative configurations to be considered. In Sect. 3 we describe the simulations used for the present analysis, and in Sect. 4 we briefly survey the component-separation methods used in this paper; algorithmic details are deferred to the Appendices. In Sect. 5, we present the results derived from

Table 1. Summary of LiteBIRD HFT instrument configurations considered in this paper in terms of frequency, beam size, and sensitivity.

Parameter	Model	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5
Frequency (GHz)	M0	195.0	235.0	280.0	337.0	402.0
	M1	214.5	258.5	308.0	370.7	442.2
	M2	234.0	282.0	336.0	404.4	482.4
	M3	253.5	305.5	364.0	438.1	522.6
	M4	273.0	329.0	392.0	471.8	562.8
	M5	292.5	352.5	420.0	505.5	603.0
Beam FWHM (arcmin)	M0	28.6	24.7	22.5	20.9	17.9
	M1	27.7	24.3	22.0	20.5	17.5
	M2	26.0	23.1	21.2	19.9	17.1
	M3	24.6	22.1	20.6	19.4	16.5
	M4	23.5	21.4	20.1	18.6	15.6
	M5	22.6	20.8	19.7	17.1	13.9
NET array ($\mu\text{K} \sqrt{\text{s}}$)	M0	5.2	5.3	6.8	11	24
	M1	5.9	6.4	9.2	16	37
	M2	6.7	8.0	12	23	60
	M3	7.8	9.9	16	34	100
	M4	9.1	13	22	50	170
	M5	11	16	30	75	280

Notes. M0 denotes the baseline configuration, while M_i indicates a hypothetical instrument for which each HFT frequency range is scaled by a constant factor of $(1 + i/10)$.

an ideal model analysis, while in Sect. 6 we consider modeling errors. Final conclusions are drawn in Sect. 7.

2. Baseline and extended LiteBIRD instrument configurations

In this section, we discuss the LiteBIRD mission characteristics that are relevant to the present analysis. A comprehensive review of the instrument concept can be found in [LiteBIRD Collaboration \(2022\)](#). First, the LiteBIRD detector arrays are populated with a total of 4508 lenslet- or horn-coupled transition edge sensor (TES) bolometers fabricated on silicon wafers, covering 15 frequency bands between 40 and 402 GHz. The LiteBIRD bolometers provide exquisite sensitivity across this full frequency range, but the technology itself is limited to frequencies <678 GHz, corresponding to the superconducting gap of niobium – the material used for mm-wave transmission lines on the detectors ([Novotny & Meincke 1975](#)). This value therefore naturally defines a clear upper limit for the feasible frequency range of the LiteBIRD instrument in any modification to the baseline.

As a second consideration, to ensure robust separation between the actual polarization signal and a wide range of intensity-specific contaminants, LiteBIRD will employ half-wave plates (HWPs) spinning at a rate of 0.5 to 0.8 Hz. Important strengths of mesh filter-based HWPs are low mass requirements and high transmission. However, an important drawback is finite effective bandwidth, typically corresponding to a ratio between its maximum and minimum effective frequency of about 2.3 ([LiteBIRD Collaboration 2022](#)). As such, the full relative LiteBIRD frequency range of $402 \text{ GHz}/40 \text{ GHz} \approx 10$ cannot be supported within a single telescope, but rather three separate telescopes are required. For LiteBIRD these are called the Low-, Mid-, and High-Frequency Telescopes, respectively, or LFT, MFT, and HFT as a short-hand. While this organization comes at a significant cost in terms of system complexity, it also allows significant optimization of each system, and different technologies may be used as appropriate.

In total, the three telescopes span 22 independent frequency channels, covering 15 frequency bands; there is an overlap of three channels inside the LFT, three channels between the LFT and MFT, and an overlap of one channel between the MFT and HFT. These overlaps represent important cross-checks against systematic errors in any of the three telescopes. Each frequency band overlaps the adjacent bands by approximately half the bandwidth. The radiation is detected by seven different wafer types, each containing one or two types of multichromatic pixels. Each pixel, for each polarization state, distributes power through one, two, or three different bandpass filters to independent TES detectors. The HFT detectors are implemented in terms of three different detector modules; one with 3-color pixels, one with 2-color pixels, and the highest frequency band on single-color pixels ([Montier et al. 2020](#)).

The main goal of the current paper is to investigate the optimal LiteBIRD frequency range with respect to the removal of the thermal dust from the CMB polarization signal. Given the mature state of the current baseline, modifications must be realistic and the scientific improvements for doing this must be substantial. One of the parameters that may still be adjusted, however, is the effective HFT frequency range. Specifically, it is possible to shift the entire frequency range by a constant factor, which ensures that the relative factor between the highest and lowest frequency bands is unchanged, as required by the mesh-filter HWP. In this paper, we therefore defined a set of extended HFT configurations parametrized by a scaling factor, f , that varied between 1.0 and 1.5 in steps of 0.1, and shifted the entire HFT frequency range. These models were denoted M_i , such that M0 corresponded to the baseline, while M5 corresponded to an extended model with a highest frequency of $1.5 \times 402 \text{ GHz} \approx 603 \text{ GHz}$. For each modified frequency channel, we calculated new instrumental parameters, such as sensitivity or beam size, using the same methodology as used for the baseline in [LiteBIRD Collaboration \(2022\)](#). This included evaluating the system temperature through ray tracing from the feed to terminals, either inside the telescope or the sky. The beam full width at half maximum (FWHM) values were estimated via the

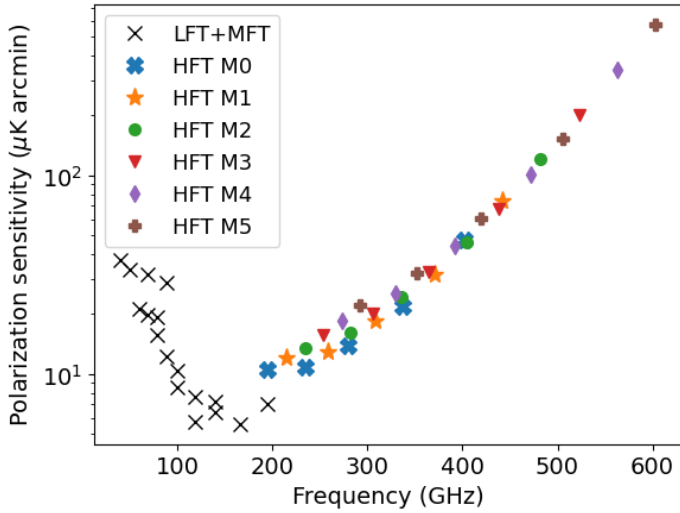


Fig. 1. Sensitivity versus frequency for the various LiteBIRD configurations listed in Table 1. The black crosses indicate LFT or MFT frequencies that are not modified in this paper, while colored symbols indicate the various modified HFT frequencies.

edge taper at the telescope aperture, as calculated from the feed far-field pattern. The resulting values are summarized numerically in Table 1, while Fig. 1 shows the sensitivity as a function of frequency for the extended bands.

It is interesting to note that the sensitivity of two channels in different configurations can vary significantly, despite the fact that both their center frequencies and detector counts are nearly identical; a typical example of this is channel 1 in M2 and channel 2 in M0. The reason for this is that the optics for the HFT are optimized for the telescope’s full frequency bandwidth, not any single channel. Frequencies near the edge of the telescope bandwidth typically have poorer optical absorption and reflection properties than frequencies nearer the center of the band.

We also note that the LFT and MFT are left unchanged in this exercise, and their instrumental parameters are not listed in Table 1. The primary motivation for this choice is programmatic; changing the baseline at this late stage will carry a nonnegligible cost, and it is therefore highly desirable to make as few changes as possible. At the same time, it is important to note that this choice does come at a cost in the form of weaker systematics cross-check between the MFT and HFT, since the two telescopes no longer have overlapping frequency ranges. Thus, these modified configurations may be considered as trading off robustness toward “unknown unknowns” (for instance, subtle HWP frequency effects) against robustness to a known unknown (polarized thermal dust emission). If this issue is considered sufficiently important at a programmatic level, it is of course also possible to modify the lower-frequency telescopes, albeit at a somewhat higher cost.

3. Simulations

The LiteBIRD simulations considered in this paper are generated directly in pixel space, and account only for sky signal and white noise. The sky emission components for all 22 frequency channels are generated using the PySM (Python Sky Model; Thorne et al. 2017) package. The LFT and MFT channels are defined according to the baseline instrument configuration in LiteBIRD Collaboration (2022), while the HFT channels are generated in six versions, called M0–M5, as tabulated in Table 1. Here, M0 is the LiteBIRD baseline, and M1–M5 are

the frequency extended data sets. The instrumental summary parameters listed in this table are derived from first principles using the same methodology as the baseline configuration in LiteBIRD Collaboration (2022), properly accounting for realistic optics, bandpass, and bolometer effects.

We simulate the polarized microwave emission considering three different components, namely a cosmological CMB signal and both diffuse synchrotron and thermal dust radiation from our own Galaxy. The CMB component is simulated as a Gaussian random realization of the *Planck* 2018 Λ CDM best-fit power spectrum (Planck Collaboration VI 2020). In the following, we set the tensor-to-scalar ratio, r , to zero. Obviously, a main goal of LiteBIRD is to actually measure a non-zero value of r , and it could therefore be of interest to also consider non-zero values of this parameter. However, this paper is primarily concerned with the relative performance of different instrument configurations. This comparison is not sensitive to the exact value of r , and then for the purpose of our analyses we decided to use $r = 0$ as reference, which also provides the benefit of not having to deal with the cosmic variance noise term. Future work focusing on absolute detection level performance should obviously revisit this with a range of alternative cosmological parameter values, including r , the optical depth of reionization τ , and others.

Galactic synchrotron radiation is physically generated by cosmic-ray electrons that are accelerated by the Galactic magnetic field. This signal is highly polarized with a polarization fraction that can reach 20% at intermediate and high Galactic latitudes (Planck Collaboration XXV 2016; Kogut et al. 2007). Synchrotron emission represents the dominant foreground at frequencies below about 70 GHz (Krachmalnicoff et al. 2016; Planck Collaboration IX 2016). As a first approximation, the synchrotron SED follows a power law with a spatially varying spectral index (Lawson et al. 1987; Platania et al. 1998; Bennett et al. 2003). Recent analyses support a spectral index of $\beta_s \approx -3$ (Fuskeland et al. 2014, 2021), with variations on the order of 10% across the sky on degree angular scales (Krachmalnicoff et al. 2018). In the following, we adopt the s1 PySM model for synchrotron emission, corresponding to a strict power-law spectral behavior with spatially varying spectral index. We note that synchrotron emission is largely irrelevant for the highest LiteBIRD frequencies, and any conclusions regarding the optimal HFT frequency range will be largely independent of the synchrotron model.

The same does, however, not hold true for thermal dust emission, which is the dominant foreground component at frequencies above 70 GHz (Krachmalnicoff et al. 2016; Planck Collaboration IX 2016). Thermal dust emission is generated by vibrating nonspherical dust grains in the interstellar medium (ISM), and is polarized because the small axes of the grains are preferentially aligned parallel to the local magnetic field (Hoang & Lazarian 2016; Hensley & Draine 2023). The detailed physical processes involved in this are, however, more complicated than for synchrotron emission, and the predicted emission from any region of the ISM will depend on a wide range of local properties, including dust grain composition, the local radiation field and magnetic field structure. For robustness, we therefore consider two different thermal dust SED models in the following, namely; (1) a phenomenologically motivated single modified blackbody parametrized by a spatially varying spectral index and temperature (1 MBB; PySM model d1); and (2) a physically motivated silicate-iron-carbon model (SiFeC; PySM model d8; Hensley & Draine 2017). The 1 MBB model is very often used in the literature. The second model is more exotic, and is physically motivated rather than data-driven.




Fig. 2. SEDs of the two dust models considered in this work, the 1 MBB, (red line) and the physical dust model SiFeC, (green line). The black line shows the SED of the CMB, while the gray refers to synchrotron emission. All the SEDs are in brightness temperature units and are computed considering the rms of Stokes Q and U maps on 70% of the sky. Sky masks are obtained by retaining the cleanest fraction of the sky in polarized intensity at 100 GHz. Vertical gray lines show the central frequency of the baseline LiteBIRD frequency channels, and in blue dashed lines we report the central frequency of the HFT M5 extension. The bottom panel shows the percentage difference of SiFeC with respect to MBB.

They also differ in complexity in terms of spatial variation. The d1 model has both a spatially varying spectral index and temperature, while d8 has spatially constant spectral parameters. Future work should also consider models that include physically motivated frequency decorrelation, for instance by implementing proper multi-layer 3D models. This is however beyond the scope of the current paper.

Figure 2 shows the dust SEDs for the two different models computed on 70% of the sky. Sky masks are obtained by retaining the cleanest fraction of the sky in polarized intensity at 100 GHz. As seen in the bottom panel, the relative difference between the two models can reach tens of percent across the LiteBIRD frequency range, and properly accounting for these variations will be essential for making a robust CMB extraction. For comparison, the gray and black curves show the synchrotron and CMB SEDs. The near exponential drop-off of the CMB spectrum is clearly seen above about 300 GHz, leading to a rapidly increasing ratio between thermal dust and CMB emission between 400 and 600 GHz.

Instrumental effects are modeled under highly idealized assumptions, since these are considered to be subdominant regarding the central question of frequency range versus thermal dust reconstruction, and rather only add computational complexity and cost. In particular, we model all instrumental bandpasses in terms of Dirac δ functions, all instrumental beams as azimuthally symmetric Gaussians, and the noise is assumed to be uncorrelated, Gaussian, and spatially isotropic. Center frequencies, beam FWHMs, and array noise equivalent temperatures (NETs) are all listed in Table 1. Under these assumptions, the simulation procedure is very straightforward, and does not require any low-level time-domain processing.

The data model for a given frequency band ν may be described as

$$\mathbf{d}_\nu = \mathbf{s}_\nu + \mathbf{n}_\nu, \quad (1)$$

where \mathbf{d} represents a given (simulated) sky map, \mathbf{s} denotes the sky signal, and \mathbf{n} is instrumental noise. The simulated sky maps analyzed in this paper are based on two different thermal dust emission models. For the baseline model, we only fit the one-component MBB model, while the SiFeC model is used to assess modeling errors and is therefore only present in the simulations, and not in the data model in the methods. Including also CMB and synchrotron emission, Eq. (1) may then be written in the following explicit form, adopting thermodynamic CMB temperature units,

$$\mathbf{d}_\nu = \mathbf{a}_{\text{CMB}} \quad (2)$$

$$+ \mathbf{a}_s(\nu) \gamma(\nu) \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s} \quad (3)$$

$$+ \mathbf{a}_d(\nu) \gamma(\nu) \frac{e^{\frac{h\nu_{0,d}}{kT_d}} - 1}{e^{\frac{h\nu}{kT_d}} - 1} \left(\frac{\nu}{\nu_{0,d}} \right)^{\beta_d+1} \quad (4)$$

$$+ \mathbf{n}_\nu, \quad (5)$$

where $\{\mathbf{a}_{\text{CMB}}, \mathbf{a}_s, \mathbf{a}_d\}$ are the signal amplitudes relative to the reference frequency ν_0 for each signal, $\gamma(\nu)$ is the conversion factor between the Rayleigh-Jeans brightness temperature and the CMB thermodynamic temperature, $\{\beta_s, \beta_d\}$ are the spectral indices, T_d is the temperature for thermal dust, and h and k are the Planck and Boltzmann constants, respectively. In principle, this sky model is valid for both intensity and polarization, with individual parameters fitted for each Stokes parameter. In practice, however, we only include CMB, synchrotron and thermal dust emission in polarization, and we additionally assume $\beta_Q = \beta_U$.

All simulations are generated both at full angular resolution (adopting a HEALPix pixelization with a resolution parameter of $N_{\text{side}} = 512$; Górski et al. 2005) and at low angular resolution (10° FWHM, $N_{\text{side}} = 16$). The former set is used by the Moment expansion, FGBuster and NILC/cMILC analyses, while the latter is used by the Commander analysis. We note that the low-resolution maps do not include subpixel (or sub-beam) structure, but are generated natively at the target resolution. The motivation for this choice is that sub-pixel structure may excite spurious B -mode power, both from beam averaging and from parallel transport inaccuracies during HEALPix downgrading. At the same time, a future LiteBIRD Commander analysis will be performed in the time-domain with full angular resolution data, following a methodology similar to that described by BeyondPlanck Collaboration (2023), and this complication will therefore not be relevant for the final LiteBIRD analysis.

4. Component-separation algorithms

In order to ensure that the general results are robust with respect to algorithmic details, we employ a total of five different component-separation algorithms in the following. These may be divided into two main groups, namely parametric (Commander, FGBuster, and Moment Expansion) and (semi)blind methods (NILC and cMILC).

Commander (Eriksen et al. 2008) is a Bayesian Gibbs sampler that has been successfully applied to Planck, WMAP, LiteBIRD, and many other data sets (e.g., Planck Collaboration IX 2016). The defining feature of this method is an explicit joint parametric model that simultaneously accounts for cosmological and astrophysical parameters, and this global model is explored by textbook Markov chain Monte Carlo methods. A major strength of this method is the availability of well-defined

goodness-of-fit χ^2 statistics and realization-specific uncertainty estimates, while a major weakness is a high computational cost. In this paper we only apply this method to low-resolution simulations; for further details, see Appendix A.1.

FGBuster (Puglisi et al. 2022) is based on a similar statistical foundation as Commander, but uses nonlinear optimization to explore the likelihood function rather than Monte Carlo sampling. It further speeds up the estimation process by first estimating all nonlinear spectral foreground parameters marginally with respect to linear amplitude parameters; then it estimates amplitudes conditionally with respect to the spectral parameters. Finally, angular power spectra are computed from the amplitude maps. This method has been used extensively for LiteBIRD forecasting and optimization (e.g., LiteBIRD Collaboration 2022). For further details, see Appendix A.2.

The moment-expansion method (Chluba et al. 2017; Vacher et al. 2023b) also has a similar statistical foundation as Commander and FGBuster but it allows for Taylor-expansion-based distortions in the SED models for each foreground to account for nonlinear averaging effects, for instance from line-of-sight integration, beam convolution or harmonic expansion. This method has already been applied to baseline LiteBIRD simulations by Vacher et al. (2022). For further details, see Appendix A.3.

While each of the above methods is based on explicit and nonlinear parametric foreground models, the Needlet Internal Linear Combination (NILC; Delabrouille et al. 2009) method takes a fundamentally different approach, and assumes simply that the total foreground signal may be written as a linear combination of the true CMB signal and the total foreground signal at each frequency channel, and then form the linear combination of frequency maps that minimizes the variance of the final product, see Appendix B.1. To account for possible spatial variations in the foreground SED, the linear combination weights are computed separately in needlet space, allowing for localized optimization both in terms of sky position and angular scale. This method has been applied to a wide range of data sets, including *Planck* (e.g., Planck Collaboration IX 2016) and LiteBIRD. For further details, see Appendix B.2.

Finally, we consider the more recently developed constrained moments ILC method (cMILC; Remazeilles et al. 2021), which combines the moment-expansion and NILC methods. Specifically, additional constraints are imposed on the ILC weights that explicitly cancels out individual foreground components through their Taylor-expanded SEDs. Each moment corresponds to one additional linear constraint in the ILC solution, and the combination of foreground-specific and the variance constraints is solved through a single Lagrange multiplier system. For further information, see Appendix B.3.

5. Ideal model analysis: Estimation of statistical uncertainties

When contemplating a major modification of a satellite’s baseline design, there are at least two aspects that must be considered carefully. The first aspect is sensitivity – we must investigate how much stronger (or weaker) constraints on the tensor-to-scalar ratio the proposed modification will lead to. The second aspect is goodness-of-fit – we must also investigate whether the proposed modification will affect our sensitivity to modeling errors or our ability to detect such errors. These two main aspects are addressed individually in this and the next section.

Starting with the sensitivity aspect, we approach this by analyzing the simulations summarized in Sect. 3 for each of the

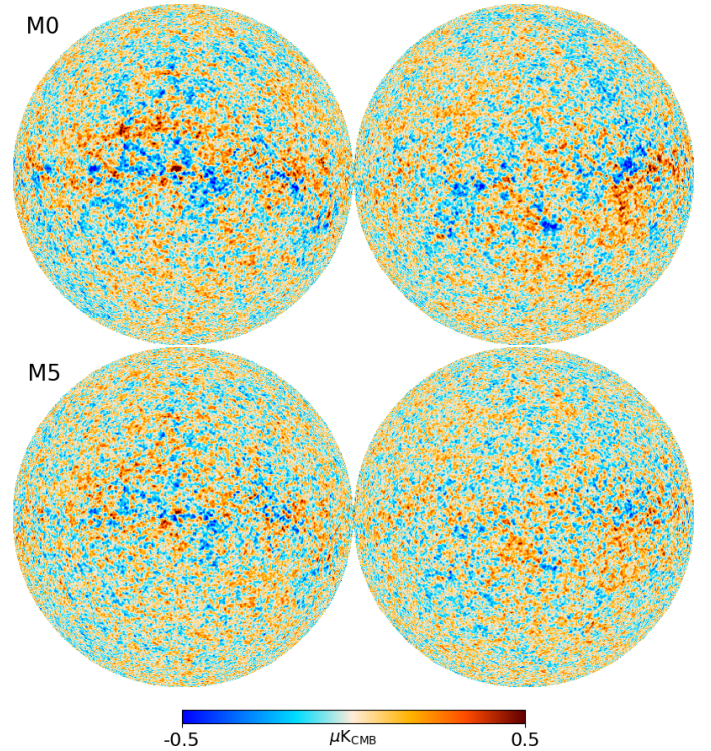


Fig. 3. Comparison of the reconstructed CMB B -mode maps from NILC for the M0 (top row) and M5 (bottom row) instrument configurations. The two panels in each row show opposite hemispheres aligned with the Galactic plane along the equator, and the north and south Galactic poles at the top and bottom, respectively. The left and right panels are centered on the Galactic center and anticenter.

instrument configurations described in Sect. 2. In this section, we primarily focus on the simplest foreground model, namely the one-component MBB thermal dust model, for which modeling errors are minor (it is worth noting that they are by no means nonexistent, since both informative priors and spatial variations in the fitted parameters can lead to biases).

We also note that each of the five component-separation methods discussed in Sect. 4 have their particular strengths and weaknesses, and our goal in this paper is not to perform a head-to-head component-separation algorithm comparison, but it is rather to assess the capabilities and limitations of the LiteBIRD frequency selection itself. In the following, we therefore present selected results from among the five methods, depending on which product is most convenient for a particular application. In general, however, we note that all five methods provide qualitatively similar results.

5.1. Map residuals

First, to build intuition regarding the overall impact of the frequency range, we show in Fig. 3 the reconstructed B -mode NILC¹ map for one arbitrarily selected simulation and each of the two most extreme instrument configurations, M0 (top row) and M5 (bottom row). While residual foreground contamination is clearly visible in the Galactic disk region of the CMB B -mode map for M0, we see that these residuals are significantly reduced for the M5 configuration, with a wider frequency range.

¹ Note that NILC, unlike the parametric methods, is defined natively in EB space rather than Stokes QU space.




Fig. 4. Commander residual (output minus input) maps for selected parameters and instrument configurations for one arbitrarily selected Monte Carlo sample. Columns show, from left to right: (1) CMB Stokes Q ; (2) CMB Stokes U ; (3) thermal dust spectral index; and (4) thermal dust temperature. Rows show M0, M2, and M5. The maps show one typical realization. Gray pixels indicate the masked pixels with $f_{\text{sky}} = 90\%$.




Fig. 5. Power spectrum of residual foreground contamination (left panel) and noise (right panel) with NILC (solid lines) and cMILC (dashed lines) on $f_{\text{sky}} = 50\%$ of the sky for the M0 (blue) and M5 (orange) configurations. For reference, we show the input CMB B -mode signal ($r = 0$, black dash-dotted line), the primordial B -mode signal expected from theory for values ranging from $r = 10^{-2}$ down to $r = 10^{-3}$ (gray-shaded area), and the residual lensing cosmic variance (yellow-shaded area) for no delensing ($A_L = 1$) up to 90% delensing ($A_L = 0.1$). All the spectra are binned with a multipole bin size of $\Delta\ell = 16$.

A complementary view from the parametric Commander code is shown in Fig. 4 in the form of residual (output-minus-input) maps for four key quantities, namely the Stokes Q and U CMB maps and the thermal dust spectral index and temperature for the M0, M2, and M5 instrument configurations². In all cases, we see that the quality of the fit improves with increasing maximum frequency. In the CMB maps, this is most clearly seen in the form of reduced scatter around the Galactic plane, while for T_d the dark blue area in the high Galactic latitudes systematically fade from M0 to M5; correspondingly, the dark red area in β_d also gradually fade away. These observations may be understood intuitively by noting that an increased frequency

² For visualization purposes, the gray pixels in this figure indicate a 90% analysis mask, which is significantly smaller than the default 73% confidence mask adopted by Commander for the likelihood analysis.

range improves the ability to break the well-known degeneracy between β_d and T_d (e.g., Juvola & Ysard 2012) that arises when fitting for them jointly.

5.2. Power spectrum uncertainties

Next, we consider the impact of changing the instrument configuration from M0 to M5 in terms of residual foreground and noise angular power spectra. This is shown in Fig. 5 for both NILC (solid lines) and cMILC (dashed lines). The left panel shows the projected foreground power spectrum, while the right panel shows the noise sample variance. First, we see that cMILC provides a significantly lower residual foreground contamination than NILC across a wide range of multipoles for both instrument configurations. This is due to the additional deprojection

of foreground moments, which is more efficient at suppressing true foreground residuals. However, this additional foreground reduction performance does come at a significant price in the form of higher noise uncertainties caused by the same additional degrees of freedom. This noise penalty is greatly reduced when extending the LiteBIRD frequency range from 402 to 603 GHz, to the point that it is comparable to the CMB cosmic variance. In addition, the overall level of residual foreground contamination is also greatly reduced for the extended frequency range, both for NILC and cMILC, across all multipoles.

5.3. Tensor-to-scalar ratio constraints

We now turn our attention to tensor-to-scalar ratio constraints, and we start with probability distributions as derived using the NILC and cMILC methods. These are summarized in Fig. 6³. These distributions are derived by averaging over an ensemble of simulations, and therefore correspond inherently to ensemble averages. The top and bottom panels show the results for the M0 and M5 configurations, respectively. Solid and dotted lines show recovered tensor-to-scalar ratio distributions for NILC and cMILC, while red and black lines show results with and without delensing.

Since the fiducial value of the tensor-to-scalar ratio in these simulations is $r = 0$, any apparent bias ($r \neq 0$, defined by the peak of the likelihood) is due to the power spectrum of the residual foreground contamination that projects into the CMB B -mode maps shown in Fig. 3. In contrast, the width of the distribution, that is σ_r , receives contributions from the cosmic variance of the lensed CMB signal, the sample variance of projected foregrounds, and noise.

Starting with the unlensed NILC results (solid black lines), we first note that both M0 and M5 have a well-defined nonzero peak, which is indicative of a nonnegligible foreground residual, consistent with the nonnegligible foreground power seen in the left panel of Fig. 5. For the delensed case (red curve), this spurious detection is statistically significant at the 5σ level for M0, while for the lensed case it is nearly consistent with zero at the 2σ level; however, this difference is simply due to the additional uncertainty added by lensing, and not lower foreground residuals as such. For M5, the overall NILC bias is reduced by about 30% in both cases, while the uncertainty remains unchanged, since this is dominated by cosmic variance.

Being a more constrained version of the ILC, the semi-blind cMILC method allows further foreground deprojection, and thereby suppresses some of the residual biases that NILC suffers from. Specifically, by assuming an MBB $f_{\text{dust}}(\nu) = \nu^{\bar{\beta}_d+1} / (e^{h\nu/k\bar{T}_d} - 1)$ with pivots $\bar{\beta}_d = 1.5$, $\bar{T}_d = 20$ K and a power-law $f_{\text{sync}}(\nu) = \nu^{\bar{\beta}_s}$ with pivot $\bar{\beta}_s = -3$ as the baseline zeroth-order SEDs for dust and synchrotron, cMILC imposes four nulling constraints to deproject the zeroth-order moments of dust and synchrotron and the first-order moments of dust, whose respective SEDs are $f_{\text{sync}}(\nu)$, $f_{\text{dust}}(\nu)$, $\partial f_{\text{dust}}(\nu) / \partial \bar{\beta}_d$, and $\partial f_{\text{dust}}(\nu) / \partial \bar{T}_d$.

The recovered distributions of the tensor-to-scalar ratio from cMILC are shown as dotted lines in Fig. 6. In contrast to NILC, cMILC shows unbiased recovery of $r = 0$ thanks to deprojection of foreground moments. However, for M0 the extra cMILC constraints also significantly increase the noise contribution to σ_r .

³ Although r posterior plots are shown in terms of a logarithmic x -axis, the actual distributions are defined in linear units; this choice is made for plotting purposes only.

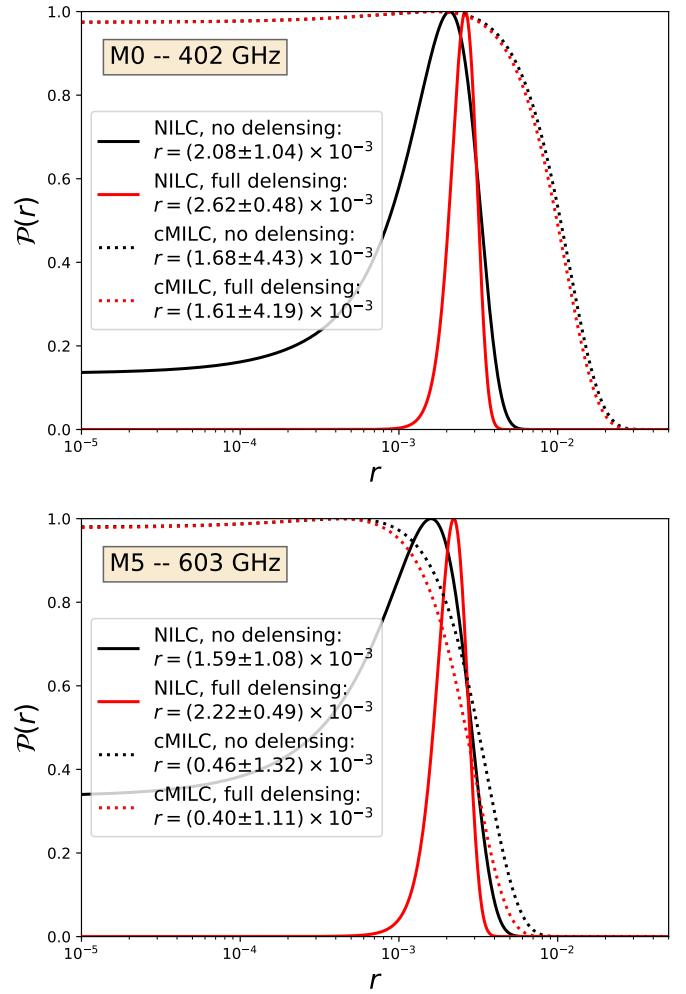


Fig. 6. Comparison of recovered tensor-to-scalar ratio distributions from NILC (solid lines) and cMILC (dotted lines) for the M0 (top panel) and M5 (bottom panel) instrument configurations. Results without delensing are shown as black curves, while results with full delensing are shown as red lines.

On the other hand, for M5 cMILC provides unbiased recovery of $r = 0$, with a negligible increase in noise. Thus, while M0 does not provide enough constraining power to allow deprojection of thermal dust temperature moments at a useful level, M5 does support this.

All the above results correspond to ensemble-averaged distributions derived with blind internal linear combination methods. To understand the behavior for individual realizations, Fig. 7 shows corresponding posterior distributions as derived with Commander for 20 realizations, plotted as thin lines for each of the six instrument configurations. No delensing is applied in any of these cases, and a uniform prior on $r > 0$ is assumed. The red vertical lines show the mean of the upper 68% confidence limits. Firstly, we note that these distributions appear (at least visually) statistically consistent with the input value of $r = 0$. Secondly, we also see that the upper limit decreases monotonically with instrument configuration, with M5 providing the tightest overall constraints. Thirdly, the internal scatter between individual realizations is notably smaller with a higher maximum frequency, resulting in an overall lower sample variance.

These general observations are summarized more quantitatively in Fig. 8, which shows σ_r as a function of instrument




Fig. 7. Tensor-to-scalar ratio posterior distributions for the single MBB component model (for both simulating and modeling) as derived with *Commander* using multipoles $\ell = [2, 12]$ and a sky fraction of $f_{\text{sky}} = 73\%$ for 20 independent simulations (black lines), each drawn from a Λ CDM spectrum with $r = 0$. The red vertical lines correspond to the mean upper 68% confidence limit, σ_r .




Fig. 8. Tensor-to-scalar ratio uncertainty, σ_r , as a function of instrument configuration for the component separation methods considered in this paper.

configuration, as derived using the different component-separation methods. We note that σ_r for NILC is dominated by the cosmic variance of the bias, so this method is omitted from the figure. We note again that each of the methods involve different masks, different multipole ranges, and different foreground models, and the relative absolute shift between the curves is therefore entirely expected. Furthermore, we also note that the baseline configuration does reach the target sensitivity of $\sigma_r < 0.001$ for both the FGBuster and *Commander* methods, as already reported by *LiteBIRD Collaboration* (2022); even lower absolute uncertainties can be achieved by exploiting additional sky coverage (which is particularly relevant for FGBuster) or multipole range (which is particularly relevant for *Commander*).

More importantly for the main topic of this paper, however, is the relative behavior between different configurations when we shift the entire HFT toward higher frequencies, and this is qualitatively similar for all methods: the uncertainty on r decreases




Fig. 9. Ratio between the tensor-to-scalar ratio uncertainties for M5 and M0 for different analysis choices and component-separation methods. The black points are *Commander* analyses using different masks and sky fractions.

monotonically for all methods over the frequency range considered in this paper.

In Fig. 9 we plot the ratio $\sigma_r^{M5}/\sigma_r^{M0}$ for different sky fractions and analysis methods. Overall, we see that M5 typically has 30–50% smaller uncertainties than M0, and the gains are larger when more sky is included in the analysis. This makes intuitive sense, since accurate component-separation is relatively more important near the Galactic plane.

5.4. Excluding the highest frequency channel

As seen in the above calculations, extending the LiteBIRD frequency range from 402 to 603 GHz leads to a substantially smaller σ_r . However, as discussed in greater detail in the next section, an extended frequency range is also associated with a potentially larger bias from modeling errors, since the overall component separation process becomes more dependent on assumptions regarding the thermal dust SED, both in terms of its specific parametric shape and its general spatial coherence between different frequency channels.

Recognizing these challenges, it is interesting to investigate how much the uncertainty on r degrades if we exclude the highest frequency channel. This represents the scenario in which we are unable to model thermal dust at frequencies above 500 GHz. In this exercise, we adopt the M2 configuration, with the highest frequency channel at 482 GHz, as a worked example. When excluding the highest frequency channel in this particular configuration, we recover an instrument model that is very similar in frequency coverage to the baseline, M0, except with slightly lower sensitivity in the remaining HFT channels.

Analyzing this truncated data set with *Commander* in the same manner as those above, we find $\sigma_r = 9.8 \times 10^{-4}$, which is to be compared with $\sigma_r = 9.2 \times 10^{-4}$ for the nominal M0 configuration with otherwise identical settings, or an increase of about 5%. The potential loss in sensitivity is thus relatively modest for this scenario. However, it is important to note that an additional effect of this modification is that there is no frequency overlap between the MFT and HFT in the extended M2 configuration, while there is such overlap in the baseline M0 configuration. Effectively, the extended instrument configurations thus trade off control of unknown systematics in the overlap between the MFT and HFT at the 195 GHz band for additional sensitivity to thermal dust emission at the highest frequencies.

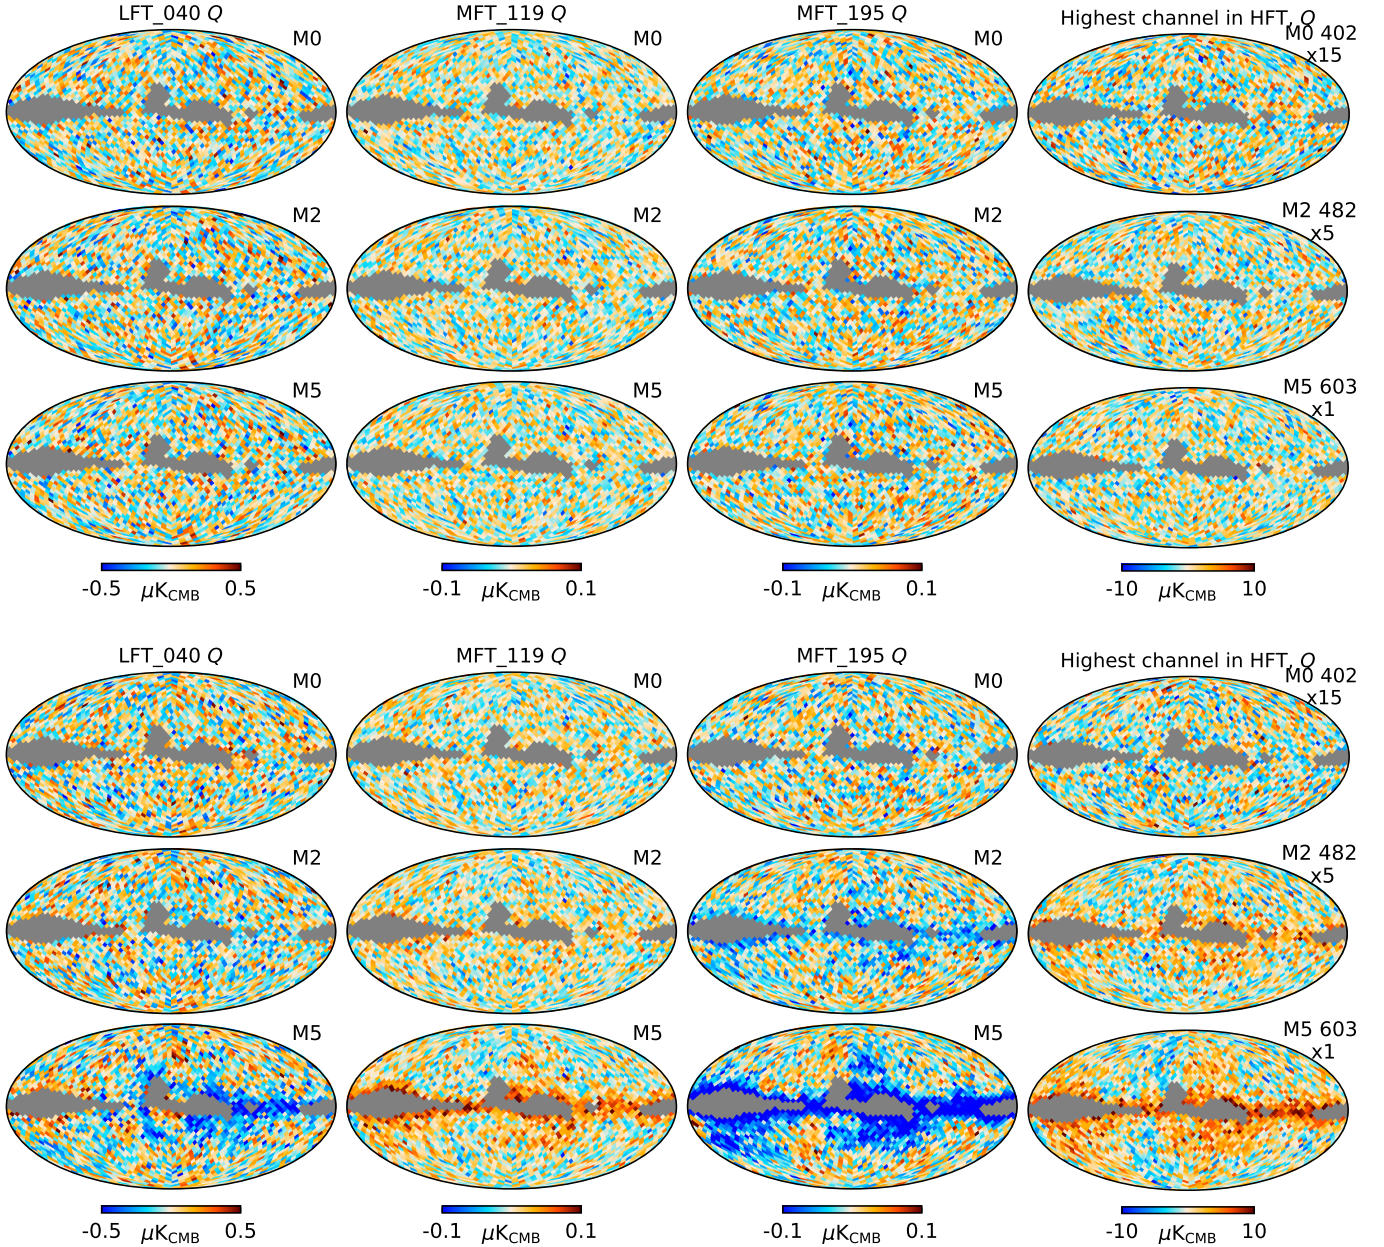


Fig. 10. Commander residual maps ($r_\nu = \mathbf{d}_\nu - s_\nu$) for four different frequency channels and three different instrument configurations. The top section show residual maps for the ideal one-component MBB model, while the bottom section show results for the nonideal SiFeC model, where a model mismatch was included.

6. Non-ideal model analysis: Sensitivity to modeling errors

We now turn our attention to the issue of modeling errors, and we consider two different types. The first type consists of SED modeling errors, in which the fitted SED model is significantly different from the SED of the true sky signal. Being able to reject a spurious foreground-induced detection is key for a robust instrument design. The second type arises from over-smoothing of spatially varying SED parameters, which also can introduce biases.

6.1. Detectability of SED model errors

To assess the ability to detect SED modeling errors as a function of instrument configuration, we employ the parametric Commander method, which provides various goodness-of-fit

quantities in terms of residual maps and χ^2 statistics among its output products. Specifically, we analyze simulations based on the complex SiFeC model summarized in Sect. 3, but using a simple one-component MBB dust model for fitting. For each instrument configuration and simulation, we compute both the tensor-to-scalar ratio posterior distributions, as in Sect. 5.3, and the map-level reduced normalized χ^2 statistic as defined by

$$\chi_{\text{red}}^2 = \frac{\sum_{\nu,p} (d_{\nu,p} - s_{\nu,p})^2 / \sigma_{\nu,p}^2 - n_{\text{d.o.f.}}}{\sqrt{2n_{\text{d.o.f.}}}}. \quad (6)$$

In this expression, the sum runs over all unmasked pixels, p , and all frequencies, ν , and $n_{\text{d.o.f.}}$ is the total number of degrees of freedom per pixel. Since the χ^2 distribution converges to a Gaussian with mean $n_{\text{d.o.f.}}$ and variance $2n_{\text{d.o.f.}}$ for large $n_{\text{d.o.f.}}$, this quantity measures statistical outliers in units of standard deviations. In this section, both the posterior distributions and




Fig. 11. Tensor-to-scalar ratio posterior distributions with a fiducial tensor-to-scalar ratio of $r = 0$, as derived with Commander when fitting SiFeC simulations using the one-component MBB model for six different instrument configurations. A χ^2 statistic is reported for each configuration, indicating the map-level goodness of fit. All results are derived from a sky fraction of $f_{\text{sky}} = 60\%$. The red dashed lines correspond to the median distribution from the ideal one-component MBB fit to an MBB simulation shown in Fig. 7 and serves as a reference.

the χ^2 statistics are evaluated on the unmasked pixels encompassing $f_{\text{sky}} = 60\%$. However, we note that the choice of sky fraction is arbitrary, as the main topic of this section is the impact of modeling errors; a smaller sky fraction will necessarily lead to smaller biases, but larger uncertainties, and vice versa. The choice adopted here is simply chosen as a representative compromise between minimizing both bias and uncertainties.

To build intuition regarding the following results, we show in the top section of Fig. 10 residual maps ($r_v = d_v - s_v$) for one arbitrarily selected realization and four frequency channels for the ideal one-component MBB model. In this case we see no significant residuals in any frequency channels or instrument configurations, simply because the fitted SED model matches the actually simulated sky. All these maps are thus consistent with white noise, and the corresponding χ^2 statistic is also consistent with the Gaussian noise expectation.

In the bottom section of Fig. 10, we show similar residual maps for the complex SiFeC simulations. Here we first see that the baseline M0 configuration appears visually consistent with the ideal MBB simulation in the top panel, and the parameter estimation algorithm will clearly not be able to identify the intrinsic model mismatch. Extending the frequency range to 482 GHz (M2), however, leaves noticeable residuals tracing high dust emission regions in the Galactic plane, and at 603 GHz (M5) the model mismatch is obvious.

The effect of these mismatches in terms of the tensor-to-scalar ratio posterior distributions is quantified in Fig. 11. Each panel shows 20 realizations as black lines, while the red line shows the median of the ideal one-component MBB simulations as a reference; since the overall foreground levels of the two sky models are comparable, and both the fitted model and the instrumental noise parameters are identical between the SiFeC and MBB models, the corresponding posterior distributions should also be roughly equal.

Starting with the baseline configuration in the top panel, we see that the tensor-to-scalar ratio is biased high by a factor of about two (as seen in the shift between the black and

red lines), even though each distribution is individually consistent with zero. The χ^2 statistic is completely consistent with the Gaussian hypothesis, with a deviation of only -1.2σ . The M1 configuration behaves very similarly.

For the case M2, we see that the bias in r is greatly increased, to the point that this configuration would have reported more than a 3σ detection of nonzero B -mode power. At this point, the χ^2 value has also increased somewhat, but is still only 2σ high compared to the Gaussian expectation. It is only with frequency ranges equal to or higher than 523 GHz that χ^2 is sufficiently deviant that we are able to conclusively detect the model mismatch, with a statistical significance of 9σ or more.

While the combination of a significant bias of r and the statistical acceptance of M2 appear disconcerting, it is important to recall that these results are obtained with a relatively conservative sky fraction of 60%. For an analysis of real data, one would also exploit the information in the high dust emission regions to identify a statistically acceptable parametric model, and then fit that also at high latitudes. As seen in the residual maps in Fig. 10, such an analysis would indicate that the one-component MBB is sufficient for M0, while for M2 and higher configurations the residuals are an obvious red flag. From a model selection point of view, frequencies higher than ≈ 500 GHz are critically important for identifying poor thermal dust models.

6.2. Resolution of spatially varying SED parameters

Next, we consider the impact of different smoothing scales for SED spectral parameters. In this case, we analyze the single MBB simulations using the FGbuster algorithm, which supports tuned pixel sizes for each spectral parameter as part of its basic implementation, quantified in terms of the HEALPix resolution parameter, N_{side} . To disentangle the SED error effects from the spatial resolution effect, we once again consider the ideal one-component MBB simulations, but fit coarsely pixelized spectral index maps, while the true spectral index maps are smoothly varying on the sky. In particular, we focus on the thermal dust temperature, which is key for the highest LiteBIRD frequencies, and we allow this to take on different resolutions at low, intermediate, and high Galactic latitudes, denoted $N_{\text{side}}^{T_d} = \{N_{\text{side}}^{\text{low}}, N_{\text{side}}^{\text{int}}, N_{\text{side}}^{\text{high}}\}$; see Appendix A.2 for details.

We are now interested in quantifying the recovery accuracy of r as a function of $N_{\text{side}}^{T_d}$, and the results from this analysis are summarized in Fig. 12 in terms of the recovered mean tensor-to-scalar ratio r (black crosses) and its uncertainty (bars). Groups along the horizontal axis indicate different $N_{\text{side}}^{T_d}$ combinations, while colors indicate different LiteBIRD configurations.

The first thing to note is that when the HFT channels are shifted to higher frequencies, the uncertainty on the spectral indices decreases, and this leads to a decrease of the statistical residuals, and therefore also to an improvement in σ_r ; this is the same effect as was illustrated in Fig. 8, but now we also see that this holds true almost independently of $N_{\text{side}}^{T_d}$.

At the same time, we see that low values of $N_{\text{side}}^{T_d}$ lead to a strong bias in r , and this effect increases rapidly with increasing frequency range. On the other hand, this bias can be mitigated by increasing $N_{\text{side}}^{T_d}$ at the expenses of an increase in the statistical uncertainty. In particular, for $N_{\text{side}}^{T_d} = [32, 32, 32]$, all LiteBIRD configurations lead to a bias $O(3 \times 10^{-4})$, with an uncertainty $\sigma_r \approx O(6 \times 10^{-4})$. In general, a maximum central frequency above 500 GHz requires at least $N_{\text{side}}^{T_d} = 16$, even for the cleanest 20% of the sky in order to keep the bias under control.




Fig. 12. Fisher uncertainty (bars) and bias (black crosses) on r for different instrumental (colors) and FGBuster component-separation configuration (bar groups). The values on the horizontal axis define the N_{side} resolution at which the temperature of the thermal dust component is fitted for in the best 20%, 20–40% and 40–60% sky fractions.

7. Discussion and conclusions

One of the most important challenges for next-generation CMB B -mode experiments is astrophysical foregrounds, and in particular polarized thermal dust emission from our own Milky Way. Properly choosing the frequency range is thus one of the key decisions to be made for any given B -mode space mission or balloon-borne experiment. In this paper, we have revisited this issue for LiteBIRD. Although the established baseline configuration has already been shown to achieve the main mission goals (LiteBIRD Collaboration 2022), it is still desirable to optimize the overall uncertainties to increase the overall system margins.

In this paper, we have addressed this issue by analyzing simulations to measure the recovered uncertainty on the tensor-to-scalar ratio using several different component-separation algorithms for six different possible instrument configurations, where the entire HFT frequency range is shifted toward higher frequencies. We find that the total statistical uncertainty on the tensor-to-scalar ratio, r , after foreground cleaning may be reduced by 30–50% by extending the current maximum frequency from the current baseline of 402 GHz to 500 GHz, or higher for the single MBB dust model, also depending on the sky fraction.

A wider frequency range also increases the instrument’s sensitivity to the shape of the foreground SED, and this has two important consequences. On the one side, a wider frequency range imposes stronger requirements on the foreground model itself, since modeling errors may more easily contaminate the CMB results. For instance, if a given component-separation algorithm (explicitly or implicitly) assumes that the high-frequency channels correlate perfectly with the thermal dust emission at lower frequencies, while the true sky exhibits frequency decorrelation due to the 3D structure of the Milky Way, some fraction of the difference between the assumed and real sky model will leak into the CMB (Tassis & Pavlidou 2015). It is therefore essential that the method of choice allows for sufficient flexibility to capture such uncertainties and model mismatches.

At the same time, a wider frequency range also provides the tools for actually identifying and constraining that model in the first place. Indeed, one of the most important concerns regarding a limited frequency range is the possibility of having low-level foreground residuals that mildly bias cosmological parameters, but still result in an acceptable goodness-of-fit, (for instance as measured by a χ^2 statistic). Higher frequency channels provide important safe-guards against this type of error. Additionally, having access to higher frequency information increases the number of basis functions that can be used to model the thermal dust emission as a function of frequencies. A particularly important example of such a mode is the spatial distribution of the thermal dust temperature, which is useful for both blind and nonblind methods.

An important intuition that underlies all the results discussed in this paper is simply the fact that the CMB SED falls nearly exponentially (relative to dust) above 300 GHz. This implies that even relatively small modifications of the frequency range can lead to significant improvements in a given instrument’s ability to separate CMB from thermal dust. In particular, at the sensitivity level of LiteBIRD, a 402 GHz channel has a non-negligible signal-to-noise ratio for CMB fluctuations, while at 500 GHz this is essentially zero. For an extended configuration, the highest frequency channel therefore represents essentially a pure thermal dust map, while for the baseline configuration it represents a weighted sum of CMB and thermal dust emission. As shown quantitatively in this paper, a clean foreground tracer substantially improves tensor-to-scalar constraints for both parametric and nonparametric (ILC) component-separation methods.

In this paper we have primarily addressed the impact of Galactic foregrounds from the point of view of CMB confusion. An important lesson learned from both WMAP and *Planck*, however, is that Galactic foregrounds also play a key role in understanding and mitigating low-level instrumental systematic effects. For instance, confusion between transmission imbalance factors and foregrounds turned out to be an important

uncertainty for WMAP (Page et al. 2007; Watts et al. 2023), while confusion between gain calibration and foregrounds was a key limitation for *Planck* LFI (Planck Collaboration I 2020; BeyondPlanck Collaboration 2023; Gjerløw et al. 2023). One may therefore argue that experience shows that the most robust approach to mitigating the full effect of astrophysical foregrounds is not by avoiding them, but rather by measuring them. In addition to improving CMB constraints, a wide frequency range will also increase the amount of Galactic science that may be derived from these observations, and it will thereby also increase the overall legacy value of the mission for other fields of astronomy.

While the various instrumental setups studied in this paper appear as possible options, we would like to stress that their technical feasibility has not been performed in depth yet. This would have to include both aspects of optical design and detection chains, but also the potential impacts on the needs for optical modeling and calibration facilities. This trade-off study would also have to deal with the impact of such modifications of the HFT design on the overall systematics effects, which should be propagated into systematics uncertainties on r . As noted in the introduction, performing such a full trade-off study lies far beyond the scope of this paper. Rather, the goal of this paper is to understand whether sufficiently significant gains are achievable under ideal conditions to motivate a proper detailed study – and with typical gains at the 30–50% level in the statistical uncertainty on r after foreground cleaning, the results do appear interesting enough to be investigated further.

Acknowledgements. This work is supported in Japan by ISAS/JAXA for Pre-Phase A2 studies, by the acceleration program of JAXA research and development directorate, by the World Premier International Research Center Initiative (WPI) of MEXT, by the JSPS Core-to-Core Program of A. Advanced Research Networks, and by JSPS KAKENHI Grant Numbers JP15H05891, JP17H01115, and JP17H01125. The Canadian contribution is supported by the Canadian Space Agency. The French LiteBIRD phase A contribution is supported by the Centre National d'Études Spatiales (CNES), by the Centre National de la Recherche Scientifique (CNRS), and by the Commissariat à l'Énergie Atomique (CEA). The German participation in LiteBIRD is supported in part by the Excellence Cluster ORIGINS, which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy (Grant No. EXC-2094 – 390783311). The Italian LiteBIRD phase A contribution is supported by the Italian Space Agency (ASI Grants No. 2020-9-HH.0 and 2016-24-H.1-2018), the National Institute for Nuclear Physics (INFN) and the National Institute for Astrophysics (INAF). Norwegian participation in LiteBIRD is supported by the Research Council of Norway (Grant No. 263011) and has received funding from the European Research Council (ERC) under the Horizon 2020 Research and Innovation Programme (Grant agreement No. 772253 and 819478). The Spanish LiteBIRD phase A contribution is supported by the Spanish Agencia Estatal de Investigación (AEI), project refs. PID2019-110610RB-C21, PID2020-120514GB-I00, ProID2020010108 and ICTP20210008. Funds that support contributions from Sweden come from the Swedish National Space Agency (SNSA/Rymdstyrelsen) and the Swedish Research Council (Reg. no. 2019-03959). The US contribution is supported by NASA grant no. 80NSSC18K0132. We also acknowledge funding from the European Research Council (ERC) under the Horizon 2020 Research and Innovation Programme (Grant agreement No. 725456 and 849169) and The Royal Society (Grant No. URF/R/191023)

References

Ade, P. A. R., Ahmed, Z., Aikin, R. W., et al. 2016, *Phys. Rev. Lett.*, **116**, 031302
 Albrecht, A., & Steinhardt, P. J. 1982, *Phys. Rev. Lett.*, **48**, 1220
 Aurlien, R., Remazeilles, M., Belkner, S., et al. 2023, *J. Cosmol. Astropart. Phys.*, **2023**, 034
 Azzoni, S., Abitbol, M. H., Alonso, D., et al. 2021, *J. Cosmol. Astropart. Phys.*, **2021**, 047
 Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2003, *ApJS*, **148**, 97
 BeyondPlanck Collaboration (Andersen, K. J., et al.) 2023, *A&A*, **675**, A1
 Chluba, J., Hill, J. C., & Abitbol, M. H. 2017, *MNRAS*, **472**, 1195

Delabrouille, J., Cardoso, J.-F., Le Jeune, M., et al. 2009, *A&A*, **493**, 835
 Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, *ApJS*, **155**, 227
 Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, **676**, 10
 Errard, J., & Stompor, R. 2019, *Phys. Rev. D*, **99**, 069907
 Errard, J., Stivoli, F., & Stompor, R. 2011, *Phys. Rev. D*, **84**, 043529
 Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Naess, S. K. 2014, *ApJ*, **790**, 104
 Fuskeland, U., Andersen, K. J., Aurlien, R., et al. 2021, *A&A*, **646**, A69
 Gjerløw, E., Ihle, H. T., Galeotta, S., et al. 2023, *A&A*, **675**, A7
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
 Guilloux, F., Faÿ, G., & Cardoso, J.-F. 2009, *Appl. Comput. Harmon. Anal.*, **26**, 143
 Guth, A. H. 1981, *Phys. Rev. D*, **23**, 347
 Hazumi, M., Borrill, J., Chinone, Y., et al. 2012, *Proc. SPIE Int. Soc. Opt. Eng.*, **8442**, 844219
 Hazumi, M., Ade, P. A. R., Adler, A., et al. 2020, *SPIE Conf. Ser.*, **11443**, 114432F
 Hensley, B. S., & Draine, B. T. 2017, *ApJ*, **834**, 134
 Hensley, B. S., & Draine, B. T. 2023, *ApJ*, **948**, 55
 Hoang, T., & Lazarian, A. 2016, *ApJ*, **831**, 159
 Juvela, M., & Ysard, N. 2012, *A&A*, **539**, A71
 Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, *Phys. Rev. D*, **55**, 7368
 Kogut, A., Dunkley, J., Bennett, C. L., et al. 2007, *ApJ*, **665**, 355
 Krachmalnicoff, N., Baccigalupi, C., Aumont, J., Bersanelli, M., & Mennella, A. 2016, *A&A*, **588**, A65
 Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. 2018, *A&A*, **618**, A166
 Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *MNRAS*, **225**, 307
 Liebe, H. J. 1981, *Radio Science*, **16**, 1183
 Linde, A. D. 1982, *Phys. Lett. B*, **108**, 389
 Linde, A. D. 1983, *Phys. Lett. B*, **129**, 177
 LiteBIRD Collaboration 2022, *Prog. Theor. Exp. Phys.*, **2023**, 042F01
 Mangilli, A., Aumont, J., Rotti, A., et al. 2021, *A&A*, **647**, A52
 Matsumura, T., Akiba, Y., Borrill, J., et al. 2014, *J. Low Temp. Phys.*, **176**, 733
 Montier, L., Mot, B., de Bernardis, P., et al. 2020, *SPIE Conf. Ser.*, **11443**, 114432G
 Narcowich, F., Petrushev, P., & Ward, J. 2006, *SIAM J. Math. Anal.*, **38**, 574
 Novotny, V., & Meincke, P. 1975, *J. Low Temp. Phys.*, **18**, 147
 Page, L., Hinshaw, G., Komatsu, E., et al. 2007, *ApJS*, **170**, 335
 Pardo, J., Cernicharo, J., & Serabyn, E. 2001, *IEEE Trans. Antennas Propag.*, **49**, 1683
 Pelgrims, V., Clark, S. E., Hensley, B. S., et al. 2021, *A&A*, **647**, A16
 Planck Collaboration XII. 2014, *A&A*, **571**, A12
 Planck Collaboration IX. 2016, *A&A*, **594**, A9
 Planck Collaboration XXV. 2016, *A&A*, **594**, A25
 Planck Collaboration I. 2020, *A&A*, **641**, A1
 Planck Collaboration IV. 2020, *A&A*, **641**, A4
 Planck Collaboration VI. 2020, *A&A*, **641**, A6
 Planck Collaboration Int. LVII. 2020, *A&A*, **643**, A42
 Platania, P., Bensadoun, M., Bersanelli, M., et al. 1998, *ApJ*, **505**, 473
 Puglisi, G., Mihaylov, G., Panopoulou, G. V., et al. 2022, *MNRAS*, **511**, 2052
 Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, *MNRAS*, **410**, 2481
 Remazeilles, M., Dickinson, C., Eriksen, H. K. K., & Wehus, I. K. 2016, *MNRAS*, **458**, 2032
 Remazeilles, M., Rotti, A., & Chluba, J. 2021, *MNRAS*, **503**, 2478
 Ritacco, A., Boulanger, F., Guillet, V., et al. 2023, *A&A*, **670**, A163
 Sato, K. 1981, *MNRAS*, **195**, 467
 Sen, A., Basak, S., Ghosh, T., Adak, D., & Sinha, S. 2022, arXiv e-prints [arXiv:2212.02869]
 Starobinsky, A. A. 1980, *Phys. Lett. B*, **91**, 99
 Stompor, R., Leach, S., Stivoli, F., & Baccigalupi, C. 2009, *MNRAS*, **392**, 216
 Tassis, K., & Pavlidou, V. 2015, *MNRAS*, **451**, L90
 Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, *MNRAS*, **469**, 2821
 Vacher, L., Aumont, J., Montier, L., et al. 2022, *A&A*, **660**, A111
 Vacher, L., Aumont, J., Boulanger, F., et al. 2023a, *A&A*, **672**, A146
 Vacher, L., Chluba, J., Aumont, J., Rotti, A., & Montier, L. 2023b, *A&A*, **669**, A5
 Watts, D. J., Galloway, M., Ihle, H. T., et al. 2023, *A&A*, **675**, A16

- 1 Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway
- 2 IRAP, Université de Toulouse, CNRS, CNES, UPS, Toulouse, France
- 3 International School for Advanced Studies (SISSA), Via Bonomea 265, 34136 Trieste, Italy
- 4 INFN Sezione di Trieste, Via Valerio 2, 34127 Trieste, Italy

- ⁵ IFPU, Via Beirut, 2, 34151 Grignano, Trieste, Italy
- ⁶ Université de Paris, CNRS, Astroparticule et Cosmologie, 75013 Paris, France
- ⁷ Instituto de Astrofísica de Canarias, 38200 La Laguna, Tenerife, Canary Islands, Spain
- ⁸ Departamento de Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain
- ⁹ Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan
- ¹⁰ NIST Quantum Sensors Group, 325 Broadway, Boulder, CO 80305, USA
- ¹¹ National Astronomical Observatory of Japan, Mitaka, Tokyo 181-8588, Japan
- ¹² Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, Roma, Italy
- ¹³ INFN Sezione di Roma, P.le A. Moro 2, 00185 Roma, Italy
- ¹⁴ University of Milano Bicocca, Physics Department, Piazza della Scienza, 3, 20126 Milano, Italy
- ¹⁵ INFN Sezione Milano Bicocca, Piazza della Scienza, 3, 20126 Milano, Italy
- ¹⁶ Instituto de Física de Cantabria (IFCA, CSIC-UC), Avenida los Castros s/n, 39005 Santander, Spain
- ¹⁷ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, School of Natural Sciences, The University of Manchester, Oxford Road, Manchester M13 9PL, UK
- ¹⁸ SLAC National Accelerator Laboratory, Kavli Institute for Particle Astrophysics and Cosmology (KIPAC), Menlo Park, CA 94025, USA
- ¹⁹ Stanford University, Department of Physics, Stanford, CA 94305-4060, USA
- ²⁰ Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK
- ²¹ Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, Via Saragat 1, 44122 Ferrara, Italy
- ²² INFN Sezione di Ferrara, Via Saragat 1, 44122 Ferrara, Italy
- ²³ INAF – OAS Bologna, Via Piero Gobetti, 93/3, 40129 Bologna, Italy
- ²⁴ Dipartimento di Fisica e Astronomia “G. Galilei”, Università degli Studi di Padova, Via Marzolo 8, 35131 Padova, Italy
- ²⁵ INFN Sezione di Padova, Via Marzolo 8, 35131 Padova, Italy
- ²⁶ INAF, Osservatorio Astronomico di Padova, Vicolo dell’Osservatorio 5, 35122 Padova, Italy
- ²⁷ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy
- ²⁸ INFN Sezione di Milano, Via Celoria 16, 20133 Milano, Italy
- ²⁹ School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, UK
- ³⁰ Dipartimento di Fisica, Università di Roma Tor Vergata, Via della Ricerca Scientifica, 1, 00133 Roma, Italy
- ³¹ INFN Sezione di Roma2, Università di Roma Tor Vergata, via della Ricerca Scientifica, 1, 00133 Roma, Italy
- ³² University of California, Berkeley, Department of Physics, Berkeley CA 94720, USA
- ³³ University of California, Berkeley, Space Sciences Laboratory, Berkeley, CA 94720, USA
- ³⁴ Lawrence Berkeley National Laboratory (LBNL), Computational Cosmology Center, Berkeley, CA 94720, USA
- ³⁵ Kavli Institute for Particle Astrophysics and Cosmology (KIPAC), Stanford University, Stanford, CA 94305, USA
- ³⁶ Centre Spatial de Liège (STAR Institute), University of Liège, Avenue du Pré-Aily, Angleur 4031, Belgium
- ³⁷ Institute of Particle and Nuclear Studies (IPNS), High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki 305-0801, Japan
- ³⁸ International Center for Quantum-field Measurement Systems for Studies of the Universe and Particles (QUP), High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki 305-0801, Japan
- ³⁹ Dpto. de Física Moderna, Universidad de Cantabria, Avda. los Castros s/n, 39005 Santander, Spain
- ⁴⁰ INFN Sezione di Bologna, Viale C. Berti Pichat 6/2, 40127 Bologna, Italy
- ⁴¹ Tohoku University, Graduate School of Science, Astronomical Institute, Sendai 980-8578, Japan
- ⁴² Japan Aerospace Exploration Agency (JAXA), Institute of Space and Astronautical Science (ISAS), Sagami-hara, Kanagawa 252-5210, Japan
- ⁴³ The Graduate University for Advanced Studies (SOKENDAI), Miura District, Kanagawa 240-0115, Hayama, Japan
- ⁴⁴ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver BC V6T1Z1, Canada
- ⁴⁵ Institut d’Astrophysique de Paris, CNRS/Sorbonne Université, Paris, France
- ⁴⁶ Lawrence Berkeley National Laboratory (LBNL), Physics Division, Berkeley, CA 94720, USA
- ⁴⁷ Laboratoire de Physique de l’École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France
- ⁴⁸ Space Science Data Center, Italian Space Agency, Via del Politecnico, 00133 Roma, Italy
- ⁴⁹ Université Paris-Saclay, CNRS, Institut d’Astrophysique Spatiale, 91405 Orsay, France
- ⁵⁰ Gran Sasso Science Institute (GSSI), Viale F. Crispi 7, 67100 L’Aquila, Italy
- ⁵¹ David A. Dunlap Department of Astronomy and Astrophysics, 50 St. George Street, Toronto, ON M5S3H4, Canada
- ⁵² Institute of Astrophysics, Foundation for Research and Technology-Hellas, Vasilika Vouton, 70013 Heraklion, Greece
- ⁵³ Department of Physics and ITCP, University of Crete, 70013 Heraklion, Greece
- ⁵⁴ INAF-Osservatorio Astronomico di Cagliari, Via della Scienza 5, 09047 Selargius, Italy
- ⁵⁵ Physics and Astronomy Dept., University College London (UCL), London, UK
- ⁵⁶ The University of Tokyo, Department of Astronomy, Tokyo 113-0033, Japan
- ⁵⁷ Suwa University of Science, Chino, Nagano 391-0292, Japan
- ⁵⁸ INFN Sezione di Pisa, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy
- ⁵⁹ Okayama University, Department of Physics, Okayama 700-8530, Japan
- ⁶⁰ Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France
- ⁶¹ National Institute of Technology, Kagawa College, 355 Chokushi-cho, Takamatsu, Kagawa 761-8058, Japan
- ⁶² Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München, Scheinerstr.1, 81679 München, Germany
- ⁶³ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

Appendix A: Parametric component-separation methods

A.1. Commander

Commander is a standard parametric Bayesian component-separation framework for CMB observations (Eriksen et al. 2004, 2008) that has been used extensively by *Planck* (e.g., Planck Collaboration XII 2014; Planck Collaboration IX 2016; Planck Collaboration IV 2020; Planck Collaboration Int. LVII 2020). In the current analysis we employ the first implementation of this framework, typically referred to as **Commander1**, which was used in the *Planck* 2013 data release. The main advantage of this framework is a very high computational efficiency, which is important when analyzing an ensemble of simulations; at the same time, the main drawback is its requirement of identical angular resolution of all frequency bands, which limits its use to low-resolution data sets. In our current analysis, all input maps are at $N_{\text{side}} = 16$, with a common resolution of 10° FWHM. Another option would be to use **Commander3** (BeyondPlanck Collaboration 2023), which supports both multiresolution component separation and low-level time-ordered data analysis, but this would increase the computational expense of the analysis by orders of magnitude, without actually providing valuable new information regarding the central question of frequency range considered in this paper.

The starting point of any parametric Bayesian analysis is to define an explicit parametric data model, described in Eqs. (1)–(5), and in **Commander1**, every parameter is fitted per pixel. The main goal is now to fit the parameters $\theta = \{\mathbf{a}_{\text{CMB}}, \mathbf{a}_s, \mathbf{a}_d, \beta_s, \beta_d, T_d\}$, which within a Bayesian framework means computing the posterior distribution $P(\theta|\mathbf{d})$, and then using this to estimate the cosmological parameter r . Using Bayes' theorem, the posterior distribution may be written as

$$P(\theta|\mathbf{d}) = \frac{P(\mathbf{d}|\theta)P(\theta)}{P(\mathbf{d})} \quad (\text{A.1})$$

$$\propto P(\mathbf{d}|\theta)P(\theta), \quad (\text{A.2})$$

where $P(\mathbf{d})$ is a normalization factor, $P(\mathbf{d}|\theta)$ is the likelihood, and $P(\theta)$ represents some set of priors. In this work, we apply loose Gaussian priors on the spectral parameters, but no priors on the amplitude parameters. Specifically, we adopt a synchrotron prior of $P(\beta_s) = N(-3.0, 0.3)$, where $N(\mu, \sigma)$ denotes a standard Gaussian distribution with mean μ and standard deviation σ . For thermal dust emission, we adopt $P(\beta_d) = N(1.54, 0.20)$ and $P(T_d) = N(23 \text{ K}, 7 \text{ K})$ where we fit for the one MBB component.

To map out the distribution in Eq. (A.2), we use Gibbs sampling as implemented in **Commander**. That is, rather than sampling directly from $P(\theta|\mathbf{d})$, which is computationally unfeasible, we iteratively sample from each of the corresponding conditional probability distributions. For the purposes of this paper, this can be written as

$$\{\mathbf{a}_{\text{CMB}}, \mathbf{a}_s, \mathbf{a}_d\}^{i+1} \leftarrow P(\mathbf{a}_{\text{CMB}}, \mathbf{a}_s, \mathbf{a}_d | \beta_s^i, \beta_d^i, T_d^i, \mathbf{d}), \quad (\text{A.3})$$

$$\{\beta_s, \beta_d, T_d\}^{i+1} \leftarrow P(\beta_s, \beta_d, T_d | \mathbf{a}_{\text{CMB}}^{i+1}, \mathbf{a}_s^{i+1}, \mathbf{a}_d^{i+1}, \mathbf{d}). \quad (\text{A.4})$$

Here \leftarrow indicates that the parameters to the left are sampled from the distribution on the right-hand side. All amplitude parameters are sampled jointly together with the CMB sky signal to avoid excessive Monte Carlo correlation lengths. The nonlinear parameters, the spectral indices and dust temperature, are sampled using a standard inversion sampler.

We run this iterative Gibbs sampler until convergence, which in the current analysis requires about 40 000 samples after rejecting a burn-in period of about 500 samples. We repeat this analysis for 20 different CMB and noise realizations.⁴ Based on the samples from each simulation, we compute a posterior mean CMB map, $\hat{\mathbf{a}}_{\text{CMB}}$, and the corresponding covariance matrix, \mathbf{N} , and these are subsequently fed into a standard Gaussian likelihood;

$$\mathcal{L}(r) \propto \frac{e^{-\frac{1}{2} \hat{\mathbf{a}}_{\text{CMB}}^t (\mathbf{S}(r) + \mathbf{N})^{-1} \hat{\mathbf{a}}_{\text{CMB}}}}{\sqrt{|\mathbf{S}(r) + \mathbf{N}|}}. \quad (\text{A.5})$$

Here $\mathbf{S}(r)$ is the CMB signal covariance matrix for a given value of r and the best-fit Λ CDM parameters described in Sect. 3. The low pixel resolution of $N_{\text{side}} = 16$ used for **Commander** in this paper is dictated by the CPU time and RAM required to evaluate this equation, as well as the number of samples needed to establish a robust noise covariance matrix.

We adopt a uniform prior for all positive values of r , and the posterior distribution is therefore numerically identical to the likelihood in Eq. (A.5) for $r > 0$. This function is mapped out by gridding Eq. (A.5) over a precomputed library of angular power spectra, $C_\ell(r)$, which are computed into covariance matrices in pixel space, $\mathbf{S}(r)$, including only multipoles in the range $\ell = [2, 12]$, as dictated by the low pixel resolution of $N_{\text{side}} = 16$. The **Commander** results thus only probe the low- ℓ reionization peak of the B -mode power spectrum, not the recombination peak around $\ell \approx 100$. However, this is also where diffuse component separation is most challenging, and the question of frequency optimization is most important.

The main analysis mask used for the **Commander** analysis is the $f_{\text{sky}} = 73\%$ polarization mask discussed in Planck Collaboration IV (2020). To visualize the impact of sky fraction, other masks have been used in Fig. 9, such as masks based on high latitude cuts, χ^2 , and by thresholding the 402 GHz polarization amplitude ($P = \sqrt{Q^2 + U^2}$) map, which primarily suppresses thermal dust emission. This latter is used for the 60% mask in Sect. 6.

A.2. FGBuster

A second implementation of parametric component separation is called **FGBuster**.⁵ This method is based on the same parametric modeling as defined by Eqs. (1)–(5), but uses a computationally efficient non-linear optimization procedure to perform the fit, as opposed to brute-force Monte Carlo sampling as implemented in **Commander**. Following Stompor et al. (2009), Errard et al. (2011), Errard & Stompor (2019), and Puglisi et al. (2022), **FGBuster** proceeds in three steps. These are: first the estimation of the spectral parameters $\{\beta_d, T_d, \beta_s\}$ through the optimization of a so-called spectral likelihood; second, the estimation of the component maps $\{\mathbf{a}_{\text{CMB}}, \mathbf{a}_d, \mathbf{a}_s\}$ introduced in Appendix A.1; and third, the estimation of the angular power spectra for the residual foregrounds (i.e., the difference between the input and recovered CMB map), as well as the estimation of the tensor-to-scalar ratio r with a Gaussian likelihood defined by the previous quantities, the theoretical BB spectrum, and the noise after component separation.

⁴ This large number of samples is the main motivation for using **Commander1** rather than **Commander3**; producing $O(10^6)$ posterior samples at full angular resolution would require tens of millions of CPU hours. While feasible for a final production analysis of real data, such a cost is not justified for the current exploratory analysis.

⁵ <https://github.com/fgbuster/fgbuster>




Fig. A.1. Example of the results from the FGBuster component-separation runs. Two instrumental configurations are considered: the baseline M0 (blue) and the high frequency extension M3 (red). The solid lines represent the post-component-separation noise and foreground residuals S_ℓ , computed by averaging 1000 simulations for the baseline, and averaging 10 simulations for the extension. Together with the power from gravitational lensing and primordial B -modes, they define the error bars on the CMB power spectrum (boxes).

Regarding the first step, we note that FGBuster supports different resolutions (as defined by the HEALPix resolution parameter, N_{side}) for each spectral parameter, while the corresponding amplitudes are fitted at full angular resolution. Choosing a low N_{side} implies that few foreground parameters are fitted for, resulting in low post-component-separation noise; however, this may also result in a high bias due to sky complexity that the model is not allowed to capture. Optimizing N_{side} for each parameter is an important point of the algorithm, and we quantify these two competing effects by splitting the foreground residuals into two categories. First, the systematic residuals, as obtained from a noiseless simulation, are driven by the mismatch between input sky and the assumed parametric model, and its power spectrum, B_ℓ , defines the bias on r . Second, the statistical residuals are computed as the average power spectrum from 10 noisy simulations, S_ℓ , which are driven by the statistical error while estimating spectral parameters, and by coaddition of the noise at different frequencies.

We consider the following approximate likelihood for the power spectrum:

$$-2 \ln \mathcal{L}(r) = f_{\text{sky}} \sum_{\ell} (2\ell + 1) \left(\frac{D_\ell}{C_\ell} + \ln C_\ell \right) + \text{const}, \quad (\text{A.6})$$

where $f_{\text{sky}} = 0.49$, $D_\ell = L_\ell + S_\ell + B_\ell$, $C_\ell = rT_\ell + L_\ell + S_\ell$, L_ℓ is the BB power spectrum from gravitational lensing, and T_ℓ is the tensor power spectrum for $r = 1$. The bias on r is estimated by maximizing this likelihood, while we define σ_r to be the Fisher estimate obtained by assuming $B_\ell = 0$. We prefer this choice over computing σ_r from the width of Eq. (A.6), because the latter procedure produces values of σ_r that are strongly dependent on the bias, while here we want it to reflect only the statistical constraining power of a given configuration. The various spectra involved in the FGBuster analysis are illustrated in Fig. A.1.

The free parameters of our setup are the N_{side} values of the spectral parameters, and the LiteBIRD instrumental configuration. Regarding the former, we follow LiteBIRD Collaboration (2022), and split the sky into three (almost) iso-Galactic-latitude regions, each covering approximately 20% of the sky. In each

region, every parameter can have different values of N_{side} , tuned according to its local signal-to-noise ratio. We define the three regions using *Planck* post-processing masks adopted for temperature component separation (Planck Collaboration IV 2020), the so-called GAL20, GAL40, and GAL60. We find this (arbitrary) choice to significantly help coping with the varying signal-to-noise, but the use of more of these separate regions is certainly possible and could further improve the post-component separation performance – especially if specialized for dust polarization. Regarding the choice of the N_{side} values of the spectral parameter, we find that an accurate estimation of β_d is particularly essential for the LiteBIRD configuration, due to a small statistical uncertainty but potentially large biases coming from spatial variability; we use $N_{\text{side}}^{\beta_d} = 64$ in all the regions. Less accurate estimation of β_s and T_d are usually tolerable as was experienced with foreground-only simulations, and following the preliminary results from Errard & Stompor (2019). For synchrotron emission, we adopt $N_{\text{side}}^{\beta_s} = [4, 2, 2]$ for the three regions (counting from lower latitudes), and this is kept constant for all instrumental configurations; we do not expect changes at the high frequencies to significantly affect the fit of the low-frequency foreground.⁶ In contrast, increasing the HFT frequencies does increase the constraining power for the dust SED, and therefore requires a more precise modeling of the spatial variability of T_d to keep the systematic residuals under control. The analysis sky mask covers $f_{\text{sky}} = 49\%$ of the sky as introduced in LiteBIRD Collaboration (2022). In addition to the Galactic plane, this mask removes regions with high residuals.

A.3. Moment expansion

A third parametric component-separation method is referred to as the “moment-expansion method,” which models the complexity of the polarized dust SED using parametric fitting with moment expansion coefficients in harmonic space. This formalism was introduced by Chluba et al. (2017) and generalized to polarization in Vacher et al. (2023b,a) and at the cross-frequency power spectra level in Mangilli et al. (2021). An application of this formalism for component separation with the LiteBIRD mission at high frequencies can be found in Vacher et al. (2022). The formalism and the pipeline developed in this last study will be applied identically in the present work.

In true experimental conditions, averaging over different SEDs emitted from distinct regions with different physical properties is unavoidable. They appear along the line of sight in the 3D sky, between different lines of sight, inside the beam of the instrument, or when performing a spherical harmonic decomposition to calculate the angular power spectra over large regions of the sky. Since SEDs depend of the frequency in a non-linear fashion, summing over them when spectral parameters are varying will lead to distortions that should be understood and modeled. The moment-expansion method aims at accounting for SED distortions from a canonical model by introducing new terms coming from a Taylor-inspired expansion of the SED around a pivot value and with respect to its spectral parameters (e.g., β_d and T_d for an MBB). One can generalize this expansion at the cross-frequency power spectra level $\mathcal{D}_\ell(\nu_i \times \nu_j)$, providing an analytical model that can be fitted over the foreground signal.

⁶ Note that the formalism does not assume or exploit possible correlations between the components. In addition to the very different spectral dependence, this should further mitigate the impact of dust and dust-synchrotron correlations on the synchrotron fit.

As discussed in Vacher et al. (2022), providing a correct model for the distortions due to spatial variations of the dust temperature is critical for a mission like LiteBIRD, while it could safely be ignored for missions like Simons Observatory (Azzoni et al. 2021), with different frequency coverage and lower sensitivity. If the dust signal in every Galactic pixel is indeed a modified black-body, reaching higher frequencies will ease component separation by breaking degeneracies between the moments in β_d and T_d , providing a better characterization of the impact of the dust temperature variations on the shape of the power spectra SED.

For the present application, we extract the cross-frequency power spectra of $N_{\text{sim}} = 500$ simulations, including different realizations of Gaussian noise and CMB lensing, with all the high-frequency configurations from M0 to M5. The only foreground component considered is thermal dust with a single MBB. The simulation maps are produced at $N_{\text{side}} = 512$ and masked in order to keep a sky fraction of $f_{\text{sky}} = 70\%$, using a raw mask made from *Planck* intensity data at 353 GHz as in Vacher et al. (2022). Only the nine highest frequencies above 100 GHz are considered, leading to 13 bands (counting the ones sharing the same frequency) and thus 91 cross-spectra. The spectra are binned up to $\ell = 200$ with bins of $\Delta\ell = 10$. Since we expect it to be the only fitting procedure allowing us to recover an unshifted r posterior for the standard LiteBIRD configuration, we shall only consider the $r\beta$ - T fitting scheme, using the moment coefficients up to order 1 in both β_d and T_d (Vacher et al. 2022). For each simulation, a χ^2 minimization is done at the cross-frequency power spectra level

$$\chi^2 = \frac{1}{N_{\text{d.o.f.}}} \mathbf{R}^T \mathbf{M}^{-1} \mathbf{R}, \quad (\text{A.7})$$

with the residual $\mathcal{R}_\ell(v_i \times v_j) = \mathcal{D}_\ell^{\text{sim}}(v_i \times v_j) - \mathcal{D}_\ell^{\text{model}}(v_i \times v_j)$ and the covariance matrix: $\mathbf{M}_{\ell, \ell'}^{i \times j, k \times l} = \text{cov}(\mathcal{D}_\ell^{\text{sim}}(v_i \times v_j), \mathcal{D}_{\ell'}^{\text{sim}}(v_k \times v_l))$, allowing us to extract one best-fit value for the tensor-to-scalar ratio \hat{r} per simulation. A histogram is built out of the N_{sim} best fit values of \hat{r} , which gives us the final posteriors, with associated standard deviation $\sigma_{\hat{r}}$ found by fitting a Gaussian curve over it.

Appendix B: Minimum variance component-separation methods

B.1. Internal linear combinations

Generally speaking, internal linear combination (ILC) methods perform a weighted linear combination of the data d_ν across frequencies:

$$\hat{s} = \sum_\nu w_\nu d_\nu = \mathbf{w}^T \mathbf{d}, \quad (\text{B.1})$$

which is designed to minimize the overall variance due to foregrounds and noise without altering the CMB signal. More explicitly, the ILC weights $\mathbf{w} \equiv \{w_\nu\}$ are the solution of a constrained minimization problem that can be formulated by the Lagrangian,

$$L(\mathbf{w}, \lambda) = \mathbf{w}^T \mathbf{C} \mathbf{w} + \lambda (1 - \mathbf{w}^T \mathbf{a}), \quad (\text{B.2})$$

in which the matrix \mathbf{C} collects the elements $C_{\nu\nu'} = \langle d_\nu d_{\nu'} \rangle$ of the data covariance matrix for all pairs of frequencies (ν, ν') and the vector $\mathbf{a} \equiv \{a_\nu\}$ collects the frequency spectrum of the CMB

component across the channels. The Lagrange multiplier λ is used to ensure that the overall variance $\mathbf{w}^T \mathbf{C} \mathbf{w}$ of the ILC estimate \hat{s} is minimized under the constraint $\sum_\nu w_\nu a_\nu = 1$ for the preservation of the CMB signal. The ILC weights are thus given by the saddle point of the Lagrangian defined in Eq. (B.2):

$$\mathbf{w}^T = (\mathbf{a}^T \mathbf{C}^{-1} \mathbf{a})^{-1} \mathbf{a}^T \mathbf{C}^{-1}, \quad (\text{B.3})$$

and as such they do not rely on any explicit parametric foreground model.

B.2. NILC

The Needlet ILC (NILC; Delabrouille et al. 2009) method is a specific version of the ILC approach that performs this optimization using a *needlet* (i.e., spherical wavelet) vector basis. In the NILC method, the ILC weights, Eq. (B.3), are computed independently in different regions of the sky and different ranges of angular scales using a wavelet decomposition of the data. Needlets have excellent localization properties, both on the sphere (pixel space) and in harmonic space (Narcowich et al. 2006; Guilloux et al. 2009), which allows the NILC algorithm to adjust the ILC weights depending on the local variations of the foreground and noise contaminations across the sky and across different scales.

B.3. cMILC

The Constrained Moment ILC (cMILC; Remazeilles et al. 2021) method is an extension of the NILC method that includes additional nulling constraints to *deproject* spectral moments of the foreground emission arising from line-of-sight integration and beam convolution. It combines the so-called constrained ILC approach (Remazeilles et al. 2011) and the moment expansion technique for dealing with foreground emission (Chluba et al. 2017). Because of spectral variations of foreground emission along the line-of-sight and/or inside the beam, the line-of-sight integration and beam convolution effects in sky observations cause spectral distortions to the baseline SEDs of thermal dust and synchrotron, and thus induce partial decorrelation across frequencies. Such higher-order corrections to the baseline foreground emission can be significant compared to the CMB B -mode signal at low tensor-to-scalar ratio values.

These effects are addressed by cMILC through moment expansion of the foreground emission $I_\nu(\hat{n})$ beyond the leading-order SED:

$$\begin{aligned} I_\nu(\hat{n}) &= A(\hat{n}) f(\nu, \bar{\boldsymbol{\beta}}) + \sum_i A(\hat{n}) (\beta_i(\hat{n}) - \bar{\beta}_i) \frac{\partial}{\partial \beta_i} f(\nu, \bar{\boldsymbol{\beta}}) \\ &+ \frac{1}{2!} \sum_{i,j} A(\hat{n}) (\beta_i(\hat{n}) - \bar{\beta}_i) (\beta_j(\hat{n}) - \bar{\beta}_j) \frac{\partial^2}{\partial \beta_i \partial \beta_j} f(\nu, \bar{\boldsymbol{\beta}}) \\ &+ \dots, \end{aligned} \quad (\text{B.4})$$

where $f(\nu, \bar{\boldsymbol{\beta}})$ is the baseline SED evaluated at fixed pivot values $\bar{\boldsymbol{\beta}} \equiv \{\bar{\beta}_i\}$ for a set of spectral parameters $\boldsymbol{\beta} \equiv \{\beta_i(\hat{n})\}$ (e.g., dust spectral index and temperature) that vary with the direction \hat{n} on the sky and along the line of sight. The expansion defined by Eq. (B.4) highlights the foreground moment components,

$$m_\alpha(\hat{n}) = A(\hat{n}) \frac{(\beta_1(\hat{n}) - \bar{\beta}_1)^{\alpha_1} \dots (\beta_n(\hat{n}) - \bar{\beta}_n)^{\alpha_n}}{\alpha_1! \dots \alpha_n!}, \quad (\text{B.5})$$

which have homogeneous SEDs given by the derivatives of the baseline SED with respect to spectral parameters:

$$b_\alpha(\nu) = \frac{\partial^{\alpha_1}}{\partial \beta_1^{\alpha_1}} \cdots \frac{\partial^{\alpha_n}}{\partial \beta_n^{\alpha_n}} f(\nu, \bar{\boldsymbol{\beta}}), \quad (\text{B.6})$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$.

cMILC enables us to *deproject* several moments of the foreground emission from the recovered CMB B -mode map by generalizing the Lagrangian in Eq. (B.2) as

$$L(\mathbf{w}, \lambda, \{\mu_\alpha\}) = \mathbf{w}^\top \mathbf{C} \mathbf{w} + \lambda(1 - \mathbf{w}^\top \mathbf{a}) + \sum_\alpha \mu_\alpha \mathbf{w}^\top \mathbf{b}_\alpha, \quad (\text{B.7})$$

in which extra Lagrange multipliers μ_α are introduced to impose several nulling constraints on the ILC weights, $\sum_\nu w_\nu b_\alpha(\nu) = 0$, with respect to the moment SEDs in Eq. (B.6). The cMILC weights are thus given by the saddle point of the extended Lagrangian in Eq. (B.7),

$$\mathbf{w}^\top = \mathbf{e}^\top (\mathbf{A}^\top \mathbf{C}^{-1} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{C}^{-1}, \quad (\text{B.8})$$

where the mixing matrix $\mathbf{A} = (\mathbf{a} \mathbf{b}_1 \cdots \mathbf{b}_m)$ contains the CMB SED vector \mathbf{a} in the first column, the moment SED vectors, $\mathbf{b}_1, \dots, \mathbf{b}_m$, as defined by Eq. (B.6) in the other columns, and $\mathbf{e}^\top = (1 \ 0 \cdots 0)$. The unconstrained foregrounds (i.e., high-order moments), which are not deprojected by cMILC, are simply mitigated through blind variance minimization, as in the NILC algorithm.

For both NILC and cMILC, we transform the Stokes Q and U full-sky simulations defined in Sect. 3 into full-sky B -mode maps

before computing the fits described above, resulting in pure B -mode constraints. We then apply these methods to the two most extreme instrument configurations defined in Table 1, namely M0 and M5. Foreground cleaning is performed on the full sky for both NILC and cMILC, while the power spectra and r estimates are computed after masking $f_{\text{sky}} = 50\%$ of the NILC and cMILC B -mode maps, where the mask is obtained by thresholding the observed 402 GHz B -mode map until 50% of the sky is masked.

The r estimates are derived from the angular power spectrum of the NILC/cMILC B -mode map using the Gaussian likelihood

$$-2 \ln \mathcal{L}(r) = \sum_\ell \frac{(\widehat{C}_\ell^{BB} - C_\ell(r, A_L) - \widehat{N}_\ell)^2}{\widehat{\Xi}_\ell^2}, \quad (\text{B.9})$$

where \widehat{C}_ℓ^{BB} is the B -mode power spectrum of the reconstructed NILC/cMILC map, \widehat{N}_ℓ is the post-component-separation noise, and $C_\ell(r, A_L)$ is the theoretical B -mode power spectrum as a function of r that we fit to the data, with lensing amplitude fixed to $A_L = 1$ or 0, depending on either no delensing or full delensing assumptions. The covariance matrix is given by

$$\widehat{\Xi}_\ell = \sqrt{\frac{2}{(2\ell + 1) f_{\text{sky}}}} \widehat{C}_\ell^{BB}, \quad (\text{B.10})$$

which implicitly includes the sample variance of the residual foregrounds and noise that are left in the reconstructed NILC/cMILC map.