

PAPER • OPEN ACCESS

On the computational feasibility of Bayesian end-to-end analysis of *LiteBIRD* simulations within Cosmoglobe

To cite this article: R. Aurvik *et al* JCAP11(2025)041

View the [article online](#) for updates and enhancements.


You may also like

- [LOCALIZATION AND BROADBAND FOLLOW-UP OF THE GRAVITATIONAL-WAVE TRANSIENT GW150914](#)
B. P. Abbott, R. Abbott, T. D. Abbott et al.
- [Multi-messenger Observations of a Binary Neutron Star Merger](#)
B. P. Abbott, R. Abbott, T. D. Abbott et al.
- [The LHCb Upgrade I](#)
R. Aaij, A.S.W. Abdelmotteleb, C. Abellan Beteta et al.

On the computational feasibility of Bayesian end-to-end analysis of *LiteBIRD* simulations within *Cosmoglobe*



The Cosmoglobe collaboration and the LiteBIRD collaboration

R. Aurvik,¹ M. Galloway,¹ E. Gjerløw,¹ U. Fuskeland ^{1,*} A. Basyrov,¹ M. Bortolami,^{2,3} M. Brilenkov,¹ P. Campeti,^{3,4,5} H.K. Eriksen,¹ L.T. Hergt,^{6,7} D. Herman,¹ M. Monelli,⁸ L. Pagano,^{2,3,9} G. Puglisi,^{10,11,12} N. Raffuzzi,² N.-O. Stutzer,¹ R.M. Sullivan,¹ H. Thommesen,¹ D.J. Watts,¹ I.K. Wehus,¹ D. Adak,¹³ E. Allys,¹⁴ A. Anand,¹⁵ J. Aumont,¹⁶ C. Baccigalupi,^{17,18,19} M. Ballardini,^{2,3,20} A.J. Banday,¹⁶ R.B. Barreiro,²¹ N. Bartolo,^{22,23,24} S. Basak,²⁵ M. Bersanelli,^{26,27} A. Besnard,⁹ T. Brinckmann,² E. Calabrese,²⁸ E. Carinos,¹⁶ F.J. Casas,²¹ K. Cheung,^{29,30,31,32} M. Citran,³³ L. Clermont,³⁴ F. Columbro,^{35,36} G. Coppi,³⁷ A. Coppolecchia,^{35,36} P. Dal Bo,³⁸ P. de Bernardis,^{35,36} E. de la Hoz,³⁹ M. De Lucia,³⁸ S. Della Torre,⁴⁰ P. Diego-Palazuelos,⁴ T. Essinger-Hileman,⁴¹ C. Franceschet,^{26,27} G. Galloni,^{2,15} M. Gerbino,³ M. Gervasi,^{37,40} R.T. Génova-Santos,^{13,42} T. Ghigna,^{43,8} S. Giardiello,²⁸ C. Gimeno-Amo,²¹ A. Gruppuso,^{20,44} M. Hazumi,^{45,46,8,47} S. Henrot-Versillé,⁷ K. Kohri,⁴⁵ L. Lamagna,^{35,36} T. Lari,³⁸ M. Lattanzi,³ C. Leloup,⁸ F. Levrier,¹⁴ A.I. Lonappan,⁴⁸ M. López-Caniego,^{49,50} G. Luzzi,⁵¹ J. Macias-Perez,⁵² B. Maffei,⁹ E. Martínez-González,²¹ S. Masi,^{35,36} S. Matarrese,^{22,23,24,53} T. Matsumura,⁸ S. Micheli,³⁵ L. Montier,¹⁶ G. Morgante,²⁰ L. Mousset,^{14,16} R. Nagata,⁴⁶ A. Novelli,³⁵ I. Obata,⁴⁵ A. Occhiuzzi,³⁵ A. Paiella,^{35,36} D. Paoletti,^{20,44} G. Pascual-Cisneros,^{8,21} F. Piacentini,^{35,36} M. Pinchera,³⁸ G. Polenta,⁵¹ L. Porcelli,⁵⁴ M. Remazeilles,²¹ A. Ritacco,⁵² A. Rizzieri,^{55,33} M. Ruiz-Granda,^{21,56} J. Sanghavi,^{57,58} V. Sauvage,⁹ M. Shiraishi,⁵⁹ S.L. Stever,^{9,60,8} Y. Takase,⁶⁰ K. Tassis,^{61,62} L. Terenzi,²⁰ M. Tomasi,^{26,27} M. Tristram,⁷ L. Vacher,¹⁷ B. van Tent,⁷ P. Vielva,²¹ G. Weymann-Despres,^{55,7} E.J. Wollack,⁴¹ M. Zannoni,^{37,40} and Y. Zhou⁴³

*Corresponding author.

- ¹*Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway*
- ²*Dipartimento di Fisica e Scienze della Terra, Università di Ferrara,
Via Saragat 1, 44122 Ferrara, Italy*
- ³*INFN Sezione di Ferrara, Via Saragat 1, 44122 Ferrara, Italy*
- ⁴*Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany*
- ⁵*Excellence Cluster ORIGINS, Boltzmannstr. 2, 85748 Garching, Germany*
- ⁶*Department of Physics and Astronomy, University of British Columbia,
6224 Agricultural Road, Vancouver, BC V6T1Z1, Canada*
- ⁷*Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*
- ⁸*Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI), UTIAS,
The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*
- ⁹*Université Paris-Saclay, CNRS, Institut d’Astrophysique Spatiale, 91405, Orsay, France*
- ¹⁰*Dipartimento di Fisica e Astronomia, Università degli Studi di Catania,
Via S. Sofia,64, 95123, Catania, Italy*
- ¹¹*INAF, Osservatorio Astrofisico di Catania, via S.Sofia 78, I-95123 Catania, Italy*
- ¹²*INFN, Sezione di Catania, via S.Sofia 64, I-95123, Catania, Italy*
- ¹³*Instituto de Astrofísica de Canarias, E-38200 La Laguna, Tenerife, Canary Islands, Spain*
- ¹⁴*Laboratoire de Physique de l’École Normale Supérieure, ENS, Université PSL, CNRS,
Sorbonne Université, Université de Paris, 75005 Paris, France*
- ¹⁵*Dipartimento di Fisica, Università di Roma Tor Vergata,
Via della Ricerca Scientifica, 1, 00133, Roma, Italy*
- ¹⁶*IRAP, Université de Toulouse, CNRS, CNES, UPS, Toulouse, France*
- ¹⁷*International School for Advanced Studies (SISSA), Via Bonomea 265, 34136, Trieste, Italy*
- ¹⁸*INFN Sezione di Trieste, via Valerio 2, 34127 Trieste, Italy*
- ¹⁹*IFPU, Via Beirut, 2, 34151 Grignano, Trieste, Italy*
- ²⁰*INAF - OAS Bologna, via Piero Gobetti, 93/3, 40129 Bologna, Italy*
- ²¹*Instituto de Física de Cantabria (IFCA, CSIC-UC),
Avenida los Castros SN, 39005, Santander, Spain*
- ²²*Dipartimento di Fisica e Astronomia “G. Galilei”, Università degli Studi di Padova,
via Marzolo 8, I-35131 Padova, Italy*
- ²³*INFN Sezione di Padova, via Marzolo 8, I-35131, Padova, Italy*
- ²⁴*INAF, Osservatorio Astronomico di Padova, Vicolo dell’Osservatorio 5, I-35122, Padova, Italy*
- ²⁵*School of Physics, Indian Institute of Science Education and Research Thiruvananthapuram,
Maruthamala PO, Vithura, Thiruvananthapuram 695551, Kerala, India*
- ²⁶*Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16 - 20133, Milano, Italy*
- ²⁷*INFN Sezione di Milano, Via Celoria 16 - 20133, Milano, Italy*
- ²⁸*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, U.K.*
- ²⁹*Jodrell Bank Centre for Astrophysics, Alan Turing Building, Department of Physics and
Astronomy, School of Natural Sciences, The University of Manchester,
Oxford Road, Manchester M13 9PL, U.K.*
- ³⁰*University of California, Berkeley, Department of Physics, Berkeley, CA 94720, U.S.A.*
- ³¹*University of California, Berkeley, Space Sciences Laboratory, Berkeley, CA 94720, U.S.A.*
- ³²*Lawrence Berkeley National Laboratory (LBNL), Computational Cosmology Center,
Berkeley, CA 94720, U.S.A.*
- ³³*Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France*

- ³⁴ *Centre Spatial de Liège, Université de Liège, Avenue du Pré-Aily, 4031 Angleur, Belgium*
- ³⁵ *Dipartimento di Fisica, Università La Sapienza, P. le A. Moro 2, Roma, Italy*
- ³⁶ *INFN Sezione di Roma, P.le A. Moro 2, 00185 Roma, Italy*
- ³⁷ *University of Milano Bicocca, Physics Department, p.zza della Scienza, 3, 20126 Milan, Italy*
- ³⁸ *INFN Sezione di Pisa, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy*
- ³⁹ *CNRS-UCB International Research Laboratory, Centre Pierre Binétruy, UMI2007, Berkeley, CA 94720, U.S.A.*
- ⁴⁰ *INFN Sezione Milano Bicocca, Piazza della Scienza, 3, 20126 Milano, Italy*
- ⁴¹ *NASA Goddard Space Flight Center, Greenbelt, MD 20771, U.S.A.*
- ⁴² *Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38206, La Laguna, Tenerife, Spain*
- ⁴³ *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*
- ⁴⁴ *INFN Sezione di Bologna, Viale C. Berti Pichat, 6/2 – 40127 Bologna, Italy*
- ⁴⁵ *Institute of Particle and Nuclear Studies (IPNS), High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki 305-0801, Japan*
- ⁴⁶ *Japan Aerospace Exploration Agency (JAXA), Institute of Space and Astronautical Science (ISAS), Sagami-hara, Kanagawa 252-5210, Japan*
- ⁴⁷ *The Graduate University for Advanced Studies (SOKENDAI), Miura District, Kanagawa 240-0115, Hayama, Japan*
- ⁴⁸ *University of California, San Diego, Department of Physics, San Diego, CA 92093-0424, U.S.A.*
- ⁴⁹ *Aurora Technology for the European Space Agency, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, Madrid, Spain*
- ⁵⁰ *Universidad Europea de Madrid, 28670, Madrid, Spain*
- ⁵¹ *Space Science Data Center, Italian Space Agency, via del Politecnico, 00133, Roma, Italy*
- ⁵² *Université Grenoble Alpes, CNRS, LPSC-IN2P3, 53, avenue des Martyrs, 38000 Grenoble, France*
- ⁵³ *Gran Sasso Science Institute (GSSI), Viale F. Crispi 7, I-67100, L'Aquila, Italy*
- ⁵⁴ *Istituto Nazionale di Fisica Nucleare-Laboratori Nazionali di Frascati (INFN-LNF), Via E. Fermi 40, 00044, Frascati, Italy*
- ⁵⁵ *Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, U.K.*
- ⁵⁶ *Dpto. de Física Moderna, Universidad de Cantabria, Avda. los Castros s/n, E-39005 Santander, Spain*
- ⁵⁷ *Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München, Scheinerstr. 1, 81679 München, Germany*
- ⁵⁸ *GRAPPA, Institute for Theoretical Physics Amsterdam, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands*
- ⁵⁹ *Suwa University of Science, Chino, Nagano 391-0292, Japan*
- ⁶⁰ *Okayama University, Department of Physics, Okayama 700-8530, Japan*
- ⁶¹ *Institute of Astrophysics, Foundation for Research and Technology – Hellas, Vasilika Vouton, GR-70013 Heraklion, Greece*
- ⁶² *Department of Physics and ITCP, University of Crete, GR-70013, Heraklion, Greece*

E-mail: unnif@astro.uio.no

ABSTRACT: We assess the computational feasibility of end-to-end Bayesian analysis of the JAXA-led *LiteBIRD* experiment by analysing simulated time ordered data (TOD) for a subset of detectors through the Cosmoglobe and **Commander3** framework. The data volume for the simulated TOD is 1.55 TB, or 470 GB after Huffman compression. From this we estimate a total data volume of 238 TB for the full three year mission, or 70 TB after Huffman compression. We further estimate the running time for one Gibbs sample, from TOD to cosmological parameters, to be approximately 3000 CPU hours. The current simulations are based on an ideal instrument model, only including correlated $1/f$ noise. Future work will consider realistic systematics with full end-to-end error propagation. We conclude that these requirements are well within capabilities of future high-performance computing systems.

KEYWORDS: Bayesian reasoning, CMBR experiments

ARXIV EPRINT: [2507.05324](https://arxiv.org/abs/2507.05324)

Contents

1	Introduction	1
2	The <i>LiteBIRD</i> experiment	3
3	Algorithm	4
3.1	Data model	5
3.2	Bayesian analysis and Gibbs sampling	6
4	Simulations	7
4.1	Galactic foregrounds	7
4.2	Extra Galactic sources	8
4.3	<i>LiteBIRD</i> instrument model and TOD simulation pipeline	8
4.4	Final data selection and compression	9
5	Results	10
5.1	Time-ordered data	10
5.2	Frequency maps	12
5.3	Computational costs	14
6	Summary and conclusions	14

1 Introduction

The Λ CDM model is the most successful model describing the Universe as we know it, with massive observational evidence [1–4] supporting its predictions. One feature of this theory is an early period of cosmic inflation during which the Universe expanded exponentially [5–10]. Fluctuations in the inflaton field would generate density perturbations, becoming the structures we see in the Universe today. An additional prediction is that tensor perturbations, or gravitational waves, would be generated during this early inflationary period. Such tensor perturbations would produce a curl pattern in the Cosmic Microwave Background (CMB), known as B modes [11, 12]. The detection of these B modes would be additional and definitive evidence for inflation, but due to their low signal strength pushes the requirements of the instruments, sensitivities, and data processing to extremely high levels.

The CMB radiation was released during the period of recombination, 380 000 years after the beginning of our Universe. Today, after the Universe has expanded for 13.8 billion years, this radiation has cooled to 2.7255 K [13]. After the first detection in the microwave regime in 1965 by ref. [14], there have been several CMB experiments, three of which were full-sky space missions. The latest state-of-the-art mission was the *Planck* space mission launched in 2009 [15], following up on the Cosmic Background Explorer (*COBE*) launched in 1989 [16, 17] and the Wilkinson Microwave Anisotropy Probe (*WMAP*) launched in 2001 [18]. *COBE*, *WMAP*, and *Planck* have all been immensely successful; *COBE* measuring the blackbody spectrum of the CMB monopole [16] and the temperature fluctuations of the CMB across the sky [17], *WMAP* giving us the age of the Universe and reducing the allowed Λ CDM parameter space by a factor of 68 000, and *Planck* giving us cosmic variance limited

observations in intensity and a sub-percent precision on all Λ CDM parameters except for the optical depth due to reionisation, τ .

However, the *WMAP* and *Planck* polarisation observations did not result in a detection of the predicted B -mode signal. The current upper limit of the tensor-to-scalar ratio r used to quantify the strength of the B -mode signal is constrained by the BICEP/Keck 2018 data at $r < 0.036$ [19], decreasing to 0.032 when combined with the latest *Planck* PR4 data [20]. For reference, some single-field slow-roll inflationary theories predict $r \simeq 0.003$ [21]. In order to set tighter constraints on r , the sensitivity of future experiments needs to improve by an order of magnitude, which sets far more stringent requirements on the sensitivity and precision of future instruments and data analysis pipelines.

One of the next-generation CMB experiments is the Japanese led Lite (Light) satellite for the studies of B -mode polarisation and Inflation from cosmic background Radiation Detection (*LiteBIRD*), due to be launched in the early 2030's [22]. *LiteBIRD*'s main goal is the detection of the inflationary B -mode signal with uncertainty of $\delta r \lesssim 0.001$ [23]. Its effective sensitivity will be $2 \mu\text{K arcmin}$, corresponding to the combined integration time of nearly 1000 *Planck* missions [22].

To process these data robustly and efficiently, the analysis pipelines required for the *LiteBIRD* data will need to be significantly improved compared to the *Planck* era. A major lesson learned from *Planck* was that the ultimate limiting factor in terms of overall sensitivity is the interplay between systematic uncertainties in the instrument and astrophysical sky models [24]. *LiteBIRD*'s much lower noise level will make it even more sensitive to systematic effects from the instruments than *Planck* was, and this underscores the need for robust data analysis pipelines that simultaneously account for all relevant effects to ensure a clean detection of the B -mode signal.

The BEYONDPLANCK collaboration developed a tool for global integrated end-to-end Bayesian data analysis, called **Commander3** [25]. This framework implements a comprehensive Gibbs sampler starting from time ordered data (TOD), and outputting high-level astrophysical and cosmological products, including cosmological parameters. This framework is sufficiently efficient to allow thousands of iterations through TOD calibrations, map-making, and component separation, and the resulting Monte Carlo samples map out a full joint posterior distribution. These samples can thereby replace traditional frequentist simulations for error propagation [26]. This framework was demonstrated with the *Planck* Low Frequency Instrument (LFI) data in ref. [27] and companion papers, resulting in new state-of-the-art LFI frequency maps between 30 and 70 GHz.

The Cosmoglobe¹ project aims to apply this global analysis process to as many radio, microwave, and sub-millimeter data sets as possible within an Open Source and supportive community. The ultimate goal of Cosmoglobe is thus to build a global joint model of the microwave sky, while properly accounting for as many instrument-specific systematic effects as possible. As a first step, ref. [28] jointly analysed *WMAP* and *Planck* LFI TOD in this framework. By using information from one data set to break degeneracies in the other, the frequency maps of both data sets were significantly improved. Particularly noteworthy was that the *WMAP* W -band polarisation sky map was excluded in the official *WMAP*

¹<https://cosmoglobe.uio.no>.

cosmological analysis due to low ℓ systematics [29], but after the analysis done by ref. [28] the improved map was, for the first time, consistent with all other comparable maps. Recently, the same framework has also been used to re-process the *COBE*/DIRBE data, resulting in improved frequency maps [30], stronger CIB monopole constraints [31], a better zodiacal light model [32], a deeper starlight model [33], and a novel three-component thermal dust model [34].

While these analyses have proved that the end-to-end joint analysis approach is indeed feasible for currently available data sets, it is not obvious whether the same approach can be scaled to the much higher computational demands of next-generation experiments. In this work, we therefore extend the Cosmoglobe `Commander3` code to support future *LiteBIRD* measurements, and apply this to the first generation of simulated *LiteBIRD* TOD generated by the *LiteBIRD* collaboration’s newly developed simulation pipeline. The main goal of this analysis is to assess the computational feasibility of global Bayesian analysis for *LiteBIRD* data. Specifically, we estimate the total computational cost of analysing the full 3-year mission with the full focal plane by extrapolating actual run times from a reduced data volume spanning one year of data and a subset of detectors.

For this analysis, we used the TOD produced by the end-to-end pipeline described in ref. [35]. This pipeline simulates the operations of the three instruments onboard the *LiteBIRD* spacecraft and saves the TOD in HDF5 files, which we then convert to be compatible with the established `Commander3` data format.

The rest of the paper is organised as follows. First, we give a brief introduction to the *LiteBIRD* experiment in section 2 before describing the algorithm in section 3. Next, the production of the time-ordered *LiteBIRD* simulations is explained in section 4. Finally, our results are presented in section 5, and conclusions are drawn in section 6.

2 The *LiteBIRD* experiment

LiteBIRD is an experiment led by the Japan Aerospace Exploration Agency (JAXA), with contributors in both Europe and North America. *LiteBIRD* is to be launched to the second Lagrange point of the Sun-Earth system in the early 2030’s. There it will map the microwave sky for three years to produce the cleanest full-sky map of the polarised sky available. For a full overview of the *LiteBIRD* experiment we refer the interested reader to ref. [22].

LiteBIRD’s main scientific goal is to measure the CMB polarisation to probe the inflationary *B*-mode signal. It aims at either making a discovery or ruling out large classes of currently viable inflationary models [e.g. ref. 36]. *LiteBIRD* also has several secondary scientific goals, including a greater understanding of the reionisation of the Universe, detection of cosmic birefringence, and detecting primordial magnetic fields.

The current *LiteBIRD* baseline design implements a total of 4508 bolometers distributed over 15 frequency bands between 34 and 448 GHz, with focal planes cooled to 0.1 K to have low thermal disturbances [22]. It consists of three different telescopes: the Low-Frequency Telescope (LFT; 34–161 GHz; [37]), the Mid-Frequency Telescope (MFT; 89–228 GHz) and the High-Frequency Telescope (HFT; 166–448 GHz; [38]). Some frequency bands are divided across two different telescopes to have a partial overlap between telescopes, and thus increase the ability to monitor systematic effects within a given frequency. The telescopes also includes

CENTER FREQUENCY [GHz]	BANDPASS $\delta\nu$ [GHz]	BEAM SIZE [ARCMIN]	TOTAL NUMBER OF BOLOMETERS	NET ARRAY [μ K RTS]	POLARISATION SENSITIVITY [μ K ARCMIN]	TELESCOPE
40	12	70.5	48	18.50	37.42	LFT
50	15	58.5	24	16.54	33.46	LFT
60	14	51.1	48	10.54	21.32	LFT
68	16	(41.6, 47.1)	(144, 24)	(9.84, 15.70)	16.87	LFT
78	18	(36.9, 43.8)	(144, 48)	(7.69, 9.46)	12.07	LFT
89	20	(33.0, 41.5)	(144, 24)	(6.07, 14.22)	11.30	LFT
100	23	(30.2, 37.8)	(144, 366)	(5.11, 4.19)	6.56	L/MFT
119	36	(26.3, 33.6)	(144, 488)	(3.8, 2.82)	4.58	L/MFT
140	42	(23.7, 30.8)	(144, 366)	(3.58, 3.16)	4.79	L/MFT
166	50	28.9	488	2.75	5.57	MFT
195	59	(28.0, 28.6)	(366, 254)	(3.48, 5.19)	5.85	M/HFT
235	71	24.7	254	5.34	10.79	HFT
280	84	22.5	254	6.82	13.80	HFT
337	101	20.9	254	10.85	21.95	HFT
402	92	17.9	338	23.45	47.45	HFT

Table 1. *LiteBIRD* parameters for instrument model V1 from July 2020 [22]. The bolometers are divided between two instruments where two numbers are listed.

a levitating continuously spinning half-wave plates (HWP) to improve polarisation angle coverage and decrease the instrumental $1/f$ noise [38].

The resolution varies from 70.5 arcmin for the lowest frequency band to 17.9 arcmin FWHM for the highest frequency band. The *LiteBIRD* simulations assume tophat bandpasses with relative widths of 0.23 or 0.30 for each band. The center frequency, bandpass width, beam size, total number of bolometers, noise equivalent temperature (NET) for the total detector array, polarisation sensitivity, and the telescope on which each frequency band is located can be found in table 1. The polarisation sensitivity takes into account an observation efficiency of $\eta = 0.77$ caused by an expected data loss from the observation duty cycle, cosmic rays, and an added margin.

The whole satellite itself will be spinning at 0.05 rotations per minute (rpm) at a spin angle $\beta = 50^\circ$ and a precession angle $\alpha = 45^\circ$ [39, 40]. The three HWPs will spin at 46, 39 and 61 rpm, respectively. The sampling rate for all three telescopes are 19 Hz. This leaves us with 8×10^{12} detector samples or almost 20 TB of uncompressed raw data in total for the sky signal, not including metadata [22].

As of today, *LiteBIRD* is the only funded next-generation full-sky CMB experiment and will observe B modes over the multipoles $2 \leq \ell \leq 200$. Therefore one of its advantages over ground based experiments is that it will be able to detect both the reionisation bump below multipoles of $\ell \simeq 10$ and the recombination peak at $\ell \simeq 80$, while ground based experiments are limited to only the recombination peak.

3 Algorithm

The code used to analyse the simulated *LiteBIRD* data in this work is called `Commander3`, and it uses an end-to-end Bayesian approach for the CMB analysis. Starting with TOD we do calibration and map-making, and component separation iteratively to sample the full parameter space. Only a brief overview of the `Commander3` algorithm will be presented here. For a full description, see refs. [27], [25] and [28].

3.1 Data model

Following the usual parametric Bayesian analysis approach, we start by writing an explicit data model, which for this paper is given for a detector j and time sample t as

$$d_{t,j} = g_{t,j} P_{tp,j} \left[B_{pp',j}^{\text{symm}} \sum_c M_{cj}(\beta_{p'}) \mathbf{a}_{p'}^c + s_j^{\text{orb}} \right] + \mathbf{n}_{t,j}^{\text{corr}} + \mathbf{n}_{t,j}^{\text{w}}. \quad (3.1)$$

Here $d_{t,j}$ is the measured data per time sample per detector. The current simulations assume ideal gain, giving us $g_{t,j} = 1$. The pointing matrix $P_{tp,j}$ contains $(1, \cos 2\psi, \sin 2\psi)$ for the column corresponding to the Stokes parameters $\{T, Q, U\}$ for pixel p and row corresponding to time t , and is zero elsewhere, where ψ is the polarisation angle, T is the intensity of the signal and $\{Q, U\}$ represent the polarised signal. The beam convolution matrix, B_j , contains both the beam and the pixel window function caused by our maps being pixelised with the HEALPix projection [41]. The simulations use perfect Gaussian beams, with no beam asymmetries. The mixing matrix, $M_{cj}(\beta_{p'})$, is dependent on the spectral parameters β , for component c and detector j relative to some reference frequency. This matrix is diagonal in pixel space. The amplitudes, $\mathbf{a}_{p'}^c$, for sky component c in pixel p is measured in brightness temperature units, and the orbital CMB dipole, s^{orb} , is measured in K_{cmb} . All bandpasses are assumed to be perfectly known. The noise is divided into correlated noise, \mathbf{n}^{corr} , and white noise, \mathbf{n}^{w} . The white noise level is caused mainly by thermal noise, and it is thus heavily temperature dependent with lower detector temperatures leading to a lower white noise level. It also decreases with the number of observations N_{obs} of pixel p , as $\mathbf{n}_p^{\text{w}} \propto 1/\sqrt{N_{\text{obs}}}$.

The astrophysical sky model adopted in this paper consists of CMB and foreground emission from our Galaxy. The relevant Galactic foregrounds for this analysis are synchrotron radiation, thermal dust emission and free-free emission. Both synchrotron and thermal dust are polarised while free-free is not, so free-free is only included in intensity.

We assume that the CMB (\mathbf{a}_{CMB}) spectral energy density (SED) can be approximated as a blackbody. Synchrotron radiation dominates the polarised signal at the lowest *LiteBIRD* frequencies, and is described by a power law model with a free amplitude per pixel, \mathbf{a}_s . Similarly, thermal dust emission dominates the highest frequencies, and is modelled with a modified blackbody spectral energy density multiplied with a free amplitude per pixel, \mathbf{a}_d . Free-free radiation (\mathbf{a}_{ff}) is based on the *Planck*-2015 model [42]. In total, our sky model SED in brightness temperature is described by

$$\begin{aligned} \text{SRJ} = & \mathbf{a}_{\text{CMB}} \frac{x^2 e^x}{(e^x - 1)^2} \\ & + \mathbf{a}_s \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s} \\ & + \mathbf{a}_d \left(\frac{\nu}{\nu_{0,d}} \right)^{\beta_d+1} \frac{e^{h\nu_{0,d}/kT_d} - 1}{e^{h\nu/kT_d} - 1} \\ & + \mathbf{a}_{\text{ff}} \left(\frac{\nu_{0,\text{ff}}}{\nu} \right)^2 \frac{g_{\text{ff}}(\nu; T_e)}{g_{\text{ff}}(\nu_{0,\text{ff}}; T_e)} \\ & + m_\nu, \end{aligned} \quad (3.2)$$

where \mathbf{a}_i are the amplitudes of the signal components at frequency ν , and $\nu_{0,i}$ are the reference frequencies for each component; $x = h\nu/kT_0$, where h and k are the Planck and Boltzmann constants respectively and $T_0 = 2.7255$ K is the CMB temperature [13]; β_i are the spectral indices for synchrotron and thermal dust and T_d is the thermal dust temperature; g_{ff} is the so-called Gaunt factor for the free-free emission described by ref. [43] dependent on the electron temperature T_e ; and m_ν is the CMB monopole. Explicitly, this model defines the sum over components in eq. (3.1) with $\beta \equiv \{\beta_s, \beta_d, T_d, T_e\}$.

The lowest frequency, 40 GHz, is chosen as the reference frequency for synchrotron, the highest frequency, 402 GHz, is set as the thermal dust reference frequency, and 50 GHz is chosen as reference frequency for free-free. Fitting spectral indices will be a minor step in terms of computational costs compared to the TOD sampling, as *LiteBIRD* has a relative low spatial resolution and a large data volume and will thus be highly TOD dominated in terms of analysis costs, see section 5.3 for more details. The spectral indices for synchrotron and dust as well as the dust temperature are therefore kept at their true input values for simplicity.

In total, our parameter space consists of $\omega = \{\mathbf{a}_{\text{CMB}}, \mathbf{a}_s, \mathbf{a}_d, \mathbf{a}_{\text{ff}}, m_\nu, \mathbf{n}_{\text{corr}}\}$ where \mathbf{a}_{ff} is unpolarised. We use a TOD processing mask with $f_{\text{sky}} = 85\%$ generated using a sky cut-off after smoothing a 337 GHz band residual map from a previous iteration (that contained the unmodelled residual signal) with a 5° beam. We denote the set of all free parameters in this data model by ω .

We note that the sky model defined by eq. (3.2) defines the parametric model that is actually fitted to the data, and this is obviously much simpler than the real sky (e.g., [44–47]), and, indeed, even than the simulations analysed in this paper; see section 4.1. Choosing the optimal fitting model is typically an iterative process informed by the χ^2 as a function of sky position, and involves a trade-off between minimising residuals and model degeneracies. For instance, the current baseline *LiteBIRD* configuration has 15 main frequencies between 40 and 402 GHz, and it can therefore at most allow 14 free parameters without becoming fully degenerate. In practice, fewer parameters can be fitted robustly. On the other hand, joint analysis with external data can in principle increase this indefinitely.

3.2 Bayesian analysis and Gibbs sampling

Commander3 is built on a Bayesian framework, with the goal of estimating the posterior distribution $P(\omega | \mathbf{d})$ using Bayes' theorem,

$$P(\omega | \mathbf{d}) = \frac{P(\mathbf{d} | \omega)P(\omega)}{P(\mathbf{d})} \propto \mathcal{L}(\omega)P(\omega), \quad (3.3)$$

where $P(\mathbf{d} | \omega) \equiv \mathcal{L}(\omega)$ is the likelihood, $P(\omega)$ is the prior for all free parameters ω , and $P(\mathbf{d})$ is called the evidence. The latter acts simply as an overall normalisation constant, since it is independent of ω , and it is therefore omitted in the following.

Since the number of free parameters in ω is very large (on the order of millions), we resort to Markov Chain Monte Carlo (MCMC) methods to map out $P(\omega | \mathbf{d})$. In particular, Gibbs sampling [48, 49] serves as the main computational engine, while other methods such as Metropolis and Metropolis-Hastings sampling are also used for individual sampling steps. In Gibbs sampling, we sample at any given time only one parameter or a subset of parameters

while the rest are kept fixed. To explore the full parameter space this way, our sampling is executed iteratively. The Gibbs sampling chain in `Commander3` for this work is as follows:

$$\mathbf{n}_{\text{corr}} \leftarrow P(\mathbf{n}_{\text{corr}} \mid \mathbf{d}, \mathbf{a}) \quad (3.4)$$

$$\mathbf{a} \leftarrow P(\mathbf{a} \mid \mathbf{d}, \mathbf{n}_{\text{corr}}), \quad (3.5)$$

where the arrows means that the value for the sampled parameter is drawn from the distribution to the right. Both of these steps are performed with Gaussian multivariate samplers, as described by refs. [50] and [51], respectively.

Finally, we note that a future full *LiteBIRD* analysis will obviously include many more parameters than those considered here, and each of those will require a conditional separate sampling step in the above Gibbs chain. However, experience from previous BEYONDPLANCK and Cosmoglobe analyses shows that the correlated noise sampler (which performs the same function as a maximum likelihood mapmaker in a traditional pipeline) and the sky signal amplitude sampler (which is the core of the component separation procedure) dominate the overall runtime. This will be even more true for *LiteBIRD*, for which the TOD volume will be much higher than for current experiments, though the resolution is actually lower. We do not anticipate new sampling steps will be more expensive than the existing ones, nor do we expect a big effect on our burn-in time from these additions. For the purposes of this paper, which focuses on computational feasibility, the two steps discussed above are therefore by far the most essential.

4 Simulations

We analyse the first generation time-ordered data produced with the dedicated *LiteBIRD* Simulation Framework (LBS)² [52]. The main instrumental properties of the *LiteBIRD* instruments are summarised in table 1. Based on these, maps of CMB and foregrounds have been co-added, tophat bandpass-integrated and beam convolved for the instrumental beam for each frequency band to represent our microwave sky. The LBS is then used to generate TOD simulations for each detector given *LiteBIRD*'s scanning strategy and the specifics of each detector. We will now go into more detail regarding the sky maps for Galactic foregrounds in section 4.1 and extra Galactic sources in section 4.2, before we present the simulation pipeline and the specifics of our simulations.

4.1 Galactic foregrounds

The simulated sky maps have been generated using the map-based simulation module (MBS) provided by LBS, which acts as a wrapper around the Python Sky Model (PySM)³ package [53, 54]. Specifically, the model considered accounts for realistic Galactic foreground emission maps using the PySM `d1s1a1f1co1` model. A thermal dust component (`d1`) is modelled as a single-component modified blackbody based on the maps from the *Planck* 2015 analysis [42] with varying temperature and spectral index across the sky. A synchrotron component (`s1`) is modeled as a power law also with a varying spectral index but no curvature,

²https://github.com/litebird/litebird_sim.

³<https://github.com/galsci/pysm>.

based on Haslam 408 MHz and WMAP data [55, 56]. An unpolarised Anomalous Microwave Emission (AME) component (**a1**) is included as two spinning dust populations from the [42] model. Unpolarised free-free emission (**f1**) is also modeled based on ref. [42] in terms of spatial morphology, but employing a constant spectral index of $\beta_{\text{ff}} = -2.14$ rather than the full non-linear model. Finally, a model of Galactic CO emission (**co1**) is modeled involving the first three CO rotational lines, i.e. $J : 1 - 0, 2 - 1, 3 - 2$, whose center frequencies are respectively at $\{115.3, 230.5, 345.8\}$ GHz [57, 58].

4.2 Extra Galactic sources

The modelling of extra Galactic emission has been recently included in the PySM models. In the *LiteBIRD* simulations we have employed consistently the models coming from the WebSky simulations [59], where the emission is obtained by populating dark matter haloes with halo-occupation distribution consistent with the observations. Starting from the dark matter distributions, ref. [59] produced cosmic infrared background (CIB), thermal (tSZ) and kinetic Sunyaev-Zeldovich emission (kSZ), and a lensing convergence map.

Radio sources have recently been included by ref. [60], with fluxes (both in total intensity and polarisation) to be consistent with the latest observations (see ref. [60, figure 5] and ref. [61]) and latest number counts [62]. Finally, we synthesise a CMB map from the *Planck* 2018 cosmological parameters, and lens it with the WebSky convergence map, so that the lensing component correlates with the other large-scale structure probes.

4.3 *LiteBIRD* instrument model and TOD simulation pipeline

The simulations used in the following have been produced using the LBS, and the current paper represents the first non-LBS application of these TOD simulations. Indeed, an important motivation for the current paper is simply to confirm that these simulations can be readily used by external codes.

The LBS interfaces with the Instrument Model database, the IMo,⁴ that contains information about the *LiteBIRD* satellite, such as parameters for the scanning strategy, HWP rotation speed, detector noises and instrumental beams. The first step of the simulation pipeline is to load the IMo information. With this, the pipeline computes the quaternions to pass from the ecliptic to the detectors reference frames. The quaternions and the HWP rotation speed are finally used to calculate the pointing information (colatitude and longitude) and the polarisation angle for each detector considered in the simulation for the whole duration of the simulation. The following TODs are then produced and saved to disk:

- **tod_cmb**, containing the CMB signal obtained by scanning the input CMB maps using the pointing information,
- **tod_fg**, containing the foregrounds signal obtained by scanning the input foregrounds maps using the pointing information,
- **tod_dip**, containing the CMB dipole signal, (the dipole model used, following the LBS definition, is **TOTAL_FROM_LIN_T**⁵),

⁴https://github.com/litebird/litebird_imo (accessible by *LiteBIRD* members upon request).

⁵<https://litebird-sim.readthedocs.io/en/master/dipole.html>.

- `tod_wn`, containing white noise obtained using the detector noise information,
- `tod_wn_1f_30mHz`, containing correlated $1/f$ and white noise obtained using the detector noise information and the realistic case with a knee frequency f_{knee} of 30 mHz,
- `tod_wn_1f_100mHz`, same as `tod_wn_1f_30mHz`, but with the pessimistic knee frequency f_{knee} of 100 mHz.

All the above components are produced individually, and later co-added for a full TOD. Due to memory and CPU hours budget limitations, not all the *LiteBIRD*'s detectors have been considered, nor the entire mission time (3 years). Specifically, for frequency channels with more than 48 detectors, only about one third of all bolometers were included, choosing a configuration that preserves the focal plane symmetry; for channels that have fewer than 48 detectors, all bolometers were simulated. Furthermore, only one third of the mission time (1 year) has been simulated.

4.4 Final data selection and compression

The expected amount of data from *LiteBIRD* far exceeds that of any previous CMB satellite mission. For instance, while *Planck* LFI had 22 radiometers observing for 4 years with a sampling rate between 32 and 79 Hz [15], *LiteBIRD* will have a total number of 4508 bolometers observing for three years with a sampling rate of 19 Hz. To comfortably fit the simulations on our nodes with 1.5 TB of memory during the Bayesian Commander-based analysis, we adopted 1 TB as the maximum data volume for our simulated *LiteBIRD* data, and included only a subset of the full simulated data volume, comprising 4 detectors at each distinct channel. The noise level per detector has been scaled down to account for both the decreased number of detectors and the reduced observation time, so that the total sensitivity per frequency band will be comparable to the full experiment.

Our analysis includes a total of 88 detectors and one year of observations. These data are stored as Hierarchical Data Format (HDF) files including pointing, flags, TOD values and polarisation per sample, as well as common meta data describing the instrument. As described by ref. [25], we compress these data to minimise the memory footprint, noting that all data are stored in RAM during processing. The total data volume for the selected data set is 1.55 TB without compression and 470 GB with compression (see table 2), compression, thus, saves a factor of 3.4 in terms of RAM requirements. In addition, we note that the current LBS TOD do not yet include an analog-to-digital conversion (ADC) step, and the resulting samples are therefore currently stored as floating point numbers rather than integers. Integers are obviously much more compressible than floats, and once a proper ADC procedure has been implemented, the compression efficiency will increase further.

Scaling the current data to the full experiment with 4508 detectors and 3 years of data results in a total data volume of 238 TB without and 70 TB with compression. The scaling should, however, be slightly better than linear. Housekeeping data, for example, does not scale with the number of detectors, and the compression implemented in this work was originally optimised for the *Planck* LFI data. Other compression alternatives are currently being considered for *LiteBIRD*, and 70 TB is thus a strict upper estimate for the final compressed data volume.




Figure 1. Time ordered data for 1 hour of samples for the 40 GHz channel. The plot includes the total signal (CMB, noise and the foreground emissions) for the simulated TOD (blue). Also plotted is the total signal with the scaled down noise levels that compensate for reducing the number of detectors from 48 to 4 (orange). Units are in millikelvin.

5 Results

We now present the results from the analysis on the simulations described in section 4 using the method described in section 3. Note that there are more foreground components present in the simulations than we currently fit for using `Commander3`. In order to keep the analysis straightforward, only the most dominant foregrounds have been included. Doing a complete analysis with more complicated foreground components, and instrumental and systematic effects is left for future work.

During the course of the following analysis, we have produced several hundreds Gibbs samples involving all *LiteBIRD* channels. However, in the following we show results from one frequency channel for each telescope, where we have chosen the lowest (40 GHz) and the highest frequency (402 GHz) for LFT and HFT respectively, and for MFT we have chosen the frequency with the lowest total noise values, the 119 GHz channel. Intermediate channels look qualitatively similar, and different Gibbs samples are indistinguishable up to instrumental noise. We first inspect the TODs before we show the resulting frequency maps and finally the data volume and the running costs needed to do the analysis.

5.1 Time-ordered data

Figure 1 shows the total TOD including noise, foreground emission and the CMB signal for one hour of data for one detector in the LFT 40 GHz channel. The blue line shows the total signal of a single detector from the simulation containing the full set of detectors. The expected noise level from real data will have a $\sqrt{3}$ stronger noise signal due to the noise reduction we applied when scaling the mission duration from three to one year. The orange line shows the same TOD segment but with the reduced noise level as described in section 4.4.

Three hours of simulated TOD for one detector each in three different frequency channels can be seen in figure 2; LFT 40 GHz in the top panel, MFT 119 GHz in the middle panel and HFT 402 GHz in the bottom panel, all for the same three hours of data. The plots show the scaled down noise including both white and correlated noise (blue), the total foreground

	88 DETECTORS	4508 DETECTORS
	1 YEAR	3 YEARS
<i>Data volume</i>		
Uncompressed TOD volume	1.55 TB	238 TB
Compressed TOD volume	470 GB	70 TB
Non-TOD-related RAM usage	230 GB	230 GB
Total RAM requirements	700 GB	70 TB
<i>Processing time (cost per run)</i>		
TOD initialisation	9.5 h	1500 h
Other initialisation	6.7 h	7 h
Total initialisation	16.3 h	1500 h
<i>Gibbs sampling steps (cost per sample)</i>		
Huffman decompression	0.6 h	90 h
TOD projection (P operation)	1.2 h	200 h
Orbital dipole (\mathbf{s}_{orb})	0.5 h	80 h
Correlated noise sampling (\mathbf{n}_{corr})	6.1 h	1000 h
TOD binning (P^t operation)	0.1 h	10 h
Sum of other TOD processing	9.8 h	1500 h
TOD processing cost per sample . .	18.3 h	3000 h
Amplitude sampling, $P(\mathbf{a} \mathbf{d}, \omega \setminus \mathbf{a})$. . .	5.9 h	10 h
Spectral index sampling, $P(\beta \mathbf{d}, \omega \setminus \beta)$		20 h
Other steps	1.1 h	1 h
Total cost per sample	25.4 h	3000 h

Table 2. Computational costs for 88 detectors and 1 year (this work) in the left column and estimations for the full mission with 4508 detectors and 3 years in the right column. The costs are split into RAM requirements and processing costs in terms of CPU hours. The processing costs are split into costs per run for initialisation steps and costs per sample for TOD processing and for component separation. The numbers are based on running the analysis on a 2.6 GHz AMD cluster node with 128 cores and 2 TB of memory.

signals (green) and the CMB signal (orange) individually. The first thing to be noticed are the green peaks corresponding to the scanning of the Galactic plane which typically happens several times each hour. The noise levels in all channels have been scaled down to preserve the noise level in the resultant maps while compensating for the reduced mission duration and number of channels in the simulations, and thus the noise shown in these three plots will not be representative of what we expect for the real data. It will not be comparable between the frequency channels either since the 119 GHz channel with its 488 detectors and the 402 GHz channel with its 338 detectors will have their detector noise level scaled down by approximately a factor of 19 and 16, respectively, while the 40 GHz channel with its 48 detectors will scale the noise down by a factor of 6. The CMB and foreground signals are however representative for how they will be observed by each of the three frequency bands.

The signal amplitudes in figure 2 are all given in millikelvin, but the strength of the signal varies for each frequency. The thermal dust dominated 402 GHz channel has a 70 times stronger foreground signal than the 119 GHz channel with the weakest foreground signal, while the synchrotron dominated 40 GHz has a 3.5 times stronger foreground signal.

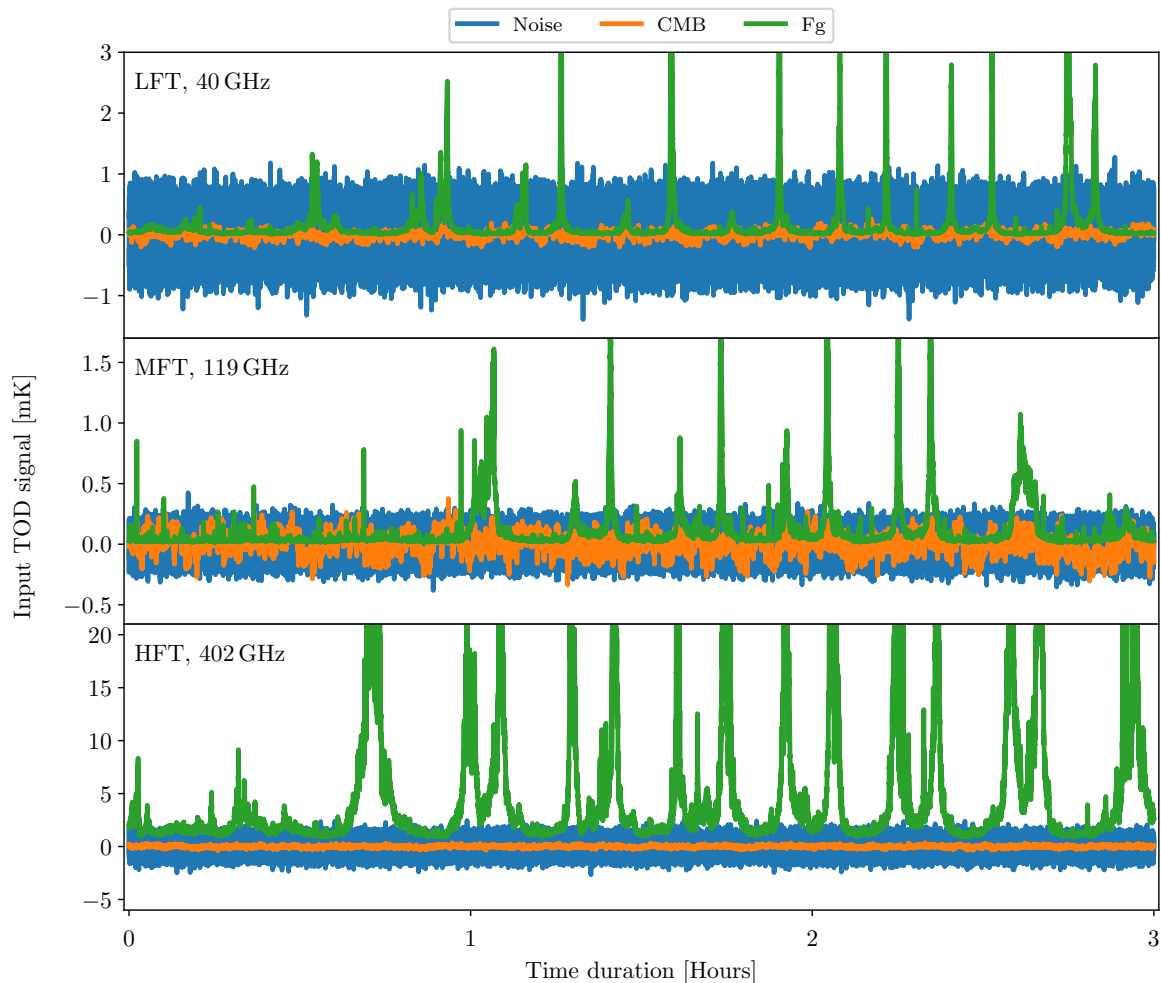


Figure 2. Time ordered data for the same three hours worth of samples for one detector for the LFT 40 GHz channel (top), the MFT 119 GHz channel (middle) and the HFT 402 GHz channel (bottom). The plots include the CMB (orange), white and correlated noise (blue) and the foregrounds (green). The peaks correspond to scanning the Galaxy plane, where synchrotron dominates the lowest frequency, the middle frequency has the lowest level of foregrounds, and the highest frequency channel foregrounds are dominated by thermal dust.

Figure 2 also shows the MFT and HFT channels scanning the Galactic plane simultaneously as these two telescopes point in the same sky direction, while the foreground peaks in the LFT channel are as expected located at different times due to the LFT channel pointing in a direction that is 100° separated from the HFT [22, 40]. This can also be observed in figure 3 where the dipole is added to the total TOD signal and the dipole varies with the opposite sign for the LFT channel as for the frequency channels on the other two telescopes.

5.2 Frequency maps

Figure 4 shows sky maps for three frequency bands, with the temperature maps (Stokes I) in the left column and polarisation (Stokes Q and U) in the middle and right columns. The LFT 40 GHz is in the top row, MFT 119 GHz in the middle row and HFT 402 GHz frequency channel in the bottom row. All maps are plotted at $N_{\text{side}} = 512$ in Galactic coordinates.




Figure 3. Time ordered data for 40 min for the 40 GHz channel (blue), the MFT 119 GHz channel (orange), and the 402 GHz channel (green) including CMB, foregrounds, white and correlated noise, and the dipole. The dipole signal causes the large sinusoidal pattern and the other components are responsible for the thickness of the lines. The spikes of the 402 GHz channel are foreground signals. The dipole in the LFT channel has the opposite sign compared to the MFT and HFT, which is caused by the telescope pointing in nearly opposite sky direction.

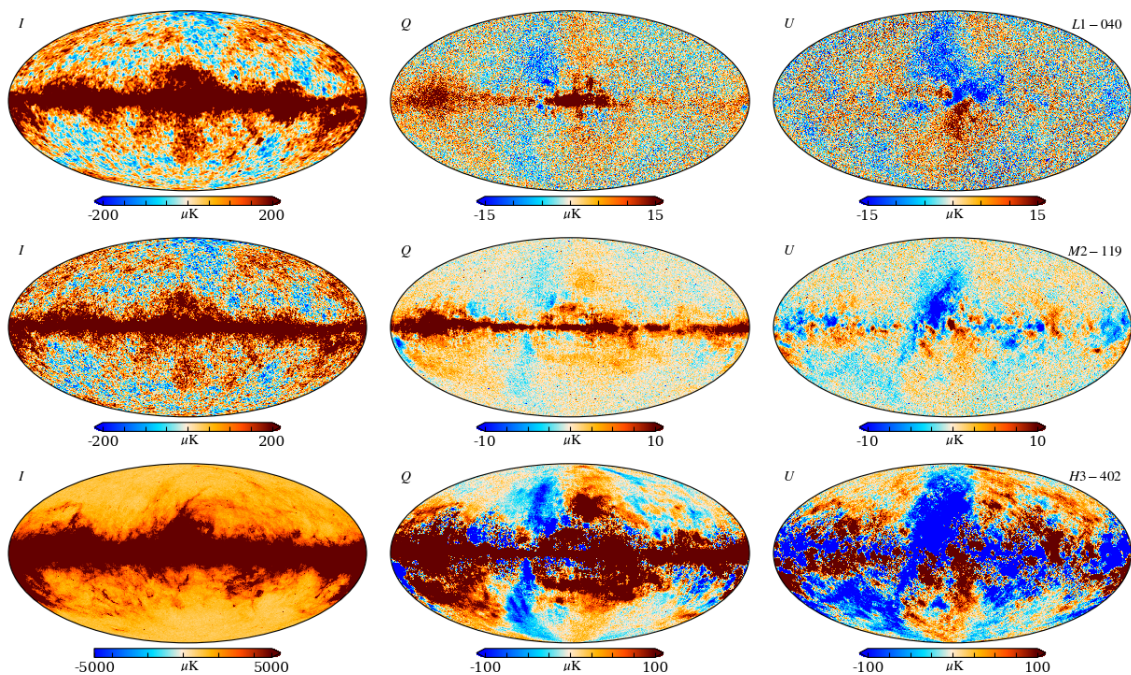


Figure 4. Frequency maps of a single Gibbs sample for three frequency bands for 40 GHz (top), 119 GHz (middle) and 402 GHz (bottom), represented by Stokes I (left column), Q (middle) and U (right column) parameters. The lowest frequency can be seen to be dominated by synchrotron while the highest one is dominated by thermal dust. The middle frequency has the lowest sky signal. All units are in microkelvin.

The polarised 40 GHz maps can be seen dominated by synchrotron and the 402 GHz maps are dominated by thermal dust, while the 119 GHz maps have a mix of both foreground signals. We clearly see the Galactic plane and we can also see the white noise in the 40 GHz maps as the sky signal is weaker than the 402 GHz sky signal and the polarisation sensitivity is significantly lower than for the 119 GHz channel, see table 1. The strength of the foreground signals also varies. Note for instance that the polarised 402 GHz signal is plotted with a 10 times higher temperature range than for the 119 GHz channel, in accordance with the TOD plots in figure 2.

5.3 Computational costs

Next, we consider the main result from this paper, namely the computational costs for running `Commander3` on the simulated *LiteBIRD* TOD. These are listed in table 2 for the main operations in units of CPU hours. Total computational costs are divided into one-time processing costs for a single run (including initialisation steps) and costs per Gibbs sample. The cost of each Gibbs sampling step has been further divided into TOD processing step and parameter sampling (component separation) step. Adding more data in time domain will not affect the cost of the component separation step as the component separation happens in map space, and the number of maps will stay fixed as long as we only work with the *LiteBIRD* data set.

In the current run, all maps are pixelised with a HEALPix [41] resolution parameter of $N_{\text{side}} = 512$ (corresponding to $7'$ pixels), but we anticipate that $N_{\text{side}} = 1024$ will be used for the highest frequency channels in a final production run. Accounting for this, we estimate that the foreground amplitude sampling step for the full *LiteBIRD* data set will increase from 6 to 10 CPU hours per Gibbs sample.

With our reduced data set the TOD processing costs are of the same order of magnitude as the component separation costs, as shown in table 2. Scaling up the amount of TOD to the full *LiteBIRD* experiment, however, makes the component separation part of the analysis strongly subdominant. The most expensive single Gibbs sample step in terms of computational cost is the correlated noise sampling with 6 CPU hours or 24% of the total CPU hours per Gibbs sample, as was seen in previous analysis [25]. Second to this are the other components grouped into “other TOD processing”, which is an upper limit on all the other steps required to sample the *LiteBIRD* TOD data model. Scaling linearly from the reduced to the full *LiteBIRD* data volume, we estimate that a full Gibbs sample will cost 3000 CPU-hours, which would put the cost of a 1000-sample production run at about 3M CPU-hours.

6 Summary and conclusions

The main goal of this paper was to assess the computational feasibility of running simulated *LiteBIRD* TOD through `Commander3` doing iterative end-to-end Bayesian analysis. By modifying `Commander3` to fit the *LiteBIRD* specific data characteristics, and run it on a partial simulation of what real *LiteBIRD* data would look like, we have demonstrated the conceptual feasibility of such an analysis. Furthermore, estimates of the expected final *LiteBIRD* data volume compared with the data volume used in this paper shows that such

an end-to-end analysis is also computationally feasible. For an analysis of the same data within a traditional simulation-based framework, we refer the interested reader to ref. [35].

The analysis for 88 detectors and one year of observations required 700 GB of RAM and 23 wall hours per full Gibbs sample on a 128 core computer node with 2 TB RAM. We have shown that scaling this up to the full mission will require approximately 70 TB of RAM with the current implemented compression algorithm and 3000 CPU hours. Given the exponential development in computational power, it is likely that a final *LiteBIRD* analysis in ten years' time will have access to much larger computational resources, which should be able to achieve the rate of tens of samples/day that we are able to reach for our current analyses.

We have showed that using **Commander3** as an analysis pipeline for simulated *LiteBIRD* data is feasible. The next step is to implement support for all the *LiteBIRD* specific features, from half-wave plate non-idealities and gain fluctuations to beam and sidelobe modelling, and then perform a full end-to-end analysis, with final r estimation as the end goal.

Based on the current analysis, however, we can now conclude that, with a cost of around 3000 CPU-hours per Gibbs sample, a Bayesian Cosmoglobe-style analysis of *LiteBIRD* is within the capabilities of currently available high-performance computing centers, although with a wall-time that is counted in years rather than days or weeks, as has been the case with previous analyses. A final *LiteBIRD* analysis will likely have to incorporate additional parameters and algorithms to account for the increased complexity of the final data model. This could increase the final CPU cost by a factor of between 1 and 10, which will become clearer as more development work occurs. Accounting for 10 more years of technological and algorithmic development still to be had, exploiting ever more computing cores and RAM, a full end-to-end Bayesian data analysis of the full *LiteBIRD* data set will be feasible. The analysis can be done on a stand-alone data set and it can also include other data sets to take advantage of the Cosmoglobe approach demonstrated in ref. [28].

In anticipation of the demands of future datasets like *LiteBIRD*, we are currently in the process of developing **Commander4**, a massively parallelised version of the core **Commander** algorithms that can be run across a huge number of nodes on a supercomputing cluster. By optimising the distribution of data and computation across many cores, we can greatly improve our abilities to store large datasets in memory as well as process time ordered data in a distributed way. **Commander4** will relax the memory and computation time constraints present in the **Commander3** code, allowing this approach to be distributed over a large number of smaller nodes, instead of the large memory nodes that it is currently optimised for. A main goal of this effort is to perform a full Bayesian end-to-end analysis of *LiteBIRD* in the course of a few weeks or months.

Acknowledgments

The current work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement numbers 819478 (ERC; Cosmoglobe), 772253 (ERC; Bits2Cosmology), and 101007633 (MSCA; CMBInflate). Some of the results in this paper have been derived using healpy [54] and the HEALPix [41] package. 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’Etudes Spatiale (CNES), by the Centre National de la Recherche Scientifique (CNRS), and by the Commissariat à l’Energie 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 351037) and has received funding from the European Research Council (ERC) under the Horizon 2020 Research and Innovation Programme (Grant agreement No. 772253, 819478, and 101141621). The Spanish *LiteBIRD* phase A contribution is supported by MCIN/AEI/10.13039/501100011033, project refs. PID2019-110610RB-C21, PID2020-120514GB-I00, PID2022-139223OB-C21, PID2023-150398NB-I00 (funded also by European Union NextGenerationEU/PRTR), and by MCIN/CDTI ICTP20210008 (funded also by EU FEDER funds). 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 UK *LiteBIRD* contribution is supported by the UK Space Agency under grant reference ST/Y006003/1 - “LiteBIRD UK: A major UK contribution to the LiteBIRD mission - Phase1 (March 25)”. The US contribution is supported by NASA grant no. 80NSSC18K0132.

References

- [1] WMAP SCIENCE TEAM collaboration, *Results from the Wilkinson Microwave Anisotropy Probe*, *PTEP* **2014** (2014) 06B102 [[arXiv:1404.5415](#)] [[INSPIRE](#)].
- [2] PLANCK collaboration, *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **641** (2020) A9 [[arXiv:1905.05697](#)] [[INSPIRE](#)].
- [3] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [4] PLANCK collaboration, *Planck 2018 results. X. Constraints on inflation*, *Astron. Astrophys.* **641** (2020) A10 [[arXiv:1807.06211](#)] [[INSPIRE](#)].
- [5] A.H. Guth, *The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems*, *Phys. Rev. D* **23** (1981) 347 [[INSPIRE](#)].
- [6] K. Sato, *First Order Phase Transition of a Vacuum and Expansion of the Universe*, *Mon. Not. Roy. Astron. Soc.* **195** (1981) 467 [[INSPIRE](#)].
- [7] A. Albrecht and P.J. Steinhardt, *Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking*, *Phys. Rev. Lett.* **48** (1982) 1220 [[INSPIRE](#)].
- [8] A.D. Linde, *A New Inflationary Universe Scenario: A Possible Solution of the Horizon, Flatness, Homogeneity, Isotropy and Primordial Monopole Problems*, *Phys. Lett. B* **108** (1982) 389 [[INSPIRE](#)].

- [9] L.P. Grishchuk, *Amplification of gravitational waves in an isotropic universe*, *Sov. Phys. JETP* **40** (1975) 409 [[INSPIRE](#)].
- [10] A.A. Starobinsky, *Spectrum of relict gravitational radiation and the early state of the universe*, *JETP Lett.* **30** (1979) 682 [[INSPIRE](#)].
- [11] M. Kamionkowski, A. Kosowsky and A. Stebbins, *A Probe of primordial gravity waves and vorticity*, *Phys. Rev. Lett.* **78** (1997) 2058 [[astro-ph/9609132](#)] [[INSPIRE](#)].
- [12] U. Seljak and M. Zaldarriaga, *Signature of gravity waves in polarization of the microwave background*, *Phys. Rev. Lett.* **78** (1997) 2054 [[astro-ph/9609169](#)] [[INSPIRE](#)].
- [13] D.J. Fixsen, *The Temperature of the Cosmic Microwave Background*, *Astrophys. J.* **707** (2009) 916 [[arXiv:0911.1955](#)] [[INSPIRE](#)].
- [14] A.A. Penzias and R.W. Wilson, *A measurement of excess antenna temperature at 4080-Mc/s*, *Astrophys. J.* **142** (1965) 419 [[INSPIRE](#)].
- [15] PLANCK collaboration, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *Astron. Astrophys.* **641** (2020) A1 [[arXiv:1807.06205](#)] [[INSPIRE](#)].
- [16] J.C. Mather et al., *Measurement of the Cosmic Microwave Background spectrum by the COBE FIRAS instrument*, *Astrophys. J.* **420** (1994) 439 [[INSPIRE](#)].
- [17] G.F. Smoot et al., *Structure in the COBE differential microwave radiometer first year maps*, *Astrophys. J. Lett.* **396** (1992) L1 [[INSPIRE](#)].
- [18] C.L. Bennett et al., *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *Astrophys. J. Suppl.* **208** (2013) 20 [[arXiv:1212.5225](#)] [[INSPIRE](#)].
- [19] BICEP2 and KECK ARRAY collaborations, *BICEP2/Keck Array x : Constraints on Primordial Gravitational Waves using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season*, *Phys. Rev. Lett.* **121** (2018) 221301 [[arXiv:1810.05216](#)] [[INSPIRE](#)].
- [20] M. Tristram et al., *Improved limits on the tensor-to-scalar ratio using BICEP and Planck data*, *Phys. Rev. D* **105** (2022) 083524 [[arXiv:2112.07961](#)] [[INSPIRE](#)].
- [21] M. Kamionkowski and E.D. Kovetz, *The Quest for B Modes from Inflationary Gravitational Waves*, *Ann. Rev. Astron. Astrophys.* **54** (2016) 227 [[arXiv:1510.06042](#)] [[INSPIRE](#)].
- [22] LITEBIRD collaboration, *Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey*, *PTEP* **2023** (2023) 042F01 [[arXiv:2202.02773](#)] [[INSPIRE](#)].
- [23] A. Suzuki et al., *The LiteBIRD Satellite Mission: Sub-Kelvin Instrument*, *J. Low Temp. Phys.* **193** (2018) 1048 [[arXiv:1801.06987](#)] [[INSPIRE](#)].
- [24] PLANCK collaboration, *Planck 2018 results. II. Low Frequency Instrument data processing*, *Astron. Astrophys.* **641** (2020) A2 [[arXiv:1807.06206](#)] [[INSPIRE](#)].
- [25] M. Galloway et al., *BEYONDPLANCK — III. Commander3*, *Astron. Astrophys.* **675** (2023) A3 [[arXiv:2201.03509](#)] [[INSPIRE](#)].
- [26] M. Brilenkov et al., *BEYONDPLANCK — IV. Simulations and validation*, *Astron. Astrophys.* **675** (2023) A4 [[arXiv:2209.04437](#)] [[INSPIRE](#)].
- [27] BEYONDPLANCK collaboration, *BeyondPlanck I. Global Bayesian analysis of the Planck Low Frequency Instrument data*, *Astron. Astrophys.* **675** (2023) A1. [[arXiv:2011.05609](#)] [[INSPIRE](#)].
- [28] D.J. Watts et al., *COSMOGLOBE DR1 results — I. Improved Wilkinson Microwave Anisotropy Probe maps through Bayesian end-to-end analysis*, *Astron. Astrophys.* **679** (2023) A143 [[arXiv:2303.08095](#)] [[INSPIRE](#)].

- [29] N. Jarosik et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results*, *Astrophys. J. Suppl.* **192** (2011) 14 [[arXiv:1001.4744](#)] [[INSPIRE](#)].
- [30] D.J. Watts et al., *Cosmoglobe DR2. I. Global Bayesian analysis of COBE-DIRBE*, [arXiv:2408.10952](#) [[INSPIRE](#)].
- [31] D.J. Watts et al., *Cosmoglobe DR2. II. CIB monopole measurements from COBE-DIRBE through global Bayesian analysis*, [arXiv:2406.01491](#) [[INSPIRE](#)].
- [32] M. San et al., *Cosmoglobe DR2. III. Improved modelling of zodiacal light with COBE-DIRBE through global Bayesian analysis*, [arXiv:2408.11004](#) [[INSPIRE](#)].
- [33] M. Galloway et al., *Cosmoglobe DR2 IV. Modelling starlight emission in DIRBE with GAIA and WISE*, *Astron. Astrophys.*, in preparation (2025).
- [34] E. Gjerløw et al., *Cosmoglobe DR2 V. A compact model of large-scale thermal dust emission in DIRBE and Planck HFI without spatially varying SEDs*, *Astron. Astrophys.*, in preparation (2025).
- [35] M. Bortolami et al., *First release of LiteBIRD simulations from an end-to-end pipeline*, *JCAP* **11** (2025) 042 [[arXiv:2507.07122](#)] [[INSPIRE](#)].
- [36] A. Linde, *Gravitational waves and large field inflation*, *JCAP* **02** (2017) 006 [[arXiv:1612.00020](#)] [[INSPIRE](#)].
- [37] Y. Sekimoto et al., *Concept Design of Low Frequency Telescope for CMB B-mode Polarization satellite LiteBIRD*, *Proc. SPIE Int. Soc. Opt. Eng.* **11453** (2020) 1145310 [[arXiv:2101.06342](#)] [[INSPIRE](#)].
- [38] L. Montier et al., *Overview of the Medium and High Frequency Telescopes of the LiteBIRD satellite mission*, *Proc. SPIE Int. Soc. Opt. Eng.* **11443** (2020) 114432G [[arXiv:2102.00809](#)] [[INSPIRE](#)].
- [39] H. Sugai et al., *Updated Design of the CMB Polarization Experiment Satellite LiteBIRD*, *J. Low. Temp. Phys.* **199** (2020) 1107 [[arXiv:2001.01724](#)] [[INSPIRE](#)].
- [40] Y. Takase et al., *Multi-dimensional optimisation of the scanning strategy for the LiteBIRD space mission*, *JCAP* **12** (2024) 036 [[arXiv:2408.03040](#)] [[INSPIRE](#)].
- [41] K.M. Górski et al., *HEALPix — A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [42] PLANCK collaboration, *Planck 2015 results. X. Diffuse component separation: Foreground maps*, *Astron. Astrophys.* **594** (2016) A10 [[arXiv:1502.01588](#)] [[INSPIRE](#)].
- [43] C. Dickinson, R.D. Davies and R.J. Davis, *Towards a free-free template for CMB foregrounds*, *Mon. Not. Roy. Astron. Soc.* **341** (2003) 369 [[astro-ph/0302024](#)] [[INSPIRE](#)].
- [44] A.M. Meisner and D.P. Finkbeiner, *Modeling Thermal Dust Emission with Two Components: Application to the Planck High Frequency Instrument Maps*, *Astrophys. J.* **798** (2014) 88 [[arXiv:1410.7523](#)].
- [45] G.M. Green et al., *A 3D Dust Map Based on Gaia, Pan-STARRS 1, and 2MASS*, *Astrophys. J.* **887** (2019) 93.
- [46] H. Liu, J.-R. Li and Y.-F. Cai, *Evaluation of the Single-component Thermal Dust Emission Model in Cosmic Microwave Background Experiments*, *Astrophys. J. Suppl.* **276** (2025) 45 [[arXiv:2411.04543](#)] [[INSPIRE](#)].

- [47] L. Vacher et al., *How bad could it be? Modelling the 3D complexity of the polarised dust signal using moment expansion*, *Astron. Astrophys.* **697** (2025) A212 [[arXiv:2411.11649](#)] [[INSPIRE](#)].
- [48] S. Geman and D. Geman, *Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images*, *IEEE Trans. Pattern Anal. Machine Intell.* **PAMI-6** (1984) 721 [[INSPIRE](#)].
- [49] A. Gelman, J.B. Carlin, H.S. Stern and D.B. Rubin, *Bayesian Data Analysis*, Chapman and Hall/CRC (2003).
- [50] H.T. Ihle et al., *BEYONDPLANCK — VI. Noise characterization and modeling*, *Astron. Astrophys.* **675** (2023) A6 [[arXiv:2011.06650](#)] [[INSPIRE](#)].
- [51] K.J. Andersen et al., *BEYONDPLANCK — XIII. Intensity foreground sampling, degeneracies, and priors*, *Astron. Astrophys.* **675** (2023) A13 [[arXiv:2201.08188](#)] [[INSPIRE](#)].
- [52] M. Tomasi et al., *A Simulation Framework for the LiteBIRD Instruments*, *JCAP* **11** (2025) 040 [[arXiv:2507.04918](#)] [[INSPIRE](#)].
- [53] B. Thorne, J. Dunkley, D. Alonso and S. Naess, *The Python Sky Model: software for simulating the Galactic microwave sky*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 2821 [[arXiv:1608.02841](#)] [[INSPIRE](#)].
- [54] A. Zonca, B. Thorne, N. Krachmalnicoff and J. Borrill, *The Python Sky Model 3 software*, *J. Open Source Softw.* **6** (2021) 3783 [[arXiv:2108.01444](#)] [[INSPIRE](#)].
- [55] J. Delabrouille et al., *The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths*, *Astron. Astrophys.* **553** (2013) A96 [[arXiv:1207.3675](#)] [[INSPIRE](#)].
- [56] M.-A. Miville-Deschênes et al., *Separation of anomalous and synchrotron emissions using WMAP polarization data*, *Astron. Astrophys.* **490** (2008) 1093 [[arXiv:0802.3345](#)] [[INSPIRE](#)].
- [57] G. Puglisi, G. Fabbian and C. Baccigalupi, *A 3D model for carbon monoxide molecular line emission as a potential cosmic microwave background polarization contaminant*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 2982 [[arXiv:1701.07856](#)] [[INSPIRE](#)].
- [58] PLANCK collaboration, *Planck 2013 results. XIII. Galactic CO emission*, *Astron. Astrophys.* **571** (2014) A13 [[arXiv:1303.5073](#)] [[INSPIRE](#)].
- [59] G. Stein et al., *The Websky Extragalactic CMB Simulations*, *JCAP* **10** (2020) 012 [[arXiv:2001.08787](#)] [[INSPIRE](#)].
- [60] Z. Li, G. Puglisi, M.S. Madhavacheril and M.A. Alvarez, *Simulated catalogs and maps of radio galaxies at millimeter wavelengths in Websky*, *JCAP* **08** (2022) 029 [[arXiv:2110.15357](#)] [[INSPIRE](#)].
- [61] G. Puglisi et al., *Forecasting the Contribution of Polarized Extragalactic Radio Sources in CMB Observations*, *Astrophys. J.* **858** (2018) 85 [[arXiv:1712.09639](#)] [[INSPIRE](#)].
- [62] G. Lagache et al., *Impact of polarised extragalactic sources on the measurement of CMB B-mode anisotropies*, *Astron. Astrophys.* **642** (2020) A232 [[arXiv:1911.09466](#)] [[INSPIRE](#)].