

The Simons Observatory: validation of reconstructed power spectra from simulated filtered maps for the small aperture telescope survey

Carlos Hervías-Caimapo^a, Kevin Wolz^b, Adrien La Posta^b, Susanna Azzoni^c, David Alonso^b, Kam Arnold^{d,e}, Carlo Baccigalupi^{f,g,h}, Simon Biquardⁱ, Michael L. Brown^j, Erminia Calabrese^k, Yuji Chinone^{l,m}, Samuel Day-Weiss^c, Jo Dunkley^{c,n}, Rolando Dünner^a, Josquin Errardⁱ, Giulio Fabbian^k, Ken Gangaⁱ, Serena Giardiello^k, Emilie Hertig^{o,p}, Kevin M. Huffenberger^q, Bradley R. Johnson^r, Baptiste Jost^{m,s}, Reijo Keskitalo^{t,u}, Theodore S. Kisner^{t,u}, Thibaut Louis^v, Magdy Morshed^w, Lyman A. Page^c, Christian L. Reichardt^x, Erik Rosenberg^j, Max Silva-Feaver^y, Wuhyun Sohnⁱ, Yoshinori Sueno^c, Dan B. Thomas^j, Ema Tsang King Sangⁱ, Amalia Villarrubia-Aguilarⁱ and Kyohei Yamada^{c,z}

^a*Instituto de Astrofísica and Centro de Astro-Ingeniería, Facultad de Física, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile*

^b*Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, U.K.*

^c*Joseph Henry Laboratories of Physics, Princeton University, Jadwin Hall, Princeton, NJ, 08544, U.S.A.*

^d*Department of Astronomy & Astrophysics, University of California San Diego, San Diego, U.S.A.*

^e*Department of Physics, University of California San Diego, San Diego, U.S.A.*

^f*The International School for Advanced Studies (SISSA), via Bonomea 265, I-34136 Trieste, Italy*

^g*The National Institute for Nuclear Physics (INFN), via Valerio 2, I-34127, Trieste, Italy*

^h*The Institute for Fundamental Physics of the Universe (IFPU), Via Beirut 2, I-34151, Trieste, Italy*

ⁱ*Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France*

^j*Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Manchester M13 9PL, U.K.*

^k*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales, CF24 3AA, U.K.*

^l*QUP (WPI), KEK, Tsukuba, Ibaraki 305-0801, Japan*

^m*Kavli Institute for the Physics and Mathematics of the Universe (WPI), UTIAS,
The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*

ⁿ*Department of Astrophysical Sciences, Peyton Hall, Princeton University,
Princeton, NJ, 08544, U.S.A.*

^o*Institute of Astronomy, University of Cambridge,
Madingley Road, Cambridge, CB3 0HA, U.K.*

^p*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, U.K.*

^q*Mitchell Institute for Fundamental Physics & Astronomy and Department of Physics & Astronomy,
Texas A&M University, College Station, Texas 77843, U.S.A.*

^r*Department of Astronomy, University of Virginia, Charlottesville, VA 22904, U.S.A.*

^s*Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS, The University of Tokyo,
Kashiwa, Chiba 277-8583, Japan*

^t*Lawrence Berkeley National Laboratory, Berkeley, CA, U.S.A.*

^u*University of California Berkeley, Berkeley, CA, U.S.A.*

^v*Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*

^w*INFN Sezione di Ferrara, Via Saragat 1, 44122 Ferrara, Italy*

^x*School of Physics, The University of Melbourne, Parkville, VIC 3010, Australia*

^y*Wright Laboratory, Department of Physics, Yale University, New Haven, CT 06511, U.S.A.*

^z*Department of Physics, The University of Tokyo, Tokyo 113-0033, Japan*

E-mail: carlos.hervias@uc.cl

ABSTRACT: We present a transfer function-based method to estimate angular power spectra from filtered maps for cosmic microwave background (CMB) surveys. This is especially relevant for experiments targeting the faint primordial gravitational wave signatures in CMB polarisation at large scales, such as the Simons Observatory (SO) small aperture telescopes. While timestreams can be filtered to mitigate the contamination from low-frequency noise, usual methods that calculate the mode coupling at individual multipoles can be challenging for experiments covering large sky areas or reaching few-arcminute resolution. The method we present here, although approximate, is more practical and faster for larger data volumes. We validate it through the use of simulated observations approximating the first year of SO data, going from half-wave plate-modulated timestreams to maps, and using simulations to estimate the mixing of polarisation modes induced by an example of time-domain filtering. We show its performance through an example null test and with an end-to-end pipeline that performs inference on cosmological parameters, including the tensor-to-scalar ratio r . The performance demonstration uses simulated observations at multiple frequency bands. We find that the method can recover unbiased parameters for our simulated noise levels.

KEYWORDS: CMBR experiments, CMBR polarisation, gravitational waves and CMBR polarization

ARXIV EPRINT: [2502.00946](https://arxiv.org/abs/2502.00946)

Contents

1	Introduction	1
2	Methods	4
2.1	Filter+bin mapmaking	4
2.2	From maps to spectra	5
3	Simulations	11
3.1	Input simulated maps	11
3.2	Simulated dataset	12
3.3	Using TOAST	13
4	Results	17
4.1	Comparison of band power window functions	17
4.2	Full multipole range using the transfer function method	20
4.3	Limitations of the transfer function method	21
5	Null test	22
5.1	Definition of splits	23
5.2	CMB-only test	24
5.3	Signal and noise test	25
6	Full time-domain simulations	26
7	Summary and conclusions	29
A	Effect of individual filters	31
B	Illustration of focal plane thinning	32

1 Introduction

The cosmic microwave background (CMB) radiation can provide a wealth of information about the very early stages of our Universe. The cosmological community is allocating significant effort to measuring the CMB polarisation with unprecedented sensitivity, as outlined in e.g., [1, 2]. Primordial tensor perturbations, in the form of a background of primordial gravitational waves (PGWs) can potentially source a large-scale parity-odd B -mode component in the CMB polarisation. In contrast, primordial scalar/density perturbations produce only parity-even E -mode fluctuations at linear order (in addition to unpolarised intensity fluctuations) [3, 4]. The amplitude of the power spectrum of the primordial tensor perturbations is parametrised in terms of its ratio with respect to that of scalar perturbations, the so-called tensor-to-scalar ratio r . To date, no detection of primordial tensor fluctuations

has been made, with the tightest upper bound at 95% confidence being $r < 0.036$ for a pivot scale of $k_p = 0.05 \text{ Mpc}^{-1}$ [5].

A detection of primordial tensor perturbations, or a sufficiently tight upper bound on r , could elucidate many of the remaining questions about the beginning of our Universe, allowing us to discriminate between different models. Cosmic inflation, which posits a short phase of exponentially accelerated expansion at early times, is the most widespread of these models (e.g. [6–8]). Within this paradigm, primordial perturbations are seeded by quantum fluctuations in the spacetime metric that are expanded to macroscopic scales during the inflationary period (e.g. [9–12]). The presence of both scalar and tensor fluctuations (i.e. PGWs), is a key prediction of this theory (e.g. [13]). Current and upcoming CMB experiments such as BICEP/KECK [5, 14], BICEP Array [15], SPIDER [16], CLASS [17], POLARBEAR [18], QUBIC [19], GroundBIRD [20], AliCPT-1 [21], CMB-S4 [22], and LiteBIRD [23] are targeting or will target the measurement of these elusive primordial B -modes in the next decade.

Polarised foreground emission from our galaxy is a substantial complication for measuring r since it produces a B -mode signal much larger than the one from primordial B -modes. Thermal dust and synchrotron emission are the two main polarised components that could pose a problem (e.g. [24]). Their B -mode emission would be equivalent to a tensor-to-scalar ratio much higher than $r \sim 0.01$ [25], the target of many upcoming CMB experiments. By exploiting the difference in spectral distributions between the CMB and foregrounds, we can clean the former, using methods collectively known as component separation [26]. However, small residuals left after cleaning may still pose a problem for measuring small values $r \sim 0.001$ in the near future, since the residuals can be as large as the primordial signal we want to measure in the first place. Another relevant foreground contamination are the coherent distortions induced by the deflection of CMB photons passing through the gravitational potential wells of the cosmic large-scale structure, known as CMB weak lensing [27]. B -modes are sourced from primordial CMB E -modes at intermediate and small scales [28], effectively creating a lensing B -mode spectrum that will dominate over the primordial B -modes in all but the largest scales, depending on the value of r . For low enough values of r , this lensing-induced noise will dominate the sample variance and limit the precision of the cosmological parameters we can measure. We can mitigate this with the technique of delensing, in which we subtract a template of the realization of lensing B -modes, reconstructed from an estimator of the lensing potential and observations of the E -modes at intermediate and high resolution [29, 30].

The Simons Observatory (SO) is a ground-based CMB experiment, located at the Cerro Toco site in the Atacama desert in Chile. It targets an assortment of science cases in two main surveys [2]. First, a 6-m aperture Large Aperture Telescope (LAT) will focus on small angular-scale CMB science, including high- ℓ power spectra and bispectra, secondary anisotropies such as CMB lensing and the Sunyaev-Zel’dovich effect, as well as extragalactic point sources, galaxy clusters, and the time-variable millimetre sky. A second deep survey is being carried out by the Small Aperture Telescopes (SATs), with the main target of measuring or constraining the PGW B -mode signal. The SAT observations are covering $\sim 10\%$ of the sky in the same six frequency bands targeted by the LAT, centred at 27, 39 GHz (LF); 93, 145 GHz (MF); and 225, 280 GHz (UHF), aiming to measure the primordial B -mode recombination bump ($\ell \sim 80$) from the ground at scales $30 \leq \ell$, constrained by low-frequency

noise. The forecasted instrumental sensitivity after 5 years of observations would translate into a measurement of r with $\sigma(r) \simeq 0.002\text{--}0.003$ after component separation [31].

The nominal SO-SAT experiment comprises 3 small refractive telescopes with a 42 cm aperture and a field of view of 35° . Each telescope has a 40 K cryogenic rotating half-wave plate (HWP), 1 K optics, and a ~ 100 mK focal plane with 7 wafers and $\sim 12,000$ Transition Edge Sensors [32]. The nominal instrument configuration will later be expanded to include 3 additional telescopes, potentially using different detector technologies. The U.K. and Japan are supporting this major enhancement through the SO:UK (two MF SATs) and SO:Japan (one LF SAT) projects, respectively. In this work, we will focus on the two nominal MF bands, 93 and 145 GHz bands (which we will refer to as f090 and f150, respectively), where most of the sensitivity to the CMB emission is concentrated.

Ground-based observatories contend with multiple sources of low-frequency noise that can affect measurements of the primordial B -mode signal. For example, temporal and spatial fluctuations in the emission of the atmosphere due to turbulence and uneven distribution of water vapour, which creates non-trivial correlations between detectors (e.g. [33]), as well as by intensity-to-polarisation (I-to-P) leakage due to factors such as thermal drifts, optical mismatch, etc. When doing detector pair differencing with orthogonal antennas, bandpass or beam mismatches arise (e.g. [34, 35]). To address these difficulties, the SO-SATs modulate the polarisation signal at frequencies where detector white noise dominates, above atmospheric fluctuations and thermal drifts, using a cryogenic HWP rotating around $f \approx 2$ Hz [36]. This technique also eliminates the need for detector pair differencing, allowing each detector to independently measure the three Stokes parameters (I, Q, U) [37]. Previous experiments that have used HWP modulation include MAXIPOL [38], ABS [39], and POLARBEAR [40], while CLASS uses a different modulation technology, the variable-delay polarization modulator (VPM [41]). A demodulation at the modulating fourth harmonics of the HWP rotating frequency, followed by a low-pass filter, is the standard procedure to recover the polarised signal (e.g. [42]). A further optional high-pass filter can be applied to the polarised timestreams to suppress the remaining low-frequency $1/f$ noise. The impact of all this filtering must be accounted for when generating maps from the filtered timestreams, or the associated bias to the maps must be taken into account and corrected when analysing them (e.g. [38, 43–46]). General filtering to remove low-frequency sources of noise is a standard practice in CMB science, even without polarised modulation, and these effects therefore need to be accounted for in cosmological analyses (e.g. [47–52]).

The set of issues described above can be stated generically as the fact that filtered maps are biased due to one or more timestream operations. The objective of this paper is to describe a procedure able to estimate angular power spectra from biased maps from e.g., simulated SO-SAT observations. Specifically, the method described below is based on the estimation of a transfer function describing the loss of power as well as mixing between different polarisation channels due to timestream filtering. This follows similar methods as used for e.g. ABS [39]. We will compare this approach to the exact, but more computationally expensive “per- ℓ method” that reconstructs the impact of filtering on the full band power window functions. This method, used in past experiments [50, 51], becomes prohibitively expensive, at least in its brute-force incarnation, for the large data volume that SO will cover. Since the transfer function approach involves non-trivial approximations, in this paper we

demonstrate its validity to obtain unbiased power spectra within the uncertainties anticipated for the SO-SAT 1-year deep survey. We also exemplify the limits where the approximations break the method. We note that the pipeline validated here is not the same and uses more aggressive filtering than we anticipate for the analysis of the real SO data.¹ The treatment of foreground complexity and component separation is left to future work (see ref. [31] for further details). We validate the methods over realistic simulated timestreams, as opposed to using map-based simulations that a forecast analysis would normally use.

In section 2, we describe in detail the procedures used to generate filtered maps and their associated unbiased power spectra. In section 3, we describe the simulations used to validate our methodology, including the procedure used to simulate SO-SAT observations. Section 4 presents our results, validating the methods presented in the previous sections and quantifying the residual biases in the measured spectra. In section 5, we demonstrate that this procedure can be used to carry out highly sensitive null tests for the target noise level by the SO-SAT 1-year survey with the target number of detectors. Section 6 presents an end-to-end analysis from simulated observations to a measurement of r after component separation. We present our conclusions in section 7. All temperatures used in this work are in thermodynamic units relative to the CMB.

2 Methods

2.1 Filter+bin mapmaking

Ground-based CMB telescopes typically observe the sky at fixed elevation, scanning back and forth in azimuth, and taking samples at a regular rate with thousands of detectors in the telescope focal plane. While each individual sample is noise-dominated, the sky signal can be teased out by accumulating samples collected over time from the same sky coordinates in a process called “mapmaking” (e.g. [53]). The relation between a given measured timestream and the true sky is given by

$$\mathbf{d} = \mathbf{P}\mathbf{m} + \mathbf{n}, \quad (2.1)$$

where \mathbf{d} is a vector containing the observed time-ordered samples, \mathbf{P} is the pointing matrix, which is sparse and relates time samples to coordinates in the sky where the telescope pointed, \mathbf{m} is the true sky map convolved by the corresponding bandpass and beam, and \mathbf{n} is the timestream noise.² \mathbf{P} also takes into account the polarisation sensitivity of each detector, i.e. the polarisation angle of the detector projected into sky coordinates, and how much Q/U polarisation a detector measures in a given time sample.

The formal least-square solution to equation (2.1) is the mapmaking equation, which is linear, minimises the variance, and gives an unbiased³ estimation of \mathbf{m} :

$$(\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P}) \hat{\mathbf{m}} = \mathbf{P}^T \mathbf{N}^{-1} \mathbf{d}, \quad (2.2)$$

¹Our filtering choices are limited by our software and their implementation of the observation matrix, as will be apparent below. However, we will have no such constraint for the first analysis of real SO data and have much more freedom to implement any required filtering.

²This model is simplified as it assumes no difference in beams or bandpasses between detectors. In this study, we will assume that the detector beams and bandpasses are identical and constant in time.

³While this solution is unbiased in the absence of model error bias, real-world experiments are prone to complex effects such as biases from sub-pixel features or mis-modelling of detector gains, among others [54].

where $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^T \rangle$ is the noise covariance matrix. In most cases, this equation cannot be solved exactly. In practice, the matrix $(\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})$ on the left-hand side (l.h.s.) is often impractical to compute, store, or invert explicitly. Firstly, the final matrix is very large: $(N_{\text{comps}} N_{\text{pixels}} \times N_{\text{comps}} N_{\text{pixels}})$, where N_{pixels} is the number of pixels in the map, and $N_{\text{comps}} = 3$ represents the I, Q, U Stokes components. Secondly, the noise covariance \mathbf{N} is several orders of magnitude larger, each of its dimensions scaling with the number of detectors and time samples, and exhibiting non-trivial correlations between different detectors and timesteps. To deal with this, optimal mapmaking techniques, like maximum-likelihood (ML) mapmakers, solve the mapmaking equation as a linear system of the form $\mathbf{A}\mathbf{x} = \mathbf{b}$ using iterative methods such as the Preconditioned Conjugate Gradient to invert \mathbf{A} without explicitly computing or storing it.

This process would be significantly simpler in the presence of a diagonal noise covariance \mathbf{N} , corresponding to the case of white noise and uncorrelated detectors. In this case, the l.h.s. matrix is also diagonal across pixels and can be expressed as N_{pixels} “binned” matrices of shape $N_{\text{comps}} \times N_{\text{comps}}$ (or equivalently one $N_{\text{comps}} N_{\text{pixels}} \times N_{\text{comps}} N_{\text{pixels}}$ matrix with zero correlation between pixels), which can be easily inverted. The right-hand side (r.h.s.) $\mathbf{P}^T \mathbf{N}^{-1} \mathbf{d}$ is also projected to map space, being a vector with shape $1 \times N_{\text{comps}} N_{\text{pixels}}$.

Although modulating and demodulating the signal (e.g., using a continuously rotating HWP) should whiten the noise, this idealised situation is never exactly realised. We may nevertheless follow this approach and correct for the impact of the associated approximations a posteriori. In this case, it is common to work with filtered timestreams, in order to deal with systematics and inverse-variance-weight the data. In other words, we apply a set of filters \mathbf{F} to the data, and project, “bin”, or “coadd” the filtered data $\mathbf{d}' \equiv \mathbf{F}\mathbf{d}$ before solving the mapmaking equation under the assumption of a diagonal noise covariance. The maps resulting from this procedure (commonly called “filter+bin” mapmaking), and their directly estimated power spectra, will be biased by an unknown multiplicative transfer function (TF). This TF can then be estimated and corrected via simulations, as described in sections 2.2.1 and 2.2.2. Our goal is to use a relatively fast method that works on the principle of the TF and some approximations.

2.2 From maps to spectra

Our objective is to compress the information from maps into efficient summary statistics, such as power spectra C_ℓ [55–59]. The pseudo- C_ℓ algorithm (often called the “MASTER” algorithm [60]) is a relatively economic procedure to achieve this goal. We summarize the method here.

We have a full-sky map $\mathbf{m}(\hat{\mathbf{n}})$ and its spherical harmonic coefficients $m_{\ell m}^\alpha$, where α labels the different polarisation channels (E and B). An unbiased and optimal estimator for the power spectra of this map is

$$\tilde{C}_\ell^{\alpha\beta} \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} m_{\ell m}^\alpha m_{\ell m}^{\beta*}. \quad (2.3)$$

In order to increase the signal-to-noise of each power spectrum measurement, it is also

common to bin the estimated C_ℓ into “band powers”:

$$\tilde{C}_b^{\alpha\beta} \equiv \sum_{\ell \in b} v_b^\ell \tilde{C}_\ell^{\alpha\beta}, \quad (2.4)$$

where v_b^ℓ are the band power weights. Here we use simple top-hat band powers, with $v_b^\ell = \Theta(\ell \in b)/N_b$, where N_b is the number of multipoles in band power b , Θ is a boxcar function, and $\ell \in b$ simply notes the multipole ℓ belonging to bin b . Note that we distinguish between the true underlying power spectrum $C_\ell^{\alpha\beta}$, the pseudo- C_ℓ or “mode-coupled” spectrum $\tilde{C}_\ell^{\alpha\beta}$, and the binned version of the latter, $\tilde{C}_b^{\alpha\beta}$.

In the actual data, we have the more complicated case of a masked map $\tilde{\mathbf{m}}(\hat{\mathbf{n}}) \equiv w(\hat{\mathbf{n}}) \mathbf{m}(\hat{\mathbf{n}})$, where $w(\hat{\mathbf{n}})$ is a sky mask, which normally will be apodised. The harmonic coefficients of $\tilde{\mathbf{m}}$ are then a convolution of the true harmonic coefficients $m_{\ell m}^\alpha$ and those of the mask $w_{\ell m}$. Likewise, applying the estimator above to this masked map leads to an estimated power spectrum that is related to that of the true underlying map via a convolution of the form:

$$\tilde{C}_\ell^{\alpha\beta} = \sum_{\alpha' \beta' \ell'} W_{\alpha' \beta' \ell'}^{\alpha \beta \ell} C_{\ell'}^{\alpha' \beta'}, \quad (2.5)$$

where $C_{\ell'}^{\alpha' \beta'}$ is the spectrum of the true map and $W_{\alpha' \beta' \ell'}^{\alpha \beta \ell}$ is the “mode-coupling matrix” (MCM). In the case of a simple local operation,⁴ such as masking, the MCM can be calculated analytically as a function of the power spectrum of the mask alone. This forms the basis of the pseudo- C_ℓ algorithm. The “mode-decoupled” power spectrum, $\hat{C}_b^{\alpha\beta}$, can be computed by inverting the MCM after binning by the band powers v_b^ℓ :

$$\hat{C}_b^{\alpha\beta} \equiv \sum_{\alpha' \beta' b'} (\mathcal{W}^{-1})_{\alpha' \beta' b'}^{\alpha \beta b} \tilde{C}_{b'}^{\alpha' \beta'}, \quad (2.6)$$

where the binned MCM is

$$\mathcal{W}_{\alpha' \beta' b'}^{\alpha \beta b} \equiv \sum_{\ell \in b} \sum_{\ell' \in b'} v_b^\ell W_{\alpha' \beta' \ell'}^{\alpha \beta \ell}. \quad (2.7)$$

Thus, the relation between the mode-decoupled spectrum and the true underlying spectrum is:

$$\hat{C}_b^{\alpha\beta} \equiv \sum_{\alpha' \beta' \ell'} B_{\alpha' \beta' \ell'}^{\alpha \beta b} C_{\ell'}^{\alpha' \beta'}, \quad (2.8)$$

where the “band power window functions” (BPWFs) [61] are given by:

$$B_{\alpha' \beta' \ell'}^{\alpha \beta b} \equiv \sum_{\alpha'' \beta'' b''} (\mathcal{W}^{-1})_{\alpha'' \beta'' b''}^{\alpha \beta b} \sum_{\ell'' \in b''} v_{b''}^{\ell''} W_{\alpha' \beta' \ell'}^{\alpha'' \beta'' \ell''}. \quad (2.9)$$

Finally, the case of a filtered and masked map is given by

$$\tilde{\mathbf{m}}(\hat{\mathbf{n}}) \equiv w(\hat{\mathbf{n}}) [\mathbf{R} \mathbf{m}](\hat{\mathbf{n}}), \quad (2.10)$$

where \mathbf{R} is the filtering operator, often called the “observation matrix” [50], relating the observed and true skies. In our case, \mathbf{R} results from the projection and binning of all

⁴This means that every coordinate in the sky (e.g., a pixel) depends on the local value of a mask and of an unfiltered map in that same coordinate.

timestream filtering operations onto the sky. This observation matrix will induce additional effects, suppressing and mixing polarisation modes in a non-trivial way beyond the effect of masking. As before, the power spectrum of $\tilde{\mathbf{m}}$ is related to the true spectrum via a mode-coupling matrix that we will label $X_{\alpha'\beta'\ell'}^{\alpha\beta\ell}$ to distinguish it from the mask-only MCM $W_{\alpha'\beta'\ell'}^{\alpha\beta\ell}$. Unlike in the simple case of only masking, the MCM for a general filter+mask operator cannot be calculated analytically. We may, however, calculate the power spectrum of $\tilde{\mathbf{m}}$ by ignoring the presence of \mathbf{R} and considering only the impact of the mask. The resulting estimator will still relate to the true power spectrum through a set of BPWFs, as in equation (2.8), although the BPWFs are given by

$$B_{\alpha'\beta'\ell'}^{\alpha\beta b} \equiv \sum_{\alpha''\beta''b''} (\mathcal{X}^{-1})_{\alpha''\beta''b''}^{\alpha\beta b} \sum_{\ell'' \in b''} v_{b''}^{\ell''} X_{\alpha'\beta'\ell'}^{\alpha''\beta''\ell''}, \quad (2.11)$$

as opposed to equation (2.9). Here \mathcal{X} is the binned version of X , in analogy to \mathcal{W} in equation (2.7). Since we cannot calculate X analytically, we must resort to a simulation-based approach in order to estimate the BPWFs. The next two sections present two different approaches to achieve this.

2.2.1 Mode coupling per- ℓ

Following ref. [50], the BPWFs can be reconstructed from the response of each band power to a signal map exciting a single multipole ℓ . See also ref. [62]. In more detail:

1. We generate a scalar full-sky simulation from a unit delta-like power spectrum at multipole ℓ_* (i.e. $C_{\ell_*} = \delta_{\ell\ell_*}$).
2. By assigning the harmonic coefficients of the above delta-like scalar simulation to either the E - or the B -mode component of a polarised map with only zeros otherwise, we effectively generate pure- E and pure- B simulations.
3. We “observe” the resulting maps, applying to them the same processing and filters used on the real data, and masking them.
4. We estimate the band powers of the resulting maps following the steps described in the previous section. Specifically, let $C_{(\ell_*)b}^{(\alpha_*\beta_*),\alpha\beta}$ be the b -th band power of the cross-spectrum between the α and β polarisation channels resulting from cross-correlating the pure- α_* and pure- β_* observed δ -like maps produced to excite the multipole ℓ_* .
5. From equation (2.8), it is clear that $C_{(\ell_*)b}^{(\alpha_*\beta_*),\alpha\beta}$ is an estimator for the BPWF $B_{\alpha_*\beta_*\ell_*}^{\alpha\beta b}$.

We repeat this procedure over simulations and average over them to increase the accuracy of the estimated BPWFs. In order to reduce the number of simulations needed to achieve a given precision in the estimated BPWFs, we generate them from distributions with a smaller 4-point function rather than a Gaussian (since the BPWFs are estimated from power spectra). Specifically, we generate them from a Bernoulli distribution (i.e. the real and imaginary elements of the resulting harmonic coefficients for $\ell = \ell_*$ are $\pm 1/\sqrt{2}$ with equal probability), which achieves the lowest possible kurtosis (and hence the lowest variance for any quantity that involves quadratic combinations of these maps) for maps with independent pixel values [63].

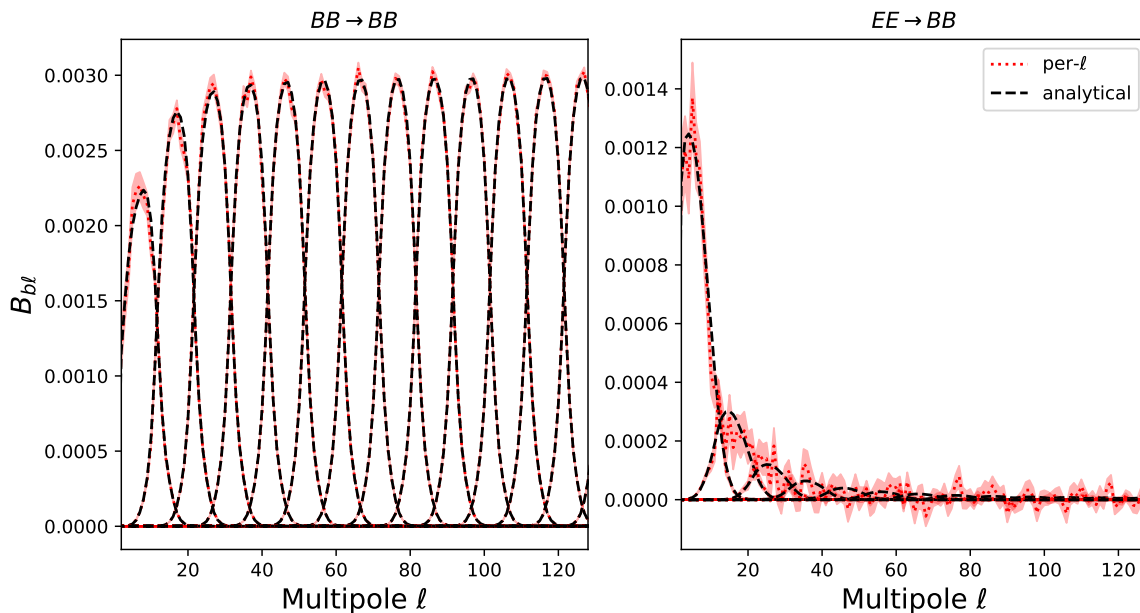


Figure 1. Comparison of the reconstructed band power window functions using the per- ℓ simulation-based method described in section 2.2.1. The dashed black line shows the analytical NaMaster-calculated BPWF for $BB \rightarrow BB$ (left hand side panel) and $EE \rightarrow BB$ (right hand side panel) for a partial sky mask at $N_{\text{side}} = 64$. The dotted red line shows the simulation-based BPWFs averaged across 100 simulations per ℓ . The red shaded area represents the standard deviation across the 100 simulations divided by $\sqrt{100}$. For this example, we use the mask constructed from the hits map in figure 2. The other BPWFs look qualitatively similar. We remind the reader this is an example of a test case with no time-domain filtering.

We can partially validate this method in a simple way, by considering the limit in which no filtering operations are applied to the data beyond masking. In this case, the MCM and the BPWFs can be estimated analytically using the standard pseudo- C_ℓ algorithm via NaMaster [64]. We can then compare this analytical solution with the estimate from δ -like simulations. Figure 1 shows the $BB \rightarrow BB$ and $EE \rightarrow BB$ BPWFs estimated through these two methods, with the analytical solution as the dashed black line, and the simulation-based estimate as the dotted red line, using 100 simulations per ℓ . Although the simulation-based estimate is intrinsically noisy, the differences between them are consistent within the sample variance due to the limited number of simulations.

The main drawback of this method is the large number of simulations needed to reconstruct the BPWF. As we will show in section 4.1, an accurate estimate of the BPWFs, particularly accounting for E/B leakage, may require $\mathcal{O}(10^5)$ simulations, on which expensive filtering operations must be carried out, in order to reconstruct the first ~ 200 multipoles for SO-SAT. This is almost impossible to achieve when doing end-to-end simulated filtering.⁵ This process

⁵As an illustration, it takes approximately 20 minutes to filter and bin one map as described in section 3.3 with one node of the PERLMUTTER computer at NERSC, using 128 cores. To reconstruct the first 200 multipoles, we would need to repeat this for $191 \times 4 \times 1000 = 764,000$ maps, assuming we do 1000 realizations per- ℓ . Often, we will need even more realizations to constrain the leakage terms. In this example, this would take 10,611 node-days.

may be accelerated through the use of observation matrices (e.g., as done by the BICEP collaboration [50]), as long as constructing this matrix is feasible given available computing resources. This may pose challenges for experiments covering large sky areas, such as SO-SAT, or reaching few-arcminute resolutions. This motivates exploring alternative approximate but less computationally demanding methods.

2.2.2 Using a transfer function

General filtering operations can lead to both mode mixing and the statistical coupling of E - and B -modes (E/B leakage). However, if we assume that mode coupling contributions are dominated by the sky mask, while mode suppression (no mixing) and leakage are driven by filtering, the number of simulations required to estimate the BPWFs is significantly reduced. None of the filtering operations studied here, at least in their idealised implementations, would lead to I-to-P leakage.

Specifically, we will assume that the binned pseudo- C_ℓ s are related to the true C_ℓ s through a combination of a mode-coupling matrix due to the mask, and a TF caused by filtering. Thus, we approximate the true mode-coupling matrix for the pseudo band powers in equation (2.11) as

$$\sum_{\ell \in b} v_b^\ell X_{\alpha' \beta' \ell'}^{\alpha \beta \ell} \approx \sum_{\gamma \delta} T_b^{(\alpha \beta)(\gamma \delta)} \sum_{\ell \in b} v_b^\ell W_{\alpha' \beta' \ell'}^{\gamma \delta \ell} \quad (2.12)$$

where $W_{\alpha' \beta' \ell'}^{\gamma \delta \ell}$ is the mode-coupling matrix due solely to the mask (and estimated quickly with **NaMaster**), and $T_b^{(\alpha \beta)(\gamma \delta)}$ is the TF. This equation implies that the masking happens first, while the filtering happens after. Of course, in a real survey, this happens in the opposite order. This simplification is done to estimate a TF depending on a reduced number of band power bins rather than a much larger number of individual multipoles. In section 4.1, we validate this approximation.

The main assumption here (that coupling between different harmonic modes is driven by the sky mask, and not by filtering) is not entirely without basis. Consider the toy scenario of a one-dimensional sky mapped by a continuously scanning telescope. Filtering the resulting timestream is, in the most common cases (e.g. Fourier-space filters, polynomial filters), a translationally-invariant linear transformation. As such, it affects each Fourier mode of the timestream (and of the 1-dimensional sky) independently, without mixing them. This is similar to beam smoothing, although most filters instead affect the largest scales. This analogy breaks down for the true two-dimensional sky, especially in regions of poor cross-linking, and when filtering occurs over finite time periods (e.g. telescope subscans), breaking translational invariance. Nevertheless, it illustrates the reasons for the good performance of the approximate transfer function approach we will demonstrate in section 4.

In the case of scalar fields, the TF has a very simple interpretation: it quantifies the level to which filtering suppresses the power of different angular scales and can be estimated from the ratio of the binned power spectra of filtered and unfiltered simulations. In the case of spin-2 polarised maps, we must also take into account the leakage between polarisation channels caused by filtering. The transfer function is a 4×4 matrix for each band power b , with estimates computed as follows:

1. We generate full-sky Gaussian realisations of a known input power spectrum, which can be anything in principle. In our case, we use a power-law (PL) spectrum with index -2 , which we denote C_ℓ^{PL} (specifically, $C_\ell^{\text{PL}} \equiv (\ell + 10)^{-2}$). Having generated a set of harmonic coefficients from this power spectrum, we create “pure- E ” and “pure- B ” spin-2 simulations by assigning these coefficients to either the E - or the B -mode component of a polarised full-sky field, respectively.
2. We filter and mask these simulations as described in the previous section for the δ -like simulations.
3. We then compute the angular power spectra of the filtered and masked simulations, and those of the same simulations with only the mask applied. For clarity, $\tilde{C}_b^{\text{PLF},(\alpha\beta)(\gamma\delta)}$ denotes the binned pseudo- C_ℓ estimate of the pure- α and pure- β filtered-and-masked simulations for the pair of polarisation channels (γ, δ) . Likewise, $\tilde{C}_b^{\text{PL},(\alpha\beta)(\gamma\delta)}$ denotes the same binned pseudo- C_ℓ for the unfiltered (but masked) simulations, which can be calculated analytically with **NaMaster**.
4. We average both sets of power spectra over a number of simulations (100 in our case). Note that $\tilde{C}_b^{\text{PL},(\alpha\beta)(\gamma\delta)}$ does not need to be calculated empirically averaging over realisations, but instead can be alternatively calculated analytically through multiplying the input spectrum by the mask MCM, as illustrated by equation (2.5).

As discussed above, the power spectra for the filtered and unfiltered simulations are related to one another by the TF:

$$\tilde{C}_b^{\text{PLF},(\alpha\beta)(\gamma\delta)} = \sum_{\alpha'\beta'} T_b^{(\alpha\beta)(\alpha'\beta')} \tilde{C}_b^{\text{PL},(\alpha'\beta')(\gamma\delta)}. \quad (2.13)$$

For a fixed combination of (α, β) , the equation above may be expressed as a matrix-vector multiplication:

$$\mathbf{b}_b^{(\alpha\beta)} = \mathbf{C}_b \mathbf{t}_b^{(\alpha\beta)}, \quad (2.14)$$

where the vectors \mathbf{b} and \mathbf{t} consist of the power spectra of filtered simulations and the TF respectively,

$$\mathbf{b}^{(\alpha\beta)} = \left(\tilde{C}_b^{\text{PLF},(\alpha\beta)(EE)}, \tilde{C}_b^{\text{PLF},(\alpha\beta)(EB)}, \tilde{C}_b^{\text{PLF},(\alpha\beta)(BE)}, \tilde{C}_b^{\text{PLF},(\alpha\beta)(BB)} \right). \quad (2.15)$$

$$\mathbf{t}^{(\alpha\beta)} = \left(T_b^{(\alpha\beta)(EE)}, T_b^{(\alpha\beta)(EB)}, T_b^{(\alpha\beta)(BE)}, T_b^{(\alpha\beta)(BB)} \right), \quad (2.16)$$

and the matrix \mathbf{C}_b is the collection of unfiltered binned pseudo- C_ℓ s

$$(\mathbf{C}_b)_{(\gamma\delta)(\alpha\beta)} = \tilde{C}_b^{\text{PL},(\alpha\beta)(\gamma\delta)}. \quad (2.17)$$

Since the matrix \mathbf{C} is not necessarily invertible, we estimate the TF by solving equation (2.14) via least-squares fitting

$$\hat{\mathbf{t}}_b^{(\alpha\beta)} = \left[\mathbf{C}_b^T \mathbf{C}_b \right]^{-1} \mathbf{C}_b^T \cdot \mathbf{b}_b^{(\alpha\beta)}. \quad (2.18)$$

With the TF estimated as we just described, we can construct our model for the mode-coupling matrix of the band powers by combining it with the mask-only mode-coupling

matrix (equation (2.12)). The method then proceeds as it does for the standard pseudo- C_ℓ approach: we bin and invert this mode-coupling matrix, and use it to estimate the mode-decoupled band powers and the associated BPWFs. Our implementation of this method⁶ includes all the cross-spectra of the IEB fields. As stated, this paper will only focus on polarisation and ignore intensity.

3 Simulations

In this section, we describe the simulations used to estimate the transfer function (TF), as well as validate the methods described. We describe the input maps we use, the scanning strategy of our simulated SO-SAT deep survey, and how we produce filter+bin maps.

3.1 Input simulated maps

All of our simulated maps are in HEALPIX pixels [65], produced at single float precision. For a comparison between the numerical per- ℓ and TF methods, shown in section 4.1, we use maps at resolution $N_{\text{side}} = 128$. The limited resolution is necessary in order to compute and store the $3N_{\text{pixels}} \times 3N_{\text{pixels}}$ observation matrix used to incorporate the effects of filtering (see section 3.3.2). Handling this matrix rapidly becomes challenging for resolutions $N_{\text{side}} \geq 256$. For the TF method results at full resolution and multipole range, shown in section 4.2, we use maps at resolution $N_{\text{side}} = 512$. All the angular power spectra in this work are calculated with **NaMaster**, which corrects for the mode coupling arising from partial sky observation. We do not use E/B -mode purification for simplicity. Incorporating the standard B -mode purification techniques that correct the E -to- B leakage caused by the sky mask [66] is trivial to implement in this formalism (one simply replaces all fields and power spectra by their purified versions). Filtering leads to additional leakage that, ideally, should also be corrected at the map level. Methods for this exist, exploiting observation matrices to construct the covariance of the filtered data [50, 67]. This is an area of active research in SO, where the larger sky area covered makes deploying these techniques challenging. The masks are constructed from normalised hits maps and with a C1 apodisation [68] with a 5° length.⁷

For the per- ℓ method, we generate 4 maps (one per each polarisation EE, EB, BE, BB) of correlated pure modes that excite one ℓ at a time, per each realisation, per each ℓ . We filter them using the $N_{\text{side}} = 128$ observation matrix of our observation schedule, which are described in section 3.3.2. We create 6144 realisations per ℓ .

For the TF method, all the maps we create have the corresponding pixel window function and a Gaussian beam smoothing with full width at half maximum (FWHM) of $30'$ for the f090 band and $17'$ for the f150 band. The power spectrum of the transfer function simulations, in general, should have the shape of the power spectrum of the unfiltered real sky signal. We estimate the TF from 100×2 input maps generated with a power-law theory spectrum $C_\ell = (\ell + 10)^{-2}$, which is similar to the angular power spectrum of galactic foregrounds (which are non-Gaussian also). Each realisation has pure- E modes with $a_{\ell m}^B = 0$ or pure- B modes with $a_{\ell m}^E = 0$. We also produce validation simulations. We use the estimated TF

⁶<https://github.com/simonsobs/SOOPERCOOL> v0.0.

⁷See https://namaster.readthedocs.io/en/latest/api/pymaster.utils.html#pymaster.utils.mask_apodization for details.

to reconstruct the polarisation spectra of these validations and check the consistency with input theory spectra. We simulate 100 CMB maps with a fiducial spectrum consistent with the *Planck* 2018 best-fit parameters [69] and lensing as the only source of B -modes (therefore with tensor-to-scalar ratio $r = 0$). To assess the limitations of the TF method, we also produce another set of 100 validation simulations with a steeper power-law with theory spectrum $C_\ell = (\ell + 10)^{-6}$.

3.1.1 Foreground simulations

In sections 5.3 and 6, we consider foregrounds and component separation. For that, we use the *Planck* 353 GHz channel as ancillary data. We use the same CMB realisation with lensing-only B -modes and smoothed to a $5'$ beam. We include two noise maps per realisation, corresponding to the detector A and B focal plane splits using the end-to-end NPIPE noise simulations [70].⁸ We observe and filter 353 GHz CMB+noise maps. By this, we mean transforming the map into a timestream, filtering it with the same SO-SAT filters, and then binning the map as described in section 2.1. We construct 100 *Planck*-like 353 GHz maps.

We use full-sky templates for synchrotron and thermal dust from PYSM [71]. We use the `s1` model for synchrotron and the `d1` model for dust, which include spatially anisotropic foreground amplitudes and, in addition, moderate spatial variation in the frequency spectral energy distributions (SED). Real observations inform these models and therefore are available as single realizations. For covariance estimation purposes, we also create 100 realisations of polarised dust emission as a random Gaussian field, with a power-law spectrum at 353 GHz given by

$$\mathcal{D}_\ell^{BB} = A_d^{BB} (\ell/80)^{\alpha_d^{BB}}, \quad (3.1)$$

where $A_d^{BB} = 35 \mu\text{K}^2$ and $\alpha_d^{BB} = -0.16$ [31].⁹ The dust maps at 353 GHz are then scaled to other frequencies using a modified black body (MBB) SED with spatially constant spectral index $\beta = 1.54$ and temperature $T_d = 20 \text{ K}$.

3.2 Simulated dataset

We simulate an approximate 1-year SO-SAT deep survey. To reduce the computational cost, we simulate the observation of the South Galactic patch performed by a single MF SAT on the first day of each month during a year. We assume no observations are made in the time block from the 15th of January to the 15th of March, since this accounts for possible bad weather conditions and downtime for hardware upgrades. Thus, in practice, the observation scanning strategy simulated here includes 10 days of observations. We chose this simplified schedule, containing a fraction of the real schedule, to reduce the computational requirements of the simulation while evenly sampling the temporal variations in the scanning pattern due to sun avoidance across the year.

As an example scanning strategy, we observe at three different fixed elevations of $50^\circ, 55^\circ, 60^\circ$. The South patch is prioritised, so it is observed from the moment it rises above the limiting elevation. The real SAT scan pattern will also include observations of the North

⁸Available at NERSC at `global/cfs/cdirs/cmb/data/planck2020/npipe`.

⁹While this reference used $A_d^{BB} = 28 \mu\text{K}^2$, $35 \mu\text{K}^2$ is a better fit to the mask used in this work.




Figure 2. Normalised hits map, in Equatorial coordinates, for the full schedule we simulate in TOAST, which correspond to 10 observing days throughout a year totalling 139 hours. In this exercise, we only observe the South Galactic patch, while the real SO-SAT deep survey will also cover parts of the North Galactic patch.

Galactic patch when the South patch is not visible, so our simulated strategy is not optimised while being fairly representative of what our real observations will look like. An individual scan is 50 minutes long on average, covering a typical azimuth range of 40° . The azimuth scan speed is 1 deg s^{-1} , while the scan acceleration is 1 deg s^{-2} . We warn the reader this describes an example scanning strategy and may not represent how SO actually scans in real observations.

In total, our schedule has 167 observations for a total integration time of 139 hours. Each observation is almost an hour long and includes the 7 detector arrays, known as wafers, covering the focal plane of one SAT, for a total of ~ 6000 detectors per each f090/f150 band (however we use only one-eighth of these as described below). The normalised hits map of our full schedule is shown in figure 2, scanning the South Galactic patch. Our schedule is run in 8 sub-schedules that produce one filter+bin map each. Sub-schedules are used to balance between the flexibility to create splits of the data and computation time. These maps are linearly co-added.¹⁰ In particular, 5 sub-schedules correspond to observations when the South patch is rising through the sky (equivalent to 102 observations and 93.5 hours) and 3 sub-schedules when the South patch is setting in the sky (equivalent to 65 observations and 45.5 hours). This separation into sub-schedules is useful when defining splits for simulating a null test or trying to estimate cross-spectra without noise bias.

3.3 Using TOAST

We simulate the observation of an input map by scanning it to time-ordered data (TOD) following a schedule like the one detailed above. Then observations are processed and mapped into filter+bin maps. In practice, this procedure is performed with the Time-

¹⁰Since the filtering and mapmaking operations are linear, coadding individual sub-schedule filter+bin maps yields the same map as applying the filter+bin mapmaker to the total schedule.

Ordered Astrophysics Scalable Tools (TOAST)¹¹ software, as well as the SOTODLIB¹² library, which provides the necessary scripting and configuration to simulate specific detectors and telescopes for SO.

The TOD is simulated for every detector for a given telescope and detector wafer. The input map is scanned into a TOD, assigning, to each time sample, the value of the pixel that the telescope was pointing at (i.e. no interpolation is performed to avoid introducing an extra window function). Also, the TOD is modulated by the spinning of an idealised cryogenic continuously rotating HWP. The HWP is rotating near 2 Hz, and therefore the polarised signal is modulated at 8 Hz.

We then implement a series of analysis steps that would simulate the filtering in a CMB survey. To test the pipeline, we filter the simulated data more aggressively than we expect to do on the real SO data. The observations are processed in the following order:

- A HWP synchronous signal filter is used to subtract the HWP synchronous signal introduced in the modulation of the TOD. This is a sines/cosines template fitted to data as a function of HWP angle and subtracted. The HWP filter is applied to one detector and one hour-long observation at a time.
- The TOD is demodulated by applying a bandpass filter centred at 8 Hz and with width 4 Hz and multiplying in Fourier space by the phasor $\exp(-4i\chi(t))$, where $\chi(t)$ is the HWP angle as a function of time [38, 42]. The real part of the result is the demodulated Q signal, while the imaginary part is the demodulated U signal. Both are then low-pass filtered at 2 Hz, since the bandpass filter will have already reduced all signals at higher frequencies. See [36, 38, 40, 72, 73] for details.
- A ground filter is used to subtract the horizontal-coordinate-fixed contamination from the ground. This is a Legendre polynomial of order 10 fitted to data as a function of azimuth and subtracted. The ground filter is applied to one detector and one hour-long observation at a time. No binning is needed: the polynomial templates are sampled at the boresight azimuth and projected out from the detector data.
- A high-pass filter is used to suppress the impact of low-frequency contamination such as the atmosphere and detector thermal drifts, among others. This is a Legendre polynomial of order 3 fitted to the data as a function of time and subtracted. This and the ground filter are applied simultaneously since they are templates deprojected in one matrix operation. The polynomial templates are projected out separately from each sub-scan (sweep across the field).

In appendix A, we illustrate the effect of individual filters. Finally, the TOD are binned into HEALPIX $N_{\text{side}} = 512$ maps according to their pointing. We make some approximations to ease the computing cost of running these simulations for this work. We use a 40 Hz sampling rate for scanning (real observations scan at 200 Hz) and we also “thin” the focal plane by a factor of 8, i.e., only one out of every eight detectors is simulated. In appendix B we

¹¹<https://github.com/hpc4cmb/toast> v3.0.0a25.

¹²<https://github.com/simonsobs/sotodlib> v0.5.0.

discuss this in detail, showing how the impact is minimal while reducing the computation time by a factor of 8. The combined effect of the three data volume reductions (one day every month; 40 Hz; 1/8th of detectors) is to reduce the computational requirements of the simulations by a factor of 1200 per frequency band per SAT. With our early data analysis, we have estimated that it takes $\sim 3\text{M}$ core-hours to filter the TF simulations (300 in total, 100 realisations and pure- T , pure- E and pure- B modes for each) required for one year of data with one SAT and one band. This is avoiding any data thinning, which can lead to a significant speed-up without compromising the quality of the estimated transfer function. It is therefore entirely feasible to deploy the methods described here with the high-performance computing resources available to the collaboration.

3.3.1 Time-domain noise simulations

While the section above describes an example of the filtering and processing steps applied to TOD, we also perform time-domain simulations of sources of noise and systematics. These sources are

- **Atmospheric noise:** we simulate an unpolarised atmosphere [74, 75] using an implementation of the model from refs. [33, 76]. This is based on a 3D structure with Kolmogorov turbulence moving through the telescope beam, with a distribution of temperature, water vapour, and wind speed. This model was fitted to match POLARBEAR observations at the Cerro Toco site. Within this model, and ignoring I-to-P leakage, this atmospheric emission does not impact the polarised signal, so the $1/f$ component of polarised noise is mostly controlled by the scanning speed and the choice of filters applied to the TOD. As such, this component has no effect on our E, B fields. While the simulated unpolarised intensity maps are not used in this work, they will be useful for future analysis work.
- **Gaussian-distributed, uncorrelated white noise:** the detector white noise is based on the forecasted levels for an SAT given in the SO science paper [2]. In the TOAST instrument model, individual detectors are modelled such that a final 5-year deep survey map has noise levels of 2.6 and $3.3\ \mu\text{K-arcmin}$ for f090 and f150, respectively. This corresponds to the baseline white map noise level in ref. [2].
- **Scan Synchronous Signal (SSS):** we use an ad hoc model where the telescope scans a fixed map in horizontal coordinates to emulate both optical pickup and signals internal to the instrument (e.g. electromagnetic interference from the azimuthal drive). All detectors observe the same map, but the map is drawn separately for each simulated observation. The map itself is Gaussian white noise, convolved with a 1° beam, with amplitude inversely proportional to the observing elevation and scaled to 1 mK RMS at 45° observing elevation. The elevation dependence is added to reflect the fact that far-side lobe pickup intensifies at lower observing elevations. The 1° smoothing is chosen to avoid sharp features in the simulated signal. Finally, the overall amplitude of the SSS is chosen to overwhelm CMB signal for the sky modes that are degenerate with it. These choices are anecdotally supported by POLARBEAR data but are not based on rigorous analysis. The model will be thoroughly calibrated against real data in the future.

- **Half-Wave Plate Synchronous Signal (HWPSS):** this is added to the TOD from a template constructed in ref. [77].

The noise simulations are run using the same schedule as described in section 3.2, and noise levels are scaled to match the SO-SAT expected noise level after a 1-year deep survey (2 telescopes and 6000 detectors on each for 4170 hours) [2].

Our overall noise model lies between the optimistic ($\ell_{\text{knee}} = 25$) and pessimistic ($\ell_{\text{knee}} = 50$) SO-SAT performance levels [2]. To evaluate this, we take the pure noise simulations, deconvolve the filtering with the corresponding transfer function, and fit a noise curve model

$$N_\ell = N_{\text{white}} [1 + (\ell/\ell_{\text{knee}})^{\alpha_{\text{knee}}}] \quad (3.2)$$

to the mean of 10 only-noise f090 realisations, finding $\ell_{\text{knee}} = 44 \pm 4$ and $\alpha_{\text{knee}} = -3.22 \pm 0.15$. We generate 100 realisations of noise observations for the f090 frequency channel and the same for the f150 channel.

3.3.2 Observation matrix

As described in section 2.2.1, the per- ℓ estimation method requires thousands of TOAST simulations. The observation matrix \mathbf{R} is such that a biased (and masked) observed map $\tilde{\mathbf{m}}$ is expressed as in equation (2.10). \mathbf{R} has a shape $3N_{\text{pixels}}$ by $3N_{\text{pixels}}$ (to capture the three Stokes components I , Q , and U). \mathbf{R} can be calculated explicitly [78] with

$$\mathbf{R} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{P}, \quad (3.3)$$

where we recognise the elements familiar from mapmaking, discussed in section 2.1, and \mathbf{Z} is the time-domain filtering operator, which includes all the operations done to an observation. It is a square matrix that has the dimensions of the time-ordered data and performs linear regression of time domain templates represented as columns of matrix \mathbf{F} :

$$\mathbf{Z} = \mathbf{1} - \mathbf{F} \left(\mathbf{F}^T \mathbf{N}_Z^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^T \mathbf{N}_Z^{-1}. \quad (3.4)$$

In place of a noise covariance matrix, \mathbf{N} , the filtering operator uses a binary diagonal matrix, \mathbf{N}_Z , that can mask out data we do not want affecting the template amplitudes: typically samples drawn near planets, point sources or intense Galactic emission.

The filtering matrix, \mathbf{Z} , operating on a TOD vector, \mathbf{d} , can be understood to perform three separate operations:

1. Perform a generalized least squares fit of \mathbf{d} , in terms of the time domain templates that form the columns of \mathbf{F} . This results in a vector of template amplitudes: $\mathbf{a} = \left(\mathbf{F}^T \mathbf{N}_Z^{-1} \mathbf{F} \right)^{-1} \mathbf{F}^T \mathbf{N}_Z^{-1} \mathbf{d}$.
2. Unroll (weight and co-add) the templates, \mathbf{F} , according to the fitted amplitudes, \mathbf{a} , to produce a model of the TOD: $\mathbf{d}_{\text{model}} = \mathbf{F} \mathbf{a}$.
3. Subtract the TOD model from the input to produce a filtered version of the input TOD: $\mathbf{d}_{\text{filtered}} = \mathbf{d} - \mathbf{d}_{\text{model}}$.

4 Results

In this section, we show the measured band power window functions (BPWFs) using the two methods presented in section 2.2 applied to the filtered maps as observed with the simulated schedule described in section 3.2. These can be used to deconvolve the effect of mode-mixing and reconstruct the power spectra of the CMB. We show the transfer function (TF) method adequately recovers the input power spectrum over the full multipole range of interest. Finally, we discuss limitations of the TF method.

4.1 Comparison of band power window functions

For both the per- ℓ estimation method and the TF method, we filter simulations with the observation matrix, which in practice is multiplying \mathbf{R} with the input map. Since we need many simulations to constrain the BPWFs, especially those associated with polarisation leakage, we simplify our analysis and calculate band power bins up to $\ell_{\max} = 255$ with width $\Delta\ell = 10$ for both methods. For the numerical per- ℓ method, we generate and filter full-sky simulations that excite every individual ℓ between 2 and ℓ_{\max} . The estimation for each multipole is run with 6144 realisations.

For the TF method, we filter our power-law simulations with the observation matrix, so we can guarantee the exact same filtering operations between both methods. We run the estimation of the TF with 100 realisations of a power-law with spectral index -2 and the same binning scheme.

The BPWFs $B_{\alpha\beta b}^{\alpha'\beta'\ell}$, shown as $4 \times 4 = 16$ $B_{b\ell}$ combinations of $\alpha'\beta' \rightarrow \alpha\beta$ polarisations and for band power b and multipole ℓ are shown in figure 3. Most band powers b are plotted together (we limit to $\ell = 30-180$ for visualisation). The numerical per- ℓ results are shown as dashed red lines for the band powers, while the TF results are shown as dotted blue for the same binning scheme.

First, we note the power suppression arising from observing and filtering in our mapmaking. The auto-BPWFs (the diagonal panels in figure 3) show excellent agreement between both methods in the higher- ℓ band powers. This implies that the approximation of equation (2.12) is valid, for example, at the level of $\sim 1\%$ at $\ell \sim 200$. In other words, the additional mixing between different angular scales caused by filtering is subdominant compared to that caused by the mask, given the steepness of the power spectra under study. However, at lower multipoles ($\ell < 100$), where the suppression is larger, there is a clear difference, highlighted by the zoom-in inset in the lower right panel. The TF method assumes that the suppression of modes is the same at all multipoles within a given band power. This assumption breaks down at larger angular scales, resulting in the difference observed in figure 3 at lower multipoles.

The relevant cross-BPWFs (e.g., the $EE \rightarrow BB$ panel in figure 3, which shows the E -mode leakage into B -modes, or $EE \rightarrow EB$, which is even smaller and noisier) are $\sim 50 - 250$ times smaller than the auto-BPWFs. Due to the signal-to-noise ratio, where the noise is dominated by the cosmic variance of a map realisation, this is difficult to constrain with fewer simulations for the numerical per- ℓ method. Hence, we need thousands of simulations and the per- ℓ method is limited by how much we are willing to shrink the cosmic variance. The reconstruction is noisy as can be seen in the figure. This means that the reconstruction of B -modes with the numerical per- ℓ method might be noisy or biased due to contamination

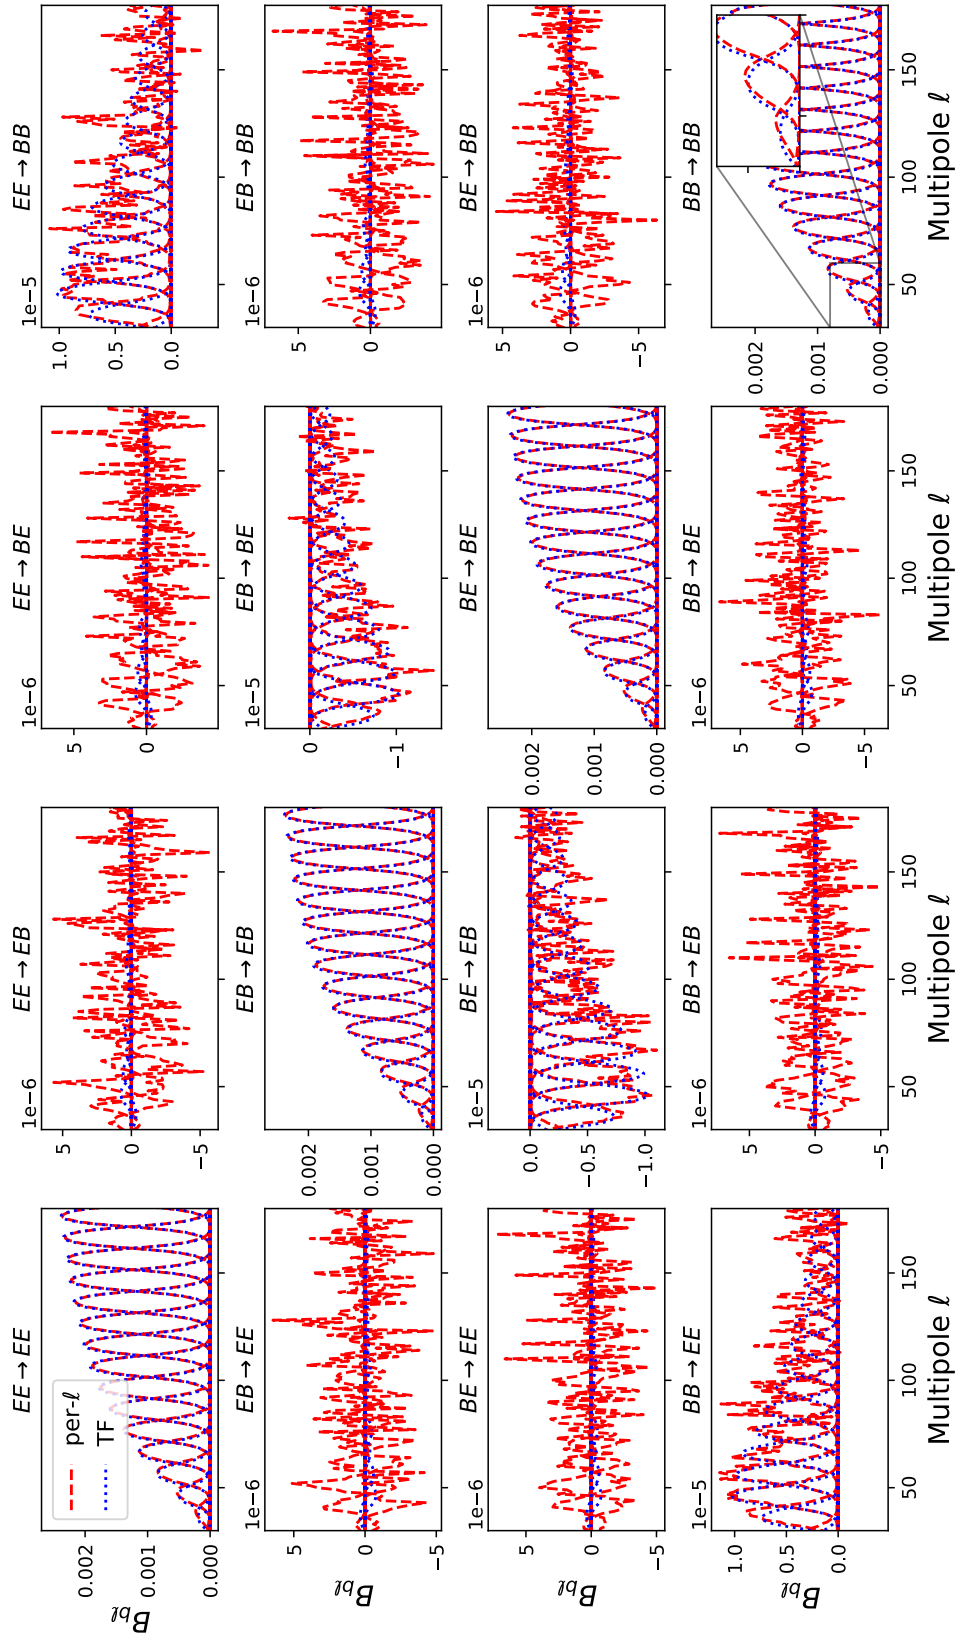


Figure 3. Reconstructed BPWFs for the full schedule and for the binning scheme $\ell = 30\text{--}255$ with size $\Delta\ell = 10$ (in the plot we limit to $\ell = 30\text{--}180$ for clarity). The titles above each subplot refer to the indices $\alpha'\beta' \rightarrow \alpha\beta$ in equations (2.8), (2.9), and (2.11). Red-dashed is the numerical per- ℓ solution estimated with 6144 simulations per ℓ . Blue-dotted is the TF approximation estimated from 100 power-law simulations. In the $BB \rightarrow BB$ panel, we show a zoom-in inset highlighting the difference between the two methods at low- ℓ . This is for the f090 frequency channel.

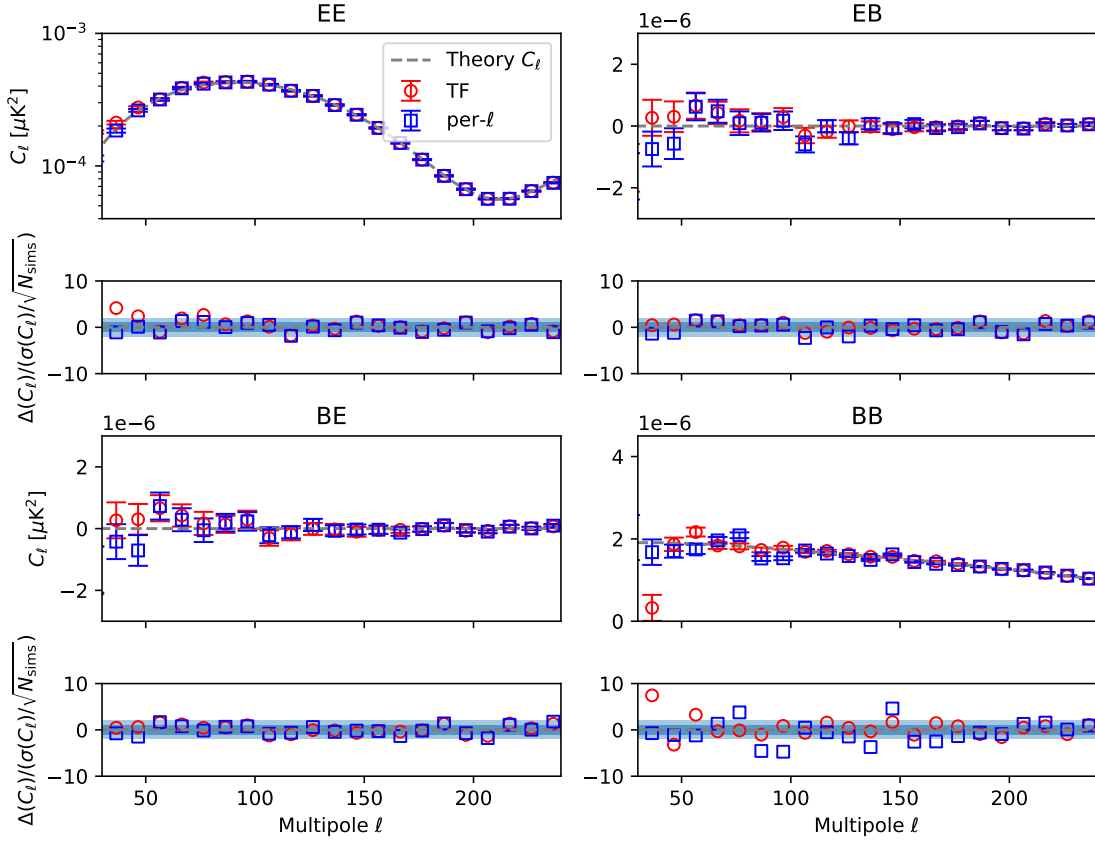


Figure 4. Comparison of the reconstructed polarisation spectra for the f090 channel using CMB simulations with only lensing B -modes. The four panels are in clockwise order the EE , EB , BB , BE spectra. The spectra are the mean over $N_{\text{sims}} = 100$ realisations, while the error bar is the standard deviation divided by $\sqrt{N_{\text{sims}}}$. The theory CMB spectra are shown as the dashed grey lines and have been convolved with the corresponding Gaussian beam. The smaller panel at the bottom of each spectrum shows the residual between the average of the reconstructed and theory spectra, $\Delta(C_\ell)$, divided by $\sigma(C_\ell)/\sqrt{N_{\text{sims}}}$. The shaded areas correspond to ± 1 and $\pm 2\sigma$. The recovered spectra are visibly biased at $\ell < 70$ in this noiseless case.

from E -modes, since in general for the CMB the B -modes are a few orders of magnitude smaller than the E -modes, and the leakage might represent a significant fraction of the signal.

We also simulate pure CMB realisations, generated with the *Planck* best-fit cosmology [69], including B -modes arising only from gravitational lensing ($r = 0$), as described in section 3.1. We generate them at $N_{\text{side}} = 128$ and filter them with the same scheme represented by the observation matrix. The reconstructed polarisation spectra with both methods are shown in figure 4. The four square panels show the EE , EB , BB , and BE spectra in clockwise order. In each, we show the mean C_ℓ across $N_{\text{sims}} = 100$ realisations. The error bar is the error of the mean, i.e. the standard deviation across realisations divided by $\sqrt{N_{\text{sims}}}$, and as such, these σ values are 10 times smaller than σ for a single realisation. The thinner lower panel on each spectrum shows the ratio between spectrum residuals and the associated error. The residual $\Delta(C_\ell)$ is the difference between the measured decoupled power spectrum and the fiducial theory spectrum binned using the BPWFs. The observed CMB maps include

beam smoothing, and so the theory spectra are convolved with a $30'$ Gaussian beam. Since these are pure CMB realizations, the only source of variance is cosmic variance.

The EE spectrum is affected the least by contamination due to leakage, since $B \rightarrow E$ leakage is very small. For example, the first bin in the EE deconvolved spectrum (upper left corner, $\ell = 30\text{--}40$) shows a $\sim 4\sigma$ bias for the TF method, while per- ℓ shows mostly unbiased band powers. This arises because the power spectrum of the TF simulations with power-law index -2 varies largely within the first multipole bin and does not match well the shape of a noiseless CMB power spectrum, leading to a bias in the TF estimate. However, in section 5 we will show that this bias is not important in the context of studying the precision of the 1-year SO SAT deep survey discussed in this paper. Another issue to note is that the cross-BPWFs for the per- ℓ method are noisy, so the leakage estimation is noisy, and if leakage is relevant, then e.g., the BB deconvolved spectrum will be vulnerable to large contamination from $E \rightarrow B$ leakage. This is very noticeable in the BB spectrum on the lower right panel of figure 4, which shows large dispersion by several sigma due to $E \rightarrow B$ leakage. Finally, the EB and BE spectra are the same in the case of the TF method, since the TF is calculated from the same realisation, but different for the per- ℓ method, where different realisations for EB and BE are used. In both cases, the residuals are consistent with zero.

The TF method can reconstruct all deconvolved power spectrum bins between $\ell = 30$ and 250 from biased filter+bin maps through $2 \times \mathcal{O}(100)$ simulations with a bias below 1σ (for a single cosmic-variance-limited CMB simulation).¹³ This is acceptable, especially given that real observations will contain additional noise. By comparison, the numerical per- ℓ method is impractical to run since it requires thousands of simulations per multipole (which is virtually impossible), or constructing an observation matrix every time that the filtering and mapmaking procedure is updated. In section 5, we will illustrate how the biases induced by the approximations made by the TF method are sub-dominant to the noise, and as such not important, through the use of null tests. Hence, we will adopt the TF for the rest of our analysis.

4.2 Full multipole range using the transfer function method

In order to evaluate the TF method across the full SAT multipole range ($\ell = 30\text{--}300$), we run timestream-to-map filtering on $N_{\text{side}} = 512$ simulated input maps. We remind the reader that estimating the TF only requires a modest number of simulations, which can therefore be filtered using brute-force methods. This is in contrast with the per- ℓ method, where the number of simulations is significantly larger, and memory-intensive observation matrices must be used to enable fast observation and filtering. Figure 5 shows the transfer function $T_b^{(\alpha\beta)(\alpha'\beta')}$ for the full observation schedule described in section 3.2 calculated from 100 realisations. Again, we note that this transfer function is estimated using more aggressive filtering than we expect to use in the real SO data analysis. The auto-terms $\alpha\alpha \rightarrow \alpha\alpha$ in the diagonal of the figure show the suppression at larger scales due to filtering, and then plateau at ~ 0.8 . In appendix A, we illustrate why this happens, where the high-pass polynomial filter of order 3 is most responsible.

¹³We note that figure 4 shows the bias on the mean of 100 CMB-only sky realisations, which must be divided by $\sqrt{N_{\text{sims}}} = 10$ to yield the bias on a single realisation.




Figure 5. The transfer function for the f090 channel, showing the 16 $\alpha'\beta' \rightarrow \alpha\beta$ polarisations estimated with equation (2.18), averaged over 100 realisations of power-law simulations. The shaded area (barely thicker than the blue line) represents the ± 1 standard deviation divided by $\sqrt{N_{\text{sims}}} = 10$. This transfer function is estimated using more aggressive filtering than we expect to use in the real SO data analysis.

The TF-deconvolved spectra for the observed CMB and for the full range $\ell = 30\text{--}300$, are shown in figure 6. For each spectrum we show the mean C_ℓ with error on the mean across $N_{\text{sims}} = 100$ realisations. The right-hand side panel shows the residuals divided by the error bar. No systematic bias in the residual is visible, except for the first band power $\ell = 30\text{--}40$. In section 5, we will show that this bias is sub-dominant and therefore not important when noise is added to these pure CMB simulations, through the use of null tests.

4.3 Limitations of the transfer function method

To illustrate the limitations of the TF method, we use the same TF estimated in section 4.2 to reconstruct a different fiducial spectrum with a very different shape, a steep spectral index $\alpha_s = -6$, which has large variation across each of the band powers. This case is not representative of what the true sky is expected to look like, but this exercise is to push the boundaries of the method.

We filter 100 power-law realisations with the same fiducial C_ℓ in the EE and BB spectra, with $\alpha_s = -6$ spectral index and the same $30'$ beam smoothing. Figure 7 shows the reconstructed spectra and the bias in the residual, analogous to figure 6. The entire multipole range in the spectra seems to be biased. The coupling of modes induced by observing/filtering within a band power bin is large enough that the TF is no longer a good approximation.




Figure 6. Deconvolved spectra for the same CMB shown in figure 4 for the full schedule and using the TF method. The fiducial EE and BB spectra are shown in grey dashed lines. We can achieve unbiased power spectra estimates for the range $\ell = 40\text{--}300$ at moderate computational cost. This is for the f090 frequency channel.




Figure 7. Deconvolved spectra for a power-law theory spectrum with a steep index $\alpha_s = -6$. This uses 100 realisations, and the full schedule. We use the same TF shown in figure 5. This figure follows the same convention as figure 6. This is for the f090 frequency channel.

5 Null test

In this section, we assess the relative bias of power spectra reconstructed from simulated data subsets (“splits”) that have undergone significantly different filtering, and hence possess different transfer functions. If the power spectrum reconstruction is unbiased and robust




Figure 8. The TF for $BB \rightarrow BB$ for the case of using rising and setting scan splits. All are normalised by the Full \times Full TF. We also include the full schedule in black pentagons, the same as shown in figure 5 but in a more restricted mask. We note how all combinations are noticeably different from each other. This is estimated from 100 power-law realisations. The error bars are the standard deviation across the realisations divided by $\sqrt{100}$. This is for the f090 frequency channel.

against different filtering prescriptions, the power spectrum corresponding to the difference map should be statistically consistent with noise. First, we take a look at pure signal with only CMB, and then we add realistic noise.

5.1 Definition of splits

In order to maximize the potential difference between transfer functions, we minimize cross-linking¹⁴ in the data splits, partitioning the scans into rising (R) and setting (S), and calculating the different TFs for each map combination: R \times R, S \times S, R \times S, and S \times R where \times represents cross-correlation between two maps. Even further, those TFs will be different from using the full mission, labelled Full and corresponding to R+S. This is illustrated in figure 8, where we show the $BB \rightarrow BB$ TF for R \times R, S \times S, R \times S, and Full \times Full, normalised by Full \times Full.

In figure 9, we show the normalised hits maps for the rising and setting scans at the left and centre, respectively. As a result of the scanning strategy used here, the setting scans observe approximately one-half of the South patch, while the rising scans almost uniformly cover that region. We construct a common hits map using the intersection of both obtained by multiplying the rising and setting hits maps, but keeping only the pixels where in both cases the hits count is greater than the median computed over all pixels. This common hits

¹⁴Cross-linking refers to the technique of scanning sky coordinates from multiple directions to mitigate potential systematics.




Figure 9. Normalised hits maps for splits. Left, normalised hits map for the rising scans. Centre, normalised hits map for the setting scans. Right, the common hits map is constructed from the intersection of the rising and setting hits maps. We build a mask from this hits map to estimate power spectra in section 5. All maps are shown in Equatorial coordinates.

map is shown on the right-hand side of the figure. Finally, we build a mask from this hits map, which we use for power spectra calculation in this section.

5.2 CMB-only test

First, we calculate the residual spectra for 100 CMB simulations with only lensing B -modes. Since we only have signal in our maps, we can calculate all spectra via direct correlation of the maps. We calculate the $R \times R$, $S \times S$, and $R \times S$ spectra ($S \times R$ in this case is the same as $R \times S$). For each type of spectra, we use the corresponding TF calculated from co-adding the pure power-law simulations in the same split combination. The residual spectrum ΔC_b^{XY} for the null test is given by

$$\Delta C_b^{XY} = C_b^{X_R Y_R} + C_b^{X_S Y_S} - C_b^{X_R Y_S} - C_b^{X_S Y_R}, \quad (5.1)$$

where $X, Y \in E, B$ and b indexes band power bins. We calculate this residual for each of the 100 CMB realisations. Then, we define

$$\chi^2(XY) = \sum_b \left[\Delta C_b^{XY} / \sigma(\Delta C_b^{XY}) \right]^2, \quad (5.2)$$

where $\sigma(\Delta C_b^{XY})$ is the standard deviation across the 100 realisations. In this case of pure signal, the statistical uncertainties are dominated by residual cosmic variance caused by differential filtering. We note that this equation ignores uncertainties correlated across different angular scales, effectively assuming a diagonal covariance matrix. We adopt this approximation because the covariance matrix estimated from only 100 simulations is too noisy to invert reliably. The matrix itself is very close to diagonal, with all off-diagonal elements consistent with zero. The sum over b goes over 27 bins in the range $\ell = 30-300$ with width $\Delta\ell = 10$.

If the residuals are consistent with noise, the χ^2 statistic of equation (5.2) should be consistent with a chi-squared distribution with 27 degrees of freedom (DOF), this being the number of band powers in the range $\ell = 30-300$. We can calculate the probability-to-exceed (PTE) of each realisation. The histogram of the 100 PTEs is shown on the left-hand side of figure 10 for the three XY polarisations: EE , EB , and BB . If the residual in this null test is consistent with noise, we would expect a uniform distribution of PTEs between 0 and 1.




Figure 10. Histograms of PTEs for the null tests. This is for the residual spectra $R \times R + S \times S - R \times S - S \times R$, where R and S represent rising and setting scans splits. (Left) Null test for the pure CMB signal. The BB spectrum in this case is very biased and a clear fail of the null test. (Right) Null test for a realistic survey, which includes noise and foregrounds on top of CMB. Each R or S split is further split into 2 half missions and the spectra are the cross between half missions to avoid a noise bias. This is for the f090 frequency channel. The dashed grey line marks $\text{PTE} = 0.05$.

Null test	EE	EB	BE	BB
pure signal	0.514	0.855	–	3×10^{-27}
realistic with noise	0.172	0.421	0.912	0.152

Table 1. PTEs for the uniform distribution K-S test over a sample of 100 realisations.

We perform a two-sided Kolmogorov-Smirnov (K-S) test, comparing our sample of PTEs to a uniform distribution.¹⁵ The PTE of each K-S test is listed in the first row of table 1. We pass the null test in the case of the EE and EB spectra, but we can reject to a very high significance the null hypothesis for the BB spectrum, which is clearly not a uniform distribution as shown in the left side panel of figure 10. As described in section 4.1, the TF is an approximation and the BPWF reconstruction can be biased. This null test failure will disappear when realistic noise is added to the simulation, as shown in the second row of table 1, and explained below.

5.3 Signal and noise test

Our survey will not be pure CMB signal but also contain anisotropic and non-Gaussian polarised foregrounds, as well as noise. We include the `PYSM d1` and `s1` foregrounds. The main issue of concern for our analysis is the spatial anisotropic and non-Gaussian nature of these templates, which could potentially induce biases in the deconvolved spectra due to a TF calculation that relies on spatially isotropic Gaussian simulations, as described in section 3.1.

¹⁵We use the implementation from SCIPY <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.kstest.html>.

We use these models with moderate levels of SED variability as our main interest is to assess the impact on mode coupling from the highly spatially anisotropic and non-Gaussian foregrounds. Using models with more SED complexity or even frequency decorrelation is beyond the scope of our paper. We leave the detailed study of foreground-induced biases in a TF reconstruction for future work.

For noise, we use the realisations scaled to a 1-year deep survey with noise $\sim 5 \mu\text{K-arcmin}$ for f090. We can co-add CMB, foregrounds, and noise in the same rising and setting scan splits, but we would be subjected to noise bias. To obtain noise-unbiased estimates of the signal power spectra, we further split each rising or setting map into a half-mission (HM) 1 and 2 split and cross-correlate between them. Since HM splits do not have significantly different hits maps, we continue the use of the mask from section 5.2. In this case, the residual spectrum is given by

$$\Delta C_b^{XY} = C_b^{X_{\text{R,HM1}} Y_{\text{R,HM2}}} + C_b^{X_{\text{S,HM1}} Y_{\text{S,HM2}}} - C_b^{X_{\text{R,HM1}} Y_{\text{S,HM2}}} - C_b^{X_{\text{S,HM1}} Y_{\text{R,HM2}}}. \quad (5.3)$$

Note that in this case the EB and BE null tests are different. In any realistic scenario such as this, the residual from a null test should be dominated by statistical noise rather than cosmic variance.

We follow the same procedure as in section 5.2. The histogram of PTEs for the chi-squared distribution for 100 realisations is shown on the right-hand side of figure 10. We perform a K-S test on each of the 4 EE , EB , BE , and BB polarisations, whose PTE is shown in the second row of table 1. In this case, we pass the null test for BB and we can state that our rising and setting splits are statistically compatible. This means that the TF method can be used for the reconstruction of noise-unbiased power spectra for filtered maps, at a $\sim 5 \mu\text{K-arcmin}$ noise level.

6 Full time-domain simulations

Since we have all the components necessary to simulate an entire SO-SAT deep survey, we could run a limited-scale test of our complete analysis, from observing data all the way to measuring a tensor-to-scalar ratio value with a spectra-based component separation pipeline. Since we only simulate two of the SO frequency bands, f090 and f150, we use *Planck* observations as ancillary data to perform component separation.¹⁶ Since this test has a limited scale, we compromise by only simulating the 353 GHz *Planck* band. On the other hand, we do not include synchrotron as a polarised foreground in our simulation, since we do not include the low-frequency bands necessary to constrain it, but only thermal dust.

The 100 CMB realisations we use only have lensing B -modes and $r = 0$. The thermal dust model is d1 from PYSM. In section 3.1 we described how these simulations are generated. We observe the 353 GHz *Planck* simulations with the same schedule and filtering described in sections 3.2 and 3.3 with TOAST. In this paper, we work under the simplified scheme that f090 and f150 have been observed and filtered in the same way. This can be easily generalised to the point where each frequency channel is assigned its own filtering recipe.

¹⁶The real SO deep survey will eventually observe 6 frequency bands in total and will not need to use ancillary data from other experiments.

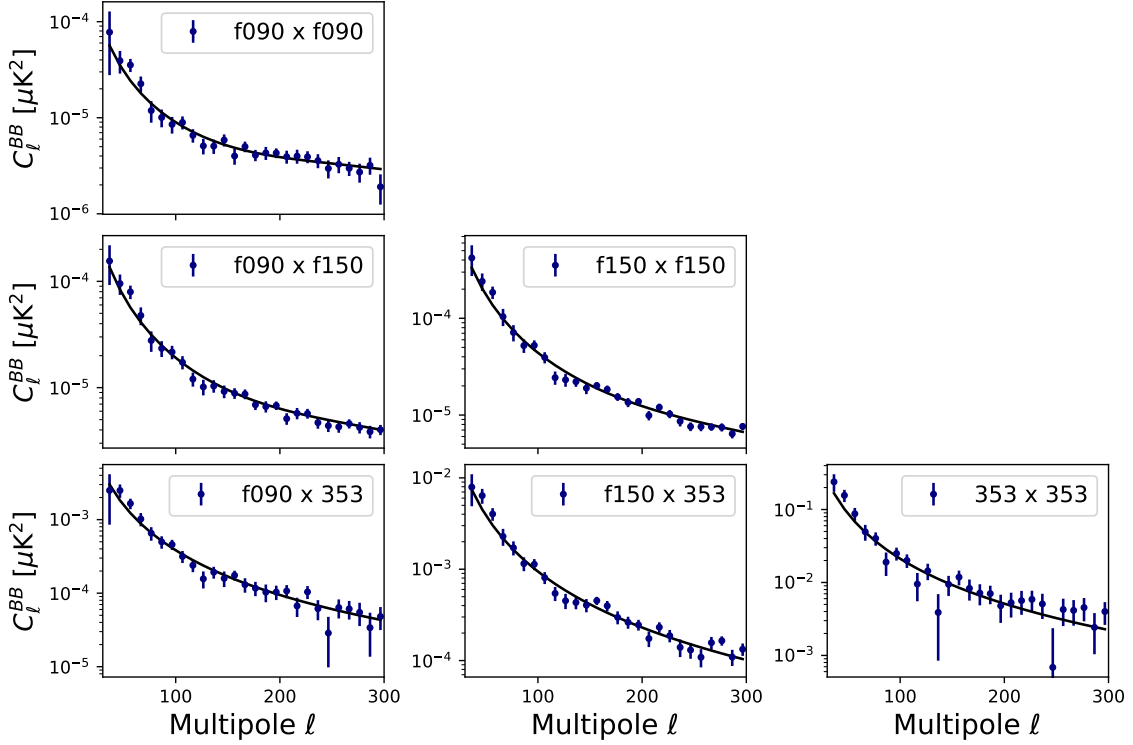


Figure 11. BB power spectra for our limited-scale 1-year deep survey simulation. Unbiased spectra are calculated from the average of the cross-spectra of splits. The spectra for one realisation are shown in filled blue dots, while the error bars are calculated from the standard deviation across 100 realisations of the simulations that include Gaussian dust. The best-fit CMB+dust model is shown as the black solid line.

This filtering scheme is applied over an entire simulated *Planck* 353 GHz map, which includes the noise. We have 100 realisations of the NPIPE 353 GHz frequency maps, with their A/B detector split. For the two SAT bands, we define four quarter-mission (QM) splits. The pure power-law simulations, used to estimate the transfer function, are co-added with these same splits, and the TFs are calculated for the corresponding split combinations and used to calculate the corresponding cross-spectra. Finally, we estimate unbiased spectra on the simulated observations of CMB+dust+noise by averaging all cross-spectra combinations, avoiding all auto-spectra. For example, the f090×f090 spectrum is calculated from the 6 cross-spectra combinations between the 4 QM splits, the f090×f150 spectrum is calculated from the 16 cross-spectra combinations, etc. For spectra, we use the full South patch mask constructed from the hits map shown in figure 2.

To estimate a covariance for error bars and parameter posteriors with Markov chain Monte Carlo (MCMC) sampling, we use 100 realisations of Gaussian dust and calculate their deconvolved spectra in the same way as described above. The total covariance is estimated from the 100 realizations of CMB + noise + Gaussian dust.

The unbiased C_ℓ^{BB} spectra are shown in figure 11 in filled blue dots. This is the spectrum for one realisation. The error bar is calculated from the standard deviation across 100 realisations of the simulations that include the Gaussian dust.

Parameter	r	A_{lens}	A_{d}^{BB}	α_{d}^{BB}	β_{d}
Units	–	–	μK^2	–	–
Prior type	TH	TH	TH	TH	G
Prior value	$[-0.1, +0.1]$	$[0, 2]$	$[0, 60]$	$[-1.0, +0.8]$	1.54 ± 0.12
Fiducial value	0	1	$36.90 \pm 1.17^\dagger$	$-0.10 \pm 0.03^\dagger$	$1.55 \pm 0.01^\dagger$

[†] These are fitted to the pure unfiltered **d1** dust model, and as such are shown here with their 1σ error bar. We used the fitted value as the fiducial value of the parameter.

Table 2. Parameters and priors used in our MCMC analysis. TH is a uniform top-hat prior in the given range, while G is a Gaussian prior with the mean \pm the standard deviation.

Our next step is to do component separation and estimate posterior probabilities for parameters, in particular r . We do this with the spectra-based BBPOWER pipeline,¹⁷ which we have used before to forecast the performance of SO [30, 79–81]. We choose this method for component separation, as it is the only foreground removal approach for which efficient and accurate methods exist to account for the impact of timestream- and/or map-level filtering (e.g. [50, 82]). We fit a model only to the BB spectrum, with five free parameters: r , a lensing amplitude A_{lens} , a dust power spectrum amplitude (at pivot scale $\ell = 80$ and pivot frequency 353 GHz) A_{d}^{BB} , a power-law index of the dust power spectrum α_{d}^{BB} , and a frequency spectral index β_{d} , corresponding to the assumed MBB spectral energy distribution of Galactic dust emission. The temperature $T_{\text{d}} = 20$ K of the dust MBB is fixed in our likelihood model. For each parameter, we have a fiducial value. The CMB realisations are generated with no primordial B -modes ($r = 0$), and $A_{\text{lens}} = 1$, while **d1** dust has spatially variable SED spectral parameters and equation (3.1) is a good approximation within small patches on the sky, but it is not exact for a larger patch such as the South Galactic patch. We wish to determine these fiducial parameters for the **d1** dust template. We fit these parameters to the pure unfiltered **d1** template within our common hits map at the three relevant frequencies: 93, 145, and 353 GHz, computing spectra in the multipole range $\ell = 30$ –300 and within our mask. As error bars for the fit, we use the standard deviation across the 100 realisations of unfiltered Gaussian dust. We use the CURVE_FIT method from SCIPY [83] to perform the fit. These fiducial values and priors are listed in table 2 for the five parameters. While we list the dust fiducial values with their 1σ uncertainties, we use the fitted value alone.

We run BBPOWER implementing MCMC from EMCEE [84] with 30 walkers and 10,000 iterations. We repeat for each of the 100 realisations. Figure 12, left, shows the histogram of the χ^2 values for the C_{ℓ}^{BB} spectrum. Since we have six distinct cross-frequency BB power spectra, each with 27 multipole bins between $\ell = 30$ –300, and five fitted parameters, the number of degrees of freedom (DOF) for our best-fit chi-squared distribution is 157, shown as the blue dashed vertical line. In blue we show the χ^2 calculated with respect to best-fit values. In orange, we show the χ^2 calculated with respect to fiducial values listed in table 2. The DOF in this case is 162, without fitted parameters. This is shown as the orange dashed vertical line.

In order to focus on the ability of our pipeline to obtain unbiased constraints from the first year of SO-SAT data, and not on the forecasted sensitivity, we present our simulated

¹⁷<https://github.com/simonsobs/BBPower>.

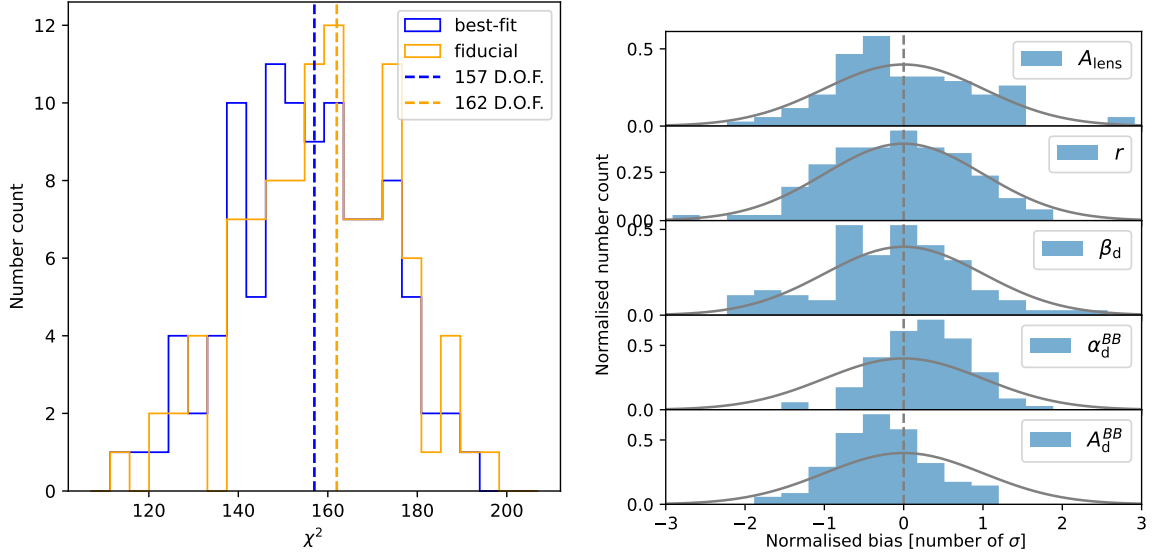


Figure 12. *Left panel:* distribution of χ^2 values for the BB power spectra with respect to the input model (orange) and with respect to the best-fit model of each simulation (blue). The vertical lines mark the number of elements in the data vector in each case. *Right panel:* distribution of the normalised bias for the best-fit parameters in each simulation. For parameter θ , these are estimated as $(\theta_{\text{BF}} - \theta_{\text{in}})/\sigma_\theta$, where θ_{BF} and θ_{in} are the best-fit value and the input value for the parameter, and σ_θ is its inferred posterior standard deviation. For reference, we also show a normal distribution with unit variance $\mathcal{N}(0, 1)$. Small deviations from this reference distribution are not unexpected, as the posterior distributions of these parameters are not exactly Gaussian.

cosmological constraints normalised by their standard deviation, referred to in this work as the normalised bias.¹⁸ The histogram for the normalised bias for the best-fit value for each parameter is shown in figure 12, right. Figure 13 shows the posterior distributions for our five parameters for the 100 realisations. For each parameter value, we subtract the fiducial value and divide by the 1σ uncertainty.

7 Summary and conclusions

In this work, we described methods to estimate pseudo- C_ℓ power spectra from biased filter+bin maps in the context of small aperture telescope surveys in CMB experiments, looking for the signature of primordial gravitational waves. The maps are biased due to timestream filtering and there is no analytical formula to account for the mode coupling or polarisation leakage induced on the pseudo- C_ℓ estimate. Hence, simulations are required, where the exact mixing can be estimated by measuring how individual ℓ s mix into one another. However, this approach is limited by the cosmic variance of the maps and therefore a huge number of realisations ($\mathcal{O}(10^4)$) is required to reduce the estimator noise down to a practical level. A much computationally cheaper alternative approach is to assume the mixing is the analytical solution from MASTER for a masked sky, suppressed by a transfer function measured in multipole band powers, under the approximations described in section 2.2.2. This transfer function can be cheaply estimated averaging over $\mathcal{O}(100)$ filtered realisations of pure E/B modes.

¹⁸We subtract the fiducial value from the parameter and divide by the 1σ error bar.




Figure 13. Triangle plot of the posterior distributions of residual parameters for all 100 realisations. We concatenated the chains for the 100 realisations and estimated the posteriors over that. For each parameter value, we subtract the fiducial value listed in table 2 and divide by the 1σ uncertainty.

The main objective of this work is to validate this method with simulations of the deep survey of the Simons Observatory. Using the TOAST software, we simulate the observing of the CMB during a one-year season, scanning into TODs and simulating the modulation by a rotating HWP. We also demodulate and apply time-domain filtering to produce filter+bin maps. However, for the real SO data, we do not expect to use filtering as aggressive as what we use in this work. Besides scanning pure signal maps, we also simulate a realistic noise with the same scanning schedule.

While the transfer function estimator can be used at smaller scales, it is biased by $> 1\sigma$ in the large-scale limit $\ell \sim 30$. We illustrate this by performing a simple null test of rising versus setting scans. Besides CMB, we include foreground emission and noise consistent with a one-year deep survey with two SO-SATs. In this case, the bias induced in the largest-scale band powers is small compared to the statistical noise introduced, so such a null test is passed.

Finally, we perform an end-to-end analysis, starting from scanning the sky, through making maps, calculating the cross-spectra, and doing likelihood fitting cosmological and

foreground parameters. By calculating a χ^2 statistic, as exemplified in figure. 12, we show that our spectra are unbiased. We also show that the constraints on the fitted parameters (especially r and A_{lens}) are unbiased. We conclude that the transfer function method described in this paper is validated for a survey with a noise level $\sim 5 \mu\text{K-arcmin}$ at 90 GHz, and thus is valid for use in the analysis of the first 1–2 years of SO data.

Acknowledgments

CHC acknowledges ANID FONDECYT Postdoc Fellowship 3220255 and BASAL CATA FB210003. KW is supported by the STFC, grant ST/X006344/1, and by a Gianturco Junior Research Fellowship of Linacre College, Oxford. ALP and DA acknowledge support from an STFC Consolidated Grant (ST/W000903/1). CB acknowledges partial support by the Italian Space Agency LiteBIRD Project (ASI Grants No. 2020-9-HH.0 and 2016-24-H.1-2018), as well as the InDark and LiteBIRD Initiative of the National Institute for Nuclear Physics, and the RadioForegroundsPlus Project HORIZON-CL4-2023-SPACE-01, GA 101135036. MLB and DBT acknowledge support from the UKRI/STFC (grant number ST/X006344/1). EC acknowledges support from the Horizon 2020 ERC Starting Grant (Grant agreement No 849169). YC acknowledges the support from JSPS KAKENHI Grant Number JP24K00667. RD thanks ANID for grants BASAL CATA FB210003, FONDEF ID21I10236 and QUIMAL240004. JE, SB, ETKS, WS,AVA acknowledge the SCIPOL project funded by the ERC under the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 101044073). GF is supported by the STFC Ernest Rutherford fellowship. SG acknowledges support from STFC and UKRI (grant numbers ST/W002892/1 and ST/X006360/1). MM is funded by the European Union (ERC, RELiCS, project number 101116027). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The Geryon cluster at the Centro de Astro-Ingenieria UC was extensively used for the calculations performed in this paper. ANID BASAL project FB21000, BASAL CATA PFB-06, the Anillo ACT-86, FONDEQUIP AIC-57, and QUIMAL 130008 provided funding for several improvements to the Geryon cluster. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award HEP-ERCAP-mp107. This work was supported in part by a grant from the Simons Foundation (Award #457687, B.K.). This research has made extensive use of NUMPY [85], SCIPY [83], MATPLOTLIB [86], NaMaster [64], HEALPY [87], EMCEE [84], PYSM [71], and GETDIST [88].

A Effect of individual filters

Here, we describe the features of the transfer function presented in section 4.2 and figure 5. Figure 14 shows transfer functions estimated for different combinations of individual filters. This is calculated from 100 realizations of power-law simulations and from a fraction of the full schedule described in section 3.2.




Figure 14. TF estimated for individual filters. The blue solid line shows the effect of only demodulating the HWP, while other lines show the effect of demodulation plus another filter. This is estimated from the average of 100 realizations of power-law simulations and on a fraction of the full schedule used in this paper. This is for the f090 frequency channel.

The first filter applied is the demodulation of the HWP, which has no effect on large scales and a minimal effect towards smaller scales, since we low-pass the timestreams at 2 Hz after demodulating, so higher frequencies (and therefore smaller scales) are completely suppressed.¹⁹ The HWP synchronous filter has effect on very large scales, while the ground filter induces a very large suppression of the very large-scale modes. Overall, demodulation, HWP synchronous signal and ground filters do not show suppression on smaller scales $\ell \sim 300$ and therefore the TF should approach $T_b \sim 1$.

On the other hand, the full TF shown in figure 5 plateaus at a lower value $T_b \sim 0.8$. As we can see here, the high-pass polynomial filter is responsible for this (shown in the solid red and dashed-dotted purple lines), which suppresses the large scales by a large amount, and suppresses the middle-to-small scales by $\sim 20\%$. Moreover, a polynomial of higher order 3 suppresses more modes than a polynomial of order 1. In this paper, we use a high-pass filter with order 3. To analyse real data, a pipeline with these filtering choices adjusted is already being used to reduce the impact of these transfer functions.

B Illustration of focal plane thinning

To ease the computation of simulations in TOAST, we thin the focal plane by 8, i.e., we only simulate one for every eight detectors. Figure 15 shows the EE TF calculated from a single

¹⁹Note that there is some interplay between the resolution of the pixelated input map and the low-pass filter of the demodulation, where there might be a fraction of power above ~ 2 Hz that is lost when low-pass filtering. This effect is more pronounced at $N_{\text{side}} = 512$, the resolution of our simulations, than at a higher N_{side} like 2048, and therefore all of the transfer functions in this plot would show an extra suppression towards higher ℓ .

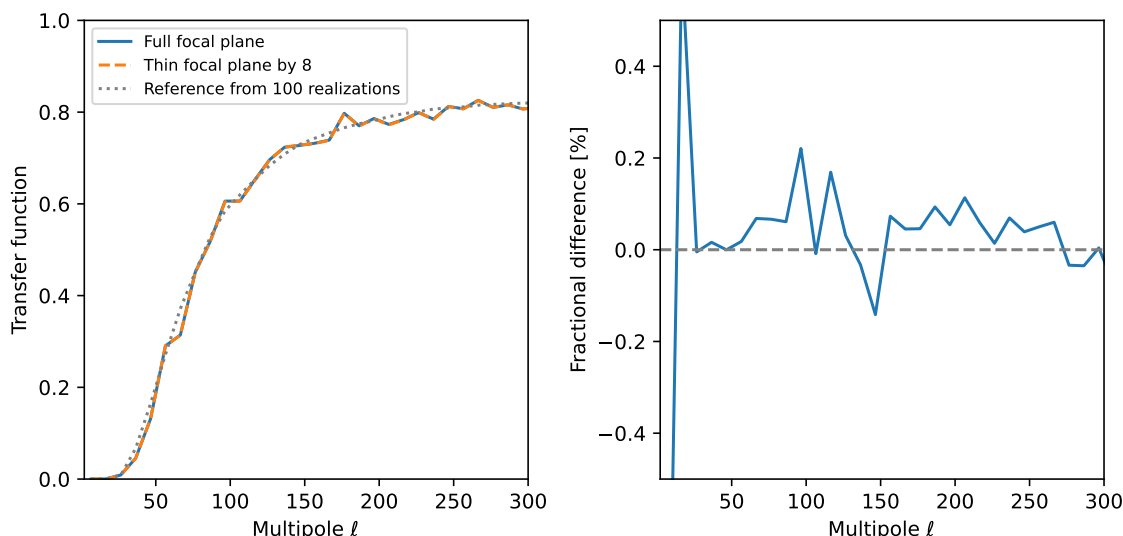


Figure 15. Illustration of the effect of focal plane thinning. EE TF estimated with a single simulated power-law map, using both the full focal plane and thinning it by a factor of 8. We also show the smooth reference TF estimated from 100 realizations in dotted grey (the same shown in figure 5). The fractional difference, as shown in the right panel, is within 0.2% for most of the multipole range.

power-law simulated map, which is filtered both using the full focal plane as well as thinning the focal plane by a factor of 8. The fractional difference is $\sim 0.1\%$ for most of the multipole range, increasing to $\sim 1\%$ only for the first couple of bins at $\ell < 30$.

References

- [1] CMB-S4 collaboration, *CMB-S4 Science Book, First Edition*, [arXiv:1610.02743](#) [[INSPIRE](#)].
- [2] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [3] 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](#)].
- [4] 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](#)].
- [5] BICEP and KECK collaborations, *Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season*, *Phys. Rev. Lett.* **127** (2021) 151301 [[arXiv:2110.00483](#)] [[INSPIRE](#)].
- [6] A.A. Starobinsky, *A New Type of Isotropic Cosmological Models Without Singularity*, *Phys. Lett. B* **91** (1980) 99 [[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] V.F. Mukhanov and G.V. Chibisov, *Quantum Fluctuations and a Nonsingular Universe*, *JETP Lett.* **33** (1981) 532 [[INSPIRE](#)].

- [10] A.H. Guth and S.Y. Pi, *Fluctuations in the New Inflationary Universe*, *Phys. Rev. Lett.* **49** (1982) 1110 [INSPIRE].
- [11] A.A. Starobinsky, *Dynamics of Phase Transition in the New Inflationary Universe Scenario and Generation of Perturbations*, *Phys. Lett. B* **117** (1982) 175 [INSPIRE].
- [12] J.M. Bardeen, P.J. Steinhardt and M.S. Turner, *Spontaneous Creation of Almost Scale-Free Density Perturbations in an Inflationary Universe*, *Phys. Rev. D* **28** (1983) 679 [INSPIRE].
- [13] L.F. Abbott and M.B. Wise, *Constraints on Generalized Inflationary Cosmologies*, *Nucl. Phys. B* **244** (1984) 541 [INSPIRE].
- [14] BICEP /KECK et al. collaborations, *Bicep/KeckXV: The Bicep3 Cosmic Microwave Background Polarimeter and the First Three-year Data Set*, *Astrophys. J.* **927** (2022) 77 [arXiv:2110.00482] [INSPIRE].
- [15] H. Hui et al., *BICEP Array: a multi-frequency degree-scale CMB polarimeter*, *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070807 [arXiv:1808.00568] [INSPIRE].
- [16] SPIDER collaboration, *A Constraint on Primordial B-modes from the First Flight of the Spider Balloon-borne Telescope*, *Astrophys. J.* **927** (2022) 174 [arXiv:2103.13334] [INSPIRE].
- [17] T. Essinger-Hileman et al., *CLASS: The Cosmology Large Angular Scale Surveyor*, *Proc. SPIE Int. Soc. Opt. Eng.* **9153** (2014) 91531I [arXiv:1408.4788] [INSPIRE].
- [18] POLARBEAR collaboration, *The POLARBEAR-2 and the Simons Array Experiment*, *J. Low Temp. Phys.* **184** (2016) 805 [arXiv:1512.07299] [INSPIRE].
- [19] QUBIC collaboration, *QUBIC I: Overview and science program*, *JCAP* **04** (2022) 034 [arXiv:2011.02213] [INSPIRE].
- [20] GROUND BIRD collaboration, *GroundBIRD: A CMB Polarization Experiment with MKID Arrays*, *J. Low Temp. Phys.* **200** (2020) 384 [arXiv:2011.07705] [INSPIRE].
- [21] S. Ghosh et al., *Performance forecasts for the primordial gravitational wave detection pipelines for AliCPT-1*, *JCAP* **10** (2022) 063 [arXiv:2205.14804] [INSPIRE].
- [22] CMB-S4 collaboration, *CMB-S4: Forecasting Constraints on Primordial Gravitational Waves*, *Astrophys. J.* **926** (2022) 54 [arXiv:2008.12619] [INSPIRE].
- [23] LITEBIRD collaboration, *Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey*, *PTEP* **2023** (2023) 042F01 [arXiv:2202.02773] [INSPIRE].
- [24] PLANCK collaboration, *Planck 2015 results. X. Diffuse component separation: Foreground maps*, *Astron. Astrophys.* **594** (2016) A10 [arXiv:1502.01588] [INSPIRE].
- [25] N. Krachmalnicoff et al., *Characterization of foreground emission on degree angular scales for CMB B-mode observations — Thermal dust and synchrotron signal from Planck and WMAP data*, *Astron. Astrophys.* **588** (2016) A65 [arXiv:1511.00532] [INSPIRE].
- [26] S.M. Leach et al., *Component separation methods for the Planck mission*, *Astron. Astrophys.* **491** (2008) 597 [arXiv:0805.0269] [INSPIRE].
- [27] A. Lewis and A. Challinor, *Weak gravitational lensing of the CMB*, *Phys. Rept.* **429** (2006) 1 [astro-ph/0601594] [INSPIRE].
- [28] M. Zaldarriaga and U. Seljak, *Gravitational lensing effect on cosmic microwave background polarization*, *Phys. Rev. D* **58** (1998) 023003 [astro-ph/9803150] [INSPIRE].
- [29] T. Namikawa et al., *Simons Observatory: Constraining inflationary gravitational waves with multitracers B-mode delensing*, *Phys. Rev. D* **105** (2022) 023511 [arXiv:2110.09730] [INSPIRE].

- [30] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Combining cross-spectral foreground cleaning with multitracer B-mode delensing for improved constraints on inflation*, *Phys. Rev. D* **110** (2024) 043532 [[arXiv:2405.01621](#)] [[INSPIRE](#)].
- [31] K. Wolz et al., *The Simons Observatory: pipeline comparison and validation for large-scale B-modes*, *Astron. Astrophys.* **686** (2024) A16 [[arXiv:2302.04276](#)] [[INSPIRE](#)].
- [32] SO collaboration, *The Simons Observatory: Design, Integration, and Testing of the Small Aperture Telescopes*, *Astrophys. J. Suppl.* **274** (2024) 33 [[arXiv:2405.05550](#)] [[INSPIRE](#)].
- [33] POLARBEAR collaboration, *Modeling atmospheric emission for CMB ground-based observations*, *Astrophys. J.* **809** (2015) 63 [[arXiv:1501.07911](#)] [[INSPIRE](#)].
- [34] WMAP collaboration, *Three-year Wilkinson Microwave Anisotropy Pr (WMAP) observations: beam profiles, data processing, radiometer characterization and systematic error limits*, *Astrophys. J. Suppl.* **170** (2007) 263 [[astro-ph/0603452](#)] [[INSPIRE](#)].
- [35] BICEP2 collaboration, *BICEP2 III: Instrumental Systematics*, *Astrophys. J.* **814** (2015) 110 [[arXiv:1502.00608](#)] [[INSPIRE](#)].
- [36] ABS collaboration, *Modulation of cosmic microwave background polarization with a warm rapidly rotating half-wave plate on the Atacama B-Mode Search instrument*, *Rev. Sci. Instrum.* **85** (2014) 024501 [[arXiv:1310.3711](#)] [[INSPIRE](#)].
- [37] ABS collaboration, *Systematic effects from an ambient-temperature, continuously rotating half-wave plate*, *Rev. Sci. Instrum.* **87** (2016) 094503 [[arXiv:1601.05901](#)] [[INSPIRE](#)].
- [38] B.R. Johnson et al., *MAXIPOL: Cosmic Microwave Background Polarimetry Using a Rotating Half-Wave Plate*, *Astrophys. J.* **665** (2007) 42 [[astro-ph/0611394](#)] [[INSPIRE](#)].
- [39] A. Kusaka et al., *Results from the Atacama B-mode Search (ABS) Experiment*, *JCAP* **09** (2018) 005 [[arXiv:1801.01218](#)] [[INSPIRE](#)].
- [40] S. Takakura et al., *Performance of a continuously rotating half-wave plate on the POLARBEAR telescope*, *JCAP* **05** (2017) 008 [[arXiv:1702.07111](#)] [[INSPIRE](#)].
- [41] CLASS collaboration, *A Measurement of the Largest-scale CMB E-mode Polarization with CLASS*, *Astrophys. J.* **986** (2025) 111 [[arXiv:2501.11904](#)] [[INSPIRE](#)].
- [42] M. Rashid, M.L. Brown and D.B. Thomas, *CMB polarization signal demodulation with a rotating half-wave plate*, *Mon. Not. Roy. Astron. Soc.* **532** (2024) 2309 [[arXiv:2307.01860](#)] [[INSPIRE](#)].
- [43] POLARBEAR collaboration, *A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales with POLARBEAR*, *Astrophys. J.* **794** (2014) 171 [Erratum *ibid.* **848** (2017) 73] [[arXiv:1403.2369](#)] [[INSPIRE](#)].
- [44] POLARBEAR collaboration, *A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales from 2 years of POLARBEAR Data*, *Astrophys. J.* **848** (2017) 121 [[arXiv:1705.02907](#)] [[INSPIRE](#)].
- [45] D. Poletti et al., *Making maps of Cosmic Microwave Background polarization for B-mode studies: the POLARBEAR example*, *Astron. Astrophys.* **600** (2017) A60 [[arXiv:1608.01624](#)] [[INSPIRE](#)].
- [46] POLARBEAR collaboration, *A Measurement of the Degree Scale CMB B-mode Angular Power Spectrum with POLARBEAR*, *Astrophys. J.* **897** (2020) 55 [[arXiv:1910.02608](#)] [[INSPIRE](#)].
- [47] QUAD collaboration, *Improved measurements of the temperature and polarization of the CMB from QUaD*, *Astrophys. J.* **705** (2009) 978 [[arXiv:0906.1003](#)] [[INSPIRE](#)].
- [48] H.C. Chiang et al., *Measurement of CMB Polarization Power Spectra from Two Years of BICEP Data*, *Astrophys. J.* **711** (2010) 1123 [[arXiv:0906.1181](#)] [[INSPIRE](#)].

- [49] BICEP1 collaboration, *Degree-Scale CMB Polarization Measurements from Three Years of BICEP1 Data*, *Astrophys. J.* **783** (2014) 67 [[arXiv:1310.1422](#)] [[INSPIRE](#)].
- [50] BICEP2 collaboration, *Detection of B-Mode Polarization at Degree Angular Scales by BICEP2*, *Phys. Rev. Lett.* **112** (2014) 241101 [[arXiv:1403.3985](#)] [[INSPIRE](#)].
- [51] BICEP2 and KECK ARRAY collaborations, *BICEP2/Keck Array V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the Keck Array*, *Astrophys. J.* **811** (2015) 126 [[arXiv:1502.00643](#)] [[INSPIRE](#)].
- [52] SPT collaboration, *Measurements of B-mode Polarization of the Cosmic Microwave Background from 500 Square Degrees of SPTpol Data*, *Phys. Rev. D* **101** (2020) 122003 [[arXiv:1910.05748](#)] [[INSPIRE](#)].
- [53] M. Tegmark, *How to make maps from CMB data without losing information*, *Astrophys. J. Lett.* **480** (1997) L87 [[astro-ph/9611130](#)] [[INSPIRE](#)].
- [54] S. Naess and T. Louis, *Large-scale power loss in ground-based CMB mapmaking*, *Open J. Astrophys.* **6** (2023) 2210.02243 [[arXiv:2210.02243](#)] [[INSPIRE](#)].
- [55] K.M. Gorski, *On Determining the spectrum of primordial inhomogeneity from the Cobe DMR sky maps. 1. Method*, *Astrophys. J. Lett.* **430** (1994) L85 [[astro-ph/9403066](#)] [[INSPIRE](#)].
- [56] J.R. Bond, *Signal-to-noise eigenmode analysis of the two year COBE maps*, *Phys. Rev. Lett.* **74** (1995) 4369 [[astro-ph/9407044](#)] [[INSPIRE](#)].
- [57] M. Tegmark, *How to measure CMB power spectra without losing information*, *Phys. Rev. D* **55** (1997) 5895 [[astro-ph/9611174](#)] [[INSPIRE](#)].
- [58] J.R. Bond, A.H. Jaffe and L. Knox, *Estimating the power spectrum of the cosmic microwave background*, *Phys. Rev. D* **57** (1998) 2117 [[astro-ph/9708203](#)] [[INSPIRE](#)].
- [59] M. Tegmark and A. de Oliveira-Costa, *How to measure CMB polarization power spectra without losing information*, *Phys. Rev. D* **64** (2001) 063001 [[astro-ph/0012120](#)] [[INSPIRE](#)].
- [60] E. Hivon et al., *Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets*, *Astrophys. J.* **567** (2002) 2 [[astro-ph/0105302](#)] [[INSPIRE](#)].
- [61] L. Knox, *Cosmic microwave background anisotropy window functions revisited*, *Phys. Rev. D* **60** (1999) 103516 [[astro-ph/9902046](#)] [[INSPIRE](#)].
- [62] J.S.-Y. Leung et al., *A Simulation-based Method for Correcting Mode Coupling in CMB Angular Power Spectra*, *Astrophys. J.* **928** (2022) 109 [[arXiv:2111.01113](#)] [[INSPIRE](#)].
- [63] M.F. Hutchinson, *A Stochastic Estimator of the Trace of the Influence Matrix for Laplacian Smoothing Splines*, *Commun. Stat. Simul. Comput.* **18** (1989) 1059.
- [64] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- C_ℓ framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [65] 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](#)].
- [66] K.M. Smith, *Pseudo- C_ℓ estimators which do not mix E and B modes*, *Phys. Rev. D* **74** (2006) 083002 [[astro-ph/0511629](#)] [[INSPIRE](#)].
- [67] E.F. Bunn, M. Zaldarriaga, M. Tegmark and A. de Oliveira-Costa, *E/B decomposition of finite pixelized CMB maps*, *Phys. Rev. D* **67** (2003) 023501 [[astro-ph/0207338](#)] [[INSPIRE](#)].
- [68] J. Grain, M. Tristram and R. Stompor, *Polarized CMB spectrum estimation using the pure pseudo cross-spectrum approach*, *Phys. Rev. D* **79** (2009) 123515 [[arXiv:0903.2350](#)] [[INSPIRE](#)].

- [69] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [Erratum *ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [70] PLANCK collaboration, *Planck intermediate results. LVII. Joint Planck LFI and HFI data processing*, *Astron. Astrophys.* **643** (2020) A42 [[arXiv:2007.04997](#)] [[INSPIRE](#)].
- [71] 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](#)].
- [72] S.M. Simon et al., *Characterizing Atacama B-mode Search Detectors with a Half-Wave Plate*, *J. Low Temp. Phys.* **184** (2016) 534 [[arXiv:1511.04760](#)] [[INSPIRE](#)].
- [73] K. Yamada et al., *The Simons Observatory: Cryogenic half wave plate rotation mechanism for the small aperture telescopes*, *Rev. Sci. Instrum.* **95** (2024) 024504 [Erratum *ibid.* **96** (2025) 019901] [[arXiv:2309.14803](#)] [[INSPIRE](#)].
- [74] S. Spinelli et al., *A template of atmospheric O₂ circularly polarized emission for CMB experiments*, *Mon. Not. Roy. Astron. Soc.* **414** (2011) 3272 [[arXiv:1103.0160](#)] [[INSPIRE](#)].
- [75] CLASS collaboration, *Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: A First Detection of Atmospheric Circular Polarization at Q Band*, *Astrophys. J.* **889** (2020) 120 [[arXiv:1911.01016](#)] [[INSPIRE](#)].
- [76] S.E. Church, *Predicting residual levels of atmospheric sky noise in ground-based observations of the cosmic background radiation*, *Mon. Not. Roy. Astron. Soc.* **272** (1995) 551.
- [77] M. Salatino et al., *Studies of Systematic Uncertainties for Simons Observatory: Polarization Modulator Related Effects*, *Proc. SPIE* **10708** (2018) 1070848 [[arXiv:1808.07442](#)] [[INSPIRE](#)].
- [78] BICEP2 and KECK ARRAY collaborations, *BICEP2/Keck Array VII: Matrix based E/B Separation applied to BICEP2 and the Keck Array*, *Astrophys. J.* **825** (2016) 66 [[arXiv:1603.05976](#)] [[INSPIRE](#)].
- [79] M.H. Abitbol et al., *The Simons Observatory: gain, bandpass and polarization-angle calibration requirements for B-mode searches*, *JCAP* **05** (2021) 032 [[arXiv:2011.02449](#)] [[INSPIRE](#)].
- [80] S. Azzoni et al., *A minimal power-spectrum-based moment expansion for CMB B-mode searches*, *JCAP* **05** (2021) 047 [[arXiv:2011.11575](#)] [[INSPIRE](#)].
- [81] S. Azzoni et al., *A hybrid map-Cl component separation method for primordial CMB B-mode searches*, *JCAP* **03** (2023) 035 [[arXiv:2210.14838](#)] [[INSPIRE](#)].
- [82] ACT collaboration, *The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz*, *JCAP* **12** (2020) 045 [[arXiv:2007.07289](#)] [[INSPIRE](#)].
- [83] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [84] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306 [[arXiv:1202.3665](#)] [[INSPIRE](#)].
- [85] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [86] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [87] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [88] A. Lewis, *GetDist: a Python package for analysing Monte Carlo samples*, [arXiv:1910.13970](#) [[INSPIRE](#)].