



HAL
open science

Accurate CMB covariance matrices: exact calculation and approximations

Étienne Camphuis, Karim Benabed, Silvia Galli, Éric Hivon, Marc Lilley

► **To cite this version:**

Étienne Camphuis, Karim Benabed, Silvia Galli, Éric Hivon, Marc Lilley. Accurate CMB covariance matrices: exact calculation and approximations. *Astronomy and Astrophysics - A&A*, 2022, 668, pp.A62. 10.1051/0004-6361/202243948 . hal-03667872

HAL Id: hal-03667872

<https://hal.science/hal-03667872v1>

Submitted on 24 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Accurate cosmic microwave background covariance matrices: Exact calculation and approximations

E. Camphuis¹ , K. Benabed¹, S. Galli¹, E. Hivon¹ , and M. Lilley²

¹ Sorbonne Université, UMR7095, Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, 75014 Paris, France
e-mail: etienne.camphuis@iap.fr

² SYRTE, Observatoire de Paris, Université PSL, CNRS, Sorbonne Université, LNE, 61 avenue de l'Observatoire, 75014 Paris, France

Received 3 May 2022 / Accepted 29 September 2022

ABSTRACT

Context. A reliable estimation of cosmological parameters from pseudo-power spectrum estimators requires accurate covariance matrices.

Aims. We focus on the analytical calculation of covariance matrices. We consider the case of observations of the cosmic microwave background (CMB) in temperature and polarization on a small footprint such as in the South Pole Telescope third-generation (SPT-3G) experiment, which observes 4% of the sky. Power spectra evaluated on small footprints are expected to have strong correlations between modes, and these need to be accurately modeled.

Methods. We present for the first time an algorithm that allows an efficient (but computationally expensive) exact calculation of analytic covariance matrices. Using it as our reference, we tested the accuracy of existing fast approximations of the covariance matrix. Furthermore, we propose a new approximation that is designed to be more precise. Finally, we derived the covariance matrices for mask-corrected power spectra estimated by the PolSpice code. In particular, in the case of a small sky fraction, we included the effect of the apodization of the large-scale modes.

Results. We find that when the power spectrum is binned in wide bandpowers, current approximations of the covariance matrix are correct up to the 5% level on the SPT-3G small sky footprint. Our new approximation improves the previous approximations and reaches a precision of 1% for the wide bandpowers. It is generally more than four times more accurate than current approaches.

Conclusions. While we considered the specific case of the CMB, our results are applicable to any other cosmological probe that requires the calculation of pseudo-power spectrum covariance matrices.

Key words. cosmic background radiation – cosmology: observations – cosmological parameters – methods: data analysis

1. Introduction

One of the most powerful probes of cosmology is the observation of the cosmic microwave background (CMB) anisotropies. The NASA WMAP and the ESA Planck satellite CMB measurements marked the entry into the era of precision cosmology, in which many Λ cold dark matter (Λ CDM) matter cosmological parameters are measured with uncertainties smaller than 1% (Hinshaw et al. 2013; Planck Collaboration V 2020). Ongoing and upcoming ground-based and satellite experiments such as the Atacama Cosmology Telescope (ACT; Aiola et al. 2020), the South Pole Telescope (SPT; Dutcher et al. 2021), the Simons Observatory (SO; Ade et al. 2019), CMB-Stage 4 (CMB-S4; Gallardo et al. 2022), and Litebird (Hazumi et al. 2012) will provide yet more information about the nature of our Universe.

Because primary CMB anisotropies in intensity and polarization are distributed as a Gaussian random field, most of the cosmological information is contained in the angular power spectrum of the CMB anisotropies. As the evolution of the primary anisotropies is linear, the multipoles of the angular power spectrum are uncorrelated when the full sky is observed. However, any realistic experiment requires masking parts of the sky, either to avoid regions that are highly contaminated by foregrounds (e.g., galactic emission or point sources) or because the scanning strategy is designed to observe specific regions of the sky. The estimation of the power spectrum on the masked

sky, the so-called pseudo-power spectrum, is biased, and different multipoles become correlated (Hivon et al. 2002). An unbiased estimator of the spectra can then be obtained through the MASTER approach (Hivon et al. 2002), as implemented in the PolSpice¹ software, for instance (Szapudi et al. 2001; Chon et al. 2004). A robust inference of cosmological parameters requires accurate covariance matrices that describe the variance of the spectra along their diagonal, as well as the correlations between multipoles in the off-diagonal terms. Pseudo- C_ℓ covariance matrices are corrected for the effect of the mask using MASTER to obtain the covariance matrices for the unbiased C_ℓ estimator. Inaccuracies in the covariance matrix estimation can lead to the misestimation of cosmological parameters and of their uncertainties (Dodelson & Schneider 2013; Sellentin & Starck 2019).

Covariance matrices can be calculated through the use of simulations. The number of simulations determines the accuracy of the estimator. As the simulations are expensive to produce, the obtained noisy realization of the covariance has to be regularized (Balkenhol & Reichardt 2022)². Alternatively, it is possible to calculate pseudo- C_ℓ covariance matrices analytically. However,

¹ <http://www2.iap.fr/users/hivon/software/PolSpice/>

² While this work focuses on covariance estimates obtained through empirical estimators, the conditioning schemes it presents can similarly be applied to estimates from simulations.

these depend on integrals whose exact numerical implementation is computationally expensive. Thus, approximations have been proposed in previous works to make these calculations efficient; see, for example, [Efstathiou \(2004\)](#), [Nicola et al. \(2021\)](#) and [Friedrich et al. \(2021\)](#).

We analyze the problem of computing accurate analytical covariance matrices. We take the specific case of the South Pole Telescope third-generation (SPT-3G) experiment, which observes the CMB anisotropies in temperature and polarization on a small sky patch that corresponds to about 4% of the sky. On such a small sky region, the calculated power spectra has strong correlations between multipoles. The existing approximations of the covariance matrix can be less accurate in these conditions. Considering this particular case is thus a particularly stringent test of the validity of analytical algorithms.

We implement for the first time the exact computationally expensive calculation of the covariance matrices, which we find to be numerically feasible at multipoles smaller than $\ell \lesssim \ell_{\text{max,ex}} \equiv 1000$ through a new algorithm that gains one order of numerical complexity over the brute-force approach, resulting in a thousandfold speed improvement. Then, we test the existing approximations, and find that they are accurate at the 5% level in the case of the SPT-3G footprint when the power spectrum is averaged in wide bandpowers. We then propose a new approximation that improves the existing algorithms to attain an accuracy of 1% in the same case. Finally, we describe how the covariance matrix of the PolSpice C_ℓ estimator can be calculated from the pseudo- C_ℓ covariance matrix.

While in this work, we focused on the specific case of the SPT-3G CMB experiment, many considerations can be applied more broadly to any probe relying on the calculation of power spectra and covariance matrices. We also highlight that in this work we only consider the signal-signal part of the covariance matrix. Analytical approaches to modeling the noise contribution have already been developed and used in CMB experiments, such as *Planck*; see Appendix C.1 of [Planck Collaboration XI \(2016\)](#). They rely on assumptions on the noise properties, such as isotropy or whiteness. When these assumptions do not hold, the noise-noise and signal-noise components can be obtained directly from the data, as was done in the SPT-3G analysis ([Lueker et al. 2010](#); [Dutcher et al. 2021](#)). The integration of the noise part in our exact formalism and new approximation is beyond the scope of this paper.

The paper is organized as follows. In Sect. 2 we introduce the pseudo-power spectrum estimator and its covariance. In Sect. 3 we perform the exact calculation of the covariance matrix. In Sect. 4 we describe the existing approximations for the calculation of the covariance matrix, and we test their accuracy against the exact computation. Section 5 presents our new approximation, which is more accurate. Section 6 describes how the covariance matrix of the PolSpice estimator can be calculated. We conclude in Sect. 7, and some detailed calculations are given in the appendices.

2. Covariance of the pseudo-power spectrum

2.1. Pseudo-power spectrum

Cosmic microwave background anisotropies in intensity and polarization can be described as maps of Stokes parameters $T(\hat{n})$, $Q(\hat{n})$, $U(\hat{n})$ for a direction \hat{n} of the sky. They are Gaussian random fields, fully characterized by their angular power spectra (C_ℓ^{TT} , C_ℓ^{EE} , C_ℓ^{BB} , C_ℓ^{TE}), which are the variances of

the harmonic coefficients $a_{\ell m}^{\text{T}}$, $a_{\ell m}^{\text{E}}$, $a_{\ell m}^{\text{B}}$ obtained by spherical harmonic decomposition of the maps. Cosmological models allow the computation of the expectation of the different power spectra in an ideal full-sky case. However, data only ever cover a part of the sky. We describe the partial coverage with the weight map $W(\hat{n})$. The power spectrum of masked maps, labeled $\tilde{C}_\ell^{\text{XY}}$, is usually called throughout pseudo-power spectrum. Its expression for temperature is given in Eq. (A.6). It can be computed from the masked harmonic coefficients $\tilde{a}_{\ell m}^{\text{X}}$, which are directly related to the unmasked ones, $a_{\ell m}^{\text{X}}$, by the application of the mode-coupling kernels ${}_s I_{\ell m \ell' m'}[W]$. In the case of temperature, we write

$$\tilde{a}_{\ell m}^{\text{T}} = \sum_{\ell' m'} a_{\ell' m'}^{\text{T}} {}_0 I_{\ell m \ell' m'}[W], \quad (1)$$

where we have defined the mode coupling kernels³

$${}_s I_{\ell m \ell' m'}[W] \equiv \int d\hat{u} {}_s Y_{\ell m}(\hat{u}) W(\hat{u}) {}_s Y_{\ell' m'}^*(\hat{u}). \quad (2)$$

These coupling kernels are an important component in the following discussions⁴. In the full-sky case, the orthonormality properties of spin-weighted spherical harmonics ensure that ${}_s I_{\ell m \ell' m'}[1] = \delta_{\ell \ell'} \delta_{m m'}$. We recall in Appendix A some summation properties of products of coupling matrices that appear in the computation of pseudo-power spectra and their covariance. In particular, they are related to the symmetric coupling kernel acting on a power spectrum \mathcal{A} , labeled $\Xi[\mathcal{A}]$, with

$$\Xi_{\ell \ell'}^{s s'}[\mathcal{A}] \equiv \sum_L \frac{2L+1}{4\pi} \mathcal{A}_L \begin{pmatrix} \ell & \ell' & L \\ s & -s & 0 \end{pmatrix} \begin{pmatrix} \ell & \ell' & L \\ s' & -s' & 0 \end{pmatrix}. \quad (3)$$

This operator, introduced in [Efstathiou \(2004\)](#), can also be seen as acting on a map A with power spectrum \mathcal{A}_ℓ . In the following, we use the notation $\Xi[A] \equiv \Xi[\mathcal{A}]$. We recall that the average of the pseudo-spectrum is related to the underlying power spectrum by the application of the asymmetric coupling kernel computed for the mask W , also known as the MASTER mode-coupling matrix M . In the case of temperature, we have

$$\langle \tilde{C}_\ell^{\text{TT}} \rangle = \sum_{\ell'} {}_0 M_{\ell \ell'}[W] C_{\ell'}^{\text{TT}} \quad (4)$$

$${}_0 M_{\ell \ell'}[W] \equiv (2\ell' + 1) \Xi_{\ell \ell'}^{00}[W]. \quad (5)$$

In this work, without loss of generality, we develop the computations for the intensity case (i.e., $s = s' = 0$), the polarization case being similar. We also assume that a single mask is used for both temperature and polarization. When required, we highlight the differences between the temperature and polarization cases and give insight into the importance of the single-mask assumption.

2.2. Covariance

Estimating the covariance of the measured power-spectrum is crucial to assess the agreement between data and model predictions and to constrain cosmological parameters from CMB maps. As discussed in [Hivon et al. \(2002\)](#) and demonstrated in Appendix A, masking breaks the statistical isotropy and induces

³ The complex conjugate is denoted with a star.

⁴ This coupling matrix is often denoted K in the literature, such as in [Hivon et al. \(2002\)](#). In this work, we modified the notation for it to be consistent with the notation of Sect. 6.

correlations between the modes of the pseudo-spectrum. The details of the derivation of the analytical expression of the pseudo-spectrum covariance can be found in Appendix B. We give here the expression in terms of the coupling matrices ${}_0I$ and the true underlying intensity power spectrum C_ℓ , for the temperature case,

$$\begin{aligned} \tilde{\Sigma}_{\ell\ell'} &\equiv \text{cov}(\tilde{C}_\ell, \tilde{C}_{\ell'}), \\ &= \frac{2}{(2\ell+1)(2\ell'+1)} \sum_{mm'} \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} C_{\ell_1} C_{\ell_2} \\ &\quad \times {}_0I_{\ell m \ell_1 m_1}[W] {}_0I_{\ell' m' \ell_1 m_1}^*[W] {}_0I_{\ell m \ell_2 m_2}^*[W] {}_0I_{\ell' m' \ell_2 m_2}[W]. \end{aligned} \quad (6)$$

As shown in Eq. (1), the mode-coupling coefficient ${}_0I$ kernels relate the underlying harmonic coefficients to the harmonic coefficients measured on the sky through the mask. In the analytical expression of the covariance, they represent the coupling between modes due to partial sky coverage. An expression similar to Eq. (6) can be written for polarization, using spin-2 spherical harmonics, that is, $s, s' = \pm 2$. These expressions mix the EE and BB power spectra.

The expression in Eq. (6) involves several convolutions, and its evaluation is computationally expensive. The full computation scales as $\mathcal{O}(\ell_{\max}^6)$, ℓ_{\max} being the largest multipole, making the exact computation of this covariance a daunting task given the currently available computation power. We have developed an algorithm that allows the computation of the covariance matrix at low multipoles with a gain of an order of magnitude in computational time. We discuss this result in Sect. 3.

This approach was previously unavailable, therefore existing work has relied on approximations of Eq. (6). In Sect. 4 we present different approximations that have been proposed in previous works, and we then validate them against our full computation. This validation was performed for a small survey footprint, where spectral modes are highly correlated. These correlations can challenge the assumptions made in the different approximations. Throughout this work, we use a test-case inspired by SPT-3G. The footprint of the first year of the survey presented in Dutcher et al. (2021) covers roughly 4% of the sky and is displayed in Fig. 1 along with the mask power spectrum \mathcal{W}_ℓ . We apodized the mask with a Gaussian window function of 30 arcmin full width at half maximum, using an algorithm similar to the one used in *Planck* (Planck Collaboration VI 2020). We also show in Fig. 1 the power spectrum of one of the masks used in the *Planck* cosmological analysis, which covers a much larger patch of the sky, around 70% before apodization. The precision of the standard approximation of the covariance was validated in the latter case, but it needs to be assessed for a smaller survey area.

3. Exact computation

An exact calculation of the pseudo-power spectrum covariance matrix can be obtained by integrating Eq. (6). We propose an algorithm that performs the computation in $\mathcal{O}(\ell_{\max}^5)$, typically gaining a thousandfold speed-up compared to the direct implementation in $\mathcal{O}(\ell_{\max}^6)$. This is achieved with the fast harmonic transform tools implemented in the HEALPix library⁵. It enables the exact computation of the covariance matrix, albeit on a limited range of multipoles. In this work, we have computed the full covariance up to $\ell_{\max, \text{ex}} \equiv 1000$, and calculated a few rows of the matrix at $\ell > \ell_{\max, \text{ex}}$. This allows the direct comparison of

⁵ <https://healpix.sourceforge.io>

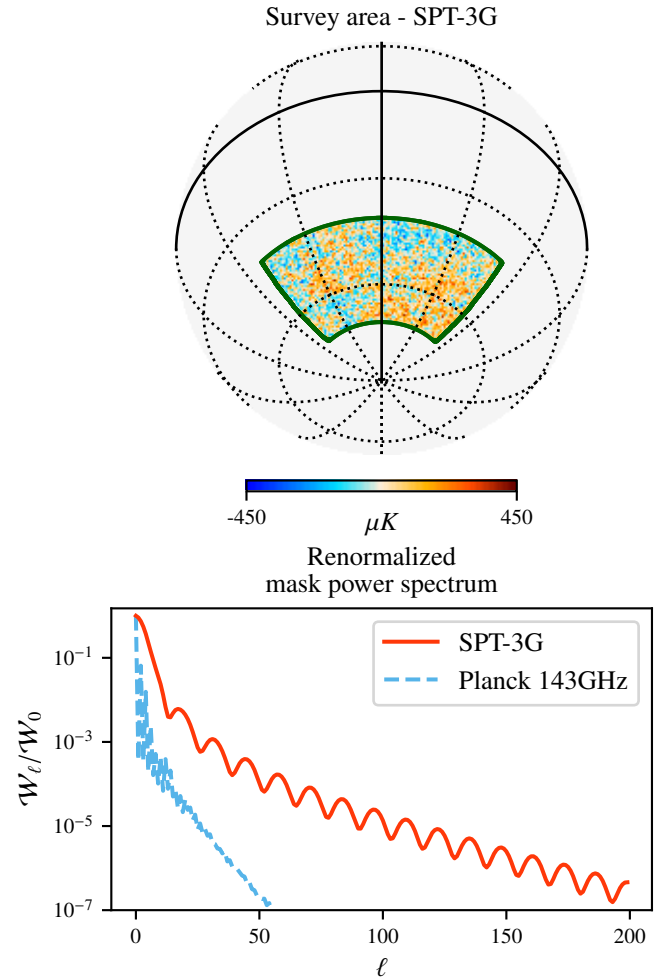


Fig. 1. Survey area and the mask power spectrum. *Top panel:* CMB temperature anisotropies on the SPT-3G patch in galactic coordinates. The dark green line delimits the survey footprint. The vertical and horizontal bold black lines are the zero-longitude and zero-latitude coordinates, respectively. The SPT-3G patch covers roughly 4% of the sky. *Bottom panel:* mask power spectra as defined in Eq. (A.4) for SPT-3G and for the 143 GHz map used in the *Planck* cosmological analysis, which covers around 70% of the sky. The spectra have been renormalized by their first value for comparison purposes. Masks corresponding to small sky fractions, such as that of SPT-3G, have a shallower power spectrum than large ones. We emphasize that the mask used here does not include point-source masking.

the various analytic covariance approximation formulae with the exact calculation.

In the following, we describe the algorithm we developed to perform this computation. We validate it with Monte-Carlo estimates of the covariance for the reference SPT-3G survey.

3.1. Algorithm description

We focus on the computation of a given row of the covariance matrix $\tilde{\Sigma}$. This allows us either to compute a full covariance matrix at low multipoles or to test our approximations on a selection of rows. We first start by describing the computation of the covariance of the intensity spectrum. A diagrammatic implementation of our calculation is presented in Fig. 2. For the polarization spectra, the calculation follows a similar pattern, and the difference between the two cases are discussed in the next section.

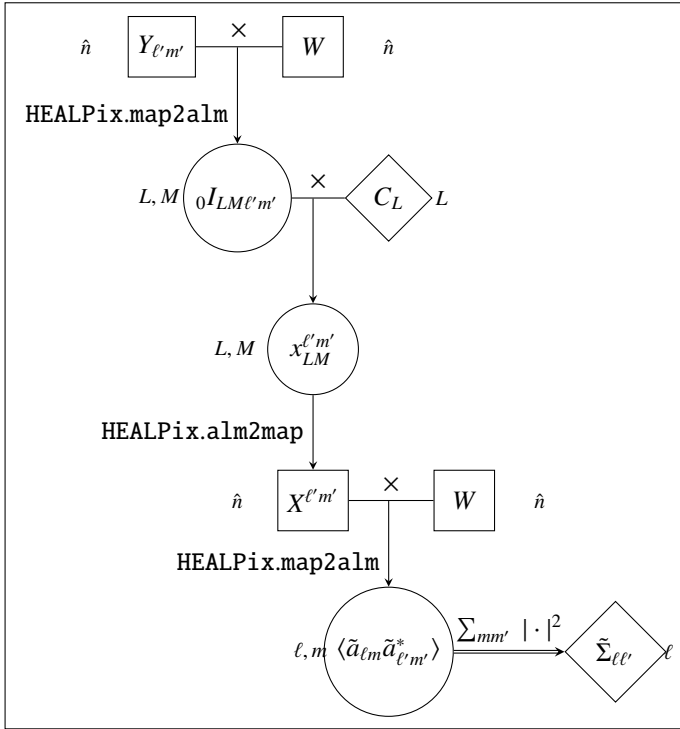


Fig. 2. Algorithm used to compute one row of the covariance $\tilde{\Sigma}_{\ell\ell'}$ using the HEALPix tools for a fixed ℓ' and varying ℓ . Square, diamond, or circle boxes are arrays representing maps, power spectra, or spherical harmonic coefficients, respectively. Operations are symbolized by arrows and described alongside. The indices of the arrays are indicated on the side of the corresponding boxes. For example, near the top of the diagram, the HEALPIX.map2alm operation applied on the product $Y_{\ell'm'}W$ produces the array ${}_0I_{LM\ell'm'}$ with indices L, M . At the bottom of the diagram, the operations before the final summation produce the array $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle$ with indices ℓ, m for a fixed ℓ', m' pair. This part of the algorithm scales as $O(\ell'^3)$ because this is the scaling of HEALPix operations (which we applied three times) when the resolution of the map is comparable to the maximum multipole index considered. The last operation, summing over the indices m, m' , requires repeating the precedent steps for all $m' \in [-\ell', \ell']$, thus repeating them $2\ell' + 1$ times. Therefore, the final complexity for producing a single row with fixed multipole index ℓ' is $O(\ell'^4)$. Computing the covariance matrix for all ℓ' then increases the computational time to $O(\ell'^5)$.

We derive in Appendix B the expression of the covariance of the pseudo-spectrum that leads to Eq. (6). In particular, we express the covariance as a sum over m and m' of the square of the correlation $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle$,

$$\tilde{\Sigma}_{\ell\ell'} = \frac{2}{(2\ell + 1)(2\ell' + 1)} \sum_{mm'} |\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle|^2. \quad (7)$$

The harmonic coefficients correlation can be written as

$$\begin{aligned} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle &= \sum_{LM} C_L {}_0I_{\ell m LM} {}_0I_{\ell'm' LM}^* \\ &= \int d\hat{u} {}_0Y_{\ell m}(\hat{u})W(\hat{u}) \sum_{LM} \{C_L {}_0I_{\ell'm' LM}^*\} {}_0Y_{LM}^*(\hat{u}), \end{aligned} \quad (8)$$

where we have used Eq. (2) to expand one of the ${}_0I$ kernels, reorganized the equation, and used the fact that the power spectrum C_L and the mask W are real quantities (we also dropped the explicit W dependences of the kernel to simplify notations). For fixed ℓ' and m' , the rightmost part of the equation can be seen

as the complex conjugate of the backward spherical harmonic transform of a set of spherical harmonic coefficients into a map $X^{\ell'm'}$, defined as

$$X^{\ell'm'}(\hat{u}) \equiv \sum_{LM} x_{LM}^{\ell'm'} Y_{LM}(\hat{u}), \quad (9)$$

where we defined the spherical harmonic coefficients with

$$x_{LM}^{\ell'm'} \equiv C_L {}_0I_{\ell'm' LM}. \quad (10)$$

Here, we emphasize that the map $X^{\ell'm'}$ is a complex map, thus it needs special care when it is decomposed into harmonic coefficients. With these definitions, Eq. (8) reduces to

$$\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle = \int d\hat{u} {}_0Y_{\ell m}(\hat{u})W(\hat{u})X^{\ell'm'*}(\hat{u}), \quad (11)$$

where we recognize the forward harmonic transform of the map $X^{\ell'm'}$, masked by W . Thus, a spherical harmonic transform of a masked map can produce the correlation $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle$ for all ℓ, m and a fixed pair of ℓ', m' . As we discussed, this $X^{\ell'm'}$ map is defined by a set of spherical harmonic coefficients whose expression is given in Eq. (10).

The computation of the $x_{LM}^{\ell'm'}$ coefficients requires the evaluation of the ${}_0I$ kernel. When Eq. (2) is used for a fixed ℓ', m' , the ${}_0I_{LM\ell'm'}$ kernel can be computed as the spherical harmonic transform of a masked ${}_0Y_{\ell'm'}$ map for all the L, M indices. When everything is added, we see that for a choice of ℓ', m' , the computation of $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle$ for all ℓ, m reduces to two forward and one backward spherical harmonic transforms, as summarized in Fig. 2.

In practice, these decompositions can be performed with HEALPix, which takes advantage of a specific pixelation scheme to make the computation more efficient. This is where the gain announced at the beginning of this section comes from, and it allows us to implement the exact computation. HEALPix decompositions typically scale as $O(\ell'^3)$ ⁶. We repeated the decompositions resulting in $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell'm'}^* \rangle$ for all $m' \in [-\ell', \ell']$ to perform the summation in Eq. (7). Thus, the computation of a single row $\tilde{\Sigma}_{\ell\ell'}$ for all ℓ and fixed ℓ' scales as $O(\ell'^4)$. Finally, the computation of a full covariance matrix for all ℓ' scales as $O(\ell'^5_{\max,ex})$.

Additional optimizations can be implemented in the algorithm by degrading maps and running the algorithm at a lower HEALPix resolution, n_{side} , for small multipoles. HEALPix computations are precise up to $\ell \sim 2n_{\text{side}}$, hence choosing a map resolution on the order of the multipole is sufficient to precisely compute the close-to-diagonal elements of the covariance. This operation requires a degraded version of the mask, which must be computed while avoiding aliasing from small-scale features. This can be done by implementing a hard cutoff of the mask harmonic coefficients before degrading its resolution. This allowed us to compute the exact covariance up to multipole $\ell_{\max,ex} = 1000$. The algorithm requires 300 h of CPU-time to compute a row of the intensity (TTTT) and polarization (EEEE) matrices at multipole $\ell = 950$ with map resolution $n_{\text{side}} = 1024$. It is also well suited to a potential GPU implementation, which could lead to more speed-ups.

⁶ Details about the computation scaling of HEALPix can be found on the website <https://healpix.sourceforge.io> or in Gorski et al. (2005).

3.2. Polarization

The polarized case is very similar to the intensity case detailed in the previous subsection. We only describe the EEEE case, which gives a general template to the other polarization and temperature×polarization cases.

The polarized version of Eq. (8) is given by Eq. (6) of Challinor & Chon (2005), which reads

$$\langle \tilde{a}_{\ell m}^E \tilde{a}_{\ell' m'}^{E*} \rangle = \sum_{LM} [C_L^{EE} + I_{\ell m LM} + I_{\ell' m' LM}^* + C_L^{BB} - I_{\ell m LM} - I_{\ell' m' LM}^*], \quad (12)$$

where we defined the Hermitian coupling coefficients

$$\pm I_{\ell m LM} = \frac{1}{2} (+2I_{\ell m LM} \pm -2I_{\ell m LM}), \quad (13)$$

with the spin-weighted coupling coefficients $\pm I$ defined as for intensity, see Eq. (2). Reordering the terms and repeating the same operations as in the previous section, the final harmonic coefficient correlation can be seen as the masked forward spherical harmonic decomposition of two maps $Z_1^{\ell' m'}$, $Z_2^{\ell' m'}$,

$$\langle \tilde{a}_{\ell m}^E \tilde{a}_{\ell' m'}^{E*} \rangle = \frac{1}{2} \left[\int d\hat{u} W(\hat{u}) Z_1^{\ell' m'}(\hat{u}) (+2Y_{\ell m}^* + -2Y_{\ell m}^*)(\hat{u}) - i \int d\hat{u} W(\hat{u}) Z_2^{\ell' m'}(\hat{u}) (+2Y_{\ell m}^* - -2Y_{\ell m}^*)(\hat{u}) \right]. \quad (14)$$

The maps $Z_1^{\ell' m'}$, $Z_2^{\ell' m'}$ are obtained using a backward spherical harmonic decomposition of the coefficients $x_{LM}^{\ell' m'; E, B}$,

$$(Z_1^{\ell' m'} - iZ_2^{\ell' m'}) (\hat{u}) \equiv \sum_{LM} \frac{x_{LM}^{\ell' m'; E} + x_{LM}^{\ell' m'; B}}{2} +2Y_{LM}(\hat{u}), \quad (15)$$

$$(Z_1^{\ell' m'} + iZ_2^{\ell' m'}) (\hat{u}) \equiv \sum_{LM} \frac{x_{LM}^{\ell' m'; E} - x_{LM}^{\ell' m'; B}}{2} -2Y_{LM}(\hat{u}). \quad (16)$$

The set of harmonic coefficients $x_{LM}^{E, B}$ is defined similarly to the temperature case (Eq. (10)), and it was obtained by filtering the coefficients computed with a masked forward harmonic decomposition of the spin-2 spherical harmonics $\pm 2Y_{\ell' m'}$,

$$\begin{cases} x_{LM}^{\ell' m'; E} = C_L^{EE} + I_{\ell' m' LM}, \\ x_{LM}^{\ell' m'; B} = C_L^{BB} - I_{\ell' m' LM}. \end{cases} \quad (17)$$

This algorithm can be extended for any combination of spectra for the other polarization cases, including the cross-correlation between temperature and polarization that we do not treat here.

3.3. Validation on simulations

We compared the results of our implementation of the exact computation with a MC estimate of the covariance, obtained with N_{sim} simulations. The MC covariance terms are expected to be Wishart distributed with N_{sim} degrees of freedom, as explained in Lueker et al. (2010). We can estimate their variance to be

$$\left\langle \left(\tilde{\Sigma}_{\ell\ell'}^{\text{sim}} - \langle \tilde{\Sigma}_{\ell\ell'}^{\text{sim}} \rangle \right)^2 \right\rangle = \frac{\tilde{\Sigma}_{\ell\ell}^2 + \tilde{\Sigma}_{\ell\ell'} \tilde{\Sigma}_{\ell'\ell'}}{N_{\text{sim}}}. \quad (18)$$




Fig. 3. Relative difference of diagonals $\tilde{\Sigma}_{\ell\ell} / \tilde{\Sigma}_{\ell\ell}^{\text{sim}} - 1$ for temperature TTTT (top) and polarization EEEE (bottom). In red we plot the relative differences every 25 multipoles until $\ell = 1500$ and a few well-chosen ones (at the locations of peaks and troughs of the spectra) up to $\ell = 2000$. In gray, the same quantity is plotted for all multipoles for $\ell \in [\ell_{\text{cut}} = 200, \ell_{\text{max,ex}} = 1000]$. We are able to compute the covariance exactly only for a limited number of rows, and it is computationally cheaper for lower multipoles, justifying our choice of full calculation at $\ell < \ell_{\text{max,ex}}$ and partial calculation for larger multipoles. This plot shows the agreement between the two approaches and validates our exact calculation.

$N_{\text{sim}} = 10\,000$ allows us to reach a percent-level accuracy on the diagonal. This is the number of realizations that we use for the MC covariance. For this validation, we used the mask shown in Fig. 1. In this idealized setting, we did not include a point source mask.

We performed an exact computation of the TTTT and EEEE covariance up to $\ell_{\text{max,ex}} = 1000$, using our algorithm and degrading the mask to smaller resolutions. Furthermore, we computed the rows every 25 multipoles of the matrix up to $\ell_{\text{max}} = 1500$, as well as at a few well-chosen multipoles that correspond to peaks and troughs of the spectra up to $\ell_{\text{max}} = 2000$.

We first focus on the diagonal of the covariance. Figure 3 presents the comparison between the exact computation of the diagonal, obtained by selecting the corresponding value in the rows we computed, and the MC evaluation. The two agree within the MC noise expected for N_{sim} .

For the off-diagonal terms, we show in Fig. 4 a few rows of the exact and MC covariance. The rows agree within the MC noise. The correlation between the modes falls to the percent level within a distance $|\ell - \ell'| \sim 25$ bands around the diagonal. The correlation matrix is defined as the covariance renormalized by its diagonal,

$$\sigma^{\ell\ell'} \equiv \frac{\Sigma_{\ell\ell'}}{\sqrt{\Sigma_{\ell\ell} \Sigma_{\ell'\ell'}}}. \quad (19)$$

We display the exact and MC correlation matrix in the same multipole range in Fig. 5. Our tests demonstrate that our implementation of the exact computation of the pseudo-spectrum

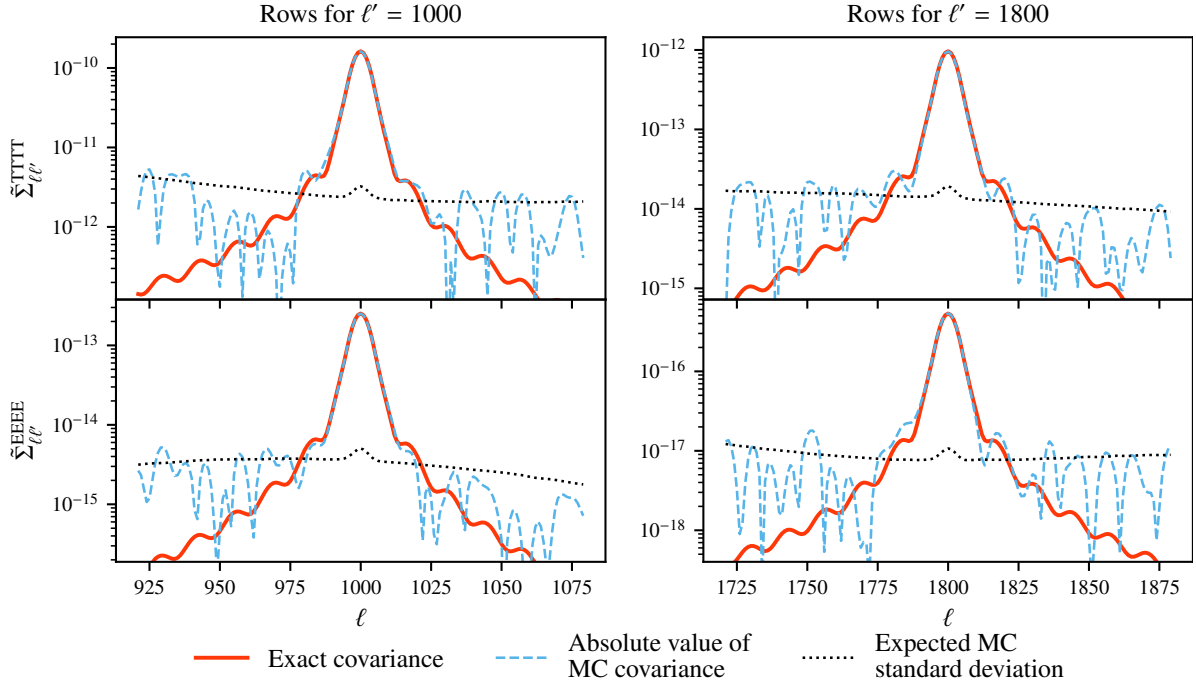


Fig. 4. Rows of the exact covariance (solid red) and of the simulated one (dashed blue) for multipole $\ell' = 1000$ (left-hand side) and $\ell' = 1800$ (right-hand side). The dotted black line shows the expected MC uncertainty on the simulated covariance from Eq. (18), using $N_{\text{sim}} = 10\,000$ realizations. The two approaches agree very well. The MC standard deviation is as large as the covariance off-diagonal terms at $|\ell - \ell'| \sim 25$.

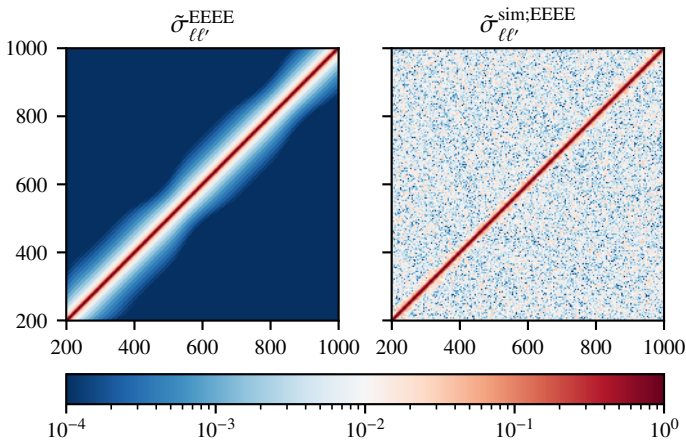


Fig. 5. EEEE correlation matrix (see Eq. (19)) obtained from the exact (left) and MC (right) calculations of the covariance matrix. The exact computation allows us to have the full correlation matrix, but the MC approach is limited by numerical noise. All terms below 10^{-4} are plotted in dark blue. Results for TTTT are comparable.

covariance is correct, at least at the level of accuracy that can be reached by MC estimation.

While massive MC estimates as we performed here in this idealized case can produce accurate off-diagonal term estimates, performing numerous MC estimates can be more challenging in the case of a realistic experiment. It requires regularization approaches, where the number of possible simulations is limited by the computational cost of mock-observations.

We also stress that our covariance matrix cannot be directly compared to the one used in Dutcher et al. (2021) because the presence of a point source mask (which we did not include in our simple example), the complications brought about by introducing realistic noise and scanning strategy, and projection effects

(the analysis in Dutcher et al. (2021) is performed using a flat-sky approach) can all yield different levels of correlations between the modes.

4. Existing approximations and their accuracy on a small patch of the sky

Our algorithm allows us to obtain the exact covariance only for $\ell < 1000$ or for a few rows at higher ℓ 's due to the expensive computing resources required. The usual analytical approach consists of using approximations of Eq. (6) to decrease the computational cost. In this section, we introduce a new framework to express the approximations of the covariance matrix and use it to list the different methods proposed in the literature. Then, we test and discuss their accuracy against our exact computation.

4.1. General framework

Before discussing the approximations of the covariance, we define a few quantities that help relate the various approximations to each other. We rewrite Eq. (6) as

$$\tilde{\Sigma}_{\ell\ell'} = \frac{2}{(2\ell + 1)(2\ell' + 1)} \sum_{\ell_1 \ell_2} C_{\ell_1} \Theta_{\ell\ell'}^{\ell_1 \ell_2} [W] C_{\ell_2}, \quad (20)$$

introducing $\Theta_{\ell\ell'}^{\ell_1 \ell_2} [W]$. The covariance coupling kernel, defined as the sum over the multipole orders (m, m', \dots) of the coupling coefficients, reads

$$\Theta_{\ell\ell'}^{\ell_1 \ell_2} [W] \equiv \sum_{mm'm'm_2} ({}_0I_{\ell m \ell_1 m_1} {}_0I_{\ell' m' \ell_1 m_1}^* {}_0I_{\ell' m' \ell_2 m_2} {}_0I_{\ell m \ell_2 m_2}^*) [W]. \quad (21)$$

The covariance coupling kernel Θ represents the coupling between the modes of the theoretical underlying power spectrum

C_{ℓ_i} for $i = 1, 2$ depending on the index of the pseudo-covariance (ℓ, ℓ') . We chose to show the two indices related to the covariance (ℓ, ℓ') as subscripts and the two summing indices (ℓ_1, ℓ_2) as superscripts. We considered a single-mask temperature case, for which the coupling kernel is symmetric with respect to the exchange of multipole indices $\ell \leftrightarrow \ell'$ or $\ell_1 \leftrightarrow \ell_2$. In the following, we write our results in this case for the sake of simplicity, but they are valid regardless of the choice of single or multiple masks. In the case of spectra obtained from maps with different masks, or in the case of cross-spectra, the kernel is not symmetric. While the results of this work apply to both cases, considering multiple masks increases computing cost.

Using the completeness relations of spherical harmonics, given in Eqs. (A.10) and (A.11), we can write

$$\sum_{\ell_1 \ell_2} \Theta_{\ell \ell'}^{\ell_1 \ell_2} [W] = (2\ell + 1)(2\ell' + 1) \Xi_{\ell \ell'}^{00} [W^2]. \quad (22)$$

We can now define the reduced covariance coupling kernel as

$$\bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] \equiv \frac{\Theta_{\ell \ell'}^{\ell_1 \ell_2} [W]}{(2\ell + 1)(2\ell' + 1) \Xi_{\ell \ell'}^{00} [W^2]}, \quad (23)$$

for which

$$\sum_{\ell_1 \ell_2} \bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] = 1. \quad (24)$$

With these notations, we can rewrite the covariance as

$$\tilde{\Sigma}_{\ell \ell'} = 2 \Xi_{\ell \ell'}^{00} [W^2] \sum_{\ell_1 \ell_2} C_{\ell_1} \bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] C_{\ell_2}. \quad (25)$$

The symmetric mode-coupling kernel $\Xi[W^2]$ provides the purely geometric coupling due to sky masking and is common to all approximations of the covariance. It only depends on the power spectrum of the squared mask, and it is easy to be convinced that up to a normalization, it corresponds to the exact covariance in the case of a constant power spectrum. Its computation scales as $\mathcal{O}(\ell_{\max}^3)$. This could be improved by noting, as was done by Louis et al. (2020), that at small enough scales, $\Xi[W^2]$ is close to a Toeplitz matrix, allowing us to further reduce the scaling to $\mathcal{O}(\ell_{\max}^2)$ for a wide range of modes. The sum on the right-hand side of Eq. (25) describes the contribution of the signal power spectrum modulated by the kernel $\bar{\Theta}$, which depends on the mask. This is the sum that all approximations try to simplify, replacing the kernel $\bar{\Theta}$ with a simpler ansatz. In the following, we describe all approximations in terms of this redefinition of the covariance matrix. Since the reduced covariance coupling kernel is normalized, every approximation formulated in this formalism yields the exact covariance for a constant underlying power spectrum $C_\ell = N$.

4.2. Approximations

4.2.1. Narrow-kernel approximation

Based on the observation that the coupling coefficients ${}_0I$ in Eq. (1) are narrow and peak at their first multipole indices ℓ or ℓ' , Efstathiou (2004) introduced the following approximation of Eq. (25), taking the convolving spectra C_{ℓ_i} , $i = 1, 2$ out of the sum, and replacing them by the power spectrum evaluated at the first multipole index of the coupling coefficients (i.e., the covariance indices of $\bar{\Theta}$), C_ℓ or $C_{\ell'}$. Following the notation introduced

in García-García et al. (2019), we refer to this approximation of the covariance as the narrow-kernel approximation (NKA),

$$\begin{aligned} \tilde{\Sigma}_{\ell \ell'} &\approx 2C_\ell C_{\ell'} \Xi_{\ell \ell'}^{00} [W^2] \sum_{\ell_1 \ell_2} \bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] \\ &= 2C_\ell C_{\ell'} \Xi_{\ell \ell'}^{00} [W^2] \equiv \tilde{\Sigma}_{\ell \ell'}^{\text{NKA}}. \end{aligned} \quad (26)$$

In terms of the reduced covariance coupling kernel, the NKA uses

$$\bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] \approx \bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2; \text{NKA}} [W] \equiv \frac{\delta_{\ell \ell_1} \delta_{\ell' \ell_2} + \delta_{\ell' \ell_1} \delta_{\ell \ell_2}}{2}. \quad (27)$$

The approximation is exact for the full sky. It provides an accurate estimator whenever the underlying power spectrum C_ℓ varies slowly as a function of ℓ compared to the typical size of the operators ${}_0I$. This condition is fulfilled when the amplitude of the mask power spectrum drops quickly with multipole ℓ , which is the case for large sky fractions observed with a mask that contains no small-scale features. This is shown in Fig. 1, for example, where we plot the power spectrum of one of the masks used in the *Planck* analysis. In this case, the above approximation holds for multipoles much larger than those for which the mask spectrum contains power.

The NKA was first introduced in intensity by Efstathiou (2004) and extended to polarization in Challinor & Chon (2005). As in the temperature case, the approximated covariances in polarization are expressed as a function of the polarization spectra EE and BB and the symmetric coupling kernels $\Xi_{\ell, \ell'}^{\pm 2, \pm 2}$. The expressions of the approximated polarization covariances mix EE and BB due to leakage that appears because the sky is masked; see Eqs. (25)–(27) of Challinor & Chon (2005).

The NKA has been widely used, for instance, in the *Planck* cosmological analysis, which masked only small portions of the full sky; see Planck Collaboration XI (2016). However, it has never been thoroughly tested on small sky fractions. As shown in Fig. 1, the mask power spectrum in the case of the small survey footprint of SPT-3G drops much more slowly than the large *Planck* one. From this observation, we expect the mode-coupling kernels ${}_sI$ to be wider, as can be deduced from Eq. (2). As a result, the theoretical underlying spectrum C_ℓ might not be treated as constant compared to the covariance coupling kernels in the sums of Eq. (25), and SPT-3G may be outside the regime of validity of the NKA assumption. This is tested at the end of this section. We now list some proposed improvements to the NKA.

4.2.2. Friedrich approximation

A straightforward extension of the NKA has been proposed in Friedrich et al. (2021). It is based on the observation that the reduced covariance-coupling kernel $\bar{\Theta}$ has four maxima at $[\ell = \ell_1, \ell' = \ell_1]$, $[\ell = \ell_1, \ell' = \ell_2]$, $[\ell = \ell_2, \ell' = \ell_1]$, and $[\ell = \ell_2, \ell' = \ell_2]$. This suggests the following form of the reduced covariance coupling matrix:

$$\bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2} [W] \approx \bar{\Theta}_{\ell \ell'}^{\ell_1 \ell_2; \text{FRI}} [W] \equiv \frac{\delta_{\ell \ell_1} + \delta_{\ell' \ell_1}}{2} \frac{\delta_{\ell \ell_2} + \delta_{\ell' \ell_2}}{2}. \quad (28)$$

Thus, the approximated covariance is

$$\tilde{\Sigma}_{\ell \ell'}^{\text{FRI}} \equiv 2 \Xi_{\ell \ell'}^{00} [W^2] \left(\frac{C_\ell + C_{\ell'}}{2} \right)^2. \quad (29)$$

We refer to this approximation as FRI in the rest of the article.




Fig. 6. Relative differences of binned approximations with respect to the exact binned covariance: $\tilde{\Sigma}_{bb'}^{\text{APP}}/\tilde{\Sigma}_{bb'} - 1$ for TTTT (left) and EEEE (right), with binning $\Delta\ell = 50$. In the *first row*, we plot the relative differences for the diagonal, i.e., $b = b'$, while in the *second row*, we plot those of the first off-diagonal, i.e., $b' = b + 1$. The NKA (light blue dashed), FRI (purple dashed-double-dot) and INKA (dark blue dashed-dot) approximations are accurate at the 5% level, whereas the ACC approximation (solid red) reaches the 1% level, both in intensity and polarization for multipoles larger than $\ell_{\text{cut}} = 200$. The relative differences are plotted for bins that include multipoles up to $\ell_{\text{max,ex}} = 1000$ because it is the maximum multipole for which we computed all the rows of the exact covariance. The third row displays the corresponding binned underlying renormalized spectrum TT or EE to show that the difference of the covariances lies in the peaks and troughs of the spectra.

4.2.3. Improved narrow-kernel approximation

Nicola et al. (2021) proposed an improved version of the NKA, the improved narrow-kernel approximation (INKA). In this approximation, the Dirac functions in Eq. (27) are replaced by ${}_0\bar{M}$, the renormalized MASTER mode-coupling kernel, as defined in Appendix A.3. It reads

$$\bar{\Theta}_{\ell\ell'}^{\ell_1\ell_2}[W] \approx \bar{\Theta}_{\ell\ell'}^{\ell_1\ell_2;\text{INKA}}[W] \equiv \frac{{}_0\bar{M}_{\ell\ell_1,0}\bar{M}_{\ell'\ell_2} + {}_0\bar{M}_{\ell'\ell_1,0}\bar{M}_{\ell\ell_2}}{2}. \quad (30)$$

The convolution in Eq. (6) indeed averages the power spectra C_{ℓ_i} , $i = 1, 2$ over multipoles close to ℓ and ℓ' . We can take advantage of this by replacing the convolution by a multiplication with a smoothed power spectrum. When $\bar{C} \equiv \bar{M}C$, the resulting covariance can be written as

$$\tilde{\Sigma}_{\ell\ell'}^{\text{INKA}} \equiv 2\bar{C}_{\ell}\bar{C}_{\ell'}\Xi_{\ell\ell'}^{\text{00}}[W^2]. \quad (31)$$

All the NKA, FRI, and INKA scale as $O(\ell_{\text{max}}^3)$, which are the computing resources needed to obtain the coupling kernels Ξ and \bar{M} . This is a significant improvement over the $O(\ell_{\text{max}}^5)$ scaling of our full computation. We now validate the approximations in the case of small surveys using the expensive exact computation of the covariance matrix.

4.3. Accuracy

We tested the accuracy of the NKA, FRI and INKA using the exact computation in the case of the SPT-3G small survey footprint shown in Fig. 1. For this mask, the correlations between modes are significant, as we showed in Fig. 4. In this case, it is customary to bin the individual multipoles into wider bandpowers. For this reason, we performed all of our tests on a binned version of the covariance. Given the shape of the power spectrum of the mask and the correlations that we expect from it, we adopted a $\Delta\ell = 50$ binning with $\ell(\ell + 1)/(2\pi)$ weights to flatten the dynamics of the spectra in each bin. With this bin size, we expect that most of the correlations between bandpowers are concentrated in the first off-diagonal bin. We also conservatively excluded the first $\ell_{\text{cut}} = 200$ multipoles from our analysis. They are more challenging to measure on a small survey footprint as they can suffer from leakage from the super-survey scales. We restrict our comparison to the multipoles between ℓ_{cut} and $\ell_{\text{max,ex}} = 1000$, where we have carried out the exact computation of all the matrix rows.

We present in Fig. 6 a comparison between the exact computation and the NKA, FRI, and INKA for the diagonal and first off-diagonal of the TTTT and EEEE binned covariances. We discuss the performance of our new accurate covariance coupling (ACC) approximation, also shown in the figure, in the next section. The existing approximations provide good

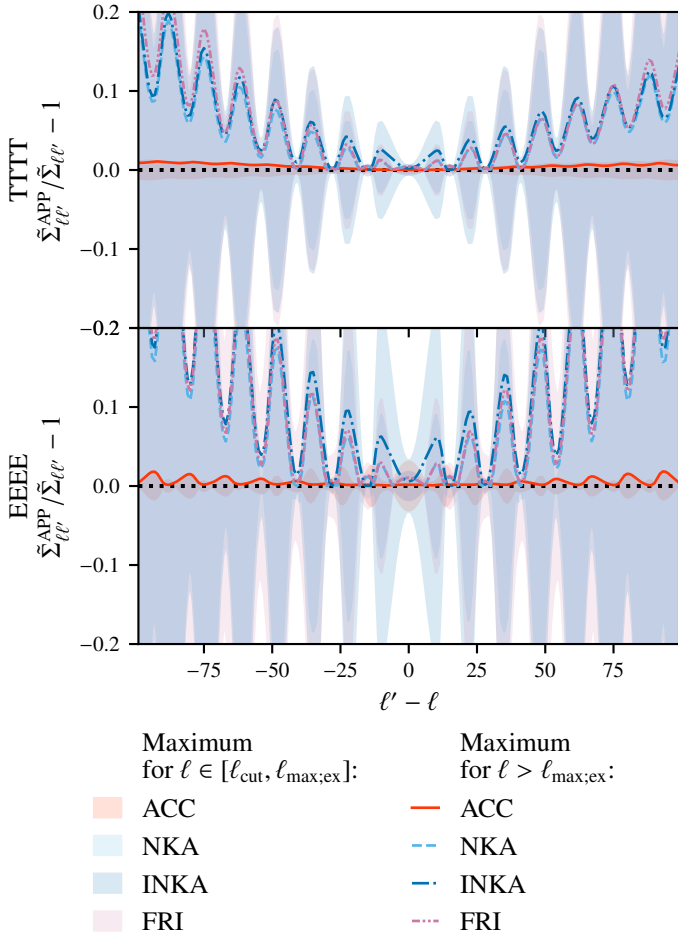


Fig. 7. Relative difference between the unbinned approximated covariance matrices compared to the exact calculation, $\tilde{\Sigma}_{\ell\ell'}^{\text{APP}}/\tilde{\Sigma}_{\ell\ell'} - 1$, as a function of $\Delta = \ell - \ell'$. We show the NKA (light blue), INKA (dark blue), and FRI (purple) approximations for TTTT (top) and EEEE (bottom). We only show the relative differences for the ℓ at which the deviation is largest, with $\ell \in [\ell_{\text{cut}} = 200, \ell_{\text{max,ex}} = 1000]$ (shaded regions) or $\ell \in [\ell_{\text{max,ex}}, 2000]$ (lines).

estimates of these elements of the covariance as they fall within 5% of the accuracy. The amplitude of the errors varies at different multipoles. Even though the FRI and INKA schemes were implemented to improve upon the simple NKA, their errors are of similar amplitude for this choice of binning. However, all approximations fail to recover the binned covariance at the percent level. On the third row, we show that the difference of the covariances lies in the peaks and troughs of the spectra. As the covariance coupling kernel acts as a symmetric convolution on the spectrum (see Eq. (25)), it is sensitive to the second-order derivative of the spectrum rather than the first. An error on the covariance coupling kernels leads to a larger relative difference in the covariance at multipoles where the curvature of the spectrum is maximum. In Sect. 5 we use the knowledge gained from the exact computation to propose an improved approximation scheme.

Because we cannot easily compute the full matrix to present binned results, we only compare some unbinned rows in Fig. 7 at higher multipoles. The shaded regions in this figure give the lowest values of the relative difference for the approximation within multipoles $\ell \in [\ell_{\text{cut}} = 200, \ell_{\text{max,ex}} = 1000]$ and $\ell' \in [\ell - 2\Delta\ell, \ell + 2\Delta\ell]$. These are the covariance terms for which we can calculate the full binned covariance; their accuracy is shown in Fig. 6.

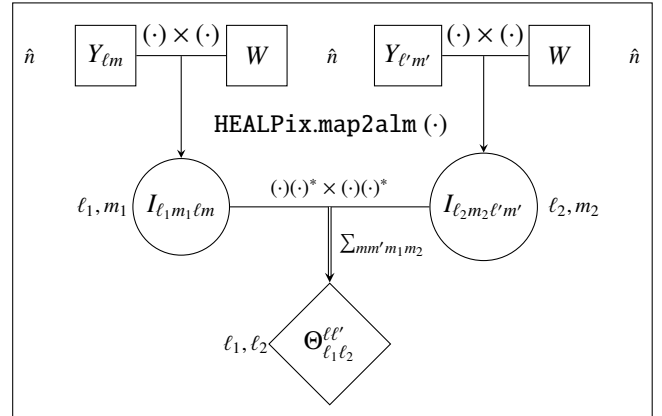


Fig. 8. Algorithm for computing the reduced covariance-coupling kernels using HEALPix tools. We use the same notation as in the diagram of Fig. 2. The HEALPix functions require $\mathcal{O}(n_{\text{side}}^3)$, with n_{side} the chosen resolution of maps W and $Y_{\ell m}$. As we choose the resolution n_{side} to be on the order of the multipoles indices ℓ, ℓ' , it is equivalent to say that they require $\mathcal{O}((\ell + \ell')^3)$. As operations are done $\mathcal{O}(\ell + \ell')$ times, the whole operation of computing $\Theta_{\ell\ell'}$ is $\mathcal{O}((\ell + \ell')^4)$. Finally, it is clear in this diagram that the computing time of this kernel is at least doubled when multiple masks are used. Different masks would be used as inputs in the first line. As a result, the coupling coefficients would need to be computed for each of the masks, as shown in Eq. (D.3).

Furthermore, the lines in Fig. 7 show the same quantity as the shaded regions, that is, the maximum relative difference, but for multipoles $\ell \in [\ell_{\text{max,ex}}, 2000]$, estimated over the sparse number of rows for which we computed the matrix exactly. The difference with the exact covariance for all approximations at large multipoles is always within the same error range as for lower multipoles. This shows that the approximations still work with the same precision at higher multipoles, both for temperature and polarization, and that the accuracy of the approximations in the unbinned case quickly falls below 20% when $\Delta = 50$.

4.4. Structure of the reduced covariance-coupling kernel

Our expression of the covariance matrix approximations in terms of the normalized coupling kernel Θ in Eq. (27) gives us a very efficient tool for examining the validity of each approximation and for better understanding their differences. We designed an algorithm to calculate this kernel exactly, similar to the one described in Sect. 3 for the exact calculation of the matrix. We show a diagram of the algorithm in Fig. 8.

The reduced covariance-coupling kernel is then displayed in Fig. 9 for the INKA, NKA, and FRI approximations compared to the exact computation for a fiducial multipole $\ell = 200$. The kernels are represented as matrices as a function of ℓ_1, ℓ_2 for different fixed choices of the indices ℓ, ℓ' . Columns show the results for different choices of $\ell' = \ell - \Delta$, with $\Delta = 0$ (i.e., the kernels for the diagonal terms of the covariance matrix, e.g., $\tilde{\Sigma}_{200,200}$), or $\Delta = 10, 50$ (i.e., the kernels for the off-diagonal terms separated by 10 or 50 multipoles, e.g., $\tilde{\Sigma}_{200,190}$). We recall that the reduced kernel is multiplied to $\mathcal{C}_{\ell_1}, \mathcal{C}_{\ell_2}$ and summed over the indices ℓ_1, ℓ_2 in Eq. (25). Hence, Fig. 9 directly shows the weight of the ℓ_1, ℓ_2 power spectra that contribute to the $\tilde{\Sigma}_{\ell\ell'}$ element of the covariance matrix.

We first focus on the kernels for the diagonal of the covariance matrix, $\Delta = 0$, shown in the first column of Fig. 9. All kernels peak at $\ell_1 = \ell_2 = \ell$, as expected. However, it is clear from the exact calculation that the kernel has a significant width




Fig. 9. Reduced covariance coupling kernels $\bar{\Theta}_{\ell\ell'}^{\ell_1\ell_2}$ for $\ell = 200$ and $\ell' = \ell + \Delta$, with $\Delta = [0, 10, 50]$ shown in the three different columns. All kernels in one column share the same color map. All of the matrices are shown as a function of ℓ_1 and ℓ_2 and are centered on $\ell_1, \ell_2 = (\ell + \ell')/2$. Each of the displayed kernels is properly normalized according to Eq. (22). The plots are restricted to the elements that have a significant value. *Top row*: approximated INKA kernels and the positions at which the delta distributions of the NKA (circles) and FRI (crosses) approximation peak. *Bottom row*: exact kernels. The comparison between the two highlights how much of the structure of the exact kernel is missed by the different approximations.

compared to the CMB power spectrum. This is more clearly shown in Fig. 10, where we plot a slice of the coupling kernels for $\ell_1 = 200$. The width of the kernel cannot be neglected compared to the slope of the CMB power spectrum. This justifies the INKA, which replaces the Dirac δ functions of NKA and FRI by renormalized mode coupling kernels, see Eq. (31). However, as shown in this figure, the INKA kernel is slightly broader than the exact calculation and its amplitude is smaller. This explains why INKA underestimates the covariance diagonal in the peaks of the power spectrum and overestimates it on the troughs, as shown in Fig. 6: it averages the underlying power spectrum in a wider range of multipoles. Conversely, the NKA/FRI kernels are much thinner than in the exact computation, and so they overestimate the diagonal in the peaks and underestimate it in the troughs of the power spectrum.

Second, we focus on the off-diagonal terms, $\Delta = 10, 50$, shown in the second and third column of Fig. 9. The difference between the exact computation and all the existing approximations is striking, and it is clear that the kernel has more structure than the simple approximated forms. For close off-diagonal terms such as $\Delta = 10$, the true kernel peaks at its central index $\ell_1 = \ell_2 = \bar{\ell} \equiv (\ell + \ell')/2$. For far off-diagonal terms such as $\Delta = 50$, there are four maxima, as predicted by the FRI approximation, which are partially missed by the INKA. Moreover, the true coupling has more dynamics and also covers negative values. Therefore the different approximations, even the INKA

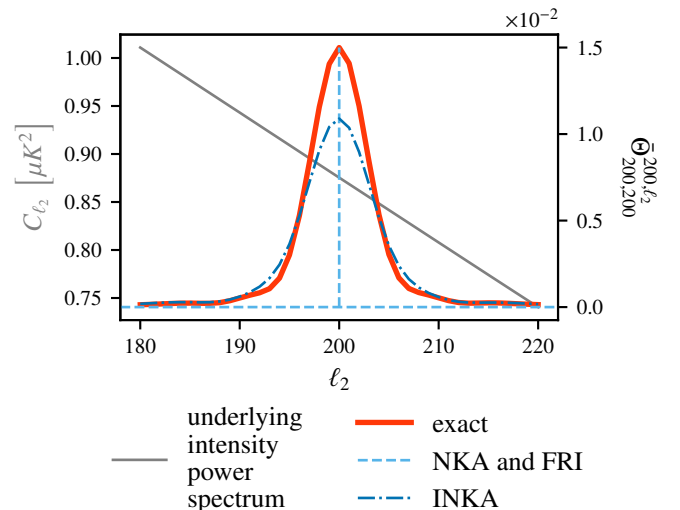


Fig. 10. Slices of the exact covariance coupling kernel (solid red) vs. the approximated NKA and FRI (dashed light blue) and INKA (dot-dashed dark blue) ones (*right scale*). We show for comparison the CMB intensity power spectrum (*left scale*, solid gray line).

approximation, fail to correctly represent the off-diagonal terms of the covariance, as observed in Fig. 7.




Fig. 11. Diagonal (*top*) and row (*bottom*) of the reduced covariance coupling kernels $\tilde{\Theta}_{\ell\ell'}^{\ell_1, \ell_2}$ for $\ell \in [150, 206, 300, 650, 750]$ and $\Delta = \ell' - \ell = 50$ as a function of $\Delta_2 = \ell_2 - \ell$. The row (*bottom*) is plotted for $\ell_1 = \ell$, i.e. $\Delta_1 = \ell_1 - \ell = 0$. The plots show that for different ℓ but the same Δ , the kernels are very similar, differing only at the 5% level. A similar result can be shown for other values of Δ . This leads us to formulate our new approximation, where we assume $\tilde{\Theta}$ to depend only on the multipole separations $\Delta, \Delta_1, \Delta_2$.

5. New approximation for the covariance

5.1. Improved approach: Approximated covariance coupling

Our ability to calculate the exact reduced covariance coupling matrix $\tilde{\Theta}$, described in Sect. 4, allows us to introduce a new approximation for the computation of the pseudo-power spectrum covariance matrix. We note that for a fixed $\Delta = \ell' - \ell$, the structure of $\tilde{\Theta}_{\ell\ell'}$ appears to be invariant. In other words, the coupling matrices contributing to the $\tilde{\Sigma}_{\ell\ell'}$ term of the covariance matrix only depend on the distance Δ from the diagonal. This is demonstrated in Fig. 11, where we plot diagonal and rows of the exact calculation of $\tilde{\Theta}_{\ell, \ell+\Delta}$ for $\Delta = 50$ and different ℓ , for $\ell_1 = \ell$. The plot reveals that the kernels are nearly identical when they are plotted as a function of $\Delta_2 = \ell_2 - \ell$. We thus infer that in general, when the $\tilde{\Theta}$ matrices are written as a function of $\Delta_1 = \ell_1 - \ell$ and $\Delta_2 = \ell_2 - \ell$, they only depend on $\Delta = \ell' - \ell$ for any ℓ and ℓ' . The difference between kernels computed at different ℓ s for same Δ is small, at the 5% percent level. We can thus assume that for any choice of multipole ℓ, λ ,

$$\forall(\ell, \lambda), \tilde{\Theta}_{\ell(\ell+\Delta)}^{(\ell+\Delta_1)(\ell+\Delta_2)} \approx \tilde{\Theta}_{\lambda(\lambda+\Delta)}^{(\lambda+\Delta_1)(\lambda+\Delta_2)}. \quad (32)$$

An analytical justification of this approximation is provided in Appendix E using the asymptotic expansion of the Wigner-3j symbols when ℓ is large.

This suggests a new approximation where the coupling kernel just has to be computed at a given fiducial ℓ for all relevant values of Δ, Δ_1 , and Δ_2 . We call this the new approximated covariance coupling method, ACC. More precisely, the ACC

kernel is given by

$$\tilde{\Theta}_{\ell\ell'}^{\ell_1, \ell_2} \approx \tilde{\Theta}_{\ell(\ell+\Delta); \text{ACC}}^{(\ell+\Delta_1)(\ell+\Delta_2)} \equiv \tilde{\Theta}_{\ell^*(\ell^*+\Delta)}^{(\ell^*+\Delta_1)(\ell^*+\Delta_2)}, \quad (33)$$

where we perform the exact and costly computation only for the $\tilde{\Theta}_{\ell^*(\ell^*+\Delta)}^{(\ell^*+\Delta_1)(\ell^*+\Delta_2)}$. We are free to choose the reference ℓ^* multipole. Because there are no significant long-range correlations in our case, we can pick a low⁷ ℓ^* and use a low n_{side} map resolution. We have to ensure, however, that ℓ^* is larger than ℓ_{cut} , the low- ℓ cutoff that was introduced to avoid issues with large-scale leakages. Close to ℓ_{cut} , the exact computation can be used. With a small mask, large-scale modes are difficult to measure and are usually excluded from the cosmological analysis. We can also restrict the range of Δ to the number of off-diagonal terms of interest in the covariance matrix. The correlation falls quickly (see Figs. 4 and 5), and in practice, we can restrict it to $|\Delta| < d_{\text{max}}$, with d_{max} being on the order of a few times the correlation length. Similarly, the kernels fall quickly in Δ_1, Δ_2 , so that we can also restrict ourselves to a small region of a similar order, and in the case of the single-mask analysis, use the symmetry around $\Delta \leftrightarrow -\Delta$ to reduce the computational cost.

While we only presented temperature coupling kernels in Fig. 9, the situation is identical in polarization, and a similar approximation can be built; see Appendix E. We used this approximation with $\ell^* = 300$, $n_{\text{side}} = 512$, and $d_{\text{max}} = 100$ to compute the ACC results in Fig. 6.

5.2. Accuracy and scaling

We validate the accuracy of our ACC approximation and compare it to the other approximations in Fig. 6. Our new approximation succeeds at estimating the covariance within an error of 1% for all multipoles larger than ℓ_{cut} in intensity and polarization, which is an improvement of a factor of ~ 4 over previous approximations. This is also shown in Figs. 7 and 12, which is just the same as Fig. 7, but focused on the EEEE ACC residuals with respect to INKA. This figure shows that the ACC approximation estimates both the diagonal and off-diagonal terms of the covariance matrix far better.

Figure 8 shows the computations needed to obtain the covariance-coupling kernels. Following the same argumentation as for the exact computation of the Sect. 3, we can show that the computation of a single kernel $\tilde{\Theta}$ scales as $O((\ell + \ell')^4)$. As a result, because we need to compute one for each diagonal index $\Delta \in [0, d_{\text{max}}]$, the final ACC approximation scales as $O(n_{\text{side}}^4 d_{\text{max}})$, where n_{side} is the map resolution chosen to compute the kernels. The computing resources needed to obtain $\tilde{\Theta}$ for all approximations are summarized in Table 1. They add up to the resources needed to compute the symmetric coupling kernel Ξ in Eq. (25). In practice, the kernel \bar{M} needed to build INKA is often already known for the sky analysis because it has the same structure as Ξ , so that the effective complexity for this approximation is $O(1)$.

5.3. Point-source mask

We did not include a point-source mask in the survey footprint. Point-source masks significantly complicate the problem as the power spectrum of the mask will have power at large multipoles, hence it will extend the correlation length. This has been an issue for all analysis thus far. Apodizing the point-source masks

⁷ In the limit where the asymptotic justification of Appendix E remains valid.

Table 1. Summary of computation methods to obtain the pseudo-power spectrum covariance matrix.

Method	Equation numbers	$\bar{\Theta}_{\ell\ell'}^{\ell_1\ell_2}$	Precision	Complexity
Exact (this work)	Eq. (25)	$\bar{\Theta}_{\ell\ell'}^{\ell_1\ell_2}$ computed $\forall \ell, \ell', \ell_1, \ell_2$	N/A	$O(\ell_{\max}^5)$ (Using HEALPix pixelation)
NKA	Eq. (26)	$(\delta_{\ell\ell_1}\delta_{\ell'\ell_2} + \delta_{\ell'\ell_1}\delta_{\ell\ell_2})/2$	4%	$O(1)$
FRI	Eq. (29)	$(\delta_{\ell\ell_1} + \delta_{\ell'\ell_1})(\delta_{\ell\ell_2} + \delta_{\ell'\ell_2})/4$	4%	$O(1)$
INKA	Eq. (31)	$({}_0\bar{M}_{\ell\ell_1,0}\bar{M}_{\ell'\ell_2,0} + {}_0\bar{M}_{\ell'\ell_1,0}\bar{M}_{\ell\ell_2,0})/2$	4%	$O(\ell_{\max}^3)$ or $O(\ell_{\max}^2)$ with Louis et al. (2020)
ACC (this work)	Eq. (33)	$\bar{\Theta}_{\ell\ell'}$ invariant for $\Delta \equiv \ell - \ell' = \text{cst}$	1%	$O(d_{\max} n_{\text{side}}^4)$

Notes. First column: name of approximation. Second column: equation to which they are referred. Third column: expression of $\bar{\Theta}$ in this approximation. Fourth column: precision determined by the maximum values of the relative difference of the EEEE binned covariance on diagonal for multipoles $\ell_{\text{cut}} \leq \ell \leq \ell_{\max,\text{ex}}$ in Fig. 6. For larger multipoles, the approximation are expected to be in this range of precision as shown in Fig. 12. Fifth column: summary of computing resources needed to obtain $\bar{\Theta}$ in each approximation. Let us here specify that for INKA, the kernel \bar{M} is often already known, thus the practical complexity is $O(1)$. ℓ_{\max} is the multipole range of the covariance, d_{\max} is the number of diagonal computed in the ACC approximation, n_{side} is the resolution chosen to compute the covariance coupling kernels in the ACC approximation (closest to ℓ_{cut} is sufficient).

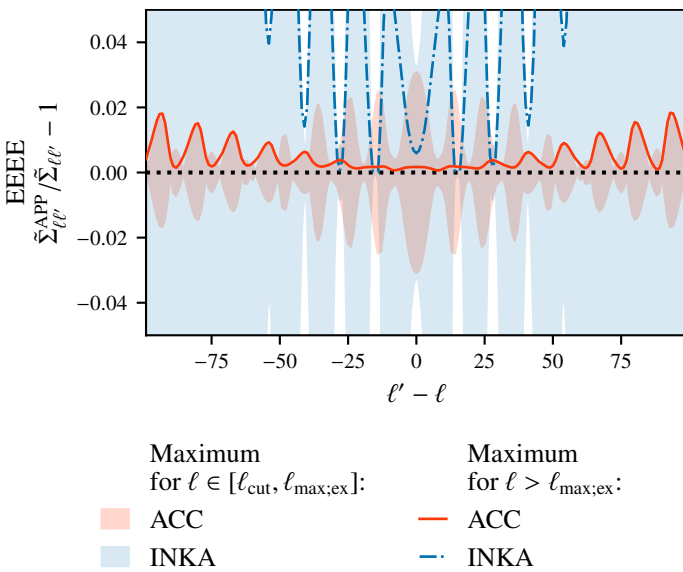


Fig. 12. Zoom of Fig. 7. We focus on the relative differences of the ACC and INKA covariance matrices with respect to the exact covariance for EEEE. The deviations found at $\ell < 1000$ (shaded regions) are similar to those found at the higher multipoles (lines), showing that the approximations work at the same level of accuracy in the two cases.

helps to alleviate the problem, but at the price of discarding a significant area of the usable sky. Even in the case of a large survey footprint, such as *Planck*, the point-source masks have been shown to break the NKA. In this case, the issue was mitigated using a simulation-based correction. For FRI, INKA, and ACC, preliminary work that we have performed also suggests that they fail when sources are included. We expect these approximations to perform poorly because the improvements over NKA are focused on the central shape of the reduced covariance coupling matrix $\bar{\Theta}$, while the sources tend to affect the far off-diagonal terms, which increases the correlation between distant multipoles. Particularly for the ACC approximation, the stronger correlations at distant multipoles will break the asymptotic behavior of the Wigner-3j symbols shown in Appendix E. We can expect that this reduces the validity of the approximated invariance by translation along the covariance diagonal, which is at the core of the ACC approximation. More work is required to assess the accuracy of ACC and other approximations in this case. However, different approaches can be adopted to mitigate the effect of a point-source mask. For example, analytical

solutions might be found (Gratton et al., in prep.), the maps might be inpainted (Benoit-Lévy et al. 2013), or a MC correction such as the one used in *Planck* might be employed (Planck Collaboration XI 2016).

6. Covariance of the PolSpice estimator

The pseudo-power spectrum is a biased estimator of the true underlying spectrum of the masked CMB maps. To recover an unbiased estimator, we can apply the MASTER (Hivon et al. 2002) formalism, which inverts the mode-coupling matrix and applies it to the biased estimator. Similarly, we can use the PolSpice (Szapudi et al. 2001; Chon et al. 2004) algorithm, which corrects the two-point correlation function for the effect of the mask in real space and then converts the result back into harmonic space. However, when the sky footprint is small and no large angular scales are observed, the unbinned mode-coupling matrix becomes singular. Analogously, the PolSpice conversion of the two-point correlation function into a power spectrum cannot be performed. In the first case, the mode-coupling matrix must be binned to allow the inversion. In the second case, we must apodize (i.e., cut gradually) the large angular scales of the two-point correlation function before calculating the corresponding power spectrum. This introduces a small bias in the final estimator that cannot be corrected for.

In this section, we explain in detail how the covariance matrix is calculated for the PolSpice estimator starting from a pseudo-power spectrum covariance matrix, which we produce through our ACC approximation. We show how the effect of the correction of the mask is included, as well as the small bias introduced by the PolSpice apodization of the two-point correlation function. In particular, we show that this apodization can be expressed in harmonic space, allowing us to relate the PolSpice spectrum covariance matrix to that of the pseudo-spectrum with a convolution.

6.1. MASTER equation

The pseudo-power spectrum is related to the true spectrum through the well-known MASTER equation introduced in Hivon et al. (2002),

$$\langle \tilde{C}_{\ell}^{\text{TT}} \rangle = \sum_{\ell'} {}_0M_{\ell\ell'} C_{\ell'}^{\text{TT}}, \quad (34)$$

with similar equations for polarization; see Appendix A.2. This bias comes from the information that is missing due to the

masked sky. Given the weighted mask $W(\hat{n})$, we can compute ${}_0M$ using Eq. (A.15). Provided that ${}_0M$ is invertible, an unbiased estimator can be constructed. These relations can be expressed in real space using the two-point correlation functions ξ , which for a statistically isotropic sky depend only on the relative angle between two directions,

$$\langle T(\hat{n}_1)T(\hat{n}_2) \rangle = \xi(\arccos(\hat{n}_1 \cdot \hat{n}_2)). \quad (35)$$

They can be related to the power spectrum C_ℓ using a Legendre series, with

$$\xi(\theta) = \sum_{\ell} \frac{2\ell + 1}{4\pi} C_\ell P_\ell(\cos \theta), \quad (36)$$

$$C_\ell = 2\pi \int_{-1}^1 d \cos \theta \xi(\theta) P_\ell(\cos \theta). \quad (37)$$

When we define the correlation function $\tilde{\xi}$ of the masked sky in the same manner, associated with the pseudo-spectrum \tilde{C}_ℓ , we obtain from Eq. (34) by applying the decomposition in a Legendre series the following relation:

$$\langle \tilde{\xi}(\theta) \rangle = w(\theta)\xi(\theta), \quad \forall \theta \in [0, \pi], \quad (38)$$

where $w(\theta)$ is the mask angular correlation function (more details can be found in Appendix C). From this relation, we can establish that the MASTER mode-coupling matrix ${}_0M$ is invertible only when the correlation function of the mask $w(\theta)$ is nonzero for all $\theta \in [0, \pi]$, which implies that the survey area explores all angular separations on the sky. While this is valid for an almost full-sky analysis such as *Planck*, it does not hold for experiments observing small patches, such as SPT-3G, where angular scales larger than $\theta \sim 30$ deg are unexplored. As a result, ${}_0M$ is not invertible. Binning allows the regularization of the MASTER matrix and thus to build a nearly unbiased estimator of the bandpowers. This approach is described in Hivon et al. (2002) and is adopted in NaMaster⁸ (Alonso et al. 2019). Similarly, we show in the next section how the unobserved large angular scales are handled in the PolSpice estimator.

6.2. Regularizing with PolSpice

6.2.1. Temperature

The pseudo-power spectrum estimator can be regularized in real space following the PolSpice approach in Szapudi et al. (2001). The pseudo-correlation function $\tilde{\xi}$ is smoothed with a scalar apodizing function $f^{\text{apo}}(\theta)$, which cuts out large θ and then corrects for the bias coming from the weighted mask described in Eq. (38). The scalar apodizing function goes smoothly from $f^{\text{apo}}(0) = 1$ to $f^{\text{apo}}(\theta_{\text{max}}) = 0$ to avoid Fourier ringing. θ_{max} should be chosen as the maximum angular size of the weighted mask. A new correlation function estimator $\hat{\xi}(\theta)$ is defined as

$$\hat{\xi}(\theta) \equiv g(\theta)\tilde{\xi}(\theta), \quad (39)$$

with

$$g(\theta) = \begin{cases} f^{\text{apo}}(\theta)/w(\theta) & \forall \theta \in [0, \theta_{\text{max}}), \\ 0 & \forall \theta \in [\theta_{\text{max}}, \pi]. \end{cases} \quad (40)$$

The function g is well defined and smooth for all angles through the apodization f^{apo} . As a consequence of Eqs. (38) and (39), the

⁸ <https://github.com/LSSTDESC/NaMaster>

PolSpice estimator of the correlation function can be related on average to the true underlying correlation function with

$$\langle \hat{\xi}(\theta) \rangle = f^{\text{apo}}(\theta)\xi(\theta) \quad \forall \theta \in [0, \pi]. \quad (41)$$

Returning to harmonic space using a Legendre transform, this operation can be expressed as

$$\langle \hat{C}_\ell \rangle = \sum_{\ell'} {}_0K_{\ell\ell'} C_{\ell'}. \quad (42)$$

$${}_0K_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{\text{00}}[f^{\text{apo}}]. \quad (43)$$

The PolSpice kernel ${}_0K$ is obtained from the scalar apodizing function f^{apo} with an extended definition of the operators Ξ (see Appendix C and Eq. (C.7) for more details). The operator acts on the Legendre transform of f^{apo} .

The advantage of PolSpice, which performs the regularization in real space rather than in harmonic space, is that it replaces an ℓ -space convolution by an integration and a multiplication, which are faster and numerically more stable. This produces an estimator for all multipoles ℓ . We denote this estimator with a hat, for instance, \hat{C}_ℓ^{XY} . This regularization (which is only required for small sky patches) introduces a bias in the PolSpice estimator that cannot be corrected for because the coupling is not invertible generally as the apodizing function f^{apo} reaches zero. The bias is small because ${}_0K$ is properly normalized, that is, $\sum_{\ell} {}_0K_{\ell\ell} = 1$. Furthermore, the regularization increases the correlations between unbinned modes. The PolSpice kernels behave as window functions, mixing multipoles of the pseudo-power spectrum. The lack of information at large scales induces the inability to distinguish multipoles that are close to each other. For this reason, the spectrum estimator is binned in ranges larger than the typical correlation between multipoles for a cosmological analysis.

6.2.2. Polarization

PolSpice allows correcting for the bias introduced by the cut sky in the same manner as for the polarized spectra. It also allows decoupling the EE and BB estimator; see Challinor & Chon (2005) or Appendix C. Similarly to the intensity case, we can express the effect of the PolSpice real-space regularization in spherical harmonics by defining the polarized PolSpice kernels, ${}_{\pm 2}K_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{2\pm 2}[f^{\text{apo}}]$. The PolSpice estimator follows for $X \in \{E, B\}$

$$\langle \hat{C}_\ell^{\text{XX}} \rangle = \sum_{\ell'} {}_{-2}K_{\ell\ell'} C_{\ell'}^{\text{XX}}. \quad (44)$$

For the temperature \times polarization case, we can show that

$$\langle \hat{C}_\ell^{\text{TE}} \rangle = \sum_{\ell'} {}_{\times}K_{\ell\ell'} C_{\ell'}^{\text{TE}}, \quad (45)$$

$$\text{with } {}_{\times}K_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{20}[f^{\text{apo}}]. \quad (46)$$

6.3. Relating POLSPICE and MASTER in harmonic space

We can translate the relations Eq. (39) into harmonic space in temperature and polarization to obtain the PolSpice estimator as a harmonic convolution of the pseudo-power spectrum estimator,

$$\hat{C}_\ell^{\text{TT}} = \sum_{\ell'} {}_0G_{\ell\ell'} \tilde{C}_{\ell'}^{\text{TT}}, \quad (47)$$

$$\hat{C}_\ell^{\text{EE}} - \hat{C}_\ell^{\text{BB}} = \sum_{\ell'} -_2G_{\ell\ell'} (\tilde{C}_{\ell'}^{\text{EE}} - \tilde{C}_{\ell'}^{\text{BB}}). \quad (48)$$

$$\hat{C}_\ell^{\text{EE}} + \hat{C}_\ell^{\text{BB}} = \sum_{\ell'} \text{dec}G_{\ell\ell'} (\tilde{C}_{\ell'}^{\text{EE}} + \tilde{C}_{\ell'}^{\text{BB}}), \quad (49)$$

$$\hat{C}_\ell^{\text{TE}} = \sum_{\ell'} \times G_{\ell\ell'} \tilde{C}_{\ell'}^{\text{TE}}. \quad (50)$$

The G kernels are constructed in the same manner as the PolSpice kernels, with the operator Ξ acting on the function $g = f^{\text{apo}}/w$ according to Eq. (C.7) (or equivalently, on the associated power spectrum of g via Legendre transform). They are given by

$${}_0G_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{00}[g], \quad (51)$$

$$-_2G_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{2-2}[g], \quad (52)$$

$$\times G_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{20}[g], \quad (53)$$

$$\text{dec}G_{\ell\ell'} = \frac{2\ell' + 1}{2} \int_{-1}^1 g(\theta) d_{22}^{\ell}(\theta) d_{2-2}^{\ell'}(\theta) d\cos(\theta). \quad (54)$$

The first three equations above reduce to the inverse of the master kernels when the PolSpice apodization function is set to 1, that is, no apodization. The last kernel, referred to as dec, is the kernel that allows the decoupling of the EE and BB spectra. Appendix C gives more details on this point. $\text{dec}G$ is associated with integral relations in real space, thus its harmonic expression is not straightforward. This expression requires more numerical resources to be computed because no closed relations exist for the Wigner d-matrix with different multipole indices. It can be obtained with PolSpice for all ℓ by setting $\tilde{C}_{\ell'}^{\text{EE}} = \tilde{C}_{\ell'}^{\text{BB}} = \delta_{\ell\ell'}$ as input.

6.4. Covariance of the POLSPICE estimator

Given the previous relations in Eqs. (47)–(50), the covariance of the PolSpice estimator can be written as a convolution of the covariance of the pseudo-power spectrum, with

$$\begin{aligned} \hat{\Sigma}_{\ell\ell'}^{\text{TTTT}} &\equiv \text{cov}(\hat{C}_\ell^{\text{TT}}, \hat{C}_{\ell'}^{\text{TT}}) \\ &= \sum_{LL'} {}_0G_{\ell L} \hat{\Sigma}_{LL'}^{\text{TTTT}} {}_0G_{L\ell'}. \end{aligned} \quad (55)$$

For polarization, there is mixing between the EE and BB components in the covariance. We write the polarized EEEE PolSpice covariance after defining $\pm G \equiv \frac{1}{2}(\text{dec}G \pm -_2G)$ as

$$\begin{aligned} \hat{\Sigma}^{\text{EEEE}} &= {}_+G\hat{\Sigma}^{\text{EEEE}}_+G^\text{T} + {}_-G\hat{\Sigma}^{\text{BBEE}}_+G^\text{T} \\ &\quad + {}_+G\hat{\Sigma}^{\text{EEBB}}_-G^\text{T} + {}_-G\hat{\Sigma}^{\text{BBBB}}_-G^\text{T}. \end{aligned} \quad (56)$$

The polarized PolSpice covariance is built on the polarized pseudo-covariance, mixing the components EE and BB, thanks to the kernel $\pm G$; see Fig. 13. This figure displays a row of the kernels that were computed on the mask SPT-3G used in our analysis. It shows the window functions that are applied to the pseudo-power spectra to produce the PolSpice spectra. They correct for the bias due to the mask, but introduce a small bias due to the lack of information at large scales. The temperature kernel ${}_0G$ and the polarization ${}_+G$ kernel are almost identical. The leakage kernel ${}_ -G$ (all negative), which accounts for the mixing of the E and B polarization pseudo-spectra in the PolSpice spectrum, is orders of magnitudes smaller than the other two. Hence, the BB covariance terms do not affect the

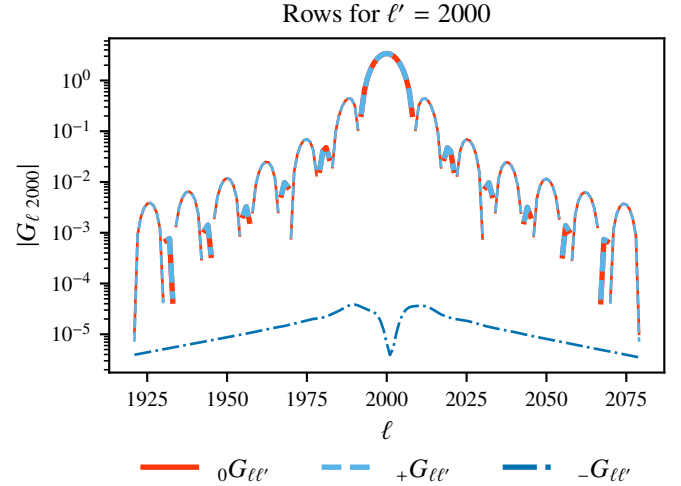


Fig. 13. Amplitude of the PolSpice convolution kernels G for the SPT-3G footprint. The negative terms are plotted with thinner lines.

EEEE covariance. On the other hand, the EE terms affect the BBBB covariance because the EE spectrum is a few orders of magnitude larger than BB. The PolSpice apodizing function of the correlation function we used is

$$f^{\text{apo}}(\theta) = \begin{cases} \frac{1}{2} \left(1 + \cos \frac{\pi\theta}{\theta_{\text{max}}} \right) & \forall \theta < \theta_{\text{max}}, \\ 0 & \text{otherwise.} \end{cases} \quad (57)$$

Here we have, $\theta_{\text{max}} = \pi/6$. Without apodization, but with partial sky, as for *Planck*, the decoupling kernel $\text{dec}G$ is not null, still resulting in a nonzero ${}_ -G$ kernel. However, it is orders of magnitude smaller than ${}_+G$, hence we can compute EEEE ignoring the leakage from covariance terms that include BB.

6.5. Accuracy of the covariance approximations for the POLSPICE estimator

We can build estimates of the PolSpice spectrum covariance by convolving the pseudo-spectrum covariance with the appropriate kernels following Eqs. (55) and (56). To calculate the pseudo-spectrum covariance, we can use the NKA, INKA, FRI, and ACC approximations or the exact computation. Figure 14 shows the accuracy of the binned PolSpice covariance calculated with the approximations compared to the exact calculation. The results are similar to those we found for the accuracy of the pseudo-spectrum covariances shown in Fig. 6. The NKA, INKA, and FRI approaches provide a good estimate of the PolSpice covariance. However, the ACC approach improves the existing approximations dramatically. This shows that our results for the accuracy of the pseudo-covariance also hold for the PolSpice.

7. Summary and conclusions

One of the key ingredients of cosmological analysis based on power spectra are covariance matrices. Accurate covariance matrices ensure precise error bars and an unbiased estimation of cosmological parameters. The analytical estimation of these matrices can be difficult when small sky fractions are observed because existing approximations might fail. We have considered the specific example of estimating accurate analytical signal-signal covariance matrices for the SPT-3G CMB experiment,




Fig. 14. Relative differences of binned PolSpice covariance matrices calculated using approximations of the pseudo-spectrum covariance with respect to the exact computation: $\hat{\Sigma}_{bb'}^{\text{APP}}/\hat{\Sigma}_{bb'} - 1$, for TTTT (left-hand side) and EEEE (right-hand side), with binning $\Delta\ell = 50$. On the first row we plot the relative differences for the diagonal, i.e. $b = b'$, while on the second row we plot the differences for the first off-diagonal, i.e. $b' = b + 1$. Similarly to the case of pseudo-covariances in Fig. 6, we find acceptable accuracy for the NKA (dash light blue), INKA (dash-dot dark blue) and FRI (dash-double-dot purple) approximations, while our ACC approximation (solid red) improves overall the others. The PolSpice covariance matrices have been calculated using Eqs. (55) and (56). The last row displays the corresponding binned underlying renormalized spectrum TT or EE, to highlight the fact that the differences in the covariances are on the peaks and in the troughs of the spectra, i.e. where the curvature of the spectrum is maximal.

whose survey covers about 4%, without masking the contribution of point sources. We considered the cases of estimating the matrix for pseudo-power spectrum and for the PolSpice power spectrum estimator.

First, we implemented in Sect. 3 an expensive exact calculation of the covariance of the pseudo-power spectrum in intensity and polarization for the first time. We used a map-based algorithm that is accelerated thanks to the HEALPix pixelation tools. We were thus able to compute the entire covariance matrix up to $\ell_{\text{max,ex}} = 1000$ exactly. We also obtained a selection of rows of the covariance of particular interest up to $\ell = 2000$.

Based on this result, we were able to estimate the accuracy of the existing approximations in Sect. 4 precisely by comparing them to the binned exact covariances of the pseudo-power spectra measured on the SPT-3G patch. The approximations were found to be precise to the 5% level.

Then, using the code we developed for an exact computation of the covariance matrix, we estimated the covariance-coupling kernel $\bar{\Theta}$, which determines how the CMB power spectrum couples into the covariance matrix. We were able to understand why the existing approximations in the literature fail to achieve a precision better than 5%. We then proposed a new approximation in Sect. 5, the ACC, which is more computationally expensive than

the existing approximations, but allows a more precise estimation of the covariance matrix at the 1% level.

Finally, we showed in Sect. 6 that we are able to build the covariance of the PolSpice power spectrum in both temperature and polarization using a harmonic correction. This computation is exact and based on the PolSpice algorithm real-space corrections that we translated into harmonic space. Through this correction, we produced estimates of the PolSpice covariance matrix based on the previous approximations of the pseudo-power spectrum covariance. The accuracy of the resulting PolSpice covariance approximations is the same as for the pseudo-power spectrum.

While this paper considered the particular example of the SPT-3G experiment, the results can be extended to a non-CMB power spectrum analysis such as that of weak-lensing shear or photometric catalogs. We would also like to stress that the accuracy of any of the approximations presented in this paper (existing or new) is reduced when a point-source mask is included in the sky footprint. Nevertheless, the exact computation of the covariance matrices still holds in this particular case. While previous experiments have included the effect of point-source masks through the use of simulations (see, e.g., Planck Collaboration XI 2016) or by inpainting the holes with

constrained realizations (see, e.g., [Benoit-Lévy et al. 2013](#)), additional work is required to find an analytical calculation of this contribution.

Acknowledgements. We are grateful to the SPT collaboration for discussions and suggestions. We thank Lennart Balkenhol for discussions and insightful comments on the manuscript. This work has received funding from the French Centre National d'Etudes Spatiales (CNES). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101001897).

References

- Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
- Aiola, S., Calabrese, E., Maurin, L., et al. 2020, *J. Cosmol. Astropart. Phys.*, 2020, 047
- Alonso, D., Sanchez, J., & Slosar, A. 2019, *MNRAS*, 484, 4127
- Balkenhol, L., & Reichardt, C. L. 2022, *MNRAS*, 512, 4394
- Benoit-Lévy, A., Déchelette, T., Benabed, K., et al. 2013, *A&A*, 555, A37
- Challinor, A., & Chon, G. 2005, *MNRAS*, 360, 509
- Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *MNRAS*, 350, 914
- Dodelson, S., & Schneider, M. D. 2013, *Phys. Rev. D*, 88, 063537
- Dutcher, D., Balkenhol, L., Ade, P. A. R., et al. 2021, *Phys. Rev. D*, 104, 022003
- Efstathiou, G. 2004, *MNRAS*, 349, 603
- Friedrich, O., Andrade-Oliveira, F., Camacho, H., et al. 2021, *MNRAS*, 508, 3125
- Gallardo, P. A., Benson, B., Carlstrom, J., et al. 2022, *Proc. SPIE*, 12190, 121900C
- García-García, C., Alonso, D., & Bellini, E. 2019, *J. Cosmol. Astropart. Phys.*, 2019, 043
- Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Hazumi, M., Borrill, J., Chinone, Y., et al. 2012, *Proc. SPIE*, 8442, 844219
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, 208, 19
- Hivon, E., Gorski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, 567, 2
- Khersonskii, V., Moskalev, A., & Varshalovich, D. 1988, *Quantum Theory Of Angular Momentum* (Singapore: World Scientific Publishing Company)
- Louis, T., Naess, S., Garrido, X., & Challinor, A. 2020, *Phys. Rev. D*, 102, 123538
- Lueker, M., Reichardt, C. L., Schaffer, K. K., et al. 2010, *ApJ*, 719, 1045
- Nicola, A., Garcia-Garcia, C., Alonso, D., et al. 2021, *J. Cosmol. Astropart. Phys.*, 2021, 067
- Planck Collaboration XI. 2016, *A&A*, 594, A11
- Planck Collaboration V. 2020, *A&A*, 641, A5
- Planck Collaboration VI. 2020, *A&A*, 641, A6
- Sellentin, E., & Starck, J.-L. 2019, *J. Cosmol. Astropart. Phys.*, 2019, 021
- Szapudi, I., Prunet, S., Pogosyan, D., Szalay, A. S., & Bond, J. R. 2001, *ApJ*, 548, L115

Appendix A: Analysis of the curved sky

This appendix describes the mathematical tools that are used on curved sky for a CMB analysis. We make use of the spherical harmonic decomposition of Gaussian fields. We introduce various operators that allow us to express the couplings and the covariance of the power spectra. We make use of some geometrical relations of spherical harmonics to obtain our results. In this appendix, we introduce formulae that can either be used in the temperature or in the polarization case.

A.1. Temperature

We first consider the case of a map of the CMB intensity anisotropies $T(\hat{n})$ observed in direction \hat{n} . The anisotropies are distributed as a Gaussian random field with a corresponding power spectrum C_ℓ^{TT} , observed through a mask $W(\hat{n})$.

Harmonic coefficients and underlying power spectrum

The intensity map can be decomposed with spin-0 spherical harmonics to obtain the harmonic coefficients and their variance, the intensity power spectrum, which fully characterizes the physical properties of the field,

$$a_{\ell m}^{\text{T}} = \int d\hat{u} T(\hat{u}) {}_0Y_{\ell m}^*(\hat{u}), \quad (\text{A.1})$$

$$\langle a_{\ell m}^{\text{T}} a_{\ell' m'}^{\text{T}*} \rangle = C_\ell^{\text{TT}} \delta_{\ell\ell'} \delta_{mm'}. \quad (\text{A.2})$$

Here the brackets $\langle \rangle$ indicate an average over many realizations of the maps.

Weighted mask $W(\hat{n})$

The weighted mask is a real map with weights from 0 to 1 that is used to taper the data on the border of the survey area in order to reduce Fourier ringing when harmonic decomposition is used. We define the mask harmonic coefficients and its power spectrum as

$$w_{\ell m} \equiv \int d\hat{u} W(\hat{u}) {}_0Y_{\ell m}^*(\hat{u}), \quad (\text{A.3})$$

$$\mathcal{W}_\ell \equiv \frac{1}{2\ell+1} \sum_m w_{\ell m} w_{\ell m}^*. \quad (\text{A.4})$$

Pseudo-power spectrum estimator

One way to obtain a biased estimator of the power spectrum in CMB experiments is to define the pseudo-harmonic coefficients and the pseudo-power spectrum,

$$\tilde{a}_{\ell m}^{\text{T}} \equiv \int d\hat{u} W(\hat{u}) T(\hat{u}) {}_0Y_{\ell m}^*(\hat{u}), \quad (\text{A.5})$$

$$\tilde{C}_\ell^{\text{TT}} \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |\tilde{a}_{\ell m}^{\text{T}}|^2. \quad (\text{A.6})$$

Relation between harmonic coefficients

We relate the masked pseudo-harmonic coefficients to the unmasked coefficient with

$$\tilde{a}_{\ell m}^{\text{T}} = \sum_{\ell' m'} a_{\ell' m'}^{\text{T}} {}_0I_{\ell m \ell' m'}[W]. \quad (\text{A.7})$$

The $I[W]$ couplings are defined below and can be expressed in terms of sums over Wigner-3j symbols and the $w_{\ell m}$, with

$${}_sI_{\ell m \ell' m'}[W] \equiv \int d\hat{u} {}_sY_{\ell m}(\hat{u}) W(\hat{u}) {}_sY_{\ell' m'}^*(\hat{u}), \quad (\text{A.8})$$

$$= \sum_{LM} w_{LM} (-1)^{m'} \left[\frac{(2\ell+1)(2\ell'+1)(2L+1)}{4\pi} \right]^{1/2} \begin{pmatrix} \ell & \ell' & L \\ -s & s & 0 \end{pmatrix} \times \begin{pmatrix} \ell & \ell' & L \\ m & -m' & M \end{pmatrix}. \quad (\text{A.9})$$

Here, we anticipated the extension of this notation to the polarized case, which deals with spin-2 fields. The mask-dependent ${}_sI_{\ell m \ell' m'}[W]$ coupling coefficients relate the underlying harmonic coefficients to the measured pseudo-harmonic coefficients. In the full-sky case, we obtain ${}_sI_{\ell m \ell' m'}[1] = \delta_{\ell\ell'} \delta_{mm'}$ through the orthonormality properties of the spin-weighted spherical harmonics.

We introduce some useful relations that are demonstrated in [Hivon et al. \(2002\)](#),

$$\sum_{\ell m} {}_sI_{\ell_1 m_1 \ell m}[W] {}_sI_{\ell_2 m_2 \ell m}^*[W] = {}_sI_{\ell_1 m_1 \ell_2 m_2}[W^2], \quad (\text{A.10})$$

$$\sum_{m_1 m_2} \frac{{}_sI_{\ell_1 m_1 \ell_2 m_2}[W] {}_sI_{\ell_1 m_1 \ell_2 m_2}^*[W]}{(2\ell_1+1)(2\ell_2+1)} = \Xi_{\ell_1 \ell_2}^{ss'}[W]. \quad (\text{A.11})$$

Here we introduced the symmetric operator $\Xi^{ss'}$ acting on a power spectrum \mathcal{A}_ℓ ,

$$\Xi_{\ell\ell'}^{ss'}[\mathcal{A}] \equiv \sum_L \frac{2L+1}{4\pi} \mathcal{A}_L \begin{pmatrix} \ell & \ell' & L \\ s & -s & 0 \end{pmatrix} \begin{pmatrix} \ell & \ell' & L \\ s' & -s' & 0 \end{pmatrix}. \quad (\text{A.12})$$

We extend this definition to an operator acting on a map $A(\hat{n})$, with

$$\Xi_{\ell\ell'}^{ss'}[A] \equiv \Xi_{\ell\ell'}^{ss'}[\mathcal{A}], \quad (\text{A.13})$$

where we defined the power spectrum \mathcal{A}_ℓ of the map A as in Eqs. (A.3) and (A.4).

MASTER relation between estimated and true spectra

By inserting Eq. (A.5) into Eq. (A.6), using the relations of Eq. (A.11) and the definition of Eq. (A.12)-(A.13), we relate the ensemble average of the pseudo-power spectrum to the underlying power spectrum using the MASTER mode-coupling kernel with

$$\langle \tilde{C}_\ell^{\text{TT}} \rangle = \sum_{\ell'} {}_0M_{\ell\ell'}[W] C_{\ell'}^{\text{TT}}. \quad (\text{A.14})$$

The MASTER mode-coupling matrix is given by

$${}_0M_{\ell\ell'}[W] \equiv (2\ell'+1) \Xi_{\ell\ell'}^{00}[W]. \quad (\text{A.15})$$

A.2. Polarization

We consider the case of the CMB intensity and polarization anisotropies, represented by maps $T(\hat{n})$, $Q(\hat{n})$, $U(\hat{n})$ in direction \hat{n} of the sky. These are Gaussian random fields, fully characterized by their power spectra (C_ℓ^{TT} , C_ℓ^{EE} , C_ℓ^{BB} , C_ℓ^{TE}) observed

through a mask $W(\hat{n})$. The definitions and relations of the previous section can be extended to polarization spectra. First we compute the pseudo-harmonic coefficients on the masked sky with spin-weighted spherical harmonics, given in the following inverse relation from [Chon et al. \(2004\)](#):

$$(Q \pm iU)(\hat{n}) = \sum_{\ell m} (\tilde{a}_{\ell m}^E \mp i\tilde{a}_{\ell m}^B) \mp_2 Y_{\ell m}. \quad (\text{A.16})$$

The pseudo-power spectrum $\tilde{C}_\ell^{\text{XY}}$ is obtained by summing over the measured pseudo-harmonic coefficients $\tilde{a}_{\ell m}^{\text{X}}$, X, Y \in [T, E, B] with the same multipole ℓ , with

$$\tilde{C}_\ell^{\text{XY}} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \tilde{a}_{\ell m}^{\text{X}} \tilde{a}_{\ell m}^{\text{Y}*}. \quad (\text{A.17})$$

More details can be found in [Challinor & Chon \(2005\)](#). Following the same approach as in the previous section, we can write the MASTER relation in polarization,

$$\langle \tilde{C}_\ell^{\text{EE}} + \tilde{C}_\ell^{\text{BB}} \rangle = \sum_{\ell'} {}_{+2}M_{\ell\ell'} (C_{\ell'}^{\text{EE}} + C_{\ell'}^{\text{BB}}), \quad (\text{A.18})$$

$${}_{+2}M_{\ell\ell'} = (2\ell' + 1) \Xi_{\ell\ell'}^{22}[W], \quad (\text{A.19})$$

$$\langle \tilde{C}_\ell^{\text{EE}} - \tilde{C}_\ell^{\text{BB}} \rangle = \sum_{\ell'} {}_{-2}M_{\ell\ell'} (C_{\ell'}^{\text{EE}} - C_{\ell'}^{\text{BB}}), \quad (\text{A.20})$$

$${}_{-2}M_{\ell\ell'} = (2\ell' + 1) \Xi_{\ell\ell'}^{2-2}[W], \quad (\text{A.21})$$

$$\langle \tilde{C}_\ell^{\text{TE}} \rangle = \sum_{\ell'} \times M_{\ell\ell'} C_{\ell'}^{\text{TE}}, \quad (\text{A.22})$$

$$\times M_{\ell\ell'} = (2\ell' + 1) \Xi_{\ell\ell'}^{20}[W]. \quad (\text{A.23})$$

A.3. Renormalized kernels

We define the renormalized MASTER kernels, which are used in the INKA. They are written as

$${}_k\bar{M}_{\ell\ell'} \equiv \frac{1}{\sum_{\ell'} {}_kM_{\ell\ell'}} {}_kM_{\ell\ell'}, \quad \forall k \in [0, -2, +2, \times]. \quad (\text{A.24})$$

Summing over ℓ' yields

$$\sum_{\ell'} {}_k\bar{M}_{\ell\ell'} = 1, \quad (\text{A.25})$$

which ensures that the approximated covariance-coupling kernel defined in Eq. (30) is properly normalized.

Appendix B: Covariance of the pseudo-power spectrum

In this appendix, we outline how the formula of the covariance matrix of the pseudo-power spectrum is obtained in the temperature case. Our goal is to introduce Eq. (B.9). The covariance matrix of the pseudo-power spectrum reads

$$\begin{aligned} \tilde{\Sigma}_{\ell\ell'} &= \langle \tilde{C}_\ell \tilde{C}_{\ell'} \rangle - \langle \tilde{C}_\ell \rangle \langle \tilde{C}_{\ell'} \rangle, \\ &= \sum_{mm'} \frac{\langle |\tilde{a}_{\ell m}|^2 |\tilde{a}_{\ell' m'}|^2 \rangle - \langle |\tilde{a}_{\ell m}|^2 \rangle \langle |\tilde{a}_{\ell' m'}|^2 \rangle}{(2\ell + 1)(2\ell' + 1)}. \end{aligned} \quad (\text{B.1})$$

As the intensity map $T(\hat{n})$ is real and the spherical harmonics follow

$${}_0Y_{\ell m}^* = (-1)^m {}_0Y_{\ell(-m)}, \quad (\text{B.2})$$

the spherical harmonic coefficients of $T(\hat{n})$ follow

$$a_{\ell m}^* = (-1)^m a_{\ell(-m)}. \quad (\text{B.3})$$

When the four-point function is computed and thanks to Wick's theorem,

$$\begin{aligned} &\sum_{mm'} \langle |\tilde{a}_{\ell m}|^2 |\tilde{a}_{\ell' m'}|^2 \rangle - \langle |\tilde{a}_{\ell m}|^2 \rangle \langle |\tilde{a}_{\ell' m'}|^2 \rangle \\ &= \sum_{mm'} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'} \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell' m'}^* \rangle + \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell' m'} \rangle. \end{aligned} \quad (\text{B.4})$$

Based on the first term of this sum and using the change of variable $m'' = -m'$, it is straightforward to show that

$$\begin{aligned} &\sum_{mm'} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'} \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell' m'}^* \rangle \\ &= \sum_{mm'} (-1)^{(2m')} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell'(-m')} \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell'(-m')}^* \rangle, \\ &= \sum_{mm'} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell' m'} \rangle. \end{aligned} \quad (\text{B.5})$$

Finally, we have

$$\begin{aligned} &\sum_{mm'} \langle |\tilde{a}_{\ell m}|^2 |\tilde{a}_{\ell' m'}|^2 \rangle - \langle |\tilde{a}_{\ell m}|^2 \rangle \langle |\tilde{a}_{\ell' m'}|^2 \rangle \\ &= 2 \sum_{mm'} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle \langle \tilde{a}_{\ell m}^* \tilde{a}_{\ell' m'} \rangle \\ &= 2 \sum_{mm'} |\langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle|^2. \end{aligned} \quad (\text{B.6})$$

Then, we have for the covariance matrix,

$$\tilde{\Sigma}_{\ell\ell'} = \frac{2}{(2\ell + 1)(2\ell' + 1)} \sum_{mm'} |\langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle|^2. \quad (\text{B.7})$$

When we use the decomposition of pseudo-harmonic coefficients, we obtain

$$\begin{aligned} \langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle &= \sum_{\ell_1 m_1 \ell_2 m_2} \langle a_{\ell_1 m_1} a_{\ell_2 m_2}^* \rangle I_{\ell m \ell_1 m_1}[W] I_{\ell' m' \ell_2 m_2}^*[W], \\ &= \sum_{\ell_1 m_1} C_{\ell_1} I_{\ell m \ell_1 m_1}[W] I_{\ell' m' \ell_1 m_1}^*[W]. \end{aligned} \quad (\text{B.8})$$

Inserting Eq. (B.8) into Eq. (B.7) gives

$$\begin{aligned} \tilde{\Sigma}_{\ell\ell'} &= \frac{2}{(2\ell + 1)(2\ell' + 1)} \sum_{mm'} \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} C_{\ell_1} C_{\ell_2} \\ &I_{\ell m \ell_1 m_1}[W] I_{\ell' m' \ell_1 m_1}^*[W] I_{\ell m \ell_2 m_2}^*[W] I_{\ell' m' \ell_2 m_2}[W]. \end{aligned} \quad (\text{B.9})$$

Appendix C: Expansion in Legendre series

In this section, we introduce the Legendre transforms of the harmonic quantities used in this work. Each relation in harmonic space has a corresponding expression in real space. The PolSpice software relies on the relations in real space.

C.1. From spin-0 spectra to two-point correlation functions

Given a spectrum \mathcal{A}_ℓ , we can associate a real two-point correlation function a with it,

$$a(\theta) = \sum_\ell \frac{2\ell+1}{4\pi} \mathcal{A}_\ell P_\ell(\cos\theta) \quad \forall \theta \in [0, \pi]. \quad (\text{C.1})$$

The inverse relation is

$$\mathcal{A}_\ell = 2\pi \int_0^\pi d\theta \sin\theta a(\theta) P_\ell(\cos\theta) \quad \forall \ell \geq 0. \quad (\text{C.2})$$

The two-point function gives the correlations between two directions of the sky, for instance, in the CMB anisotropy full-sky case, we can write

$$\langle T(\hat{n}_1) T(\hat{n}_2) \rangle = \sum_\ell \frac{2\ell+1}{4\pi} C_\ell^{\text{TT}} P_\ell(\hat{n}_1 \cdot \hat{n}_2). \quad (\text{C.3})$$

C.2. From convolution to multiplication

A convolution with a square matrix A in harmonic space, such as in Eq. (A.14), is equivalent to a multiplication in real space with the correlation function $a(\theta)$ given by

$$a(\theta) \frac{2\ell'+1}{4\pi} P_{\ell'}(\cos\theta) = \sum_\ell \frac{2\ell+1}{4\pi} A_{\ell\ell'} P_\ell(\cos\theta) \quad \forall \theta \in [0, \pi]. \quad (\text{C.4})$$

The inverse relation is

$$A_{\ell\ell'} \equiv \frac{2\ell'+1}{2} \int_0^\pi a(\theta) P_\ell(\cos\theta) P_{\ell'}(\cos\theta) \sin(\theta) d\theta. \quad (\text{C.5})$$

The last relation is equivalent to

$$A_{\ell\ell'} = (2\ell'+1) \Xi_{\ell\ell'}^{00}[\mathcal{A}], \quad (\text{C.6})$$

with \mathcal{A} the power spectrum associated with the two-point function a through a Legendre transform. We can thus extend the definition of the operator Ξ to an operator acting on a correlation function a , with

$$\Xi_{\ell\ell'}^{ss'}[a] \equiv \Xi_{\ell\ell'}^{ss'}[\mathcal{A}]. \quad (\text{C.7})$$

Here we have already extended the definition to be used in the spin-2 case, which we discuss in the next subsection.

C.3. Spin-2

Similar rules to those introduced above can be written for spin-2 quantities by replacing the Legendre polynomials by the more generally reduced Wigner d -matrix $d_{2\pm 2}^\ell$. Details can be found in [Challinor & Chon \(2005\)](#). The spin-2 relations for the \mathcal{A}_ℓ^\pm power spectra are associated with a spin-2 field,

$$a_\pm(\theta) = \sum_\ell \frac{2\ell+1}{4\pi} \mathcal{A}_\ell^\pm d_{2\pm 2}^\ell(\cos\theta), \quad (\text{C.8})$$

$$\mathcal{A}_\ell^\pm = 2\pi \int_0^\pi d\theta \sin\theta a_\pm(\theta) d_{2\pm 2}^\ell(\cos\theta). \quad (\text{C.9})$$

We can associate a spin-0 correlation function with its spin-2 convolution matrix,

$$\begin{aligned} \pm_2 A_{\ell\ell'} &= \frac{2\ell'+1}{2} \int_0^\pi a(\theta) d_{2\pm 2}^\ell(\cos\theta) d_{2\pm 2}^{\ell'}(\cos\theta) \sin(\theta) d\theta, \\ &= (2\ell'+1) \Xi_{\ell\ell'}^{2\pm 2}[a]. \end{aligned} \quad (\text{C.10})$$

We can also compute the matrix associated with the spin-0 cross spin-2 case,

$$\times A_{\ell\ell'} = (2\ell'+1) \Xi_{\ell\ell'}^{20}[a]. \quad (\text{C.11})$$

C.4. Applying this formalism to the MASTER matrix

In our case, we can write for $s \in [0, 2]$,

$$\pm_s M_{\ell\ell'} = (2\ell'+1) \Xi_{\ell\ell'}^{s\pm s}[w], \quad (\text{C.12})$$

with $w(\theta)$ the correlation function of the mask.

We apply the previous formalism and particularly the Eq. (C.12) to the MASTER relation. We use Legendre series expansion of the true power spectrum C_ℓ and the pseudo-power spectrum estimator \tilde{C}_ℓ to define the correlation functions ξ and $\tilde{\xi}$, respectively. It gives

$$\tilde{\xi}(\theta) \equiv \sum_\ell \frac{2\ell+1}{4\pi} P_\ell(\cos\theta) \tilde{C}_\ell, \quad (\text{C.13})$$

$$\xi(\theta) \equiv \sum_\ell \frac{2\ell+1}{4\pi} P_\ell(\cos\theta) C_\ell. \quad (\text{C.14})$$

Starting from the right-hand side of Eq. (34) and going to real space using a Legendre transform at an angle $\theta \in [0, \pi]$, we have

$$\begin{aligned} \sum_{\ell\ell'} \frac{2\ell+1}{4\pi} P_\ell(\cos\theta)_0 M_{\ell\ell'} C_{\ell'} \\ = \sum_{\ell'} w(\theta) \frac{2\ell'+1}{4\pi} P_{\ell'}(\cos\theta) C_{\ell'}, \\ = w(\theta) \xi(\theta), \end{aligned}$$

$$\text{which implies } \langle \tilde{\xi}(\theta) \rangle = w(\theta) \xi(\theta), \quad (\text{C.15})$$

with $w(\theta)$ the correlation function of the mask.

C.5. PolSpice in polarization

This section describes the regularization technique that is used for polarization by PolSpice. One of the main advantages of PolSpice is that it allows the possibility of eliminating EE to BB (and BB to EE) mixing, using nonlocal relations between Wigner d -matrices; see Sec. 5 of [Chon et al. \(2004\)](#). The obtained estimator \hat{C}_ℓ^{EE} (\hat{C}_ℓ^{BB}) depends only on the average of C_ℓ^{EE} (C_ℓ^{BB}) and the scalar apodizing function f .

Using the Legendre transforms of spin-2 quantities (Eq. (C.9)), we associate the correlation functions ξ_\pm with the spectra $C_\ell^{\text{EE}} \pm C_\ell^{\text{BB}}$ and $\tilde{\xi}_\pm$ to $\tilde{C}_\ell^{\text{EE}} \pm \tilde{C}_\ell^{\text{BB}}$. PolSpice builds two correlation functions $\hat{\xi}_{\text{dec}}$ and $\hat{\xi}_-$ to produce an estimator of the true underlying polarized power spectrum. This spectrum is defined similarly to ξ in Eq. (39),

$$\hat{\xi}_-(\theta) = g(\theta) \tilde{\xi}_-(\theta). \quad (\text{C.16})$$

The first is built on integral relations. PolSpice eliminates the mixing inherent in $\tilde{\xi}_+$ with the following relation in real space:

$$\hat{\xi}_{\text{dec}}(\theta) = f^{\text{apo}}(\theta) \int_{-1}^1 d\cos\theta' \frac{\tilde{\xi}_+(\theta')}{w(\theta')} \sum_\ell d_{2-2}^\ell(\theta) d_{22}^\ell(\theta'). \quad (\text{C.17})$$

This integration can be shown to depend only on ξ_+ in the range $\theta \in [0, \theta_{\text{max}}]$; see [Chon et al. \(2004\)](#). This allows decoupling the correlation functions from an incomplete range of angular separations that are missing due to the mask. This correlation function is noted with the subscript *dec* to emphasize that it is the crucial step allowing the decoupling of the polarized estima-

tor. When averaging out Eq. (C.16) and Eq. (C.17) on multiple realizations,

$$\langle \hat{\xi}_{\text{dec}}(\theta) \rangle = f(\theta) \int_{-1}^1 d \cos \theta' \xi_+(\theta') \sum_{\ell} d_{2-2}^{\ell}(\theta) d_{22}^{\ell}(\theta'), \quad (\text{C.18})$$

$$\langle \hat{\xi}_-(\theta) \rangle = f(\theta) \xi_-(\theta). \quad (\text{C.19})$$

In harmonic space, transforming Eq. (C.18) and Eq. (C.19) using d_{2-2}^{ℓ} , this gives with ${}_{-2}K_{\ell\ell'} = (2\ell' + 1)\Xi_{\ell\ell'}^{2-2}[f^{\text{apo}}]$,

$$\hat{C}_{\ell}^{\text{EE}} + \hat{C}_{\ell}^{\text{BB}} \equiv 2\pi \int_0^{\pi} d \cos \theta \hat{\xi}_{\text{dec}}(\theta) d_{2-2}^{\ell}(\theta), \quad (\text{C.20})$$

$$\hat{C}_{\ell}^{\text{EE}} - \hat{C}_{\ell}^{\text{BB}} \equiv 2\pi \int_0^{\pi} d \cos \theta \hat{\xi}_-(\theta) d_{2-2}^{\ell}(\theta), \quad (\text{C.21})$$

which implies

$$\langle \hat{C}_{\ell}^{\text{EE}} \pm \hat{C}_{\ell}^{\text{BB}} \rangle = \sum_{\ell'} {}_{-2}K_{\ell\ell'} (C_{\ell'}^{\text{EE}} \pm C_{\ell'}^{\text{BB}}). \quad (\text{C.22})$$

Summing or subtracting the last equation for + or - allows us to build unmixed estimators of the polarization power spectra. If we had chosen not to decouple the correlation functions and had built $\hat{C}_{\ell}^{\text{EE}} + \hat{C}_{\ell}^{\text{BB}}$ as the Legendre transform of $\hat{\xi}_+ = g\hat{\xi}_+$, the output PoISpice spectra would follow

$$\langle \hat{C}_{\ell}^{\text{EE}} \pm \hat{C}_{\ell}^{\text{BB}} \rangle = \sum_{\ell'} {}_{\pm 2}K_{\ell\ell'} (C_{\ell'}^{\text{EE}} \pm C_{\ell'}^{\text{BB}}), \quad (\text{C.23})$$

which would leave some mixing in the polarization spectra because ${}_{+2}K \neq {}_{-2}K$.

Appendix D: Multimask analysis

In this section, we generalize our analysis to multiple masks. This situation occurs when a cross-power spectrum analysis with maps with different masks is performed. We restrict ourselves to the study of the intensity case. Writing the expression of the covariance explicitly as in García-García et al. (2019), we obtain the following expression:

$$\begin{aligned} \text{cov}(\tilde{C}_{\ell}^{i,j}, \tilde{C}_{\ell'}^{p,q}) &= \frac{1}{(2\ell + 1)(2\ell' + 1)} \sum_{mm'} [\langle \tilde{a}_{\ell m}^i \tilde{a}_{\ell' m'}^{p*} \rangle \langle \tilde{a}_{\ell m}^j \tilde{a}_{\ell' m'}^q \rangle + (p \leftrightarrow q)], \\ &= \frac{1}{(2\ell + 1)(2\ell' + 1)} \sum_{\ell_1 \ell_2} \sum_{m_1 m_2 m m'} [C_{\ell_1}^{i,p} I_{k k_1} [W^i] I_{k_1 k'} [W^p] \\ &\quad C_{\ell_2}^{j,q} I_{k_2 k} [W^j] I_{k' k_2} [W^q] + (p \leftrightarrow q)], \\ &= \Xi_{\ell\ell'}^{00} [W^i, W^p, W^j, W^q] \sum_{\ell_1 \ell_2} [C_{\ell_1}^{i,p} \bar{\Theta}_{\ell\ell'}^{\ell_1 \ell_2} [W^i, W^p, W^j, W^q] C_{\ell_2}^{j,q} + \\ &\quad (p \leftrightarrow q)]. \end{aligned}$$

We noted $k_i = (\ell_i, m_i)$. We also extended the definition of the kernels to the case multiple masks as

$$\Xi_{\ell\ell'}^{00} [W^i, W^p] (W^j, W^q) \equiv \Xi_{\ell\ell'}^{00} [\mathcal{V}^{(ip) \times (jq)}], \quad (\text{D.1})$$

$$\text{where } \mathcal{V}_{\ell}^{(ip) \times (jq)} \equiv \frac{1}{2\ell + 1} \sum_m [W^i W^p]_{\ell m} [W^j W^q]_{\ell m}^*, \quad (\text{D.2})$$

$$\Theta_{\ell\ell'}^{\ell_1 \ell_2} [W^i, W^p, W^j, W^q] = \sum_{m m' m_1 m_2} I_{\ell m \ell_1 m_1} [W^i] \quad (\text{D.3})$$

$$I_{\ell_1 m_1 \ell' m'} [W^p] I_{\ell' m' \ell_2 m_2} [W^j] I_{\ell_2 m_2 \ell m} [W^q].$$

As long as the considered masks have similar properties, for instance, that none of them includes point sources, the computations developed in this work will hold. In Eq. (E.2), the assumptions hold as long as the mask harmonic coefficients fall quickly enough, which is the case even when the survey area varies a little from one map to the next.

Appendix E: Details of the ACC approximation

E.1. Mathematical validation

We explore the mathematical justification of the ACC approximation. From Fig. 9, the $\bar{\Theta}_{\ell\ell'}^{\ell_1 \ell_2}$ kernel appears only to depend on $\Delta \equiv \ell' - \ell$, $\Delta_1 \equiv \ell_1 - \ell$, and $\Delta_2 \equiv \ell_2 - \ell$. We recall that the normalization of the reduced coupling kernel (Eq. (22)) already approximately only depends on Δ because $\Xi_{\ell\ell'}^{00}$ is close to a Toeplitz matrix (Louis et al. 2020). The kernel itself is given by a summation of products of four coupling coefficients ${}_0I_{\ell m \ell_1 m_1}$. They are expressed as the sum of the mask window function with a product of two Wigner-3j symbols, as shown in Eq. (A.9). We remark that because the mask power spectrum falls relatively fast (Fig. 1), the terms with low L in the sum in Eq. (A.9) contribute mostly. However, we are interested in the cases in which all the other multipoles ℓ, ℓ', ℓ_1 , and ℓ_2 are significantly larger than the width of the mask spectrum. For this reason, all the Wigner-3j symbols in Eq. (A.9) can be replaced by their asymptotic behavior, where in the limit $\ell_i, \ell_j \gg L_i$, we have

$$\begin{pmatrix} \ell_i & \ell_j & L_i \\ m_i & m_j & M_i \end{pmatrix} \approx \frac{(-1)^{\ell_j + m_j}}{\sqrt{2\ell_j + 1}} d_{M_i, (\ell_j - \ell_i)}^{L_i}(\theta) \quad (\text{E.1})$$

(Khersonskii et al. 1988). Here, $\theta = \arccos(-m_j / (\ell_j(\ell_j + 1))^{1/2})$ and $d_{k,m}^j$ are reduced Wigner rotation matrices.

When this approximation is introduced in Eq. (A.9), Eq. (21) reads

$$\begin{aligned} \Theta_{\ell_3 \ell_4}^{\ell_1 \ell_2} &\approx \frac{1}{(4\pi)^2} \sum_{L_i M_i} \Pi_i w_{L_i M_i} \sqrt{2L_i + 1} d_{0, (\ell_{i+1} - \ell_i)}^{L_i}(\pi/2) \\ &\quad \left[\sum_{m_i} d_{M_i, (\ell_{i+1} - \ell_i)}^{L_i} \left(\arccos \frac{-m_{i+1}}{(\ell_{i+1}(\ell_{i+1} + 1))^{1/2}} \right) \right]. \end{aligned} \quad (\text{E.2})$$

Here we have defined $\ell_5 \equiv \ell_1$ for notational purposes. We note that when ℓ_i is large enough, which is the case because $\ell_i, \ell_j \gg L_i$, $\arccos \frac{-m_{i+1}}{(\ell_{i+1}(\ell_{i+1} + 1))^{1/2}}$ explore the $[0, \pi]$ range, and the expression in brackets in the last equation can be seen as a Riemann sum over $\theta \in [0, \pi]$. This expression can be approximated by the integral $\int d_{M_i, (\ell_{i+1} - \ell_i)}^{L_i}(\theta) d\theta$, which only depends on L_i, M_i , and $\ell_{i+1} - \ell_i$. Because L_i, M_i are summed over in Eq. (E.2), we directly see that as soon as ℓ, ℓ', ℓ_1 , and ℓ_2 are significantly larger than the width of the mask spectrum, the coupling kernel behaves as a function of their difference, and we expect that this approximation improves in accuracy with ℓ .

E.2. EE expression

For now, when we ignore BB to EE leakage, we can write

$$\tilde{\Sigma}_{\ell\ell'}^{\text{EEEE}} \approx 2\Xi_{\ell\ell'}^{22} [W^2] [C^{\text{EE}} \cdot \bar{\Theta}_{\ell\ell'}^{++++} [W] \cdot C^{\text{EE}}], \quad (\text{E.3})$$

where we defined the polarized covariance coupling

$$\Theta_{\ell\ell'}^{\ell_1 \ell_2; +++++} = \sum_{m m' m_1 m_2} +I_{\ell m \ell_1 m_1} +I_{\ell_1 m_1 \ell' m'} +I_{\ell' m' \ell_2 m_2} +I_{\ell_2 m_2 \ell m}. \quad (\text{E.4})$$

Based on the relations in [Challinor & Chon \(2005\)](#), this kernel Θ^{++++} can be approximately normalized to

$$\sum_{\ell_1 \ell_2} \Theta_{\ell \ell'}^{\ell_1 \ell_2; +++++} [W] \approx \Xi_{\ell \ell'}^{22} [W^2], \quad (\text{E.5})$$

hence the relation [\(E.3\)](#).