PHYSICAL REVIEW D **88,** 023524 (2013)

Efficiency of pseudospectrum methods for estimation of the cosmic microwave background *B*-mode power spectrum

A. Ferté,[1,2,*] J. Grain,[2,1,†] M. Tristram,[2,3,‡] and R. Stompor[4,§]

[1]*Université Paris-Sud 11, Institut d'Astrophysique Spatiale, UMR8617, Orsay F-91405, France*
[2]*CNRS, Orsay F-91405, France*
[3]*Université Paris-Sud 11, Laboratoire de l'Accélérateur Linéaire, Bâtiment 200, 91898 Orsay Cedex, France*
[4]*AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité 75205, France*

(Received 31 May 2013; published 19 July 2013)

The estimation of the *B*-mode angular power spectrum of polarized anisotropies of the cosmic microwave background is a key step towards a full exploitation of the scientific potential of this probe. In the context of pseudospectrum methods the major challenge is related to a contamination of the *B*-mode spectrum estimate with the residual power of the much larger *E*-mode. This so-called *E*-to-*B* leakage is unavoidably present whenever only an incomplete sky map is available, as is the case for any realistic observation. The leakage has to be then minimized or removed and ideally in such a way that neither a bias nor extra variance is introduced. In this paper, we compare from these two perspectives three different methods proposed recently in this context [K. M. Smith, Phys. Rev. D **74**, 083002 (2006); W. Zhao and D. Baskaran, Phys. Rev. D **82**, 023001 (2010); J. Kim and P. Naselsky, Astron. Astrophys. **519**, A104 (2010)], which we first introduce within a common algebraic framework of the so-called χ fields and then study their performance on two different experimental configurations: one corresponding to a small-scale experiment covering 1% of the sky motivated by current ground-based or balloon-borne experiments, and another to a nearly full-sky experiment, e.g., a possible cosmic microwave background *B*-mode satellite mission. We find that although all these methods allow us to reduce significantly the level of the *E*-to-*B* leakage, it is the method of Smith that at the same time ensures the smallest error bars in all experimental configurations studied here, owing to the fact that it permits straightforwardly an optimization of the sky apodization of the polarization maps used for the estimation. For a satellite-like experiment, this method enables a detection of the *B*-mode power spectrum at large angular scales but only after appropriate binning. The method of Zhao and Baskaran is a close runner-up in the case of a nearly full-sky coverage.

DOI: 10.1103/PhysRevD.88.023524 PACS numbers: 98.80.−k, 98.70.Vc, 07.05.Kf

I. INTRODUCTION

Polarized anisotropies of the cosmic microwave background (CMB) radiation come in two flavors: gradient-like, *E*, and curl-like, *B*, [1,2]. Ten years ago, the first detection of the *E*-mode anisotropies was announced by the DASI team [3]. Since then many subsequent experiments, e.g., WMAP [4], QUAD [5], or BICEP [6] have detected the *E*-mode anisotropies with high significance, deepening and confirming our understanding of the Universe's evolution and structure formation. PLANCK [7] is widely expected to shortly provide the most comprehensive and precise constraints on the *E*-mode polarization properties in a range of angular scales, extending from the largest down to a few arc minutes.

In contrast, no *B*-mode anisotropy has been detected yet and only some upper limits are currently available (see, e.g., Refs. [4–6]). This is expected given the minute amplitudes predicted for this signal. At the same time the scientific potential of the *B*-mode probe has been generally recognized as extremely promising. For instance, on the linear level the *B*-modes can be sourced by the primordial gravitational waves [8,9] and not by the scalar fluctuations, thought to be largely responsible for the observed total intensity and *E*-mode anisotropies. Consequently, a detection of the *B*-mode anisotropy at large angular scales ($\ell \lesssim 100$) in excess of what is expected from the gravitational lensing signal (see below) could be seen as a direct validation of inflationary theories, as the latter are considered to be the most likely source of the gravity waves, and could allow for discrimination between different inflationary models. It could also set useful constraints on the reionization period [10]. At smaller angular scales, *B*-modes are expected to be mainly due to gravitational lensing of CMB photons, which converts *E*-modes into *B*-modes [11] and therefore allows for their detection—a source of constraints on the matter perturbation evolution at redshift $z \sim 1$ when light massive neutrinos and elusive dark energy both play potentially visible roles.

For these reasons, many polarization experiments targeting *B*-modes have been built or proposed, including ground-based observatories, those already operating, e.g., POLARBEAR [12] or SPTPOL [13], those which are being

*agnes.ferte@ias.u-psud.fr
†julien.grain@ias.u-psud.fr
‡tristram@lal.in2p3.fr
§radek@apc.univ-paris-diderot.fr

developed, e.g., QUBIC [14] or ACTPOL [15], balloon-borne experiments such as SPIDER [16] or EBEX [17] (which flew in the winter of 2012/2013), or even a potential satellite mission, such as LiteBIRD [18], COrE [19], or PIXiE [20]. With the exception of the QUBIC experiment, all these experiments scan the sky with one or more dishes and therefore most directly produce maps of the polarized Stokes parameters, Q and U. The calculation of the E and B signals from the Q and U maps is a nonlocal operation [21] and can be done uniquely only if the full sky maps are available. However, this can hardly be the case even for the satellite missions due to the presence of heavy non-cosmological contamination due to Galactic emissions, which typically have to be masked out even after advanced and complex cleaning procedures have been applied. In the context of the pseudospectrum methods [22–24] the incomplete sky coverage leads to the so called E-to-B leakage, when the signal from E-modes is present in the reconstruction of the B-modes' power spectrum C_ℓ^B and—more problematic—in the B-modes' uncertainties. Though no bias is directly introduced, the leakage is a problem due to the much higher amplitudes of the E-modes' signal, which then inflates the overall uncertainty of the estimated B-modes' signal, potentially precluding its detection.

Several extensions of the standard pseudospectrum methods have been recently proposed that are designed to alleviate the E-to-B leakage problem. In this work we focus on the technique presented in Refs. [25–27], which works in the harmonic domain and is referred to as the SZ method hereafter, and on two other techniques operating in the pixel domain presented in Refs. [28–30], referred to as the ZB and KN techniques,[1] respectively. All these methods consist in filtering E-modes leaking into B-modes for each specific realization of the polarized anisotropies and thus potentially resolving the excessive variance problem referred to earlier.

In this article, we first describe each of these methods within a common framework of so-called χ fields and then describe our implementations of them, emphasizing differences and similarities with those proposed in the original papers. Throughout this work we compute spatial derivatives of the sky maps in the harmonic domain. This is in agreement with the original implementations of the considered techniques. We note however that an interesting, pixel-domain alternative has been recently proposed in Ref. [31] and could be exploited in future work. For spectrum estimators we consistently use cross spectra [32], rather than auto spectra, therefore avoiding a need for estimating the instrumental noise spectrum.

We use numerical experiments to test the efficiency of each of these methods in terms of the quality of the C_ℓ^B reconstruction and above all of the resulting uncertainty. The numerical experiments involve two experimental setups: one mimicking a satellite mission (loosely based on EPIC [33]), and the other a balloon-borne instrument (inspired by EBEX [34]). We note that these kinds of analyses of satellite-mission-like setups are largely absent in the literature, which predominantly has focused on small-sky cases only. Though other techniques, e.g., maximum-likelihood-based power spectrum estimators, may better address some of the problems faced by nearly full-sky observations, the performance of the pseudospectrum methods in this regime is clearly of practical importance.

The general pseudospectrum formalism, as well as its standard and extended renditions relevant for this work, are introduced in Sec. II. An overview of the methods and their implementations can be found in Sec. III. The numerical results are given in Sec. IV, which also presents the case for the SZ method as the one which gives the smallest variances while avoiding a bias. More extensive conclusions are then given in Sec. VII, while technical details are deferred to the appendices, with Appendix C treating the problem of the noise bias for the ZB and KN methods.

II. PSEUDOSPECTRUM POLARIZED POWER SPECTRUM ESTIMATORS

A. General considerations

The linearly polarized CMB polarization field is completely described by spin-2 and spin-(-2) fields, $P_{\pm 2}(\vec{n}) = Q(\vec{n}) \pm iU(\vec{n})$, with Q and U denoting two Stokes parameters. Pseudospectrum methods distill the observed information into a set of harmonic coefficients, $\tilde{a}^E_{\ell m}$ and $\tilde{a}^B_{\ell m}$, referred to as pseudomultipoles. These are related to *true* multipoles, $a^E_{\ell m}$ and $a^B_{\ell m}$, as follows:

$$\tilde{a}^E_{\ell m} = \sum_{\ell' m'} \Big[H^{(+)}_{\ell m,\ell' m'} a^E_{\ell' m'} + i H^{(-)}_{\ell m,\ell' m'} a^B_{\ell' m'} \Big], \tag{1}$$

$$\tilde{a}^B_{\ell m} = \sum_{\ell' m'} \Big[-i K^{(-)}_{\ell m,\ell' m'} a^E_{\ell' m'} + K^{(+)}_{\ell m,\ell' m'} a^B_{\ell' m'} \Big], \tag{2}$$

where $H^{(\pm)}$ and $K^{(\pm)}$ are kernels, which in general can all be different, nonvanishing, and nondiagonal in both ℓ and m. Noise terms have been neglected in these equations for shortness.

The kernels are typically singular and it is not in general possible to solve the inverse problem to recover the true multipoles, $a^X_{\ell m}$, directly. Instead the pseudospectrum approaches attempt to do so only on the power-spectrum level. This is achieved in two steps. First, owing to the statistical isotropy of CMB fluctuations, we can rewrite Eqs. (1) and (2) on the power spectrum level as

$$\langle \tilde{C}^E_\ell \rangle = \sum_{\ell'} \Big[H^{(+)}_{\ell\ell'} \langle C^E_{\ell'} \rangle + H^{(-)}_{\ell\ell'} \langle C^B_{\ell'} \rangle \Big], \tag{3}$$

$$\langle \tilde{C}^B_\ell \rangle = \sum_{\ell'} \Big[K^{(-)}_{\ell\ell'} \langle C^E_{\ell'} \rangle + K^{(+)}_{\ell\ell'} \langle C^B_{\ell'} \rangle \Big], \tag{4}$$

[1]The methods' names are based on the first letters of the names of the authors of the corresponding papers.

where the new kernels $X^{(\pm)}_{\ell\ell'}$ are given by ($X = K, H$)

$$X^{(\pm)}_{\ell\ell'} = \sum_{m'=-\ell'}^{\ell'} \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |X^{(\pm)}_{\ell m,\ell' m'}|^2, \tag{5}$$

where $\langle \ldots \rangle$ denotes an ensemble average and

$$\tilde{C}^X_\ell \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |\tilde{a}^X_{\ell m}|^2. \tag{6}$$

The kernels obtained on the power-spectrum level are clearly more manageable and easier to calculate; nevertheless, they still will be singular. To avoid this issue, the inverse problem defined in Eqs. (3) and (4) is solved only for binned spectra [24],

$$\tilde{C}^X_b \equiv \sum_\ell P_{b\ell} \tilde{C}^X_\ell, \qquad C^X_b \equiv \sum_\ell P_{b\ell} C^X_\ell, \tag{7}$$

where the binning operators are defined as

$$P_{b\ell} = \begin{cases} \frac{S_\ell}{\ell^b_{\max} - \ell^b_{\min}}, & \ell \in [\ell^b_{\min}, \ell^b_{\max}], \\ 0, & \ell \notin [\ell^b_{\min}, \ell^b_{\max}], \end{cases}$$

$$Q_{b\ell} = \begin{cases} \frac{1}{S_\ell}, & \ell \in [\ell^b_{\min}, \ell^b_{\max}], \\ 0, & \ell \notin [\ell^b_{\min}, \ell^b_{\max}], \end{cases}$$

therefore satisfying the relation $\sum_\ell Q_{b\ell} P_{b'\ell} = \delta_{bb'}$. Here, we have introduced a shape function, S_ℓ. Its role is to minimize possible binning effects by making $S_\ell \tilde{C}$ nearly flat within the bin. Hereafter, we will adopt the standard choice for it, i.e., $S_\ell = \ell(\ell+1)/2\pi$. The binned version of Eqs. (3) and (4) now reads

$$\begin{pmatrix} \tilde{C}^E_b \\ \tilde{C}^B_b \end{pmatrix} \simeq \sum_{b'} \begin{pmatrix} H^{(+)}_{bb'} & H^{(-)}_{bb'} \\ K^{(-))}_{bb'} & K^{(+)}_{bb'} \end{pmatrix} \begin{pmatrix} C^E_{b'} \\ C^B_{b'} \end{pmatrix}, \tag{8}$$

where, for $X = K$ or H,

$$X_{bb'} \equiv \sum_{\ell,\ell'} P_{b\ell} X_{\ell\ell'} Q_{b'\ell'}. \tag{9}$$

To include a correction for the presence of the instrumental noise, the pseudopower spectrum on the right-hand side of the first of Eqs. (7) needs be corrected for the noise pseudospectrum prior to the binning operations.

The estimates of the true spectra, C^X_ℓ, can then be obtained by directly solving the full system in Eq. (8). We note that by construction, and neglecting the binning effects, which are largely controllable, these will be unbiased estimates of the true binned spectra. However, as long as the polarization mode-mixing kernel, $K^{(-)}$, does not vanish[2] the power contained in the E-polarization component will contribute to the overall variance of the B-spectrum estimate—an effect referred to as the E-to-B leakage. To avoid this one should resort to methods for which $K^{(-)}$ is either zero or nearly so. We also note that if $K^{(-)} = 0$ then the estimate of the B-mode spectrum can be derived independently of the E one. This could also be the method of choice even if $K^{(-)}$ vanishes only approximately. In this case, however, a small bias in the B-spectrum estimate is to be expected.

B. Standard pseudospectrum approach

If the polarization fields are known on the entire celestial sphere, their E and B representation can be easily obtained in the harmonic domain using the spin-weighted spherical harmonics,[3]

$$\begin{aligned} a^E_{\ell m} &= \frac{-1}{2} \int [P_2(\vec{n})_2 Y^\star_{\ell m}(\vec{n}) + P_{-2}(\vec{n})_{-2} Y^\star_{\ell m}(\vec{n})] d\vec{n}, \\ a^B_{\ell m} &= \frac{i}{2} \int [P_2(\vec{n})_2 Y^\star_{\ell m}(\vec{n}) - P_{-2}(\vec{n})_{-2} Y^\star_{\ell m}(\vec{n})] d\vec{n}. \end{aligned} \tag{10}$$

If the polarization field is measured on a fraction of the sky only, the above decomposition can be most straightforwardly applied to such a case by positing that the signal over the unobserved part of the sky vanishes. This choice defines the standard pseudospectrum method, in which the resulting pseudomultipoles, $\tilde{a}^X_{\ell m}$, $X = E, B$, can be expressed as follows:

$$\begin{aligned} \tilde{a}^E_{\ell m} &\equiv \frac{-1}{2} \int M [P_2(\vec{n})_2 Y^\star_{\ell m}(\vec{n}) + P_{-2}(\vec{n})_{-2} Y^\star_{\ell m}(\vec{n})] d\vec{n} \\ &= \sum_{\ell' m'} [K^{(+)}_{\ell m,\ell' m'} a^E_{\ell' m'} + i K^{(-)}_{\ell m,\ell' m'} a^B_{\ell' m'}], \end{aligned} \tag{11}$$

$$\begin{aligned} \tilde{a}^B_{\ell m} &\equiv \frac{i}{2} \int M [P_2(\vec{n})_2 Y^\star_{\ell m}(\vec{n}) - P_{-2}(\vec{n})_{-2} Y^\star_{\ell m}(\vec{n})] d\vec{n} \\ &= \sum_{\ell' m'} [-i K^{(-)}_{\ell m,\ell' m'} a^E_{\ell' m'} + K^{(+)}_{\ell m,\ell' m'} a^B_{\ell' m'}], \end{aligned} \tag{12}$$

where M is a binary mask defining the observed patch, and where we introduced the convolution kernels, $K^{(\pm)}_{\ell m,\ell' m'}$, explicit expressions for which are well known and can be found elsewhere, e.g., Ref. [27]. We see that for the standard technique both the $H^{(\pm)}$ and $K^{(\pm)}$ kernels, Eqs. (1) and (2), coincide and that the polarization-mode-mixing kernel, $K^{(-)}_{\ell m,\ell' m'}$, does not vanish and therefore—though unbiased—the standard pseudopower spectrum estimator suffers from the E-to-B leakage. This can be quite severe. For instance, an experiment covering around 1% of the sky is essentially unable to detect a power at scales larger than $\ell \lesssim 140$ (see Fig. 16 of Ref. [27]).

[2]Strictly speaking, what is required is that the multipole kernel, $K^{(-)}_{\ell m,\ell' m'}$ vanishes, but if Eq. (5) is satisfied, exactly or approximately, it is equivalent to requiring that the power spectrum kernel, $K^{(-)}_{\ell\ell'}$, be (nearly) zero.

[3]All the integrals in this paper are taken over the *entire* celestial sphere. We therefore do not specify that the integration domain is S^2.

The above formulas can be extended to include an arbitrary weighting of the observed sky pixels as given by a window function, W. This can be done by inserting WM instead of M in all the equations above, including those for the kernels. If we further assume that the window function is always zero outside of the observed sky, i.e., if $M = 0$ then also $W = 0$, then, as a consequence, $WM = W$ and M can be dropped from the equations in favor of W. The mask, M, is then assumed to be defined implicitly by W. We will use this simplification in the following. Also, for definiteness hereafter, we assume that a field defined on the sphere, e.g., $P_{\pm 2}$, is known on the full sky and will apply a mask or an apodization explicitly to such a field to emphasize that it is known only over a limited sky area, e.g., $WP_{\pm 2}$.

C. Leakage-free pseudopower spectrum approaches

To alleviate the leakage problem within the pseudo-spectrum methods one would need to adapt a different definition of the pseudomultipoles than the one used in the standard approach. Such a new definition should not rely directly on the polarization fields, as does the standard approach, as those unavoidably incorporate contributions from both types of polarized multipoles. Instead it should be based on some other fields, which depend only on one set of the multipole coefficients, and which would therefore ensure that the polarization-mode-mixing kernels, $K^{(-)}_{\ell m,\ell' m'}$ and $H^{(-)}_{\ell m,\ell' m'}$, indeed vanish, resolving the leakage issue.

Such a construction has been proposed by Ref. [1] and the corresponding fields are called χ fields. They can be derived from the polarization fields as follows:

$$\chi^E(\vec{n}) = -\frac{1}{2}[\bar{\partial}\,\bar{\partial}\,P_2(\vec{n}) + \partial\partial P_{-2}(\vec{n})], \tag{13}$$

$$\chi^B(\vec{n}) = \frac{i}{2}[\bar{\partial}\,\bar{\partial}\,P_2(\vec{n}) - \partial\partial P_{-2}(\vec{n})], \tag{14}$$

where $\partial(\bar{\partial})$ denotes the spin-raising(lowering) operator [1]. These χ fields indeed involve either E-modes (in the case of χ^E) or B-modes (for χ^B). This can be seen directly by noting that the χ^X fields, $X = E, B$, are scalars and given by

$$\chi^X(\vec{n}) = \sum_{\ell,m} N_{\ell,2} a^X_{\ell m} Y_{\ell m}(\vec{n}), \tag{15}$$

where for future convenience we have introduced

$$N_{\ell,s} \equiv \sqrt{\frac{(\ell+s)!}{(\ell-s)!}}.$$

In the full-sky case, Eq. (15) can be readily inverted, giving

$$\chi^X_{\ell m} = \int \chi^X(\vec{n}) Y^\star_{\ell m}(\vec{n}) d\vec{n} = N_{\ell,2} a^X_{\ell m}, \tag{16}$$

which in turn can be adapted for cases of partial-sky experiments in a usual manner, rendering the following definition of the pseudomultipoles:

$$\tilde{a}^X_{\ell m} \equiv \frac{1}{N_{\ell,2}} \int M(\vec{n}) \chi^X(\vec{n}) Y^\star_{\ell m}(\vec{n}) d\vec{n}. \tag{17}$$

This definition can then be used in the general pseudo-spectrum formalism developed in Sec. II, and though it will result in a mixing of different ℓ modes it will not cause any leakage between the polarization modes as by construction the off-diagonal kernels—$H^{(-)}$ and $K^{(-)}$ in Eqs. (1) and (2)—vanish.

The major difficulty of this approach is the computation of the χ^X fields. Indeed, Eqs. (13) and (14) as they are require in principle knowledge of the full-sky polarization fields. As we will see in the next section all three methods designed to resolve the leakage problem and studied in this work rely on the χ^X-field calculation, implicitly or explicitly, and circumvent the problem of having only a limited sky coverage differently.

We note that if the χ^X fields were known exactly on the cut sky, the inverse problem in Eq. (8) could be solved separately for E and B spectra, as the off-diagonal kernels would, by construction, vanish. In more realistic circumstances the χ^X fields actually estimated on the cut sky may be imperfect giving—at least in principle—rise to nonzero off-diagonal contributions. These, if not corrected for, could lead to a bias of the estimated power spectra. Solving the full system and accounting for the nondiagonal kernels could help to trade the bias for an extra—but presumably small—variance of the spectrum estimate. Though this indeed could be possible at least for some of the methods, for others, the difficulty in calculating the off-diagonal kernels—either analytically or numerically, e.g., via Monte Carlo (MC) simulations—can be prohibitive, and an approach favored in practice is often simply to accept the bias once it is found to be sufficiently small.

III. SPECIFIC APPROACHES

A. SZ approach

1. Theoretical description

Let us start from the pseudomultipoles for B-modes as defined in Eq. (17) with the binary mask, M, replaced by an arbitrary window, W. By performing an integration by parts twice [25,26], we can rewrite this equation as

$$\tilde{a}^B_{\ell m} = \frac{i}{2N_{\ell,2}} \int d\vec{n}[P_2(\vec{n}) \times (\partial\partial W(\vec{n}) Y_{\ell m}(\vec{n}))^\star - P_{-2}(\vec{n}) \times (\bar{\partial}\,\bar{\partial}\, W(\vec{n}) Y_{\ell m}(\vec{n}))^\star], \tag{18}$$

where all the boundary terms are omitted corresponding to an assumption that the apodization window $W(\vec{n})$ and its first derivative ∂W vanish at the observed patch boundaries. This latter equation has an advantage over the former,

Eq. (17), as it does not involve any explicit calculation of derivatives of noisy sky maps. Instead, the differentiation needs to only be applied to a presumably smooth window function, W. We can therefore use Eq. (18) as a definition of the pseudomultipoles, which we will also apply from now on in cases when the apodization does not conform with the boundary conditions. Note that in these latter cases there will be no assurance that no E-to-B leakage is present.

Hereafter we will refer to this technique as a pure pseudospectrum estimator, as Eq. (18) can be interpreted as projecting the polarization field $P_{\pm 2}$ onto a basis of "pure" functions representing only B-like polarization modes on a cut sky [25,26,35].

2. Numerical implementation

Our implementation of the approach follows closely that proposed in Ref. [27] and proceeds in four steps.

Step 1:

We compute spin-0, spin-1, and spin-2 renditions of the window function W given by

$$W_0 = W, \qquad W_1 = \partial W, \qquad W_2 = \partial\partial W. \tag{19}$$

Because W is real, then $W_s^\star = W_{-s}$ for a spin $s = 1, 2$.

Step 2:

We compute pure pseudomultipoles by first constructing three apodized maps,

$$\mathcal{P}_{\pm 2} = W_0 P_{\pm 2}, \qquad \mathcal{P}_{\pm 1} = W_{\mp 1} P_{\pm 2}, \qquad \mathcal{P}_{\pm 0} = W_{\mp 2} P_{\pm 2}, \tag{20}$$

and then calculating the pure $\tilde{a}^B_{\ell m}$ as

$$\tilde{a}^B_{\ell m} = \frac{1}{N_{\ell,2}} (\mathcal{B}_{0,\ell m} + 2N_{\ell,1}\mathcal{B}_{1,\ell m} + N_{\ell,2}\mathcal{B}_{2,\ell m}), \tag{21}$$

where $\mathcal{B}_{s,\ell m}$ is a B-type mutlipole of $\mathcal{P}_{\pm s}$ defined as

$$\mathcal{B}_{s,\ell m} = \frac{i}{2} \int [\mathcal{P}_{+s}(\vec{n})\, {}_sY^\star_{\ell m}(\vec{n}) - (-1)^s \mathcal{P}_{-s}(\vec{n})\, {}_{-s}Y^\star_{\ell m}(\vec{n})] d\vec{n}. \tag{22}$$

Step 3:

In this step we compute the convolution kernels for pseudo-C_ℓ as defined in Eqs. (3) and (4). This can be done using, e.g., Eqs. (A13) and (A14) of Ref. [27]. If the applied apodization does not fulfill the boundary conditions then the off-diagonal block $K^{(-)}$ has to be included as well. In practice, the off-diagonal coupling between the polarization components will also occur due to pixelization effects. Though such effects are not accounted for in the analytic formulas for the kernels, they can be corrected for, to some extent, by a procedure described in Ref. [27], leading to a removal of the majority of the small bias induced by the residual, pixel-induced E-to-B leakage.

We note that typically, if the method is applied consistently to both E- and B-modes, the corresponding H and K kernels are identical. However, in some circumstances it may be advantageous and possible to apply hybrid approaches in which both kinds of spectra are treated differently. Such cases have been discussed recently in Ref. [36].

Step 4:

This step consists of standard operations involved in any pseudospectrum method as summarized by Eqs. (7) and (8) and discussed in Sec. II.

3. Sky apodization

As emphasized in Refs. [25–27], an appropriate sky apodization is a key element of such a construction. In the specific method discussed here the degree to which the apodization fulfills the boundary conditions will be a principal factor determining the level of suppression of the E-to-B leakage. At the same time any apodization applied to realistic (meaning noisy) data will have a direct impact on the resulting uncertainties of the spectrum estimate. In the context of the pure pseudospectrum method, systematic approaches have been developed and studied in detail, which allow for a numerical optimization of sky apodizations in order to ensure a nearly minimal value of the final spectrum uncertainty [25–27]. These are either based on MC simulations or semianalytic techniques. In the former case, MC simulations are used to tune the length of the sky apodization given by some analytic formulas. In this work, we will use the so-called C^2 function as given by Eq. (31) of Ref. [27]. In the latter case, the optimized sky apodization can be computed by solving a large linear system as proposed in Ref. [26]. We refer to these latter windows as *variance-optimized apodization*. In both cases the optimization could, and should, be applied bin-by-bin to ensure the best results. As discussed at length in Ref. [27] both of these approaches require some prior assumptions concerning, for instance, the angular power spectra of E- and B-modes; however, the results of the optimization are found to be only mildly dependent on details of the assumed B-mode spectrum.

It has been shown via numerical experiments [27] that the variance-optimized apodizations lead systematically to the lowest error bars on the reconstructed C_ℓ^B's, and therefore we will use them in this work. These variance-optimized apodizations can be computed in two ways, depending on the domain (harmonic domain or pixel domain) in which the linear system is solved. For the peculiar case of homogeneous noise, resolution can be done in the harmonic domain. In such a case, the derivative relationship $W_{s=1,2} = \partial^s W_0$ and the boundary conditions $W_0(\mathcal{C}) = W_1(\mathcal{C}) = 0$ on the contour of the observed region are fulfilled (up to pixelization effects). For more general cases, the linear system providing the variance-optimized apodization is solved in the pixel domain. In such a setting, both the derivative relationship and the boundary

conditions are relaxed (W_0, W_1, and W_2 are considered as independent). As a consequence, the final sky apodization does not strictly satisfy these conditions and the resulting pseudomultipoles will not be strictly equal to the pure pseudomultipoles. However, it was shown in Refs. [26,27] that the angular power spectra recovered in such cases consistently achieve smaller uncertainties than those of other apodization choices.

B. ZB approach

1. Theoretical description

In this approach the χ^X fields are computed directly in the pixel domain and for the cut sky. This is made possible thanks to a formula derived in Ref. [28], which reads

$$\begin{aligned} W(\vec{n})\chi^B(\vec{n}) = &\frac{i}{2}[\bar{\partial}\,\bar{\partial}(WP_2) - \partial\partial(WP_{-2})] \\ &- i\left[\frac{\bar{\partial}W}{W}\,\bar{\partial}(WP_2) - \frac{\partial W}{W}\,\partial(WP_{-2})\right] \\ &- \frac{i}{2}[(\bar{\partial}\,\bar{\partial}\,W)P_2 - (\partial\partial W)P_{-2}] \\ &+ i\left[\frac{(\bar{\partial}W)^2}{W}P_2 - \frac{(\partial W)^2}{W}P_{-2}\right]. \end{aligned} \tag{23}$$

As usual, here W is assumed to be zero outside the observed region. Moreover, if we assume that it and its first derivative vanish at the edges of the observed region, all the operations on the right-hand side of this equation can be performed with only the knowledge of the polarization field on the cut sky. Consequently, we could estimate the field χ^B consistently on the cut sky by first computing the rhs of Eq. (23) and then dividing it by the window W, and later use it to calculate pseudomultipoles via Eq. (17)—as proposed in Ref. [28]—or use some apodized rendition of the χ^B field to derive the pseudomultipoles, which are then corrected on the power spectrum level—as proposed here.[4] In either case the pseudomultipoles are in principle free of any E-to-B leakage due to cut-sky effects and the $K^{(-)}_{\ell\ell'}$ kernel should vanish. However, as emphasized by Ref. [28], both pixelization and convolution by the beam lead to some residual E-to-B leakage and ideally one would like to solve the full linear system, Eq. (8), to get the final, unbiased power-spectrum estimation.

2. Numerical implementation

An implementation of this technique was proposed in Ref. [28] and involves four steps. The implementation used in this work follows that of the original authors with the exception of the second step, as detailed below.

Step 1:

We compute the χ^B field on the observed patch of the sky using Eq. (23). This in turn requires a numerical calculation of derivatives of noisy fields, which constitutes the principal difficulty of this technique. In our implementation as well as that of Ref. [28] these calculations are performed in the harmonic domain. We emphasize that with such a choice this method becomes effectively a harmonic space approach. Yet another potential problem is related to the calculation of the terms, which involve explicit multiplication by $W^{-1}\partial W$, as W itself becomes very small at the boundary. This problem cannot be avoided by imposing more boundary conditions on W as $W^{-1}\partial W \sim |\theta - \theta_c|^{-1}$ at the boundary θ_c, and therefore necessarily diverges at the boundary.[5] This can however be dealt with in Step 2.

Step 2:

We compute the pseudomultipoles $\tilde{a}^B_{\ell m}$ of the newly constructed χ^B map. This requires effectively dividing by the window W. Though straightforward *a priori*, care has to be exercised while doing this because W vanishes at the observed area edges.

One option, adopted in Ref. [28], relies on simply trimming the troublesome boundary layer, leaving only those pixels for which the division is numerically reliable. This leads to some loss of the information but solves simultaneously the divergence problem appearing in Step 1. The amount lost due to trimming will depend on the details of how the trimming is done—a practical complication, which needs to be addressed in this approach.

An alternative way of resolving both these issues at the same time, which we propose here and which is free of such practical complications, is to define pseudomultipoles using the field $W^2\chi^B$ and then to correct for the presence of the apodization in the binned spectrum estimation step, Eq. (8). It is clear from Eq. (23) that the estimation of the $W^2\chi^B$ field does not suffer any singularities at the edges. This method is the method of choice in this work.

We note that this method is not lossless either, as the apodization it invokes will unavoidably compromise some information. Nevertheless, the information loss in this case is expected to be smaller than in the former one. For instance, it was argued in Sec. IV of Ref. [28] that to analyze a map covering 3% of the sky (a spherical cap with a radius of 20 degrees is assumed as the observed part of the sky), it is necessary to remove an external layer with a width of 2 degrees, thus reducing the effective sky coverage from 3% to 2.4% (assuming a binary mask to weight the resulting χ^B map). As shown hereafter, by focusing on $W^2\chi^B$ we are able to solve for the E-to-B leakage by using an apodization length of 1 degree. As a consequence, for a

[4]Strictly speaking, the pseudomultipoles are not divided by $N_{\ell,2}$ in the implementation of Ref. [28]. Instead, the pseudospectrum are divided by $N^2_{\ell,2}$ in the binning process. The two choices are however completely equivalent.

[5]By constraining W together with its first derivative ∂W to be continuous on the entire celestial sphere but zero outside the observed part of the sky necessarily leads to $W \sim |\theta - \theta_c|^n$ with $n \geq 1$, close to the boundary.

spherical cap with a radius of 20 degrees, the effective sky coverage is reduced from 3% to 2.9% (an explicit expression for the effective sky coverage assuming a nonbinary mask can be found in Ref. [36]).

Step 3:

The kernel $K^{(+)}_{\ell\ell'}$ is computed taking advantage of the fact that the χ^B field is a scalar—like temperature—made of B-modes. The explicit expression of $K^{(+)}_{\ell\ell'}$ is given by Eq. (39) of Ref. [28] (following what was derived for temperature [22,24,32,37]), i.e.,

$$K^{(+)}_{\ell\ell'} = \frac{(2\ell'+1)N^2_{\ell',2}}{4\pi N^2_{\ell,2}} \sum_{\ell'' m''} |w^{(2)}_{\ell'' m''}|^2 \begin{pmatrix} \ell & \ell' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2, \qquad (24)$$

with $w^{(2)}_{\ell'' m''}$ being the multipoles of the W^2 function.[6]

Step 4:

The linear system in Eq. (8) is inverted, neglecting the off-diagonal block $K^{(-)}_{\ell\ell'}$, and therefore also the residual E-to-B leakage.

3. Sky apodization

In this approach we could either use analytic windows or the variance-optimized windows obtained from the optimization procedure developed within the framework of the SZ method. In this former case, we will always use the C^2 family of windows from Ref. [27] and use MC simulations to determine their optimal apodization length.

In the case of the variance-optimized apodizations computed in the harmonic domain, it may appear that to ensure their optimality we should use a window given by a square root of the actual optimized one, i.e., $W_{\rm ZB} \equiv \sqrt{W_{\rm SZ}}$, to compensate for the fact that it is a square of the window which is used as the apodization in our implementation of the ZB approach. Whether such a window could be a viable option will depend on whether it does not cause any problems in the calculation of the rhs of Eq. (23) at the patch edges. It is straightforward to show that this is always the case for windows, which are forced to obey the boundary conditions strictly. This is because such windows scale at the boundary as $W_{\rm SZ} \sim |\theta - \theta_c|^n$, with $n > 2$ [26], and therefore both quantities, $\partial W_{\rm ZB}$ and $\partial\partial W_{\rm ZB}$ (where $W_{\rm ZB} = \sqrt{W_{\rm SZ}}$), needed to compute the rhs of Eq. (23) are well behaved for $\theta \sim \theta_c$. However, the variance-optimized windows fulfill the boundary condition only approximately, which may lead to singularities of the derivatives of $W_{\rm ZB}$. To avoid this, we further multiply the variance-optimized windows by some analytic window with a narrow apodizaton length. This is designed to affect as little as possible the properties of the initial window but enforce the boundary conditions strictly and therefore ensure the proper behaviour of the resulting window at the boundary. In practice, we have found that using either the corrected $W_{\rm ZB}$ window or directly using $W_{\rm SZ}$ leads to comparable results, and numerical results presented hereafter use the latter ones.

It is important to notice that in such settings the variance-optimized windows computed in the pixel domain cannot be directly applied. Indeed, such windows do not typically conform with the derivative relationship between the different windows, i.e., $W_{s=1,2} \neq \partial^s W_0$ or the boundary conditions, i.e., $W_0(\mathcal{C}) \neq W_1(\mathcal{C}) = 0$. However, these conditions are essentially mandatory for the ZB method for two reasons. First, the method requires that $W\chi^B$ is related to $W_{s=0,1,2}P_{\pm 2}$ and $MP_{\pm 2}$, as e.g., it is in Eq. (23), which however without the assumptions about the window properties is rather tedious. Second, even if such an expression is found, this will lead to mixing kernels, which will not be numerically computable from the "first principles," as in, e.g., Eq. (24), as they will involve the product of three functions : $P_{\pm 2}$ multiplied by either W_0 or M, and by $W_{s=1,2}$. Therefore, this leaves time-consuming Monte Carlos as the only viable option for their estimation.

C. KN approach

1. Theoretical description

Another way of estimating the χ^B field is by generalizing its definition to the cut-sky case. This can be done straightforwardly by modifying Eq. (14) as follows:

$$\tilde{\chi}^B(\vec{n}) = \frac{i}{2}[\bar{\partial}\,\bar{\partial}\,MP_2(\vec{n}) - \partial\partial MP_{-2}(\vec{n})], \qquad (25)$$

where as usual M stands for a binary mask and the tilde over the χ symbol is used to emphasize that at least in principle this is a different object than the true χ^B field defined on the cut sky, i.e., $M\chi^B$. We note however that as long as M is constant (and for simplicity assumed to be equal to 1), i.e., in the interior of the observed patch, the two fields are indeed identical, $\tilde{\chi}^B(\vec{n}) = \chi^B(\vec{n})$. Therefore, in principle the only problem arises at the patch edges. As proposed in Ref. [29] one could use this observation to reconstruct the true χ^B field everywhere with the exception of the boundary layer. The problem then becomes technical and boils down to a question how to calculate the derivatives required by such a procedure. Reference [29] proposed doing it in the harmonic domain and using the semianalytic formulas of Ref. [21] to represent the derivatives via convolutions of some geometrical kernels. Given that the mask falls abruptly from 1 to 0 at the edges, it is not surprising that such a procedure leads to significant oscillatory behavior at the edges, which extends well within the center part of the observed patch. This is a result of the necessity of imposing a finite band limit on all harmonic decompositions performed as part of this procedure, even if the considered functions with an abrupt jump do not have such a limit. Such a band limit is directly related to the pixelization used to represent the polarization fields. This has two practical consequences. First, a robust criterion has

[6]We stress that the multipoles of W^2 are *not* equal to the square of the multipoles of W.

to be found that decides which pixels are to be retained, i.e., which are sufficiently clean of any E-mode contamination. Second, the loss of area is expected to be rather significant. We refer the reader to Ref. [29] for more details of this specific implementation.

A more robust approach would either invoke different ways of calculating the derivatives, e.g., as proposed by Ref. [31], or by introducing into Eq. (25) a smooth apodization W in place of the binary mask M. This second option was proposed by Ref. [30] and this is the one we implement in this work. The apodization could alleviate the pixelization effects described earlier by truncating the band limit of the apodized polarization field so the harmonic domain derivatives perform better. Such a window would need to have a central region where W is constant (and equal to 1) before smoothly rolling off at the edges. As in the case where the binary mask is only in this central region the reconstructed $\tilde{\chi}^B$ field would coincide with the true one and would be used for the power-spectrum estimation.

The main advantage of such a technique is that it provides a clear criterion for which pixels to retain or reject. Nevertheless, it does not completely solve the pixelization effects as pixels inside the central area can be affected by the pixel-induced leakage, but this time originating from the contour around this central area. However, and as numerical results shown in Ref. [30] suggest, the pixelization effects at the inner contour are mitigated by the fact that W is continuous as compared to the pixelization effects induced by considering the noncontinuous binary mask.

Hereafter, we will use this second approach and apply a sky apodization to the polarization field. We will then use Eq. (25), but with a mask M replaced by a window W to calculate $\tilde{\chi}^B$ and later the true $\chi^B = \tilde{\chi}^B|_{M_{\chi^B}}$, where M_{χ^B} is the binary mask built from the kept-in-the-analysis pixels, i.e., pixels for which W is essentially constant.

2. Numerical implementation

The numerical implementation of this approach consists then of two main steps, which need to be applied first to simulated data and later to actual data. The Monte Carlo simulations are employed to select optimal windows for a given problem.

Step 1:

We calculate the apodized $\tilde{\chi}^B$ field for a selected window W. This involves performing numerical derivatives of the available polarization fields $P_{\pm 2}$, and these are performed in the harmonic domain. In this work we use a family of arch-sine windows as defined in Ref. [27] with an apodization length which is to be tuned via Monte Carlo simulations. The criteria we use in the apodization-length optimization process are the level of the B-spectrum bias and variance.

Step 2:

We compute the B-mode power spectrum from the precomputed $\tilde{\chi}^B$ field. The spectrum is computed using only the trimmed, central part of the available patch, $\tilde{\chi}^B|_{M_{\chi^B}}$, which can be further apodized if needed, and follows the general pseudospectrum-method framework. Hereafter, following Ref. [29] we will neglect possible leakages from the E spectrum and use the scalar kernel as it is used in the ZB approach [Eq. (24)]. We note however that unlike in the ZB method the leakage in this approach can be more pervasive, affecting even the most central areas of the patch, and therefore is never fully removed via simple area trimming. For this reason one may ponder whether a more appropriate kernel cannot be derived that could account for these effects. The answer, which we discuss in more detail in Appendix A, is that such kernels would need to be evaluated numerically and would be necessarily very costly. We will therefore only consider the simplified case in this work.

3. Sky apodization

The sky apodization and masking needs to be performed in three different stages in this approach. First, we need to apodize the maps before computing the $\tilde{\chi}^B$ field. Then we need to mask pixels, which are expected to be contaminated by the residual E-to-B leakage. Finally, we may want to apodize the reduced $\tilde{\chi}^B$ maps to better localize the bin-to-bin correlations of the recovered B spectrum.

Unlike in the case of the SZ and ZB techniques, one cannot here derive some optimal windows from "first principles." Instead, for the sky apodization required for the computation of $\tilde{\chi}^B$ we use a family of the arch-sine analytic windows, proposed in Ref. [27], and resort to Monte Carlo simulations to optimize their apodization length. In this optimization procedure we always trim all the pixels within the boundary layer of W, i.e., where it is not constant, as these are the pixels which are unavoidably affected by the E-to-B leakage, and we use only the remaining ones for the spectrum estimation. Clearly, there will still be some level of the E-mode power in the map left over after such a trimming procedure, mostly due to pixel-induced E-to-B leakage. The level of this leakage depends on the assumed apodization length, becoming slower for its larger values, and the MC simulations are then used to find the smallest value of the latter ensuring a sufficiently low level of the leakage. This will at the same time maximize the sky area—given the acceptable leakage requirement—left for the final spectrum determination and therefore ensure that the spectrum variance is the smallest.

D. Brief appraisal

The three methods considered in this work can be introduced within a common framework based on the χ^B-field concept, as has been done in this section where it was demonstrated that all the methods rather closely related—a fact which may be potentially somewhat surprising given their original derivations.

The two first methods, SZ and ZB, in the renditions considered in this paper are clearly equivalent on the analytical level if the apodizations employed in both these cases are related to each other as $W_{\rm ZB} = \sqrt{W_{\rm SZ}}$ and if $W_{\rm SZ}$ fulfills strictly the boundary conditions. The differences between these two approaches are therefore only in their numerical implementations and approximations that they imply. Both these methods suffer due to pixelization issues, in particular arising due to a need to compute numerical derivatives, and which give rise to a residual contamination of the B spectrum with the E-mode power. The SZ method requires only derivatives of the window functions; therefore, at least in the cases when these are given analytically, it is possible to estimate the nondiagonal coupling kernel $K^{(-)}$ and correct for some of those effects. Such corrections are more difficult in the case of the ZB approach, where the nondiagonal kernel would have to be estimated completely numerically. The SZ method can potentially offer more freedom for an optimization of the B-spectrum variance as estimated for realistic noisy maps as the boundary conditions on the applied apodizations can be relaxed, leading to an increase of the signal variance that allows for some E-to-B leakage but a decrease of the total signal + noise one. At the same time the off-diagonal polarization-mode coupling kernels can be readily calculated and the estimated B spectrum will be unbiased.

The KN approach can be looked at as an approximation of the ZB method. Indeed, the first term on the rhs of Eq. (14) used by the ZB method coincides with the rhs of Eq. (25) (replacing M by a sky apodization W), which defines the first step of the KN approach. We refer to Appendix B for a detailed discussion. The contributions of the extra three terms in Eq. (14) are localized around the patch boundary and removed in the KN method by trimming the boundary layer, which is retained and used for the power-spectrum estimation in the case of the former method. For this reason we may expect that the performance of the KN method should be inferior to both the ZB and SZ approaches, which in turn we could expect to be nearly equivalent. In turn, the KN method may appear as the most straightforward approach on the implementation level and therefore attractive at least at first stages of the analysis.

IV. NUMERICAL EXPERIMENTS

A. Experimental setups

For numerical investigations, we define two fiducial experimental setups. Though idealized, they are chosen to reflect the general characteristics of forthcoming CMB experiments dedicated to B-mode detection. Those characteristics which crucially impact the angular power-spectrum reconstruction are the noise level, the beam width, and a peculiar sky coverage.

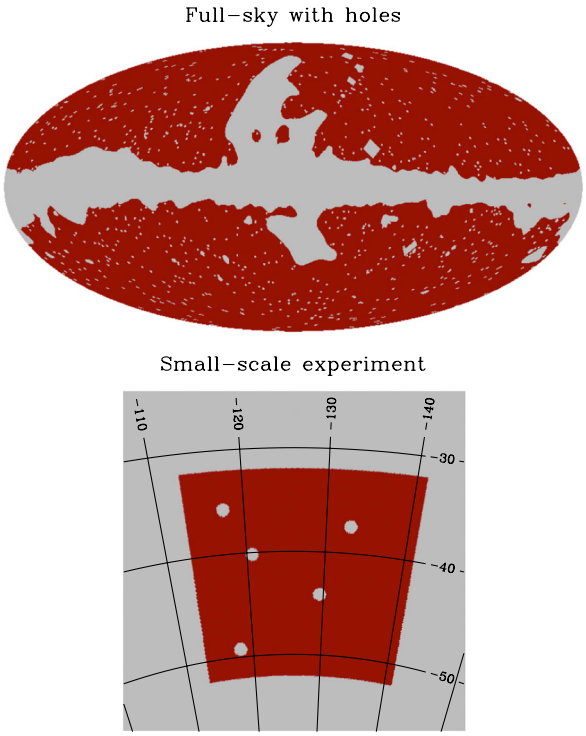

FIG. 1 (color online). Sky areas as observed by the fiducial satellite-like experiment (upper panel) and for the balloon-borne, small-scale experiment (lower panel) as considered in this work. The sky coverages are respectively ~71% and ~1% of the total celestial sphere. For the satellite experiment, the mask is a combination of the galactic mask $R9$ and the point-sources catalog used for polarized data in WMAP 7 yr release. Only the latter mask is used for the balloon-borne case.

We first consider the case of a possible satellite experiment aimed at B-mode detection. For such an experiment, we relied on the EPIC 2 m [33] specifications for the noise level and the beam width, setting these to 2.2 μK arcmin for the noise level and 8 arcmin for the beam width. For the peculiar sky coverage of such a "nearly full-sky" experiment, we consider the galactic mask r9 used for polarized data in the WMAP 7-year release (see Ref. [38]), adding the point-sources catalog mask. So we obtain a ~71% sky coverage patch, shown in the lower panel of Fig. 1. Throughout this work we use the HEALPIX pixelization scheme [39]. Here the pixel size is ~7 arc minutes, i.e., $N_{\rm side} = 512$.

Second, we consider the case of a balloon-borne experiment inspired by the ongoing EBEX experiment [34]. The noise level and the beam width are respectively set equal to 5.75 μK arcmin and 8 arcmin. The observed part of the sky covers ~1% of the total celestial sphere and its peculiar shape is displayed in the upper panel of Fig. 1. It consists of a square patch of an area of ~400 square degrees including holes to mimic polarized point-source removal. In such a case, we choose $N_{\rm side} = 1024$ corresponding to a pixel size of ~3.5 arc minutes.

B. Simulations

We numerically implement the three techniques described in the previous section and test their respective efficiency with Monte Carlo simulations. We investigated the full performances of those approaches from the perspective of B-mode power-spectrum reconstruction and therefore incorporate noise with the level, as stated in Sec. IV A. To simulate the CMB sky, the input E-mode signal is that of the cosmological model with parameters as given by the WMAP 7-year data [40] and the input B-mode includes lensing and primordial B-modes with $r = 0.05$. (Our convention for r follows the WMAP convention: $r = \mathcal{P}_{\rm T}(k_0)/\mathcal{P}_{\rm S}(k_0)$, with $\mathcal{P}_{\rm S(T)}$, the *primordial* scalar(tensor) power spectrum and $k_0 = 0.002$ Mpc^{-1} the pivot scale.)

We will assume that two identical maps are always available with the same level of the homogeneous noise in each of them—which is taken to be uncorrelated between the two maps—and use their cross spectra and their variance to compare different approaches. We calculate the latter with the help of Monte Carlo simulations and use as a common reference an estimation of the variance based on simple mode counting and given by

$$\mathbf{\Sigma}_{\ell\ell'} = \frac{\delta_{\ell,\ell'}}{(2\ell+1)f_{\rm sky}}\left[(C_\ell^B)^2 + \left(C_\ell^B + \frac{4\pi}{N_{\rm pix}}\frac{\sigma_p^2}{B_\ell^2}\right)^2\right], \qquad (26)$$

where B_ℓ is the beam function and σ_p is the noise per pixel. This formula applies to a *cross spectrum* between two maps and assumes that the noise of the two maps is uncorrelated and its level per pixel is given by σ_p. This naive mode counting is bound to underestimate the variance in our study cases and is therefore used only as a lower limit.

An effective, observed fraction of the sky, $f_{\rm sky}$, depends on an assumed apodization and therefore will be in general different for each of the methods considered here and may vary from one bin to another. For definiteness, hereafter as a reference we will use its value computed assuming only a binary mask, M. Such a choice in terms of the Fisher errors leads to the lowest variances.

V. RESULTS: SATELLITE CASE

A. Standard pseudospectrum method

The major advantage of the satellite experiments is their ability to measure the sky signals on the largest angular scales, and therefore they have the potential to constrain their power spectra all the way to the lowest multipoles. Indeed, the simple Fisher variance formula introduced earlier seems to suggest that this should be possible only if the sky coverage is sufficiently large. Though this formula neglects the leakage it seems only natural to expect that it should be small for nearly full sky maps, and therefore should lead to subdominant effects as compared to other uncertainties, e.g., cosmic variance.

In this section we confront these expectations with realistic simulations within the paradigm of pseudospectrum methods. In this context, if the leakage is indeed small, we may expect that even the standard pseudospectrum technique could perform sufficiently well, assuring a precision comparable to that of the other methods—which explicitly invoke some leakage correction—and not that far off the Fisher predictions. Below we therefore start from a discussion of the standard pseudospectrum technique.

1. Leakage

We quantify the level of the E-to-B leakage using standard pseudospectra calculated in the case of simulations with no input B-mode power, which would have been zero had there been no leakage at all. These are denoted hereafter as $\tilde{C}_\ell^{E\to B}$. We compare these pseudospectra with those calculated assuming no input E-mode power, denoted as $\tilde{C}_\ell^{B\to B}$, and therefore expressing the pseudopower of the genuine B-modes. These pseudospectra are shown in Fig. 2, which displays $\tilde{C}_\ell^{E\to B}$ (upper curve) and $\tilde{C}_\ell^{B\to B}$ (lower curves), computed for three different values of $r = 0.1, 0.05, 0.01$. Clearly, the leaked power $\tilde{C}_\ell^{E\to B}$ dominates over the true B-modes at least up to $\ell \sim 700$. We therefore conclude that the leakage is by far not insignificant even in the satellite case.

Furthermore, if we take the ratio of $\tilde{C}_\ell^{E\to B}$ and $\tilde{C}_\ell^{B\to B}$ as a measure of the magnitude of the leakage we find that its values are within a factor of 2 from those obtained for the small-scale experiment considered later on, indicating that the leakage amount in both cases is in fact comparable, even if the latter experiment covers roughly a sky area $\sim$71 times smaller than the former.

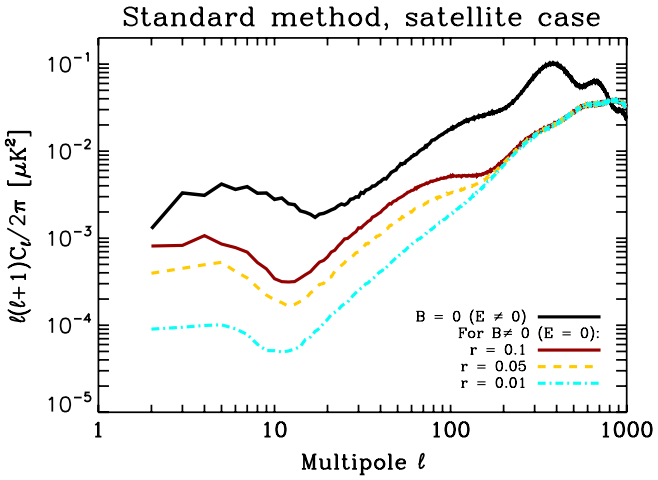

FIG. 2 (color online). Contribution of E-modes (black curve) and B-modes (colored curves) to the B-modes' pseudo-C_ℓ for the case of a satellite mission. This measures the relative amount of E-mode leaking into B if one does not correct for such leakages. The corresponding mask is depicted in the upper panel of Fig. 1.

This demonstrates that it is not merely the sky area that matters as far as the leakage is concerned. In fact, the gain in the sky area in the case of the satellite experiment considered here comes at the price of a significantly more complex and longer perimeter, the effects of which (see, e.g., Ref. [35]) offset the sky area advantage. We note that although we may attempt to simplify the boundary of the Galactic mask to suppress the leakage, this is more difficult when done with the point sources, which indeed seem to provide the major contribution to the observed level of the leakage.

2. Variance

The large leakage present on the pseudospectrum level will inevitably lead to an excess variance of the B-mode-spectrum estimate. These are depicted in Fig. 3, where variances computed assuming three different apodizations are shown. We see that in either case no meaningful constraints on the lowest multipoles, $\ell \lesssim 30$, can be set at least as long as no binning is applied. These results demonstrate that for realistic observations the standard pseudospectrum method can not ensure sufficient precision for the largest angular scales and some alternatives, explicitly correcting for the leakage, need to be considered instead, as we do so in the next section.

Figure 3 also shows a B-mode spectrum averaged over all performed MC simulations. It is unbiased, as expected, given that we include explicitly in the calculations the off-diagonal coupling kernel $K^{(-)}_{\ell\ell'}$, correcting the spectra on average for the E-mode power leaked to B. In practice, we find however that special care needs to be taken while calculating this kernel to ensure the absence of the bias. This is because the leaked power is indeed grossly dominant over the genuine B-mode (see Fig. 2), setting very demanding constraints on the precision of the kernel. For instance, the good agreement shown in Fig. 3 is only obtained when we minimize the spurious contributions due to the pixelization coming specifically from the polar caps by rotating the sky map so they have been hidden in the regions excluded by the employed mask. The residual scatter at its low-ℓ end is just a result of the insufficient number of simulations and the huge variance displayed by the standard pseudospectrum estimator on these scales.

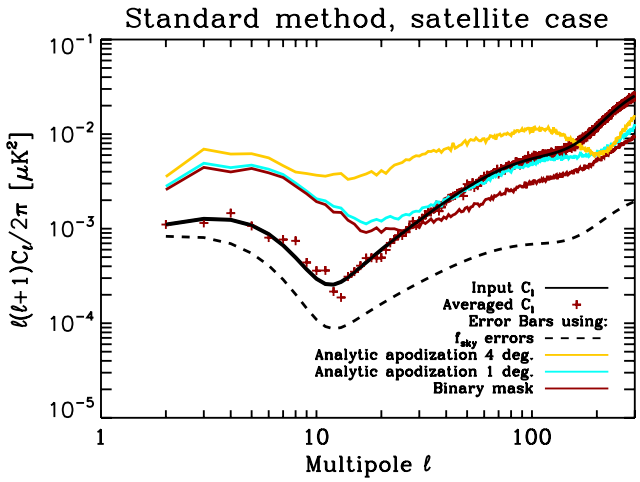

FIG. 3 (color online). Reconstruction of the B-modes' power spectrum for $r = 0.05$ in the case of a satellite mission ($f_{\rm sky} = 71\%$) using the *standard* pseudospectrum approach. The noise level is 2.2 μK arcmin and the beam is 8 arcmin. The solid black curve is the input power spectrum and the dashed black curve is the Fisher estimate of the error bars. The red crosses stand for the recovered power spectrum and the colored curves stand for the error bars from MC simulations using different apodization lengths for the sky apodization applied to the polarization maps.

The good overall agreement of the averaged spectrum with the theoretical spectrum used for the simulations validates our MC-based predictions for the variances.

B. Leakage-correcting methods

1. Apodization

The results described above demonstrate that the standard approach is not suitable for the low-ℓ recovery of the B-mode spectrum even for the nearly full-sky experiments. Therefore, if such a goal is achievable at all with a pseudospectrum method, it would have to be a method which tackles the leakage problem case-by-case, as do the three methods discussed earlier. It is important however to emphasize that the suppression of the E-to-B leakage in these methods comes at a price as the corrections they invoke may affect the variance of the recovered spectrum. Consequently, this variance will not in general be close to the variance of the B-mode spectrum obtained in the standard pseudospectrum approach in a case when the CMB E-mode power, and therefore the leakage, is set artificially to zero, as one could ideally hope for. Instead there will typically be an extra contribution to the variance, not due to the leakage anymore—as it is explicitly treated for—but rather from the removal of part of the information resulting from the leakage correction procedure.

This in principle calls for some optimization procedure between the level of the leakage and the bias (at least for some of the methods studied here) and the variance of the recovered B-mode power spectrum. As the loss of the information is related to the apodization and/or masking applied in these methods, and is used sometimes in multiple stages, such an optimization could be in general rather cumbersome to formalize and to date has been implemented in a systematic way only in the case of the SZ approach [26]. In this method the estimated power spectrum is always unbiased and the variance level is uniquely determined by one window function—if the boundary conditions and relations between different spin windows are strictly enforced—or three window functions—if the boundary conditions are relaxed and no relations between the windows is imposed. In the latter case, one admits some

level of leakage but tries to capitalize on the additional freedom to gain on the resulting variance. In past studies (e.g., Refs. [25–27]) a number of either *ad hoc* or optimized windows have been considered and shown to perform comparably at least in the simplest circumstances. In Fig. 4 we show the variances obtained with the SZ method assuming a selection of windows in the case of our satellite

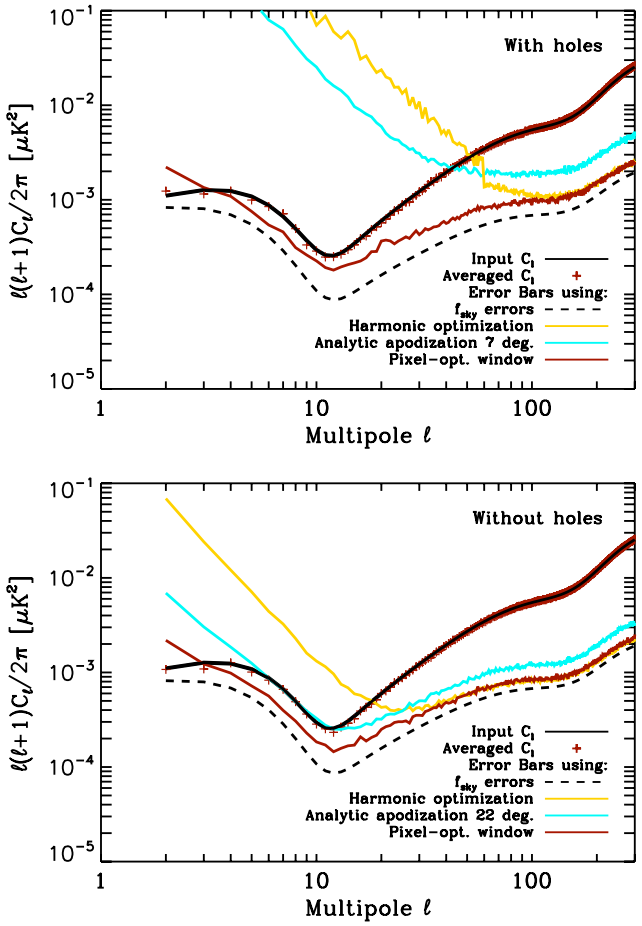

FIG. 4 (color online). Upper panel: Reconstruction of the B-modes' power spectrum for $r = 0.05$ in the case of a satellite mission ($f_{\rm sky} = 71\%$) using the SZ pseudo-C_ℓ approach and using three types of sky apodizations. The noise level is 2.2 μK arcmin and the beam is 8 arcmin. The solid black curve is the input power spectrum and the dashed black curve is the Fisher estimate of the error bars. Red crosses are the recovered power spectrum. The solid yellow curve stands for error bars on C_ℓ^B recovery using a variance-optimized apodization, forcing the boundary conditions and the relationship derivative to be fulfilled (the computation is done in the harmonic domain). The solid red curve corresponds to the error bars using a variance-optimized apodization to relax those conditions (the computation is done in the pixel domain.) The solid cyan curve shows the error bars as obtained using an analytic sky apodization with $\theta_{\rm apo} = 7$ degrees. Lower panel: Same as upper panel but considering the galactic mask only and not the holes. The sky coverage is 73%. The apodization length for the analytic sky apodization is 22 degrees.

setup, assuming the presence (upper panel) or absence (lower panel) of the masked point sources. We observe that there is a huge disparity in the performance of the different windows, in particular at the low-ℓ end of the spectrum. The windows, which tend to impose the boundary condition, i.e., the harmonic and analytic ones, perform significantly worse than the window for which these are relaxed, i.e., the pixel-domain-optimized window. Moreover, the variances in the former cases are often significantly worse than those obtained in the case of the standard approach, in particular at the low-ℓ end.

We can therefore conclude that not only do the pixel-domain-optimized windows provide the best performance—at least out of the cases we have looked at here—but also that they are unique in ensuring essentially the same performance in the cases of both of the masks considered here. For this reason we will use these windows whenever we apply the SZ approach in the following.

We note that the pixel-domain computation of the optimized windows does involve significant computational resources, which are needed to solve iteratively large linear systems [26] for a number of ℓ bins and which dominate the overall computational cost of the approach.

The situation is more complicated in the cases of the other two methods as equivalent optimization procedures have not been proposed in their context. This is in part due to technical problems related to the dimension of the parameter space, which would have to be considered. We therefore do not attempt to devise such procedures in this work. Instead, in these cases we will apply simple analytic apodizations and demonstrate the dependence of the obtained results on their parameters. As these apodizations may not be optimal, it may be in principle possible to improve on the results we derive in the following. However, we find that in general the results for these two methods are less sensitive to the apodization choices than those derived in the case of the SZ approach and therefore we do not expect the improvement to be significant enough to affect our conclusions.

We note that even with the proper optimization the determination of the low-ℓ multipoles multipole-by-multipole is burdened with a significant error. Indeed, the variance is comparable to the signal amplitude for $\ell \lesssim 20$ and even larger than the latter for $\ell \lesssim 3$–4. For this reason, in the following we will always bin the spectra even in the nearly full-sky case considered here. The choice of binning will be marked at the bottom of each plot as grey shaded boxes. The lowest bin will then span ℓ values from 2 up to 20.

The gain in using the SZ approach as compared to the standard approach—which does not correct for E-to-B leakage—is visualized in Fig. 5. It depicts the signal-to-noise ratio (SNR) of the B-mode angular power spectrum reconstruction, $C_\ell^B/\sqrt{\Sigma_{\ell\ell}}$. The red curve stands for the SNR as obtained using the SZ method while the yellow curve stands for the SNR as obtained using the standard

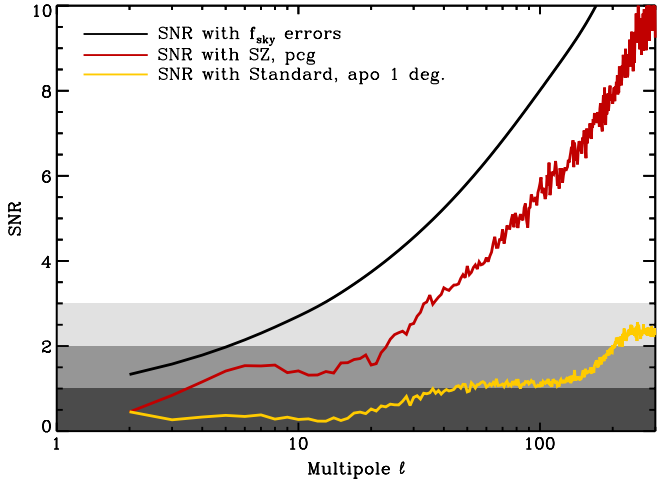

FIG. 5 (color online). Signal-to-noise ratio $C_\ell^B/\sqrt{\Sigma_{\ell\ell}}$. The red and yellow curves respectively stand for the SZ and the standard pseudo-C_ℓ estimation. The black curve shows the SNR as obtained using the Fisher estimate of the uncertainties in the reconstructed C_ℓ^B. The shaded grey areas highlight the 1σ, 2σ, and 3σ detections.

pseudo-C_ℓ method. The black curve corresponds to an idealized SNR based on the Fisher estimate of the uncertainties. The shaded grey areas highlights the 1σ, 2σ, and 3σ detections. It is clear from such a figure that detecting the primordial component of the B-modes, peaking at $\ell < 100$, for a satellite-like survey requires a correction for the E-to-B leakage.

2. Power spectrum recovery: bias and uncertainties

The reconstructed B-modes' angular power spectra and their uncertainties for each of the three above-described methods are shown in Fig. 6. The upper, middle, and lower panels, respectively stand for the SZ, ZB, and KN techniques. As explained in Sec. V B 1, the angular power spectra are estimated for $\ell \in [2, 1020]$ within multipole bands with bandwidth of $\Delta\ell = 40$. For each method, we optimize the sky apodization to obtain the lowest error bars.

The plotted solid black curve stands for the input B-modes' angular power spectrum, while the solid red curve is the estimated one, averaged over 500 simulations, which is built to be unbiased (we will discuss the results in practice for each method). The dashed black curve in each panel represents the mode-counting estimate of the uncertainties of the C_ℓ^B's, which are calculated as explained in Sec. IV B. The dashed colored curves are the MC-estimated uncertainties. The estimated binned power spectra and their associated error bars are plotted at the central value of each bandpower. The width of the bandpowers adopted here are depicted by the grey shaded rectangles.

As already mentioned, the three pseudo-C_ℓ techniques *are theoretically built to* provide *unbiased* estimations of

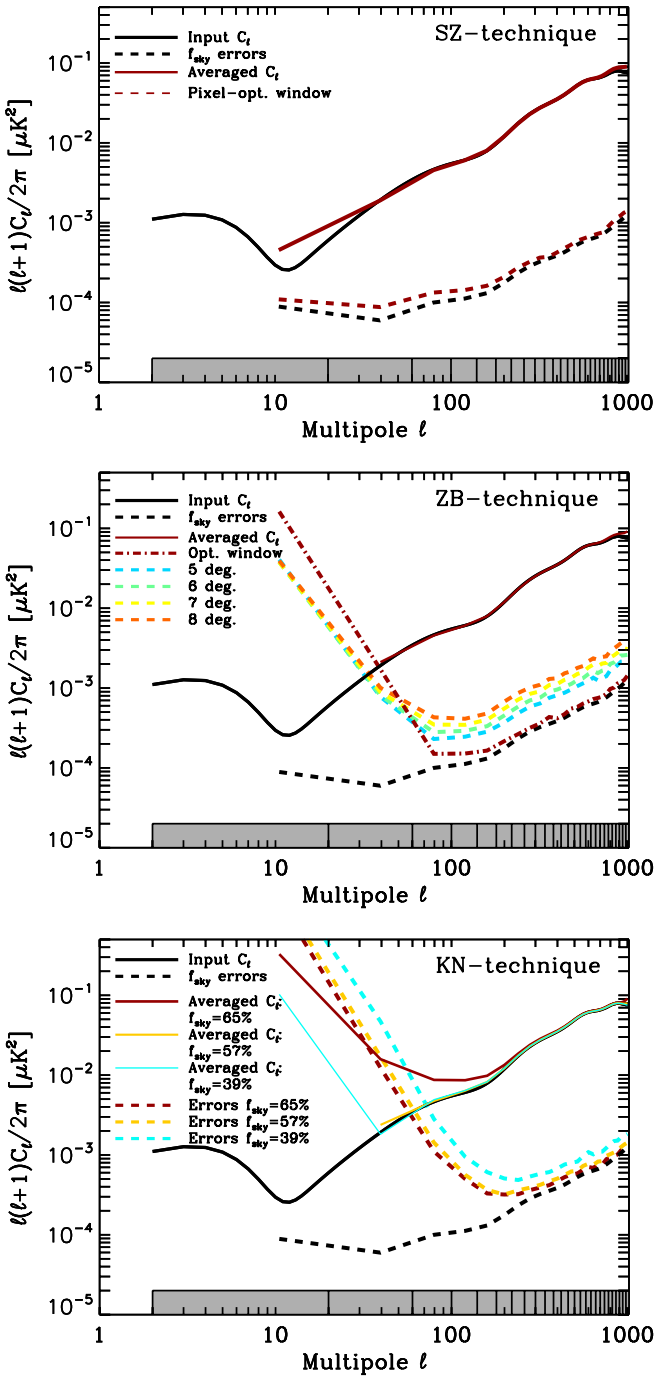

FIG. 6 (color online). Power spectrum uncertainties on B-modes using cross-spectrum estimation for the case of a satellite-like experiment with holes mimicking point-source removals ($f_{sky} \sim 71\%$). The upper, middle, and lower panels are respectively for the SZ, ZB, and KN methods. Grey shaded boxes represent the binning of the power spectra. The sky apodizations used for each technique are described in Sec. III.

C_ℓ^B. Nonetheless, due to numerical effects such as the pixelization, the reconstructed B-modes may be biased. The bias and the uncertainty behaviors for each technique are analyzed and compared hereafter.

(i) *SZ technique:* As expected, our estimation of the B-mode angular spectrum is unbiased. The window functions are optimized in the pixel domain leading to uncertainties very close to the mode-counting estimation throughout the entire range of angular scales considered.

(ii) *ZB technique:* As for the SZ technique, the B-modes' angular power spectrum C_ℓ^B is reconstructed unbiased. The dashed-dotted red curve depicts the uncertainties in C_ℓ^B via the ZB approach using harmonic-variance optimized apodizations calculated for the SZ approach, while the colored dashed curves represent the window function with different apodization lengths $\theta_{\rm apo}$ ranging from 5 to 8 degrees. We have checked that using an apodization length either smaller than 5 degrees or wider than 8 degrees systematically leads to higher uncertainties. For this technique, one cannot *a priori* apply the pixel-domain computation of the variance-optimized apodizations. We nevertheless check that this is indeed the case using numerical experiments. Our results show that weighting the maps of the Stokes parameters with the spin-0 pixel variance-optimized apodizations as derived for the SZ technique leads to very high uncertainties for $\ell < 100$. At low multipoles, a larger apodization length reduces the E-to-B leakage, lowering the uncertainties in C_ℓ^B. At high multipoles, uncertainties are driven by the sky cut, which increase with $\theta_{\rm apo}$. The harmonic-optimized window functions give the smallest uncertainties in C_ℓ^B for $\ell > 100$ but, as expected, fail to provide the smallest uncertainties for $\ell < 100$. For these large angular scales, the recovery of C_ℓ^B is only possible for $\ell > 20$ and by making use of analytic sky apodization.

(iii) *KN technique:* The estimation of the angular power spectrum appears to be biased. The solid red curve shows the estimated C_ℓ^B for an apodization length of $30'$ and is biased in the first four bins. The more we decrease the length of the apodization, the less the estimated C_ℓ^B is biased to get an unbiased estimation with $\theta_{\rm apo} = 1$ degree. This bias comes from the approximation $K_{\ell\ell'}^{(-)} = 0$, which is not verified in practice. The uncertainties derived in the KN approach are depicted in the lower panel of Fig. 6. The error bars have been obtained by first computing the map of $\tilde{\chi}^B$ using a C^2 window function with an apodization length $\theta_{\rm apo}$ and then by removing those pixels for which the sky apodization varies (that is, an external layer with a width $\theta_{\rm apo}$). The three here-adopted values for $\theta_{\rm apo}$ are 0.5, 1, and 2 degrees. As expected from the mode-counting estimation, the lowest error bars are achieved for the highest sky coverage, that is, for $\theta_{\rm apo} = 0.5$ degrees. Nonetheless, for the first two bins, the error bars for the three values of $f_{\rm sky}$ are higher than the value of the signal, which means that it is impossible to detect the primordial part. They decrease up to $\ell \simeq 200$ and then behave like the mode-counting uncertainty until $\ell = 1020$.

3. Pseudopower spectrum

A way to qualitatively describe potential bias in the methods is to study the B-modes' pseudopower spectrum $\tilde{C}_\ell^B$. Comparing these two quantities allows for a quantitative description of the leakage that bias the B-mode pseudopower spectrum. In Fig. 7, we plot the ratios $\tilde{C}_\ell^{E\to B}/\tilde{C}_\ell^{B\to B}$ for the ZB and KN methods. First of all, this ratio is not zero because of the pixelization effects. This may bias the final estimate of C_ℓ^B if such a residual leakage is not corrected for via a nonzero $K_{\ell\ell'}^{(-)}$ *and* if $\tilde{C}_\ell^{E\to B}$ cannot be safely neglected compared to $\tilde{C}_\ell^{B\to B}$. For the SZ technique, these residual leakages are corrected for via the implementation of $K_{\ell\ell'}^{(-)}$. However, such an off-diagonal block of the mode-mode coupling matrices cannot be computed in the ZB and KN techniques. The block $K_{\ell\ell'}^{(-)}$ is systematically set equal to zero, which implicitly assumes that *effectively* $\tilde{C}_\ell^{E\to B} \ll \tilde{C}_\ell^{B\to B}$. Figure 7 (the solid black curve) indicates that this assumption is valid for the ZB technique, where the ratio is approximatively equal to 10^{-2} at most. On the contrary, Fig. 7 (red curves) shows

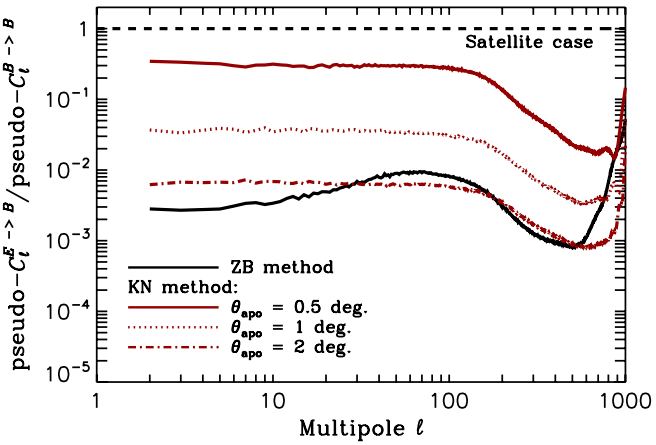

FIG. 7 (color online). Ratio between $\tilde{C}_\ell^{E\to B}$ (E-mode power spectrum is derived from the WMAP 7-year best fit) and $\tilde{C}_\ell^{B\to B}$ computed by correcting for such a leakage using the ZB and KN E/B separation techniques (respectively represented by the black curve and the red curves). This ratio amounts to the leakage of E-modes into B. The dashed black line is the benchmark to which the pseudo-C_ℓ is compared. For the KN method, the three curves are the value of the ratio for the different values of the apodization length: $\theta_{\rm ap} = 0.5$, 1, and 2 degrees from top to bottom, respectively. The sky coverage is the one expected for a satellite-like experiment with holes due to point-source removals (see upper panel of Fig. 1).

that $\tilde{C}_\ell^{E\to B}$ cannot be neglected with respect to $\tilde{C}_\ell^{B\to B}$ for the KN method inducing a bias in the B-modes' angular power spectrum, as seen in the lower panel of Fig. 6.

4. Effect of point sources in the mask

Furthermore, as already highlighted in Sec. V A 2, we confirm the importance of the point sources' holes in the mask. Indeed, we also calculated the B-modes' angular power spectra for a mask, which do not account for the polarized point sources ($f_{\rm sky} = 73\%$). The lowest achieved uncertainties for each method are depicted in Fig. 8 with holes (upper panel) and without holes (lower panel). With the difference between the two $f_{\rm sky}$'s being 2%, one could expect—from a naive mode counting—the error bars to increase by a factor of $\sim$1.01 by adding holes. Though such a scaling indeed applies to the case of the SZ method, it appears that both the ZB and the KN method are very sensitive to the presence of holes at large angular scales. Clearly, the uncertainties increase by more than $\sim$1.01 by adding holes for $\ell < 140$ for both the ZB and KN methods. Though the SZ technique can handle the impact of holes, the increase of the variance at large scales for the ZB and KN techniques shows that a dedicated treatment of holes could be mandatory.

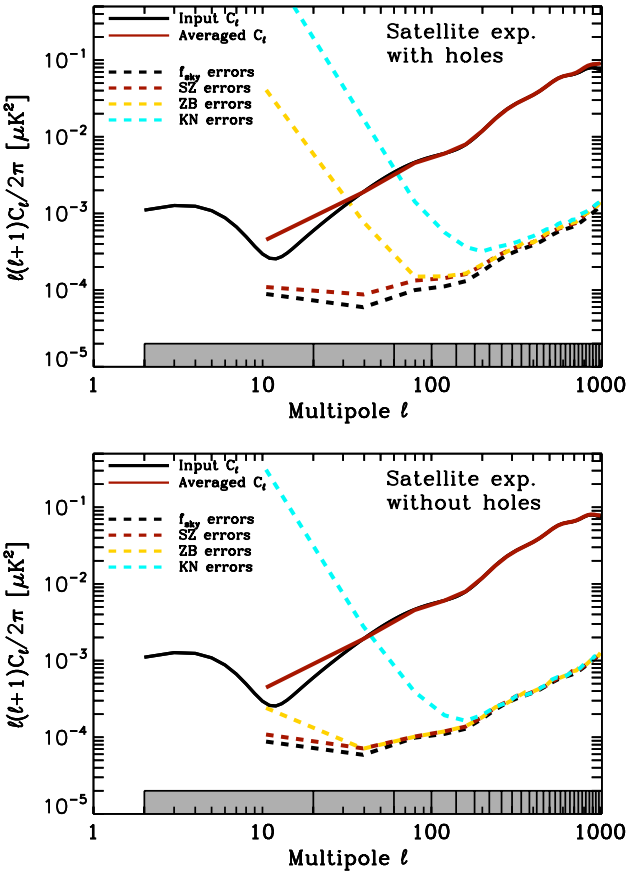

FIG. 8 (color online). Power spectrum uncertainties on B-modes using cross-spectrum estimation for the case of a satellite experiment with holes mimicking point-source removals ($f_{\rm sky} \sim 71\%$). The red dashed line represents the variance obtained via the SZ method, the blue dashed line is via the ZB method, and the yellow dashed dotted line for the KN method.

It is instructive to compare the SZ approach to the ZB approach to understand why the latter can deal with holes while the former cannot. They differ from each other by the use of two different sky apodizations: the pixel-domain, variance-optimized apodization for the SZ technique, and the harmonic-domain, variance-optimized sky apodization for the ZB technique. If one uses the harmonic-domain, variance-optimized sky apodization, the SZ approach would suffer from the high increase of the variance at large angular scales, similar to the increase of the variance observed in the ZB approach. In other words, all the additional complexity due to holes in the mask is nicely treated in the SZ approach thanks to its flexibility and a dedicated computation of the sky apodization in the pixel domain.

5. Conclusion for a satellite-like experiment

To summarize, the SZ method gives unbiased B-mode power spectra and the smallest uncertainties, close to the mode-counting one, for the case of a large sky coverage (see Fig. 4 for a reconstruction multipole-by-multipole and Fig. 8 for a reconstruction using bandpower). The results with the ZB method with the harmonic-optimized windows are similar to those of the SZ method for $\ell > 100$. For $\ell \in [20, 100]$, estimating C_ℓ^B is still possible but with a smaller significance. For $\ell < 20$, the ZB method fails to reconstruct the B-modes' angular power spectra. Our implementation of the KN method does not manage to reconstruct an unbiased C_ℓ^B for the first four bins if the apodization length is too small. For those apodizations that allow the KN method to provide an unbiased estimation ($\theta_{\rm apo} \geq 1$ degree), reconstructing C_ℓ^B is not possible for $\ell < 60$. For intermediate angular scales, $60 < \ell < 300$, the reconstruction is possible with a lower signal-to-noise ratio than the one achieved thanks to either the Sz technique or the ZB technique.

VI. RESULTS: SMALL-SCALE EXPERIMENT

In the case of a balloon-borne-like experiment, the reconstructed B-mode angular power spectra and their associated uncertainties are shown in Fig. 9 for the three techniques. The angular power spectra are estimated from $\ell = 2$ to $\ell = 1020$ with the first bin ranging from 2 to 20 and the following bins having a bandwidth $\Delta\ell$ equal to 40. We emphasize that for such a small value of the sky coverage, the amplitude of the binned C_ℓ^B in the first bin $b_1 = [2, 20]$, is $C_{b_1}^B \simeq 5.9 \times 10^{-4}$ μK^2 for $r = 0.05$. The Fisher estimate of the uncertainties for the same value of r leads to $\sqrt{\mathbf{\Sigma}_{b_1 b_1}} \simeq 7.6 \times 10^{-4}$ μK^2. Detecting a nonvanishing C_ℓ^B at angular scales between $\ell = 2$ and $\ell = 20$

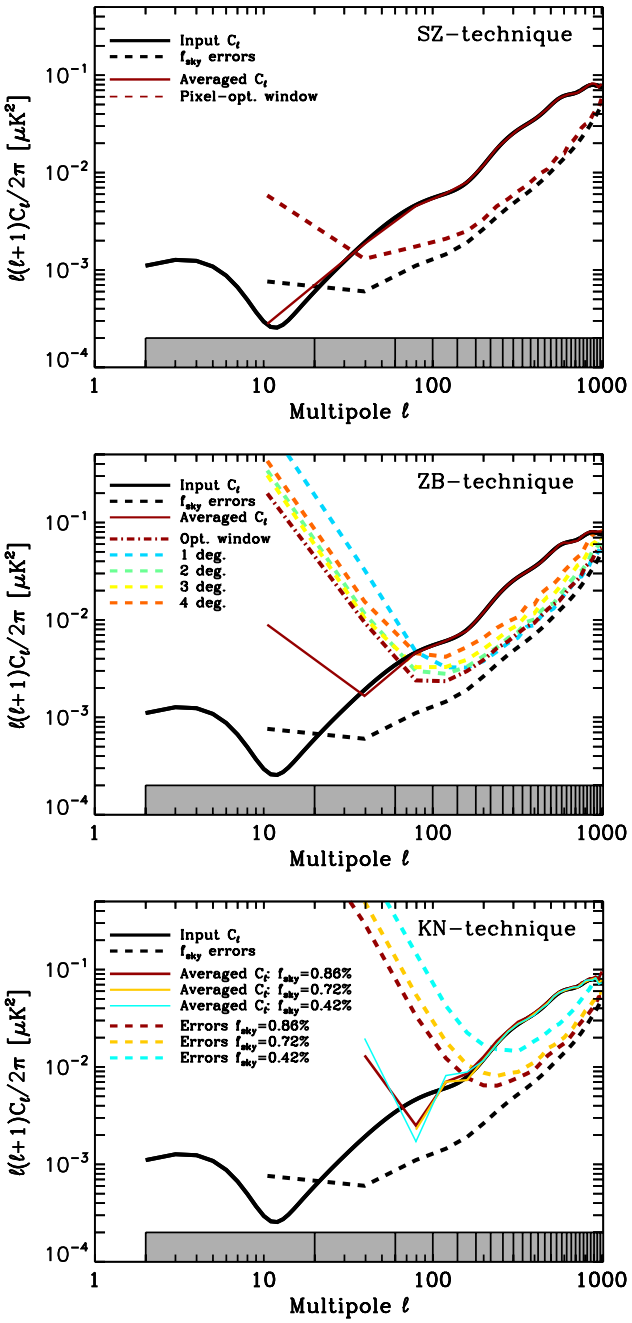

FIG. 9 (color online). Power spectrum uncertainties on B-modes using cross-spectrum estimation for the case of a balloon-borne experiment ($f_{\rm sky} \sim 1\%$). The upper, middle, and lower panels are for, respectively, the SZ, ZB, and KN methods. The sky apodizations used for each technique are described in Sec. III.

appears unfeasible for small-scale experiments since the Fisher calculation underestimates the variance on the pseudo-C_ℓ reconstruction of angular power spectra. On each of the three graphs, the solid black curve corresponds to the input B-mode power spectrum to be estimated, while the solid red curve stands for the estimated angular power spectrum averaged over 500 simulations. The dashed black curves correspond to the mode-counting estimate of power spectrum uncertainties obtained with $f_{\rm sky} = 1\%$, which serves as a benchmark. For each of the graphs, the dashed colored curves stands for MC estimations of the power spectrum uncertainties for each of the techniques.

(i) *SZ technique:* We confirm that the reconstructed angular B-mode power spectrum is unbiased for the entire range of multipoles considered here. As previously mentioned, we only use pixel-optimized window functions for the case of the SZ technique (upper panel of Fig. 9) and the displayed error bars are therefore the lowest ones to be expected in such an approach. We refer the reader to Ref. [27] for an exhaustive discussion on the performances of such a technique. The relevant conclusion in such a case is that a precise-enough estimation of C_ℓ^B is achieved for multipoles starting from $\ell = 20$ to $\ell = 1020$.

(ii) *ZB technique:* In such a case, the estimated C_ℓ^B's are also unbiased from $\ell = 2$ to $\ell = 1020$. We show the power-spectrum uncertainties for two kind of windowing. Dashed colored curves ranging from blue to orange stand for error bars derived using a C^2 window function with an apodization length varying from 1 to 4 degrees. It clearly shows that depending on angular scales the apodization length has to be adapted to reach the lowest uncertainties. For the three first bins, i.e., $2 \leq \ell < 100$, an apodization length of 3 degrees provides the lowest error bars. For higher multipoles, an apodization length of 1 degree leads to the smallest error bars. The dashed red curve corresponds to the uncertainties in the reconstructed C_ℓ^B's using an optimized window function computed in the harmonic domain.[7] This clearly shows that—unlike the case of a satellite mission—using such harmonic-variance-optimized sky apodizations provides the lowest error bars in the entire angular range. However, although it is very efficient at multipoles greater than 60, this approach fails to reconstruct the B-mode angular power spectrum for the two first bins comprised in $2 \leq \ell < 20$ and $20 \leq \ell < 60$.

(iii) *KN technique:* The KN technique provides an unbiased B-mode angular power spectrum, although it is highly scattered because of the high level of the variance at low ℓ. As for the discussed case of a satellite-like experiment, the lowest uncertainties are obtained for the highest sky coverage, i.e., for $\theta_{\rm apo} = 0.5$ degrees, though the KN technique is able to estimate C_ℓ^B only for ℓ values greater than

[7]Because the contour of the mask is rather simple for the small-scale experiment, the harmonic computation of the variance-optimized apodizations leads to very similar results to the pixel-domain computation for the SZ technique.

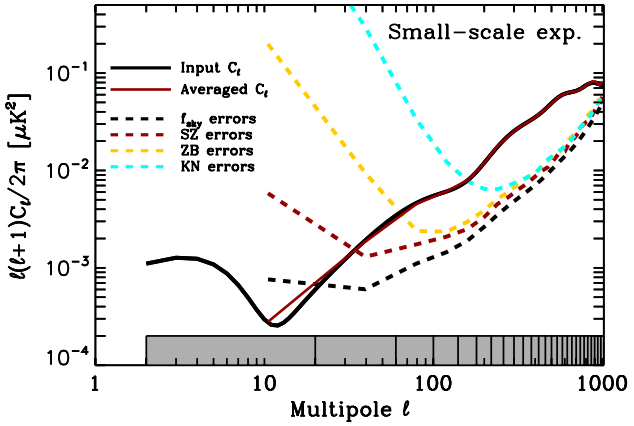

FIG. 10 (color online). Power spectrum uncertainties for each of the three techniques for the case of a small-scale experiment with $f_{\rm sky} \simeq 1\%$, $\sigma_Q = 5.75$ μK arcmin and $\theta_{\rm Beam} = 8$ arcmin. Dashed red, dashed cyan, and dashed yellow curves are for, respectively, the SZ, ZB, and KN techniques. The dashed black curve stands for the $f_{\rm sky}$ estimate of the error bars.

$\sim$150 and therefore "misses" the bump at $\ell \sim 100$ due to the primordial component of the B-mode angular power spectrum.

Figure 10 summarizes our results, depicting the lowest error bars on the B-mode estimation for each of the three techniques. From these results, it is rather obvious that the SZ technique performs the best for power-spectrum reconstruction from the viewpoints of both bias and uncertainties. This approach allows for an accurate enough estimation of C_ℓ^B for $\ell \geq 20$, while the ZB technique and the KN technique allow for such a reconstruction for $\ell \geq 60$ and $\ell \geq 150$, respectively. Those differences may drastically affect our ability to set constraints on the cosmological parameters probing the inflationary phase such as, e.g., the tensor-to-scalar ratio r. We remind the reader that the primordial component of C_ℓ^B—from which the constraint on r can be set—is dominant for ℓ values lower or equal to $\sim$100 while the lensing-induced B-modes start to dominate the angular power spectrum for $\ell > 100$. With our binning, this means that with the SZ technique one can detect the primordial B-modes in two bins (i.e., $\ell \in [20, 60]$ and $\ell \in [60, 100]$). With the ZB technique, the primordial component of C_ℓ^B can be detected in only one bin, $\ell \in [60, 100]$, while a detection of the primordial component seems impossible with the KN approach.[8]

VII. CONCLUSION AND DISCUSSION

We first presented three different pseudospectrum estimators designed to remove–or at least reduce—the E-to-B leakage, which may compromise any detection of the B-modes and especially its primordial part. We then tested the relative efficiency of these estimators to reconstruct the B-modes' angular power spectrum through Monte Carlo simulations. Two different kinds of sky coverage have been chosen for our analysis: a small-scale coverage (where the observed part of the sky is $\sim$1%) and a large coverage of the celestial sphere as motivated by a future satellite mission dedicated to B-mode detection with $f_{\rm sky} \sim 71\%$. Both sky coverages incorporate holes mimicking point-source removals.

All three techniques studied here try to reconstruct, implicitly or explicitly, the χ^B field, which is known to contain only the B-modes. We first described the so-called SZ method whose efficiency lies in an adapted choice of basis to decompose the E- and B-modes optimizing the apodization of the applied mask. Then, the ZB-technique principle was developed. It consists in calculating the masked χ^B with an adapted apodized mask; it implies derivation operations of the masked polarization field which are actually done in the harmonic space. Finally, the KN method is based on the fact that applying a mask on the reconstructed B-modes reduces significantly the level of E-to-B leakage. In this article, we do not claim to exactly implement the methods as they were described in the referred articles. Slight changes have been made in their implementation in order to minimize as much as possible the effective E-to-B leakage.

We compare the results of these methods in each of our simulation sets.

First, we found that correcting for E-to-B leakages at the levels of both the mean and variance is mandatory in the case of a satellite mission covering $\sim$71% of the sky for an efficient recovery of the primordial component of B-modes, $\ell < 30$. Moreover, we have shown that the intricate shape of the galactic mask makes the uncertainties of the reconstructed C_ℓ^B using methods correcting for E-to-B leakages very sensitive to sky apodization applied to Q and U maps for $\ell < 60$. An efficient computation of variance-optimized sky apodization is therefore crucial for the applicability of these methods. From this practical perspective, the SZ method appears to be better armed as it offers some flexibility in the computation of the sky apodization.

Second, we computed the pseudo-C_ℓ, which amounts to the E-modes leaking into B while applying the three different techniques. Each are able to significantly decrease the E-to-B leakage, though none manage to exactly cancel it because of the pixelization effects. Nonetheless, the value of the uncertainties in the C_ℓ^B reconstruction is the key issue because it tells us if a detection is possible or not. As shown by our numerical results, the final uncertainties in the estimated B-modes' power spectra can overwhelm the signal even when the E-to-B leakage is well controlled. The SZ method gives the smallest error bars on the B-modes' angular power spectra for both the large- and

[8]Strictly speaking, some constraint can be set on r even by using the KN approach (at least some upper limit). But this may probably prevent any *measurement* of r.

small-scale experiments as they follow quite well the mode-counting uncertainties. Though we cannot recover the largest angular scales ℓ-by-ℓ for $\ell \leq 5$ with the SZ approach, we can reach a detection for such scales using the appropriate binning. The ZB method, as explained in Sec. III D, is *theoretically* equivalent to the SZ one. From the numerical results, we showed that *practically* this method is less efficient at large angular scales ($\ell < 60$ for a satellite-like mission and $\ell < 100$ for a small sky survey), allowing us to reconstruct C_ℓ^B starting at $\ell \sim 20$ for a satellite-like mission (starting at $\ell = 60$ for a small sky survey). For smaller angular scales, $\ell > 60$, these two methods provide similar results. The KN method is by construction expected to be less efficient than the other methods in our implementation, as described in Sec. III D. Indeed, the sky coverage decreases according to the apodization length, leading to a higher variance compared to SZ estimator. The power-spectrum analysis shows that this method is reliable for high ℓ, but the error bars overwhelm the signal for the two first bins (i.e., $\ell < 60$) in the case when the fraction of the sky is 71% (for the three first bins, $\ell < 100$, for a small sky survey).

Figures 10 and 8 sum up the errors made on the estimated C_ℓ^B via the SZ, ZB, and KN methods in the two experimental configurations. In the way we have implemented these techniques, the SZ method is the most efficient one. For both types of experimental setups, it makes possible the estimation of C_ℓ^B with uncertainties on par with the most optimistic Fisher estimates. The key step making the SZ approach more efficient is its flexibility in terms of sky apodization. This is highlighted in Fig. 4: relaxing the derivative relationship relating the spin-1 and spin-2 windows to the spin-0 window is mandatory for computing variance-optimized sky apodizations, drastically lowering the final uncertainties on the estimated C_ℓ^B's. However, neither the ZB nor the KN approaches are currently designed to offer such a flexibility. We have checked that if one uses the same sky apodization (for example an analytic window function with a given apodization length) the SZ and ZB methods leads to similar uncertainties. Inversely, we have also checked that one cannot use the pixel-domain, variance-optimized sky apodization (relaxing the derivative relationship) in the ZB approach as it systematically leads to an increase of the final error bars as compared to, e.g., using analytic windows with an appropriate choice of the apodization length. This shows that the applicability of these pseudo-C_ℓ estimators, which do not mix E- and B-modes, is highly conditioned by the precomputation of variance-optimized sky apodization.

ACKNOWLEDGMENTS

This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Some of the results in this paper have been derived using the S^{2}HAT [41–44], HEALPIX [39], and CAMB [45] software packages.

APPENDIX A: CONVOLUTION KERNELS FOR THE KN METHOD

We show in this appendix that a complete derivation of the convolution kernels in the KN approach is computationally prohibitive.

In such an approach, a map of the masked χ^B field is first built by applying Eq. (B3). As a function of the "true" CMB E and B multipoles, this masked χ^B field reads

$$\tilde{\chi}^B(\vec{n}) = -\sum_{\ell' m'} [\mathcal{K}^{(-)}_{\ell' m'}(\vec{n}) a^E_{\ell' m'} + i \mathcal{K}^{(+)}_{\ell' m'}(\vec{n}) a^B_{\ell' m'}]. \quad \text{(A1)}$$

The above convolution kernels should be viewed as scalar functions in the pixel domain parametrized by some harmonic indices. They measure the amount of (ℓ, m) multipoles of $B(E)$ types contributing to the *masked* χ^B field in the direction $\vec{n}$. In principle, these coupling functions are given by

$$\mathcal{K}^{(\pm)}_{\ell m}(\vec{n}) = \lambda_\pm [\bar{\partial}\,\bar{\partial}(M\,{}_2Y_{\ell m}) \pm \partial\partial(M\,{}_{-2}Y_{\ell m})], \quad \text{(A2)}$$

with $\lambda_\pm$ a complex-valued numerical constant. Expanding the spin-raising and spin-lowering operation, we obtain

$$\begin{aligned}\mathcal{K}^{(+)}_{\ell m}(\vec{n}) = {} & N_{\ell,2} \times M \times Y_{\ell m} + \lambda_{+,1,\ell} [{}_1Y_{\ell m} \bar{\partial} M + {}_{-1}Y_{\ell m} \partial M] \\ & + \lambda_{+,2,\ell} [{}_2Y_{\ell m} \bar{\partial}\,\bar{\partial}\, M + {}_{-2}Y_{\ell m} \partial\partial M] \end{aligned} \quad \text{(A3)}$$

and

$$\begin{aligned}\mathcal{K}^{(-)}_{\ell m}(\vec{n}) = {} & \lambda_{-,1,\ell} [{}_1Y_{\ell m} \bar{\partial} M - {}_{-1}Y_{\ell m} \partial M] \\ & + \lambda_{-,2,\ell} [{}_2Y_{\ell m} \bar{\partial}\,\bar{\partial}\, M - {}_{-2}Y_{\ell m} \partial\partial M], \end{aligned} \quad \text{(A4)}$$

where the explicit expression of the λ_i's are of no importance here. It is clear from the above computation that where the mask is constant, i) $\mathcal{K}^{(-)}_{\ell m}(\vec{n})$ is zero and there is no E-to-B leakage, and ii) $\tilde{\chi}^B = \sum_{\ell m} N_{\ell,2} \times M \times Y_{\ell m}$, which is just the definition of the χ^B field on the mask M.

Here, we have just reconfirmed that the derivation of χ^B proposed in Ref. [29] is exact *on the part of the sky* where the mask is *constant*. The above result is made possible if and only if the convolution kernels $F_\pm$ computed in the harmonic domain [see Eq. (B4)] are an effectively precise enough representation of the operator $[\partial\partial \pm \bar{\partial}\,\bar{\partial}\,]$. However, the truncation in the (ℓ, m) summation and the pixelization show that it is not the case. Indeed, if it was the case, the $F_\pm(\vec{n}, \vec{n}')$ would be completely local and the map of the leaking E-modes would be concentrated on the boundaries of the observed sky. But the results displayed in Ref. [29] show that $F_\pm$ is not local—though well peaked—and that leaked E-modes extend inside the observed patch. As a consequence, the $\mathcal{K}^{(\pm)}_{\ell m}$ functions are not strictly equal to the above expressions leading to

residual E-to-B leakages as well as potential ℓ-to-ℓ' aliasing. These functions should be computed differently in order to keep track of, at least, the (ℓ, m) truncation, and to subsequently derive an *unbiased* pseudo-C_ℓ estimator by correcting for the different residual leakages. For this purpose, we propose here an alternative expression for the $\mathcal{K}^{(-)}_{\ell m}(\vec{n})$ which can then be plugged into the final expression for the pseudo-C_ℓ estimator.

The (ℓ, m)-to-pixel convolution kernels are expressed as functions of Wigner-3j symbols and the multipoles of the binary mask M describing the observed sky,

$$\begin{aligned}\mathcal{K}^{(\pm)}_{\ell' m'}(\vec{n}) = \frac{i}{2}\sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^{m_1} N_{\ell_1,2} F(\ell', \ell_1, \ell_2) \\ \times \begin{pmatrix} \ell' & \ell_1 & \ell_2 \\ m' & -m_1 & m_2 \end{pmatrix} \mathcal{M}_{\ell_2 m_2} Y_{\ell_1 m_1}(\vec{n}) \\ \times \left[\begin{pmatrix} \ell' & \ell_1 & \ell_2 \\ 2 & -2 & 0 \end{pmatrix} \pm \begin{pmatrix} \ell' & \ell_1 & \ell_2 \\ -2 & 2 & 0 \end{pmatrix} \right], \end{aligned} \tag{A5}$$

with

$$F(\ell, \ell', \ell'') = \sqrt{\frac{(2\ell+1)(2\ell'+1)(2\ell''+1)}{4\pi}}. \tag{A6}$$

Being scalar functions, their multipoles are obtained by projecting them onto the spherical harmonic basis,

$$\begin{aligned}\mathcal{K}^{(\pm)}_{\ell' m'}(\ell, m) &= \int \mathcal{K}^{(\pm)}_{\ell' m'}(\vec{n}) Y^\star_{\ell m'}(\vec{n}) \\ &= \frac{i}{2}\sum_{\ell_2 m_2}(-1)^m N_{\ell_1,2} F(\ell, \ell', \ell_2) \\ &\quad \times \begin{pmatrix} \ell' & \ell & \ell_2 \\ m' & -m & m_2 \end{pmatrix} \mathcal{M}_{\ell_2 m_2} \\ &\quad \times \left[\begin{pmatrix} \ell' & \ell & \ell_2 \\ 2 & -2 & 0 \end{pmatrix} \pm \begin{pmatrix} \ell' & \ell & \ell_2 \\ -2 & 2 & 0 \end{pmatrix} \right]. \end{aligned} \tag{A7}$$

Secondly, the reconstructed $\tilde{\chi}^B$ field is masked again with the M_{χ^B} from which pseudomultipoles, denoted by $\tilde{\chi}^B_{\ell m}$ hereafter, are derived. It is easily shown that

$$\tilde{\chi}^B_{\ell m} = \sum_{\ell' m'} [K^{(+)}_{\ell m, \ell' m'} a^B_{\ell' m'} - i K^{(-)}_{\ell m, \ell' m'} a^E_{\ell' m'}], \tag{A8}$$

with

$$\begin{aligned}K^{(\pm)}_{\ell m, \ell' m'} = (-i)\sum_{\ell_1 m_1}\sum_{\ell_4 m_4}(-1)^m F(\ell, \ell_1, \ell_3) \mathcal{K}^{(\pm)}_{\ell' m'}(\ell_1, m_1) \\ \times \mathcal{M}^{(\chi^B)}_{\ell_3 m_3} \begin{pmatrix} \ell & \ell_1 & \ell_3 \\ -m & m_1 & m_3 \end{pmatrix} \begin{pmatrix} \ell & \ell_1 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned} \tag{A9}$$

Finally, at the level of power spectra, the convolution kernels are in principle derived using

$$K^{(\pm)}_{\ell\ell'} = \frac{1}{2\ell+1}\sum_{m,m'} |K^{(\pm)}_{\ell m, \ell' m'}|^2.$$

In a more convential approach, the above azimuthal averaging is done analytically and allows us to greatly simplify the expression of $K^{(\pm)}_{\ell\ell'}$. However, it is easily understood by first plugging the expression of $\mathcal{K}^{(\pm)}_{\ell' m'}(\ell, m)$ into $K^{(\pm)}_{\ell m, \ell' m'}$, and second by plugging the expression of $K^{(\pm)}_{\ell m, \ell' m'}$ into $K^{(\pm)}_{\ell\ell'}$, that such simplifications cannot be applied in the KN approach. As a consequence, the computation of $K^{(\pm)}_{\ell m, \ell' m'}$ implies three summations over (ℓ, m) indices and the intricate multiplication of four Wigner-3j symbols. It is therefore obvious that the complete derivation of the convolution kernels in the KN technique cannot be performed numerically.

APPENDIX B: COMPARING THE KN AND ZB METHODS

We show in this appendix that the KN method *approximates* the ZB technique if W satisfies the Dirichlet and Neumann boundary conditions. Our starting point is the first term of the rhs of Eq. (23),

$$\begin{aligned}\mathcal{B}(\vec{n}) &= \frac{i}{2}[\bar{\partial}\,\bar{\partial}(W(\vec{n})P_2(\vec{n})) - \partial\partial(W(\vec{n})P_{-2}(\vec{n}))] \\ &= \frac{i}{2}\int d\vec{n}' \sum_{\ell m} Y_{\ell m}(\vec{n}) Y^\star_{\ell m}(\vec{n}')[\bar{\partial}\,\bar{\partial}(W(\vec{n}')P_2(\vec{n}')) \\ &\quad - \partial\partial(W(\vec{n}')P_{-2}(\vec{n}'))]. \end{aligned} \tag{B1}$$

The second line is obtained by inserting the closure properties of the spherical harmonics. By performing two integrations by parts and using the boundary conditions verified by W to cancel the contour integrals, one obtains

$$\begin{aligned}\mathcal{B}(\vec{n}) = \frac{i}{2}\int d\vec{n}' \sum_{\ell m} Y_{\ell m}(\vec{n})[W(\vec{n}')P_2(\vec{n}')\bar{\partial}\,\bar{\partial}\,Y^\star_{\ell m}(\vec{n}') \\ - W(\vec{n}')P_{-2}(\vec{n}')\partial\partial Y^\star_{\ell m}(\vec{n}')]. \end{aligned} \tag{B2}$$

We recall that

$$\bar{\partial}\,\bar{\partial}\,Y^\star_{\ell m} = N_{\ell,2} \times {}_{+2}Y^\star_{\ell m}, \qquad \partial\partial Y^\star_{\ell m} = N_{\ell,2} \times {}_{-2}Y^\star_{\ell m}.$$

By inserting the above expression into Eq. (B2), one easily recognizes the convolution kernels used in the KN method to finally get

$$\mathcal{B}(\vec{n}) = \int d\vec{n}' W(\vec{n}')[F_+(\vec{n}, \vec{n}')P_2(\vec{n}') - F_-(\vec{n}, \vec{n}')P_{-2}(\vec{n}')], \tag{B3}$$

with

$$F_\pm(\vec{n}, \vec{n}') = \frac{i}{2}\sum_{\ell m} N_{\ell,2} \times Y_{\ell m}(\vec{n}) \times {}_{\pm 2}Y^\star_{\ell m}(\vec{n}'). \tag{B4}$$

This finishes our proof that the KN method applied to $W \times P_{\pm 2}$ is equal to the first term of the central equation,

i.e., Eq. (23), of the ZB approach, as Eqs. (B3) and (B4) are exactly the numerical starting point of the KN method (see Eqs. (11) and (12) of Ref. [29]).

APPENDIX C: NOISE BIAS

We provide in this appendix the explicit calculation of the noise bias for the ZB and KN techniques. Indeed, as underlined in Sec. III D, computing the noise bias in these two techniques may be problematic, specifically for the case of inhomogeneous noise which has not been addressed in neither Ref. [28] nor in Ref. [29]. Forthcoming data sets as provided by balloon-borne or ground-based experiments are plagued by *inhomogeneous* noise and it is therefore of primary importance to have formulas of the noise bias that are applicable to inhomogeneous noise. For this purpose, we will suppose that the noise in the Q and U maps is potentially inhomogeneous but still uncorrelated from pixel-to-pixel, translating into the two-point correlation functions

$$\langle n_Q(\vec{n}) n_Q(\vec{n}')\rangle = \sigma_Q^2(\vec{n})\delta^2(\vec{n}-\vec{n}'),$$
$$\langle n_U(\vec{n}) n_U(\vec{n}')\rangle = \sigma_U^2(\vec{n})\delta^2(\vec{n}-\vec{n}'),$$

and

$$\langle n_Q(\vec{n}) n_U(\vec{n}')\rangle = 0,$$

leading to the following correlations in the harmonic space:

$$\langle n^B_{\ell m} n^{B\,\star}_{\ell' m'}\rangle = \frac{1}{4}\int d\vec{n} M^2(\vec{n})[\sigma_Q^2(\vec{n}) + \sigma_U^2(\vec{n})] \times [{}_2Y_{\ell m}(\vec{n})\,{}_2Y^\star_{\ell' m'}(\vec{n}) + {}_{-2}Y_{\ell m}(\vec{n})\,{}_{-2}Y^\star_{\ell' m'}(\vec{n})].$$

The main challenge in computing the noise bias in the inhomogeneous case is to find an expression of $\tilde{N}^B_\ell$ which is a function of $\langle n_Q(\vec{n}) n_Q(\vec{n}')\rangle$ and $\langle n_U(\vec{n}) n_U(\vec{n}')\rangle$ *only*—as we cannot know *a priori* how noise is modified by taking derivatives—and which is numerically tractable.

From the above correlation, it is easily checked that on the full sky the noise is described by a power spectrum if it is homogeneous, i.e., $\sigma_Q^2(\vec{n}) = \text{const}$ and $\sigma_U^2(\vec{n}) = \text{const}$, with

$$\langle n^B_{\ell m} n^{B\,\star}_{\ell' m'}\rangle = \frac{1}{2}[\sigma_Q^2 + \sigma_U^2]\delta_{\ell,\ell'}\delta_{m,m'}.$$

1. The SZ-technique case

For such an approach, the noise bias is easily computed, assuming only that the pixel-to-pixel correlation is vanishing. The pseudo-$a_{\ell m}$'s are the ones of $W \times \chi^B$ resulting in the following noise bias of the pseudopower spectrum:

$$\tilde{N}^B_\ell = \frac{1}{8\pi}\int d\vec{n}(\sigma_Q^2 + \sigma_U^2)(N^2_{\ell,2}W^2 + 4N^2_{\ell,1}|\partial W|^2 + |\partial\partial W|^2). \quad \text{(C1)}$$

We refer the reader to Ref. [26] or to the Appendix of Ref. [36] for a detailed derivation of such a noise bias.

2. The ZB-technique case

The noise bias as computed in Ref. [28]—given by their Eq. (50)—clearly assumes that the noise is described by a power spectrum and is therefore homogeneous. Their computation proceeds as follows. First, one assumes that second-order moments of the noise statistics are completely described by a power spectrum, denoted by N_ℓ in Ref. [28], and are valid on the entire celestial sphere. As a consequence, this noise bias is valid at the level of power spectra and not at the level of pseudopower spectra. Second, one computes the noise bias at the level of pseudo-C_ℓ—denoted by $\mathcal{N}_\ell$ in Ref. [28] and denoted by $\tilde{N}_\ell$ in this paper—using the convolution kernel, i.e.,

$$\tilde{N}_\ell = \sum_{\ell'} K_{\ell\ell'} N_{\ell'}.$$

For the above relation to be valid, assuming that $\langle n_{\ell m} n^\star_{\ell' m'}\rangle = N_\ell \delta_{\ell,\ell'}\delta_{m,m'}$ is mandatory. In other words, the noise properties should be such that the instrumental noise, as reprojected on the celestial sphere, is statistically isotropic. To our knowledge, there is no experimental setup leading to such properties of the noise.

For the case of inhomogeneous noise, one can easily obtain the noise bias by noticing that the resulting map is equivalent to the map of the pure pseudo-$a_{\ell m}$'s as computed in the SZ approach by replacing W by W^2. Our purpose is to derive the noise of the power spectrum estimated from $W^2 \times \chi^B$ as a function of the noise power per pixel of the Q and U maps. Our starting point is the pseudomultipoles given by

$$\tilde{\chi}^B_{\ell m} = \int W^2 \times \chi^B \times Y^\star_{\ell m} d\vec{n}. \quad \text{(C2)}$$

Since the χ^B field is defined by $\chi^B = i[\bar{\partial}\,\bar{\partial}\,P_2 - \partial\partial P_{-2}]/2$, it is easily shown by performing two integrations by parts and using the fact that W^2 and $\partial(W^2)$ are vanishing at the contour that

$$\tilde{\chi}^B_{\ell m} = \frac{i}{2}\int [P_2\bar{\partial}\,\bar{\partial}(W^2 Y^\star_{\ell m}) - P_{-2}\partial\partial(W^2 Y^\star_{\ell m})] d\vec{n}, \quad \text{(C3)}$$

which is exactly the definition of the pure pseudomultipoles. As a consequence, the noise bias for the ZB method is given by the noise bias as computed in the SZ method.

However, $W^2\chi^B$ is effectively computed using the following expression:

$$\begin{aligned} W^2(\vec{n})\chi^B(\vec{n}) = {} & \left(\frac{i}{2}\right) W[\bar{\partial}\,\bar{\partial}(WP_2) - \partial\partial(WP_{-2})] \\ & - i[\bar{\partial}W \times \bar{\partial}(WP_2) - \partial W \times \partial(WP_{-2})] \\ & - \left(\frac{i}{2}\right) W[(\bar{\partial}\,\bar{\partial}\,W)P_2 - (\partial\partial W)P_{-2}] \\ & + i[(\bar{\partial}W)^2 \times P_2 - (\partial W)^2 \times P_{-2}]. \end{aligned} \quad \text{(C4)}$$

Let us show that we obtain the same expression for $\tilde{\chi}^B_{\ell m}$ using the right-hand-side of the above formulas. By

plugging the above expression into the expression for $\tilde{\chi}^B_{\ell m}$, one can perform some integrations by parts in order to replace terms like $F \times \partial(WP_{\pm 2})$ by terms like $(WP_{\pm 2}) \times \partial F$. For the first line of the rhs of the above expression, two integrations by parts are required, and only one is needed for the second line. With such a procedure, some contour integrals should appear. However, the integrand in these contour integrals are always proportional to W and/or ∂W. As these functions are required to vanish at the contour, all the contour integrals are equal to zero and we are left with

$$\tilde{\chi}^B_{\ell m} = \frac{i}{2}\int [P_2 \times {}_2\mathcal{W}^\star_{\ell m} - P_{-2} \times {}_{-2}\mathcal{W}^\star_{\ell m}] d\vec{n}, \quad \text{(C5)}$$

with

$$\begin{aligned} {}_2\mathcal{W}^\star_{\ell m} &= W \times \bar{\partial}\,\bar{\partial}(WY^\star_{\ell m}) + 2W \times \bar{\partial}(Y^\star_{\ell m}\bar{\partial}W) \\ &\quad - W \times Y^\star_{\ell m} \times \bar{\partial}\,\bar{\partial}\,W + 2 \times Y^\star_{\ell m} \times (\bar{\partial}W)^2, \end{aligned} \quad \text{(C6)}$$

$$\begin{aligned} {}_{-2}\mathcal{W}^\star_{\ell m} &= W \times \partial\partial(WY^\star_{\ell m}) + 2W \times \partial(Y^\star_{\ell m}\partial W) \\ &\quad - W \times Y^\star_{\ell m} \times \partial\partial W + 2 \times Y^\star_{\ell m} \times (\partial W)^2. \end{aligned} \quad \text{(C7)}$$

Expanding the derivatives and appropriately rearranging the different terms, one easily sees that $\tilde{\chi}^B_{\ell m}$ is given by Eq. (C3). From Eq. (C3), and defining the noise bias as

$$\tilde{N}^B_\ell = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} \langle |\tilde{\chi}^B_{\ell m}|^2 \rangle,$$

with $\tilde{\chi}^B_{\ell m}$ containing noise only, it is straightforward to apply the noise bias calculation performed for the pure pseudo-C_ℓ techniques (see, e.g., Appendix C of Ref. [36]) to get

$$\begin{aligned} \tilde{N}^B_\ell &= \frac{1}{8\pi}\int d\vec{n}(\sigma^2_Q + \sigma^2_U)(N^2_{\ell,2}W^4 + 4N^2_{\ell,1}|\partial W^2|^2 \\ &\quad + |\partial\partial W^2|^2). \end{aligned} \quad \text{(C8)}$$

3. The KN-method case

For such a method, computing the noise bias is more involved. We list here three possible approaches, but none of them allows a numerical calculation of $\tilde{N}^B_\ell$ from $\sigma_{Q/U}$ to be implemented.

(i) *First method*: The pseudo-$a_{\ell m}$'s are derived via

$$\tilde{\chi}^B_{\ell m} = \int M_{\chi^B} \times \tilde{\chi}^B \times Y^\star_{\ell m} d\vec{n}, \quad \text{(C9)}$$

with

$$\tilde{\chi}^B = \frac{i}{2}[\bar{\partial}\,\bar{\partial}(W \times P_2) - \partial\partial(W \times P_{-2}]. \quad \text{(C10)}$$

The noise bias can therefore be expressed as a function of the correlation function of $\tilde{\chi}^B$, denoted by $\mathcal{C}_\chi(\vec{n},\vec{n}') \equiv \langle \tilde{\chi}^B(\vec{n})\tilde{\chi}^B(\vec{n}')\rangle$, and inserting only the noise in $P_{\pm 2}$ when computing such a correlation function,

$$\begin{aligned} \tilde{N}^B_\ell &= \iint_{4\pi} M_{\chi^B}(\vec{n}) \times M_{\chi^B}(\vec{n}') \times \langle \tilde{\chi}^B(\vec{n})\tilde{\chi}^B(\vec{n}')\rangle \\ &\quad \times \sum_{m=-\ell}^{\ell} \frac{Y^\star_{\ell m}(\vec{n})Y_{\ell m}(\vec{n}')}{2\ell+1} d\vec{n}d\vec{n}'. \end{aligned} \quad \text{(C11)}$$

By using the definition of $\tilde{\chi}^B$ as a convolution, i.e.,

$$\begin{aligned} \tilde{\chi}^B(\vec{n}) &= \int W(\vec{n}')[F_+(\vec{n},\vec{n}')P_2(\vec{n}') \\ &\quad - F_-(\vec{n},\vec{n}')P_{-2}(\vec{n}')]d\vec{n}', \end{aligned}$$

and the fact that for noise

$$\begin{aligned} \langle P_{\pm 2}(\vec{n})P_{\mp 2}(\vec{n}')\rangle &= (\sigma^2_Q(\vec{n}) + \sigma^2_U(\vec{n}))\delta(\vec{n}-\vec{n}'), \\ \langle P_{\pm 2}(\vec{n})P_{\pm 2}(\vec{n}')\rangle &= (\sigma^2_Q(\vec{n}) - \sigma^2_U(\vec{n}))\delta(\vec{n}-\vec{n}'), \end{aligned}$$

it is shown that[9]

$$\begin{aligned} \mathcal{C}_\chi(\vec{n},\vec{n}') &= \int W(\vec{n}''))\{[\sigma^2_Q(\vec{n}'') + \sigma^2_U(\vec{n}'')]\mathcal{F}^+_{\vec{n},\vec{n}'}(\vec{n}'') \\ &\quad + [\sigma^2_Q(\vec{n}'') - \sigma^2_U(\vec{n}'')]\mathcal{F}^-_{\vec{n},\vec{n}'}(\vec{n}'')\}d\vec{n}'', \end{aligned}$$

with

$$\mathcal{F}^+_{\vec{n},\vec{n}'}(\vec{n}'') = 2\,\mathrm{Re}[F_+(\vec{n},\vec{n}'')F^\star_+(\vec{n}',\vec{n}'')], \quad \text{(C12)}$$

$$\mathcal{F}^-_{\vec{n},\vec{n}'}(\vec{n}'') = 2\,\mathrm{Re}[F_+(\vec{n},\vec{n}'')F_+(\vec{n}',\vec{n}'')]. \quad \text{(C13)}$$

The above correlation function cannot be further simplified unless we assume full-sky coverage and that the noise is homogeneous. As a consequence, computing the noise bias directly from the noise properties of the Stokes-parameter maps is numerically prohibitive (at least in the not-so-general case of inhomogeneous noise).

(ii) *Second method*: A possible way out—inspired by the computation of the noise bias in the ZB approach—is to start from Eqs. (C9) and (C10) and subsequently perform two integrations by parts in order to *transfer* the derivative operators from $M \times P_{\pm 2}$ to $M_{\chi^B} \times Y^\star_{\ell m}$. Defining Ω_{χ^B} and $\partial\Omega_{\chi^B}$ as the portion of the sky and the contour of such a portion defined by the binary mask M_{χ^B}, we are lead to evaluate one integral on Ω_{χ^B} (denoted as the *domain integral*) and two other integrals on $\partial\Omega_{\chi^B}$ (denotedas *contour integrals*). The integranf of the domain integral is

$$\left(\frac{i}{2}\right)W[P_2 \times \bar{\partial}\,\bar{\partial}(M_{\chi^B}Y^\star_{\ell m}) - P_{-2} \times \partial\partial(M_{\chi^B}Y^\star_{\ell m})].$$

[9]We recall that we fix the noise in the Q map to not be correlated to the noise in the U maps, and that F_- is the complex conjugate of $(-F_+)$.

Because M_{χ^B} is constant-valued on the domain and because $M_{\chi^B} \times W = M_{\chi^B}$—the domain covered by M_{χ^B} is at most the sub-part of the domain covered by W, such as $W = \text{const} = 1$—the integrand is simply given by

$$\left(\frac{i}{2}\right) M_{\chi^B}[P_2 \times \bar{\partial}\,\bar{\partial}(Y^\star_{\ell m}) - P_{-2} \times \partial\partial(Y^\star_{\ell m})].$$

However, the two integrands relative to the contour integrals are of the forms

$$\left(\frac{i}{2}\right)[\bar{\partial}(WP_2) - \partial(WP_{-2})] M_{\chi^B} Y^\star_{\ell m}$$

and

$$\left(\frac{i}{2}\right) W[P_2 \times \bar{\partial}(M_{\chi^B} Y^\star_{\ell m}) - P_{-2} \times \partial(M_{\chi^B} Y^\star_{\ell m})].$$

Because $M_{\chi^B} = 1$ and $W = 1$ on $\partial\Omega_{\chi^B}$, the two contour integrals are *not* vanishing. Although the second contour integral can be expressed as a function of $\langle n_Q(\vec{n}) n_Q(\vec{n}')\rangle$ and $\langle n_U(\vec{n}) n_U(\vec{n}')\rangle$, the first contour integral is still a function of the derivative of $P_{\pm 2}$, preventing us from computing the noise bias.

(iii) *Third method*: This issue of contour-integrals can be naturally circumvented by replacing the binary mask M_{χ^B} with an apodized mask W_{χ^B} satisfying the Dirichlet and Neumann boundary conditions. With such a trick and using the fact that $W = 1$ on Ω_{χ^B}, the multipoles of $\tilde{\chi}^B$ are subsequently given by

$$\tilde{\chi}^B_{\ell m} = \frac{i}{2} \int [P_2 \times \bar{\partial}\,\bar{\partial}(W_{\chi^B} Y^\star_{\ell m}) - P_{-2} \times \partial\partial(W_{\chi^B} Y^\star_{\ell m})].$$

However, the above-defined pseudomultipoles are no more than the definition of the *pure* pseudomultipoles, but now computed on a reduced domain. In other words, replacing M_{χ^B} by an appropriately apodized window function reduces the KN method to the SZ and Zb techniques, but on a *smaller* portion of the sky. With such an implementation of the KN approach, part of the information is inherently lost as compared to the two other approaches, and there would be no reason to use the KN method.

[1] M. Zaldarriaga and U. Seljak, Phys. Rev. D **55**, 1830 (1997).
[2] M. Kamionkowski, A. Kosowsky, and A. Stebbins, Phys. Rev. Lett. **78**, 2058 (1997).
[3] E. M. Leitch, J. M. Kovac, C. Pryke, J. E. Carlstrom, N. W. Halverson, W. L. Holzapfel, M. Dragovan, B. Reddall, and E. S. Sandberg, Nature (London) **420**, 763 (2002); J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel, Nature (London) **420**, 772 (2002).
[4] C. L. Bennett *et al.*, Astrophys. J. Suppl. Ser. **148**, 1 (2003); G. Hinshaw *et al.*, Astrophys. J. Suppl. Ser. **170**, 288 (2007); E. Komatsu *et al.*, Astrophys. J. Suppl. Ser. **192**, 18 (2011).
[5] M. L. Brown *et al.*, Astrophys. J. **705**, 978 (2009).
[6] H. C. Chiang *et al.*, Astrophys. J. **711**, 1123 (2010).
[7] http://www.esa.int/esaSC/120398_index_0_m.html.
[8] U. Seljak and M. Zaldarriaga, Phys. Rev. Lett. **78**, 2054 (1997).
[9] D. N. Spergel and M. Zaldarriaga, Phys. Rev. Lett. **79**, 2180 (1997).
[10] M. Zaldarriaga, Phys. Rev. D **55**, 1822 (1997).
[11] M. Zaldarriaga and U. Seljak, Phys. Rev. D **58**, 023003 (1998).
[12] http://bolo.berkeley.edu/polarbear/?q=science.
[13] http://pole.uchicago.edu/.
[14] http://www.qubic-experiment.org/.
[15] http://www.princeton.edu/act/.
[16] http://cmb.phys.cwru.edu/ruhl_lab/spider.html.
[17] http://groups.physics.umn.edu/cosmology/ebex/.
[18] http://cmbpol.kek.jp/litebird/.
[19] http://www.core-mission.org/.
[20] A. Kogut *et al.*, J. Cosmol. Astropart. Phys. 07 (2011) 025.
[21] M. Zaldarriaga, Phys. Rev. D **64**, 103001 (2001).
[22] M. G. Hauser and P. J. E. Peebles, Astrophys. J. **185**, 757 (1973).
[23] F. Hansen and K. M. Górski, Mon. Not. R. Astron. Soc. **343**, 559 (2003).
[24] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, Astrophys. J. **567**, 2 (2002).
[25] K. M. Smith, Phys. Rev. D **74**, 083002 (2006).
[26] K. M. Smith and M. Zaldarriaga, Phys. Rev. D **76**, 0043001 (2007).
[27] J. Grain, M. Tristram, and R. Stompor, Phys. Rev. D **79**, 123515 (2009).
[28] W. Zhao and D. Baskaran, Phys. Rev. D **82**, 023001 (2010).
[29] J. Kim and P. Naselsky, Astron. Astrophys. **519**, A104 (2010).
[30] J. Kim, Astron. Astrophys. **531**, A32 (2011).
[31] J. Bowyer, A. Jaffe, and D. I. Novikov, arXiv:1101.0520.
[32] M. Tristram, J. F. Macías-Pérez, C. Renault, and D. Santos, Mon. Not. R. Astron. Soc. **358**, 833 (2005).
[33] J. Bock *et al.*, arXiv:0805.4207.
[34] B. Reichborn-Kjennerud *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **7741**, 77411C (2010).
[35] E. F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa, Phys. Rev. D **67**, 023501 (2003).

[36] J. Grain, M. Tristram, and R. Stompor, Phys. Rev. D **86**, 076005 (2012).
[37] G. Hinshaw *et al.*, Astrophys. J. Suppl. Ser. **148**, 135 (2003).
[38] http://lambda.gsfc.nasa.gov/product/map/current.
[39] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005).
[40] D. Larson *et al.*, Astrophys. J. Suppl. Ser. **192**, 16 (2011).
[41] http://www.apc.univ-paris7.fr/APC_CS/Recherche/Adamis/MIDAS09/software/s2hat/s2hat.html.
[42] http://www.apc.univ-paris7.fr/APC_CS/Recherche/Adamis/MIDAS09/software/pures2hat/pureS2HAT.html.
[43] I. O. Hupca, J. Falcou, L. Grigori, and R. Stompor, Lect. Notes Comput. Sci. **7155**, 355 (2012).
[44] M. Szydlarski, P. Esterie, J. Falcou, L. Grigori, and R. Stompor, arXiv:1106.0159.
[45] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).