



HAL
open science

Making maps of Cosmic Microwave Background polarization for B-mode studies: the POLARBEAR example

Davide Poletti, Giulio Fabbian, Maude Le Jeune, Julien Peloton, Kam Arnold, Carlo Baccigalupi, Darcy Barron, Shawn Beckman, Julian Borrill, Scott Chapman, et al.

► **To cite this version:**

Davide Poletti, Giulio Fabbian, Maude Le Jeune, Julien Peloton, Kam Arnold, et al.. Making maps of Cosmic Microwave Background polarization for B-mode studies: the POLARBEAR example. *Astronomy & Astrophysics - A&A*, 2017, 600, pp.A60. <10.1051/0004-6361/201629467>. <hal-01553872>

HAL Id: hal-01553872

<https://hal.science/hal-01553872v1>

Submitted on 31 Oct 2020

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

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



HAL Authorization

Making maps of cosmic microwave background polarization for *B*-mode studies: the POLARBEAR example

Davide Poletti¹, Giulio Fabbian^{2,3}, Maude Le Jeune¹, Julien Peloton⁴, Kam Arnold⁵, Carlo Baccigalupi^{2,3}, Darcy Barron⁶, Shawn Beckman⁶, Julian Borrill^{7,8}, Scott Chapman⁹, Yuji Chinone^{6,10}, Ari Cukierman⁶, Anne Ducout¹¹, Tucker Elleflot¹², Josquin Errard^{13,14}, Stephen Feeney¹¹, Neil Goeckner-Wald⁶, John Groh⁶, Grantland Hall¹⁵, Masaya Hasegawa^{16,17}, Masashi Hazumi^{16,10,17}, Charles Hill⁶, Logan Howe¹², Yuki Inoue^{18,16}, Andrew H. Jaffe¹¹, Oliver Jeong⁶, Nobuhiko Katayama¹⁰, Brian Keating¹², Reijo Keskitalo^{7,8}, Theodore Kisner^{7,8}, Akito Kusaka¹⁹, Adrian T. Lee^{6,19}, David Leon¹², Eric Linder^{8,19}, Lindsay Lowry¹², Frederick Matsuda¹², Martin Navaroli¹², Hans Paar¹², Giuseppe Puglisi^{2,3}, Christian L. Reichardt²⁰, Colin Ross⁹, Praween Siritanasak¹², Nathan Stebor¹², Bryan Steinbach²¹, Radek Stompor¹, Aritoki Suzuki^{6,22}, Osamu Tajima¹⁶, Grant Teply¹², and Nathan Whitehorn⁶

(Affiliations can be found after the references)

Received 3 August 2016 / Accepted 18 December 2016

ABSTRACT

Analysis of cosmic microwave background (CMB) datasets typically requires some filtering of the raw time-ordered data. For instance, in the context of ground-based observations, filtering is frequently used to minimize the impact of low frequency noise, atmospheric contributions and/or scan synchronous signals on the resulting maps. In this work we have explicitly constructed a general filtering operator, which can unambiguously remove any set of unwanted modes in the data, and then amend the map-making procedure in order to incorporate and correct for it. We show that such an approach is mathematically equivalent to the solution of a problem in which the sky signal and unwanted modes are estimated simultaneously and the latter are marginalized over. We investigated the conditions under which this amended map-making procedure can render an unbiased estimate of the sky signal in realistic circumstances. We then discuss the potential implications of these observations on the choice of map-making and power spectrum estimation approaches in the context of *B*-mode polarization studies. Specifically, we have studied the effects of time-domain filtering on the noise correlation structure in the map domain, as well as impact it may have on the performance of the popular pseudo-spectrum estimators. We conclude that although maps produced by the proposed estimators arguably provide the most faithful representation of the sky possible given the data, they may not straightforwardly lead to the best constraints on the power spectra of the underlying sky signal and special care may need to be taken to ensure this is the case. By contrast, simplified map-makers which do not explicitly correct for time-domain filtering, but leave it to subsequent steps in the data analysis, may perform equally well and be easier and faster to implement. We focused on polarization-sensitive measurements targeting the *B*-mode component of the CMB signal and apply the proposed methods to realistic simulations based on characteristics of an actual CMB polarization experiment, POLARBEAR. Our analysis and conclusions are however more generally applicable.

Key words. cosmic background radiation – cosmology: observations

1. Introduction

Cosmic microwave background (CMB) experiments strive to obtain accurate constraints on cosmological parameters from their surveys. Nowadays, in the quest for *B*-modes (the curl-like pattern in CMB polarization), their time-ordered data (TOD) consists of tens of trillions of samples. Map-making, that is reconstructing a map of the observed sky, is one of the major steps in the data analysis pipeline of any CMB experiment. It aims to compress the volume of the data by many orders of magnitudes, while preserving all relevant cosmological information (e.g., Janssen & Gulkis 1992; Wright et al. 1996; Tegmark 1997).

Map-making is usually couched as a linear problem, and the generalized least square (GLS) technique gives a closed form solution which is unbiased and consistent for any positive definite weight matrix. If the weight matrix is the inverse covariance matrix of the time-domain noise, this solution is also optimal as it has the lowest possible uncertainty. In general the principal difficulty in the GLS estimator is in solving the inverse problem, which requires an inversion of a large system matrix. Performing this inversion explicitly (e.g., Borrill 1999; Stompor et al. 2001) has become very difficult as the number

of sky pixels in the surveys increases: the iterative solution has therefore become the standard practice (e.g., Wright et al. 1996; Doré et al. 2001; de Gasperis et al. 2005; Cantalupo et al. 2010; Dünner et al. 2013; Szydlarski et al. 2014).

Essentially every real CMB dataset requires some kind of time-domain filtering. The purpose is usually to remove systematic effects such as drifts of often-unknown origin, correlated noise components, or, in the case of ground based experiments, contributions from the ground and atmosphere. Some of these effects can readily be incorporated into the GLS solution (e.g., Stompor et al. 2001; Cantalupo et al. 2010; Dünner et al. 2013), resulting in an unbiased, or nearly unbiased, map estimate. However, it has become common practice to perform the map-making procedure on the filtered data as if the filtering had not been applied (e.g., Culverhouse et al. 2010; QUIET Collaboration 2011; Schaffer et al. 2011; BICEP2 Collaboration 2014; The POLARBEAR Collaboration 2014). In general this approach is bound to yield an incorrect estimate of the sky signal. However, if the filters are well matched to the data so that the time-domain covariance matrix of the filtered data can be treated as diagonal, this technique usually leads to an enormous simplification of the inverse problem present in the GLS estimator, simplifying the implementation

and reducing the computational cost. This approach hinges on the hope that the bias present in the sky signal in the resulting map can be robustly estimated and corrected for at power spectrum level by studying signal-only simulations (Hivon et al. 2002), and that the loss of statistical precision is minor.

In this work we study the former route and consider map-making procedures that explicitly incorporate and correct for time-domain filtering. We present a general, self-consistent, formalism developed for this purpose and discuss its properties. We point out that even if the time-domain noise is white, such filtering unavoidably leads to correlated noise in the map domain, and in extreme cases to the presence of modes whose amplitude can not be reliably estimated from the data. We then discuss how maps of this kind can be further analyzed using widely popular pseudo power spectrum estimators. Our focus throughout this paper is on the B -mode polarization and our principal figure of merit employed here is the uncertainty on the B -mode polarization power spectrum obtained by the different approaches.

To demonstrate the formalism in a realistic setting, we use simulations based on the scan strategy and noise typical of the first observational campaign of the POLARBEAR instrument (The POLARBEAR Collaboration 2014).

This paper is organized as follows. In Sect. 2 we review the map-making formalism and present an extension to the standard procedure that accounts for time-domain filtering. In Sect. 3, we discuss the effect the filtering may have on the maps reconstructed from the filtered data. In Sect. 4, we introduce specific filters typical of ground-based experiments, while in Sect. 5 we give a worked example demonstrating the effects of such filters on the analysis of mock POLARBEAR-like data sets. We present the main results of this work in Sect. 6 including a performance comparison of the different map estimators, Sect. 6.3.2.

2. Map-making in CMB experiments

This section describes the map-making formalism, emphasizing new features related to the presence of time-domain filtering. For the time being, we assume that the problem is algebraically solvable and leave discussion of potential degeneracies to the next section, Sect. 3.

2.1. The standard problem

The starting point of map-making is the calibrated time-ordered data recorded by the detectors. We collect all these time samples in a vector, \mathbf{d} , which contains thus \mathcal{N}_t elements. The scanning strategy of the telescope and the polarization modulation define how the sky signal contributes to each measured sample, d_t , which can be then represented as,

$$d_t = I_{p_t} + \cos(2\varphi_t)Q_{p_t} + \sin(2\varphi_t)U_{p_t} + n_t. \quad (1)$$

Here n_t is the noise; I , Q and U are the Stokes parameters of the incoming light for sky pixel p being observed at time t ; and φ is the orientation of the linear polarization sensitive detector projected on the sky. We assume throughout that the instrumental beams are axially symmetric and are therefore convolved with the sky signals, which we aim at estimating. There are two important, specific cases of Eq. (1) that we will find useful in this paper. One is that of a total intensity measurement,

$$d_t = I_{p_t} + n_t, \quad (2)$$

and the other of polarization-only measurement,

$$d_t = \cos(2\varphi_t)Q_{p_t} + \sin(2\varphi_t)U_{p_t} + n_t. \quad (3)$$

All the Stokes parameters characterizing the signal in \mathcal{N}_p observed sky pixels can be arranged in a single signal \mathcal{N}_s -vector, \mathbf{s} , in such a way that the Stokes parameters for one pixel are followed by the Stokes parameter for a subsequent one. The entire data vector, \mathbf{d} , can then be represented in a compact way as,

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{n}. \quad (4)$$

Here, \mathbf{n} is the noise vector with covariance matrix \mathbf{C}_n . The *pointing matrix* \mathbf{A} is a \mathcal{N}_t by \mathcal{N}_s known matrix. Each row of \mathbf{A} defines a linear combination of the signal, which contributes to the measurement at the time corresponding to that row. A column of \mathbf{A} is a “time domain signature” of the corresponding entry of \mathbf{s} , telling us when a given sky pixel was observed and with what weight it contributed to the measured signal.

In this form, map-making is a linear statistical problem whose solution is given by the well known GLS estimator,

$$\hat{\mathbf{s}} = (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{M} \mathbf{d}, \quad (5)$$

which renders an unbiased estimate of the map for any choice of positive definite weight matrix, \mathbf{M} . If $\mathbf{M} = \mathbf{C}_n^{-1}$ and the noise is Gaussian, then $\hat{\mathbf{s}}$ defines the minimum variance and maximum likelihood solution. The linear independence of the columns of \mathbf{A} is for the time being taken as given. This is a necessary and sufficient requirement to ensure that $\mathbf{A}^\top \mathbf{M} \mathbf{A}$ is invertible and thus Eq. (5) is uniquely solvable. We discuss this issue in more detail later.

2.2. Mapmaking with time-domain filtering

2.2.1. The filtering operator

Let us collect in a single template matrix, \mathbf{T} , all the temporal templates we want to filter from the data, with each template corresponding to one column. For example, a template can be a \mathcal{N}_t -vector equal to one in some time interval and zero everywhere else. Such a template would stand for the removal of a temporal offset in this interval. The size of \mathbf{T} is considerable: in typical ground-based CMB experiments there are hundreds of templates per detector for every scan period (15–90 min of observation), easily reaching hundreds of millions of templates. The unwanted contributions in the temporal data can be represented as some linear combinations of the columns of \mathbf{T} ,

$$\mathbf{T}\mathbf{x} \equiv \hat{\mathbf{T}}\hat{\mathbf{x}}, \quad (6)$$

where $\hat{\mathbf{T}}$ denotes a column-orthonormal matrix that spans the same subspace as matrix \mathbf{T} and \mathbf{x} and $\hat{\mathbf{x}}$ are the corresponding sets of coefficients. In general, we have,

$$\hat{\mathbf{T}} = \mathbf{T}(\mathbf{T}^\top \mathbf{T})^{-1/2}, \quad (7)$$

so,

$$\hat{\mathbf{T}}^\top \hat{\mathbf{T}} = (\mathbf{T}^\top \mathbf{T})^{-1/2 \top} \mathbf{T}^\top \mathbf{T} (\mathbf{T}^\top \mathbf{T})^{-1/2} = \mathbf{1}, \quad (8)$$

as required.

We note that the number of columns of $\hat{\mathbf{T}}$ may be smaller than that of \mathbf{T} if some of the original templates are not linearly independent. In such cases, inverting matrix $(\mathbf{T}^\top \mathbf{T})^{1/2}$ would require some pseudo-inverse and the result, $(\mathbf{T}^\top \mathbf{T})^{-1/2}$ will be effectively a rectangular matrix.

We can now define the filtering as an operator that projects out from the data all temporal modes defined by the columns of $\hat{\mathbf{T}}$,

$$\mathbf{d}' \equiv \mathbf{d} - \hat{\mathbf{T}}\hat{\mathbf{T}}^\top \mathbf{d} = (\mathbf{1} - \mathbf{T}(\mathbf{T}^\top \mathbf{T})^{-1} \mathbf{T}^\top) \mathbf{d}. \quad (9)$$

We note that if our goal is to filter all modes that belong to the subspace spanned by the columns of \mathbf{T} , and at the same time to weight all the modes that are orthogonal to this subspace by some symmetric weight matrix, \mathbf{M} , then we can generalize the filtering operator in the right hand side of Eq. (9) so it performs both these functions simultaneously. If we further require that the combined operator, \mathbf{F}_T , is also symmetric, such a generalization is unique and the operator reads,

$$\mathbf{F}_T \equiv \mathbf{M} - \mathbf{M}\mathbf{T}(\mathbf{T}^\top\mathbf{M}\mathbf{T})^{-1}\mathbf{T}^\top\mathbf{M}. \quad (10)$$

As desired, this operator obviously filters all the modes from the space spanned by \mathbf{T} as,

$$\mathbf{F}_T\mathbf{T} = 0, \quad (11)$$

and weights by \mathbf{M} any mode, \mathbf{t}_\perp , that is orthogonal to \mathbf{T} , in the sense of a scalar product weighted by \mathbf{M} (i.e., $\mathbf{T}^\top\mathbf{M}\mathbf{t}_\perp = 0$), as

$$\mathbf{F}_T\mathbf{t}_\perp = \mathbf{M}\mathbf{t}_\perp. \quad (12)$$

We emphasize that in general the filtering operator, Eq. (10) is not equivalent to filtering the templates one after another, as is often implemented in practice. Indeed, this would be the case only if the templates happen to be orthogonal from the outset. In such a case, $\mathbf{T}^\top\mathbf{M}\mathbf{T}$ is diagonal (i.e., $\mathbf{t}_i^\top\mathbf{M}\mathbf{t}_j \propto \delta_{ij}$, where \mathbf{t}_i denotes the i th column of \mathbf{T}) and therefore

$$\mathbf{F}_T = \mathbf{M} \prod_i \mathbf{M}^{-1} \mathbf{F}_{t_i} = \mathbf{M} \prod_i \left(1 - \mathbf{t}_i(\mathbf{t}_i^\top\mathbf{M}\mathbf{t}_i)^{-1}\mathbf{t}_i^\top\mathbf{M}\right). \quad (13)$$

If the filters are not orthogonal, and the filtering is performed sequentially, then removing any given template will typically reintroduce some contribution to templates previously filtered, and the final result may depend on the order of the filters. These effects are often small but can sometimes be relevant.

By contrast, the filter proposed in Eq. (10) resolves all ambiguities of this kind. With help of this operator we can now generalize the map-making equation, Eq. (5), as

$$\hat{\mathbf{s}} = (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\mathbf{d}, \quad (14)$$

where \mathbf{F}_T acting upon \mathbf{d} removes all the unwanted modes and weights the others as required, while the matrix operator, $(\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}$, ensures that the estimator is unbiased. Indeed, inserting Eq. (4) for \mathbf{d} we get,

$$\hat{\mathbf{s}} = \mathbf{s} + (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\mathbf{n}, \quad (15)$$

and thus,

$$\langle\hat{\mathbf{s}} - \mathbf{s}\rangle = \langle(\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\mathbf{n}\rangle = 0, \quad (16)$$

where the average is taken over a statistical ensemble of instrumental noise realizations and $\langle\mathbf{n}\rangle = 0$.

The map-domain noise covariance matrix of the unbiased map estimator is (see, e.g., Tegmark 1997 or Stompor et al. 2001)

$$\begin{aligned} \mathbf{C}_{\hat{\mathbf{s}}} &\equiv \langle(\hat{\mathbf{s}} - \mathbf{s})(\hat{\mathbf{s}} - \mathbf{s})^\top\rangle \\ &= (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\langle\mathbf{n}\mathbf{n}^\top\rangle\mathbf{F}_T\mathbf{A}(\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1} \\ &= (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\mathbf{C}_n\mathbf{F}_T\mathbf{A}(\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}. \end{aligned} \quad (17)$$

If $\mathbf{M} = \mathbf{C}_n^{-1}$ then this can be rewritten in a compact way as,

$$\mathbf{C}_{\hat{\mathbf{s}}} = (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}, \quad (18)$$

owing to the fact that, $\mathbf{F}_T\mathbf{M}^{-1}\mathbf{F}_T = \mathbf{F}_T$. This simply generalizes the standard expression for the covariance derived in the case of the maximum likelihood map-making with no filtering, which reads,

$$\mathbf{C}_{\hat{\mathbf{s}}} = (\mathbf{A}^\top\mathbf{C}_n^{-1}\mathbf{A})^{-1}. \quad (19)$$

2.2.2. The meta-pixel approach

We note that an equation analogous to Eq. (14) can be derived from the meta-pixel approach (e.g., Stompor et al. 2001; Cantalupo et al. 2010). In order to do so, let us generalize the data model in Eq. (4) to incorporate the presence of contributions other than the CMB and noise. This can be done as follows,

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{T}\mathbf{y} + \mathbf{w}. \quad (20)$$

where \mathbf{T} is, in general, the set of time-domain templates that we want to filter out, \mathbf{y} is their unknown amplitude, and \mathbf{w} is the noise. We can rewrite Eq. (20) as,

$$\mathbf{d} = [\mathbf{A}, \mathbf{T}]\begin{bmatrix} \mathbf{s} \\ \mathbf{y} \end{bmatrix} + \mathbf{w}. \quad (21)$$

Estimating the unknown parameters for this data model and for the one in Eq. (4) are formally equivalent. We will assume that $[\mathbf{A}, \mathbf{T}]$ is full rank: the time-domain signature of the unknown parameters are all linearly independent, postponing discussion of degeneracies to Sect. 3. The GLS estimator for Eq. (20), written in a block fashion is

$$\hat{\mathbf{s}} = (\mathbf{A}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{F}_T\mathbf{d} \quad (22)$$

$$\hat{\mathbf{y}} = (\mathbf{T}^\top\mathbf{F}_T\mathbf{A})^{-1}\mathbf{T}^\top\mathbf{F}_T\mathbf{d} \quad (23)$$

where \mathbf{F}_T and \mathbf{F}_A are filtering operators as in Eq. (10).

As Eq. (22) shows, $\hat{\mathbf{s}}$ can be computed without explicitly solving for the amplitude of the templates $\hat{\mathbf{y}}$. This is the *direct* approach, adopted in this paper. Moreover, this is exactly the same expression as was derived in the previous section, Eq. (14). There are however several advantages to this derivation. For instance, it shows that if the weights are chosen appropriately (i.e., $\mathbf{M} = \mathbf{C}_n^{-1}$), then the map estimated via Eq. (14) is both maximum likelihood and minimum variance. Moreover if the Bayesian perspective is adopted, then the posterior probability distribution for the combined vector of unknowns in Eq. (21) (i.e., $[\mathbf{s}, \mathbf{y}]$) is Gaussian and the second term in the expression for the filtering operator, \mathbf{F}_T , Eq. (10), arises simply as a result of a marginalization of this posterior over the unknowns contained in \mathbf{y} , assuming flat priors (Stompor et al. 2001).

This derivation of the estimator also suggests that the map can be solved for in two steps rather than one. Indeed, instead of using Eq. (22), if one first gets the estimate $\hat{\mathbf{y}}$, then $\hat{\mathbf{s}}$ can be estimated as,

$$\hat{\mathbf{s}} = (\mathbf{A}^\top\mathbf{M}\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{M}(\mathbf{d} - \mathbf{T}\hat{\mathbf{y}}), \quad (24)$$

which, for simple \mathbf{M} may be much easier to solve than Eq. (22).

This latter approach provides a basis for the destriping technique (e.g., Poutanen et al. 2004; Keihänen et al. 2004, 2010; Tristram et al. 2011). We emphasize that the one- and two-step methods are formally equivalent: they lead to the same estimate of the sky signal. Equation (24) is usually easy and cheap to implement as the weight matrix, \mathbf{M} , is typically taken to be diagonal in this context. Choosing between the two approaches is therefore driven by the cost of Eq. (22) compared to Eq. (23). These two equations require handling dense algebraic objects of a dimension equal to the number of columns of \mathbf{A} and \mathbf{T} respectively. The former is proportional to the number of observed pixels and the latter to the number of temporal templates. As this last number typically grows with the number of detectors and the presence of spurious signals in the observations, the direct method provides a potentially attractive approach for modern experiments, which employ arrays of thousands to tens of thousands of detectors.

This includes both ground-based observatories, such as POLARBEAR (The POLARBEAR Collaboration 2014), and planned CMB satellite missions, like LiteBIRD (Matsumura et al. 2016) or COre (The COre Collaboration 2011). By contrast, the two-step approach is very well suited to experiments with a limited number of detectors observing large sky areas, which produce maps with large numbers of pixels but have relatively few templates to be removed. For these reasons it has played particularly prominent role in the analysis of the *Planck* data (Planck Collaboration VIII 2016; Planck Collaboration VI 2016).

If \mathbf{M} is diagonal then $\mathbf{A}^\top \mathbf{M} \mathbf{A}$ and $\mathbf{T}^\top \mathbf{M} \mathbf{T}$ are usually sparse and structured and therefore easy to compute and invert. Conversely $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$ and $\mathbf{T}^\top \mathbf{F}_A \mathbf{T}$ are dense and potentially large. The size of the former matrix is proportional to the number of pixels, ranging from 10^5 for ground-based experiments to 10^9 for satellites covering the whole sky. The size of the latter matrix also typically exceeds 10^8 for kilo-pixel arrays. These matrices are not only computationally expensive to manipulate but also often too large to explicitly compute and/or store. Consequently, the solution of the inverse problem either in Eqs. (22) or in (23) would ideally be found using some iterative linear equations solvers such the preconditioned conjugate gradient method (PCG; e.g., de Gasperis et al. 2005; Cantalupo et al. 2010; Szydlarski et al. 2014). However, convergence of these iterative solvers may be hard to attain if the matrices are not well conditioned. We discuss the PCG approach in the context of the extended map-making equation, Eq. (14), in Sect. 3.3, from the formal point of view, and in Sects. 5.3.1 and 6.1, in the specific context of POLARBEAR.

2.2.3. The biased map estimator

Performed directly or iteratively, the inversion of $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$ (or $\mathbf{T}^\top \mathbf{F}_A \mathbf{T}$) is the bottleneck in both implementation and execution time. This is why many experiments prefer to use instead the biased map estimator

$$\hat{\mathbf{s}} = (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \mathbf{d}. \quad (25)$$

As \mathbf{F}_T still acts upon the data vector, \mathbf{d} , the templates are still explicitly filtered out of the data. However, this is not accounted for in the system matrix, $(\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1}$, which therefore does not correct for the filtering but only for the weighting. If \mathbf{M} is diagonal, this choice enormously simplifies the implementation and drastically reduces the computational cost. The price to pay is a bias in the estimator. This bias is usually evaluated and removed at the power spectrum level using Monte Carlo simulations and typically requires some additional assumptions (e.g., Hivon et al. 2002). This approach is thus most frequently considered to be a part of the power spectrum estimation pipeline.

The presence of bias in the map is apparent as we have,

$$\text{Bias} = \langle \hat{\mathbf{s}} - \mathbf{s} \rangle = \left[(\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \mathbf{A} - \mathbf{1} \right] \mathbf{s}, \quad (26)$$

which does not vanish in general. The covariance of the estimated map then reads,

$$\begin{aligned} \mathbf{C}_{\hat{\mathbf{s}}} &\equiv \langle (\hat{\mathbf{s}} - \langle \hat{\mathbf{s}} \rangle) (\hat{\mathbf{s}} - \langle \hat{\mathbf{s}} \rangle)^\top \rangle \\ &= (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \langle \mathbf{n} \mathbf{n}^\top \rangle \mathbf{F}_T \mathbf{A} (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \\ &= (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \mathbf{C}_n \mathbf{F}_T \mathbf{A} (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1}. \end{aligned} \quad (27)$$

Therefore, if the filters are chosen such that the matrix $\mathbf{F}_T \mathbf{C}_n \mathbf{F}_T$ is nearly diagonal for any diagonal weights, \mathbf{M} , then we can take them to be, $\mathbf{M} = \text{diag}(\mathbf{F}_T \mathbf{C}_n \mathbf{F}_T)$, yielding,

$$\mathbf{C}_{\hat{\mathbf{s}}} \approx (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1}, \quad (28)$$

and therefore the covariance can have a particularly simple structure. However this is rarely the case, and instead even if the true noise is uncorrelated and $\mathbf{M} = \mathbf{C}_n^{-1}$, the covariance reads

$$\mathbf{C}_{\hat{\mathbf{s}}} = (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \mathbf{A} (\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1}, \quad (29)$$

and thus is a dense matrix with potentially non-negligible, off-diagonal correlations.

We note that the information content of both the unbiased and biased maps is the same, since one can be derived from the other via an invertible linear operation. Indeed,

$$\hat{\mathbf{s}}_{\text{unbiased}} = \mathbf{R} \hat{\mathbf{s}}_{\text{biased}}, \quad (30)$$

where $\mathbf{R} \equiv (\mathbf{A}^\top \mathbf{F}_T \mathbf{A})^{-1} (\mathbf{A}^\top \mathbf{M} \mathbf{A})$ is an invertible matrix. However, the key to the biased approach is that no attempt is ever made to estimate the matrix \mathbf{R} . Consequently, although the same information is contained in both maps it is encoded differently in each of them, and whenever equivalent biased and unbiased maps are subsequently processed their information is compressed differently, giving rise to different statistical properties in the resulting estimator.

3. Map-making in the presence of degeneracies

So far we have assumed that the generalized map-making equation can be robustly solved, implicitly assuming that the system matrix, $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$, is invertible. However this may not always be the case, and in this section we elaborate on this and discuss what can be done in such circumstances.

We first note that invertibility of the system matrix is ensured if the matrix $[\mathbf{A}, \mathbf{T}]$ is full column rank. This can be seen immediately by noting that $(\mathbf{A}^\top \mathbf{F}_T \mathbf{A})^{-1}$ stands for an upper left diagonal block of the inverse of the matrix,

$$\begin{bmatrix} \mathbf{A}^\top \\ \mathbf{T}^\top \end{bmatrix} \mathbf{M} [\mathbf{A}, \mathbf{T}] = \begin{bmatrix} \mathbf{A}^\top \mathbf{M} \mathbf{A} & \mathbf{A}^\top \mathbf{M} \mathbf{T} \\ \mathbf{T}^\top \mathbf{M} \mathbf{A} & \mathbf{T}^\top \mathbf{M} \mathbf{T} \end{bmatrix}, \quad (31)$$

which is invertible only when $[\mathbf{A}, \mathbf{T}]$ is full column rank. In practice, since all the operations have to be performed numerically, what really matters is not strict linear independence in the mathematical sense but rather linear independence sufficient to ensure stable and robust, finite-precision numerical calculations, as exemplified by Eqs. (22), (23), (10), (24) and (25).

In general, given a matrix \mathbf{B} and vector \mathbf{z} laying in its range, if \mathbf{B} is singular we can solve the equation $\mathbf{B} \mathbf{x} = \mathbf{z}$ for \mathbf{x} only down to an unknown contribution from the nullspace of \mathbf{B} . Typically, the component of the solution parallel to the nullspace will be arbitrarily set to zero and its true value unavoidably lost. In practice it could be obtained by regularizing the matrix, which involves first calculating its inverse via computing and inverting its eigenvalues. The regularization is then applied to all eigenvalues that are smaller than some predefined threshold by setting to zero the corresponding eigenvalue of the inverse.

$[\mathbf{A}, \mathbf{T}]$ may not be full column rank for three different reasons, leading to three possible types of degeneracies:

3.1. The columns of \mathbf{A} are not independent

This degeneracy would affect standard map-making as much as our extension, but we include it for completeness. In this case, the scanning strategy does not allow the reconstruction of some sky mode. A typical example is a polarization pixel that was not observed with sufficient redundancy in the polarization angle.

This case is easy to avoid because $\mathbf{A}^\top \mathbf{M} \mathbf{A}$ is easy to build. If \mathbf{M} is diagonal the cost is $O(N_r)$ operations and $\mathbf{A}^\top \mathbf{M} \mathbf{A}$ is block

diagonal (a block for each pixel); its eigendecomposition (costing $O(N_p)$ operations) then enables us to evaluate the condition number of each of the blocks. After pixel selection based on their condition numbers (and the removal of the corresponding samples from the TOD) the new $\mathbf{A}^T \mathbf{M} \mathbf{A}$ can be safely inverted.

3.2. The columns of \mathbf{T} are not independent

This reflects the fact that there is redundancy in the templates, and basically corresponds to an attempt to filter the same template twice. For example, this happens in practice when two different sets of templates (e.g., the polynomial and the ground template filters) both remove the global offset of the TOD.

Since the final goal is to estimate $\hat{\mathbf{s}}$ and not $\hat{\mathbf{y}}$, Eqs. (22) and (23), this degeneracy does not pose any fundamental problem. We merely need to construct a basis of $\text{span}(\mathbf{T})$ and use it to define a new (smaller) set of independent templates $\hat{\mathbf{T}}$ as described in Sect. 2.2.1.

In practice, the situation is also quite straightforward. By construction, there are usually many known orthogonality relations between the templates. As a consequence, $\mathbf{T}^T \mathbf{M} \mathbf{T}$ is typically reasonably easy to compute and is sparse and structured. Its inverse can then be computed explicitly using standard matrix inversion techniques. The condition number of this matrix provides an easy test of the linear independence of the templates, and if it is too high the matrix has to be regularized while being inverted. Once such a regularized inverse $(\mathbf{T}^T \mathbf{M} \mathbf{T})^{-1}$ is computed, it should be used in the projector of Eqs. (22) and (25), to take care of the redundancies and therefore the degeneracies.

3.3. Some columns of \mathbf{A} are not independent from the columns of \mathbf{T}

This is the most insidious type of degeneracy in the map-making problem, and its presence reflects a degeneracy between the sky signal and the templates. In this case the reconstruction of some sky component is not possible if the templates have been filtered. As a trivial example, by systematically filtering the mean of the total intensity TOD we create a degeneracy with the global offset of the temperature map.

This type of the degeneracy manifests itself as a singularity of both $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and $\mathbf{T}^T \mathbf{F}_A \mathbf{T}$. This can be seen immediately by noting that, if the columns of \mathbf{A} and \mathbf{T} are not independent, then there exists at least two modes, one in the map domain, $\tilde{\mathbf{s}}$, and one in the template domain, $\tilde{\mathbf{y}}$, such as,

$$\mathbf{A} \tilde{\mathbf{s}} = \mathbf{T} \tilde{\mathbf{y}}, \quad (32)$$

and therefore

$$\mathbf{A}^T \mathbf{F}_T \mathbf{A} \tilde{\mathbf{s}} = \mathbf{A}^T \mathbf{F}_T \mathbf{T} \tilde{\mathbf{y}} = 0, \quad (33)$$

$$\mathbf{T}^T \mathbf{F}_A \mathbf{T} \tilde{\mathbf{y}} = \mathbf{T}^T \mathbf{F}_A \mathbf{A} \tilde{\mathbf{s}} = 0, \quad (34)$$

given that $\mathbf{F}_T \mathbf{T} = \mathbf{F}_A \mathbf{A} \equiv 0$, Eq. (11). The two modes, $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$, constitute a pair of degenerate modes, which, while residing in different domains, lead to the same effects in the time-domain data and therefore can not be distinguished from each other.

In this case the best one can do to solve the map-making problem is to regularize the inversion of the singular matrix and to compute all the modes of the map for which the solution can be obtained and determine the modes for which it cannot. These latter modes will be missing from the reconstructed

map. We note that this is not due to the regularization procedure but because these modes are removed from the data by the filters. Indeed,

$$\mathbf{A} \mathbf{F}_T \mathbf{d} = \mathbf{A} \mathbf{F}_T \mathbf{A} \mathbf{s} + \mathbf{A} \mathbf{F}_T \mathbf{n} = \mathbf{A} \mathbf{F}_T \mathbf{A} \mathbf{s}_\perp + \mathbf{A} \mathbf{F}_T \mathbf{n}, \quad (35)$$

where the subscript \perp denotes the part of the sky signal orthogonal to the nullspace of $\mathbf{A} \mathbf{F}_T \mathbf{A}$, as the part parallel to it is unrecoverably lost. The information about these removed modes then needs to be propagated to next steps of the analysis and properly taken into account to ensure that the final results are statistically meaningful.

This route is only straightforward in practice if all the matrices appearing in Eq. (31), can be constructed and decomposed explicitly. However, because of their sizes this may be a formidable and often unfeasible task, even with help of the largest massively parallel platforms and parallel software packages. If this is indeed the case, and the solution can be only derived using some iterative technique, then the singular modes may not only be impossible to compute and correct for, but indeed it may not be clear from the outset whether the matrices are singular or not. In such cases this may need to be inferred post hoc from the behavior of the solver. We discuss these issues in more depth in Sect. 6.1.

We also emphasize that, in the presence of this kind of degeneracy, the maps computed by the direct and two-step approaches may not be identical. Indeed, in the direct case the unconstrained sky modes will be missing from the estimated map, while in the two-step case the situation is different as the regularized inversion has to happen when the estimate of the template amplitudes is performed. Consequently, these will be degenerate ‘‘modes’’ of the templates, which will be missing in $\hat{\mathbf{y}}$, while the template-corrected data vector, $\mathbf{d} - \mathbf{T} \hat{\mathbf{y}}$, will retain a time-domain component that should have been filtered out. Hence, the map estimated in the second step, Eq. (24), will contain some of the degenerate sky modes, which were set to zero in the direct approach. Obviously, if the fact that these modes can not be estimated is properly taken into account in the covariance of the maps, both maps will be statistically equivalent. We also note that to compute which map modes are non-constrainable one may use the singular modes of $\mathbf{T}^T \mathbf{F}_A \mathbf{T}$, found while performing the first step of the two-step method, and use them in the second step by replacing the template-corrected data vector, $\mathbf{d} - \mathbf{T} \hat{\mathbf{y}}$ by $\mathbf{T} \mathbf{z}$, where \mathbf{z} stands for one of the singular eigenvectors. Map-domain templates resulting from this calculation will have to be then orthogonalized using, for example, the Gram-Schmidt procedure.

In the case of the biased map estimator, Eq. (25), this degeneracy does not pose any numerical issue. As in the unbiased cases the map estimate will have no contribution from the sky signal modes residing in the nullspace of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ due to the filtering applied in the time domain, Eq. (35). Nevertheless, the amplitudes of the filtered modes in the biased map will usually be non-zero as a result of the power leaking to them from the other pixel-domain modes.

4. The case of ground-based experiments

As a specific application of the above formalism let us consider the map-making problem for a modern, ground-based, CMB experiment, which scans the sky with a kilo-pixel array of polarization sensitive detectors in the presence of both atmospheric and ground emission. Commonly, in order to minimize the impact of the atmospheric contribution, the scans are performed at constant elevation for relatively short periods of time. The elevation is then changed to track the sky patch, which moves due to

the Earth's rotation. The scan amplitude, the choice of the scan elevations, and the elevation dwelling time are specific to each observation. In the following, we consider only cases that conform to this general description.

The data of such an experiment are typically more complex than the simple model exemplified by Eq. (4), whether due to the atmospheric, instrumental, or ground contamination. Below we describe common contributions of this type and explain how they fit within the extended data model, Eq. (20), and how they can be treated with the map-making technique introduced earlier.

4.1. Noise correlations

The noise is usually correlated between different detectors and typically displays significant low frequency excess, dubbed $1/f$ noise, which arises either as an instrumental effect, or from atmospheric emission, or from some other effects. As a result the time-domain noise correlation matrix, \mathbf{C}_n , is dense, making its inversion and multiplication computationally demanding. Moreover, the matrix, \mathbf{C}_n , is unknown and has to be estimated from the data themselves (e.g., Ferreira & Jaffe 2000). This procedure is usually difficult, especially as far as the low frequency modes and detector-detector correlated modes are concerned. In addition, the low frequency contributions may not be even stationary or Gaussian, and thus can not be properly described merely by a covariance. Consequently, using the optimal weight, \mathbf{C}_n^{-1} , in the map-making process may be difficult. Using a simpler weight matrix, \mathbf{M} , does not lead to a bias in the standard map estimator, Eq. (5), however its choice does affect the quality of the final map. This is because for different weights, the time-domain data, \mathbf{d} , are coadded differently on the first step of the map-making procedure, when the right hand side of the map-making equation, $\mathbf{A}^T \mathbf{M} \mathbf{d}$, is calculated. For instance, diagonal weights (in the map domain) can not selectively suppress some temporal frequency bands over others. Consequently, if $1/f$ noise is present, even if it is Gaussian and stationary, the low frequency modes will not be properly down-weighted as compared to the high frequency ones. This may result in stripes appearing in the direction of the scans. The effect is particularly apparent if the scanning strategy does not provide good cross-linking. Non-Gaussian/non-stationary features can be even more difficult to deal with.

Instead of down-weighting such long modes one may prefer to filter them out, as in the filtered map-making, Eq. (14). The time-domain data are then explicitly filtered while being compressed to the pixel-domain object, $\mathbf{A}^T \mathbf{F}_T \mathbf{d}$. The long term trends present in the time-ordered data can be removed at this stage. Such removal is blind to the origin and nature of the trends, potentially removing the true sky signal together with the noise and some unwanted spurious contributions. However, the signal-to-noise ratio for these modes is usually very low, as the $1/f$ noise quickly dominates, and the information loss due to the potential removal of the sky signal is typically negligible. If this is the case, then filtering can provide a useful alternative to weighting.

The modes to be filtered out are typically assumed to be arbitrary linear combinations of some sufficiently rich family of temporal templates, which has to be defined case-by-case. For long term trends, the templates are often taken to be piece-wise low order polynomials or harmonic functions defined for all samples of the data, and are represented by columns of some template matrix, \mathbf{B} . If the templates are well matched to the problem at hand, then the residual noise, \mathbf{w} , defined as,

$$\mathbf{w} \equiv \mathbf{n} - \mathbf{B}\mathbf{x}, \quad (36)$$

is expected to be “whiter” than the actual time-domain noise, \mathbf{n} , and thus approximately characterized by a diagonal noise covariance matrix. Here, \mathbf{x} stands for a set of a priori arbitrary parameters to be determined, similar to the sky signal, \mathbf{s} , which define the amplitudes of the corresponding templates (Stompor et al. 2001). We point out that in general one could introduce some additional information about the long term modes of the noise by setting some constraints on \mathbf{x} . As such constraints are typically hard to identify and might not lead to a significant improvement in the sensitivity, we will not consider them in this paper (see, e.g., Keihänen et al. 2010, for an implementation of this idea in the context of the two-step map-making). As we will see below, in their absence the map-making will discard all the information in the data that matches the time-domain signature of the templates, effectively filtering all template-like modes out of the data. As previously explained, this causes a loss of a sky mode, $\tilde{\mathbf{s}}$, only if its time-domain projection $\mathbf{A}\tilde{\mathbf{s}}$ is a template-like mode. In general, subsequent observations break the degeneracy that a template might have with a subset of the dataset. There are two well-known exceptions. First, the global offset of the time streams is always filtered, as a consequence the global offset of the temperature map is unconstrained. Second, when observing from the Earth's poles, scans at constant elevation always probe constant declination stripes, with different elevations corresponding to different declinations. Filtering the offset from the time streams of each of these scans prevents the reconstruction of the offset of each of these constant declination stripes. These can be partially recovered if the stripes share some of the sky pixels owing to the assumption that the sky signal is constant across a pixel. The resulting constraints would typically be weak, in particular for the relative offset of two stripes that are not directly adjacent, which could be constrained only via the intermediary ones. Consequently, in such cases one should expect poorly constrained long modes in declination.

4.2. Ground pickup

Ground-based experiments usually have non negligible ground-synchronous signal contaminating their TOD. The most common source is ground pickup in the far side-lobes of the beam, although the magnetic response of the experiment to the Earth's magnetic field may also be a concern. Experiments try to minimize these side lobes but, as the ground is very bright by CMB standards, their contribution typically cannot be neglected. This ground-synchronous signal could be thought of as a two-dimensional template in Earth-bound coordinates. However, in practice the situation is more complex as the signal can vary in time and may be different for different detectors, as the level and structure of their side-lobes may be very different (and poorly known). Therefore we typically model the ground signal as a one-dimensional template specific to each constant elevation scan and to each detector, or at the very least to a group of detectors located sufficiently close together in the focal plane.

Such a one-dimensional ground template can be parameterized with the azimuth of the observation and represented by one dimensional discretized map with entries standing for the amplitude of the ground signal in each of the disjoint, consecutive azimuth bins. Following our previous argument, we need to introduce such a template for each detector and each constant elevation scan separately. These, concatenated together, are then denoted as \mathbf{g} . In the presence of ground pickup, the time-domain data can be then modeled as a sum of three terms: the sky signal term, $\mathbf{A}\mathbf{s}$, the ground-pickup term, $\mathbf{G}\mathbf{g}$, and the

noise, n . We can then write,

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{G}\mathbf{g} + \mathbf{n}. \quad (37)$$

This is merely a specialized version of Eq. (20). The role played by the matrix \mathbf{G} is analogous to that played by the sky pointing matrix, \mathbf{A} . Adopting the simple binned ground template model introduced above, each column of \mathbf{G} is associated with some specific azimuthal bin assigned to some specific detector and some constant elevation scan. This column will be composed of ones and zeros, with 1 indicating that the given measurement was made by the specific detector and was performed within the specific scan, when the observation's azimuth falls within the specific azimuthal bin. Therefore applying \mathbf{G} to the template, \mathbf{g} , will add the same value of the ground pick up to all these measurements. In this section we summarize the effects the modeling of the ground pick-up may have on the quality of the map recovered with the unbiased map estimator, leaving more detailed discussion to Appendix A.

We start by considering a site far from the Earth's poles. For a single constant-elevation scan, samples having the same ground pickup will correspond to measurements taken by a single detector with the azimuthal position falling into one of the azimuthal bins. As time progresses and the sky rotates in Earth-based coordinates (see Fig. 1 for a sketch of the geometry of the problem) the measurements will cover the sky area extending along the RA direction in the sky coordinates. The size in declination of the area depends on the scan elevation and the azimuth of the bin. The sky areas covered by measurements corresponding to two different azimuthal bins will be disjoint and as both these subsets are affected by a different ground pickup amplitude, their relative offsets will be unconstrained. As a consequence, a single detector map will have multiple degenerate modes, each corresponding to a sky patch swept by a different azimuthal bin. The degeneracies can be in part removed if data of another detector are used, but only if the azimuthal bins of the latter are shifted with respect to those of the first detector in such a way that their corresponding sky patches are also shifted and each patch of the second detector includes pixels from two patches adjacent to the first one. Yet another factor breaking the degeneracy here is the sky pixelization, as the sky pixels crossing the boundary of two adjacent bin sky patches will constrain their relative amplitude. However, in all these cases the degeneracy-breaking may be quite weak because the overlaps typically involve only a limited number of neighboring sky patches corresponding to single azimuthal bins. For observations taken from the poles the situation is different. Since the sky movement is in the azimuthal direction the same sky can be measured in different azimuthal bins. Consequently, there is a significant constraining power on the relative offsets, leaving the overall offset of the observed constant declination stripe as the only truly degenerate mode. We already encountered this degenerate mode in Sect. 4.1: the offset of the TOD is filtered also when removing correlated noise. The consequences of this degenerate mode and the possible degeneracy breaking effects were already discussed at the end of Sect. 4.1.

4.3. Recap

The following useful conclusions have been drawn out in this section (and are borne out further by a more detailed analysis in Appendix A). In the presence of time-domain filtering of the kinds discussed here, only relatively few sky modes are expected to be genuinely degenerate. Nevertheless, a few extra ill-conditioned modes with large variance should also be

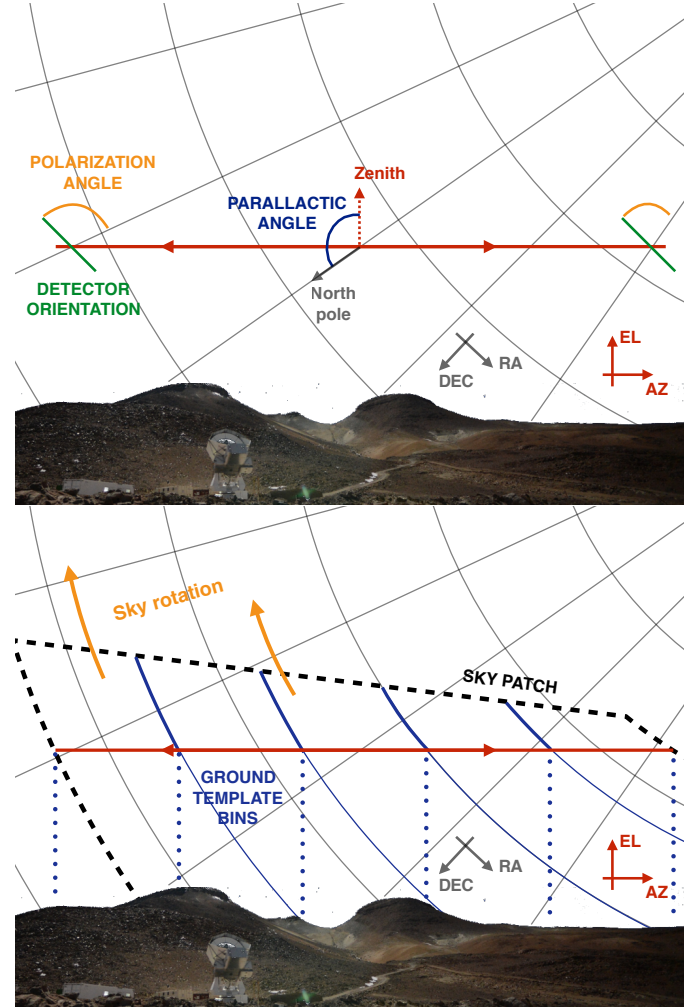


Fig. 1. *Top panel:* geometry of a constant elevation observation. Grey lines represent the equatorial coordinate system at some fixed time instant. The red line shows the scan in the horizontal coordinate system with the telescope assumed to chop back and forth along a constant elevation direction. The orientation of a polarization sensitive detector is shown with green lines. It is assumed to be fixed in horizontal coordinates and thus varies somewhat with the observation's azimuth due to changes in the parallactic angle, which is marked in blue. *The bottom panel* emphasizes the effects due to the sky rotation. This is marked with orange arrows. As the sky rotates, the constant elevation scan progressively covers the sky area delineated by a dashed black line. The area above the red line has been already observed. Also marked are its sub-areas, which are observed when the telescope's azimuth falls into one of the azimuthal bins. These for instance may be used to discretize the ground-pickup and are shown here with blue dotted lines. These sub-areas remain disjoint as the sky rotates whenever the azimuthal bins are disjoint, and are separated in the figure with blue solid lines.

expected. They are mainly related to the filtering of the ground signal, which leaves poorly constrained modes in the declination direction. The details of these modes will depend on choices made regarding pixelization, definition of the ground template bins, and scanning strategy, but they will be more prominent for observations conducted from the poles. Given this, biased map-making – by construction oblivious to the presence or absence of such modes – may be seen as providing a more convenient and adaptable way to perform the analysis. Indeed, it has been the approach of choice for multiple past analyses of this kind of data sets (e.g., Culverhouse et al. 2010; Schaffer et al. 2011; BICEP2 Collaboration 2014). However a

relevant question, which we discuss in more detail in the remainder of this paper, is whether it is feasible for a ground-based experiment to produce an unbiased estimate of the sky signal, and if so, with what fidelity.

5. Worked example: the POLARBEAR experiment

This section applies the map-making formalism to realistic simulations broadly based on the POLARBEAR experiment. We begin by reviewing the experimental characteristics of POLARBEAR, and then describe details of the map-making process. We conclude by considering how these different approaches impact the final map properties and, most importantly, the measurement of the B -mode polarization power spectrum.

Though the results and conclusions are strictly speaking specific to the POLARBEAR experiment, they are expected to be qualitatively relevant for the other on-going and planned ground-based experiments, which typically implement a similar data reduction procedure.

5.1. POLARBEAR experiment and its observations

POLARBEAR is a dedicated CMB B -mode experiment operating from the James Ax Observatory in the Atacama Desert in Northern Chile. It is composed of a cryogenic receiver mounted on the *Huan Tran* telescope (HTT; Tran et al. 2008). HTT is a two mirror off-axis Gregorian telescope designed to ensure low cross polarization. A 4 K aperture stop in the receiver creates a 2.5 m illumination pattern on the primary mirror, resulting in a beam size of 3.5' $FWHM$. The receiver hosts 637 focal plane pixels (1274 transition edge bolometers) sensitive to a spectral band centred at 148 GHz with 26% fractional bandwidth (Arnold et al. 2012). Each focal plane pixel contains two detectors, henceforth called a “detector pair”, with each detector sensitive to a different orthogonal polarization direction. The diameter of the field of view of the focal plane is 2.4°. The optics of the receiver includes a stepped half-wave plate that can modulate the polarization state of the incoming radiation. Since the optics and the receiver are installed on a moving mount, the azimuth and elevation pointing direction can be controlled during the observations.

The observation considered in the following analysis are based on the first POLARBEAR campaign, providing a high level of realism to the pointing and polarization information in our simulations. However, the simulations do not reflect all of the properties of the data sets used in the actual data analysis. Though POLARBEAR observed three patches, we restrict ourselves to the RA23 patch (RA 23^h1^m48^s Dec -32°48'). The nominal area of the patch, as used in the analysis of The POLARBEAR Collaboration (2014) is 8.8 deg², but the area actually observed and considered here is much larger, amounting to approximately 43 deg². Due to the sky rotation, the patch rises and sets, reaching a maximum elevation of 82 deg while the minimum elevation for observation is 30 deg. The scanning consists of 15-min constant elevation scans (CES; $\mathcal{N}_t^{\text{CES}} \sim 25\,000$ measurements at a rate of 31.8 Hz) in which the telescope moves at a constant velocity of 0.75 deg/s in an azimuth range of 3°. We refer to a single left-to-right sweep of the telescope as subscan; there are typically $\mathcal{N}_{\text{sub}} \sim 150$ subscans in a CES. The CES ends when the patch leaves the field of view of the telescope, both the azimuth and the elevation are then modified and a new CES is performed at the new position of the patch. The half-wave plate position is constant during a CES. During the first season, its orientation was rotated by 11.25 deg every one-two days during the first half and occasionally during

the second half. More details about the observation can be found in (The POLARBEAR Collaboration 2014).

5.2. Time domain data model and simulations

In this section we illustrate in detail our data model, providing a concrete example for the general considerations presented earlier in Sect. 2 and Sect. 4. Our goal is to investigate the effects that are inherently due to the filtering, rather than the ones that stem from a failure of the applied filtering to remove unwanted contributions. Consequently, our simulations employ an idealized data model, neglecting important properties that real data usually have: imperfect polarization efficiency, noise correlated between time samples and/or detectors, temperature to polarization leakage due to differential gain, beam or bandpass between the two orthogonal detectors, etc., implicitly assuming that all such effects can be removed by the filters. We adopt the filters defined by the proposed data model and used for the actual POLARBEAR data.

For convenience, we first consider data collected during a single CES by a single pair of detectors in the same focal plane pixel. This is the fundamental unit of our data model and we have a total of $\mathcal{N}_{\text{pair}} \times \mathcal{N}_{\text{CES}} \sim 1.5 \times 10^6$ of them. The two detectors in such a pair are sensitive to two orthogonal polarizations and are referred to as \parallel and \perp . We model their TOD to include signal, ground pickup, and templates related to the correlated noise and an actual noise term,

$$\begin{aligned} \begin{bmatrix} d^{\parallel} \\ d^{\perp} \end{bmatrix} &= \mathbf{A}s + \mathbf{G}g + \mathbf{B}x + w \\ &= \begin{bmatrix} \mathbf{A}^{\parallel} \\ \mathbf{A}^{\perp} \end{bmatrix} s + \begin{bmatrix} \mathbf{G}^d & 0 \\ 0 & \mathbf{G}^d \end{bmatrix} \begin{bmatrix} g^{\parallel} \\ g^{\perp} \end{bmatrix} + \mathbf{B}x + \begin{bmatrix} w^{\parallel} \\ w^{\perp} \end{bmatrix}, \end{aligned} \quad (38)$$

where we have arranged the data vector so that all the measurements of the first detector of the pair are gathered together and followed by the measurements taken by the other one.

The pointing matrices, $\mathbf{A}^{\parallel/\perp}$, are as given by Eq. (1), and thus have only three non-zero elements for each row, which correspond to three Stokes parameters of a given sky pixel, p , observed at the time assigned to the row. These are equal to one, $\cos(2\varphi_t^{\parallel})$ and $\sin(2\varphi_t^{\parallel})$ for detector \parallel and $1, \cos(2\varphi_t^{\perp}) = -\cos(2\varphi_t^{\parallel})$ and $\sin(2\varphi_t^{\perp}) = -\sin(2\varphi_t^{\parallel})$ for detector \perp , and for the I, Q and U signal components, respectively. The block-diagonal structure of matrix \mathbf{G} is due to the fact that we have introduced two independent ground templates, one for each detector of the pair. Since we will always use the boresight azimuth to define the azimuthal bins, each of the blocks is the same for each detector and all focal plane pixels. Matrix \mathbf{B} describes the time-domain filtering and thus can have more complex structure. In particular it needs to account for two types of contributions: the ones that are correlated between detectors and the ones that are independent. This can be achieved by assuming that matrix \mathbf{B} has the following structure,

$$\mathbf{B} \equiv \begin{bmatrix} \mathbf{B}_{\text{corr}} & \mathbf{B}_{\text{uncorr}} & 0 \\ \mathbf{B}_{\text{corr}} & 0 & \mathbf{B}_{\text{uncorr}} \end{bmatrix}, \quad (39)$$

where we assumed that we use the same filtering of the uncorrelated part for each of the detectors. This corresponds to the following breakdown of vector x ,

$$x \equiv \begin{bmatrix} x_{\text{corr}} \\ x_{\text{uncorr}}^{\parallel} \\ x_{\text{uncorr}}^{\perp} \end{bmatrix}. \quad (40)$$

Here, \mathbf{x}_{corr} collects the amplitudes of all the time-domain modes common to both the detectors, while $\mathbf{x}_{\text{uncorr}}^{\parallel/\perp}$ those specific to only one of them.

Owing to the orthogonality of the two polarization directions for the two detectors in a pair, we can represent their data with summed and differenced data streams, \mathbf{d}^+ , \mathbf{d}^- , which contain the information about total intensity and polarized sky signals, respectively. These are defined as,

$$\mathbf{d}^+ \equiv \frac{1}{2}(\mathbf{d}^{\parallel} + \mathbf{d}^{\perp}) \quad (41)$$

$$\mathbf{d}^- \equiv \frac{1}{2}(\mathbf{d}^{\parallel} - \mathbf{d}^{\perp}). \quad (42)$$

Using Eq. (38) and introducing quantities specific to each of the new data streams, defined as,

$$\mathbf{A}^{\pm} \equiv \frac{1}{2}(\mathbf{A}^{\parallel} \pm \mathbf{A}^{\perp}) \quad (43)$$

$$\mathbf{B}^+ \equiv \begin{bmatrix} \mathbf{B}_{\text{corr}} & \mathbf{B}_{\text{uncorr}} \end{bmatrix}, \quad (44)$$

$$\mathbf{B}^- \equiv \mathbf{B}_{\text{uncorr}} \quad (45)$$

$$\mathbf{x}^+ \equiv \begin{bmatrix} \mathbf{x}_{\text{corr}} \\ \frac{1}{2}(\mathbf{x}_{\text{uncorr}}^{\parallel} + \mathbf{x}_{\text{uncorr}}^{\perp}) \end{bmatrix}, \quad (46)$$

$$\mathbf{x}^- \equiv \frac{1}{2}(\mathbf{x}_{\text{uncorr}}^{\parallel} - \mathbf{x}_{\text{uncorr}}^{\perp}), \quad (47)$$

$$\mathbf{g}^{\pm} \equiv \frac{1}{2}(\mathbf{g}^{\parallel} \pm \mathbf{g}^{\perp}), \quad (48)$$

$$\mathbf{w}^{\pm} \equiv \frac{1}{2}(\mathbf{w}^{\parallel} \pm \mathbf{w}^{\perp}), \quad (49)$$

we can express these new data in a concise way as,

$$\mathbf{d}^{\pm} = \mathbf{A}^{\pm} s_{T/QU} + \mathbf{G}^d \mathbf{g}^{\pm} + \mathbf{B}^{\pm} \mathbf{x}^{\pm} + \mathbf{w}^{\pm}. \quad (50)$$

Here s_T and s_{QU} denote sky signal vectors made of the total intensity and interleaved, pixel-by-pixel, Q and U Stokes parameters, respectively. These expressions emphasize that as intended each of the new data streams contains information either about the total intensity, \mathbf{d}^+ , or polarization, \mathbf{d}^- . Moreover, as all the amplitudes appearing on the right hand side of these equations are specific for each data set, each of these two data sets can be analyzed completely separately and, under the assumptions specified earlier, without any loss of accuracy. Specifically, the maps of the total intensity on the one hand, and the Q and U Stokes parameters on the other can be estimated independently. This is the approach we follow in this work.

We point out that a perfect separation of the total intensity and polarization information is strictly speaking only possible if the two detectors of each pixel pair are perfectly calibrated and have identical beams. Otherwise, some residual total intensity contribution may be present in \mathbf{d}^- and, less harmfully, some polarization in \mathbf{d}^+ . The beams of two detectors are more likely to be similar if they belong the same focal plane pixel. Therefore, using \mathbf{d}^- to constrain the polarization may be less susceptible to total intensity leakage due to beam differences than performing global separation of the three Stokes parameters directly using $\mathbf{d}^{\parallel/\perp}$ as inputs. In any case, if needed, leakage from the total intensity to the differenced data, \mathbf{d}^- , can be modeled as a total intensity-like template and used in the map-making process. Though such tests were indeed performed as part of the analysis of the actual POLARBEAR data set (The POLARBEAR Collaboration 2014), we do not consider them in the present work. Leaving aside this kind of systematic

effect, it is mathematically equivalent whether we use one or the other data representation, as long as the filtered temporal templates, defined by \mathbf{B}_{corr} and $\mathbf{B}_{\text{uncorr}}$, are used consistently.

In Eq. (50) the pointing matrices, \mathbf{A}^+ and \mathbf{A}^- , are given as in Eqs. (2) and (3) and therefore are composed of zeros and ones for the summed data and have two non-zeros per row given by $\cos(2\varphi_i^{\parallel})$ and $\sin(2\varphi_i^{\parallel})$ for the differenced data. The ground template operator, \mathbf{G}^d , is the same for the summed and differenced data. As an independent ground template is used for each detector pair and for each CES, the \mathbf{G}^d matrix has as many columns as the number of bins used to discretize the ground signal and as many rows as the number of samples in a given CES. We bin the observed azimuths in intervals of the width of 0.08 deg and thus have $\mathcal{N}_G^{CES,p} \sim 100$ bins per template. At any given time t the corresponding row of \mathbf{G}^d has only one non-zero entry (equal to one) in a column corresponding to the ground bin observed at this time.

The \mathbf{B}^{\pm} matrices define the time-domain filtering applied to both data streams in order to suppress long term correlations. In the POLARBEAR case the filtering is done subscan-by-subscan (The POLARBEAR Collaboration 2014). Consequently the \mathbf{B}^{\pm} are block diagonal with one block per subscan, and each block displaying the same structure as in Eqs. (39) and (44). We denote such an elemental block with a subscript, s , to emphasize that we are referring to a single subscan. Each of these blocks removes from a given subscan time-domain trends given by time domain templates defined as polynomials up to some order, selected to ensure that the noise after filtering is nearly white. In our analysis, $\mathbf{B}_{\text{corr},s}$ contains four templates (the Legendre polynomials up to the 3rd order, appropriately rescaled to become orthonormal over the time interval given by the subscan) and $\mathbf{B}_{\text{uncorr},s}$ contains only the constant and linear templates. Consequently, the columns of $\mathbf{B}_{\text{uncorr},s}$ are linearly dependent on those of $\mathbf{B}_{\text{corr},s}$ and we restrict \mathbf{B}^+ to the latter ones (i.e., $\mathbf{B}^+ = \mathbf{B}_{\text{corr},s}$) without any loss of generality, as discussed in Sect. 3.2.

\mathbf{w}^+ and \mathbf{w}^- are the noise terms, describing the noise in the data after the ground template and temporal trends removal. These noise terms are expected to be ‘‘prewhitened’’ with respect to the actual noise in the data. In the simulations these vectors are modeled as white with inverse variances given by ω^+ and ω^- respectively. We allow for different weights for each CES and detector pair. These weights are evaluated from the actual data as the inverse of the average of the power spectral density of the real data sum and difference, taken between 1.04 Hz and 3.13 Hz.

It is now straightforward to generalize these considerations to multiple CESs and multiple detector pairs. In both cases we stack all the TOD for every detector pair and every CES together and, for concreteness, we do so for the summed and differenced data separately. The form of Eq. (50) for the concatenated summed and differenced data remains the same but the data objects and operators on its right hand side need to be appropriately redefined. In particular, as we define a different ground template for each detector pair and each CES the global $\mathbf{G}_{\text{all}}^{\pm}$ matrix, will be block diagonal, with each block given by the detector-pair and CES specific matrix, \mathbf{G}^d . The vector of the ground template amplitudes, $\mathbf{g}_{\text{all}}^{\pm}$ will accordingly be made of the detector-pair and CES specific vectors, \mathbf{g}^{\pm} . Similar generalizations also apply to the temporal drifts term, however in this case one may need, or want, to account for effects that would be correlated between different detector pairs. Such effects could for instance be a result of contributions to the summed data due to atmospheric fluctuations. Consequently, the ultimate filtering operator, $\mathbf{B}_{\text{all}}^{\pm}$,

may not be strictly block diagonal but have rather a form resembling that of Eq. (39). In the example studied however we do not include this possibility but introduce a separate template for each detector, each subscan and each polynomial order. This adds some flexibility that may permit better accounting for systematic effects, but it may not be always advantageous as far as statistical uncertainties are concerned due to the significant number of extra independent degrees of freedom this choice implies. The number of polynomial templates per CES and detector pair is $\mathcal{N}_B^{\text{CES},p} = \mathcal{N}_{\text{sub}} \times \mathcal{N}_{\text{poly}} \sim 150 \times \mathcal{N}_{\text{poly}}$, where $\mathcal{N}_{\text{poly}}$ is four or two if we are considering the sum or the difference of the detector pair. Consequently the total number of templates per CES and detector pair sum (resp. difference) is $\mathcal{N}_T^{\text{CES},p} = \mathcal{N}_G^{\text{CES},p} + \mathcal{N}_B^{\text{CES},p} \sim 700$ (resp. 400).

5.3. Map-making: estimators and implementation

We estimate the sky signals using both the unbiased, Eq. (22), and the biased, Eq. (25), estimators. The weight matrices, \mathbf{M} , are assumed to be block diagonal, with blocks corresponding to different CESs for every detector pair. The blocks are in turn taken to be diagonal and proportional to a unit matrix with the proportionality coefficient given by the noise weights as introduced at the end of the previous section. Consequently, the weight matrix block corresponding to the i th detector pair and c th CES reads,

$$\mathbf{M}_{c,i}^{\pm} = \omega_{c,i}^{\pm} \mathbf{1}. \quad (51)$$

We define sky pixels for which the signal estimation is performed prior to map-making. This is done in two steps. First, we remove pixels for which a two-by-two diagonal block of $\mathbf{A}^T \mathbf{M} \mathbf{A}$ matrix has a condition number higher than 10^6 . This ensures that we retain only the pixels for which there is sufficient observation redundancy to allow for numerically disentangling two Stokes parameters, Q and U . In addition, we also remove pixels that have not been observed sufficiently frequently. The threshold is chosen as a minimal number of CESs during which the pixel was observed, taken to be 120 here. We have found empirically that this criterion helps avoid strongly degenerate modes localized at the lightly-observed patch boundaries. As a result our sky maps are composed of roughly 5.2×10^4 , $n_{\text{side}} = 2048$ HEALPix pixels covering approximately 43 deg^2 . Once the pixels are designated for removal we propagate this information back to the TOD, flagging out all the samples which fall in one of those pixels.

Map-making requires an application of the filtering operator, \mathbf{F}_T , to the time-ordered data vectors, \mathbf{d}^{\pm} . This has to be preceded by a computation of the filtering operator itself. All the matrix operators involved in the computation (i.e., \mathbf{B}^{\pm} , \mathbf{G} , $\mathbf{T} \equiv [\mathbf{B}^{\pm}, \mathbf{G}]$, and \mathbf{M}) are block diagonal, as is \mathbf{F}_T . We pre-compute the orthonormalization kernel $\mathbf{K} \equiv (\mathbf{T}^T \mathbf{M} \mathbf{T})^{-1}$ via its explicit construction and inversion. We first build the template coupling kernel, $\mathbf{T}^T \mathbf{M} \mathbf{T}$, which requires $\mathcal{N}_t \times \mathcal{N}_{\text{poly}}^2 \sim 10^{12} - 10^{13}$ operations thanks to the sparsity and structure of the time-domain templates. The cost of the inversion (with regularization) of the kernel is $(\mathcal{N}_T^{\text{CES},p})^3 \times \mathcal{N}_{\text{CES}} \times \mathcal{N}_{\text{pair}} \sim 10^{13} - 10^{14}$ operations. This number is considerable, but performing a large number of small matrix inversions is very efficient because of the locality of the data to be processed by the CPUs, resulting in a considerable advantage on modern massively parallel computing systems. The kernel's rank is approximately equal to the total number of templates, $\mathcal{N}_T^{\text{CES},p} \times \mathcal{N}_{\text{CES}} \times \mathcal{N}_{\text{pair}} \sim 10^8 - 10^9$. Storing this object in the memory is demanding, even when its block diagonal structure is explicitly taken into account, since it

amounts to $(\mathcal{N}_T^{\text{CES},p})^2 \times \mathcal{N}_{\text{CES}} \times \mathcal{N}_{\text{pair}} \sim 10^{11} - 10^{12}$ double precision numbers.

We note that by construction \mathbf{B}^{\pm} and \mathbf{G} are individually column-orthogonal, but when considered together their columns are in general not orthogonal and possibly not even linearly independent. Indeed, one degeneracy of each CES and detector-pair block is readily expected. This is the one between the constant mode filtered by the time-domain templates and the constant offset of the ground template. Consequently, the kernel, \mathbf{K} , is in general non-trivial and its inversion needs to be regularized, see Sect. 3.2. In doing so, we set a threshold of 10^6 on the condition number of each diagonal block of $\mathbf{T}^T \mathbf{M} \mathbf{T}$. With the kernel pre-computed and stored in the computer memory, we apply the filter to the data without ever computing it explicitly. Instead, we perform the operations included in its implicit form, Eq. (10), from the right to left, computing first $\mathbf{M} \mathbf{d}$, followed by $\mathbf{T}^T (\mathbf{M} \mathbf{d})$, and then $\mathbf{K} (\mathbf{T}^T \mathbf{M} \mathbf{d})$. We then loop over the TOD again to compute $\mathbf{A}^T \mathbf{F}_T \mathbf{d}$. Both these operations have a $O(N_t)$ cost. This approach facilitates the entire operation: storing the filtering matrix, \mathbf{F}_T , would require a prohibitive amount of memory and its application to a time-domain vector would take too much computational time.

5.3.1. Unbiased map estimator

We have implemented the unbiased map estimator Eq. (22) in two different ways. In the first case, we perform an explicit construction and inversion of the $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ matrix, while in the second we employ an iterative solver instead.

In the explicit implementation, we first compute $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$. This is done as follows. Consider a time sample t and call p the observed sky pixel. Since the only non-zero entries of the t th row of \mathbf{A} are the columns corresponding to pixel p , the t th column of \mathbf{F}_T contributes only to the p th column of $\mathbf{F}_T \mathbf{A}$. Analogously the t' th row of \mathbf{F}_T contributes only to the p' th row of $\mathbf{A}^T \mathbf{F}_T$. Therefore, in order to build the $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ matrix we loop over the elements of \mathbf{F}_T : for the (t', t) entry we compute its contribution to the (p', p) block of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$. As mentioned earlier the \mathbf{F}_T matrix is not stored in the memory, but rather its elements are computed on the fly from the matrices \mathbf{K} and \mathbf{T} (stored in the memory in a compressed form). The non-zero entries of \mathbf{F}_T are the $\mathcal{N}_{\text{CES}} \times \mathcal{N}_{\text{pair}}$ blocks ($\mathcal{N}_t^{\text{CES}}$ by $\mathcal{N}_t^{\text{CES}}$) on the diagonal. These blocks are scattered across a number of processors that ranges between about 1300 and 5800, depending on the computational platform we used. Each processor is responsible for computing the contribution of its blocks to $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and the result is reduced at the end. Since the matrix, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$, can not be stored in the memory of a processor (and not even on an entire compute node) we divide it into blocks (typically 576) and we compute them one by one, at each step only considering entries (t', t) of \mathbf{F}_T such that (p', p) is inside the block of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ being computed. The computational cost scales as the number of non zero entries of \mathbf{F}_T : $\mathcal{N}_{\text{CES}} \times \mathcal{N}_{\text{pair}} \times (\mathcal{N}_t^{\text{CES}})^2 \sim 10^{15}$ operations.

We then perform the eigendecomposition of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ representing it as,

$$\mathbf{A}^T \mathbf{F}_T \mathbf{A} = \mathbf{V} \text{diag}(\mathbf{e}) \mathbf{V}^T. \quad (52)$$

This is done with help of a ScaLAPACK routine, `pdsyevr`, Blackford et al. (1997). The numerical cost is $O(N_p^3) \sim 10^{15}$ operations. This scaling relation is the main obstacle in the application of the explicit implementation to maps with a larger number of pixels. By construction the eigenvalues, \mathbf{e} , are all non-negative numbers though numerically some small eigenvalues may turn out to be negative. The inversion of this matrix is then performed

by inverting its eigenvalues. Since the condition number of the matrix is typically very large, the inversion needs to be regularized by employing a (pseudo)inverse defined as,

$$(\mathbf{A}^\top \mathbf{F}_T \mathbf{A})^{-1} = \mathbf{V} \text{diag}(\tilde{\mathbf{e}}) \mathbf{V}^\top, \quad (53)$$

where,

$$\tilde{\mathbf{e}}_i \equiv \begin{cases} e_i^{-1}, & \text{if } e_i > 10^{-6} \max_j e_j; \\ 0, & \text{otherwise.} \end{cases} \quad (54)$$

One of the advantages of this estimator is that once the computationally heavy objects are evaluated, we can efficiently produce multiple simulated realizations of the reconstructed sky maps directly in the map domain, see Sect. 5.4 for more details. We emphasise that the construction and inversion of the system matrix are the most complex and expensive parts of the algorithm: their memory requirements force an intrinsically parallel implementation and a non-negligible fraction of the execution time is spent in communication between the compute nodes.

In the case of the iterative solver, the system matrix, $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$, is never explicitly constructed. Instead, we follow the general blue-print of maximum likelihood map-making code implementations (e.g. Cantalupo et al. 2010) and apply the factors defining the matrix from right to left. The filtering operator is applied as described above, again without being ever explicitly constructed. We use a preconditioned conjugate gradient (PCG) technique (e.g., Golub & van Loan 1996) with the preconditioner set to be $(\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1}$, which is either diagonal, for the total intensity, or block-diagonal, for the polarization, and can therefore be straightforwardly computed, stored in the memory, and applied to a vector whenever needed. This is again the standard choice (e.g. Cantalupo et al. 2010). To quantify the convergence one typically uses the norm of the map level residuals,

$$r^{(i)} \equiv \frac{|\mathbf{A}^\top \mathbf{F}_T \mathbf{A} \mathbf{m}^{(i)} - \mathbf{A}^\top \mathbf{F}_T \mathbf{d}|}{|\mathbf{A}^\top \mathbf{F}_T \mathbf{d}|}, \quad (55)$$

where $\mathbf{m}^{(i)}$ stands for the solution after the i th iteration. We notice that r is a dimensionless quantity. A typical requirement for convergence would then be $r^{(i)} \leq 10^{-6}$.

Unlike the explicit implementation, the iterative solver is not applied to the full dataset. The constant elevation scans are divided into 267 nearly even groups and, using the iterative solver, one map $\hat{\mathbf{s}}_\alpha$ is obtained independently for each group α . Afterwards we coadd the maps as follows

$$\hat{\mathbf{s}} = \left[\sum_\alpha (\mathbf{A}^\top \mathbf{M} \mathbf{A})|_\alpha \right]^{-1} \sum_\alpha (\mathbf{A}^\top \mathbf{M} \mathbf{A})|_\alpha \hat{\mathbf{s}}|_\alpha, \quad (56)$$

where as before $|_\alpha$ denotes a quantity computed using only the data belonging to subset α .

5.3.2. Biased map estimator

The biased map estimator Eq. (25) requires a construction of $\mathbf{A}^\top \mathbf{M} \mathbf{A}$, its inversion and its multiplication by a vector, $\mathbf{A}^\top \mathbf{F}_T \mathbf{d}$. All these operations pose no issues given the block-diagonal structure of the matrix and the pixel selection procedure applied to the data beforehand, which ensures that each of its blocks is invertible. The computational cost is then driven by the construction of the kernel \mathbf{K} ($\sim 10^{13} - 10^{14}$ operations). As mentioned earlier, this kernel is block-diagonal and can be constructed and inverted efficiently on a single modern processing unit. Consequently, the entire estimator can be implemented and executed using serial or embarrassingly parallel programming models.

5.4. Simulations

In our analysis we use signal-only, noise-only, and signal and noise simulated map reconstructions. As the time-domain data set is quite large we strive to perform the simulations in the pixel domain whenever possible. We simulate signal- and noise-only maps separately, as described below, and the total maps are then produced by co-addition of these at the map level.

We also produce specialized simulations in order to validate our implementations. These are described in Sect. 5.5.

5.4.1. Signal-only maps

For the simulations of the CMB sky signal we assume the power spectrum defined by the *Planck* best fit parameters Planck Collaboration XIII (2016). We set the tensor-to-scalar ratio, r to zero, for definiteness, as the value of r is not relevant in the case considered here, given the focus of the first POLAR-BEAR campaigns on sub-degree angular scales. We synthesize the ‘‘true sky’’ maps, denoted here as \mathbf{s}^{sim} , using the *synfast* tool of the HEALPix package (Górski et al. 2005).

To simulate the CMB-only sky maps as reconstructed with the explicit implementation of the unbiased map-maker we take the simulated true sky maps, \mathbf{s}^{sim} and remove from the simulated realization of the sky maps the eigenvectors corresponding to the eigenvalues set to zero in the $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$ inversion regularization. We have validated that this is equivalent, algebraically and numerically, to first projecting the sky signal, \mathbf{s}^{sim} , to the time-domain, $\mathbf{A} \mathbf{s}^{\text{sim}}$, and then running the map-making procedure on the derived time-ordered data. For the biased map estimator we apply the operator, $(\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}_T \mathbf{A}$, directly to the generated, true sky maps, \mathbf{s}^{sim} . Again this is algebraically and numerically equivalent to first producing the signal-only data streams, $\mathbf{A} \mathbf{s}^{\text{sim}}$, and then applying the biased map-making operator, $(\mathbf{A}^\top \mathbf{M} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{F}$, to them.

Only for the iterative implementation of the unbiased map-making do we actually run the map-making solver on simulated time-streams, in order to reproduce effects related to the convergence (or otherwise) of the iterative algorithm.

5.4.2. Noise-only maps

The noise-only maps reconstructed with the explicit implementation of the unbiased map estimator are computed directly in the pixel-domain as $\hat{\mathbf{n}} = \mathbf{V} \mathbf{z}$, where \mathbf{z} is a pixel-domain vector of Gaussian random variables with variance given by the corresponding entry of $\tilde{\mathbf{e}}$, and \mathbf{V} and $\tilde{\mathbf{e}}$ are defined in Eq. (53). This assumes that the weight matrix, \mathbf{M} , provides a correct description of the TOD noise (i.e., $\mathbf{M} = \mathbf{C}_n^{-1}$). Were this is not the case, we would have to start from simulating a noise timestream, \mathbf{n} , with the desired noise properties and then process it with the map-making algorithm.

For the biased map-making we follow the latter path even if the noise is uncorrelated. We therefore start by generating a timestream of uncorrelated Gaussian numbers of appropriate variance and projecting it into the pixel-domain using the corresponding map-making procedure.

5.5. Validation and verification of the map-making code

We have performed multiple tests in order to validate and verify the map-making code. For validation, we test whether the filtering operator, \mathbf{F}_T , removes all the unwanted modes as desired; for verification we perform a number of full, end-to-end runs of the




Fig. 2. Histogram of the whitened unbiased map $\mathbf{V}\text{diag}(\sqrt{\hat{\epsilon}})\mathbf{V}^T\hat{\delta}$, where $\hat{\delta}$ is estimated starting from a time-domain white noise simulation (see Sect. 5.5 for the details).

code, testing that the various outputs have the expected statistical properties.

Our validation test checks that the filtering operator \mathbf{F}_T satisfies the relation $\mathbf{F}_T\mathbf{T}\mathbf{y} = 0$, for some vector of template amplitudes \mathbf{y} . Since we are interested in map domain residuals, we actually test whether

$$s^{\text{res}} = (\mathbf{A}^T\mathbf{M}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{F}_T\mathbf{T}\mathbf{y} = 0. \quad (57)$$

We separately test the ground pick up filtering, $\mathbf{T} = \mathbf{G}$, and the polynomial filters, $\mathbf{T} = \mathbf{P}$. In the former case, we produce a simulated ground signal timestream as follows. For every CES, there is an index i for each ground template bin. i ranges between 0 and ~ 100 . For each CES we set the amplitude of the i th ground template to $y_i^G = 1 + 0.01 \times i$, while the entries of \mathbf{y} corresponding to the polynomial templates are set to zero. We obtain a simulated data vector $\mathbf{d}^{\text{ground}} = \mathbf{T}\mathbf{y}$, where $d_i^{\text{ground}} = y_i^G$. This construction has been devised in order to ensure that the simulated timestream is a linear combination of all the ground templates, with elements that are of order unity and are always positive. This last condition mimics the worst-case scenario of a ground signal that is coherent in time. In a very similar fashion, we test the filtering operator on the polynomial filters by setting $d_i^{\text{poly}} = 1 + 0.01 \times i$, where now i is the index of the swipe in azimuth within the CES (remember that we have a set of polynomial templates for each constant direction azimuthal glide). We find that the map domain residuals, s^{res} , never exceed 10^{-6} for temperature and 10^{-8} for polarization. These levels are expected given the precision of our filter orthogonalization procedure, which tends to be merely approximate for very short subscans. They are however negligible for any practical purposes.

As part of our end-to-end verification tests we study the statistical properties of the noise-only maps produced by our map-making code. In this case, we produce the simulated noise-only stream in time-domain, with properties described by the diagonal weight matrix, \mathbf{M} , and processed it via our code. The output map was then prewhitened using the square root of the theoretically expected covariance, $(\mathbf{A}^T\mathbf{F}_T\mathbf{A})^{-1}$. The result was then histogrammed, Fig. 2, and compared to a Gaussian with unit variance. The agreement was found to be very good, for example the Kolmogorov-Smirnov test found p -values of 0.66 for polarization and 0.27 for temperature, showing that the $(\mathbf{A}^T\mathbf{F}_T\mathbf{A})^{-1}$ matrix (explicitly computed) reproduces correctly the covariance properties of the unbiased map, including the correlations due to the time-domain filtering.

Other examples of end-to-end tests, involving a direct comparison of the known input with the output are discussed in Sect. 6.1. As emphasized there, the overall agreement is found to be excellent.

5.6. Polarized power spectrum

The sky maps reconstructed from the measurements of a CMB experiment typically serve multiple purposes. They may be the end product of the analysis whose goal is a representation of the sky signal in the observed sky area. However, they will often be only a step toward some more profound statistical analysis of the underlying signal.

In the following we will therefore not only look at the reconstructed maps as images of the true sky, but will also compare them from the point of view of the constraints on the power spectra which can be derived from their analysis. In this latter case, we focus specifically on the B -mode power spectrum and use the pseudo-power spectrum approach to its estimation (e.g. Peebles & Hauser 1974; Hivon et al. 2002). This method has gained significant popularity in the field, thanks to its flexibility and relatively straightforward implementation. As our goal is the spectra of the B -mode polarization, and the observations we consider cover only a limited sky area, we use a so-called pure pseudo spectrum approach, which explicitly corrects for the bias and variance effects of the so-called E -to- B leakage generated by the presence of the observed sky boundary. The technique was first proposed in Bunn et al. (2003), and later implemented and elaborated on in Smith (2006) and Smith & Zaldarriaga (2007). In this work, we use a numerical code, X²PURE, developed by Grain et al. (2009), which has been described, tested, validated and exploited both there and in follow-up work (e.g., Grain et al. 2012; Ferté et al. 2013, 2015; Planck Collaboration Int. XXX 2016; Krachmalnicoff et al. 2016). We refer the reader to these papers for more details.

The pure pseudo-spectrum technique removes the bias due to leakage by estimating the E and B spectra simultaneously and allowing for an off-diagonal EB block of the coupling kernel, which is then used to model and subtract the leaked E power from the B -mode spectrum. The enhancement of the power spectrum variance due to the leakage is then suppressed with the help of appropriately constructed apodizations. Ideally these have to be estimated separately for every harmonic domain bin for which the spectrum is to be computed (Smith & Zaldarriaga 2007); again we use the code implemented by Grain et al. (2009). Once the apodizations are estimated we use them for all the B -mode spectra we estimate, irrespective of the algorithm used to produce the maps.

The pure techniques are only designed to deal with E -to- B leakage due to the cut sky. However, other sources of leakage are also usually present. For instance, at small angular scales leakage typically arises as a result of the pixelization adopted for the recovered map. This is typically found to be subdominant to the uncertainty due to the noise on small scales, and thus can typically be left uncorrected with little, if any, impact on the precision of the final results.

Leakage can also be expected if the Q and U maps used for power-spectrum estimation do not faithfully reflect the true underlying sky signal. This is certainly the case for biased map-making, but can also be relevant for the unbiased approach if degeneracies are present. The leakage arising in such cases can bias the estimated spectra on all angular scales of interest, and thus must be carefully accounted for. Though solutions to this problem have been proposed (e.g., Bunn et al. 2003; BICEP2 Collaboration 2014; BICEP2 and Keck Array Collaborations 2016), they are computationally very heavy. Here, instead, we use a phenomenological approach based on

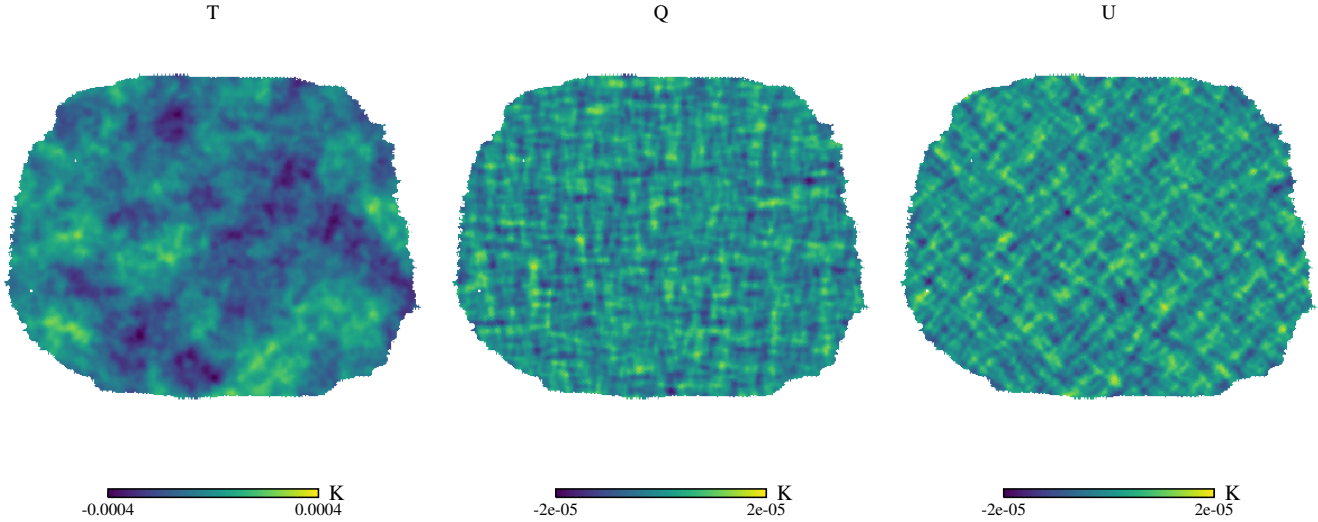


Fig. 3. Maps of the input sky used here for the reconstruction comparison of the different map estimators.

Hivon et al. (2002) and model the biased spectra as

$$\begin{pmatrix} \tilde{C}_\ell^{EE} \\ \tilde{C}_\ell^{BB} \end{pmatrix} = \begin{bmatrix} f_\ell^{EE} & f_\ell^{EB} \\ f_\ell^{BE} & f_\ell^{BB} \end{bmatrix} \begin{pmatrix} C_\ell^{EE} \\ C_\ell^{BB} \end{pmatrix} + \begin{pmatrix} N_\ell^{EE} \\ N_\ell^{BB} \end{pmatrix}. \quad (58)$$

where \tilde{C}_ℓ^{XX} represents the biased power spectrum computed from the map, while C_ℓ^{XX} is the true sky power spectrum to be estimated.

The \tilde{C}_ℓ spectra are evaluated using X²PURE, the bins in ℓ are centred at $\ell = 400, 700, 1100, 1500, 1900$ and have a width of $\Delta\ell = 400$, except for the first bin ($\Delta\ell = 200$). The transfer functions, f_ℓ^{XY} , are evaluated using CMB signal only simulations assuming the power spectrum defined by the *Planck* best fit parameters Planck Collaboration XIII (2016) and zero tensor-to-scalar ratio, $r = 0$.

In order to evaluate f_ℓ^{XY} a Y -only realization of the sky is performed using the *synfast* tool of the HEALPix package (Górski et al. 2005). Knowing the underlying algebra of the map estimator, the biased map is extracted from the simulated sky (analogously one could project the sky into the time-domain using \mathbf{A} and run the map-maker on these timestreams, see Sect. 5.4.1 for more details). A biased spectrum \tilde{C}_ℓ is computed then using X²PURE. This procedure is repeated for 100 independent sky realizations and the transfer function is evaluated as

$$f_\ell^{XY} = \frac{\langle \tilde{C}_\ell^{XY} \rangle}{C_\ell^{YY}}, \quad (59)$$

where $\langle \dots \rangle$ denotes the average over the 100 simulations. The variance on this determination of f_ℓ^{XY} is $\text{Var}(\tilde{C}_\ell^{XY}) / (100 \langle C_\ell^{XX} \rangle^2)$, in our case it corresponds to an uncertainty on f_ℓ^{XY} of the order of percent, which is sufficiently small to ignore possible biases that would arise from an inaccurate estimation of the transfer functions. Analogously the noise bias, N_ℓ , is estimated as the average of spectra produced by X²PURE on 100 noise only maps. These maps were produced running each map-maker on timestreams drawn from a Gaussian distribution with covariance \mathbf{C}_w (or through an equivalent procedure, see Sect. 5.4.2).

We now have all the ingredients needed for calculating the unbiased power spectrum estimator, which is given by

$$\begin{pmatrix} \hat{C}_\ell^{EE} \\ \hat{C}_\ell^{BB} \end{pmatrix} = \begin{bmatrix} f_\ell^{EE} & f_\ell^{EB} \\ f_\ell^{BE} & f_\ell^{BB} \end{bmatrix}^{-1} \begin{pmatrix} \tilde{C}_\ell^{EE} - N_\ell^{EE} \\ \tilde{C}_\ell^{BB} - N_\ell^{BB} \end{pmatrix}. \quad (60)$$

This extension to polarization of Hivon et al. (2002) is equivalent to the approach adopted by The POLARBEAR Collaboration (2014) if $f_\ell^{BE}, f_\ell^{EB} \ll 1$, a condition which is also fulfilled in our case since we are considering a very similar observation.

6. Results

Here we describe the maps produced from POLARBEAR-like simulations using the different map-making approaches and their implementations. We assess the maps from the point of view of the fidelity with which they reconstruct the actual sky signal. However, we also look at them as merely a step toward a statistical characterization of the signal. In this latter case, we use the B -mode polarization power spectrum as a comparison metric and specifically focus on the power spectrum estimation approach as introduced in the last section.

6.1. Reconstructed sky maps

In Figs. 3–6 we compare the maps reconstructed using both the biased and unbiased estimators, in the latter case using both the explicit and the iterative solver. The maps were computed assuming noiseless data and the same pixel size was adopted for the simulation and reconstruction to avoid any pixel effects. In none of the cases considered does the reconstructed sky correspond exactly to the input. We discuss each of the cases in turn below.

Unbiased maps via the explicit solver. The residual present in this case, Fig. 4, is due to the presence of singularities in the system matrix, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$. As expected from our earlier discussion, Sect. 3, and confirmed by our results in Sect. 6.2, there are two such singular modes for polarization and one for temperature. The inversion regularization procedure removes these two modes from the estimated map, even if they contain actual sky signal. These singular modes result from the interplay between the scanning strategy and the filtering. They make the filtered data insensitive to these modes, which are unavoidably lost. Therefore, the map estimated from the unbiased estimator using the explicit solver may not be strictly speaking unbiased but provides an unbiased representation of all the modes that can be constrained from the available data given the chosen filtering.




Fig. 4. Maps derived with the explicit implementation of the unbiased map estimator. *Top row:* reconstructed maps. *Middle row:* difference between the reconstructed and the input maps. *Bottom row:* difference between the reconstructed and the input maps with the singular modes removed from the input maps.

Unbiased maps via the iterative solver. The residual in this case is clearly more pronounced and complex (see Fig. 5, middle panel). As we mentioned earlier we use a PCG solver and

adopt as the preconditioner the matrix $(\mathbf{A}^T \mathbf{M} \mathbf{A})^{-1}$. One might expect that the result of the unbiased map estimator should be the same, whichever solver is applied. However, the result shown

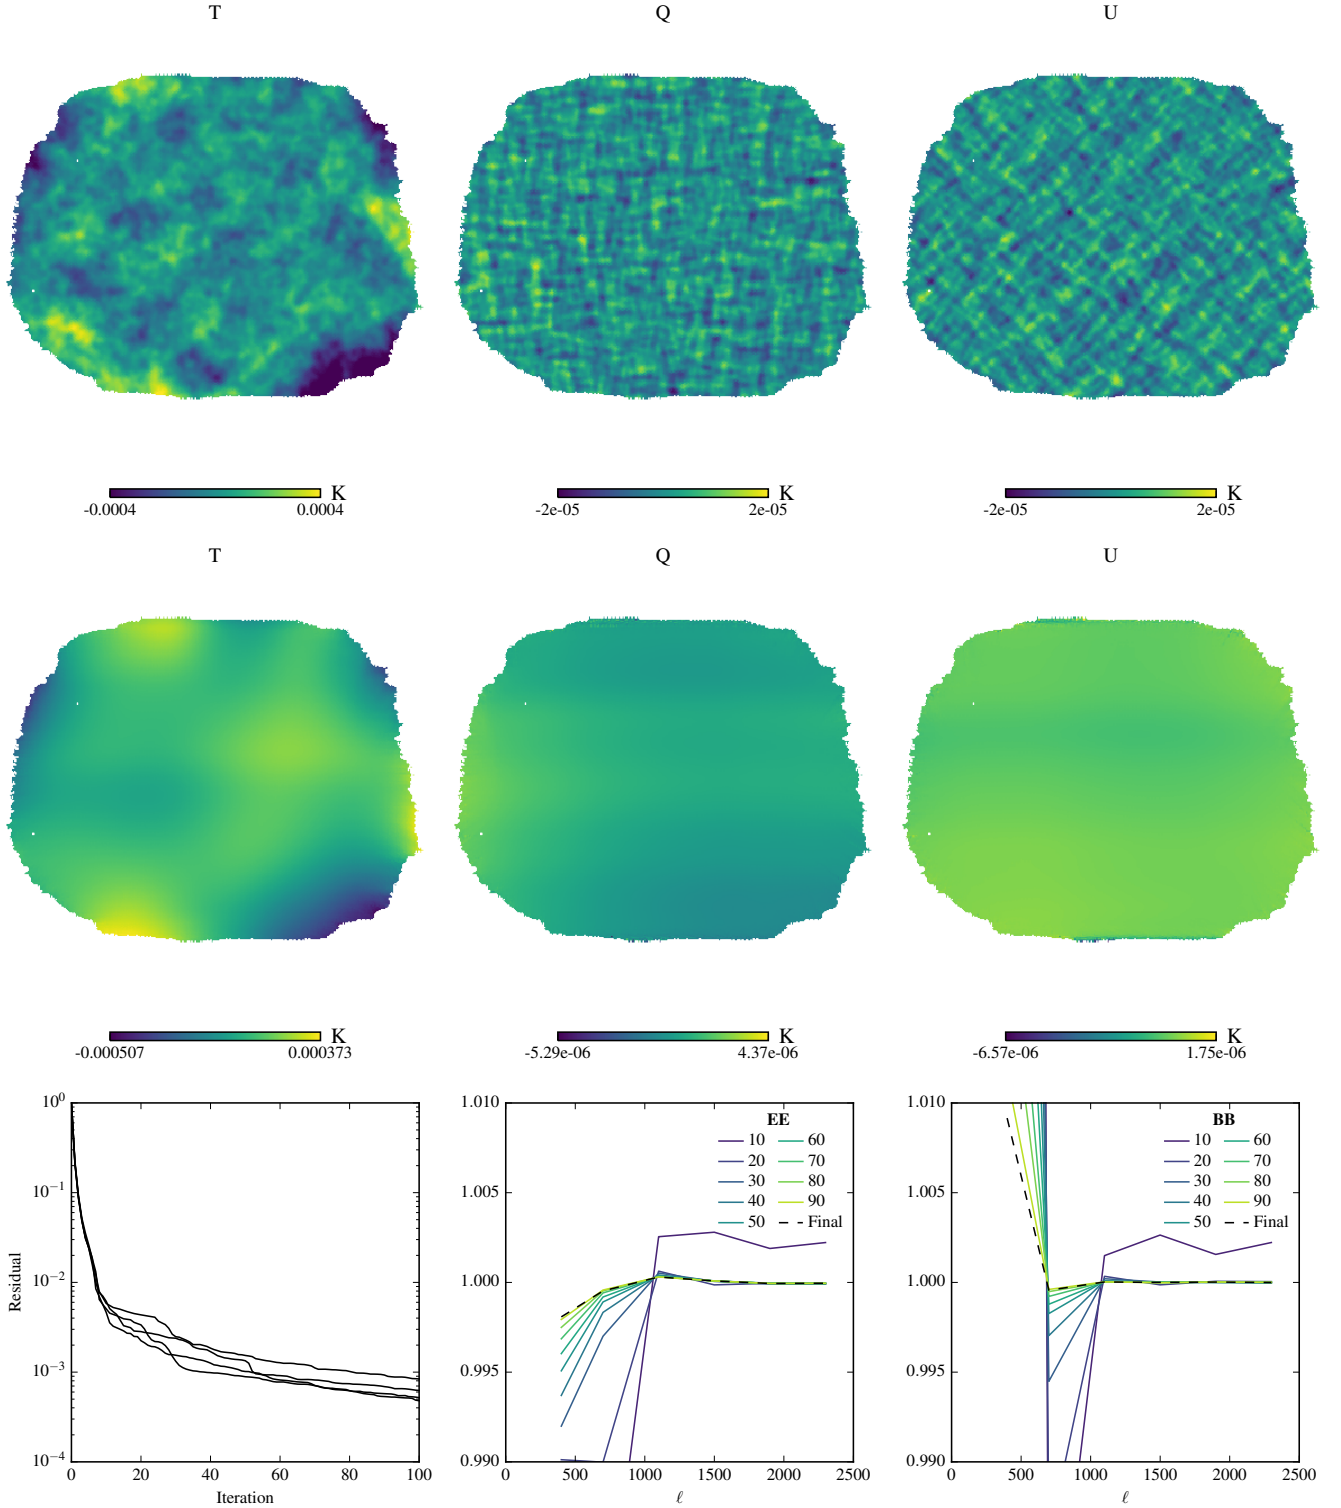


Fig. 5. Results from the iterative implementation of the unbiased map estimator. *Top row:* reconstructed maps. *Middle row:* difference between the reconstructed and the input maps. *Bottom row:* a study of convergence of our iterative solver. The *left panel* shows the standard residual, as defined in Eq. (55), which saturates and does not converge to our fiducial level of 10^{-6} in as many as 100 iterations. The *middle and right panels* show that this lack of convergence is due to the largest angular modes as the fractional difference between the power spectrum of the input map and the power spectrum of the i th map estimate in the multipole range $\ell \in [500, 2100]$ becomes quickly very small and reaches the level of better than 0.1% in fewer than ~ 100 iterations. This last observation has been used to set the convergence criterion used in the analysis of the first year POLARBEAR data set [The POLARBEAR Collaboration \(2014\)](#).

in the figure corresponds to an incompletely converged iterative solution. Indeed, we have found that the iterative solution residuals, Eq. (55), do not decay to zero, see the bottom left panel

of Fig. 5, but instead asymptote to about 10^{-3} , which is roughly three orders of magnitude above our fiducial convergence criterion of 10^{-6} . This is the case even if we allow as many as a few




Fig. 6. Map estimates derived using the biased map estimator, Eq. (25). *Upper row:* reconstructed maps. *Bottom row:* difference between the reconstructed and the input maps.

thousand iterations. Such a behavior is indeed expected in linear systems for which the system matrix is (numerically) nearly singular (Hanke 1995; Szydlarski et al. 2014).

To understand the effects of this lack of convergence of the solver on the estimated maps, we have computed the (pseudo) power spectra of the estimated map after i th iteration, bottom right panels of Fig. 5. We see that although the very low ℓ part of the spectrum does indeed fail to converge, convergence is quickly reached in the intermediate and high ℓ -range. This again is consistent with the singular modes found in the explicit solver having only large angular scales. If the singular modes are known, we could readily remove them from the solution, and thus from the residuals, at each step of the iteration and restore the proper convergence. However, this typically would require as many computations as the explicit solver, undermining the most important advantage of the iterative one.

We can still use the PCG solver in such circumstances by using this practical workaround: instead of monitoring a single residual as given by Eq. (55), we track the behavior of residuals at the scales of interest for the power spectrum, as shown in Fig. 5. Nevertheless, we have to be aware of the fact that some of the power in the final solution may be compromised.

In the case of the map shown in the figure, convergence can indeed be reached in fewer than 100 iterations in the so-called

science band defined in The POLARBEAR Collaboration (2014), which was the band of interest for the first round of the POLARBEAR papers.

Biased maps. The biased map estimator, as expected, leads to the largest residuals. These are particularly pronounced in the outskirts of the map, where the pixels crossing (and thus the filtering) may be highly anisotropic, but are still readily visible in the central part of the map, where the cross-linking and pixel sampling are better.

6.2. Eigenstructure of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$

The eigenstructure of $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ not only determines which modes are missing from the final unbiased sky estimate, it also provides information about the modes, which, while not singular, are not well constrained by the data. This is because the matrix is closely connected to the noise covariance in the pixel domain, N_p , as shown by Eqs. (17) and (18).

In the case with no filtering at all the matrix, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$, reduces to the well-known $\mathbf{A}^T \mathbf{M} \mathbf{A}$, which, for the diagonal weights assumed here, is block diagonal with two-by-two blocks describing the (weighted) coupling between Q and U Stokes parameters

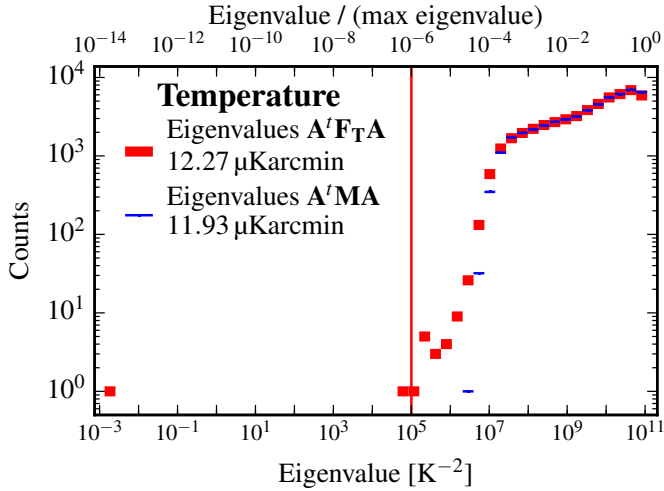


Fig. 7. Eigenvalues of the temperature block of the $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and $\mathbf{A}^T \mathbf{M} \mathbf{A}$ matrices. The spectra of the two matrices are very similar. The major difference is a group of poorly constrained modes, including one that is formally degenerate corresponding to the offset of the map.

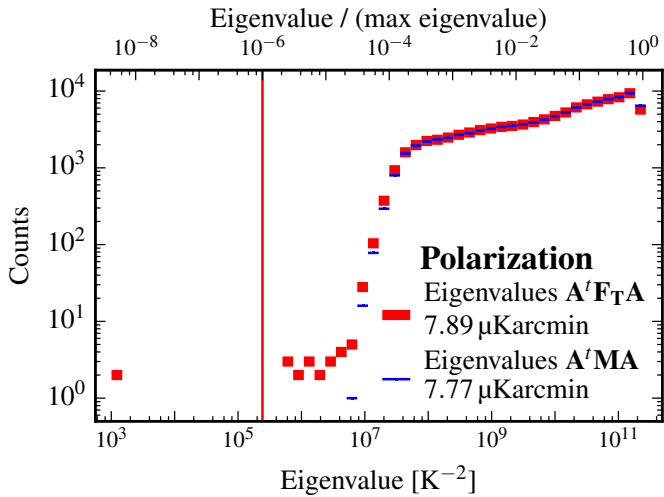


Fig. 8. Eigenvalues of the polarization block of the $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and $\mathbf{A}^T \mathbf{M} \mathbf{A}$ matrices. As with temperature, the spectra of the two matrices are very similar. In this case we also have tens of poorly constrained modes. We also expect to have two nearly degenerate modes, one for each Stokes parameter (cf. Fig. 7). Though they are treated as singular because of numerical reasons, their degeneracy is partially broken, as explained in Sect. 4.2.

in each pixel. The off-diagonal elements of these blocks will typically be negligible for pixels observed with a sufficiently homogeneous distribution of polarization angles, while the diagonal elements will be approximately equal to the eigenvalues of the two-by-two block. These are essentially given by the number of observations per pixel, and their corresponding eigenvectors are spatial modes equal to zero everywhere but in the given pixel. Departure from such behavior would then indicate the presence of strong off-diagonal coupling in some of the pixels.

The spectrum of $\mathbf{A}^T \mathbf{M} \mathbf{A}$ can therefore be used as a good reference for assessing the impact of the filtering on the map domain noise spectrum. We compare the eigenstructure of the matrices, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and $\mathbf{A}^T \mathbf{M} \mathbf{A}$, in Figs. 7 and 8, for temperature and polarization respectively. In both cases, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ and $\mathbf{A}^T \mathbf{M} \mathbf{A}$ have very similar spectra with the exception of tens of poorly constrained modes, defined as those with eigenvalues smaller by at least five orders of magnitude than the maximum eigenvalue.

The corresponding eigenvectors, Fig. 10, are long modes and, in many cases, they exhibit a striped structure close to the boundaries. We interpret these modes as the result of the ground signal filtering discussed in Sect. 4.2 as their number roughly corresponds to the number of bins in the ground-bin. Indeed, following the discussion of Sect. 4.2, we expect that with every ground-template bin there should be an associated ill-constrained mode, corresponding to an offset of the constant declination strip swept by the azimuth range of the bin during the time of a single constant elevation scan. The fact that the recovered eigenvalues are not numerically zero demonstrates that these degeneracies are weakly broken as expected given that the presence of sky pixels observed with the telescope orientation corresponding to two different ground-template bins. Consequently, this leads to only one truly degenerate mode per each Stokes parameter map, see Fig. 9.

We also see that the two most singular eigenvalues of the polarized case are significantly larger than the most singular eigenvalue in the case of temperature. The difference is at least in part due to the numerical precision of the computations, however it is also consistent with our earlier expectation that the polarizer angle change across the constant elevation sweep can break the degeneracy between the sky signal offset for each strip and the amplitude of the ground template in the corresponding bin. In the case under study, given the limited primary mirror chop, the effect is very weak.

While regularizing the inversion of the matrix, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$, we remove these modes from the solution together with all the modes which are smaller than 10^{-6} of the maximal eigenvalue. This still leaves significant number of ill-constrained modes in the estimated maps. Although they do not give rise to any discernible artefacts in the case of the noiseless results as shown Fig. 4, when noise is included in the time-domain data, these modes may dominate the maps visual appearance. However, despite being noisy these modes are correctly estimated, and are neither artefacts of the estimator or its implementation, nor remnants of any incompletely-filtered parasitic signal such as atmosphere. Rather they reflect the actual uncertainty that the observation and filtering incur. These modes are typically missing in the unbiased maps derived with the iterative solver as well as in the biased maps. This is because these modes are either the most difficult to converge (in the case of the iterative solver) or are explicitly filtered out (in the biased map-maker). Consequently, although these maps may occasionally – and somewhat deceptively – look better, they will nonetheless be missing information which is correctly included in the unbiased map computed using the explicit solver.

The main sequence of eigenvectors are “pixel-like” modes, Fig. 11, in the sense that in each of these modes the most relevant structure involves a very limited number of pixels. As one intuitively expects, these pixels move from the boundary toward the center of the patch as the eigenvalue of the mode grows (i.e., as the mode is better constrained). If there were no filtering, $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ would be block diagonal (it would be equal to $\mathbf{A}^T \mathbf{M} \mathbf{A}$) and therefore each eigenvector would correspond to exactly one pixel. Setting a threshold on the magnitude of the eigenvalues would be then equivalent to performing a selection of the best observed (here innermost) pixels. In the presence of filtering, the main effect of setting a threshold is still selecting the innermost pixels. However, because of the correlations that the filtering introduces, signal is also removed from all over the map, affecting areas relatively far from the boundary pixels removed. This effect is visible at map-level in the lower panel of Fig. 11 and it is investigated at the power-spectrum level in Sect. 5.6: removing

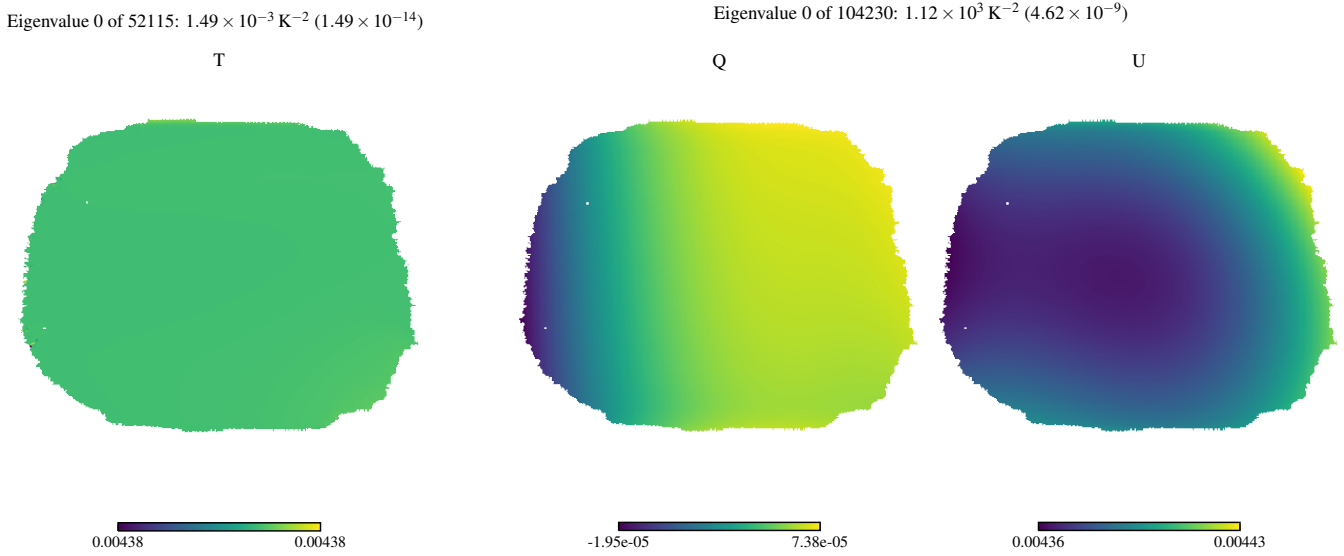


Fig. 9. Degenerate modes (the parenthesis in the title reports the corresponding eigenvalue divided by the largest eigenvalue). The degenerate mode of the temperature (*left*) map reconstruction is the global offset. The Q and U maps on the right are one of the two singular modes of polarization, the other mode is very similar: Q and U maps are swapped and the sign of one of them is flipped. The degenerate modes of the polarization are more complex because of the modulation of the polarization angle during the observation (see Sect. 4.2). For the same reason the degeneracy of these modes is partially broken (though they have to be treated as singular in order to preserve the numerical stability of the inversion of the $\mathbf{A}^T \mathbf{F}_T \mathbf{A}$ matrix).




Fig. 10. Examples of nearly degenerate modes (the parenthesis in the title reports the corresponding eigenvalue divided by the largest eigenvalue). These modes are composed of some prominent feature at the boundaries and a (usually) weaker long mode. The structures at the boundaries correspond to sets of pixels that are heavily affected by filtering. For polarization the dominant effect is the ground removal. At the high and low declination ends of the observed area the redundancy of the observations is low and therefore the degeneracy breaking effects discussed in Sect. 4.2 are mild. For temperature, the high order of the polynomial filtering plays a significant role, adding prominent features at the boundaries at intermediate declinations and increasing the complexity of the long modes. We emphasize that the prominent features at the boundaries saturate the color scale.

modes that mainly involve pixels outside of the power-spectrum mask has an important impact on the power spectrum uncertainty and bias.

6.3. Power-spectrum analysis

We now investigate how the map domain properties for the different estimators affect our pure pseudo-power spectrum estimator.

In Sect. 6.3.1, we consider the spectrum of noise-only simulations as computed by X²PURE. The qualitative results support

filtering the noisiest modes from maps produced with the unbiased map estimator.

Finally, in Sect. 6.3.2, we compare the unbiased power spectra of the maps produced by the biased and unbiased map estimators (before and after map domain filtering).

6.3.1. The noise bias

The maps generated by the different map-makers have different noise properties, in terms of both the noise amplitude and its correlations.

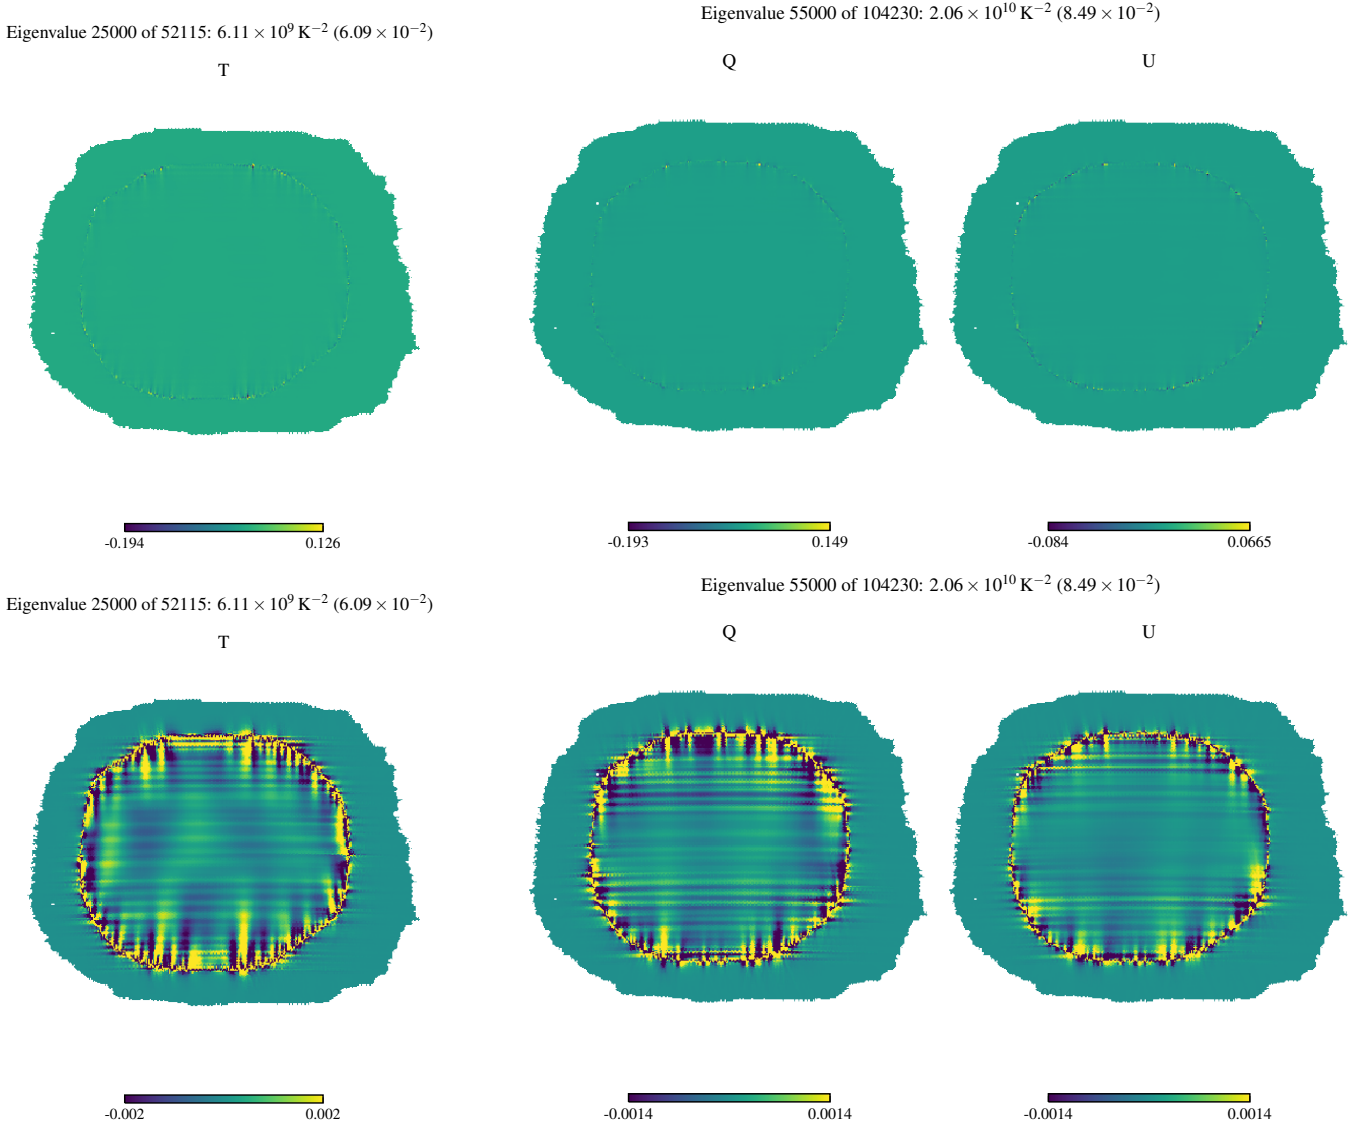


Fig. 11. Examples of pixel-like modes (the parenthesis in the title reports the corresponding eigenvalue divided by the largest eigenvalue). These modes involve mainly a “ring” composed of a very limited number of pixels (see the *first row*). In the *second row*, the same eigenvectors are displayed with a color scale squeezed by roughly two orders of magnitude, emphasizing structures inside and outside of the ring. Outside of the ring the structures quickly fade away. Also moving inwards the structures decrease their amplitude. However, compared to the outward structure, their typical length is much larger and, most important, they do not completely disappear: they have a relevant amplitude in the whole inner region. We emphasize that the “horizontal” structures represent the correlations induced by the ground template filtering.

Figure 12 shows the mean of the spectra produced by X^2 PURE from the noise-only simulations, measuring the noise bias. We then take the noise simulations in pairs and evaluate the uncertainty on the noise bias as the standard deviation of the cross spectrum of the two noise maps within a pair.

In order to interpret these results consider first the unbiased case. The analysis of the eigenstructure of $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$ – which is the covariance matrix of the map estimator – has shown that the noisiest modes involve large scales. Since the power spectrum estimator can down-weight pixels but not modes, the large noise power carried by these modes can dominate the power spectrum at large scales. This consideration suggests that the noisiest modes are the cause of the noise increase at large scales compared to the usual ℓ^2 trend.

We therefore filter our unbiased map, progressively removing the noisiest eigenvectors of the $\mathbf{A}^\top \mathbf{F}_T \mathbf{A}$ matrix: if the eigenvalue of a given eigenvector is less than α times the maximum

eigenvalue, the mode is filtered out of the map. We consider several values of α : 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 0.05, 0.08, 10^{-1} . We note that this procedure is similar to the rejection of the noisiest pixels in a map, the only difference is that we remove modes, because our noise is correlated.

As the noisiest modes are removed from the unbiased map (i.e., as α increases) the power spectrum slowly converges to ℓ^2 behavior. We note that this low ℓ noise increase is not caused only by the “long modes” related to the ground template, it is actually dominated by poorly constrained pixel-like modes: using the 10^{-4} threshold the long modes are removed but we still observe an important noise excess. This suggests that the cause is the not the ground template marginalization but the polynomial filtering (or a combination of the two).

On the contrary, in the biased map the way noise is correlated does not cause any noise increase at large scales, both the mean and the standard deviation of the noise power spectra follow the




Fig. 12. Mean (*left*) and standard deviation (*right*) of the B -mode power spectra (as computed by $X^2\text{PURE}$) of 100 noise only simulations for different input maps. “Biased” refers to the biased map estimator defined in Eq. (25). The other lines refer to different maps derived from the unbiased map estimator Eq. (22). The value in the legend is the α parameter that quantifies the amount of map domain filtering applied to the unbiased map. Higher values of α correspond to more aggressive filtering, see Sect. 6.3.1 for more details.

usual ℓ^2 behavior. We stress that in this section both the spectra derived from the biased map and the ones derived from filtered maps are biased: we have to debias them before making quantitative statements about the uncertainty on their spectra, see next section.

6.3.2. Performance comparison

As described in Sect. 5.6, in order to debias the spectra, signal only simulations are used to evaluate the transfer functions f^{XY} (Fig. 13) and noise only simulations for evaluating the noise bias Fig. 12. We use these quantities to get an unbiased power spectrum estimator for each map estimator and apply it to three sets of 100 simulations: E modes only for evaluating the uncertainty due to E to B leakage; noise only for evaluating the uncertainty due to the noise; E and B modes and noise for evaluating the total uncertainty on the BB spectrum (E to B leakage, BB cosmic variance and noise uncertainty).

In the first of these three cases the simulations do not contain noise and therefore we evaluate the uncertainty as the standard deviation of the auto-spectra of each simulation. In Sect. 6.3.1 we have shown that, for the biased power spectrum estimates, the mean value of the noise simulations is comparable with their dispersions, the auto-spectrum of noisy simulations would then result in an asymmetric distribution. Consequently, in order to evaluate the uncertainty of noisy simulations we prefer to group them in pairs and evaluate the uncertainty as the standard deviation of the cross-spectrum of the two simulations of each pair. We note that, when the simulations contain also signal, the signal is the same in the two simulations.

E to B leakage

As expected from the pure formalism implemented in $X^2\text{PURE}$, the unbiased map estimator has basically no uncertainty due to E to B leakage (the leakage is only due to pixelization effects). However the leakage quickly increases as the threshold on the

eigenvalues increases and more modes are filtered from the map. The leakage becomes relevant at large scales for any threshold higher than 0.001. The biased map estimator too has a significant amount of leakage at large scales but performs better than any threshold greater than 0.001.

These considerations can also be made by observing the transfer functions in Fig. 13. The BE transfer function quantifies the average contribution of the EE power in the sky to the BB power in the reconstructed map: the departure from zero is relevant for the biased map estimator only for the first bin and it is considerably more pronounced for $\alpha \geq 0.001$.

Noise uncertainty

In Sect. 6.3.1 we have shown that, for different map estimators and different thresholds α , we get different dispersions of the raw spectra of noise simulations. However, the BB transfer function in Fig. 13 shows that they also have different loss of BB power. Restoring this power boosts the spectrum of the noise too. The interplay between the two effects can be non-trivial.

However, Fig. 14 shows that the latter effect has minor impact: the dispersion of the unbiased spectrum of noise only simulations is still the higher the lower the threshold α and the lowest for the biased map estimator.

Because of the correlated nature of the noise in the estimated maps, the power spectrum estimator can not properly down-weight the noisy modes. Their large fluctuations dominate the power at large scales, boosting the noise uncertainty. Filtering these modes out of the map alleviates this noise excess at large scales and the lowest noise uncertainty is reached by the most aggressive filtering ($\alpha = 0.1$). The biased map estimator performs extremely well in this respect: despite the fact that its noise is correlated too, the uncertainty due to noise is comparable with the ℓ^2 trend expected by the uncorrelated noise case. We stress that the unbiased and biased map estimators preserve the same amount of information (we can convert one into the other anytime using an invertible linear operator). The disparity in their




Fig. 13. Transfer functions of the different map estimators. For the explanation of the legend, see Fig. 12.




Fig. 14. Uncertainty on the unbiased estimation of the B -mode power spectra based on 100 simulations containing E only (left panel), noise only (central panel) and E , B and noise (right panel), see Sect. 6.3.2 for more details. The last case estimates the total uncertainty on the BB power spectrum estimation. The spectra are estimated with help of cross-spectra of simulated pairs of maps. For the explanation of the legend, see Fig. 12. In the case of simulations containing noise we also display the “No filters” case, in which the simulated white noise TOD are not filtered: the simple map-making Eq. (5) is adopted.

power spectrum noise uncertainty is purely due to the fact that we are using a suboptimal weighting for the power spectrum estimation.

Total uncertainty on the BB spectrum

Finally we consider the total uncertainty on the BB spectrum for the different map estimators. In Fig. 14 we show its spectrum while in Fig. 15 we express it as ratio of the lensing BB spectrum and total uncertainty.

As far as the unbiased map estimator is concerned, given the specific noise level of these simulations, the noise plays a dominant role. Controlling its large scale excess is more important than controlling the E to B leakage and consequently we find that the higher the threshold on the eigenvalues the lower the overall uncertainty is. We note that the most aggressive threshold removes more than 50% of the modes but retains more than 90% of the information (computed as the sum of the eigenvalues retained over the sum of all the eigenvalues). The biased map estimator has good noise level over the entire spectrum and, even if the estimator produces E to B leakage at large scales, the resulting uncertainty is below the noise level. As a result it performs

substantially better than any other estimator studied here in the low ℓ part of the spectrum.

We also investigate how the situation would change if the noise level was lower by extrapolating the total uncertainty assuming the observation time was x times longer. For each power spectrum bin, this total uncertainty is evaluated as

$$\sqrt{\sigma_S^2 + \left(\frac{\sigma_N}{x}\right)^2 + 4\frac{N}{x}S}, \quad (61)$$

where S and N are the mean power of the signal only (E and B) and noise only simulations respectively and the σ s are their standard deviations. For the biased map estimator the uncertainty due to E to B leakage is smaller than the BB cosmic variance. Therefore, the unbiased map estimator has superior leakage control but it is not the limiting factor in the case we are considering. Consequently, the factor x required for the unbiased map estimator to have better performance than the biased one is very large (about 10). We emphasize that this statement depends strongly on the specific case we are considering. The situation might be very different if the E to B leakage were to provide a more significant contribution to the overall uncertainty, as can happen when larger scales are probed. In such cases the unbiased map-making approach may be more readily favored.




Fig. 15. Ratio between the signal and the total uncertainty for B -mode power spectra derived with the different power spectrum estimation choices. The results are quoted relative to the highest ratio obtained from the unbiased map is apparent and related to the strong correlated noise modes present in the map. The loss can be mostly recovered with help of a progressively more aggressive removal of the noisiest modes as expressed by the increasing value of parameter α as defined in the text and given in the legend.

7. Conclusions

Time-domain filtering is unavoidable in the analysis of CMB datasets. In this work we have addressed the issue of producing unbiased sky-signal estimates from filtered time-ordered data. We have presented a general map-making formalism that strictly accounts for the presence of time-domain filtering and which is capable of producing as faithful estimates of the sky signal as can be ever obtained in such circumstances. We have shown, however, that some modes of such estimates may be unconstrained or only poorly constrained. This happens whenever there are sky signal modes that are degenerate (or nearly degenerate) with the filtered time-domain modes, and which therefore cannot be fully disentangled from these in the map-making.

Subsequently, we have focused on ground-based experiments and two specific classes of filters – polynomial and ground-synchronous – which are frequently applied to remove long temporal modes and ground-signal pick-up. We have presented general considerations relevant to this case and then demonstrated them on specific, simulated data sets based on the first observational campaigns of the POLARBEAR experiment.

In this latter context we have shown that nearly unbiased maps of all three Stokes parameters can indeed be derived. Nonetheless, these contain a number of poorly constrained sky modes which have to be properly accounted for in the ensuing analysis of the maps in order to fully capitalize on the advantages of the proposed map-making. In particular, we have shown that the performance of simple pseudo-spectrum-based methods, commonly used to estimate the power spectra of the derived maps, may be significantly affected by the presence of such modes. This is because this approach is not naturally capable of weighting different sky modes differently and therefore these

modes tend to either lead to excess noise at low multipoles if they are left in the map, or to enhance the E -to- B leakage if they are excised. We stress that this excess noise at low multipoles should not be interpreted as a downside of the map estimation technique but rather a demonstration that the map-domain correlation induced by popular time-domain filtering operations can severely affect the performances of power-spectrum estimators that use only a pixel-by-pixel weighting.

We have developed a practical approach to compensate for such deficiency but found out that, at best, we can only match the performance of the quicker and simpler biased map-making, based on a simple, noise-weighted binning of the filtered time-domain data. Consequently, more involved and resource-consuming techniques, such as those based on maximum likelihood principles, may need to be employed to exploit at the full potential of the unbiased maps. If, on the contrary, pseudo-power spectrum estimators are used, the unbiased maps can lead to lower E -to- B leakage than the biased ones. Consequently, some improvements in the uncertainty of the recovered B -mode spectrum can be expected whenever E -to- B leakage, and not the noise, is a dominant source of uncertainty and aggressive mode removal is unnecessary.

We note that the availability of nearly unbiased maps should also be important whenever multiple maps (e.g., corresponding to different frequency bands or coming from different experiments), need to be combined, as in pixel-based component separation approaches or map-level cross-analyses involving multiple data sets. In some applications, some of the ill-constrained modes may have to be removed from the unbiased map prior to further processing, as in the power spectrum estimation used in this work. Even though this effectively leads to biased maps, the removed sky modes will be known and the effects of their removal can be kept track of. This has to be further investigated in more detail and on a case-by-case basis, and is therefore left for future work.

Since the specific filters studied in this work are common for ground based experiments, the conclusions derived in this work should be of importance for many operating and planned ground-based experiments.

Acknowledgements. The POLARBEAR project is funded by the National Science Foundation under Grants Nos. AST-0618398 and AST-1212230. The James Ax Observatory operates in the Parque Astronómico Atacama in Northern Chile under the auspices of the Comisión Nacional de Investigación Científica y Tecnológica de Chile (CONICYT). This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the US Department of Energy under Contract No. DE-AC02-05CH11231. KEK authors acknowledge the support of MEXT KAKENHI Grant Number JP15H05891 and JSPS KAKENHI Grant Number JP26220709. In Japan, this work was supported by JSPS Core-to-Core Program, A. Advanced Research Networks and used computational resources of the HPCI system (Project ID:hp150132). In Italy, this work was supported by the RADIOFOREGROUNDS grant of the European Union’s Horizon 2020 research and innovation programme (COMPET-05-2015, grant agreement number 687312) as well as by the INDARK INFN Initiative. J.P. acknowledges support from the Science and Technology Facilities Council (grant number ST/L000652/1) and from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement No. [616170]. C.R. acknowledges support from a Australian Research Council’s Future Fellowship (FT150100074).

References

- Arnold, K., Ade, P. A. R., Anthony, A. E., et al. 2012, *Proc. SPIE Int. Soc. Opt. Eng.*, 8452, 1
- BICEP2 Collaboration 2014, *Phys. Rev. Lett.*, 112, 241101
- BICEP2 and Keck Array Collaborations 2016, *ApJ*, 825, 66
- Blackford, L. S., Choi, J., Cleary, A., et al. 1997, *ScaLAPACK Users’ Guide* (Philadelphia, PA: Society for Industrial and Applied Mathematics)
- Borrill, J. 1999, ArXiv e-print [[arXiv:astro-ph/9911389](https://arxiv.org/abs/astro-ph/9911389)]

- Bunn, E. F., Zaldarriaga, M., Tegmark, M., & De Oliveira-Costa, A. 2003, *Phys. Rev. D*, **67**, 023501
- Cantalupo, C. M., Borrill, J. D., Jaffe, A. H., Kisner, T. S., & Stompor, R. 2010, *ApJS*, **187**, 212
- Culverhouse, T., Ade, P., Bock, J., et al. 2010, *ApJ*, **722**, 1057
- de Gasperis, G., Balbi, A., Cabella, P., Natoli, P., & Vittorio, N. 2005, *A&A*, **436**, 1159
- Doré, O., Teyssier, R., Bouchet, F. R., Vibert, D., & Prunet, S. 2001, *A&A*, **374**, 358
- Dünner, R., Hasselfield, M., Marriage, T. A., et al. 2013, *ApJ*, **762**, 10
- Ferreira, P. G., & Jaffe, A. H. 2000, *MNRAS*, **312**, 89
- Ferté, A., Grain, J., Tristram, M., & Stompor, R. 2013, *Phys. Rev. D*, **88**, 023524
- Ferté, A., Peloton, J., Grain, J., & Stompor, R. 2015, *Phys. Rev. D*, **92**, 083510
- Golub, G. H., & van Loan, C. F. 1996, *Matrix computations* (Johns Hopkins University Press)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Grain, J., Tristram, M., & Stompor, R. 2009, *Phys. Rev. D*, **79**, 123515
- Grain, J., Tristram, M., & Stompor, R. 2012, *Phys. Rev. D*, **86**, 076005
- Hanke, M. 1995, in *Conjugate gradient type methods for ill-posed problems* (CRC Press), 327
- Hivon, E., Gorski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, **567**, 2
- Janssen, M. A., & Gulkis, S. 1992, in *NATO Advanced Science Institutes (ASI) Series C 359*, eds. M. Signore, & C. Dupraz, 391
- Keihänen, E., Kurki-Suonio, H., Poutanen, T., Maino, D., & Burigana, C. 2004, *A&A*, **428**, 287
- Keihänen, E., Keskitalo, R., Kurki-Suonio, H., Poutanen, T., & Sirviö, A.-S. 2010, *A&A*, **510**, A57
- Krachmalnicoff, N., Baccigalupi, C., Aumont, J., Bersanelli, M., & Mennella, A. 2016, *A&A*, **588**, A65
- Matsumura, T., Akiba, Y., Arnold, K., et al. 2016, *J. Low Temperature Phys.*, **184**, 824
- Peebles, P. J. E., & Hauser, M. G. 1974, *ApJS*, **28**, 19
- Planck Collaboration VI. 2016, *A&A*, **594**, A6
- Planck Collaboration VIII. 2016, *A&A*, **594**, A8
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Planck Collaboration Int. XXX. 2016, *A&A*, **586**, A133
- Poutanen, T., Maino, D., Kurki-Suonio, H., Keihänen, E., & Hivon, E. 2004, *MNRAS*, **353**, 43
- QUIET Collaboration 2011, *ApJ*, **741**, 111
- Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, *ApJ*, **743**, 90
- Smith, K. M. 2006, *Phys. Rev. D*, **74**, 083002
- Smith, K. M., & Zaldarriaga, M. 2007, *Phys. Rev. D*, **76**, 043001
- Stompor, R., & White, M. 2004, *A&A*, **419**, 783
- Stompor, R., Balbi, A., Borrill, J., et al. 2001, *Phys. Rev. D*, **65**, 022003
- Szydlarski, M., Grigori, L., & Stompor, R. 2014, *A&A*, **572**, A39
- Tegmark, M. 1997, *Phys. Rev. D*, **56**, 4514
- The CoRE Collaboration 2011, ArXiv e-prints [[arXiv:1102.2181](https://arxiv.org/abs/1102.2181)]
- The POLARBEAR Collaboration 2014, *ApJ*, **794**, 171
- Tran, H., Lee, A., Hanany, S., Milligan, M., & Renbarger, T. 2008, *Appl. Opt.*, **47**, 103
- Tristram, M., Filliard, C., Perdereau, O., et al. 2011, *A&A*, **534**, A88
- Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, *ApJ*, **458**, L53
-
- ¹ AstroParticule et Cosmologie, Univ. Paris Diderot, CNRS/IN2P3, CEA/Irfu, Obs. de Paris, Sorbonne Paris Cité, 75205 Paris Cedex 13, France
e-mail: dpoletti@apc.univ-paris7.fr
- ² International School for Advanced Studies (SISSA), via Bonomea 265, 34136 Trieste, Italy
- ³ INFN, Sezione di Trieste, via Valerio 2, 34127 Trieste, Italy
- ⁴ Department of Physics & Astronomy, University of Sussex, Brighton BN1 9QH, UK
- ⁵ Department of Physics, University of Wisconsin, Madison WI 53706, USA
- ⁶ Department of Physics, University of California, Berkeley, CA 94720, USA
- ⁷ Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
- ⁸ Space Sciences Laboratory, University of California, Berkeley, CA 94720, USA
- ⁹ Department of Physics and Atmospheric Science, Dalhousie University, Halifax, NS, B3H 4R2, Canada
- ¹⁰ Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, 277-8583 Chiba, Japan
- ¹¹ Department of Physics, Blackett Laboratory, Imperial College London, London SW7 2AZ, UK
- ¹² Department of Physics, University of California, San Diego, CA 92093-0424, USA
- ¹³ Sorbonne Universités, Institut Lagrange de Paris (ILP), 98bis boulevard Arago, 75014 Paris, France
- ¹⁴ LPNHE, CNRS-IN2P3 and Universités Paris 6 & 7, 4 place Jussieu, 75252 Paris Cedex 05, France
- ¹⁵ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA
- ¹⁶ High Energy Accelerator Research Organization (KEK), Tsukuba, 305-0801 Ibaraki, Japan
- ¹⁷ SOKENDAI (The Graduate University for Advanced Studies), Hayama, Miura District, 240-0115 Kanagawa, Japan
- ¹⁸ Institute of Physics, Academia Sinica, 128, Sec.2, Academia Road, Nankang, 11529 Taipei, Taiwan
- ¹⁹ Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
- ²⁰ School of Physics, University of Melbourne, Parkville, VIC 3010, Australia
- ²¹ Department of Physics, California Institute of Technology, Pasadena, CA 91125, USA
- ²² Radio Astronomy Laboratory, University of California, Berkeley, CA 94720, USA

Appendix A: Filtering the ground-synchronous signal

In this appendix we study in detail the filtering of ground-synchronous signals, elaborating on and justifying the conclusions presented in Sect. 4.2.

We start by considering the data recorded from a single detector, d , during a constant elevation scan, s . According to the data model in Eq. (37), the data recorded in a given azimuthal bin, ψ , is

$$\begin{aligned} d|_{d,s,\psi} &= \mathbf{A}|_{d,s,\psi} s|_{d,s,\psi} + \mathbf{G}|_{d,s,\psi} g|_{d,s,\psi} + \mathbf{n}|_{d,s,\psi} \\ &= \mathbf{A}|_{d,s,\psi} s|_{d,s,\psi} + \mathbf{1}|_{d,s,\psi} g_\psi + \mathbf{n}|_{d,s,\psi}, \end{aligned} \quad (\text{A.1})$$

where $\mathbf{1}|_{d,s,\psi}$ is a vector of ones of the appropriate length and g_ψ denotes a ground template amplitude common to all selected samples.

Let us now focus on the shape of the sky patch corresponding to $s|_{d,s,\psi}$. The geometry of the problem is depicted in Fig. 1 and, for the time being, we neglect the role of the sky pixels. We consider a small azimuthal change of the pointing direction at a point on the sky at which the parallactic angle is η . The change in horizontal coordinates $\Delta(\text{Az}, \text{El}) = (\delta, 0)$ corresponds to an interval in the equatorial ones given by,

$$\Delta(\text{RA}, \text{Dec}) = (-\delta \cos \eta, -\delta \sin \eta). \quad (\text{A.2})$$

In particular, at the South Pole the horizontal coordinates correspond to the equatorial ones after flipping the y (El) axis, thus we have always $\eta = \pi$ and $\Delta(\text{RA}, \text{Dec}) = (\delta, 0)$. For a given elevation, the parallactic angle is always the same for a given azimuth but depends on the azimuth's value, so although it changes across a single ground template bin the changes are very small. Consequently, each constant elevation scan crossing a bin will draw a line interval on the sky given approximately by Eq. (A.2). Because of the Earth's rotation, if we keep on crossing the bin multiple times the intervals will cover a trapezoidal shape in sky coordinates, as shown in the bottom panel of Fig. 1. The size of the trapezoid depends on the bin width but also on the parallactic angle, Eq. (A.2). However, the lines traced by the azimuthal bin end points always follow the constant declination direction on the sky. We note that if $|\sin \eta| = 0$, which is always the case at the South Pole or whenever the instrument is pointed straight to the South or the North, the width of the trapezoid in declination is zero.

These patches, narrow in declination and elongated in azimuth, are degenerate (or ill-conditioned) modes. In the following we discuss this statement in detail for both temperature and polarization and investigate possible degeneracy-breaking effects.

A.1. Total intensity measurements

Let us start with the total intensity measurements and consider a data subset that has the same ground contribution. It is described by Eq. (A.1), with the pointing matrix, \mathbf{A} , merely composed of ones and zeros. For this single scan the sky modes, $\tilde{\mathbf{s}}$, that are degenerate with the ground template signal have to fulfil the following relation, stemming from Eq. (32),

$$\mathbf{A}|_{d,s,\psi} \tilde{\mathbf{s}} = \mathbf{G}|_{d,s,\psi} \tilde{\mathbf{g}} = \mathbf{1}|_{d,s,\psi} g_\psi. \quad (\text{A.3})$$

Thus for this subset of measurements, the ground template can only give rise to a constant offset in the time-domain for all samples of the scan. Given that \mathbf{A} here simply assigns the pixel amplitudes to the respective time samples without changing their

values, there is only one sky mode that reproduces this behavior: a constant offset in the corresponding sky map,

$$\tilde{\mathbf{s}} \propto \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \equiv \mathbf{1}_{s|_{d,s,\psi}}. \quad (\text{A.4})$$

This demonstrates that, as intuitively expected, the absolute offset of the map produced from these measurements is unavoidably lost as it is degenerate with the ground signal, g_ψ .

Let us consider another data subset taken by detector d' , during scan s' and with the azimuth coordinate corresponding to bin ψ' . The data for this scan can be expressed by a relation analogous to Eq. (A.1). If either the scan, the detector or the azimuthal bin is different between these two data subsets, then the subsets will have independent, and a priori different, ground-pickup amplitudes g_ψ and $g_{\psi'}$. If the sky observed during these scans overlaps, then the combined data set, $d|_{d,s,\psi} \cup d'|_{d',s',\psi'}$, will have again only one degenerate mode pair,

$$\tilde{\mathbf{s}} \propto \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} = \mathbf{1}_{s|_{d,s,\psi} \cup s'|_{d',s',\psi'}} \quad \text{and} \quad \tilde{\mathbf{g}} \propto \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} = \mathbf{1}_{g|_{d,s,\psi} \cup g|_{d',s',\psi'}}. \quad (\text{A.5})$$

This reflects the fact that the relative offset of the two sky maps recovered from each of the subsets can be constrained internally owing to the fact they overlap on the sky and have to recover the same sky signal in each common pixel.

Otherwise, if no overlap exist, the data set made of two subsets will have two degenerate pairs of modes, which can be cast as either the absolute offsets of both of the sky patches or as an absolute offset of both of them and a relative offset between them.

This latter situation can happen if the two subsets correspond to the same detector, d , and the same scan, s , but to two different though adjacent azimuthal bins. As described above, their corresponding sky areas will be indeed strictly speaking disjoint. However, as the sky maps are necessarily pixelized, there will be some pixels on the border between two patches which will straddle both of them, constraining their relative offset. The constraint will not be very strong though, and the corresponding degeneracy only weakly broken. An upshot of this is that a map produced from the single constant elevation scan data of a single detector will have as many ill-constrained modes as there are azimuthal bins used to represent the ground pick-up. These modes will become even more ill-constrained if the number of bins increases because the relative offset uncertainty for sky patches corresponding to two extreme bins increases in proportion to \sqrt{n}_ψ , while the uncertainty on the offset of two adjacent bins remains roughly unchanged. This scaling reflects the fact that the relative offset between two non-overlapping sky patches with n intermediaries is a result of a random walk of the adjacent patches' offsets, each subject to the same uncertainty (Stompor & White 2004). Consequently ill-constrained modes can be suppressed if fewer azimuthal bins are used. Similarly, the uncertainty on the offset of two adjacent bins can be decreased if larger sky pixels are adopted, and more samples from both azimuth bins fall into them. However, the pixel size is typically set by the beam size, while the size of the azimuthal bins is driven by our preconceptions about the ground signal and the structure of the far side lobes. Consequently, whatever freedom is left should be used with care, as potential improvements in statistical uncertainty can be translated into increased ground pick-up residual.

There are two reasons why these ill-constrained modes may be further suppressed in the final maps, combining the data of all

the detector and all of the scans. First, for a single scan the additional constraints on the relative offsets of these patches typically also come from the data collected by different detectors. This is because the azimuthal bins are often defined differently for different detectors so the sky patch corresponding to an azimuthal bin of one detector will often overlap with two sky patches corresponding to two different bins of the other detector, thus providing an extra leverage on their relative offset. Second, the sky patches corresponding to fixed azimuthal bins of different scans can have different width because the parallactic angle is different in the two scans (see previous section). This introduces additional overlaps between the patches of different scans, helping to constrain their relative offsets.

In general the global offset of the final map is expected to be the only truly degenerate mode in the total intensity maps derived from data that are contaminated by ground pickup that requires explicit filtering. Notwithstanding this, the constraints that can be set between the relative offsets of different patches with the same ground pick up are usually inferior to those between different parts of the same patch, and some ill-constrained large scale modes should be expected, predominantly in the declination direction.

We note that these conclusions apply qualitatively to any observational site on Earth, including the poles, with the difference that as the single ground bin patches become very narrow in declination the relative offset degeneracies in this direction are broken only by the pixel effects. By contrast, the bin size only plays a role in breaking the degeneracies in the azimuthal direction.

It is important to appreciate the role of the assumptions in breaking these (near) degeneracies. The choices made about the sky pixel size, the pixelization itself, the binning, and the size of the bins, all impact the degeneracies and can be used, or abused, to break them. In addition, the offset degeneracies can be broken if a less flexible model for the ground signal is used, for instance if we impose a prior constraint on the relative change of the ground signal from one bin to the next. The key parameter in such cases would be the assumed coherence length for the ground signal.

A.2. Polarization-sensitive measurements

The situation for polarization-sensitive observations is potentially more complex due to more complex form of the pointing matrix. However, it is qualitatively similar to the total intensity case. For concreteness, we discuss the case with three Stokes parameters contributing to the measurements, and thus with the pointing matrix as defined in Eq. (1). Since the polarization orientation may change during the operations because of some polarization modulator, we introduce a different ground pick-up amplitude not only for each azimuthal orientation of the instrument but also for each different position of the polarizer, as defined in the instrument coordinates. For simplicity, we will however keep on using a single azimuthal bin number, ψ , to distinguish between the ground signal amplitudes.

We again focus on a single constant elevation scan and a single detector. Our data model is then again given by Eq. (A.1), where the same ground pick up is added to each measurement. The major qualitative difference from the total intensity case is that in the polarization-sensitive case the pointing matrix elements may be different from sample to sample, even for samples falling into the same sky pixel on a single crossing, as could be the case if fast rotating half-wave plate were employed to modulate the signal. However, for the data subset selected above the angle of the polarizer is fixed in the instrument frame, so the

change of the pointing matrix elements, defined by the polarizer orientation but with respect to the sky coordinates, can be only related to the parallactic angle change with the azimuth of the observation. Typically the angle change within a range of azimuths corresponding to a single sky pixel can be safely neglected and we may assign a single polarizer angle as measured with respect to the sky coordinates for each pixel observed with the data subset. These angles may be somewhat different for two different pixels if these are observed at different azimuths, but as the latter have to fall within a single ground template bin, the bins would need to be rather broad to make such an effect important. Nonetheless, henceforth we assume that for the data subset as defined earlier and characterized by the same ground-pick up amplitude, the polarizer's angle in the sky coordinates, and thus the mixing matrix elements, may at most depend on the observed sky pixels and will have a unique value for all observations falling within the same pixel. We note that such small angle variations do not appear if the observations are conducted from the Earth's poles.

As in the total intensity case, the degenerate sky modes have to be able to mimic an offset in time-domain data. This can be the case for three linearly independent sky defined as,

$$\tilde{\mathbf{s}}_I \propto \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{s}}_Q \propto \begin{bmatrix} 0 \\ c_p^{-1} \\ 0 \\ \vdots \\ 0 \\ c_p^{-1} \\ 0 \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{s}}_U \propto \begin{bmatrix} 0 \\ 0 \\ s_p^{-1} \\ \vdots \\ 0 \\ 0 \\ 0 \\ s_{n_p-1} \end{bmatrix}, \quad (\text{A.6})$$

where,

$$s_p \equiv \sin 2\varphi_p \quad (\text{A.7})$$

$$c_p \equiv \cos 2\varphi_p, \quad (\text{A.8})$$

and φ_p stands for a polarizer angle in the sky coordinate in pixel p ($=0, \dots, n_p - 1$). Each of these sky modes is a vector of n_p triples where the elements of each triple correspond to the I , Q , and U Stokes parameters. We note that if the sky rotation is negligible across the sky patch covered by the scan these three modes correspond to map offsets of the maps of the respective Stokes parameters.

Within our data subset, each pixel is observed with only a single orientation of the polarizer and we thus cannot estimate all three Stokes parameters separately, but merely their linear combination, $I + Q \cos 2\varphi_p + U \sin 2\varphi_p$, even if no ground pick up is considered. The corresponding two-dimensional degeneracy space is a subspace of the three-dimensional space spanned by the vectors defined in Eq. (A.6). Consequently, adding the ground pick-up merely adds one degenerate vector to the map-making problem, corresponding to the total offset of the linear combination of the Stokes parameters, $I + Q \cos 2\varphi_p + U \sin 2\varphi_p$ ¹.

To recover all the Stokes parameters from data modeled as in Eq. (1), we need at least three visits to each pixel with a different orientation of the polarizer. These can be provided by other detectors in the focal plane during the same or different constant elevation scans, or come from the same detector if its polarizer direction is modulated either on short or long timescales. In all these cases the new data will have not only a different polarization angle but also potentially a different ground-pickup. Each

¹ We also note that even if one of the cosines, c_p , or sines, s_p , happens to be zero, and therefore only two of the three modes in Eq. (A.6) are indeed degenerate, the latter statement remains true and the loss of information is the same in all these cases.

of these extra data sets can likewise have up to three degenerate sky modes, which for data subset i ($=0, 1, 2$), we denote as $\tilde{\mathbf{s}}_I^{(i)}$, $\tilde{\mathbf{s}}_Q^{(i)}$ and $\tilde{\mathbf{s}}_U^{(i)}$, respectively. For these data sets considered together, however only the intensity offset, $\tilde{\mathbf{s}}_I = \tilde{\mathbf{s}}_I^{(0)} = \tilde{\mathbf{s}}_I^{(1)} = \tilde{\mathbf{s}}_I^{(2)}$, always leads to degeneracy, while $\tilde{\mathbf{s}}_Q^{(i)}$ and $\tilde{\mathbf{s}}_U^{(i)}$ will only do so if the polarizer angles for each data subset are effectively the same for all observed pixels, and therefore

$$\tilde{\mathbf{s}}_Q^{(i)} \propto \tilde{\mathbf{s}}_Q \equiv \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \\ 1 \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{s}}_U^{(i)} \propto \tilde{\mathbf{s}}_U \equiv \begin{bmatrix} 0 \\ 0 \\ 1 \\ \vdots \\ 0 \\ 0 \\ 1 \end{bmatrix}, \quad \text{for } i = 0, 1, 2. \quad (\text{A.9})$$

In this case each of the recovered maps of the Stokes parameters will have an arbitrary offset corresponding to three degenerate vectors, $\tilde{\mathbf{s}}_I$, $\tilde{\mathbf{s}}_Q$, and $\tilde{\mathbf{s}}_U$, as defined.

If the angles do change somewhat from pixel to pixel within a single data subset (i.e., when the change of the parallactic angle within the azimuthal bin is not negligible) only the total intensity map will have an arbitrary offset. This is because in this case $\mathbf{A}^{(i)} \tilde{\mathbf{s}}_Q^{(j)}$ and $\mathbf{A}^{(i)} \tilde{\mathbf{s}}_U^{(j)}$ are time-domain vectors with elements which depend on time in a non-trivial way. Here, $\mathbf{A}^{(i)}$ is a pointing matrix specific to subset i , while the combined pointing matrix for the three subset is given by,

$$\mathbf{A} \equiv \begin{bmatrix} \mathbf{A}^{(0)} \\ \mathbf{A}^{(1)} \\ \mathbf{A}^{(2)} \end{bmatrix}. \quad (\text{A.10})$$

We therefore also see that the time-domain vectors, $\mathbf{A} \tilde{\mathbf{s}}_Q^{(j)}$ and $\mathbf{A} \tilde{\mathbf{s}}_U^{(j)}$, are non-trivial and therefore cannot typically be mimicked by three ground template offsets and the degeneracy condition in Eq. (32) can not be fulfilled.

However, as the angle change due to the sky rotation is typically small, $\tilde{\mathbf{s}}_Q$ and $\tilde{\mathbf{s}}_U$ may be potentially ill-constrained, even if not strictly singular, and the offsets of the Q and U maps may be very uncertain.

The offsets between sky patches corresponding to adjacent ground template bins during the same constant elevation scan can be further constrained as in the case of the total intensity only measurements. The potential degeneracies can then be suppressed with the help of data from the other detectors and/or different scans, although again a natural expectation is that there will be long sky modes in the declination direction which may be ill-constrained.

If the observation is taken from the Earth's poles, the maps recovered from the three subsets of the data taken at the same elevation will have all three degenerate offsets, which will propagate to the final maps combining all the data. In addition, the relative offsets between the sky patches taken at different elevation will only be set by the presence of pixels common to both patches and therefore will lead to long modes in declination which will be ill-constrained.