



HAL
open science

Wiener filtering and pure \mathcal{E}/\mathcal{B} decomposition of CMB maps with anisotropic correlated noise

Doogesh Kodi Ramanah, Guilhem Lavaux, Benjamin D. Wandelt

► **To cite this version:**

Doogesh Kodi Ramanah, Guilhem Lavaux, Benjamin D. Wandelt. Wiener filtering and pure \mathcal{E}/\mathcal{B} decomposition of CMB maps with anisotropic correlated noise. Monthly Notices of the Royal Astronomical Society, 2019, 490 (1), pp.947-961. 10.1093/mnras/stz2608 . hal-02193940

HAL Id: hal-02193940

<https://hal.science/hal-02193940>

Submitted on 4 Jul 2023

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

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

Wiener filtering and pure \mathcal{E}/\mathcal{B} decomposition of CMB maps with anisotropic correlated noise

Doogesh Kodi Ramanah ^{1,2,★} Guilhem Lavaux ^{1,2} and Benjamin D. Wandelt^{1,2,3}

¹*Sorbonne Université, CNRS, UMR 7095, Institut d’Astrophysique de Paris, 98 bis bd Arago, F-75014 Paris, France*

²*Sorbonne Université, Institut Lagrange de Paris (ILP), 98 bis bd Arago, F-75014 Paris, France*

³*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA*

Accepted 2019 September 8. Received 2019 August 26; in original form 2019 June 30

ABSTRACT

We present an augmented version of our dual messenger algorithm for spin field reconstruction on the sphere, while accounting for highly non-trivial and realistic noise models such as modulated correlated noise. We also describe an optimization method for the estimation of noise covariance from Monte Carlo simulations. Using simulated Planck polarized cosmic microwave background (CMB) maps as a showcase, we demonstrate the capabilities of the algorithm in reconstructing pure \mathcal{E} and \mathcal{B} maps, guaranteed to be free from ambiguous modes resulting from the leakage or coupling issue that plagues conventional methods of \mathcal{E}/\mathcal{B} separation. Due to its high speed execution, coupled with lenient memory requirements, the algorithm can be optimized in exact global Bayesian analyses of state-of-the-art CMB data for a statistically optimal separation of pure \mathcal{E} and \mathcal{B} modes. Our algorithm, therefore, has a potentially key role in the data analysis of high-resolution and high-sensitivity CMB data, especially with the range of upcoming CMB experiments tailored for the detection of the elusive primordial \mathcal{B} -mode signal.

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

1 INTRODUCTION

Cosmological inference from observations of the cosmic microwave background (CMB) polarization necessitates the separation of the contributions of the gradient and curl (or \mathcal{E} and \mathcal{B}) components of the polarization signal to the data. These scalar \mathcal{E} and pseudo-scalar \mathcal{B} modes correspond to the spin-2 analogues of curl-free and divergence-free vector fields, respectively, with a polarization map being represented as the sum of both components (Zaldarriaga & Seljak 1997; Seljak & Zaldarriaga 1997; Kamionkowski, Kosowsky & Stebbins 1997a,b). The next generation of CMB experiments is focused on measuring the polarization of the CMB, with the \mathcal{E} -mode power spectrum providing an independent probe of the scalar modes measured via the temperature anisotropies (e.g. Abazajian et al. 2016; Suzuki et al. 2016; Louis et al. 2017; Henning et al. 2018). The \mathcal{B} -mode component of the polarization is the focus of growing interest in the community. First, they are an independent confirmation of the lensing effect detected in the temperature and \mathcal{E} -mode anisotropies, as \mathcal{B} modes are produced from the gravitational lensing of \mathcal{E} modes by the dark matter distribution along the line of sight (Zaldarriaga & Seljak 1998). These lensing-induced \mathcal{B}

modes have been observed by high-resolution ground-based CMB experiments (e.g. Hanson et al. 2013; POLARBEAR Collaboration et al. 2017). Secondly, and most importantly, the detection of larger angular scale \mathcal{B} modes, directly sourced by primordial gravitational waves, remains a key but elusive objective of modern cosmology (e.g. Guzzetti et al. 2016).

Observations of CMB polarization have attracted much interest due to the significance of the cosmological information encoded (e.g. Hu & White 1997; Hu & Dodelson 2002; Hu 2003). The inclusion of \mathcal{E} -mode polarization data in parameter inference pipelines allows us to derive more stringent constraints (e.g. Galli et al. 2014), whilst the scientific potential of \mathcal{B} -mode anisotropy observations is extremely promising. A measurement of \mathcal{B} -mode signal on large angular scales ($\ell \lesssim 100$), after discarding the expected lensed signal, would be regarded as a direct validation of the inflationary paradigm as the precursor of this stochastic background of gravitational waves (Kamionkowski & Kovetz 2016). The amplitude of this background would directly constrain the energy scale of inflation, thereby ruling out some inflationary models (e.g. Zaldarriaga, Spergel & Seljak 1997; Kinney 1998; Tegmark et al. 2000), while also constraining the reionization period (Zaldarriaga 1997). This would dramatically improve our understanding of the very early Universe.

The decomposition of the \mathcal{E} and \mathcal{B} modes on a partial sky is highly non-trivial due to the induced leakage between the two modes.

* E-mail: ramanah@iap.fr

Masked regions produce a discontinuity at the edges of the map and these result in a mixing of \mathcal{E} and \mathcal{B} modes, yielding ambiguous modes. Such modes can be sourced by either \mathcal{E} or \mathcal{B} components and cannot be uniquely assigned (Zaldarriaga 2001; Bunn 2002a,b; Lewis, Challinor & Turok 2002; Bunn et al. 2003; Lewis 2003; Bunn 2011). This is a highly prevalent issue due to most observations being made on an incomplete sky or full-sky maps being subjected to additional masking to reduce foreground contamination. Since the \mathcal{E} -mode power spectrum is much larger than that of \mathcal{B} modes, the ambiguous modes significantly increase the variance of the \mathcal{B} modes, resulting in a spurious measurement of the \mathcal{B} -mode power spectrum. The detection of the inflationary gravitational waves is especially challenging due to their relatively small amplitude, and hence, efficient methods for pure \mathcal{E}/\mathcal{B} decomposition are essential for extracting the cosmological information from CMB polarization data.

Several approaches for pure \mathcal{E}/\mathcal{B} decomposition are described in the literature. While some techniques yield real-space maps of the derivatives of the polarization maps (e.g. Kim & Naselsky 2010; Zhao & Baskaran 2010; Kim 2011; Bowyer, Jaffe & Novikov 2011), others are limited to power-spectrum estimation via the construction of an eigenbasis for the pure-ambiguous decomposition (e.g. Challinor & Chon 2005; Smith 2006a,b; Smith & Zaldarriaga 2007; Grain, Tristram & Stompor 2009; Alonso et al. 2019). Nevertheless, such approaches do not result in pure \mathcal{E} and \mathcal{B} maps in the real space and are computationally intensive. Ferté et al. (2013) provides a quantitative comparison of the efficiency of the above techniques for power-spectrum reconstruction. Wavelet-based techniques have also been proposed, but they must be carefully adapted for the problem under investigation (Cao & Fang 2009; Rogers et al. 2016; Leistedt et al. 2017).

Bunn & Wandelt (2017, hereafter **BW17**) have recently shown that the \mathcal{E}/\mathcal{B} decomposition can be approached from a Wiener filtering (Wiener 1949) viewpoint, resulting in faster implementation as compared to the above methods, while providing real-space maps of the \mathcal{E} and \mathcal{B} modes. Another key advantage of such an approach is that it can be naturally extended to treat more interesting cases such as providing \mathcal{E} maps free from the temperature anisotropy contributions by accounting for temperature and polarization correlations.

In this work, we present an augmented version of our dual messenger algorithm (Kodi Ramanah, Lavaux & Wandelt 2017, 2018) for pure \mathcal{E}/\mathcal{B} decomposition on the sphere, based on the principle of the Wiener filter. We adapt the algorithm to encode the **BW17** prescription for reconstruction of pure \mathcal{E}/\mathcal{B} maps and naturally extend the dual messenger framework to account for complex and realistic noise models. We demonstrate the application of this enhanced algorithm, designated as DANTE (DuAl messeNger filTEr), on a simulated CMB polarization data set that emulates the features of the actual Planck data.

The remainder of this paper is structured as follows. In Section 2, we provide a brief description of the dual messenger algorithm and illustrate a Jacobi relaxation method to account for the non-orthogonality of spherical harmonic (SH) transforms. We describe how it can be augmented to deal with non-trivial noise covariance in Section 3. We subsequently illustrate how the algorithm can encode the **BW17** prescription for pure \mathcal{E}/\mathcal{B} reconstruction, followed by an outline of the numerical implementation in Section 4. We then present a new optimization scheme for the estimation of noise covariance from Monte Carlo simulations and showcase the capabilities of DANTE in reconstructing pure \mathcal{E} and \mathcal{B} maps from simulated Planck data in Section 5. Finally, we summarize our main

findings in Section 6. In Appendix A, we provide a brief description of SH transforms, as employed in this work. Appendix B outlines the main steps in the derivation of the essential dual messenger equations to account for anisotropic correlated noise.

2 DUAL MESSENGER ALGORITHM

We briefly review the underlying framework of the dual messenger algorithm for Wiener filtering polarized CMB maps. A complementary description is provided in Kodi Ramanah et al. (2018).

2.1 Wiener filter

In linear data analysis, we often encounter the computation of the so-called Wiener filter on large and complex data sets. The Wiener filter originates from the following statistical problem. We assume our observed data set \mathbf{d} to be a linear combination of the signal \mathbf{s} with covariance \mathbf{S} and noise \mathbf{n} with covariance \mathbf{N} as follows:

$$\mathbf{d} = \mathcal{R}\mathbf{s} + \mathbf{n}, \quad (1)$$

where the signal and noise covariances are given by $\mathbf{S} = \langle \mathbf{s}\mathbf{s}^\dagger \rangle$ and $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\dagger \rangle$, respectively. \mathcal{R} is the complete response operator, with the inclusion of harmonic transforms, beam and mask effects, that models the instrument response to incoming signal. It effectively corresponds to the overall model of how the instrument converts, on average, an incoming signal \mathbf{s} to the observed data \mathbf{d} , with the residual being the noise \mathbf{n} , while encoding the relevant physics.

The Wiener filter solution, \mathbf{s}_{WF} , is the optimal linear filter when the signal and noise are both Gaussian random fields. For a particular realization of the data, \mathbf{s}_{WF} therefore maximizes the posterior probability distribution $\propto \exp(-\chi^2/2)$, or equivalently minimizes

$$\chi^2 = (\mathbf{d} - \mathcal{R}\mathbf{s})^\dagger \mathbf{N}^{-1} (\mathbf{d} - \mathcal{R}\mathbf{s}) + \mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s}, \quad (2)$$

leading to the Wiener filter equation,

$$\mathbf{s}_{\text{WF}} = (\mathbf{S}^{-1} + \mathcal{R}^\dagger \mathbf{N}^{-1} \mathcal{R})^{-1} \mathcal{R}^\dagger \mathbf{N}^{-1} \mathbf{d}. \quad (3)$$

\mathbf{s}_{WF} is the least-square optimal solution: the Wiener filter minimizes the mean-square deviations $\langle \boldsymbol{\epsilon}^\dagger \boldsymbol{\epsilon} \rangle$ of the reconstruction errors $\boldsymbol{\epsilon} = \mathbf{s}_{\text{WF}} - \mathbf{s}$, averaged over all signal and noise realizations.

The Wiener filter (Wiener 1949) is a particularly important signal reconstruction technique, with ubiquitous applications in cosmology and astrophysics (e.g. Elsner & Wandelt 2013, and references therein). Computing the Wiener filter solution for large and complex data sets from modern experiments, however, is numerically challenging. Indeed, the first matrix inversion above is dense in all bases and lives in a high-dimensional space. This space has typically the size of the number of elements in \mathbf{d} , which for Planck maps is $\mathcal{O}(10^9)$, when accounting for polarization components and the nine frequency bands. Due to the size of the covariance matrices scaling as the square of the number of data samples, the storage and processing of dense systems become numerically intractable. Traditional approaches of computing the Wiener filter rely on costly and highly non-trivial numerical schemes, such as preconditioned conjugate gradient (PCG) methods (Eriksen et al. 2004; Wandelt, Larson & Lakshminarayanan 2004; Smith, Zahn & Doré 2007; Seljebotn et al. 2014, 2019; Papež, Grigori & Stompor 2018; Puglisi et al. 2018), requiring a pre-conditioner to approximate the matrix inversion involved. These complex techniques suffer from various numerical limitations, as discussed extensively in Elsner & Wandelt (2013) and Kodi Ramanah et al. (2017), when dealing with high-dimensional data sets. Such PCG methods do have some merits, however, as they are conducive to fast convergence, provided an

adequate pre-conditioner tailored to the specific problem is available. A recent work by Horowitz, Seljak & Aslanyan (2018) was based on recasting the Wiener filtering problem as an optimization scheme and provides an alternative promising approach for dealing with complex noise models. Münchmeyer & Smith (2019) recently proposed another interesting approach where a neural network was trained to Wiener filter CMB maps.

We recently presented the dual messenger algorithm, an enhanced variant of the standard messenger technique developed by Elsner & Wandelt (2013), as a general purpose Wiener filtering tool, which surpasses its predecessor in execution time (Kodi Ramanah et al. 2017). We have demonstrated the efficiency and unconditional stability of the dual messenger technique in Wiener filtering high-resolution polarized CMB data with correlated noise (Kodi Ramanah et al. 2018, hereafter KLW18). As a preconditioner-free approach, this method circumvents the limitations of conventional PCG solvers in dealing with inherently ill-conditioned systems encountered in typical CMB polarization problems, as illustrated in KLW18.

The messenger method has been implemented within an efficient Gibbs sampling framework for Bayesian large-scale structure inference (Jasche & Lavaux 2015), with this Gibbs-messenger sampling scheme subsequently adapted for CMB gravitational lensing (Anderes, Wandelt & Lavaux 2015) and cosmic shear analyses (Alsing et al. 2016; Alsing, Heavens & Jaffe 2017). The messenger technique is being further developed for data analysis involving dense noise covariance matrices (Huffenberger 2018) and has emerged as a promising CMB map-making tool (Huffenberger & Naess 2018). The dual messenger algorithm has also been implemented in the field of optical and information engineering. For instance, it has been adapted for the removal of atmospheric haze from images (Fu et al. 2018), demonstrating the versatility of the tool developed. This class of messenger methods can therefore be tailored to solve a range of Wiener filtering problems and is not limited to astrophysical and cosmological applications.

2.2 Dual messenger field on the sphere

Conceptually, the essence of the messenger methods lies in the introduction of an auxiliary field to mediate between the different bases where the signal and noise covariances, \mathbf{S} and \mathbf{N} , can be most conveniently expressed as sparse matrices. The addition of this messenger field allows the Wiener filter equation to be rewritten as a set of algebraic equations that must be solved iteratively, circumventing the requirement of matrix inversions or pre-conditioners.

With respect to the standard messenger technique, where a messenger field \mathbf{t} is introduced at the level of the noise, the dual messenger algorithm incorporates an extra messenger field \mathbf{u} , at the level of the signal, with corresponding covariances \mathbf{T} and \mathbf{U} , resulting in the following augmented χ^2 :

$$\chi_{T,U}^2 = (\mathbf{d} - \mathbf{t})^\dagger \bar{\mathbf{N}}^{-1} (\mathbf{d} - \mathbf{t}) + (\mathbf{t} - \mathcal{Y}\mathbf{B}\mathbf{u})^\dagger \mathbf{T}^{-1} (\mathbf{t} - \mathcal{Y}\mathbf{B}\mathbf{u}) + (\mathbf{u} - \mathbf{s})^\dagger \mathbf{U}^{-1} (\mathbf{u} - \mathbf{s}) + \mathbf{s}^\dagger \bar{\mathbf{S}}^{-1} \mathbf{s}, \quad (4)$$

where $\bar{\mathbf{N}} = \mathbf{N} - \mathbf{T}$ with $\mathbf{T} = \alpha \mathbf{1}$, where $\alpha \equiv \min(\text{diag}(\mathbf{N}))$, and $\bar{\mathbf{S}} = \mathbf{S} - \mathbf{U}$ with $\mathbf{U} = \nu \mathbf{1}$, where $\nu \equiv \min(\text{diag}(\mathbf{S}))$. When dealing with polarization fields, μ and ν are actually 3×3 matrices, corresponding to the temperature, \mathcal{E} and \mathcal{B} components. \mathcal{Y} and \mathcal{Y}^\dagger correspond to the basis operators (synthesis and analysis operators, respectively) for the SH transforms, as described in Appendix A, while \mathbf{B} indicates convolution with an instrument beam. In terms of physical significance, \mathbf{t} corresponds to a homogeneous portion of

the noise covariance while its counterpart \mathbf{u} is associated with the signal covariance. Optimizing $\chi_{T,U}^2$ yields the following system of equations to be solved iteratively:

$$\mathbf{u} = (\bar{\mathbf{S}} + \mathbf{U}) [\mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{Y} \mathbf{B} (\bar{\mathbf{S}} + \mathbf{U}) + \mathbf{T}]^{-1} \mathbf{B}^\dagger \mathcal{Y}^\dagger \mathbf{t} \quad (5)$$

$$\mathbf{t} = (\bar{\mathbf{N}}^{-1} + \mathbf{T}^{-1})^{-1} (\mathbf{T}^{-1} \mathcal{Y} \mathbf{B} \mathbf{u} + \bar{\mathbf{N}}^{-1} \mathbf{d}). \quad (6)$$

Note that this is the reduced system of equations, with one of the messenger fields made implicit. The general system of equations is described in more depth in KLW18. To improve convergence, we implement a similar scheme as in KLW18. We artificially truncate the signal covariance \mathbf{S} to some lower initial value of ℓ_{iter} , corresponding to a covariance μ . We subsequently vary \mathbf{U} to bring $\mu \rightarrow \nu$, such that in the limit $\mu = \nu$, we have $\mathbf{u} = \mathbf{s}$ and we recover the usual Wiener filter equation (3) from the above system of equations (5) and (6). This is formally valid as long as μ and ν are block matrices over harmonic space. We may therefore exploit this degree of freedom to solve the temperature and polarization signals at different rates. The above cooling scheme leads to a redefinition of $\bar{\mathbf{S}}$ using the Heaviside function as $\bar{\mathcal{S}} = \Theta(\mathcal{S} - \mathbf{U})$, where \mathcal{S} corresponds to the eigenvalues of \mathbf{S} , i.e. $\mathbf{S} = \mathfrak{R}^\dagger \mathcal{S} \mathfrak{R}$ and $\bar{\mathbf{S}} = \mathfrak{R}^\dagger \bar{\mathcal{S}} \mathfrak{R}$.

To implement such a hierarchical scheme, we vary \mathbf{U} via a cooling scheme for ξ , where $\xi = \mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{Y} \mathbf{B} \mathbf{U} + \mathbf{T}$. To obtain the appropriate Wiener filter solution, we need to reduce $\mu \rightarrow \nu = 0$, due to the continuous mode of the signal, i.e. the zero eigenvalue of \mathbf{S} . The cooling scheme for ξ entails gradually reducing ξ by a constant factor and iterating until $\xi \rightarrow \mathbf{T}$, thereby finally attaining $\mu = 0$, as required. A quantitative description of the rationale underlying the above cooling scheme is presented in our previous work (Kodi Ramanah et al. 2017).

DANTE is implemented in PYTHON and it makes use of the HEALPIX¹ (Górski et al. 2005) library, in particular the PYTHON wrapper HEALPY, to perform the spherical harmonic transforms (SHTs). HEALPIX employs an equal area projection scheme, where the SHTs are quasi-orthogonal, i.e. $\mathcal{Y}^\dagger \mathcal{Y} \mathbf{1} \approx (N_{\text{pix}}/4\pi) \mathbf{1} \equiv \beta \mathbf{1}$, where N_{pix} denotes the number of pixels. We account for this non-orthogonality of SHTs via efficient Jacobi relaxation schemes, as described in future sections. If an equidistant cylindrical projection on a grid is adopted for the SHTs (e.g. Muciaccia, Natoli & Vittorio 1997; Huffenberger & Wandelt 2010; McEwen & Wiaux 2011), the equations presented in this work are significantly simplified, as a result of $\beta = 1$. DANTE employs the NUMBA² (Lam, Pitrou & Seibert 2015) compiler for PYTHON arrays and numerical functions to yield high-performance functions for all the required matricial manipulations to boost execution speed. NUMBA generates optimized native code using the LLVM compiler (Lattner & Adve 2004) infrastructure and is used to parallelize the array operations.

2.3 Non-orthogonality of SH transforms

Unlike in the case of discrete Fourier transforms, the SH synthesis and analysis operators, i.e. \mathcal{Y} and \mathcal{Y}^\dagger , respectively, are not orthogonal and differ by more than a transposition and a scale factor. The quality of the approximation, $\mathcal{Y}^\dagger \mathcal{Y} \mathbf{1} \approx (N_{\text{pix}}/4\pi) \mathbf{1} \equiv \beta \mathbf{1}$, depends on the ℓ_{max} , N_{pix} and spherical grid considered.

¹<http://healpix.jpl.nasa.gov>

²<https://numba.pydata.org>

To account for the non-orthogonality of SH transforms, we incorporate an internal Jacobi relaxation method (Jacobi 1845; Saad 2003) in DANTE to refine the solution in harmonic space (cf. equation 5). To obtain \mathbf{U} , we need the eigenvalues of \mathbf{S} , i.e. $\mathbf{U} = \min(\mathcal{S})\mathbb{1}$, where $\mathfrak{R}^\dagger \mathcal{S} \mathfrak{R} = \mathbf{S}$. Equation (5) can be formulated as $\mathbf{s} = \mathcal{A}\mathbf{b}$, where \mathcal{A} is given by

$$\mathcal{A} = \mathfrak{R}^\dagger(\bar{\mathcal{S}} + \mathbf{U}) [\mathfrak{R}\mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{Y}\mathbf{B}\mathfrak{R}^\dagger(\bar{\mathcal{S}} + \mathbf{U}) + \alpha\mathbb{1}]^{-1} \mathfrak{R}, \quad (7)$$

after including the basis transformations, and $\mathbf{b} = \mathbf{B}^\dagger \mathcal{Y}^\dagger \mathbf{t}$.

An approximation to \mathcal{A} can be obtained as follows:

$$\tilde{\mathcal{A}} = \mathfrak{R}^\dagger(\bar{\mathcal{S}} + \mathbf{U}) [\beta \mathfrak{R}\mathbf{B}^\dagger \mathbf{B}\mathfrak{R}^\dagger(\bar{\mathcal{S}} + \mathbf{U}) + \alpha\mathbb{1}]^{-1} \mathfrak{R}, \quad (8)$$

after using the approximate orthogonality relation $\mathcal{Y}^\dagger \mathcal{Y}\mathbb{1} \approx \beta\mathbb{1}$. The application of the operator \mathcal{A} is not well-defined but we nevertheless can apply its inverse \mathcal{A}^{-1} to a vector, by applying the relevant operators sequentially,

$$\mathcal{A}^{-1} = \mathfrak{R}^\dagger [\mathfrak{R}\mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{Y}\mathbf{B}\mathfrak{R}^\dagger + \alpha(\bar{\mathcal{S}} + \mathbf{U})^{-1}] \mathfrak{R}. \quad (9)$$

We therefore make use of \mathcal{A}^{-1} and $\tilde{\mathcal{A}}$ to obtain the solution for \mathbf{s} via the following Jacobi iterations:

$$\mathbf{s}_{n+1} = \mathbf{s}_n + \tilde{\mathcal{A}}(\mathbf{b} - \mathcal{A}^{-1}\mathbf{s}_n), \quad (10)$$

where n denotes the number of Jacobi iterations.

The term $(\bar{\mathcal{S}} + \mathbf{U})^{-1}$ poses a numerical predicament for the final truncation in the signal covariance, where $\bar{\mathcal{S}} = \mathcal{S} - \mathbf{U}$ and $\mathbf{U} = \mathbf{0}$, and we subsequently require the inversion of \mathcal{S} . We circumvent the corner case due to the zero eigenvalues of the continuous modes in \mathcal{S} by imposing the following constraint on the subspace \mathcal{V} where $\mathcal{S} = \mathbf{0}$: $(\bar{\mathcal{S}} + \mathbf{U})^{-1}|_{\mathcal{V}} = \mathcal{S}^+$. \mathcal{S}^+ is the pseudo-inverse, i.e. $\mathcal{S}\mathcal{S}^+ = \Pi$, where Π is a projector.

2.4 Incomplete sky coverage

CMB data analysis inevitably requires the treatment of masks, with many practical applications requiring that parts of the sky be masked out. For full-sky observations, this is mainly to avoid contamination from the galactic foreground emissions, thereby preventing spurious power-spectra measurements. In the case of ground-based or sub-orbital CMB experiments with partial sky coverage, missing data are accounted for using masks.

We provide an outline of the general procedure for solving the messenger equation (6) when dealing with temperature and polarization masks. Here, we assume correlated noise, such that the noise covariance \mathbf{N} has the following block-diagonal form for every pixel i :

$$\mathbf{N}_i = \begin{pmatrix} \langle \mathcal{I}\mathcal{I} \rangle & \langle \mathcal{I}\mathcal{Q} \rangle & \langle \mathcal{I}\mathcal{U} \rangle \\ \langle \mathcal{Q}\mathcal{I} \rangle & \langle \mathcal{Q}\mathcal{Q} \rangle & \langle \mathcal{Q}\mathcal{U} \rangle \\ \langle \mathcal{U}\mathcal{I} \rangle & \langle \mathcal{U}\mathcal{Q} \rangle & \langle \mathcal{U}\mathcal{U} \rangle \end{pmatrix}, \quad (11)$$

where \mathcal{I} , \mathcal{Q} , and \mathcal{U} are the Stokes parameters. More complex noise models will be described in Section 3.

We compute the covariance \mathbf{T} of the messenger field \mathbf{t} as follows: $\mathbf{T} = \min(\mathcal{N})\mathbb{1} = \alpha\mathbb{1}$, where $\mathcal{D}^\dagger \mathcal{N} \mathcal{D} = \mathbf{N}$. The noise covariance \mathbf{N} can be written as $\mathbf{N} = \boldsymbol{\Sigma}\mathbf{C}\boldsymbol{\Sigma} = \boldsymbol{\Sigma}\mathbf{P}^\dagger \boldsymbol{\Delta}\mathbf{P}\boldsymbol{\Sigma}$, following the orthonormal decomposition of \mathbf{C} , where $\boldsymbol{\Sigma}$ is a diagonal matrix with the eigenvalues $\{\sigma_{\mathcal{I}}^i, \sigma_{\mathcal{Q}}^i, \sigma_{\mathcal{U}}^i\}$ corresponding to the noise amplitudes for the i th pixel, with the orthonormal decomposition of \mathbf{C} resulting in the diagonal matrix $\boldsymbol{\Delta}$. We then obtain $\bar{\mathbf{N}} = \mathbf{N} - \mathbf{T}$ as follows:

$$\begin{aligned} \bar{\mathbf{N}} &= \boldsymbol{\Sigma}\mathbf{P}^\dagger \boldsymbol{\Delta}\mathbf{P}\boldsymbol{\Sigma} - \mathbf{T} \\ &= \boldsymbol{\Sigma}\mathbf{P}^\dagger (\boldsymbol{\Delta} - \alpha\mathbf{P}\boldsymbol{\Sigma}^{-2}\mathbf{P}^\dagger)\mathbf{P}\boldsymbol{\Sigma}, \end{aligned} \quad (12)$$

where α , as stated above, is the smallest eigenvalue of \mathbf{N} . To solve the messenger equation (6), we require the inverse $\bar{\mathbf{N}}^{-1}$,

$$\bar{\mathbf{N}}^{-1} = \boldsymbol{\Sigma}^{-1}\mathbf{P}^\dagger (\boldsymbol{\Delta} - \alpha\mathbf{P}\boldsymbol{\Sigma}^{-2}\mathbf{P}^\dagger)^{-1}\mathbf{P}\boldsymbol{\Sigma}^{-1}, \quad (13)$$

such that $\bar{\mathbf{N}}^{-1}$ has a block-diagonal structure in pixel space. We obtain the solution for the messenger field by simply evaluating equation (6) in pixel space,

$$\mathbf{t}_x = (\bar{\mathbf{N}}^{-1} + \mathbf{T}^{-1})_x^{-1} (\mathbf{T}_x^{-1} \mathcal{Y}\mathbf{B}\mathbf{u}_\ell + \bar{\mathbf{N}}_x^{-1} \mathbf{d}_x). \quad (14)$$

We implement the temperature and polarization masks by increasing the noise variance to infinity for masked pixels, or numerically by setting the inverse noise covariance to zero. This is achieved by setting $\boldsymbol{\Sigma}^{-1}|_{\text{mask}} = \mathbf{0}$, subsequently ensuring that data from masked regions do not contaminate the messenger field.

2.5 Constrained realizations

For full-sky coverage with parts of the sky masked out, we still seek the Wiener filter solution under the mask, constrained by the observations on the edges of the mask and determined by the prior inside the masked region. While this proposed reconstruction does not correspond to the true solution, it has the correct statistical properties, i.e. correct signal covariance. The generation of such constrained realizations is relevant for many practical CMB applications, such as exact likelihood evaluations via Gibbs sampling. A complementary conceptual discussion of the rationale underlying constrained realizations is provided in K LW18.

To draw Gaussian constrained realizations of the CMB sky, we need to simulate a reference signal $\hat{\mathbf{s}}$ in accordance with the prior signal covariance assumed for the Wiener filter. We also require a simulated data set $\hat{\mathbf{d}}$ whose signal and noise properties correspond to that of the data model. We adapt DANTE to generate constrained realizations (e.g. using the algorithm of Hoffman & Ribak 1991) in only one application of the Wiener filter via

$$\mathbf{s}_{\text{CR}} = \tilde{\mathcal{W}}(\mathbf{d} - \hat{\mathbf{d}}) + \hat{\mathbf{s}}, \quad (15)$$

where $\tilde{\mathcal{W}}$ indicates the application of the dual messenger operator at a given precision and cooling scheme. We therefore only need to modify the input data fed to DANTE to draw unbiased constrained realizations of the signal that are consistent with the observed data, i.e. having the correct covariance properties.

3 DUAL MESSENGER GENERALIZATIONS

The dual messenger approach can be extended to a broader class of problems, accounting for highly non-trivial noise covariance, as described in the following sections. In practice, the structure of the noise covariance, in pixel space, is influenced by the noise properties of the CMB time-ordered data and the scanning pattern of the telescope.

3.1 Correlated modulated noise

The first possible generalization for the noise model is the case of correlated modulated noise. The noise covariance, in pixel space, can be written as

$$\mathbf{N} = \mathcal{Y}\mathbf{F}\mathcal{Y}^\dagger \mathbf{D}\mathcal{Y}\mathbf{F}^\dagger, \quad (16)$$

where \mathbf{F} is a smoothing kernel which is diagonal in the same basis as \mathbf{S} , i.e. harmonic space, while \mathbf{D} is the noise variance that can

be easily diagonalized in some other basis, e.g. pixel space. The desired Wiener filter from equation (3) is then

$$s_{\text{WF}} = \mathbf{S}(\mathbf{S} + \mathbf{F}\mathbf{Y}^\dagger\mathbf{D}\mathbf{Y}\mathbf{F})^{-1}\mathbf{Y}^{-1}\mathbf{d}. \quad (17)$$

It turns out that this can be solved directly by the dual messenger scheme described above without any modifications. We transform the data via a simple pre-whitening step as follows:

$$\tilde{\mathbf{d}} = (\mathbf{Y}\mathbf{F}\mathbf{Y}^\dagger)^{-1}\mathbf{d}. \quad (18)$$

The Wiener filter for this model can then be computed via the following steps:

$$\begin{aligned} \tilde{s}_{\text{WF}} &= (\mathbf{F}\mathbf{Y}^\dagger)^{-1}\mathbf{S}(\mathbf{Y}\mathbf{F})^{-1}[(\mathbf{F}\mathbf{Y}^\dagger)^{-1}\mathbf{S}(\mathbf{Y}\mathbf{F})^{-1} + \mathbf{D}]^{-1}\tilde{\mathbf{d}} \\ &= (\mathbf{F}\mathbf{Y}^\dagger)^{-1}\mathbf{S}(\mathbf{S} + \mathbf{F}\mathbf{Y}^\dagger\mathbf{D}\mathbf{Y}\mathbf{F})^{-1}\mathbf{Y}^{-1}\mathbf{d} \\ &= (\mathbf{Y}\mathbf{F}\mathbf{Y}^\dagger)^{-1}s_{\text{WF}}, \end{aligned} \quad (19)$$

after plugging in the effective data vector given by equation (18).

This problem therefore reduces to one that can be solved directly using the dual messenger algorithm, requiring the same computational time as in the white noise case. The only additional steps required are a simple pre-whitening, followed by a post-smoothing operation with \mathbf{F}^{-1} and \mathbf{F} , respectively. For this particular case, we do not demonstrate the application of DANTE, as the implementation is straightforward.

3.2 Modulated correlated noise

The second generalization of the noise model corresponds to modulating the amplitude of spatially correlated noise. This is a more realistic noise model, typical of CMB experiments such as Planck, resulting from the scanning strategy of the instrument. The noise covariance, in pixel space, now takes the following form:

$$\mathbf{N} = \mathbf{D}\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger\mathbf{D}, \quad (20)$$

where \mathbf{C} is the isotropic noise covariance, encoding the inverse frequency ($1/f$) noise correlation on the large scales, typically associated with atmospheric noise, and therefore diagonal in harmonic space. The modulation, described by \mathbf{D} , is sparse in pixel space. The power spectrum C_ℓ of the non-modulated part of the noise can be expressed as

$$C_\ell = \frac{\sigma_N^2}{\beta} \left[1 + \left(\frac{\ell_{\text{knee}}}{\ell} \right)^{\alpha_{\text{knee}}} \right], \quad (21)$$

with the characteristic scale ℓ_{knee} and tilt α_{knee} of the power-law component, with σ_N being the noise amplitude per pixel.

The modulation takes into account the variation of noise amplitude due to the amount of integration time spent in any single pixel. The isotropic noise covariance in the map is expected to be a good approximation for scan strategies that cross each pixel in many directions, thereby isotropizing the way the time-ordered data noise correlations project on to the sky. But even in cases where the distribution of scan directions per pixel is not entirely isotropic, such as for the Planck data, this noise model was found to be of sufficient quality to derive the optimal primordial non-Gaussianity estimators (Planck Collaboration XVII 2016c; Planck Collaboration IX 2019).

As a result of the above dense noise covariance, equation (6) is no longer algebraically solvable due to the required inversion of a fully dense system. But since we are free to choose the covariance \mathbf{T} of the messenger field \mathbf{t} , we set $\mathbf{T} = \mathbf{D}(\mathbf{Y}\phi\mathbf{Y}^\dagger)\mathbf{D}$, where $\phi =$

$\min(\text{diag}(\mathbf{C}))$. This yields the following system of equations:

$$\begin{aligned} \mathbf{u} &= (\tilde{\mathbf{S}} + \mathbf{U}) [\mathbf{B}^\dagger\mathbf{Y}^\dagger\mathbf{D}^{-1}(\mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{Y}\mathbf{B}(\tilde{\mathbf{S}} + \mathbf{U}) + \mathbb{1}]^{-1} \\ &\quad \times \mathbf{B}^\dagger\mathbf{Y}^\dagger\mathbf{D}^{-1}(\mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{t} \end{aligned} \quad (22)$$

$$\begin{aligned} \mathbf{t} &= \mathbf{D} [(\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger - \mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1} + (\mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1}]^{-1} \\ &\quad \times [(\mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{Y}\mathbf{B}\mathbf{u} + (\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger - \mathbf{Y}\phi\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{d}], \end{aligned} \quad (23)$$

where the second equation can be simplified to the following form via straightforward linear algebraic manipulations:

$$\begin{aligned} \tilde{\mathbf{t}} \equiv \mathbf{D}^{-1}\mathbf{t} &= \mathbf{Y}(\mathbf{C} - \phi\mathbb{1})\mathbf{Y}^\dagger(\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{Y}\mathbf{B}\mathbf{u} \\ &\quad + (\mathbf{Y}\phi\mathbf{Y}^\dagger)(\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger)^{-1}\mathbf{D}^{-1}\mathbf{d}, \end{aligned} \quad (24)$$

which can now be solved trivially to obtain the messenger field. The first equation (22), however, cannot be solved directly, but it can be conveniently expanded using an extra messenger field \mathbf{v} , with covariance $\mathbf{V} = \omega(\mathbf{Y}\phi\mathbf{Y}^\dagger)\mathbb{1}$, where $\omega \equiv \min(\text{diag}(\mathbf{D}^2))$, yielding the following two trivially solvable equations:

$$\mathbf{v} = \omega\mathcal{M}\mathbf{D}^{-1}\mathcal{M}^{-1}\tilde{\mathbf{t}} + [\mathbb{1} - \omega\mathcal{M}\mathbf{D}^{-1}\mathcal{M}^{-1}\mathbf{D}^{-1}] \mathbf{Y}\mathbf{B}\mathbf{u} \quad (25)$$

$$\mathbf{u} = [\phi\omega(\tilde{\mathbf{S}} + \mathbf{U})^{-1} + \mathbf{B}^\dagger\mathbf{Y}^\dagger\mathcal{M}^{-1}\mathbf{Y}\mathbf{B}]^{-1} \mathbf{B}^\dagger\mathbf{Y}^\dagger\mathcal{M}^{-1}\mathbf{v}, \quad (26)$$

where the coupling matrix \mathcal{M} is defined as $\mathcal{M} \equiv \mathbf{Y}\mathbf{Y}^\dagger$. We therefore must solve the above system of three equations (24)–(26) when accounting for modulated correlated noise. Equation (26) can be written explicitly in terms of the relevant basis transformations as follows:

$$\begin{aligned} \mathbf{u} &= \mathfrak{R}^\dagger(\tilde{\mathbf{S}} + \mathbf{U}) [\mathfrak{R}\mathbf{B}^\dagger\mathbf{Y}^\dagger\mathcal{M}^{-1}\mathbf{Y}\mathbf{B}\mathfrak{R}^\dagger(\tilde{\mathbf{S}} + \mathbf{U}) + \phi\omega\mathbb{1}]^{-1} \\ &\quad \times \mathfrak{R}\mathbf{B}^\dagger\mathbf{Y}^\dagger\mathcal{M}^{-1}\mathbf{v}, \end{aligned} \quad (27)$$

where, as before, $\mathbf{S} = \mathfrak{R}^\dagger\mathcal{S}\mathfrak{R}$. An in-depth derivation of these equations is laid out in Appendix B. We encode the mask by doing the decomposition, $\mathbf{D} = \Sigma\tilde{\mathbf{D}}\Sigma = \Sigma\mathbf{P}^\dagger\mathbf{A}\mathbf{P}\Sigma$, as described in Section 2.4, and setting $\Sigma^{-1}|_{\text{mask}} = \mathbf{0}$. This prescription corresponds exactly to dropping the contribution of observations that are considered masked out, which may be deduced from equation (4). We apply the cooling scheme to $\xi = \mathbf{U} + \phi\omega\mathbb{1}$, as described in Section 5.3.

3.3 Nested Jacobi relaxation

As mentioned in Section 2.2, the above equations (24), (25), and (27) may be simplified significantly using the approximation $\mathbf{Y}^\dagger\mathbf{Y}\mathbb{1} \approx \beta\mathbb{1}$, thereby reducing the required number of SHTs. This approximation is not exact due to the coupling of the SHTs on the pixelized sky. In this work, we employ Jacobi relaxation to correct for the operations $\mathcal{M}^{-1} \equiv (\mathbf{Y}\mathbf{Y}^\dagger)^{-1}$ and $(\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger)^{-1}$.

We follow a similar rationale and employ the same notation as in Section 2.3, with the Jacobi iteration given by equation (10), where \mathbf{b} is any arbitrary vector in harmonic space. For the case of $\mathcal{A} = (\mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger)^{-1}$, the respective operations are as follows:

$$\tilde{\mathcal{A}} = (\beta^{-1}\mathbf{Y})\mathbf{C}^{-1}(\beta^{-1}\mathbf{Y}^\dagger), \quad \mathcal{A}^{-1} = \mathbf{Y}\mathbf{C}\mathbf{Y}^\dagger. \quad (28)$$

The correction for \mathcal{M}^{-1} is analogous to the above, with \mathbf{C} set to identity matrix, and must be embedded within the less trivial Jacobi relaxation for equation (27). The resulting nested relaxation scheme requires the following operations:

$$\mathcal{A} = \mathfrak{R}^\dagger(\tilde{\mathbf{S}} + \mathbf{U}) [\mathfrak{R}\mathbf{B}^\dagger\mathbf{Y}^\dagger\mathcal{M}^{-1}\mathbf{Y}\mathbf{B}\mathfrak{R}^\dagger(\tilde{\mathbf{S}} + \mathbf{U}) + \phi\omega\mathbb{1}]^{-1} \mathfrak{R} \quad (29)$$

$$\tilde{\mathcal{A}} = \mathfrak{R}^\dagger (\bar{\mathcal{S}} + \mathbf{U}) [\mathfrak{R} \mathbf{B}^\dagger \mathbf{B} \mathfrak{R}^\dagger (\bar{\mathcal{S}} + \mathbf{U}) + \phi \omega \mathbb{1}]^{-1} \mathfrak{R}, \quad (30)$$

and the corresponding inverse of operator \mathcal{A} given by

$$\mathcal{A}^{-1} = \mathfrak{R}^\dagger [\mathfrak{R} \mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{M}^{-1} \mathcal{Y} \mathbf{B} \mathfrak{R}^\dagger (\bar{\mathcal{S}} + \mathbf{U}) + \phi \omega \mathbb{1}] (\bar{\mathcal{S}} + \mathbf{U})^{-1} \mathfrak{R}, \quad (31)$$

with the basis operations included, and $\mathbf{b} = \mathbf{B}^\dagger \mathcal{Y}^\dagger \mathcal{M}^{-1} \mathbf{v}$.

4 PURE \mathcal{E}/\mathcal{B} DECOMPOSITION VIA WIENER FILTERING

In a finite patch of sky, the polarization field cannot be uniquely decomposed into pure \mathcal{E} and pure \mathcal{B} modes. Nevertheless, the polarization map can be uniquely decomposed into three distinct components, commonly referred to as the ‘pure \mathcal{E} ’, ‘pure \mathcal{B} ’, and ‘ambiguous’ modes (Lewis et al. 2002; Bunn et al. 2003). This new set of ambiguous modes receives contributions from both \mathcal{E} and \mathcal{B} modes. In such a framework, the signal vector space is divided into three orthogonal subspaces. The pure \mathcal{B} modes exist on the vector subspace orthogonal to that of all \mathcal{E} modes, and similarly for the pure \mathcal{E} modes. The ambiguous component, however, lies in the subspace orthogonal to both pure \mathcal{E} and \mathcal{B} subspaces. This decomposition ensures that a reconstructed pure \mathcal{B} map is not contaminated by \mathcal{E} modes.

\mathcal{E}/\mathcal{B} separation methods based on this pure-ambiguous decomposition originally involved the construction of an eigenbasis for the various orthonormal subspaces, but this is a tedious and numerically slow procedure. The \mathcal{E}/\mathcal{B} decomposition is trivial for exact methods such as Gibbs sampling (Larson et al. 2007), which infer the posterior statistics of a full-sky signal conditional on the data. Gibbs sampling requires a complete sky sample, i.e. optimally filtered data augmented to cover the whole sky via constrained generalizations (see also KLW18). This is the basis of the motivation behind the Wiener filtering approach proposed by Bunn & Wandelt (2017).

We briefly review the rationale and the formalism behind this new method, and describe how it can be incorporated in DANTE. A more comprehensive description is provided in Bunn & Wandelt (2017).

4.1 Background and notation

Considering only polarization measurements, the data set can be described as a $2N_{\text{pix}}$ dimensional vector of the Stokes parameters \mathcal{Q} and \mathcal{U} , i.e. $\mathbf{d} = (\mathbf{d}_{\mathcal{Q}}, \mathbf{d}_{\mathcal{U}})$, where N_{pix} corresponds to the dimension of the pixelized map of a given Stokes parameter. We account for the contribution from the temperature anisotropy, Stokes \mathcal{I} , in a future section. We assume a data model as given by equation (1), and Gaussian white noise, although the results presented below would still be relevant for more complex noise covariance.

The signal can be expressed as a SH expansion,

$$\mathbf{s} = \mathbf{s}^{\mathcal{E}} + \mathbf{s}^{\mathcal{B}} = \sum_{\ell, m} [a_{\ell m}^{\mathcal{E}} Y_{z, \ell m}^{\mathcal{E}}(\hat{r}_j) + a_{\ell m}^{\mathcal{B}} Y_{z, \ell m}^{\mathcal{B}}(\hat{r}_j)], \quad (32)$$

where \hat{r}_j labels the pixel corresponding to measurement j , while the index $z \in \{\mathcal{Q}, \mathcal{U}\}$ denotes the associated Stokes parameter. The functions Y can be expressed in terms of spin-2 SHs:

$$Y_{\mathcal{Q}, \ell m}^{\mathcal{E}} = Y_{\mathcal{U}, \ell m}^{\mathcal{B}} = -\frac{1}{2}(2Y_{\ell m} + {}_{-2}Y_{\ell m}) \quad (33)$$

$$Y_{\mathcal{Q}, \ell m}^{\mathcal{B}} = -Y_{\mathcal{U}, \ell m}^{\mathcal{E}} = -\frac{1}{2}(2Y_{\ell m} - {}_{-2}Y_{\ell m}). \quad (34)$$

The signal can therefore be written as

$$\mathbf{s} = \mathbf{Y}_{\mathcal{E}} \mathbf{e} + \mathbf{Y}_{\mathcal{B}} \mathbf{b}, \quad (35)$$

with the coefficients $a_{\ell m}^{\mathcal{E}}$ and $a_{\ell m}^{\mathcal{B}}$ encoded in the vectors \mathbf{e} and \mathbf{b} , respectively. The matrices \mathbf{Y}_Z , for $Z \in \{\mathcal{E}, \mathcal{B}\}$, consist of the SHs evaluated at the given pixel locations.

Under the assumption of data sourced by a statistically isotropic and random Gaussian process, the signal is uniquely described by the following covariance:

$$\mathbf{S} \equiv \langle \mathbf{s} \mathbf{s}^\dagger \rangle = \langle \mathbf{s}^{\mathcal{E}} \mathbf{s}^{\mathcal{E} \dagger} \rangle + \langle \mathbf{s}^{\mathcal{B}} \mathbf{s}^{\mathcal{B} \dagger} \rangle \equiv \mathbf{S}_{\mathcal{E}} + \mathbf{S}_{\mathcal{B}}. \quad (36)$$

For a full-sky data set, where \mathbf{d} covers the whole sky, the matrices $\mathbf{Y}_{\mathcal{E}}$ and $\mathbf{Y}_{\mathcal{B}}$ span orthogonal spaces, and hence the \mathcal{E} - \mathcal{B} coupling issue does not arise and the decomposition is straightforward. Incomplete sky coverage, however, results in the ambiguous modes that lie in both subspaces at the same time. The pure \mathcal{B} space can therefore be described as the orthogonal complement of the space spanned by $\mathbf{Y}_{\mathcal{E}}$, with an analogous definition for the pure \mathcal{E} space. The signal vector space can consequently be divided into three orthogonal subspaces, with the third ambiguous space being orthogonal to both of the pure subspaces. By projecting the data vector \mathbf{d} on the pure \mathcal{B} subspace, the \mathcal{E} -mode signal is mapped on to the null space of $\mathbf{Y}_{\mathcal{B}}$, ensuring no contamination of \mathcal{E} modes in the resulting pure \mathcal{B} map.

4.2 Purification framework

If the signal covariance \mathbf{S} employed in the Wiener filter equation (3) contains covariances of both \mathcal{E} and \mathcal{B} signals, then the resulting Wiener-filtered map would have contributions from both the scalar and pseudo-scalar components, i.e. $\mathbf{s}_{\text{WF}} = \mathbf{s}_{\text{WF}}^{\mathcal{E}} + \mathbf{s}_{\text{WF}}^{\mathcal{B}}$. Bunn & Wandelt (2017) proposed the following approach to isolate them from each other: Conceptually, the rationale is to treat one component as a source of noise. We can obtain the Wiener-filtered \mathcal{B} map, for instance, via the following replacements: $\mathbf{S} \rightarrow \mathbf{S}_{\mathcal{B}}$ and $\mathbf{N} \rightarrow \mathbf{S}_{\mathcal{E}} + \mathbf{N}$, such that the Wiener filter equation (3) becomes

$$\begin{aligned} \mathbf{s}_{\text{WF}}^{\mathcal{B}} &= [\mathbf{S}_{\mathcal{B}}^{-1} + (\mathbf{S}_{\mathcal{E}} + \mathbf{N})^{-1}]^{-1} (\mathbf{S}_{\mathcal{E}} + \mathbf{N})^{-1} \mathbf{d} \\ &= \mathbf{S}_{\mathcal{B}} [\mathbf{S}_{\mathcal{B}} + (\mathbf{S}_{\mathcal{E}} + \mathbf{N})]^{-1} \mathbf{d}, \end{aligned} \quad (37)$$

with an analogous expression for $\mathbf{s}_{\text{WF}}^{\mathcal{E}}$. Recall that $\mathbf{s}_{\text{WF}} = \mathbf{s}_{\text{WF}}^{\mathcal{E}} + \mathbf{s}_{\text{WF}}^{\mathcal{B}}$.

However, the above Wiener-filtered maps are contaminated by the ambiguous modes. Due to our prior signal covariance having a much higher \mathcal{E} -mode power, the theory assigns the ambiguous modes with high signal-to-noise mostly to the \mathcal{E} map. In order to ensure no cross-contamination, whereby a pure \mathcal{B} map should have contributions strictly from \mathcal{B} modes, Bunn & Wandelt (2017) suggested raising the signal covariance associated to the \mathcal{E} component to infinity, and provided a proof that this gives the same result as doing a costly eigenmode decomposition and projecting out the ambiguous modes. We define

$$\mathbf{S}(\lambda) = \mathbf{S}_{\mathcal{B}} + \lambda \mathbf{S}_{\mathcal{E}} \quad (38)$$

as the signal covariance with the \mathcal{E} -mode power amplified by a factor of λ . Substituting $\mathbf{S}_{\mathcal{B}} + \mathbf{S}_{\mathcal{E}} \rightarrow \mathbf{S}(\lambda)$ in equation (37) yields

$$\begin{aligned} \mathbf{s}_{\text{WF}}^{\mathcal{B}}(\lambda) &= \mathbf{S}_{\mathcal{B}} [\mathbf{S}(\lambda) + \mathbf{N}]^{-1} \mathbf{d} \\ &= \mathbf{S}_{\mathcal{B}} \mathbf{S}(\lambda)^{-1} [\mathbf{S}(\lambda)^{-1} + \mathbf{N}^{-1}]^{-1} \mathbf{N}^{-1} \mathbf{d}, \end{aligned} \quad (39)$$

such that in the limit $\lambda \rightarrow \infty$, only the pure \mathcal{B} modes survive in the null space of $\mathbf{S}_{\mathcal{E}}$. A strictly analogous procedure holds for the pure \mathcal{E} component.

4.3 Numerical implementation

To facilitate the numerical evaluation of the expressions above, it is more convenient to work with a full-sky data set, so that we can use fast transforms to move back and forth between the pixel and SH spaces. Masked pixels are assigned infinite noise covariance, i.e. we set the inverse noise covariance to zero.

Assuming isotropic and Gaussian CMB anisotropies, the signal covariance \mathbf{S} is diagonal in SH basis,

$$\mathbf{S}_{\mathcal{E}} = \text{diag}(C_2^{EE}, \dots, C_{\ell_{\max}}^{EE}, 0, \dots, 0) \quad (40)$$

$$\mathbf{S}_{\mathcal{B}} = \text{diag}(0, \dots, 0, C_2^{BB}, \dots, C_{\ell_{\max}}^{BB}), \quad (41)$$

with the ordering convention of having the \mathcal{E} -mode component first. Hence, we have

$$\mathbf{S}(\lambda)^{-1} = \text{diag}[(\lambda C_2^{EE})^{-1}, \dots, (\lambda C_{\ell_{\max}}^{EE})^{-1}, (C_2^{BB})^{-1}, \dots, (C_{\ell_{\max}}^{BB})^{-1}]. \quad (42)$$

As a result, $\mathbf{S}(\lambda)^{-1}|_{\lambda \rightarrow \infty} \rightarrow \mathbf{S}_{\mathcal{B}}^+$, where the pseudo-inverse $\mathbf{S}_{\mathcal{B}}^+$ is the inverse of $\mathbf{S}_{\mathcal{B}}$ within the subspace spanned by $\mathbf{Y}_{\mathcal{B}}$ and is zero in the orthogonal subspace of \mathcal{E} modes. We consequently obtain the operator that projects on to the pure \mathcal{E} subspace as $\mathbf{S}_{\mathcal{B}}\mathbf{S}_{\mathcal{B}}^+ = \mathbf{P}_{\mathcal{B}}$. Finally, we obtain the Wiener-filtered pure \mathcal{B} map as

$$\mathbf{s}_{\text{WF}}^{p\mathcal{B}} \equiv \lim_{\lambda \rightarrow \infty} \mathbf{s}_{\text{WF}}^{\mathcal{B}}(\lambda) = \mathbf{S}_{\mathcal{B}}\mathbf{S}_{\mathcal{B}}^+(\mathbf{S}_{\mathcal{B}}^+ + \mathbf{N}^{-1})^{-1}\mathbf{N}^{-1}\mathbf{d}. \quad (43)$$

The above formalism can be generalized to include the correlation between polarization and temperature anisotropies, where the data, $\mathbf{d} = (\mathbf{d}_{\mathcal{I}}, \mathbf{d}_{\mathcal{Q}}, \mathbf{d}_{\mathcal{U}})$, are pixelized maps of the Stokes parameters, \mathcal{I} , \mathcal{Q} and \mathcal{U} , where \mathcal{I} here corresponds to the temperature anisotropy. The signal covariance matrix now has a block-diagonal structure in SH space, with a 3×3 sub-matrix, for all multipole moments ℓ , as follows:

$$\mathbf{S}_{\ell} = \begin{pmatrix} C_{\ell}^{TT} & C_{\ell}^{TE} & 0 \\ C_{\ell}^{TE} & C_{\ell}^{EE} & 0 \\ 0 & 0 & C_{\ell}^{BB} \end{pmatrix}, \quad (44)$$

with the vanishing cross-spectra, C_{ℓ}^{TB} and C_{ℓ}^{EB} , set to zero.

To find the pure \mathcal{B} map, we proceed as before, i.e. $\mathbf{S}_{\mathcal{E}} \rightarrow \lambda \mathbf{S}_{\mathcal{E}}$ with $\lambda \rightarrow \infty$. In this limit, the temperature and polarization components decouple, yielding the following signal covariance:

$$\mathbf{S}_{\mathcal{B}} \equiv \left[\lim_{\lambda \rightarrow \infty} \mathbf{S}(\lambda)^{-1} \right]^+ = \begin{pmatrix} C_{\ell}^{TT} - \frac{(C_{\ell}^{TE})^2}{C_{\ell}^{EE}} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & C_{\ell}^{BB} \end{pmatrix}, \quad (45)$$

such that the above equation (43) for the Wiener-filtered pure \mathcal{B} map still holds.

A similar reasoning results in the following equation for the Wiener-filtered pure \mathcal{E} map:

$$\mathbf{s}_{\text{WF}}^{p\mathcal{E}} = \mathbf{S}_{\mathcal{E}}\mathbf{S}_{\mathcal{E}}^+(\mathbf{S}_{\mathcal{E}}^+ + \mathbf{N}^{-1})^{-1}\mathbf{N}^{-1}\mathbf{d}, \quad (46)$$

where $\mathbf{S}_{\mathcal{E}}$ is the signal covariance corresponding to

$$\mathbf{S}_{\mathcal{E}} \equiv \left[\lim_{\lambda \rightarrow \infty} \mathbf{S}(\lambda)^{-1} \right]^+ = \begin{pmatrix} 0 & 0 & 0 \\ 0 & C_{\ell}^{EE} - \frac{(C_{\ell}^{TE})^2}{C_{\ell}^{TT}} & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (47)$$

$\mathbf{S}_{\mathcal{E}}^+$ is then the pseudo-inverse associated to the subspace containing the temperature and \mathcal{E} modes, i.e. all the \mathcal{B} modes lie in the null space of both $\mathbf{S}_{\mathcal{E}}$ and $\mathbf{S}_{\mathcal{E}}^+$.

In this work, we encode this prescription in DANTE for optimal reconstruction of pure \mathcal{E} and \mathcal{B} maps via Wiener filtering. The numerical implementation for the pure \mathcal{B} case entails the usual Wiener filtering procedure, i.e. solving equation (3), but assuming infinite covariance for the \mathcal{E} component, followed by the application of the relevant projection operator $\mathbf{P}_{\mathcal{B}}$ to obtain the pure \mathcal{B} map. An analogous procedure yields the pure \mathcal{E} map.

The above formalism still holds in the presence of more complex noise models, such as the anisotropic correlated noise considered here. Since the signal covariance becomes fully diagonal when reconstructing the pure \mathcal{E} or \mathcal{B} map, we solve equation (26) itself since no basis transformations are required, as follows:

$$\mathbf{u} = (\bar{\mathbf{S}} + \mathbf{U})[\mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{Y}\mathbf{B}(\bar{\mathbf{S}} + \mathbf{U}) + \phi\omega\mathbb{1}]^{-1}\mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{v}. \quad (48)$$

There is, nevertheless, a caveat in the implementation of the above equation. The signal covariance \mathbf{S} , as given by equations (45) and (47), is actually the pseudo-inverse of the well-defined $\mathbf{S}(\lambda)^{-1}$ in the limit $\lambda \rightarrow \infty$. The trivial implementation is to set the relevant components of $\mathbf{S}_{\mathcal{E}}$ and $\mathbf{S}_{\mathcal{B}}$ to a numerically large value. Alternatively, equation (48) can be expressed in the following numerically convenient form:

$$\mathbf{u} = [\mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{Y}\mathbf{B} + \phi\omega\bar{\mathbf{S}}^{-1}(\bar{\mathbf{S}}^{-1} + \mathbf{U}^{-1})^{-1}\mathbf{U}^{-1}]^{-1} \cdot \mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{v}, \quad (49)$$

which, for the final step of the cooling scheme, i.e. $\mathbf{U} = \mathbf{0}$, reduces to

$$\mathbf{u} = [\mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{Y}\mathbf{B} + \phi\omega\mathbf{S}^{-1}]^{-1}\mathbf{B}^{\dagger}\mathbf{Y}^{\dagger}\mathcal{M}^{-1}\mathbf{v}. \quad (50)$$

We verified that both implementations resulted in identical solutions, within the limit of numerical errors.

5 NUMERICAL EXPERIMENTS

In this section, we demonstrate the application of DANTE to an artificially generated but realistic CMB polarization data set. We present the procedure for the mock generation, followed by a description of the different steps in the data analysis pipeline.

5.1 Mock generation

To simulate joint temperature and polarization maps on the sphere, we make use of HEALPY to generate realizations of $a_{\ell}^{\mathcal{I}}$, $a_{\ell}^{\mathcal{E}}$, and $a_{\ell}^{\mathcal{B}}$ signals with the correct covariance properties (cf. equation 44), taking into account the correlation between CMB temperature anisotropy and polarization. We employed CAMB³ (Lewis, Challinor & Lasenby 2000) to generate the input angular power spectra, C_{ℓ}^{TT} , C_{ℓ}^{EE} , C_{ℓ}^{BB} and C_{ℓ}^{TE} , from which the corresponding CMB signals are drawn. We assume a standard Λ CDM cosmology with the set of cosmological parameters ($\Omega_m = 0.32$, $\Omega_{\Lambda} = 0.69$, $\Omega_b = 0.05$, $h = 0.67$, $\sigma_8 = 0.83$, $n_s = 0.97$) from Planck (Planck Collaboration XIII 2016b). We can then construct the input \mathcal{Q} and \mathcal{U} maps by transforming realizations of \mathcal{E} and \mathcal{B} signals (cf. equation A2 in Appendix A) on the sphere, with HEALPIX resolution of $N_{\text{side}} = 128$ and $\ell_{\max} = 128$, such that the total number of pixels is $N_{\text{pix}} = 12 \times N_{\text{side}}^2 \approx 2 \times 10^5$. The input Stokes

³<http://camb.info>

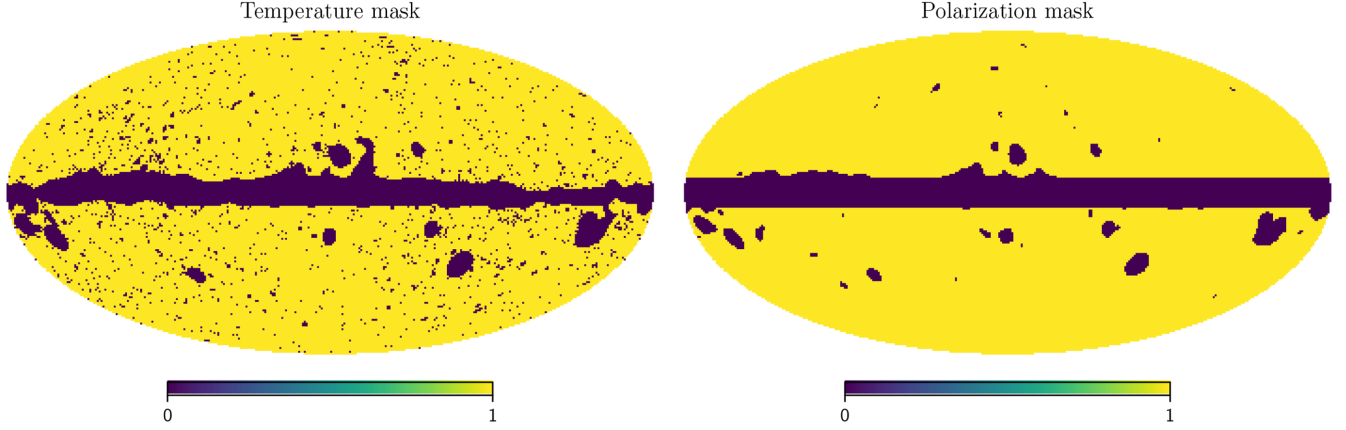


Figure 1. The temperature and polarization masks employed in the data analysis, corresponding to sky fractions of $f_{\text{sky}}^T = 0.76$ and $f_{\text{sky}}^P = 0.84$, respectively. Our artificially generated data set emulates the features of polarized Planck CMB maps.

parameters' maps are subsequently contaminated with modulated correlated noise, as described by equation (51) below, with a white noise amplitude of $\sigma_N = 40.0 \mu\text{K}$ per pixel, typical of high-sensitivity CMB experiments tailored for the detection of \mathcal{B} modes, and the corresponding $1/f$ noise parameters of $\ell_{\text{knee}} = 10$ and $\alpha_{\text{knee}} = 1.5$ (cf. equation 21). We employ the SMICA Planck temperature and polarization masks (Planck Collaboration IX 2016a), corresponding to sky fractions of $f_{\text{sky}}^T = 0.76$ and $f_{\text{sky}}^P = 0.84$, respectively, as depicted in Fig. 1. While our formalism and code account for the effect of a beam, we set the beam operator to identity for our present test cases.

5.2 Estimation of noise covariance

We now present a posterior optimization method to estimate the noise covariance using Monte Carlo (MC) simulations. For the case of modulated correlated noise, the data can be modelled as follows:

$$\mathbf{d} = \mathbf{s} + \mathbf{D}\mathcal{Y}\mathbf{C}^{1/2}\mathcal{Y}^\dagger\mathbf{n}, \quad (51)$$

following the notation of equation (1), where \mathbf{D} is diagonal in pixel space. We also use liberally the notation $\mathbf{C}^{1/2}$ to indicate the positive square root matrix of \mathbf{C} . As described in Section 3.2, the noise covariance is now given by

$$\mathbf{N} = \mathbf{D}\mathcal{Y}\mathbf{C}\mathcal{Y}^\dagger\mathbf{D}, \quad (52)$$

with \mathbf{C} being the isotropic, homogeneous, noise covariance, which incorporates the inverse frequency ($1/f$) noise correlation on the large scales. The overall aim is to estimate \mathbf{D} and \mathbf{C} using MC simulations by casting the covariance estimation problem as a two-level optimization scheme. The noise realizations can be modelled as the following linear combination:

$$\mathbf{n} = \mathbf{D}\mathcal{Y}\mathbf{C}^{1/2}\mathbf{z} + \mathbf{k}, \quad (53)$$

where \mathbf{z} and \mathbf{k} are Gaussian random fields with covariances, $\langle \mathbf{z}\mathbf{z}^\dagger \rangle = \mathbb{1}$ and $\langle \mathbf{k}\mathbf{k}^\dagger \rangle = \kappa^2\mathbb{1}$, respectively. The corresponding χ^2 , as the negative of the logarithm of the posterior distribution, with the sum over the contribution of each MC realization, can be written as

$$\chi^2 = \sum_{i=1}^{N_{\text{MC}}} \left[\frac{1}{\kappa^2} \left(\mathbf{n}_i - \mathbf{D}\mathcal{Y}\mathbf{C}^{1/2}\mathbf{z}_i \right)^\dagger \left(\mathbf{n}_i - \mathbf{D}\mathcal{Y}\mathbf{C}^{1/2}\mathbf{z}_i \right) + \mathbf{z}_i^\dagger \mathbf{z}_i \right]. \quad (54)$$

To obtain the maximum a posteriori estimate of \mathbf{D} and \mathbf{C} , we must optimize the above χ^2 with respect to \mathbf{z}_i and \mathbf{D} , in the limit $\kappa \rightarrow$

0. The χ^2 optimization with respect to \mathbf{z}_i yields, for a given MC simulation,

$$\tilde{\mathbf{z}}_i = \left(\mathcal{Y}^\dagger \mathbf{D}^2 \mathcal{Y} \mathbf{C}^{1/2} + \kappa^2 \mathbf{C}^{-1/2} \right)^{-1} \mathcal{Y}^\dagger \mathbf{D} \mathbf{n}_i, \quad (55)$$

which, in the limit $\kappa \rightarrow 0$, simplifies to

$$\begin{aligned} \lim_{\kappa \rightarrow 0} \tilde{\mathbf{z}}_i &= \left(\mathcal{Y}^\dagger \mathbf{D}^2 \mathcal{Y} \mathbf{C}^{1/2} \right)^{-1} \mathcal{Y}^\dagger \mathbf{D} \mathbf{n}_i \\ &= \mathbf{C}^{-1/2} \mathcal{Y}^{-1} \mathbf{D}^{-2} \tilde{\Pi}^\dagger \mathbf{D} \mathbf{n}_i, \end{aligned} \quad (56)$$

where $\tilde{\Pi}^\dagger = (\mathcal{Y}\mathcal{Y}^{-1})^\dagger$ is a projector on to the pixel subspace. The inversion per matrix is acceptable for the term in parenthesis because the operation \mathcal{Y}^\dagger already projects on the subspace of maps bandwidth limited to ℓ_{max} . Within that space the \mathcal{Y} operator becomes invertible, though at some cost. Optimizing the χ^2 with respect to \mathbf{D} leads to

$$\begin{aligned} \tilde{\mathbf{D}} &= \left[\sum_{i=1}^{N_{\text{MC}}} \left(\mathcal{Y} \mathbf{C}^{1/2} \mathbf{z}_i \right)^\dagger \left(\mathcal{Y} \mathbf{C}^{1/2} \mathbf{z}_i \right) \right]^{-1} \sum_{i=1}^{N_{\text{MC}}} \left(\mathcal{Y} \mathbf{C}^{1/2} \mathbf{z}_i \right)^\dagger \mathbf{n}_i \\ &\equiv \left(\sum_{i=1}^{N_{\text{MC}}} \mathbf{m}_i^\dagger \mathbf{m}_i \right)^{-1} \sum_{i=1}^{N_{\text{MC}}} \mathbf{m}_i^\dagger \mathbf{n}_i \\ &\equiv \mathcal{E}^{-1} \sum_{i=1}^{N_{\text{MC}}} \mathbf{m}_i^\dagger \mathbf{n}_i, \end{aligned} \quad (57)$$

where we defined $\mathcal{E} \equiv \sum_i \mathbf{m}_i^\dagger \mathbf{m}_i$, with \mathbf{m}_i estimated as follows:

$$\mathbf{m}_i = \mathcal{Y} \mathbf{C}^{1/2} \mathbf{z}_i = \tilde{\Pi} \mathbf{D}^{-2} \tilde{\Pi}^\dagger \mathbf{D} \mathbf{n}_i, \quad (58)$$

where $\tilde{\Pi} = \mathcal{Y}\mathcal{Y}^{-1}$ is the projector on to the SH space.

The algorithm for the noise covariance estimator proceeds according to the following iterative scheme: Compute $\tilde{\mathbf{m}}_i$ using equation (58), and subsequently \mathcal{E}^{-1} to obtain $\tilde{\mathbf{D}}$ using equation (57), followed by a power spectrum update to obtain $\tilde{\mathbf{C}}$ via $\tilde{\mathbf{C}} = \sum_i \langle \tilde{\mathbf{m}}_i \tilde{\mathbf{m}}_i^\dagger \rangle / N_{\text{MC}}$. In the above, we have defined the harmonic representation of the map with $\tilde{\mathbf{m}}_i = \mathcal{Y}^{-1} \mathbf{m}_i$. We solve equation (57) by implementing fixed point iterations, but this fixed point is not an attractor. We consequently employ the Babylonian method (e.g. Fowler & Robson 1998; Friberg 2007) to stabilize the fixed point and obtain an updated $\tilde{\mathbf{D}}$, as follows:

$$\tilde{\mathbf{D}} = \frac{1}{2} \left(\mathbf{D} + \mathcal{E}^{-1} \sum_{i=1}^{N_{\text{MC}}} \mathbf{m}_i^\dagger \mathbf{n}_i \right). \quad (59)$$

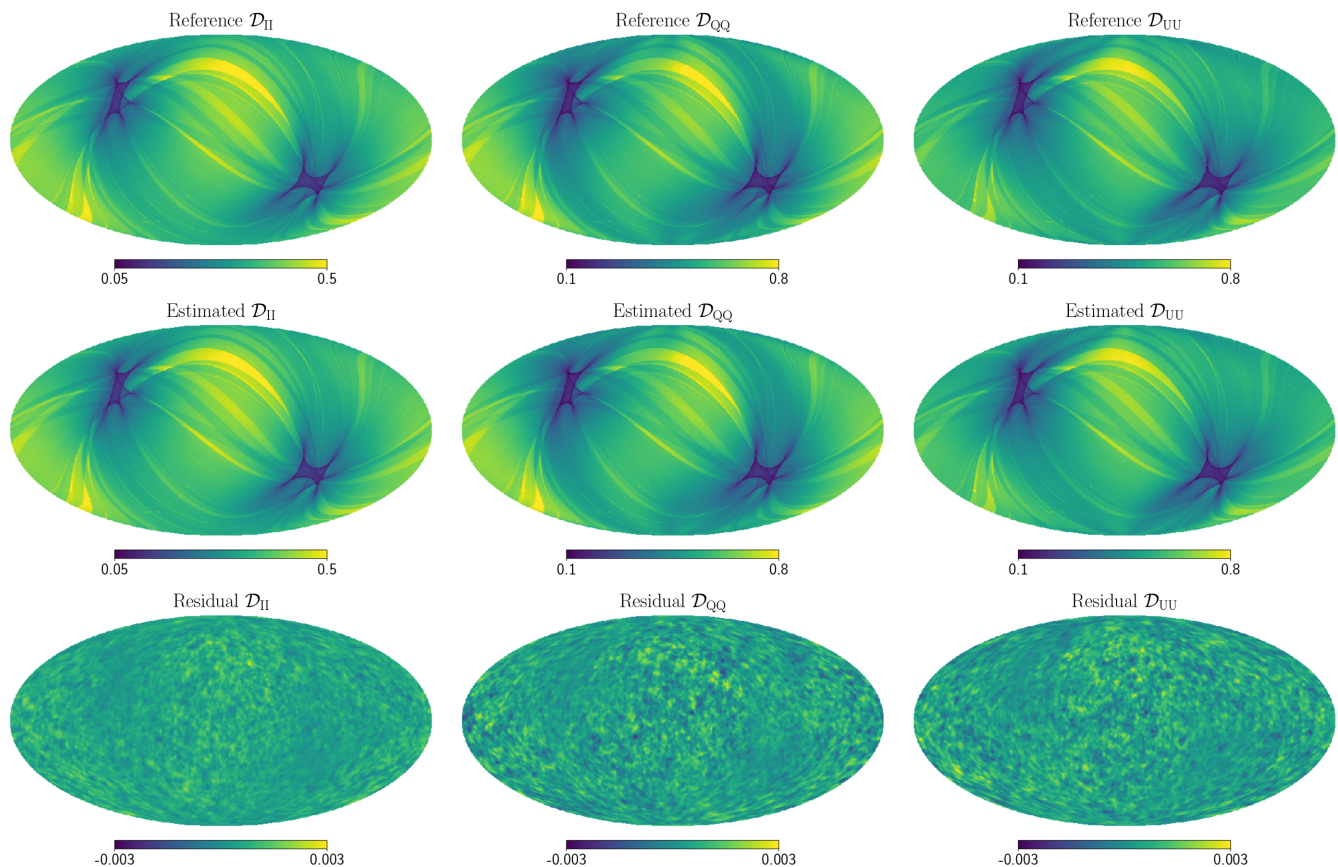


Figure 2. *Top row:* The reference modulated noise covariance used as inputs for the generation of the noise simulations. *Middle row:* The corresponding components recovered by the noise covariance estimator display the expected modulation patterns, indicating qualitatively the efficacy of our estimator. *Bottom row:* The residuals, generated by computing the difference between the reference and estimated noise covariance components, demonstrate the high-fidelity reconstructions. Note that for the relatively low residuals at the level of ~ 0.3 per cent to be visible, we employ a different colour bar scale for the residual maps.

We may verify that the fixed point of the above equation is exactly the desired \mathbf{D} matrix. Note that due to the degeneracy between the amplitudes of \mathbf{C} and \mathbf{D} , we need to anchor the amplitude of the updated $\tilde{\mathbf{C}}$ via the required re-scaling.

We therefore solve the above equation (57) iteratively, using $\mathbf{D}^2 = \sum_i (\mathbf{n}_i^\dagger \mathbf{n}_i) / N_{\text{MC}}$ as an initial guess. We generate 10^4 MC noise simulations using as template the modulated noise covariance provided by the Planck data analysis pipeline,⁴ as an estimate for the \mathbf{n}_i above. The map estimates, after only five iterations, for the diagonal and off-diagonal components of the noise covariance matrix, along with their corresponding reference and residual maps, are displayed in Figs 2 and 3, respectively. Visually, the distinct components of the covariance matrix are adequately recovered, with residuals at the level of ~ 0.3 per cent and ~ 6 per cent for the diagonal and off-diagonal components, respectively. As a quantitative diagnostic, we verify the relative deviation in the angular power spectra reconstructed from the maps, as a function of scale, with respect to their reference components: $\sqrt{C_\ell(\tilde{\mathbf{D}} - \mathbf{D}_{\text{ref}})} / C_\ell(\mathbf{D}_{\text{ref}})$, illustrated in Fig. 4. This demonstrates the accuracy of reconstruction of our noise covariance estimator across the range of scales considered, with only five iterations.

5.3 Polarization analysis

In this section, we showcase the application of DANTE in Wiener filtering polarized CMB maps contaminated with anisotropic correlated noise, and also illustrate its efficacy in generating pure \mathcal{E} and \mathcal{B} maps, guaranteed to be free from any cross-contamination. This corresponds to three distinct runs using the same realization of mock data, generated as described above in Section 5.1, labelled as ‘WF’, ‘pure \mathcal{E} ’, and ‘pure \mathcal{B} ’, respectively. The ‘pure \mathcal{B} ’ run yields a ‘pure’ temperature map as by-product, which corresponds to the map of temperature anisotropies without any contribution from the \mathcal{E} modes. We anchor the choice of hyperparameter values, described below, for all three cases.

We make an initial truncation in the power spectrum at $\ell = 50$, corresponding to a given value of ξ and the algorithm loops through the iterations until the fractional difference between successive iterations has reached a sufficiently low value, at which point ξ is reduced by a constant factor according to a given cooling scheme: $\xi \rightarrow \xi \eta$, where $\eta = 2/3$, until $\xi \rightarrow \phi \omega \mathbb{1}$. We implement a ‘weak’ criterion for convergence, $\|s_{i+1} - s_i\| / \|s_i\| < \epsilon$, where $\epsilon = 10^{-5}$, as a cheap proxy for the strong criterion that is verified a posteriori in Fig. 7.

The reconstructed angular power spectra for the temperature and polarization components are provided in Fig. 5, with the WF solution showing suppressed power on the small scales resulting from the noise and masked regions of the sky, which is

⁴Available from http://pla.esac.esa.int/pla/aio/product-action?MAP.MAP_ID=HFI_SkyMap_100_2048_R2.02_full.fits.

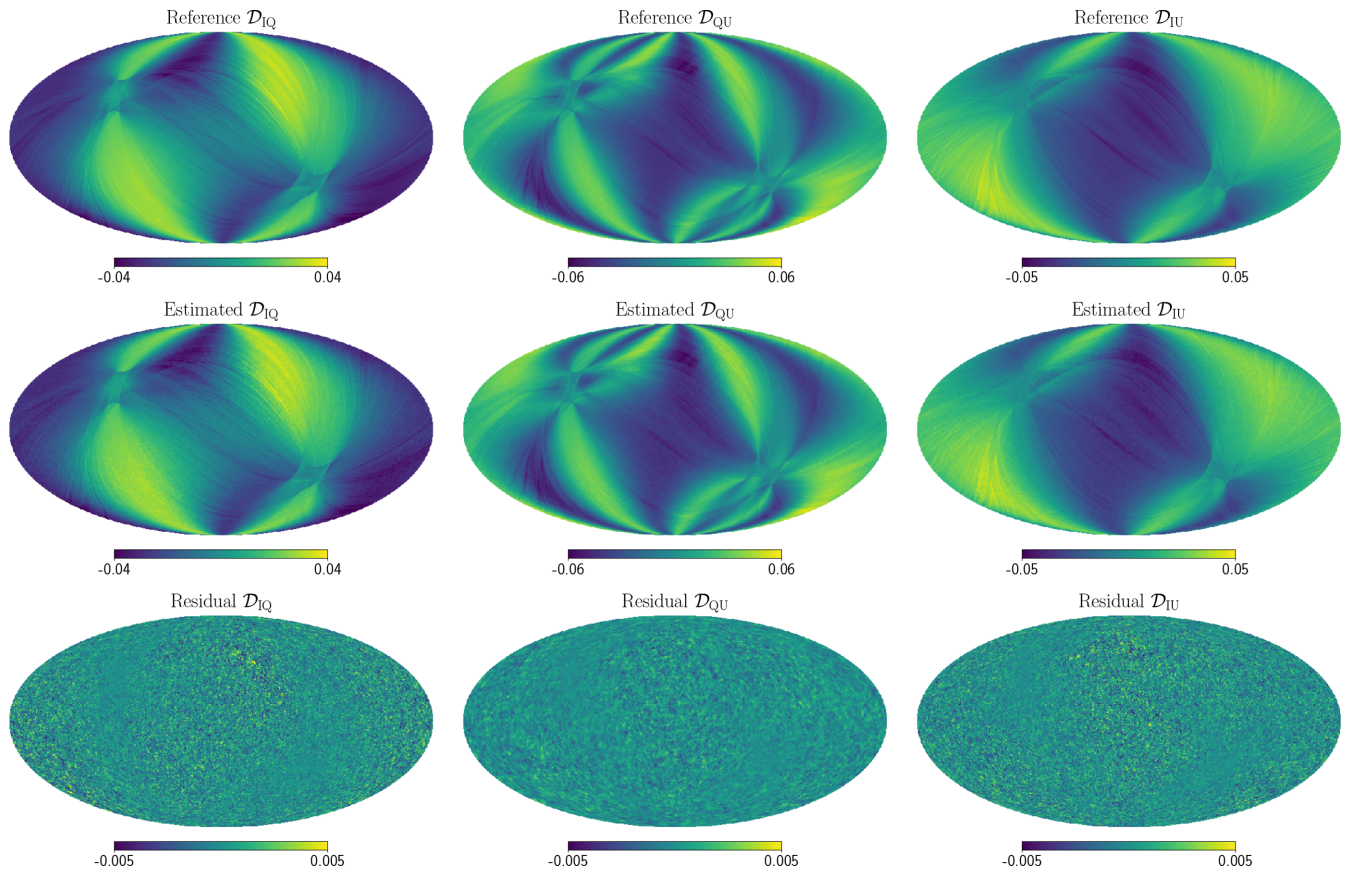


Figure 3. Same as Fig. 2, except for the off-diagonal components of the modulated noise covariance. As for their diagonal counterparts, the cross-correlations components are adequately recovered, with relatively insignificant residuals (~ 6 per cent).

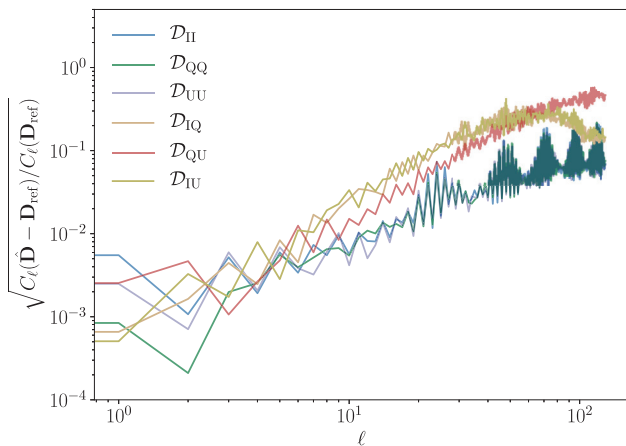


Figure 4. Relative deviation, as a function of scale, of the estimated components of the covariance matrix, with respect to their reference values. This diagnostic quantitatively demonstrates the performance of our noise covariance estimator with only five iterations.

a characteristic feature of Wiener filtering. For the temperature anisotropies, depicted in the left-hand panel, the pure \mathcal{B} run yields the temperature power spectra that has been purified with respect to the \mathcal{E} modes, and as such, corresponds to the prediction solely from the temperature data, with no contribution from the polarization

component. This pure temperature power spectrum does not display any significant difference compared to the Wiener-filtered one, as expected, due to the relatively small \mathcal{E} -mode contribution. The corresponding reconstructed power spectra for the \mathcal{E} modes are depicted in the middle panel of Fig. 5. The pure \mathcal{E} -mode power spectrum displays a smooth functional behaviour that matches the shape of the input power spectrum, although substantially suppressed because of the noisy and masked regions and since the discarded temperature contribution is significant. The right-hand panel displays the corresponding reconstructions for the \mathcal{B} -mode power spectrum. The pure reconstruction, as in the WF case, suppresses the modes in the low signal-to-noise regime, while also discarding the ambiguous modes, and shows a notable difference on the largest scales.

The real-space maps of the Wiener-filtered, pure \mathcal{E} and \mathcal{B} modes, together with their corresponding simulated maps, are illustrated in the middle and bottom rows of Fig. 6, respectively. The corresponding temperature maps are also displayed in the top row, for completeness. Note that the simulated maps are the generated signal realizations that were subsequently contaminated with anisotropic correlated noise and masked according to Fig. 1. The Wiener-filtered map, as the maximum *a posteriori* reconstruction, represents the CMB signal content of the data, with the reconstruction of the large-scale modes in the masked areas, based on the information content of the observed sky regions, being a natural consequence of Wiener filtering. Both the Wiener-filtered and pure temperature maps are similar in appearance, as expected from their reconstructed

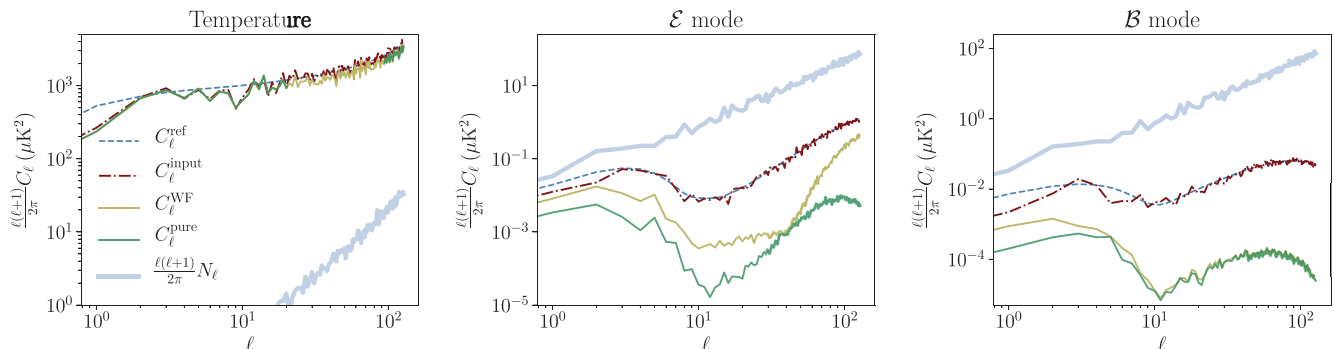


Figure 5. Reconstructed temperature, \mathcal{E} - and \mathcal{B} -mode angular power spectra from the WF and pure \mathcal{E}/\mathcal{B} runs. The simulated realizations, depicted using dash–dotted lines, were drawn from the reference power spectra, denoted by dashed lines, and were subsequently contaminated with anisotropic correlated noise and masked. For indicative purpose, we also show the power spectra of the respective noise realizations for our specific test case, taking into account the isotropic and modulated components of the noise covariance, using a thick solid line. We note that these power spectra correspond to a non-trivial averaging of the non-modulated part of our noise model specified in equation (21). *Left-hand panel:* The Wiener-filtered C_ℓ^{TT} is slightly suppressed on the small scales, as expected, due to the noise and masked regions. With the discarded \mathcal{E} -mode contribution being relatively low, the pure C_ℓ^{TT} matches the Wiener-filtered version. *Middle panel:* The pure C_ℓ^{EE} is substantially different from its Wiener-filtered counterpart, as the contribution of temperature anisotropies, by virtue of their larger power, is especially significant. *Right-hand panel:* The contrast between the Wiener-filtered and pure C_ℓ^{BB} is more significant on the largest scales, as can be seen from their real-space maps displayed in the bottom row of Fig. 6.

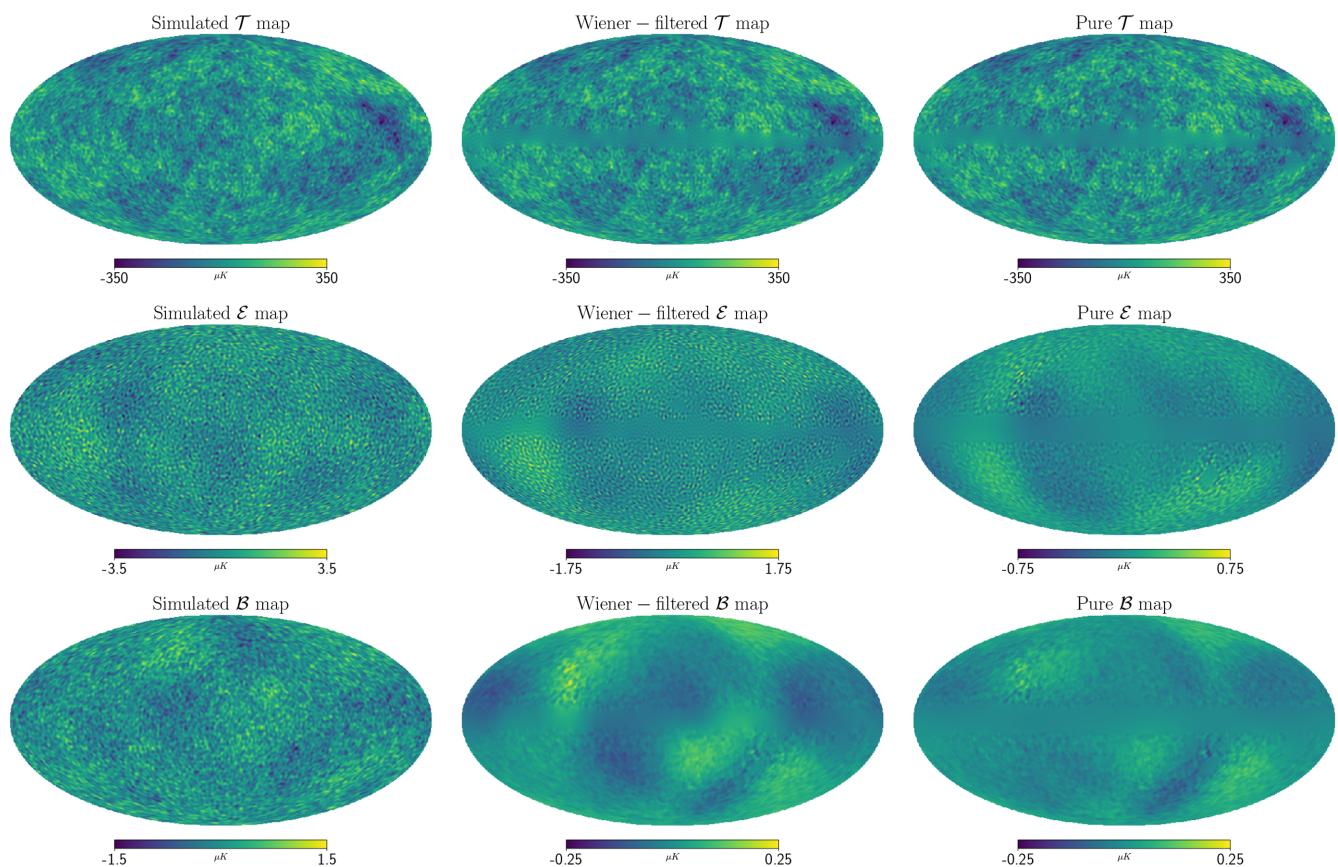


Figure 6. Simulated and reconstructed real-space maps of temperature anisotropies, \mathcal{E} and \mathcal{B} modes, from top to bottom, from the WF and pure \mathcal{E}/\mathcal{B} runs. The Wiener-filtered maps for all three components exhibit the characteristic feature, whereby the signal is extrapolated into the masked regions. The level of anisotropic correlated noise smooths out the small-scale features for the \mathcal{E} and \mathcal{B} modes, with the suppression of the small-scale power being more significant for the latter due to its low amplitude. The pure \mathcal{E} and \mathcal{B} maps, after the removal of ambiguous modes, display a reduced signal content. This difference is more striking for the \mathcal{E} map since the contribution from the temperature and \mathcal{E} -mode correlations is also discarded. The pure \mathcal{B} map has lower power close to the masked regions, relative to the Wiener-filtered one, as expected, since ambiguous modes are known to have support primarily near the mask.

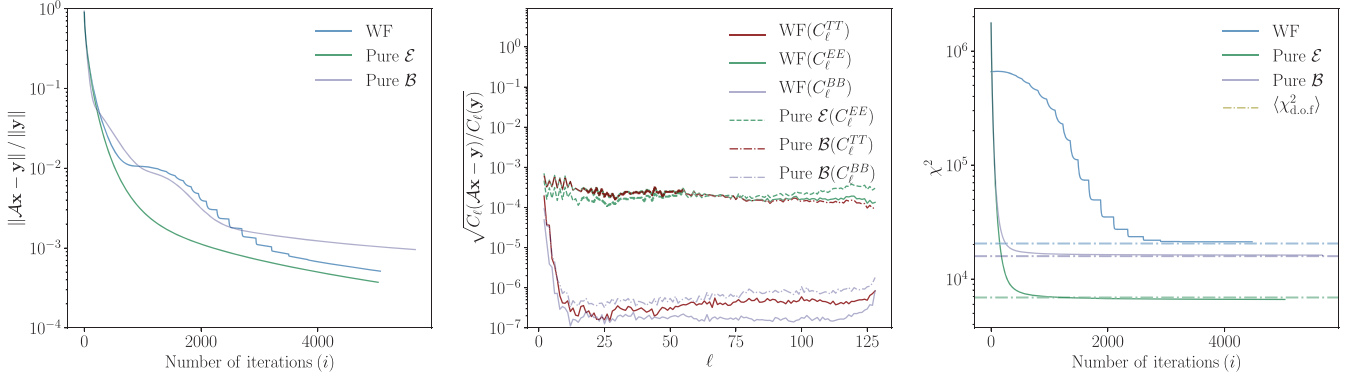


Figure 7. Convergence diagnostics for the reconstructions from the three runs. *Left-hand panel:* Variation of residual error, given by $\|\mathcal{A}\mathbf{x} - \mathbf{y}\| / \|\mathbf{y}\|$, as a function of iterations. This residual error is reduced monotonically as the iterations proceed, demonstrating the unconditional stability of the dual messenger algorithm in performing the three reconstructions. *Middle panel:* Variation of residual error as a function of angular scale for the final solutions. This error is sufficiently low for the range of scales considered, indicating the quality of the respective solutions. *Right-hand panel:* Variation of χ^2 with number of iterations for the WF and pure \mathcal{E}/\mathcal{B} solutions. Their respective χ^2 drop to a final value which is consistent with $\langle\chi^2_{\text{d.o.f.}}\rangle$ corresponding to the expectation value of the χ^2 , given by the number of degrees of freedom (d.o.f.), for the final solutions.

power spectra (cf. the left-hand panel of Fig. 5). This is not the case, however, for the reconstructed \mathcal{E} maps, purified with respect to the temperature anisotropies and ambiguous modes. The pure \mathcal{E} map, as a result, has a lower signal amplitude. Both the Wiener-filtered and pure \mathcal{B} maps show the clear reconstruction of the large-scale modes with extremely low amplitude. The striking contrast between the two maps is on the largest scales near the mask, where the pure \mathcal{B} map has lower power. This is consistent with previous work (e.g. Bunn & Wandelt 2017), where it was found that ambiguous modes have support primarily near the masked regions.

We illustrate the convergence behaviour of the three solutions via the corresponding variations of their residual error given by $\|\mathcal{A}\mathbf{x} - \mathbf{y}\| / \|\mathbf{y}\|$, for a linear system of equations given by $\mathcal{A}\mathbf{x} = \mathbf{y}$, in Fig. 7. This residual error adequately characterizes the accuracy of the final solution, with the relevant equations as follows:

$$\mathbf{y} = \mathcal{S}^{1/2} \mathfrak{R} \mathbf{y}^\dagger \mathbf{D}^{-1} (\mathcal{Y} \mathcal{C} \mathcal{Y}^\dagger)^{-1} \mathbf{D}^{-1} \mathbf{d} \quad (60)$$

and

$$\mathcal{A} = \mathbb{1} + \mathcal{S}^{1/2} \mathfrak{R} \mathbf{y}^\dagger \mathbf{D}^{-1} (\mathcal{Y} \mathcal{C} \mathcal{Y}^\dagger)^{-1} \mathbf{D}^{-1} \mathcal{Y} \mathfrak{R}^\dagger \mathcal{S}^{1/2}, \quad (61)$$

with $\mathbf{x} = \mathcal{S}^{-1/2} \mathfrak{R} \mathbf{s}$, following the notation from Section 3.2. The coupling matrix $(\mathcal{Y} \mathcal{C} \mathcal{Y}^\dagger)^{-1}$ requires Jacobi iterations for accurate evaluation of the above residual error. For the pure \mathcal{E}/\mathcal{B} runs, this is non-trivial as the signal covariance \mathbf{S} should have infinite values, as mentioned at the end of Section 4.3. To simplify the residual error evaluations in these cases, we simply set the relevant components of \mathbf{S} to a numerically large value.

As we demonstrated in our previous work, a characteristic feature of the dual messenger algorithm is the monotonic decrease in the residual error as the iterations proceed, as illustrated in the left-hand panel of Fig. 7, thereby demonstrating the unconditional stability of our method. This residual error, as a function of angular scale, for the final solutions from the three runs, are depicted in the middle panel. The Jacobi relaxation schemes employed is required to reduce this error to extremely low values across the range of scales considered.

The χ^2 is computed as follows:

$$\chi^2 = (\mathbf{d} - \mathcal{Y} \mathbf{B} \mathbf{s})^\dagger \mathbf{D}^{-1} (\mathcal{Y} \mathcal{C} \mathcal{Y}^\dagger)^{-1} \mathbf{D}^{-1} (\mathbf{d} - \mathcal{Y} \mathbf{B} \mathbf{s}) + \mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s}, \quad (62)$$

where, as for the residual error evaluations above, we employ Jacobi relaxation for the composite operation $(\mathcal{Y} \mathcal{C} \mathcal{Y}^\dagger)^{-1}$. The

corresponding χ^2 variation for the three different solutions is displayed in the right-hand panel of Fig. 7. In all three cases, the respective χ^2 of the dual messenger solutions drop to a final value which matches $\langle\chi^2_{\text{d.o.f.}}\rangle$, the expectation value of the χ^2 , given by the number of degrees of freedom (d.o.f), for the final solution. In the absence of masks, $\langle\chi^2_{\text{d.o.f.}}\rangle$ is given by the total number of harmonic modes of the temperature, \mathcal{E} and \mathcal{B} components. The computation of $\langle\chi^2_{\text{d.o.f.}}\rangle$ is, however, non-trivial when masks are involved. We estimated $\langle\chi^2_{\text{d.o.f.}}\rangle$ via Monte Carlo simulations. The convergence diagnostics discussed above, therefore, quantitatively demonstrate the efficacy of DANTE in performing the three distinct tasks.

Concerning the execution times for the WF, pure \mathcal{E} and pure \mathcal{B} runs, for the specific test case investigated, the algorithm runs to completion on four cores of a standard workstation, Intel Core i5-4690 CPU (3.50 GHz), in around 3 h. Note that a conjugate gradient method can, in principle, deal with such anisotropic noise models, provided that an adequate pre-conditioner can be found and this is the major stumbling block. Devising an appropriate pre-conditioner for such a complex problem is an extremely challenging task. For instance, the multi-grid pre-conditioner developed by Smith et al. (2007) at WMAP resolution and sensitivity is already highly non-trivial. The pre-conditioner-free approach of the dual messenger algorithm, therefore, is the key advantage.

6 CONCLUSIONS AND OUTLOOK

We present a numerically robust and fast code, DANTE, for pure \mathcal{E}/\mathcal{B} decomposition of CMB polarization maps. It accounts for complex and realistic noise models such as anisotropic correlated noise, encountered in typical CMB experiments such as Planck. DANTE is an augmented version of our dual messenger algorithm, adapted for the reconstruction of pure full-sky \mathcal{E} and \mathcal{B} maps on the sphere. The algorithm encodes a new method for the pure-ambiguous decomposition, based on a Wiener filtering approach, recently proposed by Bunn & Wandelt (2017), that guarantees no cross-contamination between the two maps. We also developed a noise covariance estimator to reconstruct the components of anisotropic noise covariance from Monte Carlo simulations, as required by the dual messenger algorithm.

We have demonstrated the capabilities of DANTE in dealing with large data sets and the associated high-dimensional covariance matrices. Moreover, as a pre-conditioner-free method, it is not hindered by ill-conditioned systems of equations inherent in CMB polarization problems, unlike standard PCG solvers, as demonstrated in [KLW18](#). DANTE also has an in-built option for drawing constrained Gaussian realizations of the CMB sky, for applications requiring homogeneous coverage of the field of observations. We have not illustrated this particular feature in this work as this was shown previously in [KLW18](#). DANTE will be rendered public in the near future.

The pure \mathcal{E}/\mathcal{B} decomposition framework implemented in this work, as a maximum a posteriori probability approach, has several advantages over traditional methods. It exploits the sparsity of the \mathcal{E}/\mathcal{B} decomposition in the SH basis, rendering the implementation extremely efficient. It is therefore much faster and straightforward than methods relying on the construction of orthonormal bases or wavelet methods that require a certain degree of fine-tuning. Moreover, \mathcal{E}/\mathcal{B} purification in the context of pseudo- C_l estimators is only feasible when the mask is differentiable up to at least its second derivatives, which is usually achieved via an appropriate apodization (Alonso et al. 2019). An interesting aspect of our approach is that it is not hindered by such limitations.

We have showcased the performance of DANTE on a realistic mock data set, emulating the features of polarized Planck CMB maps. The next step in this series of investigations is to further augment DANTE with an adaptive upgrade. Despite the improvements made to render the analysis of high-resolution CMB polarization data sets numerically feasible, the statistically optimal approach for the separation of \mathcal{E} and \mathcal{B} modes requires exact global analyses such as Gibbs sampling. This would, however, require several applications of the Wiener filter to obtain one signal realization conditional on the polarization data (e.g. Larson et al. 2007). The algorithm would therefore benefit from a further level of sophistication. A particularly interesting upgrade is to exploit the hierarchical framework of the dual messenger algorithm by adapting the working resolution progressively during execution, thereby substantially reducing the computation time. We also intend to explore the possibility of employing the dual messenger as a pre-conditioner in a standard PCG approach, in an attempt to drastically improve the convergence rate. Our algorithm could also be used to provide better examples of Wiener-filtered maps for a filtering based on machine learning, as in Münchmeyer & Smith (2019) with the J_1 loss function. These examples would be much more expensive than the purely simulation-based approach of the J_2 loss function, but would also provide a solid validation step.

Ultimately, the underlying objective is to employ this efficient tool in exact global Bayesian analyses of high-resolution and high-sensitivity CMB observations from the latest release of Planck to yield scientific products of significant value and interest. The resulting reconstructed maps may potentially be employed for various applications such as power-spectrum reconstruction, estimation of lensing potential, tests for foreground contamination and searches for non-Gaussianity and statistical anisotropy such as hemispherical power anisotropy. Another key aspect is that the features of the real-space pure \mathcal{B} maps allow the characterization of lensing-induced \mathcal{B} modes that go beyond the power spectrum.

DANTE can be easily applied to other CMB data sets in straightforward fashion without major modifications of the source code, making it a potentially powerful and robust tool for other current and future high-resolution CMB missions such as South Pole Telescope, Advanced ACTPol, Simons Observatory and CMB-S4.

The flexibility of the code can nevertheless be exploited in other cosmological contexts, due to the ubiquitous use of the Wiener filter, and even in more general scenarios involving spin field reconstruction.

ACKNOWLEDGEMENTS

We convey our appreciation to the anonymous reviewer for their remarks and suggestions that helped to improve the overall quality of the manuscript. We thank Franz Elsner and Emory F. Bunn for stimulating discussions. This work has been done within the activities of the Domaine d'Intérêt Majeur (DIM) Astrophysique et Conditions d'Apparition de la Vie (ACAV), and received financial support from Région Ile-de-France. We acknowledge financial support from the ILP LABEX, under reference ANR-10-LABX-63, which is financed by French state funds managed by the ANR within the Investissements d'Avenir programme under reference ANR-11-IDEX-0004-02. GL also acknowledges financial support from the ANR BIG4, under reference ANR-16-CE23-0002. BDW is supported by the Simons Foundation. This work is done within the Aquila Consortium.⁵

REFERENCES

- Abazajian K. N. et al., 2016, preprint ([arXiv:1610.02743](#))
 Alonso D., Sanchez J., Slosar A., LSST Dark Energy Science Collaboration, 2019, *MNRAS*, 484, 4127
 Alsing J., Heavens A., Jaffe A. H., Kiessling A., Wandelt B., Hoffmann T., 2016, *MNRAS*, 455, 4452
 Alsing J., Heavens A., Jaffe A. H., 2017, *MNRAS*, 466, 3272
 Anderes E., Wandelt B. D., Lavaux G., 2015, *ApJ*, 808, 152
 Bowyer J., Jaffe A. H., Novikov D. I., 2011, preprint ([arXiv:1101.0520](#))
 Bunn E. F., 2002a, *Phys. Rev. D*, 65, 043003
 Bunn E. F., 2002b, *Phys. Rev. D*, 66, 069902
 Bunn E. F., 2011, *Phys. Rev. D*, 83, 083003
 Bunn E. F., Wandelt B., 2017, *Phys. Rev. D*, 96, 043523(BW17)
 Bunn E. F., Zaldarriaga M., Tegmark M., de Oliveira-Costa A., 2003, *Phys. Rev. D*, 67, 023501
 Cao L., Fang L.-Z., 2009, *ApJ*, 706, 1545
 Challinor A., Chon G., 2005, *MNRAS*, 360, 509
 Elsner F., Wandelt B. D., 2013, *A&A*, 549, A111
 Eriksen H. K. et al., 2004, *ApJS*, 155, 227
 Ferté A., Grain J., Tristram M., Stomp R., 2013, *Phys. Rev. D*, 88, 023524
 Fowler D., Robson E., 1998, *Historia Math.*, 25, 366
 Friberg J., 2007, A Remarkable Collection of Babylonian Mathematical Texts: Manuscripts in the Schøyen Collection: Cuneiform Texts I. Springer Science & Business Media, New York
 Fu H., Wu B., Shao Y., Zhang H., 2018, *IEEE Photonics J.*, 10, 2837010
 Galli S. et al., 2014, *Phys. Rev. D*, 90, 063504
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Grain J., Tristram M., Stomp R., 2009, *Phys. Rev. D*, 79, 123515
 Guzzetti M. C., Bartolo N., Liguori M., Matarrese S., 2016, preprint ([arXiv:1605.01615](#))
 Hanson D. et al., 2013, *Phys. Rev. Lett.*, 111, 141301
 Henning J. W. et al., 2018, *ApJ*, 852, 97
 Hoffman Y., Ribak E., 1991, *ApJ*, 380, L5
 Horowitz B., Seljak U., Aslanyan G., 2018, preprint ([arXiv:1810.00503](#))
 Hu W., 2003, *Ann. Phys.*, 303, 203
 Hu W., Dodelson S., 2002, *ARA&A*, 40, 171
 Hu W., White M., 1997, *New A*, 2, 323
 Huffenberger K. M., 2018, *MNRAS*, 476, 3425

⁵<https://aquila-consortium.org>

Huffenberger K. M., Næss S. K., 2018, *ApJ*, 852, 92
 Huffenberger K. M., Wandelt B. D., 2010, *ApJS*, 189, 255
 Jacobi C. G., 1845, *Astron. Nachr.*, 22, 297
 Jasche J., Lavaux G., 2015, *MNRAS*, 447, 1204
 Kamionkowski M., Kovetz E. D., 2016, *ARA&A*, 54, 227
 Kamionkowski M., Kosowsky A., Stebbins A., 1997a, *Phys. Rev. D*, 55, 7368
 Kamionkowski M., Kosowsky A., Stebbins A., 1997b, *Phys. Rev. Lett.*, 78, 2058
 Kim J., 2011, *A&A*, 531, A32
 Kim J., Naselsky P., 2010, *A&A*, 519, A104
 Kinney W. H., 1998, *Phys. Rev. D*, 58, 123506
 Kodi Ramanah D., Lavaux G., Wandelt B. D., 2017, *MNRAS*, 468, 1782
 Kodi Ramanah D., Lavaux G., Wandelt B. D., 2018, *MNRAS*, 476, 2825(KLW18)
 Lam S. K., Pitrou A., Seibert S., 2015, in Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC. LLVM '15. ACM, New York, p. 7:1
 Larson D. L., Eriksen H. K., Wandelt B. D., Górski K. M., Huey G., Jewell J. B., O'Dwyer I. J., 2007, *ApJ*, 656, 653
 Lattner C., Adve V., 2004, in Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization, IEEE Computer Society, Washington, DC, USA. p. 75
 Leistedt B., McEwen J. D., Büttner M., Peiris H. V., 2017, *MNRAS*, 466, 3728
 Lewis A., 2003, *Phys. Rev. D*, 68, 083509
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Lewis A., Challinor A., Turok N., 2002, *Phys. Rev. D*, 65, 023505
 Louis T. et al., 2017, *J. Cosmol. Astropart. Phys.*, 6, 031
 McEwen J. D., Wiaux Y., 2011, *IEEE Trans. Signal Process.*, 59, 5876
 Muciaccia P. F., Natoli P., Vittorio N., 1997, *ApJ*, 488, L63
 Münchmeyer M., Smith K. M., 2019, preprint (arXiv:1905.05846)
 Papež J., Grigori L., Stompor R., 2018, *A&A*, 620, A59
 Planck Collaboration IX, 2016a, *A&A*, 594, A9
 Planck Collaboration XIII, 2016b, *A&A*, 594, A13
 Planck Collaboration XVII, 2016c, *A&A*, 594, A17
 Planck Collaboration IX, 2019, preprint (arXiv:1905.05697)
 POLARBEAR Collaboration et al., 2017, *ApJ*, 848, 121
 Puglisi G., Poletti D., Fabbian G., Baccigalupi C., Heltai L., Stompor R., 2018, *A&A*, 618, A62
 Rogers K. K., Peiris H. V., Leistedt B., McEwen J. D., Pontzen A., 2016, *MNRAS*, 463, 2310
 Saad Y., 2003, Iterative Methods for Sparse Linear Systems: Society for Industrial and Applied Mathematics, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA
 Seljak U., Zaldarriaga M., 1997, *Phys. Rev. Lett.*, 78, 2054
 Seljebotn D. S., Mardal K.-A., Jewell J. B., Eriksen H. K., Bull P., 2014, *ApJS*, 210, 24
 Seljebotn D. S., Bærland T., Eriksen H. K., Mardal K.-A., Wehus I. K., 2019, *A&A*, 627, A98
 Smith K. M., 2006a, *New A Rev.*, 50, 1025
 Smith K. M., 2006b, *Phys. Rev. D*, 74, 083002
 Smith K. M., Zaldarriaga M., 2007, *Phys. Rev. D*, 76, 043001
 Smith K. M., Zahn O., Doré O., 2007, *Phys. Rev. D*, 76, 043510
 Suzuki A. et al., 2016, *J. Low Temp. Phys.*, 184, 805
 Tegmark M., Eisenstein D. J., Hu W., de Oliveira-Costa A., 2000, *ApJ*, 530, 133
 Wandelt B. D., Larson D. L., Lakshminarayanan A., 2004, *Phys. Rev. D*, 70, 083511
 Wiener N., 1949, Extrapolation, Interpolation, and Smoothing of Stationary Time Series, Vol. 2. MIT press, Cambridge, MA
 Zaldarriaga M., 1997, *Phys. Rev. D*, 55, 1822
 Zaldarriaga M., 2001, *Phys. Rev. D*, 64, 103001
 Zaldarriaga M., Seljak U., 1997, *Phys. Rev. D*, 55, 1830
 Zaldarriaga M., Seljak U., 1998, *Phys. Rev. D*, 58, 023003
 Zaldarriaga M., Spergel D. N., Seljak U., 1997, *ApJ*, 488, 1
 Zhao W., Baskaran D., 2010, *Phys. Rev. D*, 82, 023001

APPENDIX A: SH TRANSFORMS

We provide a brief description of the transformation between pixel and SH domain in order to be precise about the notation employed in this work.

Assuming the primary CMB fluctuations to be an isotropic Gaussian random field, the CMB signal can be described as a vector of SH coefficients, with the associated signal covariance \mathbf{S} given by

$$S_{\ell m, \ell' m'} = \langle a_{\ell m} a_{\ell' m'} \rangle = \delta_{\ell \ell'} \delta_{m m'} C_{\ell}, \quad (\text{A1})$$

where C_{ℓ} is the CMB power spectrum. The proper basis to represent isotropic Gaussian random fields on the sphere is described by SHs. Given a grid on the sphere, i.e. a set of pixel positions \hat{n}_p , we can transform a field expressed in SH basis, with coefficients $s_{\ell m}$, to one sampled on the sphere via SH synthesis, as follows:

$$s(\hat{n}_p) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{n}_p) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \mathcal{Y}_p^{\ell m} a_{\ell m}. \quad (\text{A2})$$

Formally, the SH synthesis may be expressed as a matrix product, $s_{(p)} = \mathcal{Y} a_{(\ell m)}$, where \mathcal{Y} is the synthesis operator that encodes the value of the SHs evaluated at each \hat{n}_p of the grid.

Conversely, the transformation from pixel to harmonic basis is referred to as SH analysis, with the analysis operator being an integral related to the synthesis operator:

$$\begin{aligned} a_{(\ell m)} &\simeq \sum_p Y_{\ell m}^*(\hat{n}_p) s(\hat{n}_p) \delta\Omega_p = \mathcal{Y}^{-1} s_{(p)} \\ &= \sum_p \frac{4\pi}{N_{\text{pix}}} \mathcal{Y}_p^{\dagger, \ell m} s(\hat{n}_p) = \frac{4\pi}{N_{\text{pix}}} \mathcal{Y}^{\dagger} s_{(p)}. \end{aligned} \quad (\text{A3})$$

It is important to emphasize the scaling operation above and note that the last equalities are valid only for an equal-area pixelization such as HEALPIX (Górski et al. 2005). These SH transforms are the spherical analogue of Fourier transforms.

APPENDIX B: MODULATED CORRELATED NOISE COVARIANCE

We provide a more in-depth derivation of the two dual messenger equations (25) and (26) required for the treatment of modulated correlated noise covariance. The third equation (24) can be derived from equation (23) in straightforward fashion via linear algebraic simplifications.

Equation (5) can be written in its explicit form as

$$\mathbf{u} = [\mathbf{B}^{\dagger} \mathcal{Y}^{\dagger} \mathbf{T}^{-1} \mathcal{Y} \mathbf{B} + (\bar{\mathbf{S}} + \mathbf{U})^{-1}]^{-1} \mathbf{B}^{\dagger} \mathcal{Y}^{\dagger} \mathbf{T}^{-1} \mathbf{t}, \quad (\text{B1})$$

with the covariance of the messenger field \mathbf{t} being $\mathbf{T} = \mathbf{D}(\mathcal{Y} \phi \mathcal{Y}^{\dagger}) \mathbf{D}$. This equation bears a striking resemblance to the standard Wiener filter equation (3) and can therefore be solved via the introduction of an extra messenger field \mathbf{v} with covariance $\mathbf{V} = \omega(\mathcal{Y} \phi \mathcal{Y}^{\dagger}) \mathbb{1}$, where $\omega \equiv \min(\text{diag}(\mathbf{D}^2))$, resulting in the following χ^2 :

$$\begin{aligned} \chi_{\mathbf{v}}^2 &= (\mathbf{t} - \mathbf{v})^{\dagger} [\mathbf{D}(\mathcal{Y} \phi \mathcal{Y}^{\dagger}) \mathbf{D} - \mathbf{V}] (\mathbf{t} - \mathbf{v}) \\ &\quad + (\mathbf{v} - \mathcal{Y} \mathbf{b} \mathbf{u})^{\dagger} \mathbf{V}^{-1} (\mathbf{v} - \mathcal{Y} \mathbf{b} \mathbf{u}) + \mathbf{u}^{\dagger} (\bar{\mathbf{S}} + \mathbf{U})^{-1} \mathbf{u}. \end{aligned} \quad (\text{B2})$$

Minimizing the above χ^2 with respect to \mathbf{v} and \mathbf{u} yields the following set of equations:

$$\begin{aligned} \mathbf{v} &= \omega(\mathcal{Y} \phi \mathcal{Y}^{\dagger}) \mathbf{D}^{-1} (\mathcal{Y} \phi \mathcal{Y}^{\dagger})^{-1} \mathbf{D}^{-1} \\ &\quad \times [\mathbf{t} + \omega^{-1} \mathbf{D}(\mathcal{Y} \phi \mathcal{Y}^{\dagger}) \mathbf{D}(\mathcal{Y} \phi \mathcal{Y}^{\dagger})^{-1} - \mathbb{1}] \mathcal{Y} \mathbf{B} \mathbf{u} \end{aligned} \quad (\text{B3})$$

$$\mathbf{u} = [\phi \omega (\bar{\mathbf{S}} + \mathbf{U})^{-1} + \mathbf{B}^{\dagger} \mathbf{B}]^{-1} \mathbf{B}^{\dagger} \mathcal{Y}^{\dagger} \mathcal{M}^{-1} \mathbf{v}, \quad (\text{B4})$$

where, as before, we employ the definition of the coupling matrix $\mathcal{M} \equiv \mathcal{Y}\mathcal{Y}^\dagger$. It is more convenient to work with $\tilde{\mathbf{t}} \equiv \mathbf{D}^{-1}\mathbf{t}$, such that we can rewrite equation (B3) as

$$\mathbf{v} = \omega\mathcal{M}\mathbf{D}^{-1}\mathcal{M}^{-1}\tilde{\mathbf{t}} + [\mathbb{1} - \omega\mathcal{M}\mathbf{D}^{-1}\mathcal{M}^{-1}\mathbf{D}^{-1}] \mathcal{Y}\mathbf{B}\mathbf{u}. \quad (\text{B5})$$

The preference for the above form is that masked regions do not pose any numerical issue, as $\mathbf{D}^{-1}|_{\text{mask}} = \mathbf{0}$, such that $\mathbf{v}|_{\text{mask}} \rightarrow \mathcal{Y}\mathbf{B}\mathbf{u}$.

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.