



HAL
open science

MAPPRAISER: A massively parallel map-making framework for multi-kilo pixel CMB experiments

Hamza El Bouhargani, Aygul Jamal, Dominic Beck, Josquin Errard, Laura Grigori, Radek Stompor

► **To cite this version:**

Hamza El Bouhargani, Aygul Jamal, Dominic Beck, Josquin Errard, Laura Grigori, et al.. MAP-PRAISER: A massively parallel map-making framework for multi-kilo pixel CMB experiments. *Astron.Comput.*, 2022, 39, pp.100576. 10.1016/j.ascom.2022.100576 . hal-03508345

HAL Id: hal-03508345

<https://hal.science/hal-03508345>

Submitted on 22 Jul 2024

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

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



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

MAPPRAISER: A massively parallel map-making framework for multi-kilo pixel CMB experiments

Hamza El Bouhargani^{a,b}, Aygul Jamal^a, Dominic Beck^{c,d}, Josquin Errard^a, Laura Grigori^e, Radek Stompor^{a,b}

^aUniversité de Paris, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

^bCNRS-UCB International Research Laboratory, Centre Pierre Binétruy, IRL2007, CPB-IN2P3, Berkeley, CA 94720, USA

^cKavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Rd, Menlo Park, California 94025, USA

^dDepartment of Physics, Stanford University, Stanford, California 94305, USA

^eINRIA Paris, Sorbonne Université, Université Paris-Diderot SPC, CNRS, Laboratoire Jacques-Louis Lions, ALPINES Team, France

Abstract

Forthcoming cosmic microwave background (CMB) polarized anisotropy experiments have the potential to revolutionize our understanding of the Universe and fundamental physics. The sought-after, tale-telling signatures will be however distributed over voluminous data sets which these experiments will collect. These data sets will need to be efficiently processed and unwanted contributions due to astrophysical, environmental, and instrumental effects characterized and efficiently mitigated in order to uncover the signatures. This poses a significant challenge to data analysis methods, techniques, and software tools which will not only have to be able to cope with huge volumes of data but to do so with unprecedented precision driven by the demanding science goals posed for the new experiments.

A keystone of efficient CMB data analysis are solvers of very large linear systems of equations. Such systems appear in very diverse contexts throughout CMB data analysis pipelines, however they typically display similar algebraic structures and can therefore be solved using similar numerical techniques. Linear systems arising in the so-called map-making problem are one of the most prominent and common ones.

In this work we present a massively parallel, flexible and extensible framework, comprised of a numerical library, MIDAPACK, and a high level code, MAPPRAISER, which provide tools for solving efficiently such systems. The framework implements iterative solvers based on conjugate gradient techniques: enlarged and preconditioned using different preconditioners. We demonstrate the framework on simulated examples reflecting basic characteristics of the forthcoming data sets issued by ground-based and satellite-borne instruments, executing it on as many as 16,384 compute cores. The software is developed as an open source project freely available to the community at: <https://github.com/B3Dcmb/midapack>.

Keywords: Numerical methods, linear systems solvers, high performance computing, cosmic microwave background, data analysis, map-making

1. Introduction

Observations of cosmic microwave background (CMB) anisotropies have played a major role in shaping up our current view of the Universe, its composition and evolution. This has been epitomized by the results obtained by the Planck satellite [37, 40], which, with complementary information provided by other probes, have established the current standard cosmological model determining its parameters with a few percent precision. The scientific potential of the CMB has not been exhausted yet and only experiments forthcoming on the timescale of this decade are expected to exploit the full scientific information contained in the polarization properties of the CMB anisotropies. This is expected to open a new window on the

evolution and physics of the very early Universe providing unique insights about physical laws at extremely high energies exceeding by many orders of magnitude energies accessible by any current and projected human-made experiments. The potential impact of these new experiments is therefore difficult to overestimate.

Given the minute amplitudes of the sought-after signatures, data sets collected by these experiments and containing the cosmologically-relevant information will have to surpass in volume and complexity any of the CMB data sets available to date by orders of magnitude. CMB anisotropies constitute a stochastic, weakly correlated, sky-stationary signal, which, in the CMB data analysis process, has to be differentiated from other stochastic and usually correlated contributions. This typically disfavors any *divide-and-conquer* approaches and instead huge volumes of the CMB data have to be processed simultaneously. This sets demanding requirements on computer memory which can be only met by the largest massively-parallel, distributed-memory supercomputers. A key operation of this

Email addresses: elbouha@apc.in2p3.fr (Hamza El Bouhargani), ajamal@apc.in2p3.fr (Aygul Jamal), dobeck@stanford.edu (Dominic Beck), josquin@apc.in2p3.fr (Josquin Errard), Laura.Grigori@inria.fr (Laura Grigori), radek.stompor@in2p3.fr (Radek Stompor)

type is referred to as a map-making procedure and it attempts to compress the raw, time-domain data as obtained by the experiments, into smaller, and potentially more manageable, map-like objects. In fact, actual map-making procedures are more general and map-making-like operations are ubiquitous in the CMB data analysis, implicitly or explicitly, on other processing stages emphasizing the key role map-making solvers play in the entire analysis process. Depending on a specific application, a solution to the same map-making system may need to be derived only once or multiple times for many different right hand sides, potentially requiring different application-specific solvers.

The anticipated data volumes and the diversity of the contexts and applications define the challenge for any successful map-making software. In this work, we present a unified numerical framework designed to meet this challenge. The framework is massively parallel, flexible and extensible to permit processing very large data volumes while accounting for the specificity of different data sets leading to different models of the processed data.

This paper is organized as follows. In Sect. 2 we introduce the map-making problem, and provide a general approach for reconstructing sky maps while deprojecting several time-domain systematics that may be present in the data. In Sect. 3 we discuss the implementation details of the framework proposed in this work, and in Sect. 4 we present some applications on simulated data from ground-based and satellite-borne experiments demonstrating the methods implemented in the framework.

2. Map reconstruction

In this section, we lay out the basic formalism of sky maps reconstruction. We start from a generic data model of CMB experiments, from which we derive a general expression for unbiased sky maps estimators. We then discuss different types of these unbiased map estimators which employ different models of the noise and are adapted to deal with selected time-domain systematics.

2.1. Data model

In the map-making problem we typically assume that the raw data of the detectors, have been properly calibrated and cleaned in the pre-processing stage. The data are taken at a fixed time step and for each detector we order them by the time at which they were collected. Conceptually, the data from different detectors are typically assumed to be concatenated together in a single data vector. We denote this data vector as \mathbf{d} . We use a linear model to describe the data representing them therefore as,

$$\mathbf{d} = \sum_i \mathbf{T}_i \mathbf{x}_i + \mathbf{n}. \quad (1)$$

Here, the index i goes over different identified contributions to the measurement, whatever their origins: **these can describe for example sky stationary signals, scan synchronous signals, HWP spin synchronous signals, or long term drifts.**

The columns of each matrix \mathbf{T}_i define a time-domain template describing the effect. \mathbf{n} denotes then a stochastic contribution, which is characterized only statistically and assumed for most of the application to be piece-wise time-stationary. We refer to this term generically as noise.

A key assumption behind the model is that all relevant contributions can be described by a limited number of amplitudes stored in the vectors \mathbf{x}_i , and the corresponding templates, i.e., columns of the matrices \mathbf{T}_i , which are assumed to be known. For continuous signals, such as the sky signal itself, it implies that they can be discretized with sufficient precision.

For definiteness we assume that the $i = 0$ term is the actual sky signal. The corresponding template matrix, \mathbf{T}_0 , is then typically referred to as a pointing matrix and denoted as \mathbf{P} . The vector \mathbf{x}_0 stands then for a discretized sky signal, typically a pixelized sky map, which we denote as \mathbf{m} , ($= \mathbf{x}_0$). On collecting all the remaining template matrices, \mathbf{T}_i , in a single matrix $\mathbf{T} \equiv [\mathbf{T}_1, \dots, \mathbf{T}_i, \dots]$ and concatenating all the amplitude vectors, \mathbf{x}_i , together, in a single vector, \mathbf{x} , we can rewrite Eq. (1) as,

$$\mathbf{d} = \mathbf{P} \mathbf{m} + \mathbf{T} \mathbf{x} + \mathbf{n} \equiv \mathcal{P} \mathbf{y} + \mathbf{n}. \quad (2)$$

Here, for future convenience we have introduced a generalized pointing matrix, \mathcal{P} , and a generalized map, \mathbf{y} , which combine all the template matrices and amplitude vectors, respectively, including those describing the sky signal.

In the case of modern multi-kilo pixel experiments the size of the time ordered data can reach $\mathcal{N}_t = O(10^{12-15})$, while the number of the amplitudes is typically much smaller. In particular, the number of sky amplitudes, **that is the total number of Stokes parameters we are solving for**, is usually on order of $\mathcal{N}_{\text{pix}} = O(10^{5-8})$. The data are therefore hugely redundant reflecting the fact that features of the actual sky signal of current interest are very subdominant as compared to the noise of the measurements.

2.2. Map-making solution

The map-making problem consists of estimating the sky signal, \mathbf{m} , by compressing the redundant information contained in the raw data and deprojecting the undesirable effects defined by the combined template matrix \mathbf{T} . This is done given some assumptions about the statistical properties of the noise, \mathbf{n} . The most general estimator achieving all these goals can be written down in a concise form as [49, 42],

$$\hat{\mathbf{m}} = \left(\mathbf{P}^T \mathbf{F}_T \mathbf{P} \right)^{-1} \mathbf{P}^T \mathbf{F}_T \mathbf{d}, \quad (3)$$

where

$$\mathbf{F}_T \equiv \mathbf{W}^{-1} \left(\mathbf{I} - \mathbf{T} \left(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T} \right)^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right) \quad (4)$$

is a filtering and weighting operator, which simultaneously filters all the contributions contained in the subspace spanned by the columns of the template matrix, \mathbf{T} , i.e., $\mathbf{F}_T \mathbf{T} = 0$, and weights modes orthogonal to it by the weight matrix, \mathbf{W} . Consequently, the estimator given by Eq. (3) is explicitly unbiased as long as the system matrix, $\mathbf{P}^T \mathbf{F}_T \mathbf{P}$, is non-singular and the

assumed templates span the subspace affected by the unwanted signals. Indeed,

$$\begin{aligned}\hat{\mathbf{m}} &= (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T (\mathbf{P} \mathbf{m} + \mathbf{T} \mathbf{x} + \mathbf{n}), \\ &= \mathbf{m} + (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T \mathbf{n},\end{aligned}\quad (5)$$

and the last term vanishes when averaged over the statistical ensemble of noise realizations as $\langle \mathbf{n} \rangle_{\text{noise}} = 0$.

We note for completeness that if some prior information concerning the sought-after amplitudes of the templates is available, **assuming those are given by a multivariate Gaussian distribution**, it can be incorporated in the map estimation modifying the filtering operator, which in such a case reads,

$$\tilde{\mathbf{F}}_T = \mathbf{W}^{-1} \left(\mathbf{I} - \mathbf{T} (\mathbf{\Theta}^{-1} + \mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right), \quad (6)$$

where $\mathbf{\Theta}$ is the covariance matrix of the sought after amplitudes, \mathbf{x} . This matrix may be rank deficient if the prior information is available only for some of the amplitudes and not others. No prior information at all corresponds to $\mathbf{\Theta} \rightarrow \infty$ and $\mathbf{\Theta}^{-1} \rightarrow 0$ and we recover Eq. (4). We note that the templates for which the prior is available are not completely filtered out from the data during each map-making procedure, i.e., $\tilde{\mathbf{F}}_T \mathbf{T} \neq 0$, rather they are Wiener-filtered instead. The sky signal estimates are therefore not unbiased when averaged over the noise ensemble, however, their overall uncertainty over the ensemble of the noise and template amplitudes is minimized given the assumed priors [5].

Eq. (3) can be implemented implicitly, for instance, by solving for the full amplitude vector, \mathbf{y} , weighting the data with help of the weight matrix, \mathbf{W} , i.e.,

$$\hat{\mathbf{y}} = (\mathcal{P}^T \mathbf{W}^{-1} \mathcal{P})^{-1} \mathcal{P}^T \mathbf{W}^{-1} \mathbf{d}, \quad (7)$$

and truncating it afterwards to retain the sky signal part only [49].

Alternately, one can opt for a two step approach where the amplitudes of the undesirable signals, \mathbf{x} , are first estimated as,

$$\hat{\mathbf{x}} = (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{F}_P \mathbf{d}, \quad (8)$$

and the sky map is then recovered on the second step via,

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} (\mathbf{d} - \mathbf{T} \hat{\mathbf{x}}). \quad (9)$$

This approach is then essentially equivalent to the destriping technique originally proposed for Planck [6, 9, 32, 39, 38, 41], and later studied also in the context of ground-based experiments [52, 51, 42].

Both these approaches have been successfully employed in the past. In the realm of the multi-kilo pixel experiments, the number of instrumental, detector-specific contributions in the data is typically exceeding by far the number of the observed sky pixels. Thus, the system matrices in Eq. (7) and (8) are significantly larger than the one appearing explicitly in Eq. (3). Consequently, solving directly Eq. (3) may look as potentially more computationally attractive. For this to be indeed the case, there must be a computationally-efficient way to construct the template orthogonalization kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$, which in general can be rather cumbersome to calculate. Fortunately, in

many relevant applications the kernel is usually very structured and sparse. This stems from the fact that most of the templates of interest have compact and disjoint support, which is moreover known a priori. This can be capitalized on to save on the calculations making the estimator in Eq. (3) not only computationally feasible but also preferred in many applications. We therefore implement this approach in the proposed framework.

Eq. (3) generalizes and includes as special cases many of the standard map-making techniques. Those correspond to some specific choices of either the weight matrix or the template matrix or both. These particular cases include: (1) the standard maximum likelihood (minimum variance) map-making, where the templates are often dropped and the weights are taken to be given by the inverse covariance of the noise term, \mathbf{n} , **which is assumed to be Gaussian**; (2) the binned map-making where the weights are taken to be diagonal and equal to the noise RMS and the templates neglected; (3) the standard destriper, where the templates are taken to be piece-wise constant offsets and the weights diagonal and set by the noise RMS; (4) the generalized destriper, where the templates can be more general and the covariance of their amplitudes could be included as prior information.

Eq. (3) is very flexible allowing for diverse choices for the pointing, the template, and the weight matrices with the only practical constraint being numerical considerations related to their application. As we discuss in the following section this permits applications of this equation in many cases of practical interest.

2.3. Pointing matrix

The sky signal measured by an instrument is a convolution of the actual sky signal, expressed with the three Stokes parameters I, Q and U , with an instrument response matrix, $\mathcal{R}(\delta\gamma, \psi, \nu)$. This matrix will in general depend on both the sky angle, $\delta\gamma$, measured with respect to the observation direction, γ , the orientation of the instrument with respect to the sky coordinates, ψ , and the frequency, ν . Therefore, a measurement made in ν_c frequency band, defined by a bandpass, $\mathcal{W}(\nu, \nu_c)$, can be written in general as,

$$\mathbf{d}_{\nu_c}(t) = \int d\nu \int d\delta\gamma \mathbf{e}_i^T \mathcal{R}(\delta\gamma, \psi(t), \nu) \times \mathcal{W}(\nu, \nu_c) \mathbf{s}(\gamma(t) + \delta\gamma, \nu) + \mathbf{n}(t), \quad (10)$$

where $\gamma(t)$ is the observation direction at time t and \mathbf{e}_i is a unit vector which defines whether we ultimately measure total power only, $i = 0$, or Stokes Q , $i = 1$, or U , $i = 2$, parameter. $\mathbf{s}(\gamma, \nu)$ is a vector of the three Stokes parameters as defined in the direction $\gamma(t)$ and at frequency ν .

It is frequently possible to represent the general instrument response matrix as a product of a Mueller matrix describing the detection chain of the instrument and a matrix defining the beam effects,

$$\mathcal{R}(\delta\gamma, \psi(t), \nu) \equiv \mathcal{M}(t(t), \nu) \mathcal{B}(\delta\gamma, \nu) \mathbf{R}(\psi(t)), \quad (11)$$

where the matrices, \mathcal{M} and \mathcal{B} , are defined with respect to instrument coordinates and the matrix $\mathbf{R}(\psi(t))$, denotes a rotation

between the sky and instrument coordinates. We note that the Mueller matrix, \mathbf{M} , may in general depend on time via some of the instrument parameters, such as, phase of a polarization modulator, as denoted by parameter $\boldsymbol{\iota}(t)$.

Assuming that the beam matrix is diagonal, axially symmetric, i.e., depending only on the magnitude of the angle, $\delta\boldsymbol{\gamma}$, and not on its orientation, and that the beams for the Stokes parameters Q and U are identical, we can commute the rotation operator and the beam matrix, rewriting Eq. (10), as,

$$\begin{aligned} \mathbf{d}_{\nu_c}(t) &= \int d\nu \mathcal{W}(\nu, \nu_c) \mathbf{e}_i^T \mathbf{M}(\boldsymbol{\iota}(t), \nu) \mathbf{R}(\boldsymbol{\psi}(t)) \\ &\times \underbrace{\int d\delta\boldsymbol{\gamma} \mathcal{B}(\delta\boldsymbol{\gamma}, \nu) \mathbf{s}(\boldsymbol{\gamma}(t) + \delta\boldsymbol{\gamma}, \nu) + \mathbf{n}(t)}_{\equiv \bar{\mathbf{s}}(\boldsymbol{\gamma}(t), \nu)} \end{aligned} \quad (12)$$

where $\bar{\mathbf{s}}(\boldsymbol{\gamma}(t), \nu)$ stands for a beam-smoothed sky signal, which can be therefore characterized by a finite set of pixelized sky signal amplitudes, $\bar{\mathbf{s}}(p(t), \nu)$, where $p(t)$ is a sky pixel observed at time t .

For the frequency bandpass convolution, the simplest case is when the bands defined by $\mathcal{W}(\nu, \nu_c)$ are sufficiently narrow so the beam-convolved sky signal is essentially constant across the band. This allows us to write,

$$\begin{aligned} \mathbf{d}_{\nu_c}(t) &\approx \left[\int d\nu \mathcal{W}(\nu, \nu_c) \mathbf{e}_i^T \mathbf{M}(\boldsymbol{\iota}(t), \nu) \mathbf{R}(\boldsymbol{\psi}(t)) \right] \\ &\times \bar{\mathbf{s}}(p(t), \nu_c) + \mathbf{n}(t), \\ &\equiv \mathbf{w}(t) \bar{\mathbf{s}}(p(t), \nu_c) + \mathbf{n}(t). \end{aligned} \quad (13)$$

The vector, $\mathbf{w}(t)$, defines the pointing weights with which the different Stokes parameters at the band-center frequency, ν_c , are coadded together to give the measured signal, while $p(t)$ defines a sky pixel number observed at time t . This equation also holds when the bands are arbitrary but the Mueller matrix is not dependent on the frequency, ν . However, the recovered sky signal represents then the actual sky signal convolved with the experimental bandpass.

In a more general case when the bands are arbitrary and the Mueller matrix depends on the frequency, then as long as the time-dependence of all the relevant elements of the Mueller matrix is known [59], the algebraic structure of Eq. (13) is preserved. In such cases the recovered sky signal will be however in general a linear combination of Stokes parameters convolved with frequency response functions defined by the bandpasses and the Mueller elements. Moreover, for each pixel the number of the amplitudes, required to describe the data may be significantly larger than the number of the Stokes parameters themselves [59].

In all these cases the information about the pointing weights and the pixel as observed at time t can be encoded in a single matrix, the pointing matrix \mathbf{P} , which is tall and skinny with the number of rows given by the total number of measurements and the number of columns given by the total number of observed pixels multiplied by the number of observables assigned to each pixel. The data model can be then expressed conveniently in a vector form as,

$$\mathbf{d} = \mathbf{P}\mathbf{m} + \mathbf{n}. \quad (14)$$

In a row of the matrix \mathbf{P} corresponding to time t , the only non-zero elements are those in columns corresponding to pixel p observed at the time t . These elements are given by the respective elements of the pointing weight vector \mathbf{w} and are known given our knowledge of the instrument and the observation. In the most straightforward applications, there are therefore either 1 – for total intensity measurements, or 2 – for Q and/or U polarization measurements, or 3 – for polarization-sensitive measurements, non-zero elements in each row. Consequently, the pointing matrix is typically very sparse and thus manageable in spite of its huge size.

Important examples include the case of a fixed single-slot polarizer, when the measurements can be expressed as,

$$d_t = I_{p_t} + \cos(2\varphi_t)Q_{p_t} + \sin(2\varphi_t)U_{p_t} + n_t, \quad (15)$$

where φ denotes an angle of the polarizer with respect to the sky coordinates. Similarly, for an idealized frequency-independent, half-wave plate, used to modulate the incoming signal, the data model reads,

$$d_t = I_{p_t} + \cos(2\varphi_t + 4\phi_t)Q_{p_t} + \sin(2\varphi_t + 4\phi_t)U_{p_t} + n_t, \quad (16)$$

where ϕ_t stands for the HWP orientation with respect to the instrument coordinates. In both these cases, there are three non-zero entries per row of the pointing matrix, \mathbf{P} . As mentioned earlier this number can be lower, i.e., for total intensity measurement, or in the case of pair-differencing approaches. It can also be larger, for instance, in the case of more realistic models of the half-wave plate, where the induced rotation of the incoming polarization depends on the frequency and one may need to introduce as many as 5 different amplitudes to characterize the signal in each sky pixel. This results therefore in 5 non-zeros per row of \mathbf{P} [59]. The number of non-zeros per row may be even larger in the cases of the generalized pointing matrix, \mathcal{P} , as it has to incorporate information about the non-sky degrees of freedom stored in the matrix \mathbf{T} . In all these cases the pointing matrix is nevertheless extremely sparse and the number of non-zero values per row is very limited.

A notable exception is the case of asymmetric beams, i.e., when \mathcal{B} depends explicitly on a vector $\delta\boldsymbol{\gamma}$ and not just its length. If beams happen to be strongly asymmetric to the extent that it is necessary to correct for it on the map-making stage, the data model in Eq. (14) would still apply and though the pointing matrix would be in such a case dense, it would be also structured, permitting specialized algorithms for the map-making problem as have been indeed proposed [4, 20, 25]. In the framework discussed here the focus to date has been specifically on the cases with an arbitrary but sparse pointing matrix. And while the framework readily allows for extensions which would accommodate also cases with asymmetric beams, this requires however further development that is left for the future and is not considered in this work.

2.4. Weights

The estimator in Eq. (3) is unbiased whenever the templates are appropriately chosen, given the anticipated contaminants,

and the system matrix is non-singular. This property is independent of the choice of the assumed weighting. The weighting is however key in obtaining an estimate with as high signal-to-noise ratio as possible.

A maximum likelihood estimate (which is as well a minimum variance one) is indeed given by Eq. (3) with the weights corresponding to the noise covariance,

$$\mathbf{W} = \mathbf{N} \equiv \langle \mathbf{n} \mathbf{n}^T \rangle. \quad (17)$$

Given the sizes of the current and forthcoming data sets, such covariance matrices are readily unmanageable in their fully general form. However, the noise properties of any single detector data can be typically adequately characterized as piece-wise stationary. The corresponding noise covariance is then block-diagonal where each block corresponds to a different stationary interval and has the structure of a Toeplitz matrix. This means that for the i th interval, the correlation between the noise at times t and t' depends merely on the time interval, $|t - t'|$, and not actual values of t and t' . Consequently, the i th block of the covariance matrix satisfies the condition,

$$\mathbf{N}^{(i)}(t, t') = \mathbf{N}^{(i)}(0, |t - t'|). \quad (18)$$

The noise in the CMB measurements is typically correlated all the way to very long timescales and the Toeplitz blocks of the noise covariance are in general dense. However, at least for well-optimized scanning strategies, accounting for the correlations on the largest time scales is usually not necessary, as they affect predominantly only signals varying on angular scales too large to be well-constrained by the experiment anyway and therefore manifest themselves merely as an offset of the map (or a submap) [53]. Consequently, it is justifiable, and desirable, to assume that the correlations are band limited and set to zero beyond some characteristic correlation length parameter, λ (also called half-bandwidth). The value of λ should be fixed according to properties of the noise and the adopted scanning strategy [53].

In the map-making problem, Eq. (3), we need an inverse of the weight matrix. The inverse of a general Toeplitz matrix is not Toeplitz and it is not easily computable. The inverse Toeplitz blocks once computed would have to be stored in the computer memory, which would quickly undermine the numerical feasibility of the procedure. However, for a banded-Toeplitz matrix, in particular with a rather narrow band, its inverse can be very well-approximated by a Toeplitz matrix. The approximate inverse is given by,

$$\mathbf{N}^{(i)-1}(t, t') \simeq \begin{cases} \mathbf{C}^{(i)-1}(t, t'), & \text{if } |t - t'| \leq \lambda \\ 0, & \text{otherwise} \end{cases} \quad (19)$$

Here $\mathbf{C}^{(i)}$ is a circulant matrix embedding the Toeplitz block $\mathbf{N}^{(i)}$. The inverse of a circulant matrix is also circulant and can be easily calculated with the help of Fast Fourier Transforms (FFTs) at low numerical cost [14] as discussed later in Sect. 3.

The noise covariance has to be derived typically from the same data as the data used later for the map-making [12, 49, 60]. As a consequence, even for the case of Gaussian noise in the time-domain, the noise of the resulting map is not strictly

Gaussian and some care may need to be taken in minimizing or accounting for it in particular at the largest angular scales. For most of the applications this effect is however thought to be largely irrelevant.

One specific choice which is of interest, in particular because of its numerical efficiency, is that of a diagonal weight matrix. This is a limiting case of the block-Toeplitz case discussed above with $\lambda = 0$, however, the application of the weights in such cases can be done directly in the time-domain without any need for Fast Fourier Transforms. From the point of view of the quality of the estimated maps this is however often a rather bad choice which will frequently result in the estimates being dominated by the low frequency noise modes. These modes can be often treated as templates as it is for instance the case of the destriper techniques [6, 9, 32, 46, 26]. We discuss choices of templates suitable for this task in Appendix A.

Whenever a prior can be imposed on the amplitude of the templates, as in Eq. (6), introducing the templates is equivalent to adding a low-rank correction of the weight covariance, so in such a case the weights, \mathbf{W} , are replaced by [49],

$$\mathbf{W} \rightarrow \mathbf{W} + \mathbf{T} \mathbf{\Theta} \mathbf{T}^T. \quad (20)$$

Conversely, any low rank correction to the weight matrix can be rephrased as a set of templates and implemented via the projection operator, \mathbf{F}_T . A successful set of templates will correspond to a low rank approximation of the full noise covariance, which reflects all of its characteristics which are relevant for the problem at hand. The challenge is therefore in finding out what those characteristics are.

The need for templates is not limited to the cases when the weight matrix is diagonal. In particular, it is often desirable to use templates to mitigate some systematic effects in the maximum likelihood map-making, when the weights are block-Toeplitz. In practice, the number of such templates is expected to be lower than the number of templates needed to model the low-frequency noise correlations in the case of diagonal weights. It is therefore computationally advantageous to extend the solution vector by adding the amplitudes of the required templates and solve this extended system using standard maximum likelihood technique with $\mathbf{F}_T = \mathbf{W}$ as in Eq. (7). This is equivalent to the meta-pixel approach of [49]. Consequently, the proposed framework implements two options: (1) a maximum likelihood map-making with an arbitrary (but sparse) pointing matrix and block-Toeplitz weights, and, (2) cases of diagonal weights with templates which are directly included in the operator \mathbf{F}_T , as in Eq. (4) or Eq. (6). However, other cases could be implemented using the functionality provided by the proposed libraries.

All the considerations above apply directly to data of a single detector. In the realm of multi-kilo pixel arrays, we have to deal with data collected concurrently by many detectors. It is therefore overly optimistic to assume no correlations between the measurements of different detectors. As pointed out earlier, neglecting those will not bias our sky signal estimates however it may easily render them very noisy. Sources of the correlations can be diverse. They could be either due to some common noise source seen by multiple detectors, such as atmospheric

420 signal unavoidably present in the case of the ground-based ex-
periments, or can be induced in the detection chain of the in-
strument, for instance, in its readout system. In the case of the
maximum likelihood map-making such correlations could be
425 in principle included as an off-diagonal block of a full weight
(noise correlation) matrix, which includes weights for all de-
480 tectors simultaneously. However, this would lead to significant
numerical overhead rendering such a map-making hardly use-
ful in practice. Consequently, the proposed framework does not
provide this functionality at this time as a default.

430 Some workaround approaches which are more manageable
but can render comparable performance are however frequently
available. One is to use templates to account for those correla-
tions or to capitalize on differencing scheme for data of some
of the detectors devised to remove the correlated component.

435 2.5. Templates 490

The templates can be in principle arbitrary reflecting the di-
versity of possible non-cosmological contributions which may
be present in the data. Whether those can be deprojected from
the data on the map-making stage will in general depend on
440 three facts: (1) whether we can define efficiently a subspace,
as spanned by the columns of the matrix \mathbf{T} , which contains all
possible undesirable contributions; (2) whether this can be done
while retaining limited dimensionality of the subspace; and (3)
whether this subspace overlaps with the subspace spanned by
445 the sky signal. If (1) is not met then some unaccounted for
residuals will be present in the estimated sky signal. Failing
on (2) may affect the quality, signal-to-noise, of the map, but
the estimate will not be biased. Higher the dimensionality of
the deprojected subspace increases typically the chance that (3)
450 is not met either. If (3) is not fulfilled the system matrix in
Eq. (3) is singular and therefore some of the modes of the sky
signal may not be recoverable. We note that if such modes are
properly identified the map-making procedure can still be often
applied in a meaningful way estimating all the modes for which
455 this can be achieved. The information about the non-recovered
modes would have to be then passed on and accounted for on
all subsequent stages of the analysis [42]. 510

As mentioned earlier, from the computational point of view
some templates may be difficult to implement as they may lead
460 to significant computational overhead, in particular in the con-
struction of the orthogonalization kernel. This issue can be of-
ten sidestepped by introducing multiple independent copies of
the same template specific to different, disjoint time intervals.
Crucially, the procedure makes all the computations more lo-
465 cal in the computer memory, thus avoiding extra need for ex-
tensive global communication between the compute nodes. At
the same time the increase of the size of the template orthog-
onalization kernel is mitigated by the fact that it is limited to
470 additional diagonal blocks. There is a clear downside to this
procedure as it unavoidably increases the dimensionality of the
deprojected subspace, therefore leading to a loss of precision.
This however may be often acceptable. Moreover, in actual ap-
475 plications many templates of interest are naturally specific to
limited ranges of the observation time. As we comment later,
the proposed framework is therefore restricted to this kind of

templates. Within such limits the framework allows in princi-
ple for any user-defined templates. In addition, some templates
are implemented directly in the package. They are described
in [Appendix A](#).

2.6. Degeneracy issues

So far our pointing or generalized pointing matrices were
considered to be full column rank, therefore ensuring that the
system matrices are invertible. In practice, this assumption can
be violated for different reasons which we will explore in the
following subsections.

2.6.1. Scanning strategy sourced degeneracies

In order to solve for three Stokes parameters I , Q and U , in
each pixel in the sky, the scanning strategy should ensure that
each pixel is observed at least three times with sufficiently dif-
ferent orientations. The polarization modulation induced by the
half-wave plate helps in this regard as it ensures that in each
pixel crossing, sufficient samples are taken with different orien-
tations to allow for the reconstruction of the pixel Stokes param-
eters. However, for some pixels, particularly in the periphery of
the observed sky patch, it may happen that the number of revis-
its of the pixel is very low and on each revisit the pixel crossing
time is very short therefore not allowing for sufficient sampling
of the polarization angles.

In practice this issue is solved by examining the condi-
tion numbers of each block in the block diagonal matrix
 $\mathbf{P}^T \text{diag} \mathbf{W}^{-1} \mathbf{P}$. Each block corresponds to an observed sky
pixel, which if observed with sufficient redundancy, should ide-
ally have a condition number close to 2. The higher the condi-
tion number, the more the pixel noise gets boosted. We there-
fore set a threshold on the condition number of each pixel,
above which the pixel, and its corresponding TOD samples, get
effectively removed from the data.

2.6.2. Templates linear dependencies

The second type of degeneracies stems from linear depen-
dencies between templates. One prominent example, is the 0th
order polynomial template and the ground template which both
filter the global offset of the TOD (see [Appendix A](#), and [42]).

Such degeneracies prevent the kernel matrix $\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T}$ from
being invertible. In practice, this does not pose any problem for
the sky signal recovery and is easily addressed: we orthogo-
nalize the templates by computing the Moore-Penrose pseudo-
inverse of the kernel via a singular value decomposition (SVD).
This is made possible because of the simple structure of the
kernel matrix as discussed earlier in this section. See also [Ap-
pendix A](#) for more details in the case of specific templates.

2.6.3. Templates and sky degeneracies

This last type of degeneracies is due to linear dependencies
between the filtered templates and some sky modes. When such
degeneracies are present, the concerned sky modes are impos-
sible to reconstruct and would therefore be missing from the
maps, because they are filtered from the data by the deprojec-
tion operator \mathbf{F}_T .

In practice, this means that the matrix $\mathbf{P}^T \mathbf{F}_T \mathbf{P}$ would have singularities. Given the size, and the fact that this matrix is dense in general, the direct computation of these singular modes to regularize the inversion is not feasible. In an iterative approach such as the one we adopt in this work, the presence of such degeneracies manifests itself as a significant slow down in the convergence rate. This is caused by the small eigenvalues in the eigenspectrum of the system matrix introduced by the singular or nearly singular modes. One way to address this issue, is to approximate these modes with an iterative method such as the Lanczos procedure, and deproject them from the system. This provides the basis for the *a posteriori* two-level preconditioning technique, that we will address in more detail in Sect. 3.

3. Parallel mapping framework

In this section, we provide a detailed description of the software framework and highlight key aspects of its implementation. The software is publicly accessible at the following link: <https://github.com/B3Dcmb/midapack>. We start from an overview of the code with its overall architecture, then describe each of its building blocks and how they fit into the global framework.

3.1. Code overview

A summary of the high-level code architecture, is given in figure 1. The software is comprised of two main C-libraries:

- The first library, called **M**icrowave Data Analysis **P**ACKage (MIDAPACK), provides low-level operations needed to perform high-level procedures in the data analysis pipelines. Such operations include, sparse linear algebraic operations for generalized pointing and template operators, customized communication schemes for distributed data reduction, operators for structured matrices, such as Toeplitz matrices, or block-diagonal templates kernels, necessary to perform some time-domain operations such as templates filtering or noise weighting. We note that these operations can be relevant for all main stages of the data analysis pipeline, but in accordance with the scope of this paper we mainly focus on map-making applications.
- The second library, called **M**idAPack **P**aRAllel Iterative Sky EstimatorR (MAPPRAISER), builds on the first library to provide different numerical techniques, solvers and map-making methods to compute unbiased maps, and filter time-domain systematics concurrently with the map projection.

The routines conform with the MPI programming model and are therefore able to use massively parallel platforms with memory distributed between a large number of compute nodes. Some routines also make use of multithreading, but this option still needs to be optimized to be fully functional in the code. However these routines may provide the basis for a future extension of the framework to a full hybrid OpenMP/MPI model.

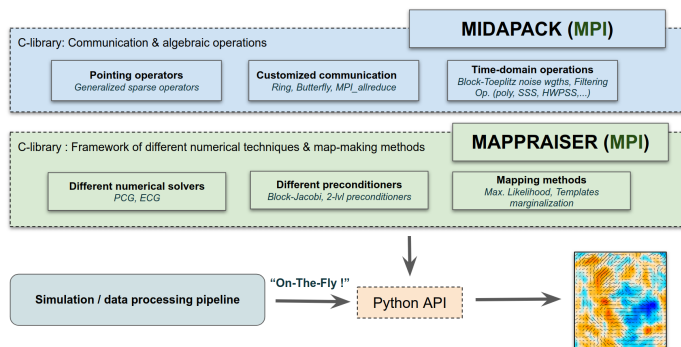


Figure 1: Overview of the software top-level architecture.

On top of this two-level core layer of the code, there is a Python wrapper, which allows the map-making code to interface easily with simulation or real data analysis pipelines. In particular, the Python wrapper can be used to perform the map-making run “on-the-fly” chaining the map-making code with simulation software or with a previous step of an actual data processing pipeline, therefore reducing I/O cost.

3.2. The microwave data analysis package (MIDAPACK)

The MIDAPACK library consists of three main sub-packages:

- **The pointing operations package MAPMAT:** this package contains sparse matrix vector products operations to perform the projections from time-domain data to map-domain data and vice-versa. In their most general form, these matrices are rectangular with considerably more rows than columns, given the relative sizes of the time-domain and map-domain data and are referred to as *tall-and-skinny* matrices. Their most important property however is their sparsity as they feature at most few non-zero elements per row.
- **The Toeplitz algebra package:** this package provides routines to perform the product of Toeplitz-structured matrices with an arbitrary data matrix. We refer hereafter to matrices which are symmetric, block-diagonal with banded Toeplitz blocks, or matrices with all of the above properties along some rows and corresponding columns missing as “Toeplitz-structured”. These special matrices are common in CMB data analysis, as they encode the correlation properties of stationary time-domain processes such as instrumental 1/f-noise or atmosphere extending up to the correlation length corresponding to the band width of the matrix. This can be adjusted to provide the best trade-off between the computational time and quality of the result. These routines are used in noise weighting operations, i.e., weighting the different frequency modes in the time-domain data with their corresponding inverse noise power.
- **The templates algebra package:** this package provides routines to handle multiple sets of templates modeling

different sorts of systematics. The routines allow to per-
form deprojections of templates, i.e., time-domain filter-
ing operations as well as to orthogonalize a set of dif-
ferent templates to regularize the deprojection operation.
The package allows for and seamlessly manipulates het-
erogeneous templates which may be defined by different
means. This is an important functionality given the di-
versity of effects which the template formalism aims at
describing.

This is facilitated by grouping the templates into template
classes. A general template corresponds to a single col-
umn of any of the matrices \mathbf{T}_i , Eq. (1). Templates cor-
responding to columns of the same matrix \mathbf{T}_i aim at de-
scribing the same physical effects and often share some
properties. Hereafter, we uniquely assign, a template
class to each matrix \mathbf{T}_i and all columns of the matrix de-
fine therefore templates which belong to the same tem-
plate class.

The templates can be defined individually or as a class.
They can be defined explicitly via explicitly stored el-
ements of the corresponding column of the matrix \mathbf{T}_i .
They can also be defined implicitly and computed only as
needed using some predefined analytic prescription for a
set of parameters. This is for instance the case of the
polynomial templates, which are uniquely defined by the
order of the polynomial. The templates can be also often
stored in a compressed form. This is the case when the
modelled contribution is synchronous with some instru-
mental or observation parameter (or parameters), which
in turn is a function of the observation time. Such con-
tributions are common including, for example, scan syn-
chronous or ground-pickup as discussed in Appendix A.
To be manageable such contributions are discretized by
binning (or pixelizing) the synchronous parameter into
disjoint bins. This procedure introduces as many inde-
pendent signals as bins. In the template formulation, this
means that there is a separate template needed for each
bin. Given that each template is a time-domain object,
the memory needed to store all such templates becomes
quickly prohibitive and at the very least restricting im-
portantly the number of bins we can consider.

However, such templates can be defined as a class in
a memory efficient fashion capitalizing on the fact that
there is only one bin (or more generally, a limited num-
ber of bins) contributing to a measurement at any given
time. The contribution of such signal to a measurement
made at some time t depends on a uniquely assigned bin-
number $b(t)$, which corresponds to this time. This infor-
mation can be readily stored as a single time-domain vec-
tor, $b(t)$. In addition we may need at most another time-
domain vector which defines with what weight the signal
amplitude corresponding to that bin should be added to
a measurement at time t . These two objects are sufficient
to characterize any binned contributions independently on
the number of assumed bins, therefore providing a con-
venient and common compressed form encoding infor-

mation relevant to the template class.

A template is in general an arbitrary time-domain-like
vector. However, in the proposed package we restrict our-
selves to templates which have a compact support fully
contained within the length of a single full scan of one of
the detectors. The templates considered here are there-
fore specific to scans of each detector.

We note that while more complex situations may arise,
scan specific templates are indeed the most typical occur-
rence in actual applications, and are sufficient to model
most, if not all, relevant systematic effects. This restric-
tion allows to significantly reduce the computational and
communication overhead.

One consequence of it is that while templates assigned to
a given process may or may not be orthogonal, templates
assigned to two different processes, even if belonging to
the same template class are always mutually orthogonal.

While the package allows for a broad set of user-defined
templates it also implements directly some of them.
These are described in detail in Appendix A.

In the following subsections, we will discuss the implemen-
tation details of each one of these sub-packages. Before doing
so we need to define the distribution scheme of the time domain
data. A typical, modern experiment features multiple detectors
observing the sky over an extended period of time. The full
observation time can be broken into a number of independent
time-intervals common for all detectors in a way that for each
detector the noise can be considered stationary within each in-
terval and uncorrelated between them. We hereafter refer to
these time-intervals as “scans”. The timestreams of all detec-
tors are stacked together vertically in a two dimensional array
with columns corresponding to the detectors and rows to times
at which data were taken. This full data array can be broken
down into disjoint sub-arrays each made of full-length rows cor-
responding to one of the scans.

Typically, we can consider two regimes depending on the
size of the data. In the small data sets regime (small number
of detectors, short scans, low sampling rate), each MPI process
can hold the timestreams of all detectors for a given number of
scans. In the large data sets regime (large number of detectors,
long scans, high sampling rate), each MPI process can hold only
the timestreams of a single scan and only for a subset of detec-
tors. To illustrate these schemes, we show a toy example of this
data distribution in figure 2. The choice of not splitting scans
of a single detector between multiple processes avoids unnec-
essary communication between neighbouring processes either
during noise weighting or time-domain filtering operations.

3.2.1. MAPMAT: pointing algebra package

Pointing data distribution & pixelization schemes: The
pointing data distribution follows from the distribution scheme
adopted for the timestreams: each process is assigned the
pixels and pointing weights corresponding to the detector
timestreams it is processing. This also defines the distribution
of pixel-domain objects, such as maps. As a result, given the




Figure 2: Distribution scheme of the TOD depending on the processed data set size. Each rectangular block corresponds to one detector time stream for the full duration of a scan. The colors represent the assignment of data to a given MPI process.

high redundancy of the typical scanning strategy, information concerning a given pixel may be stored in the memory of multiple processes. The pixels indices are defined via two indexing schemes. A global pixel indexation, corresponding to the global pixel index on the sky, and a local pixel indexation, corresponding to a reordering of the pixels observed at least once by the local timestreams of each MPI process. The package is independent of the particular choice made for the sky pixelization.

Data structure: The pointing matrix is defined with a C-struct called `Mat`. Each MPI process holds an instance of this structure in memory. One integer, `nnz`, specifies the number of non-zero elements per row. The values and columns locations of these elements are specified with two arrays as is done in the ELL¹ storage format, for efficient memory consumption. It also contains the ordered set of locally observed pixels, and a precomputed set of arrays and parameters necessary to conduct the customized collective communication operations in an efficient way, along with a flag specifying the particular communication scheme chosen for these operations.

Projection operations & collective communication: The package features two main core operations:

- **The un-pointing operation:** This operation generates a time-domain vector, \mathbf{t} , from a given map-domain object, \mathbf{m} , and a pointing matrix \mathbf{P} : $\mathbf{t} = \mathbf{P}\mathbf{m}$. Given the data distribution scheme, this operation can be performed locally in each MPI process, as a simple sum of the products of the pointing weights with the corresponding map entries.
- **The pointing operation:** This is a projection operation, where given a time-domain vector, \mathbf{t} , and a pointing matrix \mathbf{P} , we compute a map vector \mathbf{m} , where each pixel is the sum of the time samples observing it: $\mathbf{m} = \mathbf{P}^T \mathbf{t}$. Given that the data are distributed in time-domain, the

time samples observing a given pixel, are typically distributed over many MPI processes. Hence this operation is performed in two steps: we first compute local products of the pointing weights with the local time samples, \mathbf{m}_i . This results in submaps of the full map which contain only a subset, \mathcal{P}_i , of pixels, which is relevant to the time-domain data assigned to this process. We then compute the sum of these local submaps over all processes obtaining the global map projection, i.e.,

$$\mathbf{m}(p) = \sum_{i: p \in \mathcal{P}_i} \mathbf{m}_i(p),$$

where for a pixel p of the full map, the sum goes over those processes, i , for which p belongs to the local set of pixels, \mathcal{P}_i .

The second step requires a global, irregular communication operation. Given that these types of operations are not scalable, this represents a bottleneck in the map-making procedure. In addition to the standard `MPI_Allreduce`, the `MAPMAT` package implements several customized communication algorithms such as `Ring` and `Butterfly` [35].

3.2.2. Toeplitz algebra package

Data distribution: For each detector timestream, we build a Toeplitz block that is completely defined by one array containing a number of elements equal to the half-bandwidth λ . These blocks are therefore distributed following the time-domain data layout: each MPI process holding in memory the Toeplitz blocks corresponding to the detectors and the scans it has been assigned. In case, we want to include correlations between detectors, we would need to also build cross-correlation Toeplitz blocks, and hold them in memory, and eventually need a significant amount of communication depending on the spatial scale of the cross-correlations.

Data structure: The Toeplitz-structured noise covariance is defined with a C-struct called `Tp1tz`. Each MPI process holds an instance of this structure in memory. It is parametrized with the total number of rows, the number of Toeplitz blocks both in total and in the local memory of the MPI process, the localization of the local segment in the global matrix, in addition to the size of the local timestreams, and an array of instances of a C-struct called `Block`. This latter structure defines each Toeplitz block, and contains the actual elements of the first row, the half bandwidth, the total size of the block and its localization in the global data structure. The `Tp1tz` structure is also assigned an MPI communicator, since the product operations can also handle data of a single scan that is split between many MPI processes, hence requiring some point-to-point communication. However, in practice, it is the best to avoid distributing the data in such a way, reducing the communication overhead.

Toeplitz product operations & numerical algorithms: The package provides routines to perform the product of a Toeplitz-structured matrix with a general data matrix, which is stored as

¹a storage format originally used by the ELLPACK package, storing sparse matrices only via their non-zero entries and their corresponding column numbers.

vectors in the column-wise order. Three product operations can be performed:

- 805 • **stmm routines:** these concern the multiplication of a symmetric banded Toeplitz matrix by a data matrix of the same size. If the data are distributed over several MPI processes, then each process should have the same and full Toeplitz matrix stored in memory. 860
- 810 • **stbmm routines:** for the multiplication of a symmetric, block-diagonal matrix, with Toeplitz banded blocks by a data matrix. As explained above, such structure is defined with a list of instances of the C-struct `Block`, each representing a Toeplitz block in the block-diagonal matrix. 870 Each block is then multiplied with the corresponding sub-block of the data matrix. We also note that the blocks need not be contiguous, i.e., we can have rows in the data matrix which correspond to no given block in the Toeplitz structure, such rows are copied unchanged implicitly assuming the corresponding entries in the Toeplitz structured matrix to be equal to 1. Given the data layout adopted in practice, each MPI process will have a number of full detector timestreams, and therefore all the corresponding Toeplitz blocks. However, in general, as was discussed before the package also handles, detector 825 timestreams being split between multiple MPI processes, in this case each of those MPI processes should have the corresponding full Toeplitz blocks stored in memory, and communication is used to copy necessary data to enable the product operation locally. 880
- 830 • **gstbmm routines:** these are similar to stbmm routines, but with some sets of rows and corresponding columns removed from the Toeplitz-structured matrix, we refer to these rows and corresponding columns as gaps. 890 The routines perform a product in a similar way to the stbmm routines but neglecting the contributions from the columns in the gaps and setting the corresponding rows in the output to zero.

840 So far we have discussed, which operations may be called according to the data structure, but we have not discussed how the product operation with the noise covariance blocks is performed in itself. The algorithm is based on a *shift-and-overlap* approach [43], where a product of a single band-diagonal Toeplitz matrix by a data matrix is done by dividing 900 them into a set of overlapping blocks, and performing a sequence of products of the overlapping Toeplitz subblocks by the corresponding overlapping segments of the data matrix. Each of the latter products is performed in turn by embedding the Toeplitz subblock in a minimal circulant matrix and performing 905 the multiplication via Fast Fourier transforms. The size of the subblock is set appropriately to optimize the computation and is typically a multiple of the half-bandwidth λ . The overall complexity of the operation is $\mathcal{O}(n \ln \lambda)$, where n is the size of the initial Toeplitz matrix. Performing the operation directly by 910 embedding the full Toeplitz matrix into a circulant one that is twice as large and computing the product via FFTs would yield a complexity of $\mathcal{O}(2n \ln 2n)$.

When communication is needed as in the cases we discussed before, it remains local, involving only neighboring processes. Each process needs to send to and receive from a neighboring process a vector of data of the length defined by the half bandwidth of the Toeplitz block, λ , shared between them. This is sufficient to compute a part of the Toeplitz-vector product corresponding to the input data of each process locally without any need for further data exchanges. In particular we note that all the FFT calls used by the package are either sequential or threaded. The point-to-point communication pattern enabled by the package can be either *non-blocking*, using `MPI_Isend` and `MPI_Irecv` calls, or *blocking* with `MPI_Sendrecv` calls. The choice is specified via a flag `FLAG_BLOCKINGCOMM` determined when constructing the `Tp1tz` structure, and set by default to 0 for non-blocking communication.

3.2.3. Templates algebra package

Data distribution: As described in Sect. 3.2 the templates considered here are specific to scans and detectors. Consequently each MPI process stores only templates relevant to the scans and detectors assigned to it as a result of the overall time-domain distribution scheme described earlier. This data distribution is both memory and computation efficient.

Data structure: At the onset of the computation, each MPI process is assigned its specific set of template classes. For each template class, each process stores information only about the templates corresponding to the detector time streams it is processing.

Different template classes are defined with a C-struct called `TemplateClass`. The structure defines the time-domain interval for which this class is relevant as well as flags defining the type of templates in the template class as well as a stationary interval to which it corresponds.

For the templates defined implicitly, additional flags store the information about the parameters necessary to calculate the templates. For the templates defined in a compressed form, the structure defines the range of respective bins or pixels as well as two time-domain arrays, defining the mapping between the time indices and bin/pixel numbers and the corresponding weights. Given that templates are time-domain objects, they can have a significant memory footprint. Therefore, we implement two operating modes for the templates: the weights and bins can either be fully stored in memory at the moment of templates construction, if the processed data set is sufficiently small with a manageable memory footprint, or they can be computed on-the-fly while conducting the product operations with the templates matrix, for large data sets, hence trading speed for memory efficiency. Each MPI process holds an array of instances of `TemplateClass` in memory, and as such would have all necessary information to perform the time-domain filtering operations locally, without any need for communication with other processes.

Templates algebraic operations:

The main utility of the routines provided by this package is to allow the application of the filtering operator, \mathbf{F}_T , to a time-

domain vector. Three core routines form the basis for such an operation:

- **Templates projection in time-domain:** This operation generates a time-domain vector, \mathbf{t} , from a given set of templates amplitudes, \mathbf{x} , and a templates matrix, \mathbf{T} : $\mathbf{t} = \mathbf{T}\mathbf{x}$. Each process loops over the set of scans it is assigned, and for each scan, loads a set of templates weights, which typically depend only on boresight pointing parameters and therefore do not vary for individual detectors in the same scan. Higher order templates weights are constructed recursively, by using for example recursive relations between Legendre polynomials, or trigonometric functions. This allows us to avoid repeating the recursive computation for each detector time stream, which can significantly slow down the operation. Once the weights are stored in memory, we loop over the array of template classes: for each template we construct the array of bins on-the-fly, point to the already loaded weights, perform the series of products with the corresponding segment of templates amplitudes, and finish with freeing the memory segment occupied by the bins, and move on to the next template class. Once all the templates of the scan have been spanned, we remove the loaded weights from memory, and repeat the process for the next scan, until all scans are processed.

- **Time-domain projection in templates space:** This operation projects a time-domain vector, \mathbf{t} , in the templates amplitudes space, using a templates matrix \mathbf{T} : $\mathbf{y} = \mathbf{T}^T \mathbf{t}$. The operation is performed in a similar way to the one above, with loops over scans and template classes, and templates bins and weights constructed on-the-fly before computing the products with the relevant segments of data.

- **Templates kernel construction and inversion:** What we refer to as templates kernel here, is the block-diagonal matrix $\mathbf{T}\mathbf{W}^{-1}\mathbf{T}$. Each block of this matrix is constructed solely from the set of templates that apply to a single detector time stream, as long as the templates do not model any common detector modes. As such the construction of this kernel as well as its inversion are considerably simplified. We loop over the scans, and the detectors, and for each detector in a given scan, we perform a double nested loop on the template classes associated with the detector time stream, weights and bins are computed on-the-fly depending on the template class, and we compute the kernel entries with a series of products between the different templates weights. The inverse of the kernel is simply the inverse of each block. The blocks are inverted with a singular value decomposition (SVD), where singular values below a certain threshold, typically 10^{-12} , are set to zero, effectively regularizing the inversion. This is equivalent to an orthogonalization procedure, and therefore, spurious effects due to linear dependencies between templates are neutralized. Note that it is possible to explicitly exploit the inner structure of the kernel blocks

to speed-up both their inversion and construction. For example, we know that Legendre polynomials or cosine and sine harmonics over a full period are orthogonal by definition, and would therefore yield a diagonal kernel, therefore we can ignore some off-diagonal entries while constructing the kernel, and exploit the structure of the kernel when performing the inversion. As in the current applications the kernel inversion did not constitute a significant fraction of the total runtime such optimizations are not included in the current release and are left for future work.

3.3. The MIDAPACK parallel iterative sky estimator (MAP-PRAISER)

3.3.1. Mapping methods

As discussed in Section 2, Eq. (3) represents a broad class of plausible map-making approaches which result in an unbiased estimate of the sky signal. However, implementing it in its full generality is cumbersome and the resulting code likely to be rather inefficient. Instead, in the proposed framework we have implemented two highly optimized, specific cases of this general estimator, which taken together cover most of the cases of practical interest.

These two specific map-making methods are,

1. the case where the weights \mathbf{W}^{-1} are assumed to have the form of a Toeplitz block-diagonal matrix and the de-projection/weighting operator, \mathbf{F}_T , contains no templates, i.e., $\mathbf{F}_T = \mathbf{W}^{-1}$. The templates can be included as part of the generalized pointing matrix, see Eq. (7). We note that this method allows for a (nearly) maximum likelihood solution for the map in cases when the time-domain noise is piece-wise stationary and uncorrelated between different detectors. Indeed, the imposed structure of the weights is flexible enough to allow for a very good approximation to the actual inverse noise covariance as discussed in Section 2.4. The MAP-PRAISER package computes the weights using information derived directly from the collected data as in [49, 44, 60]. To do so, we consider that over each scan, every detector timestream is dominated by stationary noise with a power spectral density (PSD) of the form:

$$P(f) = \sigma_0^2 \left(1 + \left(\frac{f + f_0}{f_k} \right)^\alpha \right) \quad (21)$$

Assuming at first order that the data is only made of noise, we compute the average periodogram of each detector time stream, and fit it to the above model. The obtained PSD is then inverted and inverse Fourier transformed to obtain the inverse noise autocorrelation which is later smoothed with a gaussian window $\mathcal{N}(0, \lambda/2)$, and cut at the *a priori* specified half bandwidth λ . The smoothing avoids any ringing effects, and the resulting autocorrelation represents the first row of the corresponding Toeplitz block of the inverse noise covariance matrix. In general, even if efforts are made to use as good an approximation as possible, the weights do not have to reflect the actual inverse noise covariance. In such cases the

estimated map will be still unbiased but not optimal in the sense of ensuring the lowest possible uncertainty. Nevertheless, for simplicity, we refer to this approach hereafter as a maximum likelihood map-making; and,

2. a variant of the method which allows for arbitrary templates but assumes the weights to be diagonal. This method includes all the predefined templates provided in the software as described in [Appendix A](#). We call this method in the following a template map-making.

3.3.2. MAPPRaiser solvers

Linear systems can be solved using either direct or iterative methods. Direct methods typically compute some decomposition of the system matrix and consequently are time consuming scaling in general as a cube of the size of the system. Their advantage is that they give a (numerically) precise solution. Iterative methods form successive approximations that converge to the exact solution, and are particularly well-adapted to solving large sparse systems. They require only efficient routines for a product of the system matrix and an arbitrary vector, and are very efficient in terms of any extra storage. This first property is particularly relevant in the map-making case where the inverse of the system matrix cannot be typically explicitly computed due to both memory and computational cycle issues. Instead, the form of the system matrix allows for an efficient computation of a product of the matrix by a vector. Indeed, this can be implemented [7] by performing the operations from right to left. In the case of the maximum likelihood method, this involves a calculation of the matrix-vector product in the form,

$$\mathbf{P}^T \mathbf{W} \mathbf{P} \mathbf{x}, \quad (22)$$

which can be done by performing three subsequent operations solely on vectors: (1) unpointing; (2) a block-diagonal Toeplitz product; (3) pointing, all of which are implemented in the MIDAPACK library described earlier. In the case of the template map-making the respective product reads,

$$\mathbf{P}^T \mathbf{F}_T \mathbf{P} \mathbf{x}, \quad (23)$$

and again requires three operations performed one after another with the Toeplitz product above replaced now by an application of the deprojection/weighting operator, \mathbf{F}_T . This in turn requires a precomputation of the inverse kernel and then merely a series of template pointing and depointing operations available in MIDAPACK.

MAPPRaiser provides two iterative solvers for the map-making equation. Both are based on the Conjugate Gradient approach. These are Krylov projection methods [14] suitable for symmetric (Hermitian) positive definite matrices adapted to the specificity of the map-making problem. We discuss them below.

We note that other iterative solvers could be and have been applied in this context, e.g. [10], and some alternatives are actively investigated [22, 8].

Preconditioned conjugate gradient solver

The first one is the Preconditioned Conjugate Gradient (PCG),

which is used to solve Symmetric Positive Definite (SPD) systems. This solver has already been extensively used in the literature to solve the map-making problem [34, 61, 7, 53]. We refer the reader to these papers for the technicalities of the method. Below we illustrate the algorithm of the PCG solver, where \mathbf{M} is representing a preconditioner. The objective of applying a preconditioner is to reduce the condition number of the system matrix. In MAPPRaiser we implemented two kinds of preconditioners. These are the Block Diagonal preconditioner (*BD*) and the two level preconditioner (*2lvl*). The block-diagonal preconditioner defines the current standard in the field. It is easy to compute and in many cases performs already very well. The two-level preconditioner has been proposed and demonstrated for the map-making problem in [53] and extended to component separation in [36]. It is more involved and requires a precomputation. To the best of our knowledge the MAPPRaiser implementation of this preconditioner is the very first and only existing implementation in actual map-making software.

```

Compute  $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$ ,  $\mathbf{z}_0 = \mathbf{M}\mathbf{r}_0$  and  $\mathbf{p}_0 = \mathbf{z}_0$ ,
for  $k = 0, 1, \dots, k_{max}$  do
     $\alpha_k = \langle \mathbf{r}_k, \mathbf{z}_k \rangle / \langle \mathbf{A}\mathbf{p}_k, \mathbf{p}_k \rangle$ 
     $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k$ 
     $\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{A}\mathbf{p}_k$ 
    if  $\|\mathbf{r}_{k+1}\|_2 < \varepsilon \|b\|_2$  then stop
     $\mathbf{z}_{k+1} = \mathbf{M}\mathbf{r}_{k+1}$ 
     $\beta_k = \langle \mathbf{r}_{k+1}, \mathbf{z}_{k+1} \rangle / \langle \mathbf{r}_k, \mathbf{z}_k \rangle$ 
     $\mathbf{p}_{k+1} = \mathbf{z}_{k+1} + \beta_k \mathbf{p}_k$ 
end

```

Algorithm 1: The PCG algorithm for the system $\mathbf{A}\mathbf{x} = \mathbf{b}$ with a preconditioner \mathbf{M} [47].

Enlarged conjugate gradient solver

The second iterative solver is the Enlarged Conjugate Gradient (ECG) [16]. In this approach the matrix \mathbf{A} is partitioned into \mathcal{N} subdomains. Then the unknown vector \mathbf{x} is split into t vectors, $\mathbf{X}^{1 \leq i \leq t}$, such that we can still retrieve the original vector by summing them: $\mathbf{x} = \sum_{i=1}^t \mathbf{X}^{(i)}$. The parameter t is called the enlarging factor, and the method is in practice not sensitive to the choice of the splitting scheme [18]. Following these considerations, the method can be derived in a similar way to the standard CG solver, replacing the residuals and search directions by $\mathcal{N} \times t$ matrices, and the optimal step which is a scalar in CG becomes a $t \times t$ matrix. **The algorithm is given as Algorithm 2. Two different variants are implemented: Orthomin and Orthodir. They are equivalent in well behaved cases, however, the Orthodir method is more robust than Orthomin in cases where $\mathbf{Z}_k^T \mathbf{A} \mathbf{Z}_k$ is (nearly) singular. However, the Orthodir method is also more expensive than Orthomin as it requires twice as much memory and flops [18]. The square root of $\mathbf{Z}_k^T \mathbf{A} \mathbf{Z}_k$ is computed using a Cholesky factorization as it is faster than alternatives. The Cholesky decomposition can be preceded by a QR factorization of \mathbf{Z}_k to avoid break downs of the Orthomin method in singular**

cases [31]. This is not implemented in the present software release however.

Compute $\mathbf{r}_0 := \mathbf{b} - \mathbf{A}\mathbf{x}_0$, and split it into a matrix with t columns, \mathbf{R}_0^e , s.t. $\mathbf{r}_0 = \sum_{i=1}^T \mathbf{R}_0^{e(i)}$,
 $\mathbf{Z}_1 = \mathbf{M}\mathbf{R}_0^e$,
for $k = 1, 2, \dots, k_{max}$ **do**
 $\mathbf{P}_k = \mathbf{Z}_k (\mathbf{Z}_k^T \mathbf{A} \mathbf{Z}_k)^{-1/2}$
 $\boldsymbol{\alpha}_k = \mathbf{P}_k^T \mathbf{R}_{k-1}$ $\triangleright t \times t$ matrix
 $\mathbf{X}_k = \mathbf{X}_{k-1} + \mathbf{P}_k \boldsymbol{\alpha}_k$
 $\mathbf{R}_k = \mathbf{R}_{k-1} - \mathbf{A} \mathbf{P}_k \boldsymbol{\alpha}_k$
 if $\|\sum_{i=1}^T \mathbf{R}_k^{(i)}\|_2 < \varepsilon \|\mathbf{b}\|_2$ **then stop**
 Orthomin:
 $\mathbf{Z}_{k+1} = (\mathbf{I} - \mathbf{P}_k \mathbf{P}_k^T \mathbf{A}) \mathbf{M} \mathbf{R}_k$
 Orthodir:
 $\mathbf{Z}_{k+1} = (\mathbf{I} - \mathbf{P}_k \mathbf{P}_k^T \mathbf{A} - \mathbf{P}_{k-1} \mathbf{P}_{k-1}^T \mathbf{A}) \mathbf{M} \mathbf{A} \mathbf{P}_k$
end
 $\mathbf{x}_k = \sum_{i=1}^T \mathbf{X}_k^{(i)}$.

Algorithm 2: The ECG algorithm for the system $\mathbf{A}\mathbf{x} = \mathbf{b}$ with a preconditioner \mathbf{M} [18].

The rate of convergence of the ECG is given by the following result demonstrated in [18]: if \mathbf{x}_k is the approximate solution given by the Enlarged Conjugate Gradient with an enlarging factor t at step k , and \mathbf{x}_* the true solution then we have

$$\|\mathbf{x}_k - \mathbf{x}_*\|_{\mathbf{A}}^2 \leq C \left(\frac{\sqrt{\kappa_t} - 1}{\sqrt{\kappa_t} + 1} \right)^{2k}, \quad (24)$$

where $\kappa_t = \lambda_n / \lambda_t$, the ratio of the largest eigenvalue of \mathbf{A} , and its t smallest eigenvalue, and C a constant independent of k . As such the solver behaves as a deflated-CG [48], and converges faster than CG. However, the rate of convergence is not a sufficient metric to assess the performance of the method. In practice, a more relevant figure of merit is instead time to solution. In the case of ECG the expanded map space leads to an increase of the number of operations required on each iteration of the solver by the enlarging factor t , therefore a trade-off needs to be found between the gain in the rate of convergence and the linear increase in the algorithm's complexity. In the scaling experiments we conducted which results are given in Sect. 4.1.3, we adopt a naive implementation, where the product of the matrices \mathbf{A} and \mathbf{X} is computed by looping on the columns of the enlarged map, and performing the standard CG $\mathbf{A}\mathbf{x}$ operation for each column. This baseline implementation of the solver demonstrates the improvement in the number of iterations with increasing the enlarging factor as theoretically predicted. The gain in terms of the number of iterations is however sublinear and does not compensate for the increased cost of each iteration. Consequently, the overall run time increases. We expect this to change with an improved approach to performing the product of the system matrix and expanded map. One particularly promising avenue is to capitalize on GPU-accelerators. This is left for future work. As is, the main interest of the ECG solver is in the fact that it does not require sophisticated preconditioning and is nevertheless very efficient in the cases when the system matrix includes a number of small eigenvalues.

3.3.3. Preconditioners for PCG

The objective of preconditioning is to improve the convergence and stability of the iterative solvers. The preconditioner, \mathbf{M} , can be thought of as an approximate to the inverse of the system matrix, \mathbf{A} . In preconditioned solvers, instead of the original problem we solve a modified system with the system matrix replaced by a product of the system matrix and the preconditioner, i.e., $\mathbf{M}\mathbf{A}$. Unlike the actual inverse of the system matrix the preconditioner should be easy to compute, and the preconditioned system matrix, $\mathbf{M}\mathbf{A}$ should be better conditioned and its eigenspectrum more clustered. Both of these properties make the system more stable and speed-up the convergence. The MAPPRASER package offers two preconditioning options.

Block-Jacobi preconditioner:

The first preconditioner is the standard Block-Jacobi preconditioner defined as,

$$\mathbf{M}_{BD} = (\mathbf{P}^T \text{diag} \mathbf{N}^{-1} \mathbf{P})^{-1}, \quad (25)$$

where the inner most matrix denotes a diagonal matrix given by the diagonal of the inverse noise covariance matrix. This preconditioner is block-diagonal with square diagonal blocks sizes of which are determined by the number of non-zero elements per row of the pointing matrix. The construction of this preconditioner is done in two steps. First, for each pixel we sum over the (non-zero) elements of all the rows of the pointing matrix corresponding to this pixel while weighting them by the diagonal noise weights. On the second step, we invert each of the blocks using a direct method. On this second step, we also evaluate a condition number of each block, and use it as a criterion to excise pixels with ill-conditioned blocks from the map-making.

Two level preconditioner:

The first level of the two-level preconditioners considered in MAPPRASER is based on the deflation technique. Its purpose is to suppress (deflate) an unwanted subspace of a matrix, so as a result it is contained in its null space. In the context of preconditioning this is ideally the subspace which contains all the smallest eigenvalues of the system matrix. The second level then consists of adding an extra contribution to the deflated matrix in order to shift the eigenvalues of the deflation subspace to unity. There are different variants of the two-level preconditioners which have been shown to be efficient in different applications. Here, we follow the proposal of [53] and focus on the so called ‘‘Adapted Deflation Variant 1’’ (A-DEF1) option in [55]. The preconditioner is defined as follows,

$$\mathbf{M}_{2lvl} = \mathbf{M}_{BD} (\mathbf{I} - \mathbf{A} \mathbf{Q}) + \mathbf{Q}, \quad (26)$$

$$\mathbf{Q} = \mathbf{Z} (\mathbf{Z}^T \mathbf{A} \mathbf{Z})^{-1} \mathbf{Z}^T, \quad (27)$$

where \mathbf{Z} is a deflation subspace matrix. Its columns span a subspace to be deflated. This is a tall and skinny matrix with

a limited number of columns each of which are pixel-domain vectors. The number of columns of \mathbf{Z} defines the volume of the deflation subspace. \mathbf{A} stands for our system matrix. The first term in Eq. (26) defines the first level of the preconditioner. Indeed, we have that $(\mathbf{I} - \mathbf{A}\mathbf{Q})\mathbf{A}\mathbf{Z} = 0$. The second term is the second level correction.

A key element determining the performance of the two level preconditioner is a construction of the deflation subspace matrix, \mathbf{Z} . The MAPPRASER implementation of the PCG solver with the two level preconditioner accepts any user provided matrix \mathbf{Z} . In addition, MAPPRASER includes functionality which allows for calculating the deflation matrix based on two different approaches. We describe those in the following.

We note that the two-level preconditioner is in principle a dense, square matrix with the dimension determined by the number of the observed sky pixels. For this reason we never explicitly compute it as in many applications this would be prohibitive both from the point of view of the required computations and the computer memory. Instead, for any available deflation matrix, \mathbf{Z} , we only precompute matrix $\mathbf{Z}' \equiv \mathbf{A}\mathbf{Z}$, the so-called coarse matrix, $\mathbf{E} \equiv \mathbf{Z}^T \mathbf{A}\mathbf{Z} = \mathbf{Z}^T \mathbf{Z}'$, and its inverse. Computing \mathbf{Z}' is computationally demanding as we have to apply the system matrix to as many vectors as the number of columns of the matrix \mathbf{Z} . As each such product is also a dominant computational cost of each iteration of the PCG solver, this computation is as costly as performing as many iterations as the volume of the deflation subspace. As we discuss later this cost can be however effectively hidden in some of the procedures aiming at the calculation of the deflation matrix, \mathbf{Z} . This is indeed the case in one of the specific constructions implemented in MAPPRASER. The matrix \mathbf{E} is then computed by calculating $\mathbf{Z}^T \mathbf{Z}'$, which entails negligible cost. The coarse matrix is in most of the cases a small, dense matrix. We therefore can compute its inversion with help of standard, highly efficient, dense linear algebra routines. If it is known a priori that the columns of \mathbf{Z} are all linearly independent, and therefore the coarse operator is positive definite this could be done very efficiently using the Cholesky decomposition. Otherwise, one has to use a pseudo-inverse, which relies on the calculation of the full eigenspectrum decomposition of the coarse matrix and is therefore appropriately more costly. In our implementation we opted for the second approach as the calculation of the inverse is in any case subdominant. We use the *dgels* LAPACK routine in the MKL implementation for this purpose. All these operations have to be performed only once at the beginning of the solution. In addition we need to precompute the block-Jacobi preconditioner following the procedure outlined earlier at negligible cost.

The preconditioner has to be applied once at every step of the PCG iterations. This involves a computation of a product of the preconditioner and some arbitrary pixel-domain vector. For this purpose we represent the preconditioner in Eq. (26) as,

$$\mathbf{M}_{2vl} = \mathbf{M}_{BD} - \mathbf{M}_{BD} \mathbf{Z}' \mathbf{E}^{-1} \mathbf{Z}^T + \mathbf{Z} \mathbf{E}^{-1} \mathbf{Z}^T. \quad (28)$$

We apply the preconditioner by multiplying all factors of each of its terms sequentially, from right to left to a vector capitalizing on precomputed and stored in memory objects. All these

objects are either pixel-domain or deflation-subspace objects and the required operations are respective matrix-vector products which require at most $\mathcal{O}(\mathcal{N}_{pix} \dim \mathbf{Z})$ floating point operations. Unless $\dim \mathbf{Z}$ is very large this extra cost is usually subdominant to the cost of applying the system matrix to a vector. Each column of the matrix \mathbf{Z} is distributed between the MPI processes in the same way as any other map. This saves the memory but requires one extra global communication call of the `allreduce` type per iteration. The size of the reduced data is given by the dimension of the deflation space and the time overhead therefore negligible. Consequently, the total cost of the application of the preconditioner is subdominant and the application of the system matrix to a vector remains the dominant cost of a single iteration. Therefore, the time of a single iteration is essentially the same for the two-level and block-Jacobi preconditioners.

MAPPRASER provides two approaches to calculating the deflation subspace matrix, \mathbf{Z} . These are,

- **2vl a priori construction**

This construction is based on the structure of the data set, which is assumed to be divided into a number of time intervals (such as noise stationary intervals corresponding to the blocks of the noise covariance, \mathbf{N}). This construction is applicable only if total intensity is part of the data model. In this case we calculate the number of observations of pixel p_i during the j -th interval, denoted as s_i^j . s_i is then the total number of observations of pixel p_i . The (i, j) element of the deflation matrix, \mathbf{Z} , which corresponds to the total intensity signal is given by the fractional number of observations of the pixel p_i in the j -th interval, i.e., $\mathbf{Z}_{ij}^{int} = s_i^j / s_i$. Consequently every such row of \mathbf{Z}^{int} represents a partition of the unity, since $\sum_j \mathbf{Z}^{int}(p, j) = 1$. The rows corresponding to Q and U Stokes parameters are set to 0 [53].

- **2vl a posteriori construction**

In this construction we attempt to estimate the relevant eigenpairs of the system matrix, \mathbf{A} . We denote these as (θ, \mathbf{y}) standing for an eigenvalue and the corresponding eigenvector, respectively. We use the Lanczos algorithm [14, 47]—given explicitly below in algorithm 3—to calculate an m -by- m tridiagonal matrix, \mathbf{T}_m and an \mathcal{N}_{pix} -by- m column orthonormal matrix \mathbf{V}_m such as,

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m \mathbf{T}_m. \quad (29)$$

The Lanczos algorithm is an iterative procedure and m defines the number of the iterations. It is bound to be smaller than the size of the system matrix. In our case, $m \ll \mathcal{N}_{pix}$. The algorithm is applicable to any symmetric positive-definitive (SPD) matrix and requires only matrix-vector products. Here, these are implemented using the same routines as used by the PCG algorithm. Given the structure and the limited size of the matrix \mathbf{T}_m , its eigen-structure can be straightforwardly calculated. Given Eq. (29), if (θ, \mathbf{x}) is an eigenpair of \mathbf{T} then $(\theta, \mathbf{V}_m \mathbf{x})$

is a corresponding eigenpair of \mathbf{A} . Therefore $\mathbf{V}_m \mathbf{x}$ define the columns of the deflation matrix, \mathbf{Z} . We take all the precomputed vectors, given that the cost of applying the two level preconditioner is subdominant.

We note that the Lanczos algorithm computes internally products of the system matrix and columns of the matrix \mathbf{V}_m , which are usually discarded. We store those, i.e., the matrix $\mathbf{A} \mathbf{V}_m$, together with the matrix \mathbf{V}_m and use them to compute directly the matrices $\mathbf{A} \mathbf{Z}$ and \mathbf{Z} . This saves significant precomputation time needed to construct the preconditioner as discussed earlier.

This procedure is a special case of a more general procedure proposed in [53] and employed in [45], which uses the so-called GMRES algorithm, which is more numerically stable and applicable to a general matrix. We find the Lanczos procedure to be more computationally efficient and sufficiently robust in our applications. The main difference between the two approaches is that our procedure is applied to the system matrix and not to the system matrix preconditioned by the block-Jacobi preconditioner as in [53, 45].

```

Compute  $\mathbf{r}_0 = \mathbf{b} - \mathbf{A} \mathbf{x}_0$ ,  $\beta_1 = \|\mathbf{r}_0\|_2$ , and  $\mathbf{v}_1 = 1/\beta_1 \mathbf{r}_0$ 
for  $k = 1, 2, \dots, m$  do
     $\mathbf{w}_k = \mathbf{A} \mathbf{v}_k - \beta_k \mathbf{v}_{k-1}$  (if  $k = 1$  set  $\beta_1 \mathbf{v}_0 \equiv \mathbf{0}$ )
    store  $\mathbf{A} \mathbf{v}_k$ 
     $\alpha_k = \langle \mathbf{w}_k, \mathbf{v}_k \rangle$ 
     $\mathbf{w}_k := \mathbf{w}_k - \alpha_k \mathbf{v}_k$ 
     $\beta_{k+1} = \|\mathbf{w}_k\|_2$ 
    if  $\beta_{k+1} < \varepsilon$  then  $m := k$  and break loop

     $\mathbf{v}_{k+1} = 1/\beta_{k+1} \mathbf{w}_k$ 
end
return  $\mathbf{T}_m = \text{tridiag}(\beta_i, \alpha_i, \beta_{i+1})$ ,  $\mathbf{V}_m = [\mathbf{v}_1, \dots, \mathbf{v}_m]$ ,
and  $\mathbf{A} \mathbf{V}_m$ .

```

Algorithm 3: The Lanczos algorithm [47] for the matrix $\mathbf{A} = \mathbf{P}^T \mathbf{W}^{-1} \mathbf{P}$. The threshold, ε , is set to machine precision, and was never reached in our applications.

We compare advantages of different solvers in detail in the following sections. To quantify the quality of the current iterate, \mathbf{x}_i , we first define a residual, which for a linear system, $\mathbf{A} \mathbf{x} = \mathbf{b}$, is given by,

$$\mathbf{r}_i = \|\mathbf{b} - \mathbf{A} \mathbf{x}_i\|. \quad (30)$$

If $\mathbf{x}_0 = \mathbf{0}$, as it is always the case in our test runs, then $\mathbf{r}_0 = \mathbf{b}$. For the convergence we require that the relative residual, $\|\mathbf{r}_i\|/\|\mathbf{b}\| = \|\mathbf{r}_i\|/\|\mathbf{r}_0\|$, is smaller than some positive threshold value, ϵ . A typical value of ϵ used in the map-making applications, and therefore also adopted here, is $\epsilon = 10^{-6}$. A key performance metric is consequently a number of the PCG iterations needed to convergence. We denote it, hereafter, by n_{iter} with a superscript added to highlight the preconditioner choice.

As a rule of thumb, we note that each iteration of the PCG algorithm is as costly as applying the system matrix to a single column of the deflation matrix, \mathbf{Z} , and as costly as a single iteration of the Lanczos procedure. Solving the map-making problem using the block-Jacobi preconditioner requires n_{iter}^{BD} products of the system matrix times a vector. For two-level preconditioners, this amounts to $n_{iter}^{2lvl} + \dim \mathbf{Z}$ where the second term is due to the Lanczos iterations for the a posteriori preconditioner and the application of the system matrix to the deflation matrix in the case of the a priori preconditioner. All the other costs are then typically subdominant. If the same system or similar systems (see [45]) need to be solved multiple times the respective costs are: $n_{solves} n_{iter}^{BD}$ and $n_{solves} n_{iter}^{2lvl} + \dim \mathbf{Z}$. Obviously the key parameter is the number of iterations required to reach desired precision for different types of the preconditioners and different dimensions of the deflation space. It can be argued on theoretical grounds that $n_{iter}^{BD} \geq n_{iter}^{2lvl}$ independently of the deflation matrix [53]. This is because the two-level preconditioner cannot make worse the condition number of the preconditioned system matrix. The gain can be however very significant if \mathbf{Z} encompasses the smallest eigenvalues of the system matrix. In such cases, the precomputation time can be offset very quickly if multiple, similar map-making problems have to be solved, leading to substantial overall net gain for the full computation. In fact, whenever the number of iterations per solve is known, we can estimate the minimal number of the map-making solves required to reach this regime. We investigate this in detail in the next section.

We also note that the deflation matrix computed with the Lanczos algorithm encodes useful information about the system matrix and therefore the error covariance of the estimated map. This information can be useful on the subsequent steps of the analysis.

4. Applications

In this section, we present some of the applications that we have applied our framework to, and demonstrate the computational and scaling performance of the software. We emphasize that it is not our goal here to discuss the scientific impact of the obtained results. This is left for future work.

We consider two cases: one with partial sky coverage as expected from ground-based experiments and another with a full sky coverage as expected from a satellite mission.

We run the simulations on two different HPC platforms:

- **Cori:** This is a Cray XC40 machine based at NERSC². It can reach a peak performance of about 30 petaflops. It has two partitions, the first one is made of 2,388 Intel Xeon “Haswell” processor nodes, and the second one is made of 9,688 Intel Xeon Phi “Knight’s Landing” (KNL) nodes. In our work we use the “Haswell” partition. Each Haswell node has two sockets, each of them populated with a 2.3 GHz 16-core Haswell processor, therefore allowing the use of 32 cores per node. Each core supports 2

²National Energy Research Scientific Computing Center: <https://www.nersc.gov>

hyper-threads, and has two 256-bit-wide vector units, allowing a theoretical peak of 36.8 Gflops/core. In terms of available memory, each node has a 128 GB DDR4 2133 MHz memory.

- **Joliot-Curie:** This is a Bull Sequana machine based at TGCC³. The supercomputer features four partitions, however in our work we use a single one; the “SKL” partition, which is made of 1,656 nodes, each populated with 2.7 GHz Intel Skylake 8168 bi-processors, each processor hosting 24 cores, for a total of 79,488 cores and a total peak performance of 6.86 Pflops. Each node has a 192 GB DDR4 memory.

4.1. Case of a ground-based experiment

4.1.1. Simulations description

We simulate CMB observations using the TOAST⁴ framework. In this section we describe the results obtained with TOAST simulations of a typical ground-based experiment with specifications roughly following those of the Simons Observatory (SO) [13, 2]. We simulate two sets of observations each comprised of data collected by roughly 6,000 detectors of an SO-like, Small Aperture Telescope (SAT) operating at 150 GHz,

- **S1:** corresponds to the observation of a $20 \times 10 \text{ deg}^2$ sky patch. The observations consist of 33 constant elevations scans (CESSs), each between 45 min and 1 h long, totaling to about 24 hours of observation. The sampling rate is set to 132 Hz, yielding a total of $O(10^{11})$ time samples, i.e., ~ 10 M time samples per detector. **The time-domain objects required for this run have a total size of about ~ 5 TB worth of memory. During the data staging process from TOAST to MAPRAISER, the different detector data are concatenated at the level of each MPI process, thereby further constraining the minimum memory requirements as the high-water mark reaches around twice the total size at about ~ 10 TB.** We use this simulation to evaluate the scaling performance of the solvers.
- **S2:** corresponds to the observation of a smaller sky patch of $10 \times 10 \text{ deg}^2$. The observations consist of 95 CESSs, spanning 5 consecutive days, each about 15 min long, for a total of 24 hours of observation. The sampling rate is set to 37 Hz, yielding a total of $O(10^{10})$ time samples, i.e., ~ 1 M time samples per detector. **The memory requirement for this run is about 10 times lower than the previous case at ~ 1 TB.** We use this simulation to study the convergence of the two-level preconditioners and to validate the templates marginalization procedure.

The scanning speed is 1 deg /s in both simulations, and the half wave plate rotation frequency is set to 2 Hz. The input map of the simulations is pixelized in HEALPix⁵ format [15, 62] with a resolution corresponding to $n_{\text{side}} = 512$, and includes a CMB realization from best-fit Planck cosmology [40] and a Gaussian lensing potential. The resulting sky patch from **S1** counts around $\sim 250,000$ pixels, while the one from **S2** counts about $\sim 140,000$ pixels. In addition, we simulate instrumental $1/f$ noise and atmosphere. The instrumental noise has a characteristic knee frequency of 50 mHz (f_k in Eq. (21)), a slope of ~ -1 (α in Eq. (21)), and a detector NET of $400 \mu\text{K} \sqrt{s}$. The native atmospheric simulation built in the TOAST framework is based on a 3d turbulence model [11] sensitive to parameters such as wind properties, weather conditions, and turbulence characteristics, calibrated on real data from experiments in Chile such as POLARBEAR [28] and ACT [58]. A 3d volume is generated and propagated according to the wind direction and speed as it is scanned by the detectors, thereby generating correlations between time samples as well as cross-correlations across the focal plane between detectors.

4.1.2. Strong scaling tests and collective communication performance

We perform several strong scaling tests, using the **S1** simulation on Cori and Joliot-Curie to demonstrate the ability of the code to run on a large number of cores and handle the data reduction of a large data set. Given that the Skylake processors (Joliot-Curie) are more powerful than the Haswells (Cori) we expect to see longer runtimes with the second machine for the same set of the sky reconstruction parameters. In the runs presented here we assume the maximum-likelihood map-making approach with a Toeplitz-structured noise covariance and a half-bandwidth set to 2,048 samples. This corresponds to a time-domain correlation length of ~ 15 s, given the adopted sampling rate. The solver converges in around $\sim 1,200$ iterations for the case considered. Given that the two machines had different versions of TOAST installed, the noise simulations were not exactly the same resulting in slight changes in the total number of iterations to convergence between the two machines. However, these changes do not affect the computational performance metric adopted here, which is the mean time per iteration. We choose this metric because it depends only on the half-bandwidth and the total size of the problem which are exactly the same for all the tests performed on both machines. The results are shown in figure 3. **For reference, the total wall clock time is shown in figure 4, for the Butterfly case executed on Cori, along with the best-fit scaling power law.**

The Ring runs do not scale properly as we increase the number of MPI tasks handling the problem, and completely diverge when using a large number of processes, therefore we exclude it from any further analysis. The customized Butterfly scheme exhibits better mean times per iteration compared with the standard MPI_Allreduce on the Cori machine, particularly as we increase the number of MPI processes. On the

³Très Grand Centre de calcul du CEA (The French Alternative Energies and Atomic Energy Commission Computing Center): <http://www-hpc.cea.fr/fr/complexe/tgcc.htm>

⁴<https://github.com/hpc4cmb/toast>

⁵<http://healpix.sourceforge.net>




Figure 3: Strong scaling of the mean time per iteration comparing 5 sets of runs: (1) runs executed on Cori at NERSC using the `MPI_Allreduce` communication scheme; (2) runs executed on Cori at NERSC using the Butterfly communication scheme; (3) runs executed on Cori at NERSC using the Ring communication scheme; (4) runs executed on Joliot-Curie at TGCC using the `MPI_Allreduce` communication scheme; (5) runs executed on Joliot-Curie at TGCC using the Butterfly communication scheme.




Figure 4: Total wall clock time of the runs executed on Cori at NERSC using the Butterfly communication scheme.

Joliot-Curie machine, the Butterfly scheme shows similar or slightly better performance than the standard `MPI_Allreduce` depending on the number of MPI tasks. To get a more quantitative view of the scaling performance, we measure the scaling efficiency given by

$$\varepsilon = \frac{\delta t_1}{N \times \delta t_N}, \quad (31)$$

where δt_1 is the reference mean time per iteration to which the scaling efficiency is compared, theoretically this should correspond to the iteration time in a serial run (with a single process), and δt_N is the mean time per iteration for a run with N MPI pro-

cesses. In practice, the size of the problem does not fit into a single node and thus we cannot perform serial runs, hence we set the reference to be the run with the lowest number of MPI processes, and N to the factor by which we increase the number of processes with respect to the reference. In these tests, this corresponds to 396 MPI processes for the `MPI_Allreduce` tests on Joliot-Curie and to 512 MPI processes for the remaining cases. The change of reference point should not affect our conclusions: given the proximity between the two references one can reasonably assume that the scaling between the two is close to 100%. Results are shown in figure 5.




Figure 5: Strong scaling efficiency comparing 4 sets of runs: (1) runs executed at the Cori machine at NERSC using the `MPI_Allreduce` communication scheme; (2) runs executed on Cori at NERSC using the Butterfly communication scheme; (3) runs executed on Joliot-Curie at TGCC using the `MPI_Allreduce` communication scheme; (4) runs executed on Joliot-Curie at TGCC using the Butterfly communication scheme.

The Butterfly scheme performs better than `MPI_Allreduce` by a few percents on both machines. The scaling efficiency at ~ 16 K processes reaches around 40% for the Butterfly scheme, while it is below 30% for `MPI_Allreduce` on Cori. Nevertheless, in the considered cases and for the number of MPI processes used, the communication overhead is small as compared to the computation time, and the current implementation of the Butterfly scheme is restricted to numbers of processes which are powers of 2. While the algorithm can be straightforwardly extended, this requires more work, and is left for future extensions of the software. Instead, given the satisfactory performance of the standard `MPI_Allreduce` scheme, we select it as the default option for the rest of this work.

We point out that both schemes, `MPI_Allreduce` and Butterfly, show good scaling for these types of applications on both machines, demonstrating good overall portability of the software.

4.1.3. Solvers and preconditioners performance

Enlarged CG:

In figure 6, we show a selection of tests of the ECG solver for different enlarging factors. **We use the Orthomin method, as the case studied is well behaved and does not warrant the use of the more involved Orthodir method.** The data used here constitute a 1 h subset of the S1 simulation. This corresponds to a single constant elevation scan. We also exclude the atmospheric noise from the input, and instead artificially increase the instrumental noise f_{knee} to 1 Hz. We assume a half bandwidth, $\lambda = 2^{18}$, corresponding to a correlation length of about 33 minutes. This configuration results in a degraded conditioning of the system matrix showcasing the ability of the ECG solver to deal with such cases. We execute multiple ECG runs varying the enlarging factor from 1 to 16, and compare it to a standard PCG run. The comparison between the PCG convergence and the ECG one with enlarging factor equal to unity provides the basis for the validation of the latter: the theoretical expectation is that the two curves should coincide as it is indeed the case in figure 6. In addition to this, we observe that the number of iterations required to reach convergence (set in the present case by a tolerance parameter of 10^{-6}) is monotonically decreasing with the enlarging factor, in particular we gain by about a factor 3 with an enlarging factor of 16 compared to standard PCG. This behavior is also theoretically expected from Eq. (24). However as previously mentioned this does not translate automatically to a net gain in the run time. For that purpose, one should ensure that the implemented algebraic routines used by the solver are sufficiently optimized to compensate for the increase in the cost of each iteration. This is not the case for the naive implementation we used in this test. Our results suggest that this solver could be indeed an interesting, precomputation-free alternative for the map-making problem, in particular given the advent of new computer architectures, motivating future work directly targeting the optimization of the underlying operations.




Figure 6: Dependence of ECG iterations on the enlarging factor, using 1 hour of the SAT simulated data.

PCG with two-level preconditioners:

In the present subsection, we compare the performance of the three preconditioners offered by MAPPRASER. For this purpose we run an adapted version of the S2 simulation on Cori, where we cut down the number of detectors by a factor 14, i.e., to 430 detectors, to run on a small number of processes (64 MPI tasks) and we scale down the noise to emphasize the sky signal and study the angular convergence as well. The convergence properties should not be affected by such changes as they don't depend on the noise normalization in general. We use the block-diagonal (BD) preconditioner, the *a priori* and *a posteriori* two-level preconditioners varying the size of the deflation space, and then compare the performance of the three preconditioners in terms of the number of iterations and wall clock time necessary to converge (figures 7, 8 and 13). In addition, we also show the TT and EE angular power spectra of the reconstructed maps at all iterations to assess the convergence of the different angular scales (figures 9 and 10). The BB spectra are noise dominated in this regime and therefore not considered here. The power spectra are estimated using the NaMaster library⁶ [3], correcting for the mask, however not for the noise bias. The *a priori* preconditioner tests summarized in figure 7 show only marginal gain on the number of iterations to convergence, demonstrating that, at least in the case studied, there is no significant benefit in using the *a priori* preconditioner over the standard BD preconditioner. In fact, the *a priori* preconditioner is included mostly for historical reasons. Indeed, the main early interest in this preconditioner was driven by the possibility of saving on the precomputation time needed to construct \mathbf{Z} [53, 17], which is essentially guessed and therefore computationally cheap. However, constructing the preconditioner still requires a computation of the $\mathbf{A}\mathbf{Z}$ matrix. This is often as costly as the computation of \mathbf{Z} itself. As we discuss in the following in the novel *a posteriori* construction we propose in this work, the cost of the former can be hidden in the calculation of the latter. Consequently, this new class of the *a posteriori* constructions is not only performing better but also requires as much precomputations as the *a priori* technique. Notwithstanding these considerations, the *a priori* construction may still be of interest, in particular, in the cases where there is a limited number of nearly singular eigenvectors spanning a known low-dimensionality subspace. In the following we focus solely on the convergence properties of the *a posteriori* preconditioner.

Figure 8 shows that, for the *a posteriori* preconditioner, the number of iterations required for convergence monotonically decreases with the size of the deflation subspace, $\dim \mathbf{Z}$. For any sufficiently high $\dim \mathbf{Z}$ (here $\gtrsim 128$), the *a posteriori* two-level preconditioner performs consistently better than the *a priori* two-level preconditioner. This is expected since the former uses a more involved procedure to construct the deflation subspace, yielding better estimates of the eigenvectors of the system matrix, once the Lanczos procedure runs for a sufficient number of iterations.

⁶<https://github.com/LSSDESC/NaMaster>




Figure 7: Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the *a priori* two-level preconditioned CG, for different dimensions of the deflation subspace.




Figure 8: Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the *a posteriori* two-level preconditioned CG, for different dimensions of the deflation subspace.

Additionally, the convergence gain as shown in the figure mainly stems from a sharp drop of the residual at the very first CG steps of the solver. This can be understood by examining the convergence of the different angular scales in figures 9 and 10. The colorbar encodes the wall clock time from the start of the solver until convergence is reached, not including pre-computation time. The colors are in log scale and the curves correspond to the power spectra at different iterations. Comparing the various TT and EE plots, we can see that the higher the size of the deflation subspace, the wider the range of the angular scales recovered from the first CG step. As a matter of fact, the precomputation allows the solver to compress information encoded in the system matrix, \mathbf{A} , to a small set of modes contained in the deflation matrix, \mathbf{Z} . When running the PCG steps, that information allows to recover the smallest angular

scales immediately in the first step, while subsequent steps are spent correcting the largest angular scales. In the extreme cases, $\dim \mathbf{Z} \geq 512$, we can see that the power spectra of the different iterations coincide with each other, and therefore that we practically recover the full range of angular scales of the solution from the first step. Indeed, in the case of $\dim \mathbf{Z} = 2048$, the formal convergence is reached in two steps, meaning that we compress practically all of the information of the system matrix, \mathbf{A} , in the relatively low number of modes of the deflation subspace. The *a posteriori* two level preconditioner thus not only improves convergence, it also gives us some insight into the noise covariance structure of the computed maps. However, for $\dim \mathbf{Z} = 512$ or 1024, while the spectra reach their final shape in at most a couple of steps, e.g., figures 9 and 10, the iterations continue for nearly 200 or 100, respectively, steps more in order to attain the formal convergence. This is related to the fact that the formal residual defined in Eq. (30) is weighted towards the noisiest pixels often located at the outskirts of the observed sky patch and which matter little in the power spectrum computation. This suggests that if the power spectra are of ultimate interest for the analysis, a more suitable convergence metric could be proposed leading to significant performance gain. The software discussed here is straightforwardly extensible to allow for such tailored, user-defined convergence metrics.

To complement these results, we also show the TT and EE angular power spectra of the difference between the (noiseless) input map and the (noisy) reconstructed maps at each iteration. These are given in the bottom rows of figures 9 and 10 for both the BD case and the $\dim \mathbf{Z} = 512$ case. The obtained spectra are compared to the power spectrum of the input map and can be seen to stabilize on the level corresponding to the expected noise of the recovered maps. This level is several orders of magnitude below the input signal in all the cases but the smallest angular scales for the EE spectrum where, near the pixel scale, the signal becomes comparable to the noise. These results demonstrate the proper convergence of the maps both in terms of amplitudes and phases of the sky signal and on all angular scales down to the level of the observational precision. We observe that as before, the two-level preconditioner case considered here allows to reconstruct the full range of angular scales essentially in a single step. These observations hold for the other deflation space dimensions as well.

The cost of precomputation for different sizes of the deflation subspace constructed with the Lanczos procedure is shown in figure 11. This cost scales linearly with $\dim \mathbf{Z}$, and can be as costly as a few PCG runs. This makes this method mostly advantageous when we want to solve multiple systems, with roughly the same system matrix but different realizations of the right hand side, as it is the case for example when performing null tests, or Monte Carlo simulations. Then the preconditioner needs to be constructed only once, stored in memory, and all subsequent runs can benefit from the same convergence boost virtually for free. The cost of each iteration does not substantially change in a way that may impact the net gain in runtime as shown in figures 12 and 13.

The slight increase in the average iteration cost seen for the cases $\dim \mathbf{Z} = 1024$ or 2048, is due to the fact that oper-

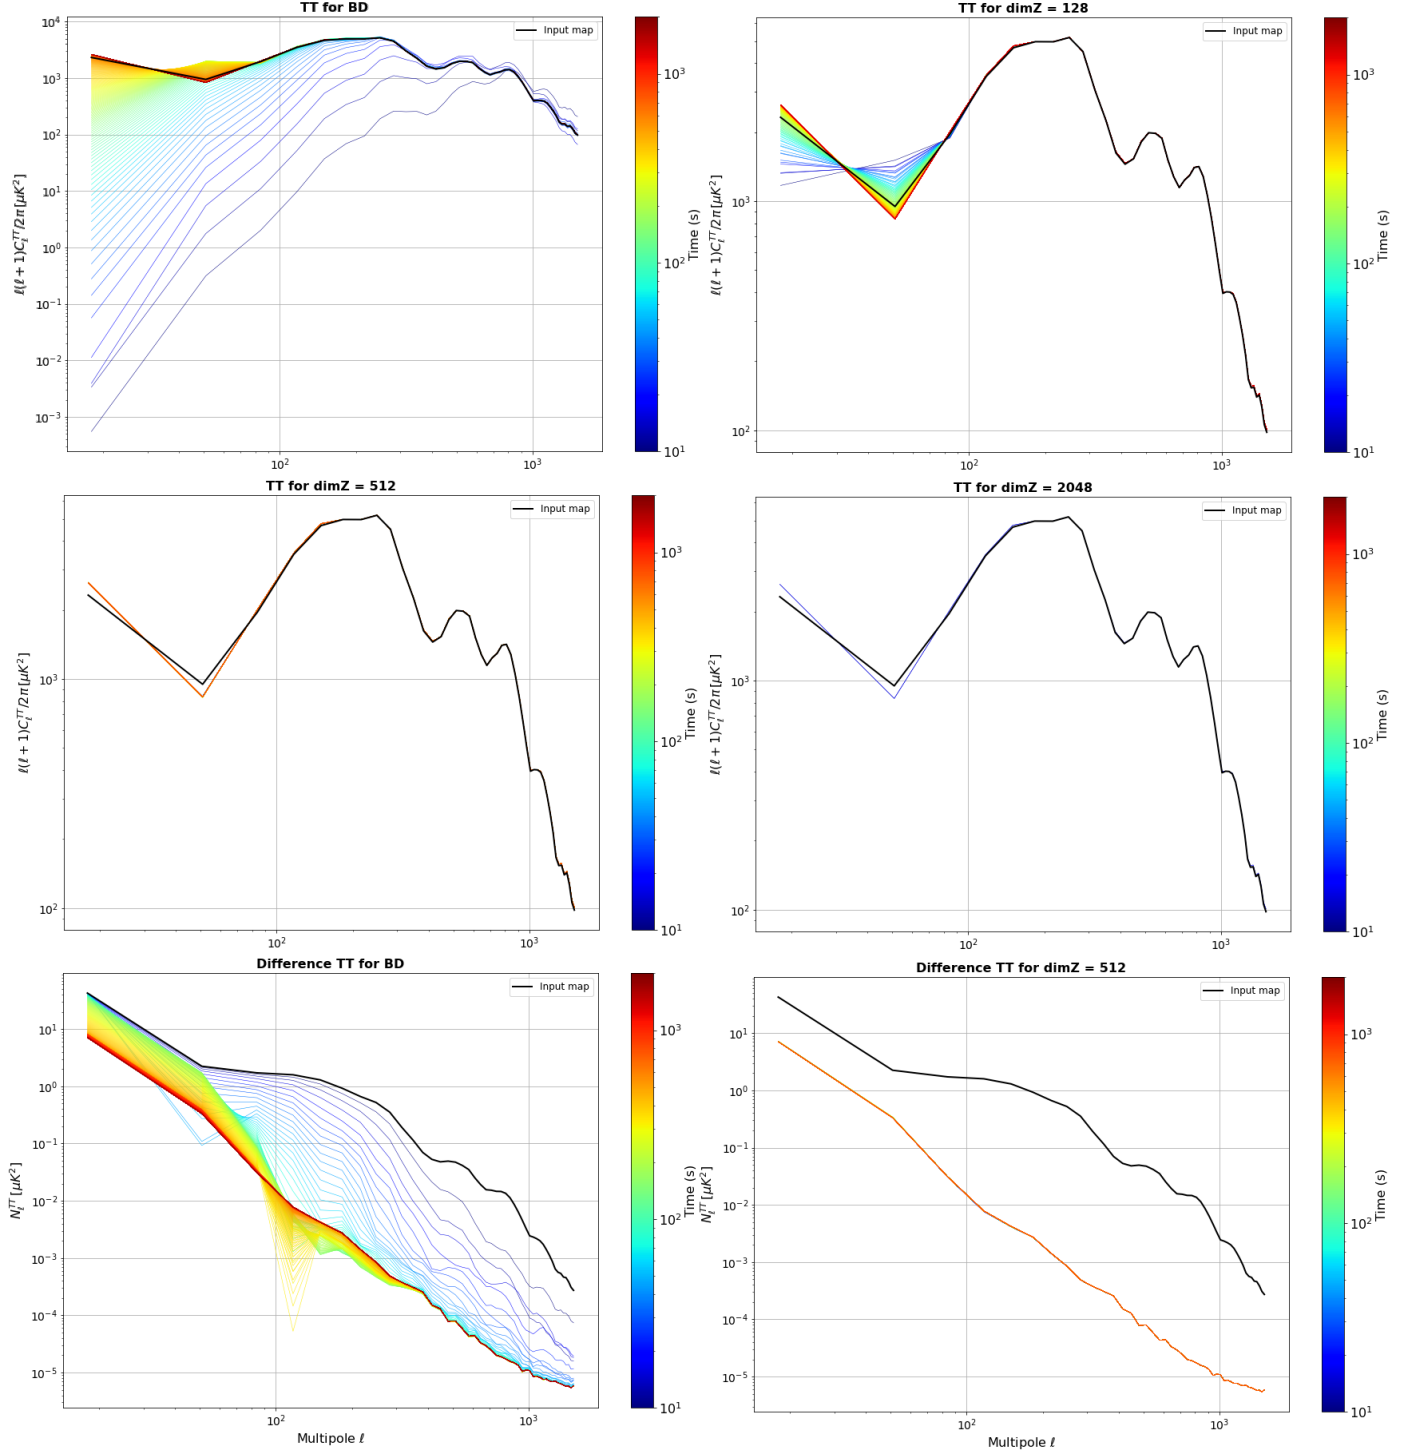


Figure 9: (Top and middle rows): Temperature angular power spectra of the maps at different stages of CG convergence for various preconditioning settings. (Bottom row): Temperature angular power spectra of the difference between the input map and the reconstructed maps at different stages of CG convergence.

1675 ations start getting non negligible contributions from the $\dim \mathbf{Z}$ map-domain operations. However, for the case considered, this₈₆₀ increase does not affect the results in any significant way, but may be of concern for high resolution maps, in the context of LAT observations for instance. In this case, memory space is also a limiting factor and one may be forced to downsample the columns of \mathbf{Z} for the method to work. In such a case, the₈₆₅

small scale modes would not be captured by the deflation subspace, but these typically converge very well regardless. This is however beyond the scope of the present work, and may be the object of future investigation.

We can consider some numbers to illustrate the effective gain in runtime enabled by this technique. In the present case studied, the precomputation cost with $\dim \mathbf{Z} = 2048$




Figure 10: (Top and middle rows): EE angular power spectra of the maps at different stages of CG convergence for various preconditioning settings. (Bottom row): EE angular power spectra of the difference between the input map and the reconstructed maps at different stages of CG convergence.

is roughly ~ 5.5 PCG runs with the standard block-diagonal preconditioner, and the solving time is roughly 10 s. Solving 100 similar systems for example, would then cost around⁶⁹⁵ $\sim 9,500 + 10 \times 100 = 10,500$ s, while it would cost around $1,650 \times 100 = 165,000$ s running with the block-diagonal preconditioner. We gain about a factor ~ 16 solving 100 similar systems. Note that the gain is only going to increase as we in-

crease the number of similar systems to solve, hence alleviating the cost of map-making in Monte Carlo runs for example, as we become only limited by the cost of the simulations.

4.1.4. Validation of templates marginalization

In this subsection, we validate the templates marginalization procedure at the map level. For this purpose, we select




Figure 11: Precomputation cost of the *a posteriori* preconditioner as a function of the dimension of the deflation subspace matrix $\dim \mathbf{Z}$. In practice this dimension is the number of iterations performed by the Lanczos procedure.




Figure 12: Average iteration cost of PCG runs with the *a posteriori* preconditioner as a function of the dimension of the deflation subspace matrix $\dim \mathbf{Z}$.

the input signals in the **S2** simulation such that we include only the cosmological CMB signal and, each time, one specific, time-domain systematic signal modelled by a given template for which we know the exact parameters. We then reconstruct the sky maps assuming 3 different noise models:

- A white noise model (Binned mapper).
- An incorrect template model: we assume the correct template class but with the wrong parameters.
- A correct template model: we assume both the correct template class and the correct parameters.

The binned mapper allows us to see the full impact of the systematic signal on the map if not taken into account in the reconstruction. Using an incorrect template model, allows us




Figure 13: Convergence in wall clock time of the BD preconditioned CG, compared with the *a posteriori* two-level preconditioned CG, for different dimensions of the deflation subspace. The precomputation time is not included here.

to qualitatively assess the impact of model errors, and the last configuration using the correct template model allows us to validate the procedure by verifying that we recover a correct estimate of the CMB map.

Scan synchronous signal subtraction:

We inject a ground-pickup signal using the scan synchronous signal template which takes the form of a step function in azimuth $g(\varphi)$; that is the ground map is binned in azimuth, with some fixed azimuth step $\Delta\varphi$, and in any given constant elevation scan,

$$\forall \varphi \in [\varphi_{min} + i\Delta\varphi, \varphi_{min} + (i+1)\Delta\varphi], g(\varphi) = g_i = \text{const.}$$

Therefore at any given instant t , if the telescope's boresight is pointing towards the i -th bin in azimuth, we can simply write: $SSS(t) = g_i$. We use a linear ground map of the form $g_i = 1 + 0.01i$ (K), with an azimuth bin $\Delta\varphi = 0.33$ deg. For the map reconstruction, given that the azimuth range of the constant elevation scans is of 10 degrees, the correct template model accounts for 30 bins in azimuth for each CES, and we choose to account for only 10 bins in azimuth for each CES in the incorrect model run. The reconstructed temperature maps are shown in figure 14.

The total offset of the map cannot be recovered, since it is fully degenerate with the ground template. Therefore, it is removed from the difference maps shown in the right panel of the figure. In the binned map, the coherent ground structure completely dominates the sky signal as one would expect, and we essentially see the projection of the offset of the ground map. The middle row maps show an improvement, as the ground signal is partially filtered, and the ground offset removed, however since we are assuming larger azimuth bins, we still get power from the sub-bin fluctuations leaking to our reconstructed sky signal, and given that the corresponding amplitudes fluctuations is of order 10 mK, the ground leakage still dominates




Figure 14: Temperature maps of CMB+SSS mapped with respectively, a binned map-maker, a 10 azimuth bins, and 30 azimuth bins SSS templates map-makers (top to bottom). From left to right, we show respectively the input, the output, and the difference maps.

Figure 15: Temperature maps of CMB+HWSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom). From left to right, we show respectively the input, the output, and the difference maps.

over the CMB signal. In the bottom row, we recover the correct CMB map effectively filtering the ground signal. The difference map in the bottom right panel, shows a low amplitude (below $\sim 0.02 \mu\text{K}$) smooth residual structure, which essentially comes from degeneracies between the sky and the ground template. These degenerate modes correspond to the projections of the azimuth bins of the ground template on the sky, which, as the sky rotates, leave stripes which are constant in right ascension (such as spherical harmonics with $m = 0$), with unconstrained relative offsets. Some of these degeneracies can be broken when data from multiple detectors, with shifted and overlapping azimuth bins, is combined, or with the sky pixels shared between adjacent bins. Nonetheless, these constraints are typically weak, leaving a number of ill-constrained modes in declination which characteristics depend on the pixelization and the size of the azimuthal bins and the scanning strategy [42, 33].

With these results we can validate the SSS filtering operators and the corresponding templates marginalization procedure.

HWP synchronous signal subtraction:

We inject a HWP synchronous signal modelled by a set of harmonics of the HWP rotation frequency, coupled with a linear drift term [24, 30, 54]. Therefore, for each CES and each detector, it takes the following form:

$$\mathcal{H}(t) = \sum_n A_n(1 + \varepsilon_n t) \cos(n\phi_t) + B_n(1 + \varepsilon_n t) \sin(n\phi_t), \quad (32)$$

where ϕ_t is the HWP angle at time t . The harmonics amplitudes A_n and B_n are drawn from a normal distribution $\mathcal{N}(\mu, \sigma)$, with $\mu = 100 \text{ mK}$, and a dispersion σ/μ of 1%. The linear drift coefficient ε_n is set such that we get a 1% drift every 1000 s (~ 16 min). We inject a total of six harmonics. For the reconstruction, we consider a HWPSS template model for which the baseline length (i.e., the time interval over which the harmonics amplitudes are considered to be constant) is 10 s. We consider two template configurations, one which takes into account only four harmonics of the HWP rotation frequency, f_{HWP} , and another one which handles all of the six harmonics. The reconstructed temperature and Q polarization maps are shown, respectively, in figures 15 and 16.

The binned maps, in the top rows, show the direct projection of the contaminated TOD into the map, and given the amplitude of the HWPSS, the CMB signal is completely dominated by the spurious signal. In the middle row of figure 15, the filtering of the first four harmonics reduces the noise level of the temperature map at large scales, as we start to see the sky signal emerging, however the remaining $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ harmonics still leak as small scale noise given that these components reside in the high frequency tail of the signal. For polarization, the signal is modulated at $4f_{\text{HWP}}$ and the filtering of the first four harmonics removes the large scale stripy structure we see in the top row of figure 16, caused by spurious signals varying more slowly than the polarization modulator, but the leakage from the $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ harmonics is still sufficiently strong for the CMB signal to remain subdominant. This is due to two




Figure 16: Q polarization maps of CMB+HWSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom). From left to right, we show respectively the input, the output, and the difference maps.

aspects: (1) the $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ harmonics are closer to the center of the science band in the case of polarization than in temperature because of the HWP modulation, and (2) the polarization signal is fainter in amplitude than temperature. In the bottom rows, we recover the correct input CMB signal, and the HWPSS is effectively filtered thereby validating the HWPSS filtering operators and the full templates marginalization procedure. The residual structure seen in the difference map is below $\sim 0.02 \mu\text{K}$ in amplitude, and is due to the linear drift term that is not perfectly modelled in the map reconstruction, since the harmonics amplitudes are considered to be constant over each 10 s interval.

4.2. Case of a satellite mission.

The purpose of this section is to demonstrate our software in the context of full sky, satellite-like observations. Our study case is based roughly on a satellite mission, LiteBIRD [21]. The presented results complement those discussed in the previous sections, which considered partial-sky observations. In addition, we employ here idealized simulations free of the presence of systematic effects and focus the discussion on statistical uncertainties of maps derived in different map-making set-ups.

4.2.1. Simulations description

We use TOAST to simulate full sky observations assuming roughly specifications of the LiteBIRD mission [50]. We simulate one year of observations at multiple frequency bands corresponding to those of LiteBIRD. The input sky includes CMB

and galactic signal assuming the d1s1 foreground model from PySM [57]. We set n_{side} to 256, i.e., the total number of pixels in the maps is 786,432. We do not include beams in the simulations. The scan strategy from the second Lagrange Point (L2) is described by the rotation of the satellite around the spin axis, with a period of 20 min, and the spin angle—the angle between the spin axis and the boresight direction—is $\beta = 50$ deg. In addition, the spin axis also precesses around the Sun-Earth axis with a period of ~ 200 min and a precession angle $\alpha = 45$ deg. Finally, the yearly revolution of the Sun-Earth axis around the Sun is also modeled.

We also simulate $1/f$ instrumental noise, with characteristic knee frequencies of about 20 mHz (f_k in Eq. (21)), and a slope of ~ -1 (α in Eq. (21)). A new realization from the detector noise PSDs is drawn every six hours. The sampling rate is ~ 20 Hz, which means that each detector data amounts to about ~ 600 M time samples. Depending on the instrument and the frequency channel, we need to process data from tens (~ 20) to hundreds (~ 500) detectors, which implies that the sizes of the data sets generated in the simulations range from $O(10^{10})$ to $O(10^{11})$ time samples, with the minimum required memory from $O(1)$ TB to $O(10)$ TB. In the following, we present results only for three selected channels with center frequencies of ~ 40 , ~ 140 , and ~ 235 GHz and assuming 50, 360 and 250 detectors per channel, respectively. The half-wave plate rotation frequency is taken to be different for each of the channels, ranging from ~ 40 to ~ 60 rpm, as are the white noise levels, which are taken to be ~ 20 , ~ 3 , and $\sim 5 \mu\text{K} \sqrt{\text{s}}$, respectively for the low, medium and high frequency channels considered here. All runs are performed on Cori at NERSC.

4.2.2. Full-sky maximum likelihood maps

In the simulated data described above, we expect no spatial noise correlations across the focal plane since the instrumental noise is drawn independently for each detector. This is a simplification of the actual case where the minor detector-detector correlation can, and will be, present due to systematic effects, such as read-out induced cross-talk or common thermal modes. In this case, we can get very close to the true maximum likelihood maps by using the banded block-Toeplitz noise model in the map reconstruction. In figure 17, we show the Q polarization full-sky maps derived from the data simulated respectively at 40, 140, and 235 GHz frequency band and processed with a map-maker assuming a Toeplitz-structured noise covariance with a half bandwidth of $\lambda = 8,192$ time samples, corresponding to a correlation length of about ~ 430 s. Given the expected proximity of the resulting maps to the optimal solution, we hereafter refer to these as maximum likelihood when comparing them to other maps.

4.2.3. Full-sky templates marginalization maps

In the templates marginalization approach, we use polynomial templates to mitigate the $1/f$ noise due to the instrument. We fix the polynomial order to 0 but vary the baseline length from 20 to 3,600 s in order to demonstrate its impact on the statistical uncertainties of the derived maps. **The procedure is**

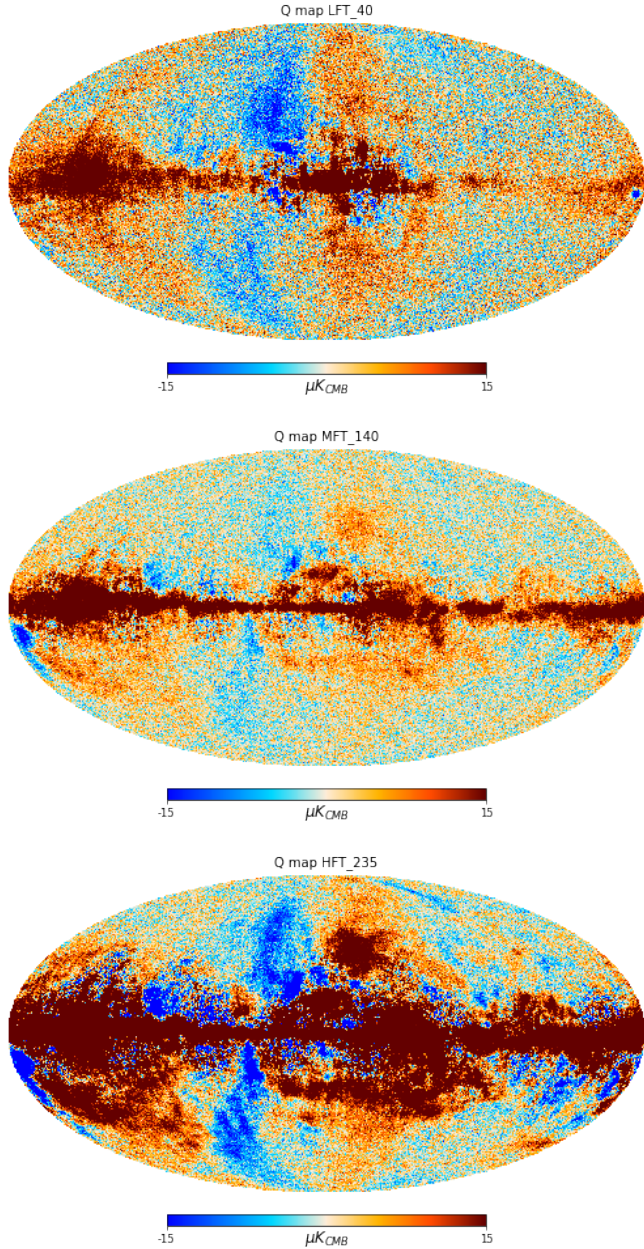


Figure 17: Full-sky maximum likelihood maps from 1 year of simulated observations of the LiteBIRD space mission. From *top* to *bottom* we show the Q polarization maps obtained in the 40, 140, and 235 GHz frequency bands, respectively.

fully equivalent to standard destriping, allowing us to compare our findings with other results in the literature. In all cases the derived maps are essentially indistinguishable by eye, between themselves and from the ML maps, for all Stokes parameters given their high signal-to-noise ratio. To make the differences more evident, we consider here instead difference maps of the output (noisy) and input (noiseless) maps. Given the unbiased character of all the maps discussed here, the difference maps reflect the noise present in the maps derived from the simulated data allowing us to quantify and compare their statistical uncertainties. The results of these tests are discussed




Figure 18: Comparison of the noise power spectra obtained via the different map-making methods available in MAPPRaiser. The baseline for the Template marginalization run is set to 120s. The maps are computed from the simulated data of the 40 GHz channel.

in the following section, where we present the results obtained using data of the 40 GHz channel, but similar observations can be made about the other channels as well.

4.3. Noise spectra

We use the case of the maximum likelihood maps as a reference to which we compare the noise structure in the map domain obtained with other mapping methods. We pay most of the attention to temperature maps, as the noise is essentially white in polarization in all studied cases.

Figure 18 shows the noise power spectra of the temperature and polarization maps computed from the difference between the output and input maps of the 40 GHz channel. The underlying maps are obtained with multiple map-making methods: maximum likelihood with the parameters mentioned above, templates marginalization including only polynomial templates with a baseline length fixed to 120 s, and binned map-making. The EE and BB noise spectra are both shown in grey and are consistent with the polarization white noise level expected from detector NETs. They are essentially flat, due to the HWP modulating the signal away from the low frequency $1/f$ noise, hence we do not see any difference in polarization between the different mapping approaches in this case. However, as mentioned previously, the simple configuration considered in the input simulation is not realistic, and one should expect a more complex noise structure to arise for the polarization signals when introducing systematics. The TT noise power spectra however show differences between different mapping approaches. In particular, the maximum likelihood estimator yields the least noisy maps as expected. The templates marginalization is closer to the maximum likelihood solution only at small scales, while the noise is boosted at large scales yielding worst estimates than the binned map-maker at $l \lesssim 20$.

Figure 19 shows the difference between the binned noise map and the maximum likelihood noise map, as well as the differences between noise maps obtained with templates marginalization assuming baselines of 20 s or 120 s, and the noise maps

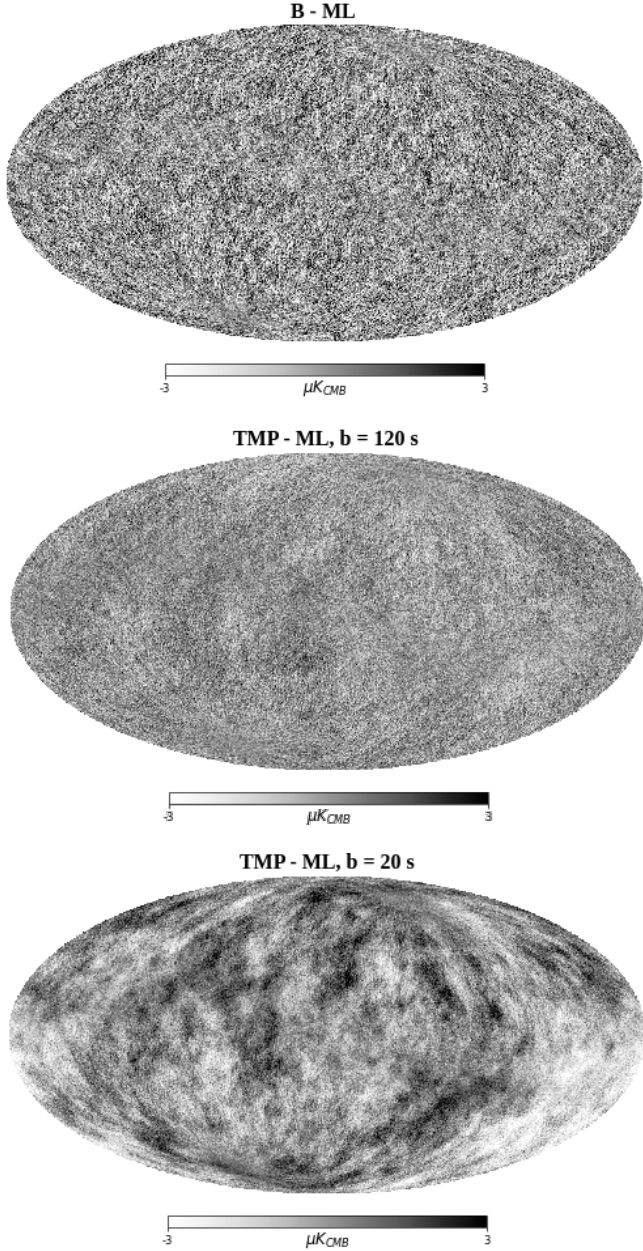


Figure 19: The upper map shows the difference between the binned temperature noise map and the maximum likelihood one. The middle and bottom maps show the difference between the templates marginalization temperature noise maps, respectively, for a baseline of 120 s and 20 s, and the maximum likelihood one. The maps are computed from simulated data of the 40 GHz channel₁₉₇₀

obtained with the maximum likelihood estimator. Given that the binned mapper does not correct for the noise correlations, we see stripes in the difference maps. However, these features are not as strong as the usual stripe patterns seen in the binned Planck maps, for example Figure 6 of [27]. This is due to both a lack of sufficient noise power at low frequencies, given the smaller knee frequencies and milder slope of the detector noise assumed here and reflected in the simulation, and to the more optimal cross-linking of the LiteBIRD-like scanning strategy [21, 56]. Indeed we see in figure 18, that the TT noise₁₉₇₅

power spectra of the binned case and the maximum likelihood case at the largest scales are roughly of the same order of magnitude, with the major difference being a highly anisotropic noise properties in the former case manifested by the stripes. These stripes are suppressed in the case of the maximum likelihood map-making by down weighting the long temporal modes and in the case of the template marginalization by their explicit filtering. This comes at the price of adding more degrees of freedom, thereby increasing the overall level of the statistical noise. This effect should be particularly noticeable for short baselines, and is clearly visible by comparing the maps of figure 19. Processing the maps with short baselines for the polynomial filters boosts the large scale features in the noise map (bottom map), although the map obtained with an intermediate baseline length of 120 s (middle map) shows less stripy patterns than the binned map (top map).₁₉₃₀
₁₉₃₅
₁₉₄₀

These effects can be seen more clearly in the power spectra. Figure 20 shows the TT noise power spectra for multiple baselines compared with the binned and maximum likelihood map-makers. The right panel shows the TT noise power spectra at multipoles $\ell \in [30, 200]$. As the baseline shortens towards 20 s, the lowest value tested in these runs, the noise level gets closer to the maximum likelihood estimate, hence the increase of the statistical noise at these scales does not dominate over the correction of the systematic errors induced by the noise correlations. Conversely, as we increase the baseline lengths, we get closer to the binned map estimate, which is expected since the offsets span durations close to the length of the stationary intervals defined in the simulation, and the map-maker does the exact same processing as a binned map-maker to any modes shorter than the offsets baselines, thus not properly handling the $1/f$ noise.₁₉₄₅
₁₉₅₀
₁₉₅₅

The left panel shows the TT noise power spectra at lower multipoles, $\ell \in [2, 30]$. For clarity, only three baselines (20, 60 and 3,600 s) are selected. In this regime, the statistical noise dominates as the number of sky modes becomes low and the maps processed with filter baselines shorter than typically $\sim 1/f_k$ become noisier than the binned maps at these scales. These findings are similar to observations made in [29, 27] in applications of destripers to Planck simulations as well as for ground-based observations [52]. In order to be able to process the maps with short baselines ($\lesssim 1/f_k$), one needs to include more information to bring these noise levels down. This could be done by introducing priors on the templates amplitudes, computed from the noise PSDs, as is done for instance in the MADAM or Descart destripers [27, 52], this also typically speeds up the convergence of the map-maker.₁₉₆₀
₁₉₆₅

5. Conclusion and future work

In this paper we provide a detailed description of the MAP-PRAISER software framework, which provides tools for computing faithful estimates of sky maps from the time-domain data collected by CMB experiments, in the presence of spurious signals which contribute linearly to the data. These tools benefit from a massively parallel implementation which is also easily extensible and portable. The framework allows the user₁₉₈₀




Figure 20: The temperature noise power spectra obtained via the templates marginalization approach for different baselines. The *left* figure shows the $\ell \in [2, 30]$ range, and the *right* figure shows the $\ell \in [30, 200]$ range. We also include the noise power spectra from the binned and maximum likelihood maps for comparison. The maps are computed from the simulated data of the 40 GHz channel.

to choose between different linear solvers and preconditioning techniques to maximize the computational efficiency according to the specific cases that need to be processed. In addition, one can also choose between different map-making methods according to the case at hand. The framework is flexible and extensible and can deal with cases of generalized map-making problems as those arising in the presence of frequency dependent systematic effects [59]. The architecture is modular such as to allow other solvers, preconditioners, map-making approaches, or different templates for handling time-domain systematics to be easily integrated in the framework.

We demonstrate all these features on simulated data from the TOAST software framework, using ground based and satellite experiments specifications. We first test the scalability and portability of the code by running our map-making code on multiple platforms using different collective communication schemes and show that the code exhibits good scaling up to our tests limit consisting of 16,384 MPI processes demonstrating its ability to handle large data sets. We then show a proof of concept of how convergence can be improved with the ECG solver, which to our knowledge, is the first application of the method in the context of CMB analysis. We conduct an in-depth study of the convergence properties of the two-level preconditioners and demonstrate a significant speedup in the context of multiple solves such as required for null tests or Monte Carlo simulations. We validate the templates marginalization procedure in the case of ground based experiments using two specific templates for ground pickup and half-wave plate synchronous signal. Finally, we also demonstrate the ability of the code to handle data from a full year of simulated observations from a satellite experiment such as LiteBIRD, and explore the noise properties in the recovered maps with the different map-making methods available.

Future work can focus on further optimization of the operations by capitalizing on GPU based architectures such as to

exploit the potential of the ECG solver, or to enable efficient implementation of time-domain operations using a noise model which takes into account spatial correlations between detectors induced, for instance, by atmospheric emission, common thermal modes in the focal plane or other systematic effects such as detector crosstalk. Such GPU accelerated operations will also allow efficient implementations of spherical harmonic transforms useful for Wiener filtering or beam deconvolution map-making. The present software framework can also be used as a forecasting tool to assess the impact of different systematics naturally modeled in time-domain and their mitigation using realistic simulations of actual CMB experiments such as Simons Observatory, CMB-S4 and LiteBIRD. Finally, the framework can also be incorporated within broader non-linear schemes integrating multiple steps of the CMB data analysis such as map-making and component separation into a single one [36].

Acknowledgement

The work was supported by the French National Research Agency (Agence National de Recherche) grant, ANR-B3DCMB, (ANR-17-CE23-0002). We salute members of the B3DCMB collaboration for their multidisciplinary support and insights throughout the implementation of the project. We thank Julian Borill, Reijo Keskitalo, Theodore Kisner for their relevant comments and help with running the simulation software and interfacing it with MAPPRASER. We also thank Sigurd Næss, Matthew Hasselfield, and all members of the Simons Observatory Pipeline Working Group for their useful comments and discussions. We acknowledge help of Jan Papež and Olivier Tissot with the integration of the ECG solver within the framework and the work of Pierre Cargemel and Frédéric Dauvergne on the development of the baseline version of MIDAPACK on which this work has built and improved upon. We acknowledge the use of the NumPy [19], Matplotlib [23], TOAST⁴, NaMas-

ter [3], healpy [62] and HEALPix [15] software packages. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02²¹⁰⁵05CH11231. In France the HPC resources were provided by GENCI–TGCC (allocation A0080411472).

References

- [1] S. Adachi et al. “A Measurement of the Degree-scale CMB B-mode Angular Power Spectrum with POLARBEAR”. In: *The Astrophysical Journal* 897.1 (July 2020), p. 55.
- [2] Peter Ade et al. “The Simons Observatory: science goals and forecasts”. In: *Journal of Cosmology and Astroparticle Physics* 2019.02 (Feb. 2019), pp. 056–056.
- [3] D. Alonso, J. Sanchez, and A. Slosar. “A unified pseudo- C_ℓ framework”. In: *Monthly Notices of the Royal Astronomical Society* 484.3 (Jan. 2019), pp. 4127–4151.
- [4] C. Armitage and B. Wandelt. “Deconvolution map-making for cosmic microwave background observations”. In: *Physical Review D* 70.12 (Dec. 2004).
- [5] E. F. Bunn, Y. Hoffman, and J. Silk. “The Wiener-filtered COBE DMR Data and Predictions for the Tenerife Experiment”. In: *The Astrophysical Journal* 464 (June 1996), p. 1. arXiv: [astro-ph/9509045](https://arxiv.org/abs/astro-ph/9509045) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [6] C. Burigana et al. “A preliminary study on destriping techniques of PLANCK/LFI measurements versus observational strategy”. In: *arXiv e-prints*, astro-ph/9906360 (June 1999), astro-ph/9906360. arXiv: [astro-ph/9906360](https://arxiv.org/abs/astro-ph/9906360) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [7] C. M. Cantalupo et al. “MADmap: A massively parallel maximum likelihood cosmic microwave background map-maker”. In: *The Astrophysical Journal Supplement Series* 187.1 (Mar. 2010), pp. 212–227.
- [8] B.-C. Chiang and K. M. Huffenberger. “Cooling Improves Cosmic Microwave Background Map-Making When Low-Frequency Noise is Large”. In: *arXiv e-prints*, arXiv:2109.11622 (Sept. 2021), arXiv:2109.11622. arXiv: [2109.11622](https://arxiv.org/abs/2109.11622) [[astro-ph](https://arxiv.org/abs/astro-ph).CO].
- [9] J. Delabrouille. “Analysis of the accuracy of a destriping method for future cosmic microwave background mapping with the PLANCK SURVEYOR satellite”. In: *Astronomy & Astrophysics Supplement Series* 127.3 (1998), pp. 555–567.
- [10] O. Doré et al. “MAPCUMBA: A fast iterative multi-grid map-making algorithm for CMB experiments”. In: *Astronomy & Astrophysics* 374 (July 2001), pp. 358–370. arXiv: [astro-ph/0101112](https://arxiv.org/abs/astro-ph/0101112) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [11] J. Errard et al. “Modeling atmospheric emission for CMB ground-based observations”. In: *The Astrophysical Journal* 809.1 (Aug. 2015), p. 63.
- [12] P. G. Ferreira and A. H. Jaffe. “Simultaneous estimation of noise and signal in cosmic microwave background experiments”. In: *Monthly Notices of the Royal Astronomical Society* 312.1 (Feb. 2000), pp. 89–102.
- [13] Nicholas Galitzki et al. “The Simons Observatory: instrument overview”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 10708. International Society for Optics and Photonics. SPIE, 2018, pp. 1–13.
- [14] G. H. Golub and C. F. van Loan. *Matrix Computations*. Fourth. JHU Press, 2013.
- [15] K. M. Górski et al. “HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere”. In: *The Astrophysical Journal* 622.2 (Apr. 2005), pp. 759–771.
- [16] L. Grigori, S. Moufawad, and F. Nataf. “Enlarged Krylov Subspace Conjugate Gradient Methods for Reducing Communication”. In: *SIAM Journal on Matrix Analysis and Applications* 37.2 (2016), pp. 744–773.
- [17] L. Grigori, R. Stompor, and M. Szydlarski. “A parallel two-level preconditioner for Cosmic Microwave Background map-making”. In: *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. 2012, pp. 1–10.
- [18] L. Grigori and O. Tissot. “Scalable Linear Solvers Based on Enlarged Krylov Subspaces with Dynamic Reduction of Search Directions”. In: *SIAM Journal on Scientific Computing* 41.5 (2019), pp. C522–C547. eprint: <https://doi.org/10.1137/18M1196285>.
- [19] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362.
- [20] D. L. Harrison, F. van Leeuwen, and M. A. J. Ashdown. “A deconvolution map-making method for experiments with circular scanning strategies”. In: *Astronomy & Astrophysics* 532 (July 2011), A55.
- [21] M. Hazumi et al. “LiteBIRD satellite: JAXA’s new strategic L-class mission for all-sky surveys of cosmic microwave background polarization”. In: *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave* (Dec. 2020). Ed. by Makenzie Lystrup et al.
- [22] K. M. Huffenberger and S. K. Næss. “Cosmic Microwave Background Mapmaking with a Messenger Field”. In: *The Astrophysical Journal* 852.2, 92 (Jan. 2018), p. 92. arXiv: [1705.01893](https://arxiv.org/abs/1705.01893) [[astro-ph](https://arxiv.org/abs/astro-ph).IM].
- [23] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95.
- [24] B. Johnson et al. “MAXIPOL: Cosmic Microwave Background Polarimetry Using a Rotating Half-Wave Plate”. In: *The Astrophysical Journal* 665 (Dec. 2008), p. 42.
- [25] E. Keihänen and M. Reinecke. “ArtDeco: a beam-deconvolution code for absolute cosmic microwave background measurements”. In: *Astronomy & Astrophysics* 548 (Nov. 2012), A110.
- [26] E. Keihänen et al. “A maximum likelihood approach to the destriping technique”. In: *Astronomy & Astrophysics* 428 (Dec. 2004), pp. 287–298. arXiv: [astro-ph/0304411](https://arxiv.org/abs/astro-ph/0304411) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [27] E. Keihänen et al. “Making cosmic microwave background temperature and polarization maps with MADAM”. In: *Astronomy & Astrophysics* 510 (Feb. 2010), A57.

- [28] Zigmund D. Kermish et al. “The POLARBEAR experiment”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*. Ed. by Wayne S. Holland²¹⁵ and Jonas Zmuidzinas. Vol. 8452. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Sept. 2012, p. 84521C. arXiv: [1210.7768](https://arxiv.org/abs/1210.7768) [[astro-ph](https://arxiv.org/archive/astro).[IM](https://arxiv.org/archive/astro)].
- [29] H. Kurki-Suonio et al. “Destriping CMB temperature and polarization maps”. In: *Astronomy & Astrophysics* 506.3 (Aug. 2009), pp. 1511–1539.
- [30] A. Kusaka et al. “Modulation of cosmic microwave background polarization with a warm rapidly rotating half-wave plate on the Atacama B-Mode Search instrument”. In: *Review of Scientific Instruments* 85.2 (Feb. 2014), p. 024501.
- [31] Bradley R. Lowery and Julien Langou. “Stability Analysis of QR factorization in an Oblique Inner Product”. In: *arXiv e-prints*, arXiv:1401.5171 (Jan. 2014), arXiv:1401.5171. arXiv: [1401.5171](https://arxiv.org/abs/1401.5171) [[math](https://arxiv.org/archive/math).[NA](https://arxiv.org/archive/math)].
- [32] D. Maino et al. “The Planck-LFI instrument: Analysis of the $1/f$ noise and implications for the scanning strategy”. In: *Astronomy & Astrophysics Supplement Series* 140.3 (Dec. 1999), pp. 383–391.
- [33] S. Naess et al. “The Atacama Cosmology Telescope: CMB polarization at $200 < \ell < 9000$ ”. In: *Journal of Cosmology and Astroparticle Physics* 2014.10 (Oct. 2014), pp. 007–007. ²²³⁵
- [34] P. Natoli et al. “A Map-Making algorithm for the Planck Surveyor”. In: *Astronomy & Astrophysics* 372.1 (June 2001), pp. 346–356.
- [35] P. Pacheco. *An Introduction to Parallel Programming*. Morgan Kaufmann, 2011. ²²⁴⁰
- [36] J. Papež, L. Grigori, and R. Stompor. “Accelerating linear system solvers for time-domain component separation of cosmic microwave background data”. In: *Astronomy & Astrophysics* 638 (June 2020), A73.
- [37] Planck Collaboration et al. “Planck 2018 results - I. Overview and the cosmological legacy of Planck”. In: *A&A* 641 (2020), A1. ²²⁴⁵
- [38] Planck Collaboration et al. “Planck 2018 results - II. Low Frequency Instrument data processing”. In: *A&A* 641 (2020), A2. ²²⁵⁰
- [39] Planck Collaboration et al. “Planck 2018 results - III. High Frequency Instrument data processing and frequency maps”. In: *A&A* 641 (2020), A3.
- [40] Planck Collaboration et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641, A6 (Sept. 2020), A6. arXiv: [1807.06209](https://arxiv.org/abs/1807.06209) [[astro-ph](https://arxiv.org/archive/astro).[CO](https://arxiv.org/archive/astro)]. ²²⁵⁵
- [41] Planck Collaboration et al. “Planck intermediate results - LVII. Joint Planck LFI and HFI data processing”. In: *A&A* 643 (2020), A42.
- [42] D. Poletti et al. “Making maps of cosmic microwave background polarization for B-mode studies: the POLARBEAR experiment”. In: *Astronomy & Astrophysics* 600 (Mar. 2017), A60. ²²⁶⁰
- [43] W. H. Press et al. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. 3rd ed. USA: Cambridge University Press, 2007.
- [44] S. Prunet et al. “Noise estimation in CMB time-streams and fast map-making. Application to the BOOMERanG98 data”. In: *arXiv e-prints*, astro-ph/0101073 (Jan. 2001), astro-ph/0101073. arXiv: [astro-ph/0101073](https://arxiv.org/abs/astro-ph/0101073) [[astro-ph](https://arxiv.org/archive/astro)]. ²²⁶⁵
- [45] G. Puglisi et al. “Iterative map-making with two-level preconditioning for polarized cosmic microwave background data sets”. In: *Astronomy & Astrophysics* 618 (Oct. 2018), A62.
- [46] B. Revenu et al. “Destriping of polarized data in a CMB mission with a circular scanning strategy”. In: *Astronomy & Astrophysics Supplement Series* 142.3 (2000), pp. 499–509.
- [47] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Second. Other Titles in Applied Mathematics. SIAM, 2003.
- [48] Y. Saad et al. “A Deflated Version of the Conjugate Gradient Algorithm”. In: *SIAM Journal on Scientific Computing* 21.5 (2000), pp. 1909–1926. eprint: <https://doi.org/10.1137/S1064829598339761>.
- [49] R. Stompor et al. “Making maps of the cosmic microwave background: The MAXIMA example”. In: *Physical Review D* 65.2 (Dec. 2001).
- [50] H. Sugai et al. “Updated Design of the CMB Polarization Experiment Satellite LiteBIRD”. In: *Journal of Low Temperature Physics* 199.3-4 (Jan. 2020), pp. 1107–1117.
- [51] D. Sutton et al. “Fast and precise map-making for massively multi-detector CMB experiments”. In: *Monthly Notices of the Royal Astronomical Society* 407.3 (Aug. 2010), pp. 1387–1402.
- [52] D. Sutton et al. “Map making in small field modulated CMB polarization experiments: approximating the maximum likelihood method”. In: *Monthly Notices of the Royal Astronomical Society* 393.3 (Feb. 2009), pp. 894–910. eprint: <https://academic.oup.com/mnras/article-pdf/393/3/894/18569414/mnras0393-0894.pdf>.
- [53] M. Szydlarski, L. Grigori, and R. Stompor. “Accelerating the cosmic microwave background map-making procedure through preconditioning”. In: *Astronomy & Astrophysics* 572 (Nov. 2014), A39.
- [54] S. Takakura et al. “Performance of a continuously rotating half-wave plate on the POLARBEAR telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2017.05 (May 2017), pp. 008–008.
- [55] J.M. Tang et al. “Comparison of Two-Level Preconditioners Derived from Deflation, Domain Decomposition and Multigrid Methods”. In: *Journal of Scientific Computing*, 39 (3), 2009 39 (June 2009).
- [56] The LiteBIRD collaboration. In: *PTEP* (2021, in prep).
- [57] B. Thorne et al. “The Python Sky Model: software for simulating the Galactic microwave sky”. In: *Monthly Notices of the Royal Astronomical Society* 469.3 (Aug. 2017), pp. 2821–2833. arXiv: [1608.02841](https://arxiv.org/abs/1608.02841) [[astro-ph](https://arxiv.org/archive/astro).[CO](https://arxiv.org/archive/astro)].
- [58] R. J. Thornton et al. “The Atacama Cosmology Telescope: The polarization-sensitive ACTPol instrument”. In: *The Astrophysical Journal Supplement Series* 227.2 (Dec. 2016), p. 21.
- [59] C. Vergès, J. Errard, and R. Stompor. “Framework for analysis of next generation, polarized CMB data sets in the presence of Galactic foregrounds and systematic effects”. In: *Physical Review D* 103.6 (Mar. 2021).
- [60] I. K. Wehus, S. K. Næss, and H. K. Eriksen. “Bayesian noise estimation for non-ideal cosmic microwave background experiments”. In: *The Astrophysical Journal Supplement Series* 199.1 (Feb. 2012), p. 15.

- [61] D. Yvon and F. Mayet. “Mirage: A new iterative map-making code for CMB experiments”. In: *Astronomy & Astrophysics* 436.2 (May 2005), pp. 729–739.
- [62] A. Zonca et al. “healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python”. In: *Journal of Open Source Software* 4.35 (Mar. 2019), p. 1298.

Appendix A. Pre-defined template classes.

Here we describe some of the template classes which are directly implemented in the proposed framework. These are constructed following previous CMB data analysis experience, and selected given potential needs of forthcoming and future experiments. We also highlight the orthogonalization kernel structure these templates imply.

• Atmosphere and $1/f$ instrumental noise

In a typical experiment the noise is correlated due to low frequency excess noise of instrumental origin, and the atmospheric emission in the case of ground experiments. We model this low frequency noise by a set of Legendre polynomials, and hence build a filtering operator which deprojects polynomial trends from the data. This can be seen as the most straightforward generalization of the standard destriper.

We define a set of time intervals, $\cup [t_i, t_{i+1}]$, over which the amplitudes of the polynomial trends for each detector are considered to be constant. In practice, these time intervals can span the duration of each sweep in the constant elevation scans (CESs) of a ground telescope for example. The number of polynomial orders, n_{poly} , is fixed. We do not consider here common detector modes, which would model correlations between detectors. We can define each template as a compact support function, which represents one polynomial order for a given detector during one time interval $[t_i, t_{i+1}]$ (e.g., one sweep). In summary, for a detector j , there exists a set, $(\alpha_{j,k})_{k \in \llbracket 0, n_{poly}-1 \rrbracket}$ of real amplitudes, such as the $1/f$ noise component in the i -th time interval would be given by:

$$n_j(t) = \sum_{k=0}^{n_{poly}-1} \alpha_{j,k} P_k(t), \quad (\text{A.1})$$

where P_k is the Legendre polynomial of degree k rescaled so it is orthonormal to others on the relevant time interval. Each of these templates represents one column of the global template matrix \mathbf{T} . By virtue of this construction and the orthogonality of Legendre polynomials, the columns would be orthogonal making the polynomial templates kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$ diagonal, simplifying the computational cost of the filtering operator. This approach to modeling the atmospheric signal explicitly ignores correlations between contributions to data registered by different detectors. Such common modes can be straightforwardly modeled in the proposed framework, though may come at additional computational cost. Whether such an extension is justified and beneficial will depend on other factors such as relative calibration errors between the detectors and should be decided on case-by-case.

• Scan synchronous signal

These templates model any spurious signals which are synchronous with the scanning strategy. For ground experiments, these would be dominated by ground pickup sourced by the far side lobes of the beam which sense the surrounding terrain [1]. We model this signal with a template parametrized by the azimuth angle of the boresight pointing.

We define an azimuth grid, $\cup_{i=1}^{n_{az}} [\varphi_i, \varphi_{i+1}]$ for each CES, fixing *a priori* a number of azimuth bins n_{az} . Then for each detector j we consider a ground structure, $(g_{j,i})_{i \in \llbracket 1, n_{az} \rrbracket}$, which is stable over the duration of the full CES. The template marginalization performance is in such a case contingent upon having enough redundancy in the scanning strategy to separate between the ground and sky signal over the duration of the CES. The scan synchronous signal is given by the following,

$$\forall t : \varphi(t) \in [\varphi_i, \varphi_{i+1}] \Rightarrow SSS_j(t) = g_{j,i}. \quad (\text{A.2})$$

In the template matrix formulation as in Eq. (1) the amplitude vector stores the bin amplitudes, $g_{j,i}$ and there is one column of the corresponding scan-synchronous template matrix for each azimuthal bin. The columns have consequently non-zero values only in rows corresponding to times at which the telescope azimuth falls into the corresponding bin. As only one bin contributes to the overall measurement at any given time the template matrix has only one non-zero element per row and is therefore very sparse and the columns are by construction orthogonal. The non-zero values of the matrix can be typically set to unity, however, they may be replaced by some uniquely defined function of time, such as one describing gain drifts. For the diagonal noise weights, the corresponding diagonal block of the orthogonalization kernel, $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$, will be diagonal. However, the scan-synchronous templates may not be and, in general, will not be orthogonal to the other templates, such as the polynomial templates discussed earlier. As discussed in [42] in general the scan-synchronous templates will not be also orthogonal to the sky signal, leading to near-degeneracies in the system matrix broken only by the assumed sky pixelization.

We note that more involved models of the scan synchronous signal are readily possible. For instance, for a given grid of azimuths, $\{\varphi_i\}$, rather than binning we can linearly interpolate between the signal estimated in the grid nodes. This could be helpful whenever the scan-synchronous signal is changing rapidly with the azimuth and when the binning may require a very large number of azimuth bins in order to minimize the signal variation across each bin. In this case, the interpolation may be more efficient. There will be two non-zero values per each row of the corresponding template matrix corresponding to two azimuth grid points, φ_k, φ_{k+1} bracketing the telescope azimuth, $\varphi(t)$, and the two non-zero elements would be given by,

$$\mathbf{T}_{tk} = \frac{\varphi(t) - \varphi_k}{\varphi_{k+1} - \varphi_k}, \quad \mathbf{T}_{t(k+1)} = \frac{\varphi_{k+1} - \varphi(t)}{\varphi_{k+1} - \varphi_k}. \quad (\text{A.3})$$

The correspond kernel in such a case would be tridiagonal and therefore could be inverted efficiently.

- *Half-wave plate synchronous signal*

Many of the modern CMB experiments feature a half-wave plate allowing to modulate the incoming polarized signal. While they play a key role in reducing the noise correlations, reducing beam effects, and minimizing leakages, they also tend to introduce their own specific effects, which are generally synchronous with the position of the HWP itself [24, 30]. They arise from differential transmission or reflection of unpolarized light as it propagates through the optical chain, or from polarized emission of the HWP itself and can be modulated at different harmonics of its rotation speed. We refer to them as HWP synchronous signals.

We could define templates for these signals in a model-independent way as in the previous subsection this time relying on binning in the HWP angle. As the HWP synchronous signals tend to vary rather abruptly with the HWP angle, this may however require a large number of bins even if some interpolation is invoked. Alternately we can model these spurious signals as a sum of harmonics of the HWP rotation frequency. This is the approach we have chosen in this work.

We define a set of time intervals, $\cup [t_i, t_{i+1}]$, with an *a priori* fixed length Δt , except for the remainder interval at the end of each CES. The amplitudes of the HWP harmonics in each detector are considered to be constant over each Δt . The number of HWP harmonics, n_{HWP} is fixed. In a similar manner to the polynomial templates, we can define each template as a compact support function, which represents either the cosine or the sine part of HWP harmonic for a given detector on time interval $[t_i, t_{i+1}]$. So, for a detector j , there exists two sets, $(\alpha_{j,k})_{k \in \llbracket 1, n_{HWP} \rrbracket}$ and $(\beta_{j,k})_{k \in \llbracket 1, n_{HWP} \rrbracket}$ of real amplitudes, such as the HWP synchronous signal in the i -th time interval would be given by:

$$\mathcal{H}_j(t) = \sum_{k=1}^{n_{HWP}} \alpha_{j,k} \cos(k\phi(t)) + \beta_{j,k} \sin(k\phi(t)), \quad (\text{A.4})$$

where $\phi(t)$ is the HWP angle as a function of time. The interval length Δt can be chosen so that we get an integer number of rotations of the HWP, in the ideal case where the rotation frequency is not fluctuating. This would simplify the structure of the HWP segment of the templates kernel. This is because each of the templates as we defined them above represents a column of the HWP segment of the templates matrix, and the proposed choice of the time intervals would make the columns orthogonal or near-orthogonal (except the last one in each CES span) by virtue of the orthogonality of cosines and sines in an integer number of periods (HWP rotations). Therefore the corresponding part of the kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$ would be nearly diagonal (last column and row of each CES excluded), again speeding up its inversion and improving the performance of the filtering operator.

We note that one could further generalize Eq. (A.4) by introducing time dependent coefficients α and β accounting on potential gain drifts. This would typically result in a full block-diagonal structure of the kernel, with each block corresponding to a different time interval and different detectors, and would therefore require a full numerical inversion.