



# CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Maximum a posteriori CMB lensing reconstruction

Julien Carron and Antony Lewis

Phys. Rev. D **96**, 063510 — Published 14 September 2017

DOI: [10.1103/PhysRevD.96.063510](https://doi.org/10.1103/PhysRevD.96.063510)

# Maximum a posteriori CMB lensing reconstruction

Julien Carron<sup>1</sup> and Antony Lewis<sup>1</sup>

<sup>1</sup>*Department of Physics & Astronomy, University of Sussex, Brighton BN1 9QH, UK*

(Dated: August 16, 2017)

Gravitational lensing of the CMB is a valuable cosmological signal that correlates to tracers of large-scale structure and acts as an important source of confusion for primordial  $B$ -mode polarization. State-of-the-art lensing reconstruction analyses use quadratic estimators, which are easily applicable to data. However, these estimators are known to be suboptimal, in particular for polarization, and large improvements are expected to be possible for high signal-to-noise polarization experiments. We develop a method and numerical code, LENSIT, that is able to find efficiently the most probable lensing map, introducing no significant approximations to the lensed CMB likelihood, and applicable to beamed and masked data with inhomogeneous noise. It works by iteratively reconstructing the primordial unlensed CMB using a deflection estimate and its inverse, and removing residual lensing from these maps with quadratic estimator techniques. Roughly linear computational cost is maintained due to fast convergence of iterative searches, combined with the local nature of lensing. The method achieves the maximal improvement in signal to noise expected from analytical considerations on the unmasked parts of the sky. Delensing with this optimal map leads to forecast tensor-to-scalar ratio parameter errors improved by a factor  $\simeq 2$  compared to the quadratic estimator in a CMB stage IV configuration.

## I. INTRODUCTION

The large-scale structure of the Universe deflects Cosmic Microwave Background (CMB) photons by a few arcminutes, introducing a characteristic signature in the fluctuations in the CMB temperature and polarization [1]. The statistical homogeneity and isotropy of the CMB gets distorted locally, and sizeable higher-order statistics are produced. Lensing estimators use these higher-order statistics to construct an integrated measure of the linear mass fluctuations in the Universe that cross-correlates to all traditional large-scale structure tracers. After the first direct detection of lensing in the CMB by the ACT team [2], the SPT [3], POLARBEAR [4], SPTpol [5], BICEP2-KECK [6] and Planck [7] collaborations have also reported the detection of the lensing signal and published band-powers. The most decisive detection yet was by the Planck satellite [8]: its full-sky coverage comes with a statistical power that simply cannot be matched at the present time.

Current measurements all use quadratic estimator techniques, first devised in optimized form by Refs. [9, 10]. The quadratic estimator uses optimally-weighted two-point statistics of the data maps to reconstruct the deflection field. At current noise levels, this estimator is nearly optimal. The science returns from the use of more sophisticated techniques are expected to be small, and no other type of estimator has been applied to data so far. However, the situation will have changed by the time of CMB stage IV (CMB-S4) [11], if not before. At this point the polarization instrumental noise is expected to become smaller than the  $\sim 5\mu K$  arcmin lensing  $B$  mode. Barring welcome detections, lensing will become the most relevant cosmological source of confusion in the search for primordial  $B$  modes [12], and more optimal delensing methods will become critical.

If the noise and primordial polarization  $B$  mode is negligible, a well-known variable-counting argument [13] suggests that as long as the lensing is fully described by a gradient deflection remapping of the unlensed fields, the observed lensed  $E$  and  $B$  fields should contain enough information to recon-

struct essentially perfectly both the lensing potential and the unlensed  $E$  field. Fundamental limits are well below near-future sensitivities, including corrections to the remapping approximation from emission angle, time delay and polarization rotation [14, 15], lensing curl modes from second order post-Born lens-lens couplings [13, 16], intrinsic non-linearities of the CMB at recombination [17, 18], and second-order sourced vector and tensor modes [17, 19–21]. With the last science release from the *Planck* team in sight, it therefore seems timely to revisit alternative, more optimal CMB lensing estimation. This paper presents and discusses a new implementation of a maximum a posteriori estimate of the lensing potential from CMB data.

Motivation for this work is not limited to primordial  $B$  modes. The CMB lensing kernel peaks at  $z \sim 2$  and overlaps the galaxy and weak lensing surveys targeting the dark sector of the Universe. The correlated information is expected to contribute to breaking important degeneracies and to help with systematics, so optimal CMB lensing mass maps will also be useful for use with large-scale structure observations. Iterative estimates may also be useful even at higher noise levels, in particular in the presence of sky cuts or wildly inhomogeneous noise maps where the analytic response of the quadratic estimator is inaccurate. An optimal estimate of the potential map might prove better than the current simpler practice of sweeping these deviations into Monte-Carlo (MC) corrections to the spectrum estimate. Finally, also looking a bit ahead, successful exploration of the lensed CMB likelihood may prove useful more generally, opening a path towards optimal joint estimation of the primary CMB and lensing potential.

It is clear how iterative estimates should work in practice, at least intuitively: delens the data using the quadratic deflection estimate, then again apply a quadratic estimator on the resulting maps, with possibly modified weights, and iterate until convergence [22]. Of course, a world of potential complications lurks in the details, and no canonically best approach is known at present. Formally, the code we present finds a maximum of the posterior probability density function (PDF) of the lensing potential. As such, it is similar in spirit

to the first iterative estimator proposed for temperature reconstruction [23] and polarization [13]. A maximum likelihood approach to lensing reconstruction is also discussed, without implementation, in the review article Ref. [24]. In contrast to Refs. [13, 23], our implementation can be considered exact, in the sense that it maximizes the relevant functions without introducing approximations, under the assumption of Gaussian unlensed CMB, noise and deflection fields. It can also account for beams, sky cuts and other non-ideal effects.

The quadratic estimator has the convenient property of being relatively straightforward. It can be implemented using a small number of harmonic transforms [8, 25], keeping the overall numerical cost under control (dominated in the *Planck* analysis by the cost of the inverse-variance filtering step). It seems unavoidable that alternative more-optimal approaches must be substantially more costly, and our implementation is no exception. At each iteration step, maximum a posteriori unlensed CMB maps are produced, under the assumption that the current deflection estimate is the correct one. This operation in effect solves a large  $N_{\text{pix}} \times N_{\text{pix}}$  set of linear equations, and must be itself performed via an iterative method, each step involving a fair number of lensing operations, even in the absence of sky cuts or other non-ideal effects.

Nevertheless, lensing and lensing reconstruction have the advantage of being very local in position space. All operations scale linearly with the number of resolution elements, or follow the cost of an harmonic transform, and the good convergence properties of the iterative searches proposed here keep the total computational burden under control. We also provide GPU implementations of the most expensive steps.

We use the flat sky approximation throughout the main text. Appendix A describes the implementation on the curved sky, using the machinery of spin-weight spherical harmonics. The implementation is otherwise identical in all respects, though we have so far only thoroughly tested everything on the flat sky where the numerical implementation is faster. We expect the same convergence properties of iterative estimator on the curved sky: empirically, the only effect we observe increasing the area is to rescale the total execution time, which is reasonable given that lensing distortions are very much localized. Furthermore, iterative delensing will probably initially be most useful on deep observations of a small patch of sky where the flat sky approximation is accurate [11].

Sections II and III describe the algorithm and details of its implementation respectively. Section IV provides tests and applications. We summarize and conclude in Section V.

## II. DESCRIPTION

Let us first establish some notation. Let  $\mathbf{x}, \mathbf{y}$  be points on a patch of the flat sky of area  $V$ , and  $\mathbf{r} = \mathbf{x} - \mathbf{y}$  be the separation vector. The primary, unlensed Stokes CMB fields  $T, E$  or  $B$  are written as  $X(\mathbf{x})$ , and  $X^{\text{dat}}$  denotes the observed Stokes data  $T, Q$  and  $U$  on the data pixels, inclusive of noise and transfer function. We use  $a, b$  in  $(0, 1)$  to denote the two cartesian axes of the flat sky, and use the symmetric Fourier convention, which is closest to the traditional curved sky nor-

malization. We use the notation  $\ell$  for multipoles of the CMB maps and  $\mathbf{L}$  for the lensing maps.

We denote the primordial, unlensed CMB modes  $\{T, E, B\}$  as a column matrix  $X$ , with primordial spectral matrix  $C_\ell^{\text{unl}}$

$$\langle X_\ell X_{\ell'}^\dagger \rangle = \delta_{\ell\ell'} C_\ell^{\text{unl}}. \quad (2.1)$$

This matrix is diagonal with respect to multipole index, but not necessarily across  $T, E, B$  indices. Also let  $D$  be the deflection operation that maps these unlensed CMB modes to the real space, lensed, Stokes parameters. For instance, in temperature we may write explicitly on the flat sky

$$D_\ell^{TT}(\mathbf{x}) \equiv \frac{1}{\sqrt{V}} e^{i\ell \cdot (\mathbf{x} + \boldsymbol{\alpha}(\mathbf{x}))}, \quad D_\ell^{TE}(\mathbf{x}) = D_\ell^{TB}(\mathbf{x}) = 0. \quad (2.2)$$

The polarization components are similar but involves the spin-2 flat-sky harmonics. Here, and throughout, we use approximation that the lensed fields are entirely defined by a remapping of the unlensed fields, where  $\boldsymbol{\alpha}(\mathbf{x})$  is the lensing deflection angle that relates the observed lensed field at  $\mathbf{x}$  to the unlensed fields at  $\mathbf{x} + \boldsymbol{\alpha}(\mathbf{x})$ .

### A. Model

The model for the CMB data  $X^{\text{dat}}$  on the observed pixels is given by a linear response matrix  $B$  operating on the lensed sky (which includes, for example, the effect of the instrumental beam and pixel window function), plus independent noise  $n_i$ , so that

$$X^{\text{dat}} = BDX + n. \quad (2.3)$$

The pixel-pixel covariance can be written in compact notation using a series of linear operators as follows:

$$\text{Cov}_\alpha \equiv \langle X^{\text{dat}} X^{\text{dat},\dagger} \rangle = BDC^{\text{unl}} D^\dagger B^\dagger + N, \quad (2.4)$$

where  $N$  is the noise covariance matrix, which we assume is diagonal in pixel space. Unlensed CMB fields and the noise on each pixel are assumed to obey Gaussian statistics, so the likelihood is also Gaussian. The log-likelihood is then

$$\ln p(X^{\text{dat}} | \alpha) = -\frac{1}{2} X^{\text{dat}} \cdot \text{Cov}_\alpha^{-1} X^{\text{dat}} - \frac{1}{2} \det \text{Cov}_\alpha. \quad (2.5)$$

We need to use a prior on the statistics of the deflection field to regularize the large number of poorly-constrained small-scale modes. The  $\Lambda$ CDM CMB lensing potential  $\phi$  is expected to be nearly linear, so choosing Gaussian field statistics for  $\phi$  is a natural choice, and will likely remain accurate in the foreseeable future on the scales where the lensing potential can accurately be reconstructed. Using a Gaussian prior on the signal does not prevent reconstruction of any non-Gaussian signal that may actually be present (as expected from non-linear structure growth and post-Born lensing [16, 26]).

We assume pure gradient lensing deflections, in which case the log-posterior becomes, up to irrelevant constants,

$$\ln p(\phi|X^{\text{dat}}) = \ln p(X^{\text{dat}}|\phi) - \frac{1}{2} \sum_{\mathbf{L}} \frac{\phi_{\mathbf{L}}^2}{C_{\mathbf{L}}^{\phi\phi}}, \quad (2.6)$$

where the likelihood is given by Eq. 2.5 with  $\alpha = \nabla\phi$ . Curl and joint curl / gradient reconstruction is very analogous, but should be of limited physical relevance in the foreseeable future, mainly serving as a consistency check on the gradient reconstruction analysis. An interesting first prospect would be the detection of the post-Born curl signal, forecast to be marginally detectable in the bispectrum with CMB-S4 [16], to which this methodology could also be applied.

## B. Gradients

To maximize the log-posterior we consider the derivative of the log-posterior with respect to the deflection. The total gradient,  $g$ , splits naturally into three pieces:

$$g_a^{\text{tot}} \equiv \frac{\delta \ln p(\alpha|X^{\text{dat}})}{\delta \alpha_a(\mathbf{n})} = g_a^{\text{QD}} - g_a^{\text{MF}} + g_a^{\text{PR}}, \quad (2.7)$$

one from the quadratic part of the likelihood ( $g^{\text{QD}}$ ), one from the likelihood covariance determinant ( $g^{\text{MF}}$ , the mean-field), and one ( $g^{\text{PR}}$ ) from the prior. The choice of the odd sign of  $g^{\text{MF}}$  is more natural and becomes clear later on. The prior gradient is straightforward to evaluate assuming Gaussian statistics.

The gradients of the likelihood are first calculated in real space, with the gradients with respect to the two cartesian components of the deflection giving

$$\frac{\delta \ln p(X^{\text{dat}}|\alpha)}{\delta \alpha_a(\mathbf{n})} = g_a^{\text{QD}}(\mathbf{n}) - g_a^{\text{MF}}(\mathbf{n}). \quad (2.8)$$

These are then rotated to harmonic space to give the gradient and curl components. The piece quadratic in the data

$$g_a^{\text{QD}}(\mathbf{n}) = [V_{\alpha} X^{\text{dat}}]^i(\mathbf{n}) [W_{\alpha}^a X^{\text{dat}}]_i(\mathbf{n}), \quad (2.9)$$

is made up of two legs with data weights

$$V_{\alpha} = B^{\dagger} \text{Cov}_{\alpha}^{-1}, \quad W_{\alpha}^a = D \nabla_a C^{\text{unl}} D^{\dagger} B^{\dagger} \text{Cov}_{\alpha}^{-1}. \quad (2.10)$$

The gradient matrix  $\nabla_a C^{\text{unl}}$  is block diagonal in harmonic space with blocks  $i\ell_a C_{\ell}^{\text{unl}}$ . These weights are identical, in the absence of deflection, to the (unnormalized) traditional Minimum Variance (MV) lensing quadratic estimators evaluated with unlensed spectra.

Both legs of the quadratic estimator (the two terms in Eq. (2.9)) can be written in terms of reconstructed unlensed CMB modes, as follows. Consider the most probable primordial CMB modes  $X_{\alpha}^{\text{WF}}$  given the data, under the assumption that  $\alpha$  is the true deflection field, and that they are Gaussian

fields with power  $C_{\ell}^{\text{unl}}$ . The maximum a posteriori (MAP) unlensed CMB maps are formally given by the Wiener-filtered data<sup>1</sup>

$$X_{\alpha}^{\text{WF}} \equiv C^{\text{unl}} D^{\dagger} B^{\dagger} \text{Cov}_{\alpha}^{-1} X^{\text{dat}}. \quad (2.11)$$

The leg  $W_{\alpha}^a S^{\text{dat}}$  of the quadratic piece is then simply the deflected gradient of these maps

$$W_{\alpha}^a X^{\text{dat}}(\mathbf{x}) = D \nabla_a X_{\alpha}^{\text{WF}}(\mathbf{x}). \quad (2.12)$$

The other leg can be written as the inverse-noise-weighted residual between the data and how the inferred primordial modes are predicted to appear:

$$\begin{aligned} V_{\alpha} X^{\text{dat}} &= B^{\dagger} \text{Cov}_{\alpha}^{-1} X^{\text{dat}} \\ &= B^{\dagger} N^{-1} (\text{Cov}_{\alpha} - BDC^{\text{unl}} D^{\dagger} B^{\dagger}) \text{Cov}_{\alpha}^{-1} X^{\text{dat}} \\ &= B^{\dagger} N^{-1} [X^{\text{dat}} - BD X_{\alpha}^{\text{WF}}]. \end{aligned} \quad (2.13)$$

The calculation of these two terms is simple once the maps  $X^{\text{WF}}$  are reconstructed. Our implementation is discussed in Sec. III A.

The second part of the likelihood gradient is the contribution from the mean field

$$g_a^{\text{MF}}(\mathbf{n}) = \frac{1}{2} \frac{\delta \ln \det \text{Cov}_{\alpha}}{\delta \alpha_a(\mathbf{n})} = \langle g_a^{\text{QD}}(\mathbf{n}) \rangle. \quad (2.14)$$

The average here is over realizations of the data, with displacement  $\alpha$  held fixed. The second equality follows from observing that the first variation of a log-likelihood always vanish in the mean, and that  $g_a^{\text{MF}}(\mathbf{n})$  itself is independent of the data (and hence is equal to its expectation). The mean field serves the same purpose here as for the traditional quadratic estimator: to subtract the known sources of anisotropy from the quadratic estimate. It depends on the current estimate of the deflection, because  $\alpha$  at this iteration acts as a known source of anisotropy when measuring residual lensing at the next iteration. Implementations are discussed in Sec. III C and Appendix B.

## III. IMPLEMENTATION

This section presents some details of our implementation. The main numerical difficulty lies in the calculation of the Wiener-filtered modes  $X_{\alpha}^{\text{WF}}$  in the presence of the deflection  $\alpha$  (and anisotropic noise, masks, etc.). This is discussed in Sec. III A. The Wiener-filtering operation requires the inversion of the deflection field, described in Sec. III B, and the mean-field evaluation is discussed in Sec. III C. We also make use of curvature information to improve convergence, as discussed in Sec. III D. Finally, we describe our choice of starting point in Sec. III E, and summarize the workflow of the method in Sec. III F.

<sup>1</sup> For a signal seen under a linear response  $s^{\text{dat}} = R s^{\text{true}} + n$  the maximum a posteriori reconstructed signal  $s$  is given assuming Gaussian statistics by  $s^{\text{WF}} = (S^{-1} + R^{\dagger} N^{-1} R)^{-1} R^{\dagger} N^{-1} s^{\text{dat}} = S R^{\dagger} (S + N)^{-1} s^{\text{dat}}$

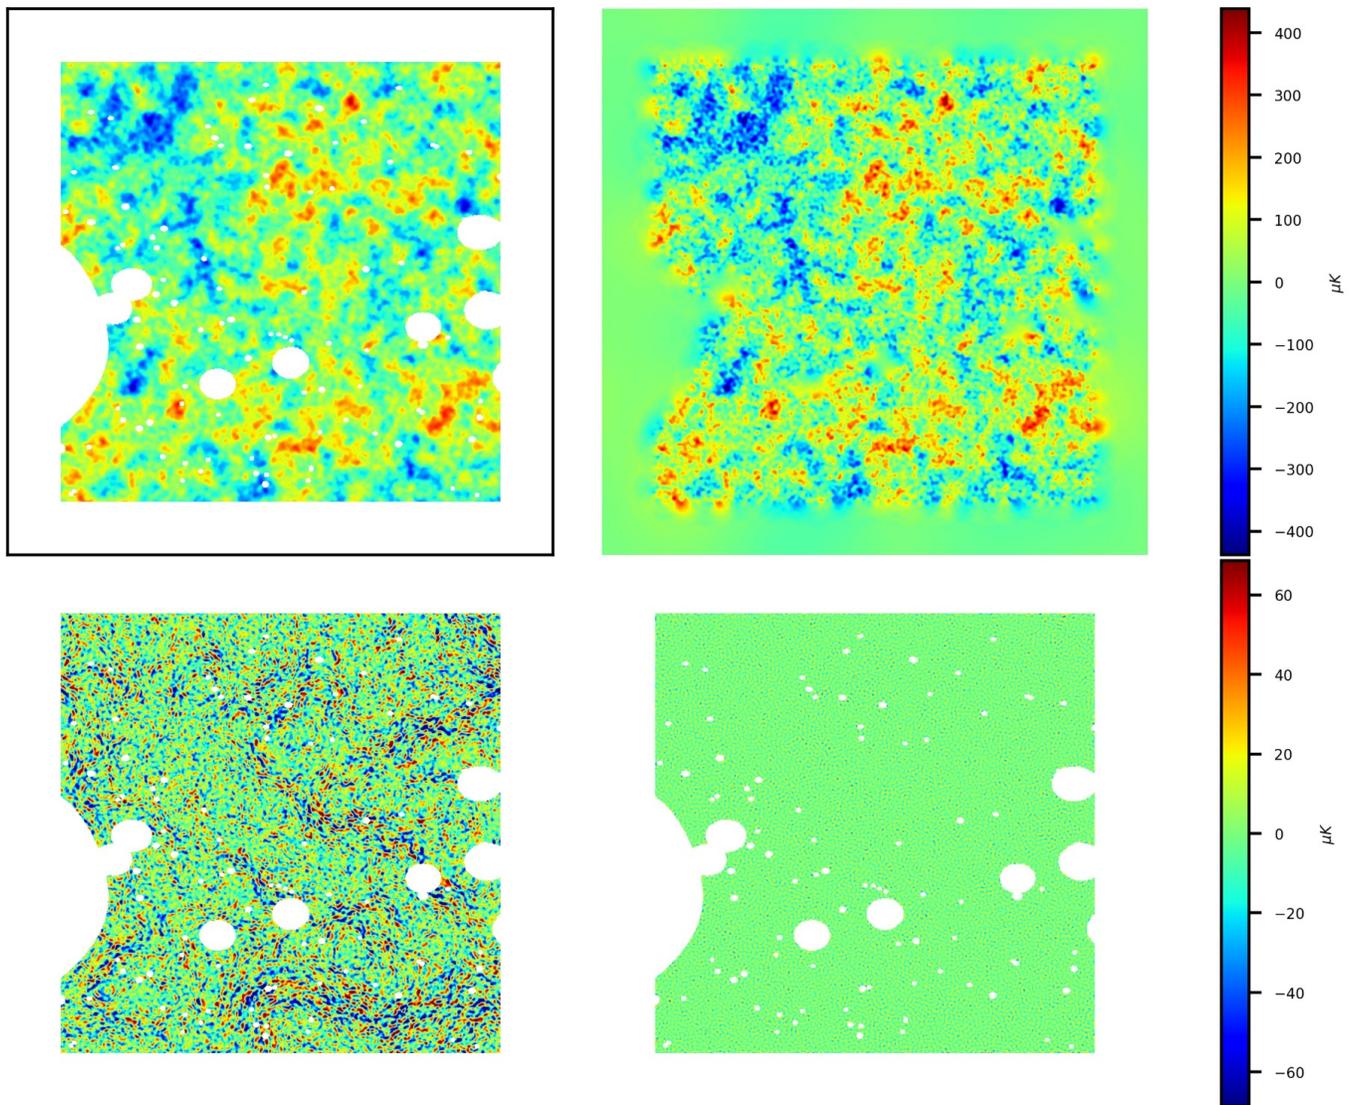


FIG. 1. A demonstration of how our Wiener filter produces optimal (maximum a posteriori) estimates on the unlensed CMB maps from masked data. The simulated temperature map is comparable to a *Planck* configuration, and we use the exact input simulated deflection and exact input unlensed  $C_l^{TT}$  spectrum in the filter. The upper-left panel shows the simulated masked temperature data map, with a homogeneous noise level of  $35 \mu\text{K-arcmin}$  and a beam FWHM of  $7\text{-arcmin}$ . The (unapodized) mask is built out of a portion of the public *Planck* lensing mask, to which we have added a band surrounding the patch on all sides. The upper right panel shows the reconstructed unmasked map  $T^{\text{WF}}$ . The residual to the true input CMB map ( $T^{\text{WF}} - T^{\text{input}}$ ) is shown on the lower-right panel. The lower-left panel shows the residual (on the unmasked pixels) of the result obtained when the Wiener filter instead uses no deflection but the lensed CMB spectrum in place of the unlensed spectrum (as in the standard quadratic estimator). These residuals are several times larger in magnitude (the same colour scale is sometimes saturated), and display the anisotropic swirly patterns generated by the pure gradient deflection field.

### A. Reconstruction of the unlensed CMB

The MAP estimate of the unlensed CMB given a deflection field  $\alpha$  is formally given by Eq. 2.11, and typically requires solving a large system of linear equations. This is not straightforward even in the absence of sky cuts or other non-ideal effects, since the deflection field breaks isotropy so that the harmonic transforms do not diagonalize the system. However, for all realistic situations we have investigated, we found

that the additional complication of the deflection field was minor in comparison to (typically highly anisotropic) sky cuts.

Our implementation is as follows. We first transform Eq. 2.11 to the following form

$$X_{\alpha}^{\text{WF}} = \left[ (C^{\text{unl}})^{-1} + D^{\dagger} B^{\dagger} N^{-1} B D \right]^{-1} D^{\dagger} B^{\dagger} N^{-1} X^{\text{dat}}, \quad (3.1)$$

and solve for the large inverse in brackets with conjugate gradient descent. Here one should understand the brack-

eted matrix to act on the space of non-zero unlensed CMB modes. There is no ambiguity regarding the unlensed spectra  $(C^{\text{unl}})^{-1}$ , since all modes in  $X^{\text{WF}}$  are exactly zero when they correspond to fiducial  $C_\ell^{\text{unl}}$  that are zero: the Wiener filter builds the maximum a posteriori  $X$  maps, hence these vanish whenever the prior variance  $C^{\text{unl}}$  does. We found Eq. (3.1) to be more efficient than other possible ways to perform the mask deconvolution, and it also is more efficient when the lensing operations are the only source of complications.

For the noise matrix  $N$  we use an input variance map that is diagonal in pixel space. The noise can be inhomogeneous, and we can also add to the noise matrix a set of templates which are projected out. This is useful for instance to account for poorly understood low- $\ell$  noise, or to project out any templates for galactic dust. Using conjugate gradient descent requires a reasonably fast way to apply  $(C^{\text{unl}})^{-1} + D^\dagger B^\dagger N^{-1} B D$  to vectors  $X$ . The first term is diagonal in harmonic space and poses no problem. The second term requires application of the lensing operators  $D$  and  $D^\dagger$ , and of the inverse noise matrix. From Eq. 3.2, applying the lensing operator to a map is simply achieved by harmonic transforms followed by lensing of the resulting map. As discussed in more detail in Sec. III B,  $D^\dagger$  also involves the inverse harmonic transform and a delensing operation (lensing with the inverse deflection field). It therefore has the same complexity as forward lensing, provided the inverse deflection field has been precomputed. The inverse noise matrix is simple under the assumption that it is diagonal in pixel space. The inclusion of templates is only a minor complication as long as there are only a reasonable number of them. Finally, assuming isotropic beams, the beam operations are fast in harmonic space.

All in all, application of the bracketed matrix in Eq. 3.1 requires 4 harmonic transform and 2 lensing operations, multiplied by 1, 2 or 3 for temperature only, polarization only or joint reconstruction respectively. Lensing of maps (or the displacement inversion) is not a cheap operation, but the cost scales linearly with the number of pixels and can easily be parallelized. Our implementations, including a GPU implementation, are discussed in Sec III B.

The use of a good preconditioner is mandatory for convergence in acceptable time, especially when dealing with masked maps. We use a multigrid preconditioner following Ref. [27], where a set of working resolutions is set up so that lower-resolution inverses are used to precondition those at higher resolution. Specifically, we extend the *qcin*v package<sup>2</sup> by Duncan Hanson to include the lensing operations. At the lowest resolution stage, we use a dense preconditioner. We offer no unique recipe of a good multigrid chain as performance appears to depend substantially on the specific configuration. The solution  $X_{\alpha_N}^{\text{WF}}$  obtained at iteration  $\alpha_N$  can however be used as starting point for iteration  $N + 1$ , which significantly speeds whole process as  $\alpha_N$  settles down to the converged estimate.

Finally, we note that in ideal situations where  $\alpha$  is the only

source of anisotropy, the use of a simple diagonal preconditioner is much faster, with no need to resort to a multigrid solution, at least up to noise levels of a CMB-S4 configuration we have been testing.

## B. Lensing and delensing operations

Lensing of maps is done at a resolution of 0.7 arcminutes, using a standard bicubic spline interpolation. Lensing is an expensive operation, even if easily computed in parallel, and there are a large number of maps to process until convergence is reached, and this can dominate the overall computational cost in typical runs. We found that porting the lensing on GPU, using a GPU-optimized implementation [28, 29] can provide substantial speed-up. This is one of the implementations that we provide.

In addition to the forward lensing operation, the filtering step also requires applying  $D^\dagger$ . This is equivalent to applying the inverse deflection together with multiplication by the magnification. To see this, consider the temperature part only. From the explicit form of the operator  $D$  in Eq. 3.2 we have

$$\begin{aligned} [D^\dagger T]_\ell &= \frac{1}{\sqrt{V}} \int d^2x e^{-i\ell \cdot (\mathbf{x} + \alpha(\mathbf{x}))} T(\mathbf{x}) \\ &= \frac{1}{\sqrt{V}} \int d^2x e^{-i\ell \cdot \mathbf{x}} |M_{\alpha^{-1}}|(\mathbf{x}) T(\mathbf{x} + \alpha^{-1}(\mathbf{x})). \end{aligned} \quad (3.2)$$

The second line follows from the first after the obvious change of variable  $\mathbf{x} \rightarrow \mathbf{x} + \alpha(\mathbf{x})$ , where  $|M|$  is the magnification matrix determinant that accounts for the change of volume element in these new coordinates:

$$[M_\alpha]_{ab}(\mathbf{x}) = \frac{\partial \alpha_a}{\partial x_b}(\mathbf{x}). \quad (3.3)$$

The inverse deflection  $\alpha^{-1}$  is defined by the condition that points deflected by  $\alpha$  are remapped to themselves

$$\mathbf{x} + \alpha^{-1}(\mathbf{x} + \alpha(\mathbf{x})) \equiv \mathbf{x}. \quad (3.4)$$

Eq. 3.2 has the simple form of the harmonic transform of the delensed temperature map, multiplied by the magnification of the inverse deflection. The generalization to polarization is immediate.

We therefore need to obtain the inverse deflection field. While some approximation to the inverse deflection is possible given the noise levels of current data [30, 31], we found the exact inversion is always well-behaved for a  $\Lambda$ CDM displacement, is always in the weak-lensing regime, and is not a bottleneck for our reconstruction. The inversion is a very localized operation which is easily parallelized, for which we use a simple real space Newton-Raphson scheme on a high-resolution grid. Specifically, following Ref. [31], we solve iteratively for  $\alpha^{-1}(\mathbf{n})$  using

$$\begin{aligned} \alpha^{-1}_{N+1}(\mathbf{n}) &= \alpha^{-1}_N(\mathbf{n}) \\ &- M_\alpha^{-1}(\mathbf{n} + \alpha^{-1}_N(\mathbf{n})) \cdot (\alpha^{-1}_N(\mathbf{n}) + \alpha(\mathbf{n} + \alpha^{-1}_N(\mathbf{n}))). \end{aligned} \quad (3.5)$$

<sup>2</sup> <https://github.com/dhanson/qcin>

In practice, we use the same 0.7 arcmin grid spacing that we use for the lensing operations, in which case 3 iterations starting from  $\alpha^{-1} = 0$  are enough for essentially exact inversion of a typical  $\Lambda$ CDM deflection field. Typical resulting r.m.s. fractional residuals on the deflection amplitude are as low as  $2 \cdot 10^{-5}$ . For lensing reconstruction in a realistic situation, the forward deflection is much smoother in comparison to a typical  $\Lambda$ CDM deflection owing to the prior effectively filtering out many small-scale modes, and coarser resolutions may also be used.

### C. Mean field evaluation

Provided with a large number of data simulations, the mean field may be evaluated using

$$\mathbf{g}^{\text{MF}}(\mathbf{n}) = \langle \mathbf{g}^{\text{QD}}(\mathbf{n}) \rangle, \quad (3.6)$$

i.e., by repeating the quadratic estimate on a number  $N_{\text{MC}}$  of independent simulations of the data maps and averaging to get the mean field. In practice, each of these naive estimates of  $\mathbf{g}^{\text{MF}}$  has spectrum  $|\mathbf{g}^{\text{QD}}|^2$  and a large resulting Monte-Carlo (MC) noise, containing the signal and noise parts of  $\mathbf{g}^{\text{QD}}$ . On small scales this noise is typically much larger than  $\mathbf{g}^{\text{MF}}$ , so for a reasonable number of MC simulations the mean field subtraction would effectively be adding noise with power  $|\mathbf{g}^{\text{QD}}|^2/N_{\text{MC}}$  to the estimation of the gradient on these scales. It is therefore desirable to obtain better ways to estimate the mean field. We suggest two types of trick to accelerate convergence of the mean-field estimation.

The first simply subtracts some of the Monte Carlo noise by subtracting a mean field calculated using an isotropic approximation to the data likelihood, using the same random phases for the simulations. This introduces no bias, since by isotropy the correction vanishes in the mean, but has the virtue of cancelling part of the MC noise where the anisotropy is mild. For example, the isotropic approximation could consist of recalculating the same quadratic estimate but setting  $\alpha$  to zero in the weights (in the absence of other non-ideal effects), or using a simulation extended to full sky in the presence of sky cuts.

The second trick is to modify the weights of the quadratic estimator, in a way that keeps its expectation value (i.e., the mean field) constant. This is discussed in more detail in Appendix B. Combined, these tricks can lead to orders of magnitude decrease of the MC noise on the mean-field estimator, drastically reducing the number of simulations required for the same target accuracy.

In principle, different random phases must be used for the simulations at each iteration step: usage of the same phases at each step causes artificial convergence of the iteration towards what is an approximation to the true posterior. This approximation might still be fairly good, however, if enough simulations are used. At any given scale it is the mean field MC noise that sets the accuracy at which the true maximum a posteriori deflection solution can be determined. We refer to this later on as the MC noise floor.

### D. Curvature

Finally, to perform an efficient search for the optimal point we need the curvature of the likelihood as well as the gradient. Specifically, to perform efficient Newton-type iteration across parameter space, the inverse curvature is needed. Curvature matrices such as

$$[H^{-1}]_{\mathbf{L}\mathbf{L}'}^{ab} \equiv -\frac{\delta^2 \ln p(X^{\text{dat}}|\boldsymbol{\alpha})}{\delta\alpha_{\mathbf{L}}^a \delta\alpha_{\mathbf{L}'}^{b,*}} \quad (3.7)$$

can never be evaluated exactly in reasonable time. We proceed as follows: starting with an initial isotropic guess,  $H_0$ , we perform a rank two update to  $H$  every time we move across parameter space. At each iteration, two maps are saved to disk and can be used to apply recursively the inverse curvature matrix to any vector. We use the limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) update [32], for which

$$H_{N+1} = (1 + \rho s y^t) H_N (1 + \rho y s^t) - \rho s s^t, \quad (3.8)$$

with

$$s(\mathbf{n}) = \phi_{N+1}(\mathbf{n}) - \phi_N(\mathbf{n}), \quad y(\mathbf{n}) = g_{N+1}^{\text{tot}}(\mathbf{n}) - g_N^{\text{tot}}(\mathbf{n}), \quad (3.9)$$

and  $\rho = 1/y^t s$ . Built in this way, the inverse curvature takes into account non-Gaussian and realization-dependent aspects of the likelihood, and we found can dramatically improve the convergence properties of the iterative search. Since the inverse curvature approximates the covariance, it can also be used to assess the width of the posterior density function, giving us approximate confidence regions for free at the end of the iterative process.

### E. Starting point

If the posterior density were exactly Gaussian, a single Newton step starting from  $\alpha \equiv 0$  would bring us directly to the optimal solution. This solution matches (neglecting the difference between the realization-dependent curvature and its average) the Wiener-filtered quadratic estimator calculated with unlensed weights [33]. However, we use lensed weights since they provide a better quadratic reconstruction [34], especially on large scales. Explicitly, we use

$$\alpha_0(\mathbf{L}) = \frac{C_L^{\phi\phi}}{C_L^{\phi\phi} + N_L^{0,\text{lens}}} i\mathbf{L} \hat{\phi}^{\text{qest}}(\mathbf{L}), \quad (3.10)$$

where  $N_L^{0,\text{lens}}$  is the Gaussian reconstruction noise of the quadratic estimator, calculated with the lensed CMB spectra.  $N_L^{0,\text{lens}}$  is calculated using its real space flat-sky representation (see Appendix B). The quadratic estimator is the minimum variance (MV) estimator built with the set of maps considered:  $T$  alone,  $Q$  and  $U$  polarization alone, or the three Stokes maps in combination. The filtering step is described in Sec. III A, with the ( $\alpha = 0, C_\ell^{\text{lens}}$ ) Wiener filter using an input noise variance map and fiducial beam transfer function. In temperature, there are known ways to optimize the weights further [35], but Eq. (3.10) works well for our purposes.

## F. Summary of the workflow

We are now in position to summarize and describe the workflow of the iterative search for the maximum a posteriori point. Initially, the displacement  $\alpha_0$  is set at the Wiener-filtered quadratic estimator as described in Sec. III E. To get the optimal reconstruction we apply the following steps recursively until satisfactory convergence is reached:

1. The displacement  $\alpha_N$  is inverted to give  $\alpha_N^{-1}$  and cached.
2. With the deflection and its inverse, the delensed maps  $X_\alpha^{\text{WF}}$  are obtained with the  $(\alpha_N, C_\ell^{\text{unl}})$  Wiener filter using multigrid-preconditioned conjugate gradient inversion. This is the most expensive step of the whole process by some margin.
3. With the delensed CMB at hand, the quadratic part of gradients are calculated from Eqs. (2.12) and (2.13).
4. In parallel (unless neglected, or some other approximation scheme is used), the mean field contributions to the gradient are calculated by repeating steps 2-3 on a number of simulated maps, using the tricks discussed in Sec. B.
5. The total gradient  $g_N$  is then obtained, the inverse curvature updated according to the BFGS scheme of (3.8), and the displacement  $\alpha_{N+1}$  is found along the Newton descent direction:

$$\alpha_{N+1} = \alpha_N + \lambda H_N g_N. \quad (3.11)$$

The parameter  $\lambda$  helps improve convergence. For CMB-S4-like configurations, we picked  $\lambda = 1/2$  at all steps, whereas the full Newton step  $\lambda = 1$  can safely be used at higher noise levels.

## IV. RESULTS

For this preliminary investigation we report three tests of our reconstruction method. First, using a simulated lensed map similar current *Planck* public data, we demonstrate use of the Wiener-filtering procedure to extract the maximum a posteriori estimate of the unlensed CMB. Second, we simulate a lensing reconstruction from polarization on a masked field, with noise levels of next-generation CMB experiments. **In its simplest incarnation, the algorithm does not directly provide a delensed primordial  $B$ -mode map, but the iteratively reconstructed potential map can still be used to delens the observed  $B$  modes.** Third and finally, while explicit delensing of  $B$ -modes with these improved lensing maps will be demonstrated in upcoming work, we discuss the increase in correlation coefficient (and hence delensing efficiency) that the method can achieve with upcoming CMB data.

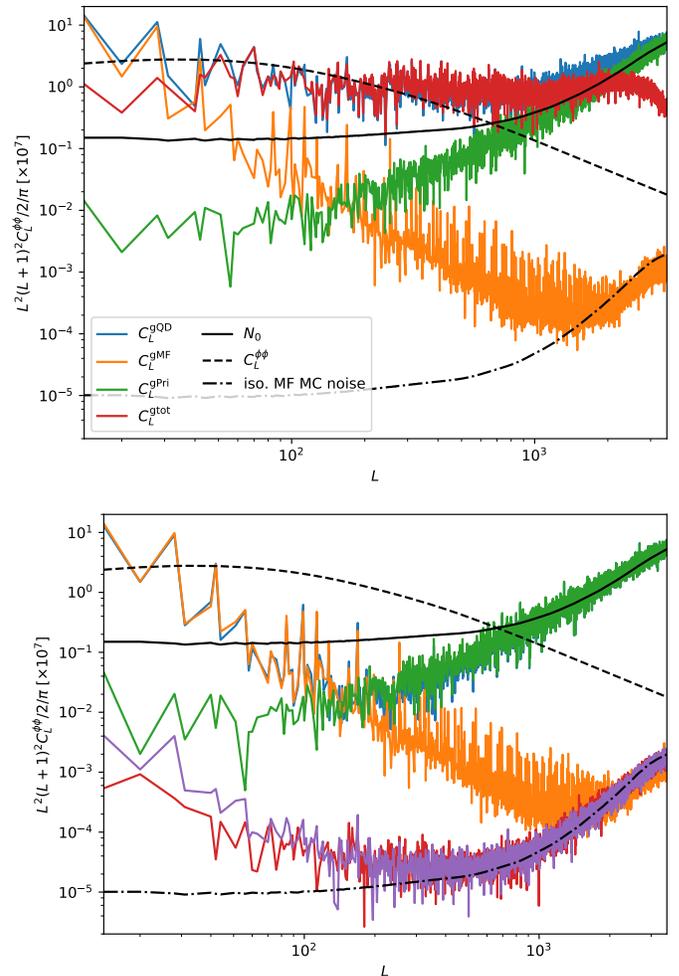


FIG. 2. Power spectra of the three gradients  $g^{\text{QD}}$  (blue),  $g^{\text{MF}}$  (orange) and  $g^{\text{PR}}$  (green) for lensing reconstruction from polarization on a masked patch, together with the total gradient spectrum (red). The algorithm works by reducing the red curve as much possible to find the most probable lensing map. The mask is shown in Fig. 1, and causes the large contribution from the mean field at low multipoles. All curves are normalized to the quadratic estimator normalization  $N_L^0$ . The upper panel shows the gradient spectra at the first iteration step, where the deflection is the Wiener-filtered quadratic estimator, and the lower panel shows the result after 20 iterations. At this point, the gradient has hit the mean-field MC noise floor (purple on the lower panel) on all scales and the solution cannot be improved by more iterations. The MC noise floor is mean-field estimator dependent and inversely proportional to the number of simulations used (here 511 per iteration). The dot-dashed line shows predictions for the MC noise floor built from an isotropic likelihood, which are inaccurate at low multipoles because of the sky cuts. At low multipoles, the improved reconstruction relies on accurately cancelling the mean field contribution, but on intermediate scales the decrease in the quadratic estimate is immediately visible.

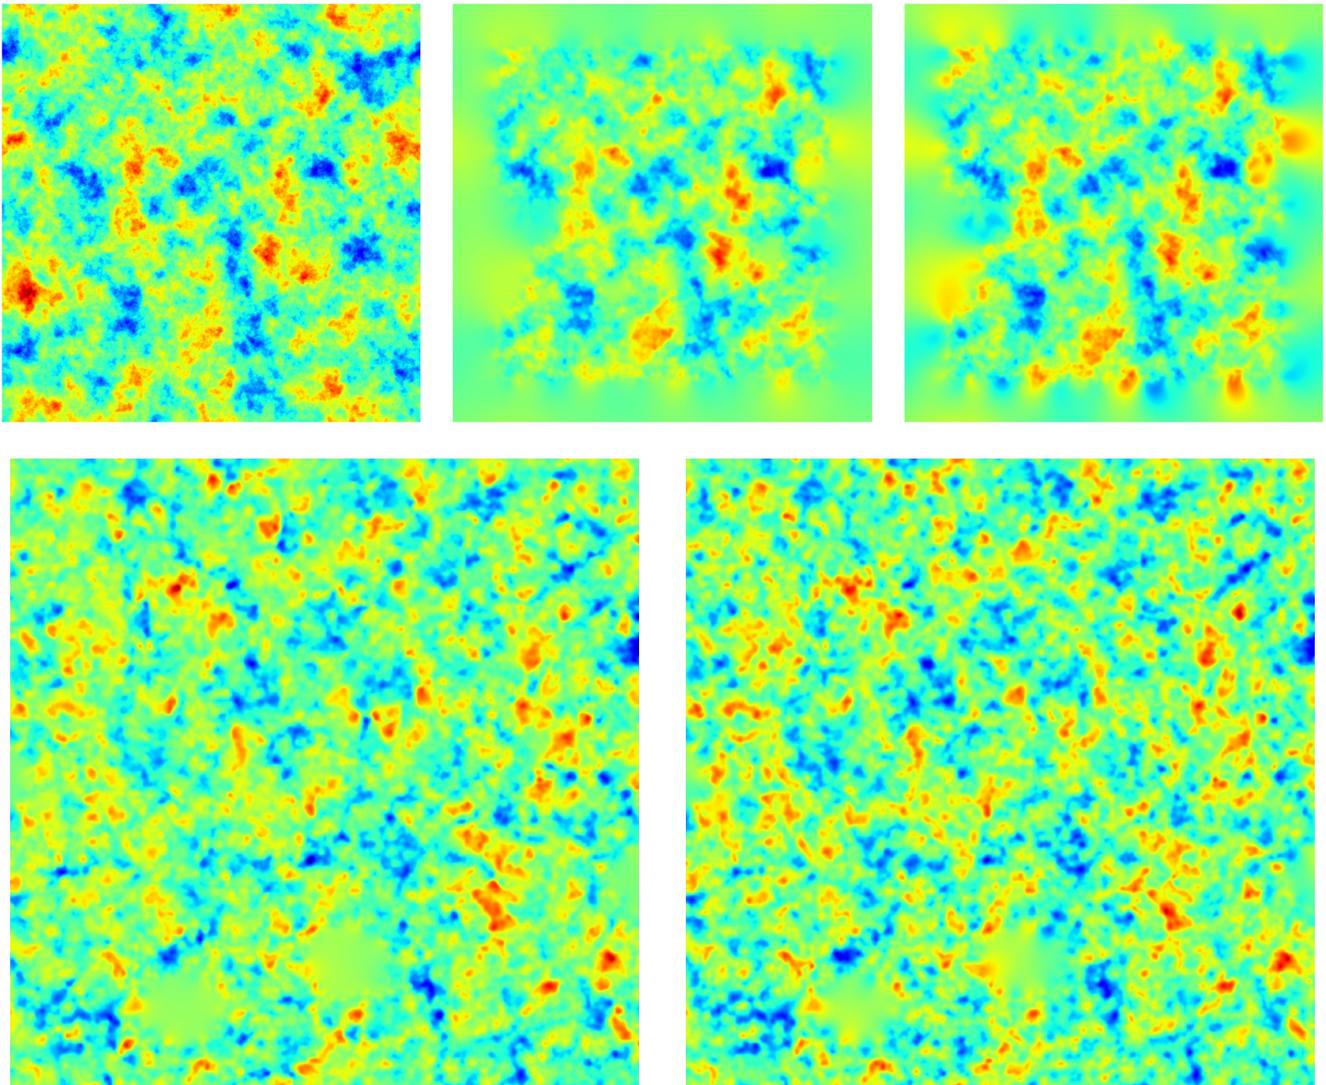


FIG. 3. The top-left panel shows a simulated lensing field used as input to a lensing reconstruction analysis. The input Stokes  $Q$  and  $U$  maps are masked with the same mask shown on Fig. 1, and we assume a polarization-sensitive experiment having polarization noise level of  $1.5 \cdot \sqrt{2} \mu\text{K-arcmin}$  and 3 arcmin FWHM beam. The middle panel shows on the same colour scale the Wiener-filtered quadratic estimate, which is the starting point of the iterative solution for the maximum a posteriori solution, Eq. 3.10. The top-right panel shows the converged solution. The top three panels show the displacement-like scalar field with transform  $|\ell|\hat{\phi}_\ell$ , Eq. (4.5). The lower two panels show the convergence maps  $\kappa(\ell) = -\frac{1}{2}\ell^2\hat{\phi}(\ell)$  (only the central regions covering one fourth of the map) for the quadratic estimator (left) and iterative solution (right). The iterative solution can resolve structure down to smaller scales, and improvement can also be seen in the masked regions.

### A. Wiener filtering

Before turning to lensing reconstruction, we first demonstrate our Wiener filtering technique to accurately estimate the unlensed CMB from a simulated masked temperature CMB map. We convolved an input simulated lensed CMB sky with a beam of 7-arcmin FWHM, and added homogeneous  $35 \mu\text{K-arcmin}$  noise. Using the exact input lensing potential and input unlensed spectrum, we obtained the reconstructed unlensed CMB modes  $T^{\text{WF}}$ , up to  $\ell \leq 3500$ . The mask was chosen arbitrarily as a piece of the public *Planck* 2015 analysis lens-

ing mask, and we have further excised a band along all sides of the  $600 \text{ deg}^2$  patch so that the boundary of the unmasked patch is non-periodic. This leaves  $f_{\text{patch}} = 60\%$  of the patch unmasked.

The upper panels of Fig. 1 show a comparison of simulated lensed data and the estimated unlensed CMB ( $T^{\text{WF}}$ ) using the known input deflection field. The lower panels compare the residuals  $T^{\text{WF}} - T^{\text{input}}$ , using either the deflection-dependent Wiener filter ( $\alpha = \nabla\phi^{\text{input}}, C_\ell^{\text{unl}}$ , right), or the usual quadratic estimate filter that we use for the initial estimate ( $\alpha = 0, C_\ell^{\text{len}}$ , left). The former has small nearly-

isotropic, near-uniform residuals set by the noise map, the latter shows the swirly patterns characteristic of lensing, with residuals of much larger amplitude.

It takes a couple of minutes on a modern laptop to reconstruct these modes, up to the point where the residual norm of the solution to the linear system of equations has decreased across the full patch by 5 orders of magnitude. At this point the values inside the masked regions have converged to sub-percent level.

## B. Lensing reconstruction on a masked patch

Next, we demonstrate lensing potential reconstruction from a simulated futuristic polarization-based experiment, including masking. We consider, on the same mask, the iterative reconstruction of the lensing potential from Stokes  $Q, U$  simulated maps with noise level of  $1.5 \cdot \sqrt{2} \mu\text{K-arcmin}$ , and a beam of 3 FWHM. This corresponds roughly to noise levels expected for the baseline, widefield CMB-S4 configuration. We use a vanishing fiducial gravitational wave amplitude,  $C_\ell^{BB, \text{unl}} = 0$ , in which case the Wiener filter reconstructs the optimal unlensed  $E$  mode map from  $Q$  and  $U$ , assuming the  $B$  mode is due to lensing and noise exclusively. We reconstruct multipoles of the  $E$ -mode map up to  $\ell \leq 3500$ , and the lensing potential over the same multipole range. The search converges just as well to higher multipoles, but there is little information there since the reconstruction becomes completely noise dominated. At no point do we apply low multipole cuts to the lensing potential, demonstrating that the iterative search can handle masking and the resulting large mean field adequately.

Convergence of the iterative search towards the optimal solution can be explicitly checked on all scales by monitoring the change in the posterior gradient and its components. The upper panel of Fig. 2 show the power spectra of the different pieces of the posterior gradient  $g$  at the starting point of the iteration, for which the lensing reconstruction is given by the MV estimate. These gradient maps are smooth across the mask, so to estimate the spectra we simply rescale naive spectra estimates by  $1/f_{\text{patch}}$  for the purpose of this figure. Assigning Fourier modes to  $L$ -bin according to  $L = |\ell| - 1/2$ , resulting in  $n_L \sim (2L + 1)V/4\pi$  modes per bin, we build

$$\hat{C}_L^g = \frac{1}{f_{\text{patch}}} \frac{1}{n_L} \sum_{\ell \text{ in } L \text{ bin}} |g_\ell|^2. \quad (4.1)$$

Shown are the quadratic piece (blue), the mean field (orange) and the prior (green) spectrum. The total gradient spectrum is shown as the red line. For easier comparison, all gradients have been normalized with the quadratic estimator normalization  $N_L^0$ , so the quantity plotted is actually

$$(N_L^0)^2 \hat{C}_L^g, \quad \text{for } g = g^{\text{QD}}, g^{\text{MF}}, g^{\text{PR}} \text{ and } g^{\text{tot}}. \quad (4.2)$$

With this normalization, the reconstruction noise in the quadratic estimate before any iteration is  $N_L^0$  itself, shown as the solid black line. On scales with large signal to noise

( $N^0 \ll C_L^{\phi\phi}$ ) the inverse curvature is (to a crude approximation)  $\sim N_L^0$ . Thus, on these scales,  $(N_L^0)^2 C_L^{g^{\text{tot}}}$  is also roughly the spectrum of the Newton increment in Eq. 3.11 added to the potential estimate at the corresponding iteration step.

The quadratic part of the gradient is pure reconstruction noise at small scales and is cancelled by the prior gradient. On large scales, the quadratic piece is dominated by the mean-field contamination from the mask. Since the mask mean field is basically independent from  $\hat{\phi}$ , it will vary little from iteration to iteration, and improving the potential estimate (by reducing the total gradient) at low multipoles demands precise evaluation of this term. On the other hand, on intermediate scales we can see that  $g^{\text{QD}}$  is the dominant contribution to the total gradient at the start of the iterations. Hence, on these scales the lensing map can be improved without relying on cancellation of the mean field.

We used 511 simulations at each step to estimate the mean field. The upturn of the orange curve at  $L \simeq 1000$  shows the onset of the MC noise dominated regime, where the mean-field estimate becomes pure MC noise. The dot-dashed black line shows an analytic prediction for the expected MC noise neglecting sky cuts, calculated with the tools from Appendix B. The MC noise spectrum is smaller than  $N_L^0$  by four orders of magnitude, so iterations should be able to reduce the gradient amplitude by a similar amount. On large scales, the MC noise stays smaller than the prior gradient, and thus the iterative procedure will have exhausted information from the data before it hits the MC noise floor. However, we will see below that the isotropic prediction for the MC noise is inaccurate on large scales, where there is a substantial contribution from sky cuts. A small contribution also comes from the deflection field. It is possible to predict this contribution perturbatively, since the mean field just follows the spatial distribution of the deflection field at each step: see Appendix B. This contribution from the deflection is comparatively larger for temperature reconstruction, and also for temperature in combination with polarization.

We start the posterior inverse curvature  $H^0$  with the isotropic estimate

$$H_L^0 = \left( \frac{1}{N_L^{0, \text{unl}}} + \frac{1}{C_L^{\phi\phi}} \right)^{-1}. \quad (4.3)$$

The second term is the prior curvature, and for the first term (the likelihood curvature) we used the unlensed weights. The choice of initial curvature is not critical as long as the BFGS scheme is used to update it. However, using the lensed weights can lead to the algorithm taking steps that are too large and hence give poor convergence, so we use unlensed weights instead (which gives a curvature that is slightly too large, but works well): steps that are too large should be avoided, as the search relies on the displacement being invertible at each step. More optimal curvature estimates might be built using partially lensed weights.

Convergence is acceptably quick, and after  $\sim 9$  iterations the bulk of the improvement has been gained, with the gradient spectrum reduced by 2–3 orders of magnitude. Only small

variations on the large-scale modes are visible in the deflection maps after this point. The lower panel of Fig 2 shows the spectra after 20 iterations. The purple curve show an empirical estimate of the mean-field MC noise. This is calculated by splitting our set of simulations into two independent sets of size  $N_1$  and  $N_2$  with corresponding mean-field prediction  $g^{\text{MF}_1}$  and  $g^{\text{MF}_2}$ , and building

$$\hat{C}_L^{\text{MF}} = \frac{N_1 N_2}{(N_1 + N_2)^2} \hat{C}_L^{\text{MF}_1 - \text{MF}_2}. \quad (4.4)$$

This is much larger than the analytic isotropic prediction on large scales because of the mask contribution to the MC noise. The total gradient closely follows the MC noise curve, and no further improvement can be achieved after the 20 iterations. At intermediate scales, the quadratic gradient is visibly much reduced and is now in equilibrium with the prior.

Finally, we show the reconstructed lensing map in Fig. 3. From left to right in the top row we show the input lensing map, the quadratic estimate, and the converged iterative solution. Here we plot the displacement-like but isotropic spin-0 transforms

$$d(\mathbf{x}) \equiv \frac{1}{\sqrt{V}} \sum_{\mathbf{L}} L \hat{\phi}_{\mathbf{L}} e^{i\mathbf{L}\cdot\mathbf{x}}. \quad (4.5)$$

The reconstruction is visibly improved, both by large-scale modes filling in the masked regions and by the presence of finer-grained structure well inside the patch. The bottom row of Fig. 3 shows a zoom in of the central area, showing instead the lensing convergence where the improvement on small scales is more clearly visible.

### C. Delensing efficiency

How can our lensing reconstruction method help with measurement of primordial tensor modes? In the absence of a fiducial non-vanishing  $C_\ell^{BB}$ , for which there is at present no preferred choice, no delensed  $B$ -mode map is directly produced by the algorithm as the prior sets it to zero. However, it is well known that the lensing map can be used to remove some of the lensing signal in the observed  $B$ -mode map. Reduction of lensing signal in the  $B$ -mode map will result in some degree of improvement on tensor constraints, since the lensing  $B$  modes act as a source of noise for any primordial signal.

Delensing of  $B$ -mode polarization has recently been demonstrated on *Planck* data by remapping the Stokes maps [31], and by the SPT team [36] using a template subtraction method. In both cases, the expected reduction of lensing-like power is approximately set by the squared cross-correlation coefficient of the measured lensing map to the true lensing map, which we call the delensing efficiency:

$$\epsilon_L \equiv \frac{\left(C_L^{\hat{\phi}\hat{\phi}}\right)^2}{C_L^{\phi\phi} C_L^{\hat{\phi}\hat{\phi}}}. \quad (4.6)$$

Fig. 4 shows this cross-correlation coefficient for simulated reconstructions. We built these curves using 128 idealized simulations, with homogeneous input noise maps and no sky cuts. In this case, the deflection field is the only source of mean field at each step, which is sufficiently well described by the perturbative predictions derived and discussed in the Appendix B. In all cases, we have considered joint temperature and polarization (MV) reconstruction, with sharp multipole cuts  $10 < \ell \leq 3000$ . For the *Planck* curve, we cut at 2048, following the public analysis [37]. Shown are the delensing efficiencies expected for the MV quadratic estimator (dashed colour) and the iterated, converged solution (solid). Besides *Planck*, we also show curves for the *Simons Observatory*<sup>3</sup>, and two distinct CMB-S4-like configurations: one for a wide but shallow coverage, and one for a deep survey with sensitivity increased by a factor of about four. The assumed beam and noise levels are shown in Table I, ignoring all experimental complications, and are not meant to be necessarily very accurate representations of the experiment label. We have used Gaussian beams of 3 arcmin FWHM in all cases, again with the exception of *Planck* where we used 6.5 arcmin. The S4-wide configuration is identical to those of the optimal reconstruction performed on the masked sky in Sec. IV B. We show an estimate of the efficiency of this reconstruction as the green data points. The points are obtained from Eq. 4.6, using pseudo- $C_L$  estimates after enlarging the mask conservatively near the mask boundaries, leaving  $f_{\text{patch}} \sim 35\%$ , in order to avoid any edge effects. The points stand in very good agreement to expectations, demonstrating that masking does not substantially affect the reconstruction quality away from the mask boundaries.

Fig. 4 also shows the delensing efficiency reachable using the publicly available<sup>4</sup> GNILC Cosmic Infrared Background (CIB) map [38] as an external tracer of the lensing map (purple points). We used the GNILC reconstruction at 545 GHz, covering 60% of the sky, and *Planck* 2015 lensing potential map to build these points. We estimated the CIB auto spectrum and the lensing-CIB cross-spectrum on the union of their released masks after apodization on a scale of 12 arcmin, deconvolving the pseudo- $C_\ell$  estimates from the mask coupling matrix. We show the CIB efficiency estimate

$$\epsilon_L^{\text{CIB}} \equiv \frac{\left(\hat{C}_L^{\text{CIB}\hat{\phi}}\right)^2}{\hat{C}_L^{\text{CIB CIB}} C_L^{\phi\phi, \text{fid}}} \quad (4.7)$$

as the brown data points. The fiducial lensing spectrum is based on the *Planck* 2015 cosmology. This estimate is justified in so far as the lensing map is an unbiased tracer of the true lensing, and in the absence of spurious cross-correlation between the two maps. Comparison to previous works on CIB delensing (Fig. 2 of Ref. [39], using *Planck* cleaned 545 GHz map, and Fig. 1 of Ref. [36] from the SPT team, using *Her-*

<sup>3</sup> [www.simonsobservatory.org](http://www.simonsobservatory.org)

<sup>4</sup> <http://pla.esac.esa.int>

schel 500  $\mu\text{m}$  map) shows good consistency<sup>5</sup> over the relevant scales.

On this large sky fraction (60%), contamination by galactic dust may reduce somewhat the cross-correlation to the lensing in the GNILC CIB map, and the brown points are slightly lower than expected for clean maps [40, 41]. Larger CIB efficiencies might be possible on cleaner regions of the sky, or with improved dust cleaning from future observations. For comparison we also show as the purple data points the efficiency on a smaller but cleaner area, using a mask built by thresholding the GNILC dust map at 545 GHz, keeping only 4% of the sky unmasked.

Not all multipoles are equally important for the purpose of  $B$ -mode delensing. To a good approximation, the  $B$  power  $C_\ell^B$  depends linearly on the lensing deflection spectrum, hence we may write the delensed  $B$ -mode power as

$$C_\ell^{B,\text{delens}} \sim \sum_L (1 - \epsilon_L) \frac{\partial C_\ell^B}{\partial \ln C_L^{\phi\phi}}. \quad (4.8)$$

If the very-low  $\ell$  reionization peak cannot be probed or is discarded, the tensor-mode recombination peak (at roughly  $40 \leq \ell \leq 100$ ) determines the scale where delensing is most important. The black line on Fig. 4 shows  $LdC_\ell^{BB}/d\ln C_L^{\phi\phi}$ , after averaging over this  $\ell$ -multipole range, and normalized such that it  $L$ -integrates to unity. By construction, weighting the efficiency curves on this figure against this line gives the delensing efficiency relevant for primordial  $B$  modes around the recombination peak.

Effective residual delensed  $B$ -mode noise amplitudes are listed on the second set of rows of Table I. The delensed  $B$ -mode lensing power is calculated from the unlensed  $E$  spectrum and a reduced lensing spectrum given by

$$C_L^{\phi\phi,\text{delens}} = (1 - \epsilon_L) C_L^{\phi\phi}, \quad (4.9)$$

where the efficiencies are those shown in Fig. 4. The numbers in the table are the mean power over  $40 \leq \ell \leq 100$ . The partially lensed spectra as well as the coupling matrix  $dC_\ell^{BB}/d\ln C_L^{\phi\phi}$  are obtained with the Python `camb` package<sup>6</sup>. The *Planck* number matches well the result of the  $B$ -mode delensing analysis performed by Ref. [31] on data. We also give predictions for the iterated solution. These predictions are obtained following Ref. [22] by iteratively producing delensed power spectra and MV reconstruction noises  $N_L^0$ , using at each step the reduced lensing power in Eq. 4.9 with efficiencies

$$\epsilon_L = \frac{C_L^{\phi\phi}}{C_L^{\phi\phi} + N_L^0} \quad (4.10)$$

to calculate the partially delensed  $B$ -mode power used when calculating  $N_L^0$  for the next step. The predictions stand in excellent agreement with our simulated reconstructions.

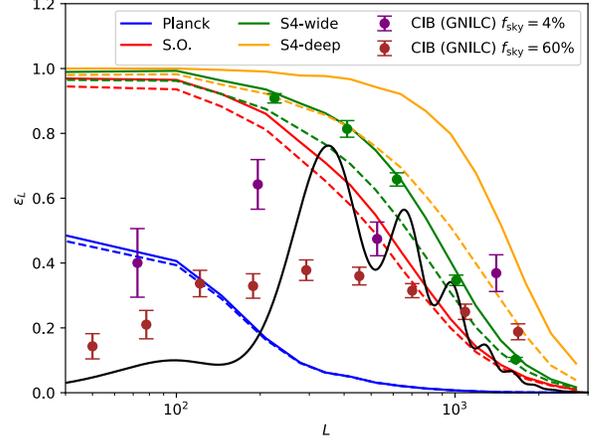


FIG. 4. The delensing efficiencies as function of lensing  $L$  multipole with either the quadratic estimator (dashed coloured lines) or the iterative solution (solid coloured lines), for current and futuristic noise levels. The curves were obtained from 128 simulated spectra and cross-spectra of the input lensing with idealized quadratic and iterative reconstructions. The green points show for comparison an estimate of the efficiency from the reconstruction on the masked sky described in Sec. IV B, which has identical noise level to the corresponding green curve. Only data far away from the mask edges were used to produce these points. Also shown are estimates of the efficiency using CIB maps, obtained as discussed in the main text from the public *Planck* GNILC maps at 545 GHz. The brown points show the efficiency obtained from 60% of the sky, while the purple points were obtained on the cleanest (according to the GNILC dust map) 4% of the sky. The black line shows the contribution per log-multipole bin of  $C_L^{\phi\phi}$  to the total  $B$ -mode power on the scales relevant for a primordial  $B$ -mode measurement, see Eq. 4.8. The average of the coloured curves weighted by the black line gives the approximate delensing efficiency relevant to each observation. The residual lensing  $B$  power is listed in Table I together with the noise levels and expected improvement on tensor-to-scalar ratio constraints.

Table I also shows the expected delensing improvement of constraints on the tensor-to-scalar ratio  $r$ , comparing results using the iterative lensing estimator to those using the quadratic estimator. These are calculated from a toy  $r$  estimator variance estimate, assuming  $r = 0$ ,

$$\frac{1}{\sigma^2(\hat{r})} = \frac{f_{\text{sky}}}{2} \sum_{\ell \geq 40} (2\ell + 1) \left( \frac{C_\ell^{B \text{ tensor}, r=1}}{C_\ell^{B \text{ delens}} + C_\ell^{B \text{ noise}}} \right)^2. \quad (4.11)$$

The delensed  $B$  power is the one calculated according to Eq. 4.9. More realistic forecasts, for example including foreground cleaning, are well beyond the scope of this paper, but we note that these ratios stand in agreement with expectations [11].

<sup>5</sup> Note that both references show the cross-correlation coefficient  $\rho_L$ , while we show the efficiency  $\rho_L^2$

<sup>6</sup> `camb.readthedocs.io`

TABLE I. The first three rows give the temperature and polarization noise levels and beam width input to the simulations used in Sec. IV C to obtain the delensing efficiencies shown on Fig. 4. The next two rows show the effective  $B$ -mode lensing power achievable on the scales relevant for primordial  $B$ -mode measurement, using the quadratic (MV) estimator and the iterated solution respectively. The latter results accurately match predictions using an iterated Gaussian noise level as described in the main text, shown on the sixth row. The last two rows show the fractional improvement on the error bar of the tensor to scalar ratio  $r$  (compared to the case of no delensing, assuming  $r = 0$ ). These results are for the idealized case of no foreground or mean-field contamination, and show some sensitivity to the largest multipole  $L$  that can be delensed. Ratios calculated using a lens multipole cut at  $L_{\min} = 100$  instead of  $L_{\min} = 40$  are shown in parentheses.

	Planck	S.O.	S4-wide	S4-deep
$N_{\text{lev}}^T / (\mu\text{K arcmin})$	35	3.0	1.5	0.38
$N_{\text{lev}}^P / (\mu\text{K arcmin})$	55	4.2	2.1	0.53
Beam FWHM / (arcmin)	6.5	3.0	3.0	3.0
$B^{\text{len}} / (\mu\text{K arcmin})$ (quadratic)	4.5	3.2	2.8	2.3
$B^{\text{len}} / (\mu\text{K arcmin})$ (iterative)	4.5	3.0	2.4	1.4
pred. (iter. $N_L^0$ )	4.5	3.0	2.4	1.4
$\sigma(r)$ impr. (quadr.)	1.0	1.7(1.5)	2.5(2.2)	4.1(3.3)
$\sigma(r)$ impr. (iter.)	1.0	1.8(1.6)	3.1(2.6)	10.6(6.3)

## V. SUMMARY

We presented an iterative method for CMB lensing reconstruction, and showed that for future high-sensitivity observations it can produce substantially better results than the quadratic estimator. Even with non-trivial masking, the method remains numerically tractable and produces results in agreement with naive expectations. For low noise levels the large-scale lensing modes are all reconstructed with high signal to noise, even by the quadratic estimator, so the cosmological information is limited by cosmic variance. The main information gain from the iterative estimator comes on smaller scales where the quadratic estimator reconstruction noise starts to be substantial. However, for  $B$ -mode delensing, even small errors on the reconstruction of the large-scale lensing realization can lead to residual lensing  $B$ -mode power, so the improvement in signal to noise is important on all scales.

The algorithm works by extracting residual lensing from optimally reconstructed unlensed CMB maps. As such, it produces both estimates of the lensing potential and the delensed CMB maps. Note, however, that the method does not directly produce a delensed  $B$ -mode map, unless a prior spectrum is adopted for the unlensed  $B$ -mode spectrum. Nevertheless, the resulting deflection estimate, alone or in combination with the delensed  $E$  map, may be used to delens the observed polarization map, giving improved delensing efficiency compared to using a quadratic estimator reconstruction.

The algorithm maximizes the posterior probability for the lensing potential, assuming Gaussianity of the unlensed maps and noise. A solution to this same problem was first attempted in Refs. [13, 23]. These references, working in the absence of non-ideal effects, introduced several approximations to reduce the computational burden, avoiding in par-

ticular the anisotropic inverse variance filtering step. For similar reasons involving the difficulty of a global analysis, Ref. [42] introduced a local likelihood reconstruction method, where the lensing map is approximated as quadratic in small neighborhoods, and large wavelengths are ignored. We have demonstrated how a conjugate gradient inversion can handle the global inversion very efficiently, and, crucially, can also successfully be applied in the presence of sky cuts and other realistic non-idealities. Our solution is the first that does not rely on approximations once the fiducial ingredients of the likelihood and prior have been chosen. This means that, given enough computational resources, the resulting lensing potential map is optimal and cannot be improved upon.

In practice, one limiting factor is the mean-field calculation. Unless some approximation is used, the mean field is calculated with a finite number of simulations, and this sets a Monte-Carlo noise floor that cannot be improved upon by further iterations. However, we demonstrated that for realistic situations reconstructions can be successfully performed on masked data for current and next-generation CMB experiments. We also showed how a perturbative approximation to the mean field is adequate in the absence of non-ideal effects, allowing very fast iterative reconstructions in this case. We expect the methods and codes described and tested here to be useful for the planning and execution of future CMB lensing analyses.

The modular, fully parallelized pipeline (using MPI) is written in Python, internally calling parts written in C, and/or sending these to a GPU device using the pyCUDA interface [43]. The flat-sky code is publicly available<sup>7</sup>. We also described the curved-sky algorithm; this will be tested on data and reported elsewhere.

Our successful exploration of the lensed CMB likelihood suggests several interesting possibilities for future investigation and improvement. The iterative estimate takes as an input the fiducial unlensed CMB spectra, which we have taken to include no primordial  $B$  modes so that our posterior (MAP) estimate of the unlensed  $B$  modes is exactly zero. This prevents us directly obtaining an optimal measurement of a delensed gravitational wave signal, which must be obtained afterwards using a more standard template subtraction or point remapping method. By allowing for non-zero unlensed  $B$  modes in the prior, we could also allow direct joint estimation of the lensing together with the primordial signal. Exactly how best to do this, given the unknown amplitude of the primordial signal and complications with delensing biases, is worth careful future consideration. Another important future direction is to go beyond estimation of the lensing map to also provide optimal lensing power spectrum estimates (and estimates of the delensed CMB power spectra). Within the maximum posterior density framework, building a posterior density for the lensing power spectrum formally requires an intractable marginalization over the deflection field, though approximations can certainly be built [13] and are worth further study.

<sup>7</sup> <https://github.com/carronj/LensIt>

Finally, this also opens exciting prospects for cluster CMB lensing [44, 45], by allowing non-parametric cluster mass profile measurements from the full likelihood.

#### ACKNOWLEDGMENTS

We thank Duncan Hanson for making his *qcin* code publicly available, on which our filtering code is partly based. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement No. [616170]. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Part of this paper is based observations obtained with Planck, an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

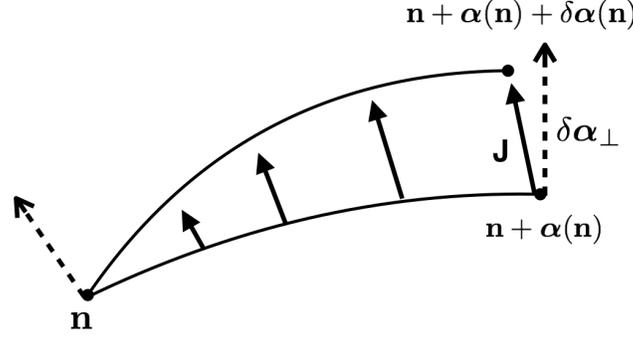


FIG. 5. Schematic aid to the curved-sky lensed CMB likelihood gradient calculation. The variation of a lensed tensor  $\mathcal{T}(\mathbf{n} + \boldsymbol{\alpha}(\mathbf{n}))$  with respect to the deflection must be calculated at the deflected position  $\mathbf{n} + \boldsymbol{\alpha}(\mathbf{n})$ . This point is defined by following the geodesic from  $\mathbf{n}$  in the direction  $\boldsymbol{\alpha}$  for length  $|\boldsymbol{\alpha}|$ . A variation  $\delta\boldsymbol{\alpha}(\mathbf{n})$  in the deflection vector at  $\mathbf{n}$  shifts the geodesic slightly. The exact first order change in position is given by the vector field  $\mathbf{J}$  (solid arrows), proportional to  $\delta\boldsymbol{\alpha}$  infinitesimally close to  $\mathbf{n}$  and evolving along the original geodesic according to the Jacobi equations set by the curvature tensor. For simplicity, we instead evaluate the variation using the parallel-transported  $\delta\boldsymbol{\alpha}$  (dashed arrows) instead of  $\mathbf{J}$ . This neglects the focussing effect of the sphere curvature, slightly overestimating the geodesic deviation. The relative error in the gradient normal component is quadratic in the deflection angle and equal to  $(1 - \sin \alpha/\alpha) \sim 6 \cdot 10^{-8}$  for 2 arcmin deflections. This is completely negligible and of similar order as the small-angle approximation and neglected physical effects such as polarization rotation [? ]: it is safe to neglect sky curvature on the scale of the deflection angles.

### Appendix A: Curved sky gradients

We give here the curved sky likelihood gradients, analogous to the flat sky version given in the main text. We first state the results and describe the implementation, and then provide a derivation. This requires only repeated use of the gradient and curl decomposition of a complex spin  $s$  field, which are readily available in widespread packages. With real part  $\mathcal{R}$  and imaginary part  $\mathcal{I}$ , we may write  ${}_{\pm|s|}f$  as

$${}_{\pm|s|}f(\mathbf{n}) = (\mathcal{R}(\mathbf{n}) \pm i\mathcal{I}(\mathbf{n})). \quad (\text{A1})$$

Then the gradient curl component are defined as ( $|s| > 0$ )

$$\int d^2n {}_{\pm|s|}Y_{lm}^*(\mathbf{n}) {}_{\pm|s|}f(\mathbf{n}) \equiv -(\pm 1)^{|s|} (G_{lm} \pm iC_{lm}), \quad \leftrightarrow {}_{\pm|s|}f(\mathbf{n}) = -(\pm 1)^{|s|} \sum_{lm} (G_{lm} \pm iC_{lm}) {}_{\pm|s|}Y_{lm}(\mathbf{n}). \quad (\text{A2})$$

We follow here for convenience the sign conventions adopted e.g. by the relevant spin harmonic transform routines of the widespread `HEALPIX` package [46]. The definitions of the spin harmonics follow e.g. Ref. [47]. The  $E, B$  decomposition of the spin-2 polarization field is

$${}_{\pm 2}P(\mathbf{n}) \equiv Q(\mathbf{n}) \pm iU(\mathbf{n}) = -\sum_{lm} (E_{lm} \pm iB_{lm}) {}_{\pm 2}Y_{lm}(\mathbf{n}). \quad (\text{A3})$$

We give the results for the joint  $T, Q, U$  analysis, the restriction to temperature only or polarization only is straightforward.

We aim to obtain the likelihood gradients with respect to the displacement modes, in analogy to the flat sky derivation in the main text. We only need to derive the quadratic part of the gradient: the mean field is, as before, its average, and the Gaussian prior is straightforward. We define it by projecting the gradient vector onto the spin basis  $e_{\pm}$ . This basis is associated to the cartesian orthonormal frame  $e_1, e_2$  orthogonal to  $\mathbf{n}$ , given by  $e_{\pm} \equiv e_1 \pm ie_2$ . The spin  $\pm 1$  quadratic gradient is then defined as

$${}_{\pm 1}g^{\text{QD}}(\mathbf{n}) \equiv e_{\pm}^a \frac{\delta}{\delta\alpha^a(\mathbf{n})} \left[ \frac{1}{2} X^{\text{dat}} \cdot \text{Cov}_{\alpha}^{-1} X^{\text{dat}} \right]. \quad (\text{A4})$$

The unnormalized potential ( $\phi$ ) and curl potential ( $\Omega$ ) quadratic estimators (as could have been obtained directly by taking gradients with respect to  $\phi$  and  $\Omega$ ) are then simply given by the harmonic expansion of  ${}_{\pm 1}g$ :

$${}_{\pm 1}g^{\text{QD}}(\mathbf{n}) = -(\pm 1) \sum_{LM} \left( \frac{\phi_{LM}^{\text{QD}} \pm i\Omega_{LM}^{\text{QD}}}{\sqrt{L(L+1)}} \right) {}_{\pm 1}Y_{LM}(\mathbf{n}). \quad (\text{A5})$$

Postponing the derivation, the end result is as follows: with  $\mathbf{n}$  and  $\mathbf{n}'$  the undeflected and deflected points,

$${}_1g^{\text{QD}}(\mathbf{n}) = - \sum_{s=0,\pm 2} {}_{-s}\text{Res}(\mathbf{n})[\partial_s X^{\text{WF}}](\mathbf{n}'). \quad (\text{A6})$$

In this equation, the left leg of the quadratic product is the inverse noise weighted residual

$$\text{Res}(\mathbf{n}) \equiv [B^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}}](\mathbf{n}) = [B^\dagger N^{-1} (X^{\text{dat}} - BD X^{\text{WF}})](\mathbf{n}), \quad (\text{A7})$$

and the right leg is given by deflected gradients of the Wiener-filtered maps. Explicitly,

$$\begin{aligned} [\partial_0 X^{\text{WF}}](\mathbf{n}') &= \sum_{lm} \sqrt{l(l+1)} T_{lm}^{\text{WF}} {}_1Y_{lm}(\mathbf{n}') \\ [\partial_{-2} X^{\text{WF}}](\mathbf{n}') &= - \sum_{lm} \sqrt{(l+2)(l-1)} [E_{lm}^{\text{WF}} - iB_{lm}^{\text{WF}}] {}_{-1}Y_{lm}(\mathbf{n}') \\ [\partial_2 X^{\text{WF}}](\mathbf{n}') &= - \sum_{lm} \sqrt{(l-2)(l+3)} [E_{lm}^{\text{WF}} + iB_{lm}^{\text{WF}}] {}_3Y_{lm}(\mathbf{n}'). \end{aligned} \quad (\text{A8})$$

The only difference between Eq. A6 and traditional position-space curved-sky implementation of the quadratic estimator (as stated above, without the  $N_0$  normalization) are the use of the unlensed spectra instead of the lensed spectra in producing  $X^{\text{WF}}$ , together with the presence of the deflection operations, both in the filter and explicitly in Eq. A6.

We now justify Eq. A6. The model for the observed signal with noise  $n$  is

$$X^{\text{dat}} = BD X^{\text{unl}} + n, \quad (\text{A9})$$

where on the curved sky the operator  $D$  sends the unlensed  $T, E$  and  $B$  CMB modes to the deflected Stokes map with definite spin  $0, \pm 2$  (i.e.  ${}_{\pm 2}P$  and not  $Q, U$ ):

$$D_{lm}(\mathbf{n}) = \begin{pmatrix} {}_0Y_{lm} & 0 & 0 \\ 0 & -{}_2Y_{lm} & -i{}_2Y_{lm} \\ 0 & -{}_2Y_{lm} & i{}_2Y_{lm} \end{pmatrix}(\mathbf{n}'). \quad (\text{A10})$$

Similarly,  $B$  projects the spin maps  $T, {}_{\pm 2}P$  to the observed  $X^{\text{dat}} = T^{\text{dat}}, Q^{\text{dat}}$  and  $U^{\text{dat}}$ . From the definitions given in Eqs. A4 and A7, we have

$$\begin{aligned} {}_1g^{\text{QD}}(\mathbf{n}) &= X^{\text{dat}\dagger} \text{Cov}_\alpha^{-1} B \left[ -e_+^a \frac{\delta D}{\delta \alpha^a(\mathbf{n})} \right] C^{\text{unl}} D^\dagger B^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}} \\ &= \int d^2n' \sum_{s=0,\pm 2} {}_{-s}\text{Res}(\mathbf{n}') \left[ -e_+^a \frac{\delta D}{\delta \alpha^a(\mathbf{n})} X^{\text{WF}} \right]_{s+1}(\mathbf{n}'). \end{aligned} \quad (\text{A11})$$

On the second line we used the spin  $s$  as the index for the different components of the residual and gradient maps. How to make sense of and evaluate the variations of  $D$ ? From Eq. A10, all elements are spin-weighted harmonics at the deflected position, hence we need to understand how this position changes under a variation of the deflection. The geometry is sketched on Fig. A. On the curved sky, the notation  $\mathbf{n} \rightarrow \mathbf{n} + \alpha(\mathbf{n})$  indicates displacement of length  $|\alpha(\mathbf{n})|$  from  $\mathbf{n}$  along the geodesic in direction  $\alpha$ . The polarization axes are parallel transported along the geodesic, leading to some small change in  $Q, U$  from the resulting misalignment with the coordinate vectors at the new point [48, 49]. Varying  $\alpha(\mathbf{n})$  by a small amount  $\delta\alpha(\mathbf{n})$  give rises to a slightly different geodesic. On the flat sky, the end separation vector between the points will be  $\delta\alpha(\mathbf{n})$ , but this is not so on the curved sphere. Since  $\alpha$  is typically a few arcminutes, the difference is very small, and we will neglect it. We now justify this, by deriving the exact but less practical result.

On any Riemannian manifold, the separation between the geodesics is described by the Jacobi vector  $\mathbf{J}$ .  $\mathbf{J}$  is initially zero, has initial velocity  $\delta\alpha(\mathbf{n})$ , and its acceleration is set by the Riemann curvature tensor through the Jacobi equations. The positive curvature of the sphere will reduce the separation vector compared to the flat sky. The covariant, first order change in a tensor  $\mathcal{T}$  on the manifold is

$$(\delta\mathcal{T})(\mathbf{n}') = (J^a \nabla_a \mathcal{T})(\mathbf{n}') \quad (\text{A12})$$

On the sphere, using a parallel orthonormal frame, with one vector  $e_{\parallel}$  initially aligned with  $\alpha(\mathbf{n})$ ,  $\mathbf{J}$  is given by

$$\mathbf{J}(\mathbf{n}') = \delta\alpha^{\parallel}(\mathbf{n}) e_{\parallel}(\mathbf{n}') + \frac{\sin \alpha(\mathbf{n})}{\alpha(\mathbf{n})} \delta\alpha^{\perp}(\mathbf{n}) e_{\perp}(\mathbf{n}'), \quad (\text{A13})$$

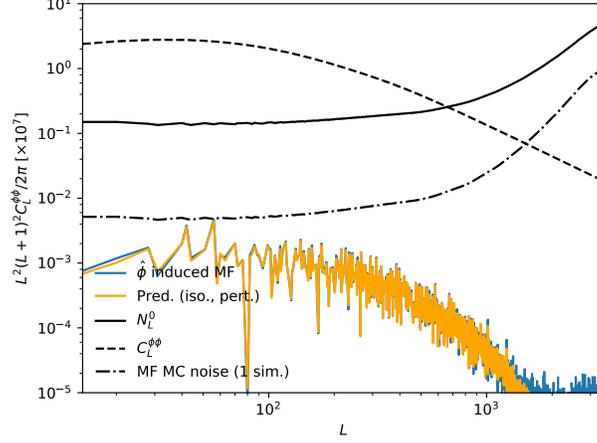


FIG. 6. The expected contribution of the deflection-induced mean field for the first step of iterative reconstruction from polarization with noise level of  $1.5 \cdot \sqrt{2} \mu\text{K-arcmin}$  and no sky cut. The prediction (orange line) is calculated using an input deflection with spectrum equal to that expected for the first iteration estimate of the maximum a posteriori solution (i.e. the spectrum of the Wiener-filtered deflection, Eq. (B8)). Also shown as the (barely visible) blue curve is the measured mean field, with MC noise visible on small scales. The accuracy of the prediction is at least percent-level. Both curves were normalized by  $N^0$ , as the gradients on Fig 2, and can be directly compared. This shows that the  $\phi$ -induced mean field plays very little role in this configuration on the first iteration, where the bulk of the reconstruction improvement is performed. The dotted-dashed black line shows (with the same normalization) the single-simulation Monte-Carlo noise of the mean-field estimator used in this work, Eq. B12. For this configuration, it improves upon the naive  $N_0$  noise (solid black) by more than an order of magnitude.

Hence, in this frame,

$$(\delta\mathcal{T})(\mathbf{n}') = \delta\alpha_{\parallel}(\mathbf{n})\nabla^{\parallel}\mathcal{T}(\mathbf{n}') + \frac{\sin\alpha(\mathbf{n})}{\alpha(\mathbf{n})}\delta\alpha_{\perp}(\mathbf{n})\nabla^{\perp}\mathcal{T}(\mathbf{n}'). \quad (\text{A14})$$

This differs by  $\sin\alpha/\alpha$  in the perpendicular component from the approximation that we use, where instead we use the parallel-transported  $\delta\alpha(\mathbf{n})$  (and not  $\delta\alpha(\mathbf{n}')$ ),

$$\frac{\delta\mathcal{T}(\mathbf{n}')}{\delta\alpha^a(\mathbf{n})} \approx \delta^D(\mathbf{n} + \boldsymbol{\alpha}(\mathbf{n}) - \mathbf{n}')\nabla_a\mathcal{T}(\mathbf{n}'). \quad (\text{A15})$$

This is extremely accurate, since  $1 - \sin\alpha/\alpha \sim 10^{-7}$  for  $\sim 2$  arcminutes deflections. With this approximation, we can make use of the spin lowering and raising form of covariant derivatives for spin weighted functions [47, 48]. The equivalent of Eq. (A14) for spin-weight quantities  ${}_s\mathcal{T}$  is

$$e_+^a \frac{\delta_s\mathcal{T}(\mathbf{n}')}{\delta\alpha^a(\mathbf{n})} = 2 \frac{\delta_s\mathcal{T}(\mathbf{n}')}{\delta_{-1}\alpha(\mathbf{n})} \approx -\delta^D(\mathbf{n} + \boldsymbol{\alpha}(\mathbf{n}) - \mathbf{n}')\eth_s\mathcal{T}(\mathbf{n}'). \quad (\text{A16})$$

Using repeatedly

$$\eth_s Y_{lm} = \sqrt{l(l+1) - s(s+1)}_{s+1} Y_{lm} \quad (\text{A17})$$

on all the  $D$  matrix entries of Eq. (A10) gives the result in Eqs. A6 and A8.

## Appendix B: Mean field

In this appendix we first discuss a perturbative analytic expression for the deflection-induced contribution to the mean field and then the tricks we used to reduce the number of simulations needed to calculate the mean field. We use the  $\star$  operator for multiplication of infinite dimensional matrices across the survey area (continuous sky indices). In the isotropic limit this multiplication reduces to a standard convolution. Sums over discrete indices (CMB pixels, or Stokes fields) are indicated by juxtaposition. In particular, the weight matrices  $W$  of the quadratic estimators act with field and pixel indices on the right, and with field and sky indices on the left. The beam operation  $B$  maps the Stokes sky onto the data Stokes pixelization. Hence it acts on field and sky indices on the right, and field and pixel indices on the left.

### 1. Deflection-induced mean-field contribution

A good handle on the deflection-induced contribution to the mean field can be obtained perturbatively. We start from

$$g_a^{\text{MF}}(\mathbf{x}) = \frac{1}{2} \frac{\delta \ln \det \text{Cov}_\alpha}{\delta \alpha_a(\mathbf{x})}, \quad (\text{B1})$$

where unlike the previous appendix, we do not need to distinguish between upper and lower indices on the flat sky.

Recall that  $\mathbf{x}$  refers to an arbitrary point on the sky, unrelated to the pixelization: we view each element of the covariance matrix as a functional of the deflection field. The linear response of the mean field is the second variation of the log-determinant functional. Hence,

$$\frac{\delta g_a^{\text{MF}}(\mathbf{x})}{\delta \alpha_b(\mathbf{y})} = -\frac{1}{2} \text{Tr} \text{Cov}_\alpha^{-1} \frac{\delta \text{Cov}_\alpha}{\delta \alpha_a(\mathbf{x})} \text{Cov}_\alpha^{-1} \frac{\delta \text{Cov}_\alpha}{\delta \alpha_b(\mathbf{y})} + \frac{1}{2} \text{Tr} \text{Cov}_\alpha^{-1} \frac{\delta^2 \text{Cov}_\alpha}{\delta \alpha_a(\mathbf{x}) \delta \alpha_b(\mathbf{y})}. \quad (\text{B2})$$

The first term (identical to the likelihood Fisher matrix), when evaluated at zero displacement, is simply minus the inverse  $N_L^0$  lensing quadratic estimator response (evaluated with unlensed weights). It is convenient to introduce  $\xi$ , the real-space two-point function of the unlensed CMB fields, and  $\xi_{,a}$ , its derivative with respect to coordinate axis  $a$ . For each element of the covariance matrix, we may then write

$$\left. \frac{\delta \text{Cov}_\alpha}{\delta \alpha_a(\mathbf{x})} \right|_{\alpha=0} = B(\mathbf{x})(\xi_{,a} \star B^\dagger)(\mathbf{x}) - (B \star \xi_{,a})(\mathbf{x})B^\dagger(\mathbf{x}). \quad (\text{B3})$$

In this equation a sum over Stokes-field indices is implicit, and the pixel and further field indices are omitted on both sides in order to prevent visual cluttering. The second variation becomes

$$\left. \frac{\delta^2 \text{Cov}_\alpha}{\delta \alpha_a(\mathbf{x}) \delta \alpha_b(\mathbf{y})} \right|_{\alpha=0} = \delta^D(\mathbf{r}) [B(\mathbf{x})(\xi_{,ab} \star B^\dagger)(\mathbf{y}) + (B \star \xi_{,ab})(\mathbf{y})B^\dagger(\mathbf{x})] - B(\mathbf{x})\xi_{,ab}(\mathbf{r})B^\dagger(\mathbf{y}) - B(\mathbf{y})\xi_{,ab}(\mathbf{r})B^\dagger(\mathbf{x}), \quad (\text{B4})$$

where  $\delta^D$  is the Dirac delta function. Introducing the isotropic operator

$$K(\mathbf{x} - \mathbf{y}) \equiv [B^\dagger \text{Cov}_{\alpha=0}^{-1} B](\mathbf{x} - \mathbf{y}), \quad (\text{B5})$$

all explicit dependence on the pixelization has disappeared, and a **short calculation gives**

$$R_{ab}(\mathbf{r}) \equiv \left. \frac{\delta g_a^{\text{MF}}(\mathbf{x})}{\delta \alpha_b(\mathbf{y})} \right|_{\alpha=0} = +\text{Tr} [(\xi_{,a} \star K)(\mathbf{r})(\xi_{,b} \star K)(\mathbf{r}) + K(\mathbf{r})(\xi_{,a} \star K \star \xi_{,b})(\mathbf{r}) - K(\mathbf{r})\xi_{,ab}(\mathbf{r}) + \delta^D(\mathbf{r})(K \star \xi_{,ab})(\mathbf{r})]. \quad (\text{B6})$$

In harmonic space, the last term is a constant, and ensures the response to the unobservable deflection monopole  $R_{ab}(\ell = 0)$  vanishes as it should. The inverse of minus the first two terms is the usual  $N^0$  lensing bias (here displayed with unlensed weights, and before projection onto gradient and curl components). For low noise experiments,  $\xi K \xi \sim \xi$  on most scales, causing large cancellations between the second and third terms.

By design, we have thus

$$g_a^{\text{MF}}(\mathbf{L}) = \sum_b R_{ab}(\mathbf{L})\alpha_b(\mathbf{L}) + O(\alpha^2), \quad (\text{B7})$$

where all terms can easily be calculated with a series of Fourier transforms.

Fig. 6 shows the expected contribution of the deflection mean field for the polarization reconstruction considered in Sec. IV B, but with no sky cuts so that the deflection is the only mean-field source. The blue line shows the mean-field spectrum estimate obtained by averaging 500 simulations. As a test case we used an input deflection field  $\nabla\phi_0$ , with spectrum

$$C_L^{\phi_0\phi_0} = \frac{(C_L^{\phi\phi})^2}{C_L^{\phi\phi} + N_L^0}, \quad (\text{B8})$$

equivalent, from Eq. 3.10, to the spectrum of the reconstruction expected at the first iteration step. The orange curve shows the predicted spectrum from Eq. B6, which is in very good agreement given the MC noise of the mean-field estimation, visible at high multipoles.

## 2. Mean field tricks

To improve the calculation of the mean field, we introduce two tricks that together can reduce the MC noise power in the mean-field estimate by orders of magnitude. The first, and often the most powerful, is to modify the weights used in the quadratic estimator that is averaged to calculate the mean field from the simulations. The usefulness of this trick, however, depends on the adequacy of  $\text{Cov}_\alpha$  to represent the true covariance of the data. Specifically, it makes use of

$$\widetilde{\text{Cov}}_\alpha^{-1} \equiv \text{Cov}_\alpha^{-1} \langle X^{\text{dat}} X^{\text{dat},\dagger} \rangle \text{Cov}_\alpha^{-1} = \text{Cov}_\alpha^{-1}. \quad (\text{B9})$$

In practice,  $\text{Cov}_\alpha$  input to the likelihood is always going to be only an approximation to the true unknown data covariance. Accurate timestream simulations  $X^{\text{dat}}$  can sometimes be used to quantify the mean field (and other biases) for the standard quadratic estimator. But moderate numbers of simulations do not directly provide the means to accurately calculate  $\text{Cov}_\alpha$  or apply its inverse. The simulations themselves may of course also only capture only parts of the complexity entering the relevant systematics and the data processing. These are difficulties affecting the quadratic estimator as well, not specifically the iterative scheme proposed in this paper, that requires some amount of testing in a realistic situation.

From Eq. (2.14), the mean field satisfies

$$g_a^{\text{MF}}(\mathbf{n}) = \langle g_a^{\text{QD}}(\mathbf{n}) \rangle = \langle [(V_\alpha X^{\text{dat}})]^i(\mathbf{n}) [W_\alpha^a X^{\text{dat}}]_i(\mathbf{n}) \rangle. \quad (\text{B10})$$

Performing the average using the normal weights of Eq. (2.10) and assuming we can use Eq. B9 results in

$$g_a^{\text{MF}}(\mathbf{n}) = \text{Tr } H_\alpha^a(\mathbf{n}, \mathbf{n}) \text{ with } H_\alpha^a \equiv B^\dagger \text{Cov}_\alpha^{-1} B \star D \nabla_a C^{\text{unl}} D^\dagger, \quad (\text{B11})$$

where the trace is over field indices. Only the diagonal of  $H_\alpha^a$  is relevant for the mean field, both in respect to sky and field indices, and since  $H^a$  is a vector it vanishes for isotropic fields where there is no mean field.

It is possible to construct mean-field estimators with modified weights, or acting on different maps: as long as the expectation of the quadratic estimator remains the same they will produce unbiased estimates of the mean field. We choose weights to apply to independent unit variance Gaussian variables  $s$ , with  $\langle ss^\dagger \rangle = \text{diag}(1)$  on the unmasked pixels, and use the following pair of weights

$$W_1 = B^\dagger \quad W_2^a = D \nabla_a C_\ell^{\text{unl}} D^\dagger \star B^\dagger \text{Cov}_\alpha^{-1}. \quad (\text{B12})$$

The matrix inverse can be performed in the same way as the usual filtering. We now proceed to justify this choice and explain why it has lower variance (though is not exactly minimum variance).

Ideally, we would like to choose a pair of weights  $W_1, W_2^a$  to minimize the Gaussian MC noise on the mean field estimate while keeping the constraint  $[W_1 \langle ss^\dagger \rangle W_2^{a\dagger}](\mathbf{x}, \mathbf{x}) = \text{Tr } H_\alpha^a(\mathbf{x}, \mathbf{x})$ . The real-space MC noise covariance is

$$N_0(\mathbf{r}) = (W_1 W_1^\dagger)(\mathbf{r})(W_2^a W_2^{a\dagger})(\mathbf{r}) + (W_1 W_2^{a\dagger})(\mathbf{r})(W_2^a W_1^\dagger)(\mathbf{r}). \quad (\text{B13})$$

For simplicity, we use the constraint  $W_1 \langle ss^\dagger \rangle W_2^{a\dagger} = W_1 W_2^{a\dagger} = H^a$ , which is much more stringent than only matching the diagonal but makes things more tractable. We then consider minimizing the variance in the isotropic limit, in which case  $W_1 W_2^{a\dagger}$  becomes a real-space convolution and hence the constraint means that in harmonic space  $W_2^a(\ell) = H_\alpha^a(\ell)/W_1(\ell)$ , where  $W_1(\ell)$  is a free function of multipole. A natural measure to minimize is the integrated variance from the above equation

$$\frac{1}{V} \sum_{\mathbf{L}} N_{\mathbf{L}}^0 = N_0(\mathbf{r} = 0). \quad (\text{B14})$$

Using the constraint equation, minimizing  $N_0(\mathbf{r} = 0)$  gives

$$W_1(\ell) \propto W_2^a(\ell) H_\alpha^a(\mathbf{r} = 0), \quad W_2^a(\ell) \propto W_1(\ell) H_\alpha^a(\mathbf{r} = 0), \quad (\text{B15})$$

so the two weight functions are proportional. From the constraint this implies that  $W_1(\ell) \propto \sqrt{H_\alpha^a(\ell) H_\alpha^a(\mathbf{r} = 0)}$ .

The key message of this calculation is that to get low MC noise the scale dependence of the weight functions should be similar, though once isotropy is broken by the deflection in practice it is never possible to obtain exact square roots even in the absence of non-ideal effects. Other aspects matter as well, such as how well the unavoidable matrix inversion behaves with the chosen weights, and the contribution of the non-ideal effects to the final MC-noise floor. The form given in Eq. B12 is an empirical compromise between these considerations, that we found works well in practice. It equilibrates the weights in a very crude way, simply by having the same powers of the signal on each leg (though not the exact scale dependence). Despite being crude, it reduces the MC noise floor by more than one order of magnitude for the polarization reconstruction performed in the main text.

This is shown on Fig. 6. The dash-dotted black line the MC noise of the estimator given by Eq. B12, and should be compared to the naive estimate MC noise given by the  $N_0$  curve (black, solid). Both curves were calculated with the lensed spectra weights, which is a slightly conservative estimate choice as the iterative search converges towards the optimal solution. We found empirically that further modifying the weights to make their scale-dependence closer (for example multiplying  $W_1$  by  $\sqrt{\ell}$  and  $W_2$  by  $1/\sqrt{\ell}$  to equilibrate power in the no-beam, no-lensing, no-noise limit) did not give substantial further improvements.

The second trick is to subtract from each MC estimate the same estimate but based on an isotropic approximation to the posterior that is as close as possible to the true one, as described in Sec. III C.

- [1] A. Lewis and A. Challinor, “Weak gravitational lensing of the CMB,” *Phys. Rept.* **429**, 1–65 (2006), astro-ph/0601594.
- [2] Sudeep Das *et al.*, “Detection of the Power Spectrum of Cosmic Microwave Background Lensing by the Atacama Cosmology Telescope,” *Phys. Rev. Lett.* **107**, 021301 (2011), arXiv:1103.2124 [astro-ph.CO].
- [3] A. van Engelen *et al.*, “A measurement of gravitational lensing of the microwave background using South Pole Telescope data,” *Astrophys. J.* **756**, 142 (2012), arXiv:1202.0546 [astro-ph.CO].
- [4] P. A. R. Ade *et al.* (POLARBEAR), “Measurement of the Cosmic Microwave Background Polarization Lensing Power Spectrum with the POLARBEAR experiment,” *Phys. Rev. Lett.* **113**, 021301 (2014), arXiv:1312.6646 [astro-ph.CO].
- [5] K. T. Story *et al.* (SPT), “A Measurement of the Cosmic Microwave Background Gravitational Lensing Potential from 100 Square Degrees of SPTpol Data,” *Astrophys. J.* **810**, 50 (2015), arXiv:1412.4760 [astro-ph.CO].
- [6] P. A. R. Ade *et al.* (BICEP2, Keck Array), “BICEP2 / Keck Array VIII: Measurement of gravitational lensing from large-scale B-mode polarization,” *Astrophys. J.* **833**, 228 (2016), arXiv:1606.01968 [astro-ph.CO].
- [7] P. A. R. Ade *et al.* (Planck), “Planck 2013 results. XVII. Gravitational lensing by large-scale structure,” *Astron. Astrophys.* **571**, A17 (2014), arXiv:1303.5077 [astro-ph.CO].
- [8] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and *et al.*, “Planck 2015 results. XV. Gravitational lensing,” *Astron. Astrophys.* **594**, A15 (2016), arXiv:1502.01591.
- [9] W. Hu and T. Okamoto, “Mass Reconstruction with Cosmic Microwave Background Polarization,” *Astrophys. J.* **574**, 566–574 (2002), astro-ph/0111606.
- [10] T. Okamoto and W. Hu, “Cosmic microwave background lensing reconstruction on the full sky,” *Phys. Rev. D* **67**, 083002 (2003), astro-ph/0301031.
- [11] Kevork N. Abazajian *et al.* (CMB-S4), “CMB-S4 Science Book, First Edition,” (2016), arXiv:1610.02743 [astro-ph.CO].
- [12] Uros Seljak and Christopher M. Hirata, “Gravitational lensing as a contaminant of the gravity wave signal in CMB,” *Phys. Rev.* **D69**, 043005 (2004), astro-ph/0310163.
- [13] Christopher M. Hirata and Uros Seljak, “Reconstruction of lensing from the cosmic microwave background polarization,” *Phys. Rev.* **D68**, 083002 (2003), arXiv:astro-ph/0306354 [astro-ph].
- [14] Wayne Hu and Asantha Cooray, “Gravitational time delay effects on cosmic microwave background anisotropies,” *Phys. Rev.* **D63**, 023504 (2001), arXiv:astro-ph/0008001 [astro-ph].
- [15] Antony Lewis, Alex Hall, and Anthony Challinor, “Emission-angle and polarization-rotation effects in the lensed CMB,” (2017), arXiv:1706.02673 [astro-ph.CO].
- [16] Geraint Pratten and Antony Lewis, “Impact of post-Born lensing on the CMB,” *JCAP* **1608**, 047 (2016), arXiv:1605.05662 [astro-ph.CO].
- [17] Silvia Mollerach, Diego Harari, and Sabino Matarrese, “CMB polarization from secondary vector and tensor modes,” *Phys. Rev.* **D69**, 063002 (2004), arXiv:astro-ph/0310711 [astro-ph].
- [18] Christian Fidler, Guido W. Pettinari, Martin Beneke, Robert Crittenden, Kazuya Koyama, and David Wands, “The intrinsic B-mode polarisation of the Cosmic Microwave Background,” *JCAP* **1407**, 011 (2014), arXiv:1401.3296 [astro-ph.CO].
- [19] Daniel Baumann, Paul J. Steinhardt, Keitaro Takahashi, and Kiyotomo Ichiki, “Gravitational Wave Spectrum Induced by Primordial Scalar Perturbations,” *Phys. Rev.* **D76**, 084019 (2007), arXiv:hep-th/0703290 [hep-th].
- [20] Devdeep Sarkar, Paolo Serra, Asantha Cooray, Kiyotomo Ichiki, and Daniel Baumann, “Cosmic shear from scalar-induced gravitational waves,” *Phys. Rev.* **D77**, 103515 (2008), arXiv:0803.1490 [astro-ph].
- [21] Julian Adamek, Ruth Durrer, and Vittorio Tansella, “Lensing signals from Spin-2 perturbations,” *JCAP* **1601**, 024 (2016), arXiv:1510.01566 [astro-ph.CO].
- [22] Kendrick M. Smith, Duncan Hanson, Marilena LoVerde, Christopher M. Hirata, and Oliver Zahn, “Delensing CMB Polarization with External Datasets,” *JCAP* **1206**, 014 (2012), arXiv:1010.0048 [astro-ph.CO].
- [23] Christopher M. Hirata and Uros Seljak, “Analyzing weak lensing of the cosmic microwave background using the likelihood function,” *Phys. Rev.* **D67**, 043001 (2003), arXiv:astro-ph/0209489 [astro-ph].
- [24] Duncan Hanson, Anthony Challinor, and Antony Lewis, “Weak lensing of the CMB,” *General Relativity and Gravitation* **42**, 2197–2218 (2010), arXiv:0911.0612 [astro-ph.CO].
- [25] C. Dvorkin and K. M. Smith, “Reconstructing patchy reionization from the cosmic microwave background,” *Phys. Rev. D* **79**, 043003 (2009), arXiv:0812.1566.
- [26] Toshiya Namikawa, “CMB Lensing Bispectrum from Nonlinear Growth of the Large Scale Structure,” *Phys. Rev.* **D93**, 121301 (2016), arXiv:1604.08578 [astro-ph.CO].
- [27] Kendrick M. Smith, Oliver Zahn, and Olivier Dore, “Detection of Gravitational Lensing in the Cosmic Microwave Background,” *Phys. Rev.* **D76**, 043510 (2007), arXiv:0705.3980 [astro-ph].
- [28] Daniel Ruijters, Bart M. ter Haar Romeny, and Paul Suetens, “Efficient GPU-Based Texture Interpolation using Uniform B-Splines,” *Journal of Graphics Tools* **13**, 61–69 (2008), <http://dx.doi.org/10.1080/2151237X.2008.10129269>.
- [29] Daniel Ruijters and Philippe Thévenaz, “GPU Prefilter for Accurate Cubic B-spline Interpolation,” *Comput. J.* **55**, 15–20 (2012).
- [30] Ethan Anderes, Benjamin Wandelt, and Guilhem Lavaux, “Bayesian inference of CMB gravitational lensing,” *Astrophys. J.* **808**, 152 (2015), arXiv:1412.4079 [astro-ph.CO].
- [31] Julien Carron, Antony Lewis, and Anthony Challinor, “Internal delensing of Planck CMB temperature and polarization,” *JCAP* **1705**, 035 (2017), arXiv:1701.01712 [astro-ph.CO].
- [32] Jorge Nocedal, “Updating quasi-newton matrices with limited storage,” *Mathematics of computation* **35**, 773–782 (1980).
- [33] Takemi Okamoto and Wayne Hu, “CMB lensing reconstruction on the full sky,” *Phys. Rev.* **D67**, 083002 (2003), arXiv:astro-ph/0301031 [astro-ph].
- [34] Duncan Hanson, Anthony Challinor, George Efstathiou, and Pawel Bielewicz, “CMB temperature lensing power reconstruction,” *Phys. Rev.* **D83**, 043005 (2011), arXiv:1008.4403 [astro-ph.CO].
- [35] Julien Peloton, Marcel Schmittfull, Antony Lewis, Julien Carron, and Oliver Zahn, “Full covariance of CMB and lensing reconstruction power spectra,” *Phys. Rev.* **D95**, 043508 (2017), arXiv:1611.01446 [astro-ph.CO].
- [36] A. Manzotti *et al.* (Herschel, SPT), “CMB Polarization B-mode Delensing with SPTpol and Herschel,” (2017), arXiv:1701.04396 [astro-ph.CO].
- [37] P. A. R. Ade *et al.* (Planck), “Planck 2015 results. XV.

- Gravitational lensing,” *Astron. Astrophys.* **594**, A15 (2016), arXiv:1502.01591 [astro-ph.CO].
- [38] N. Aghanim *et al.* (Planck), “Planck intermediate results. XLVIII. Disentangling Galactic dust emission and cosmic infrared background anisotropies,” *Astron. Astrophys.* **596**, A109 (2016), arXiv:1605.09387 [astro-ph.CO].
- [39] Patricia Larsen, Anthony Challinor, Blake D. Sherwin, and Daisy Mak, “Demonstration of cosmic microwave background delensing using the cosmic infrared background,” *Phys. Rev. Lett.* **117**, 151102 (2016), arXiv:1607.05733 [astro-ph.CO].
- [40] P. A. R. Ade *et al.* (Planck), “Planck 2013 results. XVIII. The gravitational lensing-infrared background correlation,” *Astron. Astrophys.* **571**, A18 (2014), arXiv:1303.5078 [astro-ph.CO].
- [41] Blake D. Sherwin and Marcel Schmittfull, “Delensing the CMB with the Cosmic Infrared Background,” *Phys. Rev.* **D92**, 043005 (2015), arXiv:1502.05356 [astro-ph.CO].
- [42] Ethan Anderes, Lloyd Knox, and Alexander van Engelen, “Mapping gravitational lensing of the CMB using local likelihoods,” *Phys. Rev.* **D83**, 043523 (2011), arXiv:1012.1833 [astro-ph.CO].
- [43] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, B. Catanzaro, Paul Ivanov, and Ahmed Fasih, “PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation,” *Parallel Computing* **38**, 157–174 (2012).
- [44] E. J. Baxter *et al.*, “A Measurement of Gravitational Lensing of the Cosmic Microwave Background by Galaxy Clusters Using Data from the South Pole Telescope,” *Astrophys. J.* **806**, 247 (2015), arXiv:1412.7521 [astro-ph.CO].
- [45] Mathew Madhavacheril *et al.* (ACT), “Evidence of Lensing of the Cosmic Microwave Background by Dark Matter Halos,” *Phys. Rev. Lett.* **114**, 151302 (2015), [Addendum: *Phys. Rev. Lett.* 114,no.18,189901(2015)], arXiv:1411.7999 [astro-ph.CO].
- [46] K. M. Gorski, Eric Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, “HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere,” *Astrophys. J.* **622**, 759–771 (2005), arXiv:astro-ph/0409513 [astro-ph].
- [47] Antony Lewis, Anthony Challinor, and Neil Turok, “Analysis of CMB polarization on an incomplete sky,” *Phys. Rev.* **D65**, 023505 (2002), arXiv:astro-ph/0106536 [astro-ph].
- [48] Anthony Challinor and Gayoung Chon, “Geometry of weak lensing of CMB polarization,” *Phys. Rev.* **D66**, 127301 (2002), arXiv:astro-ph/0301064 [astro-ph].
- [49] Antony Lewis, “Lensed CMB simulation and parameter estimation,” *Phys. Rev.* **D71**, 083008 (2005), arXiv:astro-ph/0502469 [astro-ph].