Compiled using MNRAS ${\rm IAT}_{\rm E}\!{\rm X}$ style file v3.0

Full-sky beam convolution for cosmic microwave background applications

Adriaan J. Duivenvoorden,^{1*} Jon E. Gudmundsson,¹ Alexandra S. Rahlin^{2,3}

¹ The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, SE-106 91 Stockholm, Sweden ² Fermi National Accelerator Laboratory, Batavia, IL, USA

³Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL, USA

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We introduce a publicly available full-sky beam convolution code library intended to inform the design of future cosmic microwave background (CMB) instruments and help current experiments probe potential systematic effects. The code can be used to assess the impact of optical systematics on all stages of data reduction for a realistic experiment, including analyses beyond power spectrum estimation, by generating signal timelines that may serve as input to full analysis pipelines. The design and mathematical framework of the Python code is discussed along with a few simple benchmarking results. We present a simple two-lens refracting telescope design and use it together with the code to simulate a year-long dataset for 400 detectors scanning the sky on a satellite instrument. The simulation results identify a number of sub-leading optical non-idealities and demonstrate significant B-mode residuals caused by extended sidelobes that are sensitive to polarized radiation from the Galaxy. For the proposed design and satellite scanning strategy, we show that a full physical optics beam model generates B-mode systematics that differ significantly from the simpler elliptical Gaussian model. The code is available at https://github.com/adrijd/beamconv.

Key words: CMB – Polarization – Convolution – Optics – Satellite – Telescope

1 INTRODUCTION

Many current cosmological observing programs are focused on a conjectured imprint of primordial gravitational waves in the degree-scale polarization anisotropies of the cosmic microwave background (CMB). This, together with efforts to quantify CMB polarization on both small (arcmin) as well as the largest possible angular scales, are driving a significant increase in the sensitivity of CMB instruments (Abazajian et al. 2015; Matsumura et al. 2016; Abitbol et al. 2017; Bryan et al. 2018; Buzzelli et al. 2018; The Simons Observatory Collaboration 2018). This growth comes mainly from a surge in the number of detectors deployed per focal plane, which in turn is facilitated by telescope designs optimised for large fields of view (FOV) (Niemack 2016). Increased sensitivity requirements also motivate extensive in-situ instrument characterisation, a time consuming process for wide FOV telescopes and satellites with limited observing time. This prompts the development of advanced modelling and analysis techniques that maximise the observing duty cycle.

Traditionally, optical designs for CMB telescopes are optimized for high Strehl numbers, and other geometrical aspects such as f-number, telecentricity, and mapping speed (Page et al. 2003; Ruhl et al. 2004; Fowler et al. 2007; Aikin et al. 2010; Niemack 2016; Young et al. 2018). Although these geometrical properties are definite predictors of some optical non-idealities, we argue that the CMB telescope design process should incorporate physical optics in conjunction with fast convolution techniques, and that the need for integrating this aspect into the design is growing with the cost and sensitivity requirements of future experiments. Unfortunately, modeling and computational challenges can significantly restrict telescope design iterations that incorporate full-sky beam convolution and realistic scan strategies to assess the impact of optical non-idealities on maps, power spectra, and cosmological analyses.

Convolution algorithms for realistic beams have been discussed extensively in the CMB literature. Wandelt & Górski (2001) introduced an efficient method that takes advantage of fast inverse spherical harmonic transforms and sparsity of the harmonic representation of the beam; the generalization to the polarized case was presented in Challinor et al. (2000). An implementation of this method, described in Prézeau & Reinecke (2010), has been used for the *Planck* analysis and is closely related to the implementation discussed in this work. Parallel to these methods, algo-

^{*} E-mail: adri.duivenvoorden@fysik.su.se

rithms that work in the pixel domain have seen use in the WMAP and Planck analyses and are discussed in Wehus et al. (2009) and Mitra et al. (2011). Approaches that focus on providing an efficient convolution operator for experiments with several thousand or more detectors have been formulated in Elsner & Wandelt (2014) and BICEP2 Collaboration (2015). Finally, Wallis et al. (2014) and Hivon et al. (2017) present extensions to the pseudo- C_{ℓ} power spectrum estimation framework that take into account the effects of beam non-idealities.

In this work, we aim to address the issue of accurately simulating optical systematic effects for current and upcoming CMB polarization experiments. We describe an opensource full-sky beam convolution code library that may be used to efficiently simulate time-ordered data and probe various optical systematics. We argue that although simulating time-ordered data is computationally intensive compared to the pseudo- C_{ℓ} extensions mentioned above, it provides a useful complementary method that is uniquely capable of quantifying optical systematic effects for analyses that do not rely solely on the angular power spectrum. Important examples are foreground characterisation, lensing and non-Gaussianity estimation. Additionally, by working in the time domain, optical effects can be simulated without having to make assumptions about other systematic effects such as non-trivial noise properties and high-level analysis choices like time-domain filtering and map-making algorithms. We address some of the associated numerical challenges faced by experiments that deploy a large number of detectors coupled to large-aperture optics. The code library is publicly available and accessible on GitHub.¹

We use the code library in conjunction with physical optics simulations to demonstrate its capabilities and to quantify some of the systematics faced by a fiducial satellite experiment designed to study the polarization of the CMB on degree angular scales (see Figure 1). We discuss the relative contributions of some of the optical systematics that are intrinsic to the proposed optical design. Although we try to identify some key questions and challenges associated with the design, this paper only covers a very small set of non-idealities formed by the interplay between detectors and refractive optics. The large number and varied properties of CMB telescope optical elements, including lenses, reflectors, baffles, filters, birefringent crystals, etc., can lead to serious modeling challenges. In fact, accurate modeling of complete optical systems is still markedly limited by computation and memory requirements as well as uncertainties in material properties.

The paper is organized as follows. In Sec. 2 and 3 we introduce the mathematical formalism and present the code implementation. In Sec. 4 the fiducial instrument is described and motivated. Results are shown in Sec. 5 and, finally, discussions and suggestions for future work are presented in in Sec. 6.



Figure 1. A very rough CAD model showing the fiducial satellite design and refractive optics. The two-lens refracting telescope and sun shields are shown with a section view to emphasise some relevant components, including the two lenses (brown) and the location of the hexagonal detector tiles at the focal plane.

2 FORMALISM

2.1 Preliminaries

We describe the polarization state of quasi-monochromatic radiation at wavelength ω originating from the far-field of a telescope by a complex vector field ϵ with a redundant overall complex phase. We represent ϵ as a vector field on the celestial sphere, i.e. we have $\epsilon(\omega) = \epsilon^i(\omega)\hat{e}_{(i)}$ with $i \in \{1, 2\}$ and $\hat{e}_{(i)}$ the basis vectors of the tangent space T_x with $x \in S^2$. Most modern CMB polarization experiments use incoherent detectors, so we will restrict ourselves to this case. The incoherency of the detectors refers to their insensitivity to the phase, frequency, and polarization state of incident radiation, meaning that the vector field ϵ is not an observable. Intrinsically, these detectors are only sensitive to the total intensity of the field: $I = \langle \epsilon_i \overline{\epsilon}^i \rangle^2$. For the purpose of CMB polarimetry, some polarization sensitive interface like an antenna is coupled to the detectors. This allows them to probe the cosmologically relevant quantity: the covariance of the field: $W_{ij} = \langle \epsilon_i \overline{\epsilon}_j \rangle$. The incident radiation is thus naturally described by this tensor-valued field on the sphere. By introducing an orthonormal coordinate frame, the field can be decomposed into the four (real-valued) Stokes parameters. For example: with the standard (θ, ϕ) spherical coordinate system we get:

$$W_{ij}(\theta,\phi,\omega) = \frac{1}{2} \begin{pmatrix} I+Q & (U-iV)\sin\theta\\ (U+iV)\sin\theta & (I-Q)\sin^2\theta \end{pmatrix} (\theta,\phi,\omega) \,. \tag{1}$$

The I and V Stokes parameters represent the total intensity and circular polarized radiation component. Linear polarization is described by Q and U.³ It is important to realise that, because they correspond to the components of

 $^{^2}$ We implicitly sum over repeated indices; explicit summation will sometimes be used for clarity.

 $^{^3}$ Throughout this work whenever we work with the spherical coordinate system or Euler angles, we will conform to the 'cosmo' polarization angle convention used by the HEALPix li-

a second order tensor field, the Stokes parameters are basis dependent with transformation properties that reflect the underlying tensor transformation law. Naturally, the I parameter, being the trace of W, remains invariant. In contrast, the U and V fields behave as parity-odd under reflections; the Q and U parameters, corresponding to the symmetric, traceless part of the tensor field, transform among themselves under rotations of the coordinate system. Because Wmust be positive semi-definite, the Stokes parameters obey the following inequality:

$$I \ge \sqrt{Q^2 + U^2 + V^2},\tag{2}$$

which is saturated for purely polarized light, while unpolarized light has Q = U = V = 0.

The four Stokes parameters may be grouped into a fourvector (Stokes vector): $s^{\mu}=(I,\,Q,\,U,\,V).$ We define ${\mathcal S}$ as the set of valid Stokes parameters: $S = \left\{ S \mid I \ge \sqrt{Q^2 + U^2 + V^2} \right\}$ and identify the linear transformations: $M : S \to S$, that transform valid Stokes vectors among themselves: the socalled Mueller matrices. We will later describe the instrumental effects through the use of these transformations. An important subset of the Mueller matrices is the set of Mueller-Jones matrices. These are transformations that could equally well be described by a 2×2 complex (Jones) operator working directly on the complex polarization state ϵ . Such transformations are said to be non-depolarizing, i.e. they are unable to convert a purely polarized signal to a partly polarized or unpolarized signal. All other Mueller matrices describe fully or partly depolarizing transformations.⁴

In a manner similar to O'Dea et al. (2007), we do not directly work with the Q and U Stokes parameters. We find it more convenient to work with the complex field $P \equiv Q + iU$ and its complex conjugate, as these quantities transform under the spin-weighted representations of the rotation group (see e.g. (Zaldarriaga & Seljak 1997)). We will denote these alternative Stokes vectors $p^{\mu} = (I, P, \overline{P}, V)$. The corresponding Mueller transformations in the space \mathcal{P} of valid p^{μ} vectors are then denoted by $\mathcal{M}: \mathcal{P} \to \mathcal{P}$.

2.2Data model with beam convolution

For each of the detectors on a focal plane, the one-dimensional array of time-ordered data (TOD): $d = \{d_0, d_1, \dots, d_n\}$ is modelled as some linear transformation A of the sky s and an additive noise component n:

$$\boldsymbol{d} = \boldsymbol{A}\boldsymbol{s} + \boldsymbol{n} \,. \tag{3}$$

brary (https://healpix.jpl.nasa.gov) and use the ZYZ convention (with fixed axes) for Euler angles (ψ, θ, ϕ) . Unless noted otherwise, we use spherical coordinates with basis vectors $\hat{e}_{(\theta)} = \partial_{\theta},$ $\hat{e}_{(\phi)} = \partial_{\phi}/{\sin\theta} \text{ and metric } g_{ij} = \text{diag}(1,\sin^2\theta).$

A sufficient and necessary condition for a matrix \boldsymbol{M} to represent a Mueller transformation is the positive semidefiniteness of the associated coherency matrix: $\boldsymbol{H} = \frac{1}{4} M^{\mu}_{\nu} \left(\boldsymbol{\sigma}_{(\mu)} \otimes \boldsymbol{\sigma}^{(\nu)} \right)$, with $\boldsymbol{\sigma}_{\mu} =$ $\{1, \sigma_3, \sigma_1, \sigma_2\}$ in terms of the Pauli matrices. If **H** has just one nonzero eigenvalue, M is a Mueller-Jones transformation. (Cloude 1986; Anderson & Barakat 1994). Note that we simply use $G_{\mu\nu} =$ diag(1, 1, 1, 1) as metric.

In the following, we will mostly ignore the noise component and focus on the transformation A by working towards a data model that includes optical effects (Eq. 10).

We start by describing the detector positioned at the instrument-side of the optical system with a Mueller transformation \mathcal{M} . We approximate the detector as infinitesimally small and place it at the centre of the spherical coordinate system. Without the coupling to a polarization sensitive interface, the incoherent detector is described by a perfectly depolarizing transformation, i.e. $\mathcal{M}^{\nu}{}_{\mu} \propto \delta^{\nu}_{0} \delta^{0}_{\mu}$. However, when the interface is included, all elements of the top row of \mathcal{M} are allowed to be nonzero: $\mathcal{M}^{\nu}_{\ \mu} \propto \delta^{\nu}_{0}$ (Jones et al. 2007). The data model thus becomes:

$$d_t \propto \int_{\Delta\omega} d\omega \int_{S^2} dx \left(\mathcal{M}_t \right)^0{}_{\mu}(x,\omega) p^{\mu}(x,\omega), \qquad (4)$$

where the integrals are over the frequency passband $\Delta \omega$ and the sky S^2 . The sky signal is denoted by the (complex) Stokes vector p. The subscript in d_t reminds us of the discrete nature of the data.⁵ We denote the nonzero elements of \mathcal{M} : $(\mathcal{M}_t)^0_{\ \mu} = (\widetilde{I}, \ \overline{\widetilde{P}}, \ \widetilde{P}, \ \widetilde{V})$ as they transform like a complex Stokes vector \overline{p}_{μ} . Together they should be interpreted as the beam of the detector. These elements obey the Mueller transformation requirement:

$$\widetilde{I} \ge \sqrt{|\widetilde{P}|^2 + \widetilde{V}^2}$$
, or equivalently: $\widetilde{I} \ge \sqrt{\widetilde{Q}^2 + \widetilde{U}^2 + \widetilde{V}^2}$. (5)
By expanding Eq. 4, we obtain:

$$d_t \propto \int_{S^2} dx \left[\widetilde{I}_t(x) I(x) + \Re\left(\overline{\widetilde{P}}_t(x) P(x)\right) + \widetilde{V}_t(x) V(x) \right], \tag{6}$$

where we have suppressed the integral over and dependence on the wavenumber ω . We now express the elements of the instrument's Mueller matrix in Eq. 6 in terms of (spinweighted) spherical harmonic (SWSH) coefficients (see Appendix A1). We do the same for the Stokes parameters of the sky and make use of the orthonormality of the SWSHs to arrive at:

$$d_t \propto \sum_{\ell,m} \left[\overline{b_{\ell m,t}^{\tilde{I}}} a_{\ell m}^I + \Re \left(\overline{2b_{\ell m,t}^{\tilde{P}}} \, 2a_{\ell m}^P \right) + \overline{b_{\ell m,t}^{\tilde{V}}} a_{\ell m}^V \right] \,. \tag{7}$$

We now impose that the only difference between the optical response at samples t and t' is the direction and orientation of the telescope with respect to the sky. Under a generic rotation $g^{-1} \in SO(3)$ of the coordinate system, the spinweighted harmonic coefficients transform among themselves as:

$$sf_{\ell m} \mapsto \sum_{m'=-\ell}^{\ell} sf_{\ell m'} D^{\ell}_{mm'}(g), \qquad (8)$$

where $D^{\ell}(g)$ are the $(2\ell + 1) \times (2\ell + 1)$ Wigner D-matrices. We may thus compute the harmonic coefficients of the beam in some fiducial reference frame — the instrument frame and transform to a coordinate system fixed on the sky using the above relation. Note that g is continuously changing due

⁵ In reality, the convolution of the continuous sky signal with the finite detector time-response can lead to significant systematic effects if not taken into account (see e.g. Planck Collaboration (2016a)).

to the scanning motion of the telescope. Doing so, we obtain the final expression for the beam-convolved TOD:

$$d_{t} \propto \sum_{\ell,m,s} \left[b_{\ell s}^{\widetilde{I}} a_{\ell m}^{I} + \frac{1}{2} \left(-2 b_{\ell s}^{\widetilde{P}} 2 a_{\ell m}^{P} + 2 b_{\ell s}^{\widetilde{P}} -2 a_{\ell m}^{P} \right) + b_{\ell s}^{\widetilde{V}} a_{\ell m}^{V} \right]$$

$$\times (-1)^{m} D_{-ms}^{\ell}(g_{t}), \qquad (9)$$

$$= \sum_{\ell,m,s} \left[b_{\ell s}^{\widetilde{I}} a_{\ell m}^{I} + \frac{1}{2} \left(-2 b_{\ell s}^{\widetilde{P}} 2 a_{\ell m}^{P} + 2 b_{\ell s}^{\widetilde{P}} -2 a_{\ell m}^{P} \right) + b_{\ell s}^{\widetilde{V}} a_{\ell m}^{V} \right]$$

$$\times q_{\ell} e^{-is\psi_{t}} {}_{s}Y_{\ell m}(\theta_{t},\phi_{t}), \qquad (10)$$

where we have defined:

$$q_{\ell} \equiv \sqrt{\frac{4\pi}{2\ell+1}},\tag{11}$$

and where $(\psi_t, \theta_t, \phi_t)$ are Euler angles parametrizing the rotation g_t . To arrive at the second line, we have used the relation between the Wigner *D*-matrices and the spin-weighted spherical harmonics in terms of Euler angles (see Eq. A9). When formulated like Eq. 9, a useful interpretation of the TOD emerges: the expression is simply an inverse Wigner transform with harmonic coefficients given by the terms in the square brackets, implying that the TOD are just discrete samples from a scalar field $d(g_t)$ on the manifold given by the rotation group SO(3) (Wandelt & Górski 2001). Intuitively, the derived expression is simply a generalisation of the standard convolution theorem. The formulation in terms of Euler angles in Eq. 10 allows for an efficient numerical implementation of the operation (see Sec. 3.1).

2.3 Beams

The three fields on the sphere: $\{\widetilde{I}, \widetilde{P}, \widetilde{V}\}$, that describe the instrumental beam in the above discussion are allowed to be independent as long as they conform to the constraint in Eq. 5. Of course, in a realistic case the fields are highly dependent; here we will discuss less general, but useful beam parameterizations.

2.3.1 Co- and cross-polarized beams

Polarized receivers are commonly characterised by their response to an electric field ϵ_{co} aligned to a reference direction (the co-polar response) and their response to the orthogonal field ϵ_{cx} (the cross-polar response). Clearly, co- and crosspolar responses are coordinate-dependent properties; in the case of linear polarization the co- and cross-polar basis is, by convention, the Ludwig-III basis (Ludwig 1973). In terms of the standard spherical basis, the unit vectors of this frame are given by:

$$\hat{e}_{(\mathrm{co})} = \sin(\phi) \,\hat{e}_{(\theta)} + \cos(\phi) \,\hat{e}_{(\phi)},\tag{12}$$

$$\hat{e}_{(\mathrm{cx})} = \cos(\phi) \,\hat{e}_{(\theta)} - \sin(\phi) \,\hat{e}_{(\phi)} \,. \tag{13}$$

The Ludwig-III basis has just a single coordinate singularity that can be placed in opposite direction to the beam centre in the detector's frame of reference. The beam centre is then in the \hat{z} direction where the coordinate system resembles a Cartesian system.

In the case where the optical response is completely described by the co- and cross-polar response, the instrumental Mueller transformation in Eq. 4 is the top row of a Mueller-Jones transformation. Simulations of the optical system, like the ones described in Sec. 4.3, can be used to estimate the optical response in this regime (see Appendix B).

Linear polarization instruments are generally designed to have minimal cross-polar response, and thus instrumental beams are often approximated by just the co-polar response. In this case, the response to circular polarization \tilde{V} vanishes while the polarized beam \tilde{P} is completely determined by the unpolarized beam \tilde{I} and a reference angle γ to the co-polar direction: the polarization angle. Using the Ludig-III basis (indicated by subscript \underline{r}), we then have:

$$\widetilde{I}_{\mathcal{L}}(x)e^{\pm 2i\gamma} = \widetilde{P}_{\mathcal{L}}(x)$$
 (co-pol. approx.). (14)

Using Eq. 12-13, one can show that the harmonic coefficients of the \tilde{P} beam in the (θ, ϕ) basis are related to those of the \tilde{I} beam by convolution with a harmonic kernel K (Hivon et al. 2017). When the support of the beam is small compared to the curvature of the celestial sphere, K is well approximated as diagonal per azimuthal mode m (see Appendix A1):

$${}_{\pm 2}b_{\ell m}^{\widetilde{P}} = e^{\pm 2i\gamma} \sum_{l'} b_{\ell'(m\pm 2)}^{\widetilde{I}} K_{\ell\ell' m}, \qquad (15)$$

$$\approx e^{\pm 2i\gamma} b_{\ell(m\pm 2)}^{\tilde{\ell}} \,. \tag{16}$$

2.3.2 Azimuthally-symmetric beams

The main beam of a well-behaved polarimetric instrument is often well approximated as being azimuthally symmetric; the beam can be described as a function of angular distance to the beam centre only. For the harmonic modes of the spin-0 \tilde{I} and \tilde{V} fields, this means that only the m = 0azimuthal modes are nonzero when the beam is placed on either pole of the (θ, ϕ) coordinate system. The case for the spin-2 field \tilde{P} is less obvious due to the coordinate singularities at the poles. In Appendix A3 we demonstrate why the only nonzero modes of the \tilde{P} field are $m = \pm 2$. As a result, the harmonic coefficients of the Stokes parameters on the (θ, ϕ) basis for an azimuthally symmetric beam centred on the pole obey:

$$b_{\ell m}^{\tilde{I}/\tilde{V}} \propto \delta_{m0}, \qquad (17)$$

$${}_{\pm 2}b^P_{\ell m} \propto \delta_{m\mp 2} \,. \tag{18}$$

This holds independently of approximating the beams as non-depolarizing or co-polar only.

In cases where the azimuthal symmetry is weakly broken, e.g. for detectors on the corners of a focal plane, or when a symmetric beam is not centred exactly on the pole due to detector pointing miscalibration, only a limited number of azimuthal (m) modes are usually needed to accurately describe the beam. In such cases, the data model in Eq. 10 still makes use of the relative sparsity of the harmonic representation, as the sum over s does not need to run over all $2\ell_{\rm max} + 1$ formally required values.

2.3.3 Gaussian and Elliptical Gaussian beams

At first order, the co-polar beam is generally well approximated by the diffraction pattern from a circular aperture. The centre region of the resulting Airy beam pattern is in turn shaped closely like an azimuthally symmetric Gaussian function with harmonic coefficients given by (Challinor et al. 2000):

$$b_{\ell m}^{\widetilde{I}} = \sqrt{\frac{2\ell+1}{4\pi}} \exp\left[-\frac{\ell(\ell+1)\sigma^2}{2}\right] \delta_{m0} \,. \tag{19}$$

In the same vein, the main beam of a detector placed far off-axis on a focal plane could be approximated by an elliptical Gaussian. Closed form expressions for the corresponding harmonic coefficients can be found in Souradeep & Ratra (2001) and Mitra et al. (2004).

2.3.4 Ghosting response

Internal reflections in a receiver, for example between lenses and focal plane, can create so-called ghost beams; mirror images of the main beam rotated away from the main beam centre (Fowler et al. 2007; Aikin et al. 2010). Optical ghosting is partially worrisome for high index-of-refraction materials such as silicon, necessitating advanced anti-reflective (AR) coating and detailed modelling and characterisation programs. Some CMB experiments using refracting telescopes have developed simulations to probe systematics caused by this effect (MacTavish et al. 2008; BICEP2 Collaboration 2015). The ghost contribution can simply be added to the fields describing the main beam and used in Eq. 10. While this method is conceptually convenient, the large number of azimuthal modes required to accurately describe the resulting azimuthally asymmetric beam make it numerically inefficient. An alternative approach wherein the ghost beam is effectively treated as a separate detector with its own pointing coordinates is therefore generally more efficient.

2.4 Modulation techniques

To reduce their dependence on accurate instrumental characterisation, current and future CMB experiments often incorporate modulation techniques that reduce the degeneracies between spurious systematic signal and the sky signal. We briefly discuss how to incorporate two common techniques: boresight rotation and half-wave-plate modulation, into the data model in Eq. 10.

2.4.1 Boresight rotation

Boresight rotation refers to physically rotating the telescope (stepwise) around the optical axis, or boresight. Having access to redundant observations made at different boresight angles is beneficial in many aspects. See e.g. the BICEP experiment and its successors (Takahashi et al. 2010; BI-CEP2 and Keck Array Collaborations 2015; Karkare et al. 2016). In terms of optical systematics, its main purpose is to suppress temperature-to-polarization leakage due to azimuthally asymmetric modes of the \tilde{I} beam (see Sec. 2.5). Boresight rotation is most naturally included in the data model by including it in the pointing: $(\psi_t, \theta_t, \phi_t)$ (in Eq. 10), while leaving the beam coefficients unchanged.

2.4.2 Half-wave plate modulation

A half-wave plate (HWP) is a birefringent material that changes the polarization state of incoming radiation of a specific wavelength and incidence angle by introducing a phase difference of π between the radiation component aligned along a direction intrinsic to the material (the fast axis) and the orthogonal component. Notably, for incident linearly polarized light, the effect is to mix Q and U by an amount based on the orientation of the HWP's fast axis in reference to the coordinate frame defining the Stokes parameters (see Sec. 2.5). Rotating the HWP thus results in a controlled modulation of the incoming linear polarization.

Half-wave plate modulation in the context of CMB polarimeters has been discussed in e.g. O'Dea et al. (2007), MacTavish et al. (2008) and Brown et al. (2009). In terms of suppressing optical systematics it differs qualitatively from the boresight rotation discussed in the above; both techniques effectively result in a controlled modulation of the linearly polarized signal of the sky, but boresight rotation does not leave the intensity signal unchanged in the case of azimuthally asymmetric beams. In contrast, an (ideal) HWP placed skywards of the telescope will leave the signal induced by the intensity beam unchanged, regardless of its shape, thus decoupling it from the modulated linearly polarized sky signal.

At subleading order, non-idealities in the HWP will spoil this behaviour by making the \tilde{I} and \tilde{V} beams weakly dependent on the HWP angle. See e.g. (Savini et al. 2006; Bryan et al. 2010; Essinger-Hileman et al. 2016). In terms of the data model in Eq. 10, HWP modulation is thus most generally described by beam coefficients that depend on the HWP angle. In case of an ideal skywards HWP however, the coefficients may be factored into two terms: one that does and one that does not depend on the HWP modulation angle:

$${}_{\pm 2}b^{\widetilde{P}}_{\ell m} \to {}_{\pm 2}b^{\widetilde{P}}_{\ell m}e^{\pm 4i\alpha} \qquad \text{(ideal HWP modulation)}, \tag{20}$$

where α is the HWP angle. The \tilde{V} coefficients simply pick up a minus sign $(b_{\ell m}^{\tilde{V}} \rightarrow -b_{\ell m}^{\tilde{V}})$ when the HWP is introduced and the \tilde{I} coefficients remain unchanged.

2.5 Systematics arising from map-making

After data acquisition, the polarized sky signal is reconstructed by solving the inverse problem associated with Eq. 3. Generally this is done by calculating a point estimate of the sky signal (a pixelized map) in a process called mapmaking. Commonly, the map-making estimator is a variation on the generalized least squares statistic, given by:

$$\hat{\boldsymbol{s}} = \left(\boldsymbol{A}^{\dagger} \boldsymbol{N}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{\dagger} \boldsymbol{N}^{-1} \boldsymbol{d} \,. \tag{21}$$

In case of Gaussian noise with an a priori known covariance N and uniform signal prior in the specified basis, this corresponds to the maximum a posteriori estimate. The (Gaussian) posterior around the maximum is then described by the $(A^{\dagger}N^{-1}A)^{-1}$ covariance matrix. In realistic analyses, this matrix is too large and dense to be available for regular ma-

trix calculations, but its operation on a map-sized array (as in Eq. 21) can be calculated iteratively. 6

Reconstruction of the three- or four-dimensional signal Stokes vector⁷ from the one-dimensional data crucially relies on knowledge of the linear transformation (A) that maps the sky signal onto the time domain. Any beam related systematics will come from an incorrectly assumed transformation. In the above, we have shown how, besides telescope pointing, calibration, sample flagging and other instrumental effects, the transformation should contain the beam-convolution. However, this aspect is often ignored in the map-making stage because the a priori knowledge of the beams is too poor to include them in a point estimate \hat{s} . Furthermore, for azimuthally symmetric beams, it is simpler to forward-propagate the effects by convolving a model of interest (sky map, power spectrum, etc.) with the beam.

We can gain some intuition for the effects of different beam non-idealities by considering limiting cases of the mapmaking procedure. We start by approximating the noise covariance as diagonal in the time sample domain (white noise): $N = \langle n_t n'_t \rangle \propto \delta_{tt'}$. Secondly, we describe the problem in the space \mathcal{P} of complex Stokes vectors p^{μ} (see Sec. 2.1). The map-making estimate then reduces to: $\hat{p} \propto \left(A^{\dagger}A\right)^{-1}A^{\dagger}d$ per pixel on the sky; we will refer to this as binning mapmaking. An ansatz for A that ignores the beam but is otherwise valid for a single co-polarized detector with vanishing polarization angle can be derived from Jones et al. (2007):

$$A = (A_{t,x})^{\mu} \propto \left(1, \frac{1}{2}e^{-2i\psi_t}, \frac{1}{2}e^{+2i\psi_t}, 0\right) \mathbf{1}_X(t).$$
(22)

The position angle ψ_t is included in Eq. 10 and the indicator function $\mathbf{1}_X(t)$ is defined to be 1 for time samples t in the set X of samples that hit pixel x and zero otherwise. We will restrict ourselves to the estimate $\hat{p}^1 = \hat{P}_x$: the linearly polarized signal in pixel x.

$$\hat{P}_x \propto [(A^{\dagger}A)^{-1}]^1_{\nu} \sum_{t \in P} (A_{t,x})^{\nu} d_t$$

The normalisation is then given by the inverse of:

$$\left[\boldsymbol{A}^{\dagger}\boldsymbol{A}\right]_{\nu}^{\mu} \propto \frac{1}{2} \sum_{t \in \boldsymbol{X}} \begin{pmatrix} 2 & e^{-2i\psi_{t}} & e^{+2i\psi_{t}} & 0\\ e^{+2i\psi_{t}} & \frac{1}{2} & \frac{1}{2}e^{+4i\psi_{t}} & 0\\ e^{-2i\psi_{t}} & \frac{1}{2}e^{-4i\psi_{t}} & \frac{1}{2} & 0\\ 0 & 0 & 0 & 0 \end{pmatrix}.$$
(23)

⁶ This model can be extended by jointly inferring the noise covariance (Prunet et al. 2001; Natoli et al. 2002; Wehus et al. 2012) or signal covariance (Wandelt et al. 2004; Eriksen et al. 2008; Taylor et al. 2008). Including statistical inference on part of the transformation A (e.g. the beam) in such approaches is relatively unexplored.

⁷ The fourth Stokes parameter V is often ignored in CMB data analysis as the cosmological signal is not expected to be significantly circularly polarized (King & Lubin 2016) and instruments are designed to be insensitive to it. Still, in the context of optical systematics, V cannot be entirely ignored. For instance, Zeeman splitting of oxygen in the Earth's magnetic field at 60 and 118.8 GHz provides a significant source of circular polarization for ground-based experiments (Hanany & Rosenkranz 2003; Hanany et al. 2013) which could be converted to linear polarization by non-ideal half-wave plates or other significant cross-polar responses (Nagy et al. 2017). We focus on a scan strategy that visits the pixel with a large uniformly distributed set of position angles ψ . Summing over *t* then diagonalizes the above, which results in a diagonal inverse matrix after we project out the singular *V* part by taking the pseudoinverse. After inserting the expression for d_t (Eq. 10), the estimate for *P* in pixel *x* becomes solely proportional to the m = 2 modes of the beam:

$$\hat{P}_{x} \propto \frac{1}{2} \sum_{\ell,m} \left\{ b_{\ell 2}^{\tilde{I}} a_{\ell m}^{I} + \frac{1}{2} \left(-2 b_{\ell 2}^{\tilde{P}} a_{\ell m}^{P} + 2 b_{\ell 2}^{\tilde{P}} -2 a_{\ell m}^{P} \right) + b_{\ell 2}^{\tilde{V}} a_{\ell m}^{V} \right\} q_{\ell} \, _{2}Y_{\ell m} \Big|_{x} \, .$$

$$(24)$$

Here $|_x$ indicates evaluation at the θ, ϕ coordinates of the pixel centre and q_{ℓ} is given by Eq. 11. This expression makes explicit how a nonzero quadrupole (m = 2) mode of the unpolarized beam biases the \hat{P} estimate by modulating the dominant unpolarized sky signal just like the linearly polarized signal. As we are already in the limit of perfectly uniform position angle coverage, it is clear that boresight rotation cannot modulate this bias away. Additionally, in the limit of uniform sky coverage, one can show (Hu et al. 2003; O'Dea et al. 2007) that the real part of $b^I_{\ell 2}$ purely sources temperature-to-*E*-mode $(I \to E)$ leakage, while the imaginary part is responsible for $I \to B$ leakage.⁸ In more practical terms, the real and imaginary parts of $b_{\ell^2}^I$ correspond to the components of the unpolarized beam with azimuthal parts proportional to $\cos 2\phi$ and $\sin 2\phi$ respectively. Only in this limit (perfect uniform coverage in θ , ϕ and ψ), the leakage may be described independently per mode; leakage from ℓ to ℓ' and m to m' will occur in more general cases (Hu et al. 2003; Hanson et al. 2010; Hivon et al. 2017). In the case of azimuthally symmetric beams (Eq. 17-18), the estimate reduces to a symmetrically smoothed version of the signal that is independent from the I and V sky.

The situation changes when, instead of relying on rotating the instrument (or Earth's rotation), the angular information needed to solve for p^{μ} is obtained by half-wave plate modulation. In this case, the data model assumed for map-making would be expanded as follows:

$$\mathbf{A} = (A_{t,x})^{\mu} \propto \left(1, \frac{1}{2}e^{-2i(\psi_t + 2\alpha_t)}, \frac{1}{2}e^{+2i(\psi_t + 2\alpha_t)}, 0\right) \mathbf{1}_X(t), \quad (26)$$

where α denotes the HWP angle (see Sec. 2.4.2). The estimate for the linearly polarized component (at fixed position angle ψ) then becomes proportional to:

$$\hat{P}_{x} \propto \frac{1}{2} \sum_{t \in X} \sum_{\ell,m} \left\{ b_{\ell s}^{\widetilde{I}} a_{\ell m}^{I} + \frac{1}{2} \left(-2 b_{\ell s}^{\widetilde{P}} 2 a_{\ell m}^{P} + 2 b_{\ell s}^{\widetilde{P}} -2 a_{\ell m}^{P} \right) + b_{\ell s}^{\widetilde{V}} a_{\ell m}^{V} \right\} q_{\ell} e^{4i\alpha_{t}} {}_{s} Y_{\ell m} \Big|_{x}.$$

$$(27)$$

⁸ The E- and B-mode harmonic coefficients are given by:

$$a_{E,\ell m} = -\frac{1}{2} (_2 a_{\ell m}^P + _{-2} a_{\ell m}^P), \quad a_{B,\ell m} = \frac{i}{2} (_2 a_{\ell m}^P - _{-2} a_{\ell m}^P).$$
(25)

In terms of covariant derivatives of the symmetric traceless (ST) part of W (Eq. 1) we have $a_{E,\ell m} = N_\ell \int dx \nabla^i \nabla^j (W_{\rm ST})_{ij} \overline{Y}_{\ell m}$ and $a_{B,\ell m} = N_\ell \int dx \, \widetilde{\epsilon^i}_k \, \nabla^k \, \nabla^j (W_{\rm ST})_{ij} \overline{Y}_{\ell m}$ ($\widetilde{\epsilon}$ is the Levi-Civita symbol and $N_\ell \equiv \sqrt{2(\ell-2)!/(\ell-2)!}$) (Kamionkowski et al. 1997), illustrating that $a_{E/B,\ell m}$ are harmonic modes of a rotational scalar and pseudoscalar field.



Figure 2. Left: log-log plot of required CPU time for the inverse spin-weighed spherical harmonic transforms needed for a single (linearlypolarized) beam as function of band-limit ℓ_{max} . The different marker types refer to the azimuthal band-limit (s_{max}) of the beam. The black dots correspond to convolution with an azimuthally symmetric beam. The dashed line shows the expected asymptotic scaling (with arbitrary normalisation). The results conform relatively well with the expected scaling, but show small, step-like deviations due to changes in the pixelisation scheme (we let N_{side} be the smallest power of 2 that is larger than $\ell_{max}/2$). Results are from a single thread on an Intel Xeon E5-2697 v2 core running at 2.7 GHz. Right: log-log plot of required CPU time for producing time-ordered data as function of the scan duration (on the same single core setup). We use a sampling frequency of 100 Hz with pointing quaternions and beamconvolved maps preloaded in memory. The scaling follows the expected linear relation with the number of time-samples n_{samp} (illustrated by the arbitrarily normalised dashed line). The required CPU time is largely independent of the number of pixels ($n_{pix} = 12N_{side}^2$) of the convolved maps and weakly linearly dependent on the azimuthal band-limit s_{max} of the beam. Again, the black dots denote the azimuthally symmetric case. No interpolation is used while sampling the data.

If the set of HWP angles is large and uniformm the only non-vanishing terms are proportional to $\exp(-4i\alpha_t)$. With an ideal (skyward) HWP that is true for $_{-2}b_{\ell s}^{\tilde{P}}$ (see Eq. 20), while the \tilde{I}, \tilde{V} coefficients remain constant with α . The \hat{P} estimate is then only sourced by the linearly polarized sky (regardless of the shape of the beam). Of course, subleading $E \leftrightarrow B$ leakage due to azimuthal asymmetry ($m \neq \pm 2$ modes) of the linearly polarized beam is not suppressed by HWP modulation, but requires sky or boresight rotation to be suppressed. Finally, any cross-polar components (e.g. a miscalibrated polarization angle) of the linearly polarized beam are not suppressed by HWP modulation, nor can its azimuthally-symmetric part be suppressed by a uniform sampling of position angle ψ .

The above examples provide intuition for the cause of some of the leading order optical systematic effects. Another leading order effect is a simple miscalibration of the (dominant) azimuthally symmetric co-polarized part of the beam, e.g. by incorrectly assuming it to be Gaussian. Such a mistake will, on its own, not mix I and P or E and B but will still result in a wrongly inferred amplitude of anisotropies. This is especially problematic at small angular scales where deviations in the amplitude of the CMB power spectra are highly degenerate with varying effective beam size.

Any realistic map-making algorithm is capable of jointly solving for the signal estimate \hat{p} using data from multiple detectors. Additionally, more sophisticated algorithms than those used in the examples above exist. One common choice is the so-called pair differencing method (Jones et al. 2007). Here, the linearly polarized signal is directly estimated from the differenced TOD from detector pairs that share a physical location on the focal plane but are coupled to orthogonal linearly polarizing interfaces. The resulting estimate uses suboptimal noise weighting compared to using both detectors independently. However, the cancelation of common modes in the noise or unpolarized signal that may otherwise be difficult to explicitly model is advantageous. The estimate is similarly uninfluenced by common features in the \tilde{I} beams, e.g. a shared azimuthally asymmetric component. One can check that this also holds true when the mapmaker from Eq. 22 is used with these paired detectors. On the other hand, any \tilde{I} beam component that does not cancel exactly, regardless of its azimuthal dependence or spin, directly biases the \hat{P} estimate by $I \rightarrow P$ leakage. This includes miscalibrated gain or beamwidth differences between two paired detectors (BICEP2 Collaboration 2015). These two systematic effects do not result in $I \rightarrow P$ leakage when a map-making scheme like Eq. 22 is used.

Finally, two other map-making approaches that attempt to correct for beam effects are worth mentioning. The first, as proposed in Bock et al. (2009) and Wallis et al. (2014), uses an ansatz for A that is similar to Eq. 22, but contains a number of additional harmonics such as $\exp(\pm i\psi_t)$ or $\exp(\pm 3i\psi)$. The resulting map-making estimate has a higher dimension than the standard $\{\hat{I}, \hat{P}\}$ estimate and therefore projects out modes that are necessarily spurious. Of course, the method results in increased uncertainty in the $\{\hat{I}, \hat{P}\}$ estimate and is unable to project out the most problematic spurious signal: the one proportional to $\exp(\pm 2i\psi_t)$. Another method, described in e.g. Armitage-Caplan & Wandelt (2009) and Keihänen & Reinecke (2012), imposes a maximally informative prior on the beam by directly using the full beam-convolved data model from Eq. 10 as ansatz for A. Computing the point estimate by solving Eq. 21 becomes much more involved but can still be done using the conjugate gradient method and by regularising the singular part

Table 1. Optical properties of the two-lens silicon designed considered in this analysis. Note that c and k represent the inverse radius of curvature and the conic constant, respectively. The silicon lenses are assumed to have an index of refraction of $n_{\rm si} = 3.42$ and the physical separation between primary and secondary lens is 550 mm. The focal plane is located 230 mm behind the secondary lens.

Lens	surface	$c [\mathrm{m}^{-1}]$	k	f-number
Primary Secondary	Sky Focal plane Sky Focal plane	1.446 0.933 1.635 0.834	0.141 1.369 -0.052 14.841	1.56 1.56



Figure 3. Ray trace diagram of the proposed design showing the path of a few fields through a pair of convex-concave lenses. An optical stop is located right after the primary lens on the image (right) side. The outermost pixel (blue rays) is at a 15° angle relative to boresight.

of A. Still, the method has not been demonstrated to work with non-white noise or high resolution data ($\ell_{\rm max} > 2000$). Arguably, a more significant challenge associated with this method is the one alluded to in the introduction to this section: any prior uncertainty on the beam is lost in the map-making procedure. This, together with the high numerical demands, and dependence on map-making schemes, may suggest that methods relying on forward-propagating beam effects are generally more useful than those that deconvolve the beam.

3 CODE DESCRIPTION

The primary functionality of the **beamconv** code library is to compute time-ordered data (TOD) that includes spurious signal due to optical systematics. The resulting TOD may be used as input to pipelines that describe further stages of data acquisition and analysis (e.g. addition of detector noise, time stream filtering and map-making), as beam convolution is a natural first step in any simulation pipeline. Alternatively, the library provides simple map-making functionalities to help assess the systematic signal in the noiseless limit.

The code is written in Python and relies heavily on the standard scientific computing package numpy. The (inverse) spherical harmonic transforms are handled by the highly optimised libsharp library (Reinecke & Seljebotn 2013) using an interface provided by the healpy Python package.⁹ All pointing related computations are done by interfacing with the qpoint library.¹⁰ The code is setup for parallel computing on massive distributed memory systems using the MPI standard. The library is bundled with several explanatory IPython notebooks.

The capabilities of the library partially overlap with those of the Time-Ordered Astrophysics Scalable Tools (TOAST) package.¹¹ The public version of TOAST has recently been upgraded with an interface to the beam convolution library conviqt (see (Prézeau & Reinecke 2010)) and thus should be able to produce similar results as beamconv. Clearly, TOAST is a more extensive simulation package, that is also capable of reproducing instrumental effects that are not optics-related. Instead of trying to reproduce the TOAST library, we aim to have beamconv purely focused on optical systematics and hope to provide an accessible tool that can be extended to include more optical systematic effects with relative ease.

In the following sections we will briefly go over the technical details of the convolution operation, explain the input and output of the code and provide a few benchmark results. Finally, we comment on possible future additions to the code.

3.1 Implementation

The beam convolution operation is performed over the full sky as point-wise multiplication in the harmonic domain, using the expression for the data in Eq. 10. The method is thus heavily inspired by the work of Wandelt & Górski (2001) and is implemented similarly to the totalconvolver and conviqt implementations of this method described in Reinecke et al. (2006) and Prézeau & Reinecke (2010) respectively. We will briefly discuss our implementation.

Calculating beam-convolved data by evaluating Eq. 10 at each time sample is equally inefficient as evaluating an integral over the sphere at each sample (see Eq. 6). The first expression is only efficient because it allows separate treatment of the convolution and data sampling. To do so means that one, for each azimuthal mode s of the beam, first evaluates the following inverse SWSH transformation over the entire sphere using available $O(\ell_{\text{max}}^3)$ algorithms:

$${}^{(s)}f(\theta,\phi) = \sum_{\ell,m} {}^{s}f_{\ell m}q_{\ell} {}^{s}Y_{\ell m}(\theta,\phi), \quad \forall \, \theta,\phi \in S^2,$$
(28)

with harmonic modes given by Eq. 11 and:

$${}_{s}f_{\ell m} = b_{\ell s}^{\widetilde{I}}a_{\ell m}^{I} + \frac{1}{2}\left(-2b_{\ell s}^{\widetilde{P}} 2a_{\ell m}^{P} + 2b_{\ell s}^{\widetilde{P}} - 2a_{\ell m}^{P}\right) + b_{\ell s}^{\widetilde{V}}a_{\ell m}^{V}.$$

Once the ${}^{(s)}f(\theta,\phi)$ maps are computed for each *s*, the TOD may be sampled from them using the (θ_t,ϕ_t) pointing information and time-dependent phase given by the factor $\exp(-is\psi_t)$. As long as the synthesised maps can be stored in memory, data from any sort of scan strategy may be obtained. The overhead given by the inverse SWSH transforms is constant.

 9 http://healpix.sourceforge.net

¹⁰ https://github.com/arahlin/qpoint

¹¹ http://github.com/hpc4cmb/toast



Figure 4. Top left: Strehl ratio (as output by Zemax) of the proposed optical system as a function of field location for three frequencies, 90, 150, and 270 GHz, represented by top, middle, and bottom solid lines respectively. The three dashed lines correspond to the Strehl ratios for the same optical design if the silicon effective index of refraction were $\tilde{n}_{\rm si} = 3.39$ instead of $n_{\rm si} = 3.42$. Top right: Beam ellipticity as a function of Strehl ratio. Solid lines correspond to average within a fixed interval whereas the dashed lines (90 and 150 GHz) are the corresponding results for $\tilde{n}_{\rm si} = 3.39$. Lightly-coloured markers indicate the results of individual simulations (300 markers for each frequency). The markers are omitted for the case of $\tilde{n}_{\rm si} = 3.39$. Also note that we have zoomed in on the area of highest Strehl ratio which contains most of the 90- and 150-GHz detectors. It is interesting to see that the rise in ellipticity at 90 GHz appears to be dominated by diffraction effects and not a significant drop in Strehl ratio. For 150 GHz, we note a more complex behavior. Bottom left: Beam ellipticity as s function of field location. Again, the solid line traces the average and the dashed lines correspond to the case where $\tilde{n}_{\rm si} = 3.39$. Markers identify individual simulation results. Bottom right: Distribution of beam ellipticity for all pixels simulated. At 270 GHz, the beam ellipticity distribution is relatively flat from 0-0.06 and we choose to zoom in on the distributions for lower ellipticities.

Given that diffraction naturally truncates the beam coefficients at some finite $\ell_{\rm max}$, the transforms only need to be computed up to $\ell_{\rm max}$. This can be done with an asymptotic $O(\ell_{\rm max}^3)$ scaling, which will dominate the total scaling for simulation runs with large $\ell_{\rm max}$ and few data samples (see Sec. 3.3)

Note that in **beamconv** we perform separate inverse transforms for the \tilde{I} and \tilde{P} beams (ignoring \tilde{V}). This is done such that that we may modulate the linearly polarized signal independently from the total intensity component. This is, for instance, used to to incorporate time-dependent HWP modulation (see Eq. 20).

By default, the TOD are directly sampled from (equal area) HEALPix pixels. We have found that this approach suffices (as long as the $N_{\rm side}$ parameter is larger than $\ell_{\rm max}/2)^{12}$, but if needed, e.g. when high accuracy is needed at scales close to $\ell_{\rm max}$, the data may be interpolated using bi-linear interpolation. The inverse spherical transforms provided by healpy synthesise the harmonic coefficients onto the full sky,

which is wasteful for experiments that observe small patches of the sky, but we allow this small hit in efficiency and defer an improvement to future work.¹³

It might seem natural to realise the modulation by $\exp(-is\psi_t)$ with an FFT over the pixels of the synthesised maps (Eq. 28). The resulting $f(\psi, \theta, \phi)$ function will have s_{\max} samples over ψ which is typically a low ($s_{\max} \ll \ell_{\max}$) number due to the azimuthal band-limit of the beam. The TOD can then be directly interpolated from f without manually iterating over s. In practise, we have found it more efficient in terms of memory and speed as well as more accurate to simply use the ψ_t pointing data and directly apply the factor $\exp(-is\psi_t)$ when the TOD are sampled from the synthesised maps. We thus treat each value of s independently, adding to the TOD with increasing s. We use recursion of the form

¹³ Unlike the original implementation suggested in Wandelt & Górski (2001) that uses fast Fourier transforms (FFTs) in the θ and ϕ (and ψ) directions (by restating the problem on the 3-torus instead of the SO(3) rotation group), libsharp does not use an FFT over the θ direction. This allows one to skip latitude rings that are not visited by the detector pointing (θ_t, ϕ_t).

¹² The N_{side} parameter is a power of 2 that determines the number of pixels within the HEALPix pixelisation scheme ($n_{\text{pix}} = 12 N_{\text{side}}^2$).

 $\exp(i(s+1)\psi_t) = \exp(is\psi_t)\exp(i\psi_t)$ to avoid unnecessary calls to trigonometric functions.

3.2 Simulation input

To evaluate the expression for the beam-convolved data in Eq. 10, we provide the telescope pointing and the spin-weighed spherical harmonic (SWSH) coefficients of the assumed sky and beams. We will briefly detail these ingredients.

Following the analytical expression in Eq. 10, the input SWSH coefficient of the beams are the $b_{\ell m}^{\tilde{I}}$ and $_{\mp 2}b_{\ell m}^{\tilde{P}}$ coefficients as defined in Eq. A1, A2 and A4. The corresponding beams are assumed to be defined on the (θ, ϕ) coordinate system and should generally be centred on the north pole. Because of the redundant description in terms of \tilde{P} and its complex conjugate, only the $m \geq 0$ modes need to be provided. In cases where the co-polar approximation is used (Eq. 16) only the \tilde{I} coefficients are required.

The beam coefficients are associated with one or several detectors. Each detector is represented as a separate instance of a Python class that contains pointers to the beam coefficients as well as properties such as detector pointing offset coordinates, polarization angle and beam band-limits. Additionally, each detector may be linked to other detector instances that serve as ghosting beams. These ghost detectors are treated as fully independent detectors with independent beam coefficients and properties, but are automatically added to the main detector data during data sampling. See the discussion in Sec. 2.3.4.

Internally, all pointing calculations are performed with **qpoint** using the computationally efficient unit quaternion representation (Hamilton 1866), rather than the more conventional matrix/vector algebra. We separate the pointing information into boresight pointing and per-detector offset pointing coordinates. The boresight pointing, representing the pointing direction of the telescope at each time sample, is independent from the detectors. The boresight quaternions (time-ordered data) may either be loaded from disk, calculated in real-time by a user-provided function or one of few preset scanning strategies. Note that the **qpoint** library may be used to convert pointing information in Equatorial coordinates (RA, DEC and Position Angle ψ) or horizon coordinates (azimuth, elevation and roll) to a suitable time stream of unit quaternions.

The detector pointing offset is unique to each detector and is assumed to be constant with time. The offset physically reflects the different fields of view for detectors placed at different locations away from the telescope's bore axis. It is realised as an active rotation $g^{(\Delta)}$ away from the boresight pointing direction, specified by an azimuth *a* and elevation *e* angle defined relative to the boresight direction (i.e. the north pole) a = e = 0:

$$g^{(\Delta)} = g_{\hat{\mathcal{T}}}(-a)g_{\hat{\mathcal{V}}}(e)g_{\hat{\mathcal{T}}}(0) \qquad (\text{ZYZ Euler convention}). \tag{29}$$

Here $g_{\hat{Z}}$ and $g_{\hat{Y}}$ represent rotations around the fixed Z and Y axes respectively.¹⁴ The polarization angle does not correspond to a physical rotation but is considered as an intrinsic

property of the linearly polarized beams and is therefore effectively applied to the \tilde{P} coefficients. The same argument applies to the ideal skyward HWP: its effect is internally handled by modulating the TOD due to the linearly polarized beams by $\exp(\pm 4i\alpha_t)$ (with HWP angle α).

Finally, the harmonic coefficients of the sky are provided in terms of $a_{\ell m}^I$ and E- and B-mode coefficients (see Eq. 25). Again, only modes with $m \ge 0$ are required.

3.3 Benchmarks

We provide some basic benchmark results to illustrate the scaling with beam band-limits ℓ_{\max} , s_{\max} and the scan duration (see Fig. 2). The results consist of two parts: the first shows required CPU time for evaluation of the convolution without any scanning (i.e. Eq. 28 for all $|s| \leq s_{\max}$) as a function of the beam band-limits. The results show the expected total $O(\ell_{\max}^3 s_{\max})$ scaling. As the azimuthal band-limit will rarely exceed the maximum depicted value $s_{\max} = 8$, the results give a rough indication of wall time in practise. The computations are completely dominated by the libsharp inverse SWSH transforms which can be sped up with the use of OpenMP threads and/or MPI tasks (see Reinecke & Seljebotn (2013)). For this test we use sequential execution on a single Intel Xeon E5-2697 v2 core running at 2.7 GHz.

The second result illustrates that sampling the TOD is largely independent of the properties of the beam once the convolved maps (Eq. 28) are stored in memory. This is demonstrated by performing a number of scans with total duration ranging from 1.5 min to 50 days. The sample rate is set at 100 Hz. The convolved maps and pointing data are preloaded into memory to isolate the test from the inverse SWSH transforms and I/O. The test is again run sequentially on the same type of core as the previous test. The results are practically identical in case of low resolution ($N_{\text{side}} = 256$) and high resolution ($N_{\text{side}} = 2048$) convolved maps. There is a constant linear scaling with azimuthal band-limit s_{max} . In general, the CPU time for the data sampling part of the procedure scales completely linearly with the number of data samples.

As expected, the timing results for a completely azimuthally symmetric beam (the black dots) lie approximately a factor 4 lower than the $s_{\text{max}} = 2$ points in the left panel of Fig. 2. In this case only 2 SWSH transforms are used (versus 8 for the $s_{\text{max}} = 2$ case).

When the two panels in Fig. 2 are compared, it can be seen that for small-aperture experiments (i.e. $\ell_{\rm max} \lesssim 2000$) computation time needed for the SWSH transforms is sub-dominant to that of the data sampling procedure. As computation time for data sampling remains constant with increasing beam band-limit, the SWSH transforms will dominate computation time for large-aperture telescopes (i.e. $\ell_{\rm max} = O(10^4)$). We comment on this case in the next section.

3.4 Future additions

For high resolution experiments, the multiple inverse SWSH transforms required per detector become impractical due to

above would also erroneously rotate the unpolarized beam and its (possibly nonzero) azimuthally asymmetric modes.

¹⁴ Note that we do not include the detector's polarization angle γ as a first rotation. Using $g_{\hat{Z}}(\gamma)$ as the starting rotation in the

their $O(\ell_{\max}^3)$ scaling. When full sky convolution is still desired (in the case of a large observed patch of sky or wide sidelobes), an approach similar to (Elsner & Wandelt (2014); see also (Hincks et al. 2010; BICEP2 Collaboration 2015)) should be used. These approaches work by describing the detector beams as linear combinations of a number of basis functions. A small number of basis functions generally suffices due to the relatively small changes in beam properties across a focal plane. Each of the basis functions are convolved with the simulated sky map and the resulting maps can be stored in memory shared between computer cores. Due to the linearity of the convolution operator, the TOD for each detector can then be sampled from a linear combination of the precomputed maps. This approach could even be extended to simulate the effect of detector bandpass differences or beams that are not constant during data acquisition due to e.g. temperature drifts or processes dependent on pointing elevation or HWP angle. We hope to report on the feasibility of this method in future work.

4 INSTRUMENT SETUP

4.1 Overall design considerations

Using the code library presented in Section 3, we choose to study a two-lens satellite refractor telescope designed to observe the CMB at two frequencies, 90 and 150 GHz. Designs similar to the one presented here have been considered in design studies for fourth generation CMB satellites (Bock et al. 2009; Suzuki et al. 2018). A very rough CAD model is shown in Figure 1. The design includes two silicon lenses embedded in a cold optics sleeve ($\leq 4 \text{ K}$) and two concentric radiation shields which prevent direct illumination of the primary lens from the sun. We choose not to incorporate a forebaffle mounted close to the location of the primary lens for fear of polarized reflections and/or increased loading from a blackened load. This puts stringent but quantifiable requirements on internal baffling and scattering in lenses and filters, which would have to be characterised in the lab prior to deployment.

A symmetric on-axis refractor design offers relatively straightforward baffling solutions and a large active focal plane area for a fixed volume design (see Figure 3). This design also allows for extensive pre-flight optical characterisation at operational temperatures through the use of a simple test cryostat. Of course, a single optics tube refractor design is hampered by current technological inability to produce anti-reflection coatings that are effective over more than an octave in frequency Datta et al. (2013); Young et al. (2017); Defrance et al. (2018). On-axis refractor systems are also more susceptible to internal reflection (ghosting). Although AR-coating challenges of refracting telescopes might mean that reflecting telescopes will ultimately be selected for a 4th-generation CMB satellite mission observing in the primary CMB frequency bands, we choose to further explore this design because of its inherent simplicity and pre-flight characterisation potential.



Figure 5. Beam ellipticity at 270 GHz as a function of focal plane location. The layout of our proposed focal plane consists of seven detector tiles each with 331 physical pixels of 6 mm diameter, for a total of 2317 pixels. Assuming each pixel is dichroic, this focal plane could support 4634 individual channels. The 200 physical pixels that are included in these simulations are highlighted with zero transparency. Each physical pixel consists of a pair of orthogonally polarized detectors.



Figure 6. Absolute values of beam responses. Top left: Single 90-GHz detector co-polar beam map in units of dBi (forward gain over isotropic). The contours indicate -3, -10, -13, and -20 dB relative to maximum, respectively. Top right: The corresponding cross-polar beam map. This particular detector is located approximately 150 mm from the center of the focal plane (edge pixel) and has an ellipticity of $e \approx 0.025$. Bottom left: stacked co-polar beam response at 90 GHz derived by averaging individual beam maps from 100 detectors. Each individual beam is normalized to peak at unity before summing, the final sum is then normalized again. Bottom right: the corresponding cross-polar response which peaks at -33 dB relative to the co-polar response.



Figure 7. Top: Azimuthally averaged beam profiles for the 400 detectors used in these simulations. The best-fit Gaussian beam model to the corresponding stacked (focal plane averaged) beam are shown in grey. Note how the Gaussian model falls off much more quickly with angle. It is clear that significant solid angle is contained in the diffraction sidelobes predicted by GRASP. Bottom: GRASP physical optics (PO) and method of moments (MoM) predictions for the extended sidelobes of the center pixel. Note how the MoM beam profile has significantly more power at wide angles. This is partially caused by internal reflections in the silicon lenses which are not accounted for by the physical optics calculations. Extended sidelobes can couple to the Galaxy and create a fake polarized signal (see Section 5.4). The interference patterns visible in the PO curves (blue and green) are caused by the finite number of frequencies used to simulate the optical response (5 frequencies per band).

4.2 Optical components

Figure 3 shows a ray tracing diagram of the proposed twolens design. The design employs two roughly 380-mm diameter silicon lenses with a maximum zag of about 16 mm on the primary lens. The system has an average effective f-number spanning 1.5–1.7 and a telecentricity angle not exceeding 0.1° over the entire field. The corresponding Strehl ratios at 90, 150, and 270 GHz for this design are shown in Figure 4 (see also discussion in Section 4.3). Table 1 describes the key optical design parameters.

Because of the high index of refraction, the silicon lenses can support a relatively wide diffraction-limited field of view (DLFOV) of approximately 30 degrees. This corresponds to an active focal plane area with a diameter of approximately 290 mm (0.103 deg/mm plate scale) and about 2500 physical pixels, assuming a 6 mm pitch size (see Figure 5). By employing dichroic bolometers with two polarization directions, this telescope could support 5,000 bolometer channels. In comparison, publications discussing the proposed Lite-BIRD satellite have suggested that the mission will deploy approximately 2,000 channels (Matsumura et al. 2016). A similar number of detectors were proposed for the CORE satellite which employed a two-mirror reflector design (Delabrouille et al. 2018). We note that advances in AR coating technology might allow for the replacement of the centre tile with one populated with pixels spanning the 220- and 270-GHz frequency bands (Coughlin et al. 2018; Nadolski et al. 2018). At those frequencies, it would be sensible to deploy smaller pixels to reduce spillover on the cold stop.

4.3 Physical optics simulations

The spatial response of the detectors are simulated using physical optics (PO) and physical theory of diffraction (PTD) simulations as provided by GRASP in results provided by the method of moments (MOM) module (GRASP User's Manual 2018).¹⁵ The physical optics simulations propagate pixel illumination patterns in succession through the two lenses and out into the far field. The pixel beam illumination pattern is based on a model of a photolithographed bolometer array coupling to corrugated feedhorns, similar to those designed by NIST for ACTPol and Advanced ACTPol (Niemack et al. 2010; Koopman et al. 2016). Given the relative simplicity of the optical system, the PO simulations are sufficiently fast that they can be generated for hundreds of detectors in a reasonable amount of time (few days) on a workstation computer.

In order to capture the focal plane distribution of the beam response, while also providing sufficient coverage to adequately capture aspects of the satellite scan strategy, we have randomly sampled 200 physical pixels spanning the entire focal plane (see Figure 5). In order to inject an additional level of realism to these simulations, we have allowed for some variation in the shape of the pixel beam used to illuminate the secondary lens. The distribution of beam size and ellipticity for the focal plane used in these simulations is shown in Figure 4. We calculate ellipticity, e, according to

$$e = \frac{\sigma_x - \sigma_y}{\sigma_x + \sigma_y},\tag{30}$$

where σ_x and σ_y are the Gaussian beamwidths along the two principal axes, with $\sigma_x > \sigma_y$. Figure 6 shows the focal plane averaged (stacked) co- and cross-polar beam response at 90 GHz. The stacking procedure washes out any azimuthal asymmetry in individual co-polarized beam maps. Note that the average geometrical cross-polar response is at -33 dB amplitude relative to the co-polar beam. This should be dominated by cross-polar effects originating in the detector architecture itself, for example through cross-talk in detector readout circuits.

The simulated detector beams are used to create 200 detector pairs. This corresponds to a scenario where two

¹⁵ GRASP is an antenna and optical modeling software capable of providing physical optics and method of moments calculations at mm-wavelenghts. See: https://www.ticra.com/



Figure 8. Left: Normalized hits map (samples per pixel) in Galactic coordinates for the typical 400-detector scan generated by following the scan strategy presented in Section 4.6 for one year at a 96.73 Hz sampling frequency. At this sampling rate, one year of scanning with 400 detectors produces a total of 1.2×10^{12} samples. Right: Condition number of I, Q, U covariance matrix for the same scan strategy.

perpendicularly linearly polarized radiation coupling devices feed optical power to separate bolometers. In this case, the Stokes \widetilde{I} and \widetilde{V} beams are shared between the two bolometers while the \tilde{O} and \tilde{U} beams only differ by a factor (-1) due to the 90° polarization angle difference. The exact common pointing, shared beams and 90° polarization angle difference for all pairs exactly cancels all $I \rightarrow P$ leakage due to differential pointing/beamwidth or azimuthally asymmetric modes of the I beams (see the discussion in Sec. 2.5). This setup allows us to focus on less explored systematic effects, such as $E \to B$ leakage due to cross-polar beam components and $m \neq \pm 2$ azimuthally asymmetric modes of the $\overline{Q}/\overline{U}$ beams. Of course, this cancellation is only approximate in realistic cases; such modifications could be included trivially in the presented framework. For example, in Sec. 5.3 we relax this condition by breaking some of the detector pairs to illustrate the $I \rightarrow P$ leakage due to the azimuthally asymmetric modes of the \tilde{I} beams.

We convert the physical optics results into (spin-weighted) harmonic modes of the corresponding beams following the method explained in Appendix B. We use a bandlimit of $\ell_{\max} = 1000$ and use $m_{\max} = 4$ as azimuthal bandlimit for each of the beams. We find the beam components with m > 4 too small (on all angular scales) to be significant for the presented analysis.

4.4 Simulation of far sidelobe response

The physical optics simulations described in Section 4.3 naturally incorporate lens and cold stop diffraction effects that cause far sidelobe response. However, those simulations do not factor in the impact of the two radiation shields and/or other passive optical components, such as internal baffling, on the far field response of the telescope. Scattering from impurities and other non-idealities in silicon lenses and filters as well as reflections internal to the optics tube are particularly challenging to model and we omit those effects in the general part of this analysis.

Figure 7 shows the 90- and 150-GHz azimuthally averaged beam profiles that are predicted by the physical optics simulations. With the exception of the beams used for analysis presented in Section 5.4, the beam profiles predicted by physical optics are apodized at a 4° angle from beam centre as part of the spherical harmonic decomposition re-

quired for beamconv input. Of course, off-axis pickup will continue past this 4-deg cutoff angle. In order to further explore far-sidelobe pickup as a candidate for degree scale B-mode systematics, we look at predictions for sidelobe response from both physical optics and method of moments simulations; these results are also shown in Figure 7. In the reciprocal sense, the method of moments calculation propagates an electric field emitted at the focal plane through the two silicon lenses and an ideal anti-reflection coating. Surface currents induced in these materials are then combined with the field sourced from the focal plane to calculate a farfield electric field distribution. The off-axis beam response from these method of moments calculations, which is only calculated for a pixel at the centre of the focal plane, is then combined with the physical optics beam maps produced on a per-pixel basis to create a hybrid beam model that is used for the analysis presented in Section 5.4. Off-centre pixels will obviously have non-symmetric sidelobes; however, we choose to only conduct a single full-sky MoM calculation in order to save computation time.

As is evident from Figure 7, the sidelobe amplitude predicted from the MoM calculation is significantly higher than that of the PO calculations. This is partially caused by the fact that the method of moments approach ignores passive optical elements such as a cold and absorbing optics tube and allows fields to freely propagate past the primary lens. In comparison, the physical optics calculations effectively ignore any power that does not propagate through the primary lens. We believe that the MoM beam profile response represents a worst case scenario for the proposed optical design. However, we also note that 1% Lambertian scattering in the silicon primary lens will create a beam sidelobe profile with a comparable amplitude.

4.5 Input maps

We use two sets of HEALPix input maps to generate the simulations presented in Section 5: a CMB-only map generated as a Gaussian random field using the **synfast** program (part of the HEALPix software library) using a standard Λ CDM cosmology, and a map that combines that CMB map with an estimate for dust contributions (*I* and *P*) in our own galaxy based on a Commander dust foreground template (Planck Collaboration 2016b).



Figure 9. Left: A 10-deg square box showing a Stokes-Q difference map between scanning simulations conducted assuming a symmetrical Gaussian beam model and an elliptical Gaussian beam model (units are μ K). Right: Same as left panel, but now differencing a Gaussian beam model with a full physical optics beam model. Note that the color scales on the two panels differ by a factor of 10 and the physical optics model shows substantially more large scale power in addition to the smaller angular scale residuals which have lower amplitude.

The power spectrum input to synfast has neither primordial nor lensing *B*-modes at all scales; this allows us to attribute residual *B*-mode power in rescanned maps to beam non-idealities. The Commander dust template allows us to assess interplay between beams and the galaxy, including the impact of $I \rightarrow P$ leakage through sidelobe coupling to low galactic latitudes. This particular beam-related systematic might constitute a significant challenge for CMB experiments hoping to constrain the epoch of reionization by measuring the associated large-scale *E*-mode signal.

4.6 Scan strategy, sampling frequency, and duration

Optical scan strategy for CMB polarimetry has been the subject of multiple publications (Delabrouille et al. 2000; Dupac & Tauber 2005; Wallis et al. 2017; Natoli et al. 2017). We consider a relatively well-studied satellite scan strategy for L2-observations where the boresight angle, β , and precession angle, α sum up to approximately 90 degrees (not to be confused with the HWP and boresight rotation angles). The strategy achieves a high degree of cross linking across the entire sky in a full year while maintaining a sun-avoidance angle of greater than 88° at all times. Using the nomenclature established in Wallis et al. (2017) we set $\alpha = 45^{\circ}$ and $\beta = 47^{\circ}$, with spin and precession periods of $T_{\rm spin} = 1 \min$ and $T_{\text{prec}} = 100 \text{ min}$, respectively. The scanning strategy and the proposed baffling solution ensure that the primary lens is only illuminated by the sun at glancing angles, if at all. These quantities are summarized in Table 2.

Figure 8 shows a map of integration time on the sky and condition number of the $A^{\dagger}A$ matrix per pixel (see Sec. 2.5) for the scan strategy used in this analysis. This scan strategy uses the 200 bolometer pairs highlighted in Figure 5, with each pixel corresponding to an orthogonally polarized detector pair. The scan strategy is implemented through **beamconv** interfacing with the publicly available **qpoint** code. We also use **qpoint** for the simple binning map-making implemented throughout this paper. The scanning strategy, de-

 Table 2. Satellite scanning parameters and overall beam properties used in these simulations.

Satellite scanning properties

Orbit	L2
Simulated duration	365 days
Sampling frequency, f_{samp}	$96.73~\mathrm{Hz}$
Spin period, $T_{\rm spin}$	60 s
Precision period, T_{prec}	6000 s
Primary aperture diameter	38 cm

Band properties					
Channel [GHz]	Count	Beam width [arcmin]	Ellipticity		
90 150	400 400	30.4 ± 0.2 18.9 ± 0.1	0.007 ± 0.005 0.006 ± 0.003		

tector counts, and sampling frequency result in a relatively Gaussian distribution of condition numbers with an average condition number of $p_{\rm cond} = 2.24 \pm 0.15$ for an $N_{\rm side} = 512$ map.¹⁶ The full simulation results in 1.2×10^{12} samples per frequency band; each map pixel is therefore visited 3.8×10^5 times in the limit of uniform coverage. See Wallis et al. (2017) for estimates of the expected suppression of leakage per pixel from azimuthally asymmetric beam components for this class of scan strategies.

Optical systematics are tightly coupled to scan strategies, field of view, and sampling frequencies. Therefore, the analysis results presented in Section 5 are only meant to

¹⁶ The map-maker explicitly solves for $\{\hat{l}, \hat{Q}, \hat{U}\}$ instead of $\{\hat{l}, \hat{P}, \hat{P}\}$ but is equivalent to the one described in Sec. 2.5. It uses $A = (A_{t,x})^{\mu} \propto (1, \cos(2\lambda_t), \sin(2\lambda_t), 0) \mathbf{1}_X(t)$ with $\lambda_t = \psi_t + 2\alpha_t + \gamma$ in terms of the position angle ψ_t , HWP angle α_t and constant detector polarization angle γ . The condition number of the (perpixel) $A^{\dagger}A$ matrix is defined as the ratio of its largest and smallest singular value and thus has a minimum value of 2.

provide qualitative insight. Any real experiment collaboration would have to perform simulations more appropriate for their design.

5 RESULTS

Before comparing the results of different beam scanning simulations, we need to standardize the calibration procedures for the different experimental realizations. Here we choose to calibrate the simulations on degree scale temperature anisotropies. Using a temperature power spectrum estimate of the best fit Gaussian beam model for reference, $\hat{C}_{\ell,\text{ref}}^{II}$, we find the best fit scaling parameter, c, which minimizes

$$\sum_{\ell=\ell_1}^{\ell_2} R_\ell \equiv \sum_{\ell=\ell_1}^{\ell_2} \left(c \frac{\hat{C}_\ell^{II}}{\hat{C}_{\ell,\text{ref}}^{II}} - 1 \right),\tag{31}$$

where $\ell_1 = 100$ and $\ell_2 = 300$ and \hat{C}_{ℓ}^{II} is a temperature power spectrum calculated using maps obtained from scanning the sky with a more involved beam model. Unless the beam models differ significantly from a Gaussian model, this calibration procedure usually results in a number $c \simeq 1$. This scaling factor is then applied to all map products and subsequently propagates to all *I*-, *E*-, and *B*-mode power spectra. The power spectrum estimates are calculated using the Pol-Spice estimator (Chon et al. 2004)¹⁷. We use uniform pixel weighting, except for the Galactic mask used in Sec. 5.4.

A future satellite experiment will likely obtain its absolute calibration on the orbital dipole or cross-calibrate to the temperature anisotropies of past experiments such as *Planck*. The assumed beam model then defines the relative sensitivity of the experiment as a function of angular scale with any error in the beam model propagating to error in the measured power spectra. By choosing degree angular scales ($\ell = 100-300$) for our calibration range, we minimize the impact of beam modeling error on the overall amplitude of our derived power spectra.

In the following sections, we study the impact of assuming that the satellite beam model is correctly described by a ensemble average best-fit Gaussian model. We will find that as we increase the complexity of our beam model, the error relative to the Gaussian assumption will grow. Of course, future satellite CMB experiments will calibrate their beam models using point source observations such as planets just as *Planck* and *WMAP* have done in the past (Weiland et al. 2011; Planck Collaboration 2016a, 2017). The beam models constructed in this way will then most likely be supplemented with model predictions from pre-flight measurements as well as ray tracing and physical optics simulations. Any beam model error will then be relative to this more realistic model. In this regard, the comparison conducted in the following sections can be considered that of a worst-case scenario.

5.1 Gaussian, elliptical Gaussian, and full physical optics beam model

Simple on-axis optics and high Strehl numbers result in a relatively symmetric beam response across the focal plane



Figure 10. Binned 90- and 150-GHz *B*-mode power spectra residuals generated from rescanning the (B = 0) Λ CDM map. The two frequencies, 90 and 150 GHz, are shown with blue and red colors, respectively. Three curves, EG, PO, and PO+G correspond to an elliptical Gaussian beam model (dotted curve), a physical optics beam model (solid curve), and a physical optics beam model (solid curve), and a physical optics beam model that includes a ghosting beam response (dashed curve). The larger beam asymmetry at 150 GHz results in a correspondingly larger residual compared to a Gaussian beam model. The grey lines correspond to a primordial *B*-mode power spectrum with a range of tensor-to-scalar ratios while the black solid lines show combination of the primordial and lensing *B*-mode power spectra. Note that the PO and PO+G curves coincide. All power spectra have been deconvolved with an ensemble-average Gaussian beam window function.



Figure 11. Binned 150-GHz *B*-mode power spectra residuals generated from scanning a map formed by combining synfast output with a dust foreground map from Commander (see Section 5.4). The solid blue and orange lines correspond to a simulation with and without a continuously rotating HWP, respectively. Low Galactic latitudes are masked using a mask provided by the Planck Collaboration (see Section 5.4). The incomplete sky coverage causes oscillations in the PolSpice spectra that are manifest in more ragged power spectra (solid lines). Expected geometric $E \rightarrow B$ leakage due to the mask (not shown) is subdominant over the depicted range in multipole. All power spectra have been deconvolved with an ensemble-average Gaussian beam window function.

¹⁷ http://www2.iap.fr/users/hivon/software/PolSpice/



Figure 12. Binned 150-GHz B-mode power spectra residuals obtained by scanning the (B = 0) ACDM map with silicon lenses of different index of refraction (see Section 5.5). The optics design assumed the $n_{\rm si}=3.42$ (solid blue curve), but we explore a scenario where the effective index is $\tilde{n}_{si} = 3.39$ instead (dashed green curve). Note that the solid blue curve here corresponds to the solid red curve in Figure 10. The dotted red curve corresponds to a scenario where a continuously rotating HWP is placed in front of the primary lens (see Section 5.3). The dash-dotted purple line corresponds to a scenario (without a HWP) where a randomly selected 15% subset of the bolometers are not operating (dead). Finally, the dashed golden line corresponds to a scenario where only a single detector per polarization pair is used for analysis, resulting in maximal $I \to P$ leakage. All power spectra have been deconvolved with an ensemble-average Gaussian beam window function.

(see Figure 4). However, an elliptical Gaussian and a full 2D beam model will both capture the spatial response of this experiment more accurately than a simple Gaussian model. Using beamconv simulations, we can explore the impact of a simple Gaussian assumption at the map level. Before making this comparison, we intercalibrate the different maps using degree-scale temperature cross-correlation (see the discussion around Eq. 31). This mimics a scenario where the output of these different experiments are calibrated on degree scale power, for example if this experiment were calibrated against the degree-scale power in *Planck* temperature maps instead of deriving absolute calibration on the orbital dipole signal. Absolute calibration on degree scale power likely represents an optimal scenario for experiments focusing on primordial B-modes, as this links the calibration to the angular scales of interest.

Figure 9 shows the map-level differences (in μ K) between these different beam model assumptions. It is clear that the physical optics beam model deviates from a Gaussian assumption much more significantly than a simple elliptical Gaussian model (note the different scales on the two color bars). This suggests that an elliptical Gaussian approximation does not fully capture the beam-induced nonidealities in the case of this proposed experiment. It is likely that other experimental setups would produce qualitatively similar outcomes. The corresponding impact on *B*mode power spectra is shown in Figure 10. Both the elliptical Gaussian and physical optics beams, when combined with the proposed scan strategy and detector pair matching, create negligible *B*-mode residuals. We also see that the physical optics systematic residual is much greater than the corresponding elliptical Gaussian model residual. This can be attributed to $E \rightarrow B$ leakage from the physical optics beams' increased azimuthal asymmetry and their nonzero cross-polar response.

5.2 Ghosting beams

The impact of ghosting beams depends strongly on their amplitude, polarization fraction, and variation across the focal plane. For this publication, we ran a single case where the primary beams are mirrored to diagonally opposite locations on the focal plane with 1% of the amplitude of the original. In order to add a level of realism, we allow the amplitude of the ghosting beam relative to the main beam to vary by a small amount (relative to 1%) for every detector. The addition of the ghosting beam does not noticeably increase the beam non-ideality systematic observed in maps and power spectra (see Figure 10). We expect this simulation to provide a best-case scenario since the shape and amplitude of the ghosting beam for every detector is roughly identical, allowing for significant cancellation through symmetry. In reality, we expect ghosting beams to vary significantly across the focal plane, but the study of that effect is beyond the scope of this paper.

5.3 HWP modulation

A rotating half-wave plate reduces susceptibility to lowfrequency detector noise which can bias polarization analysis. It also breaks the degeneracy between spurious signal due to azimuthally asymmetric beams and the linear polarized sky signal and therefore reduces potential $I \rightarrow P$ leakage (see Section 2.4.2). As a result, some proposed designs for CMB polarimeters include such a device (Matsumura et al. 2016). We run simulations with and without an ideal skyward HWP to see how much the B-mode power spectrum is impacted. Figure 12 shows the B-mode power spectrum differences for the physical optics beam model scanning the sky with and without an ideal HWP spinning continuously at 1 Hz (compare solid blue line with dotted red line). It can be seen that HWP modulation reduces the B-mode residual. Due to the lack of $I \rightarrow P$ leakage, this reduction in power can be attributed to a reduction in $E \rightarrow B$ leakage due to the nearly perfect decorrelation of Q and U in the $A^{\dagger}A$ matrix from the angular information added by the HWP modulation. The HWP leaves $E \rightarrow B$ leakage due to the cross-polar and asymmetric beam components unchanged, explaining the remaining residual.

All simulations discussed so far have included perfectly orthogonal polarized detector pairs with identical beams for every physical pixel and thus have no contribution from $I \rightarrow P$ leakage due to e.g. asymmetric \tilde{I} beams. In reality, photolithographed bolometer focal planes suffer from sporadic failures and detector malfunctions. These detectors are labelled as dead in the low level analysis and ignored. To investigate the effect of non-uniform polarization coverage we randomly suppress 15% of the 400 detectors that are included in the nominal simulations. This breaks some of the polarization pairs. The corresponding *B*-mode polarization residual is represented by the purple dot-dashed line



Figure 13. Left: Stokes Q difference map (in units of μ K) at 150 GHz as obtained using a wide beam model scanning over a synfast generated map. Right: Same as the left panel, but this time the input map is formed by combining the synfast generated map with a dust template. Note that a difference map with CMB-only map results in even larger residuals. The Planck 60% mask is included to emphasize residuals away from the Galactic plane. An apodized version of this same mask is used in the calculation of polarization power spectra. It is clear from comparing the two panels that the significant power from the Galactic plane is picked up by the extended beam model.

in Figure 12. This loss of pair uniformity clearly increases the residual significantly over the nominal case where all detectors are paired. Taking this to the extreme, we also run simulations with only one detector in each pair, such that all detectors have their polarization sensitivity aligned. In this case, assuming no spinning HWP, suppression only comes from the angular coverage due to the scan strategy and the experiment is faced with maximal $I \rightarrow P$ leakage (see dashed yellow line).

5.4 Adding sidelobes

Extended sidelobes pick up radiation far from the detector beam centroids. If these sidelobes are sufficiently strong and/or partially polarized this can lead to a significant systematic for CMB polarimeters.

It is particularly interesting to compare the B-mode residual for a sidelobe beam model with and without a bright Galactic foreground. Figure 11 shows the *B*-mode spectra for these two cases. Using a mask that permits 60% of the sky (see Figure 13), the addition of a Galactic dust component significantly increases the level of B-mode residual to the point where it is comparable in amplitude to an r = 0.003primordial B-mode power at degree angular scales. Obviously, this result depends strongly on the mask and adopted sidelobe model; for this analysis we use a standard apodized mask provided by the Planck Collaboration.¹⁸ However, we argue that both the mask and the sidelobe model represent realistic scenarios that could apply to future satellite missions. We further note that the addition of a continuously rotating HWP does not ameliorate this systematic since it is not driven by instrument $I \rightarrow P$ leakage.

Figure 13 shows the 150-GHz Stokes Q difference map (relative to simulation input) for this extended beam model. The beam sidelobe generates large scale residuals regardless of the input; however, the addition of Galactic foregrounds dramatically increases the amplitude of these modes. Because of their complex morphology, Galactic foregrounds coupling to sidelobes will source significant B-mode power. We also note that the sidelobes shown in Figure 7 are assumed to be completely unpolarized; they are therefore combined with the main beams of both detectors in a polarized pair in an identical fashion. This assumption is incompatible with the fact that sidelobes generated through reflections or diffraction on sharp edges are likely partially polarized. In that sense, assuming identical sidelobes for both detectors in the pair represents an optimistic scenario.

5.5 Index of refraction tolerancing

Optical tolerancing typically involves variational analysis that incorporates error in dimensions, locations, and physical properties of different optical components (Page et al. 2003; Niemack 2016; Parshley et al. 2018). Such errors will lead to an overall deterioration of optical performance, including defocusing and a reduction in optical Strehl ratios. As a simple proof of concept, we chose to study the impact of incorrectly estimating the silicon index of refraction during the design process. We could have just as easily considered changes in the overall shape of one of the lenses, for example due to machining error or unexpected thermal contractions. The principal challenge however, is understanding how design and modeling errors propagate to beam asymmetries and subsequently to primordial B-mode residuals.

Figure 4 shows the Strehl ratio of the proposed optics design should the effective silicon index of refraction be $\tilde{n}_{\rm si}=3.39$ instead of $n_{\rm si}=3.42$ (see dashed lines). An error of this magnitude is unlikely, but it is instructive to see how the increased beam asymmetry might impact cosmological analysis. Figure 12 compares the \hat{C}^{BB}_{ℓ} spectra obtained by scanning the sky with physical optics beams obtained using these two different refractive indices at 150 GHz; the impact is less significant at 90 GHz. For these simulations we have implemented the nominal observation strategy which uses 200 bolometer pairs to scan the sky at 96.73 Hz sampling frequency over a duration of one year (see Table 2). From comparing the dashed green curve to the blue curve, we find that the relative importance of such a large modeling error is quite small. This suggests that strong reliance on Strehl ratios as an optical performance metric for 90- and 150-GHz frequency bands might not always be warranted.

 $^{^{18}}$ HFI_Mask_GalPlane-apo2_2048_R2.00.fits

In particular, there is a relatively small difference between designs with an average Strehl ratio of 0.94 and 0.99 across the entire field of view. Obviously, design choices that cause a 0.05 shift in Strehl ratios at 150 GHz will have an even stronger impact on higher frequencies such as the usual 220-and 270-GHz bands. However, the optical systematics associated with beam asymmetries at those frequencies will be pushed to larger multipoles. In the case of this mock *B*-mode experiment, the primary science goal would not necessarily be affected.

6 CONCLUSIONS

We have presented and explained the workings of a lightweight, publicly available Python code library capable of generating realistic, full-sky beam-convolved signal timelines that can be subsequently integrated into higher-level simulation pipelines or directly processed into maps and power spectra. The code can be used to inform the design of new CMB experiments and characterise existing ones.

As a proof of concept, we study optical systematics associated with a mock satellite experiment designed to study CMB polarization on degree angular scales. As part of this process, we generate realistic estimates for the beam response of the satellite's two-lens refracting telescope that demonstrate relatively large deviations away from the ubiquitous (elliptical) Gaussian beam parameterisation. In order to focus on several unexplored types of systematic effects, we null the dominant causes for temperature-to-polarization leakage (differential pointing and beam asymmetry). We then explore the remaining E-mode-to-B-mode leakage due to the cross-polar beam components and azimuthal asymmetry of the linearly polarized beam response. The results indicate that the induced spurious signal is well under control for this setup, even when a deliberate error is introduced in the effective index of refraction of the lenses. We note that none of these systematics are negated by the use of a spinning half-wave plate skyward of the primary lens. Similarly, the addition of ghosting sidelobes due to internal reflections has little effect on the performance of the instrument, but we argue that a more involved study would be appropriate. The results also highlight the relevance of polarization systematics induced by sidelobe coupling to polarized foregrounds near the Galactic plane. Finally, we quantify the impact of temperature-to-polarization leakage from reduced detector pair symmetry and demonstrate its problematic nature for a setup without a spinning half-wave plate.

We would like to stress that although these results quantify the amplitude of some optical systematic effects for this setup, general statements are hard to make due to non-trivial dependence on scan strategy and optical design. Dedicated simulations are clearly needed for each experimental setup. Furthermore, our simulations do not include the effect of telescope components such as filters or baffles, nor do they go beyond the ideal half-wave plate parameterisation. Many of these effects can already be included as input to the presented convolution algorithm, but require a more advanced understanding of e.g. material properties in optical simulations. However, including the effects of non-ideal half-wave plates and their interaction with skyward optical components also requires further development of the convolution algorithm itself. We hope to address some of these questions in future work.

We note that the presented code library is not just capable of simulating systematic effects for B-mode power spectrum studies. In fact, the method should be especially useful in studies that rely on higher order statistics of the data or studies that directly work with the full dataset (sky maps, or even the time-ordered data). We argue that for these sort of studies, forward propagating the optical systematic effects into simulated time-ordered data is the most complete and natural approach to include such effects in the analysis. As an example, it would be interesting and timely to investigate how realistic beam effects influence upcoming studies into the CMB polarization field on small-scales over large patches of sky, e.g. for lensing estimation. We suggest a path toward efficiently simulating the full convolution operation for such data and plan to explore such questions in future work.

ACKNOWLEDGEMENTS

We are grateful to Katherine Freese and Martina Gerbino for helpful comments. AJD and JEG acknowledge support by Vetenskapsrådet (Swedish Research Council) through contract No. 638-2013-8993 and the Oskar Klein Centre for Cosmoparticle Physics. JEG acknowledges support from the Swedish National Space Agency (Rymdstyrelsen). Some computations have been performed at the Owl Cluster funded by the University of Oslo and the Research Council of Norway through grant 250672. Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005) package.

REFERENCES

- Abazajian K. N., et al., 2015, Astroparticle Physics, 63, 66
- Abitbol M. H., et al., 2017, preprint, (arXiv:1706.02464)
- Aikin R. W., et al., 2010, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V. p. 77410V, doi:10.1117/12.857868
- Anderson D. G. M., Barakat R., 1994, J. Opt. Soc. Am. A, 11, 2305
- Armitage-Caplan C., Wandelt B. D., 2009, ApJS, 181, 533
- BICEP2 Collaboration 2015, ApJ, 814, 110
- BICEP2 and Keck Array Collaborations 2015, ApJ, 811, 126
- Bock J., et al., 2009, preprint, (arXiv:0906.1188)
- Brown M. L., Challinor A., North C. E., Johnson B. R., O'Dea D., Sutton D., 2009, MNRAS, 397, 634
- Bryan S. A., Montroy T. E., Ruhl J. E., 2010, Appl. Opt., 49, 6313
- Bryan S., et al., 2018, Journal of Low Temperature Physics,
- Buzzelli A., de Bernardis P., Masi S., Vittorio N., de Gasperis G., 2018, A&A, 609, A52
- Challinor A., Fosalba P., Mortlock D., Ashdown M., Wandelt B., Górski K., 2000, Phys. Rev. D, 62, 123002
- Chon G., Challinor A., Prunet S., Hivon E., Szapudi I., 2004, MNRAS, 350, 914
- Cloude S., 1986, Optik, 75, 26
- Coughlin K. P., McMahon J. J., Crowley K. T., Koopman B. J., Miller K. H., Simon S. M., Wollack E. J., 2018, Journal of Low Temperature Physics,
- Datta R., et al., 2013, Appl. Opt., 52, 8747
- Defrance F., et al., 2018, preprint, (arXiv:1803.05168)

- Delabrouille J., Puget J.-L., Lamarre J.-M., Gispert R., 2000, Astrophysical Letters and Communications, 37, 259
- Delabrouille J., et al., 2018, J. Cosmology Astropart. Phys., 4, 014
- Dupac X., Tauber J., 2005, A&A, 430, 363
- Elsner F., Wandelt B. D., 2014, A&A, 561, A88 Eriksen H. K., Jewell J. B., Dickinson C., Banday A. J., Górski
- K. M., Lawrence C. R., 2008, ApJ, 676, 10
- Essinger-Hileman T., et al., 2016, Review of Scientific Instruments, 87, 094503
- Fowler J. W., et al., 2007, Appl. Opt., 46, 3444
- GRASP User's Manual 2018, Version 10.6.0
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
- Hamilton W. R., 1866, Elements of Quaternions. London, Longmans, Green & Co.
- Hanany S., Rosenkranz P., 2003, New Astron. Rev., 47, 1159
- Hanany S., Niemack M. D., Page L., 2013, CMB Telescopes and Optical Systems. p. 431, doi:10.1007/978-94-007-5621-2_10
- Hanson D., Lewis A., Challinor A., 2010, Phys. Rev. D, 81, 103003 Hincks A. D., et al., 2010, ApJS, 191, 423
- Hivon E., Mottet S., Ponthieu N., 2017, A&A, 598, A25
- Hu W., Hedman M. M., Zaldarriaga M., 2003, Phys. Rev. D, 67, 043004
- Jones W. C., et al., 2007, A&A, 470, 771
- Kamionkowski M., Kosowsky A., Stebbins A., 1997, Phys. Rev. D, 55, 7368
- Karkare K. S., et al., 2016, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII. p. 991430 (arXiv:1607.04567), doi:10.1117/12.2231747
- Keihänen E., Reinecke M., 2012, A&A, 548, A110
- King S., Lubin P., 2016, Phys. Rev. D, 94, 023501
- Koopman B., et al., 2016, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII. p. 99142T (arXiv:1607.01825), doi:10.1117/12.2231912
- Ludwig A. C., 1973, IEEE Transactions on Antennas and Propagation, 21, 116
- MacTavish C. J., et al., 2008, ApJ, 689, 655
- Matsumura T., et al., 2016, Journal of Low Temperature Physics, 184, 824
- Mitra S., Sengupta A. S., Souradeep T., 2004, Phys. Rev. D, 70, 103002
- Mitra S., Rocha G., Górski K. M., Huffenberger K. M., Eriksen H. K., Ashdown M. A. J., Lawrence C. R., 2011, ApJS, 193, 5
- Nadolski A., et al., 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 1070843 (arXiv:1809.00030), doi:10.1117/12.2315674
- Nagy J. M., et al., 2017, ApJ, 844, 151
- Natoli P., Marinucci D., Cabella P., de Gasperis G., Vittorio N., 2002, A&A, 383, 1100
- Natoli P., et al., 2017, preprint, (arXiv:1707.04224)
- Niemack M. D., 2016, Appl. Opt., 55, 1686
- Niemack M. D., et al., 2010, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V. p. 77411S (arXiv:1006.5049), doi:10.1117/12.857464
- O'Dea D., Challinor A., Johnson B. R., 2007, MNRAS, 376, 1767
- Page L., et al., 2003, The Astrophysical Journal, 585, 566
- Parshley S. C., et al., 2018, The optical design of the six-meter CCAT-prime and Simons Observatory telescopes, doi:10.1117/12.2314073, https://doi.org/10.1117/ 12.2314073
- Planck Collaboration 2016a, A&A, 594, A7
- Planck Collaboration 2016b, A&A, 594, A9
- Planck Collaboration 2017, A&A, 607, A122
- Prézeau G., Reinecke M., 2010, ApJS, 190, 267
- Prunet S., et al., 2001, ArXiv Astrophysics e-prints,
- Reinecke M., Seljebotn D. S., 2013, A&A, 554, A112

- Reinecke M., Dolag K., Hell R., Bartelmann M., Enßlin T. A., 2006, A&A, 445, 373
- Rosset C., et al., 2010, A&A, 520, A13
- Ruhl J., et al., 2004, in Bradford C. M., et al., eds, Proc. SPIEVol. 5498, Z-Spec: a broadband millimeter-wave grating spectrometer: design, construction, and first cryogenic measurements. pp 11–29 (arXiv:astro-ph/0411122), doi:10.1117/12.552473
- Savini G., Pisano G., Ade P. A. R., 2006, Appl. Opt., 45, 8907
- Sekera Z., 1966, J. Opt. Soc. Am., 56, 1732
- Souradeep T., Ratra B., 2001, ApJ, 560, 28
- Suzuki A., et al., 2018, preprint, (arXiv:1801.06987)
- Takahashi Y. D., et al., 2010, ApJ, 711, 1141
- Taylor J. F., Ashdown M. A. J., Hobson M. P., 2008, MNRAS, 389, 1284
- The Simons Observatory Collaboration 2018, preprint, (arXiv:1808.07445)
- Wallis C. G. R., Brown M. L., Battye R. A., Pisano G., Lamagna L., 2014, MNRAS, 442, 1963
- Wallis C. G. R., Brown M. L., Battye R. A., Delabrouille J., 2017, MNRAS, 466, 425
- Wandelt B. D., Górski K. M., 2001, Phys. Rev. D, 63, 123002
- Wandelt B. D., Larson D. L., Lakshminarayanan A., 2004, Phys. Rev. D, 70, 083511
- Wehus I. K., Ackerman L., Eriksen H. K., Groeneboom N. E., 2009, ApJ, 707, 343
- Wehus I. K., Næss S. K., Eriksen H. K., 2012, ApJS, 199, 15
- Weiland J. L., et al., 2011, ApJS, 192, 19
- Young K., Wen Q., Hanany S., Imada H., Koch J., Matsumura T., Suttmann O., Schütz V., 2017, Journal of Applied Physics, 121, 213103
- Young K., et al., 2018, preprint, (arXiv:1808.01369)
- Zaldarriaga M., Seljak U., 1997, Phys. Rev. D, 55, 1830

APPENDIX A: HARMONIC REPRESENTATIONS

A1 Spin weighted spherical harmonic decomposition

Given generic $\{I, P, V\}$ fields on the sphere, we define the corresponding (spin-weighted) spherical harmonic (SWSH) coefficients as follows (see e.g. Zaldarriaga & Seljak (1997)):

$$c_{\ell m}^{I} = \int_{S^2} \mathrm{d}x \, I(x) \,\overline{Y}_{\ell m}(x), \tag{A1}$$

$${}_{2}c^{P}_{\ell m} = \int_{S^{2}} \mathrm{d}x \, P(x) \, {}_{2}\overline{Y}_{\ell m}(x), \tag{A2}$$

$$c_{\ell m}^{V} = \int_{S^2} \mathrm{d}x \, V(x) \, \overline{Y}_{\ell m}(x), \tag{A3}$$

Due to the reality of the Stokes parameters, the *I* and *V* coefficients obey $\overline{c_{\ell m}^{I/V}} = c_{\ell-m}^{I/V} (-1)^m$ while the *P* coefficients obey $\overline{2c_{\ell m}^P} = {}_{-2}c_{\ell-m}^{\overline{P}} (-1)^m$. The ${}_{-2}c_{\ell m}^{\overline{P}}$ coefficients are the SWSH coefficients of the complex conjugate of *P*:

$${}_{-2}c^{P}_{\ell m} = \int_{S^2} \mathrm{d}x \,\overline{P}(x) \,{}_{-2}\overline{Y}_{\ell m}(x) \,. \tag{A4}$$

Note that, for brevity, we will write $_{-2}c^P_{\ell m}$ instead of the more correct $_{-2}c^{\overline{P}}_{\ell m}$.

The corresponding forward transforms are given by:

$$I(x) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} c_{\ell m}^{I} Y_{\ell m}(x),$$
(A5)

$$P(x) = \sum_{\ell=2}^{t_{\text{max}}} \sum_{m=-\ell}^{t} {}_{2}c_{\ell m}^{P} {}_{2}Y_{\ell m}(x),$$
(A6)

$$\overline{P}(x) = \sum_{\ell=2}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} {}_{-2}c_{\ell m}^{P} {}_{-2}Y_{\ell m}(x),$$
(A7)

$$V(x) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} c_{\ell m}^{V} Y_{\ell m}(x),$$
(A8)

where ℓ_{\max} denotes the band-limit of the field. Note that in the main body of the text we use $\sum_{\ell,m}$ instead of $\sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell}$.

In terms of Euler angles (ψ, θ, ϕ) , we have $dx = \sin \theta \, d\theta d\phi$ and express the spin-weighted spherical harmonics in terms of the Wigner D^{ℓ} matrices as:

$$D^{\ell}_{-ms}(\phi,\theta,\psi) = (-1)^m \sqrt{\frac{4\pi}{2\ell+1}} {}_{s}Y_{\ell m}(\theta,\phi) e^{-is\psi} .$$
(A9)

where the coefficients of the D^{ℓ} matrices are expressed in terms of the real Wigner- d^{ℓ} matrices as:

$$D_{ms}^{\ell}(\phi,\theta,\psi) = e^{-im\phi} d_{ms}^{\ell}(\theta) e^{-is\psi} .$$
(A10)

A2 Converting from Ludwig-III to spherical coordinates

We may convert the Stokes parameters defined on the Ludwig-III basis to those on the (θ, ϕ) basis as follows:

$$\widetilde{I} = \widetilde{I}_{\mathcal{L}}, \tag{A11}$$

$$\widetilde{P} = \widetilde{P}_{\mathcal{L}} e^{-2i\phi}, \qquad (A12)$$

$$\widetilde{V} = -\widetilde{V}_{\mathcal{L}} \,. \tag{A13}$$

We have found that using the above relations to convert beam maps defined on the Ludwig-III basis to the (θ, ϕ) basis leads to inaccurate harmonic modes of the \tilde{P} beam. This is ultimately due to the incomplete description of spinweighted fields on the sphere.¹⁹ A workaround for this issue is obtained by first calculating the spin-0 spherical harmonic coefficients of the $\tilde{P}_{\mathcal{L}}$ (a well-defined operation at the pole) field and use an analytic expression for the spin-±2 spherical harmonic coefficients of the transformation factor $e^{\pm 2i\phi}$ (see (Hivon et al. 2017)). By doing so, we may rewrite the above relation for \tilde{P} in the harmonic domain:

$${}_{\pm 2}b_{\ell m}^{\widetilde{P}} = \sum_{l'} b_{\ell'(m\pm 2)}^{\widetilde{P}_{\mathcal{L}}} K_{\ell\ell'm}, \qquad (A14)$$

¹⁹ Recall that a spin-2 field, such as \widetilde{P} , defined on the tangent space T_x with $x \in S^2$ picks up a factor $e^{-2i\psi}$ under a rotation of the frame through an angle ψ about x. Describing this using Euler angles at the pole leads to counterintuitive results, as ψ and ϕ become degenerate. A more complete description of the fields we consider would be as functions on the rotation group SO(3), parameterised in terms of unit quaternions (which unlike the Euler angles provide an injective mapping to SO(3) at $\theta = 0$). with spin-0 coefficients given by:

$$b_{\ell'm}^{\widetilde{P}_{\mathcal{L}}} = \int_{S^2} \mathrm{d}x \, \widetilde{P}_{\mathcal{L}}(x) \, \overline{Y}_{\ell m}(x), \tag{A15}$$

and the kernel in terms of Wigner-3j symbols:

$$K_{\ell\ell'm} = \frac{2}{\sqrt{\pi}} \sum_{\ell'' \ge 2} \frac{\sqrt{2\ell'' + 1}}{\ell''(\ell'' + 1)} (-1)^{\ell''} I_{\pm 20\mp 2}^{\ell\ell'\ell''} \times \begin{pmatrix} \ell & \ell' & \ell'' \\ -m & m \pm 2 & \mp 2 \end{pmatrix},$$
(A16)

with:

$$I_{\pm 20\mp 2}^{\ell\ell'\ell''} = \sqrt{\frac{(2\ell+1)(2\ell'+1)(2\ell''+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell'' \\ \pm 2 & 0 & \pm 2 \end{pmatrix}.$$
 (A17)

The sum over ℓ'' is formally unbounded. However, the $1/\ell''$ scaling of the factor in front of the 3j symbols suppresses any large deviations of ℓ from ℓ' , which in practise means that $b_{\ell m}^{\tilde{P}_{\ell}}$ and ${}_{\pm 2}b_{\ell m}^{\tilde{P}}$ share band-limits. As explained in Hivon et al. (2017), for a sufficiently

As explained in Hivon et al. (2017), for a sufficiently localised beam, the kernel $K_{\ell\ell'm}$ may be approximated as diagonal per azimuthal mode m, i.e. $K_{\ell\ell'm} \approx \delta_{\ell\ell'} \forall m$. We have used this approximation for the presented analysis.

A3 Azimuthally symmetric beams

Naively demanding that the harmonic modes for the \tilde{P} beam placed on the pole should also be zero for $m \neq 0$ leads to the unphysical result that the beam must vanish at the pole as $\tilde{P}|_{\theta=0} \propto _{2}Y_{\ell m}|_{\theta=0} \propto \delta_{m2}$. As noted before, a way toward a correct expression would be to explicitly represent \tilde{P} as a scalar field on SO(3) using e.g. the unit quaternions, but a simpler approach is to make use of the method presented in Appendix A2.

We start by placing the (localised) azimuthally symmetric \tilde{P} beam on the pole, i.e. centred around the \hat{z} axis. The \tilde{P} beam, with Stokes parameters defined with respect to the (θ, ϕ) coordinate system, is related to the same beam defined on the Ludwig-III basis $(\tilde{P}_{\mathcal{L}})$ through the relation in Eq. A12. We use this relation to rewrite the beam as $\tilde{P}_{\mathcal{L}}e^{-2i\phi}$. We may consider $\tilde{P}_{\mathcal{L}}$ as a spin-0 field, as long as $e^{-2i\phi}$ is considered a spin-2 field (to ensure that the product keeps the correct transformation properties). We have seen that the only nonzero spin-0 harmonic modes of $\tilde{P}_{\mathcal{L}}$ will be those with m = 0 (see Eq. 17). Inserting these coefficients, i.e. $b_{\ell'm}^{\tilde{P}_{\mathcal{L}}} \propto \delta_{m0}$, into Eq. A14 demonstrates that $\pm 2b_{\ell m}^{\tilde{P}} \propto \delta_{m\mp 2}$.

APPENDIX B: ESTIMATING BEAMS WITH OPTICAL SIMULATIONS

As mentioned in the main text (Sec. 2.3.1), optical simulations may be used to estimate the instrumental response when it is completely described by its co- and cross-polar response.²⁰ The instrument is then described as a non-depolarizing transformation (excluding the perfectly depolarizing incoherent detector for now); we may either use the

²⁰ In this regime one is unable to probe the depolarizing properties of the instrument as this would additionally require the response to an unpolarized source.

Jones or Mueller-Jones transformations to describe such a system. For conciseness we pick the Jones formalism.

Following (Rosset et al. 2010), we describe the Jones matrix of the instrument as an imperfect linear polarizer coupled to a generic Jones matrix describing the telescope:

$$J_{\mathcal{L}}(x,\omega) = \begin{pmatrix} J_{11} & J_{12} \\ \sqrt{\eta}J_{21} & \sqrt{\eta}J_{22} \end{pmatrix} (x,\omega),$$
(B1)

where $\eta \to 0$ in case the linear polarizer becomes perfect. J is defined on the Ludwig-III basis, as indicated by the subscript $_{\mathcal{L}}$. The co- and cross-polar responses are due to the J_{11} , $\sqrt{\eta}J_{21}$ and J_{12} , $\sqrt{\eta}J_{22}$ elements respectively. Note that the parameter η is sometimes referred to as the crosspolar leakage. This label might cause confusion as both the co- and cross-polar response depend on η .

In the $\eta = 0$ case, incident co-polar radiation $(\epsilon = \begin{pmatrix} 1 \\ 0 \end{pmatrix})$ probes the J_{11} component while incident cross-polar radiation probes the J_{12} component. Optical simulations like those described in Sec. 4.3, generally work in the reciprocal sense: they simulate the propagation of electric fields emitted from the location of the detector through the optical system out to the sky. As long as the system is reciprocal, the simulation results may be used to estimate the instrument response to radiation with reversed direction of propagation. In terms of Jones matrices, the reciprocal Jones matrix for such a system is given by $J^{\text{recpr.}} = \text{diag}(1, -1)J^T \text{diag}(1, -1)$ (Sekera 1966). This implies that the nonzero elements of the instrumental Jones matrix may be estimated from the farfield response of electric fields with known amplitude and polarization state that propagated through the system. For example, by emitting a purely co-polar field from the detector, the far-field vector components are proportional to the $\begin{pmatrix} J_{11} \\ -J_{12} \end{pmatrix}$ elements from Eq. B1. Probing the J_{21} and J_{22} components in case where $\eta \neq 0$ amounts to repeating the simulation with a purely cross-polar initial field. For the examples presented in this work we have omitted this last step: assuming perfect linear polarizing elements at the end of the optical chain.

With the instrumental Jones matrix estimated, the last steps are to incorporate the depolarizing incoherent detector and to calculate the three fields that describe the instrumental beam: $\{\tilde{I}, \tilde{P}, \tilde{V}\}$. One way to achieve both is to convert the Jones matrix to the associated Mueller-Jones matrix and setting all the elements but those in the top row to zero.²¹ After conversion to the spherical coordinate system, the remaining elements are then equal to the $\{\tilde{I}, \tilde{Q}, \tilde{U}, \tilde{V}\}$ beams in the instrument frame and have harmonic coefficients that can be used as input to Eq. 10.

This paper has been typeset from a $T_{\ensuremath{E}}X/\ensuremath{I\!A}T_{\ensuremath{E}}X$ file prepared by the author.

²¹ The associated Mueller-Jones matrix is related to the Jones matrix by: $M^{\mu}_{\nu} = \frac{1}{2} \text{Tr} \{ \sigma^{\mu} J \sigma_{\nu} J^{\dagger} \}$ with $\sigma_{\mu} = \{ \mathbf{1}, \sigma_{3}, \sigma_{1}, \sigma_{2} \}$ in terms of the Pauli matrices.