

Combining ILC and moment expansion techniques for extracting average-sky signals and CMB anisotropies

Aditya Rotti * and Jens Chluba *

Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

Accepted 2020 October 19. Received 2020 October 17; in original form 2019 November 26

ABSTRACT

The method of weighted addition of multifrequency maps, more commonly referred to as *Internal Linear Combination* (ILC), has been extensively employed in the measurement of cosmic microwave background (CMB) anisotropies and its secondaries along with similar application in 21cm data analysis. Here, we argue and demonstrate that ILC methods can also be applied to data from absolutely calibrated CMB experiments to extract average-sky signals in addition to the conventional CMB anisotropies. The performance of the simple ILC method is, however, limited, but can be significantly improved by adding constraints informed by physics and existing empirical information. In recent work, a moment description has been introduced as a technique of carrying out high precision modelling of foregrounds in the presence of inevitable averaging effects. We combine these two approaches to construct a heavily constrained form of the ILC, dubbed MILC, which can be used to recover tiny monopolar spectral distortion signals in the presence of realistic foregrounds and instrumental noise. This is a first demonstration for measurements of the monopolar and anisotropic spectral distortion signals using ILC and extended moment methods. We also show that CMB anisotropy measurements can be improved, reducing foreground biases and signal uncertainties when using the MILC. While here we focus on CMB spectral distortions, the scope extends to the 21 cm monopole signal and B-mode analysis. We briefly discuss augmentations that need further study to reach the full potential of the method.

Key words: cosmic background radiation – early Universe – inflation – cosmology: observations.

1 INTRODUCTION

The anisotropies of the cosmic microwave background (CMB), beyond doubt, have greatly helped in establishing the standard Λ CDM concordance model, with the key cosmological parameters being known to percent-level precision or better (Bennett et al. 2003; Planck Collaboration XVI 2014b; Planck Collaboration XIII 2016c). To reach this unprecedented precision, many obstacles had to be overcome, including an extreme control of systematic, calibration uncertainties and foreground separation (Planck Collaboration X 2016a; Planck Collaboration XI 2016b; Planck Collaboration XLVI 2016e). To mitigate foregrounds, many independent methods have been developed (e.g. Delabrouille, Cardoso & Patanchon 2003; Tegmark, de Oliveira-Costa & Hamilton 2003; Eriksen et al. 2008; Remazeilles, Delabrouille & Cardoso 2011). One of them is the *Internal Linear Combination* (ILC; Tegmark et al. 2003), which has proven highly successful in extracting foreground-cleaned CMB maps with current (e.g. Planck Collaboration XII 2014a; Planck Collaboration X 2016a; Planck Collaboration XXII 2016d) and future experiments (e.g. Remazeilles et al. 2016, 2018).

In this work, we are mainly interested in studying spectral distortions of the CMB. While the CMB energy spectrum has been shown to be extremely close to that of a blackbody (Mather et al. 1994;

Fixsen et al. 1996), minor deviations from this spectrum are expected even within standard Λ CDM (e.g. Chluba 2016). The largest spectral distortions are present in the monopole sky. These signals are created through energy exchange and photon production in the early phases of cosmic history (Zeldovich & Sunyaev 1969; Sunyaev & Zeldovich 1970a; Illarionov & Sunyaev 1974; Burigana, Danese & de Zotti 1991; Hu & Silk 1993; Sunyaev & Chluba 2009; Chluba & Sunyaev 2012; Chluba 2015). A distortion dipole is furthermore induced due to our motion with respect to the CMB rest-frame, very much like the CMB temperature dipole (Danese & de Zotti 1981; Balashev et al. 2015; Burigana et al. 2018).

Spectral distortion require absolutely calibrated measurements of the CMB spectrum. Building on the heritage of *COBE/FIRAS*, this can be achieved with instrument concepts similar to *PIXIE* (Kogut et al. 2011a, 2016, 2019). This could allow constraining many standard and non-standard processes occurring in the early Universe, at phases inaccessible by any other means (for recent overview, see Chluba et al. 2019). While ILC techniques have been used to measure components with spatial anisotropies, so far they have not been applied to average-sky (i.e. monopolar) signals, a generalization we explore here.

In this paper, we will demonstrate that ILC methods can be directly applied to absolutely calibrated maps, allowing an extraction of the average-sky signal. The simple (i.e. blind) ILC method is, however, limited and has to be augmented by moment expansion methods (Chluba, Hill & Abitbol 2017) to capture the inevitable

* E-mail: adityarotti@gmail.com (AR); jens.chluba@manchester.ac.uk (JC)

foreground averaging effects, yielding the *Moment ILC* method (MILC, see Section 2). Without this extension, large biases and enhanced uncertainties in the recovered signals remain (see Fig. 3). This conclusion extends to anisotropic signals, as we demonstrate here (see Fig. 1).

We focus on a *proof-of-concept* study, demonstrating the main aspects of the MILC and how it can be applied. Further studies are required for providing concrete CMB distortion forecasts, extending first studies (Desjacques et al. 2015; Sathyanarayana Rao et al. 2015; Abitbol et al. 2017), which largely neglected spatial information. These first studies are simplistic and the demonstrations we make in this work will pave the way for more realistic forecasts in the near future. In addition, we can already anticipate that the MILC method can be applied more broadly, e.g. to the extraction of global 21 cm signals (e.g. Pritchard & Loeb 2010) and CMB polarization B modes, as we briefly discuss.

2 METHODS

Before diving into details of the analysis, we provide a brief introduction to the methods we employ in this work. In particular, we present some details of standard ILC methods in Section 2.1 and the preliminaries of the methods of moment modelling of foregrounds in Section 2.2 needed for the MILC.

2.1 Standard ILC methods

ILC methods have been extensively used in the analysis on multifrequency microwave maps originally to extract maps of the CMB temperature and polarization anisotropies (Tegmark et al. 2003; Adam et al. 2016) and more recently for the measurement of y -distortions (Planck Collaboration XXI 2014c; Remazeilles & Chluba 2020). The multifrequency maps, $d_{\nu i}$, can be expressed as,

$$d_{\nu i} = \sum_c s_{\nu}^c \tau_i^c + n_{\nu i}, \quad (1)$$

where ν denotes the observing frequency and index ‘ i ’ denotes the pixel index; τ^c denote spatial map of the component ‘ c ’, which has the spectral energy distribution (SED) s_{ν}^c ; n denotes the measurement noise. Note that this equation is not written in any particular basis and hence the index ‘ i ’ can either be interpreted as the address of a pixel in real or in harmonic space.

Using some prior information on the spectral components of the data, and given that the data is measured with sufficient frequency sampling, we can solve for the map of each component of interest as follows:

$$\hat{\tau}_i^{c_0} = \sum_{\nu} w_{\nu}^{c_0} d_{\nu i} = \mathbf{w}_{c_0}^T \cdot \mathbf{d}. \quad (2)$$

This can be achieved by constructing suitable weights $w_{\nu}^{c_0}$ such that, they have unit response to the SED of the component of interest: $\sum_{\nu} w_{\nu}^{c_0} s_{\nu}^{c_0} = \mathbf{w}_{c_0}^T \cdot \mathbf{s}_{c_0} = 1$. Note that we have introduced the hat notation to distinguish the separated component map $\hat{\tau}^c$ from the true component map τ^c . On using these weights it is easy to see that equation (2) reduces to the following form:

$$\hat{\tau}^{c_0} = \tau^{c_0} + \sum_{c \neq c_0} (\mathbf{w}_c^T \cdot \mathbf{s}^c \tau^c + \mathbf{w}_c^T \mathbf{n}) = \tau^{c_0} + \mathcal{B}^{c_0} + \mathbf{n}^{c_0}, \quad (3)$$

where we have suppressed the spectral and spatial indices for brevity. It is important to note that the solution generally has an additive bias \mathcal{B}^{c_0} , and one can understand this as originating from the existence of non-zero projections of the SED of c_0 on the SEDs of other




Figure 1. This figure depicts the component-separated signal maps returned by the MILC algorithm, for a varying number of SEDs passed to it. The annotations to the left of the image detail the set of SED vectors passed to MILC. These solutions are found for simulations with $50 \text{ Jy sr}^{-1} \text{ px}^{-1}$ noise RMS. The input component maps are shown in the last row.

spectral components of the map. It is also important to note that the excess variance on the reconstructed map has contributions from the bias, the measurement noise and the chance correlation between the two. Therefore, it is important to realize that an assessment on the precision of the reconstructed map necessarily requires a thorough understanding of the interplay between the additive bias and the variance in the map and their relative amplitudes.

One can optimize the solution of the component map by demanding that the weights must minimize the variance of the reconstructed component map, leading to the standard ILC solution (Tegmark et al. 2003). One can further optimize the solution by demanding that the weights in addition to having unit response to s^{c_0} and minimizing the variance, must simultaneously show zero response to the SEDs of selected components in the map. This can be concisely written using a vector matrix formulation as follows:

$$\mathbf{w}_{c_0}^T \cdot [s^{c_0}, s^{c_1}, s^{c_2} \dots s^{c_n}] = [1, 0, 0 \dots 0],$$

$$\Leftrightarrow \mathbf{w}_{c_0}^T \mathbf{V} = \mathbf{e}_{c_0}^T. \quad (4)$$

This generalized minimization problem can be solved using the method of Lagrange multipliers and it can be shown that the weights are given by the following compact expression:

$$\mathbf{w}_{c_0}^T = \mathbf{e}_{c_0}^T [\mathbf{V}^T \mathbf{C}^{-1} \mathbf{V}]^{-1} \mathbf{V}^T \mathbf{C}^{-1}, \quad (5)$$

where \mathbf{C} denotes the data covariance matrix and all other symbols have the same meaning as before. Since these are matrix operators it is important to note that the order in which the elements appear is critical. Equation (5) presents a more general form of the ILC referred to as the constrained ILC (cILC) and was introduced in Remazeilles et al. (2011) and has been used to construct Sunyaev–Zeldovich (SZ) free CMB maps and vice versa. More recently this method has been shown to allow measuring the relativistic electron temperature of galaxy clusters, introducing *semiblind*¹ constraints on the dust SED (Remazeilles & Chluba 2020).

The general solution for the weights given in equation (5) reduces to the simple ILC when $\mathbf{V} = s^{c_0}$ and $e_{c_0} = 1$. The solution written in the form given in equation (5), gives the impression that it is only possible to solve for one component at a time. Presenting the solution in the following form:

$$[\hat{t}^{c_0}, \hat{t}^{c_1}, \hat{t}^{c_2} \dots \hat{t}^{c_n}] = [\mathbf{V}^T \mathbf{C}^{-1} \mathbf{V}] \mathbf{V}^T \mathbf{C}^{-1} \mathbf{d}, \quad (6a)$$

$$\Leftrightarrow \hat{\mathbf{t}} = \mathbf{w}^T \mathbf{d}, \quad (6b)$$

makes it clear that the filter actually returns a vector $\hat{\mathbf{a}}$ of minimum variance maps with mutually orthogonal SEDs, the projection operator being defined with respect to the data covariance matrix.

2.2 Moment modelling of the observed sky

The moment method can be used to efficiently model the effect of averaging (i.e. along the line of sight, in the beam and averaging operations in the data processing) for known fundamental SEDs (Stolyarov et al. 2005; Chluba et al. 2017). The limitations in our ability to model the foregrounds to arbitrary precision is a consequence of the limited sensitivity and frequency coverage of observations and, additionally, our ignorance about the presence of totally ‘new’ foregrounds components, which one may discover

¹Here, semiblind means that partial constraints on the dust have been used, admitting that we only have incomplete knowledge of the true dust SED.

on making more refined measurements.² In this section, we briefly introduce the basic idea and refer the reader to Chluba et al. (2017) for a more detailed and discussions specific to certain foregrounds (dust, synchrotron, etc.) and the broader scope of the method.

Let us assume that we know that the emission from each volume element along a given line of sight is described by an SED $s_\nu(\mathbf{p})$ with some known functional form, parametrized by a vector of parameters: $\mathbf{p} = [p_0, p_1, p_2 \dots p_n]$. Assuming there are multiple elements emitting along a given line of sight, \hat{n} , the net observed intensity from that direction is merely a sum of emission from each of the elements,

$$I_\nu(\hat{n}) = \int A(\hat{n}, l) s_\nu(\mathbf{p}(\hat{n}, l)) dl, \quad (7)$$

where $\mathbf{p}_i(\hat{n})$ denotes that the parameters characterizing (defining the spectral shape) the emission change along \hat{n} and A_l sets the overall amplitude of the emission. It is very likely that the emission from the different elements is characterized by different vectors \mathbf{p} . Given that there are great number of emitting elements, one may, equivalently, construct a statistical model for the observed intensity, in which equation (7) can be re-cast in the following form:

$$I_\nu(\hat{n}) = A_{\nu_0}(\hat{n}) \int s_\nu(\mathbf{p}) \mathcal{P}(\hat{n}, \mathbf{p}) d\mathbf{p}, \quad (8)$$

where $\mathcal{P}(\hat{n}, \mathbf{p})$ denotes the multidimensional probability distribution function of the component of \mathbf{p} , and the parameter $A_{\nu_0}(\hat{n})$ is amplitude of the observed SED at some pivot frequency ν_0 , introduced to allow for modelling the overall amplitude of the SED. In practice, we never have access to this probability distribution, but we have a reasonable idea about the variety of SEDs contributing to the observed sky. One can now Taylor expand the emission law s_ν around some pivot parameter vector $\bar{\mathbf{p}}$,

$$I_\nu(\hat{n}) = A_{\nu_0}(\hat{n}) \int \mathcal{P}(\hat{n}, \mathbf{p}) d\mathbf{p} \left\{ s_\nu(\bar{\mathbf{p}}) + \sum_i \frac{(p_i - \bar{p}_i)}{1!} \frac{\partial s_\nu}{\partial p_i} \Big|_{\bar{\mathbf{p}}} + \sum_{i,j} \frac{(p_i - \bar{p}_i)(p_j - \bar{p}_j)}{2!} \frac{\partial^2 s_\nu}{\partial p_i \partial p_j} \Big|_{\bar{\mathbf{p}}} + \sum_{i,j,k} \frac{(p_i - \bar{p}_i)(p_j - \bar{p}_j)(p_k - \bar{p}_k)}{3!} \frac{\partial^3 s_\nu}{\partial p_i \partial p_j \partial p_k} \Big|_{\bar{\mathbf{p}}} + \dots \right\}. \quad (9)$$

Since the SED and its derivatives are computed at the fixed pivot parameters $\bar{\mathbf{p}}$, these spectral functions are constants and can be pulled out³ of the integral and each of the integrals over the PDF can be understood as the corresponding amplitude weighted moment η of the distribution. This allows us to model the observed intensity in the language of the moments of the parameter distribution,

$$I_\nu(\hat{n}) = A_{\nu_0}(\hat{n}) s_\nu(\bar{\mathbf{p}}) + \sum_i \eta_i(\hat{n}) s_\nu^i(\bar{\mathbf{p}}) + \frac{1}{2!} \sum_{ij} \eta_{ij}(\hat{n}) s_\nu^{ij}(\bar{\mathbf{p}}) + \frac{1}{3!} \sum_{ijk} \eta_{ijk}(\hat{n}) s_\nu^{ijk}(\bar{\mathbf{p}}) + \dots, \quad (10)$$

where η denotes the ‘amplitude-weighted moments’ of the of the parameter distribution characterizing the total SED. Here, it is

²The moments of currently known foreground might span a broader set of SEDs but may not be optimal for these new components. So this can also be cast as large increase in sensitivity and frequency coverage, if we wanted to model new foregrounds into the existing set of SED basis functions.

³We neglect the effect of frequency-dependent beams, which can complicate this step further (Chluba et al. 2017).

important to mention that since the derivative operator with respect to different parameters commute: $s_v^{ij} \equiv s_v^{ji}$, $s_v^{ijk} \equiv s_v^{ikj}$, ..., one needs to appropriately take care of these degenerate SED vectors. One of the beautiful aspect of equation (10) is that the spectral and spatial parts are written in separable form unlike equation (7). A similar language was previously applied to the modelling of the SZ effect, also highlighting this aspect (Chluba et al. 2013). One can now think of these moment maps, $\eta(\hat{n})$, as direct astrophysical observables. The simplest first-order moments inform us about how the parameter along various lines of sight differ from the pivot parameter, while the second-order moments inform us about the variance in the emission characteristics and so on.

In the final step, the generalized foreground modelling presented in equation (10) is cast into a form that can be easily incorporated into the ILC machinery, yielding the MILC approach. By merely introducing the various SED derivatives as spectral constraints (as discussed in Section 2.1) one can then simultaneously solve for the foreground moments maps as well as the cosmological observables like the CMB temperature, y and μ distortion maps as we present in the following sections.

Some of the benefits of this approach for anisotropic signals were recently illustrated for the relativistic SZ temperature mapping (Remazeilles & Chluba 2020) and the extraction of μ - T correlations (Remazeilles & Chluba 2018). In this work, we unveil the real potential of this method, extending it to the extraction of tiny monopolar signals and also demonstrate that introducing many moment constraints can ensure an unbiased recovery of components while maintaining and even improving uncertainties. We note that the standard ILC and constrained ILC are special cases of MILC.

3 SIMULATIONS

In this section, we summarize the details of our simulations. We assume a wide frequency coverage of 30–3000 GHz, with 30 logarithmically spaced channels and assume that the measurements are absolutely calibrated. Our multifrequency simulated maps consists of the following components: blackbody, spectral distortion components y and μ that characterize the deviation from the blackbody, and foreground contamination in the maps and measurement noise. Below we provide specific details of how each of these components are injected in our simulations.

Blackbody: In our simulations we include the CMB monopole and its anisotropies $\Delta_T(\hat{n}) = \Delta T(\hat{n})/T_0$, characterized by the fiducial CMB power spectrum. The CMB SED is given by

$$I_v^{CMB}(\hat{n}) = [\Delta_T^0 + \Delta_T(\hat{n})] \frac{\partial B_v}{\partial T} T \Big|_{T_0}, \quad (11)$$

where B_v denotes the Planck function. We assume that the electromagnetic spectrum of these fluctuations is described by a blackbody at a temperature $T_0 = 2.7255$ K,⁴ set to the currently best-known measurement (Fixsen et al. 1996, 2011). Note that we do not assume the CMB temperature to be known to arbitrary precision and solve for the monopole temperature Δ_T^0 which in our simulations is set to a value of 10^{-4} , about one order of magnitude below the error on the current CMB monopole temperature measurement. A temperature anisotropy map is constructed assuming a power spectrum for the fiducial cosmological model. When analysing the simulations, we assume no prior knowledge about anisotropies of the CMB.

⁴Since we assume that the temperature of 2.7255 K has been subtracted from the data, only Δ_T^0 , the correction to this temperature, is included in our simulations.

y -distortions: These distortions are sourced via inverse Compton scattering of CMB photons off energetic electrons. The dominant contribution to this signal is sourced at low redshifts, by scattering of hot electron gas inside galaxy clusters (Zeldovich & Sunyaev 1969; Mroczkowski et al. 2019). To inject y -distortions in our simulations, we use the Compton y -map of Sehgal et al. (2010) available at LAMBDA simulations. We only use the non-relativistic thermal SZ spectrum for this work, as given by

$$I_v^y(\hat{n}) = y(\hat{n}) \left[x \coth\left(\frac{x}{2}\right) - 4 \right] \frac{\partial B_v}{\partial T} T \Big|_{T_0}, \quad (12)$$

where $x = hv/k_B T$. We ignore the relativistic corrections to the average SZ spectrum (Hill et al. 2015), however, for detailed forecasts of the distortion sensitivities this should be included. The Compton- y parameter field is a positive and non-Gaussian field and hence has non-zero positive monopole. The ensemble averaged y parameter has been predicted to be $\langle y(\hat{n}) \rangle = \bar{y} \simeq 2 \times 10^{-6}$ in the fiducial cosmological model (e.g. Refregier et al. 2000; Hill et al. 2015). On a chosen patch, the sky averaged mean \bar{y} may not necessarily match the ensemble-averaged mean. In our simulations we thus ensure that the sky-averaged value is close to this expected monopole, by directly adding/subtracting a monopole component to the y -map.

μ -distortions: The dissipation of acoustic modes in the early Universe ($z > 50,000$), introduces a μ -distortion in the CMB radiation (Sunyaev & Zeldovich 1970b; Hu & Silk 1993; Chluba, Khatri & Sunyaev 2012). While this signal also has spatial fluctuations, these are usually significantly smaller than the monopole distortions, in complete analogy with the CMB monopole temperature and hundreds of micro Kelvin fluctuations as a function of direction.⁵ In this work, we thus only include the spectral distortion to the monopole spectrum in our simulations, while ignoring spatial fluctuations. The spectrum of the μ -distortions is given by,

$$I_v^\mu = \mu \left[\frac{1}{2.1923} - \frac{1}{x} \right] \frac{\partial B_v}{\partial T} T \Big|_{T_0}, \quad (13)$$

where μ denotes the amplitude of the monopole signal and all the symbols have the same meaning as before. Unlike the y -distortions, for which we have measurements from *Planck*, we only have upper limits on the μ distortions and upcoming experiments will aim at improving these limits and potentially making a detection. For this demonstration study we inject a μ -distortion monopole signal with an amplitude of $\mu = 10^{-6}$, roughly two orders of magnitude below the current best limit set by *COBE/FIRAS* (Fixsen et al. 1996).

Foregrounds: For our demonstration, we considered the following foreground components: synchrotron, free-free and thermal dust in our simulations, although we focus our discussion mainly on dust-only simulations. We generate the foreground skies using the Planck Sky Model (PYSM; Thorne et al. 2017). In particular, we use the ‘d2’ model for thermal dust, ‘s2’ model for the synchrotron, and the nominal free-free model in our simulations. The details of these component models can be found in the original PYSM paper. The skies simulated using PYSM include spectral parameter variations across the sky as well as along different lines of sight, that naturally lead to generation of moments. We emphasize that we do not use any details relating to spectral parameter variations when carrying out the component separation analysis on the simulated maps.

⁵More significant μ -distortion anisotropies can be caused by primordial non-Gaussianity (e.g. Ganc & Komatsu 2012; Pajer & Zaldarriaga 2012; Emami et al. 2015; Ravenni et al. 2017).

Finally, we include some simple instrumental effects in our simulations: beam smoothing and measurement noise. For the instrumental beam we assume a uniform FWHM = 30 arcmin for all the channels. We take care to carry out the smoothing operation after duly converting the anisotropy maps returned by PYSM from Rayleigh–Jeans temperature units to intensity units of Jy sr^{-1} . This small detail is important when we want to properly propagate the moments generated from this smoothing operation.⁶

We assume a constant noise RMS for all the channels, however, we do carry out the analysis for different values⁷: [5000, 500, 50, 5] $\text{Jy sr}^{-1} \text{px}^{-1}$. In our simulations, the pixel size is 13.74 arcmin. We use this information to translate these noise RMS values into all-sky sensitivities, yielding: [5.6, 0.56, 0.056, 0.0056] Jy sr^{-1} , respectively. Note that while 5 Jy sr^{-1} is representative of all-sky *PIXIE* sensitivity (Kogut et al. 2011b, 2016), 0.56 Jy sr^{-1} is representative of *SuperPIXIE* (Kogut et al. 2019) and finally 0.056 Jy sr^{-1} is a little better than the full sky sensitivity for *Voyage 2050* (Chluba et al. 2019). Gains in angular resolution may be conceivable with concepts like *Millimetron*⁸ or by combining traditional CMB imagers with spectrometers (André et al. 2014; Delabrouille et al. 2019). Since we use only a part of the sky, our effective sensitivity on the sky patch is worse than the full sky sensitivity by roughly an order of magnitude. A more detailed study covering more complex sky models (e.g. including AME, CIB and CO emission) and instrumental aspects (such as optimization of angular resolution and spectral parameters) will be carried out elsewhere (Rotti et al., in preparation).

3.1 Analysis strategy

For speedy analysis we evaluate the MILC algorithm on flat sky patches, however, this method is easily extended to the full sky analysis. All our simulations are generated at the HEALPIX (Gorski et al. 2004) resolution of NSIDE = 256. From each of the frequency maps, we then extract sky tiles of dimensions $25^\circ \times 25^\circ$ centred on the galactic coordinate $(\ell, b) = (0^\circ, 30^\circ)$, using the HEALPIX gnomonic projection functionality. We emphasize that identical analyses were carried out on different sky locations and we found qualitatively similar results, however here we choose to present results at this randomly chosen sky location.

Our MILC filters are implemented in Fourier space and specifically the covariance is estimated as,

$$C_{\nu\nu'}^{kbin} = \frac{1}{(2\pi)^2} \int_{kbin} \tilde{d}_\nu(\mathbf{k}) \tilde{d}_{\nu'}(\mathbf{k}) d^2k, \quad (14)$$

where $\tilde{d}_\nu(\mathbf{k})$ denotes the Fourier coefficient of the multifrequency data. The covariance is estimated by integrating over the complete Fourier plane to minimize the ILC bias. Note that we do not remove the monopole term from the Fourier coefficients when making the binned frequency–frequency correlation matrix estimate.⁹ In a full sky analysis, a global monopole subtraction does not guarantee a local monopole subtraction, and we use this as the motivation for

⁶This aspect is not important when translating between antenna temperature to Jy sr^{-1} . However, the conversion between the commonly used μK_{CMB} units and intensity, strictly speaking, is a non-linear operation as it involves the inversion of the *Planck* function and in this case the smoothing and unit conversion operations do not commute, a detail that is potentially important.

⁷This denotes the standard deviation of the measured intensity in each pixel.

⁸www.millimetron.ru/

⁹It is important to bear in mind this detail when comparing the statistics estimated from the recovered component maps to those analytically estimated.

making this analysis choice of not subtracting the monopole when estimating the data covariance. We emphasize that the trends seen in our results do depend on this analysis choice and we will return to discussing these in the conclusions.

When constructing the moment SED vectors, one has to make a choice for the pivot parameters $\bar{\mathbf{p}}$ and pivot frequencies ν_0 . In principle, the solution to the pivot parameter changes for the simulations with different sensitivities and frequency coverage, however, in favour of simplicity we choose the pivot parameter from the highest sensitivity simulations and thereafter keep them fixed. Furthermore, the solution to the pivot parameters also changes when including higher order moment SEDs to fit the sky SED and in principle would require updates, however, this again is not essential. In our analysis, to find reasonable pivot parameters we fit the base SED to the sky-averaged spectrum and we choose the pivot frequency ν_0 such that it coincides with the peak of the dust emission. After this initial step, $\bar{\mathbf{p}}$ is fixed and used to add higher order moments. It is important to note that the construction of moment SEDs does not rely on knowing precisely the parameters that characterize the foregrounds, on the contrary, the moment maps(η) are part of the solution and these inform us about the statistical properties of parameters characterizing the foregrounds.

The multifrequency simulated data, along with a subset of SEDs, are passed to the MILC algorithm, which then returns a set of component-separated maps with a one to one correspondence to the SEDs. At the first stage of the analysis, we progressively increase the number of signal SEDs (i.e. CMB, γ and μ) passed to MILC, which then returns the respective signal component maps. At the second stage, in addition to the signal SEDs we also pass the moment SEDs, progressively increasing the number of foreground vectors that are passed to the algorithm. MILC then returns all the signal maps along with the moment maps. In summary, the very first iteration of MILC only solves for the map of the CMB and at the final iteration the algorithm returns maps of all the signals, $[\Delta_T, \gamma, \mu]$, and multiple foreground moment maps, η . This analysis procedure is repeated on each set of multifrequency simulations, where the only variable is the level of noise in the maps.

4 RESULTS

In this section, we discuss the results of applying the MILC to the simulated data. We restrict the presentation and most of the discussion to results from analysis on simulations that only include thermal dust emission as foregrounds, since our primary aim is to demonstrate that MILC can not only be used to make robust (bias- and noise-reduced) measurements of anisotropic signals but also to extract tiny average-sky (monopolar) signals.

The thermal dust emission is characterized by a modified black-body: $S_\nu^d = \nu^\alpha B_\nu(1/\beta)$, where β is the inverse of the dust temperature. Following equation (10), the moment expansion is performed by taking different order derivatives with respect to the parameters ($\mathbf{p} = [\alpha, \beta]$) characterizing the SED. The SED vectors characterizing the different moments are simply given by $S_\nu^{ij} = \mathcal{N}_{ij} \partial_\alpha^i \partial_\beta^j S_\nu$, where¹⁰ \mathcal{N}_{ij} is some normalization and the moment order is given by: $o = i + j$. The moment maps corresponding to each of these moment SED vectors are denoted by $\eta_{\alpha^i \beta^j}$. Note that $\eta_{\alpha^0 \beta^0}$ refers to the amplitude of dust emission and is denoted by A_D . We include a maximum of 10

¹⁰This normalization only changes the overall amplitude of the moment SED and not its shape and hence it is only important when we care about the amplitude of the moment maps. Given some normalization there is no ambiguity on the amplitude of the moment maps.

moment SED vectors in our analysis which corresponds to maximum moment order of $o = 3$.

Fig. 1 depicts the component-separated signal maps returned at each evaluation of MILC on simulations that include noise at the level of 50 Jy px^{-1} . Recall that our simulations are absolutely calibrated and hence retain information about the monopoles in the respective component maps (unlike experiments like *WMAP* and *Planck*, which carried out differential measurements). To diagnose the measurement of the monopole signal we study the one-point probability distribution function (1PPDF) of the recovered component maps. Fig. 2 depicts the 1PPDF of a subset of these maps. Fig. 3 depicts the evolution of the 1-point and 2-point statistics of maps, for varying number of vectors passed to MILC and for different levels of noise in the simulations. Fig. 4 depicts the first 6 foreground moments (i.e. up to second order for dust) estimated on simulation with the lowest noise. With the help of the visual aid provided by these figures we now discuss the salient features of solutions delivered by different iterations of the MILC.

The first row of Fig. 1 depicts a map of the recovered CMB, and this iteration of the MILC corresponds to the standard ILC, which only projects out the CMB and is completely blind to the presence of other components in the map, subsuming these into the variance minimization process. Comparing to the true CMB anisotropy map (right-hand column of Fig. 1) one notes that many of the features in the true CMB map are indeed recovered, however, there are residuals, particularly noticeable as negative holes that coincide with the position of the brightest galaxy clusters also seen in the injected y -map. We also note from Figs 2 and 3 that the mean is consistent with zero, in spite of a non-zero monopole in the simulation. This recovered map of CMB anisotropies is also significantly noisier and this can be understood (in hindsight) as the cost of not projecting out other components in the map. The second row in Fig. 1 depicts the simultaneous separation of the CMB and the Compton- y map, constructed by mutually de-projecting the respective SEDs. This iteration of the MILC now correspond to the standard cILC, which is one of the methods used by the *Planck* collaboration. Note that this achieves a reasonable recovery of both the CMB and the Compton- y map. The SZ residuals in the CMB maps, seen in the previous iteration of the MILC, have now disappeared. There are more subtle biases in these maps that are not as visible in the considered simulation but become more prominent at reduced experimental sensitivity. The recovered component-separated maps continue to have excess noise, which results from not de-projecting other map components, namely, foregrounds. The means of both the recovered fields are still consistent with zero as seen in Figs 2 and 3.

The third row in Fig. 1 depicts the simultaneous separation of the CMB, y and the μ distortion map, constructed by mutually de-projecting the respective SEDs. This iteration of the MILC corresponds to the cILC, with three different but well-known SEDs. Many of the large-scale features in the CMB, and y -maps at the location of some of the brightest galaxy clusters are still recovered. However, the recovery of the component maps has severely degraded as compared to the previous two iterations of MILC. This is also apparent from the increase in the variance of the CMB and y maps as seen in Fig. 3. It is interesting to note that the peak of the 1PPDF for the CMB and y -parameter maps start showing a shift towards the input monopole values as seen in Fig. 2.

The subsequent rows of Fig. 1 depict the simultaneous separation of the signal components, when projecting out an increasing number of foreground moment vectors (1 in the fourth row, 3 in the fifth, 6 in the sixth, 8 in the seventh, and finally 10 in the eighth row). With these iterations we are truly entering a new regime, where

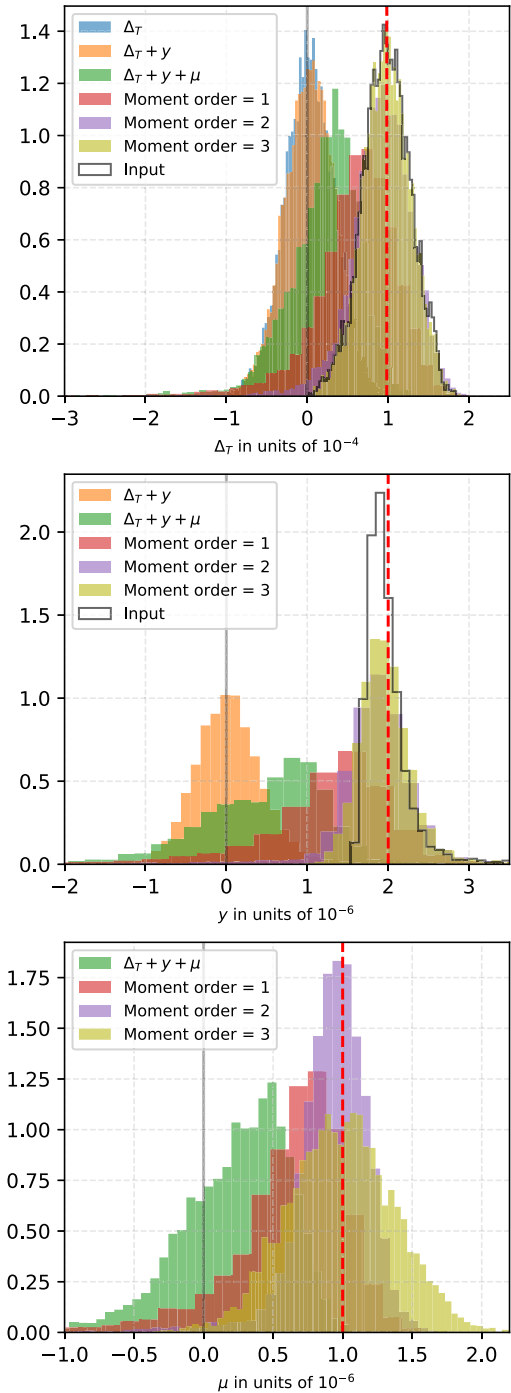


Figure 2. This figure depicts the evolution of the 1PPDF of the estimated component maps from analysis performed on simulations with a noise RMS = $50 \text{ Jy sr}^{-1} \text{ px}^{-1}$. The dashed line indicates the value of the monopole estimated directly from the input component maps.

foreground-averaging effects are gradually being captured. In fact, the foreground along different lines of sight cannot be characterized by a single SED, this is in stark contrast with the previous iterations of the MILC, where all the passed signal SEDs nearly perfectly characterize the respective components along different lines of sight. The recovered component maps depicted in rows 4–6 show a slow reduction in the residuals and then in rows 7 and 8 the recovery of the components maps is nearly perfect and one cannot discern

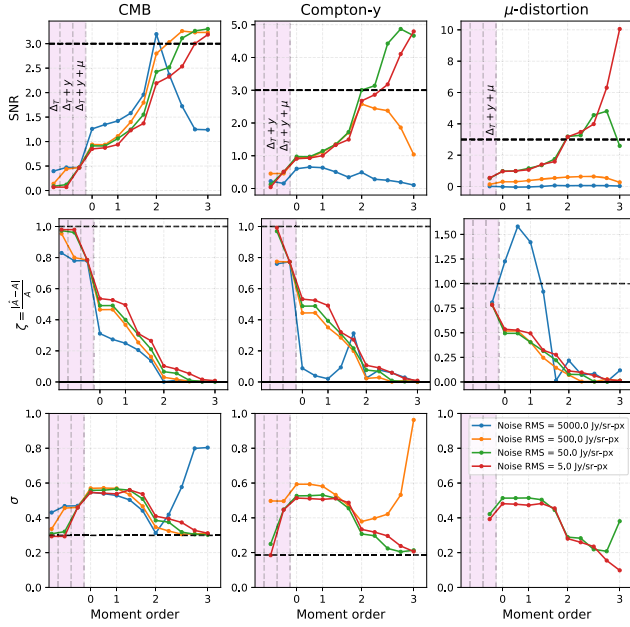


Figure 3. Evolution of the 1- and 2-point statistics of the recovered component maps, for varying number of vectors passed to the MILC at different noise levels. The top panels depict the evolution of the detection significance of the estimated monopole amplitude per pixel. The middle panels depict evolution of the estimated bias (ζ). The bottom panels depict the evolution of the standard deviation of the map in units of the true monopole of the respective component map. Here, the dashed black line marks the standard deviation of the input component map.




Figure 4. This panel of images depicts the dust foreground moments up to second order from analysis on simulations with a noise of $5 \text{ Jy sr}^{-1} \text{ px}^{-1}$. Higher order moment maps are correlated but in detail exhibit slightly different spatial morphologies.

the differences between the maps. This already shows that given sufficient frequency coverage and sensitivity, addition of moment terms does not come with unavoidable penalties in the recovery of signals but has a more complicated behaviour. We can also use the variance of the maps for a joint assessment of the quality of the recovered components. Studying the bottom panel of Fig. 3 reveals a much clearer picture of the evolution: the variance of the

map increases until one includes the zeroth-order vector, thereafter every additional foreground vector passed to the MILC results in an improved recovery, reaching near ideal recovery once all foreground moment vectors up to the second-order moment are projected out. Even in rows 7 and 8, where it is hardly possible to discern map-level differences, the variance statistics indicates that there is continual improvement. It is also very important and exciting to note that on including these higher order moments the peak of the 1PPDF of the maps converges to the true monopole amplitude of the respective components as seen in Fig. 2.

We note that a statistically significant detection of the monopole amplitude only becomes possible on nearly including all the second-order moment SED vectors and higher, however this trend also depends on the sensitivity of the measurements as seen in the top row of Fig. 3. For CMB and y fields, which are inherently anisotropic, the σ estimated from the map receives significant contribution from these intrinsic anisotropies. This is the reason for the saturation seen in the SNR plots for these respective components. Therefore, when measuring the monopoles of the CMB and y fields, it is essential to accurately model and subtract the error arising from the intrinsic variance of these fields, in order to fairly assess the statistical significance of the monopole detections (our assessments of the SNR for the monopole detection of these fields is biased low by roughly a factor of ~ 5 as we do not subtract the variance sourced by the anisotropies in the field). On the other hand, since our simulations included no anisotropic μ distortions, even though the μ distortions signal is smaller it is detected at a higher statistical significance for simulations with the lowest noise. The CMB and y signals being intrinsically larger is reflected in the observation that a statistically significant detection of the monopole amplitudes of these fields is possible even at lower sensitivities.

The details of the evolution of these statistics with the number of moment vectors passed to the MILC is of course not generic and depends on the sensitivity and frequency coverage of the experiment and the foregrounds model that is injected. Further these details also depend on how the data correlation matrix is defined, however a discussion on these details is beyond the scope of this work and will be discussed in Rotti et al. (in preparation). It will be superfluous and redundant to show the maps and 1PPDFs for analysis carried out on each simulation set with different noise. The evolution of the detection significance of the monopole amplitude, its bias and variance of the recovered component maps from analysis on simulations with different noise are neatly summarized in Fig. 3.

Here, we would like to highlight one generic trend that in the first four iterations of the MILC (these are methods employed for various work in current literature), the variance of the recovered component map *increases* when adding constraints, as seen from the bottom panel of Fig. 3. This gives rise to the misleading expectation that adding constraints invariably comes with a noise penalty. However, as the bottom panels of Fig. 3 demonstrate, this is a faulty extrapolation. We generically find that after projecting out an optimal number of foreground moment vectors, the variance of the recovered component map reaches a minimum, which is below or comparable to the error when these de-projections are not carried out (see Fig. 3).

It is also very interesting to note that, when considering the highest sensitivity ($\text{RMS} = 5 \text{ Jy sr}^{-1} \text{ px}^{-1}$, red line in Fig. 3), we find the simple ILC or cILC (only CMB and CMB+ y , respectively) do perform *on par* with the MILC, indicating that signal SED orthogonality is a direct function of sensitivity. However, for all our lower sensitivity simulations, the MILC with inclusion of higher order foreground moments has the best performance. In hindsight,

this behaviour is easy to understand: when one does not project out foregrounds, these are included in the noise budget of the recovered component maps. By projecting out an ‘optimal number’ of foreground components, the noise budget of the component-separated maps is reduced. The ‘optimal number’ of foreground components can be assessed by studying when the variance of the component maps begins to increase after reaching the minimum (see the blue line for CMB variance and orange for y -map variance in bottom panels of Fig. 3). To be able to make these projections, one requires enough channels, and this diagnostic could be used to optimize future CMB missions.

The foreground moment maps are themselves astrophysical observables and these can be used to gain deeper understanding of the galactic foregrounds. These moment maps are naturally returned by the MILC and an example set of the first measurement of foreground moment maps on simulations and are shown in Fig. 4. These moment maps are related to the spectral parameters characterizing the IGM, but elucidating this connection requires more work. Current and future analysis on multifrequency microwave data will deliver these observables as by-products.

5 DISCUSSIONS AND CONCLUSIONS

We have demonstrated that the scope of ILC approaches can be significantly extended by using the language of moments, resulting in the MILC method introduced here (Section 2). Using sky simulations, we have demonstrated that the MILC can not only be used to improve the recovery of anisotropic components but also allows an extraction of monopole signals. This opens new ways of thinking about mitigating foreground challenges that are essential to future CMB spectral distortion studies.

Specifically, we have demonstrated that it is important to consider higher order moments into account, without which both the monopole and the anisotropy measurements can be significantly biased (see Figs 1 and 2). Furthermore, we have clarified that increasing the number of moments does not generically lead to more noise in the recovered component maps (see Fig. 3). On the contrary, we have demonstrated that projecting out an ‘optimal number’ of foreground moments will invariably lead to a more robust and less noisy recovery of the component maps.

Our study also shows that the MILC can have more optimal performance as compared to simple ILC and cILC methods. Only for very high sensitivity measurements are the performances of the ILC and cILC on par with MILC. These lead to a revision in the understanding of how to apply ILC methods. The variance of the recovered component maps can be used to estimate the ‘optimal number’ of foreground vectors to reach best performance. This number can be directly translated into the requirement of the number of frequency channels at some given sensitivity and has the potential of being a quantitative metric for designing future CMB experiments (under some assumptions about the sky foregrounds).

While the discussion presented in this work is focused on simulations which include only thermal dust as foregrounds, we have also confirmed that this method works when all the foregrounds (synchrotron, free-free and dust) are included. This turns into a high-dimensional problem, and one has to worry about the most relevant SED vectors and the order in which they are included. This is an ongoing study (Rotti et al., in preparation), the details of which are beyond the basic idea being presented in this work. The foreground moment maps are a natural by-product of MILC and we foresee these becoming an important measurable with future, high-sensitivity multifrequency microwave measurements. MILC is

currently being applied to *Planck* maps and the measurement of foreground moment maps from this analysis will be reported in a future publication. However, several questions call for further investigation.

Effect of un-accounted signals: The work presented here assumes that the base SEDs, characterizing the all the foregrounds, are known and can be supplied to the MILC. This allows MILC to capture all the effects from averaging processes. We have shown that *not* de-projecting the higher order moments can come at the cost of excess noise and biases in the recovered component maps. This aspect will become even more apparent as we work with higher sensitivity data to target low signal-to-noise components like for CMB spectral distortions and B modes. Encountering foregrounds with unmodelled SEDs will lead to a poorer component separation. However, having even an approximate SED for these components can improve the signal recovery. Applied to real data, MILC should thus be thought of as a *semiblind* component separation method. Adding Synchrotron, CIB and free-free moments is straightforward (Chluba et al. 2017), however, developing a moment description for some of the other known foregrounds (e.g. anomalous microwave emission, extragalactic CO) will be an important step in generalizing MILC.

Low multipole leakage to monopole and sky-coverage: In this work, we carried out the analysis on a flat-sky patch and implicitly assumed that local patch monopole is the same as that of the global monopole. For the CMB and y , which do have significant amount of anisotropies, the local patch monopole can differ significantly from the global monopole owing to contributions from large-scale ($\lambda \geq$ patch size) fluctuations. For this reason the local patch monopole measurement of these components could be noticeably biased. However, in standard cosmology the μ -distortions are expected to be monopolar, and hence for this component, the amplitude of the local monopole is expected to be the same as the global monopole. This suggest that if one wants to measure μ distortion monopole, making deep measurements on a flat-sky patch could be a viable strategy. On the other hand, an ultimate test of isotropy of the primordial distortion signal (e.g. also including the cosmological recombination lines Sunyaev & Chluba 2009; Chluba & Ali-Haïmoud 2016) will require all-sky measurements. A detailed study with the goal to optimize for CMB spectral distortion science goals will be presented in a future publication.

Truncation and ordering of vectors: The variance of the recovered component maps was seen to have a non-monotonic behaviour, reaching a minimum at some optimal number of moment vectors (see Fig. 3). It will be important to study if the moment ordering could be altered to obtain a monotonic reduction in variance of the component-separated maps. This will help in devising the optimal strategies for taming this high-dimensional optimization problem. Additional benefits could stem from orthogonalizing moment SEDs (e.g. Gram–Schmidt or principle component analysis) prior to the MILC analysis. This will enable us to more clearly rank the various moment vectors according to their expected signal-to-noise levels. Ultimately, a comparison of the target signal level (i.e. μ) to the level of a moment SEDs has to be used to define truncation criteria for the analysis. Other diagnostics, based on Bayesian evidence, Shannon entropy or χ^2 -gains, could also be used for this purpose.

Optimality for monopole recovery: It is imperative to discuss the subtle detail of covariance estimation which we have found to be highly relevant to the question of optimality. As mentioned in Section 3.5, the results presented here were derived under the assumption that the local monopole cannot be subtracted accurately. This was motivated by the fact that in a full sky analysis it is unlikely

that all local monopoles are accurately subtracted, owing to highly anisotropic and non-Gaussian nature of galactic foregrounds.¹¹ However, were it possible to subtract the local monopole from the multifrequency maps perfectly for the covariance estimation step, then we find that the ILC forecasts are in fact more promising. Notably, from analysis on the highest sensitivity simulations, we find μ distortions to be measured at nearly six times higher significance (i.e. $\simeq 60\sigma$), without significant bias, that too without performing any foreground moment de-projections.

Further it is important to note that MILC performs a pixel-by-pixel recovery of the signal field. While this approach has been (empirically) proven to be ‘optimal’ for anisotropic signal, it is not immediately obvious that it is optimal for the recovery of monopole signals. With single pixel sky-averaged SED measurements, one can make significant gains ($\simeq \sqrt{N_{pix}}$) in sensitivity, however, possibly at the cost of increasing the foreground complexity due to generation of moments, a natural consequence of the averaging process. In principle this questions is relevant to measurements of all the long wavelength modes on the sky (i.e. low- ℓ CMB anisotropies). Working with full resolution maps, will invariably lead to lower sensitivity per pixel but provides access to additional information about pixel–pixel correlations for different foreground components and, possibly, a lower foreground complexity in each pixel. It can be anticipated that the answer lies somewhere in the middle. It will thus be essential to compare and contrast the performance of these two analysis approaches, to understand how to best combine them for extraction of spectral distortion signals. Bearing all these details in mind, it is therefore important to consider the results presented in this work as being pessimistic estimates, of the monopole measurements.

Combination of data sets: High-sensitivity and high-resolution anisotropy measurements of the CMB will soon become available. The MILC method makes it clear how we can blend CMB anisotropy measurement to help with cleaning of the absolutely calibrated measurements planned for the future (e.g. Chluba et al. 2019; Kogut et al. 2019). Increased frequency coverage at low frequencies ($\nu \lesssim 10$ GHz) has also been shown to be important for the recovery of small μ -distortion signals (Abitbol et al. 2017). Again the MILC method can be extended to include external information from low-frequency data sets. This also means that a detailed study of how to propagate of systematic effects when combining various data sets will be required.

Application to other problems: The MILC method is not limited to studies of primordial CMB distortion signals, but similarly can be applied to the extraction of the global 21cm signal from the cosmic dark ages (e.g. Pritchard & Loeb 2008; Fialkov, Barkana & Cohen 2018). At low frequencies, the dominant foregrounds are due to radio sources and galactic synchrotron, which can both be modelled using a power-law moment expansion (Chluba et al. 2017). In a similar manner, it is straightforward to extend the MILC approach to CMB B-mode searches, ensuring that foreground residuals caused by averaging effects remain under control (Remazeilles, Rotti & Chluba 2020). In this case, two moment hierarchies have to be introduced, demanding high-sensitivity and broad spectral coverage. De-correlation effects across frequency can also be modelled in this way. Moment expansions at the power-spectrum level can provide further insight and leverage (Mangilli et al. 2019). We plan to explore these possibilities in future works.

¹¹All currently planned spectral distortion measurement missions will measure the full sky.

The component-separated maps depicted in the third column of Fig. 1 represents a study analogous to that presented in Remazeilles & Chluba (2018). We see that there are significant biases at this iteration of the MILC, hinting that one could make significant improvements in forecasts for μ - T correlations by including higher order moments into the analysis. Of course these studies will need to be repeated with all the relevant foregrounds and the instrument configurations considered in Remazeilles & Chluba (2018). The utility of MILC for the extraction of γ - T correlations to study primordial non-Gaussianity (e.g. Ravenni et al. 2017) should also be carefully considered.

Finally, our study also indicates that even for the CMB temperature recovery, the simple ILCs and cILCs are not fully optimal and may have significant residuals. This could relate to some of the large-scale anomalies seen by *Planck* and *WMAP* (namely lack of power and isotropy violation on large-angular scales). Revisiting this question using MILC on real data could thus prove highly instructive and may also inform us about expected foreground complexities for future B-mode searches, both questions we plan to address in the future.

ACKNOWLEDGEMENTS

This work was supported by the ERC Consolidator Grant *CMBSPEC* (No. 725456) as part of the European Union’s Horizon 2020 research and innovation program. JC was supported by the Royal Society as a Royal Society URF at the University of Manchester, UK. We thank Mathieu Remazeilles, Eiichiro Komatsu, and Kevin Huffenberger for many useful discussions. We also thank the anonymous referee for probing us on many subtle but important aspects of our analysis which helped improve the clarity and presentation of the material significantly.

DATA AVAILABILITY

No new data was generated for this work. The simulations used in this work were generated using the publicly available code PYSM.

REFERENCES

- Abitbol M. H., Chluba J., Hill J. C., Johnson B. R., 2017, *MNRAS*, 471, 1126
 Adam R. et al., 2016, *A&A*, 594, A9
 André P. et al., 2014, *J. Cosmol. Astropart. Phys.*, 2, 6
 Balashev S. A. et al., 2015, *ApJ*, 810, 131
 Bennett C. L. et al., 2003, *ApJS*, 148, 1
 Burigana C., Danese L., de Zotti G., 1991, *A&A*, 246, 49
 Burigana C. et al., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 021
 Chluba J., 2015, *MNRAS*, 454, 4182
 Chluba J., 2016, *MNRAS*, 460, 227
 Chluba J., Ali-Haïmoud Y., 2016, *MNRAS*, 456, 3494
 Chluba J., Sunyaev R. A., 2012, *MNRAS*, 419, 1294
 Chluba J., Khatri R., Sunyaev R. A., 2012, *MNRAS*, 425, 1129
 Chluba J., Switzer E., Nelson K., Nagai D., 2013, *MNRAS*, 430, 3054
 Chluba J., Hill J. C., Abitbol M. H., 2017, *MNRAS*, 472, 1195
 Chluba J. et al., 2019, preprint (arXiv:1909.01593)
 Danese L., de Zotti G., 1981, *A&A*, 94, L33
 Delabrouille J., Cardoso J.-F., Patanchon G., 2003, *MNRAS*, 346, 1089
 Delabrouille J., Cardoso J.-F., Le Jeune M., Betoule M., Fay G., Guilloux F., 2009, *A&A*, 493, 835
 Delabrouille J. et al., 2019, preprint (arXiv:1909.01591)
 Desjacques V., Chluba J., Silk J., de Bernardis F., Doré O., 2015, *MNRAS*, 451, 4460
 Emami R., Dimastrogiovanni E., Chluba J., Kamionkowski M., 2015, *Phys. Rev. D*, 91, 123531
 Eriksen H. K., Jewell J. B., Dickinson C., Banday A. J., Górski K. M., Lawrence C. R., 2008, *ApJ*, 676, 10

- Fialkov A., Barkana R., Cohen A., 2018, *Phys. Rev. Lett.*, 121, 011101
- Fixsen D. J., Cheng E. S., Gales J. M., Mather J. C., Shafer R. A., Wright E. L., 1996, *ApJ*, 473, 576
- Fixsen D. J. et al., 2011, *ApJ*, 734, 5
- Ganc J., Komatsu E., 2012, *Phys. Rev. D*, 86, 023518
- Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelman M., 2004, *ApJ*, 622, 759
- Hill J. C., Battaglia N., Chluba J., Ferraro S., Schaan E., Spergel D. N., 2015, *Phys. Rev. Lett.*, 115, 261301
- Hu W., Silk J., 1993, *Phys. Rev. D*, 48, 485
- Illarionov A. F., Sunyaev R. A., 1974, *Astron. Zh.*, 51, 1162
- Kogut A. et al., 2011a, *J. Cosmol. Astropart. Phys.*, 7, 25
- Kogut A. et al., 2011b, *ApJ*, 734, 4
- Kogut A., Chluba J., Fixsen D. J., Meyer S., Spergel D., 2016, in MacEwen H. A., Fazio G. G., Lystrup M., Batalha N., Siegler N., Tong E. C., eds, Proc. SPIE Conf. Ser. Vol. 9904, Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave. SPIE, Bellingham, p. 99040W
- Kogut A., Abitbol M. H., Chluba J., Delabrouille J., Fixen D., Hill C. J., Patil S. P., Rotti A., 2019, preprint ([arXiv:1907.13195](https://arxiv.org/abs/1907.13195))
- Mangilli A., Aumont J., Rotti A., Boulanger F., Chluba J., Ghosh T., Montier L., 2019, preprint ([arXiv:1912.09567](https://arxiv.org/abs/1912.09567))
- Mather J. C. et al., 1994, *ApJ*, 420, 439
- Mroczkowski T. et al., 2019, *Space Sci. Rev.*, 215, 17
- Pajer E., Zaldarriaga M., 2012, *Phys. Rev. Lett.*, 109, 021302
- Planck Collaboration XII, 2014a, *A&A*, 571, A12
- Planck Collaboration XVI, 2014b, *A&A*, 571, A16
- Planck Collaboration XXI, 2014c, *A&A*, 571, A21
- Planck Collaboration X, 2016a, *A&A*, 594, A10
- Planck Collaboration XI, 2016b, *A&A*, 594, A11
- Planck Collaboration XIII, 2016c, *A&A*, 594, A13
- Planck Collaboration XXII, 2016d, *A&A*, 594, A22
- Planck Collaboration XLVI, 2016e, *A&A*, 596, A107
- Pritchard J. R., Loeb A., 2008, *Phys. Rev. D*, 78, 103511
- Pritchard J. R., Loeb A., 2010, *Phys. Rev. D*, 82, 023006
- Ravenni A., Liguori M., Bartolo N., Shiraishi M., 2017, *J. Cosmol. Astropart. Phys.*, 9, 042
- Refregier A., Komatsu E., Spergel D. N., Pen U.-L., 2000, *Phys. Rev. D*, 61, 123001
- Remazeilles M., Chluba J., 2018, *MNRAS*, 478, 807
- Remazeilles M., Chluba J., 2020, *MNRAS*, 494, 5734
- Remazeilles M., Delabrouille J., Cardoso J. F., 2011, *MNRAS*, 418, 467
- Remazeilles M., Dickinson C., Eriksen H. K. K., Wehus I. K., 2016, *MNRAS*, 458, 2032
- Remazeilles M., Rotti A., Chluba J., 2020, *MNRAS*, preprint ([arXiv:2006.08628](https://arxiv.org/abs/2006.08628))
- Remazeilles M. et al., 2018, *J. Cosmol. Astropart. Phys.*, 4, 023
- Sathyanarayana Rao M., Subrahmanyam R., Udaya Shankar N., Chluba J., 2015, *ApJ*, 810, 3
- Sehgal N., Bode P., Das S., Hernandez-Monteagudo C., Huppenberger K., Lin Y.-T., Ostriker J. P., Trac H., 2010, *ApJ*, 709, 920
- Stolyarov V., Hobson M. P., Lasenby A. N., Barreiro R. B., 2005, *MNRAS*, 357, 145
- Sunyaev R. A., Chluba J., 2009, *Astron. Nachr.*, 330, 657
- Sunyaev R. A., Zeldovich Y. B., 1970a, *Ap&SS*, 7, 20
- Sunyaev R. A., Zeldovich Y. B., 1970b, *Ap&SS*, 9, 368
- Tegmark M., de Oliveira-Costa A., Hamilton A. J. S., 2003, *Phys. Rev. D*, 68, 123523
- Thorne B., Dunkley J., Alonso D., Naess S., 2017, *MNRAS*, 469, 2821
- Zeldovich Y. B., Sunyaev R. A., 1969, *Ap&SS*, 4, 301

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.