

Lawrence Berkeley National Laboratory

Physics

Title

Detecting Patchy Reionization in the Cosmic Microwave Background

Permalink

<https://escholarship.org/uc/item/9302k424>

Journal

Physical Review Letters, 119(2)

ISSN

0031-9007

Authors

Smith, Kendrick M

Ferraro, Simone

Publication Date

2017-07-14

DOI

10.1103/physrevlett.119.021301

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Detecting patchy reionization in the CMB

Kendrick M. Smith¹ and Simone Ferraro^{2,3}

¹*Perimeter Institute for Theoretical Physics, Waterloo ON N2L 2Y5, Canada*

²*Berkeley Center for Cosmological Physics, University of California, Berkeley CA 94720, USA*

³*Miller Institute for Basic Research in Science, University of California, Berkeley CA 94720, USA*

(Dated: July 7, 2016)

Upcoming cosmic microwave background (CMB) experiments will measure temperature fluctuations on small angular scales with unprecedented precision. Small-scale CMB fluctuations are a mixture of late-time effects: gravitational lensing, Doppler shifting of CMB photons by moving electrons (the kSZ effect), and residual foregrounds. We propose a new statistic which separates the kSZ signal from the others, and also allows the kSZ signal to be decomposed in redshift bins. The decomposition extends to high redshift, and does not require external datasets such as galaxy surveys. In particular, the high-redshift signal from patchy reionization can be cleanly isolated, enabling future CMB experiments to make high-significance and qualitatively new measurements of the reionization era.

I. INTRODUCTION

On large angular scales ($l \lesssim 2000$), anisotropy in the cosmic microwave background is mainly sourced by fluctuations at redshift $z \approx 1100$. On smaller angular scales ($l \gtrsim 2000$), this “primary” anisotropy is exponentially suppressed, and CMB fluctuations are mainly a mixture of several “secondary” or late-time effects.

Among secondary effects with the same blackbody spectrum as the primary CMB, the largest are gravitational lensing, and the kinematic Sunyaev-Zel’dovich (kSZ) effect. The kSZ effect refers to Doppler shifting of CMB photons as they scatter on radially moving inhomogeneities in free electron density [1–3]. The kSZ anisotropy can be roughly decomposed into a “late-time” contribution from redshifts $z \lesssim 3$, when inhomogeneities are large due to gravitational growth of structure, and a “reionization” contribution from redshift $z \sim 7$, when the ionization fraction is expected to be inhomogeneous during “patchy” reionization [4–8].

In Fig. 1 we compare contributions to the temperature power spectrum C_l^{TT} from weak lensing of the CMB, late-time kSZ, and reionization kSZ, in a fiducial model to be described shortly. Individually, these three contributions are very interesting. Gravitational lensing depends on cosmological parameters such as neutrino mass [9], late-time kSZ probes the distribution of electrons in dark matter halos as well as the large-scale velocity field, and reionization kSZ may provide the first observational window on patchy reionization, which will shed light on the formation of first stars and other sources of ionizing photons. Although the total power spectrum will soon be measured very precisely at high l [10], it is unclear how well these signals can be disentangled, since all three components have large astrophysical modelling uncertainties, and the two kSZ contributions are essentially degenerate at the power spectrum level.

In this paper we will propose a higher-order statistic which isolates the kSZ signal, and moreover gives information about its source redshift dependence, allowing the late-time and reionization kSZ to be separated. This




FIG. 1: Fiducial model for the CMB temperature power spectrum C_l^{TT} used throughout this paper, split into primary, lensing, late-time kSZ, and reionization kSZ contributions.

will complement measurements from future 21cm experiments [11, 12]. We describe the intuitive idea here, with a more formal description in the next section.

We first recall that to a good approximation, the kSZ power spectrum may be written as an integral [13]:

$$C_l^{\text{kSZ}} = \int dz Q(z) \langle v_r(z)^2 \rangle P_e \left(\frac{l}{\chi(z)}, z \right) \quad (1)$$

where $P_e(k, z)$ is the free electron power spectrum, $\langle v_r^2 \rangle = \langle v^2 \rangle / 3$ is the mean squared radial velocity, and $\chi(z)$ is the comoving distance to redshift z . The radial weight function $Q(z)$ is given by

$$Q(z) = T_{\text{CMB}}^2 \frac{H(z)}{\chi(z)^2} \left(\frac{d\bar{\tau}}{dz} \right)^2 e^{-2\bar{\tau}(z)} \quad (2)$$

where $d\bar{\tau}/dz$ is the optical depth per unit redshift.

Throughout the paper we will use the following notation frequently. Let \bar{K} be the sky-averaged small-scale power spectrum in a fixed high- l band (say $3000 \leq l \leq 5000$). For each direction $\hat{\mathbf{n}}$ on the sky, let $K(\hat{\mathbf{n}})$ be the

locally measured small-scale power spectrum near sky location $\hat{\mathbf{n}}$. The precise definitions of \bar{K} and $K(\hat{\mathbf{n}})$ will be given in the next section.

It is intuitively clear that $K(\hat{\mathbf{n}})$ will be better approximated by using the actual realization of $v_r^2(\hat{\mathbf{n}}, z)$ along the line of sight in direction $\hat{\mathbf{n}}$ in the integral in Eq. (1), rather than the cosmic average $\langle v_r^2 \rangle$. This leads to anisotropy in $K(\hat{\mathbf{n}})$ on large angular scales.

To estimate the level of anisotropy, suppose we divide the line of sight into segments of size 50 Mpc (the coherence length of the velocity field), and roughly model the radial velocity v_r as an independent Gaussian random number in every segment. Since the line of sight is 10^4 Mpc in length, $K(\hat{\mathbf{n}})$ can be roughly modelled as the sum of squares of $N = 200$ independent Gaussians. This suggests that fluctuations in $K(\hat{\mathbf{n}})$ between different lines of sight are of fractional size $\sqrt{2/N} \approx 0.1$. Rephrasing, if we measure the CMB in two regions of sky separated by more than ~ 1 degree, so that the lines of sight sample independent realizations of the velocity field, the kSZ power spectra will differ by $\approx 10\%$. This is a large non-Gaussian effect which is not present for lensing and other secondaries, allowing statistical separation of the kSZ signal.

In fact we can go further by considering C_L^{KK} , the angular power spectrum of $K(\hat{\mathbf{n}})$. Suppose we write C_L^{KK} as a sum of contributions from multiple source redshift bins. In the next section we will show (Fig. 2) that the contribution from redshift z has a broad peak at wavenumber $L_* \sim k_* \chi(z)$, where $k_* \approx 0.01 h \text{ Mpc}^{-1}$. Thus the shape of C_L^{KK} is source redshift dependent. In the general case where C_L^{KK} is a sum over redshift bins, we can “deconvolve” the observed C_L^{KK} to infer the contribution from each bin, thus separating the late-time and reionization kSZ signals. The main advantage of this method (compared to an analysis based on C_l^{TT}) is its robustness: we can make statements about reionization which do not depend on precise modelling of the other contributions.

II. MODELLING THE SIGNAL

We use the following fiducial model for the kSZ power spectrum and its source redshift distribution. We model the late-time kSZ using Eq. (1) with $P_e(k, z) = W(k, z)^2 P_{\text{nl}}(k, z)$, where P_{nl} is the nonlinear matter power spectrum from CAMB [14], and we have defined

$$W(k, z) = 0.85 \left(1 + \frac{kD(z)}{0.5 h \text{ Mpc}^{-1}} \right)^{-1/2} \quad (3)$$

where $D(z)$ is the growth function normalized to $D(z) = 1/(1+z)$ at high z . This form of $W(k, z)$ is a simple fitting function which gives approximate agreement with the “cooling + star formation” model from [15], for both C_l^{kSZ} and dC_l^{kSZ}/dz . The prefactor 0.85 assumes that at late times, 15% of electrons are in the neutral medium or stars [16].

We model the reionization kSZ by assigning a Gaussian redshift distribution to the simulated kSZ power spectrum from Battaglia et al [4]:

$$\left(\frac{dC_l^{\text{kSZ}}}{dz} \right)_{\text{rei}} = \frac{e^{-(z-z_{\text{re}})^2/2\sigma_{\text{re}}^2}}{(2\pi\sigma_{\text{re}}^2)^{1/2}} C_l^{\text{Battaglia}} \quad (4)$$

where $(z_{\text{re}}, \sigma_{\text{re}}) = (8.8, 1.0)$.

We now give a formal definition of the quantities \bar{K} and $K(\hat{\mathbf{n}})$ from the introduction. First fix a filter $W_S(l)$, and define a high-pass filtered CMB in Fourier space by $T_S(\mathbf{l}) = W_S(\mathbf{l})T(\mathbf{l})$. Unless otherwise specified, we choose $W_S(l) \propto (C_l^{\text{kSZ}})^{1/2}/C_l^{\text{tot}}$, where C_l^{tot} is the total CMB power spectrum, including instrumental noise. We then define $K(\hat{\mathbf{n}}) = T_S(\hat{\mathbf{n}})^2$ by squaring in real space, and define \bar{K} to be the all-sky average $\bar{K} = \langle K(n) \rangle$. We note that

$$\bar{K} = \int \frac{d^2\mathbf{l}}{(2\pi)^2} W_S(l)^2 C_l^{\text{tot}} \quad (5)$$

so that \bar{K} can be interpreted as average high- l power (with W_S^2 -weighting) and $K(\hat{\mathbf{n}})$ can be interpreted as “locally measured high- l power near $\hat{\mathbf{n}}$ ”.

Our main statistic will be C_L^{KK} , the “power spectrum of the power spectrum”. Note that there are two scales, a small scale $l \gtrsim 3000$ selected by the filter W_S where the CMB is measured, and a large scale $L \lesssim 300$ where clustering in the small-scale power is measured. Viewed as a four-point estimator in the CMB, C_L^{KK} sums over quadruples $T(\mathbf{l}_1)T(\mathbf{l}_2)T(\mathbf{l}_3)T(\mathbf{l}_4)$ which are “collapsed”, in the sense that the CMB wavenumbers $|\mathbf{l}_i|$ are large, but the intermediate wavenumber $\mathbf{L} = (\mathbf{l}_1 + \mathbf{l}_2)$ is small. This is similar to CMB lens reconstruction, where the lensing potential $\phi(\mathbf{L})$ and its power spectrum $C_L^{\phi\phi}$ are estimated on large scales using CMB temperature fluctuations on scales $l \gg L$.

Continuing the analogy with lens reconstruction, we define the reconstruction noise N_L^{KK} to be the value of C_L^{KK} that would be obtained if the small-scale temperature were a Gaussian field. A short calculation gives:

$$N_L^{KK} = 2 \int \frac{d^2\mathbf{l}}{(2\pi)^2} W_S^2(\mathbf{l}) W_S^2(\mathbf{L} - \mathbf{l}) C_l^{\text{tot}} C_{\mathbf{L}-\mathbf{l}}^{\text{tot}}. \quad (6)$$

In the regime $L \ll l$ of interest, N_L^{KK} is nearly constant in L .

Now we would like to model the effect described in the previous section: large-scale non-Gaussian contributions to C_L^{KK} due to correlated radial velocities along the line of sight. We introduce a simple model, the “ η -model”, as follows.

We write the sky-averaged small-scale power spectrum \bar{K}_{kSZ} as an integral $\bar{K}_{\text{kSZ}} = \int dz (d\bar{K}/dz)$, where

$$\frac{d\bar{K}}{dz} = \int \frac{d^2\mathbf{l}}{(2\pi)^2} W_S(l)^2 \frac{dC_l^{\text{kSZ}}}{dz} \quad (7)$$

and our fiducial model for dC_l^{kSZ}/dz was given in Eqs. (1), (3), (4). The η -model is the ansatz that, in

a fixed realization of the radial velocity field $v_r(\hat{\mathbf{n}}, z)$, the locally measured small-scale power $K(\hat{\mathbf{n}})$ can be modelled as

$$K(\hat{\mathbf{n}}) = \int dz \frac{d\bar{K}}{dz} \eta(\hat{\mathbf{n}}, z) \quad (8)$$

where $\eta(\hat{\mathbf{n}}, z) = v_r(\hat{\mathbf{n}}, z)^2 / \langle v_r(z)^2 \rangle$. In other words, we assume that the locally generated kSZ power along the line of sight is proportional to v_r^2 , but neglect additional non-Gaussian effects. Fully characterizing kSZ non-Gaussianity on all scales is outside the scope of this paper. Here we are simply claiming that at minimum, the non-Gaussian signal predicted by Eq. (8) must exist.

In the Limber approximation, the contribution to C_L^{KK} predicted by the η -model is

$$C_L^{KK} = \int dz \frac{H(z)}{\chi(z)^2} \left(\frac{d\bar{K}}{dz} \right)^2 P_\eta^\perp \left(\frac{L}{\chi(z)} \right) \quad (9)$$

where $P_\eta^\perp(k)$ is the power spectrum of the field η evaluated at a wavenumber \mathbf{k} perpendicular to the line-of-sight direction. We will compute P_η^\perp in linear theory, where it is independent of z , but depends on the direction of \mathbf{k} since η is an anisotropic field. By a short calculation using Wick's theorem, P_η^\perp is given by

$$P_\eta^\perp(k) = \frac{2}{\langle v_r^2 \rangle^2} \int \frac{d^3 \mathbf{k}'}{(2\pi)^3} \frac{(k_r')^2 (k_r - k_r')^2}{k^2 (\mathbf{k} - \mathbf{k}')^2} P_v(k') P_v(\mathbf{k} - \mathbf{k}') \quad (10)$$

where P_v is the linear velocity power spectrum and $k_r = 0$ has been assumed.

In Fig. 2 we show the contributions to C_L^{KK} from late-time and reionization kSZ, computed using the η -model. Note that the reionization kSZ makes a larger contribution to C_L^{KK} than the late-time kSZ, even though the two are comparable in the CMB power spectrum C_l^{TT} . This is because the late-time line-of-sight integral is more extended in comoving distance χ , so it samples more coherence lengths of the velocity field, making the signal more Gaussian.

A crucial property of the η -model is that the contribution to C_L^{KK} from source redshift z is proportional to $P_\eta^\perp(L/\chi(z))$, with no additional L dependence. Thus the “shape” in L depends only on large-scale linear theory, but the overall amplitude depends on small-scale physics (via $d\bar{K}/dz$).

This independence of small-scale physics means that we can test the η -model using simplified simulations. We construct an ensemble of 3D simulations neglecting baryonic physics and using the 2LPT approximation to the N -body equations of motion. Rather than using a light-cone geometry, we simply project a $z = 2$ snapshot onto the 2D periodic “sky” formed by one of the box faces. The agreement with the η -model is excellent (Fig. 3 top panel). We plan to extend this simulation pipeline in future work, but expect that more accurate simulations will simply change the overall amplitude, and capture small




FIG. 2: Modelling the “power spectrum of the kSZ power spectrum” C_L^{KK} , assuming $2 \mu\text{K-arcmin}$ noise and $\theta_{\text{FWHM}} = 1'$. *Top panel.* Contributions to C_L^{KK} from reconstruction noise (Eq. (6)) and kSZ (Eq. (9)). *Bottom panel.* The kSZ contribution to C_L^{KK} per unit source redshift. The distribution is strongly bimodal, justifying a decomposition into low- z and reionization contributions.

effects such as curved-sky corrections and deviations from the Limber approximation.

So far we have considered contributions to C_L^{KK} from the η -model (our signal), and from Gaussian mode-counting (the noise N_L^{KK}). In order to claim that the signal is robust, it is important to understand how other non-Gaussian effects may contribute to C_L^{KK} .

Some non-Gaussian signals do not cluster on large scales. For example, even after multifrequency analysis, the CMB maps will be contaminated by residual thermal SZ clusters at some level. This signal can be modelled very accurately as a sum of unclustered Poisson sources with angular profile F_l . On large scales L where the profile F_L is nearly constant, a short calculation shows that the contribution to C_L^{KK} is nearly constant in L . We will account for this type of contribution by marginalizing an arbitrary constant δC_L^{KK} in our signal-to-noise forecasts.

Other non-Gaussian signals do cluster on large scales, most importantly gravitational lensing. In the middle panel of Fig. 3, we show the bias to C_L^{KK} obtained from simulated CMB lensing maps, using a pipeline similar to [17]. The lensing bias is comparable to the reconstruction noise and non-constant on large scales.

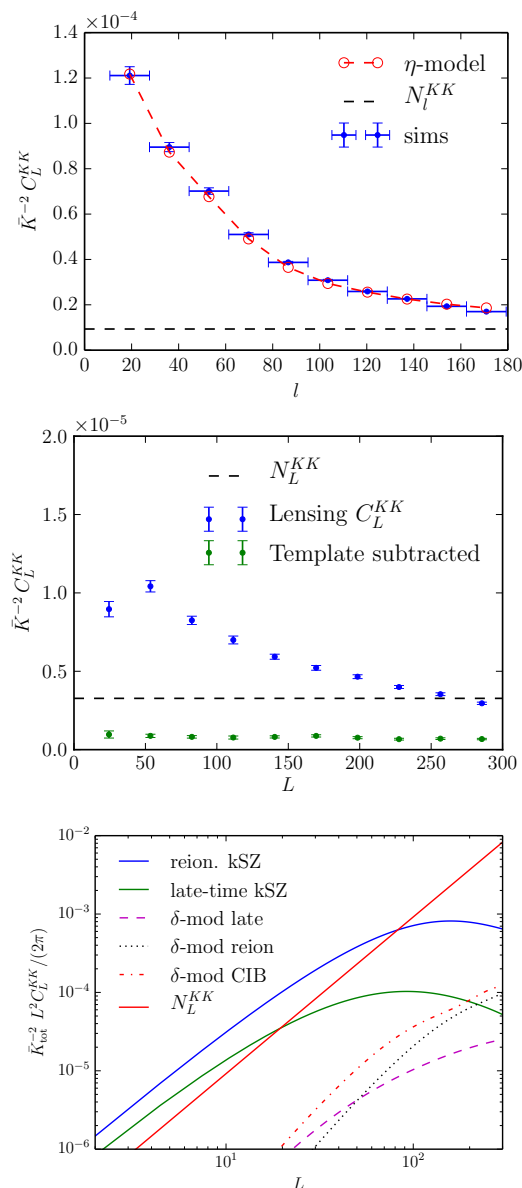


FIG. 3: *Top panel*: Comparison between our model for C_L^{KK} and 3D large-scale structure simulations described in the text. The kSZ maps were filtered to $3000 \leq l \leq 5000$ with no instrumental noise. *Middle panel*: Lensing contribution to C_L^{KK} from simulations, before and after template subtraction. *Bottom panel*: Estimated non-Gaussian contribution to C_L^{KK} from large-scale clustering of small scale modes, computed using a model (“ δ -modulation”) described in the text.

However, we have found that C_L^{KK} can be “lens-cleaned” as follows. Let us make the ansatz that on large scales, lensing produces terms in $K(\hat{\mathbf{n}})$ proportional to the local lensing convergence κ and squared CMB gradient $(\nabla T)^2$:

$$K(\hat{\mathbf{n}}) \supset \alpha \kappa(\hat{\mathbf{n}}) + \beta (\nabla T(\hat{\mathbf{n}}))^2 \quad (11)$$

We lens-clean K on large scales by constructing template maps $(\nabla T)^2$ and κ , and subtracting best-fit multi-

ples of the templates from $K(\hat{\mathbf{n}})$ before computing C_L^{KK} . Our template $(\nabla T)^2$ map is made by low-pass filtering to $\ell < 2000$ and squaring. We assume that a template map of κ is also available in the survey region from CMB polarization lens reconstruction (a detailed forecast shows that κ has high signal-to-noise for noise levels $\lesssim 6 \mu\text{K-arcmin}$).

Remarkably, this simple template-cleaning procedure removes nearly all lensing power in simulation (Fig. 3 middle panel). Furthermore the residual power is constant in L , so that it is removed by our previously mentioned marginalization. We therefore expect that CMB lensing bias to C_L^{KK} can be made negligible.

Next we would like to consider clustered secondaries. One type of non-Gaussian contribution is “ δ -modulated power”: power along the line of sight whose amplitude is linear in the local overdensity $\delta(\hat{\mathbf{n}}, z)$. Schematically, we write:

$$K(\hat{\mathbf{n}}) \supset \int dz \frac{d\bar{K}}{dz} \left(1 + \beta(z) \delta(\hat{\mathbf{n}}, z) \right) \quad (12)$$

where $\beta(z)$ is a linear bias parameter which relates the small-scale CMB *power* along the line of sight to the local density. This type of model, in conjunction with the Poisson and Gaussian terms previously considered, is often used to model large-scale clustering of small-scale modes (e.g. [18–20]).

We will study δ -modulated power from three sources: the late-time kSZ, reionization kSZ, and residual CIB. The model predicts the following contribution to C_L^{KK} :

$$C_L^{KK} = \int dz \frac{H(z)}{\chi(z)^2} \beta(z)^2 \left(\frac{\partial \bar{K}}{\partial z} \right)^2 P_\delta \left(\frac{l}{\chi(z)}, z \right) \quad (13)$$

so in each of the three cases, we will need to know the redshift distribution $d\bar{K}/dz$ of the small-scale power, and the bias-like parameter β . Considering $d\bar{K}/dz$ first, our fiducial model for the kSZ has already been described, and for the CIB we conservatively assume residual contamination with total power spectrum equal to the late-time kSZ, and Gaussian redshift distribution with $(\bar{z}, \sigma_z) = (2, 1)$.

Considering β next, assigning precise values would require dedicated simulations beyond the scope of this paper, but we will make rough estimates as follows. In the limit of high l , the kSZ power spectrum is 1-halo dominated (or during reionization, 1-bubble dominated). In this limit, the locally generated kSZ power is simply proportional to the number density of sources, and therefore the parameter β is equal to the usual linear bias b . For the late-time kSZ and CIB, we will take $\beta \approx 1.5$ as a fiducial value. For the reionization kSZ, we will take $\beta \approx 6$, a typical bubble bias from simulations [21].

In Fig. 3, bottom panel, we show the clustering contributions to C_L^{KK} from δ -modulated late-time kSZ, reionization kSZ, and residual CIB. In all three cases, the modelling is approximate but should give a rough estimate for the size of the non-Gaussian clustering effect. We find that the clustering terms are small compared to our main

signal plus reconstruction noise, and have fairly different L -dependence so that there is little degeneracy. This result makes intuitive sense: clustering terms proportional to P_v dominate on large scales, since terms proportional to P_δ are suppressed by positive powers of k . It is also consistent with the excellent agreement between the η -model and the 3D simulations seen in the top panel of Fig. 3.

In summary, the η -model predicts a large kSZ signal in C_L^{KK} , that this prediction agrees with simulations, and that it is robust to a wide range of possible contaminants.

III. FORECASTS AND DISCUSSION

In this section, we will present signal-to-noise forecasts. The first type of forecast we will consider is “single-bin detection”: total signal-to-noise of the kSZ C_L^{KK} summed over all source redshifts, marginalized over an arbitrary constant δC_L^{KK} as previously described. In our fiducial model, a single-bin detection would get 86% of its signal-to-noise from reionization and 14% from the late-time kSZ. Therefore a single-bin detection of C_L^{KK} at the expected level would be strong evidence for patchy reionization in the CMB.

The next observational milestone would be a “two-bin detection”, in which we fit for the overall amplitude of the high- z signal ($z \geq 4$), marginalized over an arbitrary multiple of the low- z contribution ($z \leq 4$). A two-bin detection would measure the amplitude of the patchy reionization signal without any assumptions on the low- z amplitude.

Finally, we consider a “three-bin detection”: detection significance of the $z \geq 5$ contribution, marginalized over independent redshift bins with $0 \leq z \leq 2.5$ and $2.5 \leq z \leq 5$. A three-bin detection would establish the bimodal redshift dependence of the kSZ sources, with peaks at late time and during reionization, and little or no power in between, which would also provide a powerful check on systematics.

The precise definitions are as follows: given N contributions $(\delta C_L^{KK})_1, \dots, (\delta C_L^{KK})_N$ to C_L^{KK} such that $C_L^{KK} = \sum_i A_i (\delta C_L^{KK})_i$ (for example N redshift bins), we forecast signal-to-noise on the amplitudes A_i by computing the N -by- N Fisher matrix

$$F_{ij} = \frac{f_{\text{sky}}}{2} \sum_{L=L_{\text{min}}}^{L_{\text{max}}} (2L+1) \frac{(\delta C_L^{KK})_i (\delta C_L^{KK})_j}{(C_L^{KK})_{\text{tot}}^2} \quad (14)$$

The signal-to-noise of $(\delta C_L^{KK})_i$, marginalizing over signals $j \neq i$, is given by $S/N = (F_{ii}^{-1})^{-1/2}$. The significance of our “ N -bin detection”, where $N = 1, 2, 3$, is defined to be the signal-to-noise of the highest redshift bin, taking $(C_L^{KK})_{\text{tot}}$ to be the sum of contributions from the lower redshift bins plus reconstruction noise, and marginalized over the other redshift bins plus a contribution of the form $\delta C_L^{KK} = \text{constant}$. The maximum multipole L_{max} in Eq. (14) will depend in practice on

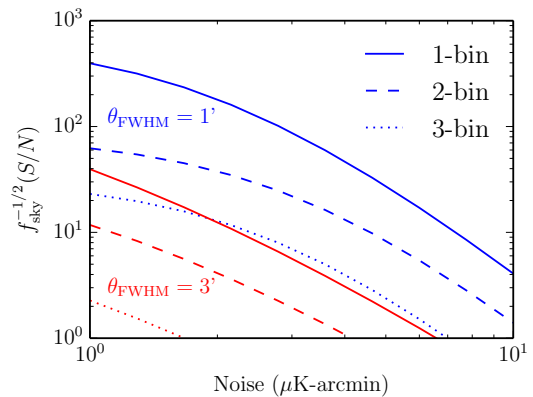


FIG. 4: Forecasted S/N for 1-bin, 2-bin, and 3-bin detections as defined in the text, for varying noise level, and beam size $\theta_{\text{FWHM}} = 1'$ (blue/upper curves) or $\theta_{\text{FWHM}} = 3'$ (red/lower curves).

the extent to which secondary contributions to C_L^{KK} can be modelled as L increases. As a fiducial value, we have used $L_{\text{max}} = 300$ here.

In Fig. 4 we show forecasted signal-to-noise as functions of instrumental noise level and beam size. These results include improvements from a generalization of the Fisher matrix in Eq. (14) in which multiple K -fields are defined corresponding to bins in CMB wavenumber l . However, we find that the only case where this significantly improves signal-to-noise is the three-band detection with $\theta_{\text{FWHM}} \lesssim 2'$.

It is seen that the signal-to-noise is a steep function of noise level, favoring a deep small-field observing strategy. However, for surveys smaller than a few hundred square degrees, there is a signal-to-noise penalty beyond the simple f_{sky} scaling in Fig. 4, since C_L^{KK} cannot be measured on super-survey scales $L \leq L_{\text{min}} = (2\pi/\theta_{\text{surv}})$. This signal-to-noise penalty ranges from 5–10% for a 1000 deg^2 survey, and 25–50% for a 100 deg^2 survey, depending on the noise level and forecast chosen.

Including this penalty, some example surveys which achieve N -bin detections are as follows. A 3σ one-bin detection can be achieved by a survey with area $A = 500 \text{ deg}^2$, noise $\Delta_T = 4 \mu\text{K-arcmin}$, and beam $\theta_{\text{FWHM}} = 1.4 \text{ arcmin}$. Likewise two-bin and three-bin 3σ detections can be achieved by surveys with $(A, \Delta_T, \theta_{\text{FWHM}}) = (900, 3, 1)$ and $(2400, 2, 1)$ respectively. An ambitious future survey with $1 \mu\text{K-arcmin}$ noise, 1 arcmin beam, and $f_{\text{sky}} = 0.5$ can achieve 1-bin, 2-bin, and 3-bin detections with significance 279σ , 44σ , and 16σ .

In this paper, we have identified a new non-Gaussian signal in the CMB which is a distinctive observational signature of the kSZ effect. It should soon be detectable, and an exciting milestone will be a “clean” detection of patchy reionization, with minimal assumptions on modelling of other CMB secondaries. Future experiments such as CMB-S4 will have sufficient signal-to-noise to measure the signal with more granularity and constrain

the redshift and wavenumber dependence of the kSZ sources, opening up a qualitatively new observational window on the epoch of reionization.

Acknowledgements. We thank Joel Meyers and Alex van Engelen for discussions and initial collaboration. We are also grateful to Nick Battaglia, Emmanuel Schaan and David Spergel for discussion. KMS was supported by an NSERC Discovery Grant and an Ontario Early

Researcher Award. SF was supported by the Miller Institute at the University of California, Berkeley. Some computations were done on the Scinet cluster at the University of Toronto. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation.

-
- [1] J. P. Ostriker and E. T. Vishniac, *Astrophys. J.* **306**, L51 (1986).
- [2] R. A. Sunyaev and Ya. B. Zeldovich, *Ann. Rev. Astron. Astrophys.* **18**, 537 (1980).
- [3] R. A. Sunyaev and Ya. B. Zeldovich, *Comments Astrophys. Space Phys.* **4**, 173 (1972).
- [4] N. Battaglia, A. Natarajan, H. Trac, R. Cen, and A. Loeb, *Astrophys. J.* **776**, 83 (2013), 1211.2832.
- [5] M. McQuinn, S. R. Furlanetto, L. Hernquist, O. Zahn, and M. Zaldarriaga, *Astrophys. J.* **630**, 643 (2005), astro-ph/0504189.
- [6] H. Park *et al.*, *Astrophys. J.* **769**, 93 (2013), 1301.3607.
- [7] M. A. Alvarez, *Astrophys. J.* **824**, 118 (2016), 1511.02846.
- [8] O. Zahn *et al.*, *Astrophys. J.* **756**, 65 (2012), 1111.6386.
- [9] A. Lewis and A. Challinor, *Phys. Rept.* **429**, 1 (2006), astro-ph/0601594.
- [10] E. Calabrese *et al.*, *JCAP* **1408**, 010 (2014), 1406.4794.
- [11] S. R. Furlanetto, (2015), 1511.01131.
- [12] M. F. Morales and J. S. B. Wyithe, *Ann. Rev. Astron. Astrophys.* **48**, 127 (2010), 0910.3010.
- [13] C.-P. Ma and J. N. Fry, *Phys. Rev. Lett.* **88**, 211301 (2002), astro-ph/0106342.
- [14] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000), astro-ph/9911177.
- [15] L. D. Shaw, D. H. Rudd, and D. Nagai, *Astrophys. J.* **756**, 15 (2012), 1109.0553.
- [16] M. Fukugita and P. J. E. Peebles, *Astrophys. J.* **616**, 643 (2004), astro-ph/0406095.
- [17] A. Benoit-Levy, K. M. Smith, and W. Hu, *Phys. Rev.* **D86**, 123008 (2012), 1205.0474.
- [18] M. Takada and W. Hu, *Phys. Rev.* **D87**, 123504 (2013), 1302.6994.
- [19] I. Mohammed and U. Seljak, *Mon. Not. Roy. Astron. Soc.* **445**, 3382 (2014), 1407.0060.
- [20] T. Baldauf, U. Seljak, L. Senatore, and M. Zaldarriaga, (2015), 1511.01465.
- [21] M. A. Alvarez, E. Komatsu, O. Dore, and P. R. Shapiro, *Astrophys. J.* **647**, 840 (2006), astro-ph/0512010.