

Simulated multitracer analyses with HI intensity mapping

A. Witzemann,¹★ D. Alonso,^{2,3} J. Fonseca^{1b} and M. G. Santos^{1,4}

¹*Department of Physics and Astronomy, University of the Western Cape, Cape Town 7535, South Africa*

²*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff CF24 3AA, UK*

³*Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

⁴*SKA SA, The Park, Park Road, Pinelands 7405, South Africa*

Accepted 2019 March 12. Received 2018 December 21; in original form 2018 August 9

ABSTRACT

We use full sky simulations, including the effects of foreground contamination and removal, to explore multitracer synergies between a SKA-like 21cm intensity mapping survey and an LSST-like photometric galaxy redshift survey. In particular we study ratios of auto and cross-correlations between the two tracers as estimators of the ratio of their biases, a quantity that should benefit considerably from the cosmic variance (CV) cancellation of the multitracer approach. We show how well we should be able to measure the bias ratio on very large scales (down to $\ell \sim 3$), which is crucial to measure primordial non-Gaussianity and general relativistic effects on large-scale structure. We find that, in the absence of foregrounds but with realistic noise levels of such surveys, the multitracer estimators are able to improve on the sensitivity of a cosmic variance contaminated measurement by a factor of 2–4. When foregrounds are included, estimators using the 21cm autocorrelation become biased. However, we show that cross-correlation estimators are immune to this and do not incur in any significant penalty in terms of sensitivity from discarding the autocorrelation data. However, the loss of long-wavelength radial modes caused by foreground removal in combination with the low-redshift resolution of photometric surveys, reduces the sensitivity of the multitracer estimator, albeit still better than the cosmic variance contaminated scenario even in the noise free case. Finally we explore different alternative avenues to avoid this problem.

Key words: large-scale structure of Universe, methods: data analysis, techniques: photometric, radio lines: general.

1 INTRODUCTION

Probing the physics of the primeval Universe is one of the main drivers for observational studies of the cosmos. The Gaussianity of the primordial cosmological perturbations remains an open question, answering it will provide great insight into the details of the dynamics of the very early Universe. The current state of the art are the Planck bounds derived from the Cosmic Microwave Background (CMB), for local-type primordial non-Gaussianity (PNG) they read $f_{\text{NL}} \simeq 0.8 \pm 5.0$ (Planck Collaboration 2016). Furthermore, local PNG introduces a scale dependence in the bias between the dark matter (DM) haloes and the astrophysical objects used as tracers of the matter distribution (Dalal et al. 2008; Matarrese & Verde 2008).

This scale dependence becomes relevant on large cosmological (horizon) scales. At the same time, general-relativistic effects become important on such ultralarge scales (past the matter-radiation equality peak), opening the possibility of testing the theory of gravity in this new regime and finding possible hints of deviations

to General Relativity (GR; for a comprehensive review on ‘GR effects’ see e.g. Bonvin & Durrer 2011; Challinor & Lewis 2011; Bonvin 2014). The next generations of Large Scale Structure (LSS) surveys such as the Square Kilometer Array (SKA),¹ Euclid², and the Large Synoptic Survey Telescope (LSST),³ promise to be able to target such effects by observing ever larger volumes of the Universe. Several forecasts show that they will improve on the Planck constraint on PNG (see e.g. Giannantonio et al. 2012; Camera et al. 2013; Alonso et al. 2015b; Camera, Santos & Maartens 2015; Raccanelli et al. 2016). Despite the improvements, errors on local PNG from single tracers of the matter distribution will still be unable to push $\sigma(f_{\text{NL}})$ below 1, which is vital to distinguishing between single field and multifield inflation (see e.g. de Putter, Gleyzes & Dor 2017).

The crucial limitation on these surveys comes from cosmic variance, which limits measurements on the largest scales. A decade

¹www.skatelescope.org.

²www.euclid-ec.org.

³www.lsst.org.

* E-mail: amadeus.witzemann@gmx.at

ago Seljak (2009) proposed a statistical method, often referred to as the multitracer technique, to overcome cosmic variance (see also McDonald & Seljak 2009; Hamaus, Seljak & Desjacques 2011; Abramo & Leonard 2013). The basic idea is that, by comparing two tracers, we can measure the ratio of their biases without requiring to measure the underlying DM distribution they trace. This avoids cosmic variance, caused by the stochasticity in the particular realization of the matter distribution we observe. This possibility also shifts the target set-up of future surveys to probe these large-scale effects, since smaller volumes thanks to lower noise (e.g. large integration times or higher number densities) may be preferred over larger volumes that provide sufficient samples of the modes of interest (as long as such smaller volumes include the target scales).

Several authors have extensively used the technique to forecast how combinations of future surveys and different DM tracers will impact on the prospects of measuring f_{NL} as well as other horizon-scale GR effects (Yoo et al. 2012; Ferramacho et al. 2014; Yamauchi, Takahashi & Oguri 2014; Alonso & Ferreira 2015; Fonseca et al. 2015; Abramo & Bertacca 2017; Fonseca, Maartens & Santos 2017, 2018; Schmittfull & Seljak 2018). While some combinations do not break the $\sigma(f_{\text{NL}}) < 1$ threshold, others have the potential to provide transformational constraints. Such a technique thus greatly extends the potential to probe the physics of inflation and GR with near-future experiments.

Despite the plethora of works studying the potential and applicability of the multitracer technique, little has been done to test and assess the performance of the technique within realistic observational settings for future surveys (although the technique has been employed in some analysis of current data, Blake et al. 2013; Ross et al. 2014; Marín et al. 2016, with an emphasis on redshift-space distortions). Questions on what estimators to use and whether they will be biased by contaminants still remain unanswered. This paper attempts to examine the potential of the multitracer technique in a realistic analysis. We will focus on the combination of an HI intensity mapping (IM) survey carried out by a SKA-like facility (Santos et al. 2015) with an LSST-like photometric galaxy survey (LSST Science Collaboration 2009). This combination is a natural choice since both surveys will observe the largest cosmological volumes in an overlapping region of the sky in both the radio and optical/infrared regimes. Moreover, such surveys will be affected by different sky systematics.

The detection of the notoriously faint HI 21cm signal is observationally challenging. While the emission of individual sources is only likely to be detectable at comparably low redshifts, a technique called IM promises to probe the HI content of large volumes (Battye, Davies & Weller 2004; Chang et al. 2008; Switzer et al. 2013). This method removes the constraint of resolving sources, and instead measures the integrated emission of all objects within a wide angular pixel. While trading angular resolution for larger survey volumes, intensity maps of the 21cm emission are also contaminated by several sources of foreground radiation, which are often many orders of magnitude brighter than the cosmological signal. Much research has already been done on this topic (e.g. Di Matteo et al. 2002; Oh & Mack 2003; Santos, Cooray & Knox 2005; Jelić et al. 2010; Moore et al. 2013; Wolz et al. 2014; Shaw et al. 2015) and the main contributors have been identified as galactic synchrotron emission, free-free emission from galactic and extragalactic origin and point sources. A discussion of cleaning methods for all of these foreground sources can be found in Alonso, Ferreira & Santos (2014). While the residuals after foreground removal should be small, much of the large-scale radial

information is removed from the cosmological signal along with the foregrounds.

On the other hand, optical galaxy surveys will be affected by galactic dust extinction and star contamination, as well as several observational systematics, which can affect the observed clustering on large scales (Ross et al. 2011). Cross-correlations between HI intensity and galaxy number count maps are free of systematic effects relevant for only one of the surveys. This is expected to be the case for most foreground residuals, but there could be non-negligible effects of foreground point sources due to thermal dust, which is discussed in the context of intensity maps of carbon monoxide and ionized carbon fine structure emission in Switzer (2017) and Pullen et al. (2018). Apart from improving the reliability of measurements (e.g. Masui et al. 2013; Poursidou, Bacon & Crittenden 2017), cross-correlations can also be used to constrain HI properties (Poursidou et al. 2016) or to calibrate photometric redshift estimates (Schulz 2010).

We explore the multitracer technique in the presence of foregrounds in the HI intensity maps using realistic simulations of the observational process. For this purpose we construct estimators of the bias ratios and assess their performance and error. For simplicity we neglect the presence of PNG on the tracer biases, making the bias ratios scale independent, which is expected to have a negligible impact on the estimator errors. We will focus on IM foregrounds, which are likely to be the main contaminant, and neglect the effects of possible systematics in the optical data. Crucially, we wish to determine how sensitive the cancellation of cosmic variance is to IM foreground cleaning and to the observational specifications of each experiment.

The paper is organized as follows: in Section 2 we discuss possible multitracer estimators that can be used to extract the bias ratio of the two tracers and in particular focus on estimators that can be free from foreground or systematic contamination. In Section 3 we describe the simulations done for both experiments (SKA1-MID and LSST) and the foreground cleaning method. In Section 4 we discuss the results, addressing the performance and error sources of the estimators and possible biases. In particular, we discuss the limitations of the current approach and show possible avenues to improve on this technique. We conclude in Section 5.

2 MULTITRACER ESTIMATORS

2.1 Signal modelling

Our basic observable is the projected fluctuation of a given tracer of the matter distribution on the sky $\Delta(\hat{n})$. Under the assumption that, on sufficiently large scales, $\Delta(\hat{n})$ is linearly related to the matter overdensity $\delta_{\text{M}}(t, \mathbf{x})$, the relation between both quantities can be modelled as

$$\Delta(\hat{n}) = \int_0^\infty dz b(z) \phi(z) \delta_{\text{M}}(t(z), \chi(z)\hat{n}), \quad (1)$$

where χ is the comoving radial distance, and $b(z)$ and $\phi(z)$ are the bias and selection functions associated with this tracer. For simplicity we have neglected the contributions from redshift-space distortions, magnification, and other relativistic effects (Bonvin & Durrer 2011; Challinor & Lewis 2011; Bruni et al. 2012; Jeong, Schmidt & Hirata 2012; Yoo et al. 2012; Hall, Bonvin & Challinor 2013; Yoo & Desjacques 2013). This simplifying approximation should not have any significant impact on the final results presented here, since Redshift Space Distortions (RSD) are suppressed by the broad redshift kernels used and all other terms are highly

subdominant (Alonso & Ferreira 2015; Alonso et al. 2015a; Fonseca et al. 2015; Yoo & Seljak 2015).

Given two tracers a and b , the angular cross-power spectrum is defined as the two-point function of their harmonic coefficients, and can be related to the matter power spectrum $P(k, z)$ as (Di Dio et al. 2014)

$$C_\ell^{ab} \equiv \langle \Delta_{\ell m}^a \Delta_{\ell m}^{b*} \rangle = \frac{2}{\pi} \int_0^\infty dk k^2 W_\ell^a(k) W_\ell^b(k), \quad (2)$$

$$W_\ell^a(k) = \int dz b_a(z) \phi_a(z) j_\ell(k\chi(z)) \sqrt{P(k, z)}. \quad (3)$$

Under Limber's approximation (Limber 1954; Loverde & Afshordi 2008), this expression can be simplified to

$$C_\ell^{ab} = \int d\chi \frac{b_a \phi_a b_b \phi_b H^2(\chi)}{\chi^2} P\left(z(\chi), k = \frac{\ell + 1/2}{\chi}\right), \quad (4)$$

where H is the expansion rate.

In this analysis we have used two different types of tracers: the overdensity of galaxy number counts, which we will label as Δ^g , and the temperature fluctuations in the 21cm line emission caused by neutral hydrogen (HI), Δ^H . In the case of galaxy clustering, we approximate the linear galaxy bias as $b^g = 1 + 0.84z$ (LSST Science Collaboration 2009), which is an estimate of the results from Weinberg et al. (2004). On the other hand, as described in Section 3.2, the presence of spectrally smooth radio foregrounds makes it infeasible to measure the average 21cm brightness temperature \bar{T}_{21} , and it is therefore completely degenerate with the linear bias function associated with this tracer: $b_H(z) = \bar{T}_{21}(z) b_{H1}(z)$, where b_{H1} is the linear clustering bias associated with the cosmic overdensity of neutral hydrogen.⁴ We model both quantities after Bull et al. (2015).

Finally, the observed fluctuations Δ^a are inevitably contaminated by noise. In the case of galaxy clustering, this is associated with shot noise due to the discrete nature of the sources used to reconstruct the true underlying distribution. In this case, the noise power spectrum is simply given by the inverse number density of tracer sources in units of Sr^{-1} ,

$$N_\ell^{gg} = \frac{1}{\bar{n}}. \quad (5)$$

For 21cm, the combination of instrumental noise and beam smoothing, caused by the telescope's finite size, effectively erases all modes below the telescope resolution. For an angular Gaussian beam, the harmonic coefficients of the beam, multiplying the signal in harmonic space, can be simply modelled as

$$B_\ell = \exp\left(-\frac{\ell(\ell+1)\theta_{\text{FWHM}}^2}{16 \log 2}\right), \quad (6)$$

where θ_{FWHM} is the beam full width at half-maximum (FWHM) at a given frequency. The instrumental noise can then be modelled as an additive Gaussian random field with flat power spectrum. For single-dish observations, this is simply given by (Bull et al. 2015)

$$N_\ell^{\text{HH}} = \frac{T_{\text{sys}}^2 4\pi f_{\text{sky}}}{N_{\text{dish}} \Delta\nu t_{\text{tot}}}. \quad (7)$$

Here T_{sys} is the system temperature, f_{sky} is the total observed sky fraction, N_{dish} is the number of dishes in the instrument, t_{tot} is the

total integration time, and $\Delta\nu$ is the frequency bandwidth for the particular sky map under consideration.

It is worth noting that we assume no cross-noise term between galaxies and HI. This is expected to be present if the HI-emitting star-forming galaxies form a significant fraction of the galaxy sample, however we assume this shot noise contribution to be subdominant. We also neglect any correlated $1/f$ -like noise component for intensity mapping. We refer the reader to Harper et al. (2018) for a more detailed discussion of correlated noise in the context of foreground contamination and removal.

2.2 The surveys

Our forecasts focus on the combination of 21cm intensity maps, constructed from the SKA data, with optical observations of the galaxy distribution as could be achieved by LSST. We describe the models used for both data sets here.

We assume the first phase of SKA (in particular SKA-1 MID; Santos et al. 2015) to consist of 197 dishes, which will use a total of $t_{\text{tot}} = 10,000$ h integration time to produce intensity maps covering ~ 60 per cent of the sky. We assume a combination of single-dish surveys carried out with band 1 and band 2 receivers, and we use a frequency range of $\nu \in (390, 1300)$ MHz, corresponding to a redshift interval $0.1 \leq z \leq 2.65$. Since we work with individual redshift bins at a time, our results are always valid for the receiver type that covers the relevant redshift range. We will assume single-dish observations, which are limited in angular resolution by a beam that we model as Gaussian with an FWHM given by $\theta_{\text{FWHM}} = 1.22 \lambda / D_{\text{dish}}$, where λ is the observed wavelength and D_{dish} is the dish diameter. We assume a diameter $D_{\text{dish}} = 14.5$ m,⁵ which corresponds to an angular resolution of $\ell \lesssim 90$ in multipoles and a beam FWHM of ~ 2 deg at redshift $z = 1$. Finally, we add white noise as described in the previous section, with a smoothly varying system temperature T_{sys} following the values given in Square Kilometre Array Cosmology Science Working Group (2018). Further particulars regarding the specific simulated intensity maps used in this analysis are described in Section 3.1

For LSST, we use the redshift distribution modelled in Alonso et al. (2015a), which yields an integrated number density of 43 galaxies per arcmin², in agreement with LSST Science Collaboration (2009). As described in Section 3.1, we do not make a precise modelling of the photometric redshift accuracy that LSST will achieve, and instead work with redshift bins wide enough ($\Delta z = 0.1$) to simulate the loss of small radial scales. We do this in order to facilitate the interpretation of the autocorrelation and cross-correlation estimators presented in the next section. A more realistic treatment would either account for the difference in radial window function between the 21cm and optical bins, or re-weight the 21cm frequency channels contributing to each bin to mimic the photo- z window function as closely as possible.

We assume almost complete overlap between SKA and LSST, given their common observable sky. After accounting for contamination from galactic synchrotron (radio) and dust (optical), the final common footprint, displayed in Fig. 1, covers 41 per cent of the sky.

⁴Note that although this is the case in our analysis, there do exist ways to extract the average HI brightness temperature, for example using cross-correlations, or HI galaxy surveys (Wolz, Blake & Wyithe 2017).

⁵SKA-1 MID will consist of a combination of 15 and 13.5 m dishes, and we use 14.5 as an approximation to the mean dish diameter. This choice should not affect the final results of this study.




Figure 1. Sky mask used in our analysis, shown in Mollweide’s projection and equatorial coordinates. The masked area is shown in grey. The footprint corresponds to the sky observable from the LSST and SKA with the regions of highest galactic emission (both in synchrotron and dust) removed. The total unmasked area is 16900 deg^2 ($f_{\text{sky}} = 0.41$).

2.3 The estimators

Under the assumption that the bias functions vary slowly over the support of the selection functions, and in the limit where the selection functions for both tracers are the same ($\phi^g = \phi^H \equiv \phi$), the three different auto and cross-power spectra described in Section 2.1 can be written as

$$\begin{aligned} C_\ell^{gg} &= b_g^2 C_\ell + N_\ell^{gg}, \\ C_\ell^{Hg} &= b_g \bar{T}_{21} b_{H1} B_\ell C_\ell, \\ C_\ell^{HH} &= \bar{T}_{21}^2 b_{H1}^2 B_\ell^2 C_\ell + N_\ell^{HH}, \end{aligned} \quad (8)$$

where C_ℓ is the angular power spectrum of the matter overdensity projected along the line of sight with ϕ . At smaller scales both b_g and b_{H1} are generally thought to be scale dependent and an additional cross-shot noise term would have to be included in C_ℓ^{Hg} , (Wolz et al. 2017; Anderson et al. 2018; Villaescusa-Navarro et al. 2018). The focus on large scales motivates our neglect of these effects throughout this work.

On a realization-by-realization basis, the measured values of these quantities will be subject to sample variance, due to the stochastic nature of both the underlying matter fluctuations and the instrumental and shot noise. For signal-dominated modes, the realization-dependent fluctuations will coincide for the three power spectra, and therefore it is possible to constrain certain parameters beyond the limit imposed by sample variance if only a single tracer was available (Seljak 2009). One obvious example of this is the ratio of the tracer bias functions, which in an ideal noiseless case could be measured exactly by taking ratios of the power spectra above. In this work we will focus on the quantity

$$\epsilon \equiv \frac{b_{H1} \bar{T}_{21}}{b_g}, \quad (9)$$

for which we propose two different estimators:

$$\hat{\epsilon}_{A,\ell} \equiv \sqrt{\frac{\hat{C}_\ell^{HH} - N_\ell^{HH}}{B_\ell^2 (\hat{C}_\ell^{gg} - N_\ell^{gg})}}, \quad (10)$$

$$\hat{\epsilon}_{X,\ell} \equiv \frac{\hat{C}_\ell^{Hg}}{B_\ell [\hat{C}_\ell^{gg} - N_\ell^{gg}]}, \quad (11)$$

where all hatted quantities (e.g. \hat{C}_ℓ^{HH}) are measurements in a given realization. In addition to this, we will also consider a third estimator making use of both the auto and cross-correlation, which combines

$\epsilon_{A,X}$ in an inverse variance-weighted manner:

$$\epsilon_{\text{opt}} = \frac{\sum_{i,j} C_{ij}^{-1} \epsilon_j}{\sum_{ij} C_{ij}^{-1}}, \quad (12)$$

where C is the covariance matrix of the two previous estimators computed from simulations.

These three estimators can be understood as different limits of a more general maximum-likelihood estimator combining the three cross-correlations simultaneously, which allow us to explore the impact of foreground contamination in the 21cm maps.

3 SIMULATED FORECASTS

3.1 The simulations

We produce synthetic signal simulations of both the galaxy distribution and 21cm maps using the publicly available code COLORE.⁶ COLORE efficiently generates intensity maps for any arbitrary line-emitting species and source catalogues tracing the same DM distribution (with their respective biases $b_{H1}(z)$ and $b_{\text{gal}}(z)$). COLORE first generates a Gaussian realization of the linearized density field at $z = 0$ along with the corresponding linear radial velocity field. It then linearly evolves density and velocity to the redshift of each grid point in the simulation and produces a 3D cube of the physical matter density in the light-cone using a lognormal transformation (described in e.g. Coles & Jones 1991).⁷ For the galaxy sample the density field is biased and then Poisson-sampled using the galaxy number density $N(z)$. For 21cm, the density field is used to generate a biased HI density, which we then interpolate into spherical shells that we output as sky maps. For simplicity, we switch off the effect of redshift-space distortions, and therefore the redshift of each source is calculated without accounting for the local velocity field. We simulate a cubic box with 2048^3 Cartesian grid points and a length large enough to encompass the comoving volume to redshift $z = 2.7$. This yields a grid resolution of $\Delta x \simeq 4 h^{-1} \text{ Mpc}$. The initial Gaussian density field is smoothed with a Gaussian kernel of size $R_G = 5 h^{-1} \text{ Mpc}$ to avoid grid artefacts as well as the non-linear distortions induced by the lognormal transformation. This scale is significantly smaller than those we focus on, or than the SKA beam, and therefore the impact of this smoothing on our results is negligible.

We generate 21cm intensity maps with a frequency resolution of $\Delta\nu = 1 \text{ MHz}$. To each of these maps we first add the simulated foreground maps, smooth them using the SKA Gaussian beam and add the instrumental noise as described above. To study the case of ideal noise-free cosmic variance cancellation we also simulate equivalent maps of the galaxy overdensity without shot noise. We simulate these as an alternative IM species with unit mean temperature and a bias given by the galaxy bias. The foregrounds are simulated using FORGET, part of the publicly available CRIME package⁸ (Alonso et al. 2014). We consider four unpolarized foreground sources, including galactic synchrotron, galactic and extragalactic free-free emission, and extragalactic point sources.

⁶<https://github.com/damonge/CoLoRe>.

⁷Note that COLORE is also able to produce physical density fields through other more accurate methods (e.g. first- and second-order Lagrangian perturbation theory, but we chose the lognormal for simplicity and performance reasons). This choice should be irrelevant given that our analysis focuses on relatively large scales.

⁸<https://github.com/damonge/CRIME>.

From these outputs we produce maps of the 21cm temperature fluctuations and of the galaxy overdensity on thin radial bins with an equivalent frequency width $\Delta\nu = 1$ MHz. After the foreground cleaning stage, described in Section 3.2, the resulting 21cm maps are merged to thicker bins with a width of $\Delta z = 0.1$, and the same is done to estimate the galaxy overdensity in bins of the same width.

Finally, in order to study the statistical properties of our estimators, we generate $N_{\text{sim}} = 200$ simulations of the DM background, using different seeds for the Gaussian density field. Each simulation is populated with the H I and galaxy distributions, using different seeds for the noise realization and foreground maps. All simulations assume a Λ CDM cosmological model with parameters $(\Omega_M, \Omega_b, n_s, \sigma_8, h) = (0.3, 0.05, 0.96, 0.8, 0.7)$.

Fig. 2 shows simulated maps of the galaxy overdensity (left) and the H I temperature (right) using this procedure at a redshift $z \simeq 0.35$. Both maps are very strongly correlated, and display similar structures. This tight correlation is the basis for the cosmic variance cancellation implicit in multitracer studies.

3.2 Foreground removal

Foreground removal methods for 21cm intensity mapping (Chapman et al. 2013; Shaw et al. 2014, 2015; Wolz et al. 2014; Alonso et al. 2015a; Zuo et al. 2019) try to separate the cosmological and foreground signals by making use of their different spectral properties: while foregrounds are expected to have a smooth dependence with frequency, which should also be highly correlated across the sky, the cosmological signal follows the large-scale structure, and therefore contains power across a large range of Fourier scales (both in frequency and angles).

Let \mathbf{d} be a vector containing our measurements of the brightness temperature along a fixed line of sight. In general it will contain contributions from foregrounds \mathbf{f} , cosmological signal \mathbf{c} and instrument noise \mathbf{n} :

$$\mathbf{d} = \mathbf{f} + \mathbf{c} + \mathbf{n} = \mathbf{f} + \mathbf{s}, \quad (13)$$

where we have grouped all noise-like components into $\mathbf{s} \equiv \mathbf{c} + \mathbf{n}$. Most foreground removal methods recover an estimate of \mathbf{s} by linearly filtering the data:

$$\mathbf{s}_c = \mathbf{W} \cdot \mathbf{d}, \quad (14)$$

using a filter \mathbf{W} that minimizes the presence of foreground residuals on \mathbf{s} . For instance, principal component analysis (PCA) corresponds to a filter $\mathbf{W} = \mathbf{1} - \mathbf{U}_{\text{PC}}$, where \mathbf{U}_{PC} is the matrix of principal eigenvectors of the data covariance matrix. As another example, a linear fit to a set of smooth functions of frequency, stored in the columns of a matrix \mathbf{A} , would correspond to a choice of filter

$$\mathbf{W} = \mathbf{1} - \mathbf{A} (\mathbf{A}^T \mathbf{S} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{S}^{-1}, \quad (15)$$

where \mathbf{S} is the covariance of \mathbf{s} .

After filtering, the cleaned signal

$$\mathbf{s}_c = \mathbf{W} \mathbf{s} + \mathbf{W} \mathbf{f} \quad (16)$$

will contain both a version of the original signal where typically the longer wavelength radial modes have been downweighted ($\mathbf{W} \mathbf{s}$), as well as foreground residuals ($\mathbf{W} \mathbf{f}$), unless a perfect knowledge of the foreground spectral behaviour can be achieved. This has two main consequences when it comes to using \mathbf{s}_c for cosmology:

(i) Unless foregrounds have been perfectly removed (which is never the case), the autocorrelation of the 21cm data will be contaminated by foreground residuals that must be marginalized

over (unless we can convince ourselves that their amplitude lies below the noise level at the relevant length-scales).

(ii) Even when cross-correlating with other tracers of the large-scale structure, the loss of radial modes implied by the filter \mathbf{W} must be taken into account and corrected for in the model for the cross-correlation.

The first effect is inherent to 21cm autocorrelations, and can only be overcome if the residual contamination is sufficiently small, or if a sufficiently accurate foreground model can be built to marginalize over their contribution. However, since we always know the filter \mathbf{W} used by the foreground cleaning pipeline, the second effect can be modelled and taken into account. In general, the action of \mathbf{W} will be to remove power from the largest radial scales, thus reducing the overall amplitude of any projected clustering statistic. Characterizing this reduction exactly requires a full model of the 3D power spectrum, however we will take a simpler approximate method here, similar to the procedure used in e.g. Masui et al. (2013) and Switzer et al. (2013). We model the impact of \mathbf{W} on the angular power spectrum as a scale-dependent, multiplicative transfer function \mathcal{T}_ℓ , i.e.:

$$\tilde{C}_\ell^{\text{HH}} = \mathcal{T}_\ell^2 C_\ell^{\text{HH}}, \quad \tilde{C}_\ell^{\text{Hg}} = \mathcal{T}_\ell C_\ell^{\text{Hg}}. \quad (17)$$

Here \tilde{C}_ℓ and C_ℓ denote power spectra computed after foreground removal and in the absence of foregrounds, respectively, and $\tilde{C}_\ell^{\text{HH}}$ does not include the contribution from foreground residuals (i.e. it is only the autocorrelation of the first term in equation 16). We estimate the transfer function from our simulations as

$$\mathcal{T}_\ell = \frac{\langle C_\ell^{\tilde{\text{HH}}} \rangle - N_\ell^{\tilde{\text{HH}}}}{\langle C_\ell^{\text{HH}} \rangle - N_\ell^{\text{HH}}}, \quad (18)$$

where C_ℓ^{HH} is the autocorrelation of a foreground-free simulation, $C_\ell^{\tilde{\text{HH}}}$ is the cross-correlation between a foreground-cleaned and a foreground-free simulation (we have subtracted the noise bias from both power spectra), and $\langle \rangle$ denotes averaging over all simulations.

After accounting for this loss of modes, the estimators $\epsilon_{\text{A}, \ell}$ in equations (10) and (11) above become

$$\epsilon_{\text{A}, \ell} \equiv \sqrt{\frac{\hat{C}_\ell^{\text{HH}} - N_\ell^{\text{HH}}}{(\mathcal{T}_\ell B_\ell)^2 (\hat{C}_\ell^{\text{gg}} - N_\ell^{\text{gg}})}}, \quad (19)$$

$$\epsilon_{\text{X}, \ell} \equiv \frac{\hat{C}_\ell^{\text{Hg}}}{\mathcal{T}_\ell B_\ell [\hat{C}_\ell^{\text{gg}} - N_\ell^{\text{gg}}]}. \quad (20)$$

4 RESULTS

4.1 Theoretical expectation

Before we set off to use our simulations to study the feasibility of multitracer methods for IM, it is instructive to produce a theoretical estimate of the expected performance of our estimators, in order to better understand the simulated results, as well as the main sources of cosmic variance cancellation.

From the expressions for ϵ_{A} and ϵ_{X} in equations (10) and (11), we can write, for one particular realization:

$$\hat{\epsilon}_{\text{A}, \ell} = \epsilon_{\text{A}, \ell} \sqrt{\frac{1 + \Delta \hat{C}_\ell^{\text{HH}} / (C_\ell^{\text{HH}} - N_\ell^{\text{HH}})}{1 + \Delta \hat{C}_\ell^{\text{gg}} / (C_\ell^{\text{gg}} - N_\ell^{\text{gg}})}}, \quad (21)$$

$$\hat{\epsilon}_{\text{X}, \ell} = \epsilon_{\text{X}, \ell} \frac{1 + \Delta \hat{C}_\ell^{\text{Hg}} / C_\ell^{\text{Hg}}}{1 + \Delta \hat{C}_\ell^{\text{gg}} / (C_\ell^{\text{gg}} - N_\ell^{\text{gg}})}, \quad (22)$$




Figure 2. Galaxy (left) and H I map at 1050 MHz ($z \simeq 0.35$) with a redshift bin width of $\Delta z = 0.1$. We choose a low-redshift bin in order to make the tight correlation between both maps more visually apparent. Both maps trace the same DM background, and show the pure cosmological signal, before inclusion of foregrounds, beam smoothing, or noise. Note that the H I maps were generated in much thinner bins of $\Delta \nu = 1$ MHz. Noise, beam, and foreground simulations were done in these thin bins, and those maps were later merged to match the thicker bins of the galaxy maps.

where, as before, all hatted quantities (e.g. $\hat{\epsilon}_{X,\ell}$) are measurements of the equivalent non-hatted observables in a given realization, and $\Delta \hat{C}_\ell^{XY}$ is the fluctuation around the mean C_ℓ^{XY} in a given realization. Linearizing with respect to these fluctuations, we obtain:

$$\frac{\hat{\epsilon}_{A,\ell} - \epsilon_{A,\ell}}{\epsilon_{A,\ell}} \approx \frac{1}{2} \left(\frac{\Delta \hat{C}_\ell^{\text{HH}}}{C_\ell^{\text{HH}} - N_\ell^{\text{HH}}} - \frac{\Delta \hat{C}_\ell^{\text{gg}}}{C_\ell^{\text{gg}} - N_\ell^{\text{gg}}} \right), \quad (23)$$

$$\frac{\hat{\epsilon}_{X,\ell} - \epsilon_{X,\ell}}{\epsilon_{X,\ell}} \approx \frac{\Delta \hat{C}_\ell^{\text{gH}}}{C_\ell^{\text{gH}}} - \frac{\Delta \hat{C}_\ell^{\text{gg}}}{C_\ell^{\text{gg}} - N_\ell^{\text{gg}}}. \quad (24)$$

To first order, the inverse-squared signal-to-noise ratio (SNR) can be found by taking the expectation value of the square of the above quantities, obtaining:

$$\left(\frac{S}{N} \right)_{A,\ell}^{-1} = \frac{1}{2} \left[\frac{\text{Cov}_\ell^{\text{HH,HH}}}{(C_\ell^{\text{HH}} - N_\ell^{\text{HH}})^2} + \frac{\text{Cov}_\ell^{\text{gg,gg}}}{(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})^2} - 2 \frac{\text{Cov}_\ell^{\text{HH,gg}}}{(C_\ell^{\text{HH}} - N_\ell^{\text{HH}})(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})} \right]^{1/2} \quad (25)$$

$$\left(\frac{S}{N} \right)_{X,\ell}^{-1} = \left[\frac{\text{Cov}_\ell^{\text{gH,gH}}}{(C_\ell^{\text{gH}})^2} + \frac{\text{Cov}_\ell^{\text{gg,gg}}}{(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})^2} - 2 \frac{\text{Cov}_\ell^{\text{gH,gg}}}{C_\ell^{\text{gH}}(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})} \right]^{1/2}, \quad (26)$$

where $\text{Cov}_\ell^{WX,YZ} \equiv \langle \Delta \hat{C}_\ell^{WX} \Delta \hat{C}_\ell^{YZ} \rangle$. For Gaussian fields, a simplified estimate of the covariance matrix (that does not account for e.g. survey geometry) is (Knox 1995):

$$\text{Cov}_\ell^{WX,YZ} = \frac{C_\ell^{WY} C_\ell^{XZ} + C_\ell^{WZ} C_\ell^{XY}}{(2\ell + 1) f_{\text{sky}} \Delta \ell}, \quad (27)$$

where f_{sky} is the survey sky fraction and $\Delta \ell$ is the width of the C_ℓ bandpowers used in the analysis.

Substituting this result into the equations above we obtain a final expression for the theoretical SNR:

$$\left(\frac{S}{N} \right)_{A,\ell}^{-1} = \frac{1}{\sqrt{2n_\ell}} \left[\frac{(C_\ell^{\text{HH}})^2}{(C_\ell^{\text{HH}} - N_\ell^{\text{HH}})^2} + \frac{(C_\ell^{\text{gg}})^2}{(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})^2} - \frac{2(C_\ell^{\text{gH}})^2}{(C_\ell^{\text{HH}} - N_\ell^{\text{HH}})(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})} \right]^{1/2}, \quad (28)$$

$$\left(\frac{S}{N} \right)_{X,\ell}^{-1} = \frac{1}{\sqrt{n_\ell}} \left[1 + \frac{C_\ell^{\text{HH}} C_\ell^{\text{gg}}}{(C_\ell^{\text{gH}})^2} - \frac{2C_\ell^{\text{gg}}(C_\ell^{\text{gg}} - 2N_\ell^{\text{gg}})}{(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})^2} \right]^{1/2}, \quad (29)$$

where $n_\ell \equiv (2\ell + 1) f_{\text{sky}} \Delta \ell$ is the number of available modes in a given bandpower.

Inspecting equations (28) and (29), the idea of cosmic variance cancellation becomes apparent: for perfectly correlated tracers ($C_\ell^{\text{gH}} \equiv \sqrt{C_\ell^{\text{HH}} C_\ell^{\text{gg}}}$), and in the absence of noise ($N_\ell^{\text{gg}}, N_\ell^{\text{HH}} \rightarrow 0$), the negative terms in these equations, originating from the covariance between numerator and denominator in the estimators, exactly cancel the positive terms, and we obtain $(S/N)^{-1} \rightarrow 0$. This cosmic variance cancellation would not be possible if the observables entering the estimators were not strongly correlated, as would be the case if, for instance, the 21cm maps and galaxy catalogue covered non-overlapping regions of the sky. In this case, $\text{Cov}_\ell^{\text{HH,gg}} = 0$, and the SNR for a cosmic variance limited version of $\epsilon_{A,\ell}$ would read:

$$\left(\frac{S}{N} \right)_{\text{CV},\ell}^{-1} = \frac{1}{\sqrt{2n_\ell}} \left[\frac{(C_\ell^{\text{HH}})^2}{(C_\ell^{\text{HH}} - N_\ell^{\text{HH}})^2} + \frac{(C_\ell^{\text{gg}})^2}{(C_\ell^{\text{gg}} - N_\ell^{\text{gg}})^2} \right]^{1/2}. \quad (30)$$

We will make use of these theoretical estimates (equations 28, 29, and 30) in the next section to validate the results of our simulated results in the absence of foregrounds.

4.2 Foreground-free results

In order to quantify the full power of the cosmic variance cancellation in the estimators described in Section 2.3, we first explore the results from simulations without foregrounds or foreground removal, while including noise, masking, and beam smoothing. In this case, all the radial modes are present in the H I data (i.e. the transfer function is $\mathcal{T}_\ell = 1$), and can be used to constrain the bias ratio. The upper panel of Fig. 3 shows the SNR of all estimators as a function of multipole ℓ for the redshift bin centred around $z = 0.8$. For concreteness, the quantity plotted is

$$\left(\frac{S}{N} \right)_\ell = \frac{\epsilon_\ell^{\text{true}}}{\sigma_\ell}, \quad (31)$$

where

$$\epsilon_{A,\ell}^{\text{true}} \equiv \sqrt{\frac{\langle C_\ell^{\text{HH}} \rangle}{\langle C_\ell^{\text{gg}} \rangle}} \Big|_{\text{FG-free}}, \quad \epsilon_{X,\ell}^{\text{true}} \equiv \frac{\langle C_\ell^{\text{gH}} \rangle}{\langle C_\ell^{\text{gg}} \rangle} \Big|_{\text{FG-free}}, \quad (32)$$




Figure 3. The SNR for all estimators in the no foregrounds case (top) and no foregrounds, no noise, and no beam case (bottom). Results from the simulations (solid lines) give slightly lower signal to noise than the theoretical predictions (dot-dashed lines) in the upper panel from Section 4.1. Naturally, ϵ_{opt} (yellow) has the smallest variance, while ϵ_A and ϵ_X perform similarly (blue and red, respectively). All of them beat the cosmic variance estimator ϵ_{CV^+} (cyan), in the foreground-free case including noise and beam by a factor of 2–4, and in the noiseless case by a factor of 3–8. It is worth noting that little or no sensitivity is lost by discarding all 21cm autocorrelation information and using only cross-correlations (red versus orange lines).

and

$$\sigma_\ell^2 = \langle \epsilon_\ell^2 \rangle - \langle \epsilon_\ell \rangle^2. \quad (33)$$

Here, angle brackets denote averaging over all simulations. Note that we define $\epsilon_\ell^{\text{true}}$ as the value of the estimator found in foreground-free simulations, and not as the bias ratio given in equation (9). This is due to the fact that the bias functions and the background 21cm temperature vary slightly within the redshift bin, giving rise to a non-negligible scale dependence of the estimators that would be interpreted as a bias when compared with averages of ϵ over redshift, even for foreground-free simulations. For comparison, the figure also shows results for an additional estimator ϵ_{CV^+} , defined as a version of ϵ_A in which the autopower spectra of 21cm and galaxies are computed from simulations with different seeds. The aim of this estimator is to show the results that would be obtained in the absence of cosmic variance cancellation (e.g. as would be the case when trying to constrain f_{NL} from a single tracer). Note that we calculate ϵ_{CV^+} in different scenarios, also including instrumental noise, therefore it is not necessarily limited by cosmic variance.

The SNR of all estimators is shown in the top panel of Fig. 3, which shows how it should be possible to significantly increase the sensitivity within the multipole range $\ell \lesssim 100$ by a factor of up to ~ 4 with respect to the CV-dominated case. This is true for both ϵ_A and ϵ_X , which achieve very similar sensitivities. The tight correlation between both estimators implies that the improvement associated with combining both into ϵ_{opt} is mild, and that very little information is lost by using only cross-correlation information and discarding the 21cm autocorrelations. For comparison, we show the theoretical predictions derived in the previous section as dashed lines. The theory lines follow the same trends as the simulated

results, although they predict an SNR that is ~ 1.3 times higher than the simulations, owing to the approximations that go into their derivation. In all cases, no significant cosmic variance cancellation can be achieved beyond the scale of the SKA beam ($\ell \sim 100$), and the overall SNR drops significantly.

The impact of noise on cosmic variance cancellation can be further explored in a more idealized scenario, by making use of noiseless maps (i.e. simulations containing no 21cm instrumental noise or galaxy shot noise, as described in Section 3.1). The results, in terms of S/N , are shown in the lower panel of Fig. 3. Even in this idealized situation it is not possible to achieve exact cosmic variance cancellation ($S/N \rightarrow \infty$), and the relative improvement with respect to the CV-dominated case asymptotes at a factor of ~ 4 –5. This is caused by two factors: the redshift evolution of the bias functions within the relatively thick redshift bins, and the non-linear lognormal transformation used by COLORE to guarantee positive-definite density fields. Both effects produce slight differences in the galaxy and HI maps that prevent exact cosmic variance cancellation. We can only expect the impact of both effects to increase in a more realistic situation, in the presence of uncertain and scale-dependent bias relationships. As expected, the absence of noise allows this level of CV cancellation to be sustained beyond $\ell \sim 100$, in comparison with the results described above.

Although the results presented here are encouraging in terms of the large relative improvement with respect to the CV limit, their validity must be verified when foregrounds are included.

4.3 Foreground removal

To remove the foregrounds from our simulations, we use the PCA method, as described in Alonso et al. (2015a). In short, the method is based on de-projecting the principal eigenmodes of the frequency–frequency covariance matrix estimated from the data, under the assumption that those modes are the ones most contaminated by foregrounds. The level of conservativeness in the foreground removal stage can be parametrized by the number of de-projected modes, which we will refer to as the number of foreground degrees of freedom N_{FG} .

In order to estimate the number of foreground degrees of freedom that must be de-projected in our simulations, we ran the foreground removal algorithm on all of them for different values of N_{FG} . For each value, we use, as a diagnostic for foreground contamination, the relative systematic deviation in the angular power spectrum as a function of frequency and angular scale, defined as

$$\zeta_\ell(\nu) = \left\langle \frac{C_\ell^{\text{clean}}(\nu)}{C_\ell^{\text{free}}(\nu)} - 1 \right\rangle. \quad (34)$$

Here C_ℓ^{free} and C_ℓ^{clean} are the power spectra for foreground-free simulations and for simulations in which N_{FG} foreground modes have been subtracted, respectively. The optimal N_{FG} was then determined as the minimum value that achieves an acceptable degree of foreground removal over a large fraction of the ℓ – ν plane. This quantity is shown in Fig. 4 for the cases $N_{\text{FG}} = 7$ and $N_{\text{FG}} = 9$. Green colours represent a higher power spectrum with respect to the true one, and are a sign of foreground contamination, while purple areas represent lower power spectrum amplitudes and denote a loss of signal-dominated modes caused by overfitting. As mentioned in Section 3.2, the latter effect can be corrected analytically once the foreground removal transformation has been established (e.g. through the transfer function \mathcal{T}_ℓ), and therefore we seek to minimize foreground contamination. In view of the results shown in this figure, we chose to use $N_{\text{FG}} = 9$ as our fiducial value. The transfer




Figure 4. We demonstrate the effectiveness of foreground cleaning with 7 (top) and 9 degrees of freedom (bottom). The relative difference ζ is defined in equation (34). For our purposes (measuring large scales), cleaning with 7 degrees of freedom is clearly not sufficient as it leaves visible residuals on scales up to $\ell \lesssim 100$. Therefore the choice of $N_{\text{FG}} = 9$ is adopted throughout this work unless otherwise stated. The horizontal lines indicate the frequency binning used in this analysis, corresponding to a fixed width in redshift of $\Delta z = 0.1$.

function associated with this choice of N_{FG} , as defined in Section 3.2, is shown in Fig. 5 for all different redshift bins as a function of scale.

4.4 Results in the presence of foregrounds

As described in Section 3.2, the effect of foregrounds is twofold:

- (i) Foreground contamination in the autocorrelation will lead to a bias in ϵ_A that can be statistically significant;
- (ii) Foreground removal will erase some of the long-wavelength modes in the signal. This reduces the number of common modes between the foreground-cleaned intensity maps and the galaxy distribution, thereby degrading the performance of the multitracer technique.

We first quantify these two effects and then elaborate on their root causes and possible ways around them.




Figure 5. The foreground removal transfer function $\mathcal{T}_\ell(z)$ for $9 \leq \ell \leq 21$ (blue) and $54 \leq \ell \leq 66$ (orange). The transfer function shows only a mild scale dependence, but drops significantly at the edges of the redshift range, where foreground removal is less efficient (see Alonso et al. 2015a).

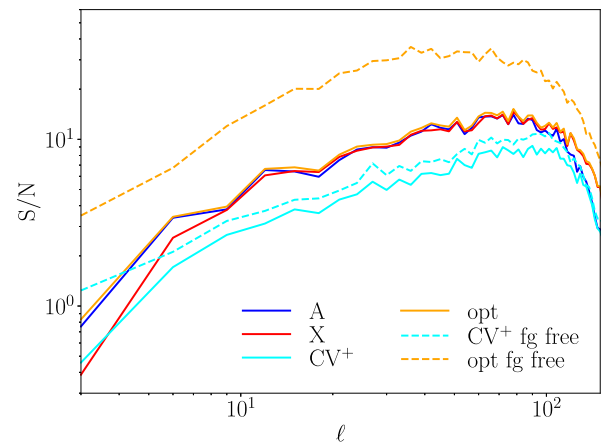


Figure 6. The SNR for all estimators in the full analysis plus ϵ_{opt} and ϵ_{CV^+} in the foreground-free case, as references. The inclusion of foregrounds in the analysis introduces a significant degradation in sensitivity, and only a slight improvement (a factor ~ 1.5) over the foreground-free cosmic variance limit (dashed cyan line) is possible.

4.4.1 Sensitivity and bias

The impact of the loss of long-wavelength modes in the method's sensitivity can be studied through the SNR defined in Section 4.2. The results are shown in Fig. 6 as solid lines for ϵ_A , ϵ_X , and ϵ_{opt} as a function of scale for a redshift bin at $z = 1$. The figure also shows the results for the cosmic variance dominated estimator ϵ_{CV^+} described in Section 4.2 as a solid cyan line. When comparing with the CV limit in the presence of foregrounds we observe that all estimators are able to improve upon ϵ_{CV^+} , although now only by a factor of ~ 2 . However, when comparing with the full constraining power in the absence of foregrounds, shown as a dashed orange line for ϵ_{opt} and as a dashed cyan line for ϵ_{CV^+} in the same figure, we observe a significant loss in S/N and that the impact of foregrounds prevents the estimators from producing a significant improvement in sensitivity with respect to an analysis without CV-cancellation (as would be the case of a single-tracer galaxy survey). We explore this effect in more detail below.

To explore the first effect described at the beginning of this section (the foreground bias), we start by defining the ‘bias-to-noise’ ratio




Figure 7. The signal-to-noise (bottom) and bias-to-noise ratio (top) as a function of redshift, for ϵ_A (blue), ϵ_X (red) and ϵ_{opt} (yellow) in the full analysis. The bias for ϵ_{opt} and ϵ_A increases for low and high redshifts, because foreground cleaning is less effective there (see also Figs 4 and 5). The bias in ϵ_X is compatible with 1σ fluctuations, thanks to lack of foreground residuals in the H I-galaxy cross-correlation. Foreground cleaning still introduces a random error in all estimators, which is highest at the upper and lower ends of the frequency range, similar for all estimators.

for a given estimator as

$$\left(\frac{B}{N}\right)_\ell \equiv \frac{\langle \epsilon_\ell \rangle - \epsilon_\ell^{\text{true}}}{\sigma_\ell}, \quad (35)$$

where $\epsilon_\ell^{\text{true}}$ and σ_ℓ are defined in Section 4.2. We compute this quantity for all redshift bins and multipoles, and then estimate a scale-averaged B/N as a quadrature sum of the ℓ -dependent ratio

$$\left(\frac{B}{N}\right)(z) = \sqrt{\sum_{\ell=0}^{\ell_{\text{max}}} \left(\frac{B}{N}\right)_\ell^2}, \quad (36)$$

where the sum is taken over the range of relevant multipoles, depending on redshift. This quantity is shown in the upper panel of Fig. 7 as a function of redshift for the three estimators considered here. While the bias of the cross-correlation-based estimator ϵ_X is compatible with $\sim 1\sigma$ fluctuations, the use of autocorrelations through either ϵ_A or ϵ_{opt} produces noticeable biases of up to 10σ , caused by foreground contamination. The lower panel of the same figure shows the integrated S/N ratio (estimated as a quadrature sum over power spectrum multipoles), and reinforces our conclusion that all estimators achieve similar sensitivities, and therefore we do not incur in any significant loss by dropping all autocorrelation information and using only cross-correlations for which foregrounds do not induce any bias. For the results in both panels of this figure $\ell_{\text{max}} = 81$.

Finally, we summarize the main findings of this section in Fig. 8. The left-hand panel shows the distribution of $\epsilon_A - \epsilon^{\text{true}}$ and $\epsilon_X - \epsilon^{\text{true}}$ across all simulations and ℓ values for a bin at $z \sim 0.8$. The distributions are close to Gaussian, and the ϵ_A shows a clear foreground bias. The right-hand panel shows the degradation in sensitivity caused by instrumental noise (dotted line to dashed line)

and by the partial removal of signal due to foregrounds (dashed line to solid line).

4.4.2 The effects of foregrounds

We have carried out a number of tests to further understand the effects of foregrounds on multitracer analyses, and to explore different avenues to mitigate these effects.

As we have seen, the cross-correlation estimator ϵ_X is immune to foreground bias and its use does not incur in any significant penalty in terms of sensitivity. Therefore, the main impact of foregrounds in 21cm observations is the loss of long radial wavelength modes present in the galaxy distribution. To quantify this effect we have studied the radial 1D power spectrum $P_{1D}(k_\parallel)$, defined as the variance of the line-of-sight Fourier coefficients of our 21cm maps. In practice we estimate this observable, as outlined in Villaescusa-Navarro, Alonso & Viel (2017), by computing the Fourier transform for a given pixel across all frequencies:

$$\Delta T(k_\parallel, \hat{n}) = \int \frac{dv}{v_{21}\sqrt{2\pi}} \exp\left[i\frac{vk_\parallel}{v_{21}}\right] \Delta T(v, \hat{n}). \quad (37)$$

Note that we use v/v_{21} , as a radial coordinate, where $v_{21} = 1420$ MHz is the frequency of the 21cm line, and therefore the radial wavenumber k_\parallel is dimensionless.⁹ The 1D power spectrum is then computed as the covariance between two fields ΔT_1 and ΔT_2 :

$$P_{1D}(k_\parallel) = \frac{\Delta v}{v_{21}2\pi} \langle \text{Re} [\Delta T_1(k_\parallel)\Delta T_2^*(k_\parallel)] \rangle, \quad (38)$$

where the average is taken across all unmasked pixels and all simulations.

Fig. 9 shows two 1D power spectra, computed from the auto-correlation of the foreground-free simulations (solid line) and from the cross-correlation of the foreground-clean and foreground-free simulations (dotted line). Although both power spectra match on small scales ($k \gtrsim 200$), the loss of long-wavelength radial modes becomes apparent on larger scales, where the amplitude of the cross-correlation becomes significantly smaller than the foreground-free power spectrum. On the other hand, the radial smearing effect of photometric redshifts will erase all structure on scales smaller than the photo- z error σ_z . Since $v/v_{21} = (1+z)^{-1}$, we can associate σ_z with a threshold wavenumber $k_{\text{ph}} \equiv \pi(1+z)^2/\sigma_z$. At $z \sim 1$ and assuming $\sigma_z = 0.03(1+z)$, we obtain $k_{\text{ph}} \sim 200$, which coincides with the scale at which the mode loss to foregrounds becomes noticeable. The bin width $\Delta z = 0.1$ would correspond to a scale $k_\parallel \sim 125$, and so effectively all the modes within the shaded region of Fig. 9 are erased in the data, due to the top-hat smoothing. The range of radial scales over which a significant overlap between 21cm observations and an LSST-like galaxy sample can be found becomes significantly reduced, which has a negative impact on the cosmic variance cancellation of the estimators studied here.

To circumvent this problem we have explored a few alternative avenues:

(i) **Foreground degrees of freedom.** To reduce the number of modes lost to foreground removal it is worth exploring the possibility of subtracting a smaller number of degrees of freedom at that stage. As discussed in Section 4.3, this will produce significant foreground residuals that will bias the autocorrelation but, since ϵ_X is immune to this bias, its sensitivity might benefit significantly

⁹In practice the Fourier transform is computed as a discrete Fourier transform (Frigo & Johnson 2005).

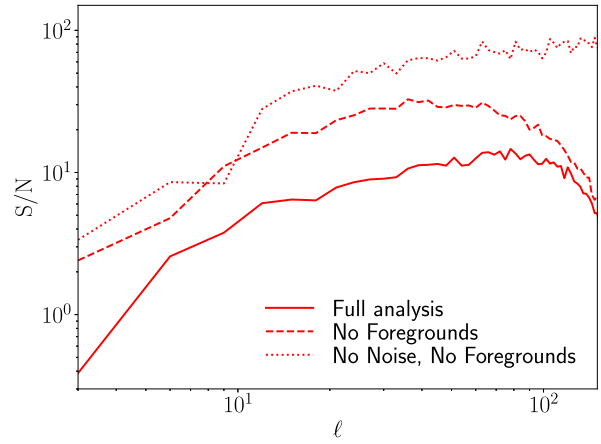



Figure 8. *Left-hand panel:* Distribution of $\Delta(\epsilon) = \epsilon - \epsilon^{\text{true}}$ for ϵ_A (blue) and ϵ_X (red) in the full analysis on the range $0 \leq \ell \leq 81$. The systematic error in ϵ_A is due to the effects of foreground cleaning combined with the low radial resolution of the galaxy maps. *Right-hand panel:* The SNR of ϵ_X for the full analysis (solid line), without foregrounds (dashed line) and without foregrounds or noise (dotted line). At large scales up to $\ell \lesssim 80$ foregrounds are the dominant source of uncertainty for ϵ_X .




Figure 9. Radial H I power spectrum, averaged over 200 simulations, showing the autocorrelation of foreground-free maps (solid line) and the cross-correlation of foreground-free and foreground-cleaned maps (dotted line). Note that we use a non-standard dimensionless radial coordinate ν/ν_{21} (see equation 37), and therefore the wavenumber $k_{||}$ is also dimensionless. The loss of long-wavelength radial modes is apparent in the drop of the dotted line for $k_{||} \lesssim 100$. The grey shaded area indicates the smoothing scale due to the redshift bin width of $\Delta z = 0.1$, associated with the LSST photo- z uncertainty. Unfortunately this is where foreground cleaning works best and the solid and dotted lines agree, limiting the scale overlap between both types of observations.

from the presence of additional signal modes. However, although the foreground residuals will not contribute to the bias of ϵ_X , they will also provide a contribution to its variance, and therefore there will be a balance between the preservation of long-wavelength modes and the contribution of foreground residuals to the noise.

Fig. 10 shows the S/N ratio of ϵ_A (blue) and ϵ_X (red) for the fiducial case, in which $N_{\text{FG}} = 9$ foreground degrees are subtracted (solid lines) and for an alternative scenario with $N_{\text{FG}} = 7$ (dashed lines). No significant improvement is obtained in both cases, and in fact we observe a reduction in sensitivity on large scales. Therefore, at least for this set-up, the contribution of foreground residuals to the




Figure 10. The SNR for ϵ_X (red) and ϵ_A (blue) with noise and foregrounds, comparing results obtained from cleaning with 9 (solid lines) and 7 foreground degrees of freedom (dashed lines). The contribution of foreground residuals to the estimator noise outweighs the potential improvement in sensitivity due to the milder subtraction of long-wavelength modes, and the $N_{\text{FG}} = 7$ case yields a poorer results than our fiducial choice of $N_{\text{FG}} = 9$.

estimator variance outweighs the impact of the additional signal degrees of freedom allowed by a more lenient removal stage. More efficient foreground removal methods preserving more information from the signal while at the same time removing all residuals on large scales could potentially improve this result.

(ii) **Thinner redshift bins.** The large photo- z uncertainties that can realistically be achieved by an experiment like LSST will make the use of redshift bins smaller than $\Delta z \sim 0.1$ pointless. This is due to the strong correlations between narrower bins induced by the photo- z scatter. None the less, it is worth exploring the possible cosmic variance cancellation gains if a better redshift precision were available. To do so, we have repeated our analysis making use of redshift bins with width $\Delta z = 0.02$.

The results of this exercise are shown in Fig. 11 for a bin centred around $z \sim 0.8$. The upper panel corresponds to the signal to noise as a function of scale, and shows an improvement of a factor ~ 2




Figure 11. *Top panel:* Similarly to Fig. 6, but working in thin bins of $\Delta\nu = 1$ MHz for both the galaxy and H I maps. We show the SNR for all estimators in the full analysis, and ϵ_{opt} in the foreground-free case. The latter is worse due to the smaller information content in the much thinner bin, but notably all solid lines actually outperform the respective results obtained from thick bins. This is due to the much bigger overlap of radial modes in the galaxy and H I power spectra when precise galaxy redshifts are assumed. Note that the full constraining power in this case would be realized by combining the information from the much larger number of redshift bins. *Bottom panel:* same as Fig. 8, but also for the thin frequency bin. While the shape and size of the distributions of ϵ_X (red) and ϵ_A (blue) hardly change, the bias in ϵ_A becomes negligible, due to the larger relative number of overlapping radial modes between the 21cm and galaxy data.

with respect to the cosmic variance-limited case for all estimators. Although this is comparable with the results shown for individual bins with $\Delta z = 0.1$, the number of uncorrelated bins available in this case is 5 times larger, and therefore the total signal to noise increases significantly. The lower panel of the same figure shows the distributions of ϵ_A and ϵ_X for the same redshift bin across all simulations and values of ℓ . The significant bias in ϵ_A observed in Fig. 8 is now gone, owing to the relative increase in radial modes on which the 21cm signal dominates over foreground residuals. As we emphasized above, although higher redshift resolution improves the performance of multitracer methods for 21cm intensity mapping, photometric redshift surveys are unlikely to achieve the required redshift accuracy. On the other hand, although spectroscopic surveys can easily reach that level of radial resolution, they can only do so for a substantially smaller number of objects, and the larger

shot noise will inevitably affect the performance of the multitracer technique. The most promising option is perhaps intensity mapping of other emission lines (e.g. the CO line; Padmanabhan 2018), as long as the instrumental noise can be reduced sufficiently.

(iii) **Matching scales.** Finally, another possibility would be to subject the galaxy overdensity data to the same linear transformation that downweights the long-wavelength modes in the 21cm maps. If this can be done with sufficient accuracy, the resulting auto and cross-power spectra should manifest the same fluctuations around the mean on a realization-by-realization basis, and a higher degree of CV cancellation could be expected from the estimators used here, based on ratios of those quantities.

Note that in a more optimal analysis, where the full data from the 21cm maps and the galaxy overdensity are used, including all the signal-dominated radial and angular modes (instead of just the power spectrum ratios of matching redshift bins), this is unlikely to provide any advantage over preserving all of the modes in the latter probe. A likelihood evaluation of the full data would automatically produce the cancellation of cosmic variance on all common modes, and would use the additional galaxy long-wavelength modes to increase the final constraints further.

Thus, to summarize: although multitracer methods applied to 21cm data in cross-correlation with photometric redshift surveys do improve the constraining power beyond the cosmic variance limit, this improvement is strongly hampered by the loss of long-wavelength modes, common to both data sets, due to the presence of foreground contamination and low z resolution. Multitracer analyses using 21cm observations are therefore more likely to achieve a better performance when combined with other IM data or spectroscopic surveys, assuming the noise amplitude of the latter (instrumental or shot noise) can be reduced sufficiently.

5 DISCUSSION

21cm intensity mapping and photometric redshift surveys are two promising techniques to study the three-dimensional distribution of matter in the Universe on large scales. A number of cosmological observables, such as the level of primordial non-Gaussianity, benefit from the combined analysis of multiple proxies of the same density inhomogeneities in what is known as the ‘multitracer’ technique. In this paper we have explored the feasibility of multitracer analyses that exploit the combination of the two aforementioned probes, for the particular case of 21cm observations to be carried out by a SKA-like instrument and an LSST-like galaxy sample. For concreteness, we have focused our analysis on two estimators of the bias ratio for both samples, ϵ_A and ϵ_X , described in Section 2.3. Since these estimators make use of the 21cm autocorrelation and its cross-correlation with galaxies, respectively, they allow us to explore both the bias induced on ϵ_A by the presence of foreground residuals, and the potential loss of information associated with dropping autocorrelation information (ϵ_X). For completeness, we also consider an optimal inverse variance combination of both estimators, ϵ_{opt} , that uses all the data available.

In the absence of foregrounds, we show that both ϵ_A and ϵ_X are able to achieve similar sensitivities, with little improvement when using ϵ_{opt} due to the tight correlation between both estimators. When compared with the cosmic variance contaminated measurement of the same bias ratio, we show that these estimators are able to improve the signal to noise by a factor of ~ 4 – 5 , even when compared to the cosmic variance-contaminated, noise-free case, showcasing the tremendous potential gains of the multitracer technique.

Table 1. Signal to noise from combining all redshift bins for all estimators and all modelling scenarios of this work. Here using 9 degrees of freedom for the foreground cleaning and a redshift bin width of $\Delta z = 0.1$. While ϵ_A uses the H I and galaxy autocorrelations, ϵ_X uses the H I-g cross-correlation and g-g autocorrelation and ϵ_{opt} is the inverse variance-weighted sum of both (equations 10–12). The estimator $\epsilon_{\text{CV}+}$ on the other hand uses autocorrelations with different DM realizations for the galaxy and H I populations and shows the constraints achievable in the absence of multitracer cosmic variance cancellation.

Case	ϵ_A	ϵ_X	$\epsilon_{\text{CV}+}$	ϵ_{opt}
No noise, no FG	1291	1292	192	1306
No FG	495	502	154	509
No noise	299	298	155	312
Full analysis	178	183	120	192

The impact of the presence of foregrounds in the 21cm data is twofold. On the one hand, residuals after foreground removal produce an offset in the H I autocorrelation which biases both ϵ_A and ϵ_X at high significance. We show however, that ϵ_X is immune to this bias, while preserving the same statistical power as the two other estimators. On the other hand, foreground removal is based on the separation of foregrounds and cosmological signal through their different spectral behaviour, effectively downweighting the radial long-wavelength modes where foregrounds dominate. Since photometric redshifts effectively erase all structure along the line of sight on all but the largest scales, the overlap between SKA and LSST in the $k_{\parallel}-k_{\perp}$ plane reduces significantly, partially spoiling the cosmic variance cancellation. We show that, in this case, the sensitivity of all estimators drops by more than a factor of ~ 2 , and that the improvement in SNR with respect to a cosmic variance-limited measurement made in the same circumstances is now only a factor ~ 2 . This drops to a smaller ~ 50 per cent improvement when we compare either estimator with the cosmic variance-limited measurement without foregrounds. These results are summarized in Table 1, which shows the cumulative signal to noise (quadrature-summed over all multipoles and redshift bins) for the three estimators as well as the CV limit in different scenarios regarding the presence of noise and foregrounds.

We have also explored two possible ways to overcome this problem. First, a less aggressive foreground removal that leaves a larger fraction of foreground residuals in the maps, would also leave a larger number of long-wavelength modes untouched, increasing the scale overlap between LSST and SKA. In practice, however, we have seen that the contribution of the foreground residuals to the estimator uncertainties in fact decrease the total SNR when a smaller number of foreground degrees of freedom are subtracted. Another way to increase the scale overlap between both experiments would be to reduce the size of the redshift bins used in the analysis. Although this is not a real possibility for photometric surveys, since structures can never be resolved on scales smaller than the photo- z uncertainty, this case allows us to explore other possible synergies with either spectroscopic surveys or IM observations of other emission lines. Our results show that in this case the gain in sensitivity associated with the multitracer technique is likely restored, with the added advantage that the foreground bias is also reduced due to the larger fraction of signal-dominated modes. An analysis pairing H I IM with a spectroscopic survey would work with maps in thin redshift slices and therefore need to take into account cross-correlations between adjacent redshift bins as well as redshift-space distortions. A large area survey would most likely suffer from high shot noise degrading the results, but the potential of

a low-volume survey with a large number density of galaxies could be interesting. Another promising way forward for these types of analyses may be the combination of IM observations for different emission lines, even though two emission line surveys might suffer from foreground residual cross-correlations.

In a follow-up work we plan to study these new avenues in more detail, first by considering constraints on f_{NL} directly, and including estimators that can deal naturally with the mismatch in the modes that are removing in different surveys does allowing a more perfect cancellation. We will also consider other foreground cleaning methods that might be less aggressive on cleaning this large scales and new tracers with higher redshift resolution.

ACKNOWLEDGEMENTS

We would like to thank Pedro Ferreira for useful comments and discussions. AW, JF, and MGS acknowledge support from the South African Square Kilometre Array Project and National Research Foundation. AW was also supported by the Newton Fund PhD Partnership Scheme, NRF grant 98555 and UK grant ES/N013956/1, part of the UK’s official development assistance (ODA). DA acknowledges support from the Beecroft Trust and from the Science and Technology Facilities Council (STFC) through an Ernest Rutherford Fellowship, grant reference ST/P004474/1. We acknowledge the support of the Centre for High Performance Computing, South Africa, under the project ASTRO945.

REFERENCES

- Abramo L. R., Bertacca D., 2017, *Phys. Rev. D*, 96, 123535
 Abramo L. R., Leonard K. E., 2013, *MNRAS*, 432, 318
 Alonso D., Ferreira P. G., 2015, *Phys. Rev. D*, 92, 063525
 Alonso D., Ferreira P. G., Santos M. G., 2014, *MNRAS*, 444, 3183
 Alonso D., Bull P., Ferreira P. G., Santos M. G., 2015a, *MNRAS*, 447, 400
 Alonso D., Bull P., Ferreira P. G., Maartens R., Santos M., 2015b, *ApJ*, 814, 145
 Anderson C. J. et al., 2018, *MNRAS*, 476, 3382
 Battye R. A., Davies R. D., Weller J., 2004, *MNRAS*, 355, 1339
 Blake C. et al., 2013, *MNRAS*, 436, 3089
 Bonvin C., 2014, *Class. Quantum Gravity*, 31, 234002
 Bonvin C., Durrer R., 2011, *Phys. Rev. D*, 84, 063505
 Brunl M., Crittenden R., Koyama K., Maartens R., Pitrou C., Wands D., 2012, *Phys. Rev. D*, 85, 041301
 Bull P., Ferreira P. G., Patel P., Santos M. G., 2015, *ApJ*, 803, 21
 Camera S., Santos M. G., Ferreira P. G., Ferramacho L., 2013, *Phys. Rev. Lett.*, 111, 171302
 Camera S., Santos M. G., Maartens R., 2015, *MNRAS*, 448, 1035
 Challinor A., Lewis A., 2011, *Phys. Rev. D*, 84, 043516
 Chang T.-C., Pen U.-L., Peterson J. B., McDonald P., 2008, *Phys. Rev. Lett.*, 100, 091303
 Chapman E. et al., 2013, *MNRAS*, 429, 165
 Coles P., Jones B., 1991, *MNRAS*, 248, 1
 Dalal N., Doré O., Huterer D., Shirokov A., 2008, *Phys. Rev. D*, 77, 123514
 de Putter R., Gleyzes J., Doré O., 2017, *Phys. Rev. D*, 95, 123507
 Di Dio E., Montanari F., Lesgourgues J., Durrer R., 2014, *Class. Quantum Gravity*, 31, 234002
 Di Matteo T., Perna R., Abel T., Rees M. J., 2002, *ApJ*, 564, 576
 Ferramacho L. D., Santos M. G., Jarvis M. J., Camera S., 2014, *MNRAS*, 442, 2511
 Fonseca J., Camera S., Santos M. G., Maartens R., 2015, *ApJ*, 812, L22
 Fonseca J., Maartens R., Santos M. G., 2017, *MNRAS*, 466, 2780
 Fonseca J., Maartens R., Santos M. G., 2018, *MNRAS*, 479, 3490
 Frigo M., Johnson S. G., 2005, *Proc. IEEE*, 93, 216
 Giannantonio T., Porciani C., Carron J., Amara A., Pillepich A., 2012, *MNRAS*, 422, 2854

- Hall A., Bonvin C., Challinor A., 2013, *Phys. Rev. D*, 87, 064026
- Hamaus N., Seljak U., Desjacques V., 2011, *Phys. Rev. D*, 84, 083509
- Harper S. E., Dickinson C., Batty R. A., Roychowdhury S., Browne I. W. A., Ma Y.-Z., Olivari L. C., Chen T., 2018, *MNRAS*, 478, 2416
- Jelić V., Zaroubi S., Labropoulos P., Bernardi G., de Bruyn A. G., Koopmans L. V. E., 2010, *MNRAS*, 409, 1647
- Jeong D., Schmidt F., Hirata C. M., 2012, *Phys. Rev. D*, 85, 023504
- Knox L., 1995, *Phys. Rev. D*, 52, 4307
- Limber D. N., 1954, *ApJ*, 119, 655
- Loverde M., Afshordi N., 2008, *Phys. Rev. D*, 78, 123506
- LSST Science Collaboration and LSST Project, 2009, LSST Science Book, Version 2.0, preprint ([arXiv:0912.0201](https://arxiv.org/abs/0912.0201))
- Marín F. A., Beutler F., Blake C., Koda J., Kazin E., Schneider D. P., 2016, *MNRAS*, 455, 4046
- Masui K. W. et al., 2013, *ApJ*, 763, L20
- Matarrese S., Verde L., 2008, *ApJ*, 677, L77
- McDonald P., Seljak U., 2009, *J. Cosmol. Astropart. Phys.*, 0910, 007
- Moore D. F., Aguirre J. E., Parsons A. R., Jacobs D. C., Pober J. C., 2013, *ApJ*, 769, 154
- Oh S. P., Mack K. J., 2003, *MNRAS*, 346, 871
- Padmanabhan H., 2018, *MNRAS*, 475, 1477
- Planck Collaboration, 2016, *A&A*, 594, A17
- Pourtsidou A., Bacon D., Crittenden R., Metcalf R. B., 2016, *MNRAS*, 459, 863
- Pourtsidou A., Bacon D., Crittenden R., 2017, *MNRAS*, 470, 4251
- Pullen A. R., Serra P., Chang T.-C., Doré O., Ho S., 2018, *MNRAS*, 478, 1911
- Raccanelli A., Montanari F., Bertacca D., Doré O., Durrer R., 2016, *J. Cosmol. Astropart. Phys.*, 1605, 009
- Ross A. J. et al., 2011, *MNRAS*, 417, 1350
- Ross A. J. et al., 2014, *MNRAS*, 437, 1109
- Santos M. et al., 2015, in Bourke T. L. et al., eds, Proc. Sci., Advancing Astrophysics with the Square Kilometre Array (AASKA14). SISSA, Trieste, Italy, p. 19
- Santos M. G., Cooray A., Knox L., 2005, *ApJ*, 625, 575
- Schmittfull M., Seljak U., 2018, *Phys. Rev. D*, 97, 123540
- Schulz A. E., 2010, *ApJ*, 724, 1305
- Seljak U., 2009, *Phys. Rev. Lett.*, 102, 021302
- Shaw J. R., Sigurdson K., Pen U.-L., Stebbins A., Sitwell M., 2014, *ApJ*, 781, 57
- Shaw J. R., Sigurdson K., Sitwell M., Stebbins A., Pen U.-L., 2015, *Phys. Rev. D*, 91, 083514
- Square Kilometre Array Cosmology Science Working Group, 2018, preprint ([arXiv:1811.02743](https://arxiv.org/abs/1811.02743))
- Switzer E. R. et al., 2013, *MNRAS*, 434, L46
- Switzer E. R., 2017, *ApJ*, 838, 82
- Villaescusa-Navarro F. et al., 2018, *ApJ*, 866, 135
- Villaescusa-Navarro F., Alonso D., Viel M., 2017, *MNRAS*, 466, 2736
- Weinberg D. H., Davé R., Katz N., Hernquist L., 2004, *ApJ*, 601, 1
- Wolz L., Abdalla F. B., Blake C., Shaw J. R., Chapman E., Rawlings S., 2014, *MNRAS*, 441, 3271
- Wolz L., Blake C., Wyithe J. S. B., 2017, *MNRAS*, 470, 3220
- Yamauchi D., Takahashi K., Oguri M., 2014, *Phys. Rev. D*, 90, 083520
- Yoo J., Desjacques V., 2013, *Phys. Rev. D*, 88, 023502
- Yoo J., Seljak U., 2015, *MNRAS*, 447, 1789
- Yoo J., Hamaus N., Seljak U., Zaldarriaga M., 2012, *Phys. Rev. D*, 86, 063514
- Zuo S., Chen X., Ansari R., Lu Y., 2019, *AJ*, 157, 4

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