

Blind foreground subtraction for intensity mapping experiments

David Alonso,^{1★} Philip Bull,² Pedro G. Ferreira¹ and Mário G. Santos^{3,4}

¹*Astrophysics, University of Oxford, DWB, Keble Road, Oxford OX1 3RH, UK*

²*Institute of Theoretical Astrophysics, University of Oslo, PO Box 1029 Blindern, N-0315 Oslo, Norway*

³*Department of Physics, University of Western Cape, Cape Town 7535, South Africa*

⁴*SKA SA, 3rd Floor, The Park, Park Road, Pinelands 7405, South Africa*

Accepted 2014 November 21. Received 2014 November 21; in original form 2014 September 30

ABSTRACT

We make use of a large set of fast simulations of an intensity mapping experiment with characteristics similar to those expected of the Square Kilometre Array in order to study the viability and limits of blind foreground subtraction techniques. In particular, we consider three different approaches: polynomial fitting, principal component analysis (PCA) and independent component analysis (ICA). We review the motivations and algorithms for the three methods, and show that they can all be described, using the same mathematical framework, as different approaches to the blind source separation problem. We study the efficiency of foreground subtraction both in the angular and radial (frequency) directions, as well as the dependence of this efficiency on different instrumental and modelling parameters. For well-behaved foregrounds and instrumental effects, we find that foreground subtraction can be successful to a reasonable level on most scales of interest. We also quantify the effect that the cleaning has on the recovered signal and power spectra. Interestingly, we find that the three methods yield quantitatively similar results, with PCA and ICA being almost equivalent.

Key words: methods: statistical – large-scale structure of universe – radio lines: galaxies.

1 INTRODUCTION

The last few decades have seen a radical improvement in the quantity and quality of observational data that has transformed cosmology into a fully-fledged data-driven science – so much so that it has now become common to refer to the current status of the field as the ‘era of precision cosmology’. The vast majority of these data come from two very distinct regimes, both observationally and physically: on the one hand, observations of the cosmic microwave background (CMB), emitted in the early Universe ($z \sim 1000$), have allowed us to measure the values of most cosmological parameters with astonishing precision (Hinshaw et al. 2013; Planck Collaboration XVI 2014). On the other, astronomical observations in the optical range of wavelengths have supplied a wealth of data with which we have been able to characterize the late-time ($z \lesssim 1$) evolution of the Universe and to make one of the most puzzling discoveries in cosmology: its accelerated expansion (Riess et al. 1998; Perlmutter et al. 1999). In this regime, galaxy redshift surveys (Contreras et al. 2013; Anderson et al. 2014) have allowed us to draw a plausible picture describing the evolution of the primordial density perturbations observed in the CMB to form the non-linear large-scale structure (LSS) that we observe today.

There exist however multiple open questions in cosmology, the study of which would greatly benefit from observational data in the remaining intermediate range of redshifts. For instance, one of the most important stages in the evolution of the Universe, the epoch of reionization (EoR) and the formation of the first stars and galaxies, must occur at $z \sim 10$, but its precise nature remains unknown due to the lack of observations (Furlanetto, McQuinn & Hernquist 2006). Furthermore, being able to observe the distribution of matter out to redshift $z \simeq 3\text{--}4$ would make a very large volume available for LSS analyses, which could significantly improve our understanding of the process of structure formation, as well as opening the field to the study of ultralarge scales, which contain substantial information about the underlying gravitational theory and the statistics of the primordial density field (Camera et al. 2013; Hall, Bonvin & Challinor 2013). Radio-astronomical observations offer an excellent tool in this regard through the measurement of the 21 cm signal from neutral hydrogen (H I). Measuring the intensity of this line emission, due to the hyperfine spin-flip transition ($\nu_{21} \simeq 1420$ MHz), as a function of frequency constitutes an ideal way to trace the H I density field as a function of redshift $z = \nu_{21}/\nu - 1$.

While the potential scientific value of these techniques is enormous, the detection of the H I signal faces several observational challenges. After reionization, most of the neutral hydrogen in the Universe is thought to reside in high-density regions inside galaxies, in what are known as damped Lyman α systems, where it is shielded from the ionizing UV photons. Unfortunately, the H I emission from

★ E-mail: david.alonso@astro.ox.ac.uk

individual galaxies is usually too weak to be measured efficiently, and thus a technique known as *intensity mapping* has been advocated (see Barkana & Loeb 2005; Pritchard & Loeb 2008; Wyithe & Loeb 2008; Peterson et al. 2009; Bagla, Khandai & Datta 2010 for a description of the theoretical principles, Battye, Davies & Weller 2004; McQuinn et al. 2006; Chang et al. 2008; Loeb & Wyithe 2008; Mao et al. 2008; Wyithe, Loeb & Geil 2008; Seo et al. 2010; Lidz et al. 2011; Ansari et al. 2012; Battye et al. 2013; Bull et al. 2014 for proposed experiments and cosmological forecasts and Masui et al. 2013; Switzer et al. 2013, for the first observations to date), in which the combined emission from all sources in wide angular pixels is observed instead. While small angular scales are lost through this method, it is possible to retain the larger scales of most cosmological relevance, such as the baryon acoustic oscillation (BAO) scale. Thus, intensity mapping constitutes a promising method to efficiently observe the distribution of matter in the Universe on cosmological scales.

The success of intensity mapping relies on being able to separate the cosmological H I signal from that of foreground radio sources emitting in the same frequency range, however. The most important of these sources, synchrotron emission from our own galaxy, is ~ 5 orders of magnitude larger than the expected 21 cm signal, even at high galactic latitudes, and so the problem of foreground subtraction is of prime importance. This topic has been widely covered in the literature (Di Matteo et al. 2002; Oh & Mack 2003; Santos, Cooray & Knox 2005; Morales, Bowman & Hewitt 2006; Wang et al. 2006; Gleser, Nusser & Benson 2008; Jelić et al. 2008; Bernardi et al. 2009, 2010; Jelić et al. 2010; Moore et al. 2013), mostly for the EoR regime, and several foreground removal techniques have been proposed (Liu et al. 2009; Liu & Tegmark 2011; Masui et al. 2013; Shaw et al. 2013, 2014; Wolz et al. 2014). Fortunately, as opposed to the cosmological signal, most relevant foregrounds have a very smooth frequency dependence, a fact that can be exploited to subtract them efficiently. Nevertheless, their large amplitude makes a thorough study of the consequences of foreground cleaning on the recovered cosmological signal indispensable.

Foreground removal techniques can be classified as being either *blind*, if they only make use of generic foreground properties like spectral smoothness, or *non-blind*, if a more detailed model of the foregrounds is required. In this paper, we have used a large set of simulations of both the cosmological signal and several different foregrounds to study and compare the performance of different blind foreground subtraction methods. The paper is structured as follows: in Section 2 we briefly describe the physics of the H I cosmological signal and the most relevant foregrounds for intensity mapping. Section 3 covers the mathematics of blind foreground cleaning within a unified formalism. Section 4 contains the details of the simulations used in this analysis, as well as the criteria used to compare the different methods. Finally, Section 5 presents the results of this study regarding the performance and limits of blind foreground cleaning, and Section 6 summarizes our findings.

2 INTENSITY MAPPING AND ITS FOREGROUNDS

The emissivity (luminosity per unit frequency and solid angle) of a cloud of neutral hydrogen due to the 21 cm line is given, in its rest frame, by

$$\frac{dL}{d\Omega dv} = \frac{h_p A_{21} N_2}{4\pi N_H} N_H v \varphi(v), \quad (1)$$

where h_p is Planck's constant, $A_{21} = 2.876 \times 10^{-15}$ Hz is the Einstein coefficient corresponding to the 21 cm line, N_2 is the number of atoms in the excited state, N_H is the total number of H I atoms and $\varphi(v)$ is the normalized line profile. From this result, it is easy to calculate the intensity measured by an observer in the lightcone and hence the brightness temperature along a particular direction \hat{n} in the Rayleigh–Jeans approximation (Abdalla & Rawlings 2005):

$$T(z, \hat{n}) = \frac{3 h_p c^3 A_{21}}{32\pi k_B v_{21}^2} \frac{(1+z)^2}{H(z)} n_{H\text{I}}(z, \hat{n}) \\ = (0.19055 \text{ mK}) \frac{\Omega_b h (1+z)^2 x_{H\text{I}}(z)}{\sqrt{\Omega_M(1+z)^3 + \Omega_\Lambda}} (1 + \delta_{H\text{I}}), \quad (2)$$

where $n_{H\text{I}} \propto (1 + \delta_{H\text{I}})$ is the comoving number density of H I atoms, k_B is Boltzmann's constant and $x_{H\text{I}}$ is the fraction of the total baryonic mass in H I. Note that to reach this result we have neglected scattering and self-absorption of the emitted photons, as well as any relativistic effects other than the Hubble expansion. The observed H I signal is also affected by redshift-space distortions (Kaiser 1987) in the same way as galaxy number counts are (Hall et al. 2013).

The most relevant foregrounds for intensity mapping can be classified as being extragalactic (caused by astrophysical sources beyond the Milky Way) or galactic. In terms of amplitude, the largest foreground by far is galactic synchrotron emission (GSE), caused by high-energy cosmic ray electrons accelerated by the Galactic magnetic field. Although at the lowest frequencies the GSE spectrum is damped by self-absorption and free–free absorption, for the relevant radio frequencies it should be possible to describe it by a single power law $T_{\text{GSE}} \propto \nu^\beta$, with a spectral index $\beta(\hat{n}) \approx -2.8$ (Dickinson et al. 2009; Delabrouille et al. 2013). GSE is expected to be partially linearly polarized, and so its polarized part will undergo Faraday rotation as it travels through the Galactic magnetic field and the interstellar medium (Rybicki & Lightman 1986; Waelkens et al. 2009). This effect gives polarized GSE a non-trivial and non-smooth frequency dependence, which would make it a very troublesome foreground if leaked into the unpolarized signal (Jelić et al. 2010; Moore et al. 2013; De & Tashiro 2014).

Another source of foreground emission within our Galaxy is the so-called free–free emission, caused by free electrons accelerated by ions, which traces the warm ionized medium. As in the case of GSE, free–free emission is predicted to be spectrally smooth in the relevant range of frequencies (Dickinson, Davies & Davis 2003). All other known potential galactic foregrounds, such as anomalous microwave emission or dust emission, should be negligible below ~ 1 GHz.

Extragalactic radio sources can be classified into two different categories: bright radio galaxies, such as active galactic nuclei, and ‘normal’ star-forming galaxies. While the distribution of the former should be Poisson-like, the clustering of the latter should be significant, which has a direct impact on the degree of smoothness of their joint emission (Zaldarriaga, Furlanetto & Hernquist 2004; Santos, Cooray & Knox 2005). Note that radio sources have been observed up to very high redshifts ($z \sim 5$), and therefore many of them will be physically behind the cosmological signal in the case of intensity mapping. However, we will refer to all sources of radio emission that need to be separated from the cosmological signal as ‘foregrounds’, trusting that the use of this term will not cause any confusion.

Other possible foreground sources are atmospheric noise, artificial radio frequency interference (which we discuss in Section 5.3) and line foregrounds, caused by line emission from astrophysical sources in other frequencies. Due to the spectral isolation of the

21 cm line, together with the expected low intensity of the most potentially harmful lines (such as OH at $\nu_{\text{OH}} \sim 1600$ MHz), the H I signal should be very robust against line confusion.

3 BLIND FOREGROUND SUBTRACTION

As has been noted above, the most relevant foregrounds for intensity mapping should be correlated (i.e. smooth) in frequency. Blind foreground subtraction methods make use of this property, modelling the sky brightness temperature along a given direction in the sky \hat{n} and at a frequency ν as

$$T(\nu, \hat{n}) = \sum_{k=1}^{N_{\text{fg}}} f_k(\nu) S_k(\hat{n}) + T_{\text{cosmo}}(\nu, \hat{n}) + T_{\text{noise}}(\nu, \hat{n}), \quad (3)$$

where N_{fg} is the number of foreground degrees of freedom to subtract, $f_k(\nu)$ are a set of smooth functions of the frequency, $S_k(\hat{n})$ are the foreground sky maps and T_{cosmo} and T_{noise} are the cosmological signal and instrumental noise components.

For a particular line of sight \hat{n} measured in a discrete set of N_ν frequencies, this model can be written as a linear system:

$$\mathbf{x} = \hat{\mathbf{A}} \cdot \mathbf{s} + \mathbf{r}, \quad (4)$$

where $x_i = T(\nu_i, \hat{n})$, $A_{ik} = f_k(\nu_i)$, $s_k = S_k(\hat{n})$ and we have grouped the cosmological signal and thermal noise in a single component $r_i = T_{\text{cosmo}}(\nu_i, \hat{n}) + T_{\text{noise}}(\nu_i, \hat{n})$. The problem of foreground subtraction then boils down to devising a method to determine $\hat{\mathbf{A}}$ and \mathbf{s} so that $\mathbf{r} = \mathbf{x} - \hat{\mathbf{A}}\mathbf{s}$ is recovered as accurately as possible. The three methods described below correspond to three different approaches to this problem.

3.1 Polynomial fitting

Probably the most intuitive (and naïve) way of approaching blind foreground subtraction is to choose an ad hoc basis of smooth functions f_k that we think can describe the foregrounds and then find the foreground maps s_k by least-squares fitting the model in equation (4) to each line of sight. This is done by minimizing the χ^2 ,

$$\chi^2 = (\mathbf{x} - \hat{\mathbf{A}} \cdot \mathbf{s})^T \hat{\mathbf{N}}^{-1} (\mathbf{x} - \hat{\mathbf{A}} \cdot \mathbf{s}), \quad (5)$$

where $\hat{\mathbf{N}}$ is the covariance matrix of \mathbf{r} . The solution is

$$\mathbf{s} = (\hat{\mathbf{A}}^T \hat{\mathbf{N}}^{-1} \hat{\mathbf{A}})^{-1} \hat{\mathbf{A}}^T \hat{\mathbf{N}}^{-1} \mathbf{x}. \quad (6)$$

In this analysis, we have neglected the correlation between different frequency bands and therefore we have assumed a diagonal covariance $N_{ij} = \sigma_i^2 \delta_{ij}$.¹

It is important to note that the deviations from the smooth foregrounds are due not only to thermal noise, but also to the cosmological clustering signal. For this reason, in order to optimally fit for the foregrounds we have used the joint variance of both components:

$$\sigma_i = \sqrt{\sigma_{\text{noise},i}^2 + \sigma_{\text{cosmo},i}^2}. \quad (7)$$

Evidently it is not possible to compute σ_{cosmo} from the cosmological signal before subtracting the foregrounds. In our analysis, we have done this by using the foreground-free maps, which would not be

¹ This assumption is not exactly correct, since even for uncorrelated instrumental noise, the cosmological signal (which we have included as a noise-like contribution in this formalism) is correlated in real space, which should ideally be taken into account.

available in a real experiment; we can propose three alternative methods, however.

(i) Assuming a particular cosmological model it would be possible to quantify the clustering variance, at least at the linear level, as

$$\sigma_{\text{cosmo},i}^2 = \int \frac{dk^3}{(2\pi)^3} |W_i(\mathbf{k}_\perp, k_\parallel)|^2 P(k), \quad (8)$$

where $P(k)$ is the power spectrum for the assumed model and $W_{\text{pix},i}$ is the Fourier-space window function describing the beam and pixel size and the radial width for the i th frequency bin.

(ii) Again assuming a particular cosmological model, another possibility would be to simulate the cosmological signal and include instrumental effects to compute σ_i^2 from the simulations.

(iii) Through a multistep analysis: first the foregrounds are subtracted using only the variance due to the instrumental noise (or even no variance weighting at all). The joint variance σ_i is then estimated from the cleaned maps and the foreground subtraction is repeated using this better estimate. The process can then be repeated until it converges. This method would be computationally more demanding, but it has the advantage of being model independent.

It is worth noting that in our analysis we observed that the differences in the final results between weighting with the full variance σ_i or only with the contribution from instrumental noise $\sigma_{\text{noise},i}$ were basically negligible. This is reasonable, since both the instrumental noise and the clustering variance have a similar frequency dependence for our instrumental setup.

Previous studies of foreground subtraction in intensity mapping (particularly for the EoR regime) have made use of this technique in log–log space, i.e. using the model:

$$\log(T(\nu, \hat{n})) = \sum_{k=1}^{N_{\text{fg}}} \alpha_k(\hat{n}) f_k(\log(\nu)). \quad (9)$$

Following what has been done by other groups (Wang et al. 2006; Ansari et al. 2012), we used polynomials as basis functions in this study, i.e. $f_k(\log(\nu)) = [\log(\nu)]^{k-1}$. Note that although very similar in spirit, this procedure does not adhere exactly to the model in equation (4), since the fitting is done in log–log space. This also implies that the fit weights σ_i^{-2} must be translated from linear space:

$$\sigma_{\log}(T) \simeq \frac{\sigma_{\text{lin}}}{T}. \quad (10)$$

3.2 Principal component analysis

Principal component analysis (PCA) applied to foreground subtraction seeks to make use of the main properties of the foregrounds (namely their large amplitude and smooth frequency coherence) to find both the foreground components s_k and an optimal set of basis functions A_{ik} at the same time. Here, we present the motivation for PCA and its connection to the main equation for blind foreground subtraction (equation 4).

It is easy to prove that the frequency–frequency covariance matrix of any component that is very correlated in frequency will have a very particular eigensystem: most of the information will be concentrated in a small set of very large eigenvalues, the other ones being negligibly small. As an extreme case, consider a completely correlated component, whose covariance matrix is $C_{ij} = 1$ ($\forall i, j \in [1, N_\nu]$), and whose eigenvalues are $\lambda_1 = N_\nu$, $\lambda_i = 0$ $i > 1$. Thus, we can attempt to subtract the foregrounds by eliminating from the

measured temperature maps the components corresponding to the eigenvectors of the frequency covariance matrix with the N_{fg} largest associated eigenvalues. The explicit method is as follows.

(i) Compute the frequency covariance matrix from the data by averaging over the available N_θ lines of sight:

$$C_{ij} = \frac{1}{N_\theta} \sum_{n=1}^{N_\theta} T(v_i, \hat{\mathbf{n}}_n) T(v_j, \hat{\mathbf{n}}_n). \quad (11)$$

(ii) Diagonalize the covariance matrix:

$$\hat{\mathbf{U}}^T \hat{\mathbf{C}} \hat{\mathbf{U}} = \hat{\Lambda} \equiv \text{diag}(\lambda_1, \dots, \lambda_{N_v}), \quad (12)$$

where $\lambda_i > \lambda_{i+1} \forall i$ are the (ordered) eigenvalues of $\hat{\mathbf{C}}$, and $\hat{\mathbf{U}}$ is an orthogonal matrix whose columns are the corresponding eigenvectors.

(iii) At this stage, we identify N_{fg} eigenvalues corresponding to the foregrounds as those that are much larger than the rest. Depending on the frequency structure of the foregrounds and the different instrumental effects this number will be more or less evident (see the discussion in Section 5.1). We then build the matrix $\hat{\mathbf{U}}_{\text{fg}}$ from the columns of $\hat{\mathbf{U}}$ corresponding to these eigenvalues and model the brightness temperature for each line of sight as

$$\mathbf{x} = \hat{\mathbf{U}}_{\text{fg}} \mathbf{s} + \mathbf{r}, \quad (13)$$

which is analogous to equation (4). The foreground maps \mathbf{s} are then found by projecting \mathbf{x} on the basis of eigenvectors of $\hat{\mathbf{C}}$, yielding

$$\mathbf{s} = \hat{\mathbf{U}}_{\text{fg}}^T \mathbf{x}. \quad (14)$$

Note that, since $\hat{\mathbf{U}}$ is orthogonal, $\hat{\mathbf{U}}_{\text{fg}}^T \hat{\mathbf{U}}_{\text{fg}} = \mathbf{1}$, and, except for the presence of the covariance $\hat{\mathbf{N}}$, equations (13) and (14) are equivalent to equations (4) and (6).

In the presence of $\hat{\mathbf{N}}$, the frequency covariance must be computed using an inverse-variance scheme:

$$C_{ij} = \frac{1}{N_\theta} \sum_{n=1}^{N_\theta} \frac{T(v_i, \hat{\mathbf{n}}_n)}{\sigma_i} \frac{T(v_j, \hat{\mathbf{n}}_n)}{\sigma_j}, \quad (15)$$

which is equivalent to carrying out the steps above using the weighted maps $x_i = T(v_i)/\sigma_i$ and multiplying by σ_i in the end to obtain the de-weighted temperature maps. It is easy to see that doing this we introduce the missing $\hat{\mathbf{N}}^{-1}$ factors that make the equivalence between equations (14) and (6) complete.

We thus see that the PCA method is in fact equivalent to the polynomial fitting method, with the set of basis functions A_{ik} given by the data themselves in the form of those that contain most of the total variance (i.e. the principal eigenvectors of the covariance matrix).

Finally, let us note that in a real experiment the situation will be more complicated, for example due to the presence of correlated instrumental noise. This problem was addressed in the pioneering analysis done by the Green Bank Telescope team (Masui et al. 2013; Switzer et al. 2013) by computing the frequency covariance matrix through a cross-correlation of temperature maps corresponding to different seasons. However, this yields a non-symmetric covariance matrix, and therefore its singular value decomposition must be used, instead of the PCA. The simulations used for this analysis however do not contain correlated noise, and hence we will not worry about these complications. We leave the analysis of such instrumental issues for future work.

3.3 Independent component analysis

Independent component analysis (ICA) tries to solve the blind equation (equation 4) under the assumption that the sources \mathbf{s} are statistically independent from each other ($P(\mathbf{s}) = \prod_i P_i(s_i)$). Explicitly this is enforced by using the central limit theorem, according to which, if \mathbf{x} is made up of a linear combination of linearly independent sources, its distribution should be ‘more Gaussian’ than those of the independent sources. Therefore, we can attempt to impose statistical independence by maximizing any statistical quantity that describes non-Gaussianity.

FASTICA (Hyvärinen 1999) is a relatively popular and computationally efficient algorithm to apply ICA to a general system. Here, we will outline the operations carried out by FASTICA as well as its similarities with PCA. As before, we label the brightness temperature measured in different frequencies along a given line of sight $\hat{\mathbf{n}}$ by a vector \mathbf{x} ; however, the reader must bear in mind that FASTICA is provided with a number of samples of \mathbf{x} (i.e. different pixels), which it uses to compute the expectation values in the equations below.

FASTICA begins by ‘whitening’ the data \mathbf{x} . This implies first of all performing a full PCA analysis of the data, decorrelating them and sorting them by the magnitude of the covariance eigenvalues. The uncorrelated variables are then divided by their standard deviations so as to impose a unit variance on all of them (this is done in order to simplify the subsequent steps in the algorithm). We are then left with the corresponding equation for the whitened data $\tilde{\mathbf{x}} \equiv \hat{\Lambda}^{-1/2} \hat{\mathbf{U}} \mathbf{x}$:

$$\tilde{\mathbf{x}} = \hat{\mathbf{A}}' \mathbf{s} \equiv \hat{\Lambda}^{-1/2} \hat{\mathbf{U}} \hat{\mathbf{A}} \mathbf{s}, \quad (16)$$

where as before $\hat{\mathbf{U}}$ and $\hat{\Lambda}$ are the eigenvectors and eigenvalues of the data covariance matrix. Note that although the data have been whitened, FASTICA still preserves their PCA order. The problem is further simplified by requiring that the sources be also ‘white’ (i.e. uncorrelated and unit variance), since that implies that $\hat{\mathbf{A}}'$ must be orthonormal.

FASTICA then tries to find the independent components by inverting equation (16):

$$\mathbf{s} = \hat{\mathbf{W}} \tilde{\mathbf{x}}. \quad (17)$$

and finding the rows of $\hat{\mathbf{W}}$ by maximizing the non-Gaussianity of the individual components. FASTICA parametrizes the level of non-Gaussianity in terms of the negentropy $J(y) = H(y_G) - H(y)$, where $H(y)$ is the entropy of the variable y and y_G is a unit-variance Gaussian random variable. Since for all variables with the same variance the entropy is maximal for a Gaussian variable, the negentropy J is always positive for non-Gaussian variables. However, computing the negentropy requires an intimate knowledge of the probability distribution, which in general we lack. For this reason FASTICA uses an approximation to J given by

$$J(y) \sim \sum_i k_i [\langle G_i(y) \rangle_\theta - \langle G_i(y_G) \rangle_\theta], \quad (18)$$

where (k_i) are positive constants, $\langle \cdot \rangle_\theta$ denotes averaging over all available samples (i.e. pixels) and G_i is a set of non-quadratic functions. FASTICA makes use of two functions in particular:²

$$G(y) = e^{-y^2/2}, \quad G(y) = \frac{1}{a} \log \cosh(a y), \quad 1 \leq a \leq 2. \quad (19)$$

² No significant difference in the final foreground-cleaned maps was found for either choice of function.

The rows of $\hat{\mathbf{W}}$ are found as follows.

- (i) First an initial unit random vector \mathbf{w} is chosen as one of the rows of $\hat{\mathbf{W}}$ and its projection on the data $\mathbf{y} \equiv \mathbf{w}^T \mathbf{x}$ is computed.
- (ii) The correct direction of \mathbf{w} is found by maximizing the negentropy of \mathbf{y} . This is done by an iterative Newton–Raphson algorithm (renormalizing the modulus of \mathbf{w} after each step).
- (iii) The process ends when a given Newton–Raphson iterations yields a vector that is close enough to the one from the previous step.

Further details regarding the algorithm and its relation to other existing ICA methods can be found in Hyvärinen (1999) and Hyvärinen & Oja (2000).

It is important to note that the second step in FASTICA’s algorithm (namely, the maximization of the negentropy) will be unable to separate Gaussian foreground components beyond the initial PCA, since their negentropy is identically 0. This is simply a consequence of the fact that ICA is equivalent to PCA for the case in which all the sources are Gaussian.³

Unfortunately, as described in Section 4.1, the foreground maps corresponding to point sources and free–free emission were generated as Gaussian random fields in our simulations, and therefore we will not be able to appreciate the full power of ICA in removing those. However, the bulk of the GSE (which is the most relevant foreground for intensity mapping) was extrapolated from the Haslam map (Haslam et al. 1982), and is therefore non-Gaussian. Thus we should still be able to evaluate the impact of ICA methods in foreground removal.

3.4 Summary

As we have shown in the previous three sections, the three blind methods under study in this paper can be understood as three different ways of approaching the least-squares problem of equations (4) and (6), each following different criteria in order to find the basis functions $\hat{\mathbf{A}}$.

- (i) *Line-of-sight (polynomial) fitting* imposes an ad hoc set of basis functions that, we think, should be able to describe the foregrounds accurately enough.
- (ii) *PCA* finds $\hat{\mathbf{A}}$ by requiring that the sources should be uncorrelated and account for the majority of the total variance in the data.
- (iii) *ICA*, like PCA, maximizes the total variance but instead of *uncorrelatedness* it imposes *statistical independence*. Since these two properties are equivalent for Gaussian variables, PCA and ICA are mathematically equivalent when all the sources are Gaussian.

4 METHOD

In order to compare the three blind methods discussed in the previous section, we have tested them on a suite of 100 fast simulations

of the cosmological signal and most relevant foregrounds. In this section, we will describe the characteristics of these simulations as well as the criteria used to compare the different methods.

4.1 The simulations

The simulations used here were generated using the publicly available code presented in Alonso, Ferreira & Santos (2014).⁴ The code generates random realizations of the cosmological signal and most relevant foregrounds, and includes a basic set of instrumental effects.

The cosmological signal. We generate the cosmological signal by producing a three-dimensional Gaussian realization of the matter density and velocity fields in a cubic grid. The Gaussian density field is transformed into a more physically motivated one using the lognormal transformation (Coles & Jones 1991). During this step, we also apply a linear redshift-dependent bias and lightcone evolution based on linear perturbation theory to the overdensity field. This field is interpolated on to a set of temperature maps at different frequencies (i.e. redshifts) using equation (2). At this point, redshift-space distortions are included using the radial velocity field. For these simulations, we used the parametrization $x_{\text{HI}} = 0.008(1+z)$ for the neutral hydrogen fraction, based on the observations of Zwaan et al. (2005) and Wolfe, Gawiser & Prochaska (2005), and for simplicity we assumed that H I is an unbiased tracer of the density field ($\delta_{\text{HI}} = \delta$). Further details regarding this method can be found in Alonso et al. (2014).

A box of size $L_{\text{box}} = 8240 \text{ Mpc } h^{-1}$ was used with a Cartesian grid of 2048 cells per side. This way we are able to generate full-sky maps to redshift $z_{\text{max}} = 2.55$ with a spatial resolution of $\Delta x \simeq 4 \text{ Mpc } h^{-1}$. The cosmological signal was generated using Planck-compatible parameters: $(\Omega_{\text{M}}, \Omega_{\text{b}}, \Omega_{\text{k}}, h, w_0, w_a, \sigma_8, n_s) = (0.3, 0.049, 0, 0.67, -1, 0, 0.8, 0.96)$, and the H I temperature maps were produced in the frequency range 400–800 MHz ($0.78 \lesssim z \lesssim 2.55$) at intervals of $\delta\nu \sim 1 \text{ MHz}$, corresponding to radial separations of $r_{\parallel} \sim 5\text{--}6 \text{ Mpc } h^{-1}$. These maps were created using the HEALPIX pixelization scheme (Górski et al. 2005) with a resolution parameter $\text{nside}=512$ ($\delta\theta \sim 0^\circ.115$), corresponding to transverse scales $3.8 \text{ Mpc } h^{-1} < r_{\perp} < 8.2 \text{ Mpc } h^{-1}$.

The foregrounds. The foregrounds included in these simulations can be classified into two categories: isotropic and anisotropic. We expect that extragalactic foregrounds, due to astrophysical sources beyond our galaxy should be isotropically distributed across the sky, while galactic foregrounds should be significantly more prominent closer to the galactic plane.

The most relevant foreground in this frequency range in terms of amplitude is unpolarized GSE. In broad terms, synchrotron emission was simulated by extrapolating the 408 MHz Haslam map (Haslam et al. 1982) to other frequencies using a direction-dependent spectral index provided by the Planck Sky model (Delabrouille et al. 2013). Structure on angular scales smaller than the ones resolved in the original Haslam map and a mild frequency decorrelation was included by adding a constrained Gaussian realization with parameters given by (Santos et al. 2005). Further details can be found in Alonso et al. (2014).

³ It is a common misconception to explain the ability of ICA to remove radio foregrounds by saying that it searches for non-Gaussianity in the frequency direction, and thus separates the ‘noisy’ components (cosmological signal and instrumental noise) from the smooth foregrounds. It is important to stress the fact that non-Gaussianity is in reality maximized in angles and not in frequency (i.e. the negentropy is estimated by averaging over pixels), and that the motivation for this is to improve on the initial PCA separation by imposing statistical independence, rather than to distinguish noisy and smooth components.

⁴ <http://intensitymapping.physics.ox.ac.uk/CRIME.html>

Table 1. Instrumental parameters used for the simulations.

D_{dish}	15 m
t_{obs}	10 000 h
$\delta\nu$	1 MHz
N_{dish}	254
T_{inst}	25 K
$(\nu_{\text{min}}, \nu_{\text{max}})$	(400, 800) MHz
f_{sky}	0.5

Other relevant foregrounds are point sources and free-free emission. We have modelled these as Gaussian realisations of the generic power spectra:

$$C_l(\nu_1, \nu_2) = A \left(\frac{l_{\text{ref}}}{l} \right)^\beta \left(\frac{\nu_{\text{ref}}^2}{\nu_1 \nu_2} \right)^\alpha \exp \left(-\frac{\log^2(\nu_1/\nu_2)}{2\xi^2} \right), \quad (20)$$

with parameters given by Santos et al. (2005, SCK from here on; see also table 1 in Alonso et al. 2014). As mentioned in Section 3.3, ICA and PCA will only yield different results for non-Gaussian foregrounds, and therefore we will not be able to observe the advantage of using ICA over PCA for point sources and free-free emission. Any difference found between both methods will then be due only to the non-Gaussian galactic synchrotron foreground.

As has been mentioned in Section 2, even though we expect the 21 cm signal to be largely unpolarised, the presence of polarized foregrounds should be a concern for intensity mapping experiments due to polarization leakage. Since our aim in this paper is to compare the usefulness of different foreground cleaning methods, and since the amplitude of polarized foregrounds depends both on instrumental design and survey strategy, we have decided to postpone the analysis of polarized foreground removal for future work. Nevertheless, in Section 5.3, we have studied the minimum degree of frequency correlation that foregrounds must have in order to be efficiently removed by blind methods.

Instrumental effects. The final observed temperature maps were generated by including the most relevant instrumental effects corresponding to a single-dish experiment (i.e. an instrument made up of a number of single-dish antennas used as a set of autocorrelators) with the parameters listed in Table 1. These are similar to those planned for the Square Kilometre Array (SKA)-MID configuration, for which a single-dish strategy is optimal for cosmological purposes (Bull et al. 2014). Two instrumental effects were implemented: a frequency-dependent beam and uncorrelated instrumental noise.

Ideally, we can parametrize the beam of such a system as being Gaussian with a simple frequency-dependent width:

$$\theta_{\text{FWHM}} \sim \frac{\lambda}{D_{\text{dish}}}, \quad (21)$$

where D_{dish} is the antenna diameter. For the frequency range under consideration this yields an angular resolution that ranges between 0.7 and 1.4 .

Furthermore, we have described the instrumental noise as being Gaussian and uncorrelated, with a frequency-dependent, pointing rms:

$$\sigma_N = T_{\text{inst}} \sqrt{\frac{4\pi f_{\text{sky}}}{\delta\nu t_{\text{obs}} N_{\text{dish}} \Omega_{\text{beam}}}}, \quad (22)$$

where T_{inst} is the instrument temperature, f_{sky} is the observed sky fraction, $\delta\nu$ is the frequency resolution, t_{obs} is the total observation time, N_{dish} is number of dishes and $\Omega_{\text{beam}} = 1.133 \theta_{\text{FWHM}}^2$ is the

beam solid angle. For the values listed in Table 1 σ_N varies between 0.025 and 0.05 mK.

The process to build the full observed temperature maps is as follows.

(i) We directly add the temperature maps corresponding to the cosmological signal and the different foregrounds. For computational ease and to avoid redundant calculations, we then degrade the angular resolution of the total map from $n_{\text{side}} = 512$ to 128, which is still perfectly compatible with the beam size ($\theta_{\text{pix}} \sim 0.45$).

(ii) We smooth the total temperature using a Gaussian beam with full width-half-maximum (FWHM) θ_{FWHM} .

(iii) We add a random realization of uncorrelated Gaussian noise with variance σ_N . Note that the per-pixel variance is similar to the one given in equation (22), with the total number of pixels N_θ substituting the factor $4\pi f_{\text{sky}}/\Omega_{\text{beam}}$.

At this point the observed maps should be cut to the desired survey mask. However, in order to simplify the analysis and to avoid the complication of accurately estimating the angular power spectrum of cut-sky maps, we decided to use full-sky maps for our fiducial simulations. The most important effect of an incomplete mask for foreground subtraction is a reduction in the number of pixels that can be used to determine the statistical properties of the foregrounds (e.g. estimating the covariance matrix for PCA or the negentropy for ICA). We have studied the relevance of this effect in Section 5.3.4.

Finally, we note that the parameters listed in Table 1 correspond to our fiducial set of simulations. However, in order to study the effect of different instrumental configurations, we have varied these parameters from their fiducial values. These variations are clearly stated in the corresponding sections below.

4.2 Clustering analysis

The most important cosmological constraints from H I intensity mapping will probably be obtained from the two-point statistics of the overdensity field, since this encapsulates the vast majority of the information regarding the underlying cosmological model. Hence, the merit of a given foreground cleaning method should be judged partly in terms of its ability to recover the true power spectrum on different radial and angular scales. We have therefore defined three quantities to characterize the efficiency of the different cleaning methods:

$$\eta \equiv \frac{\langle P_{\text{clean}} - P_{\text{true}} \rangle}{\sigma_P}, \quad \rho \equiv \frac{\langle P_{\text{res}} \rangle}{\sigma_P},$$

$$\epsilon \equiv \left\langle \frac{P_{\text{clean}} - P_{\text{true}}}{P_{\text{true}}} \right\rangle, \quad (23)$$

where $P_{\text{clean}}, P_{\text{true}}, P_{\text{res}}$ are the power-spectra of the cleaned maps, the true signal (cosmological signal and noise) and the residuals ('cleaned - true'), respectively, and σ_P is the statistical error in the power-spectrum, caused both by clustering variance and instrumental noise. Thus, ϵ and η quantify the bias induced by foreground subtraction on the power spectrum as a fraction of the true power spectrum and of the expected statistical uncertainties, respectively, while ρ is a measure of the corresponding loss in the signal itself. A perfect foreground cleaning would yield 0 for these three quantities.

We must clarify that we have used generic term 'power spectrum' here referring to both the angular power spectrum C_l and the radial power spectrum (defined below), which depend on (ν, l) and $(z_{\text{eff}}, k_{\parallel})$ respectively. This distinction is important: since the foregrounds

are assumed to be extremely smooth in frequency (i.e. the radial direction) we can expect the effects of foreground subtraction to be reduced to the largest radial scales.

It is also important to note that we have computed the ensemble averages in equation (23) by averaging over the 100 different realizations of the cosmological signal, and the foregrounds. We believe that, given the current uncertainties in the multifrequency description of the different relevant foregrounds, this approach is reasonably conservative in that, by averaging over foreground realisations, we effectively marginalize over these uncertainties. Consequently, the mean and variance of the quantities in equation (23) will depend not only on the statistics of the cosmological signal, but also on the model used to describe the foregrounds.

4.2.1 Angular power spectrum.

Assuming a full-sky survey, for a fixed frequency shell, the angular power spectrum of the brightness temperature fluctuations ΔT is estimated by first computing their harmonic coefficients:

$$a_{lm}(\nu) = \int d\hat{n}^2 \Delta T(\nu, \hat{n}) Y_{lm}^*(\hat{n}), \quad (24)$$

where $Y_{lm}(\hat{n})$ are the spherical harmonics. The angular power spectrum $C_l \equiv \langle |a_{lm}|^2 \rangle$ is then estimated by averaging over the symmetric index m

$$\bar{C}_l \equiv \frac{1}{2l+1} \sum_{m=-l}^l |a_{lm}|^2. \quad (25)$$

Unfortunately, this estimator is biased and non-optimal for cut-sky maps, and in fact the estimation of angular power spectra for incomplete maps is a non-trivial problem that has been widely treated in the literature (Tegmark 1997; Wandelt, Hivon & Górski 2001; Hivon et al. 2002). As we mentioned above, our fiducial set of simulations were generated as full-sky maps in order to avoid these complications. However, we have studied the effect of an angular mask in Section 5.3.4, and in this case the angular power spectra were computed using the POLSPICE software package (Chon et al. 2004).

4.2.2 Radial power spectrum.

Because astronomical observations are performed on the lightcone, cosmological observables such as the temperature perturbation ΔT are not homogeneous in the radial direction, since they evolve in time. For this reason, the only rigorous method to analyse the clustering pattern of the overdensity field in the radial direction is indirectly through the angular cross-correlation of maps at different redshifts (Montanari & Durrer 2012; Asorey et al. 2012). However, in practice it is possible to divide the survey into redshift shells that are wide enough to contain the most relevant radial scales and narrow enough so that we can assume a static universe at some effective redshift z_{eff} , neglecting evolution effects. This has in fact been common practice among most galaxy redshift surveys (Contreras et al. 2013; Anderson et al. 2014), and enables much more direct, easier access to important cosmological observables. In order to obtain a direct intuition about the effectiveness of foreground subtraction in the radial direction we have followed this approach here.

We have divided our full frequency range into a number of thick redshift shells of equal bandwidth $\Delta\nu \sim 100$ MHz, each containing the same number of frequency bins (see Table 2 for details). The location, width and number of these shells were chosen in order

Table 2. Frequency bins used for the analysis of the radial power spectrum. $\Delta\chi$ is the radial comoving distance between the bin edges, and $k_{\parallel, \text{max}}$ is the maximum radial wavenumber reached in that bin given the frequency resolution. For the three bins the minimum radial wavenumber $k_{\parallel, \text{min}} \equiv 2\pi/\Delta\chi$ is approximately $0.01 h \text{ Mpc}^{-1}$. Note that these bins were only used for the analysis of the radial power spectrum, while the foreground cleaning was done on the full frequency band ($400 \text{ MHz} < \nu < 800 \text{ MHz}$).

No.	$\Delta\nu$ (MHz)	z range	$\Delta\chi$	$k_{\parallel, \text{max}}$
1	436–537	1.64–2.25	$636 \text{ Mpc } h^{-1}$	$0.50 h \text{ Mpc}^{-1}$
2	537–638	1.22–1.64	$565 \text{ Mpc } h^{-1}$	$0.57 h \text{ Mpc}^{-1}$
3	638–739	0.92–1.22	$491 \text{ Mpc } h^{-1}$	$0.64 h \text{ Mpc}^{-1}$

to avoid the edge effects described in Section 5.3.2 and to include the most relevant cosmological scales. For each of these shells, every individual pixel corresponds to a different realization of the overdensity field along the line of sight. We can therefore compute the radial power spectrum by averaging over the modulus of the Fourier transform of each line of sight:

$$P_{\parallel}(k_{\parallel}) \equiv \frac{\Delta\chi}{2\pi N_{\theta}} \sum_{i=1}^{N_{\theta}} |\widetilde{\Delta T}(\hat{n}, k_{\parallel})|^2, \quad (26)$$

where $\Delta\chi = \chi(z_{\text{max}}) - \chi(z_{\text{min}})$ is the comoving width of the redshift shell under consideration. We estimate the Fourier coefficients $\widetilde{\Delta T}$ by computing the Fast Fourier Transform (FFT) of each line of sight. Since the different values of ΔT are separated by a constant frequency interval $\delta\nu = 1$ MHz, its FFT is calculated at constant interval in the conjugate variable corresponding to ν , $\delta k_{\nu} = 2\pi/\Delta\nu$. Assuming a constant effective redshift z_{eff} for each bin, k_{ν} can be related to the radial wavenumber k_{\parallel} through

$$k_{\parallel} = \frac{\nu_{21} H(z_{\text{eff}})}{(1 + z_{\text{eff}})^2} k_{\nu}. \quad (27)$$

A more rigorous definition for $P_{\parallel}(k_{\parallel})$ can be found in Alonso et al. (2014).

5 RESULTS

We have used 100 independent simulations as described in the previous section in order to compare the three blind subtraction methods and evaluate the ranges of scales in which the recovered maps can be reliably used for cosmology. We have studied the number of foreground degrees of freedom N_{fg} that need to be removed (i.e. the dimension of s in equation 4) for each method, the values of the parameters defined in equation (23) as a function of scale and method, and the limits of applicability of these techniques.⁵

5.1 Foreground degrees of freedom

By subtracting the foregrounds we are effectively removing information from the temperature maps. We should therefore be concerned about removing only the information related to the foregrounds, and minimizing the loss of cosmological information. This implies minimizing the number of degrees of freedom that we subtract. In the case of PCA, and assuming that the foregrounds are

⁵ We have made available a public version of the computer codes used to carry out this analysis at <http://intensitymapping.physics.ox.ac.uk/codes.html>.




Figure 1. Principal eigenvalues of the frequency covariance matrix for a simulation including a frequency-dependent beam (blue circles) and a constant beam (red squares), as well as for simulations containing foregrounds with different correlation lengths $\xi \in (1, 0.05)$. In the first two cases, there is a clear division between foreground eigenvalues and cosmological ones, although the presence of a ν -dependent beam makes the transition between the two smoother, and a larger number of foreground degrees of freedom must be eliminated. For smaller frequency correlation lengths, the contribution of foregrounds spreads through a larger number of eigenvalues, making foreground contamination more troublesome.

uncorrelated, this number would ideally correspond to the number of foreground sources.

In our case, we have included four different foregrounds (galactic synchrotron, point sources and galactic and extragalactic free-free). Additionally, the mean H I temperature is also a smooth function of the frequency, and will be inevitably removed by the foreground cleaning method. On top of this, a frequency-dependent instrumental beam will also act as an extra effective foreground, yielding a total of ~ 6 potential foreground sources, the cleaning of which will depend critically on them being clearly distinct from the cosmological signal we hope to measure.

PCA offers an intuitive way to address this question, and to decide on the number of foreground degrees of freedom to subtract. We know that most of the information related to the foregrounds should be contained in a few of the largest eigenvalues of the frequency covariance matrix. Thus, by looking at these eigenvalues it is possible to see whether a number of them is clearly different (much larger) from the rest, and whether there exists a clear divide between foreground and cosmological eigenvalues. Fig. 1 shows the first 25 principal eigenvalues for several of our simulations. One of them (blue circles) includes a frequency-dependent beam size given by equation (21), while another one (red squares) was generated for a constant-beam size corresponding to the median frequency of the simulation. It is possible to see that in both cases there exists a clear distinction between foreground and cosmological eigenvalues, although the presence of a frequency-dependent beam makes the transition between the two smoother. We can predict an optimal value of $N_{\text{fg}} = 5$ for the constant-beam simulation, and either 6 or 7 for the ν -dependent beam using PCA.

We have studied the total number of eigenvalues to subtract by computing the ratio of the power spectrum of the residual maps to the cosmological power spectrum in one of our simulations for the three methods and for different numbers of foregrounds. Fig. 2 shows this ratio for the angular (left-hand panel) and radial (right-hand panel)

power spectra and for the three different methods (in descending order: polynomial fitting, PCA and ICA). For this plot, we chose to show the results corresponding to an intermediate frequency bin $599 \text{ MHz} < \nu < 600 \text{ MHz}$ for the angular case, and the second bin in Table 2 for the radial case, but similar results hold in general. As anticipated above, both PCA and ICA converge for $N_{\text{fg}} = 6$ or 7, and we do not gain anything by subtracting additional degrees of freedom. For polynomial fitting however the optimal value is clearly $N_{\text{fg}} = 7$. In view of this result, we have performed the analysis below on the total ensemble of simulations using the fiducial value $N_{\text{fg}} = 7$ for the three methods.

5.2 The effects of foreground removal

In order to evaluate the extra statistical and systematic uncertainties introduced by the process of foreground subtraction, and to ascertain the range of scales where the cleaned maps can be reliably used to obtain cosmological constraints, we have computed the parameters η , ϵ and ρ introduced in equation (23) from our 100 independent simulations.

Fig. 3 shows the values of these parameters computed for the angular power spectrum in the frequency bin $599 \text{ MHz} < \nu < 600 \text{ MHz}$ (left-hand column) and for the radial power spectrum for the second bin in Table 2 (right-hand column). Each plot shows the results for the three cleaning methods, and the error bars show the statistical deviation of these quantities. In all cases, the three methods yield very similar (almost equivalent) results, all of them being able to clean the foregrounds reasonably well for the same number of foreground degrees of freedom. There is a very wide range of scales where foreground contamination is small in comparison with the expected statistical errors, and which can be reliably used for cosmology. In particular, PCA and ICA yield remarkably similar results, in spite of the latter exploiting the central limit theorem to enforce the statistical independence of the different foreground components. Thus, at least for the foregrounds simulated in this work, requiring the foregrounds to be uncorrelated (which is common to PCA and ICA) is the most relevant constraint. Furthermore, the results for polynomial fitting are very similar to the other two methods (even slightly better for certain radial scales), in spite of this method being the most naïve approach to foreground cleaning. It is also worth noting that both η and ρ become more consistent with 0 on large angular scales. This is due to the fact that these two quantities are weighed by the statistical errors (see equation 23), which are inversely proportional to the square root of the number of modes available ($\sigma(C_l) \propto (2l+1)^{-1/2}$). On large radial scales ($k_{\parallel} \lesssim 0.1 h \text{ Mpc}^{-1}$), we can clearly see the effect of foreground removal: due to the spectral smoothness of the foregrounds, any foreground leakage (or overfitting) will contribute mainly to the largest radial scales, which are catastrophically contaminated. However, in terms of the parameters η , ϵ and ρ we do not observe any clear degradation in the small- l regime (large angular scales).

It is important to study the validity of these results across different frequencies. Fig. 4 shows the values of η and ρ (upper and lower rows, respectively) computed for the angular power spectrum in the 400 different frequency channels (left-hand column) and the radial power spectrum in the three bins described in Table 2 (right-hand column) using a PCA approach. Equivalent results were found for polynomial fitting and ICA. As is evident from this figure, even though foreground removal is reasonably successful for a large range of frequencies, it breaks down close to the edges of the frequency band, where the recovered maps are highly biased. The cosmological analysis must therefore be performed on




Figure 2. Ratio of the power spectrum of the foreground cleaning residuals to the power spectrum of the cosmological signal as a function of the number of foreground degrees of freedom removed for the three different cleaning methods. The left- and right-hand panels show the results for the angular and radial power spectrum, respectively. Each plot contains three subpanels, showing the results for polynomial fitting, PCA, and ICA (in descending order). The angular power spectra shown correspond to a frequency bin $599 \text{ MHz} < \nu < 600 \text{ MHz}$, while the radial ones correspond to the second bin in Table 2. As can be seen, in most cases the efficiency of the foreground cleaning converges for $N_{\text{fg}} \sim 6\text{--}7$, although the distinction between both values is clearer for polynomial fitting, which attains a clear minimum for $N_{\text{fg}} = 7$. This result could have been anticipated visually from Fig. 1. The black lines in both plots correspond to the optimal case for the constant-beam simulation ($N_{\text{fg}} = 7$ for polynomial fitting and $N_{\text{fg}} = 5$ for PCA and ICA).

the frequency interval not affected by these edge effects. We will discuss this effect in more detail in Section 5.3.2. This effect is also evident in the case of the radial power spectrum: the first bin, which is close to the lower edge of the frequency band, is slightly affected by these edge effects and the foreground cleaning residuals are noticeably larger.

In order to quantify the performance of each foreground cleaning method with a single number, we have computed an effective bias η_{eff} for the angular power spectrum by averaging over all the values of l in the range of frequencies 450 MHz–750 MHz (i.e. omitting the frequencies close to the edges where the cleaned maps are less reliable). In all cases, we obtain an effective bias $\eta_{\text{eff}} \simeq -0.2$, i.e. the amount of signal loss in the power spectrum is on average 20 per cent of the statistical errors. The bias in the radial power spectrum drops below this 20 per cent level for scales smaller than $k_{\parallel} \simeq 0.15 \text{ h Mpc}^{-1}$. Throughout the same range of scales, the angular power spectrum of the cleaning residuals is also about one fifth of the statistical uncertainties.

The negative value of η_{eff} implies that we are in fact overfitting the foregrounds, and that there is a net leakage of the cosmological signal into the removed foregrounds which induces a non-zero (albeit small) bias in the power spectrum. We have studied the possibility of alleviating this signal loss by decreasing the number of

subtracted foreground degrees of freedom, but for any smaller N_{fg} the leakage of foregrounds into the signal became catastrophically large on certain scales.

This bias is therefore an undesirable but seemingly unavoidable effect of foreground cleaning which could, if accurately characterized, be potentially corrected for. However, the stochastic nature of this bias implies that, even if corrected for, it will induce an additional uncertainty in the measurement of the power spectrum. This additional uncertainty is given by the standard deviation of η , shown as error bars in the first row of Fig. 3. While the additional uncertainty on $P_{\parallel}(k_{\parallel})$ is relatively small (5–10 per cent for $k_{\parallel} > 0.1 \text{ h Mpc}^{-1}$), the statistical errors in the angular power spectrum must be enlarged by a significantly larger fraction (20–30 per cent). We can understand this result as due to the much higher degree of stochasticity of the foregrounds in the angular direction with respect to the radial one. These results are consistent for the three different cleaning methods and across the whole frequency range, and show that foreground removal will inevitably degrade the cosmological constraints that can be obtained from the measurement of the temperature power spectrum.

Besides this increase in the amplitude of the statistical uncertainties in the power spectrum, foreground removal could potentially also affect their correlation structure. Thus, it is important to




Figure 3. The parameters η , ϵ and ρ (rows 1–3, respectively), introduced in equation (23), for the angular power spectrum in the bin $\nu \in (599, 600)$ MHz (left-hand panel) and the radial power spectrum for the second bin in Table 2 (right-hand panel). Each plot shows the result for polynomial fitting (red circles), PCA (green squares) and ICA (blue triangles). The error bars in the plots show the variance of each of these quantities.

compare the full covariance matrices of the power spectrum for the cleaned and the true temperature maps. We have estimated the covariance matrix for a fixed frequency bin:

$$C_{ij}^v = \langle C_{l_i}(\nu)C_{l_j}(\nu) \rangle - \langle C_{l_i}(\nu) \rangle \langle C_{l_j}(\nu) \rangle \quad (28)$$

and for a fixed scale l

$$C_{ij}^l = \langle C_l(\nu_i)C_l(\nu_j) \rangle - \langle C_l(\nu_i) \rangle \langle C_l(\nu_j) \rangle \quad (29)$$

from our 100 simulations. Note that 100 independent realizations is not nearly enough for an accurate estimation of the covariance matrix (e.g. usually as many as the number of matrix elements are needed for a successful inversion); however, it is enough to get an idea of any evident effect that foreground cleaning could have on the correlation structure. The top row in Fig. 5 shows the correlation matrix ($r_{ij} \equiv C_{ij} / \sqrt{C_{ii}C_{jj}}$) for a fixed frequency $\nu = 600$ MHz for the true and foreground-cleaned maps (left- and right-hand panels, respectively), while the bottom row shows the analogous matrices for a fixed angular scale $l = 50$. According to these results, the variation in the correlation structure of the statistical uncertainties of the angular power spectrum caused by foreground subtraction is quantitatively negligible. This disagrees with the results of Wolz et al. (2014), who find a noticeable effect on the correlation matrix of

the angular power spectrum induced by foreground removal using FASTICA. This disagreement could be due to a number of reasons, including differences in the foreground simulations used in both analyses, differences in the criterion used to determine the number of foreground sources to subtract or the different sky completeness of the maps used in both analyses. After studying these covariance matrices for suboptimally cleaned maps (using $N_{\text{fg}} = 5$ instead of 7), we were not able to observe the effect reported by Wolz et al. (2014), and therefore the most plausible explanation is the different foreground simulations used in both analyses. We found analogous results for the correlation matrix of the radial power spectrum (see Fig. 6), in spite of the large bias induced by foreground subtraction on large radial scales.

Although the bias in the recovered power spectra is a good way to parametrize the effect of foreground removal in intensity mapping, certain observables, such as the BAO scale are known to be extremely robust against systematic alterations in the overall shape of the power spectrum. Therefore, it is also relevant to visualize directly the effect of foreground cleaning on the power spectrum. The left-hand panel of Fig. 7 shows the average angular power spectrum at $\nu = 600$ MHz computed for the foreground-free map (black solid line) and for the foreground-cleaned map (red solid line with




Figure 4. The parameters η and ρ from equation (23) (top and bottom rows, respectively) for the angular (left-hand column) and radial (right-hand column) power spectra for all the different frequency (or redshift) bins considered in this analysis. The results shown correspond to the PCA method, but equivalent results were found for polynomial fitting and ICA. In most cases, the bias due to foreground cleaning is much smaller than the statistical errors for a wide range of scales, although this breaks down on large radial scales and close to the edges of the frequency band.

shaded region showing the statistical uncertainties). Although the angular resolution of single-dish observations with the SKA-MID configuration is not good enough to measure the angular BAO, the overall shape of the power spectrum is not dramatically changed by foreground removal, and therefore it should be possible to constrain large-scale cosmological observables with intensity mapping. Analogously, the right-hand panel of Fig. 7 shows the radial power spectrum in the second bin of Table 2 divided by a smooth (no-BAO) fit to the overall shape of $P_{||}(k_{||})$, thus highlighting the BAO wiggles. The positions of the BAO wiggles are not significantly altered by foreground cleaning, and hence it should be possible to measure the radial BAO scale accurately with this configuration.

5.3 Other effects

As has been shown in the previous section, blind foreground cleaning methods are reasonably effective, and should allow us to recover the cosmological signal with a relatively small bias compared to the statistical uncertainties. However, it is interesting to explore how much this result depends on the assumptions and approximations used in this work. In this section, we have thus studied several

effects that could potentially affect the performance of foreground subtraction in cases that depart from the fiducial simulations used in the previous section. Since, as we have also seen, the performance of the three blind cleaning methods under study is almost equivalent, in this section we will show only results corresponding to a PCA analysis, although similar results hold for ICA and polynomial fitting.

5.3.1 Foreground smoothness

The success of foreground cleaning for intensity mapping relies heavily on the foreground sources being very correlated (smooth) in frequency. Even though we tried to be conservative in this regard when simulating the foregrounds, it is of key importance to quantify the minimum degree of smoothness required for a successful subtraction. By doing this we can, at least qualitatively, assess the effects of a potentially non-smooth foreground, such as polarization leakage.

The SCK model (equation 20) provides an ideal way to quantify this degree of smoothness in terms of the frequency correlation length ξ , which describes the number of e-folds in frequency over




Figure 5. Correlation matrix of the angular power spectrum at a fixed frequency $\nu = 600$ MHz (top row) and at a fixed scale $l = 50$ (bottom row) for the true signal (left-hand column) and the cleaned maps (right-hand column). In order to reduce the statistical fluctuations the covariance matrices have been rebinned in bins of $\Delta l = 4$ and $\Delta \nu = 6$ MHz.

which the foregrounds do not deviate significantly from a perfect correlation. We have generated foreground maps using Gaussian realisations of this model for different values of this parameter, from $\xi = 1$ (corresponding to the model used for point sources) to $\xi = 0.05$, and keeping all other parameters fixed to the values used to simulate extragalactic point sources. For each value, we generated 100 independent realizations, which were combined with the maps of the cosmological signal through the process defined in Section 4.1.

Fig. 1 shows the first 25 principal eigenvalues for simulations with different values of ξ . As ξ decreases, the foregrounds become less correlated in frequency and their contribution spreads over a larger number of eigenvalues. It is then easy to understand how foreground cleaning would fail: at some point the number of foreground degrees of freedom that must be subtracted becomes too large, and too much information about the cosmological signal is lost.

We performed a full PCA analysis on the 100 simulations for each value of ξ and computed the bias figure of merit η in each case. The result described above is explicitly shown in Fig. 8: as the foregrounds become ‘noisier’ more signal is lost in cleaning them, and

the bias in the estimated power spectra eventually becomes larger than the statistical errors. In view of this results, we estimate that an effective correlation length $\xi \gtrsim 0.25$ is necessary for a reliable foreground cleaning using a blind approach. This value was estimated as the correlation length for which the average effective bias $|\eta_{\text{eff}}|$ in the angular power spectrum becomes larger than 30 per cent.

It is a priori not easy to ascertain whether this result should be a concern for leaked polarized foregrounds. In particular, the frequency correlation of galactic polarized synchrotron depends on the description of phenomena such as the turbulent part of the galactic magnetic field, which are currently very poorly understood. We have used a realization of the polarized synchrotron emission generated with the code presented in Waelkens et al. (2009) in order to get an idea of the plausible range of values in which this ξ could lie. To do that, we computed the frequency correlation coefficient $\rho(v_1, v_2) \equiv \langle T(v_1)T(v_2) \rangle / \sqrt{\langle T(v_1)^2 \rangle \langle T(v_2)^2 \rangle}$ for the Stokes Q component from this realization in a few bands of galactic latitude. Fitting ρ to the model in equation (20) (i.e. $\rho(v_1, v_2) = \exp[-\log^2(v_1/v_2)/2\xi^2]$), we were able to determine that at high galactic latitudes ($b \gtrsim 60^\circ$) and in the high-frequency end




Figure 6. Correlation matrix of the radial power spectrum for the second bin in Table 2 for the true cosmological signal (left-hand panel) and the cleaned maps (right-hand panel).




Figure 7. Left-hand panel: angular power spectrum in the bin $\nu \in (599, 600)$ MHz for the foreground-free temperature maps (black solid line) and for the foreground-cleaned ones (red solid line). The grey-shaded region shows the variance of the cleaned power spectrum. Right panel: BAO wiggles in the radial power spectrum for the second bin in Table 2 in the same two cases.

of our band (~ 800 MHz), an effective correlation length of $\xi \sim 0.1\text{--}0.2$ can be expected. This value would become smaller at lower frequencies, and shows that polarization leakage should indeed be an important concern for intensity mapping experiments.

5.3.2 Edge effects

In order to verify that the larger cleaning bias found at high and low frequencies is not related to the specific frequency values or to a computational error in the simulation of the foregrounds, but to the fact that these regions are the boundaries of the frequency band under study, we have performed the following test: we re-analysed the simulations in a restricted frequency range, cutting out two bands of width 20 MHz at either end of it. By doing so we confirmed that the regions of unreliable foreground cleaning shift to the new band edges (see Fig. 9).

In an intensity mapping experiment, it is therefore important to allow for a buffer of frequencies at either end of the band in which the foreground cleaning is to be carried out, where the results will not be reliable. More interestingly, this justifies extending the frequency coverage of the survey beyond the values of interest for cosmology (i.e. above $\nu = 1420$ MHz) in order to improve the foreground subtraction.

5.3.3 Radio Frequency Interference

Until now we have not taken into account the possible presence of man-made radio frequency interference (RFI), which can completely prevent the usage of certain radio channels for astronomical purposes. The usual way to deal with RFI is to place the experiment in a remote location, out of reach of artificial electronic signals; however, there will inevitably be certain bands in the radio




Figure 8. Foreground cleaning bias as a function of the foreground frequency correlation length for the angular power spectrum at 600 MHz (left-hand panel) and for the radial power spectrum in bin 2 of Table 2 (right-hand panel).




Figure 9. Foreground cleaning bias for restricted frequency band. The first and last 20 MHz of the fiducial frequency band were cut out (grey, hash-marked bands) and the foreground subtraction was done in the restricted band. The region of large bias near the boundaries of the frequency range observed in Fig. 4 is now shifted to the new boundaries, showing that it is truly an edge effect that will affect blind foreground subtraction in general.

spectrum that will be rendered useless by RFI. The presence of RFI could affect foreground cleaning by limiting our ability to characterize the frequency dependence of the foregrounds.

We have attempted to quantify the effect on foreground subtraction by crudely simulating the presence of RFI in our simulations. We did so by flagging certain frequency channels as RFI and removing them entirely from the analysis. Two RFI models were simulated, which we labelled *random* and *clustered*: for a fixed fraction of flagged channels, *random* RFI is simulated by selecting those at random inside the simulated frequency range. For *clustered* RFI, on the other hand, the same number of flagged channels are collected in a small number of wide frequency bands or ‘clusters’. While the effect of *random* RFI on foreground subtraction should be minimal as long as the fraction of RFI is low enough, *clustered* RFI could have a more significant influence: if the clusters are wide enough in frequency they will act as effective ‘boundaries’ and, as shown in Section 5.3.2, the cleaned maps close to these edges will be less reliable.

A number of simulated observations were generated, varying the fraction of RFI-dominated channels and the number and width of RFI clusters. For *random* RFI, no significant degradation of the foreground cleaning process was observed even for RFI fractions of up to 20 percent, where the average effective bias rises only from -0.2 to $\eta_{\text{eff}} \simeq -0.22$ (left-hand panel in Fig. 10). However, for *clustered* RFI the boundary effect quoted above becomes relevant for clusters wider than $\Delta\nu_{\text{RFI}} \sim 20$ MHz (right-hand panel in Fig. 10). In this limiting case the average effective bias becomes $\eta_{\text{eff}} \sim -0.28$.

Another further complication caused by the presence of RFI is that a more involved treatment is necessary in order to study the clustering statistics of the cosmological signal in the radial direction, in the same way that angular masking affects the computation of the angular power spectrum C_l . The estimation of the radial power spectrum in the presence of RFI lies beyond the scope of the present work, however, and so we have not studied the corresponding effect on $P_{||}(k_{||})$.

5.3.4 Angular masking

There are a number of reasons why we cannot expect full-sky coverage for any realistic intensity mapping experiment. For example, ground-based experiments can only observe slightly more than one celestial hemisphere, depending on their location, and in most cases this maximal coverage is not reached due to technical or time limitations. We have studied the potential effects of incomplete sky coverage on foreground cleaning by performing a full PCA analysis on our fiducial set of simulations using different masking criteria. We can anticipate that angular masking should affect the efficiency of foreground removal in two different ways.

(i) On the one hand, a reduced sky fraction implies a smaller number of independent samples (pixels) that can be used to characterize the foregrounds (e.g. to calculate the frequency covariance matrix in the case of PCA), thus potentially reducing the quality of the cleaned maps. In order to quantify this effect, we created sky masks where only a region in the southern celestial hemisphere below a given declination $\delta_{\text{thr}} < 0$ is visible.

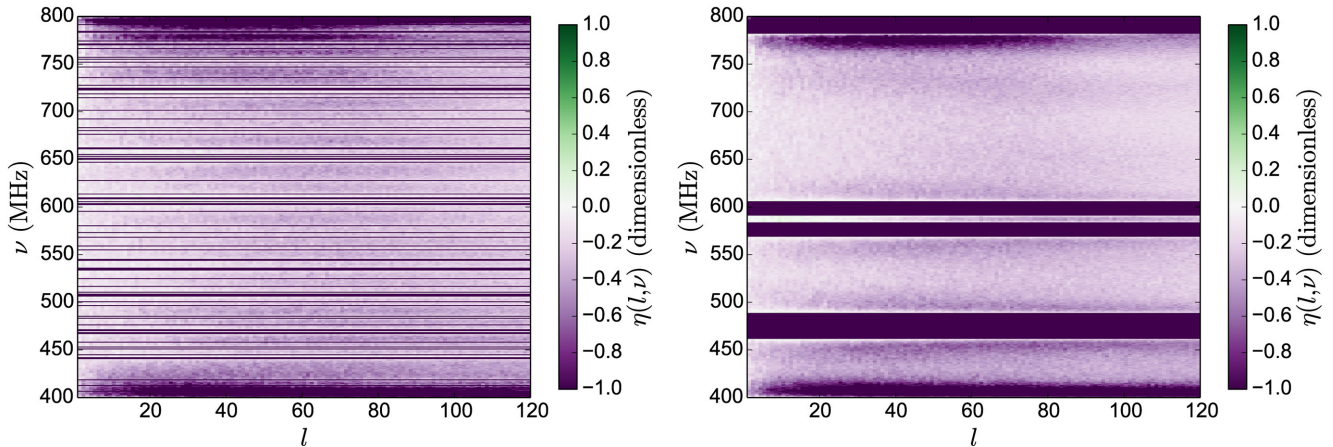


Figure 10. Bias parameter η for the angular power spectrum in the presence of *random* RFI (left-hand panel) and *clustered* RFI (right-hand panel). 20 per cent of the channels were flagged as RFI in both cases.

(ii) On the other hand, masking regions of the sky where we expect that foreground subtraction will be complicated (e.g. regions close to the galactic plane) can have the opposite effect, improving the efficiency of the method. In order to study this possibility, we masked all pixels where the synchrotron temperature at 408 MHz, given by the Haslam map, is above a given threshold T_{thr} . This threshold must be found as a compromise between covering regions of high synchrotron emission and minimizing the loss of sky coverage. After trying different values we found an optimal range of temperatures $T_{\text{thr}} \sim 30\text{--}50$ K in which the foreground residuals are minimized while retaining a large enough fraction of the sky. For instance, the choice $T_{\text{thr}} = 40$ K still leaves a sizeable fraction of the sky observable ($f_{\text{sky}} \sim 0.8$).

As mentioned before, and incomplete sky coverage also complicates the estimation of the angular power spectrum. For this reason, the software package POLSPICE (Chon et al. 2004) was used to estimate the C_l s, and we only studied scales on which we found this estimation to be unbiased for all the different masks ($l \geq 10$). Also note that the statistical uncertainties on the angular power spectrum increase for a masked sky by a factor $f_{\text{sky}}^{-1/2}$. In order to eliminate this effect from the comparison between different masks, in this case we used the ratio of the power spectrum of the cleaning residuals to that of the cosmological signal as a figure of merit.

Fig. 11 summarizes our findings. The figure shows the ratio $\langle C_{l,\text{res}} \rangle / \langle C_{l,\text{cosmo}} \rangle$ for the bin $600\text{ MHz} < \nu < 601\text{ MHz}$ and for different masks. As could be expected from the discussion above, as we decrease the observable fraction of the sky the cleaning becomes less efficient and the residuals grow, especially on large scales. However, when the appropriate parts of the sky are masked (i.e. regions with large foregrounds), the cleaned maps become more reliable, although the magnitude of this improvement is not very large.

5.3.5 Angular resolution

Due to the fiducial SKA dish size, a single-dish intensity mapping survey is only able to resolve very large scales ($\theta \gtrsim \theta_{\text{FWHM}} \simeq 1.4^\circ$). However, it is relevant to study the performance of blind foreground cleaning for an experiment with a better angular resolution. Two effects will be most relevant: on the one hand a better




Figure 11. Ratio of the power spectrum of the residuals to that of the cosmological signal for different masks: full-sky (red), half-sky ($\delta_{\text{thr}} = 0$, green), $\delta_{\text{thr}} = 60^\circ$ (blue), $T_{\text{thr}} = 40$ K (black) and $T_{\text{thr}} = 30$ K (magenta).

angular resolution provides a larger number of independent samples (pixels) that can be used to model the frequency structure of the foregrounds, thus potentially improving the cleaning. On the other, the statistical errors on smaller angular scales are also smaller ($\sigma(C_l) \propto (2l+1)^{-1/2}$), and hence the requirement on the foreground cleaning bias becomes more stringent.

We have generated observed sky maps for our 100 simulations assuming a constant angular beam size of $\theta_{\text{FWHM}} = 0.3$, and keeping all other instrumental parameters equal to their fiducial values (except for the angular resolution parameter, which was increased to $n_{\text{side}} = 256$). After applying the three blind methods, similar results to the ones quoted above for the fiducial simulations were found on all scales and frequencies (see Fig. 12), which shows that blind foreground cleaning should also be successful for more futuristic experiments. It is also worth noting that due to the fact that in this case we used a constant-beam width, only 5 foreground degrees of freedom had to be subtracted.




Figure 12. Bias parameter η for PCA in simulations with higher angular resolution ($\theta_{\text{FWHM}} = 0.3$). The cosmological signal is recovered to an accuracy similar to the one found in the fiducial case.




Figure 13. Foreground cleaning bias for a simulation with larger instrumental noise (twice as much as in the fiducial case). The degree of foreground contamination is similar to the one found for the fiducial simulations.

5.3.6 Instrumental noise

A large level of instrumental noise could also in principle prevent a correct characterization of the frequency structure of the foregrounds, and thus affect the effectiveness of foreground subtraction. In order to address this possible issue, we generated an ensemble of 100 simulations with a value of the system temperature T_{inst} twice as large as the fiducial one (quoted in Table 1), and used them to estimate the foreground cleaning bias η . This is shown in Fig. 13. In all cases, we are able to separate the smooth foregrounds from the cosmological signal and noise to the same level as in the fiducial case (Section 5.2).

6 DISCUSSION

In this work, we have studied the efficiency of blind foreground removal methods for H I intensity mapping. By ‘blind’ we refer here to methods that only assume generic characteristics about the foregrounds (in particular spectral smoothness). Due to the lack of multifrequency information about the foregrounds relevant for intensity mapping, blind cleaning methods will inevitably be necessary for the first experiments.

In particular, we have tested and compared three different methods: polynomial line-of-sight fitting, PCA and ICA. We have shown

that all of the methods can be described as different approaches to the same mathematical problem of blind source separation – all of them attempt to filter out the foregrounds by fitting their frequency dependence to a combination of smooth functions, and they only differ in their approach to finding these functions.

In order to carry out this study, we have generated a suite of 100 independent simulations of both the expected intensity mapping signal and the most relevant foregrounds using the publicly available code by Alonso et al. (2014). The simulated sky maps were then modified to simulate the angular resolution and instrumental noise level expected for the SKA-MID configuration (see Bull et al. 2014, for details).

In order to compare the different methods, we have characterized the efficiency of foreground cleaning through the quantities η , ϵ and ρ given in equation (23), which describe the bias induced on the power spectrum (both angular and radial) of the cleaned maps as well as the contamination in the signal itself. We find that the cosmological signal can be optimally recovered by removing a total of $N_{\text{fg}} = 7$ foreground degrees of freedom for the three methods. The analysis of the foreground-cleaned maps has yielded several important results.

(i) Not only can the three methods be described using the same mathematical formalism, but they also yield quantitatively similar (almost equivalent) results.

(ii) The cosmological signal can be successfully recovered across a wide range of scales and redshifts. However, the foreground removal induces a bias in the power spectrum that must be taken into account.

(iii) In the angular power spectrum, this bias is of the order of 0–20 per cent of the statistical uncertainties ($\eta \sim -0.2$) for all angular scales and most frequencies.

(iv) The radial power spectrum in turn is strongly affected by foreground removal on large scales ($k_{\parallel} \lesssim 0.1 \, h \, \text{Mpc}^{-1}$), but the bias η is rapidly suppressed on smaller ones.

(v) Foreground removal also increases the amplitude of the statistical uncertainties in the power spectrum by as much as 20–30 per cent in the case of the angular power spectrum. No relevant variation in the correlation structure of these uncertainties was observed, however.

(vi) In spite of this bias, we have verified that some of the most important cosmological features, such as the BAO wiggles, are not strongly affected by foreground removal, and therefore it should be possible to measure them with good accuracy.

(vii) Blind foreground cleaning relies heavily on the spectral smoothness of the foregrounds. In order to quantify the minimum degree of smoothness for a successful removal, we generated simulations for foregrounds with a varying frequency correlation length. We have determined that, in order to limit the average effective bias to $|\overline{\eta}_{\text{eff}}| < 0.3$, a frequency correlation length of $\xi > 0.25$ is necessary.

(viii) Close to the boundaries of the frequency band, the foreground-cleaned maps cease to be reliable. This suggests that a buffer of frequencies at either end of the band should be allowed for, and that an intensity mapping experiment could benefit from extending its observations beyond the frequency values of cosmological interest.

(ix) We have studied the effect of RFI (i.e. an incomplete frequency coverage) on the cleaned maps. Although foreground cleaning only deteriorates noticeably when up to 20 per cent of the frequency channels are lost due to RFI at random, the effect is much

stronger when the RFI-dominated channels are clustered into wider bands.

(x) We have also quantified two possible effects that an angular mask could have on foreground subtraction. On the one hand, we observe a deterioration in the cleaned maps for smaller fractions of the sky, due to the smaller number of independent samples (pixels) that can be used to characterize the foregrounds. On the other, the quality of the foreground-cleaned signal can benefit from masking regions of the sky where the foregrounds are larger.

We have left a number of analyses for future work, such as studying the effect of polarization leakage and correlated noise or characterizing the effects of foreground cleaning on the cosmological parameters inferred from the power spectrum. However, the results presented here should be relevant in order to produce more realistic forecasts for intensity mapping, optimise the design of future intensity mapping experiments and analyse the data extracted from them. The computer tools developed and used in this analysis have been made publicly available at <http://intensitymapping.physics.ox.ac.uk/codes.html>.

ACKNOWLEDGEMENTS

We would like to thank Gianni Bernardi, Thibaut Louis, Sigurd Næss and Laura Wolz for useful comments and discussions. DA is supported by ERC grant 259505. PB acknowledges support from European Research Council grant StG2010-257080. PGF acknowledges support from the Higgs Visiting Fellowship, STFC, BIPAC and the Oxford Martin School. MGS acknowledges support from the National Research Foundation (NRF, South Africa), the South African Square Kilometre Array Project and FCT under grant PTDC/FIS-AST/2194/2012.

REFERENCES

- Abdalla F. B., Rawlings S., 2005, *MNRAS*, 360, 27
 Alonso D., Ferreira P. G., Santos M. G., 2014, *MNRAS*, 444, 3183
 Anderson L. et al., 2014, *MNRAS*, 441, 24
 Ansari R. et al., 2012, *A&A*, 540, A129
 Asorey J., Crocce M., Gaztañaga E., Lewis A., 2012, *MNRAS*, 427, 1891
 Bagla J., Khandai N., Datta K. K., 2010, *MNRAS*, 407, 567
 Barkana R., Loeb A., 2005, *ApJ*, 624, L65
 Battye R. A., Davies R. D., Weller J., 2004, *MNRAS*, 355, 1339
 Battye R. A., Browne I. W. A., Dickinson C., Heron G., Maffei B., Pourtsidou A., 2013, *MNRAS*, 434, 1239
 Bernardi G. et al., 2009, *A&A*, 500, 965
 Bernardi G. et al., 2010, *A&A*, 522, A67
 Bull P., Ferreira P. G., Patel P., Santos M. G., 2014, preprint ([arXiv:1405.1452](https://arxiv.org/abs/1405.1452))
 Camera S., Santos M. G., Ferreira P. G., Ferramacho L., 2013, *Phys. Rev. Lett.*, 111, 171302
 Chang T.-C., Pen U.-L., Peterson J. B., McDonald P., 2008, *Phys. Rev. Lett.*, 100, 091303
 Chon G., Challinor A., Prunet S., Hivon E., Szapudi I., 2004, *MNRAS*, 350, 914
 Coles P., Jones B., 1991, *MNRAS*, 248, 1
 Contreras C. et al., 2013, *MNRAS*, 430, 924
 De S., Tashiro H., 2014, *Phys. Rev. D*, 89, 123002
 Delabrouille J. et al., 2013, *A&A*, 553, A96
 Di Matteo T., Perna R., Abel T., Rees M. J., 2002, *ApJ*, 564, 576
 Dickinson C., Davies R. D., Davis R. J., 2003, *MNRAS*, 341, 369
 Dickinson C. et al., 2009, *ApJ*, 705, 1607
 Furlanetto S. R., McQuinn M., Hernquist L., 2006, *MNRAS*, 365, 115
 Gleser L., Nusser A., Benson A. J., 2008, *MNRAS*, 391, 383
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Hall A., Bonvin C., Challinor A., 2013, *Phys. Rev. D*, 87, 064026
 Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, *A&AS*, 47, 1
 Hinshaw G. et al., 2013, *ApJS*, 208, 19
 Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
 Hyvärinen A., 1999, *IEEE Trans. Neural Netw.*, 10, 626
 Hyvärinen A., Oja E., 2000, *Neural Netw.*, 13, 411
 Jelić V. et al., 2008, *MNRAS*, 389, 1319
 Jelić V., Zaroubi S., Labropoulos P., Bernardi G., de Bruyn A. G., Koopmans L. V. E., 2010, *MNRAS*, 409, 1647
 Kaiser N., 1987, *MNRAS*, 227, 1
 Lidz A., Furlanetto S. R., Oh S. P., Aguirre J., Chang T.-C., Doré O., Pritchard J. R., 2011, *ApJ*, 741, 70
 Liu A., Tegmark M., 2011, *Phys. Rev. D*, 83, 103006
 Liu A., Tegmark M., Bowman J., Hewitt J., Zaldarriaga M., 2009, *MNRAS*, 398, 401
 Loeb A., Wyithe J. S. B., 2008, *Phys. Rev. Lett.*, 100, 161301
 McQuinn M., Zahn O., Zaldarriaga M., Hernquist L., Furlanetto S. R., 2006, *ApJ*, 653, 815
 Mao Y., Tegmark M., McQuinn M., Zaldarriaga M., Zahn O., 2008, *Phys. Rev. D*, 78, 023529
 Masui K. W. et al., 2013, *ApJL*, 763, L20
 Montanari F., Durrer R., 2012, *Phys. Rev. D*, 86, 063503
 Moore D. F., Aguirre J. E., Parsons A. R., Jacobs D. C., Pober J. C., 2013, *ApJ*, 769, 154
 Morales M. F., Bowman J. D., Hewitt J. N., 2006, *ApJ*, 648, 767
 Oh S. P., Mack K. J., 2003, *MNRAS*, 346, 871
 Perlmuter S. et al., 1999, *ApJ*, 517, 565
 Peterson J. B. et al., 2009, in *Astro2010: The Astronomy and Astrophysics Decadal Survey*. The National Academies Press, Washington, DC, p. 234
 Planck Collaboration XVI, 2014, *A&A*, 571, A16
 Pritchard J. R., Loeb A., 2008, *Phys. Rev. D*, 78, 103511
 Riess A. et al., 1998, *AJ*, 116, 1009
 Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*. Wiley, New York
 Santos M. G., Cooray A., Knox L., 2005, *ApJ*, 625, 575
 Seo H.-J., Dodelson S., Marriner J., McGinnis D., Stebbins A., Stoughton C., Vallinotto A., 2010, *ApJ*, 721, 164
 Shaw J. R., Sigurdson K., Pen U.-L., Stebbins A., Sitwell M., 2013, *ApJ*, 781, 57
 Shaw J. R., Sigurdson K., Sitwell M., Stebbins A., Pen U.-L., 2014, preprint ([arXiv:1401.2095](https://arxiv.org/abs/1401.2095))
 Switzer E. R. et al., 2013, *MNRAS*, 434, L46
 Tegmark M., 1997, *Phys. Rev. D*, 55, 5895
 Waelkens A., Jaffe T., Reinecke M., Kitaura F. S., Enßlin T. A., 2009, *A&A*, 495, 697
 Wandelt B. D., Hivon E., Górski K. M., 2001, *Phys. Rev. D*, 64, 083003
 Wang X., Tegmark M., Santos M. G., Knox L., 2006, *ApJ*, 650, 529
 Wolfe A. M., Gawiser E., Prochaska J. X., 2005, *ARA&A*, 43, 861
 Wolz L., Abdalla F. B., Blake C., Shaw J. R., Chapman E., Rawlings S., 2014, *MNRAS*, 441, 3271
 Wyithe J. S. B., Loeb A., 2008, *MNRAS*, 383, 606
 Wyithe J. S. B., Loeb A., Geil P. M., 2008, *MNRAS*, 383, 1195
 Zaldarriaga M., Furlanetto S. R., Hernquist L., 2004, *ApJ*, 608, 622
 Zwaan M. A., Meyer M. J., Staveley-Smith L., Webster R. L., 2005, *MNRAS*, 359, L30

This paper has been typeset from a \LaTeX file prepared by the author.