CSIRO PUBLISHING

Publications of the Astronomical Society of Australia, 2011, **28**, 46–57 www.publish.csiro.au/journals/pasa

Subtraction of Bright Point Sources from Synthesis Images of the Epoch of Reionization

B. Pindor[A,D], *J. S. B. Wyithe*[A], *D. A. Mitchell*[B], *S. M. Ord*[B], *R. B. Wayth*[C], *and L. J. Greenhill*[B]

[A] University of Melbourne, School of Physics, Parkville 3010, Australia
[B] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138-1516, USA
[C] ICRAR/Curtin Institute of Radio Astronomy, GPO Box U1987, Perth, WA 6845, Australia
[D] Corresponding author. Email: bpindor@unimelb.edu.au

Received 2010 July 2, accepted 2010 September 14

Abstract: Bright point sources associated with extragalactic active galactic nuclei and radio galaxies are an important foreground for low-frequency radio experiments aimed at detecting the redshifted 21-cm emission from neutral hydrogen during the epoch of reionization. The frequency dependence of the synthesized beam implies that the sidelobes of these sources will move across the field of view as a function of observing frequency, hence frustrating line-of-sight foreground subtraction techniques. We describe a method for subtracting these point sources from dirty maps produced by an instrument such as the MWA. This technique combines matched filters with an iterative centroiding scheme to locate and characterize point sources in the presence of a diffuse background. Simulations show that this technique can improve the dynamic range of epoch-of-reionization maps by 2–3 orders of magnitude.

Keywords: cosmology: diffuse radiation — methods: data analysis

1 Introduction

Highly redshifted 21-cm radiation emitted by neutral hydrogen gas during the epoch of reionization (EOR) contains a wealth of cosmological and astrophysical information. Several experiments aimed at detecting and characterizing this radiation are currently in progress or under construction (MWA,[1] PAPER,[2] LOFAR,[3] GMRT (Pen et al. 2008)). The EOR signal is also one of the main science targets of next-generation radio facilities such as the SKA.[4]

Following the cosmological recombination which produced the cosmic microwave background (CMB), the hydrogen in the intergalactic medium was essentially neutral. Subsequently, the first luminous sources produced a sufficient number of ionizing photons to eventually reionize the universe, leaving a small residual neutral fraction. However, the exact nature of the ionizing sources and the redshift or range of redshifts, $z_{\rm reion}$, at which this process occurred is very much an open question. Observations of the Gunn–Peterson trough in the spectra of high-redshift quasars indicate that reionization was essentially complete by redshift $z = 6$ (Becker et al. 2001). Meanwhile, the optical depth due to Thompson scattering observed for CMB photons by the *WMAP* satellite implies that for an instantaneous process $z_{\rm reion} = 10.9 \pm 1.4$ (Komatsu et al. 2009). Together, these observations support the theoretical expectation that reionization was an extended process (Pritchard et al. 2009). The expected redshift range of reionization, $z \simeq 6$–15, puts the 21-cm line at 200–90 MHz. Such low radio frequencies pose a number of observational challenges. Terrestrial transmissions from radio, television and satellite communications are all prominent at or near this band. Additionally, refraction of low-frequency radio waves by the ionosphere introduces time-variable distortions into the observed radio sky which require continuous re-calibration (Mitchell et al. 2008).

Observations of the EOR signal are further complicated by astrophysical foregrounds. These foregrounds are dominated by galactic diffuse synchrotron emission (GDSE) and extragalactic active galactic nuclei (AGN) (point sources), with a smaller contribution from galactic free–free emission (Shaver et al. 1999). The observed brightness temperature of the GDSE has a steep spectral index ($T \sim \nu^{-\beta}$; $\beta \sim 2.5$, Rogers & Bowman 2008) which makes it hundreds of times brighter at 150 MHz than in the 21-cm rest frame, while extragalactic point sources have a typical spectral index of $\beta \sim 2.7$ (Subrahmanyan 2002), with some evidence of flattening

[1] http://www.mwatelescope.org/.
[2] http://astro.berkeley.edu/~dbacker/eor/.
[3] http://www.lofar.org/.
[4] http://www.skatelescope.org/.

10.1071/AS10023 1323-3580/11/01046

at lower frequencies (Cohen et al. 2004). Interferometric observations are sensitive only to fluctuations in the sky brightness, and, taken together, these foregrounds are expected to have fluctuations at least three orders of magnitude greater than the EOR fluctuations (Santos et al. 2005). Initial observations of the diffuse foregrounds confirm these expectations (Ali et al. 2008 ;Bernardi et al. 2009; Bernardi et al. 2010).

Fortunately, the daunting task of isolating the comparatively faint EOR signal is made tractable by the spectral smoothness of the foregrounds. The GDSE exhibits an essentially featureless power-law spectrum, while the combination of many point sources of varying spectral indices can also be well reproduced as a smoothly varying function of frequency (Wang et al. 2006). In contrast, observing the EOR signal at different frequencies corresponds to probing the relatively rapidly varying intergalactic medium (IGM) density field across a substantial range of redshifts. Numerous authors have taken advantage of this distinction to demonstrate the possibility of subtracting foregrounds by fitting polynomials or other smoothly varying functions along the observing frequency/line-of-sight dimension (Wang et al. 2006; Geil et al. 2008; Gleser et al. 2008; Bowman et al. 2009; Harker et al. 2009; Liu et al. 2009a).

The effectiveness of these foreground-removal strategies assumes the prior removal of the brightest point sources. Bright point sources cannot be effectively removed by fitting spectrally smooth functions along the frequency direction/line of sight because the necessarily incomplete *uv*-coverage of any real interferometer inevitably creates a sidelobe pattern, evocatively termed 'frizz' by Liu et al. (2009b), across the plane of the sky. Changing the observing frequency changes the size of the synthesized beam and consequently moves this 'frizz' across any given point on the sky. In this way, high angular-frequency structure in the sidelobe pattern across the plane of the sky enters the observational frequency dimension, causing 'mode-mixing' in the line-of-sight power spectrum (Bowman et al. 2009). For this reason, most previous foreground removal studies assume there is some flux threshold $S_{\rm cut}$ above which all point sources can be removed. Bowman et al. (2009) and Liu et al. (2009b) agree that $S_{\rm cut}$ should be of order 10–100 mJy. However, the method by which these bright sources are to be removed remains an open question. In this paper, we introduce a new technique for subtracting bright point sources from synthesis images produced by an instrument such as the MWA.

2 The MWA EOR Experiment

The primary goal of the MWA EOR experiment is to measure the power spectrum of the cosmic neutral hydrogen density field. 21-cm radiation from neutral hydrogen will be visible against the CMB when the spin temperature, T_S, deviates from the CMB temperature, $T_{\rm CMB}$ (Wouthuysen 1952; Field 1959), with a predicted differential brightness temperature given by (Ciardi & Madau 2003):

$$\delta T_b = 26\,{\rm mK}\; x_{HI}(1+\delta)\left(1-\frac{T_{\rm CMB}}{T_S}\right) \times \left(\frac{\Omega_b h^2}{0.02}\right)\left[\left(\frac{1+z}{10}\right)\left(\frac{0.3}{\Omega_m}\right)\right]^{1/2} \quad (1)$$

where x_{HI} is the local hydrogen neutral fraction, δ is the local overdensity, and Ω_m and Ω_b are the cosmological matter and baryon densities respectively. During reionization, $T_S > T_{\rm CMB}$ (Madau et al. 1997), and hence neutral hydrogen can be detected in emission.

The detectability of the cosmological 21-cm signal is fundamentally limited by the thermal noise of the radio telescope. For an interferometer, each visibility is subject to a thermal noise contribution given by (e.g. Thompson et al. 2001)

$$V_{\rm rms} = \frac{2k_B T_{\rm sys}}{A_e\sqrt{df\,\tau}}, \quad (2)$$

where $T_{\rm sys}$ is the system temperature,[5] A_e is the effective area of a single antenna, df is the channel bandwidth, and τ is the integration time. For the MWA, the system temperature is dominated by the sky and is expected to be ~180 K at 178 MHz (Furlanetto et al. 2006). For an isotropic power spectrum,[6] this thermal noise is averaged over the thousands of baselines which sample each three-dimensional Fourier mode, leading to a formally highly significant detection.

The ability of the MWA to detect the EOR signal is in practice limited by the accuracy of the instrument calibration and foreground subtraction. As described in Mitchell et al. (2008), the high data rate output by the MWA correlator precludes the long-term storage of the raw visibilities, necessitating real-time calibration and imaging. The calibration cadence is set by the need to (i) resolve temporal variation in the ionospheric refraction; and (ii) update the direction-dependent antenna response model as sources move across the sky. The MWA will complete a calibration and imaging cycle every 8 seconds. The calibrated images formed from these 8-s integrations will then be co-added into stacks of ~8 minutes of observing time to further reduce the data-storage requirements. These co-added 8-minute dirty maps form the basic data product from which the offline data analysis of the EOR experiment will be done.

[5] As discussed in Morales (2005), this noise is implicitly averaged over two polarizations and $T_{\rm sys}$ is a factor of $\sqrt{2}$ lower than for a single dipole.

[6] The neutral hydrogen density field is expected to undergo some cosmological evolution over the full 32-MHz MWA bandpass, but this does not affect the present discussion.

3 A Description of the Method

3.1 Motivation

Before describing our subtraction technique, we briefly review the treatment of point sources in a number of observational regimes which will motivate our approach.

3.1.1 Radio Astronomy

The most well known and widely practiced method for removing point sources in radio astronomy is the CLEAN algorithm (Högbom 1974) and its many variants. CLEANing addresses one of the central difficulties in processing synthesis images: the intrinsic involvement of deconvolution. The Cotton–Schawb CLEAN variant (Schwab 1984), which subtracts the clean components from the ungridded visibilities and hence allows one to avoid aliasing and gridding errors, generally produces the best results. Nonetheless, the obtained dynamic range is usually limited to $\sim 10^3$ (Briggs & Cornwell 1992). More sophisticated approaches have been able to achieve dynamic ranges of $\sim 10^6$ (Voronkov & Wieringa 2004; Cotton & Uson 2008); however, applying these techniques to the MWA data would require real-time processing of the ungridded visibilites which would almost certainly exceed the MWA computing resources. Additionally, despite some theoretical efforts in quantifying the uncertainties associated with CLEANed images (Schwarz 1978), practical experience has shown that CLEANing often alters image statistics and leaves spatially correlated residuals (i.e. 'stripes', Cornwell et al. 1999) which would likely corrupt measurements of the EOR signal. Briggs (1995) showed that robust weighting can reduce deconvolution artifacts. Interestingly, Briggs also noted the similarities between robust weighting of the ungridded visibilities and Wiener filtering of the dirty map. By extension, these similarities would also apply to the matched filtering techniques we employ in this work. However, we note that subtraction of the spectrally smooth foregrounds strongly favours uniform over natural weighting (Liu et al. 2009b).

A second approach to subtracting radio point sources is 'peeling' (Noordam 2004; van der Tol et al. 2007; Mitchell et al. 2008; Intema et al. 2009). Peeling is essentially the sequential self-calibration and subtraction of the brightest sources in the field and is intended to overcome the calibration difficulties introduced by ionospheric refraction and direction-dependent gains which are inevitable in wide-field low-frequency radio observations. Peeling has the advantage of subtracting the point-source model from the ungridded visibilities, similar to the Cotton–Schawb CLEAN variant. However, due to the need to update the calibration solution on timescales which are short compared to the timescale for ionospheric variations, it is unlikely that more than ~ 100 sources will be peeled by the MWA real-time calibration, leaving hundreds of bright sources above the confusion limit. Additionally, the expected data rate from the MWA correlator ($\sim 40\,\mathrm{Gb\,s^{-1}}$) precludes long-term storage of the raw visibilities, implying that post-processing to remove foregrounds will likely be restricted to the gridded visibilities/dirty maps.

3.1.2 Optical Astronomy

The location and measurement of the flux of point sources are the most fundamental processes in optical astrometry and photometry. Images are usually convolved with the point-spread function (PSF) to maximize signal-to-noise. For reasonably oversampled images, it is routine to obtain astrometric centroids on the order of 10% of the pixel scale (Pier et al. 2003).

3.1.3 CMB

Removal of point sources is also an important foreground subtraction step in processing of CMB temperature maps. Matched filters are used to increase the contrast between the point sources and the CMB or other diffuse components (Tegmark & de Oliveira-Costa 1998). As a brief introduction to the method of matched filtering, suppose that we wish to identify a source of spatial profile $\tau(\mathbf{x})$ in a dataset, $D(\mathbf{x})$, so that

$$D(\mathbf{x}) = A\tau(\mathbf{x}) + N(\mathbf{x}) \tag{3}$$

where A is the amplitude of the source and $N(\mathbf{x})$ is the noise. To measure A, we apply some weighting function, $\psi(\mathbf{x})$, to the data so that our estimate of A becomes

$$A_{estimated} = \int \psi(\mathbf{x}) D(\mathbf{x}) d^2x \tag{4}$$

As in the case of optical data described above, we might simply choose $\psi(\mathbf{x}) = \tau$ in order to smooth $N(\mathbf{x})$ on the scale of τ. However, $N(\mathbf{x})$ need not be simply a thermal or photon noise which fluctuates from pixel to pixel. Instead, it can include a diffuse sky component or, more generally, it is the combination of all other components which contribute to $D(\mathbf{x})$. In this case, as shown by Haehnelt & Tegmark (1996), the optimal linear filter is one which upweights the source profile, while downweighting the scales at which the generalized noise, $N(\mathbf{x})$, is prominent. In Fourier space, this implies

$$\hat{\psi}(\mathbf{k}) \propto \hat{\tau}(\mathbf{k})/P(\mathbf{k}) \tag{5}$$

where $P(\mathbf{k})$ is the power spectrum of the generalized noise. Sources are then identified as peaks in an image of D which has been filtered (convolved) with ψ. Note that CMB beams are sufficiently compact that it suffices to identify and mask out the bright pixels associated with the main lobe. This is generally not the case for radio synthesis images.

3.2 Subtraction Procedure

Based upon the above considerations, we designed a procedure for subtracting bright point sources from synthesized radio images produced by an interferometric

instrument such as the MWA. As mentioned above, the very high MWA data rate and a finite storage capacity imply that only time-averaged dirty maps will be available for offline data analysis. Given a dirty map, our subtraction proceeds as follows:

(1) **Filtering:** We use a matched filter to optimize the signal-to-noise ratio and reduce contamination in the flux and centroid estimates caused by diffuse sky emission (GDSE). Following Equation 5, we approximate τ, the source profile, as a two-dimensional Gaussian fitted to the main lobe of the synthesized beam at the centre of the field of view (FOV). $P(\mathbf{k})$ is the power spectrum of the generalized noise. If there is no diffuse sky component, then $P(\mathbf{k}) = 1$ and matched filtering is equivalent to convolving the dirty image with our Gaussian beam model. If there is a diffuse sky component, then we estimate $P(\mathbf{k})$ iteratively from the data. Specifically, we first carry out our subtraction procedure assuming that $P(\mathbf{k})$ is constant. The result of this initial subtraction is a residual image with most of the point-source flux removed. We then measure $P(\mathbf{k})$ from this residual image and fit the measured values to a second-order polynomial in log space (i.e. $\log(P(k)) = a(\log(k))^2 + b(\log(k)) + c$). We then repeat our subtraction procedure using this model of $P(\mathbf{k})$ in Equation 5. In our simulations, we proceed to identify point sources as 5σ local maxima in the filtered images. For future real data, it is likely that the locations of most bright sources would be known *a priori* from higher frequency observations.

(2) **Centroiding:** We use the centroiding procedure developed for the Sloan Digital Sky Survey (SDSS (York et al. 2000)) photometric pipeline to measure the positions of detected sources in the filtered maps. This procedure uses Gaussian quartic interpolation to predict the centroid based on a local maximum and the eight surrounding pixels. There are two subtleties in this centroiding procedure: (i) the main lobes are not actually Gaussian, or even symmetric. Hence, the measured centroids are biased and the bias is direction-dependent. We correct for this bias in the same manner as the SDSS: by simulating the synthesized beam at the inferred centroid position and determining the offset between the expected and measured positions; and (ii) centroid and flux measurements of any given source are perturbed by the sidelobes of the other bright sources in the field. We address this problem through iteration: an initial estimate is made of the position and flux of each source neglecting the contribution of other sources, and this estimate is subsequently used to correct for the sidelobe contribution at the position of every source. In principle, this process could be repeated until the parameters for each source converged, but in practice the generation of each source with sub-pixel positional accuracy is computationally intensive and the results presented in this work are based on a single such iteration.

(3) **Aperture photometry:** We estimate the flux of a source through the use of a circular aperture at the measured centroid position in the match-filtered image. We correct for aperture bias by simultaneously measuring the aperture flux of a match-filtered image of the synthesized beam reconstructed at the centroid position.

4 Description of Simulations

We developed and tested our subtraction method through the use of a simulation pipeline which combines MAPS, the Multipurpose Array Performance Simulator (Wayth et al., in preparation), with the MWA calibration and imaging system (MWA Real Time System (MWA RTS), Mitchell 2008). MAPS is a software package for simulating the observed visibilities generated by a real interferometric array. In our case, it has been configured to simulate a 512-tile array with the antenna design and provisional layout of the MWA. Each simulation is a two-second snapshot integration at an observational frequency of 178 MHz (corresponding to the 21-cm line at $z = 7$) and over a channel bandwidth of 40 kHz. These simulations are computationally intensive as each integration contains over 10^5 visibilites. Consequently, we make two important simplifications at this stage.

First, we ignore the effects of thermal noise. Each visibility should be subjected to a thermal noise contribution given by equation 2. However, this level of thermal noise is not appropriate for EOR foreground subtraction as the foregrounds will be subtracted not from snapshots but rather from co-added maps which result from hundreds of hours of integration. Naively, we could introduce a longer integration time τ to scale the thermal noise, but this does not account for the scaling of the sidelobe level which would result from the simultaneous rotation synthesis. The correct scaling could be obtained by generating and co-adding snapshots corresponding to a $\sim$6 hr exposure; however, such a simulation is beyond our present computational resources and is a subject for future work.

Second, we assume that all MWA tiles have identical response (gain). By extension, this assumption implies that the MWA beam in our simulations is direction-independent. On the other hand, our imaging procedure and point-source subtraction technique do not make this assumption. Rather, they treat the data as in the general case of a direction-dependent beam. In future we intend to simulate a representative variation of tile gains, but in this work we only consider the simplest case of identical tiles. MAPS also has the ability to simulate the effect of the ionosphere on the observed visibilities for a given model sky, but in this work we neglect this effect.

The simulated visibilities produced by MAPS are subsequently processed into synthesized images by the MWA RTS. The MWA will make use of graphical processing unit (GPU) hardware to accelerate the imaging and calibration

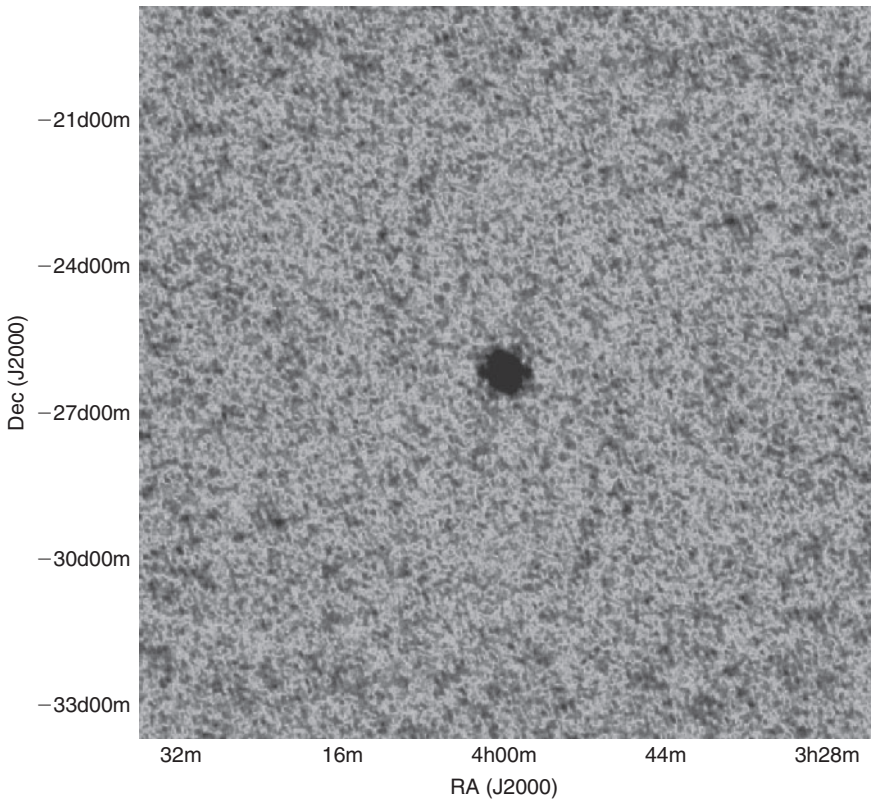

Figure 1 The synthesized beam for a 512-element MWA. A logarithmic brightness scale has been applied in order to emphasize the extended sidelobe structure.

pipeline (Edgar et al. 2010), but in this work we only used the CPU version. The resulting SIN projection images have a pixel scale of 1.9 arcmin/pixel at the field centre. We cropped the central 512 by 512 pixels of each image (corresponding to a FOV of ∼18 by 16 deg) and restrict our analysis to this region. This same simulation procedure is also used to reconstruct the synthesized beam at any point in the field. The resulting synthesized beam at the field centre is illustrated in Figure 1. The noiselike emissions seen throughout this image are in fact the sidelobes of the single, central source. These sidelobe fluctuations peak at ∼1% of the peak source flux.

5 Simple Tests

5.1 Single Point Sources

We begin by simulating 100 realizations of a randomly placed, isolated point source in an empty sky to measure the accuracy of our centroiding algorithm in the ideal case. These point sources have a flux of 100 Jy, although the flux normalization is not meaningful in this simple case. Applying our centroiding algorithm directly to the dirty maps yields an RMS centroiding error of 0.02 pixels. Measuring the centroid in the filtered images ($P(\mathbf{k}) = 1$ in this case) yields an RMS centroiding error of 0.002 pixels. Following correction for the centroiding bias described in Section 3.2, we obtain an RMS centroiding error of 10^{-4} pixels. These errors and all subsequent results are summarized in Table 1.

We also measure the change in the dynamic range of our simulated images following point source subtraction. The dynamic range is conventionally defined as the peak brightness divided by the RMS in empty parts of the image. We define empty regions as those pixels with counts less than twice the maximum value observed in our diffuse sky model (see below). In practice, this threshold is about 3% of the maximum source flux. We are most interested in the dynamic range improvement, which is the ratio of the dynamic ranges of the raw and subtracted images. Table 1 lists the dynamic range for the raw (unsubtracted) images, the images with sources subtracted from the match-filtered image, and the images with sources subtracted from the match-filtered image with the PSF bias correction also applied. For brevity, we subsequently refer to images with sources subtracted from the match filtered image as 'MF subtracted' and images with sources subtracted from the match-filtered image with the PSF bias correction also applied as 'MF+PSF subtracted'. The listed values are the logarithmic means of our 100 realizations. The dynamic range improvement for the MF+PSF subtracted images is 2.0×10^4. We illustrate the subtraction results for a single simulated source in Figure 2. Both panels in this figure show cross-sections through the brightest pixel in the field. The upper line in both panels is the raw image. We have plotted the absolute values to allow for a logarithmic scale. The point source is clearly visible near the center of the field, as are the ∼1% sidelobes. In the left panel, the lower line shows the residuals following MF subtraction. The dynamic range improvement in this case is ∼240. Note that the subtraction for this particular source happens to be substantially less accurate than the average of 100

Table 1. Point-source subtraction results

Subtraction step	Centroid error[a]	Raw dynamic range	Dynamic-range improvement
Single point source — no background			
Dirty image	0.02	5.3×10^2	1
Matched filter	0.002	4.2×10^5	7.9×10^2
PSF bias correction	10^{-4}	1.0×10^8	2.0×10^4
Single point source — diffuse background			
Dirty image	0.02	5.3×10^2	1
Matched filter	5.4×10^{-3}	1.6×10^5	2.9×10^2
PSF bias correction	5.5×10^{-4}	1.3×10^6	2.5×10^3
10 point sources — no background			
Dirty image	0.02	1.7×10^2	1
Matched filter	1.9×10^{-3}	1.2×10^5	6.9×10^2
PSF bias correction	5.9×10^{-4}	3.1×10^5	1.8×10^3
RTS peeling			2.2×10^3
10 point sources — diffuse background			
Dirty image	0.02	1.8×10^2	1
Matched filter	9.2×10^{-3}	3.4×10^4	1.9×10^2
PSF bias correction	5.5×10^{-4}	3.3×10^5	1.9×10^3
RTS peeling			1.1×10^3
100 point sources — no background			
Dirty image	0.02	2.4×10^2	1
Matched filter	9.7×10^{-3}	4.4×10^3	1.9×10^2
RTS peeling			4.6×10^3
100 point sources — diffuse background			
Dirty image	0.07	2.5×10^2	1
Matched filter	0.04	2.0×10^3	80
RTS peeling			26

[a]In pixels. 1 pixel = 1.9 arcmin or 1 arcsec $\sim 10^{-2}$ pixels.

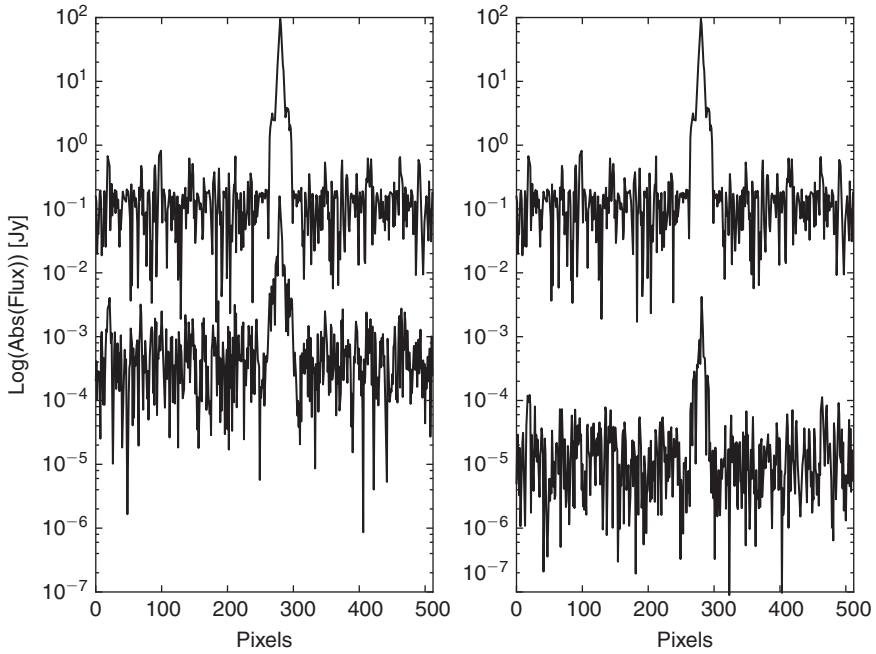

Figure 2 The results of point-source subtraction for a single source. In the left panel, the upper line represents the raw image and the lower line represents the residuals following subtraction from the match-filtered image. In the right panel, the upper line represents the raw image and the lower line represents the residuals following subtraction from the match-filtered image with the PSF bias correction applied. See text for more details.

realizations reported in Table 1 (i.e. 7.9×10^2). In the right panel, the lower line shows the residuals following MF+PSF subtraction. The dynamic range improvement in this case is $\sim 7.6 \times 10^3$. In both cases, significant residuals remain at the position of the point source. The question of whether these pixels should be masked out or if they will be further reduced by continuum foreground subtraction is a subject for future work.

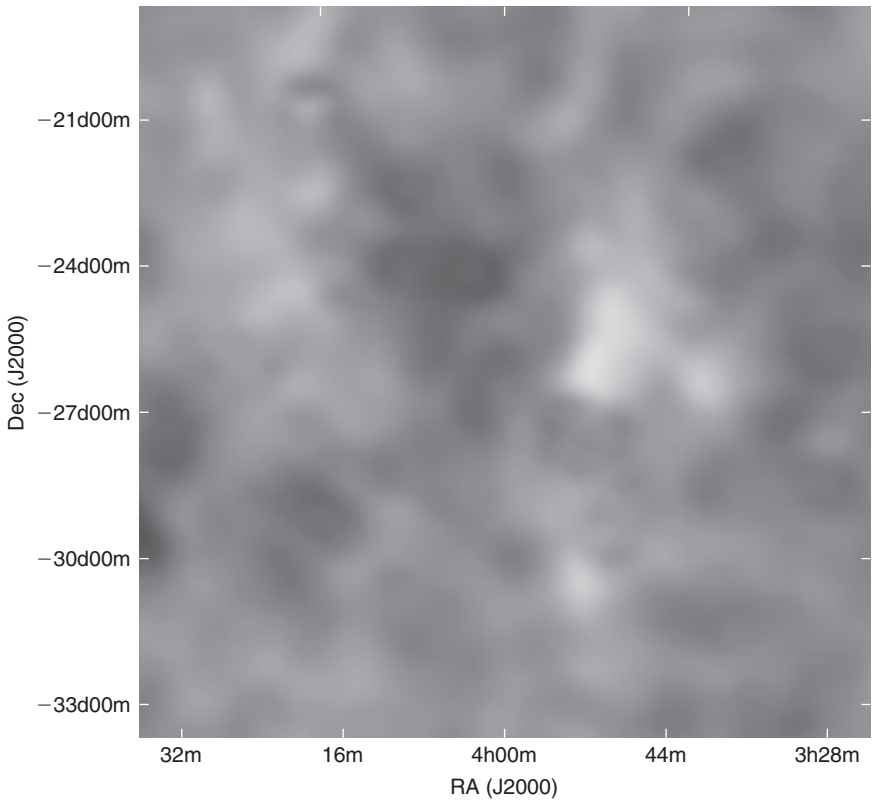

Figure 3 The diffuse sky component generated from de Oliveira-Costa et al. (2008). A linear brightness scale has been applied in this case.

5.2 Single Point Sources with a Diffuse Background

Next, we add a diffuse sky component based upon de Oliveira-Costa et al. (2008). We use a scaling appropriate to an observational frequency of 178 MHz, and locate the center of our field at the planned center of the main MWA EOR field (RA 4 h, Dec −26°). Hence, this diffuse component is identical for all of our simulations. Figure 3 illustrates the resulting diffuse emission. As before, we simulate 100 realizations of a randomly placed, isolated point source which are added to this diffuse background. The matched filter in this case uses an empirical $P(\mathbf{k})$ as described in Section 3.2. The RMS centroiding error following matched filtering and PSF bias correction is 5.5×10^{-4} pixels. In order to be able to directly compare the results of these simulations to those performed with no diffuse component, we subtracted a simulated image of the diffuse sky without any point sources (i.e. the result of a perfect point-source subtraction) from the residual image prior to calculating the dynamic range, so that the reported value is the dynamic range of the point-source residuals only. To clarify, this subtraction of the diffuse sky is performed only after we have completed our point-source subtraction procedure with the diffuse component present. The dynamic range improvement for the MF+PSF subtracted images is 2.5×10^3.

5.3 Multiple Point Sources

The centroid and flux estimates for each source are perturbed by the presence of all other sources in the field. We investigated this effect by creating simulations with 10 randomly placed sources of equal (100 Jy) flux in an empty sky. Equal flux sources correspond to the case of maximal mutual perturbation. Each realization has a different random distribution of sources within the FOV. The iterative correction for source centroids and fluxes described in Section 3.2 was applied. The RMS centroiding error following matched filtering and PSF bias correction is 5.9×10^{-4} pixels and the corresponding improvement in the dynamic range is 1.8×10^3. We note that the raw dynamic range in this case is predictably lower than in the single-source case due to the increased sidelobe noise.

5.4 Multiple Point Sources with a Diffuse Background

We added our diffuse sky model to the realizations of 10 randomly placed equal-flux point sources. The RMS centroiding error following matched filtering and PSF bias correction is 5.5×10^{-4} pixels and the corresponding improvement in the dynamic range is 1.9×10^3. Notably, the results are effectively the same with and without the diffuse component for these sufficiently bright sources.

5.5 Comparison with Peeling

We compared the results of our subtraction technique to the effectiveness of the MWA RTS peeling routine. As described above, peeling is performed on the ungridded visibilites, theoretically allowing for a higher-fidelity source subtraction. However, just as in the case of our image-based subtraction procedure, there are a number

practical considerations which limit the empirical effectiveness of peeling. Peeling requires an input catalog of sources which are to be peeled. In a real survey, this catalog would be likely be generated from higher-frequency observations (e.g. Parkes Source Catalog, Wright & Otrupcek 1996). It would also be possible to supplement the catalog with data from complementary low-frequency facilities (i.e. GMRT) or, in a bootstrapping fashion, from earlier MWA observations. In any case, the input catalog would invariably contain flux and position errors. Additionally, the observed source positions are perturbed from their true positions by the ionosphere. To account for these errors, the RTS peeling routine uses a local least-squares minimization to fine-tune the position and flux of peeled sources. For our simulations, the input catalogs used by the peeling procedure contain the true (simulated) positions and fluxes of the sources in the field, so that peeling with these values would produce a perfect subtraction. However, the local least-squares minimization in this case actually serves to slightly perturb the point-source parameters, leading to imperfect subtraction.

Further, diffuse emission complicates the peeling process. Peeling from the unweighted visibilities would produce flux errors of $\sim$1% even for the brightest sources. Since diffuse emission resolves out on longer baselines, weighting down short baselines in the least-squares minimisation — or removing them entirely — can help make peeling robust against this emission. However, the ideal weighting function is uncertain. Presently, the MWA RTS implements a simple scheme in which baselines are weighted by a factor of $1-\exp(-b^2/b_0^2)$, where b is the baseline length and b_0 is a scaling parameter, both in wavelengths. We choose $b_0 = 500$, which effectively saturates the scaling and results in only the longest baselines making a substantial contribution. Finally, due to the processing-time constraints, MWA RTS peeling is non-iterative within a cadence, meaning that the subtractions of the brightest sources are not updated to account for the subsequently inferred flux of the fainter sources. However, when processing a series of consecutive integrations, the peeling algorithm makes use of previous solutions to improve the current peel. This improvement is not represented in our single snapshot simulations.

For the simulations with ten sources in an empty sky described in Section 5.3, the measured dynamic range improvement after peeling is 2.2×10^3. With the diffuse sky component added, the measured dynamic range improvement is 1.1×10^3. We note that for these bright sources, the performance of our subtraction procedure from the dirty maps is comparable to peeling in the case of no diffuse background (1.8×10^3 compared to 2.2×10^3) and actually performs better in the case of a diffuse background (1.9×10^3 compared to 1.1×10^3). Although these examples do not constitute a definitive comparison of these techniques, they do illustrate that high-fidelity point-source subtraction is not necessarily contingent on access to the ungridded visibilities.

5.6 *A More Realistic Source Distribution*

As a further illustration of our subtraction procedure, we created simulations with and without diffuse background together with 100 sources drawn from a population whose counts are given by

$$N(>S_{\rm Jy}) = N_0 S_{\rm Jy}^{-2.5}\ {\rm Jy}^{-1}\ {\rm sr}^{-1} \qquad (6)$$

where we restricted the fluxes to be 1–100 Jy. For comparison, the 6C 151 MHz counts (Hales et al. 1988) predict $\sim$200 sources (>1 Jy) across our field of view. The brightest simulated source happens to have a flux of 34.6 Jy. The sources were randomly placed on the sky, subject to a minimum separation of 10 pixels. We found that, unsurprisingly, our perturbative approach to centroiding produces significant errors for close pairs of sources. For such sources a simultaneous source-fitting procedure such as DAOPHOT (Stetson 1987) may need to be incorporated. Alternately, higher-resolution radio data could be used to inform the centroid fits of blended sources. We have also ignored source clustering, although it will be an important consideration in more realistic sky models (Di Matteo et al. 2002). Figure 4 shows the resulting image for the case of sources with a diffuse background.

For the case of 100 sources with no background, the RMS centroiding error following matched filtering is 9.7×10^{-3} pixels and the corresponding improvement in the dynamic range is 190. The PSF bias correction is effectively irrelevant due to the magnitude of the input centroid errors. Figure 5 shows the centroiding errors as a function of source flux. The centroiding errors are largely independent of source flux, implying that the sidelobes of the brightest sources are being accounted for effectively while measuring the centroids of the fainter sources. For the case of 100 sources with a diffuse background, the RMS centroiding error following matched filtering is 0.04 pixels and the corresponding improvement in the dynamic range is 80. Again, the PSF bias correction is effectively irrelevant. The centroiding error is inversely correlated with source flux, as shown in Figure 6. The centroid fits of the fainter sources are significantly affected by the relatively bright diffuse background.

5.6.1 *Comparison with EOR Experiment Requirements*

Since we have access to the true flux and position of each simulated source, we are able to examine the dynamic range improvement for each individual source. For each source, this quantity is simply the dynamic range improvement when the selected source has been subtracted with the fitted position and flux, as above, while all of the other sources have been perfectly subtracted using the known true fluxes and positions. In other words, while the above dynamic range improvement values measure the factor by which the combined sidelobe level of all of the sources in the image has been reduced, the

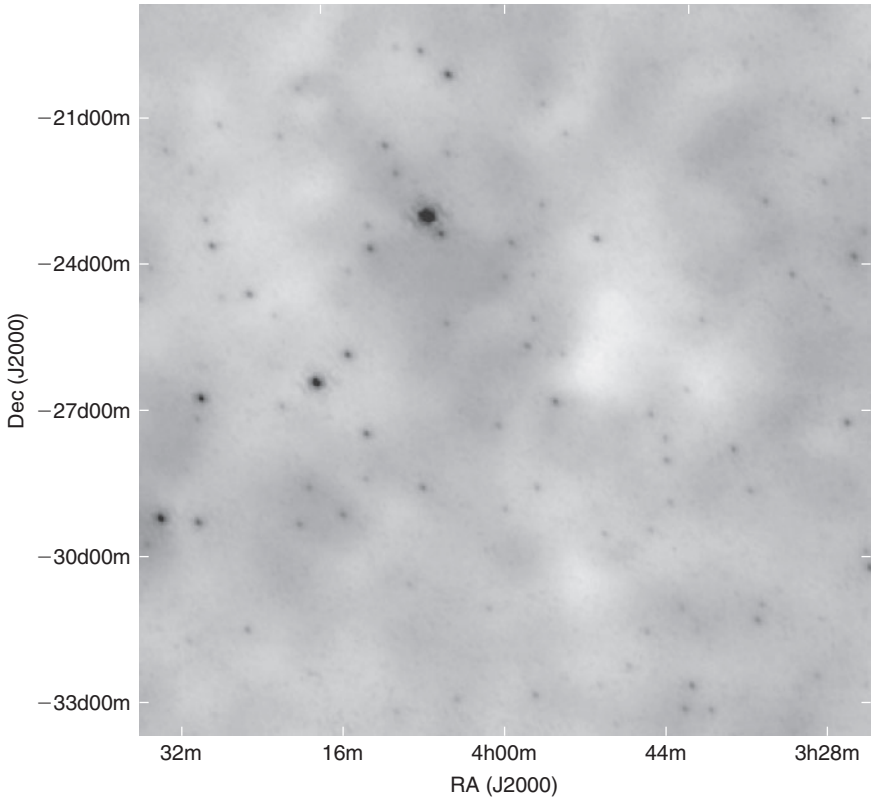

Figure 4 100 Point sources together with a diffuse background as described in Section 5.6.

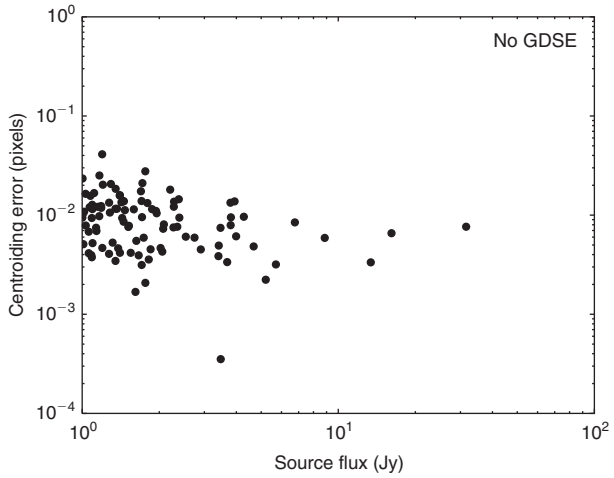

Figure 5 The centroiding errors (in pixels) as a function of simulated (input) flux for 100 sources with no diffuse background as described in Section 5.6.

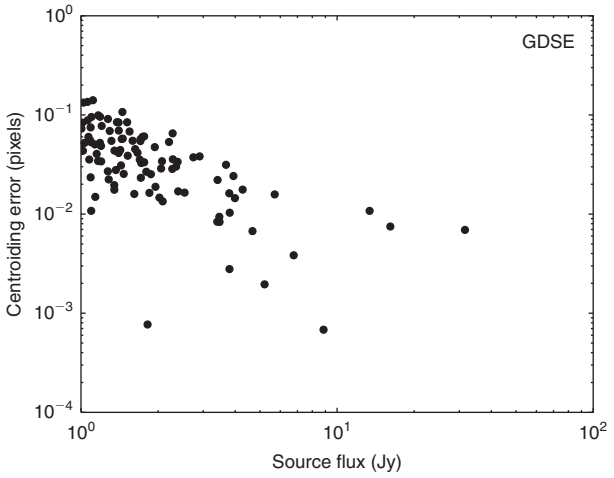

Figure 6 The centroiding errors (in pixels) as a function of simulated (input) flux for 100 sources with a diffuse background as described in Section 5.6.

individual dynamic range improvement measures the factor by which the sidelobes of a particular source have been reduced. Figure 7 shows the individual dynamic range improvement as a function of source flux in the case of no diffuse background. It is evident that the dynamic range is correlated with source flux; the brightest sources are least perturbed by the sidelobes of the other sources and their subtraction is relatively more accurate. This measurement also allows for a rough comparison to the values of S_{cut}, the flux limit to which bright sources are required to be removed, as reported by Bowman et al. (2009) and Liu et al. (2009a). For example, suppose that a 1 Jy source was subtracted with a 1% flux-only error to produce a dynamic range improvement of 100. The residuals would then correspond to an unsubtracted 10 mJy source. Such a source is sufficiently faint that the foreground continuum subtraction procedure could tolerate its presence and still achieve the sensitivity required to detect the EOR signal. Although our subtraction procedure produces both flux and centroid errors

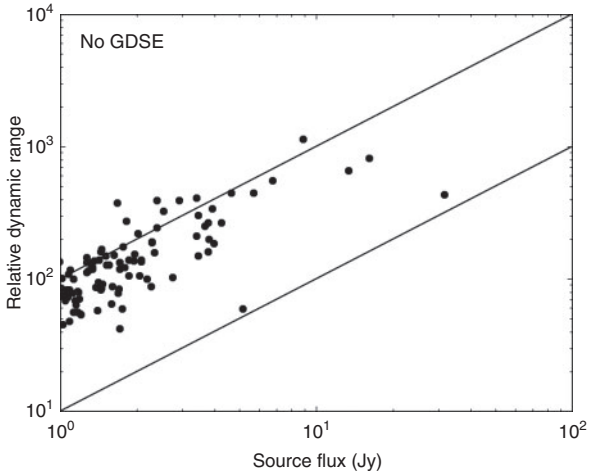

Figure 7 The individual dynamic-range improvement as a function of simulated (input) flux for 100 sources with no diffuse background as described in Section 5.6. The two solid lines represent the dynamic range improvement corresponding to $S_{\rm cut} = 10$ mJy (upper) and 100 mJy (lower).

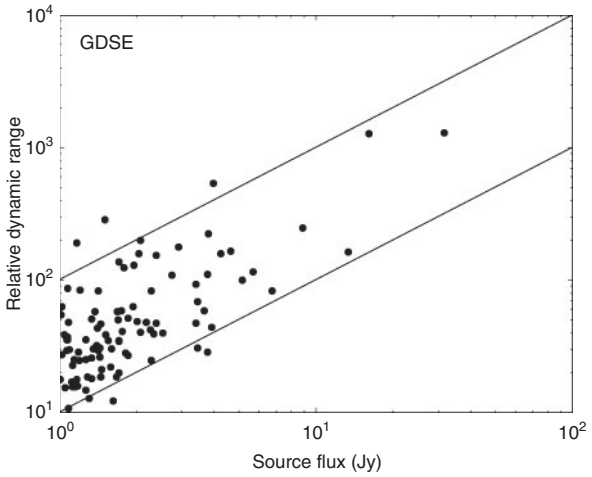

Figure 8 The individual dynamic range improvement as a function of simulated (input) flux for 100 sources with a diffuse background as described in Section 5.6. The two solid lines represent the dynamic range improvement corresponding to $S_{\rm cut} = 10$ mJy (upper) and 100 mJy (lower).

(as well as beam model errors in the case of calibration errors), the dynamic range improvement provides an aggregate measure of the magnitude of the residuals. The two solid lines in Figure 7 indicate the dynamic range improvement required to correspond to $S_{\rm cut}$ levels of 10 and 100mJy. Our subtraction exceeds the more optimistic 100mJy level for all sources and approaches the 10mJy level for most. Figure 8 shows the individual dynamic-range improvement as a function of source flux for the case with a diffuse background. As before, the dynamic range is correlated with source flux. In this case, nearly all sources exceed the 100 mJy cut level, but few exceed the 10 mJy level.

5.6.2 Comparison with Peeling

We also applied the RTS peeling, as in Section 5.5, to our 100 source images. For the case of no diffuse background, the dynamic range improvement after peeling is 4.6×10^3. Peeling is more accurate for this example than in the 10 point source case, presumably because the steep source counts imply that the brightest sources can be peeled first with relatively little sidelobe contamination. For the case of 100 sources with a diffuse background, the dynamic-range improvement after peeling is only 26. The residuals are dominated by a few intrinsically faint sources whose centroid estimates are stochastically degraded by the local least-squares minimization. This example should not be taken as demonstrating the limiting performance of RTS peeling; the peeling results could almost certainly be improved by some algorithmic fine-tuning. However, it does illustrate signal-to-noise limitations when peeling fainter sources.

6 Outstanding Issues

Our results indicate that subtraction of point sources from dirty maps can substantially improve the sidelobe noise in wide-field synthesis images. A simple estimate indicates that the accuracy of our subtraction attains or at least approaches the accuracy required for EOR detection as reported by previous foreground subtraction studies. However, a number of open questions must be resolved before it will be possible to unambiguously determine whether this approach can satisfy the requirements of an EOR experiment. In particular, outstanding issues include:

- **Calibration errors:** The accuracy of the calibration solution will be of vital importance to point-source subtraction, both for determining source fluxes and for accurately reconstructing the sidelobe pattern of each source. Datta et al. (2009) have recently investigated the calibration tolerances for an idealized peeling-like source subtraction. A comprehensive study of the calibration budget will need to account for a combination of complex effects including time-varying calibrators, non-idealized dipole response, and the effects of the ionosphere.
- **Multiple frequencies:** The MWA correlator will simultaneously generate visibilities for hundreds of frequency channels across a 32-MHz bandwidth. Recent work aimed at analyzing CMB data from the *Planck* satellite (Herranz et al. 2009) has demonstrated methods for extending matched filters to point-source detection in multi-frequency data. Application of these techniques to MWA data will effectively allow continuum fitting to be combined with the angular power spectrum to increase the contrast between point sources and the diffuse component.
- **Extended integrations:** As mentioned in Section 4, rotation synthesis will significantly alter (reduce) the sidelobes of all sources. Co-adding images produced over an extended observing campaign will also probably require an additional level of regridding. It will

also be necessary to model the time dependence of the synthesized beams. Presently, generating simulations which reproduce even a six-hour integration is a significant computational challenge.

- **Statistics of residuals:** We have used the dynamic range as a simple metric for judging the effectiveness of our subtraction procedure. In reality, a more detailed understanding of the subtraction residuals will be required to correctly ascertain the uncertainties in the measured EOR power spectrum. Bowman et al. (2009) have suggested the construction of statistical templates which could be used to distinguish foreground residuals from the EOR signal.
- **Sky model:** The sky model of (de Oliveira-Costa et al. 2008) uses input data which limit its angular resolution to ∼1 degree. Consequently, spatial power at small angular scales is underrepresented in our diffuse component. A method for introducing additional small-scale power into this sky model is currently in development (Bowman, private communication). An alternative is to build up a generic sky model from known sources of emission as done by Jelić et al. (2008).
- **Unresolved and extended sources:** As discussed in Section 5.6, our iterative centroiding technique performs poorly for marginally resolved sources. Our point-source model will also obviously be a poor fit for sources which have bright extended emission on angular scales comparable to the synthesized beam. It will likely be necessary to use higher frequency/longer baseline observations to constrain the radio morphology of the brightest sources in the MWA field, while also using simulations to investigate the level at which extended or unresolved sources can be tolerated.
- **Out-of-beam sources:** Bright sources which are located outside the primary beam can introduce significant sidelobe noise across the field of view. Our subtraction procedure, which relies upon the ability to locate the main lobe of the synthesized beam, may be inapplicable for such sources. Further, the direction-dependent antenna gains may be poorly constrained towards such sources. A satisfactory approach to dealing with out-of-beam sources remains an open question.

7 Conclusions

Bright point sources have previously been recognized as an important EOR foreground, but the method by which they should be removed has been unclear. In this work, we have presented a procedure for subtracting point sources from radio interferometric synthesis images. We are able to increase the dynamic range of our simulated images by a factor of 2–3 orders of magnitude. The efficacy of this method relies in large part on the excellent *uv* coverage of the MWA. These results are comparable to the results of the RTS peeling, but are achieved from the dirty maps, alleviating the need for long-term storage of the ungridded visbilities. Of course, peeling is an essential element of the MWA calibration strategy, and hence these two techniques will be used in a complementary fashion; peeling will remove some number of the brightest sources in real-time, and a subtraction procedure such as we have described will then be used offline to subtract a larger population of fainter sources from the time-averaged dirty maps. Significantly larger image simulations, particularly over longer integrations and multiple observing frequencies, are required to further refine and validate this approach. Nonetheless, our initial estimates indicate that this procedure will be able to remove point sources with sufficient accuracy to satisfy the foreground subtraction requirements of the MWA EOR experiment.

Acknowledgments

We thank Gianni Bernardi for helpful comments during the preparation of this work. We also thank the anonymous referee for a constructive report. We acknowledge the support of the Australian Research Council through grants LE0775621, LE0882938 and DP0877954, the U.S. National Science Foundation through grant AST-0457585 and the Smithsonian Astrophysical Observatory.

References

Ali, S. S., Bharadwaj, S. & Chengalur, J. N., 2008, MNRAS, 385, 2166
Becker, R. H. et al., 2001, AJ, 122, 2850
Bernardi, G. et al., 2009, A&A, 500, 965
Bernardi, G. et al., 2010, arXiv1002.4177
Bowman, J. D., Morales, M. F. & Hewitt, J. N., 2009, ApJ, 695, 183
Briggs, D. S., 1995, in Bulletin of the American Astronomical Society, Vol. 27, Bulletin of the American Astronomical Society, 1444
Briggs, D. S. & Cornwell, T. J., 1992, in Astronomical Society of the Pacific Conference Series, Vol. 25, Astronomical Data Analysis Software and Systems I, Eds. Worrall, D. M., Biemesderfer, C. & Barnes, J., 170
Ciardi, B. & Madau, P., 2003, ApJ, 596, 1
Cohen, A. S., Röttgering, H. J. A., Jarvis, M. J., Kassim, N. E. & Lazio, T. J. W., 2004, ApJS, 150, 417
Cornwell, T., Braun, R. & Briggs, D. S., 1999, in Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II, Eds. Taylor, G. B., Carilli, C. L. & Perley, R. A., 151
Cotton, W. D. & Uson, J. M., 2008, A&A, 490, 455
Datta, A., Bhatnagar, S. & Carilli, C. L., 2009, ApJ, 703, 1851
de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., Jonas, J., Landecker, T. L. & Reich, P., 2008, MNRAS, 388, 247
Di Matteo, T., Perna, R., Abel, T. & Rees, M. J., 2002, ApJ, 564, 576
Edgar, R. G., Clark, M. A., Dale, K., Mitchell, D. A., Ord, S. M., Wayth, R. B., Pfister, H., & Greenhill, L. J., 2010, ArXiv1003.5575
Field, G. B., 1959, ApJ, 129, 536
Furlanetto, S. R., Oh, S. P. & Briggs, F. H., 2006, Phys. Rep., 433, 181
Geil, P. M., Wyithe, J. S. B., Petrovic, N. & Oh, S. P., 2008, MNRAS, 390, 1496
Gleser, L., Nusser, A. & Benson, A. J., 2008, MNRAS, 391, 383
Haehnelt, M. G. & Tegmark, M., 1996, MNRAS, 279, 545
Hales, S. E. G., Baldwin, J. E. & Warner, P. J., 1988, MNRAS, 234, 919
Harker, G. et al., 2009, MNRAS, 397, 1138
Herranz, D., López-Caniego, M., Sanz, J. L. & González-Nuevo, J., 2009, MNRAS, 394, 510
Högbom, J. A., 1974, A&AS, 15, 417

Intema, H. T., van der Tol, S., Cotton, W. D., Cohen, A. S., van Bemmel, I. M. & Röttgering, H. J. A., 2009, A&A, 501, 1185
Jelić, V. et al., 2008, MNRAS, 389, 1319
Komatsu, E. et al., 2009, ApJS, 180, 330
Liu, A., Tegmark, M., Bowman, J., Hewitt, J. & Zaldarriaga, M., 2009a, MNRAS, 398, 401
Liu, A., Tegmark, M. & Zaldarriaga, M., 2009b, MNRAS, 394, 1575
Madau, P., Meiksin, A. & Rees, M. J., 1997, ApJ, 475, 429
Mitchell, D. A., Greenhill, L. J., Wayth, R. B., Sault, R. J., Lonsdale, C. J., Cappallo, R. J., Morales, M. F. & Ord, S. M., 2008, IEEE Journal of Selected Topics in Signal Processing, vol. 2, issue 5, pp. 707–717, 2, 707
Morales, M. F., 2005, ApJ, 619, 678
Noordam, J. E., 2004, in Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, Vol. 5489, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Ed. Oschmann, J. M., Jr., 817
Pen, U., Chang, T., Peterson, J. B., Roy, J., Gupta, Y. & Bandura, K., 2008, in American Institute of Physics Conference Series, Vol. 1035, The Evolution of Galaxies Through the Neutral Hydrogen Window, Eds. Minchin, R. & Momjian, E., 75
Pier, J. R., Munn, J. A., Hindsley, R. B., Hennessy, G. S., Kent, S. M., Lupton, R. H. & Ivezić, Ž., 2003, AJ, 125, 1559
Pritchard, J. R., Loeb, A. & Wyithe, J. S. B., 2009, ArXiv 0908.3891
Rogers, A. E. E. & Bowman, J. D., 2008, AJ, 136, 641
Santos, M. G., Cooray, A. & Knox, L., 2005, ApJ, 625, 575
Schwab, F. R., 1984, AJ, 89, 1076
Schwarz, U. J., 1978, A&A, 65, 345
Shaver, P. A., Windhorst, R. A., Madau, P. & de Bruyn, A. G., 1999, A&A, 345, 380
Stetson, P. B., 1987, PASP, 99, 191
Subrahmanyan, R., 2002, in IAU Symposium, Vol. 199, The Universe at Low Radio Frequencies, Eds. Pramesh Rao, A., Swarup, G. & Gopal-Krishna, 58
Tegmark, M. & de Oliveira-Costa, A., 1998, ApJL, 500, L83
Thompson, A. R., Moran, J. M. & Swenson, G. W., Jr., 2001, Interferometry and Synthesis in Radio Astronomy, 2nd Edition
van der Tol, S., Jeffs, B. D. & van der Veen, A.-J., 2007, IEEE Transactions on Signal Processing, 55, 4497
Voronkov, M. A. & Wieringa, M. H., 2004, Experimental Astronomy, 18, 13
Wang, X., Tegmark, M., Santos, M. G. & Knox, L., 2006, ApJ, 650, 529
Wouthuysen, S. A., 1952, AJ, 57, 31
Wright, A. E. & Otrupcek, R., 1996, VizieR Online Data Catalog, 8015, 0
York, D. G. et al., 2000, AJ, 120, 1579