

Limits on Fast Radio Bursts and other transient sources at 182 MHz using the Murchison Widefield Array

A. Rowlinson,^{1,2,3,4*} M. E. Bell,^{3,4} T. Murphy,^{4,5} C. M. Trott,^{4,6} N. Hurley-Walker,⁶ S. Johnston,³ S. J. Tingay,^{4,6} D. L. Kaplan,⁷ D. Carbone,¹ P. J. Hancock,^{4,6} L. Feng,⁸ A. R. Offringa,² G. Bernardi,^{9,10,11} J. D. Bowman,¹² F. Briggs,¹³ R. J. Cappallo,¹⁴ A. A. Deshpande,¹⁵ B. M. Gaensler,^{4,5,16} L. J. Greenhill,¹¹ B. J. Hazelton,¹⁷ M. Johnston-Hollitt,¹⁸ C. J. Lonsdale,¹⁴ S. R. McWhirter,¹⁴ D. A. Mitchell,^{3,4} M. F. Morales,¹⁷ E. Morgan,⁷ D. Oberoi,¹⁹ S. M. Ord,^{4,6} T. Prabu,¹⁵ N. Udaya Shankar,¹⁵ K. S. Srivani,¹⁵ R. Subrahmanyam,^{4,15} R. B. Wayth,^{4,6} R. L. Webster,^{4,20} A. Williams⁶ and C. L. Williams⁷

Affiliations are listed at the end of the paper

Accepted 2016 February 24. Received 2016 February 24; in original form 2015 September 18

ABSTRACT

We present a survey for transient and variable sources, on time-scales from 28 s to ~ 1 yr, using the Murchison Widefield Array (MWA) at 182 MHz. Down to a detection threshold of 0.285 Jy, no transient candidates were identified, making this the most constraining low-frequency survey to date and placing a limit on the surface density of transients of $< 4.1 \times 10^{-7} \text{ deg}^{-2}$ for the shortest time-scale considered. At these frequencies, emission from Fast Radio Bursts (FRBs) is expected to be detectable in the shortest time-scale images without any corrections for interstellar or intergalactic dispersion. At an FRB limiting flux density of 7980 Jy, we find a rate of < 82 FRBs per sky per day for dispersion measures $< 700 \text{ pc cm}^{-3}$. Assuming a cosmological population of standard candles, our rate limits are consistent with the FRB rates obtained by Thornton et al. if they have a flat spectral slope. Finally, we conduct an initial variability survey of sources in the field with flux densities $\gtrsim 0.5$ Jy and identify no sources with significant variability in their light curves. However, we note that substantial further work is required to fully characterize both the short-term and low-level variability within this field.

Key words: instrumentation: interferometers – techniques: image processing – catalogues – radio continuum: general.

1 INTRODUCTION

Until recently, little was known about the population of transient sources at low radio frequencies due to the lack of previous dedicated, sensitive surveys. Many of the known target transient populations are synchrotron sources, hence predicted to be faint and vary on long time-scales at low radio frequencies (such as afterglows from gamma-ray bursts and tidal disruption events; for a recent review see Metzger, Williams & Berger 2015). However, there are a number of different populations of sources that are expected to emit short-duration bursts of low-frequency coherent radio emission and are anticipated to be detectable in short snapshot low

radio-frequency images (e.g. giant pulses from pulsars and flares from brown dwarfs or exoplanets; Bastian, Dulk & Leblanc 2000; Berger et al. 2001; Jaeger et al. 2011; Law et al. 2011; Murphy et al. 2015).

One such coherently emitting target is the population of Fast Radio Bursts (FRBs; Lorimer et al. 2007; Thornton et al. 2013). FRBs were discovered at 1.4 GHz using high time resolution observations from the Parkes radio telescope. These sources constitute single, non-repeating, bright pulses of millisecond duration at 1.4 GHz that are highly dispersed, suggesting an extragalactic origin. A number of theories have been proposed as the progenitors of FRBs, including both extragalactic (e.g. Kashiyama, Ioka & Mészáros 2013; Totani 2013; Falcke & Rezzolla 2014; Lyubarsky 2014; Zhang 2014) and Galactic origins (e.g. Loeb, Shvartzvald & Maoz 2014). The scattering for FRBs is highly dependent upon

* E-mail: b.a.rowlinson@uva.nl

the observing frequency and is expected to smear out the pulse to much longer durations at low radio frequencies (Hassall, Keane & Fender 2013; Trott, Tingay & Wayth 2013). The pulse durations at low radio frequencies make them more difficult to detect using standard search methods at high time resolution. Instead, their durations are expected to be comparable to those attainable in short snapshot images. However, it is unclear what the rates of FRBs at low frequencies will be because the rates are still being constrained at higher frequencies and little is known about their spectral shape (e.g. Karastergiou et al. 2015; Keane & Petroff 2015). Therefore, observations at low frequencies will aid in constraining both the rates and the spectral slopes of FRBs. By more tightly constraining the rates, some progenitor mechanisms may be ruled out, including those associated with other populations with relatively low rates (such as short gamma-ray bursts; Zhang 2014). Additionally, all FRBs to date have been detected using single dish telescopes leading to large positional uncertainties (e.g. 14 arcmin; Thornton et al. 2013). By detecting FRBs in short snapshot image plane data observed using a low-frequency radio interferometer, their positions can be constrained to higher accuracy (<1 arcmin) enabling host galaxy associations and deep constraints on multi-wavelength counterparts. Additionally, an interferometer will obtain more reliable flux densities, as single-dish observations are subject to flux density uncertainties as the position of the source within the primary beam is unknown. This provides better constraints on the flux density distribution of sources ($\log N$ – $\log S$ distribution).

Over the past few years, the search for transient sources at low radio frequencies has intensified with the arrival of sensitive, wide-field instruments such as the Murchison Wide-field Array (MWA; Lonsdale et al. 2009; Tingay et al. 2013), the Low Frequency Array (LOFAR; van Haarlem et al. 2013) and the Long Wavelength Array Station 1 (LWA1; Ellingson et al. 2013). Additionally, the automated processing of very large data sets is being enabled via the long-term source monitoring capabilities of specially developed pipelines, including the LOFAR TRANSIENTS PIPELINE (TRAP; Swinbank et al. 2015) and the pipeline for the ASKAP Survey for Variables and Slow Transients (VAST; Murphy et al. 2013). Dedicated transient surveys are utilizing the improvement in instrumentation and software to constrain the surface densities of transients at these low frequencies on a range of time-scales and sensitivities (e.g. Bell et al. 2014; Carbone et al. 2014a; Obenberger et al. 2015; Stewart et al. 2016). Orders of magnitude improvement in sensitivity or search area will be required to more tightly constrain their rates. This can be attained by the next generation of radio telescopes, such as the Square Kilometre Array (SKA; e.g. Fender et al. 2015). However, obtaining the required observation time may be difficult on over-subscribed instruments and transient surveys will need to utilize commensal observations. This paper uses observations from one such data set, the MWA observations for the campaign to detect the Epoch of Re-ionization (EoR) in which hundreds of hours of observing time are required on individual fields. This data set can probe variability and transients on time-scales ranging from seconds up to years, enabling constraints to be placed on both the long time-scale incoherent emission mechanisms, e.g. synchrotron emission from active galactic nuclei (AGN), short time-scale coherent emission mechanisms such as FRBs and scintillation processes on a range of time-scales.

This paper describes a pilot transient and variability search using 78 h of the MWA EoR data set, producing highly competitive transient rates. The 28-s snapshot time-scale is chosen to specifically target the expected population of FRBs. This work complements Tingay et al. (2015), a search for FRBs using MWA observations

imaged on a much shorter integration time (2 s) and conducting an image plane de-dispersion to search for FRBs. Via this method, Tingay et al. (2015) are able to attain an improvement in sensitivity for FRBs in comparison to the standard processing strategies at the expense of processing speed and resolution, whereas a standard imaging strategy, such as that utilized in this paper, enables more observations to be processed in a comparable time-scale and the use of the data products for additional science such as longer duration transient and variability studies. Without de-dispersion, a dispersed FRB will be detected at a lower flux density in the short snapshot images as the original signal is averaged over both time and frequency. Therefore, these two approaches are complementary; Tingay et al. (2015) increase sensitivity by sacrificing surveyed area whereas the survey conducted in this paper sacrifices sensitivity to increase the amount of surveyed area. Additionally, candidate FRBs identified in this analysis can be independently confirmed as FRBs by measuring their dispersed signal using the pipeline developed (Tingay et al. 2015).

Section 2 of this paper describes the processing strategies used to make all the images and the analysis strategies implemented to conduct quality control and to search for transient sources. In Section 3, we present the limits on transients detected on a range of time-scales and focus on the implications for the rates and spectral shapes of FRBs by comparison to previous studies at other frequencies. Finally, Section 4 provides an initial analysis of variability of known sources within the field.

2 OBSERVATIONS AND PROCESSING METHOD

2.1 Data set

The data used in this paper are obtained from a commensally observed data set for the transients team and the EoR team. The full data set comprises >1000 h targeting three specific fields, well off the Galactic plane, centred on two different observing frequencies (154 and 182 MHz). Processing the full data set is a ‘big data’ scale computational challenge due to the supercomputing time required and the data volume at all of the processing stages. By targeting a subsample of the data set, we can develop automated strategies to make the data volume more manageable and quantify the supercomputing requirements for the full data set. In this study, we choose observations of a single-target field, centred on RA: 0.00 deg, Dec: $-27^{\circ}00'00''$, $-27^{\circ}00'00''$ (J2000), at the observing frequency of 182 MHz and at elevations of $>75^{\circ}$. This field is centred on the Galactic co-ordinates $l: 30^{\circ}636$, $b: -78^{\circ}553$ (30:38:08.4, $-78:33:10.6$). The observations were conducted by taking multiple pointed observations as the field drifts through five different azimuth-elevation pointing directions centred on zenith. These observations were then phase centred to RA: 0.00 Dec: $-27^{\circ}00'$ (J2000), with a primary beam half width half-maximum (HWHM) of $11^{\circ}3'$.

This leads to a sample size of 3010 individual observations of 2 min integration times, or 100 h, in the time range 2013 August 23 to 2014 September 14.

2.2 Imaging strategy

Our imaging strategy builds upon the MWA imaging pipeline developed for transient searches by Bell et al. (2014). As described in this section, we have updated and adapted this pipeline to utilize calibration and imaging tools specifically developed for the MWA,

Table 1. The WSCLEAN settings used to image all the observations presented in this analysis. All other settings were the default settings.

Setting	Value
UV range ($k\lambda$)	>0.03
Maximum number of clean iterations	20 000
Size of image (pixels)	3072
Size of one pixel (arcsec)	54
Stopping threshold for CLEAN (Jy beam^{-1})	0.2
Briggs weighting	-1

Table 2. The calibrator observations used in this analysis. Observation ID 1061661200 is the principal calibration data set, used to calibrate the rest of the data set, and was calibrated using the observation of 3C444.

Observation ID	Time (UT)	Azimuth (deg)	Elevation (deg)
1061650704 (3C444)	2013-08-27 14:58:08	63.43	74.63
1061657536	2013-08-27 16:52:00	90.00	76.28
1061659368	2013-08-27 17:22:32	90.00	83.19
1061661200	2013-08-27 17:53:04	0.00	90.00
1061663032	2013-08-27 18:23:36	270.00	83.19
1061664856	2013-08-27 18:54:00	270.00	76.28

targeting a consistent flux density scale across the data set. All data are initially processed by the MWA pre-processing tool COTTER (Version 3.3), which includes radio-frequency interference (RFI) flagging by AOFLAGGER (Version 2.6.1; Offringa et al. 2010, 2015).

We conduct an initial calibration on one of the observations at zenith (observation id: 1061661200, selected at random from a sample of good observations), with an integration time of 122 s, from the night of 2013 August 27 by transferring a time-independent, frequency-dependent calibration solution from a 2-min observation of 3C444. The calibrator 3C444 was chosen because it has a good calibration model, which is based on an extrapolation of its VLSSr (the Very Large Array Low-frequency Sky Survey Redux; Lane et al. 2014) flux density using a spectral index of -0.95 . The calibrated observation of the target field is imaged using WSCLEAN (Version 1.6 Offringa et al. 2014) using the settings in Table 1. The resulting model image is primary beam corrected using MWA specific tools (Version 1.1.0) to give a model image of the field. To optimize calibration and imaging for different Az-El (azimuth-elevation) pointing directions, we created calibration images from the same night for each of the unique lower elevation pointing directions where the pointing direction is approximately equal to that of ID 1061661200.¹ These observations, presented in Table 2, are calibrated using the 122 s image from ID 1061661200 and the phase and amplitude calibration method described below, producing a total of five calibration images for the full observations. This produces a consistent flux density scale across the calibration images for each of the pointing directions used in this data set.

For each unique Az-El pointing direction on each observing night, the target RA-Dec pointing direction changes by $\sim 3^\circ$ as the field tracks through the Az-El pointing direction. The observations are then shifted to a common RA-Dec phase centre of $0.00, -27.00$.

¹ Calibration can fail for observations on different nights, with differing Az-El and RA-Dec pointing directions due to the increased complexity and uncertainties in beam models. By choosing the same night and RA-Dec pointing we can achieve a reliable calibration image for each Az-El pointing that can be applied to other nights.

This leads to the centre of the primary beam shifting with respect to the centre of the image. To enable direct comparisons between images, avoiding the lower sensitivity regions, we use a conservative 12° radius (approximately the primary beam HWHM) from the centre of the image for the transient and variability analysis.

We calculate the MWA primary beam for each unique observation and, using this new primary beam, un-correct the relevant model calibrator Stokes-I image (taken from Table 2). This gives a non-primary beam-corrected model image of the expected sky that the MWA has observed. For the initial calibration of the non-zenith observations in Table 2 we used the zenith image (observation ID: 1061661200) and, following the creation of these calibrator images, for the imaging of the full data set we used the calibrator image corresponding to the identical Az-El pointing direction. Using the imager, WSCLEAN, this model image is converted to a clean component image which can be input into the calibration tools as a sky model. With the MWA specific tools CALIBRATE and APPLYSOLUTIONS, we complete a phase and amplitude self calibration on each observation using the newly created sky model.²

We note that the field of view will likely include different isoplanatic patches within the ionosphere, which can lead to issues when conducting self-calibration (e.g. Lonsdale 2005). However, due to the compactness of the MWA, the full array observes the same isoplanatic patches (regime 3 from Lonsdale 2005). This leads to the apparent positions of sources varying with time, but no deformations of sources (Intema et al. 2009). Due to the large number of sources spread across the sky model image, the MWA calibration strategy can account for this shift in position (e.g. Morales 2005). Additionally, as shown by Loi et al. (2015a), under normal ionospheric conditions, the typical positional shift due to the ionosphere at 182 MHz is ~ 10 arcsec, which is less than the resolution of the array. Therefore, the self-calibration strategy outlined here is not likely to be significantly affected by viewing different isoplanatic patches, although it may lead to small positional shifts in some of the sources.

Following calibration, we image each observation in four parts, corresponding to 28-s integration times, using the settings in Table 1 and a primary beam correction using the MWA specific tool BEAM.

2.3 Initial image rejection

The imaging strategy described in Section 2.2 is on the whole very successful, leading to a large number of images with consistent properties. However, on some occasions the images will not be of sufficient quality for transient and variability searches, for a variety of reasons such as calibration errors or significant ionospheric activity (e.g. Loi et al. 2015b), and we want to remove these images from the sample. The rms noise is a powerful indicator of the quality of the image, where rms values that deviate significantly from the expectations indicating that there are problems such as high RFI or calibration errors. We measure the mean rms noise in the central $\frac{1}{8}$ region of each of the images using the method described in Swinbank et al. (2015), where the sources have been excluded by rejecting pixels that are 4σ above the median rms.

A number of images had rms values that were highly deviant from the general population and some observations failed to image. These extremely low quality images were from the same nights, including the night of 2013 October 15, which was demonstrated

² Each observation ID is hence calibrated with a unique sky model which is based upon the calibration image.




Figure 1. A histogram of the rms noise measured in the images and fitted with a Gaussian distribution (solid red line). The solid black line is the mean rms observed, ~ 31 mJy beam $^{-1}$, and the black dashed lines represent the 2σ quality control rejection thresholds. A total of ~ 5 per cent of the images were rejected by this analysis.

to have significant ionospheric activity by Loi et al. (2015b). All observations from these bad nights were removed from the sample.

The rms values of the remaining 10 615 images are plotted in the histogram shown in Fig. 1 and are fitted with a Gaussian distribution.³ The rms distribution is sharply peaked with a typical rms of $31.4^{+1.9}_{-1.8}$ mJy beam $^{-1}$. MWA becomes confusion limited in images with an integration time on the order of 2 min and, hence, these images do not reach the classical confusion noise in the 30 s snapshot images. Additionally, due to their short integration times, these images are dominated by side-lobe confusion which is difficult to quantify. This result is unsurprising as all images are from the identical field with the same imaging settings, therefore we can easily identify low-quality images as outliers to this distribution. We reject all images with rms values in excess of 2σ from the Gaussian distribution (corresponding to the dashed lines in Fig. 1) as this limit corresponds to where the observed distribution is deviating from the Gaussian distribution. Following this quality control step, 10 122 images remain in the sample used in the remainder of this paper.

2.4 Correlated noise between observations

The rms noise in each of the images in the sample, characterized in Section 2.3, is made up of a number of components including the thermal noise, classical confusion noise and side-lobe confusion noise. The thermal noise component consists of a Gaussian random noise term made up of the sky noise and receiver noise, which is ~ 2 mJy beam $^{-1}$ at 182 MHz (Wayth et al. 2015). Due to the compactness of the MWA, classical confusion noise dominates the thermal noise component as multiple faint sources will be detected

within the restoring beam. This component can be estimated using source counts from surveys (Condon 1974). The density of faint sources at 182 MHz is currently poorly known and we need to extrapolate from surveys at other frequencies. Here, we calculate the classical confusion noise by extrapolating the source counts from the VLSS survey at 74 MHz (Cohen et al. 2007) assuming a spectral index of -0.7 , using

$$\sigma_{\text{conf}} = 29 \left(\frac{\theta}{1 \text{ arcsec}} \right)^{1.54} \left(\frac{\nu}{74 \text{ MHz}} \right)^{-0.7} \mu\text{Jy beam}^{-1} \quad (1)$$

where σ_{conf} is the classical confusion noise, θ is the size of the restoring beam and ν is the observing frequency. At our survey frequency of 182 MHz, the classical confusion noise is estimated to be 7 mJy. The 28 s images presented in this survey are therefore dominated by side-lobe confusion and other sources of noise. The side-lobe confusion noise is caused by imperfect deconvolution of sources either within the image, which we have minimized by using a self-calibration strategy, or in the side-lobes of the primary beam. Both of the sources of confusion noise are examples of correlated noise, where the noise for specific positions within the images between consecutive images is related to each other.

MWA images are subject to having correlated noise caused by sources drifting through the side-lobes of the primary beam (side-lobe confusion) and these images are likely to have correlated noise caused by the classical confusing sources below the detection threshold. The correlated noise component caused by side-lobe confusion will most strongly affect images that are close together in time and particularly those with similar local sidereal times (LSTs), because the position of sources in the side-lobes will be essentially identical. Correlated noise can be quantified by measuring the correlation coefficient for pixels⁴ between two different images. To investigate the characteristic time-scale of this correlated noise, we focus on a single data set centred on the zenith pointing direction and compare the first image to every other observation in that data set. In Fig. 2, we plot the correlation coefficient between the first image and every other image as a function of the difference between the LST for the two observations. Fig. 2 clearly shows there is significant correlated noise between two images separated by <90 s in LST. Additionally, there is a small level of correlated noise present in all the images; this is caused by a number of sources in the field that are just below the detection threshold.

The presence of correlated noise means that two observations that are close to each other in time are not statistically independent. Therefore, when the light curve of sources is processed, data points that are close in time are not statistically independent and can bias the reduced weighted χ^2 (which we define as η) that is typically used to identify variable sources causing issues with variability studies (Bell et al. 2014). This can be corrected for in the variability statistics by reducing the number of degrees-of-freedom. Correlated noise dominates for images separated by <90 s; this corresponds to three snapshot images and reduces the number of degrees-of-freedom used in computing η by a factor of the order of 3. As this paper focuses on a data set monitoring the identical field and sources typically have the same number of data points, this will lead to a systematic shift in the distributions (Bell et al. 2014) but variable sources will still be anomalous to the distribution. In conclusion, we note that there is correlated noise in this data set, which will be important for detailed intrinsic variability studies, but does not

³ To obtain the optimal Gaussian fit to the histogram, we do not fit images with rms $>3\sigma$ from the mean rms.

⁴ Again removing sources by filtering out the pixels $>4\sigma$ from the median value, as utilized in the analysis in Section 2.4.




Figure 2. A plot showing the correlation coefficient for pixels between the first zenith snapshot image and all other zenith snapshot images as a function of the change in LST between the two images. For images less than 90 s apart in LST, shown by the black dashed line, correlated noise will dominate.

hinder the identification of significantly variable sources as targeted in Section 4 of this paper.

2.5 Median image

A simple method to confirm the detection of transient sources is to compare the candidate to a deeper image of the same region. To obtain this deep image, we use all the 28 s images that passed the quality control strategies described in Section 2.3. As the images are already on a common coordinate grid, centred at an RA and Dec. of 0.0, -27.0 , we can simply average the pixel values by taking the mean or the median value of each pixel. We choose to produce the median image, as poor quality images will not bias the median image but could affect the mean image. A bright transient source that is ‘on’ for a small number of images may leave a residual source in a mean image. However, when computing the median image, we may lose some of the extended flux density as the broader regions of the point spread function (PSF) may not combine neatly. This does not significantly affect the point sources and so will have a negligible effect on our analysis.

The final median image produced, shown in Fig. 3, reaches an rms noise of ~ 8 mJy beam $^{-1}$ in the inner $\frac{1}{8}$ region, which is consistent with that expected from the combination of thermal noise and classical confusion noise (9 mJy; assuming a typical spectral slope and source counts from the VLSS at 74 MHz; see Section 2.4). Therefore, this method has significantly reduced the background noise as a number of noise components such as side-lobe confusion are averaged out due to the large number of images at different LSTs and pointing directions.

We identify 5548 unique sources in the median image (using the settings in Table 3; see Section 2.6) and we can compare this to the number of sources that are expected to be detected in this image using the VLSSr (Lane et al. 2014) catalogue number counts. This is the ideal survey for comparison as it is at a reasonably

close observing frequency and has a comparable resolution to our images. VLSSr detected 92 964 sources in a sky area of 9.38 steradians, equivalent to ~ 3 sources per square degree, at a sensitivity of ~ 450 mJy. Here, we assume that the median image has a consistent rms of 8 mJy, which is a reasonable assumption due to the observing strategy for this field and our choice to focus on the inner 12° radius, so much of the primary beam response has been averaged out. By assuming a spectral slope of -0.7 , we can scale the sensitivity of the median image at 182 MHz to the frequency of the VLSSr, giving a sensitivity of 90 mJy at 74 MHz. We can then utilize the relation

$$N \propto S^{0.9}, \quad (2)$$

where N is the number of sources expected in the survey and S is the sensitivity of the survey. We find that we expect to detect 5810 sources in the median image, therefore we detect 95 per cent of the sources expected. The number of sources detected is slightly lower than expected as we have not accounted for any residual differences in the primary beam response across the image and we have assumed a single spectral slope of -0.7 can represent the full population of sources. We also compare the number of sources detected in the median image to the 7C catalogue (McGilchrist et al. 1990) and determine that we detect an equivalent number of sources down to the completeness sensitivity of the 7C catalogue.

2.6 The LOFAR TRANSIENTS PIPELINE

To process the images, we used the LOFAR TRANSIENTS PIPELINE (TRAP, Release 2.0; Swinbank et al. 2015). In the following sections, we utilized TRAP default settings unless stated otherwise.⁵

The source finder used in TRAP has been optimized for transient searches (Spreeuw 2010). Gaussians are fitted to each detected source and the fitted parameters (size and orientation) can significantly vary for a point source between consecutive images, due to the noise properties in the surrounding region, leading to artificial variability in the flux density of the source. Variable sources are expected to be point sources, so we can assume they take the shape of the restoring beam to mitigate against this problem. This strategy leads to underestimation of the flux densities of extended sources, but this is a reasonable sacrifice for variability searches as their flux densities will remain stable. A further underestimation of the flux densities can be caused by the ionosphere making point sources slightly larger than the restoring beam size; typically this effect is negligible but can become larger on nights of high ionospheric activity (these nights are typically rejected from the data set; see Section 2.3).

However, by forcing all the sources to take the shape of the restoring beam, the source finder can fail to converge to a solution and affect the completeness of the sample, especially at low signal-to-noise ratios (SNRs) and in images with a small number of pixels across the restoring beam shape. For the images used in this analysis, we found the completeness was significantly reduced when fitting sources where the Gaussian shape parameters in the least-squares fit were fixed to take the shape of the restoring beam.

Therefore, to maximize both the completeness of the survey for faint transient sources and have reliable flux densities, we have used two different strategies to process the images.

(i) Standard: TRAP conducts a blind and unconstrained source extraction on each individual image and produces a light curve for

⁵ For further details about these capabilities, refer to TRAP documentation at <http://tkp.readthedocs.org/en/release2.0/> (TRAP Contributors, 2015).




Figure 3. The median image created using 10 122 images (excluding the bad images following the quality control in Section 2.4) of the field with a J2000 coordinate grid overlaid (in degrees), using a logarithmic colour scale. The dashed red circle shows the 12° radius source extraction region used in the analysis conducted in this paper.

Table 3. The TRAP source finder settings used in the transient search (Section 3). All other settings were the default settings.

Setting	Transient search
Detection threshold	6σ
Analysis threshold	4σ
Grid size	100 pixels
Deblend thresholds	1
Gaussian fit using restoring beam	False
Source extraction radius	12° (\sim primary beam HWHM)

every source detected. This setting was used for the transient search conducted in Section 3 with the parameters given in Table 3.

(ii) **Monitoring:** a list of positions is given as an input to TRAP, which then fits a Gaussian using a least-squares method, with the Gaussian shape being held constant at fixed values equal to that

of the restoring beam, at each position in every image to produce a reliable light curve for each source. The position is also held approximately constant during the fitting procedure, with a 10 pixel variation allowed. This setting was used for the flux density stability tests in Section 2.9 and the variability search conducted in Section 4.

2.7 Source finder performance

We tested the completeness and reliability of the TRAP source finder settings used for the transient search. The source finder used in this section is the stand-alone, command-line version of the source finder used by TRAP (known as PYSE).

We extracted sources from 250 randomly chosen images from the data set (those passing the quality control in Section 2.3), using the settings summarized in Table 3 and compared them to the sources




Figure 4. The precision (solid line) and recall (dashed line) of the source finder settings used in this paper as a function of the flux density of the sources. The precision is roughly constant at ~ 100 per cent whereas the recall drops to below 90 per cent for flux densities < 0.5 Jy.

detected in the significantly deeper median image using a source association radius equivalent to 10 pixels. We raised the detection threshold for the median image to 18σ such that a source detected at 6σ in the lowest rms image in the sample (mean rms -2σ ; see Fig. 1) would be easily detected in the median image. By raising the analysis threshold from 3σ to 4σ , we were able to reduce the number of spuriously large ellipses fitted to noise artefacts by the source finder. For evenly spaced SNR bins, we counted the number of true positive (TP) detections where we correctly found the source, the false positive (FP) detections which are spurious sources and the false negative (FN) detections which are the missed sources. For each flux density bin we calculated the precision given by

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \equiv 1 - \text{FDR}, \quad (3)$$

where the FDR is the False Detection Rate, and the recall (commonly also referred to as completeness) given by

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \quad (4)$$

The precision is quantifying the reliability of the source extractions, i.e. the likelihood of spurious sources, whereas the recall is giving the probability that a transient will be detected. Calculating the correct number of FN detections is complex as the flux densities of sources will be fluctuating within their uncertainties and hence the source may randomly fall below the detection threshold in the 28 s images. Therefore, we counted the number of FNs assuming no flux density uncertainties and the number of FNs excluding any source in the median image that would be undetected if its flux density had dropped to the 3σ lower limit.

In Fig. 4, the precision and recall are plotted as a function of the flux density of the sources. We find that the precision is consistently ~ 100 per cent, therefore we are not dominated by a large number of spurious transient sources. The recall is consistent with > 90 per cent for sources above a flux density of ~ 0.5 Jy, and is in excess of

70 per cent for flux densities above 0.21 Jy (the sensitivity used in this analysis). This is a pessimistic lower limit on the recall for new point sources as the source may fall below the detection threshold as the noise level in the image can fluctuate for reasons such as residual calibration errors.

2.8 Simulations: transient recovery

Conducting a phase and amplitude self-calibration of the data set using a model image of the field, as described in Section 2.2, leads to significantly improved image quality; the image noise is lower and artefacts surrounding bright sources are reduced. Additionally, this method ensures that the flux density scale is more consistent from image to image across the entire data set. However, the method has the implicit assumption that the resultant image will be identical to the input image (i.e. the total flux density is constant and is distributed in the same places). If a transient occurs during the observation, or a known source is significantly variable, the output data set is contrary to this assumption. The expectation is that the calibration step will find the solution which best fits the model and, as long as the majority of the flux density in the model is well known, the transient and variable behaviour will be recovered. The anticipated limit to this method is when the transient flux density significantly exceeds the flux density of the brightest source in the field and then calibration will fail.

In order to confirm that transients within the field are recovered, we simulate a range of transient flux densities within the image. As a transient is a source which is not in the model image used to calibrate the data set, we can simply simulate a transient by removing a source from the model image (by setting the relevant pixels equal to zero). Using the strategy outlined in Section 2.2, we then apply a phase and amplitude self-calibration, using this edited model image, on a previously un-calibrated data set (observation ID 1061661200). We create 50 new models, by removing 50 sources with a range of flux densities pseudo-logarithmically distributed in the range ~ 0.3 – 20 Jy (corresponding to roughly the faintest and the brightest sources in the field) and at random positions within the image, as shown in Fig. 3. Following imaging of the calibrated data sets, we measure the recovered flux density of the simulated transient using the source finding settings used in the transient search (as described in Section 2.6 and Table 3).

We confirm that all the simulated transient sources are successfully recovered via this calibration and the source extraction method. In Fig. 5, we plot the ratio between the observed and original source flux density as a function of the flux density of the source. The observed flux density is typically within the uncertainties of the flux density measurements; however it is typically lower than the original flux density by a few per cent and is proportional to the flux density of the transient source. The only exceptions are for sources where the source finder fits the source as an extended source and may be affected by local structured noise in the image. The reduction in flux density is unsurprising as the calibration assumes that the majority of the flux density in the field is contained in known sources. Therefore, as the brightness of the unknown source becomes a significant contribution of the total flux density in the sky, the calibration is underestimating the total flux density in the field and leads to the reduction in the flux density of the transient source.

This calibration strategy is expected to fail when the flux density of the transient source dominates the total flux density in the sky model, leading to an extremely poor quality image with few recognisable sources. Therefore, images that significantly fail quality




Figure 5. The ratio between the observed flux density of the simulated transient to the original flux density when the source was included in the sky model as a function of the original flux density. The black dashed line marks where the two flux densities are equal giving a ratio of 1.

control are a possible signature of a bright transient within the field (although these are likely to have a RFI origin). The transient surface densities quoted in this paper are for transients that are not significantly brighter than the brightest source in the field, i.e. $\gg 20$ Jy. In the event of failed calibrations, recalibration via the transfer of calibration solutions from a calibrator can be used to search for bright transients. As the occurrence of this is expected to be ex-

tremely rare, the self calibration strategy used in this paper will give the optimal image quality required for fainter transient sources.

2.9 Flux density stability between images

The variability analysis conducted in Section 4 requires a well-understood flux density calibration between all the images in the data set. In this section, we use all sources in the field with average flux densities in excess of 0.5 Jy, as measured by TRAP using the monitoring strategy (as described in Sections 2.6 and 4). As stated in Section 2.2, there are residual primary beam effects in the observed flux densities, when moving between different Az-El pointing directions. In this section and Section 4, we will only consider source variability between detections at the same Az-El pointing directions.

For every source detection, we measure the ratio between the observed flux density and the average flux density of the source and the observed flux density uncertainty. We plot the histograms of these parameters from the zenith Az-El pointing directions in Fig. 6. For each distribution we calculate the mean and 1σ standard deviation. We find that the extracted source flux densities are typically in excellent agreement with their average flux density with a deviation equivalent to ~ 10 per cent of the average flux density at the 1σ level. The typical flux density uncertainties are 76 ± 24 mJy, or ~ 2.5 times the typical image rms in the inner part of the image, although this is strongly dependent on radius from the pointing centre due to the decreased sensitivity of the primary beam (the image rms in the source extraction region varies from a minimum of 23 mJy beam^{-1} to a maximum of $153 \text{ mJy beam}^{-1}$) and varies significantly for extended sources.

We find no dependence of the source flux densities on radial position within the image or the image rms throughout the entire data set. Therefore, the images from specific Az-El pointing directions have a reasonably consistent flux density scale, enabling initial variability searches for the identification of significantly variable sources and is sufficient for this pilot survey.




Figure 6. Here, we show the histograms for the zenith observations showing the typical distributions of the $\frac{\text{Observed flux}}{\text{Average flux}}$ (left) and the flux density uncertainties (right). The solid line marks the mean of each distribution with the scattered lines giving the 1σ standard deviation.

To confirm the absolute flux density scale, we compare the observed flux densities to the following catalogues.

(i) 180 MHz Murchison Commissioning Survey (MWACS; Hurley-Walker et al. 2014) that fully overlaps with this field. This survey was observed approximately 1 yr prior to the start of the observations used in this analysis.

(ii) Sydney University Molonglo Sky Survey (SUMSS; Mauch et al. 2003) at 843 MHz, covering declinations $< -30^\circ$.

(iii) VLSSr (Cohen et al. 2007; Lane et al. 2014) at 74 MHz, covering declinations $\gtrsim -25^\circ$ with sparse coverage of the field in the range -35 to -25° .

(iv) The NRAO VLA Sky Survey (NVSS; Condon et al. 1998) at 1.4 GHz, covering declinations $> -40^\circ$.

We cross matched all the sources from each of the Az-El pointing directions separately and compare the ratio between the average source flux density and the catalogue flux density (extrapolated to 182 MHz assuming a typical spectral slope of -0.7). We find the observed flux densities are consistent with the extrapolated catalogue flux densities with an average ratio of 1.1 ± 0.2 (MWACS), 1.5 ± 0.5 (SUMSS), 1.4 ± 0.4 (VLSSr) and 1.6 ± 0.8 (NVSS). These are all consistent within the 1σ uncertainties. However, we note that these ratios are all larger than unity, which is most likely a systematic offset caused by uncertainties in the primary beam model and the difference in elevation between the observation of 3C444 (used to calibrate the flux scale) and the observations.

3 TRANSIENT ANALYSIS

This data set constitutes one of the largest sky areas, with a field of view of $\Omega = 452 \text{ deg}^2$, surveyed with excellent sensitivities on a very wide range of time-scales (30 s \sim 1 yr) at low radio frequencies. The field is well off the Galactic plane and hence likely to be dominated by extragalactic sources. This enables us to place tight constraints on the surface density of extragalactic low-frequency radio transients.

We processed all the images with TRAP utilizing the transient search source finder parameters given in Table 3. To aid with source association, we use a systematic position uncertainty of 270 arcsec, corresponding to 5 pixels, which is added in quadrature to the position uncertainties for each source. As we are only searching for transient sources in this section, we can assume that they will not be detected in the deep median image (see Section 2.5 for more details) so we forced TRAP to process the median image first to create a deep source catalogue for source association. Therefore, any sources identified as new sources from the subsequent, shallower 30 s images are all transient candidates. As we focused on the detection of new sources and not their variability parameters in this section, we were able to subdivide the data set by time into smaller, more manageable, chunks for processing.

The source finder settings used will occasionally lead to large elongated Gaussians fitted to noise structures in the image. These fitting errors can be easily mitigated against as we expect all point transient sources to be roughly circular in these images (the restoring beam is close to circular), so we rejected all fitted sources with Gaussian shapes with $\frac{\text{major axis}}{\text{minor axis}} > 2$.

Although the self-calibration strategy has reduced the number of artefacts around bright sources, there remained an over density of new source detections around the bright sources. Visual inspection of a randomly chosen sample confirmed these are caused by side-lobes of the bright sources. Due to this, we rejected all candidates that occurred within 0.4 of a source with flux densities in excess of 8 Jy, leading to a reduction in surveyed area of $\sim 3 \text{ deg}^2$.

A further over-density of transient candidates occurred on the source extraction boundary at a radius of 12° . This is caused by the source finder not modelling the rms noise in the region beyond the source extraction region. We visually inspected a randomly chosen sample of these candidates and, by increasing the source finder radius to better model the rms noise in this region, showed that these were not significant sources. Therefore, we rejected all candidates occurring within 0.2 of the edge of the source extraction region, leading to a reduction in surveyed area of $\sim 15 \text{ deg}^2$.

Following this, we overlaid the transient candidates on the median image and rejected all candidates that had a counterpart in the median image.⁶ We identified three candidates requiring further analysis and visually inspected the detection image. One of the candidates was rejected following visual inspection, as it was consistent with an artefact that was a deconvolution error of a source in close proximity to a $\sim 15 \text{ Jy}$ source.

We developed a number of further tests using one of the most convincing of the two remaining candidates (with duration $\sim 30 \text{ s}$, detected at 6.45σ).

(i) The images were processed using different source finder settings and a different source finder, AEGEAN (Hancock et al. 2012), to confirm the detection significance.

(ii) We re-imaged the region with a range of different imaging parameters, such as changing pixel scale, weighting, image phase centre, and UV range.

(iii) We produced new images on a range of additional time-scales (2 min, 10 s and 4 s).

(iv) The transient candidate remained detected following these tests, so we also processed the observation using the de-dispersion pipeline in Miriad developed by Tingay et al. (2015), resulting in a faint detection of the source in the Miriad images but a non-detection of dispersed signals.

(v) We conducted an image subtraction using an image with the identical LST from the previous night. The transient candidate was not significantly detected in the subtracted image, suggesting the source is related to a correlated noise artefact (see Section 2.4).

(vi) A median image was created using all the images in the data set with the identical LST and confirmed the presence of a noise peak at this position for this LST. The transient candidate had a flux density in excess of the noise feature in the median image; however, given the noise feature and the low SNR, this is an unconvincing source.

The remaining transient candidate was also consistent with a noise peak at a specific LST and, hence, is unconvincing.

The processing strategy outlined in this section may fail to detect bright long-duration (\sim months) transient sources on year time-scales, as they could leave a residual source in the median image that could be detected by the source finder (e.g. a 300 mJy transient source with a duration of 2 months could potentially lead to a $\sim 50 \text{ mJy}$ source, i.e. a $\sim 6\sigma$ detection, in the median image). However, the targeted population of transient sources will not have a counterpart in the existing radio catalogues covering this region. Therefore, we cross-match all the sources detected in the median image with the NVSS catalogue (the only catalogue covering the entire region). The peak flux densities for sources that were not detected in NVSS were extrapolated to 1.4 GHz to determine if we expected to detect them. 18 faint, uncatalogued radio sources were

⁶ We note that these may be variable sources and will investigate further in future analysis.

Table 4. The number of FRBs expected to be detected in the 28 s images as a function of the spectral index of the FRB, assuming the observed rates from Thornton et al. (2013) and 75 per cent lower rates consistent with the analysis by Keane & Petroff (2015).

Spectral index	Thornton et al. (2013)		Keane & Petroff (2015)	
	Number predicted	Null detection probability	Number predicted	Null detection probability
-2	27.1	1.8×10^{-12}	6.8	1.2×10^{-3}
-1	4.8	8.2×10^{-3}	1.2	3.0×10^{-1}
0	0.7	4.8×10^{-1}	0.2	8.3×10^{-1}

identified in the median image; however they are not expected to be detectable in previous surveys of this region. We used the monitoring capability in TRAP (see Sections 2.6 and 4 for more details) to obtain a light curve at the position of these faint sources to confirm that they were not residual sources from bright transients. Using this strategy, we identified no long duration bright transient sources and analysis is ongoing to determine if these sources are variable.

3.1 FRB limits

This data set, with an snapshot time-scale of 28 s, has a very good cadence for the detection of FRBs in the image plane (without de-dispersion) at 182 MHz. By scaling from known rates for FRBs detected at 1.4 GHz, we can predict the number of FRBs we expect to observe in the 10 122 images included in this analysis.

Using the methods presented in Trott et al. (2013), we predicted the number of FRBs that we expect to detect in the 28 s images, with no image plane de-dispersion, as a function of a range of spectral slopes assuming that FRBs are a standard candle (given in Table 4). In a 28 s image we are able to probe dispersion measure (DM) values up to 700 pc cm^{-3} ; within this component we estimate a Galactic DM of $\sim 10 \text{ pc cm}^{-3}$ (field is well off the Galactic plane, so the Galactic component is low.) and a host galaxy component of 100 pc cm^{-3} . The rates predicted are normalized to the whole sky rate of FRBs observed by Thornton et al. (2013). For the Thornton et al. (2013) rates, we rule out spectral indices ≤ -1 at >95 per cent confidence, while for the 75 per cent lower rates reported by Keane & Petroff (2015) we are able to rule out spectral indices ≤ -2 at >95 per cent confidence.

In this data set, we expect all FRBs to only be detected in one image, as their flux densities will likely be too low to be detected if spread over multiple images. To determine if there are any FRB candidates in this data set, we queried all of the new sources detected by TRAP to identify all sources which are only detected in one image and that were not found in the first, deep median image. This resulted in a small number of candidates that were visually inspected using the median image, the detection image and other 30 s images

close by in time. The majority of the remaining sources had visible counterparts in the median image, some of which are related to failed source associations and others are candidate variable sources warranting future investigation but do not meet our requirements for candidate FRBs.

Our non-detection of any FRB candidates can place tight constraints on the rates of FRBs at low radio frequencies. Following the method adopted by Trott et al. (2013), where we are in the regime that the snapshot duration is greater than the signal width, we can estimate the minimum FRB flux densities $S_{\text{FRB,min}}$ that we are sensitive to using:

$$S_{\text{FRB,min}} = S_{\text{min,28s}} \left(\frac{\Delta t}{w} \right), \quad (5)$$

where $S_{\text{min,28s}} = 0.285 \text{ Jy}$ is the sensitivity in one snapshot image, $\Delta t = 28 \text{ s}$ is the snapshot integration time and w is the intrinsic width of the FRB. For consistency with Tingay et al. (2015), we assume that the intrinsic width is 1 ms. Therefore, this experiment is sensitive to FRBs with flux densities in excess of 7980 Jy. We can estimate an upper limit on the rate of FRBs per sky per day observable using the standard transient rate for the 28 s snapshot rate, $3 \times 10^{-6} \text{ deg}^{-2}$ (as calculated in Section 3.2). To observe the equivalent of 1 d would require 3085 snapshots of 28 s integration time, therefore our whole sky FRB rate (ρ_{FRB}) is

$$\rho_{\text{FRB}}(S_{\text{FRB}} > 7980 \text{ Jy}) < 82/\text{sky/d}. \quad (6)$$

In Table 5 and Fig. 7 (adapted from Coenen et al. 2014, assuming a flat spectral slope), we show this FRB rate in comparison to previous surveys at a range of frequencies, assuming FRBs have a flat spectrum. Here, we assume that FRBs are a standard candle and the observed FRB population have been shown to be consistent with this (Dolag et al. 2015). Assuming a flat spectrum and a cosmological population and using

$$N \propto S^{-\frac{3}{2}}, \quad (7)$$

where N is the number of transients and S is the flux density of the transient, we can determine that the rates we obtain are broadly

Table 5. The current rate constraints for FRBs. * The peak flux density sensitivity given by Rane et al. (2016) is converted to an observed flux density for a pulse that is on for $\sim 10 \text{ ms}$ (the average duration observed by Keane & Petroff 2015).

Survey	Frequency (MHz)	Sensitivity (Jy)	Rate (day/sky)	DM range (pc cm^{-3})	Method	Citation
MWA	182	7980	<82	<700	28 s images	This Work
MWA	156	700	<700	170–675	2 s de-dispersed images	Tingay et al. (2015)
ARTEMIS	145	62	<29	<320	High time resolution	Karastergiou et al. (2015)
Parkes (Bayesian, All)	1400	1*	$3.3_{-2.5}^{+5} \times 10^3$	~ 2000	High time resolution	Rane et al. (2016)
Parkes (9 FRBs)	1400	2	~ 2500	375–1103	High time resolution	Keane & Petroff (2015)
LOFAR	142	107	<150	2–3000	High time resolution	Coenen et al. (2014)
VLA	1400	1.2	$<7 \times 10^4$	0–3000	High time resolution	Law et al. (2015)
Parkes (4 FRBs)	1400	3	$1_{-0.5}^{+0.6} \times 10^4$	553–1103	High time resolution	Thornton et al. (2013)
Parkes (1 FRB)	1400	30	~ 400	375	High time resolution	Lorimer et al. (2007)




Figure 7. This plot, based on plot from Coenen et al. (2014), shows the FRB rate per sky per day as determined from the surveys given in Table 5, assuming a flat spectral slope. This plot assumes that the only limiting factor is the survey sensitivity. Symbols are as in the legend and arrows denote upper limits. Assuming a cosmological population, where the only constraining factor is the sensitivity of the survey (although note this is not a valid assumption for FRBs, see the text for further details), the data points are expected to be consistent with a straight line of slope -1.5 . We plot a solid black line representing this and normalized to Thornton et al. (2013) with the uncertainties bounded by the dashed black lines. For reference, we also plot this for the lower rate given by Keane & Petroff (2015) using a blue dotted line. Finally the red dash-dotted lines represent the significantly steeper and lower rate model proposed by Macquart & Johnston (2015).

consistent with the rate obtained by Tingay et al. (2015). However, we note that our rate limit is higher than that expected when extrapolating the Thornton et al. (2013) rate to our sensitivity and is unconstraining for flat spectral slopes. Additionally, a recent calculation by Keane & Petroff (2015) has shown that the observed FRB rate may be four times lower than that determined by Thornton et al. (2013). Recently, Macquart & Johnston (2015) have postulated that the rate of FRBs does not follow that of a standard cosmological population, explaining the lack of FRBs at low Galactic latitudes, and instead is given by

$$N \propto S^{-\frac{7}{2}}. \quad (8)$$

A $S^{-\frac{7}{2}}$ distribution is no longer considering a cosmological population or a standard candle, i.e. this would imply that either the population or the luminosity is strongly dependent upon the redshift. Alternatively, there may be further selection effects that have not been accounted for. Additionally, they show that the whole sky rates will be a factor of 3 lower than the observed rates, therefore normalizing by a third of the Thornton et al. (2013) rate, we plot this constraint in Fig. 7 and our upper limits are consistent with this underlying population. We note that the rates determined by Lorimer et al. (2007) are significantly higher than this model, even when the expected uncertainties on this rate are taken into account.

Fig. 7 assumes that only the survey sensitivity is required to account for the volume probed by each of the surveys. However, for FRBs this is not strictly the case as it also strongly depends on the DM range searched over. The DM search range constrains the volume that can be searched for FRBs, irrespective of their luminosities. In this analysis, the 28 s images produced can be used to probe DMs up to 700 pc cm^{-3} , whereas other surveys can exceed DMs of 1000 pc cm^{-3} or can be much lower than this. The DM can be converted to a redshift (z) using the relationship (e.g. Ioka 2003; Inoue 2004; Lorimer et al. 2007; Karastergiou et al. 2015)⁷:

$$\text{DM} \approx 1200 z \text{ cm}^{-3} \text{ pc}. \quad (9)$$

However, before conversion to a redshift, we want to remove the Galactic component of the DM and also a contribution from the host galaxy of the FRB. The host galaxy contribution is unknown and, for consistency with Karastergiou et al. (2015), we assume it is 100 pc cm^{-3} . Using the model produced by Cordes & Lazio (2002)⁸ we find a DM of 29 pc cm^{-3} ; however we note that this model is based on the extrapolation of DM measurements from pulsars and only one of the pulsars used is within our field.⁹ This is a small amount relative to the DM searched, unsurprising due to the Galactic latitude of the target field (centred on the Galactic co-ordinates l, b : $30.6, -78.5$), and is therefore negligible in our analysis. Therefore, the DM component that is expected to be from the intergalactic medium (IGM) is $\sim 600 \text{ pc cm}^{-3}$, corresponding to a redshift of ~ 0.5 .

In conclusion, we find a rate of <82 FRBs per sky per day that are brighter than 7980 Jy at a frequency of 182 MHz out to a maximum redshift of ~ 0.5 . We show our non-detection is consistent with the lower rates calculated by Keane & Petroff (2015), if FRBs have a spectral slope > -1 , or the non-standard cosmological population suggested by Macquart & Johnston (2015).

3.2 Transient rates

In addition to the non-detection of FRB candidates (as described in Section 3.1), no convincing transient sources were detected in this data set. We can calculate the standard transient surface density limit using Poisson statistics via:

$$P = \exp^{-\rho(N-1)\Omega}, \quad (10)$$

where $(N-1)\Omega$ is the total area surveyed by N snapshots of a field each with an area of Ω , ρ is the surface density limit and P is the confidence interval. Following Bell et al. (2014), we utilize $P = 0.05$ to give a 95 per cent confidence limit.

The sensitivity to transients depends upon the location within the image, such as the increase in rms noise with radius due to the decreasing sensitivity of the primary beam (e.g. Croft, Bower & Whysong 2013; Bell et al. 2014), and on the specific images used as some are of higher quality (e.g. Carbone et al. 2014a). Here, we can characterize the area surveyed for a given sensitivity by assuming that the sensitivity as function of the position (radius, r) within the primary beam, Sensitivity (r), can be approximated as a

⁷ See also Dolag et al. (2015) for an alternative method to constrain the redshift.

⁸ Using the web interface here: <http://www.nrl.navy.mil/rsd/RORF/ne2001/>

⁹ Gaensler et al. (2008) find a DM of 25 pc cm^{-3} for similar Galactic latitudes but in different directions to this field.

Table 6. The low-frequency radio surveys undertaken to date (<1 GHz) and the current surface density constraints.

Survey	Frequency (MHz)	Sensitivity (Jy)	Time-scale	Surface density (minimum time-scale) (deg ⁻²)
This Work	182	0.285	28 s	<6.4 × 10 ⁻⁷
This Work	182	0.285	5 min	<6.6 × 10 ⁻⁶
This Work	182	0.285	10 min	<1.1 × 10 ⁻⁵
This Work	182	0.285	1 h	<2.5 × 10 ⁻⁵
This Work	182	0.285	2 h	<9.5 × 10 ⁻⁵
This Work	182	0.285	1 d	<1.2 × 10 ⁻⁴
This Work	182	0.285	3 d	<2.4 × 10 ⁻⁴
This Work	182	0.285	10 d	<3.9 × 10 ⁻⁴
This Work	182	0.285	30 d	<9.5 × 10 ⁻⁴
This Work	182	0.285	90 d	<3.3 × 10 ⁻³
This Work	182	0.285	1 yr	<6.6 × 10 ⁻³
Stewart et al. (2016)	60	36.1	30 s–4 months	<4.1 × 10 ⁻⁷
Stewart et al. (2016)	60	7.9	4 min–4 months	1.4 × 10 ⁻⁵
Obenberger et al. (2015)	38	1440	5 s	<1.9 × 10 ⁻¹¹
Bell et al. (2014)	154	5.5	minutes–1 yr	<7.5 × 10 ⁻⁵
Carbone et al. (2014a)	150	0.5	minutes–months	<1 × 10 ⁻⁴
Cendes et al. (2014b)	149	0.5	minutes–months	<2.2 × 10 ⁻²
Jaeger et al. (2012)	325	0.0021	1 d–3 months	1.2 × 10 ⁻¹
Bannister et al. (2011)	843	0.014	days–years	1.3 × 10 ⁻²
Lazio et al. (2010)	74	2500	300 s	<9.5 × 10 ⁻⁸
Hyman et al. (2009)	235–330	0.03	days–months	3.4 × 10 ⁻²

one-dimensional Gaussian distribution given by:

$$\text{Sensitivity}(r) = 6 \text{ rms}_c \exp\left(\frac{r^2}{2 \text{ HWHM}^2}\right) \text{ mJy}, \quad (11)$$

where rms_c is the rms in the central 100 pixels of the image and the factor of 6 is because we extract sources which are 6σ above the detection threshold. By measuring the rms in the central 100 pixels of all the images, we find an average $\text{rms}_c = 27 \text{ mJy beam}^{-1}$. For the full source extraction region, with a radius of 12° and an area of 452 deg^2 , we have a sensitivity of $\sim 285 \text{ mJy}$. We note this is still an approximation to the sensitivity as the primary beam response may not take a simple Gaussian shape and does not account for other causes of variation in rms throughout the field. However, due to the large number of images utilized in this analysis, a detailed characterization of the sensitivities and surface densities (such as that conducted by Bell et al. 2014), for all the surface densities on each of the different time-scales probed, would take a disproportionate amount of compute time and not significantly affect the results presented.

In this data set, we extract sources within a radius of 12° which are 6σ above the detection threshold corresponding to a detection limit of $\leq 0.285 \text{ Jy}$ (see equation 11), with a field of view of $\Omega = 452 \text{ deg}^2$. On the shortest time-scale, 28 s, we used $N = 10\,122$ images which corresponds to a surface density of $< 6.4 \times 10^{-7} \text{ deg}^{-2}$. We compare our observations to the other transient surveys conducted at low radio frequencies to date in Table 6.

Using the method presented in Carbone et al. (2014a), where only statistically independent images are used to calculate the rates for a specific transient time-scale, the transient surface density limit is calculated for a range of unique time-scales probed by this data set. In Fig. 8, we plot the surface densities obtained for each of the time-scales in comparison to the existing surveys at low frequencies. This highlights the orders of magnitude decrease in the transients surface density limits on a range of time-scales that this survey provides. The only surveys with surface densities on faster time-scales or lower surface density constraints are for flux density sensitivities

that are orders of magnitude higher than this survey. We note that we do not have sensitivity on time-scales between $\sim 2 \text{ h}$ and $\sim 1 \text{ d}$ due to the EoR observing strategy. Also, as the observations occur when the field is optimally located in the observable sky, we have no sensitivity for transients between time-scales of $\sim 3 \text{ months}$ to $\sim 1 \text{ yr}$. In addition to these gaps in sensitivity, we note that we will also have some sensitivity to transients with time-scales $< 28 \text{ s}$ where very short lived coherent transients from Galactic sources, such as pulsars, may be anticipated (the sensitivity to these transients is dependent upon the flux density and duration of the transient; see equation 5). Finally, we have not probed time-scales $\gtrsim 1 \text{ yr}$ where we expect to observe the longer duration synchrotron sources at this observing frequency. These parameter spaces remain to be explored in the future.

In Fig. 8 we plot the standard surface densities versus the sensitivity (flux density detection limit) for the sample of low-frequency radio transient surveys; the published surface densities from higher radio frequencies are plotted for reference. This survey is typically greater than an order of magnitude more sensitive than previous studied or, conversely, orders of magnitude more constraining for a given sensitivity. To further constrain these rates, future surveys will require significantly increased sensitivity (for instance via SKA-low) or have at least an order of magnitude increase in surveyed area (e.g. via commensal observations).

To date, few of the transient surveys <1 GHz have detected transient sources, with most detections at an order of magnitude higher sensitivity (e.g. Hyman et al. 2009; Bannister et al. 2011; Jaeger et al. 2012) on the days–months time-scales making them unlikely to be detected in this survey. Stewart et al. (2016) have identified a bright transient source at 60 MHz on the minutes time-scale using LOFAR (see also Fender et al. 2015). Assuming a flat spectrum, with two orders of magnitude improvement in sensitivity, this survey would be naively predicted to detect hundreds of these transients. As no transients were detected on this time-scale, we investigate the implications for the Stewart et al. (2016) transient source. Assuming these sources are a standard candle distributed




Figure 8. The observed transient surface densities plotted as a function of the time-scales probed (left) and the limiting flux density (right). The different surveys are given in the legend and down arrows denote upper limits. The red solid line represents the typical rate constraints taking into account the differing sensitivity across the images due to the primary beam response (assuming a Gaussian primary beam), starting at a radius of 2° and ending at the source extraction radius of 12° .

Table 7. The number of transients expected to be detected on the minutes time-scale, as a function of the spectral index of the transient, scaled from the transient detection by Stewart et al. (2016).

Spectral index	Number predicted	Null detection probability
-3	1.0	3.7×10^{-1}
-2	5.2	5.5×10^{-3}
-1	27.6	1.0×10^{-12}
0	146	0

isotropically, we can scale the sensitivity of this survey to that of Stewart et al. (2016) via equation (7). In Table 7, we predict, for a range of spectral indices, the number of transients expected in this survey and calculate the probability of null detection. Assuming the surface density is correct, we can rule out spectral indices ≥ -2 at >95 per cent confidence. We conclude that the most likely scenarios are: the transient surface density is much lower than observed by Stewart et al. (2016) and/or the spectral index of these transients is very steep (< -2). A steep spectral index may be consistent with this emission being from coherently emitting sources such as pulsars. This is consistent with the very steep spectrum < -4 proposed by Stewart et al. (2016).

Finally, we note that this field is close to the Galactic poles and does not constrain the Galactic population of transient sources. This will be best determined by up-coming whole sky transient surveys.

4 VARIABILITY SEARCH

4.1 Method

For this section, we put all the images through TRAP utilizing the monitoring strategy described in Section 2.6 and monitor the vari-

ability of sources detected in the median image with flux densities in excess of 0.5 Jy. As stated previously, there are residual primary beam issues when comparing images at different Az-El pointing directions so we process each pointing direction separately. Variable candidates are then compared between each of the pointing directions.

TRAP measures two key variability parameters for every unique source in the data set. The first parameter is the reduced weighted χ^2 at a given observing frequency, η , given by:

$$\eta = \frac{N}{N-1} \left(\frac{\overline{w I^2}}{\overline{w} \overline{I^2}} - \frac{\overline{w I}^2}{\overline{w}^2} \right) \quad (12)$$

where N is the number of data points, I_i is the flux density of a data point, $\overline{w} = \frac{1}{N} \sum_{i=0}^N w_i \equiv \sum_{i=0}^N \frac{1}{\sigma_i^2}$, σ_i is the error on the i th flux density measurement and over bars represent the mean values (the full derivation of this from the standard reduced weighted χ^2 is given in Swinbank et al. 2015). The second parameter is the coefficient of variation (also known as the modulation index) at each observing frequency, V , given by

$$V = \frac{s}{\overline{I}}, \quad (13)$$

where s is the standard deviation of the observed flux densities. These parameters are measured for each time-step that the source is observed. In the following analysis, we focus on the variability parameters for each unique source from the final time-step.

4.2 Results

4.2.1 Blindly detected variability

In Fig. 9, we show these variability parameters from the end of each TRAP run for each of the five unique azimuth-elevation pointing




Figure 9. This plot shows the two variability parameters: the reduced weighted χ^2 , η , and the coefficient of variation, V (given by equations 12 and 13, respectively) for all the sources tracked in the variability search. The colour scheme corresponds to the five separate unique azimuth-elevation phase centres, as given in Table 2 and the legend. In the top and right panels, we plot histograms of the two variability parameters and fit them with a Gaussian distribution (dotted line). We utilize a 2σ threshold on both parameters, represented by the black dashed line. Significantly variable sources would reside in the top right corner bounded by the two thresholds.

directions. As the time-scales probed by each of the different pointing directions is roughly the same, it is expected that the typical source parameters for the different pointing directions will be in good agreement and this is clearly the case in Fig. 9. The V distribution is well fitted with a Gaussian distribution, with a typical value of $V \sim 0.1$, consistent with the typical flux density uncertainties measured in Section 2.8. We note that the absolute η values should be a factor of ~ 3 lower due to the correlated noise observed in these images (see Section 2.4); this does not affect the analysis in this section as we are only considering sources that are anomalous to the distribution. This factor will need to be considered when quantifying low-level variability in future analysis. The η distribution is clearly right-skewed with an excess of sources at higher values, suggesting that there may be variability in some of the source light curves. However, as the V parameters of these sources are comparable to the rest of the population, this is unlikely (variable sources have anomalously high values for both variability parameters and are expected to reside in the top right corner of this plot.). By visual inspection, we note that many sources show variation on specific nights, pointing to a possible ionospheric origin or residual calibration issues. Further analysis is ongoing.

Using additional source parameters (particularly the maximum flux density that a source attains and the ratio between that maximum flux density and the average flux density of the source) can aid in understanding the population of sources and can be used to more clearly separate the variable sources from a stable population. In Fig. 10, we plot these four parameters for each of the

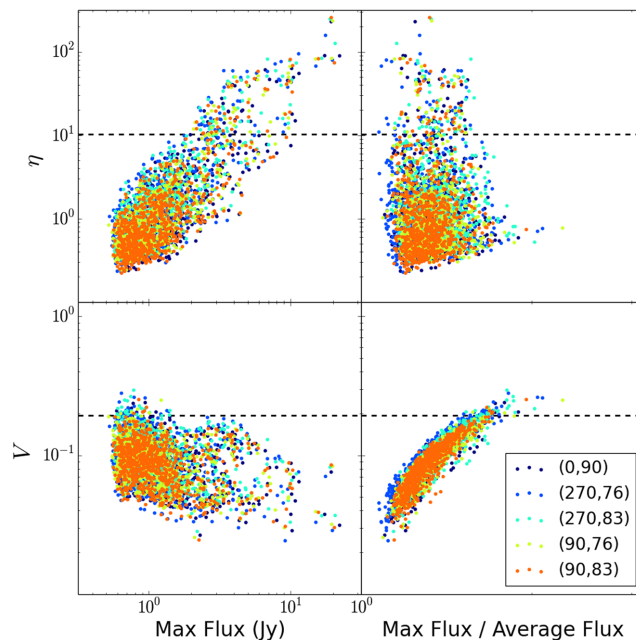


Figure 10. Here we show the two variability parameters against the maximum flux density that each source attains and the ratio between this maximum flux density and the average flux density of the source. The colour scheme is as in Fig. 9 and the dashed lines are the 2σ thresholds used to identify variable sources.

different pointing directions. We see a clear correlation between the maximum flux density and η as expected (this is caused by the measurement accuracy of flux densities for bright sources, which does not take into account systematic uncertainties.). Additionally, we observe a negative trend between the maximum flux densities and V ; however there is a clear diversion from this trend at flux densities >2 Jy. This diversion is caused by a large number of images, coincident in time, having a systematic flux density scale offset (of the order of 10 per cent) from the rest of the images. This is likely caused by uncertainties in the primary beam model and is expected to be resolved in future work when we address the issues between different pointing directions.

To avoid the region where the systematic uncertainties are dominating, we utilize a 2σ threshold on both the η and V parameters to identify variable sources. This corresponds to variable sources requiring $\eta \geq 10.6$ and $V \geq 0.20$, equivalent to a flux density variation in excess of 20 per cent. These thresholds will only identify sources which are significantly more variable than the typical population and do not address any intrinsic low-level variability, such as the phenomenon of ‘low frequency variability’ likely caused by refractive interstellar scintillation (although this is not expected to occur in our data set as it typically has time-scales of $\gtrsim 1$ yr and is more prevalent at lower Galactic latitudes; e.g. Mitchell et al. 1994). From the skewed η distribution and visual inspection of a large number of light curves, we note that there is low-level variability on specific observing nights which may correspond to ionospheric activity. No significantly variable sources were identified via this method.

4.2.2 Known variables within the field

Two sources within this field have observed variability, with time-scales of 2 s at ~ 150 MHz, which has been interpreted as Interplanetary Scintillation (IPS; Kaplan et al. 2015). Using the zenith observations, we identify these two sources within our data set and determine their variability parameters. PKS 2322–275 has $V = 0.047$ and $\eta = 6.2$ and PKS 2318–195 has $V = 0.050$ and $\eta = 2.0$. Both of these sources are well below the variability thresholds and their parameters are highly typical for the source population. Kaplan et al. (2015) show that these sources have a typical variability time-scale which is significantly shorter than the 30 s integration time-scale used in this analysis. Therefore, any IPS event would be statistically averaged to the mean value in these images.

Additionally, there is one pulsar within the field, PSR 2327–20. This pulsar has a low DM and, hence, may undergo diffractive and refractive scintillation. In our median image, we note that this source is very faint, with a flux density of ~ 0.07 Jy, and is unlikely to be detected in our images. We monitored the position of this pulsar to see if this pulsar scintillates above the detection threshold. PSR 2327–20 is not detected in any of the snapshot images.

4.3 Future work

Although this analysis has not identified any significant variability, we note that the variability analysis needs a significant amount of further work to be able to identify variability of sources corresponding to $\lesssim 20$ per cent of their flux densities. For future analysis:

(i) we intend to resolve remaining systematic primary beam uncertainties within the images. This will enable all the pointing directions to be processed at once, giving a much larger data set for characterizing the sources. Additionally, it will resolve the

deviation in V at flux densities in excess of 2 Jy, which will lead to an increase in sensitivity.

(ii) The variability parameters for sources can change significantly as the number of data points in the light curve increases. For instance, if a source emits a single flare at early times this variability may not be apparent using the variability parameters from the final time-step in the data set. This means that we may be missing interesting variability on short time-scales due to processing large numbers of images at once. TRAP records these variability parameters as a function of snapshot and, in future, we aim to develop methods to study the variability of sources as a function of time.

5 CONCLUSIONS

From our analysis of $\sim 10\,000$ images, we note that the EoR0 field is remarkably stable at 182 MHz. There are no convincing transient candidates and all sources have flux density variations of $\lesssim 20$ per cent. In future work, we will target remaining systematic flux density uncertainties to enable us explore low-level variation within the field.

The transient surface densities obtained are more constraining than previous surveys by orders of magnitude in time-scale, sensitivity and snapshot rates; although we note that this field is not sensitive to a Galactic population of transient sources due to being well off the Galactic plane (Galactic latitudes $\lesssim -65^\circ$). Despite expecting to observe transients comparable to the source observed by Stewart et al. (2016), we instead place a constraining limit on their surface densities and/or spectral indices. On the shortest time-scale, predictions scaled from the observed populations suggested that this survey would identify a small number of FRBs. Again, there are no detections, which are consistent with suggestions of lower rates and flat spectral indices.

To further pursue these elusive transients at low radio frequencies, we need to conduct similar surveys at a range of frequencies, whilst also increasing the sensitivities and surveyed area by an order of magnitude or more. Finally, there are a range of time-scales that this survey does not explore, most notably the very short and > 1 yr where a range of transient sources are anticipated.

ACKNOWLEDGEMENTS

This scientific work makes use of the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the MWA comes from the U.S. National Science Foundation (grants AST-0457585, PHY-0835713, CAREER-0847753, and AST-0908884), the Australian Research Council (LIEF grants LE0775621 and LE0882938), the U.S. Air Force Office of Scientific Research (grant FA9550-0510247), and the Centre for All-sky Astrophysics (an Australian Research Council Centre of Excellence funded by grant CE110001020). Support is also provided by the Smithsonian Astrophysical Observatory, the MIT School of Science, the Raman Research Institute, the Australian National University, and the Victoria University of Wellington (via grant MED-E1799 from the New Zealand Ministry of Economic Development and an IBM Shared University Research Grant). The Australian Federal government provides additional support via the Commonwealth Scientific and Industrial Research Organization (CSIRO), National Collaborative Research Infrastructure Strategy, Education

Investment Fund, and the Australia India Strategic Research Fund, and Astronomy Australia Limited, under contract to Curtin University. We acknowledge the iVEC Petabyte Data Store, the Initiative in Innovative Computing and the CUDA Center for Excellence sponsored by NVIDIA at Harvard University, and the International Centre for Radio Astronomy Research (ICRAR), a Joint Venture of Curtin University and The University of Western Australia, funded by the Western Australian State government.

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI), which is supported by the Australian Government. This work was supported by the Flagship Allocation Scheme of the NCI National Facility at the ANU.

We gratefully acknowledge discussions with members of the LO-FAR Transients Key Science Project when determining the optimal strategy for TRAP processing.

REFERENCES

- Bannister K. W., Murphy T., Gaensler B. M., Hunstead R. W., Chatterjee S., 2011, *MNRAS*, 412, 634
- Bastian T. S., Dulk G. A., Leblanc Y., 2000, *ApJ*, 545, 1058
- Bell M. E. et al., 2014, *MNRAS*, 438, 352
- Berger E. et al., 2001, *Nature*, 410, 338
- Carbone D. et al., 2014a, *MNRAS*, preprint ([arXiv:1411.7928](https://arxiv.org/abs/1411.7928))
- Cendes Y. et al., 2014b, preprint ([arXiv:1412.3986](https://arxiv.org/abs/1412.3986))
- Coenen T. et al., 2014, *A&A*, 570, A60
- Cohen A. S., Lane W. M., Cotton W. D., Kassim N. E., Lazio T. J. W., Perley R. A., Condon J. J., Erickson W. C., 2007, *AJ*, 134, 1245
- Condon J. J., 1974, *ApJ*, 188, 279
- Condon J. J., Cotton W. D., Greisen E. W., Yin Q. F., Perley R. A., Taylor G. B., Broderick J. J., 1998, *AJ*, 115, 1693
- Cordes J. M., Lazio T. J. W., 2002, preprint ([arXiv:astro-ph/0207156](https://arxiv.org/abs/astro-ph/0207156))
- Croft S., Bower G. C., Whysong D., 2013, *ApJ*, 762, 93
- Dolag K., Gaensler B. M., Beck A. M., Beck M. C., 2015, *MNRAS*, 451, 4277
- Ellingson S. W. et al., 2013, *IEEE Trans. Antennas Propagation*, 61, 2540
- Falcke H., Rezzolla L., 2014, *A&A*, 562, A137
- Fender R., Stewart A., Macquart J.-P., Donnarumma I., Murphy T., Deller A., Paragi Z., Chatterjee S., 2015, preprint ([arXiv:1507.00729](https://arxiv.org/abs/1507.00729))
- Gaensler B. M., Madsen G. J., Chatterjee S., Mao S. A., 2008, *PASA*, 25, 184
- Hancock P. J., Murphy T., Gaensler B. M., Hopkins A., Curran J. R., 2012, *MNRAS*, 422, 1812
- Hassall T. E., Keane E. F., Fender R. P., 2013, *MNRAS*, 436, 371
- Hurley-Walker N. et al., 2014, *PASA*, 31, e045
- Hyman S. D., Wijnands R., Lazio T. J. W., Pal S., Starling R., Kassim N. E., Ray P. S., 2009, *ApJ*, 696, 280
- Inoue S., 2004, *MNRAS*, 348, 999
- Intema H. T., van der Tol S., Cotton W. D., Cohen A. S., van Bemmell I. M., Röttgering H. J. A., 2009, *A&A*, 501, 1185
- Ioka K., 2003, *ApJ*, 598, L79
- Jaeger T. R., Osten R. A., Lazio T. J., Kassim N., Mutel R. L., 2011, *AJ*, 142, 189
- Jaeger T. R., Hyman S. D., Kassim N. E., Lazio T. J. W., 2012, *AJ*, 143, 96
- Kaplan D. L. et al., 2015, *ApJ*, 809, L12
- Karastergiou A. et al., 2015, *MNRAS*, 452, 1254
- Kashiyama K., Ioka K., Mészáros P., 2013, *ApJ*, 776, L39
- Keane E. F., Petroff E., 2015, *MNRAS*, 447, 2852
- Lane W. M., Cotton W. D., van Velzen S., Clarke T. E., Kassim N. E., Helmboldt J. F., Lazio T. J. W., Cohen A. S., 2014, *MNRAS*, 440, 327
- Law C. J., Jones G., Backer D. C., Barott W. C., Bower G. C., Gutierrez-Kraybill C., Williams P. K. G., Werthimer D., 2011, *ApJ*, 742, 12
- Law C. J. et al., 2015, *ApJ*, 807, 16
- Lazio T. J. W. et al., 2010, *AJ*, 140, 1995
- Loeb A., Shvartzvald Y., Maoz D., 2014, *MNRAS*, 439, L46
- Loi S. T. et al., 2015a, *MNRAS*, 453, 2731
- Loi S. T. et al., 2015b, *Geophys. Res. Lett.*, 42, 3707
- Lonsdale C. J., 2005, *ASP Conf. Ser. Vol. 345, Configuration Considerations for Low Frequency Arrays*. Astron. Soc. Pac., San Francisco, p. 399
- Lonsdale C. J. et al., 2009, *The Murchison Widefield Array: Design Overview*. Proc. IEEE, Vol. 97, p. 1497
- Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, *Science*, 318, 777
- Lyubarsky Y., 2014, *MNRAS*, 442, L9
- McGilchrist M. M., Baldwin J. E., Riley J. M., Titterton D. J., Waldram E. M., Warner P. J., 1990, *MNRAS*, 246, 110
- Macquart J.-P., Johnston S., 2015, *MNRAS*, 451, 3278
- Mauch T., Murphy T., Buttery H. J., Curran J., Hunstead R. W., Piestrzynski B., Robertson J. G., Sadler E. M., 2003, *MNRAS*, 342, 1117
- Metzger B. D., Williams P. K. G., Berger E., 2015, *ApJ*, 806, 224
- Mitchell K. J., Dennison B., Condon J. J., Altschuler D. R., Payne H. E., O'dell S. L., Broderick J. J., 1994, *ApJS*, 93, 441
- Morales M. F., 2005, *ASP Conf. Ser. Vol. 345, Design principles of the Mileura Wide-field Array Low Frequency Demonstrator (MWA-LFD)*. Astron. Soc. Pac., San Francisco, p. 452
- Murphy T. et al., 2013, *PASA*, 30, e006
- Murphy T. et al., 2015, *MNRAS*, 446, 2560
- Obenberger K. S. et al., 2015, *J. Astron. Instrum.*, 4, 1550004
- Offringa A. R., de Bruyn A. G., Biehl M., Zaroubi S., Bernardi G., Pandey V. N., 2010, *MNRAS*, 405, 155
- Offringa A. R. et al., 2014, *MNRAS*, 444, 606
- Offringa A. R. et al., 2015, *PASA*, 32, e008
- Rane A., Lorimer D. R., Bates S. D., McMann N., McLaughlin M. A., Rajwade K., 2016, *MNRAS*, 455, 2207
- Spreuw J. N., 2010, PhD thesis, University of Amsterdam, URL <http://hdl.handle.net/11245/1.324881>
- Stewart A. J. et al., 2016, *MNRAS*, 456, 2321
- Swinbank J. D. et al., 2015, *Astron. Comput.*, 11, 25
- Thornton D. et al., 2013, *Science*, 341, 53
- Tingay S. J. et al., 2013, *PASA*, 30, e007
- Tingay S. J. et al., 2015, *AJ*, 150, 199
- Totani T., 2013, *PASJ*, 65, L12
- TRAP Contributors 2014, *ascl.soft*, 1412.011
- Trott C. M., Tingay S. J., Wayth R. B., 2013, *ApJ*, 776, LL16
- van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Wayth R. B. et al., 2015, *PASA*, 32, e025
- Zhang B., 2014, *ApJ*, 780, L21

¹Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, NL-1090 GE, Amsterdam, the Netherlands

²Netherlands Institute for Radio Astronomy (ASTRON), PO Box 2, NL-7990 AA Dwingeloo, the Netherlands

³CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia

⁴ARC Centre of Excellence for All-sky Astrophysics (CAASTRO), Australia

⁵Sydney Institute for Astronomy, School of Physics, The University of Sydney, NSW 2006, Australia

⁶International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley, WA 6102, Australia

⁷Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁸Department of Physics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

⁹SKA SA, 3rd Floor, The Park, Park Road, Pinelands, 7405, South Africa

¹⁰Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa

¹¹Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

¹²*School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA*

¹³*Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia*

¹⁴*MIT Haystack Observatory, Westford, MA 01886, USA*

¹⁵*Raman Research Institute, Bangalore 560080, India*

¹⁶*Dunlap Institute for Astronomy and Astrophysics, University of Toronto, ON, M5S 3H4, Canada*

¹⁷*Department of Physics, University of Washington, Seattle, WA 98195, USA*

¹⁸*School of Chemical & Physical Sciences, Victoria University of Wellington, Wellington 6140, New Zealand*

¹⁹*National Centre for Radio Astrophysics, Tata Institute for Fundamental Research, Pune 411007, India*

²⁰*School of Physics, The University of Melbourne, Parkville, VIC 3010, Australia*

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