

Interferometric Imaging with the 32 Element Murchison Wide-Field Array

S. M. ORD,¹ D. A. MITCHELL,¹ R. B. WAYTH,^{1,2} L. J. GREENHILL,¹ G. BERNARDI,¹ S. GLEADOW,³ R. G. EDGAR,⁴
M. A. CLARK,^{1,4} G. ALLEN,⁵ W. ARCUS,² L. BENKEVITCH,⁶ J. D. BOWMAN,⁷ F. H. BRIGGS,⁸ J. D. BUNTON,⁵
S. BURNS,⁹ R. J. CAPPALLO,⁶ W. A. COLES,¹⁰ B. E. COREY,⁶ L. DESOUZA,⁵ S. S. DOELEMEN,⁶ M. DEROME,⁶
A. DESHPANDE,¹¹ D. EMRICH,² R. GOEKE,¹² M. R. GOPALAKRISHNA,¹¹ D. HERNE,² J. N. HEWITT,¹²
P. A. KAMINI,¹¹ D. L. KAPLAN,¹³ J. C. KASPER,¹ B. B. KINCAID,⁶ J. KOCZ,⁸ E. KOWALD,⁸
E. KRATZENBERG,⁶ D. KUMAR,¹¹ C. J. LONSDALE,⁶ M. J. LYNCH,² S. R. MCWHIRTER,⁶
S. MADHAVI,¹¹ M. MATEJEK,¹² M. F. MORALES,¹⁴ E. MORGAN,¹² D. OBEROI,⁶
J. PATHIKULANGARA,⁵ T. PRABU,¹¹ A. E. E. ROGERS,⁶ A. ROSHI,¹¹
J. E. SALAH,⁶ A. SCHINKEL,⁵ N. UDAYA SHANKAR,¹¹
K. S. SRIVANI,¹¹ J. STEVENS,⁵ S. J. TINGAY,²
A. VACCARELLA,⁸ M. WATERSON,^{2,8}
R. L. WEBSTER,³ A. R. WHITNEY,⁶
A. WILLIAMS,¹⁵ AND C. WILLIAMS¹²

Received 2010 July 17; accepted 2010 September 03; published 2010 October 6

ABSTRACT. The Murchison Wide-Field Array (MWA) is a low-frequency radio telescope, currently under construction, intended to search for the spectral signature of the epoch of reionization (EOR) and to probe the structure of the solar corona. Sited in western Australia, the full MWA will comprise 8192 dipoles grouped into 512 tiles and will be capable of imaging the sky south of 40° declination, from 80 MHz to 300 MHz with an instantaneous field of view that is tens of degrees wide and a resolution of a few arcminutes. A 32 station prototype of the MWA has been recently commissioned and a set of observations has been taken that exercise the whole acquisition and processing pipeline. We present Stokes I , Q , and U images from two ~ 4 hr integrations of a field 20° wide centered on Pictoris A. These images demonstrate the capacity and stability of a real-time calibration and imaging technique employing the weighted addition of warped snapshots to counter extreme wide-field imaging distortions.

Online material: color figures

1. INTRODUCTION

The MWA is intended to be a 512 station low-frequency radio telescope, as described in detail in Lonsdale et al.

(2009). In contrast to other connected element interferometers, the MWA will not store visibilities, instead performing calibration and imaging in real time with little or no human interaction. The MWA will be a compact radio telescope, with the longest baselines intended to be ~ 1.5 km. This compact size is a significant advantage for treating ionospheric effects, which are the dominant source of phase error. Under typical conditions the physical extent of the array is within an ionospheric isoplanatic patch, permitting image-based calibration of the ionospheric phase terms (Kassim et al. 2007). The baseline distribution and the time variability of the ionospheric calibration solutions requires that visibilities be averaged for no longer than 8 s before calibration (Lonsdale et al. 2009), and this is the fundamental cadence of the imaging pipeline. Data volume considerations preclude the storage of 8 s data cubes; the output is reduced by integrating for several minutes in the case of the EOR or restricting the field of view to a narrow range about the Sun for heliospheric science. The wide-field nature of the MWA experiment presents a significant challenge to image integration, and we employ image-based transformations on individual snapshots

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138; sord@cfa.harvard.edu.

² Curtin University of Technology, GPO Box U1987, Perth, WA 6845, Australia.

³ University of Melbourne, School of Physics, Parkville, VIC 3010, Australia.

⁴ Initiative for Innovative Computing, 29 Oxford Street, Cambridge, MA 02138.

⁵ CSIRO Astronomy and Space Science, ATNF, P.O. Box 76, Epping, NSW 1710, Australia.

⁶ MIT-Haystack Observatory, Westford, MA.

⁷ California Institute of Technology, Pasadena, CA 91125.

⁸ The Australian National University, Weston Creek, ACT 2611, Australia.

⁹ Burns Industries, 141 Canal Street, Nashua, NH 03064.

¹⁰ University of California, San Diego, La Jolla, CA 92093.

¹¹ Raman Research Institute (RRI), Bangalore, India.

¹² MIT Kavli Institute, 70 Vassar St., Cambridge, MA 02139.

¹³ University of California Santa Barbara, CA 93106-4030.

¹⁴ University of Washington, Seattle, WA 98195.

¹⁵ University of Western Australia, Nedlands, WA 6009, Australia.

as discussed in Brouw (1971), Bracewell (1984), and Cornwell & Perley (1992) in order to mitigate these effects; we consider this to be the first time an extensive integration employing such a technique has appeared in the literature.

A 32 tile (32T) engineering prototype intended to demonstrate technologies required for the operation of the full array has recently been commissioned, and a processing pipeline has been developed that is capable of calibrating the array in real time as well as tracking, imaging, and integrating in time and frequency. The outline of this article is as follows: first, we briefly discuss the real-time instrument calibration scheme; then we discuss the wide-field imaging problem as it pertains to the MWA telescope in general and the 32 station prototype in particular; following that, image weighting and the removal of instrumental polarization are detailed; finally, the observations are presented.

2. CALIBRATION

The real-time calibration of the full 512 element MWA-512T is described in Mitchell et al. (2008) and Lonsdale et al. (2009). The snapshot beam of the MWA-512T allows the simultaneous calibration of both the instrument gain and the ionospheric phase contribution toward many calibrator sources every 8 s. The former is in the form of a 2×2 complex Jones matrix for each antenna toward each calibrator, and the latter is a fit for a λ^2 -dependent refractive shift of each calibrator. The individual calibrator measurements are used to constrain all-sky fits for use in the wide-field calibration of the instrument and in the removal of strong sources before gridding and imaging.

The snapshot sensitivity of 32T is insufficient, however, for making calibration measurements toward many sources, so a partial solution has been developed. A few strong unpolarized point sources are tracked, continual measurements of antenna Jones matrices are made, and the refractive effects of the ionosphere toward each source are measured and assumed to be direction-dependent angular displacements that are constant across the field of view. We designate one calibrator—contributing the most power—as the *primary calibrator* and use it to set the bulk direction-independent calibration and bandpass of each tile. Further direction-dependent gain is assumed to be equal for each antenna and given by our default beam model. The measurement process is then repeated for the rest of the calibrators, so that they can be subtracted if desired. We have discovered that this method is sufficient to calibrate the amplitude and phase of the tile gain at any given pointing, but have found that the amplitude of this complex gain does not transfer well to other telescope pointing positions. The phase solution is sufficiently good to permit application of one phase solution to another pointing, separated by many degrees and considerable time, but the amplitude calibration does not transfer well, being in error by several percent.

The longest 32T baseline is ~ 300 m and since we expect ionospheric displacements to be small compared to this rela-

tively large 32T pixel size (Erickson, 1984) we do not apply an ionospheric solution that is a function of position within the field of view. However, if a field contains multiple calibrators, they each have an independent gain and ionospheric fit.

3. WIDE-FIELD IMAGING

3.1. The Problem

By way of introduction, the imaging transformation performed by an interferometer can be described by the van Cittert-Zernike equation (see Clark 1973, and Thompson, Moran and Swenson, 2001, for a detailed treatment):

$$V(u, v, w) = \iint \frac{A(\ell, m)I(\ell, m)}{n} \times \exp(-i2\pi(\ell u + mv + (n-1)w))d\ell dm. \quad (1)$$

Here, the u , v , and w are the spacing (the distance between antenna pairs) measured in wavelengths, and ℓ , m , and n are direction cosines of the brightness distribution with respect to this coordinate frame ($n = \sqrt{1 - \ell^2 - m^2}$). The coordinate system is defined such that the w -axis points in the direction of the phase center. The van Cittert-Zernike equation indicates that the complex-valued visibility function, $V(u, v, w)$, is a Fourier-like integral of the sky brightness, $I(\ell, m)$, multiplied by the primary-beam response of an interferometer, $A(\ell, m)$, and $1/n$.

In *standard* imaging it has been conventional to make the van Cittert-Zernike equation independent of n and w to and reduce it to a two-dimensional Fourier transform. This is conventionally achieved by the *small-field* approximation ($n \approx 1$), which results in $(n-1)w \approx 0$ and a projection from the celestial sphere onto a plane linear in ℓ and m , presenting an orthographic projection of the sky. However, this approximation effectively neglects w , implying that the direction of the projection is always perpendicular to the (ℓ, m) plane, and not perpendicular to the tangent plane. This manifests as a phase error that can only be kept to a tolerable level by ensuring the small-field approximation holds, thus restricting the size of the field of view. The phase error to first order is given by

$$\phi_w \sim \pi w(\ell^2 + m^2). \quad (2)$$

The small-field approximation does not hold for the MWA and therefore a mechanism must be applied to reduce this phase error; there are a number of approaches in the literature: three-dimensional Fourier transforms (Perley & Cornwell 1991), faceting (Cornwell & Perley 1992), and w -projection (Cornwell et al. 2005). But the MWA intends to take advantage of the fact that snapshot response of a coplanar array results in an image of the sky that is related to the true sky via a coordinate transformation (Brouw 1971; Bracewell 1984).

3.1.1. Snapshot Imaging with the MWA

As discussed in Brouw (1971), Bracewell (1984), Cornwell & Perley (1992), and Cornwell et al. (2005), the coplanar assumption amounts to saying that although w may not be zero, it is at all times some linear combination of u and v , and therefore that all the u and v lie in a plane given by

$$w = au + bv, \quad (3)$$

where

$$a = \tan Z \sin \chi, \quad (4)$$

$$b = -\tan Z \cos \chi, \quad (5)$$

Z is zenith distance, and χ is parallactic angle, as defined by

$$\tan \chi = \frac{\cos \phi \sin H}{\sin \phi \cos \delta - \cos \phi \sin \delta \cos H}, \quad (6)$$

$$\cos Z = \sin \phi \sin \delta + \cos \phi \cos \delta \cos H, \quad (7)$$

where H is hour angle, δ is declination, and ϕ is latitude. We can then eliminate w from Equation (1) to obtain

$$V(u, v) = \iint \frac{A(\ell, m)I(\ell, m)}{n} \times \exp(-i2\pi(\ell'u + m'v))d\ell dm, \quad (8)$$

revealing an instantaneous Fourier relationship between sky brightness and visibility, at the expense of a position-dependent coordinate distortion of the sky; that is now described by ℓ' and m' . The relationship between the measured ℓ' and m' and the “orthographic” ($w = 0$) ℓ and m is given by

$$\ell' = \ell + a(\sqrt{1 - \ell^2 - m^2} - 1), \quad (9)$$

$$m' = m + b(\sqrt{1 - \ell^2 - m^2} - 1), \quad (10)$$

$$d\ell dm = \left(1 - \frac{(a\ell + bm)}{n}\right)d\ell' dm'. \quad (11)$$

The effect of this is to warp the coordinate grid, and the coordinates ℓ and m are generally transformed to the sky coordinates of right ascension (α) and declination (δ) by the following relationships:

$$\ell = \cos \delta \sin \Delta\alpha, \quad (12)$$

$$m = \sin \delta \cos \delta_0 - \cos \delta \sin \delta_0 \cos \Delta\alpha, \quad (13)$$

$$n = \sin \delta \sin \delta_0 + \cos \delta \cos \delta_0 \cos \Delta\alpha, \quad (14)$$

where (α_0, δ_0) is the sky position of the tracking center and $\Delta\alpha = \alpha - \alpha_0$. These relations are then inverted to give sky coordinates:

$$\delta = \arcsin(m \cos \delta_0 + \sqrt{1 - \ell^2 - m^2} \sin \delta_0), \quad (15)$$

$$\Delta\alpha = \arctan\left(\frac{\ell}{\sqrt{1 - \ell^2 - m^2} \cos \delta_0 - m \sin \delta_0}\right). \quad (16)$$

If no account is taken of the difference between measured ℓ' , m' and orthographic ℓ , m , then this coordinate distortion results in an incorrect labeling of the sky coordinates in the image. Conversely, this wide-field error can be negated by transforming the coordinate space of the snapshot images to account for this distortion.

It has been noted (Calabretta & Greisen 2002) that the coordinate projection described previously is a slant orthographic projection or generalized SIN projection. This is incorporated into the Flexible Image Transport System (FITS) standard with the projection code SIN, but with extra FITS keywords to implement the slant. The SIN projection is extended with the use of the PV2_1 and PV2_2 FITS keywords set to a and $-b$, respectively. The approach we have followed in the real-time imaging pipeline is to define a world coordinate system (WCS) for each snapshot that incorporates this projection and then utilize the WCSLIB coordinate transformation library to cast input pixels and their vertices into a common output frame. The individual snapshots are accumulated in this stable frame.

The MWA does not lie on a perfect plane; therefore, a phase error will still be present due to noncoplanarity. The standard deviation from planar sampling over the tiles of the MWA-32T in the north-south direction is 0.75 m and in the east-west direction is 0.44 m. The topology of the array being dominated by a slope in the east-west direction of gradient -0.006 mm^{-1} . The wavelength range of the MWA is approximately 1–3 m. The deviation from planar sampling is indicative of the rms deviation from zero w at zenith and this corresponds to a worst-case rms phase error at the edge of a 30° field of view (FOV) of $\phi_w = 0.32 \text{ rad}$. This phase error maps to a decorrelation factor (r) given by the approximation

$$r = \left(1 - \frac{1}{2}\phi_w^2\right), \quad (17)$$

which is 0.95, or 5%. For the observations presented in this article, which have been taken at a longer wavelength than this

limit and are over a narrower FOV, the decorrelation factor is 2.5%.

3.1.2. Defining the Output Frame, Resampling, and Integration

The resampling of the measured brightness distribution is the method by which the wide-field distortions are removed in the 32T Real-Time System (RTS) imaging pipeline and the mechanism for projecting the instrument coordinates into a library frame. Due to the equal-area nature of its pixels, the readily available indexing tools, and the fact that it is an all-sky pixelization, we have chosen to regrid into the HEALPix pixelization of the sphere (Górski et al. 2005). The HEALPix scheme is commonly used by the cosmic microwave background community, and many tools exist to perform analysis tasks on this data format. Calabretta & Roukema (2007) present the family of projections to which HEALPix belongs and also a new FITS projection (HPX) that permits imaging of HEALPix data without any interpolation onto a regular grid. To perform the real-time resampling we have implemented a distance-weighted interpolation scheme. Each output pixel value in the HEALPix frame is formed by a distance-weighted average of the nearest input instrument-frame pixels. Integration can be performed by the addition of the HEALPix pixels and the application and storage of the snapshot image weights.

3.1.3. The Weighting Scheme

The MWA tiles are phased arrays of 16 crossed dipoles on a ground plane, and they display significant instrumental polarization. The changing projection of the dipoles on the sky results in a polarization response that is a strong function of source position in the telescope beam. Furthermore, the tiles are not identical; different tiles display different direction-dependent complex gains, due to antenna manufacturing tolerances and variations in the performance of electrical components. In combining the warped snapshots we attempt to produce a stable polarimetric measurement of the sky brightness distribution as detected by the interferometer. The changing projection of the dipoles on the sky as a function of hour angle results in different measurements of sky polarization that have varying direction-dependent weight. For example, at the beginning of a track, a field rises in the east and only a fraction of the Y -polarization dipole is projected onto the sky in that direction, whereas the full X -polarization is projected. At zenith both polarizations are presented equally to the sky. The beam response is not just a function of this projection, but the combination of 16 phased dipoles and the ground screen. These differing contributions result in polarimetric measurements of varying fidelity as a function of hour angle, which is a feature that must be incorporated in the combination of snapshots by a suitable weighting scheme.

For each output image pixel, we are attempting to form the best estimate of the sky-polarized brightness ($\bar{\mathbf{b}}_{\text{pix}}$) from the

combination of, n individual measurements ($\mathbf{b}_{\text{pix},j}$), each of different weight ($w_{\text{pix},j}$). The best such combined estimate is the weighted mean:

$$\bar{\mathbf{b}}_{\text{pix}} = \frac{\sum_{j=1}^n w_{\text{pix},j} \mathbf{b}_{\text{pix},j}}{\sum_{j=1}^n w_{\text{pix},j}}, \quad (18)$$

where the weights ($w_{\text{pix},j}$) are generally the reciprocal of the pixel variance. The inverse variance of the MWA image pixels is best represented by the direction-dependent polarimetric tile response in the direction of that pixel (the power response). We are therefore weighting by the power signal-to-noise ratio (S/N), under the assumption that the thermal noise is direction-independent. Following the formalism developed in the series of papers by Hamaker et al. (1996); Sault et al. (1996); Hamaker & Bregman (1996); Hamaker (2000), and Hamaker (2006), the brightness of an image pixel in the measured, or instrument frame is

$$\mathbf{b}_{\text{pix}}^{\mathbf{m}} = (\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*) \mathbf{S} \mathbf{b}_{\text{pix}}^{\mathbf{s}}, \quad (19)$$

where \mathbf{b} are the coherency 4-vectors, the superscript \mathbf{m} indicates measured and should be taken to imply an instrument polarization basis (p, q), and \mathbf{s} represents a Stokes ($I, Q, U, \text{ and } V$) basis. The matrix \mathbf{S} transforms polarization from the instrument basis (orthogonal linear feeds) to the Stokes basis on the sky and is

$$\mathbf{S} = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & i \\ 0 & 0 & 1 & -i \\ 1 & -1 & 0 & 0 \end{pmatrix}. \quad (20)$$

The MWA-32T is an interferometer comprising 32 different elements, but, as is described in § 2, each visibility is calibrated by an equalization process that allows us to approximate the instrument polarimetric response by an *instrument Jones matrix*. The 2×2 instrument Jones matrix \mathbf{J} determines how the generally nonorthogonal projected receptors (p, q) respond to the incident orthogonal sky polarization (X, Y). As the projection of the tile receptors changes as a function of position, the Jones matrix is different for every pixel.

It is this instrument Jones matrix that is used to form the variance weight matrix for each input pixel. As we are not accounting for the different primary-beam shape, we are not forming the best estimate of the array response in a given pixel direction; it is this simplification that makes the problem tractable in the 32T domain. But it should be noted that the weighting scheme will evolve as the MWA develops. As more tiles are added, significantly overconstraining the parameters of the array, a direction-dependent calibration scheme as described in Mitchell et al. (2008), with a more accurate weighting scheme incorporating

the differences between tiles, will be applied in the visibility domain.

3.1.4. Building and Applying the Weights

The 2×2 Jones matrix of a radio telescope is generally decomposed into

$$\mathbf{J} = \mathbf{GDCPF}, \quad (21)$$

where \mathbf{G} is the electronic gain, \mathbf{D} represents the feed errors, \mathbf{C} is the configuration of the feed response, \mathbf{P} is parallactic rotation, and \mathbf{F} is Faraday rotation. We have not attempted to correct for Faraday rotation in the experiment, the ramifications of which are explored in § 4. In our application the dipoles do not rotate relative to the source, but the projection of the dipoles changes on the sky:

$$\mathbf{J} = \mathbf{GDJ}_p. \quad (22)$$

where \mathbf{J}_p is the dipole projection matrix given by

$$\mathbf{J}_p = \begin{pmatrix} \cos \varphi \cos \delta + \sin \varphi \sin \delta \cos H & -\sin \varphi \sin H \\ \sin \delta \sin H & \cos H \end{pmatrix}, \quad (23)$$

where ϕ , δ , and H are latitude, declination, and hour angle, respectively. In the calibration stage of the imaging pipeline (Mitchell et al. 2008), the Jones matrices of the individual antenna elements are solved for, which essentially allows \mathbf{G} and \mathbf{D} to be obtained, but they are incorporated in the fitted Jones matrix and not decomposed.

The 32T calibration scheme attempts to solve for the elements of the diagonal \mathbf{G} matrix, and the rest of the matrices are assumed to be the same for all tiles. Thus, individual beams can be scaled and shifted (the elements are complex) but not changed in shape. This is required by the limited sensitivity of the 32T array.

The weight matrix must be constructed for every input sky pixel and in the 32T RTS is precomputed at the beginning of each ~ 3 minute integration. For each pixel the image coordinates (l' and m') are converted into a topocentric hour angle and declination. In this process, account has to be taken of the initial orthographic projection and the required spherical coordinates.

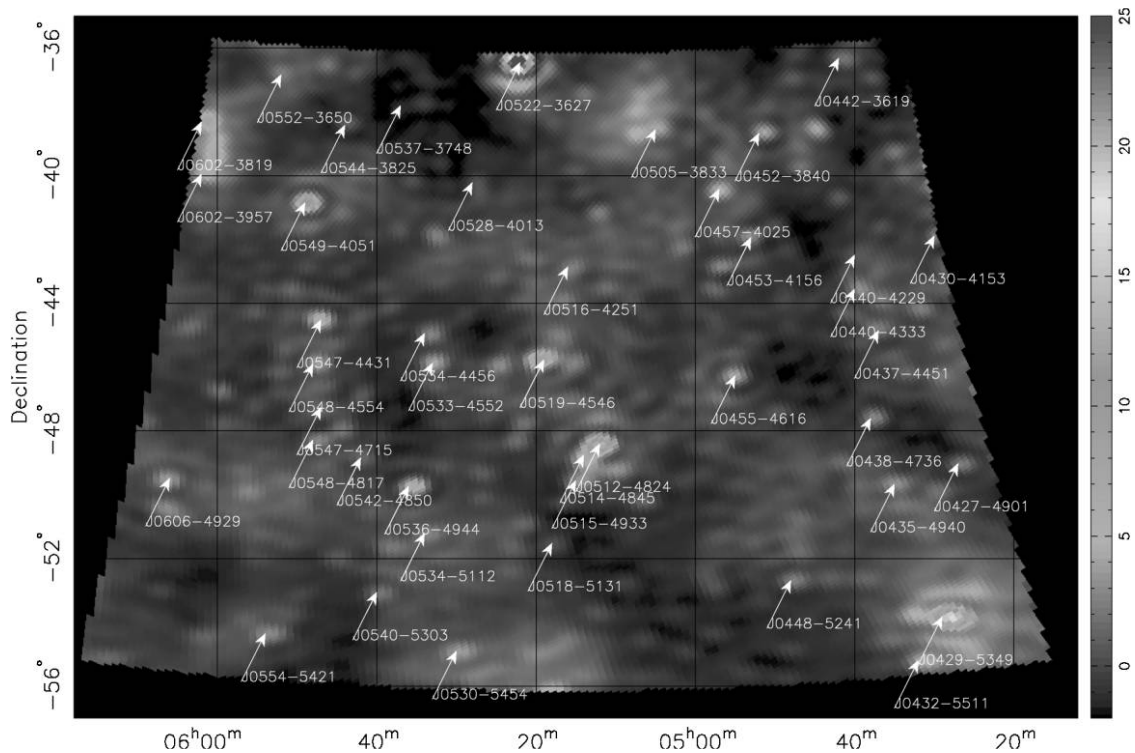


FIG. 1.—Total intensity (Stokes D) image in Jy beam^{-1} with a central frequency of approximately 103 MHz, a bandwidth of 30 MHz, and an integration time of approximately 4 hr. The pixelization is HEALPix, but the pixels have been rendered into a plate carrée projection for display purposes. A selection of sources from the Parkes catalog (Wright & Otrupcek [1990]) predicted to be above 5 Jy in this field are labeled. Many more sources are visible than are predicted—the predominant reason for which is that they are predicted to have fluxes less than the 5 Jy limit for the labeling. The presence of Pictoris A (J0519-4546) residual in this image, despite an attempt to remove this 400 Jy source from the visibilities, is indicative of an imperfect calibration. The grayscale has been clipped at 25 Jy beam^{-1} to enhance contrast. The true peak pixel flux is approximately 60 Jy beam^{-1} and is associated with the point source in the north of the map, J0522-3627. See the electronic edition of the *PASP* for a color version of this figure.

The model of the beam is then formed as the sum of 16 dipoles with a phase delay commensurate with the required tile pointing to the center. The complex gain of the tile in the desired pixel direction is then calculated and the Jones matrix for that pixel (\mathbf{J}_{pix}) is formed.

The next step is the formation of the image weights; ideally these would variance-weight the image so as to provide close to a maximum likelihood estimation of the sky. We do not have the true power response of the instrument, but we have the fitted model response \mathbf{J}_{pix} . We therefore calculate the power beam response from the fitted voltage response, and this is done by forming the outer product ($\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*$). We then form the Hermitian transpose of this matrix. And each input pixel polarimetric vector is formed from the measured instrumental polarization ($\mathbf{b}_{\text{pix}}^m$) multiplied by the weight matrix:

$$\mathbf{b}_{\text{pix}}^{\text{m,w}} = (\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*)^\dagger (\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*) \mathbf{S} \mathbf{b}_{\text{pix}}^s. \quad (24)$$

This is performed for every pixel in each snapshot. Then each input pixel is resampled into the output HEALPix frame.

4. STOKES CONVERSION

Accounting for the changing response of the tiles to correctly weight the instrumental polarization maps permits the production of a best estimate of the Stokes parameters when the weighted images are correctly normalized. However, the observations that are presented here were not taken in a manner that maximizes polarization fidelity. For the purposes of clarity, we will emphasize the following points:

1. We are assuming a direction-dependent response for the instrument for which a complex gain has been determined, but it is the same for each tile. This is not the case, and each tile will make a different contribution to the array response. This mimics a leakage between Stokes parameters at a level commensurate with the departure from the *mean* beam shape.

2. The bandwidth of each channel of the presented images is 1.92 MHz wide and will inflict considerable bandwidth depolarization on celestial sources; from Gardner & Whiteoak (1966) the depolarization factor is given by $\sin \Delta\phi / \Delta\phi$, where $\Delta\phi$ is given by

$$\Delta\phi \sim -RM\lambda_0^2 \frac{2\Delta\nu}{\nu}. \quad (25)$$

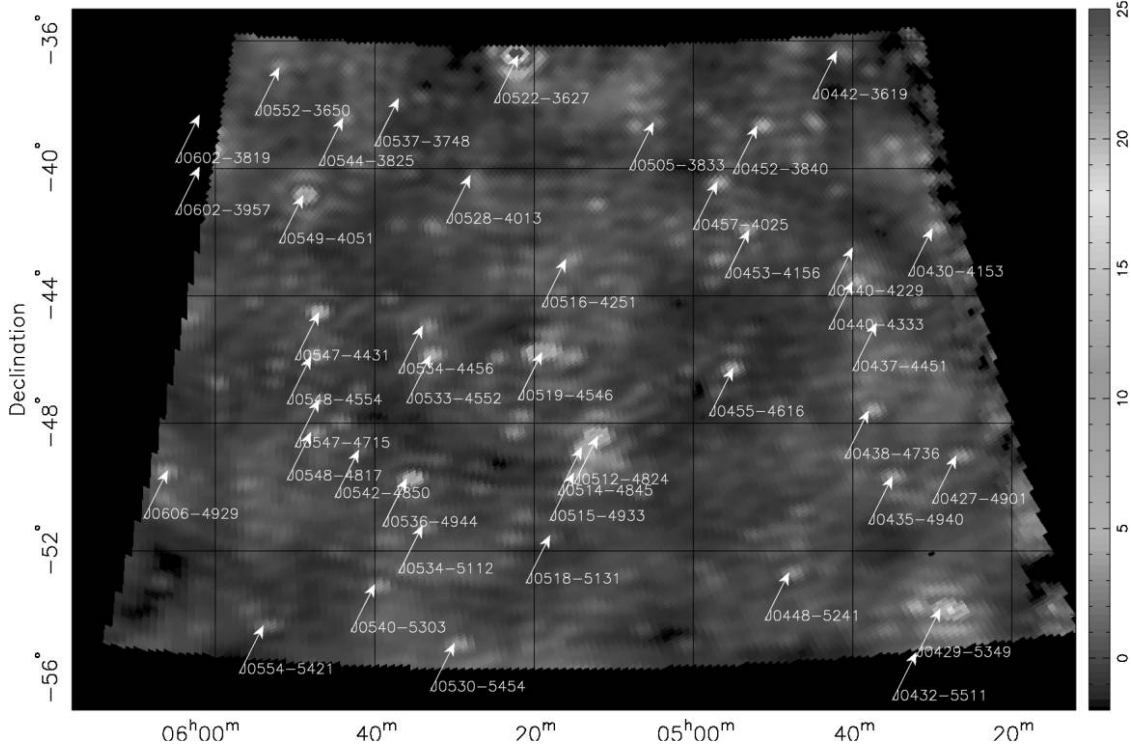


FIG. 2.—Total intensity (Stokes I) image in Jy beam^{-1} with a central frequency of approximately 134 MHz, a bandwidth of 30 MHz, and an integration time of approximately hr. The pixelization is HEALPix, but the pixels have been rendered into a plate carrée projection for display purposes. Sources from the Parkes catalog predicted to be above 5 Jy in this field are shown. The presence of a Pictoris A (J0519-4546) residual in this image, despite an attempt to remove this 400 Jy source from the visibilities, is indicative of an imperfect calibration. The grayscale has been clipped at 25 Jy beam^{-1} to enhance contrast. The true peak pixel flux is approximately 60 Jy beam^{-1} and is associated with the point source in the north of the map, J0522-3627. See the electronic edition of the *PASP* for a color version of this figure.

At a central wavelength of 2 m, this factor is 1% at a rotation measure (RM) of 1 rad m^{-2} , but 100% at an RM of $\sim 10 \text{ rad m}^{-2}$.

But we are also not correcting for any Faraday rotation between these channels, which will effectively ($\sim 80\%$) depolarize a signal with an RM as low as 1 rad m^{-2} .

3. Another depolarization factor results from not accounting for any variations in Faraday rotation within and between the subintegrations. The integrations run for several hours and are made up of 3 minute subintegrations, Erickson et al. (2001) present compelling observation at 327 MHz with the VLA that demonstrate Faraday rotation of tens of degrees throughout several hours of daytime observations. Our observations are at lower frequencies, so we would expect even more rotation (a factor of ~ 4 more) rotating flux between Stokes Q and Stokes U throughout the observation. The observations presented here were taken at nighttime, which should mitigate this effect; however, observations of the pulsar J0218 + 4232 at the Westerbork Synthesis Radio Telescope found typical nighttime RM variations of $0.2\text{--}0.4 \text{ rad m}^{-2}$: about 70° of variation at our frequencies (G. Bernardi 2008, unpublished).

Even a benevolent ionosphere with RM deviations considerably less than 1 rad m^{-2} would still be sufficient to significantly depolarize any polarized source or diffuse background during

the long integrations presented here. Given the combination of these various depolarization effects, any residual polarization in the maps presented here are most likely due to instrumental polarization leaking between the Stokes parameters, due to imprecise beam models and polarized side lobes from bright, intrinsically unpolarized, diffuse emission from outside the field of view.

4.1. Reconstructing the Weights and Converting to Stokes Parameters

Each output sky pixel is the result of the combination of many pixels with different weights, applied at different times. The 32T prototype solution to this problem is to store the parameters that are used to analytically calculate the weights. This information is then used off-line to reconstruct the direction-dependent instrument Jones matrix for each constituent weight, at the time and sky position that each weight was applied.

4.1.1. The Stokes Conversion as Part of Image Normalization

The weight normalization factor would normally be the sum of the variance weights as per equation (18); however, in order

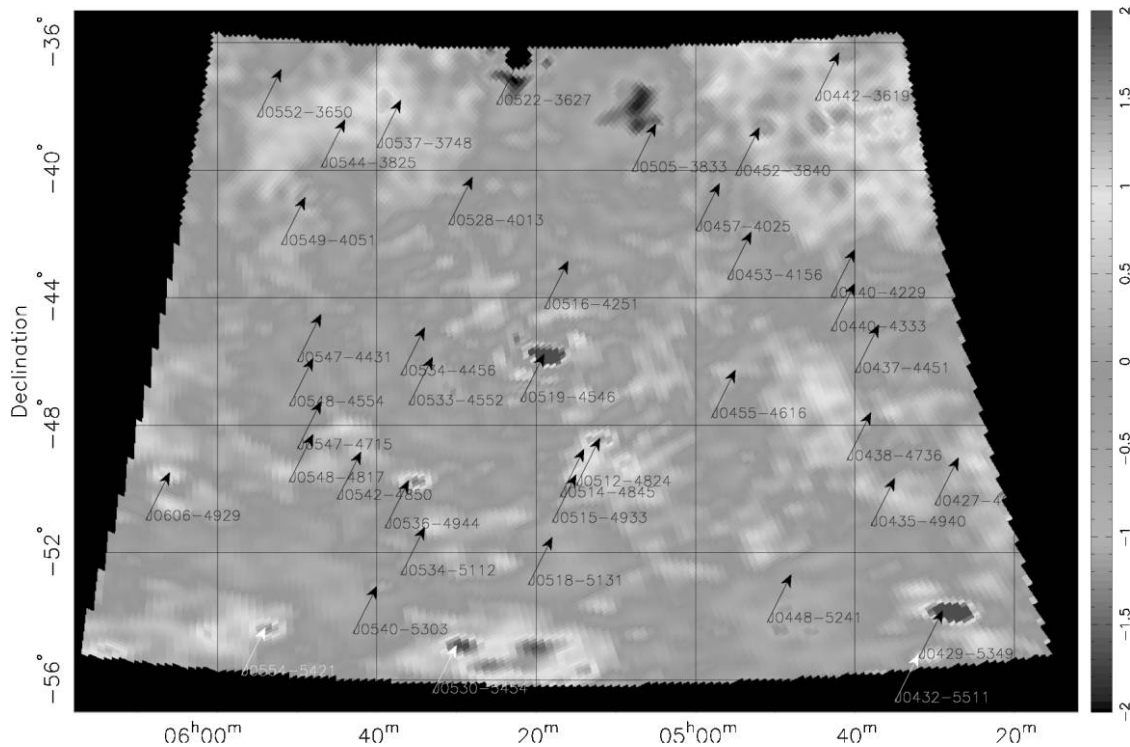


FIG. 3.—The Stokes Q image associated with the 103 MHz Stokes I image in Fig. 1. The average brightness of the image is significantly lower than the Stokes I ($0.14 \text{ Jy beam}^{-1}$ versus 0.9 Jy beam^{-1} ; see Table 1), but clear structure is evident associated with areas of high Stokes I . This is also evidence of calibration errors at the level of a few percent of Stokes I . As discussed in the text, the low frequency and wide bandwidth make it unlikely that the polarization structure is intrinsic to the objects. See the electronic edition of the *PASP* for a color version of this figure.

to simultaneously normalize the images and convert to Stokes parameters from instrument polarization, the following product is formed:

$$\mathbf{N}_{\text{HPX}} = \sum_{j=1}^n (\mathbf{J}_{j,\text{pix}} \otimes \mathbf{J}_{j,\text{pix}}^*)^\dagger (\mathbf{J}_{j,\text{pix}} \otimes \mathbf{J}_{j,\text{pix}}^*). \quad (26)$$

It is the formation of this product that permits the transformation to Stokes parameters. Essentially, the images are being stored in weighted-instrument polarization, but the information required to convert the instrumental polarization to the X-Y basis is also being stored. This is reasonable, as the weights are formed from the beam model, which has a limited number of parameters; if the weight matrix had to be stored, then four times more weight information would have to be stored than actual measurements.

The weight and transformation product is evaluated, summed, and inverted for every pixel. We then multiply the weighted pixels by this normalization matrix to produce an estimate of the true sky brightness $\hat{\mathbf{b}}_{\text{HPX}}^S$ in the following manner:

$$\hat{\mathbf{b}}_{\text{HPX}}^S = \sum_{j=1}^n [(\mathbf{J}_{j,\text{pix}} \otimes \mathbf{J}_{j,\text{pix}}^*)^\dagger (\mathbf{J}_{j,\text{pix}} \otimes \mathbf{J}_{j,\text{pix}}^*)]^{-1} \cdot \sum_{j=1}^n [(\mathbf{J}_{j,\text{pix}} \otimes \mathbf{J}_{j,\text{pix}}^*)^\dagger (\mathbf{J}_{j,\text{true}} \otimes \mathbf{J}_{j,\text{true}}^*)] \mathbf{Sb}_{\text{HPX}}^S. \quad (27)$$

As outlined previously, in practice, the summations occur during the image integration step and the weight integration step, so dropping the summation and the j subscript for clarity we have

$$\hat{\mathbf{b}}_{\text{HPX}}^S = [(\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*)^\dagger (\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*)]^{-1} (\mathbf{J}_{\text{pix}} \otimes \mathbf{J}_{\text{pix}}^*)^\dagger \times (\mathbf{J}_{\text{true}} \otimes \mathbf{J}_{\text{true}}^*) \mathbf{Sb}_{\text{HPX}}^S. \quad (28)$$

4.1.2. On the Equality of \mathbf{J}_{true} and \mathbf{J}_{pix}

The operation described previously effectively cancels the instrument response, but only to the extent that the \mathbf{J}_{pix} represents \mathbf{J}_{true} , as discussed in § 3.1.3. It is the level to which this does not hold that constrains the polarization fidelity of this experiment. Under this assumption, the application of \mathbf{S}^{-1} will

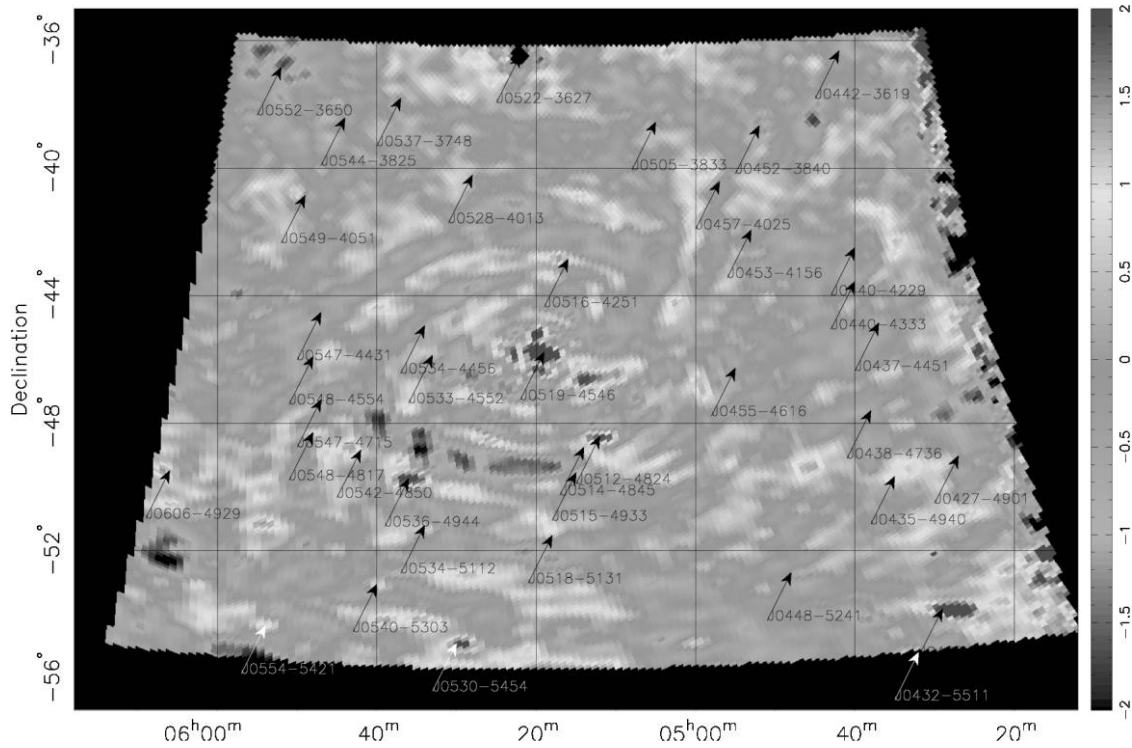


FIG. 4.—The Stokes Q image associated with the 134 MHz Stokes I image in Fig. 2. The average brightness of the image is significantly lower than the Stokes I ($0.18 \text{ Jy beam}^{-1}$ versus $0.76 \text{ Jy beam}^{-1}$; see Table 1), but clear structure is evident associated with areas of high Stokes I . This is also evidence of calibration errors at the level of a few percent of Stokes I . As discussed in the text, the low frequency and wide bandwidth make it unlikely that the polarization structure is intrinsic to the objects. See the electronic edition of the *PASP* for a color version of this figure.

convert the normalized sky coherency vector from the X - Y basis to the Stokes basis.

The tile model is a simplification of the true tile response, although informed by electromagnetic simulations, there is surely an appreciable difference in the polarized response of the true tile and that predicted by the current tile model. This is another aspect of 32T processing that will be less of a factor for the full array, which will attempt to fit more parameters of the tile model than have been employed here. Over and above model errors and the tile-to-tile variations, there are other possible differences between \mathbf{J}_{true} and \mathbf{J}_{pix} , an important one being Faraday rotation. As discussed in both §§ 3.1.3 and 4, the Faraday rotation is included in the Jones decomposition of the tile response; however, it is not included in \mathbf{J}_{pix} . As we know that there is an unknown Faraday component to \mathbf{J}_{true} , flux will be redistributed between Q and U as a function of time, which will effectively depolarize the images over long integrations. This effect can be corrected by including the direction-dependent Faraday component in \mathbf{J}_{pix} , and thus the normalization matrix, but this would involve fitting a time- and direction-dependent Faraday rotation, which we have not attempted here. Another consideration is that \mathbf{J}_{pix} is not calculated for every time step; instead, it is assumed to be valid for a length of time (typically, a few minutes). Simulations indicate that the validity of this approximation is a strong function of hour angle, with observa-

tions taken at zenith being immune to this effect for all reasonable integration times, but observations at 30° zenith angle will suffer 1% of polarization leakage for integrations of 2 minutes. To counter this effect in the 512T system, the \mathbf{J}_{pix} will be updated every cadence (8 s).

5. THE OBSERVATIONS

5.1. The Real-Time System Constraints

As discussed in the Introduction, the MWA when fully deployed will run in real time and will therefore require considerable computing resources. The MWA real-time computer (RTC), as implemented in the MWA-32T prototype, presages the architecture of the planned full RTC. It consists of a cluster of four quad-core Xeon E54 servers, each containing an nVidia GeForce 8800 GTS graphics card (graphics processing unit [GPU]). Visibilities from the correlator are presented to the cluster via a Gigabit Ethernet network. Each server receives 192 consecutive 40 kHz channels. These are processed individually and the entire imaging pipeline is run on the GPU. The processing load can be reduced by combining gridded visibilities across consecutive channels before imaging. Although the processing load on the calibration is a function of the number of baselines, the imaging load is a function of the number of pixels; therefore, the imaging pipeline is required to do a similar

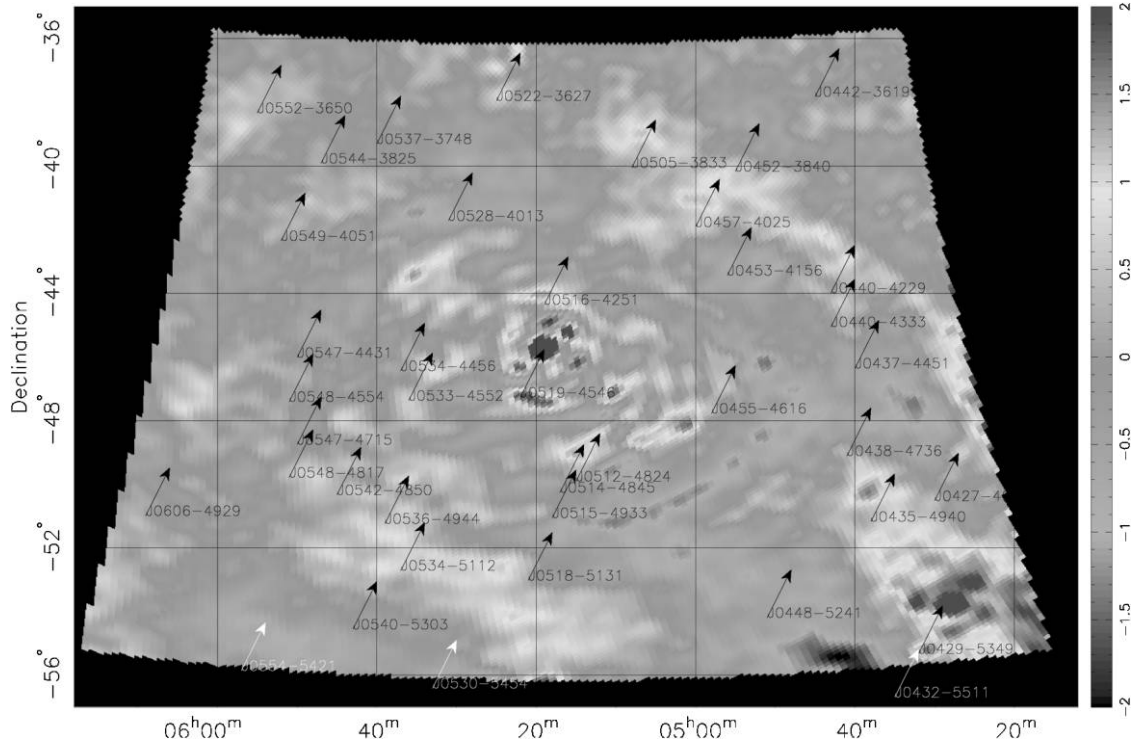


FIG. 5.—The Stokes U image associated with the 103 MHz Stokes I image in Fig. 1. The rms brightness of the image is significantly lower than the Stokes I ($0.16 \text{ Jy beam}^{-1}$ versus 0.9 Jy beam^{-1} ; see Table 1). There is considerably less association between bright Stokes I points and features in the Stokes U map than is evident in the Stokes Q map, although J0429-5349 is clearly present. See the electronic edition of the *PASP* for a color version of this figure.

amount of work in the prototype as the full MWA. The full MWA will be at a higher resolution, resulting in more pixels, but in actuality, the imaging load on an individual processing node in the 32T is higher than the full MWA, because 192 channels are sent to each node instead of 12; in order to perform in real time, it was necessary to make some adjustments to the pipeline. First, each imaging channel is ~ 2 MHz wide, which is much wider than the 40 kHz required for 32T and was chosen to be the widest that was possible without suffering more than 1% bandwidth decorrelation, and is integrated for just over 3 minutes in memory before each snapshot is output to disk.

As part of an MWA engineering test program in 2010 January, multiple tracks of a 20° field around Pictoris A were calibrated and imaged by the MWA real-time system. We present two 20° wide observations of Stokes I , Q , and U centered on Pictoris A, across 30 MHz and over ~ 4 hr, which were taken on 2010 January 18 and 2010 January 19. The RTS generated a calibration data stream and wide-field sky images with an 8 s cadence; calibration, imaging, and integration to 3 minutes were done in real time. The RTS output was written in the specific projection and all-sky storage frame that has been selected for the MWA (HEALPix) and has been weighted for optimal averaging in time and frequency. The output images have been subsequently integrated off-line, primary-beam-corrected, and

converted to Stokes parameters to produce the images presented in Figures 1–6.

5.2. Dynamic Range and Peeling

The noise in the integrated maps presented here, both simulated and actual, is dominated by side-lobe contamination and not thermal noise. Pictoris A is approximately 400 Jy at these frequencies and its side lobes are the main contributor to noise in unprocessed images. We have measured the root-mean-square brightness in Jy beam^{-1} for each of the images and present them in Table 1. As discussed in § 2, the 32T version of the RTS is based around measuring a Jones matrix for each tile toward a primary calibrator. This serves two purposes: since it is assumed that the tiles have known and identical beam shapes, the measured Jones matrices can be used with a source model to calibrate the visibilities over the entire primary beam, and a calibrated model of the source and its side lobes can be subtracted (or *peeled*) from the visibilities before gridding and imaging. We have subtracted a point-source model of Pictoris A directly from the visibilities during the calibration process in real time. The low S/N for a set of instantaneous (8 s) 32T visibilities limits the accuracy of our calibrator model and, as a result, a clear residual can be seen after the image noise has been integrated down. However, this process has increased

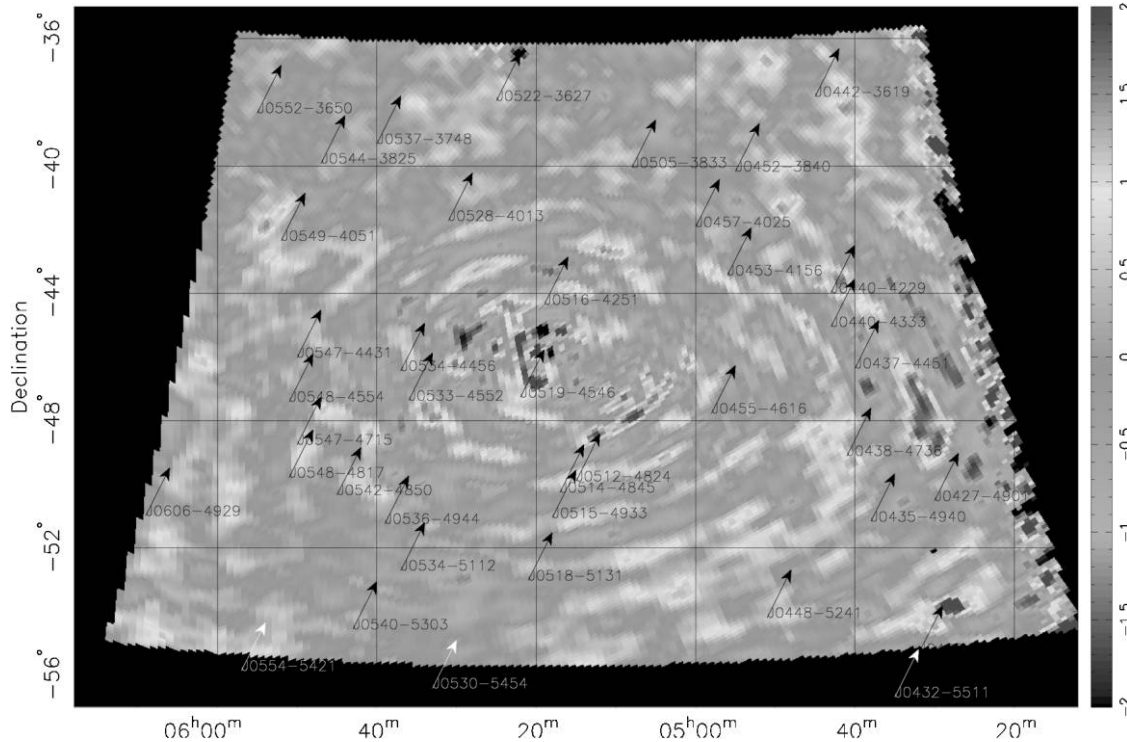


FIG. 6.—The Stokes U image associated with the 134 MHz Stokes I image in Fig. 2. The rms brightness of the image is significantly lower than the Stokes I ($0.19 \text{ Jy beam}^{-1}$ versus $0.76 \text{ Jy beam}^{-1}$; see Table 1). There is considerably less association between bright Stokes I points and features in the Stokes U map than is evident in the Stokes Q map, although J0429-5349 is clearly present. See the electronic edition of the *PASP* for a color version of this figure.

TABLE 1
ROOT-MEAN-SQUARE BRIGHTNESS IN Jy Beam^{-1} OF THE
OBSERVATIONS PRESENTED

Stokes parameter	Frequency (MHz)	RMS (Jy beam^{-1})
<i>I</i>	103	0.9
<i>I</i>	134	0.76
<i>Q</i>	103	0.14
<i>Q</i>	134	0.18
<i>U</i>	103	0.16
<i>U</i>	134	0.19

NOTE.— The integration time is approximately 4 hr and bandwidth is approximately 30 MHz. There is some variation between integrations, due to a variable amount of frequency and time data flagging, which fractionally lowers the effective bandwidth and integration time.

the dynamic range displayed in these images by a factor of 10, to about 400; the subtraction of adjacent channels, which suppresses the side-lobe noise very effectively, results in a noise floor consistent with a dynamic range of 1800.

Iterative deconvolution as it is routinely employed in interferometric imaging is not possible within this pipeline. The point-spread function of the instrument is spatially variant, and there is no access to the visibilities. It is for this reason that

we have employed peeling to remove point sources from the ungridded visibilities. The 32T prototype lacks the sensitivity to remove many sources this way, and the full array will be limited not by sensitivity, but by how many sources can be peeled in the allotted time. As a result, there will be residual point sources and polarized diffuse emission that will need to be removed by postprocessing, and it is likely that sky modeling plus simulations of the instrument response will be required to achieve this (Bernardi 2010; Pindor 2010).

5.3. Comparison to Simulations

The MAPS software package (Wayth et al. 2010, in preparation) has been used to provide a well-understood simulated input data set to the software pipeline. Figures 7, 8, and 9 are the companions to the actual observations presented, although at a slightly higher frequency (160.02 MHz). The simulation contains a subset of the Parkes catalog of point sources and a diffuse background in Stokes *I* from Haslam et al. (1981). The simulator generates visibilities as they would be produced by the MWA correlator, and the array parameters can be controlled by the user. In this case the calibration is perfect, there is no ionospheric refraction, and all the point sources and diffuse background are intrinsically unpolarized. The integration time was approximately 2 hr and only a single 40 kHz channel

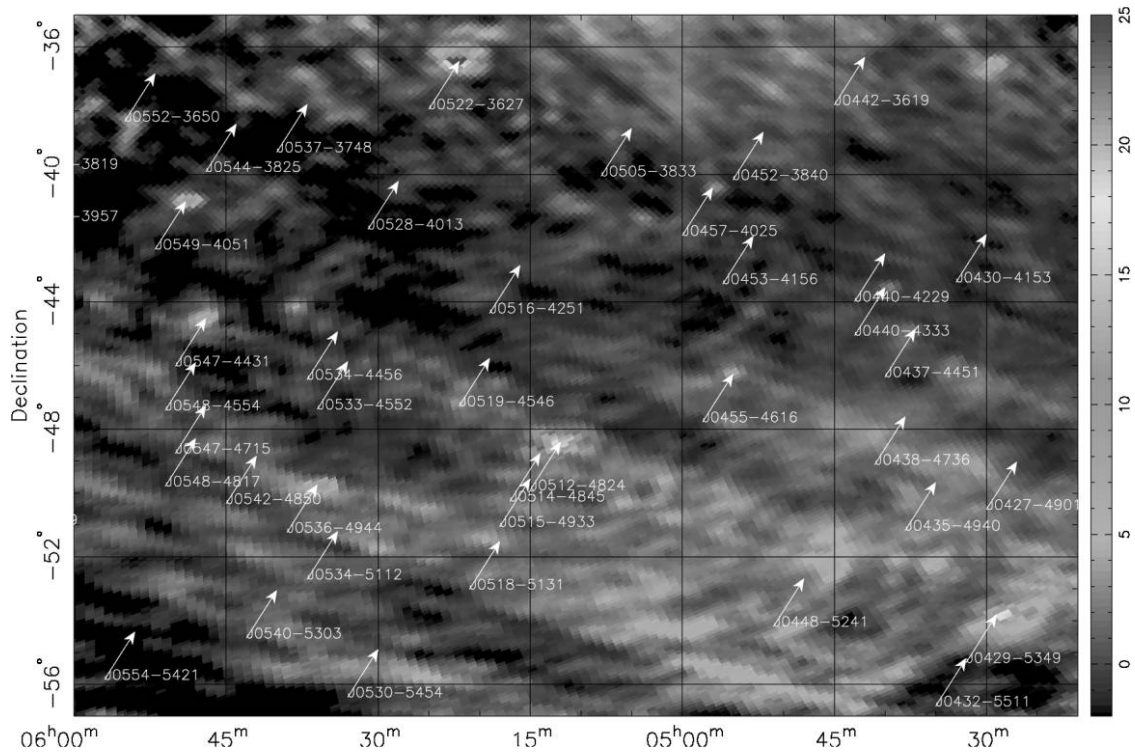


FIG. 7.—Total intensity (Stokes *I*) image of a simulated sky at a central frequency of approximately 160 MHz, with a bandwidth of 40 kHz and an integration time of approximately 2 hr. The pixelization is HEALPix, but the pixels have been rendered into a plate carrée projection for display purposes. The color scale has been clipped at 25 Jy beam^{-1} to enhance contrast. The true peak pixel flux is approximately 60 Jy beam^{-1} and is associated with the point source in the north of the map, J0522-3627. See the electronic edition of the *PASP* for a color version of this figure.

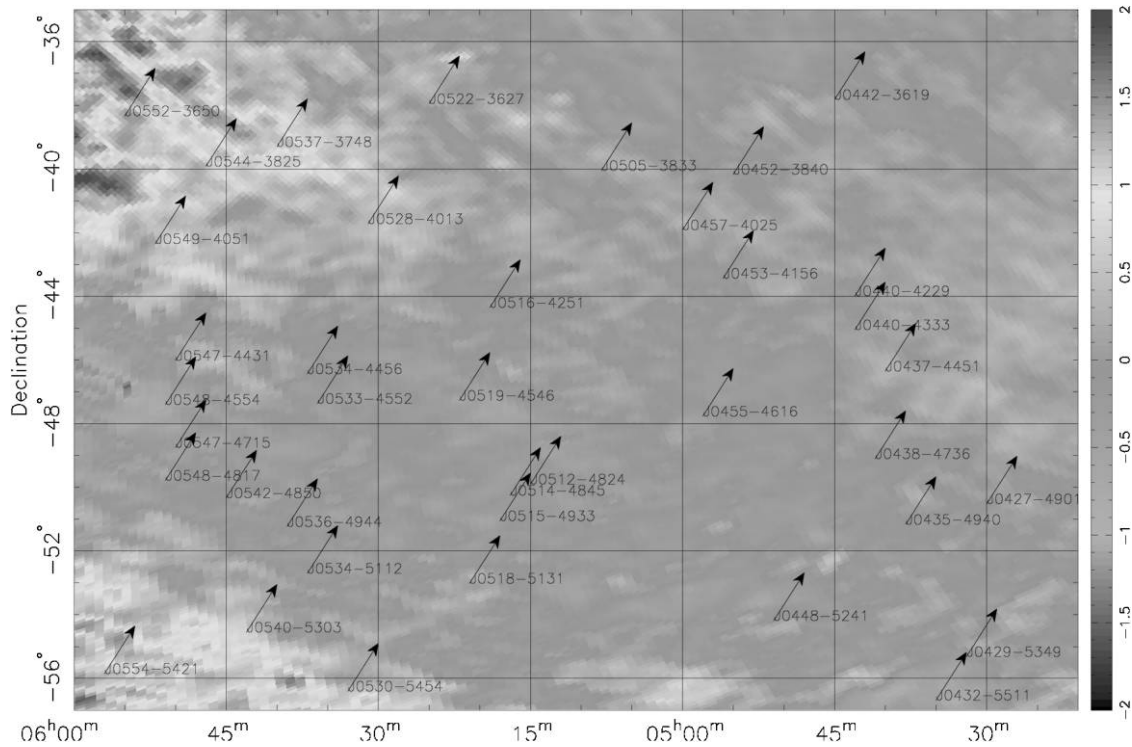


FIG. 8.—The Stokes Q image of a simulated sky at approximately 160 MHz, with a bandwidth of 40 kHz and an integration time of approximately 2 hr. The rms brightness of the image is significantly lower than the Stokes I image shown in Fig. 7. This image shows considerably less association with high Stokes I features, but does display considerable diffuse polarization. This is due to the side lobes of bright diffuse emission containing significant instrumental polarization. See the electronic edition of the *PASP* for a color version of this figure.

centered at 160.02 MHz was simulated. The full observation has not been simulated, as the computational demands are large and a representative integration is required to test the fidelity of the Stokes conversion. The 2 hr integration is sufficient to ensure that the dominant noise source in the simulated map is from side lobes. The images were produced by exactly the same pipeline that produced the real sky images, except that these images were generated off-line and not in real time. The FOV imaged was also slightly wider: closer to 30° . Pictoris A has been peeled out of the simulated images; due to perfect calibration, there are no residuals remaining.

Figures 1 and 2 contain significantly more point sources than the simulation (Fig. 7). These sources are actually in the catalog. However, the simulation is very conservative when including catalog sources and has only included those sources with flux measurements near the MWA observing band and that are predicted to be larger than 5 Jy. Nevertheless, this simulation gives a good indication of the performance of the system under more controlled conditions than the measured data set.

5.3.1. Polarization Properties

The rms brightness of the polarization observations is approximately 16% of the Stokes I observations; the point-source polarized residual is considerably less than this: 9% of J0522-

3627 at the very edge of the Q map and only 2% at the same point in the U map. The residual is also a function of position within the primary beam; near the center of the map, the residual in Q is nearer 2%, and in the U maps there is almost no measurable association between Stokes I sources and polarized structure in the beam center—except as polarized residuals from the peeling of Pictoris A. The polarized residuals are a feature of imperfect calibration of the individual antennas. The fact that the residuals deteriorate away from beam center lends further weight to this claim, as differences between the primary beams become more apparent further from the beam center.

The simulated point sources have no intrinsic polarization and, as the calibration is perfect, there should be no polarized point sources in the integrated maps, unlike the images of the true sky, which do display some polarized point-source signals, due to imperfect calibration of the primary beams. Figures 8 and 9 are featureless in this regard, with little evidence of the existence of any point-source features in the maps. The maps from simulation possess diffuse polarization, despite the fact that the input sky is unpolarized. This is due to the polarized side lobes of the Stokes I emission. The polarized brightness in a given pixel contains contributions from not only the polarized sky at that position, but from all the side lobes of the diffuse and point-source emission in the sky. The side lobes are generated

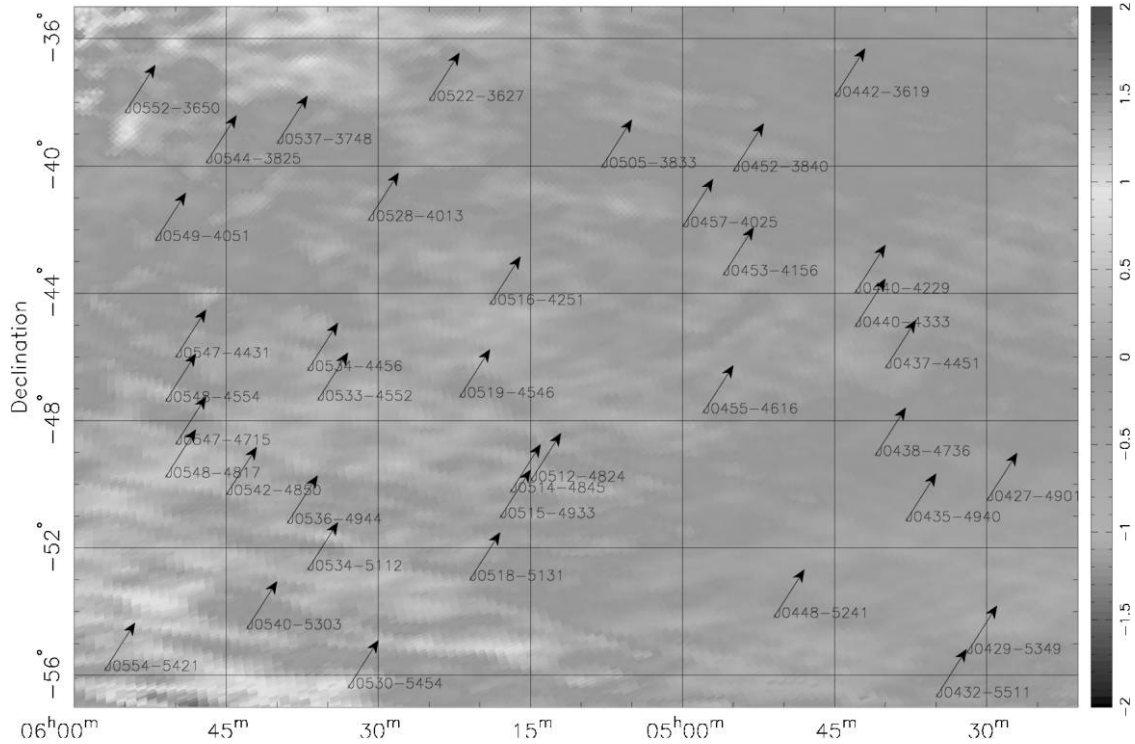


FIG. 9.—The Stokes U image of a simulated sky at a central frequency of approximately 160 MHz, with a bandwidth of 40 kHz and an integration time of approximately 2 hr. The rms brightness of the image is significantly lower than in the Stokes I image, shown in Fig. 7. This image shows considerably less association with high Stokes I features, but does display considerable diffuse polarization. This is due to the side lobes of bright diffuse emission containing significant instrumental polarization. See the electronic edition of the *PASP* for a color version of this figure.

with instrumental polarization commensurate with the position of the generating source, not the position at which the side lobe is measured. This response is modulated by the antenna gain and so sources outside the primary beam will be suppressed; however, bright point sources and patches of diffuse emission will still display polarized side lobes across the images; these will generate an apparent polarized signal where none exists. This produces an error in the Stokes conversion even if the array is perfectly calibrated. This is the prevailing polarized signal in images of the simulated sky and is, no doubt, also present in the images of the true sky, but the significant frequency synthesis over the 30 MHz bandwidth mitigates the effect, and therefore calibration errors dominate those images.

This effect will severely limit the polarization fidelity of integrated images unless an accurate model of the radio sky can be obtained. A significant project is underway with the MWA to obtain an initial model of the sky through an all-sky survey project that will be continually improved through the lifetime of the instrument.

6. CONCLUSIONS

We have demonstrated the performance of a real-time calibration and imaging pipeline for the MWA-32T. Wide-field polarimetric imaging has successfully been performed via the

integration of warped and weighted snapshots. The simulations indicate that the calibration and imaging pipeline works with sufficient fidelity to remove instrumental polarization from a simulated point-source population. But the observed point-source population displays residual polarization: at the $<10\%$ level in the worst case and at the $<5\%$ level in general; this is likely due to an imperfect calibration. Calibration precision is limited by low snapshot sensitivity and the fact that we cannot account for the differing primary-beam shapes of array elements. The simulations also indicate that a major component of measured polarized emission from this and other wide-FOV instruments will be due to the polarized side lobes of bright unpolarized features observed off-axis, necessitating the construction of comprehensive sky models applicable to these instruments. The removal of these off-axis sources, preferably in the visibility domain or via an iterative forward modeling scheme, is required to limit this instrumental polarization.

The MWA elements are crossed dipoles on a ground plane and electrically phased to a pointing direction, and they have considerable instrumental polarization (up to 100%). We have removed instrumental polarization down to $<5\%$ in real time, with limited sensitivity and a common scaled model for each array element response. The 512 tile MWA, with better constrained individual tile responses and much greater snapshot

sensitivity, should perform substantially better. The key improvement lies not simply in improving sensitivity but in accounting for the different direction-dependent polarimetric responses of the constituent antennas.

Subsequent publications in this series will aim to perform point-source subtraction from these images and others like them at different frequencies, in order to obtain spectral index information and to constrain the behavior of the MWA antenna beam.

This work uses data obtained from the Murchison Radio-Astronomy Observatory. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support

came from the U.S. National Science Foundation (grants AST-0457585 and PHY-0835713), the Australian Research Council (grants LE0775621 and LE0882938), the U.S. Air Force Office of Scientific Research (grant FA9550-0510247), the Smithsonian Astrophysical Observatory, the MIT School of Science, the Raman Research Institute, the Australian National University, the iVEC Petabyte Data Store, the Initiative in Innovative Computing and NVIDIA-sponsored Center for Excellence at Harvard, and the International Centre for Radio Astronomy Research, a Joint Venture of Curtin University of Technology and The University of Western Australia, funded by the Western Australian State government.

REFERENCES

- Bernardi, G. 2010, MNRAS, in press
- Bracewell, R. N. 1984, in *Indirect Imaging*, ed. J. A. Roberts (Cambridge: Cambridge Univ. Press), 177
- Brouw, W. N. 1971, Ph.D. thesis, Leiden Univ.
- Calabretta, M. R., & Greisen, E. W. 2002, A&A, 395, 1077
- Calabretta, M. R., & Roukema, B. F. 2007, MNRAS, 381, 865
- Clark, B. 1973, VLA Scientific, Memo 107 (Charlottesville:NRAO)
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2005, in ASP Conf. Ser. 347, *Astronomical Data Analysis Software and Systems XIV*, ed. P. Shopbell, M. Britton, & R. Ebert, (San Francisco: ASP), 86
- Cornwell, T. J., & Perley, R. A. 1992, A&A, 261, 353
- Erickson, W. C. 1984, J. Astrophys. Astr., 5, 55
- Erickson, W. C., Perley, R. A., Flatters, C., & Kassim, N. E. 2001, A&A, 366, 1071
- Gardner, F. F., & Whiteoak, J. B. 1966, ARA&A, 4, 245
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759
- Hamaker, J. P. 2000, A&AS, 143, 515
- . 2006, A&A, 456, 395
- Hamaker, J. P., & Bregman, J. D. 1996, A&AS, 117, 161
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, A&AS, 117, 137
- Haslam, C. G. T., Klein, U., Salter, C. J., Stoffel, H., Wilson, W. E., Cleary, M. N., Cooke, D. J., & Thomasson, P. 1981, A&A, 100, 209
- Kassim, N. E., Lazio, T. J. W., Erickson, W. C., Perley, R. A., Cotton, W. D., Greisen, E. W., Cohen, A. S., Hicks, B. et al. 2007, ApJS, 172, 686
- Lonsdale, C., Cappallo, R., Morales, M., Briggs, F., Benkevitch, L., Bowman, J., Bunton, J., Burns, S. et al. 2009, *Proceedings of the IEEE*, 97, 1497
- 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*, 2, 707
- Perley, R. A., & Cornwell, T. J. 1991, in IAU Colloq. 131, *Radio Interferometry. Theory, Techniques, and Applications*, (ASP Conf. Ser. 19; San Francisco: ASP) 184
- Pindor, B. 2010, PASA, in press
- Sault, R. J., Hamaker, J. P., & Bregman, J. D. 1996, A&AS, 117, 149
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2001, *Interferometry and Synthesis in Radio Astronomy* (2nd ed.; New York: Wiley)
- Wright, A. E., & Otrupcek, R. E. eds. 1990, *Parkes Catalogue* (Sydney: ATNF)