Publications of the Astronomical Society of Australia (PASA), Vol. 34, e026, 10 pages (2017).

doi:10.1017/pasa.2017.20

Observing Pulsars with a Phased Array Feed at the Parkes Telescope

X. Deng[1,2], A. P. Chippendale[1,4], G. Hobbs[1], S. Johnston[1], S. Dai[1], D. George[1], M. Kramer[2,3], R. Karuppusamy[2], M. Malenta[3], L. Spitler[2], T. Tzioumis[1] and G. Wieching[2]

[1]CSIRO Astronomy and Space Science, Australia Telescope National Facility, P.O. Box 76, Epping NSW 1710, Australia
[2]Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
[3]Jodrell Bank Centre for Astrophysics, The University of Manchester, Alan Turing Building, Manchester M13 9PL, UK
[4]Email: Aaron.Chippendale@csiro.au

(Received January 13, 2017; Accepted May 23, 2017)

Abstract

During 2016 February, CSIRO Astronomy and Space Science and the Max-Planck-Institute for Radio Astronomy installed, commissioned, and carried out science observations with a phased array feed receiver system on the 64-m diameter Parkes radio telescope. Here, we demonstrate that the phased array feed can be used for pulsar observations and we highlight some unique capabilities. We demonstrate that the pulse profiles obtained using the phased array feed can be calibrated and that multiple pulsars can be simultaneously observed. Significantly, we find that an intrinsic polarisation leakage of −31 dB can be achieved with a phased array feed beam offset from the centre of the field of view. We discuss the possibilities for using a phased array feed for future pulsar observations and for searching for fast radio bursts with the Parkes and Effelsberg telescopes.

Keywords: methods: observational, pulsars: individual (PSRs J0742–2822, J0835–4510, J1559–4338, J1644–4559, J0437–4715, J1740–3015, J1741–3016 and J1739–3023)

1 INTRODUCTION

The Parkes 64-m diameter telescope has been used for pulsar research since the discovery of pulsars. As the receivers and backend instruments have continued to be upgraded the telescope still produces high quality pulsar observations. The telescope has been used to discover more than half of the known pulsars. The majority of these pulsars were found using a 13-beam, 20-cm multibeam receiver (Staveley-Smith et al. 1996). Observations of known pulsars are usually carried out with the central feed of the multibeam receiver, or with a dual-band, single-pixel receiver. A new ultra-wideband single-pixel receiver is currently being designed that will provide continuous coverage from 0.7 to 4 GHz and is expected to be installed on the telescope during 2017 (Dunning et al. 2015; Manchester (for the PPTA Team) et al. 2013a).

The multibeam receiver has been used extensively for pulsar surveys over the last decade (Keith et al. 2010; Ng et al. 2015; Lorimer et al. 2015) and has been the main instrument for detecting fast radio bursts (FRBs) (Thornton et al. 2013; Champion et al. 2016) to date. The multibeam receiver is sensitive, but the current FRB searches are limited by the field of view (FoV) being observed. In 2016 February, a phased array feed (PAF) receiver was installed at Parkes that more than doubled the number of simultaneous beams on the sky compared with the multibeam system. This was the first long-term installation of a PAF on a high-gain, single-dish telescope with significant direct access by astronomers. The Commonwealth Scientific and Industrial Research Organisation (CSIRO) and the Max-Planck-Institute for Radio Astronomy (MPIfR) used this unique opportunity to collaborate on wide FoV projects relating to pulsars and transient sources with this novel receiver system. We report on our results relating to studies of known pulsars in this paper. The PAF was located in the focus cabin of the Parkes telescope from 2016 February to October. It was then removed for shipment to Germany where it will be installed in early 2017.

A PAF is a dense array of antenna elements at the focus of a reflector telescope and the output of these elements can be combined to form beams on the sky. The direction of these beams is controlled by varying the weighting of individual elements of the PAF. The PAF described here (Hampson et al. 2012; Hay & O'sullivan 2008) was designed for the Australian Square Kilometre Array Pathfinder (ASKAP) telescope (DeBoer et al. 2009; Schinckel & Bock 2016), but slightly modified for use on the Parkes and Effelsberg telescopes (Chippendale et al. 2016). At present, the PAFs on ASKAP are used in standard synthesis imaging mode, with

correlations between corresponding PAF beams on different antennas produced every 10 s. This is sufficient to see slow transients (Hobbs et al. 2016) but not sufficient to observe pulsars with high time resolution.

Pulsars have been observed in a traditional pulsar observing mode with PAFs on the Westerbork Synthesis Radio Telescope. Initial profiles for two simultaneously observed pulsars were given by van Cappellen, Bakker, & Oosterloo (2011). An earlier generation PAF has also been installed on a 12-m diameter test-bed antenna at Parkes. This system is currently being used to monitor PSR J0835–4510 (the Vela pulsar) and the results will be published elsewhere (Sarkissian et al. 2017).

In this paper, we (1) demonstrate that a PAF on a high-gain single dish telescope can be used for pulsar observations, (2) show that such observations can be calibrated, (3) show that three pulsars can be simultaneously observed for extended periods without rotating the receiver by updating the beam positions in real time, and (4) discuss the expected future use of PAFs for pulsar and transient astronomy.

2 THE PAF INSTALLED AT PARKES

The CSIRO PAF system (Hampson et al. 2012) was designed for the ASKAP telescope to carry out fast astronomical surveys with a wide FoV in a frequency band between 0.7 and 1.8 GHz. One of these PAFs was installed on the Parkes radio telescope to test and commission the PAF and the corresponding backend instrumentation (Chippendale et al. 2016), to study the use of PAFs for spectral line observations, to develop software for pulsar and transient observations, and to explore the use of a PAF for studying known pulsars.

Hampson et al. (2012) described the PAF and its associated digital processing hardware in detail. In brief, an array of 188 connected 'chequerboard' antenna elements (Hay & O'sullivan 2008) is distributed over approximately a 1.2-m diameter circle. It is a dual-polarisation receiver and each polarisation has 94 elements. The analogue signals from all elements, each of up to 600-MHz bandwidth, are transmitted to the digital receiver via RF-over-fibre links and sampled there by 12 'Dragonfly' digital receivers (Brown et al. 2014). The digital receivers also channelise the data to 1 MHz via a multistage oversampled filterbank (Tuthill et al. 2012). With 16 ports per receiver, this results in a 192 port digital system, with four spare ports beyond the 188 connected to the PAF. The digitised signals are processed by eight 'Redback' beamformers (Hampson et al. 2014) to form up to 36 dual-polarisation beams of 384-MHz bandwidth (48 MHz per beamformer) in 1-MHz frequency channels.

For ASKAP, the outputs of the beamformers from different antennas are correlated at a central site for calibration and synthesis imaging. At Parkes, we stream 336 MHz of the beamformed data (42 MHz per beamformer) at the full sampling rate into Graphics Processing Unit (GPU) nodes via Ethernet switches in 7-MHz frequency chunks. The beamformed data are scaled from 16 bit down to 8 bit on these

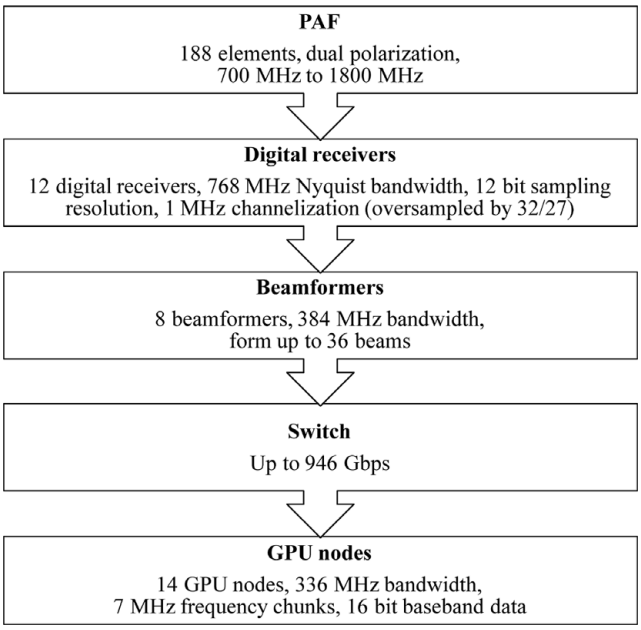

Figure 1. Schematic showing the data flow for the PAF installed at the Parkes radio telescope.

GPU nodes. We can then (1) record 8 bit baseband data on the GPU nodes and fold it offline at the period of a known pulsar with the DSPSR software package (van Straten & Bailes 2011) or (2) fold the streaming data in ring buffers using DSPSR in real time. Figure 1 summarises the data flow.

As there is more radio-frequency interference (RFI) at Parkes and Effelsberg than at the ASKAP site, we used narrower bandpass filters in the PAF to reject interference from mobile phones and lower frequency digital television services. The PAF system used at Parkes was optimised to cover the quieter 1.2 to 1.74 GHz band with two frequency bands covering 1.2 to 1.48 GHz and 1.34z to 1.74 GHz (Chippendale et al. 2016). In these modified bands, the system was able to operate at Parkes with the same attenuator settings and analogue-to-digital converter input levels used for ASKAP at the radio-quiet Murchison Radio-astronomy Observatory (MRO). A further two observing bands are available with unmodified ASKAP filters covering 0.6 to 0.7 GHz and 0.7 to 1.2 GHz. The RFI in these unmodified bands required at least 10 dB more attenuation than at the MRO and varied significantly with the orientation of the 64-m antenna and PAF relative to local transmitters. In this paper, we only discuss observations in the 1.2 to 1.48 GHz band, which is relatively quiet with the exception of satellite navigation signals that come and go as satellites overfly the observatory.

2.1. Dedispersing oversampled time series

In this work, we applied offline, incoherent dedispersion via PSRCHIVE (Hotan, van Straten, & Manchester 2004) on data recorded from the oversampled polyphase filterbank designed for ASKAP (Tuthill et al. 2012). The incoming data

from each 1-MHz sub-band of the filterbank are folded at the period of the pulsar and time averaged online by the DSPSR package as described above.

Aliased signals from adjacent filterbank sub-bands distort integrated full-band pulse profiles by adding attenuated and time-shifted copies of the correct pulse profile (van Straten & Bailes 2011). The attenuation is approximately the mean level of adjacent-channel leakage and the time shift is approximately the dispersion delay between adjacent channels. Figure 4 in van Straten & Bailes (2011) shows a −12 dB feature at a time shift of 100 μs for a critically sampled FFT filterbank[1] with 500 kHz channel width and a dispersion delay between channels of 88 to 154 μs over the observing band for a pulsar with $DM = 71$ cm^{-3} pc. The −12 dB level of the time-shifted copy in (van Straten & Bailes 2011) is very close to the value of −12.3 dB given by averaging the spectral leakage calculated for an FFT with rectangular window over one adjacent channel.

ASKAP's oversampled coarse filterbank provides a flatter overall passband response and reduces aliasing in its 1-MHz sub-bands (Tuthill et al. 2012, 2015). Calculating the average spectral leakage over one adjacent 1-MHz sub-band of the ASKAP filterbank suggests we would expect a time-shifted copy of the pulse profile at the level of −27 dB in the averaged pulse profiles presented in this paper. For an observation of pulsar with $DM = 71$ cm^{-3} pc in the 1 200 to 1 480 MHz MPIfR PAF band range we would expect this additive distortion to appear at a time shift equal to the inter-channel dispersion delay of 182 to 341 μs.

In fact, near artefact-free dedispersion of pulsar measurements made with the oversampled ASKAP filterbank should be possible if the 1-MHz channels are further channelised. This would allow the unwanted band edges of the oversampled 1-MHz sub-bands to be discarded before coherent dedispersion and frequency averaging of the pulse profile over the full observing band. This can be accomplished by channelising each 1 MHz sub-band into N fine channels then, for an oversampling ratio of $\rho = 32/27$, discarding $N(1 - \rho^{-1})/2$ fine channels from both ends of each 1 MHz sub-band before further processing. By averaging the spectral leakage from one adjacent sub-band over only the retained fine channels, we expect the erroneous copy of the pulse profile to appear at −74 dB.

3 RESULTS

Throughout the PAF commissioning period, we observed PSRs J0437–4715, J0742–2822, J0835–4510, J1559–4338, and J1644–4559. We also identified a non-globular-cluster sky region in which multiple pulsars, with flux density greater than 0.5 mJy, are in the PAF FoV.

First, we tested the entire observing system with the bright pulsars, PSRs J0742–2822, J0835–4510, J1559–4338, and J1644–4559. This allowed us to test the beam positions and the correct ordering of the frequency channels in the resulting data files. We then observed the bright millisecond pulsar PSR J0437–4715 to study the achievable timing precision and calibration method. Finally, we simultaneously observed three pulsars to test the timing quality including beam updates to track sources drifting through the FoV.

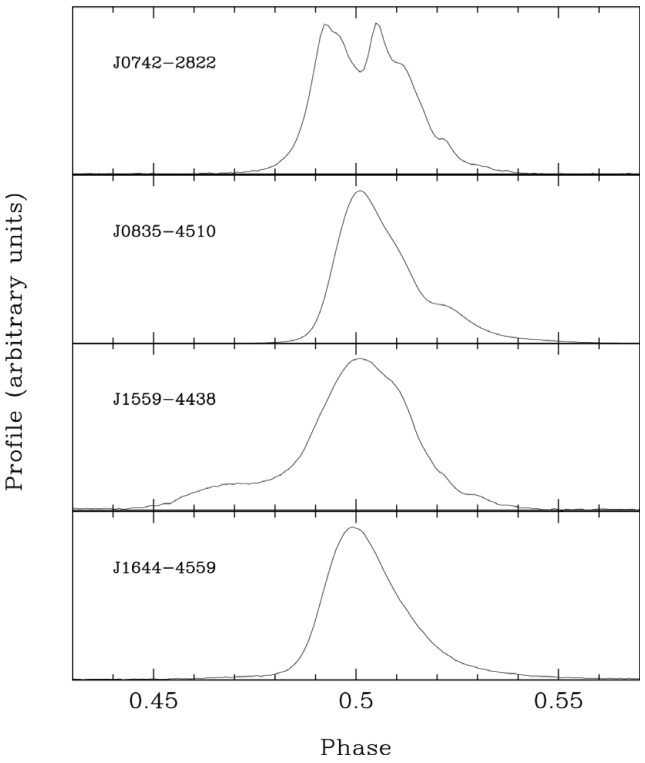

Figure 2. Total intensity (Stokes I) profiles of four pulsars we measured as a test of the entire PAF observing system. The integration times of these observations are 1 h, 10 min, 40 min, and 10 min, respectively. We centred these profiles and zoomed them to the pulse-phase range from 0.43 to 0.57.

For all of this work, we used maximum signal-to-noise ratio (SNR) beam weights calculated via the same method used for performance evaluation of the ASKAP Boolardy Engineering Test Array (BETA, McConnell et al. 2016). We applied this beam weights calculation algorithm independently for each set of 94 PAF elements with common polarisation. Thus, as with the ASKAP BETA work, we only weighted 94 elements with matching polarisation into each beam. This makes the beam polarisations line up with the native PAF element polarisations, which are inherently linear with low and stable leakages (Sault 2014). This also provides a clear link between the beam polarisation and the physical orientation of the PAF, which is convenient during commissioning. Unlike the BETA work, we used strong extragalactic sources such as Virgo A for beam weight calculation instead of the Sun which is too extended for the narrower beam of the 64-m Parkes telescope.

Figure 2 shows the uncalibrated total intensity profiles for PSRs J0742–2822, J0835–4510, J1559–4438, and J1644–4559. We compared them with the profiles observed with

[1] Referred to as a 'deprecated' filterbank by van Straten & Bailes (2011).

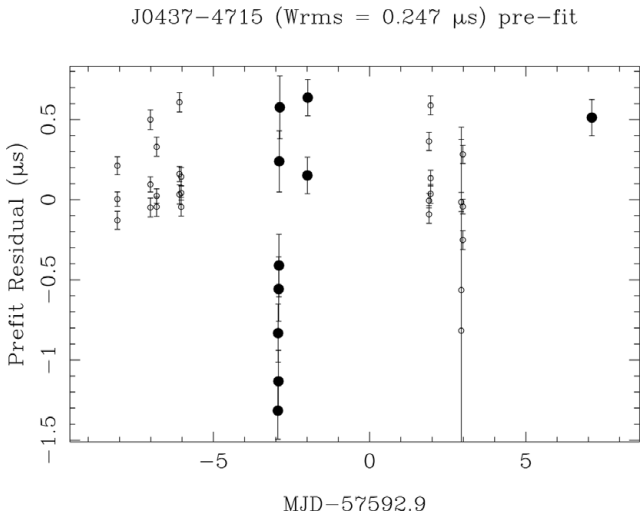

Figure 3. Timing residuals of PSR J0437–4715. The open and solid circle symbols are for the timing residuals obtained with a traditional receiver and the PAF system respectively.

Parkes project P574 (see Weltevrede & Johnston 2008) in 2014 January and the maximum relative deviation is less than 10% for the latter three of the four pulsars[2].

Our observed profile for PSR J0742–2822 is slightly different from that given in the archive. This pulsar has a profile, which is variable on short timescales (see Keith, Shannon, & Johnston 2013) and observations taken with a traditional receiver only 1 d after the observation obtained with the PAF show the same structure in the profile. We therefore conclude that the PAF functions well in Stokes I and will discuss polarisation calibration in the next section.

In order to demonstrate that pulse arrival times can be measured over days, we made multiple observations of the brightest millisecond pulsar, PSR J0437–4715. This pulsar can be timed with an rms timing precision of 100 ns over years using existing receiver systems at the Parkes telescope (e.g., Manchester et al. 2013b). We used the PAF to make multiple observations of this pulsar on 2016 July 21 and then follow-up observations on 2016 July 22 and 2016 July 31. In Figure 3, we compare the timing residuals obtained using the PAF (10 min integrations, solid circle symbols) with those being obtained for the Parkes Pulsar Timing Array (PPTA) project using the traditional instrumentation (1 h integrations, open circle symbols, see Manchester et al. 2013b for detail). We note that the timing residuals from the PAF are consistent with those from the PPTA data.

We did not apply any polarisation calibration to the PAF data in Figure 3 and therefore expect to see increased scatter. The scatter in the PAF timing is approximately 2 μs, but we show in the next section that polarisation calibration can reduce this by more than a factor of two. For comparison, the jitter noise of PSR J0437–4715 with 10 min integration time is approximately 100 ns (Shannon et al. 2014) and the PPTA time of arrival (ToA) uncertainties are approximately 100 ns.

A simple calculation based on the radiometer equation suggests that the PAF ToA uncertainties, shown by the error bars in Figure 3, should be ∼5.5 times larger due to the increased system temperature and slightly wider bandwidth of the PAF compared to the multibeam receiver. This is consistent with the uncertainties in our observed arrival times.

3.1. Calibrating the pulsar data

Pulsars are often highly polarised. In order to form stable pulse profiles for high-precision timing, to study the emission mechanism, or to analyse the scintillation properties of a pulsar, it is necessary to carry out polarisation and/or flux calibration. The polarisation properties of a signal are fully described using the four Stokes parameters: I, Q, U, and V. The measured Stokes parameters will have been affected by the observing instrument and telescope. Calibrating traditional, single-pixel receivers is non-trivial and often relies on the injection of a pulsed-calibration signal (e.g., van Straten 2004). The PAF does not include such a noise calibration system. Instead a broad-band radiator was installed on the telescope surface and this has been used to study methods for forming the PAF beams and making them stable with time.

In the work described here, we solve for the instrumental polarisation parameters of a PAF beam that minimise the mean square difference between new PAF measurements of the polarisation profile of PSR J0437–4715 and a corresponding archived reference measurement that is known to be accurately calibrated. Our approach is very similar to that of Johnston (2002) and we use a reference measurement from Dai et al. (2015) that was made using the central feed of the 20-cm multibeam receiver. The multibeam receiver is carefully calibrated and in regular use for precise measurements of pulsars.

The relationship between the Stokes parameters of the incident wave and beamformed voltages ($\mathbf{S}_i$ and $\mathbf{S}_o$, respectively) can be modelled as (Warnick et al. 2012)

$$\mathbf{S}_i = \mathbf{M}^{-1}\mathbf{S}_o, \tag{1}$$

where $\mathbf{S} = [I, Q, U, V]^T$ is a column vector of Stokes parameters, $\mathbf{M} = \mathbf{T}\,(\mathbf{J} \otimes \mathbf{J}^*)\,\mathbf{T}^{-1}$ is the Mueller matrix of the PAF system and $\mathbf{T}$ is the transformation matrix between the Stokes parameters and the coherency vector

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 1 & 0 \\ 0 & -i & i & 0 \end{bmatrix}. \tag{2}$$

$\mathbf{J} = \mathbf{J}_{\rm cal}\mathbf{J}_\psi$ is the Jones matrix of the combined PAF and 64-m antenna that comprises a factor due to the rotation of the feed by the parallactic angle, ψ, as the altazimuth mounted

[2] To assess pulse profile similarity in phase with PSRCHIVE for each pulsar, we normalised the aligned profiles and got the deviation between them. We repeated that for the latter three pulsars and the maximum relative deviation is the maximum deviation given by the whole procedure.

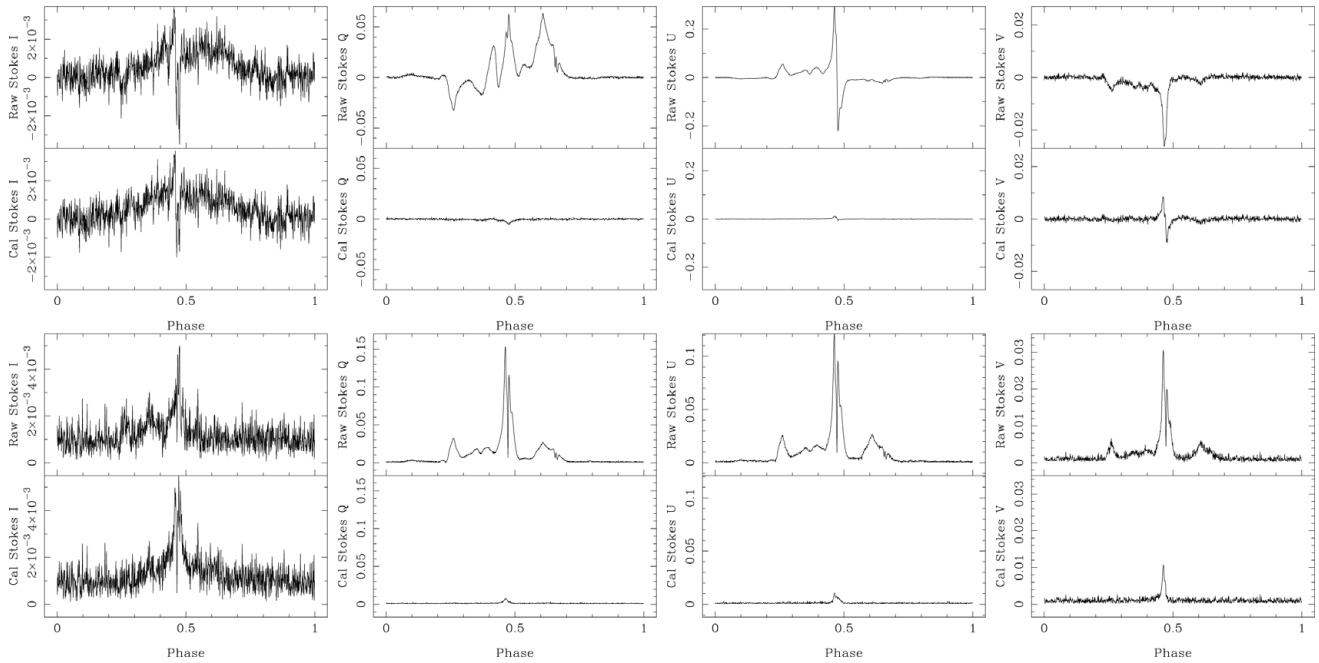

Figure 4. Error in Stokes polarisation parameters of PSR J0437–4715 for PAF observations with respect to PPTA reference observations (with Stokes *I*, *Q*, *U*, and *V* from left to right). The figures on the top row are for the mean of the difference (between these four observed profiles and the reference template) and the figures on the bottom row are for the standard deviation of the difference. The upper panel of each figure shows the uncalibrated result and the bottom panel of each figure represents the calibrated result.

antenna tracks the source

$$\mathbf{J}_{\psi} = \begin{bmatrix} \cos\psi & \sin\psi \\ -\sin\psi & \cos\psi \end{bmatrix} \tag{3}$$

and a generalised Jones matrix that represents the instrumental polarisation of the PAF

$$\mathbf{J}_{\rm cal} = \begin{bmatrix} A & B \\ C & D \end{bmatrix}. \tag{4}$$

Here, B, C, and D are complex, but we force A to be real to account for the fact that this technique cannot calibrate a phase or delay difference between the PAF and the receiver used to make the archived measurement. $\mathbf{J}^*$ is the complex conjugate of $\mathbf{J}$ and $\otimes$ is the Kronecker product operator.

The calibration problem can be expressed as

$$\underset{\mathbf{J}_{\rm cal}}{\arg\,\min} \left\langle |\mathbf{S}_i - \mathbf{S}_{\rm ref}|^2 \right\rangle, \tag{5}$$

where $\mathbf{S}_{\rm ref}$ are the 'true' Stokes parameters provided by the archived measurement and $\langle\cdot\rangle$ denotes averaging over pulse-phase bins and PAF observation epochs. This equation can be uniquely solved if we include measured profiles of the pulsar from multiple observation epochs with different parallactic angles.

We followed these steps to calibrate the data:

1. The weighted profile of PSR J0437–4715 in the 20-cm observing band was downloaded from the PPTA pulsar profile collection (Dai et al. 2015)[3] of the Parkes Observatory Pulsar Data Archive (Hobbs et al. 2011).
2. The band between 1 304 and 1 465 MHz was extracted to avoid RFI contamination from satellite navigation services at lower frequencies and to avoid exceeding the PAF observing band edge at higher frequencies. This extracted profile was used as our polarisation template[4].
3. The Stokes parameters of this reference template were normalised by its peak Stokes I value.
4. The PAF was used to observe PSR J0437–4715 four times on 2016 June 6 with each observation lasting 1 h[5].
5. The PAF data were binned to 10 min sub-integrations and the data were summed in frequency between 1 304 and 1 465 MHz.
6. The Stokes parameters of each 10 min PAF measurement were normalised to its peak Stokes I value.
7. The observed profiles were aligned in pulse phase with the reference template[6].

[3] Available via the CSIRO Data Access Portal at http://doi.org/10.4225/08/54F3990BDF3F1.

[4] This profile was based on the multibeam receiver and calibrated using the PCM calibration method within the PSRCHIVE pulsar data processing software package (Hotan et al. 2004).

[5] To confirm that the calibration works for a pulsar that is not directly in the pointing direction of the telescope, we used a beam with 0.215° offsets in both elevation and cross-elevation. We tracked PSR J0437–4715 with complementary offsets in elevation and cross-elevation to keep this pulsar in the offset beam.

[6] The alignment shift was determined by cross-correlating the Stokes I profile from the PAF observation with that of the reference. The same alignment shift was applied to all Stokes parameters of the measured profile.

8. The minimisation problem of Equation (5) was solved to calculate $\mathbf{J}_{\rm cal}$ using all pulse-phase bins for all of these 10 min PAF profiles.
9. The 10 min PAF profiles were calibrated using Equation (1).

Figure 4 shows the difference between the observed profiles[7] and template (with and without calibration). The uncalibrated Stokes I profile is very close to the true value and that the uncalibrated Stokes Q, U, and V profiles are significantly deviant. Post-calibration, the standard deviation of the differences between the observed profiles and the reference profile, for all Stokes parameters, is better than 1% of the peak amplitude of Stokes I. We note that the profiles are significantly improved after calibration, but they are not perfect. Such imperfection could arise from intrinsic pulse shape changes since the PPTA measurement, imperfect calibration, or ionospheric effects.

We tested for possible Faraday rotation in the ionosphere by repeating the calibration and simultaneously solving for an independent rotation measure, RM, at each 10 min observation interval. We forced RM = 0 at the first 10 min interval so this defined a rotation of the received polarisation plane by $\mathrm{RM}\lambda^2$ at each interval with respect to the first (where λ is the wavelength for each observation channel). The fit suggests there was a change of RM of approximately 1 rad m^{-2} over the first 3 h of the observation. We explored the sensibility of this fit by using RMEXTRACT[8] to calculate expected ionospheric RM from Global Positioning System measurements of total electron content, and the World Magnetic Model[9]. This also suggested a 1 rad m^{-2} change in RM, but over the fourth to ninth hours of the observation and so our results are not conclusive, but changes in RM of this size are not unexpected.

We formed an analytic, noise-free standard template for PSR J0437–4715 using the PSRCHIVE package PAAS and determined pulse arrival times using PAT for the uncalibrated and calibrated 10 min profiles. Figure 5 represents the resulting timing residuals with respect to the timing model we used for Figure 3. The polarisation calibration reduces the weighted rms of the timing residuals from 0.653 to 0.236 μs, an improvement of more than a factor of two.

3.2. Simultaneous observations of multiple pulsars

As noted above, the primary advantage of using a PAF, instead of a single-pixel receiver, is that the PAF provides a wide FoV. We used the Australia Telescope National Facility (ATNF) pulsar catalogue (Manchester et al. 2005) to select regions of the sky in which multiple pulsars can be observed in a single observation with the PAF at Parkes. We removed regions that contained globular clusters as pulsars in those clusters could be observed with a single pixel receiver and therefore using the PAF is not advantageous. We then selected a region in which three pulsars could easily be detectable with the PAF.

Table 1. Simultaneously observed pulsars with the PAF system.

PSR name	RAJ2000 (hh:mm:ss.s)	DECJ2000 (+dd:mm:ss)	DM (cm^{-3}pc)	RD (°)	S1400 (mJy)
J1740–3015	17:40:33.8	−30:15:43	152.15	0.00	6.40
J1741–3016	17:41:07.0	−30:16:31	382.00	0.12	2.30
J1739–3023	17:39:39.8	−30:23:12	170.00	0.23	1.00

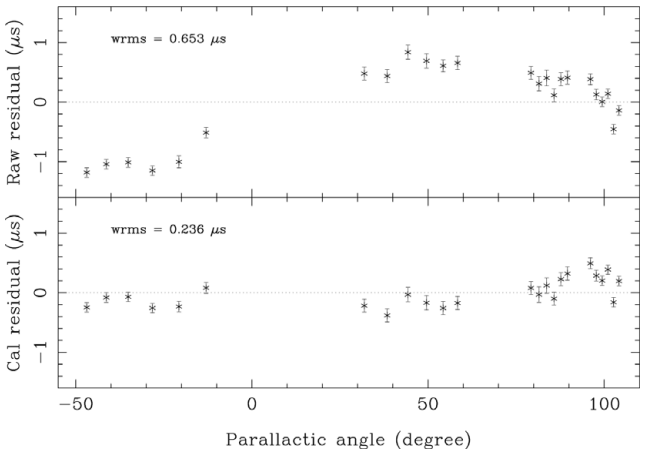

Figure 5. Timing residuals of PSR J0437–4715. The upper and lower panels represent the timing residuals before and after polarisation calibration.

The pulsars in this region are listed in Table 1 and their positions in the FoV are indicated in Figure 6. For each pulsar, the table contains the pulsar name, its right ascension and declination, dispersion measure (DM), the radial distance (in degrees) from the central pulsar position and the flux density in the 20-cm observing band. Figure 6 shows that we could have observed five pulsars simultaneously in the PAF FoV, but we were restricted to three due to the limited backend configurations available during commissioning. We were also limited to observe with 112-MHz bandwidth for each of the three pulsars but should be able to achieve 336-MHz bandwidth per pulsar in the future.

On 2016 September 22 from UTC 06:41, we observed the three pulsars simultaneously. We tracked PSR J1740–3015 and so this pulsar remained in the centre of the FoV throughout the observation. However, as the Parkes telescope has an altazimuth mount and we are unable to rotate the PAF, the positions of the other two pulsars in the FoV changed throughout the observation. They will drift out of non-central beams if these beams are fixed. Figure 7 shows the effective observing time of PSR J1739–3023 in an outer beam as a function of the hour angle. The pulsar drifts quickly out of the beam near transit (and also the observing duration becomes small when the source is close to setting).

[7] These calibrated 10 min profiles were normalised (as at Step 6) and aligned (as at Step 7). We formed four calibrated 1 h profiles by adding calibrated 10 min profiles in groups of six.
[8] Available at https://github.com/maaijke/RMextract.
[9] See https://www.ngdc.noaa.gov/geomag/WMM/DoDWMM.shtml.

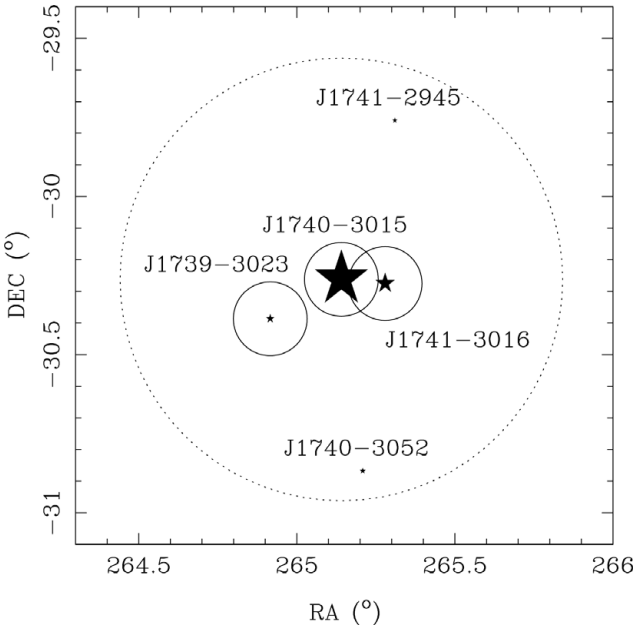

Figure 6. Region around PSR J1740–3015. Pulsars are indicated with star symbols and the names of these pulsars are given in the figure. The size of each star symbol indicates the flux density of that pulsar (the larger the symbol, the higher the flux density). Pulsars surrounded by 7 arcmin radius solid circles were observed simultaneously. The dotted circle gives an approximate indication of the sky region that could be observed using all 36 beams and shows that the PAF could observe five pulsars simultaneously in this field with appropriate backend configuration.

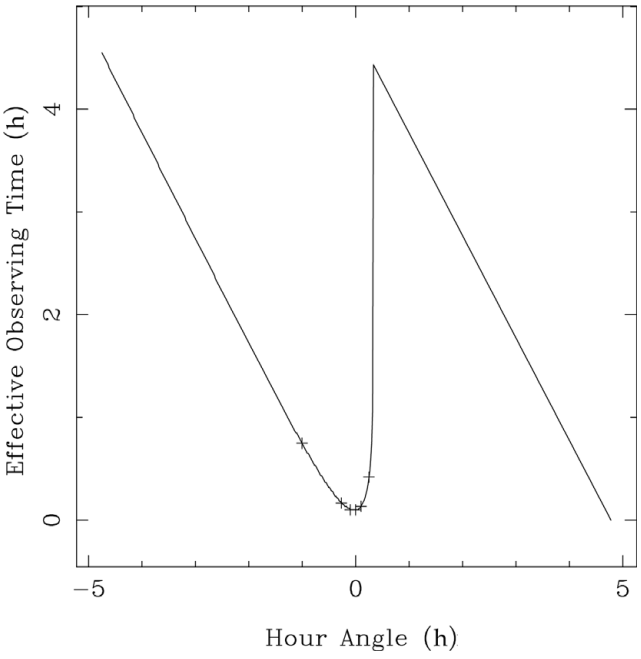

Figure 7. Effective observing time of PSR J1741–3016 with an outer beam when the central beam is tracking PSR J1739–3023. We commence observing PSR J1739–3023 with the peak of an outer beam and record the end of effective observing time when the pulsar crosses the half-power point of that beam or sets below the elevation limit of the telescope. We calculate the effective observing time for different hour angles from rise to set at 1-min intervals. Plus symbols in the figure indicate the hour angles at which we uploaded beam weights.

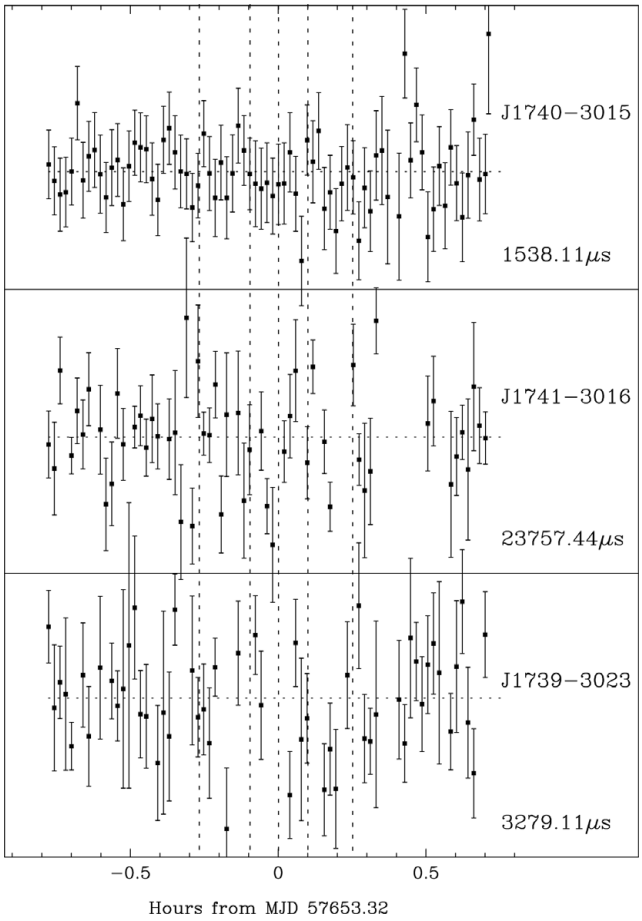

Figure 8. Long-duration simultaneous pulsar observation with beam tracking. Timing residuals for different pulsars are shown in different panels and the range from the minimum to the maximum residual are 1.5, 23.8, and 3.3 ms, respectively. Vertical dashed lines indicate when beamformer weights were updated.

We formed beams pointing at PSRs J1739–3023 and J1741–3016 at the beginning and updated the beam positions by uploading new beam weights when the pulsar would have drifted through the beam in order to keep tracking each source (the hour angles at which we uploaded beam weights are shown with plus symbols in Figure 7). We folded each 10 s of data for each pulsar online to form a folded profile. We finally averaged six adjacent profiles together to form a high S/N profile every minute. Our observation lasted approximately 1.5 h.

We formed an analytic, noise-free standard template for each pulsar using the PSRCHIVE package PAAS and determined pulse arrival times using PAT. Figure 8 shows the resulting timing residuals with respect to the timing model for each pulsar in the ATNF pulsar catalogue (Manchester et al. 2005), obtained using the timing package TEMPO2 (Hobbs, Edwards, & Manchester 2006; Edwards, Hobbs, & Manchester 2006). The vertical lines in Figure 8 indicate the times we updated beam positions. We clearly see that we can simultaneously observe three pulsars and their ToA uncertainties

do not significantly change for the pulsars in the outer beams throughout the observation.

4 DISCUSSION

4.1. Timing pulsars precisely with PAFs

Although we are far from recommending that the current generation of PAFs be used for high precision pulsar timing, we note that the PAF design has no intrinsic hardware or other technical issues that would prohibit high precision timing. For instance, some of the most sensitive millisecond pulsar timing data sets are currently being achieved by the PPTA project (see, e.g., Manchester et al. 2013b for an overview and Shannon et al. 2014 for an application of such data sets to gravitational wave searches). The most precise data sets in the PPTA are obtained with an incoherent dedispersion system similar to that used with the PAF in this paper. The various backend instruments have timing offsets that can vary during upgrades, but it is now possible to measure such time offsets either by comparing different backend receivers or by observing a known stable source.

The most commonly used PPTA receiver is the central beam of the Parkes 13-beam receiver (Staveley-Smith et al. 1996). However, it was not specifically designed for high-precision pulsar timing. In particular, the polarisation purity is poor and van Straten (2004) clearly identifies significant cross-coupling effects within the receiver. None of this affects the achievable timing precision assuming that the data are properly calibrated. We therefore highlight that high precision pulsar timing does not require excellent polarisation purity before calibration.

The polarisation purity of −40 dB required for precision timing as reported by Cordes et al. (2004) refers to a post-calibration result and not the intrinsic purity of the receiver (see also Foster et al. 2015). Such calibration is carried out for the PPTA project using observations of PSR J0437–4715 from rise-to-set (e.g., van Straten 2004). The calibration observations described in Section 3.1 are restricted in number and parallactic-angle coverage, yet we were able to use a single Jones matrix $\mathbf{J}_{\rm cal}$ to describe the PAF instrumental polarisation over a 161 MHz band and four 1 h observations spread over 6 h. This indicates good instrumental stability. Analysing the resulting Jones matrix $\mathbf{J}_{\rm cal}$ shows that the *intrinsic polarisation leakage*[10], for a PAF beam offset by 0.215° in both elevation and cross-elevation, is −31 dB. We note that Foster et al. (2015) found that only a limited fractional improvement in pulsar timing capabilities is achieved by pushing a feed's intrinsic polarisation leakage below −30 dB.

In summary, we have successfully timed a millisecond pulsar despite the relatively high system temperature of the current PAF. With sufficient observations, all the techniques used for the PPTA could be applied to improve the precision of the PAF observations.

Table 2. Efficiency of pulsar observation with PAF.

Telescope/project	Pulsars	PAF pointings	Reduction (%)
Parkes	2 018	1 365	32.4
Effelsberg	1 404	1 011	28.0
P574	321	260	19.0

4.2. Future opportunities

4.2.1. Speeding up pulsar monitoring

Current long-term monitoring of pulsars at Parkes uses a single pixel receiver system with a 14 arcmin full-width-at-half-maximum beamwidth. Here, we calculate the reduction in pointings that would be obtained if a PAF was used for these observations. We first consider all the pulsars that can be observed by Parkes (i.e., declination <20°) and are not in globular clusters. The total number of pulsars and the number of pointings (to cover all of these pulsars if we observe with a PAF) are listed in Table 2. We then repeat this calculation for the Effelsberg telescope and assume a declination limit of $>-30°$. Finally, we repeat the calculation for the 321 pulsars that are currently being observed for the young-pulsar timing project at Parkes (P574, Kerr et al. 2014). We can see from Table 2 that a PAF could reduce pointings by about 32% at Parkes, by 28% at Effelsberg, and by 19% for the on-going young pulsar timing project at Parkes. Further work is required to translate the reduction in the number of pointings into a saving in observing time given the difference between the system temperature and bandwidths of the various observing systems. However, for the P574 project, the observing time is dictated by the time taken for the pulsar profile to become stable (i.e., several thousand rotations) rather than to reach a certain S/N threshold. Hence for this project, the reduction in the number of pointings translates directly to an overall saving in observing time.

Millisecond pulsars are observed for high-time precision experiments. For instance, pulsar timing array projects (e.g., Manchester et al. 2013b; Verbiest et al. 2016; Babak et al. 2016) observe these pulsars with the primary goal of detecting ultra-low-frequency gravitational waves. Such a detection would require the unambiguous signature of a quadrupolar angular correlation in the pulsar timing residuals. Tiburzi et al. (2016) demonstrated that without sufficient sky coverage it is possible to falsely detect a gravitational wave background signal. They showed that it was necessary to observe pulsars that were close together (as well as pulsars separated far apart).

For the PPTA project, the closest pulsar pair is PSR J2129–5721 and PSR J2241–5236. Even with the PAF, these pulsars

[10] Foster et al. (2015) defines intrinsic polarisation leakage as $-10\log_{10}$ (IXR) where IXR is the intrinsic cross-polarisation ratio defined by Carozzi & Woan (2011). These metrics are defined independently of the coordinate system and relate closely to polarimetric errors after full calibration.

are too widely spaced for simultaneous observations. However, for the International Pulsar Timing Array (IPTA), two pulsars (PSRs J1910+1256 and J1911+1347) are separated by only 0.8° (Verbiest et al. 2016). Not only would the use of a PAF to observe these pulsars reduce the required observation time, it would also enable simultaneous measurements of instrumental offsets with two precisely timed pulsars.

4.2.2. Detecting and localising FRBs

One of the main science goals for the PAF on Effelsberg is the discovery of FRBs. Currently, Effelsberg is equipped with a 7-beam multibeam receiver and consequently the PAF with its 36 beams yields a five-fold increase in sky coverage for FRB searches. Current FRB results (Ravi et al. 2016) seem to indicate that sky area is more important than sensitivity for FRB detection.

In addition to the increased sky coverage, the PAF fully samples the focal-region field unlike the multibeam which is highly undersampled. This implies that the PAF can measure FRB positions with sub-arcminute accuracy, unlike the multibeam that yields positions with errors greater than 10 arcmin. Determining accurate positions is crucial for understanding the nature of FRBs (e.g., Keane & SUPERB Collaboration 2016).

5 CONCLUSION

We have demonstrated that pulsars can be observed using a PAF on a high-gain, single-dish telescope. We can produce calibrated profiles, form timing residuals without any unexpected instrumental offsets, and observe multiple pulsars simultaneously whilst accounting for pulsars drifting through the FoV.

ACKNOWLEDGEMENTS

The Parkes radio telescope is part of the ATNF, which is funded by the Commonwealth of Australia for operation as a National Facility managed by the CSIRO. The MPIfR financed the PAF discussed in this paper and its modification to cope with a less radio-quiet site. Installation and operation of the PAF at Parkes has benefited from the contributions of R. J. Beresford, M. Leach, M. Marquarding, S. Broadhurst, D. Craig, J. Crocker, R. Kaletsch, B. Preisig, K. Reeves, Dr J. E. Reynolds, T. Ruckley, M. Smith, E. Troup, and Dr J. Tuthill. Exploration of ionospheric Faraday rotation was suggested by Dr S. Ord and supported by Dr E. Lenc. The calibration work benefits from discussions with Dr K. J. Lee.

References

Babak, S., et al. 2016, MNRAS, 455, 1665
Brown, A. J., et al. 2014, in 2014 International Conference on Electromagnetics in Advanced Applications (Piscataway: IEEE), 268
Carozzi, T., & Woan, G. 2011, ITAP, 59, 2058
Champion, D. J., et al. 2016, MNRAS, 460, L30
Chippendale, A. P., Beresford, R. J., Deng, X., Leach, M., Reynolds, J. E., Kramer, M., & Tzioumis, T. 2016, in 2016 International Conference on Electromagnetics in Advanced Applications (Piscataway: IEEE), 909
Cordes, J., Kramer, M., Lazio, T., Stappers, B., Backer, D., & Johnston, S. 2004, NewAR, 48, 1413
Dai, S., et al. 2015, MNRAS, 449, 3223
DeBoer, D. R., et al. 2009, Proceedings of the IEEE, 97, 1507
Dunning, A., Bowen, M., Bourne, M., Hayman, D., & Smith, S. L. 2015, in IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communications (Piscataway: IEEE), 787
Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549
Foster, G., Karastergiou, A., Paulin, R., Carozzi, T., Johnston, S., & van Straten, W. 2015, MNRAS, 453, 1489
Hampson, G., et al. 2012, in 2012 International Conference on Electromagnetics in Advanced Applications (Torino: COREP), 807
Hampson, G. A., Brown, A., Bunton, J. D., Neuhold, S., Chekkala, R., Bateman, T., & Tuthill, J. 2014, in 2014 XXXIth URSI General Assembly and Scientific Symposium (URSI GASS), 1
Hay, S., & O'sullivan, J. 2008, RaSc, 43, RS6S04
Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, MNRAS, 369, 655
Hobbs, G., et al. 2011, PASA, 28, 202
Hobbs, G., et al. 2016, MNRAS, 456, 3948
Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, PASA, 21, 302
Johnston, S. 2002, PASA, 19, 277
Keane, E. F., SUPERB Collaboration 2016, preprint, arXiv:1602.05165
Keith, M. J., et al. 2010, MNRAS, 409, 619
Keith, M. J., Shannon, R. M., & Johnston, S. 2013, MNRAS, 432, 3080
Kerr, M., et al. 2014, Young Pulsar Timing and the Fermi Mission, ATNF Proposal P574
Lorimer, D. R., et al. 2015, MNRAS, 450, 2185
Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993
Manchester, R. N., (for the PPTA Team), Caretti, E., Norris, R. P., & Phillips, C. J. 2013a, ATNF Technical Memo 40.3.2/002, Development of an Ultra-Wideband (UWL) Receiver System at Parkes
Manchester, R. N., et al. 2013b, PASA, 30, 17
McConnell, D., et al. 2016, PASA, 33, 42
Ng, C., et al. 2015, MNRAS, 450, 2922
Ravi, V., et al. 2016, Science, 354, 1249
Sarkissian, J., Reynolds, J., Hobbs, G., & Harvey-Smith, L. 2017, preprint, arXiv:1705.08355
Sault, R. J. 2014, ACES MEMO 2, Initial characterisation of BETA polarimetric response
Schinckel, A. E. T., & Bock, D. C.-J. 2016, in Proc. SPIE, Vol. 9906, Ground-based and Airborne Telescopes VI, ed. H. J. Hall, R. Gilmozzi, & H. K. Marshall (Edinburgh: SPIE), 99062A
Shannon, R. M., et al. 2014, MNRAS, 443, 1463
Staveley-Smith, L., et al. 1996, PASA, 13, 243
Thornton, D., et al. 2013, Science, 341, 53
Tiburzi, C., et al. 2016, MNRAS, 455, 4339
Tuthill, J., Hampson, G., Bunton, J., Brown, A., Neuhold, S., Bateman, T., de Souza, L., & Joseph, J. 2012, in 2012 International Conference on Electromagnetics in Advanced Applications (Torino: COREP), 1067

Tuthill, J., Hampson, G., Bunton, J., Harris, F., Brown, A., Ferris, R., & Bateman, T. 2015, in 2015 IEEE Signal Processing and Signal Processing Education Workshop (Piscataway: IEEE), 255
van Cappellen, W. A., Bakker, L., & Oosterloo, T. A. 2011, in 2011 30th URSI General Assembly and Scientific Symposium (Piscataway: IEEE), 1
van Straten, W. 2004, ApJS, 152, 129
van Straten, W., & Bailes, M. 2011, PASA, 28, 1
Verbiest, J. P. W., et al. 2016, MNRAS, 458, 1267
Warnick, K. F., Ivashina, M. V., Wijnholds, S. J., & Maaskant, R. 2012, ITAP, 60, 184
Weltevrede, P., & Johnston, S. 2008, MNRAS, 391, 1210