

A pilot ASKAP survey of radio transient events in the region around the intermittent pulsar PSR J1107–5907

G. Hobbs,¹ I. Heywood,^{1,2*} M. E. Bell,^{1,3} M. Kerr,¹ A. Rowlinson,^{1,3,4,5}
S. Johnston,¹ R. M. Shannon,¹ M. A. Voronkov,¹ C. Ward,^{1,6} J. Banyer,^{3,6}
P. J. Hancock,^{3,7} Tara Murphy,^{3,6} J. R. Allison,¹ S. W. Amy,¹ L. Ball,¹ K. Bannister,¹
D. C.-J. Bock,¹ D. Brodrick,¹ M. Brothers,¹ A. J. Brown,¹ J. D. Bunton,¹
J. Chapman,¹ A. P. Chippendale,¹ Y. Chung,¹ D. DeBoer,^{1,8} P. Diamond,^{1,9}
P. G. Edwards,¹ R. Ekers,¹ R. H. Ferris,¹ R. Forsyth,¹ R. Gough,¹ A. Grancea,¹
N. Gupta,^{1,10} L. Harvey-Smith,¹ S. Hay,¹¹ D. B. Hayman,¹ A. W. Hotan,¹ S. Hoyle,¹
B. Humphreys,¹ B. Indermuehle,¹ C. E. Jacka,¹ C. A. Jackson,^{7,1} S. Jackson,¹
K. Jeganathan,¹ J. Joseph,¹¹ R. Kendall,¹¹ D. Kiraly,¹ B. Koribalski,¹ M. Leach,¹
E. Lenc,^{1,3,4} A. MacLeod,¹ S. Mader,¹ M. Marquarding,¹ J. Marvil,¹
N. McClure-Griffiths,^{1,12} D. McConnell,¹ P. Mirtschin,¹ S. Neuhold,¹ A. Ng,¹
R. P. Norris,¹ J. O’Sullivan,¹ S. Pearce,¹ C. J. Phillips,¹ A. Popping,^{1,13}
R. Y. Qiao,^{11,14} J. E. Reynolds,¹ P. Roberts,¹ R. J. Sault,^{1,15} A. E. T. Schinckel,¹
P. Serra,¹ R. Shaw,¹ T. W. Shimwell,^{1,16} M. Storey,¹ A. W. Sweetnam,^{1,17}
A. Tzioumis,¹ T. Westmeier,^{1,13} M. Whiting¹ and C. D. Wilson¹

Affiliations are listed at the end of the paper

Accepted 2015 December 7. Received 2015 December 7; in original form 2015 April 12

ABSTRACT

We use observations from the Boolardy Engineering Test Array (BETA) of the Australian Square Kilometre Array Pathfinder (ASKAP) telescope to search for transient radio sources in the field around the intermittent pulsar PSR J1107–5907. The pulsar is thought to switch between an ‘off’ state in which no emission is detectable, a weak state and a strong state. We ran three independent transient detection pipelines on two-minute snapshot images from a 13 h BETA observation in order to (1) study the emission from the pulsar, (2) search for other transient emission from elsewhere in the image and (3) to compare the results from the different transient detection pipelines. The pulsar was easily detected as a transient source and, over the course of the observations, it switched into the strong state three times giving a typical time-scale between the strong emission states of 3.7 h. After the first switch it remained in the strong state for almost 40 min. The other strong states lasted less than 4 min. The second state change was confirmed using observations with the Parkes radio telescope. No other transient events were found and we place constraints on the surface density of such events on these time-scales. The high sensitivity Parkes observations enabled us to detect individual bright pulses during the weak state and to study the strong state over a wide observing band. We conclude by showing that future transient surveys with ASKAP will have the potential to probe the intermittent pulsar population.

Key words: pulsars: individual: J1107–5907.

* E-mail: george.hobbs@csiro.au

1 INTRODUCTION

The Australian Square Kilometre Array Pathfinder (ASKAP) telescope, currently being built in Western Australia, is designed to be a high-speed survey instrument (see Johnston et al. 2007). Each antenna in the array operates with a Phased Array Feed (PAF) receiver system (Hay & O’Sullivan 2008; Hampson et al. 2012) allowing observations with a field of view (FoV) of 30 deg². The PAFs have a frequency range from 700 MHz to 1.8 GHz within which the correlator can process ~ 300 MHz of instantaneous bandwidth. The full ASKAP telescope is not yet ready for scientific observations. In this paper, we describe results obtained using six of the ASKAP antennas. This small array is known as the Boolardy Engineering Test Array (BETA; Hotan et al. 2014) and has been created to demonstrate the effectiveness of PAFs in enabling high-speed radio surveys. BETA has a reduced capability in the number of beams that are able to be formed and therefore has a reduced instantaneous FoV (the actual FoV that was processed is described below).

One of the main goals for ASKAP is to explore the transient and variable radio sky via highly sensitive, wide-field observations. The Commensal Real-time ASKAP Fast Transients survey (Macquart et al. 2010) will search for fast (< 5 s) transient radio sources such as fast radio bursts. The survey for Variables and Slow Transients (VAST; Murphy et al. 2013) will search for more slowly varying sources (> 5 s) such as flare stars, intermittent pulsars and radio supernovae. However, it is likely that unexpected transient emission will also be detected. For instance, previous surveys (e.g. Hyman et al. 2002; Levinson et al. 2002; Bannister et al. 2011; Frail et al. 2012; Jaeger et al. 2012) identified a number of transient and variable sources; the origin of many of these is still unknown.

Pulsars are usually discovered through the detection of their regular sequence of pulses. However, pulsars are also known to suddenly switch off (a recent example was described in Kerr et al. 2014) or to emit giant pulses. Rapidly rotating radio transient sources (RRATS; McLaughlin et al. 2006) are now thought to be a small number of bright pulses from otherwise undetected neutron stars. Pulsar flux densities can also vary significantly because of interstellar scintillation or because the pulsar signal becomes eclipsed by a companion object. The Compact Objects with ASKAP: Surveys and Timing (COAST) survey team¹ plans to carry out traditional-style pulsar surveys using ASKAP, but the transient surveys may also find pulsars through bright individual pulses or, as described later in this paper, by monitoring their variability in the image plane.

Our survey with BETA is not as sensitive as many previous transient surveys. However, we can probe areas of parameter space not covered by previous surveys (e.g. short time-scales) and we can also target fields with known bright transient sources. For example, the pulsar PSR J1107–5907 has an extremely variable flux density and, in its bright state, is easily detectable in our image-plane survey. This provides us with a source that can be used to study one of the many phenomena that will arise in the main ASKAP transient surveys. It is also scientifically valuable to study this exotic pulsar in more detail.

PSR J1107–5907 is thought to have three emission states: strong, weak and off. The only in-depth study of this pulsar has been carried out by Young et al. (2014) who used the Parkes radio telescope to show that

- (i) the pulsar remains in the strong-emission state for ~ 200 –6000 pulses with apparent nulls over time-scales of up to a few pulses at a time,
- (ii) during the weak state a few bursts of up to a few pulses are detectable, and
- (iii) the off-state may actually be the bottom end of the pulse-intensity distribution for the source.

Young et al. (2014) emphasized that their statistical analysis of the bright emission was limited because of their small number of detections of the strong state. We therefore conducted BETA early-science observations that were centred on PSR J1107–5907. We also obtained some coincident observations of the pulsar with the Parkes 64-m radio telescope. The goals of this work were to

- (i) identify issues relating to the search for transient radio emission in a complex region of the Galactic plane,
- (ii) study the image quality obtained using the PAFs,
- (iii) monitor PSR J1107–5907 with much longer observation times than previous observations,
- (iv) search for unexpected, bright transient sources elsewhere in the field,
- (v) compare the results from different transient detection pipelines, and
- (vi) consider the possibility for discovering pulsars in wide-field imaging surveys.

All previous observations of this pulsar were obtained with the Parkes telescope. In contrast, our survey (1) uses an interferometer with a wide field of view allowing us to study not just the pulsar, but the Galactic sources in the pulsar’s vicinity, (2) is based on analysis of images instead of using traditional pulsar search methods, (3) is at a site with extremely low levels of radio-frequency interference (RFI), (4) allows us to study the pulsar with a 304 MHz bandwidth between 700 and 1000 MHz and (5) provides long-tracks on source (~ 13 h, compared with ~ 3.6 h for the longest low-frequency observations yet carried out).

The use of a telescope that was still being commissioned does lead to problems that are unlikely to arise in future surveys. We note, and will discuss later, that the current array has a poor instantaneous coverage of the (u, v) plane. When imaging complex regions of the Galactic plane this makes CLEANING the image challenging. In this paper, we therefore apply the transient pipelines to non-CLEANED images. Our survey can therefore be thought of as a worst-case scenario for future transient surveys in terms of sensitivity, calibration artefacts and snapshot image quality.

In Section 2, we describe our observing setup with BETA and Parkes. Our results are presented in Section 3. We discuss the implications of our results for studies of the pulsar and for wide-area transient surveys in Section 4. We conclude in Section 5.

2 OBSERVATIONS

Our observations were centred on the position of PSR J1107–5907. The basic parameters of this pulsar, as obtained from the ATNF pulsar catalogue (Manchester et al. 2005), are listed in Table 1.

2.1 The beta observations

BETA has been described in detail by Hotan et al. (2014). In brief, six 12 m antennas provide a total, geometric collecting area of 678 m². The maximum baseline length is 916 m giving an angular

¹ <http://pulsar.ca.astro.it/pulsar/COAST/Welcome.html>

Table 1. Basic parameters from Lorimer et al. (2006) for PSR J1107–5907.

Parameter	Value
Right ascension (hh:mm:ss)	11:07:34.46(4)
Declination (dd:mm:ss)	−59:07:18.7(3)
Dispersion measure (cm^{-3}pc)	40.2(11)
Pulse frequency (Hz)	3.95611366927(8)
Pulse frequency time derivative (s^{-2})	$-1.41(15) \times 10^{-16}$
Epoch of frequency determination (MJD)	53089
Epoch of position determination (MJD)	53089
Derived parameters	
Galactic longitude (deg)	289.94
Galactic latitude (deg)	1.11
Distance (kpc)	1.8
Pulse period (s)	0.25
Pulse period derivative	9×10^{-18}
Characteristic age (yr)	4×10^8
Surface magnetic field strength (G)	5×10^{10}
Spin-down energy loss rate (erg s^{-1})	2×10^{31}

Uncertainties in parentheses are the 1σ error on the last decimal place given for each parameter.

resolution of 1.3 arcmin. The beamformers were configured to deploy nine beams in a diamond pattern with a 1-2-3-2-1 row layout in declination, however for our project only the central, on-axis beam was processed. The array was configured to observe between 711.5 and 1015.5 MHz with the correlator delivering $16\,416 \times 18.5$ kHz frequency channels. The integration time per visibility point is 5 s.

PSR J1107–5907 was observed from UTC 01:38:00 to 14:38:00 on 2014 July 14 giving a total of 13 h of on-source time.² After the observation completed the array was repointed to place the strong radio source PKS B1934–638 at the centre of the on-axis beam and 15 min of data were recorded from UTC 14:44:00 for calibration purposes. The data were subsequently averaged to form a 304×1 MHz channel ‘continuum’ data set. The data are provided in standard measurement sets, and the *CASA*³ package was used for editing, calibration and imaging. The *FLAGDATA* task was used on both the target and the calibrator scans in order to remove the autocorrelations, and to clip the data (with a manually determined threshold) in order to remove occasional spurious instrumental artefacts (that give unrealistically high amplitudes or amplitudes of exactly zero). A single pass of the *RFLAG* algorithm with time and frequency thresholds of 5σ was also applied to both data sets.

The PKS B1943–638 scans were then averaged in time and the *BANDPASS* task was used to solve for the instrumental bandpass shape against the model derived by Reynolds (1994). This set of corrections has the dual purpose of correcting the bandpass shape and setting the absolute flux-density scale of the data. The derived corrections were then applied to the target data set.

At this stage the usual procedure is to apply one or more passes of self-calibration to correct the data further, with calibration of the complex antenna-beam gains being performed against a sky model derived from the image. For typical extragalactic fields away from the Galactic plane this approach works extremely well. However, self-calibration was not applied to the data presented here. With

² Since the observations for this work were carried out it has been found that the time stamps were incorrect by ~ 10 s. The times given in this paper are therefore incorrect by approximately this time (the exact offsets are not yet fully determined). Such an offset does not change the scientific results of this paper.

³ <http://casa.nrao.edu>




Figure 1. The (u, v) plane coverage of the PSR J1107–5907 observations. Each baseline is individually coloured. The (u, v) plane coverage for one of the two minute snapshot images is overlaid in green. The total observing time was 13 h with an observing bandwidth of 304 MHz.

only six antennas and a maximum baseline of approximately 1 km, BETA is a relatively sparse array. The (u, v) plane coverage of the target observation is shown in Fig. 1. Even with the benefits of multifrequency synthesis imaging (e.g. Sault & Wieringa 1994; Rau & Cornwell 2011), significant gaps remain in the Fourier plane coverage, particularly in the inner region. Strong emission from the Galactic plane is present over a broad range of low spatial frequencies, and since the Fourier domain is somewhat sparsely sampled by the interferometer the large-scale emission is not faithfully reproduced in the image domain. This precludes self-calibration: using the image to derive a sky model results in one that is highly incomplete and does not accurately transform into an approximately true representation of the visibilities. We trialled different methods in which we used subsets of the available baselines and only using compact features, but these methods made the image noticeably worse, suppressing many of the real diffuse structures. Thus, the images that are presented in this paper have been produced with the complex gain corrections derived from the single observation of PKS B1934–638 as their sole correction. However, as will be shown below and in the next sections, the images produced accurately reproduce features known in the field and allow us to successfully detect the pulsar.

The data from the beam centred on PSR J1107–5907 were imaged using the *CASA CLEAN* task. A 2240×2240 pixel image was produced with each square pixel spanning 12 arcsec, giving a total sky area of $7.5 \times 7.5 \text{ deg}^2$. The image was deconvolved interactively employing a multifrequency synthesis algorithm that models each component with a third-order Taylor series (Rau & Cornwell 2011), manually improving the cleaning mask in order to avoid spurious features with each major cycle. The final deconvolved wide-field image is presented in Fig. 2 and described in Section 3. No primary beam correction has been applied.

Searches for transient radio emission can be based on the raw visibilities obtained from the telescope or on snapshot images. It is likely that transient searches (for events lasting more than a few seconds) with the full ASKAP array will be based on studies of




Figure 2. BETA image of the area around PSR J1107–5907 formed from the central on-axis beam with a central frequency of 863.5 MHz. No primary beam correction has been applied. The array pointing position is that of the pulsar, which is marked with the cross. The dashed circle shows the approximate half power point of the formed PAF beam at the centre frequency. Inset are two examples of complex structures in the Galactic plane that BETA images (left) with excellent correspondence to the existing MGPS-2 images (right). The intensity scale of the main image is given by the grey-scale bar at the top of the figure. Note that the inset images are not shown on the same brightness scale due to the primary beam attenuation. The Carina nebula is observed through the first sidelobe of the primary beam.

snapshot images and so we use the same procedure here. Before these images were produced, the spectral clean component model produced as a result of the imaging step was inverted into a set of model visibilities. This model was then subtracted from the observed data to produce a residual visibility data base (note that the pulsar was not included in this model). The residual visibilities were then split into two-minute intervals (390 epochs) and dirty (i.e. not cleaned) images were formed to be processed by the transient search pipelines. The choice of two-minute intervals was based on the need to study the shortest time-scale variability possible, whilst ensuring that the pulsar can be detected with high S/N in its bright state.

In summary, our transient search is based on two-minute snapshot images. In those images the brightest sources have been removed as well as possible using the CLEAN component model that was derived from the full-length image. The CLEAN procedure was not applied to the snapshot images. Any transient source that lasted up to a few minutes would therefore be detectable in the snapshot images with the shape of the dirty beam at that time. The images described in this paper are available for download from the CSIRO data archive.⁴

⁴ <http://hdl.handle.net/102.100.100/25153?index=1>. Relevant files have file-names beginning with SB220 (this nomenclature is from the scheduling block #220).

2.2 The Parkes observations

The main aim of our work is to use PSR J1107–5907 to demonstrate the effectiveness of transient searches with ASKAP. In order to confirm our results, and to obtain multifrequency observations of the pulsar, we carried out simultaneous observations with the Parkes telescope. The Parkes telescope has a long history of detecting and studying pulsars and fast transient events and relatively straightforward and mature procedures are in place to record and process the observations. The Parkes observations described here were carried out as part of a program to study the stability of the pulsed emission (project code P863⁵) through long, multifrequency observations of a small sample of pulsars (which includes PSR J1107–5907). These observations were obtained in ‘search-mode’, in which the data are averaged over a short integration time (256 μ s) from UTC 09:53 to 10:53 on 2014 July 14. We observed using a dual-band receiver providing 64 MHz of bandwidth centred on 732 MHz and simultaneously 1 GHz of bandwidth centred on 3100 MHz.

⁵ After the standard embargo period of 18 months these data will be publically available for download through the CSIRO Data Access Portal; <http://data.csiro.au> (Hobbs et al. 2011). The relevant files are s140714_095334 and t140714_09333.

Before the pulsar observation commenced we observed a pulsed-calibrator signal for subsequent polarization and flux-density calibration. The band edges have low gain and are often corrupted and so, following standard procedure, we removed the outermost 5 per cent of the channels. A recent flux calibration file for the appropriate receiver and backend was selected from archival data. A `PSRSH` script with instructions to load these files and to undertake frontend and flux calibration was produced. This script was loaded into the `DSPSR` software package along with an ephemeris for the pulsar to output single calibrated pulse profiles with 256 phase bins and 512 frequency channels. The `DSPSR` package was similarly used to produce time averaged profiles with two minute integrations. The two minute integrated profiles were frequency-averaged. The `PSR-FLUX` package with self-standardization was then used to estimate the average flux density of each two minute profile and its uncertainty. We have also studied the entire time series (i.e. without splitting into individual pulses) using the `PFTS` software package. For this processing we summed the two polarization states, dedispersed at the known dispersion measure of the pulsar and then summed in frequency. We then do not flux calibrate these time series.

3 RESULTS

As this paper reports some of the first observations carried out by ASKAP, we first confirmed that the image contains known sources and that the flux and astrometric calibration was adequate for our purposes. An overview of the area around PSR J1107–5907 is shown in Fig. 2. The position of the pulsar is marked, and the dashed circle represents the approximate half power point of the formed beam at the centre of the band. The Galactic plane runs diagonally across the image. The pixel scale is given by the grey-scale bar at the top of the image and saturates black at 100 mJy beam⁻¹.

The broad ripples introduced by the undersampled Galactic emission are evident, however smaller scale individual features in the plane of the Galaxy are reproduced remarkably well. Inset into the upper left of Fig. 2 are two such examples, featuring the BETA image (left-hand panels) and comparison images from the second epoch Molonglo Galactic Plane Survey (Murphy et al. 2007, MGPS-2; right-hand panels). The upper row shows the supernova remnant SNR G289.7–00.3 (e.g. Green 2009) and the lower row shows the bright Carina nebula (e.g. Preibisch et al. 2011). There is excellent correspondence between the features recovered by ASKAP and those seen in the MGPS-2 images. Note however that these images are not on the same intensity scale. For SNR G289.7–00.3 the ASKAP and MGPS-2 images saturate at 13 and 50 mJy beam⁻¹, respectively, and the corresponding values for the Carina nebula are 0.1 and 1.6 Jy beam⁻¹. We note that the Carina nebula has been imaged through the first sidelobe of the primary beam.⁶

As the position of the pulsar is known, we simply measured the flux density corresponding to the known source position as a function of time using the 2 min snapshot images. This is shown in panel (A) of Fig. 3. In panel (B), we show the flux density at a nearby position in the sky. The pulsar was undetectable for most of the observation. However, three events occurred (around 200, 500 and 700 min from the start of the observation) in which the pulsar switched into its bright state and flux densities of around 200 mJy

were observed. Panel (B) shows that remaining imaging artefacts do vary the baseline noise level, but the pulsar state changes are easily detectable.

3.1 Transient detection pipelines

Even though the pulsar state changes were detected using the simple method described above, we also applied various transient pipelines to the data set. The aim was to (1) study the effectiveness of those pipelines with this challenging data set and (2) to search for unexpected events elsewhere in the image. We trialled three pipelines. The first pipeline was developed specifically for this data set. The second pipeline was developed for the LOFAR telescope (the Transient Pipeline; `TRAP`; Swinbank et al. 2015). The third pipeline was developed to search for transient events with the completed ASKAP telescope (the VAST Transients Pipeline; Murphy et al. 2013).

Both the `TRAP` and VAST pipelines have been designed to work with `CLEANED` images. Sources that have not been de-convolved are therefore problematic for the pipelines. However, because a sky model has been subtracted from these snapshots, any point-like flux found in these images is either a transient object or an imaging artefact. In either case, the pipelines offer convenient ways to identify and explore these objects.

3.1.1 Pipeline 1

The first pipeline uses a straightforward matched-filter methodology in which the point spread function (PSF) of the beam is matched with every pixel of the image. To mitigate edge effects, we considered only the central quadrant (1024 pixels). At each position within this quadrant, we computed the log likelihood difference, $\delta \log \mathcal{L}$, between the null and alternative hypotheses representing the absence and presence, respectively, of a source. If the image noise were Gaussian, the significance of excesses would be given by $\sqrt{2\delta \log \mathcal{L}}$. However, the noise in the image is dominated by sidelobes of the PSF, thus is highly correlated, non-Gaussian and time-dependent. We thus adopted a ‘normalized’ statistic, $S = \delta \log \mathcal{L} / \sigma^2$ where σ^2 is the root-mean-square deviation for each residual image.

To determine a threshold for detection, we ‘self-calibrated’ by examining the distribution of S in an 80×70 pixel subset of the images. Interestingly, we found that, when normalized to the mean value of S within this region, the overall distribution of S follows the theoretical χ^2 distribution well to significances $>5\sigma$. In panel (C) of Fig. 3, we plot the flux density of the pulsar as measured using this pipeline as a function of time. The three burst events are clearly detectable. This pipeline does not naturally provide an error bar for each flux-density measurement. Instead we obtained the significance of the event being real. These significance values are shown in Fig. 4. The first bright state is detected with high significance ($\sim 10\sigma$) in each two minute bin. The second and third bursts are detected with significance values of $\sim 5\sigma$.

In order to search for other burst events in the field we choose $S > 25$, nominally 5σ , as the threshold for source detection. Using the ‘known’ flux of the pulsar, this threshold translates to a flux sensitivity of about 65 mJy for a two-minute snapshot. This sensitivity rises slightly as the (u, v) coverage becomes poorer at low source elevations. We summed the S images obtained from each snapshot to produce a significance map of the field. We next counted the number of pixels surpassing the 5σ threshold and show the results in Fig. 5. The pulsar (which is significant in ~ 20 snapshots) and the artefacts produced by the incompletely subtracted bright sources dominate the field. We therefore conclude that the pulsar can

⁶ The ASKAP antennas can rotate to keep the parallactic angle fixed as an observation progresses, thus the sidelobes do not rotate with respect to the sky, and a typical ASKAP observation generally detects a significant number of sources in the sidelobes without distortion.




Figure 3. Panel (A) flux density versus time for PSR J1107–5907 and Panel (B) for a sky position close to the pulsar as obtained directly from the original images. Panel (C) shows the results from Pipeline 1. Panels (D) and (E) show the blind search and monitored-search, respectively, output from the TRAP pipeline. Note that panel (D) shows the direct output from the TRAP pipeline, which only provides data after the transient is detected. Panel (F) contains the output from the VAST pipeline. Panel (G) shows the flux density obtained in the 40 cm band from the Parkes telescope.

easily be detected using this algorithm and that no other unexpected transient sources exist in the data.

The application of a matched filter, pixel-based technique to the dirty images is not common in radio transient searches, but was well-suited to the quality of these commissioning data. We expect image quality with the full ASKAP array to be substantially better, both due to improved UV coverage and improved calibration and beamforming techniques. Searches for transients in such data will certainly benefit from cleaned images, particularly those in regions requiring complicated sky models. On the other hand, fully cleaning images is substantially more expensive from a computational standpoint. For transient searches of data based on short (minute time-scale) integrations, a pixel-based method like the one used here could provide an efficient, quasi-real time transient search ca-

pability that is optimally matched to instantaneous UV coverage of the array.

3.1.2 The LOFAR transients pipeline

The LOFAR Transients Pipeline (TRAP) has been developed by the LOFAR Transients Key Science Project in order to search automatically for transient and variable sources (Swinbank et al. 2015).⁷ We applied release 2.0 of this pipeline to our data set. Initially the

⁷The TRAP software (TraP contributors 2014) is publicly available: <https://github.com/transientskp/tpk> with documentation at <http://docs.transientskp.org>.




Figure 4. Detection significance as a function of time at the position of the pulsar from Pipeline 1.




Figure 5. The grey-scale bar indicates the number of times a given pixel in the image has been significantly detected by Pipeline 1. The pulsar is at the centre of the image. The sources to the top-right and to the bottom-left are related to the incompletely subtracted bright sources (MGPS J110538–584950 and MGPS J111245–594729).

pipeline determines the RMS noise in the inner quarter of each residual image. The noise was found to be roughly Gaussian with an RMS noise of 21 mJy for the 2 min residual images.

We assumed that any transient event will occur from a point source and, in these images, a point source will take the shape of the PSF. We performed an unconstrained source extraction on the PSF of the dirty beam and use the result from a Gaussian fit to the most significant part of the PSF as the ‘restoring beam shape’. Sources were blindly extracted from the images using a 5σ detection threshold and a 3σ analysis threshold,⁸ fitted to the

⁸ The detection threshold is the threshold above which sources are considered ‘found’. The pipeline then extracted all the pixels around a detection to the 3σ level. A Gaussian function was fitted to those 3σ and above pixels to determine the source flux density. Full details are given in Swinbank et al. (2015).

shape of the most significant part of the PSF. The position of the sources can vary significantly in these images because of systematic calibration errors. Therefore, we used a 5 arcmin systematic position uncertainty to aid with source association. All other TRAP parameters were kept as the default values and we searched a FoV of $6.1 \times 6.1 \text{ deg}^2$ centred on the phase centre of the image.

Newly detected sources were identified within TRAP. A source is considered to be a transient if the measured flux density was greater than eight times the worst RMS noise region from the previous best image. This 8σ threshold was determined from the 5σ detection threshold and a 3σ margin.

As our images are ‘dirty images’, TRAP finds an excess of imaging artefacts extracted around the location of the two, incompletely subtracted, brightest sources. With the exception of PSR J1107–5907, which is blindly extracted in multiple images, no convincing transient or variable sources were identified. Four sources were labeled as confirmed transients by the pipeline: PSR J1107–5907, an artefact from a bright sidelobe of PSR J1107–5907 and two imaging artefacts from one of the subtracted sources. The results from the blind search are shown in panel (D) of Fig. 3. Additionally, TRAP has a monitoring functionality where the fit is always conducted at the position of interest. We re-ran the pipeline utilizing the average position from TRAP of PSR J1107–5907. The resulting light curves obtained using the two minute snapshot images are shown in panel (E) of Fig. 3. This pipeline clearly identifies the initial large burst occurring around UTC 05:53:04 which lasts for ~ 40 min. As expected, the pipeline also detects statistically significant bursts at UTC 10:47:04 and 13:35:04. These two panels (D and E) are almost identical. The exception being the second ‘on’ period around 500 min after the observation start. The flux density measured in panel (D) (the blind-search method) is significantly higher than that in panel (E). This is because the apparent position of PSR J1107–5907 slightly varies throughout the observation (with a maximum deviation of ~ 13 arcsec). This variation is caused by a combination of instrumental and calibration effects, the low S/N determination of the actual source position in the 2-min snapshot images and ionospheric variations. In panel (D), a fit has been made for the flux density of the closest source that can be associated with the pulsar. This is therefore a better measure of the pulsar’s flux density.

3.1.3 The VAST pipeline

The VAST detection pipeline is fully described by Bell et al. (2014). In brief, for a time-sequence of images the VAST pipeline first performs a source finding procedure via the AEGEAN algorithm (see Hancock et al. 2012 for details). We use a background estimation tool to efficiently determine local measurements of the noise in regions surrounding the sources. Sources are then cross-matched as a function of time (and frequency) and variability statistics are calculated for the resulting light curves.

This pipeline was applied to our data set and we followed the same procedure of calculating the PSF as described in the previous subsection. The pipeline identifies the pulsar as a transient object and the resulting light curve is shown in panel (F) of Fig. 3. This pipeline clearly shows that all three burst events from the pulsar are detected. No other sources in the field were identified as being transient. A number of artefacts were identified within the images owing to poorly subtracted bright sources, lack of (u, v) coverage and deconvolution: these artefacts were easily mitigated by visual inspection.

3.2 The Parkes observations

Panel (G) in Fig. 3 shows the flux density measured at 40 cm (732 MHz) from the Parkes telescope for observations that have been averaged in 2 min intervals. The total observation time is much shorter than that obtained using ASKAP. However, the Parkes data clearly show that the pulsar switched to its bright state at a time that agrees with the ASKAP observations and confirms that the second ‘on’ state is real. The flux density during the ‘on’ state when averaged over 2 min is 150 mJy. This is slightly lower than that obtained with ASKAP, but consistent with the uncertainties on the ASKAP flux density measurements. We note that the Parkes data is at a slightly lower frequency than the ASKAP observations and therefore it is not surprising that the flux density measurements do not perfectly agree.

4 DISCUSSION

4.1 The pulsar

The observations described here were undertaken to study the effectiveness of transient pipelines for identifying strong radio burst events. The ASKAP observations provide the longest observations to date of the pulsar. We have 11.2 h of data in which we can search for transient emission (for this calculation we do not account for the data near the start and near the end of the observation in which some data are missing). During this time the pulsar was in the strong state for ~ 44 min giving an ‘on’-fraction of ~ 6.5 per cent. This agrees

well with the result of Young et al. (2014) of 6 per cent obtained with Parkes data alone. These switches give a typical time-scale for the bursting phenomenon of $11.2/3 = 3.7$ h although we emphasize that the event times are not periodic.

Young et al. (2014) noted that the strong-mode pulses are only emitted during relatively short burst periods (they report intervals of ~ 60 s to ~ 24 min). The ASKAP observations show that the interval can be even longer. The initial strong-mode lasted around 44 min. However, there is at least one 2 min subimage during that time in which the measured flux density is close to the baseline level suggesting that the pulsar switched off for up to 2 min before returning to the bright state.

The high sensitivity of the Parkes telescope and the use of a dual-band receiver system provides the opportunity to study the pulsar in more detail than before. The Parkes sensitivity is such that we can detect individual pulses from the pulsar in its bright state and, via averaging over a few minutes, can detect the pulsar in its weak state. For instance, we show in Fig. 6 the time series recorded with Parkes in both observing bands. We have not flux calibrated these data streams, but it is clear that the pulsar switched into the bright state near the end of the Parkes observation.

In Fig. 7, we re-plot the Parkes observations around the region of the bright burst event. Individual pulses from pulsars are known to be complex and exhibit significant pulse-to-pulse variability. This is clearly true for PSR J1107–5907, but with the high S/N Parkes observations of the single pulses we can study the pulse-to-pulse variability across a wide bandwidth. We have highlighted three pulse regions in the figure. Pulse region #1 is significantly brighter in the




Figure 6. The Parkes observation of the single pulses from PSR J1107–5907. The top and bottom panels show the 40 and 10 cm observations, respectively.




Figure 7. The Parkes observation of the single pulses from PSR J1107–5907 during the time that the pulsar remained in the strong emission state. The top and bottom panels show the 40 and 10 cm observations, respectively.

40 cm band than in the 10 cm band. This is typical of most pulsar pulses – pulsars typically have steep spectral indices and are brighter at lower frequencies. However, in contrast pulse #2 is significantly brighter in the 10 cm band and almost undetectable in the 40 cm band. Pulse #3 has a similar flux density in the two bands. These data sets have not been flux calibrated and so we cannot determine an absolute spectral index, but it is clear that the spectral index is significantly changing between the individual pulses.

We therefore conclude:

- (i) The pulsar switched into the bright state at exactly the same time in the two observing bands.
- (ii) The pulsar was in the bright state for around 81 s (322 pulses), but the pulses are not constant during the bright state. Many pulses seem to be missing and the flux density of the observed pulses vary dramatically. We have inspected the dynamic spectra for some of the brightest single pulses. These dynamic spectra show that the diffractive bandwidth for J1107–5907 is smaller than the observing bandwidth in these observing bands and so the flux density variations are not caused by the interstellar medium.
- (iii) The spectral index for individual pulses varies significantly.

The 40 cm data set in Fig. 6 contains a few individual bright pulses that occur outside of the strong emission state. However, this observing band is affected by RFI. The RFI can be easily distinguished from pulses originating at the pulsar by identifying dispersed pulses. In Fig. 8, we show two examples. The first is clearly a dispersed pulse originating from the pulsar and the second

is RFI. In the 40 cm observing band, the dedispersed time series indicated 67 significant pulses in the ‘weak’ or ‘off’ states. We inspected each of these by eye and found that 14 were clearly RFI, 13 were definitely pulses from the pulsar and the remainder were too weak to make a definitive statement. We note that no clear bright pulses from the pulsar were detected after the end of the strong state (for the final 380 s of our Parkes observation). This may suggest that the pulsar became ‘quieter’ after the strong emission state ceased, or this may simply be a co-occurrence (there are earlier times in which the pulsar exhibited no detected pulses over a similar duration).

In this section, we have provided an initial look at the high-time resolution data being obtained by the Parkes P863 observing program. This program is continuing and a detailed study of the individual pulses from this pulsar using Parkes observations that span many years will be published elsewhere.

4.2 Comparison of the pipelines

Pipeline 1 was designed for this data set. The other two pipelines were designed for much cleaner data sets. The first pipeline searches for pixels that change between different images. The TRAP and VAST pipelines are based around identifying and tracking sources in the image. If the flux density of any source significantly changes then that source is identified as a possible transient. Whereas pipeline 1 was not optimized for speed, the other pipelines are required to run on very large data sets and therefore are required to run quickly on a given data set.




Figure 8. Example pulses occurring when the pulsar is in the weak state. The left-hand pulse originated at the pulsar. The right-hand pulse was caused by RFI.

Within the `TRAP` and `VAST` pipelines there are subtle differences in the underlying algorithms and techniques used to produce results. Both pipelines, however, have a workflow that is fundamentally similar and the final science goals are the same. So it is unsurprising that both of these pipelines were able to clearly detect PSR J1107–5907 and that the light curves in Fig. 3 (D), (E) and (F) are similar.

There are differences in how the `VAST` pipeline and `TRAP` deal with a newly detected transient source such as PSR J1107–5907. After detection of a new source within an image, both pipelines operate in the following manner for subsequent images (in time). If a given source is blindly detected in multiple images it will be cross-matched but, if it is undetected, then a flux measurement will be taken at the location of the source using the restoring beam parameters (or, as used for this paper, the PSF shape). The difference between the pipelines lies in the treatment of images prior to the transient event. Upon the first detection of a new source, the `VAST` pipeline will go back through all the previous images and measure the flux at the location, thus building up a complete light curve of the source and confirming its transient nature. However, refitting the source in all the previous images can be very time consuming when there are hundreds of images. Additionally, in some cases the previous images may have been discarded (because of data storage restrictions), so they will not be available for analysis. `TRAP` has been designed to counter these two issues by assuming that the images are not available and uses statistics from the previous images, stored in the image data base, to determine if the new source is a transient (as explained in Swinbank et al. 2015). Therefore, `TRAP` outputs the light curve of transient sources from the time of first detection (as shown in panel D of Fig. 3). In addition to the blind searches, `TRAP` has the functionality of a monitoring list, in which the flux is always measured at the position of specified sources irrespective of the blind detections. As these observations were specifically targeting the variability of PSR J1107–5907, we were able to take advantage of this monitoring list functionality to extract the full light curve of PSR J1107–5907 (as shown in panel E of Fig. 3), which can be directly compared to the outputs of Pipeline 1 and the `VAST` pipeline.

We note that the flux density uncertainties obtained by the `VAST` pipeline and `TRAP` are different, with larger uncertainties quoted by `TRAP`. This is likely related to the differences in fitting approaches and background characterization used by the source finders within each pipeline and the constraints enforced on each of the fits (position and Gaussian shape). For instance, `TRAP` always assumed a point source

and all blind fits by the source finder were constrained to Gaussians taking the shape of the PSF, while the `VAST` pipeline used fully unconstrained fits on blind detections. Both source finders assumed the identical constraints for fits when the source was undetected. As neither source finder has been designed to deal with `UNCLEANED` images of this type, a detailed comparison would be best conducted with an appropriately `CLEANED` data set and this is beyond the scope of this paper.

4.3 Implications for transient sources

Even though `BETA` has relatively poor sensitivity compared with other transient surveys (e.g. Carilli, Ivison & Frail 2003; Mooley et al. 2013), we have studied a relatively large sky area on a time-scale that has not been probed in great detail before in the 20 cm observing band (from 2 min to 12 h). For instance, large single dish telescopes such as the Parkes telescope have probed millisecond to second time-scales over periods of a few hours, often for a single object. The majority of transient imaging surveys focus on studying time-scales greater than minutes, typically days to months (see table 5 in Carbone et al. 2014, for details). These imaging surveys, so far, have typically operated between 1.4 and 5 GHz and quite rarely achieve sensitivities below 1 mJy. Note, a number of blind transient surveys at frequencies <500 MHz have opened up the low end of the radio spectrum to transient detection (see Lazio et al. 2010; Jaeger et al. 2012; Bell et al. 2014; Carbone et al. 2014); ASKAP will be able to detect transients at intermediate frequencies from 700 to 1800 MHz.

In this section, we consider the rate of transients events as none were detected. We exclude the pulsar from this analysis as it was a targeted source. Presenting and comparing results from transient surveys is non-trivial as different surveys probe different time-scales, observing frequencies and sky areas. To ascertain the area and sensitivity this survey has probed we perform the following analysis (note this differs from the detection threshold reported for pipeline 1).

- (i) For each image, we calculate an RMS map where each of the pixels represents the 1σ sensitivity in that image and region.
- (ii) A Gaussian primary beam correction is applied to the RMS map values. Note, the ASKAP-`BETA` primary beam response has not been fully characterized yet. We therefore work under the




Figure 9. Log sensitivity σ_{cut} ($\mu\text{Jy beam}^{-1}$) versus cumulative area A (deg^2) normalized over all epochs. The cumulative area gives the amount of sky we have observed with sensitivity $\sigma < \sigma_{\text{cut}}$. For example, approximately 2.0 deg^2 was observed with sensitivity better than 25 mJy .

assumption that each antenna is a fully illuminated circular aperture with Gaussian profile.

(iii) Over all the images we calculate the cumulative area that the images probe for a number of different sensitivity bins. We normalize this cumulative sky area by the number of images in our sample and we plot the results in Fig. 9. We plot both the 1σ values and also the 8σ values of sensitivity. The 8σ curve represents the detection threshold of this survey. We will discuss further below caveats to this analysis.

We can use Poisson statistics to calculate an upper limit on the expected number of transient events (per deg^2) in any given snapshot at this frequency, sensitivity and cadence (see also Bell et al. 2014). This snapshot surface density (ρ) is defined as

$$\rho < \frac{\ln(0.05)}{(n-1) \times \Omega}, \quad (1)$$

where Ω is the total field of view (or sky area) and n is the number of independent images searched ($n = 389$).

From equation (1), we therefore obtain the upper limit on the snapshot surface density of transient events calculated as function of area and sensitivity, listed in Table 2. At the most sensitive part of the images ($< 2 \text{ deg}^2$) we place a limit of $\rho < 3.9 \times 10^{-3} \text{ deg}^{-2}$ above a detection threshold of 0.2 Jy beam^{-1} . Over the full field of view that was searched (13.3 deg^2) a limit of $\rho < 5.8 \times 10^{-4} \text{ deg}^{-2}$ above a detection threshold of 8 Jy beam^{-1} .

The half-power beam width is defined as $\theta = 1.028(\lambda/D)$ and assuming $\lambda = 0.34 \text{ m}$ (centre frequency of the observations) and $D = 12 \text{ m}$ we find an area of $A = \pi(\theta/2)^2 = 2.28 \text{ deg}^2$. Our limit of $\rho < 3.9 \times 10^{-3} \text{ deg}^{-2}$ covers an area of $< 2 \text{ deg}^2$ in the most sensitive part of the beam, which is also the most appropriate place to search for transients. As we imaged a larger field of view than the main beam we place limits on these sky areas as well, albeit it with decreased sensitivity. There are however caveats to searching in

Table 2. The sensitivity calculated over seven discrete sky area bins normalized over all epochs. The corresponding upper limits on the snapshot rates of transients at those areas and sensitivities are given in the final column.

Cumulative area (deg^2)	Sensitivity (Jy beam^{-1})	8σ sensitivity (Jy beam^{-1})	ρ (deg^{-2})
2.0	0.025	0.2	3.9×10^{-3}
4.0	0.037	0.3	1.9×10^{-3}
6.0	0.053	0.4	1.3×10^{-3}
8.0	0.08	0.6	9.7×10^{-4}
10.0	0.14	1.1	7.7×10^{-4}
12.0	0.31	2.5	6.4×10^{-4}
13.3	1.0	8.0	5.8×10^{-4}

these far out regions, for example the beam is not well characterized and the ability to detect off-axis transients is not well investigated.

Furthermore within the main beam region we did not perform deconvolution and so certain areas will be contaminated by bright sources and side lobes and not suitable for transient searches. Our sky area calculations should therefore be considered as conservative upper limits. We note that in calculating the actual noise properties of the images we do account for regions of increased noise due to sources, side lobes and imaging artefacts.

These upper limits are time independent, i.e. they do not factor in a characteristic time-scale of transient phenomena. From the Pipeline 1 results, we can place a 95 per cent confidence upper limit (Feldman & Cousins 1998) on the rate of transients with a $\sim 2 \text{ min}$ time-scale of $7 \times 10^{-6} \text{ s}^{-1} \text{ deg}^{-2}$.

The most comparable work to this with regards to the choice of observing frequency is Bannister et al. (2011). In that study, 22 yr of archival The Sydney University Molonglo Sky Survey (SUMSS) images at 843 MHz were searched for transients on time-scales of 12 h to years above a flux density threshold of 14 mJy . Two transient sources were detected in their survey yielding a snapshot surface density of $\rho = 1.5 \times 10^{-3} \text{ deg}^{-2}$. The surface density of transients found by Bannister et al. (2011) is comparable to the limits placed via this work albeit at a lower detection threshold. The typical cadence of the two surveys is however very different (minutes versus years). Bannister et al. (2011) reported that the detected transients were extragalactic in origin and had long time-scales \sim years, that plus the lower detection threshold of their survey means that it would be impossible for them to be detectable with our observations. However, it is noteworthy that we can achieve a very competitive limit on the surface density of events with such a small amount of observing time and with a commissioning array with 1/6th of the final collecting area and 1/30th of the field of view of the final ASKAP telescope.

4.4 Implications for pulsar discovery

Traditional pulsar surveys are based on recording the signal from the telescope averaged over a relatively short period ($\sim 100 \mu\text{s}$) with sufficient frequency resolution. Off-line processing is carried out to de-disperse the time series with a large number of trial dispersion measures. For each dispersion measure, a Fourier transform is taken to search for a periodic signal. This method is affected by RFI and telescope gain variations. The sensitivity of the survey to a particular pulsar significantly degrades if the emission is intermittent or if the pulsar has a binary companion. These traditional searches are also generally only sensitive to pulsars with periods $> 1 \text{ ms}$ and $< 10 \text{ s}$. As all the pipelines tested in this paper managed to clearly identify PSR J1107–5907, it is possible that intermittent pulsars could also

be discovered in future ASKAP transient searches and we briefly discuss this prospect here.

We do not know the properties of the population of intermittent pulsars in detail and therefore it is not possible to provide a detailed study of what new pulsars a given survey may find. We know that pulsars can be detected from a single pulse (the RRATS), they can be ‘on’ for minutes and then ‘off’ on time-scales of hours (such as PSR J1717–4054; Kerr et al. 2014), days (such as PSR B1931+24 which has a 30 d time-scale; Kramer et al. 2006) or that they can be ‘on’ for years and then ‘off’ for years (e.g. PSR J1832+0029; Lorimer et al. 2012).

We have simulated pulsars switching between ‘on’ and ‘off’ states on various time-scales and then determined whether different survey strategies could detect them. For traditional pulsar surveys we simply require that the signal, when averaged over the entire observation is above the limiting flux density sensitivity for the survey (~ 30 min and 0.2 mJy for the Parkes multibeam pulsar survey; Manchester et al. 2001). We also consider the detectability of a pulsar as a continuum point source in an ASKAP continuum survey. We assume a survey similar to that proposed for the Evolutionary Map of the Universe (EMU) survey which will allow a point source detection limit of ~ 50 μ Jy (Norris et al. 2011). Finding a pulsar in such a survey will involve being able to distinguish a pulsar candidate from all other point sources.

We also simulate two transient searches. First, we have a simulated survey similar to the actual survey carried out in this paper; we assume that the EMU 12 h observations are divided into 5 min snapshot images and searched for transient emission. Secondly, we consider a transient search in which a specific area of sky is only imaged for 5 min, but that observation is repeated once per week over a period of five years (labeled here as a VAST-style survey).

The detectability of a pulsar in these surveys depends upon how long the pulsar remains in its ‘on’ and ‘off’ states. If the pulsar switches on a very short time-scale (i.e. less than a minute) then all surveys will simply identify the pulsar as a non-transient object and the detectability simply depends upon the limiting flux density of the survey. Such pulsars will not be found in the transient searches. However, for pulsars similar to J1107–5907, that are on for minutes to hours and then off for a few hours, the chance of detection in the traditional observations decreases as the ‘on’ time-scale decreases.

We created a simulated data set in which a pulsar with a given flux density switched ‘on’ for a certain time and then ‘off’ for another time. We then considered 1000 trial searches with specified survey parameters and, in each trial, determined whether the pulsar was detectable or not. We then increased the pulsar’s flux density until it could be detected in at least 50 per cent of the trials. This gave us the limiting flux density for 50 per cent detection for a given ‘on’ and ‘off’ time-scale. These results are plotted in Fig. 10. The top panel is for a pulsar with an ‘on’ time-scale of 18 min (the typical time-scale for PSR J1717–4054). The bottom panel is for a pulsar with an ‘on’ time-scale of 30 d (much larger than the observing time for any simulated observation). In this figure, the red dashed line indicates the Parkes multibeam pulsar survey, the blue dashed line the EMU ASKAP survey, the green solid line is the transient search based on a single 12 h observation and the blue solid line is for repeated, short observations of the same sky position.

The top panel shows that the transient searches cannot detect pulsars with very short (< 5 min) ‘off’ time-scales as they always switch ‘on’ during each 5 min snapshot image and therefore are not observed as transient sources. These pulsars are detected with the traditional pulsar survey and the EMU survey if their flux densities are greater than the survey sensitivity. The VAST-style survey




Figure 10. Sensitivity of various surveys to intermittent pulsars with different ‘off’ time-scales. The top panel is for a pulsar that is ‘on’ for 18 min. The bottom panel is for a pulsar that is ‘on’ for 30 d. The dashed red line indicates a survey similar to the Parkes Multibeam Pulsar Survey. The dashed blue line an EMU-style continuum survey. The solid green line represents a transient survey based on 5 min snapshots of the 12 h EMU observations. The blue line represents a VAST-style survey in which 5 min observations are repeated each week for 5 yr. The pulsars indicated have the correct ‘off’ time-scales, but all have different ‘on’ time-scales and are therefore simply included to give the currently known range of switching time-scales.

provides sensitivity to pulsars with relatively long ‘off’ time-scales, but the survey is only sensitive to relatively bright pulsars.

The bottom panel is more extreme. For a pulsar that is ‘on’ for 30 d, it is unlikely that it switches ‘off’ during an observation. Our simulations therefore indicate that such pulsars are unlikely to be detected in a transient search of the 12 h EMU observations. No green, solid line is therefore shown in this panel. The EMU continuum survey and the traditional pulsar survey are limited only by the survey sensitivity until the ‘off’ time-scale reaches > 30 d. The VAST-style survey is sensitive to such pulsars with a much greater range of ‘off’ time-scales.

We therefore conclude that repeated, relatively short observations of the same sky region provides the best survey strategy for discovering unknown intermittent pulsars. We note that the advantages of such a survey with ASKAP would also include (1) observations at a very radio-quiet site (RFI mitigation is currently a major challenge for finding such pulsars) and (2) is at a lower frequency than most Parkes surveys (and therefore pulsars are likely to be brighter).

5 CONCLUSION

We have shown that BETA can produce high-quality images of complex regions of the Galactic plane that are consistent with previous work. The observations allow searches for transient events on time-scales between minutes and ~ 12 h. These are time-scales that have not previously been studied in detail. We have shown that pipelines developed for LOFAR and ASKAP can successfully detect PSR J1107–5907, but no other transient sources were detected.

This work has just touched on the possibilities that will arise with the full ASKAP telescope. We have shown that it will be an ideal telescope for discovering intermittent pulsars that spend significant amounts of time in their ‘off’ state. This work therefore demonstrates the role of future, wide-area survey telescopes to detect and identify transient sources of radio emission. It also demonstrates the requirement for a sensitive telescope, like Parkes, that can follow-up on the detected events with wide frequency coverage, high time resolution and the ability to carry out long-term monitoring programs. Together such telescopes will be able to discover and study both expected and unexpected transient sources.

ACKNOWLEDGEMENTS

The Australian SKA Pathfinder is part of the Australia Telescope National Facility which is managed by CSIRO. Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. Establishment of the Murchison Radio-astronomy Observatory was funded by the Australian Government and the Government of Western Australia. ASKAP uses advanced supercomputing resources at the Pawsey Supercomputing Centre. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Parts of this research were conducted by the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020. The Parkes radio telescope is part of the Australia Telescope, which is funded by the Commonwealth of Australia for operation as a National Facility managed by the Commonwealth Scientific and Industrial Research Organization (CSIRO). This paper includes archived data obtained through the CSIRO Data Access Portal (<http://data.csiro.au>). The work was supported by iVEC through the use of advanced computing resources located at The Pawsey Centre. We gratefully acknowledge the LOFAR Transients Key Science Project for providing us with access to the LOFAR Transients Pipeline prior to the public release.

REFERENCES

Bannister K. W., Murphy T., Gaensler B. M., Hunstead R. W., Chatterjee S., 2011, *MNRAS*, 412, 634
 Bell M. E. et al., 2014, *MNRAS*, 438, 352
 Carbone D. et al., 2014, preprint ([arXiv: 1411.7928](https://arxiv.org/abs/1411.7928))
 Carilli C. L., Ivison R. J., Frail D. A., 2003, *ApJ*, 590, 192
 Feldman G. J., Cousins R. D., 1998, *Phys. Rev. D*, 57, 3873
 Frail D. A., Kulkarni S. R., Ofek E. O., Bower G. C., Nakar E., 2012, *ApJ*, 747, 70
 Green D. A., 2009, *Bull. Astron. Soc. India*, 37, 45
 Hampson G. et al., 2012, International Conference on Electromagnetics in Advanced Applications (ICEAA), 807
 Hancock P. J., Murphy T., Gaensler B. M., Hopkins A., Curran J. R., 2012, *MNRAS*, 422, 1812
 Hay S. G., O’Sullivan J. D., 2008, *Radio Sci.*, 43
 Hobbs G. et al., 2011, *PASA*, 28, 202

Hotan A. W. et al., 2014, *PASA*, 31, 41
 Hyman S. D., Lazio T. J. W., Kassim N. E., Bartleson A. L., 2002, *AJ*, 123, 1497
 Jaeger T. R., Hyman S. D., Kassim N. E., Lazio T. J. W., 2012, *AJ*, 143, 96
 Johnston S. et al., 2007, *PASA*, 24, 174
 Kerr M., Hobbs G., Shannon R. M., Kiczyński M., Hollow R., Johnston S., 2014, *MNRAS*, 445, 320
 Kramer M., Lyne A. G., O’Brien J. T., Jordan C. A., Lorimer D. R., 2006, *Science*, 312, 549
 Lazio T. J. W. et al., 2010, *AJ*, 140, 1995
 Levinson A., Ofek E. O., Waxman E., Gal-Yam A., 2002, *ApJ*, 576, 923
 Lorimer D. R. et al., 2006, *MNRAS*, 372, 777
 Lorimer D. R., Lyne A. G., McLaughlin M. A., Kramer M., Pavlov G. G., Chang C., 2012, *ApJ*, 758, 141
 McLaughlin M. A. et al., 2006, *Nature*, 439, 817
 Macquart J.-P. et al., 2010, *PASA*, 27, 272
 Manchester R. N. et al., 2001, *MNRAS*, 328, 17
 Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, *AJ*, 129, 1993
 Mooley K. P., Frail D. A., Ofek E. O., Miller N. A., Kulkarni S. R., Horesh A., 2013, *ApJ*, 768, 165
 Murphy T., Mauch T., Green A., Hunstead R. W., Pietrzynska B., Kels A. P., Sztajer P., 2007, *MNRAS*, 382, 382
 Murphy T. et al., 2013, *PASA*, 30, 6
 Norris R. P. et al., 2011, *PASA*, 28, 215
 Preibisch T. et al., 2011, *A&A*, 530, A34
 Rau U., Cornwell T. J., 2011, *A&A*, 532, A71
 Reynolds J., 1994, AT Technical Memos, 39.3/040
 Sault R. J., Wieringa M. H., 1994, *A&AS*, 108, 585
 Swinbank J. D. et al., 2015, *Astron. Comput.*, 11, 25
 TraP contributors 2014, *Astrophysics Source Code Library*, record ascl:1412.011
 Young N. J., Weltevrede P., Stappers B. W., Lyne A. G., Kramer M., 2014, *MNRAS*, 442, 2519

- ¹ *CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia*
- ² *Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa*
- ³ *ARC Centre of Excellence for All-sky Astrophysics (CAASTRO)*
- ⁴ *Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, the Netherlands*
- ⁵ *Netherlands Institute for Radio Astronomy (ASTRON), PO Box 2, NL-7990 AA Dwingeloo, the Netherlands*
- ⁶ *Sydney Institute for Astronomy, School of Physics, University of Sydney, NSW 2006, Australia*
- ⁷ *International Centre for Radio Astronomy Research (ICRAR), Curtin University, GPO Box U1987, Perth, WA 6845, Australia*
- ⁸ *Radio Astronomy Laboratory, University of California Berkeley, 501 Campbell, Berkeley, CA 94720-3411, USA*
- ⁹ *SKA Organisation, Jodrell Bank Observatory, Lower Withington, Macclesfield, Cheshire SK11 9DL, UK*
- ¹⁰ *Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune University Campus, Pune 411 007, India*
- ¹¹ *CSIRO Digital Productivity, PO Box 76, Epping, NSW 1710, Australia*
- ¹² *Research School of Astronomy and Astrophysics, Australian National University, Mount Stromlo Observatory, Cotter Road, Weston Creek, ACT 2611, Australia*
- ¹³ *International Centre for Radio Astronomy Research (ICRAR), University of Western Australia, 35 Stirling Highway, Crawley, WA 6009, Australia*
- ¹⁴ *Sonartech ATLAS Pty Ltd, Unit G01, 16 Giffnock Avenue, Macquarie Park, NSW 2113, Australia*
- ¹⁵ *School of Physics, University of Melbourne, VIC 3010, Australia*
- ¹⁶ *Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*
- ¹⁷ *31 Ellalong Road, North Turrumurra, NSW 2074, Australia*

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