










## Research Article

# The Southern-sky MWA Rapid Two-metre (SMART) pulsar survey—I. Survey design and processing pipeline

N. D. R. Bhat<sup>1</sup>, N. A. Swainston<sup>1</sup>, S. J. McSweeney<sup>1</sup>, M. Xue<sup>2</sup>, B.W. Meyers<sup>1,3</sup>, S. Kudale<sup>4</sup>, S. Dai<sup>5</sup>, S. E. Tremblay<sup>6</sup>, W. van Straten<sup>7</sup>, R. M. Shannon<sup>8</sup>, K. R. Smith<sup>1</sup>, M. Sokolowski<sup>1</sup>, S. M. Ord<sup>9</sup>, G. Slep<sup>1</sup>, A. Williams<sup>1</sup>, P. J. Hancock<sup>10</sup>, R. Lange<sup>10</sup>, J. Tocknell<sup>11</sup>, M. Johnston-Hollitt<sup>10</sup>, D. L. Kaplan<sup>12</sup>, S. J. Tingay<sup>1</sup>, and M. Walker<sup>1</sup>

<sup>1</sup>International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia, <sup>2</sup>National Astronomical Observatories, Chinese Academy of Sciences, Datun Road, Chaoyang District, Beijing 100101, China, <sup>3</sup>Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada, <sup>4</sup>National Centre for Radio Astrophysics, Tata Institute of Fundamental Research, Pune 411 007, India, <sup>5</sup>Western Sydney University, Locked Bag 2751, Penrith South DC, NSW 1797, Australia, <sup>6</sup>National Radio Astronomy Observatory, 1003 Lopez Road, Socorro, NM 87801, USA, <sup>7</sup>Institute for Radio Astronomy & Space Research, Auckland University of Technology, Private Bag 92006, Auckland 1142, New Zealand, <sup>8</sup>Centre for Astrophysics and Supercomputing, Swinburne University of Technology, P.O. Box 218, Hawthorn, VIC 3122, Australia, <sup>9</sup>CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia, <sup>10</sup>Curtin Institute for Computation, Curtin University, GPO Box U1987, Perth, WA 6845, Australia, <sup>11</sup>Australian Astronomical Optics Macquarie, Macquarie University, Sydney, NSW, Australia and <sup>12</sup>Department of Physics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA

### Abstract

We present an overview of the Southern-sky MWA Rapid Two-metre (SMART) pulsar survey that exploits the Murchison Widefield Array's large field of view and voltage-capture system to survey the sky south of 30° in declination for pulsars and fast transients in the 140–170 MHz band. The survey is enabled by the advent of the Phase II MWA's compact configuration, which offers an enormous efficiency in beam-forming and processing costs, thereby making an all-sky survey of this magnitude tractable with the MWA. Even with the long dwell times employed for the survey (4800 s), data collection can be completed in <100 h of telescope time, while still retaining the ability to reach a limiting sensitivity of ~2–3 mJy (at 150 MHz, near zenith), which is effectively 3–5 times deeper than the previous-generation low-frequency southern-sky pulsar survey, completed in the 1990s. Each observation is processed to generate ~5000–8000 tied-array beams that tessellate the full ~610 deg<sup>2</sup> field of view (at 155 MHz), which are then processed to search for pulsars. The voltage-capture recording of the survey also allows a multitude of *post hoc* processing options including the reprocessing of data for higher time resolution and even exploring image-based techniques for pulsar candidate identification. Due to the substantial computational cost in pulsar searches at low frequencies, the survey data processing is undertaken in multiple passes: in the first pass, a shallow survey is performed, where 10 min of each observation is processed, reaching about one-third of the full-search sensitivity. Here we present the system overview including details of ongoing processing and initial results. Further details including first pulsar discoveries and a census of low-frequency detections are presented in a companion paper. Future plans include deeper searches to reach the full sensitivity and acceleration searches to target binary and millisecond pulsars. Our simulation analysis forecasts ~300 new pulsars upon the completion of full processing. The SMART survey will also generate a complete digital record of the low-frequency sky, which will serve as a valuable reference for future pulsar searches planned with the low-frequency Square Kilometre Array.

**Keywords:** surveys: sky surveys – instrumentation: interferometers – methods: observational – pulsars: general – techniques: interferometric

(Received 22 March 2023; accepted 23 March 2023)

## 1. Introduction

Even after five decades of productive research, pulsars continue to enable us to push the frontiers of physics and astrophysics. These compact dense stars harbour physical conditions that are non-existent elsewhere in the universe (e.g. ultra-strong gravitational and magnetic fields and supra-nuclear matter densities),

which make them invaluable tools for studying extreme physics. They are arguably amongst the most widely-exploited astrophysical objects, with applications ranging from probing the state of ultra-dense matter to testing strong-field gravity (e.g. Demorest et al. 2010; Kramer et al. 2006; van Straten et al. 2001), and from probing micro-arcsecond structure and turbulence in the interstellar medium (ISM) to complex stellar evolutionary scenarios (e.g. Bhat et al. 2004; Archibald et al. 2009; Bailes et al. 2011). The phenomenal impact and high-profile scientific applications (e.g. pulsar timing arrays for the detection of nanohertz-frequency gravitational waves) has elevated pulsar science to the ranks of a key science for the Square Kilometre Array (SKA; e.g. Keane et al. 2015; Janssen et al. 2015; Shao et al. 2015).

**Corresponding author:** N. D. R. Bhat, Email: ramesh.bhat@curtin.edu.au.

**Cite this article:** Bhat NDR, Swainston NA, McSweeney SJ, Xue M, Meyers BW, Kudale S, Dai S, Tremblay SE, van Straten W, Shannon RM, Smith KR, Sokolowski M, Ord SM, Slep G, Williams A, Hancock PJ, Lange R, Tocknell J, Johnston-Hollitt M, Kaplan DL, Tingay SJ, and Walker M. (2023) The Southern-sky MWA Rapid Two-metre (SMART) pulsar survey—I. Survey design and processing pipeline. *Publications of the Astronomical Society of Australia* 40, e021, 1–20. <https://doi.org/10.1017/pasa.2023.17>

The backbone that enables this is the net result of a series of large pulsar surveys conducted over the past five decades (e.g. Manchester *et al.* 2001; Cordes *et al.* 2006; Keith *et al.* 2010; Stovall *et al.* 2014). Invariably, most of them involved tessellating large parts of the sky of the instrument and recording data at high time and frequency resolutions (i.e. large data rates) and performing sensitive searches over the vast parameter space that is practically feasible. Many of them were prompted by the advent of new instrumentation or technology, and often exploited the computing affordable at the time. They have also proven invariably rewarding in the longer term, and often yielded a substantial increase in the pulsar population. For instance, the Molonglo pulsar survey in the 1970s found 150 pulsars, practically doubling the known pulsar population at the time (Manchester *et al.* 1978), while the Parkes multibeam survey from the 1990s (Manchester *et al.* 2001) found 742 pulsars, and discovered exotica such as the double pulsar system J0737–3039A/B and the eccentric neutron star-white dwarf binary J1141–6545, both of which have proven to be unique laboratories for testing general relativity and alternate theories of gravity (Kramer *et al.* 2006; Bhat, Bailes, & Verbiest 2008; Venkatraman Krishnan *et al.* 2020; Kramer *et al.* 2021). This success led to next-generation multibeam surveys at Parkes and Arecibo, and more recently with the Five-hundred-metre Aperture Spherical radio Telescope (FAST). Already these have collectively discovered 600 pulsars to date. The landmark discovery of fast radio bursts (FRBs) in the Parkes high-time resolution radio universe (HTRU) survey (Thornton *et al.* 2013) even opened up an entirely new field of research. Large pulsar surveys have a proven track record of their ability to return significant scientific dividends in the long run, with the majority of the discoveries and spin-off science emerging from follow-up processing over the years.

These multibeam surveys have largely been at frequencies  $\gtrsim 1$  GHz. The past decade also witnessed a number of successful low-frequency pulsar surveys, most of which were prompted by the advent of new-generation low-frequency facilities (e.g. Low-Frequency Array; LOFAR), or new receivers or pulsar instrumentation at the more traditional facilities such as the Green Bank Telescope (GBT) and the Giant Metre-wave Radio Telescope (GMRT). Notable among these are the drift-scan surveys with the Arecibo Telescope and GBT, and the ongoing surveys at the GMRT and GBT. The drift-scan surveys, in the 300–350 MHz range, despite their non-traditional nature, have led to  $>100$  pulsar discoveries, while the highly successful Green Bank Northern Celestial Cap (GBNCC) survey has, to date, found 160 pulsars. The net tally from the low-frequency surveys of the past decade alone is  $>400$  pulsars, including 73 pulsars by the LOFAR Tied-Array All-Sky (LOTAAS) survey (Sanidas *et al.* 2019). Additionally, targeted searches have been undertaken toward unidentified Fermi gamma-ray sources (mostly at low frequencies), leading to  $>80$  pulsars (Deneva *et al.* 2021, and references therein). The LOTAAS survey, the processing of which is still ongoing, also discovered the longest-period (23.5 s) pulsar known until recently (Tan *et al.* 2018b), when a 76-s pulsar was discovered with MeerKAT (Caleb *et al.* 2022). In essence, surveys at low frequencies have proven to be highly effective, particularly in uncovering the local population of pulsars, and mapping out the high-Galactic latitude ( $b$ ) parts of the sky.

Surveys at low frequencies offer several benefits but they also have their limitations. An appealing factor is the generally steep spectral nature of most radio pulsars, where the flux density at

frequency  $\nu$  is  $S_\nu \propto \nu^\alpha$ , where  $\alpha$  is the spectral index. The spectral index is known to vary over a wide range for pulsars,  $-4 \lesssim \alpha \lesssim 0$ , but the average spectral index  $\langle \alpha \rangle = -1.6 \pm 0.03$  for long-period pulsars (Jankowski *et al.* 2018), and is somewhat steeper ( $\langle \alpha \rangle = -1.9 \pm 0.1$ ) for millisecond pulsars (MSPs; Toscano *et al.* 1998; Dai *et al.* 2015), with a  $1-\sigma$  dispersion of  $\sim 1$ . While this suggests most pulsars are significantly brighter at low frequencies, this is more than offset by the even steeper dependence of the sky background noise ( $T_{\text{sky}} \propto \nu^{-2.55}$ ). The sky background is also highly direction-dependent and is typically significantly reduced toward higher Galactic latitudes. The main benefit is the inherently larger fields-of-view of the low-frequency telescopes, which substantially increase the efficiency in telescope time required and hence the time for completion of large surveys.

Amongst the multitude of other considerations are interstellar medium (ISM) propagation effects, which tend to majorly influence low-frequency pulsar searches; the most familiar (and significant) one is the dispersion that manifests as frequency-dependent time delays in arrival times  $\Delta t \propto \text{DM} \nu^{-2}$ , where the dispersion measure (DM) is the line-of-sight integral of the electron density  $n_e$ . This non-linear, inverse dependence in frequency implies very large delays at low frequencies ( $\lesssim 200$  MHz); that is, a pulsar with a  $\text{DM} = 100 \text{ pc cm}^{-3}$  will have its signal spread over  $\sim 7.5$  s in observations made over a 30 MHz band centred at 150 MHz, as opposed to  $\lesssim 0.1$  s across a similar (i.e. 20%) fractional bandwidth around 1.4 GHz. Circumventing this necessitates much finer frequency resolution ( $\Delta\nu$ ) so the residual dispersive smearing across the finite channel width can be minimised, and consequently requires many more channels across the recording bandwidth, and hence a much larger data rate and substantial processing needs.

The other significant effect is pulse broadening resulting from multipath propagation as a consequence of scattering in the ISM, the characteristic time for which is a non-linear function of both DM and frequency, that is,  $\tau_d \propto \text{DM}^{-2.2} \nu^{-4.4}$ , under the assumption of a pure Kolmogorov form of electron density spectrum (Cordes, Weisberg, & Boriakoff 1985). This poses a significant limitation in low-frequency pulsar searches, especially when the pulse broadening time exceeds the pulsar's spin period, that is,  $\tau_d \gtrsim P$ , as it results in a significant degradation or even a loss of sensitivity to periodic emission. As with the sky background, scatter broadening is also highly line-of-sight dependent; it is much larger in the plane, or toward the Galactic Centre, compared to high- $|b|$  sight lines. Empirical relations exist to guide expected broadening times as a function of DM and frequency (e.g. Bhat *et al.* 2004; Geyer *et al.* 2017), and can be used to guide the observing/search strategies, that is,  $\tau_d \gtrsim 100$  ms at  $\text{DM} \gtrsim 300 \text{ pc cm}^{-3}$ , for a line of sight as far off as  $|b| \sim 5^\circ$  and  $\sim 30^\circ$  away from the Galactic Centre (GC) in longitude. This implies, at low frequencies, the search volume is largely limited to a few kiloparsecs in the plane. However, this is not a serious limitation at higher Galactic latitudes, where the DM tends to saturate at  $\sim 20\text{--}50 \text{ pc cm}^{-3}$  for  $|b| > 15^\circ$ . In other words, the higher survey speeds of low-frequency surveys can be optimally exploited for covering high- $|b|$  parts of the sky, without compromising detection sensitivity.

Yet another relevant ISM effect, especially at low frequencies, is the modulation of apparent pulsar intensities due to scintillation effects. As with the pulse broadening, the observable effects strongly depend on frequency and the line of sight, as it is essentially another manifestation of multipath propagation. For relatively nearby pulsars ( $\text{DM} \lesssim 50 \text{ pc cm}^{-3}$ ) this often manifests as rapid (and very large) modulations in both time and frequency

with characteristic scales in the range  $\sim 0.1\text{--}5$  MHz and  $\sim 1\text{--}100$  min at  $\sim 150$  MHz; this is diffractive scintillation (e.g. Rickett 1990). Refractive scintillation also leads to intensity modulations, but on much longer timescales of days to weeks (at low frequencies), and the observed variations in mean flux densities can be as much as by a factor  $\sim 5\text{--}6$  for low to moderate DM pulsars (e.g. Bell et al. 2016; Bhat et al. 2018). From the perspective of candidate detection in low-frequency searches, this sometimes results in fortuitous brightening (or inauspicious dimming) of pulsars, which provides the opportunity to detect pulsars that were missed earlier (e.g. owing to scintillation dimming), or to detect a pulsar that might be below the sensitivity limit of a survey. This further strengthens the case for low-frequency surveys.

Despite these challenges, pulsars were originally discovered at low frequencies (at 81.5 MHz; Hewish et al. 1968) and much of the early years of pulsar astronomy were focused at low frequencies. The eventual quest to find them in large numbers and timing them at high precision pushed much of pulsar astronomy (searches and timing in particular) to frequencies  $\gtrsim 1$  GHz. However, the advent of several low-frequency telescopes over the past decade and advances in affordable high-performance computing are effectively leading to a resurgence of low-frequency astronomy including large sky surveys, many of which are conducted at frequencies  $\lesssim 500$  MHz.

The success of these northern surveys strongly motivates an all-sky pulsar survey with the Murchison Widefield Array (MWA) that operates in the 80–300 MHz range in the Southern Hemisphere. The MWA, which was originally built as an array of 128 tiles (where each tile is a  $4\times 4$  dipole array) with a maximum baseline of 3 km, is also Australia's precursor for the low-frequency SKA (i.e. SKA-Low; Tingay et al. 2013). Even though the MWA was not initially designed for pulsar science, the eventual addition of a voltage-capture system (VCS; Tremblay et al. 2015) and the development of software-defined instrumentation (for offline processing) equipped it as a pulsar-capable facility. Notwithstanding the limitations of large data rates ( $28\text{ TB h}^{-1}$ ) and the associated data management/processing challenges, the VCS has been exploited for wide-ranging science from studies of millisecond pulsars to sporadic emission from pulsars (e.g. Bhat et al. 2016; Meyers et al. 2018; Kaur et al. 2019), and from investigating the pulsar emission physics to studying propagation effects caused by the interstellar medium (e.g. McSweeney et al. 2017; Kaur et al. 2022). The progress in this area, along with the array's upgrade to Phase II (Wayth et al. 2018), whereby a compact configuration of 128 tiles within 300 m was possible on a semi-regular basis, has made all-sky pulsar searches tractable with this telescope.

The SMART survey described in this paper has two main objectives: (1) performing sensitive searches for pulsars and fast transients in the sky south of  $+30^\circ$  in declination at 140–170 MHz; and (2) mapping the sky for low-frequency detection of already known pulsars in the southern sky. The main novelty of the survey is the use of a voltage-capture mode for data recording (as opposed to the filterbank data format that has been adopted for all past and ongoing surveys), and hence an astonishingly high survey speed for data collection, that is,  $\sim 450\text{ deg}^2\text{ h}^{-1}$  in 100- $\mu\text{s}/10\text{-kHz}$  resolutions). However, the computational cost of processing (i.e. beamforming and searching) are substantial at low frequencies, and thus drive the feasible strategies for data processing, especially at early stages.

With the large survey speed substantially reducing the demand for telescope time for survey completion, longer dwell times

become affordable, which also increases the sensitivity to the detection of sporadic or intermittent class of objects such as rotating radio transients (RRATs; McLaughlin et al. 2006), intermittent or state-switching pulsars, extreme nullers etc. (e.g. Kerr et al. 2014) among the classes of radio-emitting neutron stars, and even the enigmatic fast radio bursts (FRBs; e.g. Thornton et al. 2013). The detectability all of these transient class of objects is dictated by the 'on-sky' time metric  $\Sigma = \Omega T$  where  $\Omega$  is the instantaneous field-of-view (FoV) and  $T$  is the time spent on sky (dwell time in the case of an all-sky survey). Following the discussion in Sanidas et al. (2019) in the context of LOTAAS,  $\Sigma_{\text{SMART}} = 52735\text{ deg}^2\text{ h}$ , which is a factor of two more than that of LOTAAS for which  $\Sigma_{\text{LOTAAS}} = 23400\text{ deg}^2\text{ h}$  (at 135 MHz), and indeed much larger than  $\Sigma_{\text{GBNCC}} = 1430\text{ deg}^2\text{ h}$ ,  $\Sigma_{\text{GHRSS}} = 835\text{ deg}^2\text{ h}$  and  $\Sigma_{\text{AO327}} = 132\text{ deg}^2\text{ h}$  (all at 300–350 MHz).

Here we present an overview of the Southern-sky MWA Rapid Two-metre (SMART) pulsar survey. In Section 2 we outline the main science goals, and describe the observing strategy adopted for sky tessellation. Procedures for data processing and analysis are described in Section 3, and the strategies for confirmation and initial follow-up in Section 4. In Section 5 we describe the survey simulations and the expected yield. Future processing plans are outlined in Section 6, followed by a summary in Section 7.

## 2. Survey description

### 2.1. Science goals and motivation

The broader goals of the SMART survey are similar to most other large sky surveys, that is, exploring the new parameter space that is opened by a leap in instrumentation, technology, or sensitivity and to uncover a large population of previously undetected pulsars. The fact that the currently known pulsar population ( $\sim 3300$ , cf. the Australia Telescope National Facility (ATNF) pulsar catalogue<sup>a</sup> v1.67; Manchester et al. 2005) represents only a small fraction ( $\lesssim 10\%$ ) of the total expected (i.e. beamed in our direction) Galactic population (e.g. Keane et al. 2015, and references therein) strongly motivates such large sky surveys. Indeed, conducting a full Galactic census of pulsars is a high-priority science objective for the SKA. Further, given the number of broader questions surrounding the neutron-star population (e.g. birth rates, and comparison with rates of supernovae), the detectable pulsar population is largely guided by the known population of pulsars at any given time. It is therefore imperative to explore every possible avenue and steadily refine our knowledge of pulsar population. Furthermore, the detection prospects of pulsars in a given frequency band strongly depends on the emission and propagation properties at those frequencies; however, the current forecast of a detectable population in the SKA-Low band is largely guided by the pulsar population uncovered by high-frequency surveys.

Obtaining a large body of measurements such as DM, scattering and Faraday rotation, by using pulsars as probes of the ISM, will also enable mapping out the distribution of magneto-ionic (and turbulent) plasma in the Galaxy, which is steadily refined with a larger sample of measurements (e.g. Cordes & Lazio 2002; Bhat et al. 2004; Deller et al. 2016; Yao, Manchester, & Wang 2017).

Finally, an underlying goal of any large sky pulsar survey is to discover exotic objects; while it is hard to design any particular survey specifically for this, historical examples are abundant, for

<sup>a</sup><https://www.atnf.csiro.au/research/pulsar/psrcat/>.



**Table 1.** Parameters of large pulsar surveys over the past decade.

Survey	Telescope	Frequency band (MHz)	Sky coverage	Time resolution ( $\mu$ s)	Frequency resolution (kHz)	Dwell time (s)	$S_{\min}^a$ (mJy)	Reference
LOTAAS	LOFAR	119–151	$\delta > 0^\circ$	491.52	12.21	3600	1–2	SCB+19
SMART	MWA	140–170	$\delta < +30^\circ$	100	10	4800	2–3	This work
GBNCC	GBT	300–400	$\delta > -40^\circ$	81.92	24	120	1.1	SLR+14
GHRSS	GMRT	306–338	$-40^\circ > \delta > -54^\circ$	30.72–61.44	15.625–31.25	900, 1200	1.0	BCM+16
HTRU	Parkes	1182–1522	$\delta < +30^\circ$	64	390	240, 540, 4200	0.2–0.6	KJvS+10
GPPS	FAST	1100–1500	$-10^\circ < b < +10^\circ$	49.152	244.14	300	0.005	HWW+21

<sup>a</sup> Minimum detectable flux density for a  $10\text{-}\sigma$  detection, for long-period pulsars ( $P \gtrsim 0.1$  s), with small duty cycle ( $W/P \sim 0.05$ ), and at DMs  $\lesssim 50$  pc  $\text{cm}^{-3}$ .

Notes: Survey description reference – SCB+19: Sanidas *et al.* (2019) for LOTAAS; SLR+14: Stovall *et al.* (2014) for GBNCC; BCM+16: Bhattacharyya *et al.* (2016) for GHRSS; KJvS+10: Keith *et al.* (2010) for HTRU; HWW+21: Han *et al.* (2021) for GPPS.

example, the discovery of the double pulsar in the Parkes multi-beam (PMB) survey (Lyne *et al.* 2004), the 23.5-s period pulsar in LOTAAS (Tan *et al.* 2018b), and the transitional millisecond pulsar (MSP) in the Arecibo drift-scan survey (Archibald *et al.* 2009). All such broader and high-profile objectives are certainly applicable for the SMART survey.

The SMART pulsar survey also perfectly complements ongoing northern-sky surveys in sky and frequency coverage (Table 1). Surveys at low frequencies will likely be sensitive to a different pulsar population, and therefore an all-sky survey at low frequencies is also essential to develop a comprehensive picture of neutron-star/pulsar populations in the Galaxy. Bearing this in mind (and as we detail in Section 2.4), the survey is designed to reach a final sensitivity comparable to that of LOTAAS, that is, the use of long dwell times (4800 s) to attain a limiting sensitivity ( $10\sigma$ ) of  $\sim 2\text{--}3$  mJy for long-period pulsars with small duty cycles, and assuming a spectral index  $\alpha = -1.5$  and no turnover down to  $\sim 150$  MHz. This is  $\sim 3\text{--}5$  times deeper than the previous-generation low-frequency (70 cm) survey (Manchester *et al.* 1996) in the south (and thence an accessible search volume  $\sim 5\text{--}10$  times larger), and  $\sim 2\text{--}3$  times deeper than the high-latitude segment of the Parkes HTRU survey (Keith *et al.* 2010).

The SMART survey will also serve as a reference survey for future deeper surveys at low frequencies, such as those planned with SKA-Low (Keane *et al.* 2015). While the success of (and the lessons learned from) all ongoing low-frequency surveys will indeed inform SKA-Low pulsar surveys, the SMART survey will potentially play an additional important role, since the MWA is also the official low-frequency precursor for SKA-Low, and is located at the same site where SKA-Low will be built. Specifically, the sky coverage of the SMART survey is identical to that of SKA-Low, which means a higher degree of synergistic overlap in calibration and beamforming methodologies, than most northern facilities. The role of reference surveys is vividly demonstrated by the later generation multibeam surveys in the south; for example, the PMB survey for its successors, the HTRU pulsar survey (Keith *et al.* 2010) and the SURvey for Pulsars and Extragalactic Radio Bursts (SUPERB; Keane *et al.* 2018), which can now play a similar role for the planned surveys with MeerKAT. However, aside from the Parkes 70 cm survey of the 1990s, the low-frequency southern-sky remains essentially unexplored for pulsar searches, especially at  $\lesssim 300$  MHz.

Aside from the aforementioned primary science goals, there are also some auxiliary goals for the SMART survey, largely enabled by the novelty of the data recording strategy, that is, the use of voltage capture system and post-processing, as opposed to

the beamformed data in the filterbank format. These not only facilitate a number of additional strategies for confirmation and follow-up, but they can also be potentially exploited for developing and trialling alternate strategies for pulsar searches; for example, image-based techniques for the identification of promising candidates that take advantage of pulsar properties such as steep spectrum, variability or circular polarisation (e.g. Sett *et al.* 2022). These, in principle, also offer some advantages over traditional search methods, especially for extreme pulsars like those with sub-millisecond periods, or distant pulsars whose pulse shapes will be significantly broadened due to multi-path scattering, but will be sensitive primarily to very bright sources.

Notwithstanding the anticipated scientific merits of the SMART survey, computational requirements are substantial, especially given the large data rate of the VCS and searching at low frequencies, thereby necessitating a multi-pass processing strategy. In the first-pass processing, we perform a shallow survey, where 10 min of data from each observation are processed, and the search is limited to basic periodicity, and DMs up to 250 pc  $\text{cm}^{-3}$ . In this paper, we outline the observing strategies employed for the survey, and processing strategies adopted for the initial phase, and present analysis and results to date, as well as plans and strategies for future processing. A companion paper (hereafter Paper II) will describe the survey status, pulsar census to date and more details on follow-up strategies including timing and imaging follow-ups.

## 2.2. Survey strategy

The novel strategy employed for the SMART survey, that is, the use of VCS recording from 128 tiles, which allows high-time resolution (and instantaneous) sampling of a very large patch of the sky (but at the expense of a large data rate of 28 TBh<sup>-1</sup>), necessitates substantial processing to enable large-scale pulsar searching applications. Most importantly, the voltage data from the tiles need to be coherently combined to generate thousands of tied-array beams prior to any search processing. The undertaking of the SMART survey is particularly enabled by the Phase II upgrade, whereby a compact configuration of 128 tiles within  $\sim 300$  m became available on a semi-regular basis. The compact configuration of Phase II brings an enormous efficiency in terms of beamforming cost; specifically, the number of tied-array (i.e. phased array) beams required to fill the full FoV (at a gain level down to half power point) is reduced from  $2.7 \times 10^5$  for the Phase I array to  $3.9 \times 10^3$  for the Phase II compact array. This reduction of more than two orders of magnitude in the computational cost makes an all-sky high-sensitivity pulsar search tractable (and affordable) with an

interferometric array like the MWA. Thus, with the beamforming step integrated into software-defined instrumentation, this effectively translates into an impressively large survey speed of  $\sim 450 \text{ deg}^2 \text{ h}^{-1}$ , that is, the full visible sky of the MWA ( $\delta < +30^\circ$ ) can be surveyed in a modest number of VCS pointings.

The first-pass survey strategy of processing only 10 min of data from each observation (hence reaching about one-third of the full-search sensitivity) was adopted also to boost the prospects of early pulsar discoveries. Even though the combination of the VCS mode and the FoV provides a large survey speed, practical considerations such as the availability of the compact configuration necessitated multiple observing campaigns to advance the survey. Further details including the survey status and completion plans are described in Paper II.

### 2.3. Beamforming and Sky tessellation

The signal processing chain of the MWA including the high time resolution system is described in a number of earlier papers (e.g. Tingay et al. 2013; Prabu et al. 2015; Tremblay et al. 2015), and is briefly reiterated here. In the legacy system that was employed for survey campaigns to date, the VCS sub-system follows the second stage of channelisation in the signal path. Each element of the array is a  $4 \times 4$  dipole array, called a ‘tile’, the signals from which are fed to an analogue beamformer that defines the FoV. The beamformed signals are Nyquist-sampled at 655.56 Msps and channelised (after signal conditioning) using a polyphase filterbank (PFB) to generate  $256 \times 1.28$ -MHz signal outputs (i.e. coarse channelisation), 24 of which are transported to the central processing facility, where a second-stage PFB operation is performed, resulting in  $128 \times 10$ -kHz time series for each coarse channel, that is, 3072 channels across the recording 30.72 MHz bandwidth. These voltage time series are written to an array of RAID disks by the VCS as 4+4-bit complex voltage samples. These data are recorded (up to a maximum duration of 100 min) and transported to the Pawsey Supercomputing Centre where further processing (including calibration and beamforming) is carried out.

VCS-recorded data can be processed offline for calibration and tied-array beamforming (Ord et al. 2019) and, optionally, can also be reprocessed to reconstruct a higher time resolution voltage data at the native coarse-channel resolution of  $0.78 \mu\text{s}$  (McSweeney et al. 2020). To realise the SMART pulsar survey, this beamformer functionality was further enhanced to optionally generate several dozens of tied-array beamformed outputs simultaneously—i.e. the so-called *multi-pixel* beamformer, which is essentially the front end of the pulsar search processing chain. The implementation details and benchmarks are described in Swainston et al. (2022). This software tied-array beamformer has been benchmarked on Pawsey’s Garrawarla and Swinburne’s OzSTAR supercomputers. It performs  $3 \times$  faster on the latter, which has been the primary high-performance computing (HPC) platform for much of our SMART data processing.

Thanks to the large FoV of the MWA ( $\sim 610 \text{ deg}^2$  at 155 MHz, near zenith), the entire sky south of declination  $\delta < +30^\circ$  can be covered in a modest number of telescope pointings. The sky tiling strategy is shown in Figure 1. In short, we adopted pointings similar to that of the GaLactic and Extragalactic All-sky MWA survey (Wayth et al. 2015), that is, meridian drift scans optimised for maximum sensitivity at each declination as well as for more reliable calibration (referred to as ‘sweet spots’). In this case, the number of pointings depends on the degree of overlap in right

ascension (RA), with a minimum of 58 pointings for minimal ( $1^\circ$ ) overlap and 78 for a  $15^\circ$  overlap. A large overlap is more optimal as it effectively serves as a two-pass strategy, which is desirable at low frequencies where intermittency (from effects such as scintillation) tends to be more pronounced. After exploring the full range of options, and also factoring in the available resource constraints, we converged on a  $10^\circ$  overlap as an acceptable choice.<sup>b</sup> As shown in Figure 1, this amounts to a total 70 pointings, that is, 93 h of telescope time for the full SMART survey.

For each pointing, many tied-array beams (TABs) are formed to maximise the sensitivity across the FoV. The tied-array beams are pointed towards fixed right ascension and declination, with the necessary adjustments to the tile phases made for every second of data (Ord et al. 2019). Thus, although the observations themselves are drift scans, sources can be tracked by the same TAB for up to the full duration of the observation.

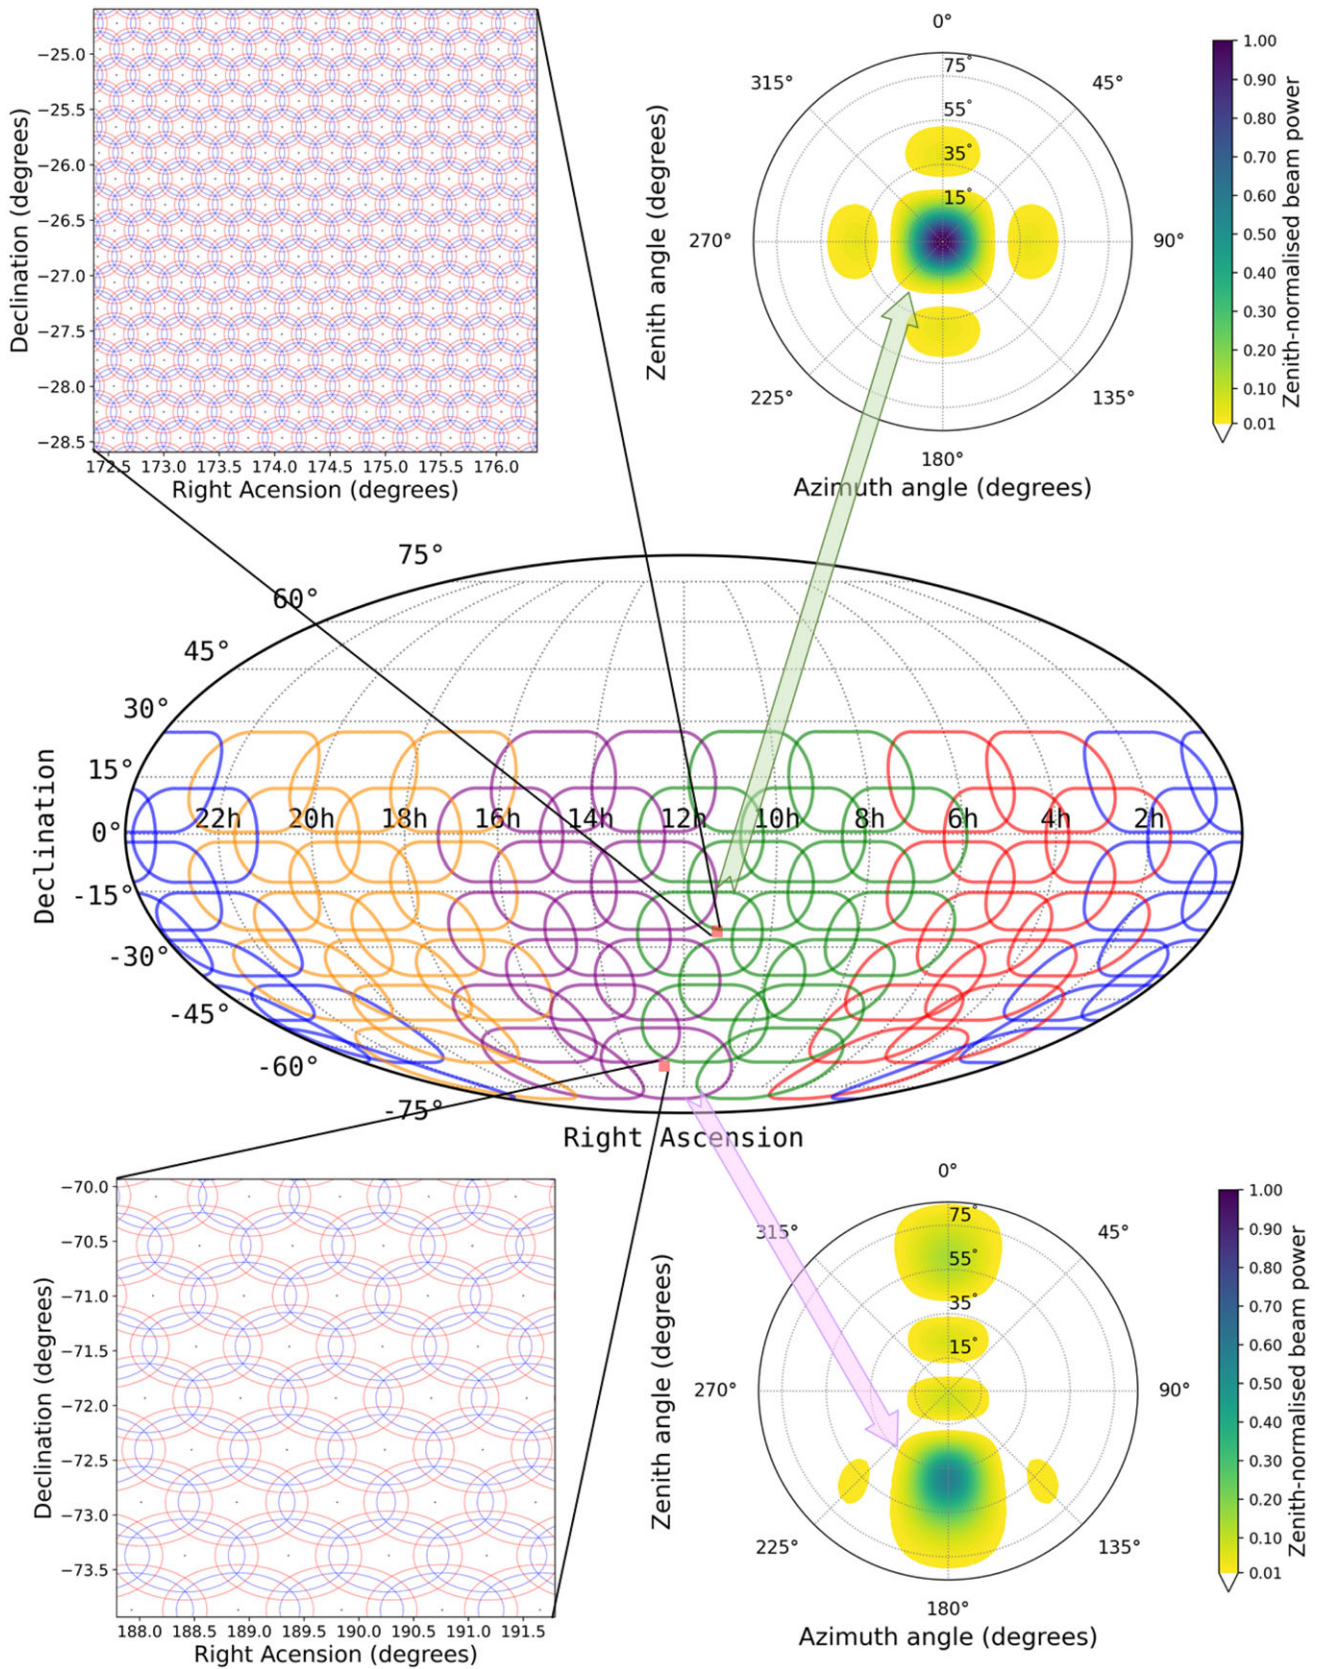
The precise size and shape of the TABs is a non-trivial function of the tile layout of the compact configuration, equivalent to the ‘Compact robust 1’ synthesised beam whose cross section is presented in Figure 7 of Wayth et al. (2018) and discussed in Swainston et al. (2022) and in Section 3. Due to the compact configuration’s redundant baselines (in the two ‘hexes’), the most sensitive parts of the TAB consist of a main lobe whose full width half maximum (FWHM) at 155 MHz is  $23'$ , surrounded by a pattern of discrete grating lobes of similar width. Although these grating lobes can be exploited for candidate confirmation (further discussed below), we choose the TAB pointings to form a dense (hexagonal) grid such that the main lobes overlap by  $\sim 20\%$ , as shown in Figure 1. This effectively Nyquist-samples the sky at a gain of the half-power level or more. The beam shape used for this calculation assumes that all 128 tiles are functioning, whereas, in reality, up to  $\sim 10\%$  of tiles may be flagged in any given observation. Unless the flagged tiles preferentially result in a reduction of the longest baselines, the effect on the beam shape is negligible.

Tiling the FoV in this way translates to  $\sim 6300$  TABs for an observation pointed toward the zenith. For pointings away from the zenith, where the beam shape develops a significant ellipticity (e.g. at zenith angle  $15^\circ$ , ellipticity  $\epsilon = \theta_{\text{maj}}/\theta_{\text{min}} = 1.36$  where  $\theta_{\text{maj}}$  and  $\theta_{\text{min}}$  are the major and minor axes of the TAB), the number of TAB pointings are in the range  $\sim 4200$ – $4500$ . Further, the beam size also varies across the 20% fractional bandwidth of our survey observations; for example, for a pointing toward the zenith (where the TAB is nearly symmetrical), the FWHM is  $25.3'$  at 140 MHz but reduces to  $20.7'$  at 171 MHz. This further justifies our rationale for a 20% overlap, as it ensures every single spot in the sky is covered at a gain near or above the half power level even at the high end of the observing band.

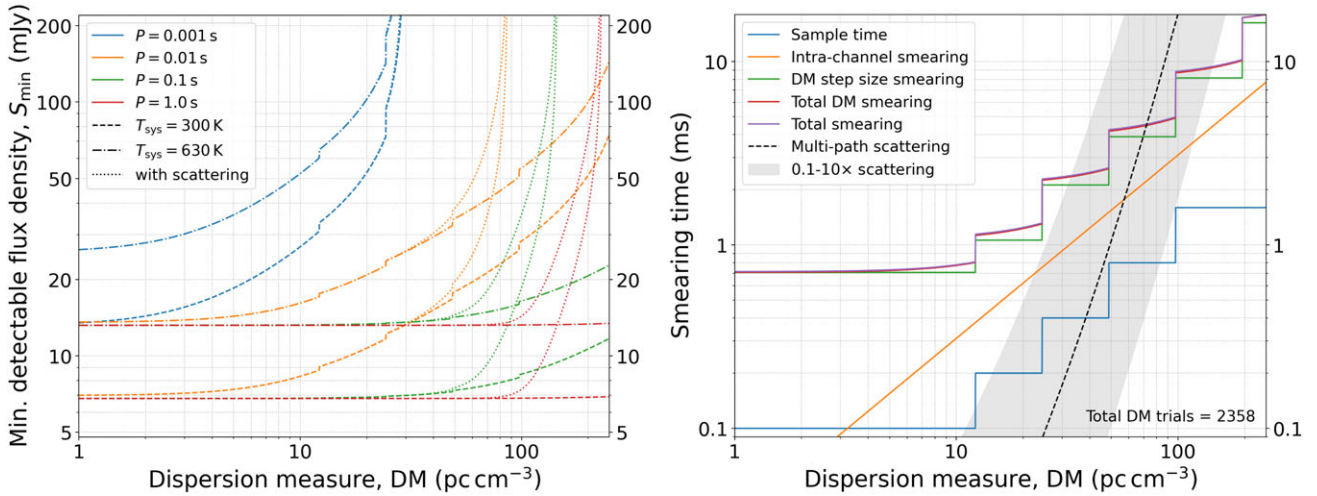
Finally, as with any other aperture array, the sensitivity is not uniform across the sky and is strongly declination-dependent; to first order, the loss in sensitivity is by a factor  $\cos(\theta_z)$  where  $\theta_z$  is the zenith angle. In principle, this can be compensated to a certain extent by longer integrations, though in practice, the inherent limitations of our data recording system (VCS) limits this to no more than 90 min, and we therefore use 80 min recordings for all pointings. As such, the sensitivity will not be uniform across the sky due to other factors; for example, the sky background temperature  $T_{\text{sky}}$  is direction dependent, and the loss in sensitivity from severe pulse broadening for distant pulsars, which applies to the

<sup>b</sup>The operational constraints of the MWA limited VCS mode observations to a maximum of 25 h per observing semester, with the legacy system.





**Figure 1.** Sky tessellation of the SMART survey. The left panels show beam tiling patterns for two select pointings: top one a near-zenith pointing ( $\delta = -28^\circ$ ), the bottom one a far southern pointing ( $\delta = -70^\circ$ ). The number of tied-array beams vary from  $\sim 6000$  to  $\sim 8000$  from near-zenith to far-zenith pointings, and the beam shape becomes elliptical at large offsets from the zenith. The size of the circle/ellipse indicates half power tied-array beam size; the red and blue circles correspond to the low and high ends of the SMART band (140–170 MHz). The right panels show the primary beam response for the same declination pointings, at the central frequency of 155 MHz.



**Figure 2.** *Left:* Minimum detectable flux density,  $S_{\min}$ , for the first-pass processing of the SMART survey as a function of DM. Sensitivity limits, assuming a 10-min integration time, are plotted for different pulse periods,  $P = 1.0, 0.1, 0.01, 0.001$  s, and for two different system temperature values  $T_{\text{sys}}$ ; one corresponding to mean  $T_{\text{sky}}$  for regions away from the Galactic plane, and the other for a mean  $T_{\text{sky}}$  in the plane, but excluding the region toward the Galactic Centre. The effect of pulse broadening due to interstellar scattering (Bhat et al. 2004) is shown by the dotted lines. *Right:* Pulse broadening (smearing) incurred by using the first-pass processing dedispersion plan (Table 2) due to various factors such as the finite sampling time, dispersive smearing due to the incoherent de-dispersion algorithm used, and the effects of multi-path scattering based on the  $\tau_d$ -DM relation from Bhat et al. (2004). The grey shaded region denotes one order of magnitude larger or smaller range in the predicted scattering.

sight lines within the Galactic plane or toward the Galactic centre. Some of these are considered in detail in Section 2.4.

### 2.4. Survey sensitivity

The sensitivity of a pulsar survey is determined by the combination of some instrumental and processing parameters and a variety of broadening effects to pulsar signals. Following Dewey et al. (1985), the minimum detectable flux density for a pulsar with period  $P$  and effective pulse width  $W_{\text{eff}}$ , down to detection significance  $(S/N)_{\min}$ , that is, minimum detectable signal-to-noise ratio, is related to the telescope gain  $G$  and system temperature  $T_{\text{sys}}$ , which is the sum of the receiver and sky background temperatures, that is,

$$S_{\min} = \frac{(S/N)_{\min}(T_{\text{recv}} + T_{\text{sky}})}{G\sqrt{n_{\text{pol}}t_{\text{obs}}B_{\text{obs}}}} \sqrt{\frac{W_{\text{eff}}}{P - W_{\text{eff}}}} \quad (1)$$

where  $n_{\text{pol}}$  is the number of polarisations summed,  $t_{\text{obs}}$  is the integration time and  $B_{\text{obs}}$  is the recording bandwidth. As evident from this equation, the sensitivity is maximum for long-period pulsars with a small duty cycle, i.e. when  $W_{\text{eff}} \ll P$ . The gain  $G = A_{\text{eff}}/2k_B$ , where  $A_{\text{eff}}$  is the effective collecting area and  $k_B$  is the Boltzmann constant. At 150 MHz,  $A_{\text{eff}} \approx 2750 \text{ m}^2$  for a 128-tile MWA (Tingay et al. 2013), which may imply  $G \sim 1 \text{ K Jy}^{-1}$ , however, for an aperture array such as the MWA, it is a strong function of the zenith angle, i.e.  $G(\theta_z) = G_{\text{max}}\cos(\theta_z)$ , where  $\theta_z$  is the zenith angle and  $G_{\text{max}}$  is the gain at  $\theta_z = 0$ . Moreover, for drift-scan type observations that we employ for the SMART,  $G$  depends on the offset from the phase centre, and can be  $\sim 0.5 G_{\text{max}}$  at the half power point. We therefore assume a conservative  $G \sim 0.5 \text{ K Jy}^{-1}$  for all our sensitivity calculations. This is assuming a full coherent beam sensitivity, that is, perfect calibration for TAB formation and no loss of sensitivity due to flagged tiles. In practice, a small number of tiles ( $\lesssim 10$ ) are typically flagged due to malfunctioning, sub-optimal performance or poor calibration solution. As we detail in Section 3, the strategy of observing multiple calibrators for SMART observation allows us to perform useful cross-checks

and maximise the achievable sensitivity using the best available calibration solutions.

At the low frequencies of the MWA, the system temperature  $T_{\text{sys}}$  is dominated by the sky background  $T_{\text{sky}}$ . Both  $T_{\text{recv}}$  and  $T_{\text{sky}}$  are frequency-dependent, and  $T_{\text{sky}}$  is also a strong function of the direction  $(l, b)$ , where  $l$  and  $b$  are the Galactic longitude and latitude, respectively. We assume a mean  $T_{\text{recv}} = 50 \text{ K}$  for the 140–170 MHz band. Excluding a  $\sim 10^\circ$  cone around the Galactic centre,  $T_{\text{sky}}$  can vary from  $\sim 200 \text{ K}$  toward  $|b| \gtrsim 60^\circ$  to as much as  $\sim 1200 \text{ K}$  in the plane, toward  $\gtrsim 10^\circ$  from the Galactic centre, where  $T_{\text{sky}}$  can be as large as  $\sim 10^4 \text{ K}$  at 155 MHz. We use the Haslam et al. (1982) map as the reference and assume  $T_{\text{sky}} \propto \nu^{-2.55}$  scaling from Lawson et al. (1987). Given this strong dependence of  $T_{\text{sky}}$  with  $(l, b)$ , we consider two cases: (1) the sky at  $|b| \lesssim 5^\circ$  where mean  $T_{\text{sky}} \sim 600 \text{ K}$  and (2) the sky at  $|b| \gtrsim 5^\circ$ , where mean  $T_{\text{sky}} \sim 270 \text{ K}$ ; that is,  $T_{\text{sys}} = 630$  and  $300 \text{ K}$ , respectively, as shown in Figure 2.

Intrinsic pulses are broadened due to a variety of effects, as discussed earlier. As detailed in Lorimer & Kramer (2012), the total smearing time  $\tau_{\text{tot}}$  is the quadratic sum of the finite sampling time  $\tau_{\text{samp}}$ , the residual dispersive smearing due to finite frequency channel  $\tau_{\text{chan}}$ , the dispersive smearing across the full recording band due to finite DM steps in trial DM values  $\tau_{\text{dm}}$ , and the dispersive smearing resulting from piece-wise linear approximation of the quadratic dispersion law in the sub-band dedispersion algorithm employed in searches  $\tau_{\text{sub}}$ . Figure 2 summarises these for our current first-pass processing. The planned second-pass search will significantly enhance the search sensitivity by processing the full observation (4800 s) and the use of more optimal DM steps, i.e. many more trial DMs than that used in current search.

As evident from the figure, for our current first-pass processing, the total smearing time is dominated by finite DM steps; this sub-optimal choice was made in an effort to maximise the number of observations that can be processed to completion toward a shallow all-sky survey within available computing resources. The dedispersion plan utility used is shown in Table 2. In effect, we progressively downsample the data five times over the DM range searched, each time making the DM step size coarser. At  $\text{DM} \gtrsim 3$



**Table 2.** Dedispersion plan for the first-pass SMART processing.

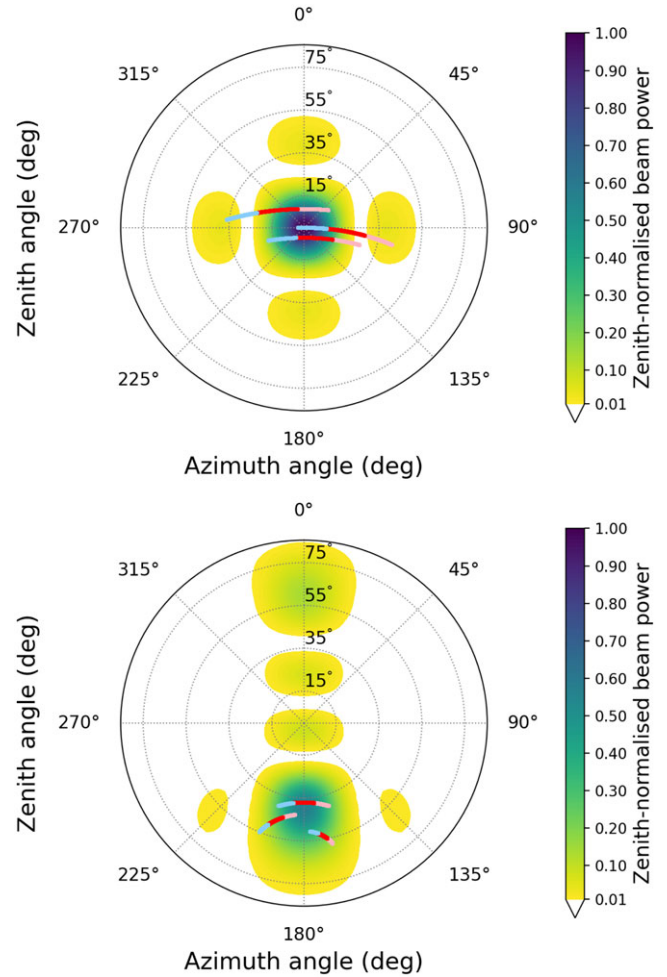
$DM_{\min}$ ( $\text{pc cm}^{-3}$ )	$DM_{\max}$ ( $\text{pc cm}^{-3}$ )	$\delta DM$ ( $\text{pc cm}^{-3}$ )	$N_{DM}$	$d_s$	$\Delta t_{\text{eff}}$ (ms)
1.0	12.2	0.02	560	1	0.1
12.2	24.4	0.03	406	2	0.2
24.4	48.8	0.06	406	4	0.4
48.8	97.6	0.11	443	8	0.8
97.6	195.2	0.23	424	16	1.6
195.2	250.0	0.46	119	16	3.2

The columns 1 and 2 denote the ranges in dispersion measure, between  $DM_{\min}$  and  $DM_{\max}$ , with a DM step size of  $\delta DM$ , resulting in  $N_{DM}$  trial DM values. The down sampling factor is denoted by  $d_s$ , i.e. the factor which the temporal resolution is averaged to yield a net resolution  $\Delta t_{\text{eff}}$ .

$\text{pc cm}^{-3}$ , the dispersive smearing time within the 10-kHz channel is larger than the native sampling time (100  $\mu\text{s}$ ) but still a smaller contribution to the total smearing time, compared to that due to the DM step size. As a result, the net smearing time  $\tau_{\text{tot}}$  displays a step-wise increase as shown in Figure 2, given our dedispersion plan. At very low DMs  $\lesssim 10 \text{ pc cm}^{-3}$ ,  $\tau_{\text{tot}} \sim 0.7 \text{ ms}$  but increases to  $\sim 10 \text{ ms}$  at  $DM \sim 100 \text{ pc cm}^{-3}$ . In essence, our first-pass search severely compromises the sensitivity to millisecond pulsars (MSPs) at larger DMs and shorter periods, that is, it is currently sensitive to MSPs at  $DM \lesssim 30 \text{ pc cm}^{-3}$  and  $P \gtrsim 10 \text{ ms}$ . As shown in Figure 2, at those larger DMs, the smearing due to scattering (i.e. pulse broadening) can also be significant. The broadening time here is based on the empirical relation in Bhat *et al.* (2004), which is mostly relevant for pulsars near the plane. As is well known, these scattering estimates can be uncertain by more than an order of magnitude, denoted by the grey shaded region.

The theoretical sensitivity is shown in Figure 2 for different periods,  $P = 1.0, 0.1, 0.01$  and  $0.001 \text{ s}$ . In all these calculations, we have assumed a duty cycle of 3%, that is,  $W_{\text{eff}}/P = 0.03$ . For each period, a pair of curves are shown: one for the best-case scenario, that is, searches away from the plane, where  $T_{\text{sys}} \sim 305 \text{ K}$ ; and the second for the sky near the plane where the mean  $T_{\text{sys}}$  is twice as high. In either case, the sensitivity is maximum for long-period pulsars, at low to moderate DMs of  $\lesssim 50 \text{ pc cm}^{-3}$ , and toward  $|b| \gtrsim 5^\circ$ .

With our first-pass processing scheme (i.e. 10-min integrations and a sub-optimal dedispersion plan), we reach a limiting sensitivity of  $S_{\min} \sim 7\text{--}12 \text{ mJy}$  for long-period pulsars, and  $\sim 12\text{--}25 \text{ mJy}$  for MSPs at low to moderate DMs. For the proposed deep-pass processing (i.e.  $\sim 80\text{-min}$  integrations and a more granular dedispersion plan), we can achieve a limiting sensitivity of  $S_{\min} \sim 2\text{--}3 \text{ mJy}$  for long-period pulsars and  $\sim 5\text{--}10 \text{ mJy}$  for MSPs at low to moderate DMs. In this case, the SMART survey sensitivity is comparable to that of the LOTAAS survey in the northern hemisphere. While LOTAAS can be twice as sensitive as SMART for long-period pulsars, the sensitivity for  $P \lesssim 10 \text{ ms}$  is almost similar, owing to a lower degradation in sensitivity in the SMART band. Compared to the Southern Pulsar Survey of the 1990s at 430 MHz (i.e. a wavelength of 70 cm), also known as the Parkes 70cm survey (Manchester *et al.* 1996), the SMART survey is  $\sim 3\text{--}5$  times more sensitive, especially for pulsars at DMs  $\lesssim 100 \text{ pc cm}^{-3}$  and spectral index  $\alpha \lesssim -1.5$ . Even the ongoing shallow survey is comparable to the 70 cm survey in theoretical sensitivity, and if at all, slightly more sensitive to steep spectrum pulsars with no turnover down to  $\sim 100 \text{ MHz}$ . This provides a strong motivation to undertake a full-scale pulsar survey with the MWA.



**Figure 3.** Tied-array beam traces through the MWA primary beam for SMART observations. Three example pointing directions for each observation are traced including 1 h before and 1 h after the 80-min observation. The target trace (rotating clockwise as time advances) is coloured pink to represent the trajectory before the observation, red during the observation, and blue after the observation is complete. North is at  $0^\circ$  and the azimuth angle increases to the East. The colour scales are the same for each subplot, highlighting the sensitivity penalty incurred for observing away from zenith.

### 2.5. Effective dwell time and sensitivity

Unlike most other pulsar surveys, where single-dish telescopes are used to track targeted positions for small time intervals (e.g. HTRU, GBNCC), the SMART observations are drift scans, where the primary beam is pointable but static in horizontal coordinates (azimuth and zenith angle) once an observation starts and the sky moves through the FoV. When forming TABs, we track the sky position as it moves through the MWA primary beam and as a consequence not all TABs necessarily remain within a sensitive part of the primary beam for the full 80-min duration.

The amount of time spent within an individual observation FWHM depends both on the observing declination (i.e. where the primary beam is pointed) and the target source position to be tracked with a TAB. As an example, in Figure 3 we plot some representative TAB pointings along with the primary beam response for the same observations as in Figure 1 in horizontal coordinates. As already noted, our sensitivity drops substantially as we observe at larger zenith angles, which we visualise by having the colour scale represent the zenith-normalised primary beam power as a



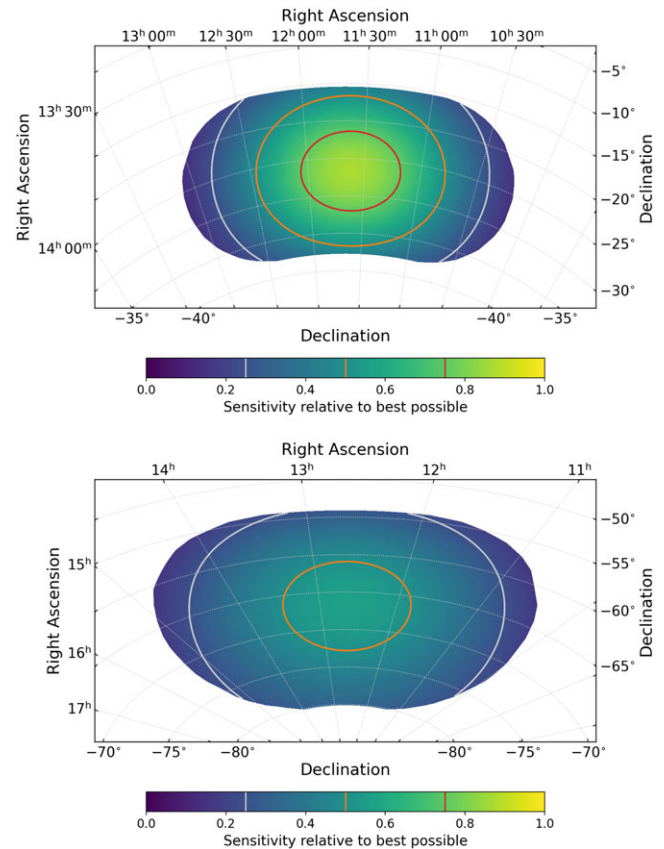
proxy for sensitivity. Secondly, the TAB pointing directions are traced before, during, and after the 80-min observation, which highlights that not all targeted positions remain in a usable part of the primary beam. These effects highlight at least three points for consideration: (1) it will be an inefficient use of resources to track certain pointings for the full observing duration, (2) tracking pointings naively for the full duration, especially if a significant fraction of the time is spent below the 10% power point, may actually reduce sensitivity to pulsars, and (3) the full-sky sensitivity will be patchy regardless of TAB forming strategies (although this is partly mitigated by having observations overlap by  $\sim 20\%$  at the central frequency). To address (1) we can estimate the time a source remains in a reasonable power range of the primary beam and only form TABs from the appropriate subset of voltages recorded (e.g. while the target source is not in a null). For (2) we must strike a careful balance between achieving maximal sensitivity (by cutting off parts of the TAB) and dwell time (which benefits searches for longer-period pulsars and single pulse events). The consequence of (3) is unavoidable given the telescope configuration and observing strategy employed, but is quantifiable.

We can evaluate the relative sensitivity (assuming a 80-min track for a given TAB) by summing the primary beam response power at discrete time steps, where we use our current best Full Embedded Element (FEE) model (Sokolowski et al. 2017), normalised to the equivalently summed power that would represent the best possible dwell time and sensitivity combination. For our purposes we define this quantity as the sum of the primary beam power at zenith for the full observing duration (i.e. imagining we can track an equatorial position with full zenith sensitivity). This is useful as it scales the effective sensitivity to a quantity close to what a single-dish steerable telescope could achieve. In Figure 4 we present these effective sensitivity maps, in equatorial coordinates, for the same example observations used in Figure 3.

### 3. Data processing and analysis

In terms of data collection and processing requirements, the SMART survey is the largest all-sky pulsar survey undertaken in the southern hemisphere, and is only the second largest after LOTAAS. The SMART survey will accrue  $\sim 3$  PB of VCS data, compared to  $\sim 1$  PB (search mode data) by the highly successful Parkes HTRU survey, and  $\sim 8$  PB (beamformed data) by LOTAAS. As outlined earlier, the survey will cover the sky in 70 VCS pointings, each VCS observation being 4800 s (42 TB). The management and processing of this volume of data is non-trivial, particularly considering the computational resources currently available. The processing software and pipelines are developed, tested and benchmarked on Pawsey's Galaxy/Garrwarla clusters, and subsequently ported and benchmarked on Swinburne's OzSTAR supercomputer. The time on OzSTAR is secured via the merit allocation scheme under Astronomy and Supercomputer time allocation, and is typically 0.5–0.6 million service units (CPU core) hours per annum. These constraints largely drive the initial processing strategies, thereby necessitating a first-pass shallow survey.

Compared to the HPC resources available at Pawsey, the processing efficiency has been relatively higher on OzSTAR, where the current benchmarks are 2 kSU for beamforming and 25 kSU for searching a 10-min observation (4.4 TB), where 1 kSU = 1000 service units (CPU core hours). The current allocation thus allows



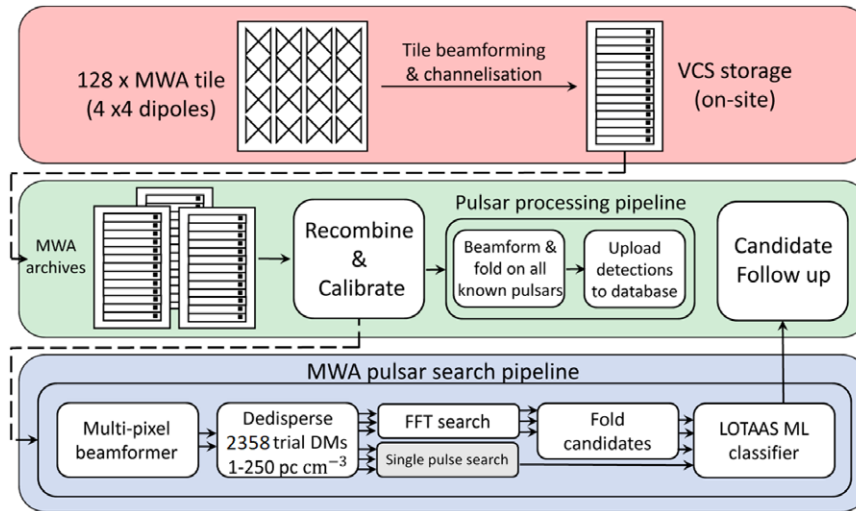
**Figure 4.** Effective sensitivity maps, assuming a full 80-min tracking and integration for a given TAB sky position. The colour map is normalised to the best possible sensitivity (described in the text), and contours at 25, 50, and 75% are drawn for clarity. Due to the drift scan nature of the observations versus the tracking TABs, we can never achieve the best possible sensitivity. Right Ascension and Declination are marked by the vertical and horizontal curved grid lines, respectively.

processing of 9 observations (fields) per semester, where each 10-min VCS observation is processed for  $\sim 6000$  tied-array beams, each of which is then searched in 2358 trial DMs, out to  $250 \text{ pc cm}^{-3}$ . The completion of first-pass processing will thus require  $\sim 2$  million core hours. Scaling from the current benchmarks, we would thus expect 1500 kSU per full observation for deeper searches, and 60 million core hours for full DM searches ( $\sim 10000$  searches, for a max DM of  $250 \text{ pc cm}^{-3}$ ), necessitating the integration of GPU-based search processing in the future.

An overview of the processing pipeline is presented in Figure 5, the details of which are described in the sections below. In essence, this involves preprocessing and beamforming of voltage data from 128 tiles of the array to generate beamformed time series, before the data can be processed through the search and detection pipelines. The main steps are outlined below.

#### 3.1. Preprocessing and beamforming

The main step in the preprocessing stage involves processing VCS data so they can be calibrated and coherently combined to produce beamformed time series at the native resolution of  $100\text{-}\mu\text{s}/10\text{-kHz}$  of the VCS. The array calibration is performed using one of the standard calibrators (e.g. 3C444), recorded in the visibility mode at the default  $0.5\text{-s}/40\text{-kHz}$  resolution, where complex gain solutions



**Figure 5.** Workflow diagram illustrating the first-pass SMART processing pipeline: voltage data at 100- $\mu$ s/10-kHz resolutions are recorded from 128 tiles of the array after tile beamforming and channelisation stages, and are subsequently ported to the Pawsey supercomputer where the initial processing including calibration, beamforming and known pulsar detections are carried out. Search processing is currently performed on the OzSTAR supercomputer, and is limited to basic periodicity searches.

(amplitude and phase) are obtained for each of the 128 tiles, for every coarse channel (1.28 MHz wide), using the Real Time System (RTS) software package. The procedure is essentially similar to those employed for other VCS observations (e.g. Swainston *et al.* 2021). The calibration solutions can then be used to coherently combine the voltage data in phase using the tied-array beamformer, the conceptual details and implementation of which are detailed in Ord *et al.* (2019). The functionality was enhanced, and GPU parallelised, in preparation for SMART data processing (Swainston *et al.* 2022).

The beamformed data are written as Stokes I at 100- $\mu$ s/10-kHz resolutions. The current implementation allows processing 120 coarse-channel beams at once, that is, 5 full-bandwidth (30.72 MHz) beams, resulting in a data rate of 87 GB beam<sup>-1</sup> for a 10-min observation. For each survey pointing, this amounts to  $\sim$ 500 TB in beamformed data. These data are equivalent to that would emerge from the standard pulsar backends and so can be processed using standard pulsar search packages. Typically, data would be processed to generate radio frequency interference (RFI) masks; however the superb radio-quiet environment at the telescope site and preferential observing during the nightly hours (and within an hour of the source transit) make this step not essential for the SMART data. In most cases, data are minimally affected by RFI, and consequently no RFI-related processing is carried out in the ongoing first-pass processing.

The large FoV of the MWA means excellent prospects for detecting multiple known pulsars within each pointing, which is also important for crucial data quality checks and initial assessment of array calibration and tied-array sensitivity. In short, each SMART observation is processed for known pulsars within the primary beam ( $\sim$  610 deg<sup>2</sup>), using a custom pulsar detection pipeline.

### 3.2. Search pipeline

The current SMART pipeline includes a GPU-based pipeline for front-end processing (beamforming) and a CPU-based pipeline for downstream (search) processing. The search pipeline is based

on the Pulsar Exploration and Search Toolkit<sup>c</sup> (PRESTO; Ransom 2001, 2011) pulsar search software suite, with the addition of machine-learning (ML) tools adopted from the LOTAAS classifier (Tan *et al.* 2018a). This was adopted as a first-pass processing strategy, to ensure an end-to-end working pipeline from the data collection and reordering stage (occurring at the observatory site) to array calibration/quality checks (Pawsey) and search processing (OzSTAR). To encapsulate the full-search workflow, we make extensive use of Nextflow<sup>d</sup> (Di Tommaso *et al.* 2017) to manage data input, output, processing tasks, and intermediate or final product creation and tracking.

In the near future, as we transition to full sensitivity searches, the search component will be replaced by a GPU-based implementation. Here we present a detailed breakdown of the current SMART search pipeline, where 10-min data (4.8 TB) are processed from each observation.

#### 3.2.1. Dedispersion and periodicity search

The beamformed data are processed to create dedispersed time series for each beam. As mentioned earlier, for the first-pass processing, maximum DM searched is 250 pc cm<sup>-3</sup>. At higher DMs, scattering can be significant; for example, pulse broadening times  $\gtrsim$ 100 ms are expected at 155 MHz for sight lines toward  $|b| \lesssim 5^\circ$ , and  $l \gtrsim 330^\circ$  or  $l \lesssim 30^\circ$ , where such high DMs can be expected. Further, even with 10-kHz channels, DM smearing can still be significant at low frequencies. For instance, at a frequency of 140 MHz (i.e. the low end of the SMART band), intra-channel dispersion smearing is  $\sim$ 1.5 ms at DM = 50 pc cm<sup>-3</sup>, and  $\sim$ 10 ms at DM  $\sim$  250 pc cm<sup>-3</sup>. The dedispersion plan was created using the PRESTO `DDp1an.py` utility, but with the caveat that sub-optimal settings were chosen (the use of coarser DM steps) to limit the number of DM trials to 2358, given the limitation of computational resources. The `prepsubband` tool from PRESTO was used to create incoherently dedispersed time series from the PSRFITS

<sup>c</sup>See <https://github.com/scottransom/presto>.

<sup>d</sup>See <https://github.com/nextflow-io/nextflow>.

(i.e. search mode) files. It makes use of the sub-band dedispersion technique, which uses a piece-wise linear approximation to the quadratic dispersion relation. The dedispersion plan employed in the first-pass search is shown in Table 2.

Searching for periodic signals involves computing the power spectra of the dedispersed time series, which is performed using the `realfft` tool within PRESTO, by applying Fourier transform techniques. These power spectra are then searched for periodicities using `accelsearch` (Ransom, Eikenberry, & Middleditch 2002), which detects the most significant periodic signals and uses harmonic summing to recover the power spectra at multiples of a given spin frequency. No acceleration searches are performed in this first pass; that is, searches are only performed at zero acceleration. Acceleration searches would require significant processing cost, given the large data rates, and the number of trial DMs required, but will be part of the second-pass search. If the significance of any spectral bin is in excess of  $2\sigma$ , it is marked as a candidate and the corresponding harmonics up to the 16th are summed to increase the detection significance.

A sifting procedure is then performed on the list of candidates from all 2358 DM trials. We adopt a fairly standard procedure, quite similar to that followed for LOTAAS, where candidates with  $P < 1$  ms or  $P > 30$  s are rejected,<sup>e</sup> as well as those with  $DM < 1$  pc cm<sup>-3</sup>. Candidates with similar DMs and harmonically related periods are then grouped, and only the instance with the highest S/N is kept. From this reduced candidate list, only those with  $\gtrsim 5\sigma$  detections are then folded.

### 3.2.2. Candidate folding

Folding of the candidates is performed using the `prepfold` tool, which creates the associated candidate files and standard diagnostic plots such as those shown in Figure 6. Since our pipeline uses the LOTAAS classifier, the folding analysis is carried out using the identical parameter setup as in the LOFAR search pipeline; that is, 100 pulse phase bins, 256 sub-bands, 120 sub-integrations for  $P > 10$  ms, whereas 50 pulse phase bins and 40 sub-integrations for  $P < 10$  ms. With this, the folded candidate information can be classified and processed using the ML classifier that we have adopted from the LOFAR search.

### 3.2.3. Single pulse search

Single pulse searches have proven to be effective for detecting the class of pulsars that emit sporadically (e.g. RRATs, and giant-pulse emitters such as the Crab). The basic algorithm involves trialling a range of box car widths,  $2^n t_{\text{samp}}$ , where  $t_{\text{samp}}$  is the sampling time resolution (100  $\mu$ s for SMART) and  $n = 0, 1, 2, \dots, N$ , where  $N$  corresponds to the maximum width searched (e.g. Cordes & McLaughlin 2003), and detecting ‘events’ that are above a set threshold. It is not computationally demanding, and is routinely performed in most pulsar searching. The pipeline has been tested using a SMART observation containing the Crab pulsar, and has also yielded a blind detection of a LOFAR-detected RRAT J0301+20 (Michilli et al. 2018). Integrating this into the processing chain is part of our second-pass search strategy.

<sup>e</sup>This period range was adopted given the minimum and maximum period of known pulsars in the ATNF pulsar catalogue when our processing commenced, which was 1.3 ms and 23.5 s, respectively.

## 3.3. Pulsar detection pipeline

### 3.3.1. ML classification of candidates

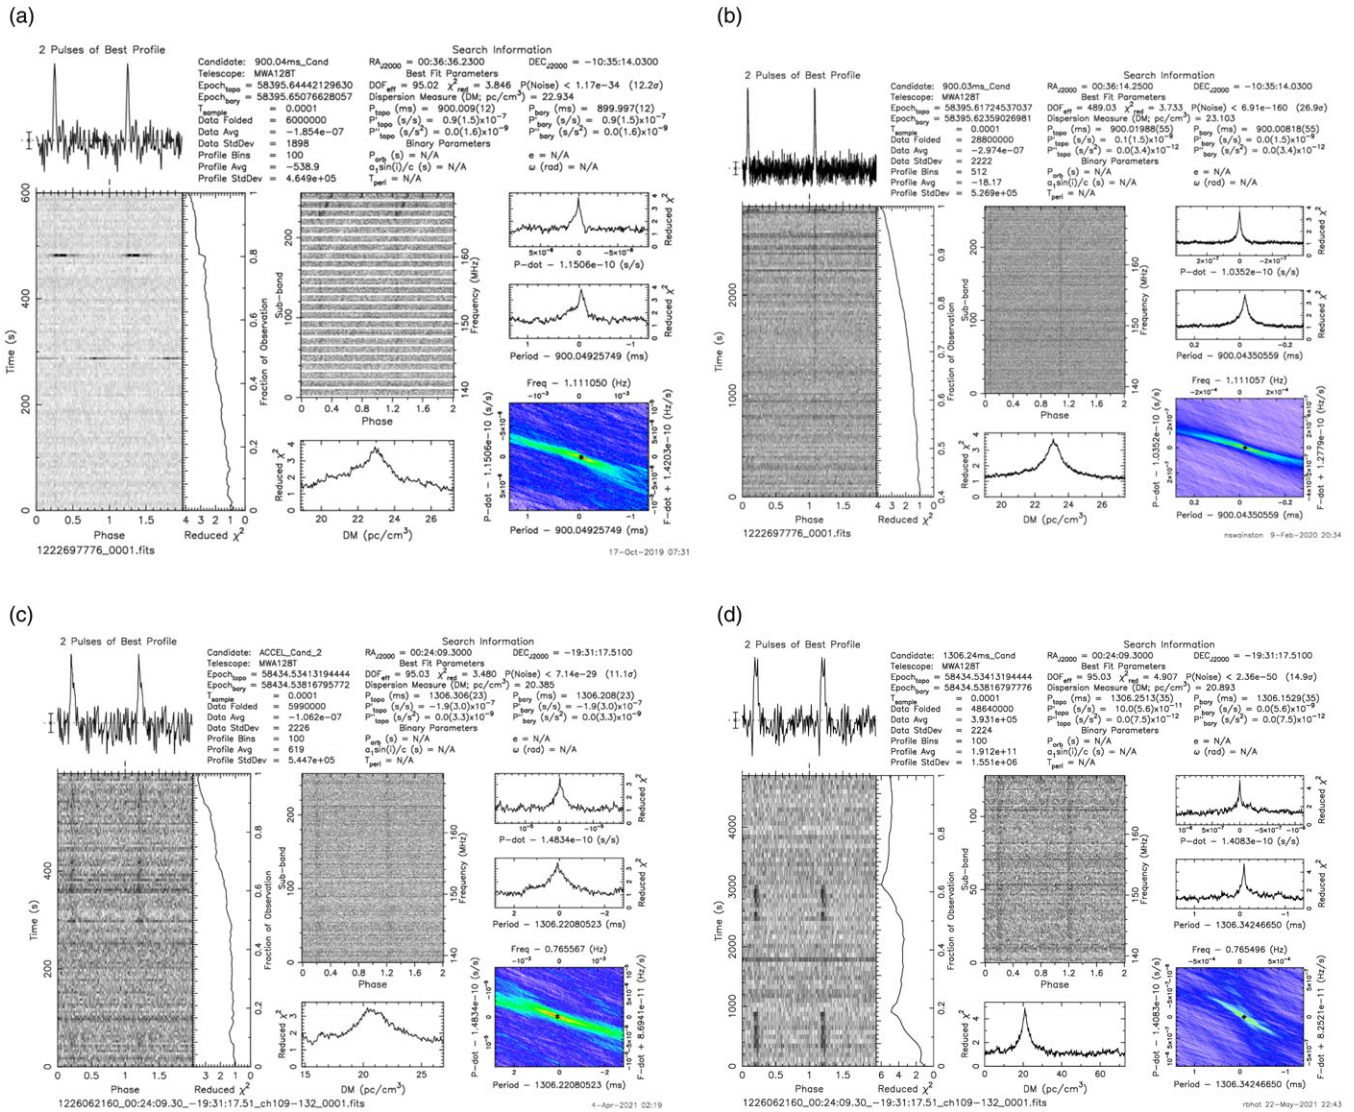
For each VCS survey pointing ( $\sim 610$  deg<sup>2</sup>, which is tessellated to  $\sim 6000$ – $8000$  beams), the processing typically results in  $\sim 135000$  candidates. Scaling for a significantly larger sensitivity ( $3\times$ ) and a larger number of DM trials ( $\sim 4\times$ ) anticipated in full-scale deep searches in the second pass, we may expect over 50 million candidates. Even for the first pass, as many as 9 million candidates can be expected, extrapolating the rate of candidates requiring scrutiny from the current pipeline. Indeed, visual inspection of that many candidates is unrealistic, thus necessitating the use of ML classifiers.

As an initial strategy, we have adopted the ML software that was developed for LOTAAS. The algorithm used is described in Lyon et al. (2016) and Tan et al. (2018a), and is summarised here. The classifiers use the statistics of the pulse profile (i.e. mean, variance, skewness and kurtosis) and the DM curve (i.e. S/N vs DM; see Figure 6). As described in Tan et al. (2018a), this basic approach is expanded by also calculating the correlation coefficient between each sub-band of the profile, as well as correlation coefficients between each sub-integration and the profile. In effect, the classifier uses the statistics of correlation coefficient distributions, in addition to the statistics of the profile and the DM curve, in order to classify the periodicity candidates. Four standard models are used for the regression: (1) decision tree algorithm, (2) multilayer perceptron, (3) probabilistic Bayes classifier, and (4) linear support vector machine.

Even without being trained on MWA data, the software performs reasonably well, with a recall rate of  $\sim 83\%$  for the worst-performing regression model. While clearly not optimised for a MWA search, it can still provide a significant cull on the number of candidates that require human scrutiny as long as the number of false negatives is kept below an acceptable threshold. To minimise the false negative rate, we use the provided ‘ensemble’ classifier, which labels candidates as positive if at least three models classify them as positive. Under this criterion, the number of candidates is cut down from the original  $\sim 135000$  per pointing down to  $\sim 20000$  that require human scrutiny, that is, an efficiency of  $\sim 85\%$ . The false negative rate can be lowered further by allowing candidates classified as pulsars by a smaller number of regression models to be passed, but this comes at the cost of also lowering the efficiency. For the first-pass processing, we find the current arrangement to be an acceptable compromise, but will be implementing an improved ML classifier for the second pass.

Of the remaining  $\sim 20000$  candidates per pointing, only a small fraction are true pulsar detections, with the vast majority of candidates consisting of noise and RFI. Here, we are extending the definition of RFI to include any artefact from the MWA signal path that may result in spurious detections. Owing to the radio-quietness of the observatory site, such candidates belong almost exclusively to this category, and almost never arise from external sources. The most common RFI candidates are those with periods of either 1 s or with a close harmonic relationship (e.g. 0.5 s, 2 s), relating to the division of data packets by 1-s boundaries. Such candidates are sufficiently few (and easily identified) that we do not apply any automatic procedure for removing them from our pool of candidates.





**Figure 6.** Examples of standard PRESTO diagnostic plots of original periodic pulsar candidate detections (left panels), and improved detection plots from follow-up processing for confirmation (right panels). Upper panels are the first pulsar discovered from the SMART, PSR J0036–1033, and the lower panels are the second pulsar, PSR J0026–1955. Initial detections are from 10-min observing durations (first-pass processing), while the confirmation ones are from longer durations of the same initial detection observations.

3.3.2. *Prioritisation and scrutiny of candidates*

The candidates that survive the initial ML cull are still mostly dominated by noise and RFI detections, with only a small minority being true pulsar detections. Although all of these candidates are intended ultimately to be visually inspected, we have developed a so-called ‘clustering’ algorithm to prioritise which candidates get inspected first, in order to accelerate the detection of sufficiently bright, new pulsars.

The clustering algorithm leverages the fact that the tied-array beam of the MWA’s compact configuration is relatively complex, with significant grating lobes located in different parts of the primary beam. Because the spacing between tied-array pointings is equal to the FWHM of the main lobe of the tied-array beam, any sufficiently bright pulsar will likely be detected in multiple beams. For instance, Figure 7 shows a map of multiple detections of PSR B2327–20 superimposed on the theoretical sensitivity of each tied-array beam towards the pulsar, as predicted by the array factor

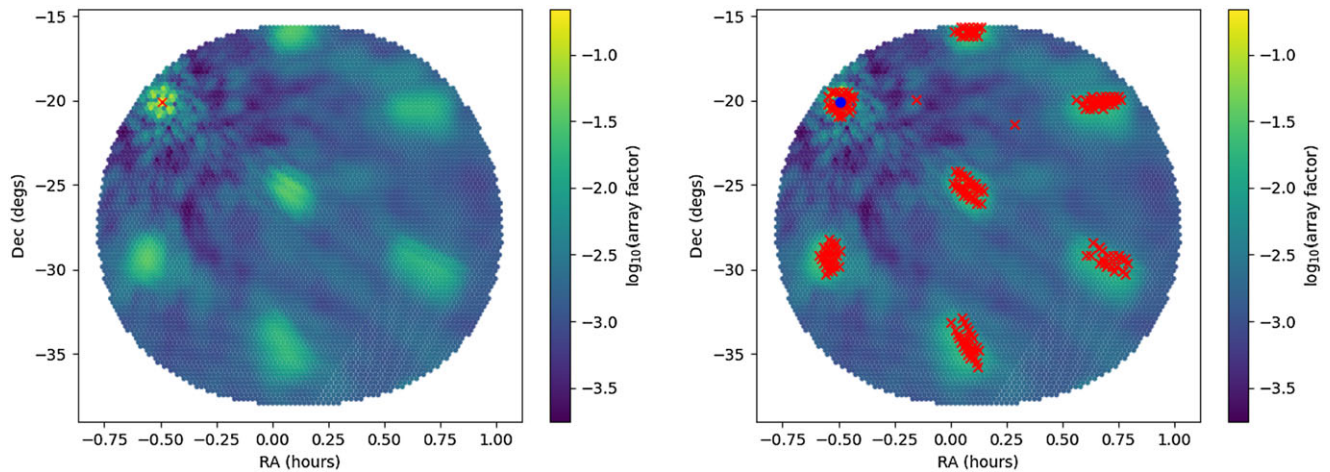
formalism developed for the MWA by Meyers et al. (2017). Since noise candidates will not be correlated across different beams, prioritising similar candidates that appear in multiple beams dramatically increases the likelihood that candidates representing true astrophysical signals will be inspected first.<sup>f</sup>

Candidates are considered similar if


1. they appear in at least two adjacent beams,
2. they have periods within 0.5% of each other, and
3. they have DMs within 3 pc cm<sup>-3</sup> of each other.

As a demonstration of the usefulness of the clustering algorithm, we show how it would detect PSR J0026–1955, the second pulsar discovery in the SMART survey (McSweeney et al. 2022). In

<sup>f</sup>This is counter intuitive to the case of multibeam surveys with Parkes-like single-dish telescopes, where similar candidates detected in multiple beams across the sky would indicate RFI.



**Figure 7.** The theoretical array factor (a proxy for sensitivity) of each tied-array beam towards the pulsar B2327–20, with the red cross marking the position of the pulsar (left panel) and the beams in which the pulsar was detected (right panel). SMART observation 1226062160 was used for the demonstration.



**Figure 8.** The theoretical array factor (proxy for sensitivity) in the vicinity of PSR J0026–1955 for observation 1226062160, assuming a true position (centre of image) derived from GMRT imaging (cf. Paper II for details). Red crosses mark the position of beams in which it was detected, and the blue dot marks the first detection. A single cross may indicate multiple detections with slightly different periods and DMs.

reality, the clustering algorithm was not implemented until *after* PSR J0026–1955 was discovered, but it is interesting to note that the first detection (chronologically) of this pulsar was a grating lobe detection (at the time, the candidates were being served up randomly), which motivated the development of the clustering algorithm in the first place.

The final set of detections of PSR J0026–1955 is shown in Figure 8, on a backdrop of the theoretical array factor (a proxy for sensitivity) towards the pulsar assuming that our current best-fit position is correct. In this case, three of the search beams contained the nominal pulsar position in the main lobe, while several others positioned the pulsar in their respective grating lobes. All of the displayed detections meet the second and third clustering criteria (similar periods and DMs). Therefore, any pair of detections in the same or adjacent beams are considered ‘clusters’, and if the clustering algorithm was in use when this observation was processed, this pulsar would have been picked up immediately in multiple clusters.

The clustering algorithm offers no advantage for relatively weak pulsars that would be detected only in a single (bore-sight) tied-array beam. Therefore, unclustered candidates are not deleted, only deprioritised.

### 3.3.3. Human inspection and ranking

Just as the clustering algorithm is a method for prioritising candidates for human inspection, so too is human inspection a method for prioritising candidates for follow-up (see Section 4). Users are served up candidates one at a time and presented with the candidate’s PRESTO diagnostic plots (e.g. Figure 6). Each candidate is given an integer rating from 1 to 5, with higher numbers corresponding to a higher confidence that the candidate is a *bona fide* pulsar detection. Clear pulsar detections are then compared to the ATNF catalogue of pulsars to check if it is a known pulsar. If a detection is unknown, candidates listed in other surveys are then checked using the Pulsar Survey Scraper tool.<sup>§</sup> If the pulsar is in either the ATNF catalogue or in another survey’s candidate list, a note is made against the candidate with the pulsar’s name, visible to all other users.

Each candidate can be ranked by multiple users (but users can only rank each candidate once). A candidate that has been rated by at least four users becomes eligible for follow-up, and the list of eligible candidates is ordered by the average rating.

Currently, as the number of users of the system is still relatively small, the rating of candidates is the primary bottleneck in the whole processing chain. This means that during first-pass processing, interesting candidates have been followed up immediately. In the future, however, as the number of users performing the task of rating candidates grows, the pool of eligible candidates may grow faster than the rate at which they can be followed up. However, the above system of candidate prioritisation means that the most interesting candidates are always followed up first.

## 3.4. Data management and web app

The large number of generated candidates, the complex metadata associated with them, and the desire to distribute the tasks of

<sup>§</sup>See <https://pulsar.cgca-hub.org/>.



data processing, candidate rating and candidate follow-up, motivated the implementation of a relational database to track the progress of the SMART survey and coordinate processing efforts. The database, implemented in PostgreSQL, is comprised of a set of tables containing metadata for

1. MWA observations (e.g. primary beams, tied-array beams, candidates);
2. software (e.g. for beamforming, searching, ML classification), including versioning information;
3. candidate ratings;
4. pulsars;
5. users; and
6. supercomputer facilities.

The users, along with their database access privileges and authentication, are managed by a subset of tables which interface with website front end implemented in Django. Both the database and the website are hosted by Data Central.<sup>h</sup>

Once an observation has been processed and the candidates have been subjected to the first-pass ML cull (Section 3.3.1), both the metadata of the remaining candidates as well as the candidates themselves (i.e. PRESTO.pfd files and the associated diagnostic plots) are uploaded to Data Central. The uploaded candidates are then available for users to rate via the web interface (Section 3.3.3).

As described above, candidates can then be sorted by their average rating, and followed up at will by any authorised user. Before following up a candidate, the user may ‘claim’ it by clicking a button in the candidate list. This feature is designed to prevent multiple users from following up the same candidate and unnecessarily duplicating effort. The decentralised design allows members of the SMART collaboration from different research institutions to work through the SMART data set without the need for someone to oversee and coordinate the different groups’ activities.

#### 4. Confirmation and initial follow-up of candidates

Confirmation and follow-up of promising pulsar candidates typically relies on multiple re-observations, often requiring a significant amount of telescope time. Fortunately, the SMART survey’s unique design, where VCS data are retained (unlike preprocessed beamformed data), offers flexible reprocessing options, allowing us to accelerate important confirmation and follow-up procedures. Furthermore, a substantial amount of archival VCS data (from past projects) are available for a large part of the MWA sky, which can also be suitably exploited for further detection and improved localisation. These features make the SMART survey distinct from other pulsar surveys.

In the following sections we outline the main strategies that are adopted for confirmation and initial follow-up, including: reprocessing of the original observation for improved detection; performing a dense grid for improved sky localisation; and polarimetry via reprocessing the survey observation for full Stokes information and rotation measure (RM) determination. Further detailed follow-ups including the use of archival data for timing analysis and imaging for improved localisation are discussed in the companion paper (Paper II).

<sup>h</sup>See <https://apps.datacentral.org.au/smart>.

#### 4.1. Improved detection

For our ongoing shallow survey, processing the full 80-min observation itself readily provides an avenue for confirmation. If the source is genuine and a steady periodic emitter, this should result in a three-fold improvement in S/N. The improvement will be reduced if it is an intermittent source; for example, a pulsar with large nulling fraction. Both these possibilities are exemplified in Figure 6, which shows the original discovery plots along with the improved detections for PSRs J0036–1033 and J0026–1955. The full 80-min observations (42 TB) containing the original detection can be processed and searched over a restricted range in  $P$  and DM using the PRESTO `prepfold` routine. The observations were also processed using the `pdmp` routine within PSRCHIVE pulsar data processing suite<sup>i</sup> (Hotan, van Straten, & Manchester 2004; van Straten, Demorest, & Osłowski 2012), to provide a cross-check and a more accurate DM. This is equivalent to undertaking a longer observation for confirmation. For many of our candidates, this readily provides effective ways of confirming or rejecting a candidate, and eliminates the need for securing additional telescope time that most other surveys typically require.


While the long dwell time of 4800 s should in principle result in an increased sensitivity to sporadic or intermittent pulsars, our current first-pass processing does not necessarily benefit from this. Given this, the discovery of PSR J0026–1955 in the first 10 min of observations, a pulsar with long-duration nulls and a nulling fraction of  $\sim 77\%$ , was remarkably fortuitous (see Figure 6). Details of the discovery, including an analysis of sub-pulse drifting, are reported in McSweeney *et al.* (2022). As mentioned therein, this pulsar turned out to have already been reported as a candidate in the GBNCC survey but was blindly (and independently) discovered in the SMART survey data.

#### 4.2. Improved positional determination

As outlined in Section 2.3, the tied-array beam size for SMART is  $\sim 23'$ . Therefore a more accurate position is essential both for improved detection (i.e. re-beamforming on a more exact sky position) and to facilitate effective follow-ups with other (and more sensitive) telescopes, particularly at higher frequencies where the beams are narrower, even with single-dish telescopes such as Parkes. This would typically involve making multiple re-observations to form a grid around the nominal candidate position. The SMART survey design where the sky is densely sampled (at a rate comparable to, or slightly better than, the Nyquist; Figure 1), allows this to be achieved via reprocessing of the original survey observation, where a dense grid of pointings encompassing the initial position is used for improved positional determination. An example is shown in Figure 9 for the case of PSR J0026–1955. In general, for an initial detection with a modest significance of  $S/N \sim 10$ , we may expect a positional accuracy  $\sim 1\text{--}2'$  through this exercise. In practice, archival VCS data, if available, can also be suitably exploited to progressively further improve the position. In an ideal scenario, where data recorded from all three different configurations are available, an improvement of the order of nearly two orders of magnitude can be achieved through this procedure, as demonstrated in Swainston *et al.* (2021).

<sup>i</sup>See <https://sourceforge.net/projects/psrchive/>.





**Figure 9.** MWA localisation of PSR J0026-1955 by performing a dense grid around the initial pulsar position from the discovery observation. The source position (RA, Dec) = (00<sup>h</sup>26<sup>m</sup>37.5<sup>s</sup>, -19<sup>°</sup>56′24.9″) is  $\approx 32''$  offset from uGMRT-determined position (cf. Paper II for further details). Observations were made using the extended MWA array (Phase II, with  $\sim 6$  km maximum baseline). The uncertainties in the MWA position is  $\sim 12''$  (i.e. about one-tenth of the tied-array beam size, shown as dashed circles on the left panel).


### 4.3. Polarimetry

The VCS recording allows the reprocessing of discovery observations to generate full polarimetric beamformed time series, which can be analysed using standard pulsar packages such as DSPSR<sup>1</sup> (van Straten & Bailes 2011) and PSRCHIVE, for full Stokes profiles. These beamformed MWA data were obtained using the procedures described in Ord et al. (2019) and Xue et al. (2019). The Faraday rotation measure synthesis technique (Brentjens & de Bruyn 2005) can then be applied to estimate the rotation measure (RM).

As an example, Figure 10 shows polarisation data for pulsar J0026-1955, obtained by reprocessing the original discovery observation. This yielded an RM estimate of  $3.65 \pm 0.09 \text{ rad m}^{-2}$ . After correcting for Faraday rotation, linear and circular polarisation was detected. The pulsar exhibits significant amount of linear polarisation but only a small amount of circular polarisation. We attempted to fit the rotating vector model (Radhakrishnan & Cooke 1969) to the position angle (PA) of the linear polarisation across the on-pulse window, in order to constrain the viewing geometry, ( $\alpha, \beta$ ), where  $\alpha$  is the angle between the magnetic and rotation axes, and  $\beta$  is the impact angle of the magnetic axis on the line of sight. In the absence of relativistic effects, the PA curve is expected to be steepest in the centre of the pulse profile, with slope  $d\psi/d\phi = \sin \alpha / \sin \beta \approx 2.4$ , where  $\psi$  is the PA at phase  $\phi$ .

### 5. Survey simulations and forecast

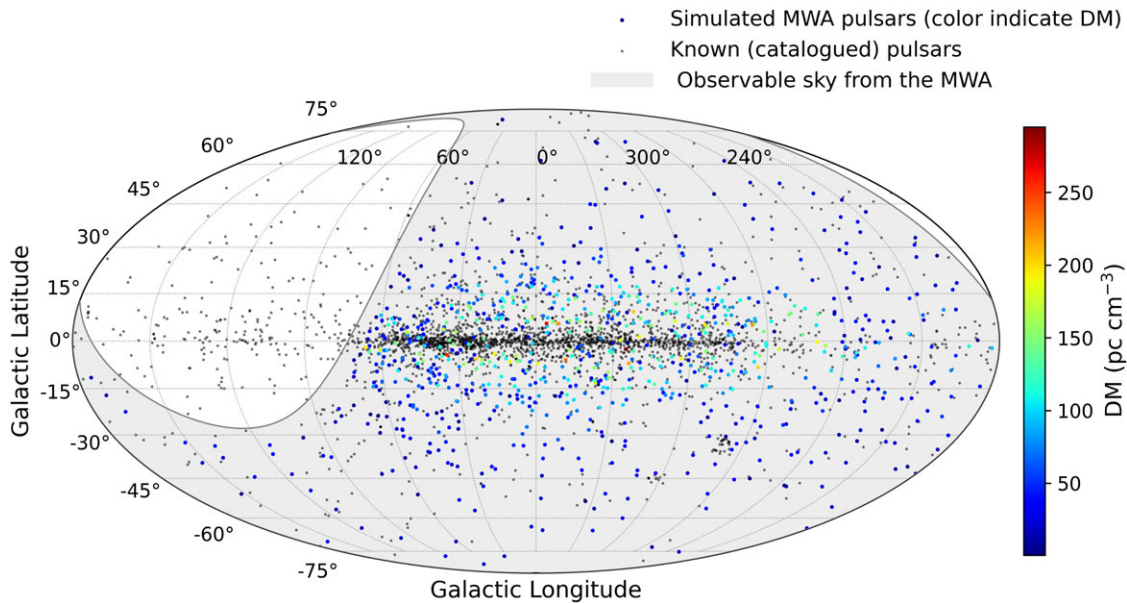
The ongoing first-pass processing (i.e. essentially a shallow survey for long-period pulsars) is limited to processing only a fraction (1/8th) of our observation time over coarser (sub-optimal) trial DM values, out to a maximum DM of  $250 \text{ pc cm}^{-3}$ , and to basic



**Figure 10.** Polarimetric profiles of PSR J0026-1955 obtained by reprocessing the discovery observation at 155 MHz. The black, red, and blue curves in the lower panels show the total intensity, linear, and circular polarisation, respectively. An RM estimate of  $3.65 \pm 0.09 \text{ rad m}^{-2}$  was obtained, and the data were corrected for Faraday rotation.

periodicity search. In the second pass we will extend this to full 80-min observations and employ more optimal DM steps. Besides a three-fold increase in sensitivity expected for long-period pulsars (by virtue of longer integration times), substantial improvements in sensitivity is also expected to millisecond pulsars via finer DM steps and optimal dedispersion plans to match our  $100\text{-}\mu\text{s}/10\text{-kHz}$  resolutions. These considerations motivated our simulation analysis to make some meaningful forecast of the expected survey yield, both for long-period pulsars and MSPs, as summarised below. They provide further justification to undertake a full-scale search processing, planned as part of second-pass processing.

<sup>1</sup>See <https://sourceforge.net/projects/dsprsr/>.



**Figure 11.** Simulated pulsars detectable (colour filled circles) in an all-sky high-time-resolution pulsar search with the MWA in the 140–170 MHz band. The shaded region represents the MWA’s visible sky, that is, the sky south of  $+30^\circ$  in declination. The black filled circles represent known pulsars in the ATNF pulsar catalogue (version 1.67). The colour scale indicates the DM in units of  $\text{pc cm}^{-3}$ .

### 5.1. Long-period pulsars


The discovery of two new pulsars from the processing of a small fraction of survey data hints at the potential for many new pulsar discoveries from a deeper survey that will take advantage of the full 80-min observation. To estimate the survey yield, we have performed survey simulations, using the formalism outlined in Xue *et al.* (2017). The analysis made use of the popular simulation package PsrPopPy (Bates *et al.* 2014) that was developed from the original pulsar simulation software PSRPOP by Lorimer *et al.* (2006). The simulations take into account the sky dependence of the system temperatures at low frequencies ( $T_{\text{sky}} \propto \nu^{-2.55}$ ), as well as the loss in the array gain ( $G$ ) expected at large zenith angles, modelled as  $G(\theta_z) = G_{\text{max}} \cos(\theta_z)$ , where  $\theta_z$  is the zenith angle and  $G_{\text{max}}$  is the gain at  $\theta_z = 0$ . We simulated a population of  $1.6 \times 10^5$  Galactic canonical pulsars, extrapolated from Parkes Multibeam Pulsar Survey (Manchester *et al.* 2001) detections. The luminosity distribution of the canonical pulsar population follows a log-normal distribution  $\langle \log_{10} L \rangle = -1.1$ ,  $\sigma[\log_{10} L] = 0.9$ , where  $L$  is the radio pseudo-luminosity in units of  $\text{mJy kpc}^2$  (Faucher-Giguère & Kaspi 2006). The Galactic radial density distribution follows the Yusifov & Küçük (2004) model. With the caveat that our understanding of the pulsar luminosity function and beaming fraction is limited, we project the deep survey to reach a limiting sensitivity of  $\sim 2\text{--}3$  mJy, with a potential net yield of  $310 \pm 100$  new pulsar discoveries (see Figure 11). This projection mainly applies to the population of long-period pulsars and does not account for other classes of pulsars such as sporadic emitters (e.g. RRATs), or millisecond and binary pulsars, whose populations are hard to model or simulate.

Assuming an isotropic distribution of our simulated local pulsar population ( $\text{DM} \lesssim 250 \text{ pc cm}^{-3}$ ), and scaling for the current (first-pass) search sensitivity (i.e. one-third of the deep-pass sensitivity), and the fraction of data for which the candidate scrutiny

has been completed ( $\sim 5\%$ ), we may expect  $\sim 3\text{--}5$  pulsars. The detection rate at this early stage of SMART thus appears to be in line with this general expectation. While this may seem fortuitous, the unique advantages of the SMART pulsar survey, especially the accessibility to the southern hemisphere, the radio-quiet environment, and the survey parameters (e.g. long dwell times and high time/frequency resolutions), offer excellent prospects for new pulsar discoveries, provided the substantial processing challenges can be addressed.

### 5.2. Millisecond pulsars

Even though the detection sensitivity to MSPs is significantly reduced in our current shallow pass of the survey (owing to the use of coarse or sub-optimal DM step sizes; see Figure 2), the second-pass processing, where we plan to employ more optimal DM searches with a finer step size in DMs, is expected to yield a substantial improvement in sensitivity, particularly at low to moderate DMs, out to  $\lesssim 50 \text{ pc cm}^{-3}$ . At DMs  $\gtrsim 70 \text{ pc cm}^{-3}$ , and especially in regions near the Galactic plane and toward the centre, scatter broadening is expected to result in sensitivity degradation, given the strong frequency dependence (pulse broadening time,  $\tau_d \propto \nu^{-3.9}$ ; cf. Bhat *et al.* 2004), due to which  $\tau_d \gtrsim 10$  ms, which, for millisecond pulsars, can be a substantial fraction of the rotation period. Using PsrPopPy, we simulated a population of  $3 \times 10^4$  MSPs with  $P$  and DM distributions essentially derived from the HTRU intermediate latitude pulsar survey (Levin *et al.* 2013), and with a luminosity limit of  $L_{1400} \sim 0.2 \text{ mJy kpc}^2$ . This corresponds to a limiting flux density  $\sim 10$  mJy at 150 MHz, assuming a spectral index of  $\alpha = -1.8$  (and a distance of  $\sim 1$  kpc), and thus in principle detectable provided there is no significant degradation from dispersive smearing or temporal broadening from scattering.



**Figure 12.** Simulated pulsars detectable in an all-sky pulsar search with the MWA’s 140–170 MHz band with a dwell time of 4800 s. The shaded region represents the MWA’s visible sky, that is, the sky south of  $+30^\circ$  in declination. The black filled circles denote the long-period pulsars, whereas millisecond pulsars detectable in high-sensitivity searches (e.g. using the CDMT) are shown as colour filled circles. The colour scale indicates DM in units of  $\text{pc cm}^{-3}$ .

As with the population of long-period pulsars, this analysis accounted for the sky dependence of  $T_{\text{sky}}$ , non-uniformity in the array gain, and strong frequency scaling of scattering, which is especially important for MSPs. For example, using some preliminary dedispersion plan estimates for the second round of processing (i.e. the deep search), where we assume a typical plan would involve DM steps of  $0.01 \text{ pc cm}^{-3}$  up to  $54 \text{ pc cm}^{-3}$  and  $0.02 \text{ pc cm}^{-3}$  out to  $107 \text{ pc cm}^{-3}$ , our simulations predict 55 detectable MSPs above our detection threshold, and hence  $\sim 15$  new MSP discoveries. However, a substantial increase is forecast in simulations that closely emulate the higher sensitivity attainable through more optimal searches that make use of *coherent dispersion measure trials* (CDMT), which is equivalent to the use of finer DM steps of  $0.002 \text{ pc cm}^{-3}$ , and will limit residual DM smearing to  $\sim 150 \mu\text{s}$  (comparable to  $\sim 100 \mu\text{s}$  native resolution of the VCS). In essence, this means that full-scale, high-sensitivity searches employing the implementation of CDMT, if feasible for SMART, can potentially lead to the discovery of as many as  $\sim 30$  MSPs.

The simulated population of  $\sim 70$  MSPs, along with the simulated population of long-period pulsars (see Section 5.1), is shown in Figure 12. Our simulation analysis did *not* consider a large population of MSPs discovered in recent (and highly successful) Fermi-directed targeted searches (Deneva et al. 2021, and references therein). Even so, the detectable population of MSPs is almost twice the currently known population within  $\text{DM} \lesssim 100 \text{ pc cm}^{-3}$ , which means a net MSP yield that is competitive to that from the highly successful Parkes HTRU survey. Indeed, as evident from Figure 12, the detectable population of MSPs is limited to  $\text{DM} \lesssim 70 \text{ pc cm}^{-3}$ , which is reconcilable given the expected pulse broadening times of  $\tau_d \gtrsim 10 \text{ ms}$  toward such moderate DM pulsars at the low frequencies of the MWA (e.g. Kirsten et al. 2019). Consequently, the vast majority of MSPs discovered will likely be suitable for high-precision timing applications such as pulsar timing arrays.

## 6. Future processing plans

The planned second-pass survey will extend the processing to the full 80-min observations and carry out more optimal searches in the DM parameter space, while incorporating searches for both long-period pulsars and millisecond pulsars. As such, the long dwell times of SMART (4800 s) can be exploited to search for pulsars with very long periods, like those discovered by LOFAR and MeerKAT (Tan et al. 2018b; Caleb et al. 2022), and provide increased sensitivity to objects that emit intermittently, for example, pulsars with long null durations such as PSR J0026–1955 (McSweeney et al. 2022). In addition, the adopted strategy to archive recorded voltages offers additional avenues for future processing; for example, searches for millisecond pulsars through the application of novel hybrid dedispersion approaches that involve the use of coherent dispersion measure trials (CDMT), which was demonstrated by the LOFAR through the discovery of PSR J0952–0607 (Bassa et al. 2017). Below we outline our processing plans and strategies in the near-term and highlight some of the computational challenges and other considerations in planning this second-pass processing.

### 6.1. Beamforming and sensitivity optimisation

As discussed earlier in Section 2.5, the tied-array beamforming strategy warrants some more careful thought in order to maximise sensitivity while also reducing needless processing. Inevitably, this produces an uneven sensitivity threshold across the sky due to both primary beam pointing effects and effective dwell time. These considerations are also important when estimating survey-wide statistics. We are formulating a more efficient beamforming scheme that takes into account these technical details, which will be presented in a subsequent paper detailing the second-pass survey processing.



## 6.2. Dedispersion planning and RFI mitigation strategies

For the first-pass survey processing described in this paper, the dedispersion plan outlined in Table 2 is adequate for all observations. In contrast, a slightly more sophisticated plan may be required for the second-pass processing to accommodate the eight-fold increase in observation length and to provide increased sensitivity to shorter-period pulsars. We are actively developing a sensible strategy that balances our sensitivity goals and the relatively large computational costs associated with dedispersing MWA VCS data, especially since we would essentially be producing  $\sim 10\times$  as many DM trials.

In addition to revisiting the dedispersion plan, we will also incorporate a more careful approach to excising or mitigating RFI (both periodic and impulsive). The observatory site is exceptionally RFI-quiet (owing to the geographical location and radio-quiet zone status), hence the first-pass processing did not include any active RFI mitigation other than what is naturally gained by forming TABs (where off-axis RFI is ‘phased out’). We are currently examining the periodic RFI environment by processing observations taken throughout the SMART observing semesters and using a standard PRESTO-based approach to find bright, common terrestrial signals by searching for periodic ‘candidates’ in the zero-DM topocentric time series data. Once we collect this information, we will apply the masks (after appropriate barycentric corrections are made) during the periodicity search pipeline. Additionally, there can occasionally be bursts of narrowband interference (e.g. air-craft and satellites in TAB grating lobes) that could severely affect our data quality for short periods of time. There are several software preprocessing solutions to this kind of RFI (e.g. Eatough, Keane, & Lyne 2009; Men *et al.* 2019; Morello, Rajwade, & Stappers 2022), which we will explore in parallel to the periodic RFI mitigation strategies. Empirically, VCS data are remarkably clear of impulsive/narrowband RFI in the SMART observing band, and data excision is  $\ll 10\%$  for a typical observation.

## 6.3. Searches for long-period pulsars and sporadic emitters

The long dwell times of SMART make it particularly amenable to the application of fast-folding algorithms that offer significantly higher sensitivity to pulsars with rotation periods  $\gtrsim 10$  s (e.g. Morello *et al.* 2020). Such slow-spinning pulsars are likely to be near the radio emission ‘death lines’ so can be invaluable in gaining useful insights into the intricacies of the pulsar radio emission process. Recent applications of this algorithm in Parkes and Arecibo searches have led to the discoveries of pulsars with  $P > 10$  s (Morello *et al.* 2020), or very weak pulsars ( $S_{1400} \sim 10$   $\mu$ Jy) with a  $\sim 2\%$  duty cycle (Parent *et al.* 2018). These, and other recent discoveries such as a 76-s pulsar with MeerKAT (Caleb *et al.* 2022), provide a strong motivation for undertaking fast-folding searches. The low levels of RFI at the observatory site are particularly advantageous for this.

The SMART survey dwell time is substantially longer than those of previous-generation southern-sky surveys, particularly at high- $|b|$  parts of the sky, where it is 20 times longer than the HTRU survey (Keith *et al.* 2010) and 40 times longer than the southern pulsar survey (Manchester *et al.* 1996). It is also 40 times longer than that of the ongoing GBNCC survey (Stovall *et al.* 2014) that covers the sky north of  $-55^\circ$  in declination (Table 1). Considering this, detection prospects are promising, especially given the  $\sim 2$ – $3$  mJy limiting sensitivity that the SMART can attain for long-period

pulsars (Section 2.4) and negligible degradation in signal strength due to dispersion and pulse broadening effects.

As described earlier, the long dwell times also increase the search sensitivity to objects that emit sporadically, such as RRATs and giant-pulse emitters (e.g. the Crab pulsar), which can be more effectively detected by searching for individual dispersed pulses, and will be part of the second-pass processing.

## 6.4. Searches for binary and millisecond pulsars

The long dwell times, and high time and frequency resolutions, of the SMART can also be exploited, in principle, to search for binary and millisecond pulsars. However, a full-scale acceleration search can be prohibitively expensive at the low frequencies, given the very large number of DM and acceleration trials that are required (e.g. typically  $\sim 10^4$  up to  $250$   $\text{pc cm}^{-3}$ , and  $\sim 2400$  across  $\pm 100$   $\text{m s}^{-2}$ ). Compared to the HTRU-south low-latitude survey, which has been successful in finding such systems (e.g. Cameron *et al.* 2020), the cost of searching SMART data can be more than an order of magnitude greater. The successful detection of several MSPs and the double pulsar in our initial census (cf. Paper II) makes such searches worthwhile.

An inherent limitation in the searches for such short-period pulsars is the significant degradation in sensitivity due to substantial dispersion smearing (relative to rotation periods) despite our 10-kHz channels. Fortunately, this can be alleviated by using CDMT-based searches (Bassa *et al.* 2017). Recording in  $24 \times 1.28$ -MHz channels makes the SMART data highly amenable to the application of CDMT searches, and can result in a substantial increase in detection sensitivity to short-period millisecond pulsars. Integration of this novel method, and benchmarking on prospective HPC clusters with significant computational resources (e.g. Pawsey’s emerging Setonix cluster) is also part of our future processing plans, although a full-scale processing may have to await access to sufficient computational resources. We are also exploring publicly available, GPU-enabled Fourier domain acceleration search software (e.g. AstroAccelerate; Armour *et al.* 2020) as a drop-in replacement for PRESTO’s CPU-based `accelsearch`.

Regardless, the high cost of such computationally-intensive searches will likely necessitate a multi-pass processing strategy; for instance, an initial pass involving acceleration searches, but limited to a modest number of acceleration trials (e.g.  $\sim 150$  to cover  $\pm 6$   $\text{m s}^{-2}$ ), thereby retaining sensitivity to short-period objects ( $P \lesssim 10$  ms) but with the binary orbital period,  $P_b \gtrsim 5$  d (i.e. with low-mass white dwarf type companions). Full-scale acceleration searches that target binary systems such as PSR J0737–3039 or PSR J1757–1854 with  $P_b \lesssim 5$  h (i.e. requiring  $\sim 2400$  trials spanning across  $\pm 100$   $\text{m s}^{-2}$ ) are hence deferred to the longer-term future. Such searches will be primarily limited to the regions around the Galactic plane, at least initially, thus processing only a fraction of the SMART data (e.g. sky within  $|b| \lesssim 5^\circ$ ). Such a multi-pass strategy is also motivated by the demonstrated success of HTRU-south, which has led to the discovery of exotic systems such as PSR J1757–1854 (Cameron *et al.* 2018) and wide-orbit double neutron-star system (Sengar *et al.* 2022). In any case, notwithstanding the high computational cost, the high-profile scientific applications of such rare systems make similar full-scale acceleration searches scientifically compelling for the SMART data. The long-term scientific dividends of such systems are vividly demonstrated by Kramer *et al.* (2021) through the

16-yr timing analysis of the double pulsar, enabling the most stringent tests of general relativity and alternative theories of gravity.

## 7. Summary and conclusions

With its novel features such as voltage recording and long dwell times, and access to the pristine radio-quiet environment in the southern hemisphere, the SMART survey is well positioned to play an impactful role in the exploration of the southern, low-frequency sky for pulsar surveys and science. Since the MWA is a precursor for SKA-Low, the SMART survey will also serve as an important preparatory step for pulsar surveys planned with SKA-Low. Additionally, it will map out the southern sky for low-frequency detections of many pulsars that were originally discovered at frequencies  $\gtrsim 400$  MHz.

The survey is enabled by the advent of the Phase II upgrade of the MWA, the compact configuration of which offers an enormous gain in the beamforming and processing cost, thereby making large all-sky pulsar surveys tractable with large-FoV interferometric arrays such as the MWA. The combination of voltage recording and the FoV brings a survey efficiency of  $\sim 450 \text{ deg}^2 \text{ h}^{-1}$ , but at the expense of large data rates of  $28 \text{ TBh}^{-1}$ . Consequently,  $\sim 3 \text{ PB}$  of (VCS mode) data for the full survey and significant processing costs.

Due to the substantial computational cost involved in searching at low frequencies, the processing is undertaken in multiple passes. In the ongoing first-pass processing, 10 min of data from each observation are processed in 2358 trial DMs, out to a maximum DM of  $250 \text{ pc cm}^{-3}$ , thereby reaching about one-third of the sensitivity that will eventually be attainable in full observation processing.

The voltage recording strategy adopted for the SMART survey enables a multitude of avenues for follow-ups and confirmations, including improved detection, initial polarimetry and arcminute-level positional determination—all by reprocessing the original observation and, where possible, also archival VCS data. This also facilitates timely follow-up studies using more sensitive telescopes such as Parkes and the upgraded GMRT (uGMRT) that operate at frequencies  $\gtrsim 300$  MHz.

With the recent development of a *web app* for facilitating efficient scrutiny of candidate analysis, including classification and ranking for identifying promising ones to follow-up, we anticipate the discovery rate to increase in the coming years. As software tools mature and the search pipelines are expanded to include acceleration trials and fast-folding based algorithms, and additional computational resources become available, it will become possible to extend the processing to include searches for binary and millisecond pulsars, and those with very long periods, or even sporadic emitters. Our simulation analysis forecasts a survey yield of  $\sim 300$  long-period pulsars and  $\sim 30$  millisecond pulsars by the completion of full processing. The SMART survey data will serve as a complete digital record of the low-frequency southern sky, and an important reference for even more ambitious surveys planned with the SKA-Low.

**Acknowledgement.** We thank an anonymous referee for several useful comments that helped to improve the content and presentation of this paper. The scientific work made use of Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory. We acknowledge the Wajarri Yamaji people as the traditional owners of the Observatory site. This work

was supported by resources provided by the Pawsey Supercomputing Centre with funding from the Australian Government and the Government of Western Australia. This work was also supported by resources awarded under Astronomy Australia Ltd's ASTAC merit allocation scheme on the OzSTAR national facility at the Swinburne University of Technology. The OzSTAR programme receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government. The development of SMART web app was facilitated by the software support scheme of ADACS. We thank Simon O'Toole for help with the migration of web app to Data Central. The GMRT is run by the National Centre for Radio Astrophysics of the Tata Institute of Fundamental Research, India. The Parkes (*Murriyang*) radio telescope is part of the Australia Telescope National Facility which is funded by the Australian Government for operation as a National Facility managed by CSIRO. We thank L. Levin for help with the simulation analysis of MSPs.

**Software:** We acknowledge the use of the following software/packages for this work: CASA (McMullin et al. 2007), DSPSR (van Straten & Bailes, 2010, 2011), PRESTO (Ransom 2001, 2011), PSRCHIVE (Hotan et al. 2004; van Straten et al. 2011), Tempo2 (Hobbs, Edwards, & Manchester 2006; Hobbs & Edwards 2012), PINT (Luo et al. 2019, 2021), PsrPopPy (Bates et al. 2014), Nextflow (Di Tommaso et al. 2017).

**Data Availability.** Data products from the SMART project (e.g., beamformed time series, pulsar detections and archives) will be made available via Data Central services (<https://apps.datacentral.org.au/smart>), which also hosts the database web app for candidate classification and rating.

## References

- Archibald, A. M., et al. 2009, *Sci*, **324**, 1411  
 Armour, W., et al. 2020, *AstroAccelerate*, doi: [10.5281/zenodo.4282748](https://doi.org/10.5281/zenodo.4282748)  
 Bailes, M., et al. 2011, *Sci*, **333**, 1717  
 Bassa, C. G., et al. 2017, *ApJ*, **846**, L20  
 Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, *MNRAS*, **439**, 2893  
 Bell, M. E., et al. 2016, *MNRAS*, **461**, 908  
 Bhat, N. D. R., Bailes, M., & Verbiest, J. P. W. 2008, *PhRvD*, **77**, 124017  
 Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, *ApJ*, **605**, 759  
 Bhat, N. D. R., Ord, S. M., Tremblay, S. E., McSweeney, S. J., & Tingay, S. J. 2016, *ApJ*, **818**, 86  
 Bhat, N. D. R., et al. 2018, *ApJS*, **238**, 1  
 Bhattacharyya, B., et al. 2016, *ApJ*, **817**, 130  
 Brentjens, M., & de Bruyn, A. 2005, *A&A*, **441**, 1217  
 Caleb, M., et al. 2022, *NatAs*, [arXiv:2206.01346](https://arxiv.org/abs/2206.01346)  
 Cameron, A. D., et al. 2018, *MNRAS*, **475**, L57  
 Cameron, A. D., et al. 2020, *MNRAS*, **493**, 1063  
 Cordes, J. M., & Lazio, T. J. W. 2002, *arXiv e-prints*, astro  
 Cordes, J. M., & McLaughlin, M. A. 2003, *ApJ*, **596**, 1142  
 Cordes, J. M., Weisberg, J. M., & Boriakoff, V. 1985, *ApJ*, **288**, 221  
 Cordes, J. M., et al. 2006, *ApJ*, **637**, 446  
 Dai, S., et al. 2015, *MNRAS*, **449**, 3223  
 Deller, A. T., et al. 2016, *ApJ*, **828**, 8  
 Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Natur*, **467**, 1081  
 Deneva, J. S., et al. 2021, *ApJ*, **909**, 6  
 Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, *ApJ*, **294**, L25  
 Di Tommaso, P., et al. 2017, *NatB*, **35**, 316–319  
 Eatough, R. P., Keane, E. F., & Lyne, A. G. 2009, *MNRAS*, **395**, 410  
 Faucher-Giguère, C.-A., & Kaspi, V. M. 2006, *ApJ*, **643**, 332  
 Geyer, M., et al. 2017, *MNRAS*, **470**, 2659  
 Han, J. L., et al. 2021, *RAA*, **21**, 107  
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, **47**, 1  
 Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Natur*, **217**, 709  
 Hobbs, G., & Edwards, R. 2012, *Tempo2: Pulsar Timing Package*, Astrophysics Source Code Library, record ascl:1210.015, ascl:1210.015

- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, **21**, 302
- Jankowski, F., et al. 2018, *MNRAS*, **473**, 4436
- Janssen, G., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 37
- Kaur, D., et al. 2022, *ApJ*, **930**, L27
- Kaur, D., et al. 2019, *ApJ*, **882**, 133
- Keane, E., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 40
- Keane, E. F., et al. 2018, *MNRAS*, **473**, 116
- Keith, M. J., et al. 2010, *MNRAS*, **409**, 619
- Kerr, M., et al. 2014, *MNRAS*, **445**, 320
- Kirsten, F., et al. 2019, *ApJ*, **874**, 179
- Kramer, M., et al. 2006, *Sci*, **314**, 97
- Kramer, M., et al. 2021, *PhRvX*, **11**, 041050
- Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *MNRAS*, **225**, 307
- Levin, L., et al. 2013, *MNRAS*, **434**, 1387
- Lorimer, D. R., & Kramer, M. 2012, *Handbook of Pulsar Astronomy* (Cambridge University Press)
- Lorimer, D. R., et al. 2006, *MNRAS*, **372**, 777
- Luo, J., et al. 2019, PINT: High-precision pulsar timing analysis package, Astrophysics Source Code Library, record ascl:1902.007, ascl:1902.007
- Luo, J., et al. 2021, *ApJ*, **911**, 45
- Lyne, A. G., et al. 2004, *Sci*, **303**, 1153
- Lyon, R. J., Stappers, B. W., Cooper, S., Brooke, J. M., & Knowles, J. D. 2016, *MNRAS*, **459**, 1104
- Manchester, R., Hobbs, G., Teoh, A., & Hobbs, M. 2005, *AJ*, **129**, 1993
- Manchester, R. N., et al. 1978, *MNRAS*, **185**, 409
- Manchester, R. N., et al. 1996, *MNRAS*, **279**, 1235
- Manchester, R. N., et al. 2001, *MNRAS*, **328**, 17
- McLaughlin, M. A., et al. 2006, *Natur*, **439**, 817
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 376, *Astronomical Data Analysis Software and Systems XVI*, ed. R. A. Shaw, F. Hill, & D. J. Bell, 127
- McSweeney, S. J., Bhat, N. D. R., Tremblay, S. E., Deshpande, A. A., & Ord, S. M. 2017, *ApJ*, **836**, 224
- McSweeney, S. J., et al. 2020, *PASA*, **37**, e034
- McSweeney, S. J., et al. 2022, *ApJ*, [arXiv:2206.00805](https://arxiv.org/abs/2206.00805)
- Men, Y. P., et al. 2019, *MNRAS*, **488**, 3957
- Meyers, B. W., et al. 2017, *ApJ*, **851**, 20
- Meyers, B. W., et al. 2018, *ApJ*, **869**, 134
- Michilli, D., et al. 2018, *MNRAS*, **480**, 3457
- Morello, V., Rajwade, K. M., & Stappers, B. W. 2022, *MNRAS*, **510**, 1393
- Morello, V., et al. 2020, *MNRAS*, **493**, 1165
- Ord, S. M., et al. 2019, *PASA*, **36**, e030
- Parent, E., et al. 2018, *ApJ*, **861**, 44
- Prabu, T., et al. 2015, *ExA*, **39**, 73
- Radhakrishnan, V., & Cooke, D. J. 1969, *ApL*, **3**, 225
- Ransom, S. 2011, PRESTO: Pulsar Exploration and Search Toolkit, Astrophysics Source Code Library, record ascl:1107.017, ascl:1107.017
- Ransom, S. M. 2001, PhD thesis, Harvard University
- Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *AJ*, **124**, 1788
- Rickett, B. J. 1990, *ARA&A*, **28**, 561
- Sanidas, S., et al. 2019, *A&A*, **626**, A104
- Sengar, R., et al. 2022, *MNRAS*, doi: [10.1093/mnras/stac821](https://doi.org/10.1093/mnras/stac821)
- Sett, S., Bhat, N. D. R., Sokolowski, M., & Lenc, E. 2022, arXiv e-prints, [arXiv:2212.06982](https://arxiv.org/abs/2212.06982)
- Shao, L., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 42
- Sokolowski, M., et al. 2017, *PASA*, **34**, e062
- Stovall, K., et al. 2014, *ApJ*, **791**, 67
- Swainston, N. A., et al. 2022, *PASA*, **39**, e020
- Swainston, N. A., et al. 2021, *ApJ*, **911**, L26
- Tan, C. M., et al. 2018a, *MNRAS*, **474**, 4571
- Tan, C. M., et al. 2018b, *ApJ*, **866**, 54
- Thornton, D., et al. 2013, *Sci*, **341**, 53
- Tingay, S. J., et al. 2013, *PASA*, **30**, e007
- Toscano, M., Bailes, M., Manchester, R. N., & Sandhu, J. S. 1998, *ApJ*, **506**, 863
- Tremblay, S. E., et al. 2015, *PASA*, **32**, e005
- van Straten, W., & Bailes, M. 2010, DSPSR: Digital Signal Processing Software for Pulsar Astronomy, Astrophysics Source Code Library, record ascl:1010.006, ascl:1010.006
- van Straten, W., & Bailes, M. 2011, *PASA*, **28**, 1
- van Straten, W., et al. 2001, *Natur*, **412**, 158
- van Straten, W., et al. 2011, PSRCHIVE: Development Library for the Analysis of Pulsar Astronomical Data, Astrophysics Source Code Library, record ascl:1105.014, ascl:1105.014
- van Straten, W., Demorest, P., & Osłowski, S. 2012, *ART*, **9**, 237
- Venkatraman Krishnan, V., et al. 2020, *Sci*, **367**, 577
- Wayth, R. B., et al. 2015, *PASA*, **32**, e025
- Wayth, R. B., et al. 2018, *PASA*, **35**, 33
- Xue, M., et al. 2019, *PASA*, **36**, e025
- Xue, M., et al. 2017, *PASA*, **34**, e070
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, **835**, 29
- Yusifov, I., & Küçük, I. 2004, *A&A*, **422**, 545