Publications of the Astronomical Society of Australia (2026), **43**, e054, 9 pages
doi:10.1017/pasa.2026.10174

CAMBRIDGE UNIVERSITY PRESS

Research Article

MWA tied-array processing V: Super-resolved localisation via amplitude-only maximum likelihood direction finding

Bradley W. Meyers[1,2] and Arash Bahramian[2]

[1]Australian SKA Regional Centre (AusSRC), Curtin University, Bentley, WA, Australia and [2]International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley, WA, Australia

Abstract

Interferometric localisation of transients and pulsars via tied-array beam processing is challenging and can be limited by the native spatial resolution achievable by the instrument, especially at low frequencies and for compact interferometers. Knowledge of the telescope primary and tied-array beam patterns allows the exploitation of the beam structures and the relationship to measured quantities, such as signal-to-noise ratio, through radio direction finding techniques. The additional information provides a 'super-resolved' localisation (i.e. where the precision is much better than the native spatial resolution) of a source when there are multiple detections in adjacent tied-array beams. We demonstrate this approach using the Murchison Widefield Array (MWA) and its voltage capture and tied-array processing capabilities, with a specific focus on how it benefits the on-going Southern-sky MWA Rapid Two-metre pulsar survey as it starts producing more candidates requiring follow-up. Examples of localisations with previously discovered MWA pulsars that were subsequently localised via imaging with higher spatial resolution interferometers are used to validate the process, along with localisations of a sample of known pulsars to demonstrate the robustness of the method and its uncertainty estimation.

Keywords: Radio interferometry (1346); astronomical techniques (1684)

(Received 2 May 2025; revised 27 January 2026; accepted 3 March 2026)

1. Introduction

The tied-array beam (TAB) processing capabilities for Murchison Widefield Array (MWA; Tingay et al. 2013; Wayth et al. 2018) has evolved substantially over the past decade. Ord et al. (2019) introduced the concept of tied-array beamforming in the context of MWA voltages, and Xue et al. (2019) verified that methodology is feasible at large offsets from zenith and produces sensible polarisation properties. McSweeney et al. (2020) developed and verified an 'inverse polyphase filterbank' (synthesis filter) approach to recover the 1.28 MHz coarse channels from the recorded 10 kHz voltages, making coherent dedispersion tractable and allowing a wider range of millisecond pulsar science. Most recently, Swainston et al. (2022) introduced a multi-pixel beamformer designed to aid in pulsar survey science, and demonstrated its usefulness in mitigating ionospheric bulk shifts of pulsars in the field of view. The development of these technical processing capabilities made it possible to consider conducting a full southern-sky pulsar survey with the MWA by leveraging the voltage capture system (VCS; Tremblay et al. 2015; Morrison et al. 2023), leading to the inception of the the Southern-sky MWA Rapid Two-metre (SMART) pulsar survey (Bhat et al. 2023a,b).

During data acquisition periods for the SMART survey the MWA configured in its 'compact' layout. In this configuration, the maximum element baseline is $B_{\rm max} \approx 740$ m although given the centrally condensed element layout the effective characteristic baseline is $\tilde{B}_{\rm c} \sim 300$ m. At the central frequency of SMART (154 MHz), this corresponds to a TAB width of $\lambda/\tilde{B}_{\rm c} \sim 20-30'$ (depending on how baselines are weighted, how the bandwidth is averaged, etc.). This makes the offline voltage processing tasks feasible for an all-sky survey, enlarging the TAB size and reducing the total number required to tessellate the observed field of view. A feature of the compact configuration is the introduction of 'Hex' clusters, where tiles are arranged in a hexagonal layout (see Figure 1 of Wayth et al. 2018), adding many redundant baselines and subsequently modifying the synthesised beam pattern to be both wider (short baseline dominated) and inherit a hexagonal spatial distribution on the sky.

Given the wider TABs, it becomes difficult to localise pulsar candidates from SMART to the level required for follow-up with higher frequency and/or higher spatial resolution facilities. Fast imaging (e.g. $\lesssim$100 ms integrations) is technically feasible due to the versatility of voltage capture observations, but can be computationally expensive (especially since the TABs have already been created). Image based pulsar or transient detection is nominally sensitive to a different population of objects than those probed by classical beamforming approaches (e.g. Sett et al. 2023, 2024), and also relies on the ability to detect the transient source in individual snapshot images, which can be challenging for moderate to weak sources. The imaging approach is further complicated when the interferometric elements are closely spaced where one suffers from a wide synthesised beam and confusion-limited images. It is possible for the brightest sources to achieve comparable or better localisation via imaging and cross-matching to calibrated astrometric survey positions at similar frequencies (e.g. Wayth et al.

Corresponding author: Bradley W. Meyers; Email: bradley.meyers@curtin.edu.au

Cite this article: Meyers BW. and Bahramian A. (2026) MWA tied-array processing V: Super-resolved localisation via amplitude-only maximum likelihood direction finding. *Publications of the Astronomical Society of Australia* **43**, e054, 1–9. https://doi.org/10.1017/pasa.2026.10174

© The Author(s), 2026. Published by Cambridge University Press on behalf of Astronomical Society of Australia. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (https://creativecommons.org/licenses/by-nc-sa/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

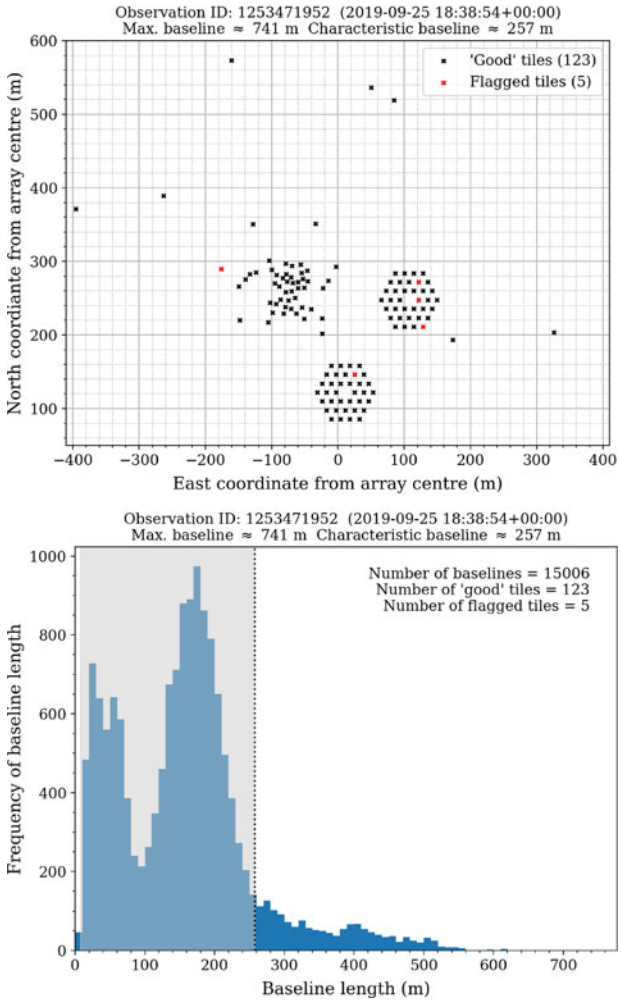


Figure 1. *Top*: An example of a typical MWA compact configuration tile layout for Observation ID 1253471952, observed on 2019-09-25 18:38:54 UTC. Here, each black cross is an operational element, or 'tile', and red crosses indicate flagged (bad) tiles. Usually, only 128 tile signals were recorded at any given time. *Bottom*: The baseline length distribution, identifying the characteristic baseline length as ∼260 m (black dotted line) as well as the 90% highest density interval (grey shaded region), i.e. the baseline length range that contain 90% of all baselines.

2015; Intema et al. 2017; Hurley-Walker et al. 2022), but there is no guarantee that pulsars, or other fast-transients, will appear in these catalogues. For the MWA, where the correlator does not run concurrently with the VCS, the trade-off favours TAB localisation using the voltage data, whereas for systems where the correlator and beamforming engines can run simultaneously it is will often be better to attempt image-based localisation (assuming the transient is detectable on the relevant accumulation timescale).

Strategic efforts utilising higher spatial resolution telescopes are the most directly effective at refining the candidate positions, where both imaging and standard high time resolution data acquisition modes can help to finalise localisation to the level of ∼10″ or better. This allows for a rapid convergence of a follow-up position for transient sources, or even a working timing solution for discovered pulsars, as demonstrated by Swainston et al. (2021). Producing a timing solution for such new pulsars would otherwise require a dedicated observing campaign over the course of ∼1 yr in order to break the covariance between the unknown pulsar position and the rotational (and possible binary) properties. Achieving a robust position as early as possible in the timing process can help to reveal binary motion and to constrain an orbital ephemeris which further reduces the amount of time required to fully model the pulsar rotation. However, the required input to these follow-up observations is a best-estimate localisation from the MWA data available (see Section 4.2 of Bhat et al. 2023a). While there is occasionally sky coverage of the candidate from MWA observations in its extended configurations (where the TAB width is ∼2′ at 154 MHz), this is not generally the case. Typically, one only has the information from the tessellated TAB detections of candidates from the initial SMART search process, or a re-processed TAB grid surrounding the nominal candidate position.

Moment invariants of digitised images can be used to calculate image centroids to sub-pixel precision, where each of the pixels (or pixel regions) are intensity weighted, however, they are biased due to undersampling of the field (e.g. Teh & Chin 1986; Thomas et al. 2006; Vyas, Prasad, & Roopashree 2009). This approach can be translated directly to using TABs (as a 'pixel') where the intensity weighting is provided by the detection metric, e.g. the signal-to-noise ratio (S/N). In this case a simplistic and identical TAB shape is often assumed, and struggles to produce meaningful uncertainties. These downfalls can be mitigated by using a realistic TAB shape and sensitivity estimate (which could change across the sampled sky) for each grid pixel position, along with the corresponding detection metric, and an appropriately tailored direction-finding algorithm to arrive at a more robust position and uncertainty estimate. We can achieve 'super-resolved' localisation by incorporating such additional information, wherein the localisation and uncertainty regions are far smaller than the nominal TAB width.

The fundamental localisation technique has been demonstrated previously in similar circumstances for the Low Frequency Array (LOFAR), Australian Square Kilometre Array Pathfinder (ASKAP) and MeerKAT interferometers (Obrocka, Stappers, & Wilkinson 2015; Bannister et al. 2017; Bezuidenhout et al. 2023), but is also deeply connected with radio direction finding and frequency estimation methodologies (e.g. Capon 1969; Schmidt 1986; Lipsky 2004), where instead of having multiple physical detectors focused on different bearings, we generate our detector signals via digital tied-array beamforming from voltage data. In practice, distributed islands of likelihood are usually identified, where the interferometric element layout (and, therefore, TAB shape) strongly impacts the complexity of the likelihood map (e.g. see Section 3.1 and Figure 1 of CHIME/FRB Collaboration et al. 2019). Nevertheless, the maximum-likelihood amplitude-based direction finding algorithms are robust (assuming individual element gains are well-known/calibrated) and relatively simple to approach, making them attractive and effective tools for current and next-generation radio telescopes.

In Section 2, we adapt and develop the framework used to localise a target by leveraging the MWA tied-array beam pattern in the context of amplitude-only maximum likelihood direction finding, and in Section 3, we verify the approach. In Section 4, we discuss the impact this approach can have on localisation for MWA candidates. Finally, in Section 5, we summarise and touch on the implications of this method in terms of follow-up efficiency for SMART and MWA VCS science more generally.

2. Methodology

While we focus on direct applications to the MWA, the methodology can in principle be applied to any similar interferometric

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

array where there is the capability to re-process voltage data, or where clusters of tied-array beams are formed in real time. An example of the latter is the TRAnsients and PUlsars with MeerKAT (TRAPUM) survey (Stappers & Kramer 2016), and their localisation software package SeeKAT (Bezuidenhout et al. 2023). The underlying localisation method can be summarised in the following steps:

- Simulate (or otherwise obtain) a tied-array beam pattern as a function of sky position and frequency. Outlined in Section 2.1.
- Consolidate and assign the target source detection significance metrics (flux density, S/N, etc.) to each respective tied-array beam pattern. Outlined in Section 2.2.1.
- Using the ratio of combinations of the tied-array beam pattern and its corresponding detection metrics, form a pseudo-likelihood and, therefore, probability map. Outlined in Sections 2.2.2–2.2.4.
- Identify the peak of the the probability map as the most-likely position.
- Calculate an error estimate based on the 2D distribution of the likelihoods. Outlined in Section 2.2.5.

The manner in which these steps are achieved can vary between telescopes, and in the case of the MWA some further tempering of the probability maps are required due to the complex beam patterns. The ASKAP localisation method approaches the problem similarly, but uses a Bayesian estimator rather than a maximum-likelihood estimator when forming the probability maps. Fundamentally, we apply the same core steps for the MWA, but specifically consider the more complex beam patterns and how this method helps to mitigate what would otherwise be an untenable localisation follow-up scheme.

2.1. Sky response pattern

Here, we will provide a summary of the relevant antenna theory principles needed to construct the information required for localisation inference (see e.g. Balanis 2016 for a thorough exploration). Meyers et al. (2017) and Lee et al. (2025) have used a similar approach in the context of tied-array beamforming to facilitate flux density calibration of MWA pulsar data, and Sutinjo et al. (2015) in the context of the MWA primary beam patterns.

2.1.1. Calculating the array factor

Consider an array with N non-interacting receiving elements, where the nth element has a sky-response pattern expressed as $D_n(\hat{p}, \nu)$, with ν as the observing frequency and $\hat{p} = (\vartheta, \varphi)$ as the desired pointing direction in horizontal coordinates, i.e. azimuth angle φ and zenith angle ϑ. We then consider an impinging planar wave from the look-direction, which we express as

$$\psi_n(\hat{p}, \nu) = \exp\left[i\frac{2\pi\nu}{c}(\hat{\mathbf{k}}_{\hat{p}} \cdot \mathbf{r}_n)\right]. \tag{1}$$

Here, $\mathbf{r}_n = [x_n, y_n, z_n]$ represents the nth element position vector[a] and $\hat{\mathbf{k}}_{\hat{p}} = [\sin\vartheta\cos\varphi, \sin\vartheta\sin\varphi, \cos\vartheta]$ is the unit wavevector corresponding to $\hat{p}$. We may then define the tied-array voltage pattern, $D_{\rm ta}(\hat{p}, \nu)$, as the weighted mean of the elements in response to that plane wave,

$$D_{\rm ta}(\hat{p}, \nu) = \frac{1}{N}\sum_{n=1}^{N} w_n D_n(\hat{p}, \nu)\psi_n(\hat{p}, \nu), \tag{2}$$

where w_n is a weighting function that we are free to define, but which will ultimately act to 'steer' the tied-array voltage pattern on the sky. If one assumes that D_n is identical for each element in equation (2), then we can write the element response (i.e. the telescope primary beam) as $D_{\rm el}$ for each antenna, and it becomes a multiplicative factor that can be brought outside of the summation. We can then express the TAB pattern as $D_{\rm ta}(\hat{p}, \nu) = A(\hat{p}, \nu)D_{\rm el}$, defining the 'array factor' to be

$$A(\hat{p}, \nu) = \frac{1}{N}\sum_{n=1}^{N} w_n \psi_n(\hat{p}, \nu). \tag{3}$$

This describes the response toward $\hat{p}$, however, the element and TAB patterns are well defined at all points on the sky above the horizon, $\hat{q} = (\theta, \phi)$ for $0 \le \theta \le \pi/2$ and $0 \le \phi < 2\pi$. To compute the response at an arbitrary sky direction, $\hat{q}$, we set w_n to be the complex conjugate of an impinging planar wave, equation (1), from the direction of $\hat{q}$, i.e. $w_n = \psi_n^{\dagger}(\hat{q}, \nu)$. Therefore, for any sky position, equation (3) becomes

$$A(\hat{q}; \hat{p}, \nu) = \frac{1}{N}\sum_{n=1}^{N} \psi_n^{\dagger}(\hat{q}, \nu)\psi_n(\hat{p}, \nu). \tag{4}$$

This weighting scheme has the effect of ensuring that the squared-modulus, $C(\hat{q}; \hat{p}, \nu) = |A(\hat{q}; \hat{p}, \nu)|^2$, is constrained to the range [0, 1] and is unity if $\hat{p} \equiv \hat{q}$. It can also conveniently be treated as a proxy for the fractional sensitivity of a TAB pointing to some arbitrary sky location, $\hat{q}$, with respect to the look-direction, $\hat{p}$.

From the perspective of a beamforming operation, $C(\hat{q}; \hat{p}, \nu)$ represents regions on the sky from which any signal will be included in the total intensity dynamic spectrum of the beam formed towards $\hat{p}$. The look-direction (and thus, the TAB pattern) is time-dependent for astrophysical sources in observations of any significant duration, i.e. $\hat{p} \to \hat{p}(t)$, but for notation simplicity we will keep the time dependence implicit until it becomes pertinent.

An example of the MWA's compact configuration layout and baseline distribution with 128 tiles is given in Figure 1, which produces the array factor power pattern as shown in the upper panel of Figure 2.

2.1.2. Applying the primary beam apodisation

To retrieve a realistic TAB response to the sky we must include the element pattern response, which can be achieved through multiplying $C(\hat{q}; \hat{p}, \nu)$ by the squared-modulus of the element pattern, often referred to as the primary beam, $B(\hat{q}; \hat{p}, \nu) = |D_{\rm el}(\hat{q}; \hat{p}, \nu)|^2$.

Initially, we have assumed that the element pattern and tied-array pattern are both pointed in the same direction. This is not necessary, nor is it typical for beamformed observations with telescopes with static primary beam patterns. The primary beam may be pointed to any arbitrary sky position, $\hat{b} = (\theta_{\rm pb}, \phi_{\rm pb})$, and is usually static in horizontal (telescope) coordinates, whereas the TAB pattern is the desired 'pixel' to evaluate for maximum sensitivity and usually moves across the sky with time (i.e. tracking a celestial source will result in a changing azimuth and zenith angle

[a] The exact tile-set in operation need not be the same throughout an observing semester, therefore, the TAB shape can change with observing epoch.

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

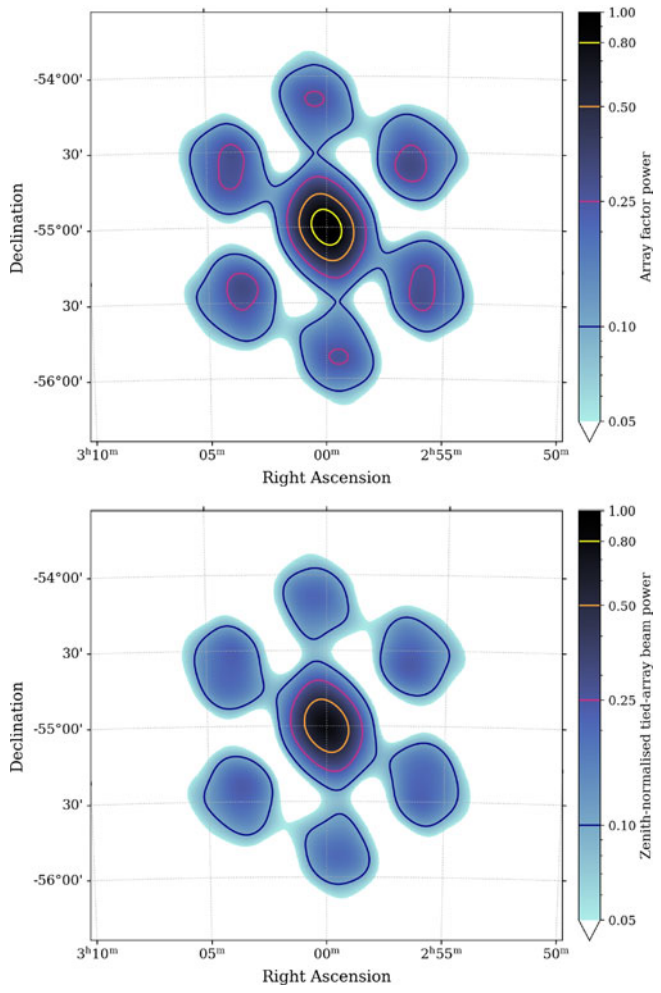


Figure 2. The tied-array beam pattern at 154 MHz based on the element layout in Figure 1. *Top*: A zoom-in of the array factor power pattern, centred on the pointing target direction (approximately 30° off zenith). The effects of the compact array layout are visible in the first side-lobe patterns; a rotated and slightly asymmetrically hexagonal pattern. *Bottom*: The same as the top panel, but it including the primary beam effects and normalised such that the primary beam power is unity at zenith. The target direction is just outside the 80% zenith-normalised power point for this observation.

for all sources). Thus, the primary beam pattern transforms via $B(\hat{q};\hat{p},\nu) \rightarrow B(\hat{q};\hat{b},\nu)$ in the general case.

For the MWA, we use the MWA_HYPERBEAM package[b] to calculate the Full Embedded Element model of the primary beam response (Sokolowski et al. 2017). Specifically, MWA_HYPERBEAM produces a 2×2 complex Jones matrix, $\mathbf{J}$, which represents the polarimetric primary beam response for each provided sky position at a given frequency. These Jones matrices can be converted into a power response for each Stokes parameter via a transformation with the polarisation measurement equation (e.g. Hamaker 2000; van Straten 2004). In our case, as we are specifically concerned with the total intensity (Stokes I) response, the measurement equation reduces to

$$B = \frac{1}{2}\mathrm{tr}\left(\mathbf{J}\mathbf{I}_2\mathbf{J}^{\mathrm{H}}\right), \tag{5}$$

[b] A Rust implementation with Python and C/C++ bindings. The pre-compiled binaries and source code can be found at https://github.com/MWATelescope/mwa_hyperbeam, or the Python package can be retrieved from PyPI.

where tr is the trace operator, $\mathbf{M}^{\mathrm{H}}$ is the Hermitian transpose of a matrix $\mathbf{M}$, and $\mathbf{I}_2$ is the 2×2 identity matrix (see equation 10 in Ord et al. 2019).

2.1.3. Creating the tied-array beam power pattern

Combining the above primary beam response with the array factor response provides a complete description of the theoretical tied-array power pattern, $T(\hat{q};\hat{b},\hat{p},\nu)$, for any given celestial source position, observing frequency, and array configuration, across the entire sky,

$$T(\hat{q};\hat{b},\hat{p},\nu) \equiv C(\hat{q};\hat{p},\nu)B(\hat{q};\hat{b},\nu). \tag{6}$$

In radio aperture synthesis imaging parlance, this is essentially the *naturally weighted* synthesised beam pattern. An example of the TAB pattern based on the element layout in Figure 1 can be seen in the lower panel of Figure 2.

As an aside, the benefits of post-correlation beamforming techniques (e.g. Roy, Chengalur, & Pen 2018; Price 2024), compared to the standard brute-force operations represented here, include: the ability to incorporate additional weighting/tapering schemes which can be applied to reduce the TAB pattern effects (much like in standard synthesis imaging procedures), and that it notionally uses only the cross-correlation terms which can reduce red noise. The downside of post-correlation beamforming is the limited time sampling, which is often driven by the practicalities of $\sim N^2$ visibility production in real-time. This alternate beamforming method allows for flexible balancing of TAB field-of-view and overall survey sensitivity limits (a vital consideration for large pulsar survey projects, especially for the next-generation interferometric radio telescope facilities).

2.2. Statistical localisation map

We adopt the 'TABLo' statistical localisation approach presented by Bezuidenhout et al. (2023). For completeness, we outline the fundamental steps and note any challenges specific to the MWA. Under the assumption of near-uniformity among TAB sensitivity (in a local grid), the starting premise is that the ratio of S/N measurements for each TAB pair is equal to the ratio of the TAB pair's sensitivity. Notionally this holds only if the primary beam does not change drastically over the sampled sky region, or that the primary beam response is included in the model TAB sensitivity pattern. There is also the consideration of the packing and overlap of adjacent TABs around the target candidates, but how to optimally decide this is beyond the scope of the work here (and, generally speaking, not obvious). The S/N measurements, or their proxies, must be generated consistently across TABs and not use dynamic optimisation algorithms (i.e. methods that search parameter spaces to optimise S/N for data sets independently will introduce a bias into the collection that is hard to model).

2.2.1. Per-beam S/N ratios and covariance

First, we construct a vector of ratios of each TAB S/N with the denominator set to the maximum measured S/N among any TAB considered (S_0),

$$\vec{S} = \frac{S_i}{S_0}, \qquad i = 1, ..., K, \tag{7}$$

where S_i is the S/N for TAB T_i, and i enumerates all TABs excluding that which contributes S_0 (i.e. there are $K+1$ total S/N measurements).

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

Creating ratios in this manner imparts a (small) covariance between entries in $\vec{S}$. To characterise this for each element in $\vec{S}$, we simulate a normal distribution where the mean is the entry value itself and the variance $\sigma^2 = 1$ (by definition of S/N). Conducting this simulation L times for each of the K elements in $\vec{S}$ results in a $K \times L$ matrix, which may be used to construct a covariance matrix, $\mathbf{C}$. Defining $\mathrm{cov}(S_i, S_j) = \mathbf{C}_{ij}$ to be the S/N covariance between TAB i and j, then the $K \times K$ covariance matrix elements are

$$\mathbf{C}_{ij} = \frac{1}{L-1}\sum_{m=1}^{L}\left(S_{im} - \bar{S}_{im}\right)\left(S_{jm} - \bar{S}_{jm}\right), \qquad i,j = 1, ..., K, \quad (8)$$

where S_{km} is the element drawn from the mth iteration distribution corresponding to the kth element in $\vec{S}$, and $\bar{S}_{km} \approx S_k$ is the mean value of that distribution.

2.2.2. Statistical metric estimation

Using the result of equation (6) and the elements of the vector in equation (7), a residual vector at each sky position, $\vec{R}(\alpha, \delta)$ is computed such that

$$R_i = S_i - T_i, \qquad i = 1, ..., K, \quad (9)$$

where T_i is the corresponding TAB which provided the measurement S_i. which then allows us to employ a generalised least squares approach to compute a statistical metric that can be used to generate a likelihood map as a function of celestial coordinates. This metric is the χ^2 statistic, the weighted sum of the squared deviations at each sky position,[c]

$$\chi^2(\alpha, \delta) = \vec{R}^T \mathbf{C}^{-1} \vec{R} \quad (10)$$

which may then be used to determine the best-fit sky localisation and corresponding uncertainty. This approach is essentially a 2D application of amplitude-only maximum likelihood radio direction finding (e.g. P_{ML} from Schmidt 1986).

2.2.3. Metric regularisation

The resulting χ^2 map may have many apparently significant points. This is especially true in the MWA's case when using the compact tile configuration, but will be less problematic with a more pseudo-random tile layout or when the field of view is smaller/tied-array beam shape is more regular. There are several reasons for these statistical ambiguities, including a sparse power distribution in the TAB patterns across the sky (which introduces large variations under the ratio operation), and that we completely ignore the absolute brightness/real flux density of the signal (because it is not always guaranteed to exist, whereas the S/N is commonly available). These two realities mean that a typical $\chi^2(\alpha, \delta)$ map has several minima yielding a degenerate and multi-modal localisation map. To mitigate this, we can re-weight the statistic map based on expectations of how far a typical source of interest can be from the TAB centre and still be significantly detected. However, the choice of weighting function is not always obvious. Regularisation in this sense suppresses information based on assumed knowledge. We recommend always inspecting the probability maps without such rescaling as it can reveal, for instance, side-lobe associations.

[c]This is a tensor operation because the residual vector is a tensor of shape (K, N_α, N_δ), where N_α and N_δ are the number of pixels in the R.A. and Dec. axes, respectively, and the transpose operation is only along the first dimension, i.e. the dimension corresponding to the TAB index.

We elect, by default, to re-weight the statistical map with the TAB pattern with the greatest S/N detection, i.e. $\chi^2 \to T_0\chi^2$. This approach favours weaker but nearby (to the T_0 pointing direction within the grid of TABs) likely localisations over the ones that are significantly stronger but further away. Essentially, this means that a signal detected in the primary lobe of the TAB would need to be commensurately brighter if it were instead located in a less-sensitive part of the TAB for it to be detected with the same S/N. On the downside, this drastically reduces the effective sky-area that is considered when computing the localisation probability by essentially down-weighting side-lobe contributions of other TABs that do not overlap with those of T_0. (See Appendix A.)

One could instead use a less restrictive regularisation such as a multivariate Gaussian weighting function centred at the position of T_0, with a width as some multiple of the maximum distance between any TAB pair. This still favours the T_0 pointing direction but allows information from scales near or beyond the first side-lobes of multiple TABs to be incorporated.

2.2.4. Localisation probability maps

Given the underlying χ^2 basis, the log-likelihood at any celestial position, (α, δ) is,

$$\ln \mathcal{L}(\alpha, \delta) \propto -\frac{1}{2}\chi^2(\alpha, \delta), \quad (11)$$

and, hence, the relative probability density at a given celestial position becomes

$$\Pr(\alpha, \delta) \propto \mathcal{L}(\alpha, \delta) \Big/ \int_{\alpha,\delta} \mathcal{L}(\alpha, \delta). \quad (12)$$

For our purpose, where the (α, δ) parameter space has been discretised on a grid, equations (11) and (12) become

$$\ln \mathcal{L}[\alpha, \delta] \approx -\frac{1}{2}\chi^2[\alpha, \delta], \quad (13)$$

and

$$\Pr[\alpha, \delta] \approx \mathcal{L}[\alpha, \delta] \Big/ \sum_{\alpha,\delta} \mathcal{L}[\alpha, \delta]. \quad (14)$$

The best localisation is then identified from equation (14) via the argmax operation.

There is an important caveat to the above approach, namely that using χ^2 as a basis implicitly assumes that the observations encoded through T_i and S_i are independent of each other and at least close to being normally distributed. Given the way in which the TAB data are formed (by applying different complex phase rotation to the same underlying voltage streams), this assumption is not accurate. While we do estimate a covariance matrix for $\vec{S}$, this covariance is mainly between the measured signal strengths and not their spatial covariance. Furthermore, another implicit assumption of the framework is that any of the 2π sr of sky not simulated as part of the TAB patterns contributes negligibly to the observed signals.

2.2.5. Localisation uncertainty estimate

The uncertainty in the localisation is computed using a similar method to that described in Bezuidenhout et al. (2023), based on the Mahalanobis distance (Mahalanobis 1936; Anderson 2003), d_M, under the assumption of a nearly-Gaussian probability distribution. By convention we use the 5σ equivalent d_M to derive a positional uncertainty.

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

In practice, there may be multiple 'islands' of probability above the desired d_M threshold, however the island of interest is that which contains the peak probability density value. To determine this and the uncertainty region, we first create a mask that highlights only pixels with a value above the d_M threshold. The mask is then passed into a feature identification and labelling algorithm,[d] and the island containing the peak localisation pixel is selected. The Euclidean distance is computed from the peak pixel to each other pixel in the island, and the maximum value is taken as the conservative symmetric localisation uncertainty. This can be extended to any arbitrary $n\sigma$ interval by simply re-evaluating the d_M and the corresponding contour boundaries.

Due to the probabilistic nature of the technique (e.g. random sampling of Gaussian distributions to estimate the covariances) the final localisation and uncertainty can also change despite with identical inputs. In part, the magnitude of this random shift is driven by the sky sampling grid, as is the absolute probability density value, hence we only make use of the relative probabilities and do not place any strong meaning on the actual numerical quantities resulting from equation (14). If the grid is too coarse, then the localisation method is insensitive to the small-scale TAB shape which is crucial to distinguishing the relative TAB powers and, therefore, the final localisation. If the grid is too fine, the compute and memory cost increases without meaningful improvement.[e] Ultimately, the localisation peak and corresponding uncertainty shift by ∼1 pixel equivalent width. Furthermore, while the assumption of well-calibrated interferometric element gains is reasonable it is not perfect, thus there will also be some systematic bias due to residual differences (though nominally the same within a given observation).

The 'pixel scale' will necessarily depend on the observing setup, decreasing as frequency and/or baselines increase. Generally, the pixel scale should correspond to no larger than ∼1/10 of the TAB FWHM which ensures oversampling and a dynamic range suitable for identifying the most restrictive uncertainty intervals. In the following analysis, we choose a pixel size of 5″, corresponding to ∼300 pixels across the nominal 24′ TAB FWHM for SMART.

3. Verification with known pulsar positions

Here we demonstrate the localisation method using pulsars with well-constrained positions, starting with two specific examples alongside a broader sample study to explore the general accuracy of the localisation method. Generally speaking, the predicted pulsar position relative to its interferometrically imaged location is offset by angles similar to those expected from random source position shifts induced by ionospheric turbulence and refractive effects (∼0.1–1′) in this frequency range (e.g. Arora et al. 2015; Jordan et al. 2017; Swainston et al. 2022; Waszewski, Morgan, & Jordan 2022). To be conservative, one can add in quadrature ∼1′ to the quoted systematic uncertainty.

3.1. SMART PSR J0026−1955

The second discovered SMART pulsar (which was an independent re-discovery of the Green Bank North Celestial Cap (GBNCC) pulsar survey candidate J0025−19) went through a number of localisation iterations before being imaged (McSweeney et al. 2022; Bhat et al. 2023b) with the upgraded Giant Metrewave Radio Telescope (uGMRT). This pulsar was in the minority where there were multiple archival MWA VCS data sets, including some with longer baseline configurations, which were used to constrain the pulsar position. Here, we demonstrate the localisation that would have been capable with just the initial SMART detections from 10-min of data in adjacent beams. We re-beamformed the data using the same grid setup as the original search data for 61 adjacent beams. Figure 3 (left) demonstrates the MWA localisation from only the discovery observation. The high angular resolution imaged pulsar position (in the J2000 epoch) is $\alpha_{\rm img} = 00^{\rm h}26^{\rm m}36^{\rm s}.3(2)$ and $\delta_{\rm img} = -19°55'59''.3(3)$ (Bhat et al. 2023b). Our localisation method predicts a position of $\alpha_{\rm loc} = 00^{\rm h}26^{\rm m}32^{\rm s}.65$ and $\delta_{\rm loc} = -19°55'49.30''$, with $(1\sigma, 3\sigma, 5\sigma)$ uncertainties of (0.3′, 0.9′, 1.6′) respectively, not including any ionospheric shift. The two positions are offset by ∼0.87′, thus match at a $\sim 3\sigma$ level.

3.2. SMART PSR J0452−3418

The fifth discovered SMART pulsar, PSR J0452−3418 (Grover et al. 2024), was initially localised via a secondary (incomplete) gridded beamforming analysis (i.e. not the SMART gridded observations themselves) which was used to estimate the pulsar position. Targeted phased array beams from the uGMRT did not initially detect pulsations, thus the visibilities recorded in tandem were processed to produce synthesis images which ultimately localised the pulsar with a positional uncertainty of a few arcseconds. The original MWA localisation had a 3′ uncertainty, ignoring ionospheric effects, from a weighted centroid combination and was ultimately shown to be offset from the image-based position by ∼2′. The pulsar position (in the J2000 epoch), again from high angular resolution imaging, is $\alpha_{\rm img} = 04^{\rm h}52^{\rm m}00^{\rm s}.7(3)$ and $\delta_{\rm img} = -34°18'42''(4)$ (Grover et al. 2024; Bhat et al. 2023b). Figure 3 (right) demonstrates the MWA localisation using the method, based on the same information initially used for the centroid-based refinement. The predicted position is $\alpha_{\rm loc} = 04^{\rm h}52^{\rm m}01^{\rm s}.15$ and $\delta_{\rm loc} = -34°18'25''.69$, with $(1\sigma, 3\sigma, 5\sigma)$ uncertainties of (0.3′, 0.9′, 1.6′) respectively, not including any ionospheric shift. The positions are offset by ∼0.29′, thus match at the $\sim 1\sigma$ level.

3.3 A sample of known pulsars

It is not straightforward to simulate and adequately capture the complexity of the tied-array beam and its impact on the S/N of fake targets. Instead, to approach the problem of general localisation accuracy, we produce several beamformed data sets of known pulsars in various places within the primary beam of the telescope and use these to empirically estimate the systematics involved.

We use a single observation (SMART G06; MWA Observation ID 1266329600) containing many detectable pulsars distributed across the primary beam, typically[f] ranging from peak S/N of ∼3 to >50 when targeted directly. In total we created TAB grids

[d] Using the ndimage.label function from SciPy, with the default structural element matrix.

[e] The scaling increases with the number of pixels, $N_\alpha N_\delta$.

[f] There are a handful of pulsars that are barely detectable and/or extremely scattered, so we did not use those pulsars for this exercise owing to the complexity of defining a sensible S/N.

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

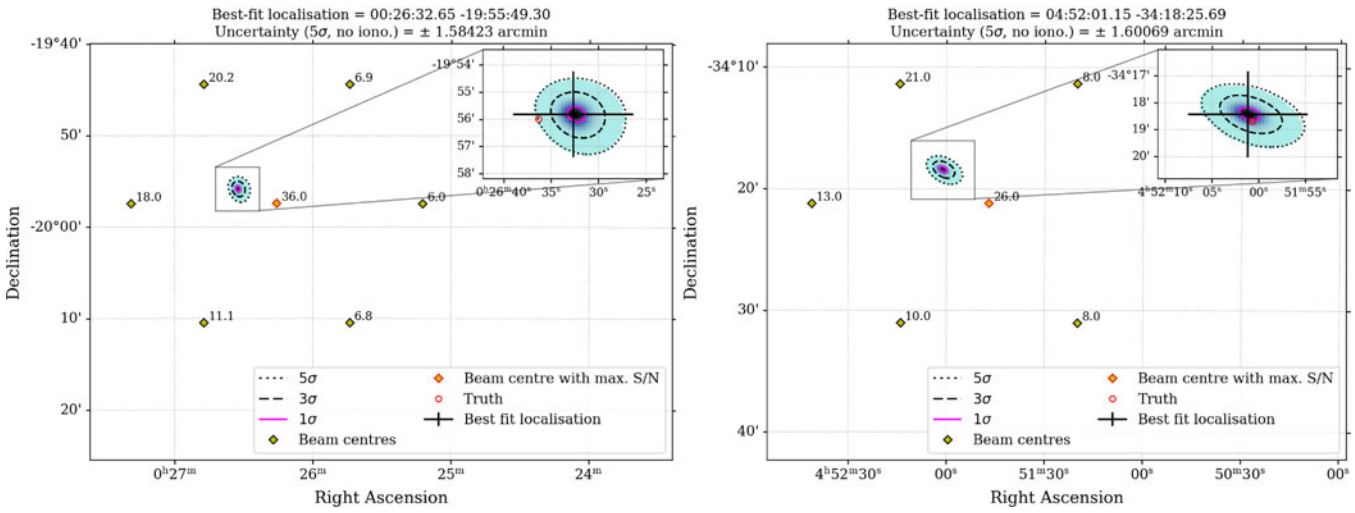


Figure 3. Localisation maps for two example SMART pulsar discoveries. Each TAB in the grid their nominal pointing coordinates marked as a yellow diamond, with their corresponding measured S/N. The TAB with the greatest S/N is also outlined in red, and the known pulsar position based on follow-up observations with higher-resolution telescopes (small red circle in the inset figure, labelled as 'Truth') is reasonably well predicted. Contours representing the 1, 3, and 5σ uncertainties based on the probability density are overlaid, and the best-fit position with its corresponding symmetrical uncertainty is plotted in the inset. *Left*: The localisation of PSR J0026−1955 with the MWA using the original follow-up grid detections. The 1, 3 and 5-σ systematic uncertainties in this localisation instance are 0.3′, 0.9′, and 1.6′ respectively, with a nominal offset from the known position of 0.87′. *Right*: The localisation of PSR J0452−3418 with the MWA given the targeted grid detections (i.e. not the original SMART grid detections). The 1, 3 and 5-σ systematic uncertainties in this localisation instance are 0.3′, 0.9′, and 1.6′ respectively, with a nominal offset from the known imaging position of 0.29′.

around, and localised, 15 known pulsars. For each pulsar position we applied an angular offset to it's known position in a random direction with a magnitude of 10–30′ (i.e. between a half and a full SMART TAB width) and formed a hexagonal grid of overlapping TABs around the offset central position, similar to how the SMART sky tessellation is achieved. Each TAB was generated from the full 80-min drift scan data and folded without S/N optimisation enabled. The folding period and dispersion measure were gathered from the SMART data release provided by Bhat et al. (2026, submitted) and passed to the prepfold routine in the PRESTO software suite (Ransom 2011). For each TAB position the S/N was computed from the profiles by first identifying the 'on-pulse' and 'off-pulse' regions in the best detection, obtaining a noise estimate from the off-pulse region and comparing it to the maxima in the 'on-pulse' region.

In Figure 4 we plot the offset from the known position of the pulsar and the best-fit localisation against the equivalent 5σ uncertainty. Ideally, the data would lie on the equivalence line, or slightly below it, indicating that the uncertainty estimate encompasses the offset from the known position. This is the case for many of the examples, especially when including a conservative error contribution of ∼1′ from ionospheric shifts. Many of the pulsar detections are marginal, and the range of S/N measurements across beams is small, which ultimately limits the information available and, therefore, limits the localisation precision. Broadly speaking the uncertainty region and the offset from known positions are relatively close, especially in the high S/N regime. As the detection S/N reduces, or if the diversity of S/N measurements over the TAB grid is small, the localisation effectiveness also reduces as the information provided loses constraining power.

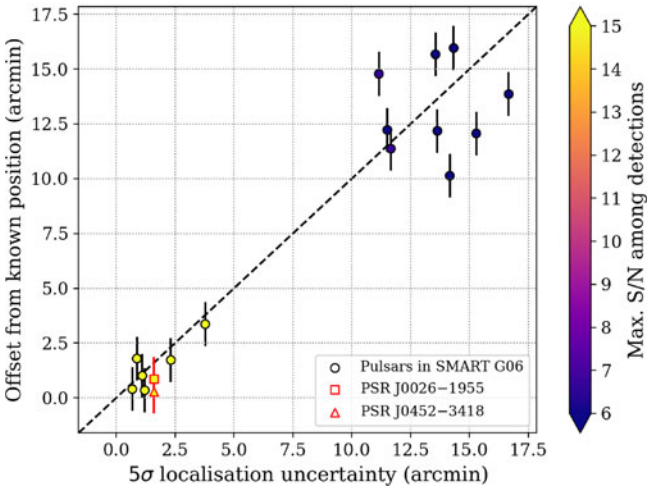


Figure 4. The localisation uncertainty vs. the offset of the known position from the best-fit localisation position for PSRs J0026−1955 (red outlined square), J0452−3418 (red outline triangle), and 15 detectable pulsars from the SMART G06 observation (circles). A one-to-one line is drawn for comparison. Generally, the localisation uncertainty matches any offset from the known pulsar position reasonably well, with increasing magnitude and spread as the S/N diminishes. Vertical error bars represent the maximum ionospheric refractive shift one might expect (∼1′) in the known position vs. observed position.

4. MWA pulsar and transient candidate follow-up

Instruments with a large native synthesised beam (and, therefore, tied-array beam) pattern compared to any plausible follow-up instrument incur a corresponding large overhead. This is essentially due to the higher-resolution follow-up instrument

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press

needing to conduct a (smaller scale) search across the broader synthesised beam area. For SMART in particular this is challenging given the 20–30′ TAB width and the available follow-up instruments with overlapping sky coverage. In the case of single dish radio follow-up, Parkes/*Murriyang* has a beam width 25.8−6.6′ with the Ultra Wide-bandwidth Low (UWL) receiver frequency coverage (0.7–4 GHz), and the Green Bank Telescope has a beam width 25.1−3.6′ with the Ultra Wideband Receiver (UWBR) coverage (0.5–3.5 GHz). For interferometric follow-up, the uGMRT synthesised beam width is 3.5′ at 350 MHz assuming only the central 1 km antennas, while for MeerKAT it is 0.28′ at 600 MHz with a maximum 8 km baseline.

An illustrative example of this challenge is provided by Grover et al. (2024), where multiple iterations of imaging and beamforming on subsequently detected targets was eventually required to localise PSR J0452−3418. Use of the methods described here would have saved much of the exploratory time (see Section 3.2), which resulted in hours of uGMRT time being used to search for rather than observe the target. Without localisation to a handful of arcminutes, targeted MeerKAT TAB gridding observations are expensive and/or infeasible. A recent SMART follow-up campaign with MeerKAT, which still required its own TAB gridding and localisation processing, led to a successful candidate localisation of PSR J0125−5854 (Tan et al. 2026, submitted) based on the initial localisation provided by the method defined here.

5. Summary

The method demonstrated here is an efficient way to obtain super-resolved localisation of events for telescopes: (1) when imaging and pulsar observing modes are not contemporaneous or easy to produce, (2) where multiple simultaneous TABs can be formed, and (3) when a reasonable model of the primary beam is available (since the TAB simply relies on interferometric element positions). The approach has been successfully adapted and demonstrated on at least two SKA precursors/demonstrators for both the Low and Mid telescopes, and may be useful for certain SKA science objectives and data processing tasks. Pathways worth exploring in future include: implementing subband localisation (where the frequency-dependence of the TABs may be exploitable to improve precision), and multi-epoch combinations (where the time-dependent refractive shifts from the ionosphere could be extracted). Both trajectories would require special attention be paid to the ionosphere, and would help to quantify the effect on localisation given their time and frequency dependence. Additionally, with the Phase 3 expansion of the MWA (where all 256 tiles may be correlated/combined simultaneously, and a real-time beamforming mode is under development) the arcminute-level positions possible with SMART VCS observations and the described localisation technique provide an excellent starting point for MWA follow-up and position refinement of it's own candidates before moving to other instruments.

For SMART, this approach will allow for rapid convergence on robust pulsar candidate positions, which will then accelerate follow-up observations with independent telescopes that have much smaller fields of view and/or are operating at higher frequencies. It will also prove useful in localising single-pulse events, such as Rotating Radio Transients (RRATs), potential Fast Radio Burst candidates or Gamma-ray Burst prompt radio emission, which have both been quite elusive thus far at low frequencies.

Acknowledgements. This scientific work uses data obtained from Inyarrimanha Ilgari Bundara, the Murchison Radio-astronomy Observatory. We acknowledge the Wajarri Yamaji People as the Traditional Owners and native title holders of the Observatory site. Establishment of CSIRO's Murchison Radio-astronomy Observatory is an initiative of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited.

This work was supported by resources provided by the Pawsey Supercomputing Research Centre's Setonix Supercomputer (https://doi.org/10.48569/18sb-8s43) and Garrawarla GPU Cluster (https://doi.org/10.48569/gskb-tp15), with funding from the Australian Government and the Government of Western Australia.

The authors thank Ramesh Bhat, Adrian Sutinjo and Sam McSweeney for useful discussions and suggestions throughout the development and writing process. We also thank the anonymous referee for their valuable feedback and suggestions.

Data availability statement. The code developed here can be found on GitHub and is publicly available under the Academic Free License v3. The data and example scripts used to generate the figures in this manuscript are also included.

Software. ASTROPY (Astropy Collaboration et al. 2013, 2018, 2022), NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), ARVIZ (Kumar et al. 2019), MATPLOTLIB (Hunter 2007), MWALIB (github.com/MWATelescope/mwalib), MWA_HYPERBEAM (github.com/MWATelescope/mwa_hyperbeam), PRESTO (Ransom 2011).

References

Anderson, T. W. 2003, An Introduction to Multivariate Statistical Analysis, Wiley Series in Probability and Statistics (3rd edn.; Hoboken, New Jersey: Wiley)
Arora, B. S., et al. 2015, PASA, 32, e029
Astropy Collaboration, et al. 2013, A&A, 558, A33
Astropy Collaboration, et al. 2018, AJ, 156, 123
Astropy Collaboration, et al. 2022, ApJ, 935, 167
Balanis, C. A. 2016, Antenna Theory: Analysis and Design (4th edn.; Hoboken, New Jersey: Wiley)
Bannister, K. W., et al. 2017, ApJ, 841, L12
Bezuidenhout, M. C., et al. 2023, RASTI, 2, 114
Bhat, N. D. R., et al. 2023a, PASA, 40, e021
Bhat, N. D. R., et al. 2023b, PASA, 40, e020
Capon, J. 1969, Proc. IEEE, 57, 1408
CHIME/FRB Collaboration, et al. 2019, ApJ, 885, L24
Grover, G., et al. 2024, ApJ, 970, 78
Hamaker, J. P. 2000, A&AS, 143, 515
Harris, C. R., et al. 2020, Natur, 585, 357
Hunter, J. D. 2007, CSE, 9, 90
Hurley-Walker, N., et al. 2022, PASA, 39, e035
Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, A&A, 598, A78
Jordan, C. H., et al. 2017, MNRAS, 471, 3974
Kumar, R., Carroll, C., Hartikainen, A., & Martin, O. 2019, JOSS, 4, 1143
Lee, C. P., et al. 2025, arXiv e-prints, arXiv:2508.10330
Lipsky, S. E. 2004, Microwave Passive Direction Finding, EngineeringPro Collection (Raleigh, North Carolina: SciTech Publishing)
Mahalanobis, P. C. 1936, PNISI, 2, 49
McSweeney, S. J., et al. 2020, PASA, 37, e034
McSweeney, S. J., et al. 2022, ApJ, 933, 210
Meyers, B. W., et al. 2017, ApJ, 851, 20
Morrison, I. S., et al. 2023, PASA, 40, e019
Obrocka, M., Stappers, B., & Wilkinson, P. 2015, A&A, 579, A69

Ord, S. M., et al. 2019, PASA, 36, e030
Price, D. C. 2024, PASA, 41, e037
Ransom, S. 2011, PRESTO: PulsaR Exploration and Search TOolkit, Astrophysics Source Code Library, record ascl:1107.017, ascl:1107.017
Roy, J., Chengalur, J. N., & Pen, U.-L. 2018, ApJ, 864, 160
Schmidt, R. 1986, IEEE TAP, 34, 276
Sett, S., Bhat, N. D. R., Sokolowski, M., & Lenc, E. 2023, PASA, 40, e003
Sett, S., Sokolowski, M., Lenc, E., & Bhat, N. D. R. 2024, PASA, 41, e045
Sokolowski, M., et al. 2017, PASA, 34, e062
Stappers, B., & Kramer, M. 2016, in MeerKAT Science: On the Pathway to the SKA, 9
Sutinjo, A., et al. 2015, RSci, 50, 52
Swainston, N. A., et al. 2022, PASA, 39, e020
Swainston, N. A., et al. 2021, ApJ, 911, L26
Teh, C.-H., & Chin, R. T. 1986, CVGIP, 33, 318
Thomas, S., et al. 2006, MNRAS, 371, 323
Tingay, S. J., et al. 2013, PASA, 30, e007
Tremblay, S. E., et al. 2015, PASA, 32, e005
van Straten, W. 2004, ApJS, 152, 129
Virtanen, P., et al. 2020, NM, 17, 261
Vyas, A., Prasad, B., & Roopashree, M. 2009, in International Conference on Advances in Recent Technologies in Communication and Computing (IEEE Computer Society), 366
Waszewski, A., Morgan, J., & Jordan, C. H. 2022, PASA, 39, e036
Wayth, R. B., et al. 2015, PASA, 32, e025
Wayth, R. B., et al. 2018, PASA, 35, e033
Xue, M., et al. 2019, PASA, 36, e025

Appendix A. Effects of differing regularisation methods

To demonstrate the effects of the different regularisation techniques, we take PSR J0026−1955 as an example. In particular we focus on the original side-lobe detection of the pulsar (see McSweeney et al. 2022; Bhat et al. 2023b). The pulsar was further localised after more candidates were collected from SMART data processing and it was realised that the first was a side-lobe detection. The follow-up grid formed thereafter is what is replicated in Section 3.1.

Here we compare what the localisation would have been with only the original side-lobe detection information. The comparison is visualised in Figure A1. As noted in the main text, it is important to always inspect the un-regularised localisation map, especially when there is no other *a priori* information known about the object sky position. Classical centroiding algorithms would also suffer from the same issue as the TAB-weighted case.

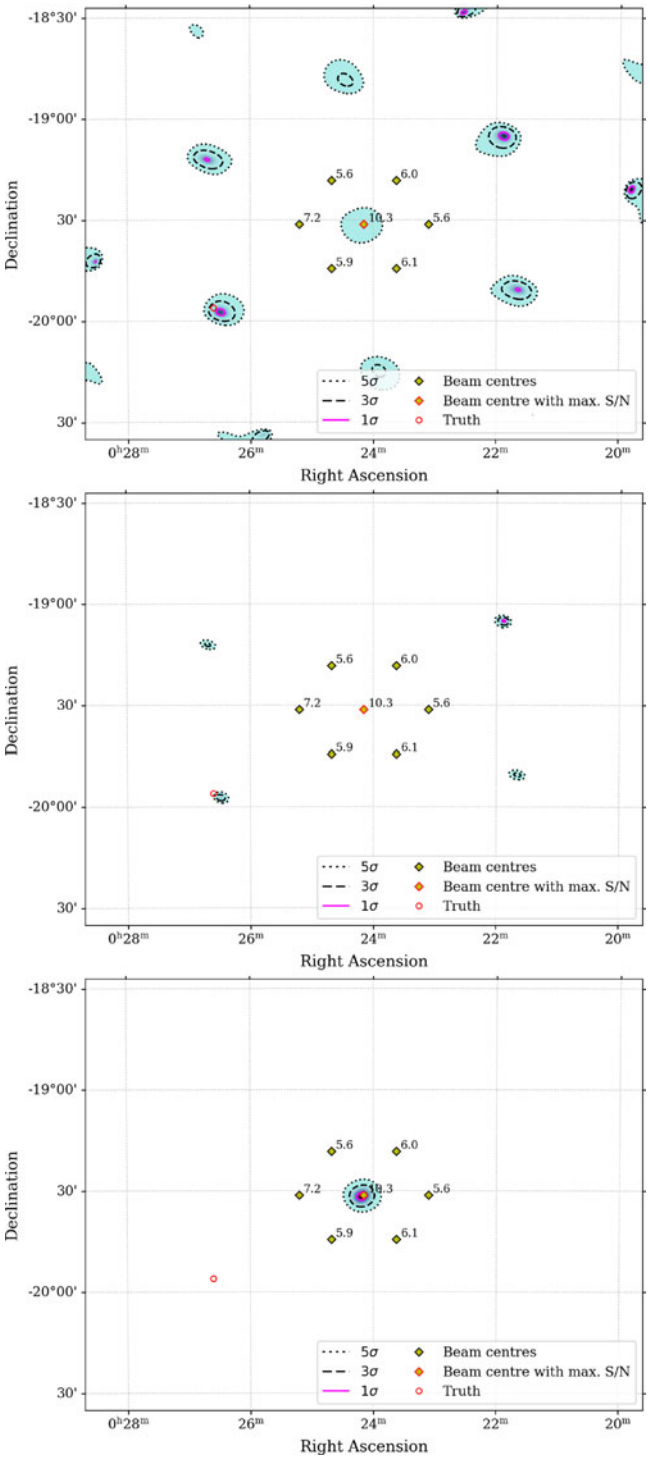


Figure A1. The localisation probability map computed over the same sky area for the same inputs based on an initial side-lobe detection of PSR J0026−1955, with three different regularisation schemes. The figure legend in the lower sub-figure applies to the others. *Top*: Without any re-weighting of the statistical map. *Middle*: Regularised with a wide Gaussian window centred at the TAB with the largest detection S/N. *Bottom*: Regularised with the TAB pattern corresponding to the TAB with the largest detection S/N. Both the Gaussian regularised and un-regularised maps suggest that follow-up is required near the (ultimately correct, based on interferometric imaging) pulsar position, which would ultimately lead to more significant detections and a more confident localisation.

https://doi.org/10.1017/pasa.2026.10174 Published online by Cambridge University Press