



CHIME All-sky Multiday Pulsar Stacking Search (CHAMPSS): System Overview and First Discoveries

Christopher Andrade¹ , P. J. Boyle² , Charanjot Brar³ , Alyssa Cassity¹ , Kathryn Crowter¹ , Davor Cubranic⁴ ,
 Abigail K. Denney⁵ , Fengqiu Adam Dong⁶ , Emmanuel Fonseca^{7,8} , Victoria M. Kaspi^{2,9} , Ajay Kumar¹⁰ ,
 Lars Kunkel¹ , Magnus L'Argent² , Dustin Lang¹¹ , Robert A. Main^{2,9} , Kiyoshi W. Masui^{12,13} , Sujay Mate^{14,15} ,
 Juan Mena-Parra^{5,16} , Bradley W. Meyers^{17,18} , Cherry Ng¹⁹ , Aaron B. Pearlman^{2,9,25,26,27} , Ue-Li Pen^{11,16,20,21,22} ,
 Scott M. Ransom⁶ , Alexander P. Roman^{11,23} , Kendrick Smith¹¹ , Reynier Squillace^{6,24} , Ingrid Stairs¹ , Chia Min Tan¹⁸ ,
 Laurent Tarabout² , Wenke Xia² , and Tarik J. Zegmott^{2,9}

(The CHAMPSS Collaboration)

¹ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada; lkuenkel@phas.ubc.ca

² Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada; robert.main@mcgill.ca

³ NRC Herzberg Astronomy and Astrophysics, 5071 West Saanich Road, Victoria, BC V9E2E7, Canada

⁴ AbCellera Biologics, 150 W 4th Avenue, Vancouver, BC V5Y 1G6, Canada

⁵ David A. Dunlap Department of Astronomy and Astrophysics, 50 St. George Street, University of Toronto, ON M5S 3H4, Canada

⁶ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

⁷ Department of Physics and Astronomy, West Virginia University, PO Box 6315, Morgantown, WV 26506, USA

⁸ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

⁹ Trotter Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada

¹⁰ National Centre for Radio Astrophysics, Post Bag 3, Ganeshkhind, Pune, 411007, India

¹¹ Perimeter Institute of Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada

¹² MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

¹³ Department of Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

¹⁴ Raman Research Institute, C. V. Raman Avenue, Sadashivanagar, Bangalore, Karnataka-560080, India

¹⁵ Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research, Mumbai, 400005, India

¹⁶ Dunlap Institute for Astronomy and Astrophysics, 50 St. George Street, University of Toronto, ON M5S 3H4, Canada

¹⁷ Australian SKA Regional Centre (AusSRC), Curtin University, Bentley WA 6102, Australia

¹⁸ International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley WA 6102, Australia

¹⁹ Laboratoire de Physique et Chimie de l'Environnement et de l'Espace-Université d'Orléans/CNRS, 45071, Orléans Cedex 02, France

²⁰ Institute of Astronomy and Astrophysics, Academia Sinica, Astronomy-Mathematics Building, No. 1, Sec. 4, Roosevelt Road, Taipei 10617, Taiwan

²¹ Canadian Institute for Theoretical Astrophysics, 60 St. George Street, Toronto, ON M5S 3H8, Canada

²² Canadian Institute for Advanced Research, 180 Dundas St West, Toronto, ON M5G 1Z8, Canada

²³ Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

²⁴ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

Received 2025 April 22; revised 2025 June 23; accepted 2025 June 24; published 2025 August 25

Abstract

We describe the Canadian Hydrogen Intensity Mapping Experiment (CHIME) All-sky Multiday Pulsar Stacking Search (CHAMPSS) project. This novel radio pulsar survey revisits the full northern sky daily, offering unprecedented opportunity to detect highly intermittent pulsars, as well as faint sources via long-term data stacking. CHAMPSS uses the CHIME/FRB datastream, which consists of 1024 stationary beams streaming intensity data at 0.983 ms resolution, with 16,384 frequency channels across 400–800 MHz, continuously being searched for single, dispersed bursts/pulses. In CHAMPSS, data from adjacent east–west beams are combined to form a grid of tracking beams, allowing longer exposures at fixed positions. These tracking beams are dedispersed to many trial dispersion measures (DM) to a maximum DM beyond the Milky Way's expected contribution, and Fourier transformed in time to form power spectra. Repeated observations are searched daily to find intermittent sources, and power spectra of the same sky positions are incoherently stacked, increasing sensitivity to faint persistent sources. The 0.983 ms time resolution limits our sensitivity to millisecond pulsars; we have full sensitivity to pulsars with $P > 60$ ms, with sensitivity gradually decreasing from 60 ms to 2 ms, as higher harmonics are beyond the Nyquist limit. In a commissioning survey, data covering $\sim 1/16$ of the CHIME sky were processed and searched in quasi-realtime over two months, leading to the discovery of 11 new pulsars, each with $S_{600} > 0.1$ mJy. When operating at scale, CHAMPSS will stack >1 yr of data along each sightline, reaching a sensitivity of $\lesssim 30$ μ Jy for all sightlines above a decl. of 10° , and off of the Galactic plane.

Unified Astronomy Thesaurus concepts: Pulsars (1306)

²⁵ Banting Fellow.

²⁶ McGill Space Institute Fellow.

²⁷ FRQNT Postdoctoral Fellow.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

There are >3700 pulsars that have been discovered to date (R. N. Manchester et al. 2005), enabling a wealth of physics and astrophysics. Relativistic binary systems revealed indirect evidence for gravitational waves (J. H. Taylor & J. M. Weisberg 1989), and they continue to be a laboratory for the most precise

tests of general relativity (e.g., M. Kramer et al. 2021). Pulsar timing arrays, which combine the timing residuals of the most precisely timed pulsars, are starting to detect evidence for \sim nHz gravitational waves from supermassive black hole binaries (e.g., G. Agazie et al. 2023; EPTA Collaboration et al. 2023; D. J. Reardon et al. 2023). By measuring pulsar masses (P. B. Demorest et al. 2010; E. Fonseca et al. 2021) and radii (M. C. Miller et al. 2021), the neutron star equation of state is constrained, thereby providing insight into how matter behaves at extreme densities (J. M. Lattimer & M. Prakash 2007).

Pulsar signals are also a powerful probe of intervening plasma. Pulses are dispersed, acquiring a frequency-dependent delay directly proportional to the total electron column along the line of sight. Moreover, pulses are often highly linearly polarized, allowing measures of the Faraday rotation, probing intervening magnetic fields. Pulsars experience multipath propagation owing to small-scale variations in electron density, leading to scattering, and as pulsars are effectively point sources, interference effects between multiple deflected paths, known as “scintillation.” Measuring the above effects on many pulsar sightlines has led to Galactic electron models (J. M. Cordes & T. J. W. Lazio 2002; J. M. Yao et al. 2017), foreground maps of Galactic magnetism (J. L. Han et al. 2006), and a holographic view of plasma substructure on $\lesssim 0.1$ au (D. R. Stinebring et al. 2022).

To study pulsars, first they must be discovered. Most pulsar surveys to date used traditional parabolic dishes. Due to the small field of view (FoV), the strategy is typically to gradually tile the sky with pointings, or to focus on regions where pulsars are a priori expected to reside. The most obvious choice is to search in the Galactic plane, although pulsars also reside in globular clusters (>340 to date)²⁸ and supernova remnants (M. I. Large et al. 1968; D. H. Staelin & E. C. Reifenstein 1968), are often associated with γ -ray sources (D. J. Thompson et al. 1994), and can sometimes be seen as compact steep-spectrum radio continuum sources (D. C. Backer et al. 1982). Both the Five-hundred-meter Aperture Spherical Telescope (FAST) and MeerKAT are performing a Galactic plane survey (P. V. Padmanabh et al. 2023; J. L. Han et al. 2025), also done as part of the Parkes Multibeam Survey (R. N. Manchester et al. 2001).

All-sky surveys such as the Green Bank Northern Celestial Cap Pulsar Survey (GBNCC; K. Stovall et al. 2014) require a much longer time to survey the full sky, and they are typically limited to one short exposure per pointing. Low-frequency telescopes that form beams digitally, such as the Low-Frequency Array (LOFAR) and Murchison Widefield Array (MWA), can have a much larger FoV and survey their full visible sky more rapidly. The LOFAR Tied-Array All-Sky Survey (LOTAAS) forms many tied-array beams in real time within the LOFAR FoV (from just the inner dense core stations) at ~ 135 MHz, and searches them offline (S. Sanidas et al. 2019). Similarly, the Southern-sky MWA Rapid Two-meter (SMART) pulsar survey utilizes the MWA FoV (in its compact configuration) at ~ 154 MHz to rapidly survey the southern sky for pulsars down to $\sim 2\text{--}3$ mJy with up to 80 minute dwell times, but at the cost of very large offline computation and storage footprints by virtue of processing the raw tile voltage data (N. D. R. Bhat et al. 2023a, 2023b).

In this paper, we describe a new pulsar survey using the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME Collaboration et al. 2022). The unique cylindrical nature of CHIME means it has a much larger field of view than a parabolic dish, and it is an ideal survey instrument. CHIME has been transformative to the field of Fast Radio Bursts (CHIME/FRB Collaboration et al. 2021), and the real-time system to find individual bursts has already led to the discovery of more than 80 pulsars²⁹ (D. C. Good et al. 2021; F. A. Dong et al. 2023). The CHIME All-sky Multiday Pulsar Stacking Survey (CHAMPSS) aims to use the large field of view of CHIME to carry out a daily full-sky fast Fourier transform (FFT) search for pulsars. Power spectra from repeated observations of the same positions will be incoherently stacked (M. van der Klis 1989), a method that has been successfully used to discover pulsars in globular clusters (S. B. Anderson 1993; Z. Pan et al. 2016; M. Cadelano et al. 2018). Through the combination of daily searching and stacking of the full northern sky, this survey will be deeper than any other full-sky survey to date, and it will be sensitive to intermittent sources that could, by chance, have been missed previously. This can occur if, e.g., pulsars are scintillating (B. J. Rickett 1990), eclipsed (A. S. Fruchter et al. 1988; S. Johnston et al. 1992), precessing (R. P. Breton et al. 2008), nulling (D. C. Backer 1970), or intrinsically intermittent (A. Lyne et al. 2010), or if observations simply are corrupted by radio frequency interference (RFI). CHAMPSS will additionally lead to pulsar detections on undersearched lines of sight, providing a better sampling of the Galactic electron structure. Compared to similar existing surveys, we will be as sensitive as or better than GBNCC within 1 month stacks for decl. $>10^\circ$, and will reach comparable sensitivity to the FAST Galactic Plane Pulsar Snapshot survey (GPSS, J. L. Han et al. 2025) above a decl. of $>30^\circ$ if we are able to successfully stack ~ 1 yr.

The distribution of the paper is as follows: In Section 2, we describe CHIME, as well as the FRB and Pulsar datastreams relevant for CHAMPSS. In Section 3, we discuss the pipeline, how power spectra are formed from the incoming data stream, and how candidates are then sifted and grouped. In Section 4, we detail the process of candidate verification through a phase-coherent search, as well as the timing of newfound sources. In Section 5, we describe our operations and real-time system for processing and analyzing data. In Section 6, we describe our commissioning survey, first discoveries, and implications. Section 7 details the current status of CHAMPSS, its planned expansion, and the forecast for the full survey.

2. CHIME Systems

Located at the Dominion Radio Astrophysical Observatory near Penticton, British Columbia, CHIME is a radio telescope comprised of 4×100 m parabolic cylinders, oriented north–south (CHIME Collaboration et al. 2022). Each cylinder has 256 dual-polarization linear feeds, operating in the frequency range 400–800 MHz. CHIME operates as a drift-scan telescope, observing the full northern sky above a decl. of -10° , with an instantaneous field of view of $\gtrsim 200$ square degrees.

We briefly overview the two crucial datastreams used for CHAMPSS, namely CHIME/FRB, which is the backbone of our search, and CHIME/Pulsar, which aids in confirming candidates and timing newly discovered pulsars.

²⁸ <https://www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html>

²⁹ <https://www.chime-frb.ca/Galactic>

2.1. CHIME/FRB

The CHIME/FRB system is described in detail in CHIME/FRB Collaboration et al. (2018); here, we summarize the first steps, before CHAMPSS accesses the data.

The CHIME correlator has two stages. First, the F-engine channelizes the incoming data of 1024 dual-polarization receivers into 1024 channels each, resulting in time sampling of $2.56 \mu\text{s}$ of the channelized data. The X-engine is a GPU correlator that forms 1024 stationary beams on the sky (256 north-south \times 4 east-west; C. Ng et al. 2017). Known as “L0,” this stage also performs an additional FFT of each 128 samples both to upchannelize the data and downsample in time. After an additional factor-of-eight averaging in frequency channels, and averaging every three successive time samples, the output is a data stream of 0.98304 ms, with 16,384 channels per formed beam, as 8-bit intensity data.

The rest of the CHIME/FRB pipeline consists of four layers dubbed “L1” through “L4.” The L1 stage of the pipeline performs initial RFI rejection, doing sigma-clipping from a series of polynomial and spline detrendings of the data in time and frequency (for details, see CHIME/FRB Collaboration et al. 2018). It is at this stage after the L1 RFI masking that CHAMPSS taps into the CHIME/FRB data stream, before the dedispersion and subsequent single-burst searches. While CHAMPSS also needs to perform a dedispersion step (Section 3.5), we require additional filtering of nonperiodic (Section 3.4) and periodic (Section 3.6) sources.

2.2. CHIME/Pulsar

Running in parallel, CHIME/Pulsar can simultaneously form up to 10 tracking beams on the sky. The system observes sources with a probabilistic scheduler, where each source has a tunable priority ranking (CHIME/Pulsar Collaboration et al. 2021). CHIME/Pulsar operates on complex baseband data with 1024 channels and $2.56 \mu\text{s}$ resolution, and it can be used to produce either fold-mode or search-mode data. Fold-mode data are coherently dedispersed to a specific DM and folded according to an ephemeris, forming an archive: a data cube with dimensions of subintegration, polarization, frequency, and pulse phase. Search-mode data can be coherently dedispersed to a specific DM, with the data produced being intensity as a function of time and frequency; this can later be search for individual pulses or folded.

The tracking beam, finer time resolution, and coherent dedispersion all make CHIME/Pulsar comparatively more sensitive for targeted observations, while the wide field of view of CHIME/FRB is better for searches. For CHAMPSS, CHIME/Pulsar is used to help confirm and time candidates, as described in Section 4.3.

3. CHAMPSS Pipeline: Acquisition, Reduction, and Search

In this section, we describe the different stages of the pipeline, starting from the CHIME/FRB data stream, resulting in power spectra (daily and stacked) and clustered/filtered candidates. A schematic flowchart is shown in Figure 1.

Our regularly updated codebase is located in our Git repository³⁰ and the subrepositories within.

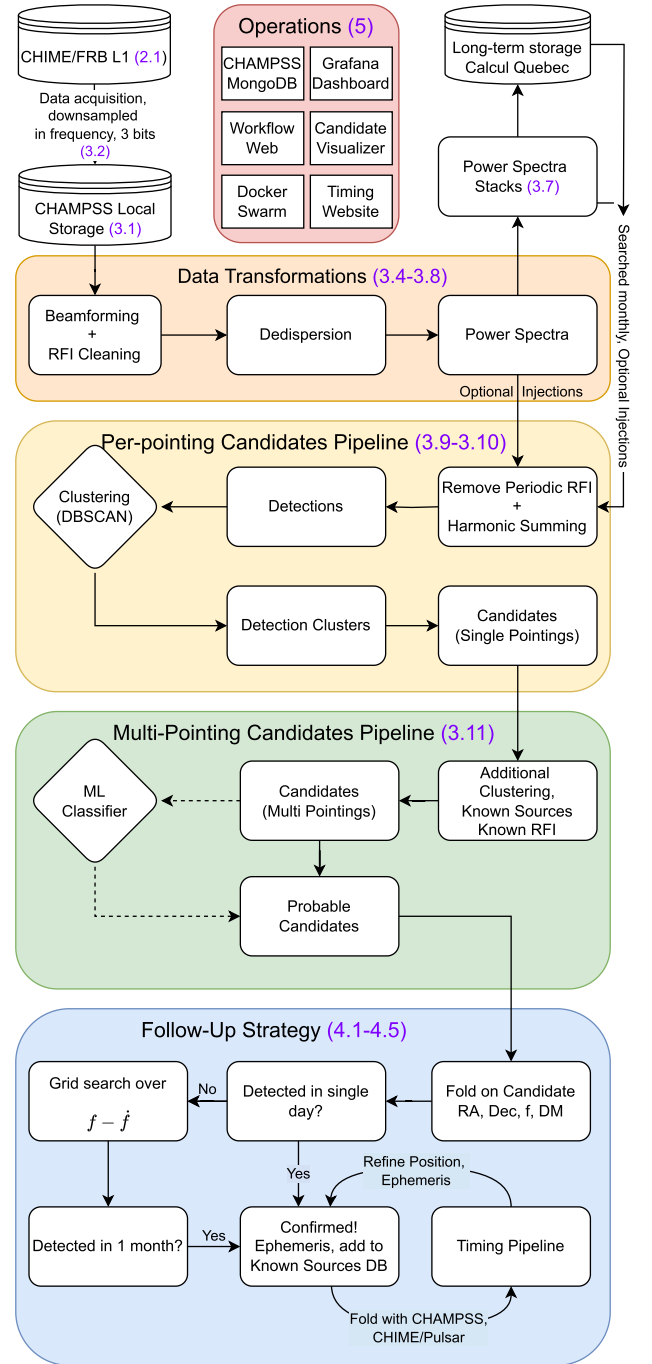


Figure 1. Simplified flowchart of the CHAMPSS pipeline, as described in Sections 3 and 4. The dotted arrow to the Machine Learning candidate classifier indicates that it is a planned component of our pipeline, which was not used in our commissioning survey.

3.1. Computing Setup

Installed in 2022, we have a commissioning cluster for CHAMPSS located at the CHIME site. There are two compute nodes, each with 128 logical cores (64 physical cores, utilizing AMD’s Simultaneous Multithreading) and 256 GB of RAM. These are connected via 10 GbE link to an archiver node with 16×14 TB hard drives, for a total 208 TB storage in a Zettabyte File System Pool. This cluster is connected to the L1 nodes of CHIME/FRB via a 40 GbE link. Our storage and compute capacities will be greatly expanded to increase the sky capacity of the search, and details will be provided in future works.

³⁰ https://github.com/chime-sps/champss_software

The aforementioned cluster is used to process incoming data in quasi-real time (described in the following sections), as the data rate is sufficiently high that it cannot be copied off-site rapidly enough. We do, however, copy long-term data products to a supercomputing cluster “Narval,” owned and operated by Calcul Québec, where we have a 2.2 PB storage allocation. The timing pipeline (Section 4.3) is run there, as will be the search and long-term storage of power-spectrum stacks (Section 3.7).

3.2. CHAMPSS Data Description and Acquisition

As mentioned above, CHAMPSS begins with the CHIME/FRB data stream, which has 16,384 channels and 0.98306 ms time resolution. This corresponds to ~ 1.5 PB/day, for the full data stream from all 1024 beams. However, the data rate for CHAMPSS can be greatly reduced; we can record with fewer channels in most parts of the sky, to be sensitive up to the DM contributions expected for the local Universe (i.e., the Milky Way and nearby galaxies). As described in Section 3.3, we base our pointing map on known Galactic electron models, with conservative errors on the max DM allowing for model uncertainties, halo contribution, and even Local Group galaxies. Additionally, while the data are stored in 8-bit integers, less dynamic range is needed for faint, periodic pulsar signals. The 256×4 grid of formed beams searched by CHIME/FRB extends from a decl. of $\delta = -10^\circ$, up to the North Pole at $\delta = 90^\circ$, and beyond the pole down to $\delta \approx 75^\circ$. We ignore the “lower transit” beams beyond the pole because they are less sensitive, due to the lower effective collecting area toward the horizon, and since they overlap decl. covered by the lower beam rows. This leaves 224 of the 256 beams.

The L1 nodes on CHIME/FRB have a continuous data buffer for each of the 1024 beams. To record the data for a specific beam, we use a remote procedure call (RPC) from our archiver, mapped to the corresponding node. The RPC requests the desired reduction in frequency resolution and bit depth, where the frequency downsampling is governed by the optimal number of channels to reach the maximum DM along the given sightline (described in Section 3.3). In blocks of 1024 time samples (~ 1 s), the intensity data (I) are downsampled on the L1 nodes and saved per channel ν_i and time sample t as $(I(\nu_i, t) - \mu(I(\nu_i)))/\sigma(I(\nu_i))$, along with a header containing the starting time stamp, the RFI mask, and the per-channel mean $\mu(I(\nu_i))$ and standard deviation $\sigma(I(\nu_i))$ used in the normalization above. The downsampled data are saved with 3-bit resolution, following Huffman coding, which is a scheme to compress data with minimal loss of information (D. A. Huffman 1952). The encoding has five levels for data values, and it gives an estimated loss of information of $\sim 5\%$, when the S/N per sample is low. With both the channel reduction and bit reduction, the daily data rate for CHAMPSS is reduced to 52 TB. We note that, at the time of writing, a bug was identified where normalizations can jump at the 1% level between blocks in downsampled data, leading to increased red noise at $f \lesssim 1$ Hz—this reduced the sensitivity of our pilot survey but will be fixed for any further data taken.

3.3. Pointing Map and D . Limits

Unlike CHIME/FRB, which is content with static beams, we wish to grid the sky into pointed observations at a fixed R.A. (α) and decl. (δ). The separation in δ matches the CHIME/FRB beams by design, and we choose the separation in α between

adjacent pointings to be $\Delta\alpha = 0.32/\cos(\delta)$ degrees, where 0.32 is approximately the half width at half maximum of the CHIME beam at 600 MHz, and the $\cos(\delta)$ factor ensures equal angular spacing across the sky. For the full CHIME sky, this amounts to 165537 independent pointings. The pointing durations differ based on the decl., due to the transit nature of CHIME.

The maximum expected DM for a Galactic pulsar is highly position-dependent, with the total model DM ranging from $\lesssim 100$ pc cm $^{-3}$ far off the Galactic plane, and > 4000 pc cm $^{-3}$ in the direction of the Galactic center for a pulsar at the far edge of the MW. For sightlines with low maximum DM, we can reduce the channelization without suffering additional DM smearing; i.e., we wish to choose an optimal channelization to minimize CPU time and storage, without introducing DM smearing larger than our sampling time.

We form a Galactic DM map using the two most used Galactic electron models, NE2001 (J. M. Cordes & T. J. W. Lazio 2002) and YMW16 (J. M. Yao et al. 2017), as follows:

1. For each α , δ , we compute the maximum DM along the line of sight as $\text{DM}(\alpha, \delta, D = 50 \text{ kpc})$, taking the maximum of NE2001 or YMW16 as $\text{DM}_{\text{max}}(\alpha, \delta)$.
2. For Galactic latitudes $|g_b| > 15^\circ$, $\text{DM}_{\text{search}} = 2\text{DM}_{\text{max}} + 50 \text{ pc cm}^{-3}$.
3. For Galactic latitudes $|g_b| < 15^\circ$, $\text{DM}_{\text{search}} = \text{DM}_{\text{max}} \times \exp(0.0313 \times |g_b| + 0.223)$, a heuristic function that rises from 1.25 at $|g_b| = 0^\circ$ to 2.0 at $|g_b| = 15^\circ$.
4. Along sightlines to M31 and M33, we add an additional 400 pc cm^{-3} .

Galactic electron models are less constrained at high Galactic latitudes, since there are fewer pulsars than in the plane. NE2001 predicts higher values of DM_{max} than YMW16 at Galactic latitudes $|g_b| < 2^\circ$, and vice versa, with fractional differences of DM_{max} between the two models in excess of 50% along some sightlines (for a comparison, see D. C. Price 2021). By taking the conservative largest expected DM of either model, we expect few Galactic pulsars to be beyond our DM search limits, with the possible exception of pulsars lying behind H II regions (S. K. Ocker et al. 2024, and see our Section 6.3). Aside from the Milky Way, M31 (Andromeda), and M33 (Triangulum) are the two most massive galaxies in the Local Group, at $D \approx 0.8$ Mpc and $D \approx 0.9$ Mpc, respectively. Their expected DM distributions have been modeled in S. K. Ocker et al. (2022), with the model reaching a maximum of $\text{DM} \approx 400 \text{ pc cm}^{-3}$ adopted above.

We use the “DDplan” utility from PRESTO (S. Ransom 2011) to determine the optimal channelization given $\text{DM}_{\text{search}}(\alpha, \delta)$, and round up to the nearest power of 2. Our pointing map is illustrated in Figure 2, with a zoom-in showing the pointing grid compared to CHIME’s beam in Figure 3. The duration and sensitivity as a function of decl. are shown in Figure 4. The sensitivity plot is illustrative for the cold sky, using the temperature and gain values outlined in Section 6.2, and the frequency-averaged CHIME beam model³¹ for the relative sensitivity across decl.

We also present the expected sensitivity as a function of pulse period at different Galactic latitudes in Figure 5. The sensitivity calculation uses the radiometer equation (R. J. Dewey et al. 1985; D. R. Lorimer & M. Kramer 2004) assuming a duty cycle of 0.024, and a digitization correction

³¹ <https://chime-frb-open-data.github.io/beam-model/>

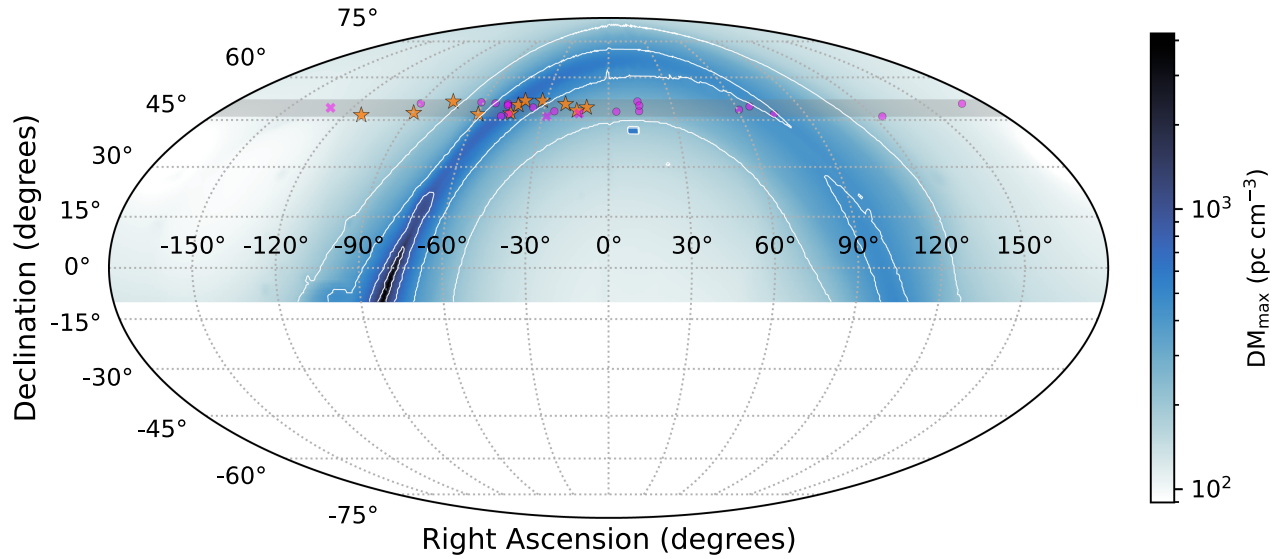


Figure 2. CHAMPSS pointing map. The color bar represents the maximum DM that we search to along a given sightline in units of pc cm^{-3} . The contours indicate increasing number of channels in powers of 2, representing 5 tiers from 1024 to 16,384. The gray shaded region denotes our commissioning survey comprising 10,310 pointings out of 165,537 for the full sky, with 14 beam rows (120–133)/56 beams; see Section 6.1. The orange stars denote the newly discovered pulsars, and the magenta points show known pulsars that our survey detected.

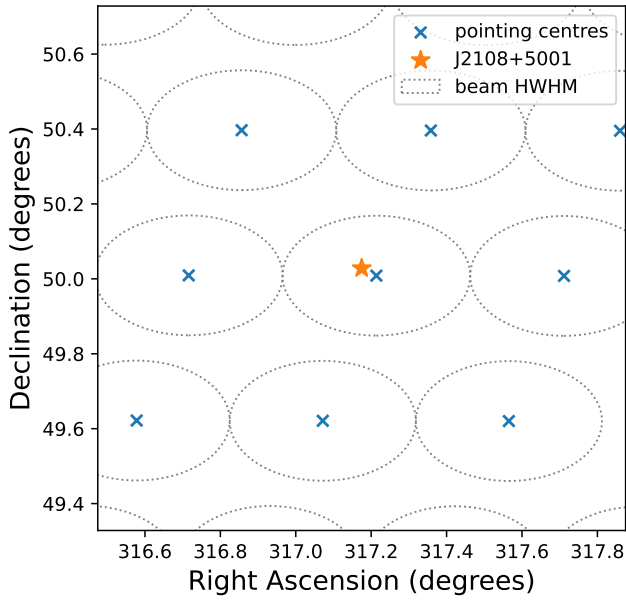


Figure 3. Zoomed-in version of a small region of the CHAMPSS pointing map surrounding the timing position (see Section 4.3) of newly discovered PSR J2108+5001. The contours of the beam half-power at the central frequency 600 MHz are shown surrounding the pointing centers.

factor of $\beta = 1.15$, based on an expected 8.5% loss from 4-bit digitization in the voltages (M. L. A. Kouwenhoven & J. L. L. Voûte 2001) and $\sim 5\%$ observed losses from five-level quantization of CHIME/FRB intensities. A typical integration time of $t_{\text{int}} = 10$ minutes is used, as well as an FFT efficiency factor $\epsilon_{\text{FFT}} \approx (1 + 4.73 \times 10^{-2} \delta^{-0.627})^{-1}$ for duty cycle $\delta \equiv W/P$, following V. Morello et al. (2020).

3.4. Beamforming and Initial RFI Cleaning

The CHAMPSS pipeline works by first providing a date and the R.A. and decl. of one of the pointings in our pointing maps. The pipeline then tests whether data were recorded during the time corresponding to this pointing. If data are present on disk,




Figure 4. Duration (top) and sensitivity (bottom) of pointings as a function of decl. Pointings at high decl. diverge in time as sources spend longer in the beam; pointings longer than 2^{22} samples ≈ 68.7 minutes are split, due to memory limitations.

then a “beamforming” step is performed; each pointing in our pointing map is traversed by four static beams of CHIME/FRB, which can be combined to a quasi-tracking beam at fixed position. A combined time series of the data is formed by appending the data from these four beams, based on when each beam is closest to the chosen pointing. This creates a single $I(\nu, t)$ data product, with corresponding RFI mask provided by the CHIME/FRB system. The appending of data from adjacent beams is done on the per-sample level, to maintain a contiguous timestream. CHIME/FRB is time-aligned between all of its beams to precision much smaller than a 0.938 ms sample, and from our timing analysis (see Section 4.3), we find no evidence of any time jumps between beams.

For further RFI cleaning, a known bad channel list that currently masks 21.6% of channels is incorporated into the RFI mask, and the remaining channels are passed through a filter based on the generalized spectral kurtosis estimator (G. M. Nita & D. E. Gary 2010a, 2010b; G. M. Nita et al. 2016a, 2016b).

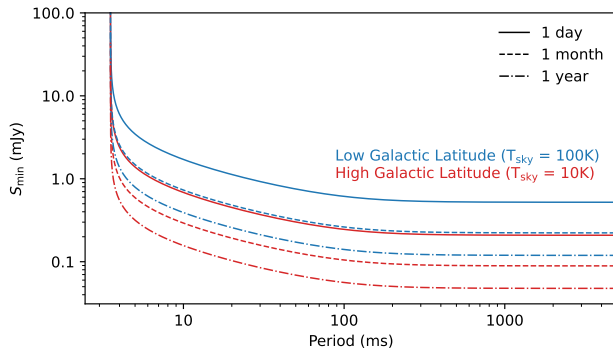


Figure 5. Stacking sensitivity as a function of period at low (blue) and high (red) Galactic latitudes, calculated at the maximum DM along the line of sight (valid for any contour boundary of our pointing map, with $DM = 2^n \times 212.5 \text{ pc cm}^{-3}$ and $n_{\text{chan}} = 2^n \times 1024$, for $0 \leq n \leq 4$). The number of days of stacking is indicated by different line styles (see legend). Sensitivity increases as more days of observation are stacked. We increased the number of frequency channels as the maximum searched DM along the line of sight increased, which kept our sensitivity the same for different points at the maximum DM. Thus, we have the same sensitivity curve for different pointings at the maximum DM. The effects of red noise are not modeled here, and they will decrease sensitivity at large periods.

This filtering step is notionally best suited to identify narrow-band, dynamic, and impulsive RFI. Each frequency channel is treated independently, on a timescale of 1024 time samples (1006.63296 ms). The estimator statistic threshold ($\approx 25\sigma$ equivalent) used to flag additional data is set by the underlying data statistics and the data chunk size, and assumes that the channel data statistics are Gaussian without the presence of RFI.

For each segment corresponding to one input beam, the values marked by the RFI mask are replaced by the median of the full segment. Linear trends from the data are then subtracted on scales corresponding to 32,768 time samples.³² The previously computed median of the segment is then re-added to the full segment. Afterward, the data are rescaled so that the data of each channel have the same median across all beam segments.

As a final step before dedispersion, all frequency channels where more than 75% of the samples are masked are set to 0. All other channels are normalized to have a mean of 0 and a standard deviation of 1. Ultimately, our individual RFI cleaning steps (the mask from the CHIME/FRB system, our known bad channel mask, the filter based on generalized spectral kurtosis estimator, and the masking of largely masked channels), result in a median mask fraction of $\approx 55\%$.

3.5. Dedispersion

We dedisperse each pointing to many DM trials using the Fast Dispersion Measure Transform (FDMT) introduced by B. Zackay & E. O. Ofek (2017), an algorithm that dedisperses in $2N_t N_f + N_t N_d \log_2(N_t)$ time, where N_t , N_f , N_d are the number of time samples, frequency channels, and DM trials, respectively. The maximum DM and the number of frequency channels for a given pointing are given by our CHAMPSS pointing map (see Figure 2 and Section 3.3), and the DM spacing is set by the input time resolution, where the increment in DM corresponds to a delay of one time sample across the band. Due to storage and computational constraints, we calculate only every two DM trials, resulting in spacing of

$\Delta DM \approx 0.1 \text{ pc cm}^{-3}$. We currently use a uniform time sampling and DM spacing for ease of use, but as we run the survey at scale, it may be useful to increase the time sampling and ΔDM at high DMs to reduce computation and storage, e.g., following DDplan in PRESTO. We use an implementation of FDMT in a parallelized CPU code written in Python with C++ bindings.³³ After FDMT, the data product is then $I(DM, t)$ per pointing.

3.6. Power Spectra

For each of the dedispersed time series, the power spectrum is computed after padding them to a length of $2^{20} = 1,048,576$ samples, resulting in a data product $I(DM, f)$.

In order to add power spectra from different days, we need to perform a barycentric correction. To achieve this in the power-spectra domain, we apply the barycentric velocity correction to the sampling time of the observation (computed using *astropy*; Astropy Collaboration et al. 2013, 2018, 2022), calculate the corresponding frequency bins, and then interpolate (using nearest neighbor interpolation) to the frequency bins from the unaltered topocentric sampling time, which will be constant in time for each pointing.

After barycentric correction, we apply a red-noise correction method from PRESTO to each of the power spectra. In this method, a logarithmically increasing window is used at low frequencies to compute the local median. A linear fit to these windowed median values then provides a local red-noise estimate for each frequency bin, which is divided into the data value.

Two schemes for RFI suppression are used in the power-spectrum domain. The first is a list of persistent “birdies” that show up regularly in our observations, and the second is a method to dynamically detect strong, unwanted RFI signals in each observation. RFI signals will usually show up in many pointings across the sky, while pulsars are more localized. This difference is used by the dynamic scheme; we use the power spectrum $I(DM=0, f)$ (which we call the “zero-DM power spectrum” throughout), find strong peaks ($>5\sigma$), and then add their frequency bins to the dynamic mask. In order to decide whether weaker peaks (2σ – 5σ) should be removed, we compare them with the peaks of other nearby observations. If a frequency bin is marked in more than 50% of the compared pointings, then it is added to the birdie list. To enable this comparison, we store the birdie peaks in our observation database (see Section 5.2), which allows other pipeline processes to access them, and the information about which frequency bins are masked is stored in order to use it in the search process. Since neighboring observations are needed for proper RFI removal, this necessitates that the available observations are processed in an order that guarantees at least some nearby observations have been processed already. All frequency bins marked by the dynamic and static birdie filters are set to zero. These filters will mask a few percent of all frequency bins, depending on the RFI situation.

3.7. Power-spectrum Stack

In order to gain sensitivity to faint pulsars, we incoherently sum the power spectra from multiple days of observations of the same pointing. The effectiveness of this technique in

³² Using `scipy.signal.detrend` (P. Virtanen et al. 2020).

³³ <https://github.com/pravirk/dmt>

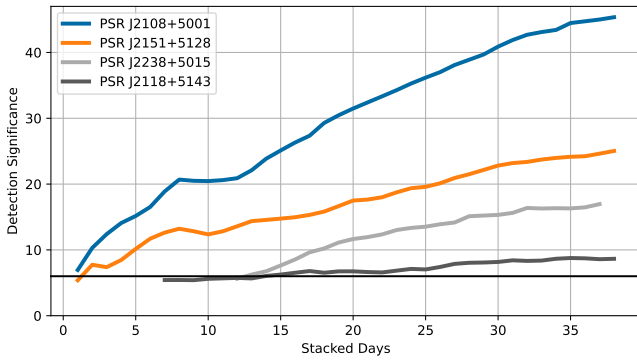


Figure 6. Candidate significance for four of our new pulsars as a function of number of days stacked. The horizontal black line shows the threshold of 6σ , which we currently use during the stack-searching process. In this run, a threshold of 5σ was used and the stack was searched after each day, which is not the case in our normal processing scheme.

allowing us to find pulsars that otherwise would have been missed in single observations is illustrated in Figure 6. Retaining all power spectra from all observations on disk would not be possible for us, due to storage constraints, which necessitates us summing the spectra from different days.

For each new daily observation, the power spectra are created as described above and first individually searched as outlined in the following sections. We compute quality metrics to determine whether a new set of power spectra should be added to the existing stack that is saved on disk. This filters out power spectra that may worsen our ability to find new pulsars in the stack, due to them having unforeseen processing errors or strong RFI that has not been fully removed. For one of our quality metrics, we perform the Kolmogorov–Smirnov test (F. J. Massey 1951) on the zero-DM power spectrum to test whether it follows the χ^2 distribution expected for our power-spectrum stack (M. van der Klis 1989). We also compute the number of outliers in a χ^2 distribution and compare it with the observed number of outliers in the zero-DM spectrum. Another quality metric that we use is the number of “detections” (see Section 3.9) that result from our search process, which will increase in the presence of a pulsar or RFI. We expect strong RFI to increase the number of detections much more than any new pulsars we expect to find.

Once these quality metrics are computed, we compare them with static and dynamic thresholds. If a quality metric lies outside the maximum of these two thresholds, the power spectra are not added to the power-spectrum stack. The static thresholds are manually set to adequate values to filter out bad power spectra. The dynamic thresholds are computed for each individual pointing; the median and median absolute deviation (MAD) of the previous values for the quality metric are computed, and the dynamic threshold is set to the median + $3 \times \text{MAD}$. The dynamic threshold allows us to use these quality metrics even if they are consistently higher than expected due to the presence of strong pulsars.

In this study, we only employed one power-spectrum stack for each pointing. In future CHAMPSS work where we will create much deeper stacks, we intend to save two different stacks on disk: a monthly stack that contains only relatively recent data and a cumulative stack that contains all previous data. The monthly stack allows a second pass of quality control to prevent the cumulative stack from containing too much RFI, and it may increase our sensitivity to intermittent pulsars. One set of power spectra for the full sky is roughly 400 TB.

3.8. Injections

In order to constrain the reliability of our pipeline, we have the capability to inject fake pulsar signals directly into the power spectrum. For this purpose, we use analytic templates derived from real pulse profiles from the MeerKAT Thousand Pulsar Array (B. Posselt et al. 2023), but we vary the injected significance, frequency, and DM. We will inject across all pointings to constrain the completeness of our pipeline, and are tracking red noise across all pointings, to map it as a function of time and beam.

The injections only interact with the local version of the stack used for a given run of the pipeline, and so do not modify the database. The injection design and subsequent results on the transmissivity of the pipeline will be detailed in a future paper.

3.9. Detections

We follow PRESTO’s methodology for searching a power spectrum, summing harmonics within a power spectrum and evaluating the significance of any particular power. For a time series of purely Gaussian noise, the powers will follow a χ^2 distribution with 2 degrees of freedom, as already mentioned in Section 3.7. If m such powers are summed together, whether in the harmonic summing procedure or power-spectrum stacking, the result will follow a χ^2 distribution with $2m$ degrees of freedom, χ^2_{2m} (M. van der Klis 1989; S. M. Ransom et al. 2002). Therefore, to determine the significance of a power, we find the probability of this power occurring by chance in a χ^2_{2m} distribution, via the cumulative distribution, and convert this to a Gaussian-equivalent sigma.

By stacking power spectra, we reach larger values of m than are typically encountered in a pulsar search. The cumulative distribution of χ^2_{2m} is given by

$$F(P; \chi^2_{2m}) = \frac{\gamma\left(m, \frac{P}{2}\right)}{\Gamma(m)} \quad (1)$$

for a power P , with γ being the lower incomplete gamma function and Γ the gamma function. In earlier stages of development, high values of m led to overflow errors in the calculation of γ . To resolve this, we use an algorithm specifically designed to avoid such errors (R. Abergel & L. Moisan 2020).

This search is performed for the power spectra at each trial DM, and at different harmonic sums (1, 2, 4, 8, 16, and 32). As an example, for a harmonic sum with four harmonics at frequency f , the bins in the power spectrum corresponding to f , $2f$, $3f$, and $4f$ are summed together. Our significance threshold for the search is currently the Gaussian-equivalent 5σ and 6σ for searched daily observations and stacks, respectively. Any points above this significance threshold are termed “detections.” Detections in a single harmonic sum search of a single DM trial that are closer than $1.1 \times$ the frequency resolution of the power spectrum are grouped together, and only the most significant detection is kept. Each detection contains information about its DM, frequency, Gaussian-equivalent σ , and the indices, frequencies, and powers of the individual bins in the power spectrum summed to produce the detection.

During the search step, we can also filter out known pulsars. This is performed by checking if a known pulsar has previously been identified in this pointing, using the known

source sifter (see Section 3.11), and comparing the previously observed sigma with a given threshold. When looking for previously detected known pulsars, one can choose to look only at candidates resulting from single-day spectra or from stacked spectra. If the threshold is surpassed, the pulsar is removed by masking all frequency bins where the pulsar and its harmonics are expected. This prevents strong pulsars from flooding the pipeline with detections, which would drastically worsen clustering performance (as described in the following Section). As the strength of the observed pulsars continuously increases as more days are stacked, this step is crucial when stacking more deeply.

3.10. Clustering

In a classic pulsar search, and earlier in the development of the pipeline, detections would be clustered in DM-frequency space to group detections from the same source together. The groups would then be assessed to determine if they were harmonically related, based on the most-significant frequency within the group, and filtered by this process. Detections at many harmonics of the fundamental frequency are expected for both RFI and pulsars, so this is a key stage in thinning the number of clusters to be made into candidates.

RFI peaks are often broad in frequency, meaning the spread of frequencies in a group of detections resulting from RFI may be quite large. When assessing for harmonic relation, only using the most-significant frequency within a cluster led to many RFI clusters not meeting that criterion. As the vast majority of detections are due to RFI, this was not ideal, and we implemented a system that clusters the detections in DM, frequency, and harmonic relation simultaneously. This is achieved via scikit-learn’s DBSCAN³⁴ algorithm using a custom metric that allows us to not only to identify clusters that are close in DM and frequency but also to identify harmonically related signals. There is no simple transformation to convert harmonic relation into a nearest neighbors problem, and calculating this for each pair of detections would be prohibitively computationally expensive. Therefore, steps are taken to reduce the number of harmonic-relation calculations required.

The first step is to reduce the number of detections. If multiple detections occur at the same frequency and DM for different harmonic sums, only the most significant is kept. We also emulate PRESTO and eliminate any detections that are dominated by individual strong harmonics above the fundamental frequency of the detection, e.g., if a detection at frequency f is dominated by a power at $7f$. The second step is to identify detections that share the same frequency, in order to avoid duplicate calculations. The last step uses the information about which bins were summed in the power spectrum. Most harmonically related detections will share some power-spectrum bins in their sum; for example, an eight-harmonic sum to produce a detection at $\frac{13}{3}f$ will sum bin power-spectrum bins

$$\frac{13f}{3}, \frac{26f}{3}, 13f, \frac{52f}{3}, \frac{65f}{3}, 26f, \frac{91f}{3}, \text{ and } \frac{104f}{3},$$

and clearly the $13f$ and $26f$ bins will be shared by an f detection that summed 32 harmonics. An example of this is




Figure 7. Example showing bins in a power spectrum shared by a 32 harmonic sum for a detection at f (highlighted in red) and an 8 harmonic sum for a detection at $13f/3$ (blue).

also shown in Figure 7. This property is utilized to form groups of detections that share power-spectrum bins. Then, only detections within the same group are evaluated for harmonic relations.

Before the actual harmonic relations are calculated, an additional step of RFI filtering can be performed. This filtering step also uses DBSCAN but only utilizes the frequency and DM as input, using the detections that were not filtered out in the previous step. The DBSCAN run results in a list of detections that are dense in frequency and DM. Clusters that have their strongest σ below a DM of 2 pc cm^{-3} are filtered. Additionally, all clusters that have a mean frequency within the frequency span of those low-DM clusters can also be filtered out. This allows filtering of RFI signals that appear along a wide range of DM values. This step was added in 2024 October and was not present when the pulsars in this paper were discovered.

During the processing of the results shown in this paper, we used two different approaches to create the final metric used in the DBSCAN clustering. In the first approach, the metric is computed based on the Euclidean distance in frequency and DM, then lowered if two detections are harmonically related. The originally calculated Euclidean distances are multiplied by a factor k , which is given by

$$k = 1 - \frac{|B_1 \cap B_2|}{\min(|B_1|, |B_2|)}, \quad (2)$$

where B_1 and B_2 are the frequency bins in the compared detections, $B_1 \cap B_2$ denotes the frequency bins that are in common between those detections, and the $||$ operator denotes counting the number of frequency bins. This was the clustering scheme used when our first pulsars were found.

In 2024 September, we changed to a new scheme for the DBSCAN metric, which will be used in future searches. In the new scheme, we not only compare the overlapping bins between detections but also use power in those frequency bins to reduce accidental clustering. In order to compare two detections that share frequency bins, we use the metric

$$M = 1 - \max\left(\frac{P_{1 \cap 2}}{P_1}, \frac{P_{2 \cap 1}}{P_2}\right) + \Delta \text{DM}. \quad (3)$$

$P_{1 \cap 2}$ is the sum of the powers in the first detection that share frequency bins with the second detection. This is then divided by P_1 , the total power in the first detection, to give the fraction of detection 1’s power that lies in the shared bins, $\frac{P_{1 \cap 2}}{P_1}$.

³⁴ <https://scikit-learn.org/stable/modules/generated/sklearn.cluster.DBSCAN.html>

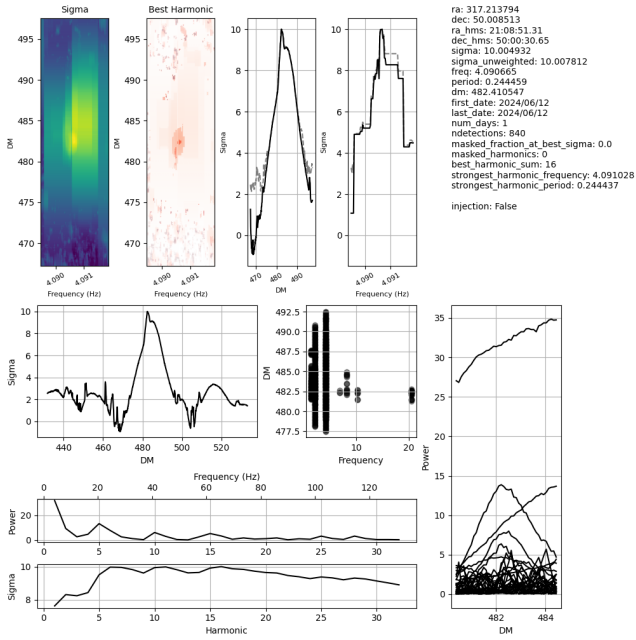


Figure 8. Candidate plot for a detection of our new pulsar PSR J2108+5001 in a single observation. The various diagnostic plots show how the signal develops as a function of frequency and DM. It also shows the signal strength at various harmonic frequencies. The clustered detections that make up this candidate also contain detections at subharmonics that result from the harmonic summing. A full description of all fields included is provided in Appendix A.

Similarly, $\frac{P_{2n1}}{P_2}$ is the fraction of detection 2’s power that lies in the shared bins. ΔDM is the DM difference between the detections. This metric allows for better clustering because, if a signal is detected at a harmonic frequency, most of the power will still be contained in the bins shared with a detection at the fundamental frequency. In that case, the metric becomes very small and the detections are clustered. During this change, we also started using `scipy`’s sparse arrays and only computed the metrics for detections that are close in DM. This allows us to save on memory, which would otherwise limit the number of detections that can be clustered.

At the end of the clustering process, the identified clusters are then passed on to the candidate creation process.

3.11. Candidate Creation

After the detections have been clustered, additional diagnostics are computed in order to save a rich representation of the detected signal on disk and create a candidate plot (example shown in Figure 8). The diagnostics show how signals compare at nearby frequencies, nearby DMs, and different harmonics, which helps us distinguish between real pulsars and RFI candidates.

The detected signals of pulsars and RFI show a fundamentally different signature across the recorded pointings, which is already used in the dynamic birdie filter outlined in Section 3.6. While pulsars show up in either only one pointing or a cluster of adjacent pointings with a clear central point, RFI candidates will be spread over a large area of the observed sky with a mostly flat signal strength. We use this difference to perform multi-pointing clustering after the pointings for one

day have been processed or one run of the monthly stack search has finished. This process uses DBSCAN to cluster candidates from multiple pipeline runs by comparing the frequency, DM, R.A., and decl. This creates “multi-pointing candidates,” which significantly reduces the number of candidates and allows us to further filter them based on their spatial distribution. Examples of multi-pointing candidate plots are shown in Figure 9. The multi-pointing candidates are saved on disk in individual files and a csv file containing a summary of all multi-pointing candidates per day is created. These new candidates files also contain all candidates that were not grouped with other candidates during the clustering process.

The created multi-pointing candidates are matched to known sources by computing the likelihood ratio (i.e., the Bayes factor) of association with neighboring known sources, based on the values and uncertainties of the candidate and catalog sources. This is the same matching scheme used in CHIME/FRB (see CHIME/FRB Collaboration et al. 2018), except with an additional parameter of the spin frequency. After known source identification, the multi-pointing candidates are saved in individual files.

We perform the multi-pointing process for the candidates resulting from the daily search and monthly stack search separately. The multi-pointing process is essential in reducing our number of candidates. In a run of our multi-pointing pipeline in 2024 December on the candidates of a search in 10,288 power-spectra stacks, which is our total number of stacks with 20 stacked days or more, we reduced the total number of candidates from 847,516 to 126,877.

4. Candidate Confirmation and Follow-up

In this section, we describe our process for candidate confirmation and follow-up of confirmed candidates.

4.1. Sifting for Promising Candidates

The best way to confirm a candidate is to see a convincing folded pulse profile. However, due to the nature of our survey, the computational cost is too high to fold every candidate; by the time single-day candidates are formed, the beamformed data have already been removed from memory, so folding requires an additional step of data reading and beamforming per pointing. To fold every candidate would incur a roughly $1.5\times$ increase in compute load.

Instead, we wish to only fold on candidates that are likely to be pulsars. Beginning with a day’s worth of multi-pointing candidates, we perform a set of heuristic cuts. First, we cut on the significance, DM, and frequency, restricting the candidates to $\text{S/N} > 7$, $\text{DM} > 2 \text{ pc cm}^{-3}$, $0.01 \text{ Hz} < f < 100 \text{ Hz}$, and position spread $\sigma_{\text{position}} < 5^\circ$ in the multi-pointing candidate. Candidates matched to known sources are excluded. Bright pulsars result in many candidates tightly spaced in DM and spread among many harmonics in f . Not all of these are matched with the known source sifter, so we add an additional cut: all candidates within 1° , $\sigma_{\text{DM}}/\text{DM} < 0.1$ of a known source are excluded. We note that, at the time of detection, most faint pulsars are seen in 1–2 pointings, although they appear in more as more data are stacked or when they are near beam edges. Sources can appear quite spread out in R.A., due to the significant power in CHIME’s sidelobes, and the $\sigma_{\text{position}} < 5^\circ$ cut is deliberately very conservative in order to not remove a bright transient pulsar, while still removing a great deal of RFI sources.

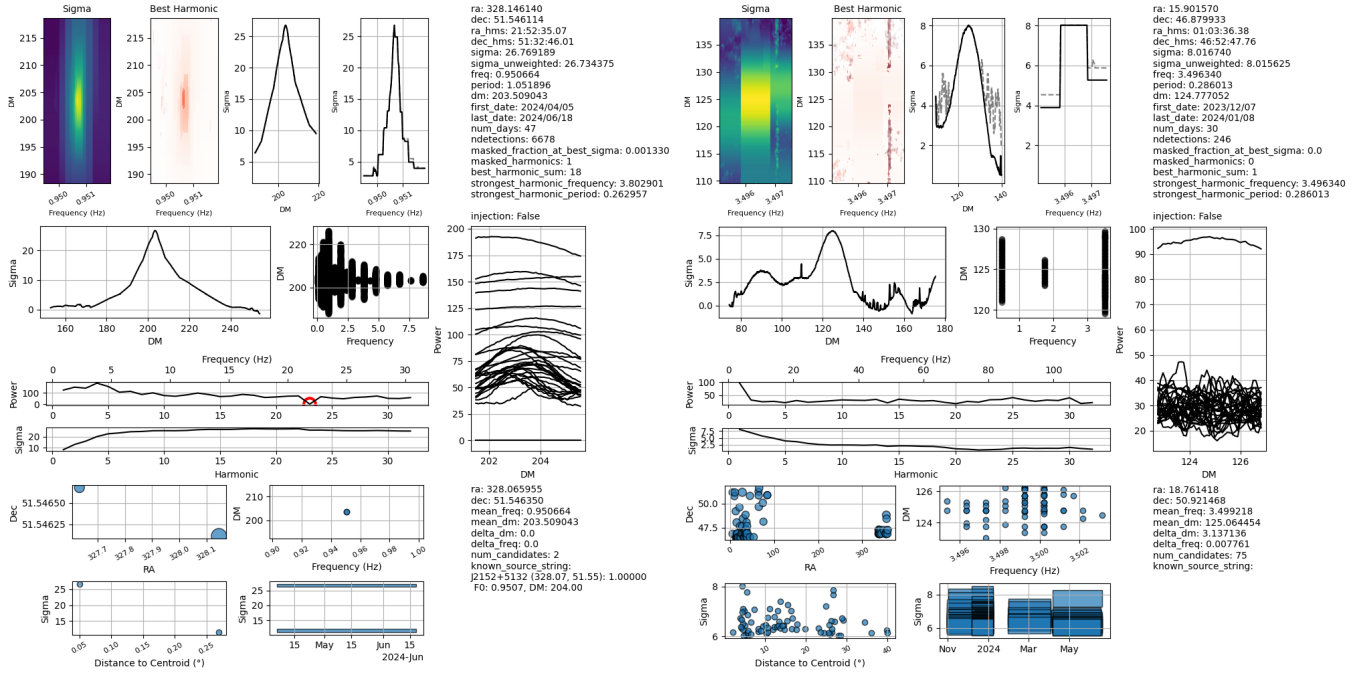


Figure 9. Examples of a pulsar (left) and an RFI candidate (right). The plots are the result of the multi-pointing clustering process, in candidates from the power-spectrum stacks. The upper parts of the plots show the data from the strongest clustered detections, and the lower parts of the plots show information derived from the multi-pointing clustering process. A full description of all fields included in these candidate plots is provided in Appendix A. While RFI signals come in many shapes and forms, often they can be discerned from pulsar candidates using the diagnostic plots. This RFI candidate, for example, shows a spatially broad distribution without a clear central point, and the power is concentrated in one harmonic. Often, RFI candidates are also very broad or show multiple clear peaks in the DM plane. The pulsar was detected only in two nearby pointings.

Additionally, unflagged periodic RFI sources often show up at a tight window in f (i.e., compared to the full search space of spin frequency, as they can be broadened due to frequency modulation and barycentring in the stacks), spread across a large DM range. We perform an additional cut for these candidates. We create a histogram $N(f)$ of candidates in logarithmic bins of 1% in f ; for bins with ≥ 2 candidates with $\sigma_{\text{DM}}/\text{DM} > 0.1$, the candidates are excluded. Finally, RFI coming from specific sources (e.g., electronics, planes, and satellites) can cluster in time—and thus R.A., which is a proxy for time. After all other filters, we take only the highest- σ candidate per pointing. After the above steps, the characteristic number of daily candidates for four beam rows is reduced from $\sim 100,000$ to ~ 200 . The above sifting steps are run automatically each day to decide which single-day candidates to fold, but it can also be run on the power-spectrum stacks. We confirmed that our newly discovered pulsar candidates are not excluded by the cuts.

The above steps are currently necessary to reduce the inflow of candidates to a reasonable amount. We reiterate that the above steps are heuristic—and in some cases, overly conservative, such that real pulsars would be lost. To better tackle this problem as we scale up the survey, we are developing a machine learning classifier, trained on many real candidates, RFI, and injected signals, to be used in conjunction with (or to replace) the above sifting steps. This will be discussed in future work.

4.2. Folding and Multiday $f - \dot{f}$ Search

When a promising candidate is identified, the data corresponding to the R.A. and decl. of the pointing are folded at the candidate f and DM. For candidates from individual

days, the folded candidates are inspected visually, in a manner similar to how previous pulsar surveys have operated (see Figure 10).

For candidates only detected in the power-spectrum stack, the pulsar is likely too faint to detect in a single day; all of the existing data on disk, and each subsequent day, are folded with the same candidate ephemeris, and summed in time to form $I(T_i, \phi)$, where ϕ is the pulse phase and T_i is the central time of the i th observation. The candidate f derived from minute-scale transits is too imprecise to phase align pulses from day to day. Moreover, a pulsar’s spin-down and uncertain position result in a time-variable spin frequency. Over 1 month, for an isolated pulsar, these effects can all be approximated sufficiently well with a spin-frequency derivative \dot{f} , which is allowed to be positive or negative, and the phase is approximated as

$$\phi(t) = f t_1 \approx f_0 t_1 + \frac{1}{2} \dot{f} t_1^2, \quad (4)$$

where $t_1 \equiv t - t_{\text{ref}}$ and f_0 is the spin frequency at the reference time t_{ref} , set to the central observation. Binary pulsars add an additional complication; see Section 4.5.

When the requisite number of days of data have been folded (a tunable parameter, set to 10 days thus far), a grid of f, \dot{f} values are searched, resulting in $\chi^2(f, \dot{f})$. The spacings in both f and \dot{f} are chosen to be uniform in phase, with $\Delta f, \Delta \dot{f}$, corresponding to a one-bin phase shift between the first and last observations. We set the maximum search frequency of $\Delta f \approx 11.6 \mu\text{Hz}$, corresponding to $1/T_{\text{sid}}$. This is guaranteed to contain a value of f that can align the pulses in phase; due to the transitory nature of CHIME observing at the same sidereal time, there will be a family of aliased solutions separated by an

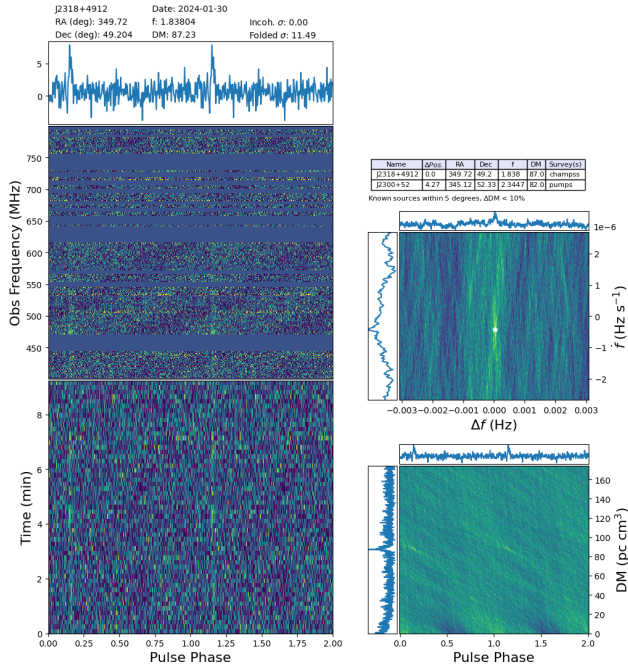


Figure 10. Example of a candidate plot for a single-day fold, on newly discovered pulsar PSR J2319+4919. The side panels (right) are akin to PRESTO plots, showing the χ^2 as a function of f and \dot{f} (top), as well as I (ϕ , DM) (bottom). The top contains information on the candidate leading to the fold, as well as nearby pulsars that might be associated and have passed through the source matching described in Section 3.11.

integer number of N pulses per sidereal day. We search to a maximum frequency derivative of $\dot{f} = 10^{-12} \text{ Hz s}^{-1}$, which is sufficiently large to contain most pulsars, and it captures any intrabeam positional uncertainty. If no signal is seen in this grid, then \dot{f}_{max} is increased by a factor of 10 and the search is rerun, up to $\dot{f}_{\text{max}} = 10^{-10} \text{ Hz s}^{-1}$. An example of the multiday search is shown in Figure 11.

The entire confirmation process described above uses the CHAMPSS data themselves (i.e., from the CHIME/FRB data stream). More rapidly rotating sources will benefit from CHIME/Pulsar follow-up (see discussion in Section 6.2).

4.3. Timing

When a new pulsar candidate is confirmed, it is observed by CHIME/Pulsar, using the ephemeris derived from the search pipeline. Since CHIME/Pulsar (mainly due to its tracking beam) provides a higher sensitivity than CHAMPSS, most timing TOAs are obtained from the CHIME/Pulsar fold-mode backend. However, as CHIME/Pulsar uses a probabilistic scheduling system (CHIME/Pulsar Collaboration et al. 2021), daily observations are not guaranteed. Thus, CHAMPSS fold-mode data are used to start the timing model (i.e., before the pulsar is queued into CHIME/Pulsar schedule) and to fill gaps between CHIME/Pulsar TOAs.

We developed a pulsar timing pipeline³⁵ to time pulsars on a daily basis. To start the pipeline, we create an initial standard template for each pulsar using `paas`; this template is later replaced by an averaged profile or a `fitburst`³⁶ (E. Fonseca et al. 2024) modeled template. In the pipeline, data are first




Figure 11. An example of our multiday search script for candidate confirmation, which searches a grid of f, \dot{f} values as described in Section 4.2. All of our newly discovered pulsars were found with high significance from this algorithm. Shown above is known pulsar PSR J2208+4610 (F. A. Dong et al. 2023), which is on a 412.5 day orbit, demonstrating how the search can phase-align systems in long orbits. Typical isolated pulsars are found with smaller \dot{f} , owing to spin-down and position uncertainties.

masked for bad channels and cleaned using the `clfd` RFI removal algorithms³⁷ (V. Morello et al. 2019). TOAs are obtained using `pat` after summing the profile in frequency, time, and polarization using `pam` (`paas`, `pat`, and `pam` are all part of `PSRCHIVE`;³⁸ A. W. Hotan et al. 2004). Finally, timing solutions are fitted using the `PINT`³⁹ (J. Luo et al. 2021) least-squares fitter, with `TT(BIPM2021)` (G. Petit 2010) as the reference clock standard and `DE421` (W. M. Folkner et al. 2009) as the solar system ephemeris (automatically used by `PINT`). For this paper, the timing solutions were refitted using the `PINT` MCMC fitter to account for parameter degeneracies and these are presented in Table 1. Parameters of the pulsars at the time of discovery are presented in Table 2.

This timing procedure also automatically adds parameters to the timing model by running F-tests, in a method similar to the “Algorithmic Pulsar Timing” scheme proposed by C. Phillips & S. Ransom (2022). Once the degeneracy in a timing model

³⁵ https://github.com/chime-sps/champss_timing

³⁶ <https://github.com/CHIMEFRB/fitburst>

³⁷ <https://github.com/v-morello/clfd>

³⁸ <https://psrchive.sourceforge.net/>

³⁹ <https://github.com/nanograv/PINT>

Table 1
Parameters of Newly Discovered Pulsars Are Shown, Including the Follow-up Timing Solutions (Left) and Derived Parameters (Right)

PSR	Timing Solution					Derived Parameters	
	R.A. (hh:mm:ss)	Decl. (hh:mm:ss)	P (s)	\dot{P} (s s ⁻¹)	PEPOCH (MJD)	DM (pc cm ⁻³)	S_{600} (mJy)
J1629+4636	16:29:52.904(3)	+46:36:51.75(2)	0.314056114549(5)	$8.51(2) \times 10^{-17}$	60301	34.8	0.51 ± 0.08
J1802+47 ^{b,c}	18:02:14.4(5)	+47:16:00(\pm_{50}^{+10})	0.346624379(4)	$1.1(3) \times 10^{-15}$	60546	30.1	0.31 ± 0.05
J1900+5106	19:00:0.92(7)	+51:06:14(1)	0.3377591590(8)	$3.2(3) \times 10^{-16}$	60404	71.6	0.46 ± 0.08
J2002+4652 ^c	20:02:06.87(9)	+46:52:43(4)	0.248260284(1)	$2.6(7) \times 10^{-16}$	60546	141.2	0.76 ± 0.12
J2100+4711 ^a	21:00:13.393(\pm_8^{+7})	+47:11:15.4(1)	1.45874256105(7)	$4.139(3) \times 10^{-15}$	60275	231.1	0.31 ± 0.05
J2108+5001	21:08:41.962(3)	+50:01:41.73(5)	0.24446137520(2)	$9.1442(4) \times 10^{-15}$	59752	482.3	0.89 ± 0.14
J2118+5143	21:18:30.26(7)	+51:43:15.4(6)	0.3702230490(5)	$3.470(2) \times 10^{-14}$	60404	146.3	0.61 ± 0.10
J2151+5128	21:51:46.60(5)	+51:28:48.9(2)	1.0519028955(6)	$9(3) \times 10^{-17}$	60404	203.5	0.73 ± 0.11
J2238+5015	22:38:16.49(6)	+50:15:52.4(3)	0.5600971676(5)	$4(2) \times 10^{-17}$	60404	28.3	0.59 ± 0.12
J2302+4807 ^{b,c}	23:02:10.9(4)	+48:07:26(4)	0.741973791(8)	$1.3(3) \times 10^{-15}$	60376	72.7	0.24 ± 0.04
J2319+4919	23:19:14.842(3)	+49:19:8.72(8)	0.54406513568(2)	$1.26(1) \times 10^{-16}$	60369	87.2	0.81 ± 0.16

Notes. The timing solution was fitted by PINT using its MCMC fitter, and the upper and lower bounds are given by 16th and 84th percentiles of the posterior distribution (a single value is shown when the distribution is symmetric), respectively. Reference epochs for spin-down measurements are shown in the PEPOCH column. Reduced χ^2 values for these timing solutions are shown in Figure 22. More technical information about timing can be found in Section 4.3. Note that EFAC was not applied to rescale the uncertainties during the fit.

^a Codiscovery of J2100+4712g, as discovered by FAST-GPPS (J. L. Han et al. 2025).

^b Likely matches to J1803+47, J2302+48, as discovered in power spectra by Puschino (S. A. Tyul'bashev et al. 2024)

^c Likely matches to J1802+4721, J2002+4653, and J2301+4809, marked in the CRAFTS website.

is sufficiently constrained (typically after ~ 1 yr of timing, as inferred from the MCMC posterior distributions), the initial ephemeris used in the CHIME/Pulsar fold-mode backend is replaced by the newly fitted ephemeris. The updated ephemeris provides an improved timing position and an unaliased period solution (which will be further introduced in Section 4.4), resulting in a higher signal-to-noise ratio for fold-mode observations. An example diagnostic plot from our automatic timing pipeline is shown in Figure 12.

4.4. Resolving Aliased Solutions

As mentioned in Section 4.2, due to the transitory nature of CHIME, the multiday phase connecting algorithm allows for a family of solutions with pulse periods $\pm N$ pulses/day, for any integer N . A period off by N rotations/day will drift across a transit by $\Delta\phi = NT_{\text{transit}}/T_{\text{sid}}$. For example, a solution off by 1 pulse/day, for a 10 minute transit, has a phase drift of $\Delta\phi \approx 0.007$ across each transit. This is a small effect, but measurable given high S/N; see, e.g., the phase errors on individual TOAs in Figure 12.

For pulsars that are bright enough to be detected in a single day, multiple TOAs are made for each observation. This allows for the frequency to be fitted based on the drift across a transit. If necessary, JUMPS are added between each observation to fit for the drift across all transits alone, without complications due to other factors such as position. This frequency can then be used to derive a new timing solution. These TOAs can also be used as a diagnostic to check whether an existing timing solution contains an aliased frequency. To do this, the frequency is adjusted by a multiple of the sidereal frequency $f_{\text{sid}} = 1/T_{\text{sid}}$ (where T_{sid} is a sidereal day), and the timing parameters are refitted. By plotting the residuals as a function of local sidereal time (LST), a clear trend can be seen for the aliased solutions, especially after averaging residuals at similar LSTs. An example of this is shown in Figure 13. For pulsars that are not bright enough to be seen in individual days, the above process can still be done. Each profile traverses the

Table 2
Parameters of New Pulsars at the Time of Discovery

PSR	Pointing		Stacked Days	Detection Significance
	(hh:mm:ss)	(hh:mm:ss)		
J1629+4636	16:29:53	46:36:53	22	19.21
J1802+47	18:01:27	47:22:50	25	6.85
J1900+5106	19:00:19	51:13:40	37	22.16
J2002+4652	20:02:00	46:56:20	25	37.40
J2100+4711	21:00:13	47:11:10	30	15.00
J2108+5001	21:08:51	50:00:30	10	15.89
J2118+5143	21:18:54	51:33:24	44	11.76
J2151+5128	21:52:15	51:32:46	47	26.80
J2238+5015	22:37:58	50:22:20	46	19.08
J2302+4807	23:02:36	48:02:33	40	15.24
J2319+4919	23:19:10	49:24:05	36	11.03

Note. Some pulsars are also be visible in daily candidates, but we identified them first in the stack candidates. For PSR J2108+5001, the candidate was not at the fundamental frequency but instead at one-seventh of the fundamental frequency. All others were found at frequencies close to the derived timing frequency shown in Table 1.

same sidereal time; after first timing with a single-time, frequency-averaged TOA per day, profiles can be stacked to form $I(t_{\text{LST}})$, and TOAs can be extracted from this stacked profile to form TOAs as a function of LST.

4.5. Binaries

We may often find that a binary pulsar is bright enough to be seen in the power-spectrum stack but not bright enough to show up as a candidate on individual days. In these cases, the candidate f value will be close to the average over the (unknown) orbit, but will be off at any given time, which will lead to phase wrapping between observations. Moreover, tight high-velocity orbits will lead to phase evolution during single days.

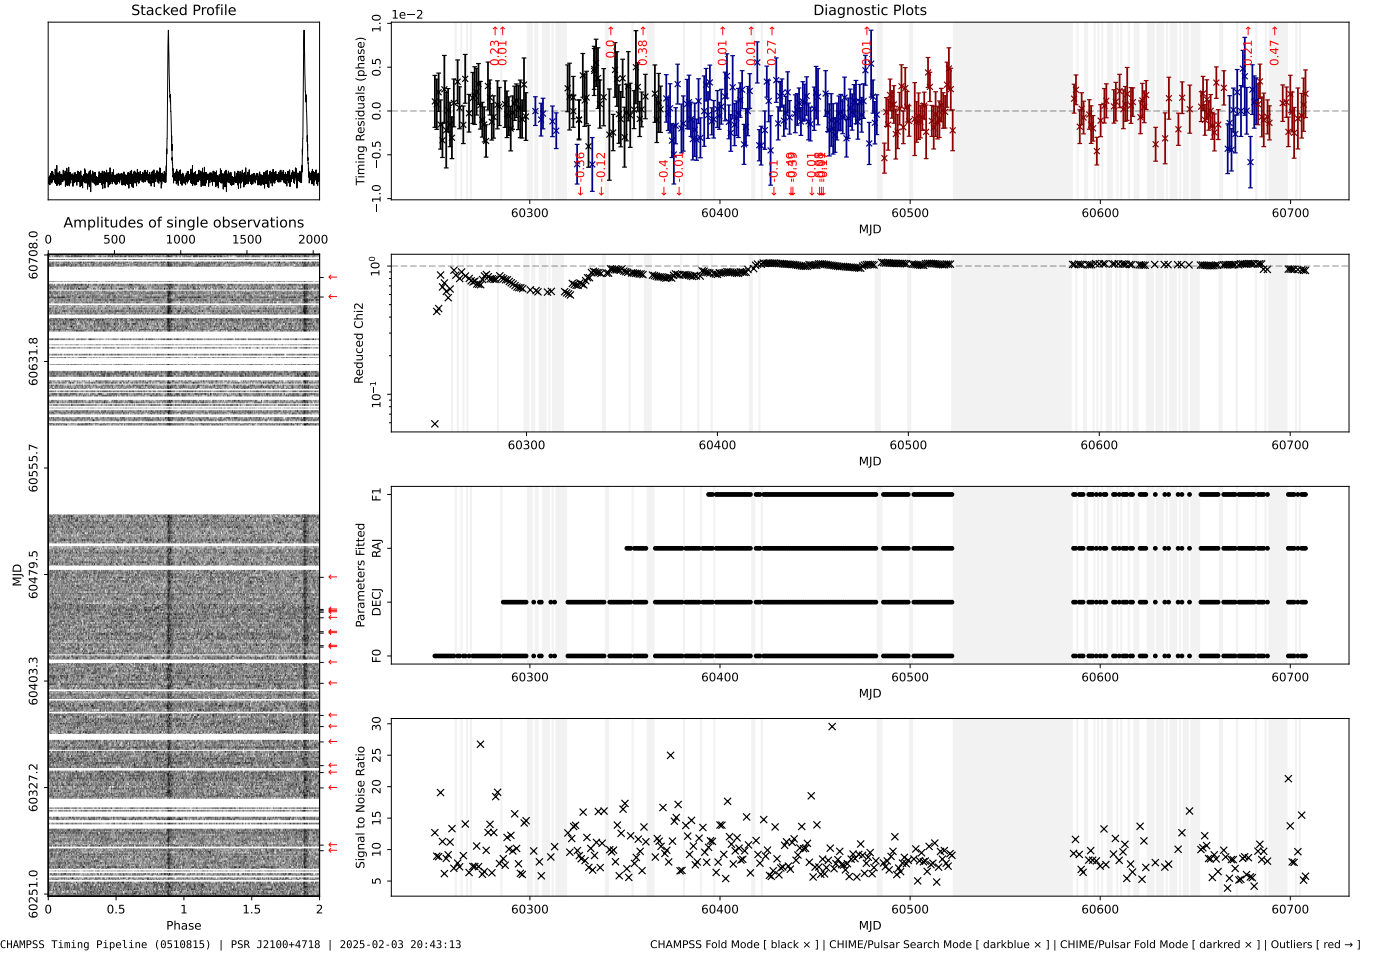


Figure 12. An example output from the timing pipeline described in Section 4.3 with PSR J2100+4711. Left: Pulse intensity (grayscale) as a function of phase and date of pulse profiles after the timing model has been applied (bottom), and their stacked profile obtained by summing over time (top). The noise level depends on whether CHAMPSS or CHIME/Pulsar data were used. The panels on the right (from top to bottom) show the timing residuals, reduced χ^2 , parameters being fitted for, and the S/N per day. For the residuals, black points correspond to CHAMPSS, blue points to CHIME/Pulsar filterbank (search-mode) data, and red points to CHIME/Pulsar fold-mode data. Red arrows indicate outlier TOAs, whose residuals are given by the numbers next to them. Omitted for space, the viewer shows a panel with the full ephemeris and uncertainties on free parameters.

We describe several ways in which we enable sensitivity to binaries.

$P_b \gtrsim 100$ days. This is the easiest case for our survey, where the orbital phase is well captured by a frequency derivative over the course of the 10 day phase-coherent search. The \dot{f} term can then capture the binary motion, allowing us to phase-connect the pulsar. For an example, see left panel of Figure 11. Using the Keplerian parameters of all known pulsars in the ATNF Catalog (R. N. Manchester et al. 2001) with binary periods greater than 100 days and rotation frequency less than 100 Hz, the maximum and mean induced acceleration \dot{f} throughout the orbit are less than $10^{-8} \text{ Hz s}^{-1}$ and $10^{-9} \text{ Hz s}^{-1}$, respectively.

$P_b \lesssim 100$ days. This case is more difficult to detect in our survey. On individual days, we search a range of $\Delta f, \dot{f}$ values, setting the maxima to

$$\Delta f_{\max} = f v_{\max} / c \quad (5)$$

$$\dot{f}_{\max} = \frac{2\pi}{P_{b,\min}} \Delta f_{\max}. \quad (6)$$

Since we fold comparatively few candidates, we can search a large range in $\Delta f, \dot{f}$ without much additional computational

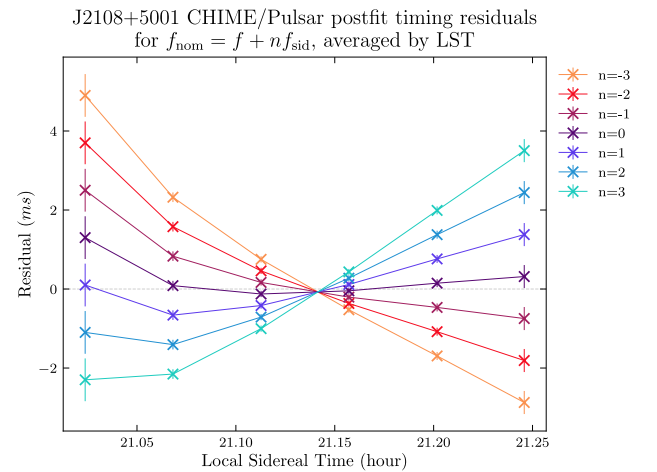


Figure 13. LST-averaged residuals for aliased timing solutions for J2108+5001. For each line, the frequency in the timing solution was adjusted by $n f_{\text{sid}}$, then TOAs were refitted, and residuals at similar local sidereal times were averaged together. A slope indicates residuals shift in phase over the course of an observation and therefore that the frequency is an alias of the true value.




Figure 14. (Top:) Varying period found from the single-day candidate folds of PSR B2303+46, in an eccentric ($e = 0.658$) 12.34 day orbit. The pulsar was treated as an unknown candidate, folded with a constant period from the peak of f in the stack, and folded daily using the standard $f - \dot{f}$ grid search for candidates. The fit orbital solution was close enough that the pulsar could be phase-aligned using `pink`, assigning jumps to phase-wrapping TOAs, with pre- and post-fit residuals shown in the bottom two panels. Note the vastly different scales: The post-fit residuals are contained within $<0.5\%$, while the pre-fit residuals are scattered across the entire pulse phase.

pressure. We set the values using $v_{\max} = 1000 \text{ km s}^{-1}$, $P_{b,\min} = 2 \text{ hr}$. With this scheme, even tight, high-velocity orbits can be detected if the pulsar is sufficiently bright to be seen in individual days. A starting estimate for the orbital parameters can be obtained by fitting $f(t)$, although this can also lead to aliased solutions due to observing at the same sidereal time each day. A significant measure of \dot{f} , or an observation offset from CHIME’s observing window, can immediately break the aliasing degeneracy.

To phase-connect binary pulsars, TOAs in single days are produced as described in Section 4.3. Then the incoherent fit for the binary orbital parameters is included in the timing model, and it is inspected for correlated structure in the TOAs and to see if the solution needs phase jumps applied. The requirement for the pulsar to be visible in individual days, along with the benefit of being able to measure \dot{f} from orbital acceleration, makes CHIME/Pulsar the preferred instrument for these techniques.

As proof of concept, we tested our algorithm on PSR B2303+46, which was in the decl. range of our commissioning survey. This pulsar has a spin frequency of $f \approx 0.938 \text{ Hz}$ and is in a 12.34 day orbit with eccentricity $e = 0.658$. The pulsar was folded blindly on its candidate peak f , DM, from the power-spectrum stack; the daily best-fit f , least-squares orbital fit, and final phase-connected timing solution are shown in Figure 14. Currently, this is done on a case-by-case basis, for sources that are bright enough to be seen on individual days but are not phase-aligned by the algorithm in Section 4.2.

Caveat on detecting binary pulsars in power-spectra stacks. The aforementioned strategies both assume that the source is sufficiently bright to be detected either in the single-day power spectra, or in the stacks. However, the binary motion can have the effect of smearing the signal across multiple power-

spectrum bins, decreasing S/N. The signal will be smeared by

$$\Delta f = \frac{v_r}{c} f_0, \quad (7)$$

scaling both with the radial velocity of the orbit and with the spin frequency. This effect will thus be most deleterious for rapidly rotating pulsars in tight binaries.

In later iterations of the survey, we will include some strategy to search for orbitally modulated signals in the power-spectra stacks. For an orbit with $P_b \ll T_{\text{stack}}$, with sufficiently many observations, the power distribution in the stack will reflect the integral of v_r/c throughout the orbit. These could be searched for using additional binning or matched filters; we leave this to future work.

5. Operations

5.1. Workflow

Originally developed for the CHIME/FRB team, Workflow is an in-house framework designed to manage the lifecycle of processes, including execution, pipelining, result saving, and querying. It is agnostic to underlying parameters such as hardware, software, compute environment, or processing constraints. Workflow provides a command-line interface (CLI) and Python modules that enable users to queue tasks (“Work”) into collections (“Buckets”) in a MongoDB database. These tasks can then be processed through the CLI or Python, with the results automatically stored in a separate MongoDB collection (“Results”).

An additional feature, Pipelines, allows users to define sequential or concurrent tasks in YAML files, specifying schedules and resource requirements such as CPU and RAM. Pipelines integrate with Docker Swarm,⁴⁰ a multi-node orchestration tool for managing containerized workloads, to automate the scheduling and execution of tasks based on available resources. This ensures efficient parallel processing across nodes meant to handle the entire workflow—task submission, container scheduling, execution, and result storage.

The Workflow Web interface provides real-time status monitoring of Buckets, Pipelines, and Results, as well as advanced querying and visualization of outputs, including plots linked from stored file paths.

We utilize Docker Swarm’s replicated mode by creating a Docker Service for each pointing, where each service corresponds to a single Docker Container (referred to as a “replica” within Docker services, which typically have many replicated containers per each service). Using one replica per service is deliberate, due to each pointing process requiring a unique CPU and memory reservation that cannot be replicated across a single service.

To coordinate operations, a manager service determines the available daily processes based on the specified raw data folder, and it uses the Docker SDK for Python to launch single-pointing jobs as services. A cleaner service monitors for completed tasks (successful or failed), outputs logs to a designated folder, and removes the completed services from the Docker Swarm state. Currently, this is designed to operate in quasi-real time, determining and running all processes for a given day and stepping through days sequentially (with the ability to process any day for which there are data on disk).

⁴⁰ <https://docs.docker.com/engine/swarm/>

Each service is assigned CPU and memory reservations according to an empirically derived formula based on the maximum DM per pointing. Docker Swarm manages job scheduling by comparing the cumulative memory reservations against the total available memory of the node, determining whether a job should be “pending” or “running.” The allocation of jobs across nodes follows a round-robin strategy based on available resources. To enforce task ordering, which is necessary for the RFI algorithm, the manager service waits if any one job is in a “pending” state.

Given that each CHAMPSS compute node has a CPU with 128 logical cores paired with 256 GB of RAM (described in Section 3.1), the number of threads per pointing process is set to the reserved memory (in GB) divided by three, as CPU usage scales proportionally with memory usage. This allocation prevents CPU overscheduling, ensuring optimal performance for other system tasks that need threads (hence why the reserved memory is not strictly divided by two) and processing tasks (that benefit from having fully available threads). Each container starts by using the Workflow CLI to poll for jobs in its Workflow Bucket and completes once the pointing process finishes. The container is booted up from an image that, on the first job for a node, is pulled from a local Docker Image registry. The registry is updated using GitHub Actions upon any modification to the main GitHub branch of our codebase.

This setup also applies to other CHAMPSS tasks, including the multi-pointing candidate grouping (Section 3.11) and the folding (Section 4), which are run daily after all individual pointings have been processed. In future iterations, a single YAML file will define the entire CHAMPSS pipeline, specifying Python modules and functions, their execution order, and jobs counts per step. This YAML configuration will be deployed to an on-site Workflow Pipelines server running within a Docker Service. Workflow will then manage this full process lifecycle automatically for the team. Currently, the team employs Workflow Buckets, Results, and the Web interface only, while integration of our custom Docker Swarm scheduling into Pipelines is ongoing.

5.2. Metadata Database

Our pipeline uses a local MongoDB database. This database contains the following collections:

1. Pointings in pointing map.
2. Available pipeline processes.
3. Processed observations.
4. Existing power-spectra stacks.
5. Known sources.
6. Sources that are followed up.

These collections are accessed and updated during the various processing steps of the pipeline. This database is separate from the the Workflow database (see Section 5.1).

The known sources are first populated using the ATNF database, and then updated by querying the Pulsar Survey Scraper (D. L. Kaplan 2022).

5.3. Candidate Viewer

In order to visualize our daily and monthly data products, a Candidate Viewer (shown in Figure 15) was developed to map the distribution of multi-pointing candidates across the sky.

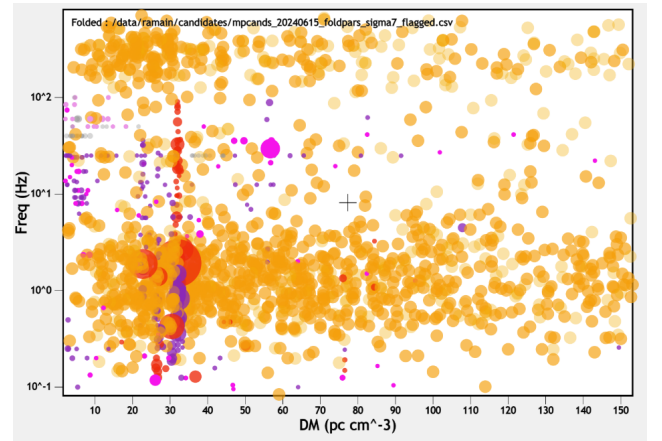


Figure 15. Example data visualization for a given day’s data product. Orange points represent known pulsars from the PSRCAT database, purple and pink points are multi-pointing candidates from the pipeline, and red points are candidates that were associated with a known source. The radius of each circle scales with the σ of that candidate.

Constructed using a browser-based framework called p5.js,⁴¹ this tool can load summary csv files to represent these candidates as points in the user-defined phase space (R.A./decl., Frequency/DM, etc.). The customizable viewing port then allows users to visually compare candidates to each other and to known pulsars with similar values. The data can also be filtered and sorted dynamically by any parameter.

If the user finds an interesting candidate, they can select it to get a more detailed overview of its properties and associated plots, with the option to give the candidate a label shared with the rest of the team. Furthermore, since the tool can interface with Workflow, users can queue a folding job on the candidate to the computing cluster, making it easy to find and follow up on promising sources manually.

While the candidate viewer was developed specifically for CHAMPSS, it can be broadly applied to any table with coordinates, numerical values, and optional metadata. The viewer will be shown in detail, along with a code release, in an accompanying research note (L. Tarabout et al. 2025, in preparation).

6. Commissioning Surveys

To test the system, we have performed multiple commissioning surveys including only a subset of the sky.

The first survey was performed in 2021 December/2022 January. Data were recorded for 34 days in beam rows 125–132 and R.A. range 313.7–322.3. The data were shipped off-site for further analysis because, at that point, the pipeline was not fast enough to run in real time and the on-site CHAMPSS processing nodes did not exist. Unfortunately, these were affected by severe RFI, so additional data were taken for 11 days in 2022 June. The analysis of these data resulted in the discovery of J2108+5001.

Our CHAMPSS commissioning cluster was installed at the CHIME site in 2022 (see Section 3.1). With this cluster, quasi-real-time processing is possible for a small fraction of the CHIME sky. We performed a real-time commissioning survey that continuously recorded data for over a month from each of beam rows 120–123, 124–127, and 128–133, sequentially,

⁴¹ <https://p5js.org/>

lasting from 2023 October to 2024 June (56 beams total). These data were automatically searched daily, in addition to the ~ 1 –2 month incoherent stack. A footprint of this commissioning survey is shown in Figure 2.

6.1. Results

The commissioning surveys have resulted in the discovery or codiscovery of 11 pulsars. Of the pulsars, seven are new and one is a nearly contemporaneous codiscovery with FAST-GPPS (J. L. Han et al. 2025). The other three appear likely to be matches to pulsars discovered in power-spectra stacks by the Pushchino Multibeam Pulsar Search (S. A. Tyul’bashev et al. 2024) or by the CRAFTS survey, as marked in their online table.⁴² These codiscoveries were all in different frequency bands, and we present the first phase-connected timing solutions in all cases.

The newly discovered pulsars are all isolated, spanning a period range of 0.24–1.46 s. Profiles of our newly discovered pulsars over 10 days are shown in Figure 16, and their basic properties are listed in Table 1. Of the 11 pulsars, six have sufficient follow-up observations to derive their timing solutions at the time of writing. Their parameters fitted by our timing pipeline are also listed in the Table 1, and their timing residuals are presented in Appendix B. New pulsars and updated ephemerides will be uploaded to our website.⁴³

To estimate the flux densities of our newly discovered pulsars, we start with CHIME/Pulsar pulse profiles, generated as described in Section 4.3. The off-pulse mean (\bar{p}) of these uncalibrated profiles corresponds to the system equivalent flux density (SEFD) of our telescope. We can calibrate our profile, $I(\phi)$, by subtracting and then dividing by \bar{p} , then multiplying by SEFD to convert it to units of Jy:

$$I(\phi)_{\text{cal}} = \text{SEFD} \left(\frac{I(\phi) - \bar{p}}{\bar{p}} \right). \quad (8)$$

The mean flux density is then simply the area under $I(\phi)_{\text{cal}}$ divided by the number of phase bins. To use this method, we need the SEFD for each pulsar, for which we use calibrator radio sources monitored by CHIME/Pulsar, as outlined in Section 3.5 of F. Dong (2024). Measuring the calibrator sources when they are on sky and then off sky and comparing their cataloged flux densities, we can get the system temperature and the gain, and thus the SEFD, for the pointings corresponding to each pulsar. We apply the frequency-averaged SEFD values to the uncalibrated profiles and calculate a mean flux density S_{600} for each pulsar, listed in Table 1. The weakest pulsar we detect is J2302+48, with a flux density of 0.24 ± 0.04 mJy. By comparing with many calibrators located at different parts of the sky, the total uncertainty is on the order of $\sim 15\%$ (F. Dong 2024). The quoted uncertainties include this value added in quadrature to the measurement error.

We also folded archival GBNCC data (2 minutes durations, spanning 300–400 MHz) toward the closest pointing of each newly discovered pulsar, finding weak detections for J1629+4636, J1802+47, J1900+5106, J2151+5128, J2238+5015,




Figure 16. Panorama of newly discovered pulsars. The pulsars have been averaged in frequency across the band, binned to 128 phase bins, and averaged over 10 days.

and J2302+4807. The observations, which occurred in 2010, were too separated in time to phase-connect to ours.

6.2. Known Pulsars and System Performance

To better quantify our sensitivity, we use detections of known pulsars that lie within the R.A./decl. range of our commissioning survey and have previously known flux densities. Using the radiometer equation, we calculate the expected minimum detectable flux density for each known pulsar,

$$S_{\min} = \beta \frac{S/N_{\min}(T_{\text{rec}} + T_{\text{sky}})}{G \sqrt{n_p t_{\text{int}} \Delta \nu}} \sqrt{\frac{W}{P - W}}, \quad (9)$$

where the minimum detection threshold $S/N_{\min} = 7$, the average gain at zenith $G = 1.16 \text{ K Jy}^{-1}$ (D. C. Good et al. 2021), $n_p = 2$ is the number of polarizations, t_{int} is the integration time, $\beta = 1.15$ is the approximate loss from digitization, $\Delta \nu = 200 \text{ MHz}$ is the effective bandwidth,⁴⁴ W is the pulse width, and P is the spin period.

The process is similar to the one outlined in Section 3.3.4 of D. C. Good et al. (2021). We assume a different receiver temperature of 50 K based on the average of the calibrator

⁴² http://groups.bao.ac.cn/ism/CRAFTS/202203/t20220310_683697.html

⁴³ https://chime-sps.github.io/pulsars_webpage/pulsars/

⁴⁴ CHIME has a native bandwidth of 400 MHz. The bad channel and the RFI cleaning in CHAMPSS remove $\sim 50\%$ of the bandwidth, leaving $\Delta f = 200 \text{ MHz}$.




Figure 17. Comparison of expected vs. actual signal-to-noise values for folded known pulsars that lie within the CHAMPSS commissioning survey R.A./decl. range. An upper limit is indicated for pulsars that are not detected by CHAMPSS, and the dashed line is the line of equality. PSR J1800+5034 (far right) is an outlier, with the expected signal-to-noise ratio being an order of magnitude greater than what is observed. It is possible the ATNF Pulsar Catalog S_{400} value from GBNC (23.5 mJy) is an overestimate, since the S_{150} value recorded from LOTAAS (6.6 mJy) is significantly lower, meaning that GBNC may have simply observed it when it was particularly bright, possibly due to scintillation.

sources used in the SEFD calibration method of F. Dong (2024), as the previous value of 30 K has consistently been shown to underestimate flux densities of known pulsars. For the sky temperature, we use a global sky model from `pygds` that combines several surveys, as recommended by D. C. Price (2021). This global sky model includes the 408 MHz all-sky map of C. G. T. Haslam et al. (1982). We use S_{600} values from the Australian Telescope National Facility (ATNF), or if only available in other bands, extrapolate to 600 MHz using an assumed spectral index value of -1.8 . Using the minimum detectable flux density and the previously recorded flux density, we can predict a signal-to-noise ratio:

$$\left(\frac{S}{N}\right)_{\text{predicted}} = \frac{S_{600}}{S_{\text{min}}}. \quad (10)$$

We compare this predicted signal-to-noise ratio to the actual signal-to-noise ratio of the pulse profile from folded CHAMPSS observations, as seen in Figure 17. The pulse profile is created by phase-connecting multiple (up to 30) days of observations using the known pulsar ephemeris. In general, we find good agreement between the predicted and detected signal-to-noise ratios.

6.3. Cygnus Region/Implications for Galactic Electron Models

Three of the newly discovered pulsars show DM in excess of the maximum predicted line-of-sight DM from one or both of the NE2001 or YMW16 Galactic electron density models (see Figure 18). Two pulsars, PSR J1629+4636 and PSR J1900+5106, are well off the Galactic plane (with $g_b = 43^\circ$, $1, 19.4^\circ$, respectively), in regions devoid of many pulsars. Such sightlines are poorly sampled for Galactic electron density models, and thus they are more uncertain.

The third pulsar is PSR J2108+5001, which is in the Galactic Plane at $g_l = 91.2^\circ$, $g_b = 1.47^\circ$. This lies on the

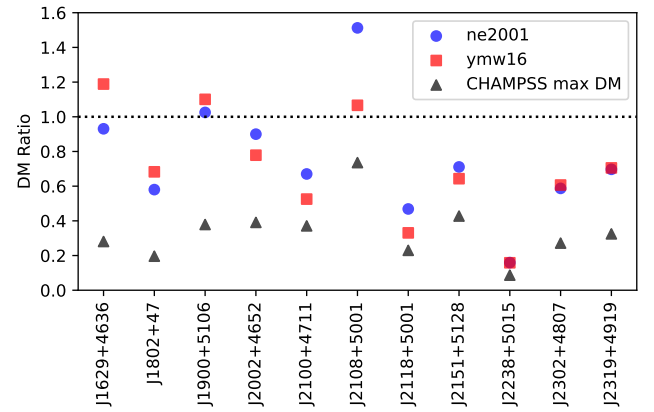


Figure 18. Ratio of the new discovered pulsar's DM to the maximum from NE2001, YMW16, and the max search DM from the pointing map. Although Galactic, three pulsars are in excess of NE2001 and/or YMW16, indicating a higher DM than predicted for the entire Milky Way. Despite the excess, all pulsars have DM well below the search limit of CHAMPSS.

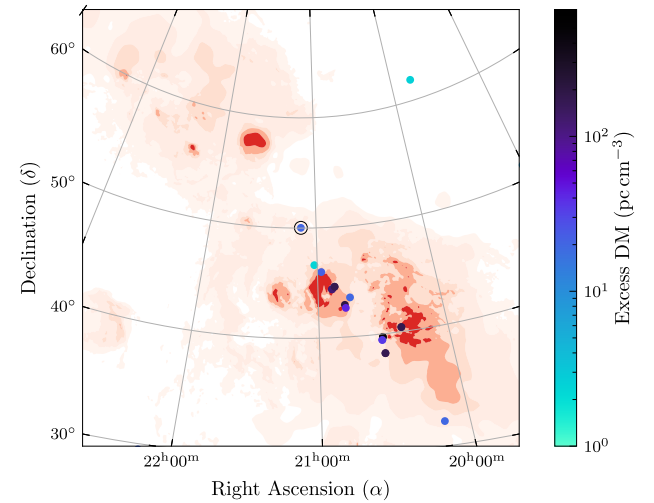


Figure 19. Excess DM beyond the Galactic maximum predicted by the YMW16 model in the vicinity of PSR J2108+5001. Pulsars with DMs within model predictions are not shown, and PSR J2108+5001 is highlighted with a black ring. Contours show the D. P. Finkbeiner (2003) $H\alpha$ map at values of 10, 20, 50, 100, 200, and 500 Rayleigh, and the pulsar colors show the DM in excess of YMW16.

outskirts of the Cygnus star-forming region, which has been found via recent pulsar discoveries by FAST to have DM far in excess of existing Galactic DM models (J. L. Han et al. 2021). S. K. Ocker et al. (2024) show how many pulsars with excess DM and scattering intersect H II regions, including the aforementioned sources. We overlay all pulsars in excess of the predicted YMW16 dispersion measure onto $H\alpha$ contours of the D. P. Finkbeiner (2003) full-sky map, shown in Figure 19. J2108+5001 is near the edge of a large region of higher $H\alpha$ emission that spans $\sim 20^\circ$ and includes the FAST-GPPS pulsars. However, this larger structure could just be due to the Orion-Cygnus arm of the Galaxy. The line of sight to J2108+5001 also intersects a known H II region, G090.856 +01.691 (L. D. Anderson et al. 2014), which could explain its excess DM. In either case, it is clear that the Cygnus region is not modeled well by either NE2001 or YMW16.

J2108+5001 is also noticeably scattered. We fit a scattering time using `fitburst` (E. Fonseca et al. 2024), after stacking the profiles of 10 days together for increased S/N. The

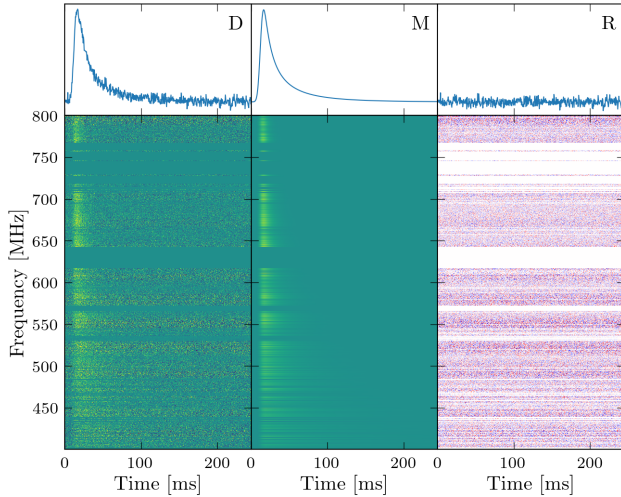


Figure 20. Fitburst (E. Fonseca et al. 2024) modeling of the pulse profile for PSR J2108+5001. The central plot shows the model, with the data to the left and residuals to the right.

resulting scattering time is $\tau = 14.7 \pm 0.3$ ms referenced to 600 MHz, or $\tau = 1.79 \pm 0.5$ ms at 1 GHz. The profile and fit results are shown in Figure 20.

While not extraordinary compared to the most scattered pulsars known, such scattering times have consequences for other transient searches. CHAMPSS will struggle to discover more rapidly rotating pulsars in this region, as the time delay will be comparable to or greater than the rotational period. CHIME/FRB has significantly reduced sensitivity to bursts with widths $\gtrsim 10$ ms (M. Merryfield et al. 2023), suggesting it may be difficult to detect an FRB through this line of sight—or any others intersecting prominent H II regions.

7. Conclusions and Future Work

We have presented an overview of CHAMPSS, a periodicity search using the CHIME/FRB data stream that will search for pulsars daily and in long-term power-spectra stacks. The pipeline runs in real time, and in a commissioning survey covering $\approx 6\%$ of the sky over 2 months, we discovered 11 new pulsars in the period range 0.2–1.5 s.

CHAMPSS will scale up in 2025, following significant hardware upgrades and software optimizations. When operating at scale, we plan to process $>1/4$ of the sky in real time, with the aim of surveying the full northern sky. This survey will be complementary to sensitive targeted surveys; repeated pointing will allow us the ability to discover many intermittent sources that could otherwise be missed in a single pointing (whether due to being intrinsically intermittent or because of propagation effects such as scintillation, lensing, and eclipsing). Discoveries of more pulsars on undersearched lines of sight will help better constrain Galactic electron models, and will be additionally useful in determining the boundary between high-DM Galactic pulsars and low-DM FRBs (e.g., A. M. Cook et al. 2023). Even with only a few discoveries, it is clear that CHAMPSS will be fruitful in this regard, with three sources having DM in excess of the existing models.

As the survey progresses, we will include extensions to the standard FFT-based periodicity search. One such search is a fast folding algorithm (FFA; D. H. Staelin 1969), which coherently sums all harmonics and is particularly well suited to

long-period pulsars, which tend to have short duty cycles (e.g., E. Parent et al. 2018). While FFAs are computationally expensive over equivalent search space of a power-spectrum search, the trial space is small if the search is restricted to long periods. Additionally, we will investigate the efficacy of searching power-spectra stacks for the change in observed frequency caused by the orbital response. These extensions will be expanded on in future work.

One major hurdle of CHAMPSS is dealing in real time with such a large volume of candidates—the scale of this survey is far in excess of the ability of a few humans to verify the candidates. We are developing a machine learning algorithm to filter candidates, trained on real pulsars and realistic injected pulsar signals. These additional features will be presented in future work.

CHAMPSS represents a new era of pulsar searches, as modern radio observatories move to large-N, small-D designs and have PB/day data volumes. Similar survey designs can be used for, e.g., CHORD and the upcoming SKA.

Acknowledgments

We acknowledge that CHIME is located on the traditional, ancestral, and unceded territory of the Syilx/Okanagan people. We are grateful to the staff of the Dominion Radio Astrophysical Observatory, which is operated by the National Research Council of Canada.

CHIME is funded by a grant from the Canada Foundation for Innovation (CFI) 2012 Leading Edge Fund (Project 31170) and by contributions from the provinces of British Columbia, Québec, and Ontario. The CHIME/FRB Project, which enabled development in common with the CHIME/Pulsar instrument, is funded by a grant from the CFI 2015 Innovation Fund (Project 33213), by contributions from the provinces of British Columbia and Québec, and by the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto. Additional support was provided by the Canadian Institute for Advanced Research (CIFAR), McGill University, and the McGill Space Institute, thanks to the Trottier Family Foundation and the University of British Columbia. The CHIME/Pulsar instrument hardware was funded by NSERC RTI-1 grant EQPEQ 458893-2014.

CHAMPSS acknowledges funding from the Canadian Initiative for Radio Astronomy Data Analysis (CIRADA). CIRADA is funded by a grant from the Canada Foundation for Innovation 2017 Innovation Fund (Project 35999) and by the Provinces of Ontario, British Columbia, Alberta, Manitoba, and Quebec, in collaboration with the National Research Council of Canada, the US National Radio Astronomy Observatory, and Australia’s Commonwealth Scientific and Industrial Research Organisation. This research was enabled in part by support provided by Calcul Québec, the BC Digital Research Infrastructure Group, and the Digital Research Alliance of Canada (<https://alliancecan.ca/en>). This work was supported in part by a Canada Excellence Research Chair in Transient Astrophysics (CERC-2022-00009).

E.F. is supported by the National Science Foundation under grant AST-2407399. V.M.K. holds the Lorne Trottier Chair in Astrophysics & Cosmology, as well as a Distinguished James McGill Professorship, and receives support from an NSERC Discovery grant (RGPIN 228738-13). K.W.M. holds the Adam J. Burgasser Chair in Astrophysics. J.M.P. acknowledges the support of an NSERC Discovery Grant (RGPIN-2023-05373).

A.B.P. is a Banting Fellow, a McGill Space Institute (MSI) Fellow, and a Fonds de Recherche du Quebec—Nature et Technologies (FRQNT) postdoctoral fellow. U.P. is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) [funding reference number RGPIN-2019-06770, ALLRP 586559-23, RGPIN-2025-06396], Ontario Research Fund-Research Excellence (ORF-RE Fund), Canadian Institute for Advanced Research (CIFAR), and AMD AI Quantum Astro. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. S.M.R. is a CIFAR Fellow and is supported by the NSF Physics Frontiers Center award 2020265. Pulsar and FRB research at UBC is funded by an NSERC Discovery Grant and by the Canadian Institute for Advanced Research.

Appendix A

Description of Candidate Plots

Figure 21 shows an exemplary multi-pointing candidate plot. This section describes the various segments of the candidate plot based on the read labels that were added to the plot. Fields A–J show information about the strongest single-pointing candidate, while fields K–O show information about all single-pointing candidates included in the multi-pointing candidate.

(A) Sigma as a function of DM and frequency. The plotted value is the maximum of the different harmonic sums performed during the search.

(B) The color shows the harmonic sum for each pixel in plot (A) with maximum sigma. A stronger red implies more summed harmonics. The alpha value of each pixel is controlled by the corresponding sigma value.

(C) Slice through plot (A) along the DM axis.

(D) Slice through plot (A) along the frequency axis.

(E) Text field describing the parameters of the single-pointing candidate.

ra: R.A. in degrees.

dec: decl. in degrees.

ra_hms: R.A. in hh:mm:ss.s format.

dec_hms: R.A. in +dd:mm:ss format.

sigma: Maximum sigma of the candidate. Masked frequency bins are not included in the number of summands in the expected distribution during sigma calculation.

sigma_unweighted: Maximum sigma of the candidate. Masked frequency bins are included in the number of summands in the expected distribution during sigma calculation.

freq: Best frequency of the candidate.

period: Best period of the candidate.

dm: Best DM of the candidate.

first_date: First date included in the power-spectrum stack.

last_date: Last date included in the power-spectrum stack.

ndays: Number of days included in the power-spectrum stack.

detections: Number of detections that were clustered to form this candidate.

masked_fraction_at_best_sigma: Fraction of masked frequency bins at maximum sigma. This will also show if some frequency bins have been masked in individual days.

masked_harmonics: Number of harmonics of the best candidate that have been masked completely.

best_harmonic_sum: Number of summed harmonics that results in the highest sigma. Maximum position of plot (J).

strongest_harmonic_frequency: Frequency of the harmonic with the most power.

strongest_harmonic_period: Period of the harmonic with the most power.

injection: Whether the candidate is the result of an injection or not.

(F) Sigma as a function of DM along the best frequency. Same plot as (C), but with a bigger span.

(G) Clustered detections included in this candidate.

(H) Raw power at best frequency across the first 32 harmonics as a function of DM. When a harmonic is completely masked (as in the left plot of Figure 9), a red circle is drawn at that harmonic. This is done to more easily discern masked harmonics from weak harmonics.

(I) Power of each harmonic at the best DM.

(J) Sigma as a function of summed harmonics.

(K) Sky positions of all single-pointing candidates. Size represents the sigma of the candidates.

(L) Frequency and DM of all candidates.

(M) Sigma as a function of the distance to the centroid for all candidates.

(N) Date span and sigma of each candidate.

(O) Text field describing the parameters of the multi-pointing candidate.

ra: R.A. in degrees of the centroid of the multi-pointing candidate. When calculating this centroid, sigma is used as the weight of each candidate.

dec: decl. in degrees of the centroid of the multi-pointing candidate.

mean_frequency: Mean frequency.

mean_dm: Mean DM.

delta_frequency: Frequency span.

delta_DM: DM span.

num_candidates: Number of single-pointing candidates that have been grouped to form this multi-pointing candidate.

known_source_string: Text field describing the output of the known source sifter. For each known source with a positive classification, the name, R.A., decl., likelihood that the candidate is the known source, frequency, and DM are shown.

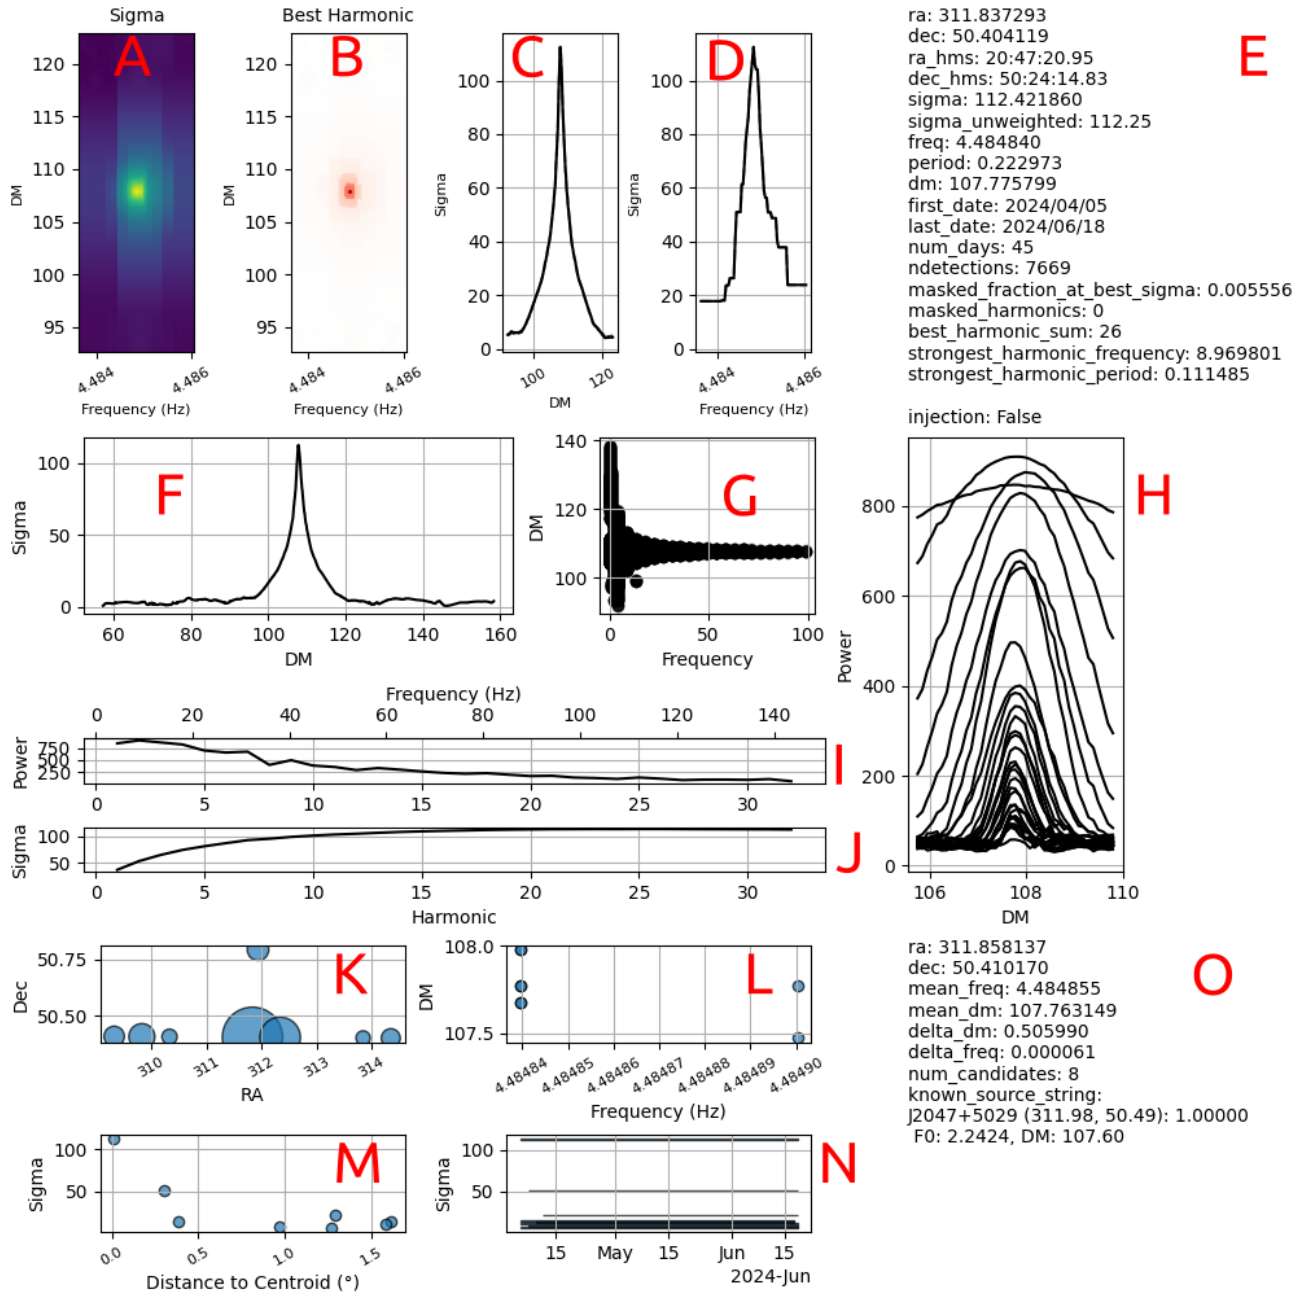


Figure 21. Example multi-pointing candidate plot of PSR J2047+5029. Red alphabetic labels have been added to the individual segments of the plot.

Appendix B

Timing Solution Residuals

In Figure 22, we present timing residuals for our newly discovered pulsars.




Figure 22. Timing residuals for 11 newly discovered pulsars. Pulsar names (top), periods, and the reduced χ^2 values of the fits (bottom) are given to the left of each panel. Black, blue, and red data points are residuals (with their errorbars) of TOAs from CHAMPSS, CHIME/Pulsar fold-mode, and CHIME/Pulsar filterbank (search-mode) observations, respectively.

ORCID iDs

Christopher Andrade <https://orcid.org/0009-0002-2429-3947>
P. J. Boyle <https://orcid.org/0000-0001-8537-9299>
Charanjot Brar <https://orcid.org/0000-0002-1800-8233>
Alyssa Cassity <https://orcid.org/0009-0007-0757-9800>
Kathryn Crowter <https://orcid.org/0000-0002-1529-5169>
Davor Cubranic <https://orcid.org/0000-0003-2319-9676>
Abigail K. Denney <https://orcid.org/0009-0004-5775-8821>
Fengqiu Adam Dong <https://orcid.org/0000-0003-4098-5222>
Emmanuel Fonseca <https://orcid.org/0000-0001-8384-5049>
Victoria M. Kaspi <https://orcid.org/0000-0001-9345-0307>
Ajay Kumar <https://orcid.org/0009-0002-0330-9188>
Lars Künkel <https://orcid.org/0000-0003-4634-5453>
Magnus L'Argent <https://orcid.org/0000-0001-5523-6051>
Dustin Lang <https://orcid.org/0000-0002-1172-0754>
Robert A. Main <https://orcid.org/0000-0002-7164-9507>
Kiyoshi W. Masui <https://orcid.org/0000-0002-4279-6946>
Sujay Mate <https://orcid.org/0000-0001-5536-4635>
Juan Mena-Parra <https://orcid.org/0000-0002-0772-9326>
Bradley W. Meyers <https://orcid.org/0000-0001-8845-1225>
Cherry Ng <https://orcid.org/0000-0002-3616-5160>
Aaron B. Pearlman <https://orcid.org/0000-0002-8912-0732>

Ue-Li Pen <https://orcid.org/0000-0003-2155-9578>
Scott M. Ransom <https://orcid.org/0000-0001-5799-9714>
Kendrick Smith <https://orcid.org/0000-0002-2088-3125>
Reynier Squillace <https://orcid.org/0000-0001-6748-5290>
Ingrid Stairs <https://orcid.org/0000-0001-9784-8670>
Chia Min Tan <https://orcid.org/0000-0001-7509-0117>
Laurent Tarabout <https://orcid.org/0009-0009-5322-0932>
Wenke Xia <https://orcid.org/0009-0009-9343-4193>
Tarik J. Zegmott <https://orcid.org/0000-0002-7076-8643>

References

- Abergel, R., & Moisan, L. 2020, *ACM Trans. Math. Softw.*, 46, 10:1
Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2023, *ApJL*, 951, L8
Anderson, S. B. 1993, PhD thesis, Caltech
Anderson, L. D., Bania, T. M., Balser, D. S., et al. 2014, *ApJS*, 212, 1
Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
Backer, D. C. 1970, *Natur*, 228, 42
Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, *Natur*, 300, 615
Bhat, N. D. R., Swainston, N. A., McSweeney, S. J., et al. 2023a, *PASA*, 40, e021
Bhat, N. D. R., Swainston, N. A., McSweeney, S. J., et al. 2023b, *PASA*, 40, e020
Breton, R. P., Kaspi, V. M., Kramer, M., et al. 2008, *Sci*, 321, 104
Cadelano, M., Ransom, S. M., Freire, P. C. C., et al. 2018, *ApJ*, 855, 125
CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022, *ApJS*, 261, 29
CHIME/FRB Collaboration, Amiri, M., Andersen, B. C., et al. 2021, *ApJS*, 257, 59
CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, *ApJ*, 863, 48
CHIME/Pulsar Collaboration, Amiri, M., Bandura, K. M., et al. 2021, *ApJS*, 255, 5
Cook, A. M., Bhardwaj, M., Gaensler, B. M., et al. 2023, *ApJ*, 946, 58
Cordes, J. M., & Lazio, T. J. W. 2002, arXiv:astro-ph/0207156
Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Natur*, 467, 1081
Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, *ApJL*, 294, L25
Dong, F. 2024, PhD thesis, Univ. British Columbia, doi:10.104288/1.0445624
Dong, F. A., Crowter, K., Meyers, B. W., et al. 2023, *MNRAS*, 524, 5132
EPTA Collaboration, InPTA Collaboration, Antoniadis, J., et al. 2023, *A&A*, 678, A50
Finkbeiner, D. P. 2003, *ApJS*, 146, 407
Folkner, W. M., Williams, J. G., & Boggs, D. H. 2009, *IPNPR*, 1, 42
Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, *ApJL*, 915, L12
Fonseca, E., Pleunis, Z., Breitman, D., et al. 2024, *ApJS*, 271, 49
Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, *Natur*, 333, 237
Good, D. C., Andersen, B. C., Chawla, P., et al. 2021, *ApJ*, 922, 43
Han, J. L., Manchester, R. N., Lyne, A. G., Qiao, G. J., & van Straten, W. 2006, *ApJ*, 642, 868
Han, J. L., Wang, C., Wang, P. F., et al. 2021, *RAA*, 21, 107
Han, J. L., Zhou, D. J., Wang, C., et al. 2025, *RAA*, 25, 014001
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, 47, 1
Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, 21, 302
Huffman, D. A. 1952, *PIRE*, 40, 1098
Johnston, S., Manchester, R. N., Lyne, A. G., et al. 1992, *ApJL*, 387, L37
Kaplan, D. L. 2022, PSS: Pulsar Survey Scraper, Astrophysics Source Code Library, ascl:2210.001
Kouwenhoven, M. L. A., & Voûte, J. L. L. 2001, *A&A*, 378, 700
Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2021, *PhRvX*, 11, 041050
Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, *Natur*, 220, 340
Lattimer, J. M., & Prakash, M. 2007, *PhR*, 442, 109
Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy*, Vol. 4 (Cambridge: Cambridge Univ. Press)
Luo, J., Ransom, S., Demorest, P., et al. 2021, *ApJ*, 911, 45
Lyne, A., Hobbs, G., Kramer, M., Stairs, I., & Stappers, B. 2010, *Sci*, 329, 408
Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993
Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, 328, 17

- Massey, F. J. 1951, *JASA*, 46, 68
- Merryfield, M., Tendulkar, S. P., Shin, K., et al. 2023, *AJ*, 165, 152
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2021, *ApJL*, 918, L28
- Morello, V., Barr, E. D., Cooper, S., et al. 2019, *MNRAS*, 483, 3673
- Morello, V., Barr, E. D., Stappers, B. W., Keane, E. F., & Lyne, A. G. 2020, *MNRAS*, 497, 4654
- Ng, C., Vanderlinde, K., Paradise, A., et al. 2017, in XXXII Int. Union of Radio Science General Assembly & Scientific Symp. (URSI GASS) 2017 (Piscataway, NJ: IEEE), 1
- Nita, G. M., & Gary, D. E. 2010a, *MNRAS*, 406, L60
- Nita, G. M., & Gary, D. E. 2010b, *PASP*, 122, 595
- Nita, G. M., Gary, D. E., & Hellbourg, G. 2016a, in 2016 Radio Frequency Interference (RFI) (Piscataway, NJ: IEEE), 75
- Nita, G. M., Hickish, J., MacMahon, D., & Gary, D. E. 2016b, *JAI*, 5, 1641009
- Ocker, S. K., Anderson, L. D., Lazio, T. J. W., Cordes, J. M., & Ravi, V. 2024, *ApJ*, 974, 10
- Ocker, S. K., Cordes, J. M., Chatterjee, S., & Gorsuch, M. R. 2022, *ApJ*, 934, 71
- Padmanabh, P. V., Barr, E. D., Sridhar, S. S., et al. 2023, *MNRAS*, 524, 1291
- Pan, Z., Hobbs, G., Li, D., et al. 2016, *MNRAS*, 459, L26
- Parent, E., Kaspi, V. M., Ransom, S. M., et al. 2018, *ApJ*, 861, 44
- Petit, G. 2010, *HiA*, 15, 220
- Phillips, C., & Ransom, S. 2022, *AJ*, 163, 84
- Posselt, B., Karastergiou, A., Johnston, S., et al. 2023, *MNRAS*, 520, 4582
- Price, D. C. 2021, *MNAAS*, 5, 246
- Ransom, S. 2011, PRESTO: Pulsar Exploration and Search Toolkit, Astrophysics Source Code Library, ascl:1107.017
- Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *AJ*, 124, 1788
- Reardon, D. J., Zic, A., Shannon, R. M., et al. 2023, *ApJL*, 951, L6
- Rickett, B. J. 1990, *ARA&A*, 28, 561
- Sanidas, S., Cooper, S., Bassa, C. G., et al. 2019, *A&A*, 626, A104
- Staelin, D. H. 1969, *IEEEP*, 57, 724
- Staelin, D. H., & Reifenstein, E. C., III 1968, *Sci*, 162, 1481
- Stinebring, D. R., Rickett, B. J., Minter, A. H., et al. 2022, *ApJ*, 941, 34
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, 791, 67
- Taylor, J. H., & Weisberg, J. M. 1989, *ApJ*, 345, 434
- Thompson, D. J., Arzoumanian, Z., Bertsch, D. L., et al. 1994, *ApJ*, 436, 229
- Tyul'bashev, S. A., Tyul'basheva, G. E., Kitaeva, M. A., et al. 2024, *MNRAS*, 528, 2220
- van der Klis, M. 1989, in NATO ASI Series, Vol 262, Timing Neutron Stars, ed. H. Ögelman & E. P. J. van den Heuvel (Dordrecht: Kluwer), 27
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, 17, 261
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, 835, 29
- Zackay, B., & Ofek, E. O. 2017, *ApJ*, 835, 11