

**Timing Pulsars and Detecting Radio Transients with  
CHIME**

by

Deborah C. Good

B.S. Colorado School of Mines, 2014

M.Sc. University of British Columbia, 2017

A THESIS SUBMITTED IN PARTIAL FULFILLMENT  
OF THE REQUIREMENTS FOR THE DEGREE OF

**Doctor of Philosophy**

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL  
STUDIES

(Physics)

The University of British Columbia

(Vancouver)

August 2021

© Deborah C. Good, 2021

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

“Timing Pulsars and Detecting Radio Transients with CHIME”

---

submitted by Deborah Good in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Physics

**Examining Committee:**

Ingrid Stairs, Physics & Astronomy, UBC

---

Supervisor

Mark Halpern, Physics & Astronomy, UBC

---

Supervisory Committee Member

Aaron Boley, Physics & Astronomy, UBC

---

University Examiner

Raymond Ng, Computer Science, UBC

---

University Examiner

**Additional Supervisory Committee Members:**

Jeremy Heyl, Physics & Astronomy, UBC

---

Supervisory Committee Member

Stefan Reinsberg, Physics & Astronomy, UBC

---

Supervisory Committee Member

# Abstract

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) is a transit telescope located at the Dominion Radio Astrophysical Observatory in Kaleden, BC. Though initially designed to map redshifted neutral hydrogen and constrain dark energy, it also supports several commensal science projects. This thesis focuses on work conducted with the CHIME/FRB fast radio burst searching backend and the CHIME/Pulsar pulsar timing backend.

This thesis focuses on pulsars and fast radio bursts. Pulsars are rapidly rotating, highly magnetized neutron stars, the remnants of massive stars following their supernova explosions. Fast Radio bursts are mysterious millisecond duration radio transients, originating from outside the Milky Way Galaxy. Although their origin is still unknown, evidence is mounting that FRBs also originate from highly magnetized neutron stars known as magnetars.

First, we discuss ongoing efforts to integrate CHIME/Pulsar daily cadence pulsar timing data into large-scale pulsar timing datasets maintained by the North American Nanohertz Observatory for Gravitational Waves (NANOGrav). NANOGrav is engaged in a long-term effort to detect gravitational wave signals from supermassive black hole mergers via pulsar timing. The full NANOGrav array consists of approximately 70 sources; in this initial work, we present timing solutions from CHIME/Pulsar data eight sources. These initial results are promising and suggest a bright future for CHIME/Pulsar-NANOGrav data combination.

We then discuss new pulsars and rotating radio transients (RRATs) discovered via detection of single pulses by CHIME/FRB. CHIME/FRB is ideally situated to detect sources with substantial periods of intermittency or high levels of transience. CHIME/Pulsar's ability to track sources digitally allows us to follow-up initial de-

tections with more conventional search mode observations. The combined effect has allowed us to discover and characterize seven new sources so far.

Finally, we discuss observations conducted with the Arecibo Observatory's 300-m single dish radio telescope, following-up low declination FRBs discovered with CHIME/FRB. This work focused on better understanding repeating FRBs by observing a small number of known repeater and some bursts with repeater-like structure in-dept. It did not result in the detection of new bursts from these sources, but it allows us to constrain the repetition rate of these sources.

# Lay Summary

Pulsars, rotating radio transients (RRATs), and most likely fast radio bursts (FRBs) are signals from neutron stars, very dense objects left behind after a large star dies. In this thesis, I use the radio telescope CHIME (the Canadian Hydrogen Intensity Mapping Experiment) to learn more about these sources. First, I combine CHIME data with data from the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) as the first step towards adding CHIME to NANOGrav's effort to learn more about gravitational waves emitted by black holes millions of times heavier than the sun. (Gravitational waves are ripples in the fabric of space caused by very large objects.) Second, I use CHIME to discover three new pulsars and four new RRATs. Third, I study FRBs discovered with CHIME with the Arecibo Observatory to learn more about how often FRBs repeat.

# Preface

This thesis is a testament to the power of collaboration science, and therefore, it is important to recognize the contributions of those individuals who contributed. All text was written by D.C.G. all figures, unless otherwise noted were created by D.C.G. Data from CHIME/Pulsar and CHIME/FRB are collected by the CHIME/Pulsar and CHIME/FRB collaboration, of which D.C.G. is a member and are used with permission. Throughout this thesis, I.H. Stairs has been an invaluable resource.

- In Chapter 1, Figures 1.2 and 1.3 are from CHIME/FRB publications CHIME/FRB Collaboration (2020a) and Pleunis et al. (2021a) on which D.C.G. is an author. Figure 1.4 is from Arzoumanian et al. (2020), on which D.C.G. is an author. All figures are used with permission.
- In Chapter 3, Figure 3.2 is a photo taken by J. Richard Shaw, and is used with permission.
- Chapter 4 uses the NANOGrav preliminary 15 year dataset generated by Paul Demorest using his code `nanopipe`. The analysis was completed using the NANOGrav 15 year dataset analysis framework, created by the NANOGrav Timing and CyberInfrastructure working groups. All data was used with permission. CHIME/Pulsar data is collected by the CHIME/Pulsar collaboration, of which D.C.G. is a member and used with permission.
- Chapter 5 reproduces a paper accepted by ApJ in August 2021. All data was collected by the CHIME/Pulsar and CHIME/FRB collaborations. Z. Pleunis developed and implemented the CHIME/FRB metadata analysis,

including generating Figures 5.2 to 5.7. K. Crowter designed and implemented `rrat_period_multiday`. B. Meyers completed analysis of PSR J1931+42 and comparisons between known RRATs and new CHIME/Pulsar sources. F. Q. Dong completed population synthesis analysis. E. Fonseca completed timing analysis of PSR J0854+5449 and B. Anderson completed timing analysis of PSR J2108+4516. All exposure data for CHIME/FRB were calculated by P. Chawla.

- The Arecibo Observatory observations used in Chapter 6 were collected by D. C. G., P. Chawla, E. Fonseca, I. H. Stairs, and S. Tendulkar. Again, exposure calculations for CHIME/FRB were conducted by P. Chawla. Figures 6.2, 6.3, and 6.4 reproduce data from the first CHIME/FRB Catalog, and are used with permission (CHIME/FRB Collaboration, 2021).

# Table of Contents

<b>Abstract</b> . . . . .	<b>iii</b>
<b>Lay Summary</b> . . . . .	<b>v</b>
<b>Preface</b> . . . . .	<b>vi</b>
<b>Table of Contents</b> . . . . .	<b>viii</b>
<b>List of Tables</b> . . . . .	<b>xii</b>
<b>List of Figures</b> . . . . .	<b>xv</b>
<b>List of Abbreviations</b> . . . . .	<b>xxvi</b>
<b>Acknowledgments</b> . . . . .	<b>xxx</b>
<b>1 Introduction</b> . . . . .	<b>1</b>
1.1 Pulsars . . . . .	2
1.2 Intermittent Pulsars, Nulling Pulsars, and RRATs . . . . .	7
1.3 Fast Radio Bursts . . . . .	10
1.4 Understanding Repeating Fast Radio Bursts in Depth . . . . .	14
1.4.1 FRB Morphology and Repetition . . . . .	15
1.4.2 FRB Repetition Properties . . . . .	17
1.5 Applications of Millisecond Pulsar Timing & Pulsar Timing Ar- rays . . . . .	19
1.6 Outline of the Thesis . . . . .	21

<b>2</b>	<b>Observing Pulsars and Repeating Radio Transients . . . . .</b>	<b>24</b>
2.1	Pulsar & FRB Observations . . . . .	24
2.2	Radio astronomy in the interstellar medium . . . . .	26
2.2.1	Dispersion Measure . . . . .	27
2.2.2	Rotation Measure . . . . .	30
2.2.3	Scattering . . . . .	31
2.2.4	Scintillation . . . . .	31
2.3	Searching for Pulsars and Repeating Radio Transients . . . . .	32
2.3.1	Pulsar & Transient Search Strategies . . . . .	32
2.3.2	Finding and Verifying Pulsar Candidates . . . . .	33
2.4	Pulsar Timing . . . . .	37
2.4.1	Timing Corrections . . . . .	40
2.4.2	Timing Isolated Pulsars . . . . .	42
2.4.3	Timing Binary Pulsars . . . . .	43
2.5	Pulsar observation in this work . . . . .	48
<b>3</b>	<b>The Canadian Hydrogen Intensity Mapping Experiment . . . . .</b>	<b>49</b>
3.1	The Canadian Hydrogen Intensity Mapping Experiment . . . . .	49
3.2	A brief introduction to interferometry . . . . .	50
3.3	CHIME technical information . . . . .	56
3.3.1	The Analog Front-end . . . . .	57
3.3.2	The F-Engine . . . . .	59
3.3.3	The X-Engine . . . . .	59
3.4	CHIME/FRB System Overview . . . . .	62
3.4.1	CHIME/Pulsar . . . . .	64
<b>4</b>	<b>Integrating CHIME/Pulsar data into the NANOGrav dataset . . . . .</b>	<b>67</b>
4.1	Introduction and Motivation . . . . .	67
4.2	The Unique Challenges of CHIME/Pulsar Timing . . . . .	69
4.3	Creating CHIME/Pulsar TOAs . . . . .	70
4.3.1	RFI Removal . . . . .	71
4.3.2	Creating TOAs . . . . .	71
4.4	Creating the Timing Model . . . . .	72

4.5	First Joint Timing Results from CHIME/Pulsar . . . . .	76
4.5.1	General Discussion . . . . .	76
4.5.2	Individual Pulsar Discussion . . . . .	82
4.6	Future Directions . . . . .	104
4.6.1	Improving Timing Solutions . . . . .	104
4.6.2	Applying the Joint Dataset . . . . .	108
<b>5</b>	<b>Detection of New Galactic Sources with CHIME/FRB . . . . .</b>	<b>109</b>
5.1	Introduction . . . . .	109
5.2	Detection, timing, and analysis methods . . . . .	111
5.2.1	Initial detection . . . . .	111
5.2.2	CHIME/FRB metadata analysis . . . . .	112
5.2.3	CHIME/Pulsar data analysis . . . . .	113
5.3	New sources detected with CHIME/FRB . . . . .	122
5.3.1	Source without spin period measurement: PSR J1931+42 . . . . .	124
5.3.2	Highly intermittent sources . . . . .	127
5.3.3	Persistent Sources . . . . .	130
5.3.4	Discussion of results . . . . .	131
5.3.5	Potential for Future Discoveries . . . . .	135
5.4	Conclusions & Future Work . . . . .	136
<b>6</b>	<b>Follow-up Observations of CHIME/FRB Repeaters with Arecibo Ob-</b>	
	<b>servatory . . . . .</b>	<b>138</b>
6.1	Introduction . . . . .	138
6.2	Motivation . . . . .	141
6.3	Technical Details of Follow-up Program . . . . .	142
6.3.1	Source information . . . . .	142
6.3.2	Search Strategy . . . . .	146
6.4	Results . . . . .	150
6.4.1	Understanding Sensitivity Limits . . . . .	150
6.4.2	Constraints on Burst Rate . . . . .	152
6.4.3	Comparison with other repeaters . . . . .	162
6.5	Discussion . . . . .	165

<b>7 Conclusion</b> . . . . .	<b>168</b>
<b>Bibliography</b> . . . . .	<b>172</b>
<b>A Parameter files from CHIME/Pulsar Timing Data</b> . . . . .	<b>189</b>

# List of Tables

Table 4.1	Summary of joint NANOGrav and CHIME/Pulsar timing model fits, including number of spin, astrometric, and binary parameters, number of DMX bins, and frequency dependent profile bins, jumps, and RMS residual values for both the total dataset and the CHIME components. . . . .	79
Table 4.2	Timing parameters and their uncertainties (in parentheses) from combined NANOGrav and CHIME/Pulsar data . . . . .	80
Table 4.3	Ratio of timing parameter uncertainties between CHIME/Pulsar solutions and combined solutions for isolated pulsars, for all parameters fit separately in the two models. The large ratios for $v_s$ and $\dot{v}_s$ uncertainties reflect the much shorter dataspan in the CHIME/Pulsar only dataset. . . . .	81
Table 4.4	Ratios of timing parameters uncertainties between CHIME/Pulsar solutions and combined solutions for binary pulsars, for all parameters fit in both models. Parameters marked “NG” are held constant at combined CHIME/Pulsar-NANOGrav values, parameters marked “–” are not included for this pulsar. . . . .	81
Table 4.5	Noise parameters for CHIME/Pulsar, calculated with the joint CHIME/Pulsar-NANOGrav dataset. . . . .	82
Table 4.6	Parameter comparison between the combined NANOGrav and CHIME/Pulsar dataset and the results of Fonseca et al. (2021). . . . .	87
Table 5.1	Period determination methods for each of the seven sources. . . . .	123

Table 5.2	Positions and dispersion measures for all new sources. Distances are determined by comparison with the NE2001 Galactic DM map (Cordes & Lazio, 2002). . . . .	125
Table 5.3	Basic timing parameters for all new sources. . . . .	126
Table 5.4	Flux density measurements for all sources. The range of peak flux densities for the intermittent sources are given in the $S_{600}^{\text{peak}}$ column, while for the more persistent sources we provide flux density estimates averaged over the entire pulse period in the $S_{600}^{\text{mean}}$ column. . . . .	126
Table 5.5	Best-fit parameters and derived quantities for PSRs J0854+5449 and J0209+5759. . . . .	133
Table 5.6	Confidence intervals of p-values from the bootstrapped Kruskal-Willis H-tests comparing the CHIME and known RRAT distributions. . . . .	135
Table 6.1	A summary of positions and dispersion measures for known repeaters discovered with CHIME/FRB and observed with AO. Reported DMs are structure optimizing. Methods for determining DM and position are explained in detail in (CHIME/FRB Collaboration, 2019) and (Fonseca et al., 2020). $N_{\text{bursts}}$ refers only to CHIME/FRB bursts; there were no detections with AO. All exposure times for CHIME/FRB data are calculated between August 30, 2018 and July 1, 2019. . . . .	147
Table 6.2	A summary of positions and dispersion measures for repeater-like FRBs discovered with CHIME/FRB and observed with AO. Reported DMs and locations are those presented in CHIME/FRB Catalog 1, and determined by the methods discussed therein. All exposure times for CHIME/FRB are calculated between August 30, 2018 and July 1, 2019. . . . .	148
Table 6.3	Basic properties of Arecibo 327 MHz Receiver and of CHIME/FRB.	149

Table 6.4 Sensitivity thresholds for AO and CHIME/FRB, calculated using the formalism of Cordes & McLaughlin (2003) and the parameters in Table 6.3. FRB 20190116A and FRB 20190117A (bolded) are known repeaters . . . . . 152

Table 6.5 Poisson repetition rates for FRBs selected for followup with Arecibo. We report observed repeat rate, repeat rate scaled to account for sensitivity, and the observed CHIME/FRB and AO joint repeat rate, both observed and scaled. For bursts without known repetition, we present upper limits. Uncertainties are set by the uncertainty in CHIME/FRB exposure times, which is substantial. . . . . 154

Table 6.6 Posterior means and for CHIME/FRB + AO repetition rates and clustering coefficients, based on a Weibull distribution of burst rates. The Weibull distribution reduces to the Poisson distribution when  $k = 1$ ; the low  $k$  values listed here suggest a high degree of clustering in bursts from these sources. Lower  $k$  values also decrease the predictive power of  $r$  estimates. Note that  $r$  is in units of  $\text{hr}^{-1}$ , whereas it is presented in units of  $\text{days}^{-1}$  in Oppermann et al. (2018). FRB 20190116A and FRB 20190117A (bolded) are known repeaters . . . . . 159

# List of Figures

- Figure 1.1 A period vs. period-derivative plot, showing all known pulsars. This figure provides a high level view of the pulsar population, highlighting key sub-populations such as magnetars and MSPs. Isolated pulsars are represented by green dots, binary pulsars are blue triangles, pulsars with known supernova remnant associations are orange x's, magnetars are red stars, and rotating radio transients are purple squares. The solid black lines are lines of constant magnetic field, from  $10^8$  to  $10^{14}$  G. The dashed lines are characteristic pulsar ages, from 1 kYr to 1 GYr. The dash-dot lines are  $\dot{E}$  estimates at  $10^{29}$ ,  $10^{33}$  and  $10^{37}$  erg  $s^{-1}$ . Pulsar data are from the ATNF Pulsar Database (Manchester et al., 2016) (<https://www.atnf.csiro.au/people/pulsar/psrcat/>), magnetar data are from the McGill Magnetar Catalog (Olausen & Kaspi, 2014), and RRAT data are from the RRATalog (<http://astro.phys.wvu.edu/rratalog/>). 8
- Figure 1.2 Dynamic spectra for four CHIME/FRB bursts, demonstrating the four major FRB morphology archetypes. From left to right, FRB 20190527C displays simple wideband structure, FRB 20190515D displays simple narrowband structure, FRB 20181117B displays complex wideband structure, and FRB 20190117A displays complex narrowband structure (Pleunis et al., 2021a). 12

Figure 1.3 SGR 1935+2154 in context, with pulsar and FRB bursts (CHIME/FRB Collaboration, 2020a). The x-axis is distance from Earth in parsec, the y-axis is fluence. Grey diagonals show assorted energy thresholds in ergs. CHIME/FRB measured the burst energy of SGR 1935+2154 as  $\sim 3 \times 10^{34}$  erg, which bridges the energy gap between Galactic magnetars and FRBs. . . . . 14

Figure 1.4 This figure from Arzoumanian et al. (2020) shows the average angular distribution of cross-correlated power in the NANOGrav dataset. Panel (a) shows the 11 year dataset; Panel (b) shows the 12.5 year dataset. The blue dashed line shows a Hellings-Downs correlation, the expected signature of a gravitational wave detection, and the orange dashed line shows a monopole distribution for comparison. Here,  $\hat{A}^2 \Gamma_{ab}(\zeta)$  takes the place of  $\alpha$  as the cross-correlation. The points are binned; they represent the same number of bins for both panels, but the bins contain different numbers of pulsar pairs. . . . . 22

Figure 2.1 Results from folding data from PSR J0209+5759 taken on MJD 58527 with a slightly incorrect period (a) and with the correct period (b). The greyscale plot shows time vs. phase within the observation, the top inset shows pulse profile vs. phase, and side inset shows reduced  $\chi^2$  vs. fraction of the observation. The tilt in the phase vs. time plot in Panel b is indicative of an incorrect period, as the pulse moves slightly out of phase as time progresses. The pulse is visible only for a short portion of the period due to nulling in the source; the source location is visible to the CHIME/Pulsar beam for the duration of the observation. . . . . 36

Figure 2.2	Output from PRESTO routine <code>single_pulse_search.py</code> , showing a single pulse detection of RRAT J0121+53, a source discussed in Chapter 5. Clockwise from the top left, the subplots show number of pulses vs. S/N, number of pulses vs. DM, S/N vs. DM, and DM vs. observation time. In the DM vs. time plot, each black dot is representative of a pulse with $S/N \geq 6$ . The size of the dot indicates the S/N. Most apparent signals in this figure are in fact RFI. This includes the signals that seem to appear at all values of DM (e.g. at 900, 1100, and 1300 s) and those signals which appear at $DM = 0$ , such as those visible between 0 and 200s into the observation. The times showing detections at all DM values and at $DM \sim 0$ are RFI signals; the RRAT is visible three times at $DM \sim 21$ and approximately 500, 650, and 750 seconds. . . . .	38
Figure 3.1	A block diagram demonstrating the path signals take through CHIME. All projects share the same front-end and F-engine; projects diverge in the GPU-based X-engine. . . . .	57
Figure 3.2	The Canadian Hydrogen Intensity Mapping Experiment, as seen from the air. Photo by J. Richard Shaw. . . . .	58
Figure 4.1	Timing residuals and DMX fit for PSR J0645+5158, from CHIME/Pulsar data only. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. . . . .	84
Figure 4.2	Timing residuals and fit for DMX fit for PSR J0645+5158, from CHIME/Pulsar and NANOGrav data. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Blue data points are collected with the GBT L-band receiver, green points are collected with the GBT 800 MHz receiver, and red data points are collected with CHIME. . . . .	85

Figure 4.3	CHIME/Pulsar Timing residuals and DMX fit for PSR J0740+6620. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. . . . .	88
Figure 4.4	NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J0740+6620. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME. . . . .	89
Figure 4.5	CHIME/Pulsar Timing residuals and DMX fit for PSR J1012+5307. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. . . . .	91
Figure 4.6	NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1012+5307. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME. . . . .	92
Figure 4.7	CHIME/Pulsar Timing residuals and DMX fit for PSR J1125+7819. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. . . . .	93

Figure 4.8	<p>NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1125+7819. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME. . . . .</p>	94
Figure 4.9	<p>CHIME/Pulsar Timing residuals and DMX fit for PSR J1713+0747. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. . . . .</p>	96
Figure 4.10	<p>NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1713+0747. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Light blue and purple data points were collected with the AO L-wide and S-wide receivers, respectively. Dark blue and green data points were collected with the GBT L-band and 800 MHz receivers, respectively. Red data points were collected with CHIME. . . . .</p>	97
Figure 4.11	<p>CHIME/Pulsar Timing residuals and DMX fit for PSR J1744-1134. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. . . . .</p>	98

Figure 4.12	NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1744-1134. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME. . . . .	99
Figure 4.13	CHIME/Pulsar Timing residuals and DMX fit for PSR B1937+21. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. The large-scale structure in the residuals is due to red noise, not poor fit. . . . .	101
Figure 4.14	NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR B1937+21. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Light blue and purple data points were collected with the AO L-wide and S-wide receivers, respectively. Dark blue and green data points were collected with the GBT L-band and 800 MHz receivers, respectively. Red data points were collected with CHIME. . .	102
Figure 4.15	Whitened NANOGrav and CHIME/Pulsar timing residuals for PSR B1937+21. The first panel shows whitened timing residuals vs. time, the second shows epoch-averaged, whitened residuals vs. time, and the histograms show the distribution of whitened residuals. The colours correspond to receivers as discussed in Figure 4.14. . . . .	103
Figure 4.16	CHIME/Pulsar Timing residuals and DMX fit for PSR J2145-0750. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. . . . .	105

Figure 4.17	NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J2145-0750. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME. . . . .	106
Figure 5.1	CHIME/FRB detections of new RRAT-like source PSR J0121+53, shown as an example, along with sensitivity to this location and exposure time at this location. The top panel shows the daily median detection S/N and standard deviation of the in-beam events with $S/N > 8$ . The second panel shows the number of single pulses that were detected. If no pulse was detected on a day with non-zero exposure the marker is a gray cross. The third panel shows the daily sensitivity of the experiment as characterized by the relative rms noise of known pulsar detections. The color bar in the inset shows how many pulsars were used to obtain the average rms noise. The bottom panel shows the exposure to the source's localization uncertainty region. These plots have been reproduced for the other six sources in Figures 5.2 to 5.7. . . . .	114
Figure 5.2	CHIME/FRB detections of new pulsar-like source PSR J0209+5759, which is detected consistently with CHIME/FRB. After an initial CHIME/FRB timing solution, this source was determined with CHIME/Pulsar data to be a nulling pulsar, with period 1.06 s. See Figure 5.1 for more details of the figure structure, and Section 5.3.3 for more discussion of the source. . . . .	115

Figure 5.3	CHIME/FRB detections of new pulsar-like source PSR J0854+5449, which is detected consistently with CHIME/FRB. This pulsar has a period of 1.23 s, and was determined by CHIME/Pulsar data to be a conventional pulsar. See Figure 5.1 for more details of the figure structure and Section 5.3.3 for more discussion of the source. . . . .	116
Figure 5.4	CHIME/FRB detections of new RRAT-like source PSR J1252+53. This source is detected only sparsely with CHIME/FRB and with CHIME/Pulsar. We were unable to determine a period estimate using CHIME/FRB data, but using CHIME/Pulsar data, we find an underlying period of 0.220 s. See Figure 5.1 for more details on the figure structure and Section 5.3.2 for more discussion of the source. . . . .	117
Figure 5.5	CHIME/FRB detections of new RRAT-like source PSR J1838+50. This source is detected regularly with CHIME/FRB and CHIME/Pulsar, allowing determination of a period with both sets of data. The period is measured to be 2.577 s. See Figure 5.1 for more details of the figure and Section 5.3.2 for more discussion for the source. . . . .	118
Figure 5.6	CHIME/FRB detections of new RRAT-like source PSR J1931+42, detected sparsely with CHIME/FRB. This is the only source presented in this work for which a period has not been determined. This is not unheard of for RRATs, but further observations may enable us to better characterize this source. See Figure 5.1 for more details on the figure structure and Section 5.3.1 for more information about the source. . . . .	119
Figure 5.7	CHIME/FRB detections of new pulsar PSR J2108+4516. Its spin period is 0.577 s. This pulsar is notable for its intermittency, with months-long timespans where the pulsar is not seen by CHIME/FRB or CHIME/Pulsar. This behaviour is due to interaction within the binary system in which the source resides. See Figure 5.1 for more details about the figure and 5.3.3 for more information about the source. . . . .	120

Figure 5.8	Single pulses detected by CHIME/Pulsar from each highly intermittent source. The single-pulse profile versus pulse phase is in the inset panels and the frequency spectrum versus pulse phase in the main panels. . . . .	129
Figure 5.9	Best-fit timing residuals of PSRs J0209+5759 and J0854+5449 estimated from CHIME/Pulsar filterbank and fold-mode data. The smaller data set for J0209+5759 is due to fewer pulses emitted by the source, in a manner consistent with other known RRATs. Best-fit parameters of the timing model for these sources are shown in Table 5.5. . . . .	132
Figure 5.10	Single observations profiles for J0209+5759, J0854+5449, and J2108+4516. J0209+5759 has S/N 11.6, J0854+5449 has S/N 21.5, and J2108+4516 has S/N 196. . . . .	134
Figure 6.1	This figure plots declination (in degrees) against CHIME/FRB exposure time (shown on a log scale) in hours, for all bursts included in the first CHIME/FRB catalog. Arecibo Observatory’s maximum declination is approximately $+38^\circ$ (shown as a blue vertical line). Sources followed-up in this work are marked as green stars; blue dots are apparently non-repeating FRBs. The additional repeating FRBs within AO’s declination range were observed after the proposal deadline for the data displayed here. . . . .	143
Figure 6.2	Time vs. frequency dynamic spectra for the two known bursts from FRB 20190116A. The red or blue traces represent fits to the profile in time (top panel) and frequency (right panel). The colour indicates whether (blue) or not (red) the final burst fit includes scattering information. Neither is a clear demonstration of downward-drifting morphology; the difference in burst widths and bandwidths highlight the variability of repeat bursts. . . . .	144

Figure 6.3	Time vs. frequency dynamic spectra for two published bursts from FRB 20190117A. Six total bursts from this source have been detected by CHIME/FRB. This source, especially in the burst shown in panel (b) demonstrates clear downward-drifting structure. . . . .	144
Figure 6.4	Time-vs-frequency dynamic spectra for morphologically interesting FRBs at declination $< +38^\circ$ detected with CHIME/FRB and followed-up with Arecibo Observatory. All plots are from CHIME/FRB data (CHIME/FRB Collaboration, 2021) . . . .	145
Figure 6.5	Poisson repetition rates for known repeaters (panel a) and upper limits on repetition for non-repeating bursts (panel b). In both cases, the left-hand panel shows observed and sensitivity scaled CHIME/FRB only rates or limits, while the right-hand panel shows combined CHIME/FRB and AO rates and limits. Note that FRB 20181030E is much narrower than other bursts and therefore has a much higher sensitivity threshold. This causes the causes the scaled rate to increase rather than decreasing . . . . .	155
Figure 6.6	Weibull distribution posterior probability distribution for clustering coefficient $k$ and burst rate $r$ for FRB 20190116A (Panel a) and FRB 20190117A (Panel b). Contours are at 68%, 95%, and 99% probability. The insets are marginalized posteriors for $\log(r)$ and $\log(k)$ . The Oppermann et al. (2018) implementation of the Weibull distribution finds a moderate amount of clustering and burst rates similar to the Poissonian confidence interval burst rates discussed earlier in the chapter. The preponderance of non-detections means that there is little distinction between maximum likelihood $k$ and $r$ values for the two sources. . . . .	160

Figure 6.7	Excess DM ( $\text{pc cm}^{-3}$ ) plotted against number of bursts detected with CHIME/FRB for repeaters published in CHIME/FRB Collaboration (2019) and Fonseca et al. (2020), as well as FRB 20121102 and AO follow-up sources. AO follow-up sources are represented by blue stars; other repeaters are represented by orange circles. There is not a significant correlation between the number of bursts detected by CHIME/FRB and excess DM. Additionally, our AO follow-up sources do not appear to be more distant than observed CHIME/FRB repeaters.	163
Figure 6.8	Average CHIME/FRB peak flux measurement for known repeaters and for AO follow-up sources, plotted against number of bursts. . . . .	164

# List of Abbreviations

AO	Arecibo Observatory; a recently decommissioned, 305 m diameter single dish radio telescope located in Arecibo, Puerto Rico.
ASP	A pulsar backend formerly used at AO; sister instrument to GASP.
CHIME	The Canadian Hydrogen Intensity Mapping experiment; a radio transit telescope located in Kaleden, BC operating at frequencies from 400-800 MHz.
CHIME/FRB	The CHIME Fast Radio Burst detection backend.
CHIME/Pulsar	The CHIME Pulsar timing backend.
DM	Dispersion Measure; a measure of the column density of electrons, which induces a frequency dependent time delay in pulsar and FRB data.
DSPSR	Digital Signal Processing for Pulsars; a C++ software library for pulsar data reduction. See van Straten & Bailes (2011).
ECORR	A pulsar timing noise parameter representing additional white noise present in a timing solution, correlated within an observation.
EFAC	A pulsar timing noise parameter representing additional multiplicative noise present in a timing solution.

ENTERPRISE	Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite; a pulsar timing analysis software, focused on noise analysis and gravitational wave searches. See Ellis et al. (2019).
EQUAD	A pulsar timing noise parameter representing additional white noise present in a timing solution, added in quadrature
FD parameters	Frequency-dependent parameters; polynomial coefficients in log-frequency space fit in pulsar timing models to account for frequency-dependent pulse shape and evolution
FRB	Fast Radio Burst; a high energy, millisecond duration radio transient.
GASP	A pulsar backend formerly used at the GBT; sister instrument to ASP.
GBT	Green Bank Telescope; a 100 m fully-steerable single dish radio telescope located in Green Bank, West Virginia, USA.
GUPPI	The Green Bank Ultimate Pulsar Processing Instrument; a pulsar backend formerly used at the GBT; sister instrument to PUPPI.
IPTA	The International Pulsar Timing Array; an ongoing international effort to combine data from regional PTAs.
ISM	Inter-stellar medium; gas and dust located between stars within the Galaxy. In the context of pulsars, we are mostly concerned with the ionized gas that leads to dispersion.
MJD	Modified Julian Day.
MSP	Millisecond Pulsar; a recycled pulsar with a period less than $\sim 30$ ms.

NANOGrav	The North American NANOHertz Observatory for Gravitational Waves; a North American pulsar timing array, searching for low-frequency gravitational waves using precision timing of about 70 MSPs.
PINT	PINT is not TEMPO 3; a python-based software package used in pulsar timing to fit timing models to TOAs. See Luo et al. (2021).
PRESTO	A python-based software package used for pulsar searching. See Ransom (2011).
PSRCHIVE	An open-source software library for pulsar data analysis, including RFI mitigation, template creation, calibration, and TOA generation.
PTA	Pulsar Timing Array; an effort to compile precision timing data for many MSPs, usually with the goal of gravitational wave detection.
PUPPI	Puerto Rico Ultimate Pulsar Processing Instrument; a pulsar backend formerly used at AO; sister instrument to GUPPI.
RFI	Radio Frequency Interference; human-made signals detected by radio telescopes.
RRAT	Rotating Radio Transient; a pulsar-like radio transient with underlying periodicity and highly intermittent emission.
$S/N$	Signal-to-noise ratio
TEMPO	A fortran based software package used in pulsar timing to fit timing models to TOAs. See Nice et al. (2015).
TEMPO2	A C-based software package used in pulsar timing to fit timing models to TOAs. See Hobbs et al. (2006); Edwards et al. (2006); Hobbs et al. (2009).

TOA

Pulse Time of Arrival; literally, the time at which a pulse from a pulsar or radio transient is observed. The primary unit of data for pulsar observations.

# Acknowledgments

They say that it takes a village to raise a child. I don't know much about raising children, but at this juncture I'm certain that it also takes a village to write a Ph.D. thesis. There are so many people who have helped me reach this point that I could never thank all of them, but I do want to mention a few.

In the interest of giving acknowledgments from the ground up, I want to first acknowledge that this work was conducted at UBC, on the traditional, ancestral, and unceded land of the Musqueam people, using an instrument located on the traditional, ancestral, and unceded territory of the Syilx/Okanagan people.

Among my scientific acknowledgements, the first must be a sincere thank you to Dr. Ingrid H. Stairs, who has done so much to make me the scientist I am today. Thank you for allowing and encouraging me to be independent and to trust myself, while also always supporting me and being ready to offer advice. I couldn't have asked for a better Ph.D. supervisor.

Thank you to my PhD committee. Little as I enjoyed some parts of our conversations, each of our meetings taught me something not to do again. But as one of my fictional heroines once said, "Have you noticed something encouraging about me? I never make the same mistake twice."<sup>1</sup> (At least, not in our committee meetings.)

Thank you to the army of early career researchers in CHIME and NANOGrav who made it possible for me to succeed. In particular, I want to thank Dr. Kiyoshi Masui and Dr. Emmanuel Fonseca. It's certain that without some timely advice from both of them, and the incredible support of Dr. Richard Shaw during my M.Sc., I would never have reached this point.

---

<sup>1</sup>Anne Shirley, in *Anne of Green Gables* by L.M. Montgomery

I've always been proud to be an Oredigger, and I view the education and support I received at the Colorado School of Mines to be determining factor in my success. In particular, I want to thank Dr. Kristine Callan, who was exactly the mentor I needed as a nervous senior, and Dr. Lincoln Carr, who was right that I just needed to put in the hours to get the hang of research.

Particularly for women in the physical sciences, we know that early mentoring is often a key factor in persistence in the field. With that in mind, I want to thank the day camp counselor who encouraged me to study physics if I was serious about astronomy. This led to years of telling stupefied grown-ups that I was going to be an astrophysicist, and it would seem the document you are currently reading is evidence that I did. Thank you as well to my high school math and physics teacher, Mr. Blake, whose classroom was the first place I ever felt the pure joy of understanding physics.

Finally, I can never thank my family enough. Locally, my Sixteenth Avenue Gospel Chapel family has borne my burdens with such faithfulness and loving kindness that it's impossible not to see Christ in them. But more literally, no one will ever support me more thoroughly than my Earthly family. I could never have reached this point without my wonderful parents, who have supported me in every possible sense. It has meant more than i can ever say to know that you were behind me 100% of the way through my grad school journey, no matter what. My brother, Daniel, has stepped ably into the draining role of "best friend to a Ph.D. student writing a thesis," even as he went through a challenging semester and job search of his own. My extended family, too, especially my grandparents have been so supportive and gracious through this process, responding to a disappearing granddaughter with a years-long string of messages telling me how proud they are. I've kept every single one of your cards, and I treasure them.

And there are so many others I haven't named. From motivational stuffed dogs, long-suffering stuffed rabbits, and gelato-loving teddy bears to officemates, human and canine, this thesis has been the work of a plethora of colleagues, friends, and relatives. I will not forget the guidance, kindness, and support I have received through this process, and I hope and pray that I will someday be able to pay it forward in the lives of friends, family, colleagues, and (if I'm lucky) students.

# Chapter 1

## Introduction

*I am among those who think that science has great beauty. A scientist in his laboratory is not only a technician: he is also a child placed before natural phenomena which impress him like a fairy tale.*

— Marie Curie

Astronomy has a great legacy of answering our most basic questions, satisfying the nigh universal human need to know what is “out there” for us to discover. Many of our great questions are long answered, and yet each time we think we may have reached the end of our ability to learn and discover new things, it becomes apparent that there are more mysteries to be explored.

In 1967, such a mystery emerged with the discovery of periodic radio signals, ultimately revealed to be pulsars. Pulsars are rapidly rotating neutron stars, which range broadly from the universe’s best natural clocks to sources with maddeningly intermittent emission, helpful in the extreme for our ability to test fundamental physics while their own fundamental physics remains infuriatingly murky.

Forty years after the discovery of pulsars, we again found ourselves faced with a new and mysterious short-duration radio signal in the form of a fast radio burst (FRB). This first FRB, dubbed the Lorimer burst in honor of its discoverer, but formally known as FRB 20010724 was an unexpected start to an entire new subfield, discussed in more detail in Section 1.3 and 1.4. Though the physical origins of FRBs remain unclear, they do appear to share characteristics (and perhaps origins) with a special sub-set of pulsars known as magnetars. At the ill-defined bound-

aries of pulsar behaviour and perhaps providing some insight into FRB origins sit intermittent pulsars and RRATs, pulsar-like sources with intermittent emission.

In this thesis, I will discuss three major projects, spanning the spectrum from precisely timed millisecond pulsars (MSPs) to only partly understood FRBs. The common thread through this work is the Canadian Hydrogen Intensity Mapping Experiment (CHIME), a modern transit telescope facility located in Kaleden, BC. CHIME enables us to study the sky between 400-800 MHz in unusual ways with large field of view and very high cadence observations.

This unique perspective makes CHIME data an addition with great potential for pulsar astronomy and a complete game changer for FRB astronomy.

In the remainder of this introduction, I will provide a brief overview of pulsars and FRBs, discuss the questions posed by repeating FRBs, and discuss applications of pulsar timing.

## 1.1 Pulsars

Pulsars are, at their simplest, rapidly rotating, highly magnetized neutron stars. These objects were first discovered in 1967, by then Ph.D. student Jocelyn Bell using radio observations at the Mullard Observatory in the United Kingdom (Hewish et al., 1968). The neutron star had first been proposed in Baade & Zwicky (1934), and the theory had been laid out in Oppenheimer & Volkoff (1939). That the pulsar could be a neutron star was posited in the initial discovery paper and by Pacini (1967) and Gold (1968) within a year. This interpretation was bolstered by the discovery of further pulsars, associated the Crab and Vela nebulae and with shorter spin periods (Staelin & Reifenstein, 1968; Large et al., 1968).

Neutron stars form following Type II supernovae, with moderately massive ( $m_{\text{star}} = 10 - 25 M_{\odot}$ <sup>1</sup>) progenitors. After the supernova explosion, most material is disseminated, but a remaining core of material collapses to form an extremely dense object. In very massive systems, this collapse results in a black hole. In moderate mass systems, this creates a star-like object supported by neutron degeneracy pressure. These objects are extremely dense; though a neutron star diameter is only 10-20 km, their masses can be as high as  $2 M_{\odot}$ . The highest reported neu-

---

<sup>1</sup>  $1 M_{\odot}$  is roughly  $2 \times 10^{30}$  kg

tron star mass is  $2.08^{+0.07}_{-0.07}M_{\odot}$  for PSR J0740+6620, which has a radius of  $3.7^{+2.6}_{-1.5}$  km (Fonseca et al., 2021; Riley et al., 2019; Cromartie et al., 2020; Miller et al., 2021).

The term pulsar is applied specifically to neutron stars which are highly magnetized, rapidly rotating, and observable from the perspective of the Earth. There are many more neutron stars than there are observable pulsars, as not all neutron stars would be expected to behave as pulsars. Pulsars are high energy objects, and some are known to emit light at optical, X-ray, and  $\gamma$ -ray wavelengths in addition to radio wavelengths. However, in this thesis, we will focus on observations of radio pulsars.

Radio pulsar signals are coherent, broadband, and highly polarized. They are also observed to be periodic, though this due to the pulsar’s rotation and not intrinsic to the emission. We do not have a single model that can clearly generate such emission, but we do have strong phenomenological understanding of pulsars. This allows us to create a working model for pulsar emission, without a full understanding of the theory of pulsar emission. This model, known as the lighthouse model, is conceptually straightforward, but allows for the observed properties of pulsar emission. In the lighthouse model, we envision the pulsar as a spinning magnetic dipole in the shape of a sphere with a beam of light extending from its magnetic axis. As the pulsar spins, the beam is periodically visible to the observer. This periodicity is equal to the the spin period of the pulsar.

It is generally accepted that pulsar emission can be modeled as a conical beam centered on the pulsar’s magnetic axis; see Radhakrishnan & Cooke (1969); Komesaroff (1970). Rankin (1983) expanded on this model by recognizing the importance of emission not just from this cone but also from a quasi-axial core as well. Lyne & Manchester (1988) further clarified that this emission should not be considered as two fully discrete processes but as a continuum, gradually changing emission characteristics between the core and the cone.

In examining pulsars, one of the most important quantities we observe is the polarization position angle (PPA;  $\Psi$ ). The PPA is observed to have an S-shaped sweep, which can be understood by considering the linearly polarized emission to be in a plane determined by the direction of the magnetic field at the point of emission. As the beam sweeps across the observer’s field of view, the PPA varies

slowly at profile edges and rapidly at the profile center. This is called the rotating vector model (RVM) and predicts the source's PPA will follow the pattern

$$\tan(\Psi - \Psi_0) = \frac{\sin \alpha \sin(\phi - \phi_0)}{\sin(\alpha + \beta) \cos \alpha - \cos(\alpha + \beta) \sin \alpha \cos(\phi - \phi_0)} \quad (1.1)$$

where  $\phi$  is rotational phase,  $\alpha$  is the magnetic inclination angle and  $\beta$  is the impact parameter.

One of the properties of a pulsar which the RVM model can explain is the linear polarization behavior. Pulsars are highly polarized; they are primarily linearly polarized but can also emit circularly polarized signals. Pulsar polarization is generally quantified using the four Stokes parameters, which together form the Stokes vector,  $\mathbf{S} = (IQUV)$ . Stokes parameter  $I$  is the total intensity,  $V$  is circular polarization, linear polarization is  $L = \sqrt{Q^2 + U^2}$ , and the position angle is defined as  $\Psi = \frac{1}{2} \arctan\left(\frac{Q}{U}\right)$ . It is important to note that the definitions used in pulsar astronomy are slightly atypical; Everett & Weisberg (2001) discuss conversion to the IAU convention.

Pulsars rotation is powered by magnetism, and pulsar behaviour is determined by consequences of magnetism and of basic rotational dynamics. The systems can be approximated as isolated rotating magnetic dipoles, but also as rotating solid spheres.

This allows us to relate the expected pulsar rotation period and its rate of change to loss of rotational energy using classical mechanics:

$$\dot{E} \equiv -\frac{dE_{\text{rot}}}{dt} = -\frac{d(I\Omega^2/2)}{dt} = -I\Omega\dot{\Omega} = 4\pi^2 I \dot{P} P^{-3}, \quad (1.2)$$

where  $\Omega$  is the rotational angular frequency ( $2\pi/P$ ) not solid angle and moment of inertia  $I$ . The moment of inertia for pulsars have not yet been measured; the canonical value is  $I = 10^{45} \text{ g/cm}^2$ .  $\dot{E}$  is known as the spin down luminosity and provides the total power output of a neutron star based on rotation (Lorimer & Kramer, 2005).

Our understanding of classical magnetic fields allows us to also determine the rate of change of the period, by equating the radiation power, as given by e.g.,

Jackson (1962) to the pulsar's power:

$$\dot{E}_{\text{dipole}} = \frac{2}{3c^2} |\mathbf{m}|^2 \Omega^4 \sin^2 \alpha. \quad (1.3)$$

where  $\alpha$  is the angle between the magnetic moment and spin axis and  $|\mathbf{m}|$  is the dipole magnetic moment.

We can combine Equations 1.2 and 1.3 and re-arrange to equate

$$\dot{\Omega} = - \left( \frac{2|\mathbf{m}|^2 \sin^2 \alpha}{3Ic^3} \right) \Omega^3, \quad (1.4)$$

and converting from angular frequency to rotational frequency pulsar  $\nu = 1/P$ ,

$$\dot{\nu} = -K\nu^n, \quad (1.5)$$

where  $K$  is constant and  $n$  is known as the braking index. In the case that the dipole model is entirely correct for the pulsar,  $n = 3$ , but external factors can alter it. If we are able to measure the second derivative of rotational frequency (possible for some systems, but not all), we can determine the braking index by differentiating Equation 1.5:  $n = \nu\ddot{\nu}/\dot{\nu}^2$ .

If we write Equation 1.5 in terms of pulsar period  $P$  and integrate the resulting differential equation, we can obtain the age of the pulsar

$$T = \frac{P}{(n-1)\dot{P}} \left[ 1 - \left( \frac{P_0}{P} \right)^{n-1} \right], \quad (1.6)$$

where  $P_0$  is the initial spin period of the pulsar when it is born. As pulsars are “spinning-down” – decreasing in spin frequency – we can assume that the pulsar birth period was much larger than the current period and simplify this expression to the characteristic age  $\tau_c$  of the pulsar

$$\tau_c \equiv \frac{P}{2\dot{P}} \approx 15.8 \text{MYr} \left( \frac{P}{\text{s}} \right) \left( \frac{\dot{P}}{10^{-15}} \right)^{-1}. \quad (1.7)$$

This method for estimating the birth period is highly approximate, and there is some tension between our understanding of supernovae and our simple model for

pulsar behaviour. The only pulsar ages we can know with certainty are those associated with historical supernovae (Lorimer & Kramer, 2005).

The last parameter we can easily determine from such simple manipulation of the pulsar’s period and period derivative is the magnetic field strength of the neutron star. From Equation 1.3 and the fact that  $B \approx |\mathbf{m}|/r^3$ , we determine

$$B_s \equiv B(r = R) = \sqrt{\frac{3c^3}{8\pi^2} \frac{I}{R^6 \sin^2 \alpha} P\dot{P}}, \quad (1.8)$$

which, assuming canonical  $I = 10^{45}$  g cm<sup>2</sup>, canonical radius  $R = 10$  km, and  $\alpha = 90^\circ$ , gives us

$$B_s = 3.2 \times 10^{19} \text{G} \sqrt{P\dot{P}} \approx 10^{12} \text{G} \left( \frac{\dot{P}}{10^{-15}} \right)^{1/2} \left( \frac{P}{\text{s}} \right)^{1/2}. \quad (1.9)$$

Much of this basic information about pulsars can be understood by looking at a  $P - \dot{P}$  diagram, such as that in Figure 1.1. Most isolated pulsars (marked in Figure 1.1 by green circles) hover between  $P = 10^{-1}$  s and  $10^1$  s in period and  $\dot{P} = 10^{-14} - 10^{-18}$  s · s<sup>-1</sup>, resulting in a characteristic magnetic field of  $B \sim 10^{12}$  G. However, there are also clear clusters near the top-right of the plot (long periods, large period derivative, high magnetic field) and near the bottom-left (very short periods, small period derivative). These two special cases, magnetars and millisecond pulsars respectively, are worthy of special attention.

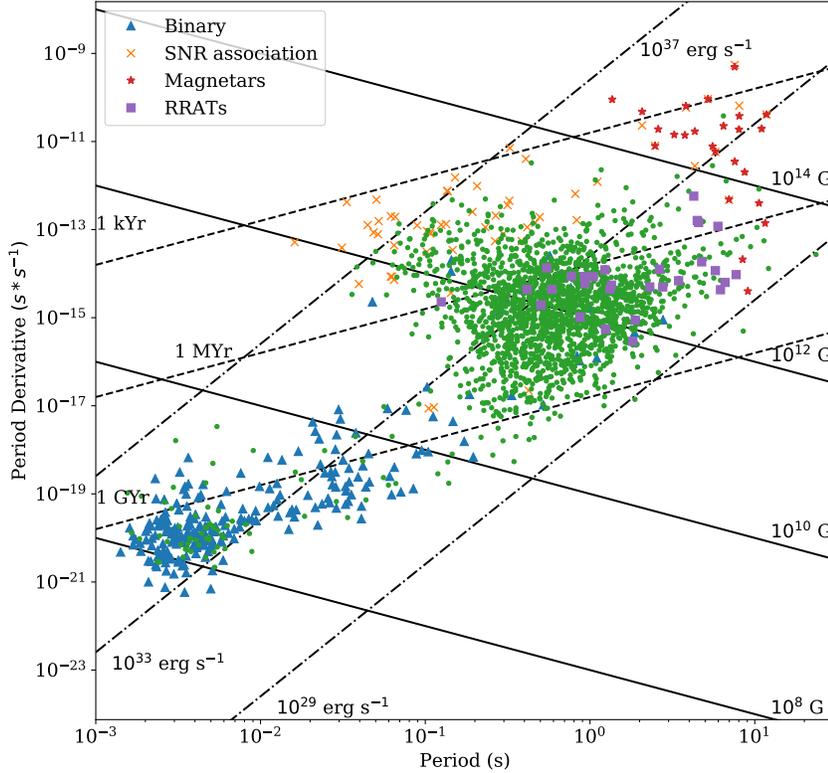
Millisecond pulsars (MSPs) are, as the name suggests, pulsars with millisecond to tens of milliseconds long periods. These periods are formed through a process known as recycling. These pulsars are or at one time were in a binary system with a star that has evolved off of the main sequence and overflowed its Roche Lobe. The pulsar siphons material from the Roche Lobe of the companion and been “spun-up,” increasing its angular velocity and decreasing its spin period (Radhakrishnan & Srinivasan, 1982; Alpar et al., 1982). One of the fastest spinning millisecond pulsars was the first discovered, PSR B1937+21, which is featured in Chapter 4 (Backer et al., 1982). Though not all MSPs are currently in binary systems, many are, accounting for the high prevalence of binary pulsars (blue triangles) in the lower-left corner of Figure 1.1.

Another special case among pulsars are magnetars. Magnetars are young pulsars with very high magnetic fields:  $B \sim 10^{13} - 10^{15}$  G (Kaspi & Beloborodov, 2017). Some magnetars are observable as radio pulsars, but many are observed primarily or exclusively at higher frequencies, as Soft Gamma Repeaters (SGRs) and Anomalous X-Ray Pulsars (AXPs). These sources were first observed in the late 1970s and early 1980s (Mazets et al., 1979a,b; Gregory & Fahlman, 1980), and the theory that they were very highly magnetized neutron stars emerged in the 1990s (Duncan & Thompson, 1992; Paczynski, 1992). The connection by Duncan & Thompson (1996) between the X-Ray sources and gamma ray sources proved to be critical in developing our understanding of magnetars as sources which emit across the electromagnetic spectrum. There are now about 30 known magnetars, and they are of particular interest due to their possible connection to FRBs. (See Section 1.3.) The McGill Magnetar Catalog (Olausen & Kaspi, 2014) maintains a record of all known magnetars.

Pulsar observations, including timing observations, are now routinely observed at both radio and X-ray frequencies. Pulsars can be readily observed with most radio telescopes, as they do not require highly precise angular resolution. Pulsar spectra actually lead to stronger emission at lower frequencies, but scattering can blur signals at these lower frequencies. Additionally, dispersion due to plasma in the interstellar medium (ISM) induces frequency dependent delays in the signal, such that signal arrives in higher frequency portions of the observing band before it arrives in the lower frequency portion of the band. The magnitude of this dispersion is known as dispersion measure (DM) and is the column density of electrons between the source and observer. Details of the procedure to search for new pulsars and to build timing models for known pulsars are discussed in Chapter 2.

## 1.2 Intermittent Pulsars, Nulling Pulsars, and RRATs

In general, pulsars emit consistently, and we can expect to see a signal from them corresponding to each pulse period. This is however, not universal. There is a continuum of intermittency for pulsars and pulsar-like sources. These distinctions are observational, determined by what we see to be occurring in our observations as opposed to being determined by theoretical predictions. (This is common for



**Figure 1.1:** A period vs. period-derivative plot, showing all known pulsars.

This figure provides a high level view of the pulsar population, highlighting key sub-populations such as magnetars and MSPs. Isolated pulsars are represented by green dots, binary pulsars are blue triangles, pulsars with known supernova remnant associations are orange x's, magnetars are red stars, and rotating radio transients are purple squares. The solid black lines are lines of constant magnetic field, from  $10^8$  to  $10^{14}$  G. The dashed lines are characteristic pulsar ages, from 1 kYr to 1 GYr. The dash-dot lines are  $\dot{E}$  estimates at  $10^{29}$ ,  $10^{33}$  and  $10^{37}$  erg  $s^{-1}$ . Pulsar data are from the ATNF Pulsar Database (Manchester et al., 2016) (<https://www.atnf.csiro.au/people/pulsar/psrcat/>), magnetar data are from the McGill Magnetar Catalog (Olausen & Kaspi, 2014), and RRAT data are from the RRATalog (<http://astro.phys.wvu.edu/rratalog/>).

pulsars and radio transients more broadly.)

The first and least intermittent category is nulling pulsars. Nulling pulsars appear as typical pulsars but within an observation, the pulsar appears to drop out or disappear for one or a few pulse periods. This phenomenon was first noted in Backer (1970), and has since been studied intermittently.<sup>2</sup> CHIME/Pulsar recently published a paper on this phenomenon and mode-changing (Ng et al., 2020). Mode-changing is a possibly related phenomenon where the pulse shape or polarization of the source appears to change during observations, often from one pulse to the next. This is somewhat common but can be challenging to observe; high cadence observations increases our likelihood of observational success.

Additionally, some pulsars emit continuously during an observation but may be seen to “disappear” for days, weeks, or even months at a time. Intermittent pulsars go dark for sustained periods, but when they are present, they behave as normal pulsars (Lyne, 2009; Kramer et al., 2006a). These sources can be “folded,” a process where we generally combine multiple pulses from a single source. An individual observation of an intermittent pulsar could not be distinguished from a non-intermittent pulsar. The exact mechanism that triggers these disappearances varies between pulsars; some may be binary interactions and some are known to be magnetospheric changes.

This is in contrast to rotating radio transients (RRATs), sources which can be observed only via single pulses, first discovered in the Parkes Multibeam Pulsar Survey (McLaughlin et al., 2006; Keane et al., 2011). RRATs appear as single pulses from pulsars and have a measurable underlying periodicity consistent with rotation. RRAT periods are hundreds of milliseconds to seconds, like pulsars, but the source does not emit every period. This means that we must analyze data entirely using single pulses, not using any folded data. Despite the observational challenges of detecting single pulses, RRATs have now been found in many pulsar surveys; see (e.g. Cui et al., 2017; Karako-Argaman et al., 2015; Patel et al., 2018).

With the exception of nulling, these are relatively new phenomena, and studies have been somewhat limited. In Chapter 5, we will discuss further the potential of CHIME/FRB and CHIME/Pulsar to improve our understanding of intermittent

---

<sup>2</sup>Pun, believe it or not, unintended.

and especially RRAT sources. High cadence search mode observations will allow us to deepen our understanding of known RRATs, and CHIME/FRB provides a straightforward mechanism for passively discovering new RRATs.

### 1.3 Fast Radio Bursts

At their most basic level, FRBs are exactly as their name says: fast, radio, bursts. They are fast: the last for only milliseconds in time. They are radio signals: they have been detected now at frequencies from 0.1 – 8 GHz (although most have been detected either at the CHIME band or at L-band<sup>3</sup>). FRBs are also known to be extragalactic, as their DMs are larger than the expected maximum Galactic DM, based on estimates of the amount of plasma in the Milky Way (Cordes & Lazio, 2002; Yao et al., 2017). The recent discovery of FRB-like emission from SGR 1935+2154 confirms the existence of FRBs with high energy counterparts (CHIME/FRB Collaboration, 2020a; Bochenek et al., 2020b). These sources are bursts; they are not continuous sources of emission, and so far only a small subset have been observed to have periodic emission.

Until 2015, each source had only been observed to emit a single burst. Work by (Spitler et al., 2016) showed repeat bursts from FRB 121102.<sup>4</sup> <sup>5</sup> The discovery of repetition was a major boon to the field, as it allowed for the addition of follow-up studies. These follow-up studies allowed FRB 20121102 to become the first localized FRB. It is located in dwarf elliptical galaxy at redshift  $z = 0.19273(8)$  and its association with a starforming region in that galaxy (Chatterjee et al., 2017; Tendulkar et al., 2017; Bassa et al., 2017). Subsequent to the localization of FRB 20121102, several more FRBs have been localized. FRB 20180916B has been localized to Sloan Digital Sky Survey galaxy SDSS J015800.28+654253.0 using the European VLBI Network Marcote et al. (2020). This is a massive spiral galaxy,

---

<sup>3</sup>1-2 GHz

<sup>4</sup>FRB naming convention has been a matter of substantial debate, but a consensus is emerging that a simple scheme following the example of supernovae is appropriate: The name of the first FRB on a given day is thus FRB YYMMDDA, the second is FRB YYMMDDB, and so on. Older FRBs often do not receive the A/B/C suffix as multiple FRB detections in a day has only been enabled by recent improved systems.

<sup>5</sup>It is worth noting that the field’s conventions surrounding burst vs. source remain an evolving and somewhat awkward linguistic tangle.

similar to the Milky Way, at redshift  $z = 0.0337 \pm 0.0002$ . Recently, real-time localization has become possible with ASKAP, allowing the realtime localization of FRBs without repeat detections for the first time (Bhandari et al., 2020).

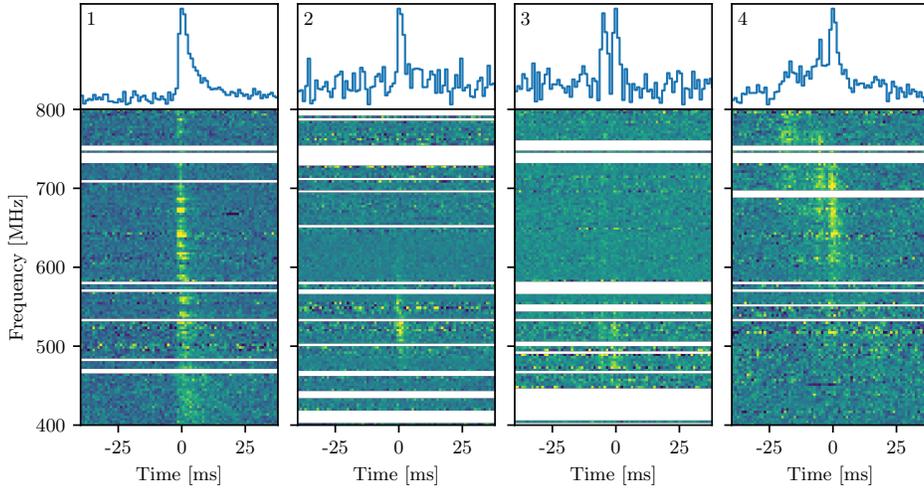
In 2020, FRB 20121102 and FRB 20180916B were both observed have periodic repetition. FRB 20180916B experiences a  $16.35 \pm 0.18$  day activity cycle, and FRB 20121102 experiences a  $161 \pm 5$  day activity cycle (CHIME/FRB Collaboration, 2020b; Rajwade et al., 2020; Cruces et al., 2021). These periodicities are too long to be spin periods like in pulsars or RRATs. These values could be consistent with binary periods, and repeating FRBs could be experiencing interactions with material in a binary system such as occurs for some intermittent pulsars. This does not preclude the possibility of underlying short period spin periodicity, as pulsars demonstrate both can occur in the same system.

Though we are beginning to understand FRBs as a population (work that will undoubtedly be advanced by the just-released CHIME/FRB catalog; CHIME/FRB Collaboration (2021)), there are certain properties that remain particularly intriguing. One is the variety of morphologies demonstrated by FRBs. The first FRBs discovered were all broadband, with limited to no time dependent structure. As our catalog has grown, we have come to realize that there are several possible morphologies for FRBs. In Pleunis et al. (2021a), we classify the bursts in the CHIME/FRB catalog into four basic morphological categories: simple broadband bursts, simple narrowband bursts, complex broadband bursts (often multi-peak), and complex narrow-band bursts (especially downward-drifting bursts). An example of each of these four types are shown in Figure 1.2. Simple broadband bursts make up about 60% of FRBs in the CHIME/FRB catalog, simple narrow-band bursts make up about 30% of FRBs, and each of the complex types make up about 5% of bursts.

The fourth category, complex narrowband or downward-drifting bursts, has sparked substantial interest. It was first discovered in repeat bursts from FRB 20121102 (Hessels et al., 2019). It is commonly referred to as the “sad trombone” structure, as it resembles the classic sound effect when sonified.<sup>6</sup> This structure cannot be explained by DM errors: the segments are piecewise. They do not move

---

<sup>6</sup>This term was coined by journalist Lisa Grossman in 2019. Of course, there are other ways to describe this sound, but the sad trombone has gained purchase in the community.



**Figure 1.2:** Dynamic spectra for four CHIME/FRB bursts, demonstrating the four major FRB morphology archetypes. From left to right, FRB 20190527C displays simple wideband structure, FRB 20190515D displays simple narrowband structure, FRB 20181117B displays complex wideband structure, and FRB 20190117A displays complex narrowband structure (Pleunis et al., 2021a).

continuously downwards but jump – each segment is properly described by the global DM and does not show signs of over-dispersion or under-dispersion. Indeed, if we were to decrease our timespan so that we isolated each sub-burst, we would no longer see this structure. This complex structure has been observed on timescales as short as  $30 \mu\text{s}$ , but it is unknown if that is the absolute limit or if we are constrained by our resolution. It is also unknown if some apparently simple bursts contain complex structure at timescales too small to be observed by search instruments such as CHIME/FRB.

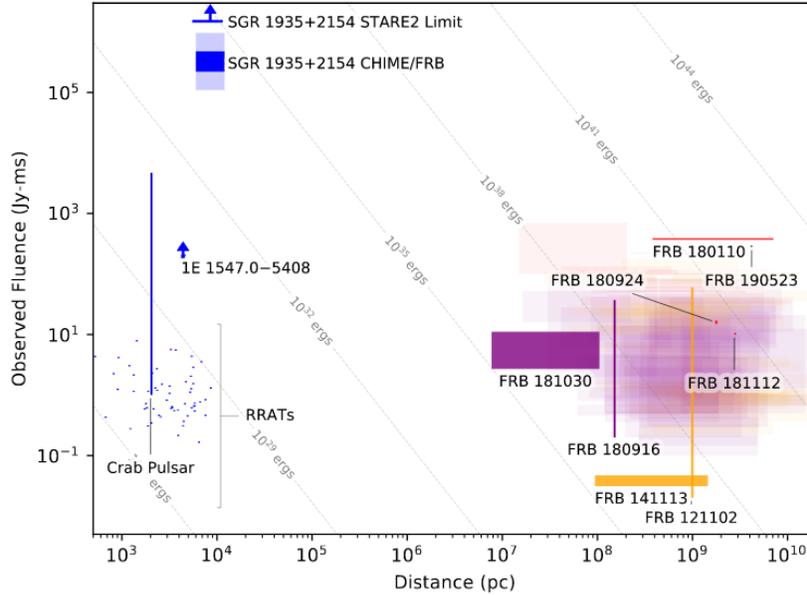
Though we are not certain of what the correct FRB model is, we do know some requirements it must meet. FRBs must broadly be high energy radio transients with approximately millisecond duration. We follow Lu & Kumar (2018) in insisting that FRB models must include bursts with certain parameters like luminosity of  $L_{\text{iso}} \sim 10^{43} \text{erg/s}$ , energy of  $E \sim 10^{40} \text{erg}$ , brightness temperature  $T_b \sim 10^{35} \text{K}$ , and

duration  $t \sim 1$  ms. The source object must be able to exist in a variety of galaxies including both dwarf spheroidal galaxies and Milky Way-like spiral galaxies, based on successful localizations to date (Tendulkar et al., 2017; Marcote et al., 2020; Bhandari et al., 2020). These properties are most consistent with compact objects, such as neutron stars, white dwarfs, or black holes.

A myriad of models have been proposed to explain FRBs; Platts et al. (2019) is one systematic attempt to document the many proposed theories. We will not attempt to summarize all proposed FRB models here, but to provide a flavour of the possibilities. Catastrophic models were among the first proposed, suggesting the FRBs could arise from the collapse of a neutron star (e.g., Zhang, 2016) or from compact object mergers (e.g., Mingarelli et al., 2015; Totani, 2013). Most such models are not viable for repeating FRBs, although there are some exceptions which attempt to include repetition in a merger model, such as Yamasaki et al. (2018). Other models appealed to pulsars, such as (Cordes & Chatterjee, 2019; Katz, 2017), or supernova (e.g., Piro, 2016). Others still went further afield, investigating the possibility of exotic phenomena like cosmic strings (e.g., Ye et al., 2017).

Increasingly, though, it appears that the best models are those invoking magnetars. There are a variety of magnetar models, (e.g., Pen & Connor, 2015; Popov & Postnov, 2010; Kumar et al., 2017; Beloborodov, 2017), so even constraining FRBs to be a consequence of magnetars does not entirely answer the question of what causes FRBs. The case for magnetars and the status of magnetar models is summarized in Margalit et al. (2020): FRBs can have large linear polarization and high RM, FRBs and magnetars both tend occur in star-forming regions, and the general repetition properties are consistent between the two are consistent. There are also sufficient magnetars to match the all-sky FRB rates (Nicholl et al., 2017).

The most powerful argument in favor of magnetar models, however, is the confirmed correspondence between an FRB and a magnetar. In April 2020, CHIME/FRB and STARE2 both detected an FRB from SGR 1935+2154 (CHIME/FRB Collaboration, 2020a; Bochenek et al., 2020b). SGR 1935+2154 is a known magnetar within the Milky Way, and it now appears that it emits FRBs (or at least FRB-like signals). CHIME/FRB results suggest that the SGR 1935+2154 burst energy was  $\sim 3 \times 10^{34}$  erg. This is three orders of magnitude larger than the brightest



**Figure 1.3:** SGR 1935+2154 in context, with pulsar and FRB bursts (CHIME/FRB Collaboration, 2020a). The x-axis is distance from Earth in parsec, the y-axis is fluence. Grey diagonals show assorted energy thresholds in ergs. CHIME/FRB measured the burst energy of SGR 1935+2154 as  $\sim 3 \times 10^{34}$  erg, which bridges the energy gap between Galactic magnetars and FRBs.

known radio-emitting magnetar bursts, but still lower than the energy of many known FRBs (CHIME/FRB Collaboration, 2020a). Figure 1.3 demonstrates this point. It is unclear if magnetars compose all FRBs or just a sub-population, but it does appear that magnetars are a source of FRBs.

## 1.4 Understanding Repeating Fast Radio Bursts in Depth

As repeating FRBs can be observed more than once, they are the ideal sources to study in detail. This has been particularly true of FRB 20180916B and FRB 20121102, both of which have proved relatively easy to detect for many telescopes. For FRB 20121102, see (Law et al., 2017; Gajjar et al., 2018; Michilli et al., 2018a;

Spitler et al., 2018; Josephy et al., 2019; Gourdji et al., 2019; Di et al., 2019; Caleb et al., 2020), and for FRB 20180916B, see (CHIME/FRB Collaboration, 2019; Marcote et al., 2020; CHIME/FRB Collaboration, 2020b; Chawla et al., 2020; Aggarwal et al., 2020; Pilia et al., 2020; Pleunis et al., 2021b). However, such study need not be limited to these sources. In Chapter 6, we present an effort to obtain detailed observations of additional repeating sources. This work and work by James et al. (2020) suggest that if all FRBs repeat, our most observed repeating FRBs may not be indicative of typical repeating FRBs.

Though there are many interesting aspects to repeating FRBs, recent results from the CHIME/FRB catalog, which is the most comprehensive FRB catalog ever assembled, suggest that repeating FRBs are largely like non-repeating FRBs (CHIME/FRB Collaboration, 2021; Pleunis et al., 2021a). They appear to share the same DM distribution and the same sky location distribution. However, there are two areas where repeating FRBs may diverge from non-repeating FRBs. The first is in their morphology, where repeating FRBs tend to have more complex time-frequency structure featuring sub-bursts and narrower bandwidth. The most obvious area in which repeating FRBs differ from non-repeating FRBs is, of course, repetition. Repetition has proved challenging to characterize for many bursts (including those presented in Chapter 6). As a community, FRB astronomers are also grappling with questions about the universality of FRBs: do all FRBs repeat or are there two populations of sources which appear very similar?

#### **1.4.1 FRB Morphology and Repetition**

The sad trombone structure discussed in Section 1.3 has become one of the signatures of FRB astronomy. It is in an absolute sense rare, representing only about 5% of bursts in the first CHIME/FRB catalog, but it is over-represented among repeaters (Pleunis et al., 2021a). It was first discovered in FRB 20121102, and later seen early in observations of FRB180814A and FRB 20180916B (Hessels et al., 2019; CHIME/FRB Collaboration et al., 2019a; CHIME/FRB Collaboration, 2019).

Due to its prevalence in repeating FRBs, downward-drifting frequency structure has become a sign of repetition. Not all repeating sources demonstrate this

structure, and there is often variation within bursts from a source – some will show this structure and some will not. This structure is also visible in bursts from sources which are not known to repeat, as in Chapter 6.

If there are distinct populations of FRBs, some repeating and some non-repeating, it is possible that this structure is unique to repeaters. This would be convenient both theoretically and observationally, as it would allow us to instantly know that a source with this structure is a repeater, and it would allow us to remove this structure from the list of features needed for theoretical models of single burst sources. However, it is premature to suggest that this is the case; many sources with morphologically complex bursts also emit simple bursts, some known repeating FRBs are entirely morphologically simple, and some sources with complex bursts have to date not been observed to repeat. The strength of the link between downward-drifting, sad-trombone structure and repetition is an open question.

Besides the much discussed and exciting trombone burst structures, we have also found that repeaters have a tendency to narrower bandwidth (Pleunis et al., 2021a). The first FRBs discovered tended to be broadband, but as the population has grown, we have found that a significant number have only a narrow emission bandwidth. This does not appear to be a selection effect (Pleunis et al., 2021a). This narrow emission bandwidth is more common amongst repeaters, but as in the case of the sad trombone structure, the strength of the link is unclear. Many narrowband FRBs have not been observed to repeat and many repeating FRBs have been observed with broad bandwidth. Interestingly, sources with narrower emission bandwidth are not necessarily confined to the specific narrow bandwidth at which they are initially observed. Other bursts from the same source may have a larger emission bandwidth or may demonstrate a different narrow emission bandwidth. This does represent an observational challenge and demonstrates the value of broad observing bandwidths – a source which bursts only between e.g., 400 and 430 MHz will not be observed by a 500–700 MHz instrument.

Likely the most thorough treatment of FRB morphology, covering both repeating and non-repeating FRBs is the recent work is Pleunis et al. (2021a) and the related Ph.D thesis. These complex structures are certainly an important part of understanding repeating FRBs and may prove to be useful in furthering our understanding whether all FRBs repeat or not. They are a major motivation for the work

presented in Chapter 6.

### 1.4.2 FRB Repetition Properties

From the first repeat detections of FRB 121102, it was clear that FRB repetition is not purely random: it is much more likely that a burst will be observed if a burst has recently been seen (Spitler et al., 2016; Scholz et al., 2016). Clustering of this type is dramatically illustrated by FRB 20190116A, discussed in more detail in Chapter 6 (CHIME/FRB Collaboration, 2019). This was observed twice within the same CHIME/FRB observation (bursts 45 seconds apart) at the time of its initial discovery but has not been observed since.

After the discovery of this clustered bursting, several models were proposed for understanding repetition, see, (e.g., Connor et al., 2016; Lawrence et al., 2017), but the most prominent is the Weibull proposed by Oppermann et al. (2018). The Weibull distribution is an extension of the Poisson distribution which allows for clustering. In their initial paper, they derive a framework for using this method and apply it to what was the FRB 20121102 dataset in 2018. This work found that pure Poisson repetition is disfavoured. Since then, this extension has become a common framework for analyzing FRB repetition rate; we apply it in Chapter 6.

We now know that at least some repeaters are periodic, but we do not know if this is a universal property or only a property of certain sources. Reliably determining a period for a repeater requires a large number of bursts, so it is unsurprising that the two sources for which we have observed the most bursts are the ones for which we have a periodicity. This highlights that we should continue to study more repeating FRBs in depth if we want to better understand repetition.

One interesting question in the study of FRBs is whether all repeat. The community remains divided on the issue, and there are serious arguments in both directions.

Perhaps the simplest argument in favour of a single population derives from Occam's Razor: it is philosophically desirable that we not have two populations for nearly identical radio emission. However, there is an obvious precedent in astrophysics for such a two-population solution. Short and long gamma ray bursts (GRBs) come from different sources, and are actually very different in physical meaning

though the signals are similar observationally (Nakar, 2007; Woosley & Bloom, 2006).

Statistically significant variations in morphology such as those we report in Pleunis et al. (2021a) are among the strongest arguments for the possibility of two populations. On the other hand, there are possible explanations for this variation in a single-population model. For example, Connor et al. (2020) propose that in certain beaming angle configurations, the morphological and bandwidth discrepancies could be explained without the need to propose two populations.

The morphological argument for two populations is also challenged by the analogy to pulsars. Pulsars are universally agreed to be all one type of object (though with several sub-populations), yet they display a wide variety of morphologies. Indeed, thanks to mode-changing and nulling phenomena, an individual pulsar can itself display a number of different morphologies. Though of course, FRBs are not pulsars, much of our understanding of FRBs is tied up with our understanding of pulsars as they share an obvious overlap in being high energy, short duration radio sources. If, like pulsars, FRBs do emerge from neutron stars, it is entirely possible we could see wide variation in their properties without needing two separate populations.

We are unlikely to determine an immediate answer to this question; time and long-term observations are likely to be key factors in understanding repetition. For example, if we see a plateau in the number of new repeaters detected with CHIME/FRB while not observing a comparable plateau in new FRBs, we might suggest that we had seen all of the existing repeaters above a certain instrumental threshold. On the other hand, if we continue to find new repeating sources within the CHIME/FRB dataset, this provides evidence that repeaters are simply the frequently-repeating or particularly easy to observe tail of a fully repeating population. A diverse set of follow-up observations that could probe potential low-luminosity tails of repetition may prove valuable in answering this question.

## 1.5 Applications of Millisecond Pulsar Timing & Pulsar Timing Arrays

Although all pulsars are interesting, there are particularly exciting applications for MSPs. They are among the most accurate natural clocks in the universe, with equivalent precision to atomic clocks, and are very stable sources. This allows us to create very high precision timing solutions describing their rotation and (for binaries) orbital motion which can in turn be used to test theories of gravity, learn about neutron star properties, and search for low-frequency gravitational waves.

Timing of MSPs in binary systems enables us to learn more about the neutron star itself, by measuring the neutron star mass precisely. For highly-inclined binary systems, we can measure the relativistic Shapiro delay (Shapiro, 1964). This enables us to precisely determine the mass of the pulsar and its companion (Lorimer & Kramer, 2005). These mass measurements in turn enable us to constrain models for the neutron star equation of state.

As very dense objects rotating rapidly (and often in high velocity binary systems), MSPs are ideal laboratories for testing general relativity in the strong-field regime; see Stairs (2003) for a detailed review. Binary neutron star systems, including the Hulse-Taylor binary (PSR B1913+16) (Hulse & Taylor, 1975), the double pulsar (PSR J0737-3039) (Burgay et al., 2003; Lyne et al., 2004; Kramer et al., 2006b) and other double neutron stars like PSR J0453+1559 (Martinez et al., 2015) are particularly valuable for these applications.

One component of GR that pulsar timing can test is the strong equivalence principle, which specifies the equivalence of gravitational and inertial mass (Damour & Schaefer, 1991; Wex, 1995; Stairs, 2003; Archibald et al., 2018). Such tests have been conducted for several sources including the triple system PSR J0337+1715 (Archibald et al., 2018; Voisin et al., 2020) and well-known MSP PSR J1713+0747 (Zhu et al., 2015a, 2019). We can provide constraints on post-Newtonian parameters by modeling motion within the binary system and of the binary system relative to a preferred rest frame. MSPs can also constrain changes in the value of gravitational constant  $G$  by measuring the intrinsic component of the rate of change of the binary period (Stairs, 2003; Zhu et al., 2015a, 2019). The most famous pulsar test of gravity was the detection of gravitational waves in the Hulse-Taylor binary

system by measuring the spin-down of that binary (Hulse & Taylor, 1975).

A major application of pulsar timing is to the search for low frequency gravitational waves with pulsar timing arrays (PTAs), including the North American Nanohertz Observatory for Gravitational Waves (NANOGrav). PTAs (Detweiler, 1979; Foster & Backer, 1990) are created by timing a set of pulsars for an extended period of time. They derive their name from the idea that each pair of pulsars represents a new “arm” of an interferometer, loosely analogous to the arms of terrestrial gravitational wave detectors. NANOGrav now includes 70 pulsars and has been observing for more than 15 years. To search for gravitational waves with PTAs, we search for a spatial correlation among the pulsars with a quadrupolar signature of the form presented by Hellings & Downs (1983):

$$\alpha_{ab} = \frac{1 - \cos \zeta_{ab}}{2} \ln \left[ \frac{1 - \cos \zeta_{ab}}{2} \right] - \frac{1}{6} \frac{1 - \cos \zeta_{ab}}{2} + \frac{1}{3} \quad (1.10)$$

where  $\alpha$  is the arrival time correlation and  $\zeta_{ab}$  is the angular separation of pulsars  $a$  and  $b$ .

Gravitational waves propagate through spacetime, contracting and expanding it as they pass. From the perspective of the PTA, this contracts and expands the distance between the Earth and the pulsar of interest. In this way, PTAs operate similarly to laser interferometer gravitational wave detectors like LIGO or LISA. In those systems, such contractions and expansion alter the path length within the arms of an interferometer. In a PTA, the “arms” become Galactic scale: the distance between pairs of pulsars.

The primary target of nanohertz gravitational wave searches is the stochastic background of supermassive black hole binaries (Sesana et al., 2004; Burke-Spolaor et al., 2019). However, recent years have seen an increase in interest in continuous wave searches, which search for gravitational wave signals from individual supermassive black hole binaries (Ellis et al., 2012; Zhu et al., 2015b; Arzoumanian et al., 2021).

Since 2005, the NANOGrav has been engaged in a gradually expanding effort to search for such gravitational waves, publishing upper limits on the stochastic gravitational wave background with each dataset (Demorest et al., 2013; Arzoumanian et al., 2016, 2018a). In the 12.5 year data set, we see a new and potentially

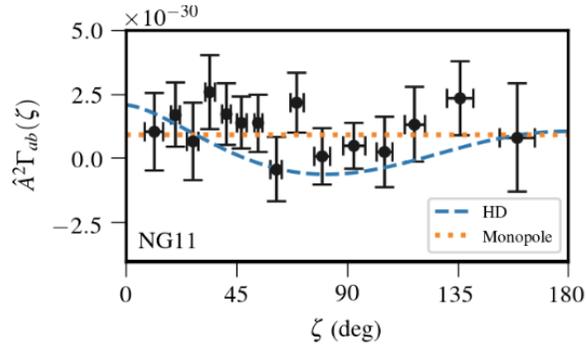
exciting signal. Previous datasets demonstrated a gradual but progressive tightening in the allowed range for stochastic gravitational wave background values. In contrast, the 12.5 year dataset allows a slightly less tight constraint on the upper limit, but provides strong evidence for a common spectrum red noise process (Arzoumanian et al., 2020). This alone is not a gravitational wave detection, but it may prove to be a precursor to such a detection.

The timescale for detecting a low-frequency gravitational wave is long. Unlike the compact binary coalescences observed by instruments like LIGO, which are observed in seconds, we would expect a gradual, growing increase in signal over the course of several years. Therefore, one potential early signature would be a common spectrum noise process among all pulsar pairs. To claim a detection, we would specifically require that this correlation be consistent with a Hellings-Downs spatial correlation. At the present, the NANOGrav result strongly suggests a common spectrum process but only slightly suggests a Hellings-Downs spatial correlation (Arzoumanian et al., 2020). Figure 1.4 shows the Hellings-Downs curve with the results from the 11 year and 12.5 year datasets (Arzoumanian et al., 2020). At present, then, we are in an intermediate state. We see what may be initial evidence of a detection, but cannot claim a true detection. This is a very exciting hint, but it is still possible that the correlated noise signal will be revealed to be unrelated to gravitational waves.

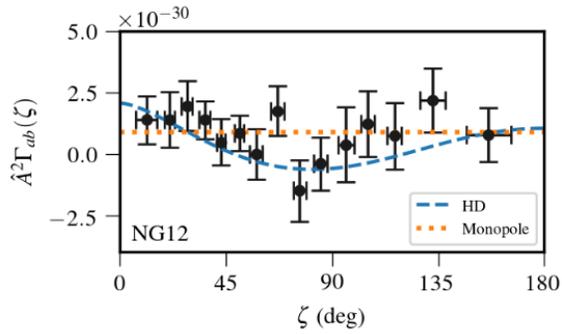
In addition to detecting low-frequency gravitational waves, PTAs are important as enable international collaboration; the European PTA (Desvignes et al., 2016), Parkes PTA (Kerr et al., 2020), NANOGrav and the Indian PTA (Joshi et al., 2018) together form the International Pulsar Timing Array (IPTA)(Perera et al., 2019).. The IPTA exists to foster collaboration between pulsar timing arrays, allowing the community to band together to the best possible dataset, incorporating the widest possible range of pulsars.

## 1.6 Outline of the Thesis

In this thesis, we present results from three projects, each incorporating data from CHIME/FRB, CHIME/Pulsar, or both. Though the science goals of the three projects span a large swath of pulsar and FRB astronomy, all projects build on



(a)



(b)

**Figure 1.4:** This figure from Arzoumanian et al. (2020) shows the average angular distribution of cross-correlated power in the NANOGrav dataset. Panel (a) shows the 11 year dataset; Panel (b) shows the 12.5 year dataset. The blue dashed line shows a Hellings-Downs correlation, the expected signature of a gravitational wave detection, and the orange dashed line shows a monopole distribution for comparison. Here,  $\hat{A}^2\Gamma_{ab}(\zeta)$  takes the place of  $\alpha$  as the cross-correlation. The points are binned; they represent the same number of bins for both panels, but the bins contain different numbers of pulsar pairs.

observations made with CHIME.

Chapters 2 and 3 present more technical introductory materials. Chapter 2 discusses modes of observation for pulsars and FRBs, the procedure employed in pulsar and repeating transient searches, and the procedure for constructing pulsar timing models. Chapter 3 provides brief introduction to CHIME, with a particular focus on the FRB detection and pulsar timing backends.

Chapter 4 presents initial results from a long-term project, focused on integrating daily CHIME/Pulsar observations of NANOGrav MSPs into the NANOGrav dataset. These early results demonstrate the utility of CHIME/Pulsar data in high-precision pulsar timing.

Chapter 5 presents the first seven Galactic sources discovered using single pulses from CHIME/FRB. This work, also presented in Good et al. (2021), incorporates both CHIME/FRB and CHIME/Pulsar data and represents an exciting new avenue for studying RRATs and detecting intermittent or highly nulled pulsars.

Chapter 6 presents the results of follow-up observations of CHIME FRBs with Arecibo Observatory. We studied two known repeating FRBs (FRB 20190116A and FRB 20190117A) and seven non-repeating FRBs with morphological structure suggestive of repetition. We did not detect repeat bursts from any of these sources, but present constraints on the repetition rate based on these observations.

Chapter 7 presents a brief reflection on the results presented in this thesis and looks forward to future work that can be conducted building on it.

## Chapter 2

# Observing Pulsars and Repeating Radio Transients

*I never am really satisfied that I understand anything; because, understand it well as I may, my comprehension can only be an infinitesimal fraction of all I want to understand about the many connections and relations which occur to me, how the matter in question was first thought of or arrived at, etc., etc. — Ada Lovelace*

### 2.1 Pulsar & FRB Observations

Chapters 4 to 6 present results from a variety of observations of pulsars and radio transients. It is important, therefore, to pause and understand the procedures for observing these sources. In this chapter, we briefly outline the types of pulsar observations, the effects of the interstellar medium on pulsar and FRB data, a method for finding new pulsars and radio transients, and the structure of pulsar timing models.

Within radio observations, pulsars are generally observed in one of two “modes”: search-mode and fold-mode. Pulsar searches make use of search mode (filterbank) data, which does not presume the properties of a potential pulsar. This mode is agnostic to the presence of a source, merely recording intensity data. Search-mode observations can detect not only new ordinary pulsars, but also single pulse sources such as RRATs and both repeating and non-repeating FRBs. They are particularly

valuable for studying any pulsar or pulsar-like object where single pulse properties are of interest. This includes rotating radio transients (RRATs) and intermittent pulsars as well as potential FRB sources. Search-mode data have been used both in following-up FRBs in Chapter 6 and in determining the properties of new pulsars and RRATs in Chapter 5.

Pulsar searches generally employ incoherent de-dispersion, not correcting for DM as the data are collected but correcting off-line during analysis. This necessitates high frequency resolution to enable accurate de-dispersion. Follow-up searches for single pulse sources such as RRATs and repeating FRBs, may employ coherent de-dispersion in a data collection mode known as coherent search mode. In coherent search mode, data are coherently de-dispersed to a given DM but otherwise behaves as search mode data. Coherent search mode data are employed in Chapter 6.

Search-mode files can be large, so data are often collected with only total intensity information, not full polarization information. This minimizes the search mode file's utility for some analyses, but is not a problem for initial detections of pulsars, RRATs, or FRBs. Additionally, time resolution may be lower for search-mode files than for fold-mode files; this is the case for CHIME/Pulsar data where search mode data have a time resolution of  $327 \mu\text{s}$  and record only total intensity (not polarization), while fold-mode data have a time resolution of  $2.56 \mu\text{s}$  and record full Stokes polarization information.

Fold-mode data in contrast requires prior knowledge of the pulsar's spin period (and is made more effective by each improvement in knowledge of the timing model), but allows known pulsars to be observed with greater precision and efficiency. A folded observation incorporates knowledge of the pulsar's spin period stack pulses as the signals arrive. Provided the period is correctly estimated, this results in a stronger signal in the observation by combining multiple pulse periods.

This method is used with some new pulsars in Chapter 5 and extensively with NANOGrav sources in Chapter 4. It not currently feasible to use this technique in FRB science, as FRBs have not yet been observed to have periodicity on the scale of a single observation. (Those FRBs with observed periodicities see active periods of a few days followed by inactive periods to days to months, not millisecond to second scale periodicity).

For many pulsars, some level of folding is required to observe a signal. In pulsar searches, this is part of the analysis process, conducted after the fact with search-mode data. However, once the period is determined, it can be used to collect folded data, improving our understanding of the new source and further confirming our period estimate. Chapter 5 employs fold-mode data for persistent sources with initial timing models determined from CHIME/Pulsar search-mode data.

FRBs can be observed using search-mode data, as was the case for the first FRB detections (Lorimer et al., 2007) and as is often the case for repeater follow-up. However, as new FRBs are unpredictable and isotropically distributed on the sky, it is preferable to use purpose-built, automated transient search systems; (see, e.g., Caleb et al., 2016; Law et al., 2018; James et al., 2019; CHIME/FRB Collaboration, 2018). We will discuss the CHIME/FRB system in more depth in Chapter 3.

## 2.2 Radio astronomy in the interstellar medium

Though in a macro sense, space is “empty,” the interstellar medium (ISM) is a cold, ionized plasma, full of gas, dust, and magnetic fields. This material in ISM interacts with the radiation from pulsars and FRBs in ways that must be understood to interpret our observations. In this section, we follow the derivations supplied by Lorimer & Kramer (2005).

The cold, ionized plasma of the ISM has a frequency dependent group velocity,

$$v_g = \frac{c}{n} = c \sqrt{1 - \left(\frac{f_p}{f}\right)^2}, \quad (2.1)$$

where  $f$  is the observing frequency,  $f_p$  is the plasma frequency in,  $c$  is the speed of light,  $n$  is the index of refraction, and

$$f_p = \sqrt{\frac{e^2 n_e}{\pi m_e}} \approx 8.5 \text{ kHz} \left(\frac{n_e}{\text{cm}^{-3}}\right)^{1/2} \quad (2.2)$$

and  $n_e$  is the number density of electrons.

### 2.2.1 Dispersion Measure

As the result of Equation 2.1 will be less than  $c$ , the speed of light in plasma will be less than the speed of light in a vacuum and there is a frequency-dependent time delay of the amount

$$t = \left( \int_0^d \frac{d\ell}{v_g} \right) - \frac{d}{c}, \quad (2.3)$$

where  $d$  is the distance to the source and the integral is along the line of sight to the source.

We substitute in Equation 2.1 and note  $f_p \ll f$ , so the time delay is

$$t = \frac{1}{c} \int_0^d \left[ 1 + \frac{f_p^2}{2f^2} \right] d\ell - \frac{d}{c} = \frac{e^2}{2\pi m_e c} \frac{\int_0^d n_e d\ell}{f^2} \quad (2.4)$$

We simplify this by defining the quantity dispersion measure as

$$DM = \int_0^d n_e d\ell \quad (2.5)$$

where  $n_e$  is the number density of electrons and  $d$  is again the distance to the source. If we know the expected column density of electrons on any given line of sight, we can estimate the distance to the source from the dispersion measure. Previous work does provide us maps of the Galactic DM (Yao et al., 2017; Cordes & Lazio, 2002). However, these maps are imprecise, so distance measurements determined in this manner are only accurate to a few tens of percent. This does not allow us to make precise distance measurements, but excess DM beyond the Galactic DM does play a key role in identifying that FRBs are extragalactic sources.

Practically, we are most concerned about the observed time delay between frequencies caused by this dispersion. The expected delay is

$$\Delta t \approx 4.15 \times 10^6 \text{ ms} \times (f_1^{-2} - f_2^{-2}) \times DM, \quad (2.6)$$

where frequency is measured in MHz and DM in  $\text{pc cm}^{-3}$ . If we know the DM of

the source before starting our observation, we can use this information to coherently de-disperse the incoming signals. Alternatively, if we do not know the DM a priori, we can use the observed time delay to calculate it.

Removing dispersion is an important step in processing both pulsar and FRB observations. There are two basic methods: incoherent de-dispersion and coherent de-dispersion. Though incoherent de-dispersion is the older method, both are still used, with incoherent de-dispersion generally more applicable to searching for new pulsars and FRBs while coherent de-dispersion is more applicable to observations of known pulsars and repetition from FRBs.

In incoherent de-dispersion, the observation’s bandwidth is divided into frequency channels, and the appropriate delay based on a known DM is applied to each channel. This requires high frequency resolution; CHIME/FRB uses a frequency resolution of 24.4 kHz. This method also has the advantage that it does not require the DM to be known prior to the observation being conducted. This is crucial for searches, as we cannot know the DM of an object that has not yet been discovered.

In pulsar searches and in small-scale FRB searches, incoherent de-dispersion is generally conducted post facto, as part of off-line analysis after an observation. However, large-scale FRB searching, like in CHIME/FRB, require real-time de-dispersion. Data volumes from such experiments are staggering, requiring us to save intensity data only for potential detections.

The algorithm used by CHIME/FRB, discussed in instrumental context in 3.4, is a tree de-dispersion algorithm. Tree de-dispersion was first proposed in Taylor (1974), and its great advantage over brute force de-dispersion is that it is an  $\mathcal{O}(N \log_2 N)$  operation. Tree de-dispersion differs from brute force de-dispersion in that it creates small clusters or “branches” of channels, building these “branches” progressively to determine the overall best DM for a source. Though still computationally expensive, this enables it to be run on CPUs, if implemented thoughtfully.

In contrast, when the target source’s DM is well known, we can establish much better observational precision, including for pulsar timing by applying coherent de-dispersion. In a coherent de-dispersion system, first proposed by Hankins & Rickett (1975), data are sampled and dispersive smearing is removed prior to detecting the pulsar.

Coherent de-dispersion depends on the concept that the dispersive effect of the ISM on the pulsar signal can be described by a transfer function  $H$ . While raw voltages are collected in the time domain, they can be converted to the frequency domain via a Fourier Transform.<sup>1</sup>

In the frequency domain, we can describe a signal centered at  $f_0$  with bandwidth  $\Delta f$  as

$$V(f_0 + f) = V_{int}(f_0 + f)H(f_0 + f), \quad (2.7)$$

where  $V$  is the Fourier transform of the measured voltage and  $V_{int}$  is the Fourier transform of the intrinsic voltage of the pulsar.

Dispersive delays can be modeled as a phase rotation dependent on frequency and distance to the source,  $\Delta\Psi = -kd$ , with  $k$  the frequency dependent wave number and  $d$  the distance to the source. The wave number is

$$k(f) = \frac{2\pi}{c}f\sqrt{1 - \frac{f_p^2}{f^2} \mp \frac{f_p^2 f_B}{f^3}}, \quad (2.8)$$

where  $f_p$  is the plasma frequency and  $f_B$  is the cyclotron frequency,

$$f_B = \frac{eB}{2\pi m_e c} \cong 3\text{MHz} \left( \frac{B_{\parallel}}{\text{G}} \right). \quad (2.9)$$

The “ $\mp$ ” prior to the final term in Equation 2.8 represents the delay between left (−) and right-hand (+) circular polarizations.

Via Taylor Expansion and substitutions, we eventually arrive at the result that

$$H(f_0 + f) = e^{+i \frac{2\pi D}{(f+f_0)f_0^2} DM f^2}, \quad (2.10)$$

where  $D$  is the dispersion constant,  $D = \frac{f_p^2}{2cn_e}$ . In practice, the CHIME FRB and Pulsar systems follow the recommendations of Manchester & Taylor (1977) and set  $D = 1/2.41 \times 10^{-4}$  (for frequency in MHz and time in seconds). This number remains fixed, even as the physical constants are redefined.

To coherently de-disperse a source, then, we Fourier Transform a voltage timestream,

---

<sup>1</sup>The last resort of scoundrels and the first resort of radio astronomers. Scoundrels are bad at radio astronomy.

multiply it by the inverse of  $H$  for each pulsar, and again Fourier Transform the data back to the time domain. This automatically provides us with a de-dispersed signal.

The idea for coherent de-dispersion is by no means new, and it was long used in specialized hardware based systems. However, it only became practical in software-based systems in the late 1990s. It is now the dominant mode of de-dispersion for fold-mode pulsar data collection, including by CHIME/Pulsar fold mode data, by all data used in Chapter 4's NANOGrav analyses, and for Arecibo followup observations presented in Chapter 6.

## 2.2.2 Rotation Measure

The presence of magnetic fields in the ISM also introduces a phase delay in our signal, again of the form  $\Psi = -kd$ .

The differential phase rotation between right and left hand circular polarizations is

$$\Delta\Psi_{\text{Faraday}} = \int_0^d (k_R - k_L) dl, \quad (2.11)$$

where  $k_L$  and  $k_R$  are the wave numbers for the left and right circularly polarized wave respectively. In the limit that  $f \gg f_p$  and  $f \gg f_B$ ,

$$\Delta\Psi_{\text{Faraday}} = \frac{e^3}{2\pi m_e^2 c^4} \int_0^d n_e B_{\parallel} dl. \quad (2.12)$$

It is more immediately applicable to write  $\Delta\Psi_{\text{Faraday}}$  in terms of the pulsar's polarization position angle, which is periodic in  $\pi$  instead of  $2\pi$

$$\Delta\Psi_{\text{PPA}} = \Delta\Psi_{\text{Faraday}}/2 \equiv \lambda^2 \times RM, \quad (2.13)$$

where  $\lambda$  is our observing wavelength and rotation measure,  $RM$ , is

$$RM = \frac{e^3}{2\pi m_e^2 c^4} \int_0^d n_e B_{\parallel} dl. \quad (2.14)$$

If we have measured both  $RM$  and  $DM$ , we have sufficient information to es-

timate the average magnetic field along the line of sight, weighted by the electron density

$$\langle B_{\parallel} \rangle = 1.23 \mu\text{G} \left( \frac{\text{RM}}{\text{rad m}^{-2}} \right) \left( \frac{\text{DM}}{\text{cm}^{-3} \text{pc}} \right) \quad (2.15)$$

We observe the phase delay  $\Delta\Psi_{\text{PPA}}$  and thus are able to extrapolate the RM and ultimately the approximate magnetic field between the source the observer. Though none of the analysis presented in this thesis makes direct use of RM measurements, such measurements are an important part of our observational understanding of pulsars and FRBs.

### 2.2.3 Scattering

Another important transmission effect, observed both in pulsars and FRBs, is scattering. This scattering is induced by inhomogeneities in the interstellar medium, which randomly broaden the observed signal.

Multi-path scattering is itself complex, but the overall effect can be represented as a one-sided exponential with a scattering timescale  $\tau_s$ . This scattering time is itself frequency dependent:

$$\tau_s = \frac{e^4}{4\pi^2 m_e^2} \frac{\Delta n_e^2}{a} d^2 f^{-4}. \quad (2.16)$$

The ultimate effect is that scattering is sharply dependent on frequency; we can expect lower frequency signals to be much more scattered than higher frequency signals. This is an important consideration for CHIME data, as 400-800 MHz is a relatively low frequency range and will lead to large amounts of scattering. We will briefly use this scaling relation in Chapter 6, to convert CHIME/FRB 600 MHz scattering times to the equivalent 327 MHz to determine sensitivity thresholds for our observations with the Arecibo Observatory 327 MHz receiver.

### 2.2.4 Scintillation

The final major transmission effect for pulsars is scintillation. Scintillation is caused by irregularities and turbulence in the ISM. It was first observed by Lyne & Rickett (1968) and explained by Scheuer (1968). The irregularities in the ISM

cause interference patterns within the observed pulsar signal, with some frequency regions within the observation having increased intensity and some having decreased intensity. Interference of this type is limited to waves where the phases differ by less than  $\sim 1$  radian. Phase is frequency dependent, so this sets a bandwidth limitation for bandwidth of

$$2\pi\Delta f\tau_s \sim 1 \quad (2.17)$$

or rather, substituting  $\tau_s$ ,

$$\Delta f \propto 1/\tau_s \propto f^4 \quad (2.18)$$

Though scintillation is observable with CHIME/Pulsar and CHIME/FRB, we do not address its effects further in this work.

## 2.3 Searching for Pulsars and Repeating Radio Transients

### 2.3.1 Pulsar & Transient Search Strategies

Pulsar searching is an important component of pulsar science. We are assisted by the existence of standardized software packages (particularly PRESTO) which allow us to follow standardized procedures (Ransom, 2011).

In a conventional pulsar search, we select an area of sky, either based on our belief that pulsars will be present (as in the case of a targeted search of a globular cluster) or based on instrumental suitability. Most searches are of the latter type, with relatively few focused on targeted observations, (e.g., Barr et al., 2013; Titus et al., 2019),

Pulsar search programs have been conducted by most major radio observatories. Important programs of the past 20 years include the Parkes Multibeam Pulsar Survey and High Time Resolution Universe Survey at the Parkes Telescope (Manchester et al., 2001; Keith et al., 2010), the Green Bank North Celestial Cap Survey at GBT (Stovall et al., 2014), and the Pulsar ALFA Survey and AO327 Drift Survey at Arecibo Observatory (Lazarus, 2012; Deneva et al., 2013). The Five Hundred Metre Aperture Synthesis Telescope (FAST) is a new observatory,

with great potential for pulsar discovery that is only starting to be realized (Han et al., 2021). In Chapter 5, we will discuss CHIME’s potential as an unconventional pulsar searching tool, but a formal slow pulsar search backend is also planned for CHIME/FRB. In the future, instruments like SKA will discover many, many new pulsars.

CHIME/Pulsar, as will be discussed in Chapter 5, employs an unusual strategy. We do not conduct an open-ended search, nor do we target a specific area of the sky. Instead, we find potential Galactic candidates from astrophysical events in CHIME/FRB data then conduct targeted observations of these locations. This method is adept at discovering highly intermittent sources, including RRATs, and the dispersion measure.

### **2.3.2 Finding and Verifying Pulsar Candidates**

There are several steps involved in searching data for pulsars and radio transients. First, we remove radio frequency interference (RFI) from the data, using a combination of automated and manual methods. Second, we de-disperse the data at many trial DMs to ascertain the correct DM for the source. Third, we search for an appropriate spin period, using one of several possible approaches. Fourth, if we find a period value, we fold data at the nominal period and iterate our timing model using pulsar timing software. Simultaneously with the period search, we can search for single pulses from pulsars or transients. This a critical step for transient searches, both RRATs and FRBs, and can be useful for pulsars as well. It is used in Chapters 5 and 6. However, not all pulsars are visible in single pulses, so this may not be useful when searching for conventional pulsars.

The software package PRESTO provides python routines to complete most of these components of the pulsar searching process. It is the standard starting point for modern pulsar searches, but it is also the case that we can make use of additional information and software packages throughout the process.

The first stage in searching for pulsar candidates is to remove RFI from search data. For many projects, including those conducted with CHIME/Pulsar, we have standardized “kill” files, which include lists of frequency channels that are perennially affected by RFI and should be removed. For CHIME/Pulsar, this includes

the LTE band and digital TV stations. In addition to known RFI, there is often transient RFI, which requires observation-specific excision. For this purpose, PRESTO includes the tool `rfifind` which searches for RFI using FFTs and removes, particularly narrow-band and short-duration RFI. After running the automated RFI search, it is best practice to visually inspect the resulting mask and ensure there are no obviously bad channels left intact. It is also best practice to use FFTs and period searching tool `accelsearch` to find periodic sources of RFI known as “birdies.” These are generally related to the environment of the observatory (an RF signal at 60 Hz due to North American AC power is a typical example), so for CHIME/Pulsar observations, the CHIME/Pulsar team has created a standardized birdies file.<sup>2</sup>

After data has been cleaned of RFI, we test different DMs. The range of possible DMs when searching for Galactic pulsars is roughly constrained between zero and the local maximum Galactic DM. For FRBs, the maximum DM is theoretically infinite, but we will have an instrumental maximum beyond which the dispersion time will be too great for us to search. For sources with possible known DMs, like Galactic sources previously detected by CHIME/FRB or repeating FRBs, we can select an acceptable error range around the nominal DM value. PRESTO routine `ddplan.py` enables us to optimize parameters for our DM search, and the de-dispersion itself is conducted with PRESTO routine `prepsubband`. This routine incoherently de-disperses the data to each test DM and outputs the results as `.dat` files.

De-dispersed data are ready to be searched for sources. Our first step is generally to search for periodicity. If we are searching for transients, we will skip this step initially and begin with searching for single pulses, but period searching is the heart of pulsar searching. The first method we employ is generally the PRESTO `accelsearch` routine. This routine conducts Fourier-domain acceleration searches for periodicity using Fourier interpolation and summing a tunable number of harmonics (Ransom, 2011). An alternative method for period searches is fast folding algorithms (FFA), first proposed in Staelin (1969). This method, as the name implies, involves rapidly folding the data at nominal periods and deter-

---

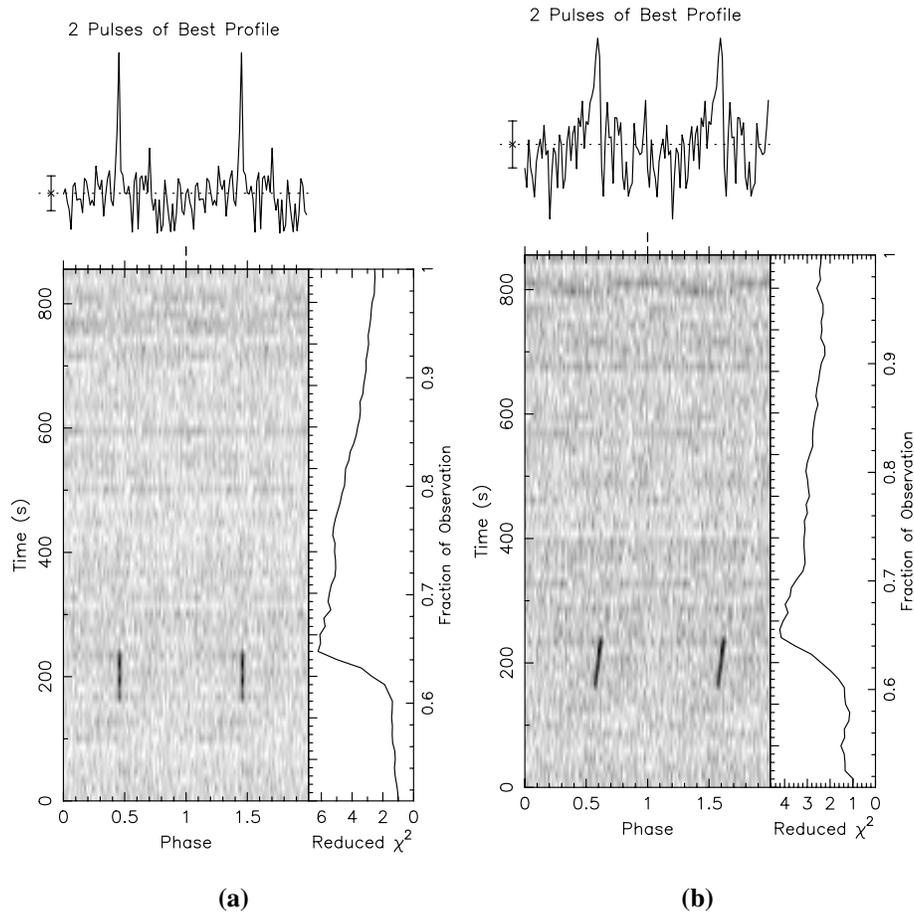
<sup>2</sup>Bradley Meyers and Kathryn Crowter, both of UBC, have been particularly key in this effort.

mining which best suits the data. It can be computationally expensive and therefore has been of limited utility, but has been used periodically throughout the past 50 years of pulsar astronomy. One recent successful implementation of FFA methods was for PALFA, which also demonstrated that FFA methods are more efficient at detecting long period sources (500 ms to 30 s) than `accelsearch` methods (Parent et al., 2018). Both methods were employed in periodicity searches for the new sources presented in Chapter 5.

If automated period searches fail to determine the periodicity for our source, we can also make use of brute force algorithms. The tool `rrat_period_search` included in PRESTO is one such tool. Brute force period determination searches for the spin period that allows the largest integer number of rotations between detections. For RRATs, this is often our best option for determining periods. For certain pulsars, particularly those with large nulling fractions, this may also be the best option for determining periodicity.

After a period is determined, we attempt to fold data. Folding is the process of coherently adding hundreds or even thousands of pulses from a pulsar together to create an integrated profile, with much greater signal to noise than an individual pulse. Most pulsars' single pulses are not bright enough to be observed and therefore the folded signal is needed to see the source. We use the PRESTO tool `prepfold` to test our initial period estimates. When data are folded to the correct period, the `prepfold` output looks like Figure 2.1a, but a period mis-estimate looks like Figure 2.1b.

When we are confident that we have correctly determined the spin period of the pulsar, we can use DSPSR (Digital Signal Processing for Pulsar Astronomy) to de-disperse all search-mode data taken for this source and to fold that data with the newly determined spin period. We can also use this basic information to develop a simple parameter (or “par”) file and begin collecting fold-mode observations. One of the advantages of CHIME/Pulsar is our flexible ability to observe many pulsars every day; in conventional pulsar searches, finding time for fold-mode follow-up to confirm detections and determine timing solutions is challenging. From either folded search-mode archives or new fold-mode archives, we are able to extract TOAs using the tools included in PSRCHIVE and begin the process of timing the pulsar. Pulsar timing is discussed in more depth in Section 2.4.



**Figure 2.1:** Results from folding data from PSR J0209+5759 taken on MJD 58527 with a slightly incorrect period (a) and with the correct period (b). The greyscale plot shows time vs. phase within the observation, the top inset shows pulse profile vs. phase, and side inset shows reduced  $\chi^2$  vs. fraction of the observation. The tilt in the phase vs. time plot in Panel b is indicative of an incorrect period, as the pulse moves slightly out of phase as time progresses. The pulse is visible only for a short portion of the period due to nulling in the source; the source location is visible to the CHIME/Pulsar beam for the duration of the observation.

With new pulsars, we often conduct gridding observations to improve our understanding of the pulsar’s position. In gridding observations, we simply split a single observation of the pulsar, switching between slightly different coordinates for the pulsar during the observation. As pulsars are observed as pulses and not resolved objects in folded profiles, comparing pulse S/N values for different coordinates is one way to refine our position. This position is ultimately improved by timing the pulsar.

For sufficiently bright pulsars and for radio transients, it is desirable or, in the case of transients, necessary to search for single pulses. PRESTO enables this with the `single_pulse_search.py` routine. This routine searches through the output of `prepsubband` analyses to find clusters of high signal-to-noise. A sample `single_pulse_search.py` output showing the detection of RRAT J0121+53 is shown in Figure 2.2.

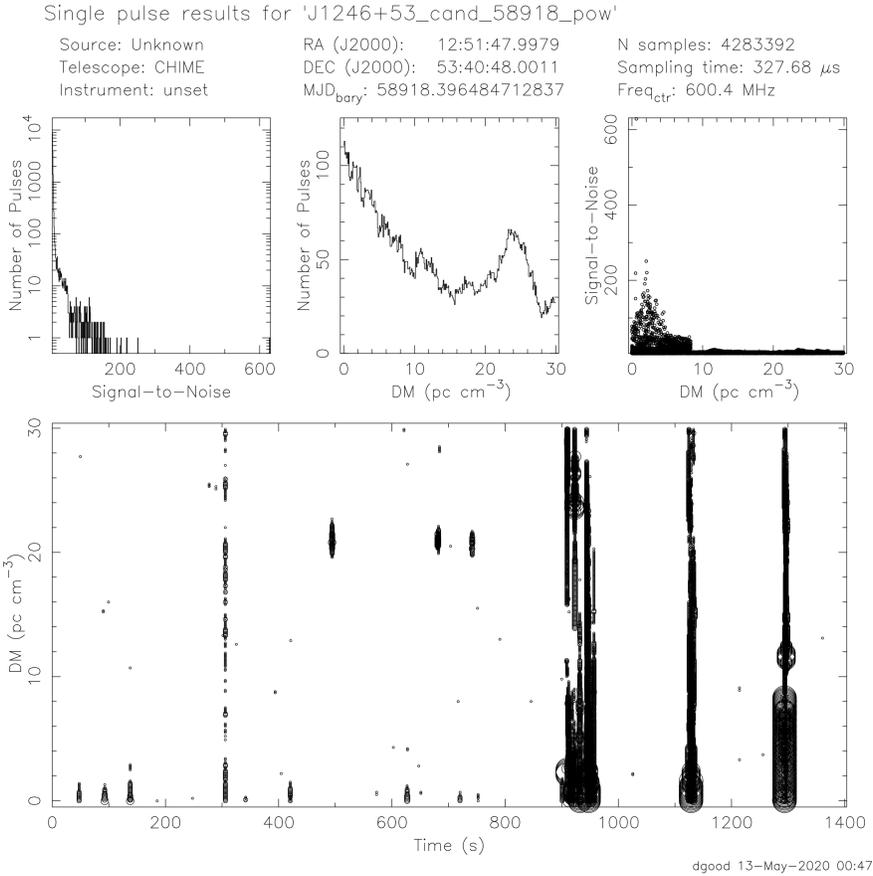
## 2.4 Pulsar Timing

In Chapter 1, we discussed the potential applications of pulsar timing, without discussing the process for creating pulsar timing datasets. In Chapter 4, we will construct such a dataset. Therefore, we now consider how to create pulsar times of arrival (TOAs) and pulsar timing models.

Building high precision models of the pulsar’s behaviour must incorporate a precise understanding of the source’s own properties, the properties of the ISM, and the solar system barycenter (SSB; the precise center of mass of the solar system). The derivation of TOAs and pulsar timing models in this section is drawn from Lommen & Demorest (2013) and Lorimer & Kramer (2005). It is also informed by other sources where listed.

### Template Generation

To create pulsar timing solutions, we must first determine a standard template pulse profile for the source. Pulsar profiles are generally stable at a given frequency, although they can vary dramatically between frequencies. This means that for a given frequency, we can create standardized profiles or templates by adding together the results of previous observations and smoothing the final profile.



**Figure 2.2:** Output from PRESTO routine `single_pulse_search.py`, showing a single pulse detection of RRAT J0121+53, a source discussed in Chapter 5. Clockwise from the top left, the subplots show number of pulses vs.  $S/N$ , number of pulses vs. DM,  $S/N$  vs. DM, and DM vs. observation time. In the DM vs. time plot, each black dot is representative of a pulse with  $S/N \geq 6$ . The size of the dot indicates the  $S/N$ . Most apparent signals in this figure are in fact RFI. This includes the signals that seem to appear at all values of DM (e.g. at 900, 1100, and 1300 s) and those signals which appear at  $DM = 0$ , such as those visible between 0 and 200s into the observation. The times showing detections at all DM values and at  $DM \sim 0$  are RFI signals; the RRAT is visible three times at  $DM \sim 21$  and approximately 500, 650, and 750 seconds.

## TOA Generation

Pulsar timing must begin by extracting TOAs from a set of observations. The standard Fourier-domain algorithm for this process was created by Taylor (1992). Working in terms of pulse phase  $\phi$ , which ranges from 0 to 1, the data are  $d(\phi)$  and the template  $p(\phi)$ . We know  $d(\phi)$  and  $p(\phi)$  are offset by a phase shift  $\Delta\phi$ . Our goal is to determine the value of  $\Delta\phi$ , which we do by a  $\chi^2$  minimization in the Fourier domain.

$$\chi^2(a, \delta\phi) = \sum_{k=1}^{k_{max}} \frac{|d_k - a p_k e^{-i2\pi k \Delta\phi}|^2}{\sigma^2}, \quad (2.19)$$

where  $d_k$  and  $p_k$  are discrete Fourier transforms of  $d(\phi)$  and  $p(\phi)$ ,  $\sigma^2$  is the noise power for each harmonic component, and  $a$  is a scale factor between the profiles (Lommen & Demorest, 2013).

This  $\chi^2$  value can be expanded as

$$\chi^2 = \sigma^{-2} \left( \sum_k d_k^2 + \sum_k p_k^2 - 2a \operatorname{Re} \left[ \sum_k d_k p_k^* e^{2\pi i k \phi} \right] \right). \quad (2.20)$$

All phase-shift information is contained in the final, cross-correlation term, and  $\chi^2$  is minimized at the phase shift that maximizes this final term (Lommen & Demorest, 2013). This  $\chi^2$  minimizing phase shift is  $\Delta\hat{\phi}$ , and is the difference between the observed and predicted pulse phases.

However, we are ultimately searching for a time, not a phase, so we must convert our  $\chi^2$  minimizing pulse phase to a TOA. We reference this time to the midpoint of the observation, not the start to avoid dependence on the initial timing model used in observing. This is important, as the models used to fold pulsars are often less accurate than the models derived from a large set of TOAs. To generate the TOA, we subtract the model-predicted phase at the observation midpoint from the measured phase shift. Then we multiply this difference by the model-predicted pulse period and add it to the midpoint. This creates our TOA. We can estimate the uncertainty in this TOA in phase space by taking the second derivative of Equation 2.20, and it is limited only by the pulsar  $S/N$  (Lommen & Demorest, 2013; Lorimer & Kramer, 2005).

In time space, the uncertainty in the TOA can be expressed in terms of the pulse width and the  $S/N$

$$\sigma_\phi \approx \frac{W}{S/N} \propto \frac{S_{\text{sys}}}{\sqrt{t_{\text{obs}}\Delta f}} \times \frac{P\delta^{3/2}}{S_{\text{mean}}}, \quad (2.21)$$

where  $S_{\text{sys}}$  is the system equivalent flux density,  $\Delta f$  is the observing bandwidth,  $t_{\text{obs}}$  is the integration time,  $P$  is the pulse period,  $W$  is the pulse width,  $\delta$  is the pulsar duty cycle, and  $S_{\text{mean}}$  is the mean flux density of the pulsar (Lorimer & Kramer, 2005). This equation leads to the conclusion that long, wide bandwidth observations of bright, short-period, narrow, stable pulsars have particularly low uncertainties, a key motivation for Chapter 4.

Once we have developed a set of TOAs with corresponding uncertainties, we are ready to build a timing model. Timing models are simply sets of physically motivated parameters fit to a given set of TOAs. The prospect of doing such a fit by hand is daunting; for this reason we use standard pulsar timing software such as TEMPO (Nice et al., 2015), TEMPO2 (Hobbs et al., 2006; Edwards et al., 2006; Hobbs et al., 2009), and PINT (Luo et al., 2021). These programs perform many-parameter least squares fits for a set of TOAs known as a “tim” file and a starting set of parameters, known as a “par” file, outputting fit parameters, residuals, and  $\chi^2$ . Par files for recently discovered pulsars begin very simply, perhaps with just a period (spin frequency), dispersion measure, and position, but they gradually become more complex as our understanding of the pulsar’s parameters grows.

### 2.4.1 Timing Corrections

Before we can begin fitting our timing model, we must address the fact that observing reference frame in our Earth-based observatories is not inertial. Therefore, we must convert our topocentric TOA ( $t_{\text{topo}}$ ) measurement to a barycentric TOA, referenced to the solar system barycenter, which is to good approximation an inertial reference frame.

We can calculate the barycentric TOA from the topocentric TOA by the expression

$$t_{\text{SSB}} = t_{\text{topo}} + t_{\text{corr}} - \Delta D/f^2 + \Delta_{R_\odot} + \Delta_{S_\odot} + \Delta_{E_\odot}. \quad (2.22)$$

The first correction,  $t_{\text{corr}}$  is the clock correction associated with the observatory

and with the time standard itself. Observatory times are usually determined by a hydrogen maser at the observatory and regularly compared to GPS time, which is in Universal Time (UTC). UTC is by definition an integral number of seconds from International Atomic Time (TAI) and not more than 0.9 s away from UT1, another time standard, this one referenced to Earth's rotation. UTC is intended for human use and therefore includes leap seconds. Therefore,  $\text{TAI} = \text{UTC} + \Delta T$  where  $\Delta T$  is the sum of the leap seconds. TAI is maintained by the Bureau International de Poids et Mesures (BIPM), which publishes the Terrestrial Time (TT) standard, the timescale used to measure TOAs (Lorimer & Kramer, 2005).

The next term is the frequency dependent correction needed to manage dispersion from the interstellar medium. Dispersion and DM were discussed in more depth in 2.2.1. This correction is dependent on the DM and the observing frequency, with  $\Delta D = \frac{e^2}{2\pi m_e c} \text{DM}$ .

The next three terms are corrections to convert from the topocentric to barycentric frame. The topocentric reference frame is the non-inertial reference frame of the Earth and the barycentric frame is the inertial reference frame of the Solar System Barycenter (SSB). First, the Römer delay  $\Delta_{R\odot}$  is the classical light travel time between the phase center of the telescope and the SSB. Second, the Shapiro delay  $\Delta_{S\odot}$  is the relativistic correction due to curvature of space-time created by the presence of masses in the solar system (Shapiro, 1964). Third, the Einstein delay  $\Delta_{E\odot}$  is the effect of both time dilation due to the motion of the earth and gravitational redshift due to other masses in the Solar System (Backer & Hellings, 1986). Mathematically, these corrections are

$$\Delta_{R\odot} = -\frac{1}{c} \mathbf{r} \cdot \hat{\mathbf{s}} = -\frac{1}{c} (\mathbf{r}_{\text{SSB}} + \mathbf{r}_{\text{EO}}) \cdot \hat{\mathbf{s}}, \quad (2.23)$$

$$\Delta_{S\odot} = -2 \sum_i \frac{GM_i}{c^3} \ln \left( \frac{\hat{\mathbf{s}} \cdot \mathbf{r}_i^E + r_i^E}{\hat{\mathbf{s}} \cdot \mathbf{r}_i^P + r_i^P} \right) \quad (2.24)$$

$$\frac{d\Delta_{E\odot}}{dt} = \sum_i \frac{GM_i}{c^2 r_i^E} + \frac{v_E^2}{2c^2} - \text{constant}, \quad (2.25)$$

where  $\mathbf{r}$  is the vector connecting the Solar System Barycenter and the observatory, split into the component from the SSB to the center of the Earth ( $\mathbf{r}_{\text{SSB}}$ ) and the

component from the center of the Earth to the observatory ( $\mathbf{r}_{EO}$ ),  $M_i$  are the masses of bodies in the solar system,  $\mathbf{r}_i^P$  and  $\mathbf{r}_i^E$  are the vectors between the pulsar and the solar system body and the Earth and the solar system body, and  $v_E$  is the velocity of the earth relative to the Sun. In Equation 2.25, we are summing over all the masses in the solar system excluding the Earth. Calculating the SSB depends on the solar system ephemeris we make use of; precision timing in Chapter 4 uses DE440 and timing of new RRATs and pulsars in Chapter 5 uses the slightly older DE436 standard Park et al. (2021).

## 2.4.2 Timing Isolated Pulsars

For all pulsars, timing solutions begin with spin parameters and astrometric parameters.

Not every parameter is measurable for every pulsar. Indeed, comparing solutions for the well-known MSPs presented in Chapter 4 to the solutions for new slow pulsars presented in Chapter 5 is a striking illustration that developing precise timing solutions is not an instantaneous task.

The most basic parameters for the timing model are the spin parameters: the period, the period derivative, and sometimes the period second derivative. Astrometric parameters are the parameters that describe our understanding of the pulsar's location in space and its interaction with the background environment. There are five major parameters: two position parameters, two proper motion parameters, and parallax.

Proper motion is, in general, a measure motion of a nearby source relative to its distant background. For pulsars specifically, the vector  $\hat{\mathbf{s}}$  is gradually altered by the transverse component of the pulsar's velocity relative to the SSB ( $v_T$ ). There is also, in principle, a radial component, but this is generally not measurable. The transverse velocity causing such proper motion for a pulsar at distance  $d$  is

$$v_T = 4.74 \text{km s}^{-1} \left( \frac{\mu_T}{\text{mas yr}^{-1}} \right) \left( \frac{d}{\text{kpc}} \right). \quad (2.26)$$

The proper motion that can be measured as a motion in declination ( $\delta$ ) is  $\mu_\delta \equiv \dot{\delta}$  and in right ascension ( $\alpha$ ) is  $\mu_\alpha \equiv \dot{\alpha} \cos \delta$ , such that the total proper motion is

$\mu_T = \sqrt{\mu_\alpha^2 + \mu_\delta^2}$ . The units are milliarcseconds per year.

Timing residuals can also demonstrate a phenomenon known as timing parallax. Timing parallax is specifically a measure of the curvature of the emitted wavefronts at different points within the Earth's orbit, but like optical parallax, it can be used to determine a precise distance to the object in question. The time delay due to timing parallax is (Backer & Hellings, 1986)

$$\Delta t_\pi = -\frac{1}{2cd} (\mathbf{r} \times \hat{\mathbf{s}})^2, \quad (2.27)$$

where  $\mathbf{r}$  is the vector connecting the observatory and SSB,  $\hat{\mathbf{s}}$  is a unit vector pointing from the SSB to the pulsar, and  $d$  is the distance to the pulsar.

The timing parallax's amplitude is dependent on the Earth-Sun distance, the pulsar distance, and the ecliptic latitude of the pulsar:

$$\text{PX Amplitude} = l^2 \cos \beta / (2cd). \quad (2.28)$$

This can be challenging to measure, but is measurable for many MSPs and can be our most accurate method for determining distance to a pulsar. This is commonly included in NANOGrav timing, but the relatively short timespan of CHIME/Pulsar data means that it is not fitted for in CHIME/Pulsar only models in Chapter 4. Timing parallax is also difficult to measure for young pulsars, which are more dynamic.

### 2.4.3 Timing Binary Pulsars

Timing solutions for binary pulsars incorporate the spin parameters and astrometric parameters described in Section 2.4.2, but add further parameters to describe the binary system.

#### Binary Parameters

With the binary parameters added, the barycentric TOA is now

$$t_{\text{SSB}} = t_{\text{topo}} + t_{\text{corr}} - \Delta D / f^2 + \Delta_{R\odot} + \Delta_{S\odot} + \Delta_{E\odot} + \Delta_{\text{RB}} + \Delta_{\text{SB}} + \Delta_{\text{EB}}. \quad (2.29)$$

These new terms are Römer Delay, Shapiro Delay, and Einstein Delay from the gravitational field within the binary system, and the aberration from orbital motion.

In our timing model, we combine both Keplerian (classical) dynamical parameters for binary systems as well as extended relativistic (post-Keplerian) parameters. The Keplerian parameters are the binary period, the semi-major axis (projected to account for inclination)  $a_p \sin i$ , orbital eccentricity  $e$ , the longitude of periastron  $\omega$ , the epoch of periastron passage  $T_0$ , and the position angle of the ascending node  $\Omega_{\text{asc}}$ . As pulsar binaries are often relativistic systems, we will often have to work with the post-Keplerian (PK) relativistic parameters which include relativistic advance of periastron  $\dot{\omega}$ , binary period derivative  $\dot{P}_b$ , amplitude parameter  $\gamma$ , and “range” parameter  $r$  and “shape” parameter  $s$ . Any combination of two post-Keplerian parameters will enable us to measure the mass of the pulsar, though measuring relativistic Shapiro delay is a popular strategy. Binary parameters each have a physical definition based on orbital dynamics, but are usually found using simpler parametrizations. A more detailed description of pulsar binary parameters can be found in Lorimer & Kramer (2005).

For a generic binary system, the Keplerian parameters can be described in terms of the eccentric  $E$  and true anomaly  $A_T$ , and the mean angular velocity  $\Omega_b = 2\pi/P_b$ :

$$E - e \sin E = \Omega_b \left[ (T - T_0) - \frac{1}{2} \frac{\dot{P}_b}{P_b} (t - T_0)^2 \right] \quad (2.30)$$

$$A_T(E) = 2 \arctan \left[ \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2} \right] \quad (2.31)$$

$$\omega = \omega_0 + \frac{\dot{\omega}}{\Omega_b} A_T(E). \quad (2.32)$$

We omit a detailed discussion of Keplerian orbits, but such derivations can be found in many astronomical dynamics books, such as Binney & Tremaine (2008).

The Keplerian Römer delay for binary system can be expressed in terms of the anomaly as

$$\Delta_{\text{RB}} = x (\cos E - e) \sin \omega + x \sin E \sqrt{1 - e^2} \cos \omega, \quad (2.33)$$

where  $x$  is the projected semi-major axis,  $x \equiv a_p \sin i/c$ .

Pulsars are very compact objects, and when our binary systems involve another

compact object, such as a white dwarf, neutron star, or black hole, we can observe relativistic components to the binary motion, due to the strong gravitational fields and high orbital velocities involved in motion.

In such systems, we modify the Römer delay, so that it is

$$\Delta_{\text{RB}} = x(\cos E - e_r) \sin \omega + x \sin E \sqrt{1 - e_\theta^2} \cos \omega, \quad (2.34)$$

where  $e_r \equiv e(1 + \delta_r)$  and  $e_\theta \equiv e(1 + \delta_\theta)$  and  $\delta_r$  and  $\delta_\theta$  are relativistic deformations to the binary orbit. These deformations are not usually measurable.

We also need to consider the effects of time dilation and redshift due to the presence of the companion, and this is encapsulated by the Einstein delay

$$\Delta_{\text{EB}} = \gamma \sin E, \quad (2.35)$$

where  $\gamma$  is a post-Keplerian parameter representing the amplitude of the delay in seconds.

For some systems (generally only those viewed close to edge-on, i.e. with inclination angle nearly  $90^\circ$ ), we can measure relativistic Shapiro delay. This allows us to measure PK parameters  $r$  and  $s$ . Within a binary system, the Shapiro delay is

$$\Delta_{\text{SB}} = -2r \ln \left[ 1 - e \cos e - s \left( \sin \omega (\cos E - e) + \sqrt{1 - e^2} \cos \omega \sin E \right) \right], \quad (2.36)$$

where the maximum occurs at superior conjunction, i.e. phase  $\Phi = \omega + A_T(E) = \pi/2$ . For small-eccentricity binaries, this simplifies to

$$\Delta_{\text{SB}} = -2r \ln [1 - s \sin \Phi]. \quad (2.37)$$

## Models for Binary Pulsars

The equations governing binary pulsars are complex and do not immediately lead to the parameters we want to fit. There are multiple approaches to parametrizing binary parameters. Here, we present two major structures, each of which is used in NANOGrav pulsar timing.

For systems with small eccentricity, we simplify our timing solutions and use

the ELL1 timing model, developed by Lange et al. (2001). In the case of small eccentricity, we can re-write Equation 2.33 as

$$\Delta_{\text{RB}} \approx x \left( \sin \Phi + \frac{\varepsilon_2}{2} \sin 2\Phi - \frac{\varepsilon_1}{2} \cos 2\Phi \right) \quad (2.38)$$

where  $\Phi$  is measured from the ascending node. We also expand Equation 2.37 as

$$\Delta_{\text{SB}} = 2r(a_0 + b_1 \sin \Phi - a_2 \cos \Phi + \dots), \quad (2.39)$$

where

$$a_0 = -\ln \left( \frac{1 + \sqrt{1 - s^2}}{2} \right), \quad (2.40)$$

$$b_1 = 2 \left( \frac{1 - \sqrt{1 - s^2}}{s} \right), \quad (2.41)$$

and

$$a_2 = 2 \left( \frac{1 - \sqrt{1 - s^2}}{s} \right) - 1. \quad (2.42)$$

Our Keplerian parameters  $T_0$ ,  $e$ , and  $\omega$  have now been replaced by the epoch of the ascending node ( $T_{\text{asc}}$ ) and the Laplace-Lagrange parameters ( $\varepsilon_1$  and  $\varepsilon_2$ ):

$$T_{\text{asc}} = T_0 - \frac{\omega}{\Omega_b}, \quad (2.43)$$

$$\varepsilon_1 = e \sin \omega \quad (2.44)$$

$$\varepsilon_2 = e \cos \omega. \quad (2.45)$$

Though we use  $\varepsilon_1$ ,  $\varepsilon_2$ , and  $T_{\text{asc}}$  as our fit parameters, we can extract the Keplerian parameters from these parameters:

$$e = \sqrt{\varepsilon_1^2 + \varepsilon_2^2}, \quad (2.46)$$

$$\omega = \arctan(\varepsilon_1/\varepsilon_2) \quad (2.47)$$

$$T_0 = T_{\text{asc}} + \frac{P_b}{2\pi} \arctan(\varepsilon_1/\varepsilon_2) \quad (2.48)$$

As the Shapiro Delay and Romer Delay are coupled in this scenario, we find that the observed and intrinsic values of the semi-major axis  $x$  and Laplace-Lagrange parameter  $\varepsilon_1$  differ slightly:

$$x^{\text{obs}} = x^{\text{int}} + 4r \left( \frac{1 - \sqrt{1 - s^2}}{s} \right) \quad (2.49)$$

$$\varepsilon_1^{\text{obs}} = \varepsilon_1^{\text{int}} + \frac{4r}{x^{\text{int}}} \left( 2 \frac{1 - \sqrt{1 - s^2}}{s^2} - 1 \right). \quad (2.50)$$

For high eccentricity systems, we cannot use the simplifying assumptions of the ELL1 model and instead make use of the DD model (Damour & Deruelle, 1986). In the DD model, the post-Keplerian parameters are expressed (to lowest post-Newtonian order) as

$$\dot{\omega} = 3T_{\odot}^{2/3} \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{1}{1 - e^2} (m_p + m_c)^{2/3}, \quad (2.51)$$

$$\gamma = T_{\odot}^{2/3} \left( \frac{P_b}{2\pi} \right)^{1/3} e \frac{m_c(m_p + 2m_c)}{(m_p + m_c)^{4/3}} \quad (2.52)$$

$$r = T_{\odot} m_c \quad (2.53)$$

$$s = \sin i = T_{\odot}^{-1/3} \left( \frac{P_b}{2\pi} \right)^{-2/3} x \frac{(m_p + m_c)^{2/3}}{m_c} \quad (2.54)$$

$$\dot{P}_b = -\frac{192\pi}{5} T_{\odot}^{5/3} \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{(1 + (73/24)e^2 + (37/96)e^4)}{(1 - e^2)^{7/2}} \times \frac{m_p m_c}{(m_p + m_c)^{1/3}}, \quad (2.55)$$

where  $m_p$  is the mass of the pulsar and  $m_c$  is the mass of the companion, both in units of solar masses.

The DDK model is a further expansion to the DD model which adds considerations of annual orbital parallax within the binary system and periodic variations in the longitude of the periastron (Kopeikin, 1995, 1996).

## Noise Modeling

Although we have a straightforward expression for the uncertainty in the TOA, it sometimes becomes clear that this uncertainty is underestimated. In these cases, we can add additional noise parameters to address underestimated or unmodeled noise in our data (Lommen & Demorest, 2013). There are three main noise components: EFAC, EQUAD, and ECORR. EFAC is a multiplicative parameter that accounts for systematic underestimation of TOA uncertainty and is calculated for each pulsar, receiver, and backend combination. A value near one indicates that our systemic understanding is accurate. EQUAD is, as the name suggests, added in quadrature. It accounts for white noise beyond the statistical uncertainty of the TOA calculation. ECORR describes short time-scale noise that is not correlated between epochs but is entirely correlated among TOAs observed simultaneously. Physically, this includes wideband noise processes such as pulse jitter.

Some pulsars also demonstrate red noise in the residuals, which requires either whitening of the residuals before analysis, as in Coles et al. (2011) or modeling of red noise parameters such as in van Haasteren & Levin (2013). In NANOGrav noise analysis, we define this as any steep-spectrum noise component and model those components with a single stationary Gaussian process with the spectrum

$$P(f) = A_{\text{red}}^2 \left( \frac{f}{1\text{yr}^{-1}} \right)^{\gamma_{\text{red}}}, \quad (2.56)$$

where  $A_{\text{red}}$  is the red noise amplitude and  $\gamma_{\text{red}}$  is the red noise index and fit for  $A_{\text{red}}$  and  $\gamma_{\text{red}}$  as part of our noise analysis.

## 2.5 Pulsar observation in this work

The techniques and considerations outlined in this chapter are key for understanding the rest of the work in this thesis. In Chapter 4, we build combined CHIME/Pulsar and NANOGrav timing models for a set of eight pulsars. In Chapter, 5, we search for new pulsars with CHIME/FRB and CHIME/Pulsar then determine initial timing solutions for some of these sources. In Chapter 6, we present the results of a search for repeating FRBs, which leverages single pulse search techniques from pulsar searching.

## Chapter 3

# The Canadian Hydrogen Intensity Mapping Experiment

*No observational problem will not be solved by more data.* — Vera Rubin

### 3.1 The Canadian Hydrogen Intensity Mapping Experiment

CHIME was initially proposed in the early 2010s as an observational cosmology experiment. One of cosmology’s great mysteries is the accelerating expansion of the universe and thus dark energy. The expansion history of the universe can be mapped by tracking the baryon acoustic oscillation (BAO) scale through different eras of the universe’s expansion. (These eras are denoted by their “redshifts” – a measure of distance based on the Doppler shift of galaxies moving away from the Milky Way.) The BAO are fluctuations in the density of baryonic matter in the universe caused by acoustic waves in the early universe plasma, which form a standard cosmological ruler. See, e.g., Chang et al. (2010) for a previous example of this technique.

Previous work on this problem has focused on detailed galaxy surveys. However, resolving galaxies at relevant redshifts is a time consuming, arduous process which ultimately provides more information than the cosmologists need anyway.

Thus, cosmologists turned to hydrogen intensity mapping. This method focuses on scanning large areas of the sky with low angular resolution at a high mapping speed. Thus, the idea for CHIME – a wide field, low angular resolution telescope at 400–800 MHz (redshift 0.8 to 2.5) was born. CHIME is a transit telescope, meaning it has no moving parts and views different sections of the sky as the Earth rotates. This allows CHIME to see the entire Northern sky every day, with maximized mapping speed (Swenson, 1969).

Shortly after CHIME was proposed, it was suggested that such an instrument would be useful for other applications. The first was high cadence pulsar timing and other observations – an idea that eventually became CHIME/Pulsar. As development was ongoing, FRB scientists also realized the value of a radio frequency all sky monitor and CHIME/FRB was born. Since this time, the scope of CHIME projects has expanded further to encompass the upcoming Slow Pulsar Search and CHIME/FRB very long baseline interferometry (VLBI) outrigger projects, both of which leverage the existing CHIME/FRB system. CHIME operations are fully commensal, with all backends receiving data simultaneously 24/7 (except when undergoing occasional maintenance and software upgrades).

## **3.2 A brief introduction to interferometry**

This introduction to interferometry follows the logic of Masui et al. (2019), which itself leverages the notation of Shaw et al. (2014) and Shaw et al. (2015). This notation is built on nearly a century of radio astronomy, but a summary of interferometry and synthesis in radio astronomy can be found in Thompson et al. (2017).

The minimal case of a radio telescope is a single antenna attached to some form of data storage mechanism. A marginally more complex and much more useful set-up is a single dish telescope, where a single antenna is affixed to a reflector and data are recorded from this set-up.

The data collected by such a single-dish telescope would be a combination of the electric field emitted by the sky and instrumental effects, codified as the antenna response function and as Gaussian noise. Though simple in principle, single dish radio astronomy has been the dominant form of radio astronomy and is sufficient for many applications. In particular, it is very well suited to pulsar astronomy.

Ultimately, we are observing electric fields in the far-field limit, which can be described as a sum of plane waves,

$$\mathbf{E}(\mathbf{x}, t) = \frac{1}{\sqrt{\epsilon_0 c}} \int \boldsymbol{\varepsilon}(\hat{\mathbf{n}}, \nu) e^{-i2\pi\nu(t - \mathbf{x} \cdot \hat{\mathbf{n}}/c)} d^2\hat{\mathbf{n}} d\nu, \quad (3.1)$$

where  $\nu$  is frequency,  $\hat{\mathbf{n}}$  is a unit vector pointing to a single direction on the sky,  $d^2\hat{\mathbf{n}}$  is a differential solid angle, and this equation defines the differential electric field in some direction at some frequency,  $\boldsymbol{\varepsilon}(\hat{\mathbf{n}}, \nu)$ .

We treat the total emission observed by our telescope as incoherent and in the far field, allowing us to create an intensity matrix by cross-correlating the electric fields in different directions

$$\langle \boldsymbol{\varepsilon}_j(\hat{\mathbf{n}}, \nu) \boldsymbol{\varepsilon}_k^*(\hat{\mathbf{n}}', \nu') \rangle = \delta^2(\hat{\mathbf{n}} - \hat{\mathbf{n}}') \delta(\nu - \nu') \frac{k_B \nu^2}{c^2} I_{jk}(\hat{\mathbf{n}}, \nu), \quad (3.2)$$

where indices  $j, k$  are directions perpendicular to the incident radiation.

$I_{jk}$  can be decomposed in terms of the Stokes parameters, using polarization matrices  $P^P$ , which are equivalent to the familiar Pauli matrices:

$$I_{jk}(\hat{\mathbf{n}}) = P_{jk}^I I(\hat{\mathbf{n}}) + P_{jk}^Q I(\hat{\mathbf{n}}) + P_{jk}^U I(\hat{\mathbf{n}}) + P_{jk}^V I(\hat{\mathbf{n}}) \quad (3.3)$$

$$P_{jk}^I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} P_{jk}^Q = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} P_{jk}^U = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} P_{jk}^V = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}. \quad (3.4)$$

We re-write Equation 3.3 as  $I_{jk}(\hat{\mathbf{n}}) = P_{jk}^P I_P(\hat{\mathbf{n}})$ , where repeated indices are summed.

We relate this intensity framework to the spectral flux density for an unresolved source by

$$\frac{dF_\nu}{d\Omega} = \frac{k_B \nu^2}{c^2} \delta^{jk} P_{jk}^P I_P(\hat{\mathbf{n}}) = \frac{2k_B \nu^2}{c^2} I^S(\hat{\mathbf{n}}, \nu), \quad (3.5)$$

where  $I^S$  is the unpolarized intensity for a single source.

The signal an individual antenna receives is a combination of the complex 2-D vector field antenna response function  $A_a^j(\hat{\mathbf{n}}, \nu)$  and the sky signal  $\boldsymbol{\varepsilon}_j(\hat{\mathbf{n}}, \nu)$ . The index  $j$  is again a direction perpendicular to the incident radiation and  $a$  is an

antenna index. This creates a general signal:

$$\eta_a = \int A_a^j(\hat{\mathbf{n}}) \varepsilon_j(\hat{\mathbf{n}}) e^{i2\pi \mathbf{u}_a \cdot \hat{\mathbf{n}}} d^2 \hat{\mathbf{n}} + n_a(\nu). \quad (3.6)$$

The quantity  $\mathbf{u}_a = \mathbf{x}_a/\lambda$  is the feed position in wavelengths and  $n_a(\nu)$  is the receiver noise for receiver  $a$ .

However, even in the early years of radio astronomy, it was realized that we could expand our observational capabilities by combining multiple antennae into an interferometer (Ryle & Vonberg, 1946). For example, interferometers allow for greater angular resolution than single dish telescopes, by creating long baselines (and even very long, continental scale baselines) between antennae. This creates the effect of telescopes far too large to be constructed as single dishes.

Though by no means a new technology, interferometry is in an exciting period of expansion. It has long been limited by the computational resources required to combine signals, but this problem has been ameliorated by Moore’s Law. Therefore, the past decade has seen an explosion of interferometry, including the Event Horizon Telescope (Event Horizon Telescope Collaboration et al., 2019), the MWA (Tingay et al., 2013), LOFAR (van Haarlem et al., 2013), HERA (DeBoer et al., 2017), HIRAX (Newburgh et al., 2016), STARE2 (Bochenek et al., 2020a), MEERKAT (Booth & Jonas, 2012), and of course CHIME. The coming decade will see the continued operation of many existing interferometers and the growth of large-scale interferometer observatory projects like DSA2000 (Hallinan et al., 2019), the ngVLA (Selina et al., 2018), and the SKA (Schilizzi et al., 2008)

We now present a brief discussion of interferometer signals, again following the example of Masui et al. (2019). Interferometers combine many individual antenna by cross-correlating them, combining two antenna. In Fourier space, this creates the quantity “the visibility.”<sup>1</sup> For two antenna  $a$  and  $b$ , the visibility is

$$V_{ab} \equiv \frac{c^2}{k_B \nu^2} \frac{1}{n_{\text{samp}}} \sum_t \eta_a[t] \eta_b[t]^*, \quad (3.7)$$

which is can be determined by estimating the covariance between the feeds over a

---

<sup>1</sup>For an in-depth discussion of image vs. Fourier space in this context, see Chapter 2 of Thompson et al. (2017).

set of time samples

$$C_{ab} \equiv \langle V_{ab} \rangle = S_{ab} + N_{ab}. \quad (3.8)$$

We define  $S_{ab}$  as the measured visibility and  $N_{ab}$  as receiver noise. We can combine the antenna response function, Equation 3.6, and the intensity generated by a single source, from Equation 3.5.

$$S_{ab}(\nu) = \int A_a^j(\hat{\mathbf{n}}, \nu) A_b^{k*}(\hat{\mathbf{n}}, \nu) * P_{jk}^P I_p(\hat{\mathbf{n}}, \nu) e^{i2\pi \mathbf{u}_{ab} \cdot \hat{\mathbf{n}}} d^2 \hat{\mathbf{n}}, \quad (3.9)$$

where  $\mathbf{u}_{ab} = \mathbf{u}_a - \mathbf{u}_b$  is the physical separation between antennae, measured in wavelengths.

For observations of a single source, Equation 3.9 reduces to

$$S_{ab}^S(\nu) = \frac{c^2}{2k_B \nu^2} F_\nu \delta_{jk} A_a^j(\hat{\mathbf{n}}_s, \nu) A_b^{k*}(\hat{\mathbf{n}}_s, \nu) e^{i2\pi \mathbf{u}_{ab} \cdot \hat{\mathbf{n}}_s} \quad (3.10)$$

The covariance of two visibilities (deriving from antenna combinations  $ab$  and  $cd$ ) is

$$\text{Cov}(V_{ab}, V_{cd}) = \frac{C_{ac} C_{bd}^*}{\Delta \nu \Delta t}, \quad (3.11)$$

where  $\Delta \nu$  is the observation bandwidth and  $\Delta t$  is the sampling time. In the case that the receiver noise term  $N_{ab}$  is dominant over the non-resolved sky noise, this equation reduces to the radiometer equation,

$$\Delta S_{\text{sys}} = \frac{T_{\text{sys}}}{G \sqrt{n_p t_{\text{obs}} \Delta \nu}}, \quad (3.12)$$

where  $T_{\text{sys}}$  is the system temperature in K,  $t_{\text{obs}}$  is integration time, and  $\Delta \nu$  is the system bandwidth (Lorimer & Kramer, 2005). The antenna forward gain,  $G$  is the maximum response of an antenna to a point source. It can be derived from the Equation 3.10 and is  $G_f = c^2 \delta_{jk} A_a^j A_a^{k*} / 2k_B \nu^2 = A_{\text{eff}} / 2k_B$ . Its units are K/Jy.

Visibilities are the basic building block for interferometric data, and for many applications, they are the desired final output. However, visibilities necessarily encompass the entire field-of-view of the telescope. For mapping applications, like CHIME's flagship cosmology experiment, or for telescopes with smaller beams, this is ideal. For CHIME/FRB and CHIME/Pulsar, where we want to observe

smaller areas than the  $\sim 2^\circ$  by  $\sim 100^\circ$  primary beam, we employ a process known as beamforming. It is important at this stage to draw a distinction between formed beams, created using the process outlined below, and the primary beam. The latter is the total field-of-view of the telescope and the beam pattern seen by the telescope when not using beamforming.

Beamforming allows us to create miniature telescopes digitally by combining visibilities  $V_{ab}$  and weights  $w_{ab}$ :

$$b = w^{ab} V_{ab}, \quad (3.13)$$

where  $w^{ab}$  is a pre-defined set of weights, generally normalized such that  $\sum_{ab} w^{ab} w^{ab*} = 1$ . The expectation value for this beam is

$$\langle b(\mathbf{v}) \rangle = \int B^{jk}(\hat{\mathbf{n}}, \mathbf{v}) P_{jk}^P I_P(\hat{\mathbf{n}}, \mathbf{v}) d^2\hat{\mathbf{n}} + w^{ab} N_{ab}, \quad (3.14)$$

with the beam response function  $B^{jk}(\hat{\mathbf{n}}, \mathbf{v})$

$$B^{jk}(\hat{\mathbf{n}}, \mathbf{v}) = w^{ab} A_a^j(\hat{\mathbf{n}}, \mathbf{v}) A_b^{k*}(\hat{\mathbf{n}}, \mathbf{v}) e^{i2\pi\mathbf{u}_{ab}\cdot\hat{\mathbf{n}}}. \quad (3.15)$$

The purpose of beamforming is usually to digitally point our telescope at some sky location  $\hat{\mathbf{n}}_p$ . We leverage the measured visibility in a given source's direction to write an expression for a pointed beam  $b_p$ .

$$b_p(\hat{\mathbf{n}}_p) = \frac{1}{\mathcal{N}} \sum_{ab} V_{ab} \delta_{jk} A_a^j(\hat{\mathbf{n}}_p, \mathbf{v}) A_b^{k*}(\hat{\mathbf{n}}_p, \mathbf{v}) e^{-i2\pi\mathbf{u}_{ab}\cdot\hat{\mathbf{n}}_p}, \quad (3.16)$$

where

$$\mathcal{N}^2 \equiv \sum_{ab} \delta_{jk} A_a^j(\hat{\mathbf{n}}_p, \mathbf{v}) A_b^{k*}(\hat{\mathbf{n}}_p, \mathbf{v}) \delta_{lm} A_a^l(\hat{\mathbf{n}}_p, \mathbf{v}) A_b^{m*}(\hat{\mathbf{n}}_p, \mathbf{v}). \quad (3.17)$$

and therefore the beam weights  $w_p^{ab}(\hat{\mathbf{n}}_p)$  are

$$w_p^{ab}(\hat{\mathbf{n}}_p) = \frac{1}{\mathcal{N}} \delta_{jk} A_a^j(\hat{\mathbf{n}}_p, \mathbf{v}) A_b^{k*}(\hat{\mathbf{n}}_p, \mathbf{v}) e^{-i2\pi\mathbf{u}_{ab}\cdot\hat{\mathbf{n}}_p}. \quad (3.18)$$

This is the basic structure we can follow to create formed beams pointing in

a particular direction on the sky, e.g. to track pulsars or to find FRBs. Equation 3.16 is optimized to produce the maximum signal, not necessarily the maximum  $S/N$ , but can be optimized to create the maximum  $S/N$  as outlined in Masui et al. (2019).

Though effective, the basic pointed formed beam is computationally expensive, scaling as  $\mathcal{O}(N^2)$ , with  $N$  the number of inputs. This makes a large number of formed beams from a large number of antennae untenable. However, for the special case of redundant arrays, we can apply Fourier Transforms to decrease our computational costs. This method was proposed by Tegmark & Zaldarriaga (2009, 2010), and expanded to allow for irregular (i.e. not square) arrays by Morales (2011).

Redundant arrays are those where antennae are identical and are evenly spaced. This means that every combination of inputs with the same spatial separation  $\delta$  between antennae nominally has the same visibility (up to a complex calibration factor).

We imagine a redundant array which is also one-dimensional so that  $\mathbf{u}_{ab} = \hat{\mathbf{x}} = d(a-b)/\lambda$  and observes a one-dimensional sky so that  $\hat{\mathbf{n}} \cdot \hat{\mathbf{x}} = \sin \theta$ . This is less over-simplified than it may initially appear; though CHIME is a 2-D array, the CHIME/FRB beamformer uses FFT beamforming in only one direction. This assumption allows us to simplify the pointed beam model and its weights to

$$b_p(\theta_p) = \frac{1}{n_{ant}} \sum_{ab} e^{-i2\pi(a-b)(d/\lambda) \sin \theta_p} V_{ab}, \quad (3.19)$$

and

$$w_p^{ab}(\theta_p) = \frac{1}{n_{ant}} e^{-i2\pi(a-b)(d/\lambda) \sin \theta_p}. \quad (3.20)$$

These weights can be factorized so that

$$w_p^a(\theta_p) = \frac{1}{\sqrt{n_{ant}}} e^{-i2\pi a(d/\lambda) \sin \theta_p}. \quad (3.21)$$

The format of Equation 3.21 is suggestive of a Fourier Transform.<sup>2</sup> Indeed, we can form  $M$  beams from our visibility data created from  $n_{ant}$  by taking a spatial FFT of

---

<sup>2</sup>And there's nothing a radio astronomer loves like a Fourier Transform.

the pre-correlation antenna signals  $\eta_a[t]$ . We zero-pad the data so that the length is  $M$  before taking the FFT.

The formed beam model is now

$$b_p(\theta_A) = \frac{1}{n_{samp}} \sum_t \left| \sum_a \eta_a[t] \frac{1}{\sqrt{n_{ant}}} e^{-i2\pi Aa/M} \right|^2 \quad (3.22)$$

and the weights are

$$w_p^a(\theta_A) = \frac{1}{\sqrt{n_{ant}}} e^{-ie\pi Aa/M}, \quad (3.23)$$

where  $M$  is the total number of formed beams,  $A$  is the index of a formed beam,  $a$  is the index of an antenna, and the sum over  $a$  is a Fourier Transform.

FFT formed beams are static and point to a fixed but tunable angle on the sky, which is referred to as the steering angle. The steering angle for beam  $A$ ,  $\theta_A$ , is

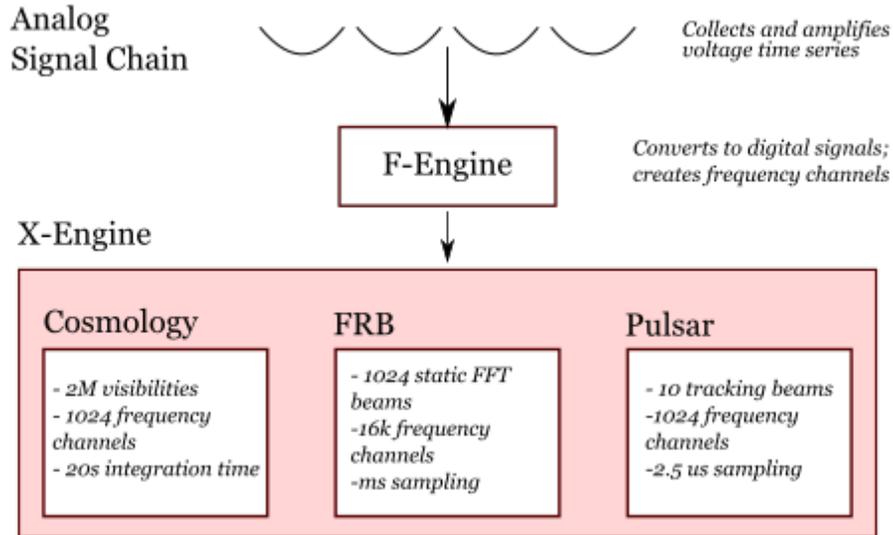
$$\sin \theta_A = \frac{A\lambda}{Md} \quad (3.24)$$

This FFT beamforming technique can also be combined with conventional beamforming to create multi-dimensional arrays. This is the technique used in CHIME/FRB to create four columns of 256 beams each.

In addition to static beamforming, we can also use beamforming to facilitate digital tracking, simulating the effect of a single dish telescope following the source across the sky. This is done using conventional (not FFT) beamforming in the CHIME/Pulsar system. We create 10 tied-array pointed beams which are rapidly re-pointed to account for the Earth's rotation and keep the desired source centered in the formed beam as it passes through the primary beam. This allows us to make more detailed observations and to have more control over our observations.

### 3.3 CHIME technical information

The three CHIME systems share a signal path, with separate backends stemming from the same frontend. Figure 3.1 shows a schematic of the path through the CHIME system, demonstrating the joint analog signal chain and F-engine as well as the multiple functions of the X-engine.



**Figure 3.1:** A block diagram demonstrating the path signals take through CHIME. All projects share the same front-end and F-engine; projects diverge in the GPU-based X-engine.

### 3.3.1 The Analog Front-end

CHIME has four cylindrical reflectors. Each is 100 m long by 20 m wide, though only 80 m of the cylinder’s length is illuminated. The cylinders are made of a durable steel mesh with 16-mm openings. This mesh allows snow and debris to fall through the reflector surface while remaining opaque to the roughly 1 m wavelength radio signals CHIME observes. As CHIME is a transit telescope, neither the cylinders nor the antennae affixed to them move; different portions of the sky become visible to CHIME as the Earth rotates. The CHIME primary beam is roughly cylindrical, with a north-south field-of-view of  $\sim 110^\circ$  and an east-west field-of-view of  $\sim 2.5 - 1.3^\circ$ . (The beam is wider at lower frequencies.)

CHIME has 256 dual polarization receivers per cylinder, for a total of 2048 spatial inputs. The antennae are cloverleaf-shaped and were designed specifically for CHIME. The two polarizations have equivalent beamwidth, and excellent behaviour throughout the CHIME band (Deng et al., 2017). These antennae do not cover the entire north-south extent of the cylinders; only 80 m of the cylinder’s



**Figure 3.2:** The Canadian Hydrogen Intensity Mapping Experiment, as seen from the air. Photo by J. Richard Shaw.

length are illuminated. They are positioned along a central focal line, at a focal length of 5 m (CHIME/FRB Collaboration, 2018).

The primary beam structure of the CHIME is complicated. It resembles a sinc function, with most of the power concentrated in the central lobe of the beam. However, the sidelobes also provide a substantial of the observed power, to a degree larger than is typical for single dish radio telescopes. This is an important consideration for FRB measurements, as seeing a source in the sidelobes of the CHIME primary beam can result in the source appearing in a different CHIME/FRB formed beam than its true sky position would suggest.

The next portion of the analog system is low noise amplifiers (LNAs), which are directly connected to the antennae, one per polarization. Amplifiers are an important component of radio telescopes, as the signal level of astrophysical signals is in general faint and requires amplification. Interestingly, the small, inexpensive LNAs used in CHIME were enabled by advances in cell phone technology.<sup>3</sup>

The cloverleaf antennae and the LNAs are the only portions of the system located on the CHIME focal line. Power is sent to the focal line and voltage data are

---

<sup>3</sup>Technology giveth and technology taketh away.

returned from it using 50 m coaxial cables from the LNAs on the focal line to the next stages of the signal chain in two receiver huts. All cables are of equal length to ensure that the thermal and loss properties of these cables are equivalent for all inputs.

The final stage of the CHIME analog system is the filter amplifiers (FLAs). These amplifiers are located in the receiver huts, which primarily house the F-engine. The FLAs provide more amplifications for the signals and apply a bandpass-filter (CHIME/FRB Collaboration, 2018).

### **3.3.2 The F-Engine**

CHIME uses a hybrid FX correlator, composed of FPGAs (in the F-engine) and GPUs (in the X-engine).

The first stage is the F-engine, where data are converted from analog to digital signals and frequency channels are introduced. The F-Engine is located in two 20-ft steel shipping containers, each managing signals from half of the array. It is composed of 128 ICE motherboards (see Bandura et al. (2016b) for more discussion). On each board, sixteen amplified and filtered inputs are digitized, then converts the voltage timestream into 1024 frequency channels with  $2.56 \mu\text{s}$  time resolution. At this stage, the data rate is 13.1 Tb/s. The F-Engine is discussed in greater detail in (Bandura et al., 2016a).

### **3.3.3 The X-Engine**

The second stage of the CHIME correlator is the GPU based X-engine portion of the correlator. The X-engine takes its name from the primary CHIME task: to spatially correlate inputs from all CHIME antennae.

The GPUs actually simultaneously complete 3 tasks: this spatial correlation, Fast Fourier Transform (FFT) beamforming for the CHIME/FRB system, and tracking beamforming for CHIME/Pulsar. From this point, CHIME data separates into the three distinct paths. The technical details of the CHIME X-engine are discussed in Denman et al. (2020).

### **CHIME Visibility Generation**

The name X-engine refers to the system’s primary task: the spatial correlation of all 2048 inputs. Data flows from the F-engine as voltages, sampled at  $2.56 \mu\text{s}$  with 1024 frequency channels. Each of the 256 GPU nodes processes four frequency channels, creating spatial correlations between all physical inputs and integrating in time to create the visibility dataset for CHIME. This is the ultimate product for the X-engine in CHIME, but further work is conducted for CHIME/FRB and CHIME/Pulsar to create formed beams.

### **CHIME/FRB FFT Beamforming**

CHIME/FRB is enabled by a large-scale implementation of FFT beamforming. Beamforming is the process of combining spatial inputs in a weighted manner to create a new, digital telescope beam within the interferometer by weighting visibilities. This is a straightforward concept, but it is computationally expensive, scaling as  $\mathcal{O}(N^2)$ , where  $N$  is the number of telescope inputs. For a system with few inputs (e.g., an array of 10 antennae) or for systems which make relatively few formed beams (e.g., the 10 tracking beams of CHIME/Pulsar), this does not pose a problem. On the other hand, for a system such as CHIME/FRB with 2048 inputs and 1024 formed beams, computational cost in beamforming is a substantial hurdle.

We overcome this challenge by using FFT beamforming, as outlined in detail in Masui et al. (2019) and briefly in Section 3.2. This modern implementation of beamforming scales as  $\mathcal{O}(N \log(N))$ , making it a much more feasible computational cost for a large-scale experiment such as CHIME/FRB. In fact, we use hybrid approach: we generate 256 beams in the north-south direction organized in four columns in the east-west direction. Beamforming in the north-south direction is FFT beamforming, but the east-west direction uses conventional beamforming. The formed beams are spaced evenly in  $\sin \theta$ . The spacing is tunable, but set to  $-60^\circ$  to  $60^\circ$ , where  $\theta$  is the angle from the zenith along the meridian CHIME/FRB Collaboration (2021).

Implementing this FFT beamforming exactly as presented in Masui et al. (2019) would lead to poor results for a telescope such as CHIME, as there is substantial

chromatic smearing over the 400 MHz of bandwidth. For a detailed discussion of this effect and its consequences for beam placement, see (Ng et al., 2017). Therefore, our implementation zero-pads and creates 256 additional beams which are redundant with the 256 desired beams. This increases the computational cost from  $\mathcal{O}(N \log_2 N)$  to  $\mathcal{O}(NP \log_2 N)$ , where  $P$  is the padding factor we apply.  $P = 2$  for CHIME/FRB. As beam width is frequency dependent, the beam coverage varies across the frequency band.<sup>4</sup>

CHIME has an intrinsic frequency resolution of 390 kHz, with 1024 frequency channels over its 400 MHz bandwidth. For many applications, this is sufficient frequency resolution, but FRB searching requires more precise frequency resolution as incoherent de-dispersion at 400-800 MHz requires very small frequency channels. Therefore, during the beamforming process, we also increase the frequency resolution from 1024 channels with 390 kHz resolution to 16k channels with 24.4 kHz resolution. We do this by collecting 128 voltage samples, each  $2.56\mu\text{s}$  long, and Fourier Transforming the set. We then downsample in frequency by a factor of 8 (Ng et al., 2017; CHIME/FRB Collaboration, 2018).

The final output of the CHIME/FRB beamforming kernel within the X-engine is 1024 static beams, with 16k frequency channels and approximately 1 ms time resolution.

### **CHIME/Pulsar Tracking Beamforming**

The CHIME/Pulsar beamformer, described in more detail in (Amiri et al., 2021), takes a different approach from the CHIME/FRB beamformer. CHIME/FRB uses FFT beamforming to create a large array of static beams, whereas the CHIME/Pulsar beamformer creates ten tracking beams, which digitally alter the linear combinations of visibilities over time to follow a source across the sky. As these beams are conventionally formed rather than FFT formed, they are more computationally expensive on a per-beam basis than CHIME/FRB beams; we are constrained to only forming 10 such beams, though upgrade efforts to add an eleventh are ongoing. The output of the CHIME/Pulsar beamforming kernel is 10 tracking beams with 1024 frequency channels and  $2.56\mu\text{s}$  resolution.

---

<sup>4</sup>This fact is briefly addressed in discussion of PSR J0209+5759 in Chapter 5.

Before CHIME/Pulsar data passes from the X-engine to the CHIME/Pulsar nodes, it is packetized. Data are scaled down and encoded as (4+4) bit complex numbers. Each packet includes 625 samples for one polarization for each of the four frequency channels processed by a single node. The total data rate from the X-engine to the CHIME/Pulsar nodes is  $\sim 64$  Gbps.

### 3.4 CHIME/FRB System Overview

The CHIME/FRB backend has one main goal: to find FRBs. To accomplish this goal, we must first de-disperse the data and find bursts, then determine if they are astrophysical, then determine if they are in fact FRBs. This section summarizes information found in CHIME/FRB Collaboration (2018).

The first and most important step in finding FRBs is to de-disperse the input data. CHIME/FRB uses a tree de-dispersion system (Taylor, 1974) known as `bonsai`. In addition to de-dispersing the data `bonsai` also performs a spectral weighting and searches for peaks. We search for DMs up to  $13,000 \text{ pc cm}^{-3}$ , pulse widths up to 100 ms, and we use two trial spectral indices. This system runs on 128 CPU-nodes. `bonsai` is the culmination of many years of de-dispersion development, and its unique features are discussed in CHIME/FRB Collaboration (2018).

The next stage of CHIME/FRB searching is the all-important process of RFI mitigation. RFI is a thorn in the side of all radio astronomers but is particularly troublesome for FRB searchers. FRBs are short duration, bright transients, albeit at high DM while RFI is theoretically at  $DM = 0$ . In practice, some kinds of RFI can appear to be de-dispersed at an astrophysical DM, but plotting the intensity data from these triggers at  $DM = 0$  provides characteristic, identifiable signals. Therefore, the process to remove RFI signals from CHIME/FRB data are necessarily both rigorous and complex. There are several stages, involving removal of known bad frequency bands (about 20% of the CHIME frequency band) and searches for transient RFI. We make extensive use of cutting edge machine learning techniques to enable the system to progressively improve its ability to find man-made signals. RFI sifting is discussed in greater detail in (CHIME/FRB Collaboration, 2018); a complete discussion of the algorithms used is beyond the scope of this work.

After candidates are identified as astrophysical (not RFI), we must determine

whether they are in fact potential FRBs. Bright candidate FRBs are often detected simultaneously in several of the 1024 static CHIME/FRB beams. We first group together these multi-beam detections so that they are considered as a single event and not as several events.

Then, we classify the astrophysical sources based on whether their DM indicates they are extragalactic. `bonsai` determines a DM and DM uncertainty for each CHIME/FRB candidate. We determine the maximum Galactic DM in the direction of the candidates using both the Cordes & Lazio (2002) and Yao et al. (2017) Galactic DM maps and calculate the difference between the two estimates. We then add in quadrature the uncertainty in measured DM and the difference between the two maximum Galactic DM estimates. This is our total uncertainty  $\sigma$ . Sources where the measured DM exceeds the both maximum Galactic DM models by at least  $5\sigma$  are classified as Extragalactic. If the measured DM is  $2\sigma$  to  $5\sigma$  in excess of the maximum Galactic DM, the source is classified as Ambiguous. If the measured DM is less than  $2\sigma$  in excess of the maximum Galactic DM, the source is classified as Galactic CHIME/FRB Collaboration (2018). Intensity data are saved only for sources with Ambiguous or Extragalactic DM.

Finally, we must determine whether these sources are in fact *new* astrophysical sources. The CHIME/FRB backend detects a number of non-FRB astrophysical sources every day, most importantly every pulsar with sufficiently bright single pulses. However, these sources are not of interest to CHIME/FRB, so the Known Sources Sifter maintains a list of known pulsars and removes from candidacy bursts which appear to derive from known pulsars. The same process allows us to flag potential repeat bursts from known FRBs.

Ultimately, if the candidate appears to be a new Extragalactic or Ambiguous astrophysical event, intensity data are saved to disk. If the source does not meet these conditions, data are not saved. For candidates above a tunable  $S/N$  parameter, we trigger a baseband “dump.” CHIME/FRB maintains a 35.5 s buffer of raw voltage data, siphoned off prior to the X-Engine. If we see a sufficiently bright FRB candidate, the system will automatically save the voltage data corresponding to that event. The CHIME/FRB baseband system is discussed in depth in Michilli et al. (2021).

Although only potentially extragalactic astrophysical sources have intensity

callback data saved, metadata for every event CHIME/FRB sees is saved, including the time it is observed, its DM, and its beam location. This can be an invaluable resource for examining system performance and Galactic sources; see Chapter 5 for much more detail on this second application.

After intensity data are saved to disk, we apply one more automated check. We use convolutional neural networks to predict whether a trigger is likely to be an FRB or an RFI based on previous triggers classified by humans (Yadav, 2020). Following automated identification, each candidate event is examined by a human “data tsar(itsa)” who classifies it based on whether or not is indeed a promising FRB candidate. In spite of our best efforts, there are false positives; sometimes a known source is detected in the beam’s sidelobe (and thus does not initially appear to be in its proper position), sometimes RFI are not caught by any of the automated filters. Finding recurring sources of RFI and unfiltered known sources can be key strategies for improving algorithmic performance.

As it is fully digital, the CHIME/FRB system is highly tunable; parameters can be changed at will (although consistency is valuable for fully understanding the system). For example, we can change the SNR cuts determining callbacks, search parameters, and even beam placement.

CHIME/FRB, since the start of science operations in July 2018, has been a tremendously productive instrument. In January 2019, we published our first 13 new FRBs, including a new repeater (the second ever), detected over the course of three weeks in the summer of 2018 (CHIME/FRB Collaboration et al., 2019b,a). Later in 2019, we released a new set of repeaters and a detection of the first repeater, FRB 121102. In 2020, we released another set of repeaters as well as a surprise detection of SGR 1935+2154, a Galactic magnetar which provides tantalizing evidence of a possible correspondence between magnetars and FRBS (CHIME/FRB Collaboration, 2020a). In 2021, we published the first CHIME/FRB catalog, including over 500 FRBs (CHIME/FRB Collaboration, 2021).

### **3.4.1 CHIME/Pulsar**

CHIME/Pulsar is a pulsar timing backend for CHIME. It is adaptable and versatile, with the ability to conduct long-term observation programs with minimal

human intervention or to quickly pivot operations when exciting sources arise. CHIME/Pulsar is pursuing several pulsar science goals, but most relevant to this work are high precision pulsar timing observations for NANOGrav (discussed in greater detail in Chapter 4) and initial pulsar timing solutions for new pulsars discovered with CHIME/FRB (discussed in greater detail in Chapter 5). However, work is also ongoing on rotation measure studies examining pulsar polarization, nulling studies, single pulse studies for RRATs, and binary neutron star systems (Fonseca et al., 2021; Good et al., 2021; Ng et al., 2020; Ng et al., 2020). The instrument details are discussed in Amiri et al. (2021).

The question of determining which sources will be observed when is handled automatically by the CHIME/Pulsar scheduler. CHIME/Pulsar can observe up to 10 sources at a time and can observe in either fold-mode or search-mode. Though observing 10 sources simultaneously allows us to observe a large percentage of Northern pulsars each day, we sometimes need to make decisions about which pulsars to observe. The scheduler is designed to manage this problem. We set priorities for each CHIME/Pulsar source, telling the scheduler how important observations of each source are. Top priority sources always include NANOGrav MSPs. The scheduler will always try to incorporate sources in order of priority, but will also account for elapsed time since previous observations. This means that we avoid the scenario where we observe the same 10 sources every day during a given block of sidereal time to the exclusion of all others. Additionally, a lower priority source, which has not been observed in a long time will automatically increase in priority to ensure that we do ultimately collect data.

The heart of the pulsar backend, and the piece which most closely compares with similar pulsar backends, is the 10 GPU nodes. When data arrives at the pulsar backend, it is de-packetized and re-ordered. The X-engine nodes each process a set of four frequencies, but this frequency breakdown is not useful for pulsar analysis. This process is discussed in Amiri et al. (2021).

From this point, the CHIME/Pulsar backend is very similar to other modern pulsar backends, making use of standardized processing software. Data are then coherently de-dispersed using DSPSR (van Straten & Bailes, 2011). From here, data are treated slightly differently based on whether they are search-mode or fold-mode data. Coherent search-mode data are integrated to  $327 \mu\text{s}$  and written out.

Fold-mode data retain  $2.56 \mu\text{s}$  time resolution and are folded before being written to disk, using PSRCHIVE. Filterbank data are written out using a custom software described in Naidu et al. (2015).

## Chapter 4

# Integrating CHIME/Pulsar data into the NANOGrav dataset

*That the study of mathematics and their application to astronomy are full of interest will be allowed by all who have devoted their time and attention to these pursuits; and they only can estimate the delight of arriving at truth, whether it be in the discovery of a world, or of a new property of numbers. — Mary Somerville*

### 4.1 Introduction and Motivation

In this chapter, we will discuss ongoing efforts to combine CHIME/Pulsar timing data with the NANOGrav long-term precision timing datasets. We first discuss the motivation for this work. We then discuss the process for making TOAs from CHIME/Pulsar data and for creating combined timing models. Finally, we present the results of this analysis for a set of eight pulsars excerpted from the NANOGrav pulsar timing array.

CHIME/Pulsar data is a natural fit for inclusion in the NANOGrav dataset for several reasons.

First, high cadence is desirable for high precision pulsar timing. This is obviously true in a general sense: the smaller the gaps are between data points, the better parameters can be fit. High cadence data has specifically been shown to be advantageous in e.g., Lam (2018). In addition to general high cadence, highly tun-

able flexible high cadence is also a powerful tool for conducting with continuous wave searches. If we want to target specific possible supermassive black hole binaries for observations, we need to be able to emphasize specific pulsars within our dataset. Having the option to go up to daily cadence on all NANOGrav sources with CHIME/Pulsar enables this in two ways: first, it obviously provides a large set of data on those sources and second, it allows us to free up resources to conduct more observations of high value sources with other telescopes.

Second, the inclusion of a large set of CHIME/Pulsar TOAs should be expected to improve the overall timing results in the NANOGrav dataset. The uncertainty in each TOA is given by Equation 2.21 (Lorimer & Kramer, 2005) and depends on the pulsar’s properties, the system sensitivity, observing bandwidth, and observing time. To date, NANOGrav has used the Green Bank Telescope (GBT), the Arecibo Observatory (AO), and the Very Large Array (VLA) for pulsar timing. These are more sensitive telescopes with lower system sensitivity thresholds, but CHIME/Pulsar is competitive and advantaged by its large fractional bandwidth and daily cadence.

CHIME is a good match to pulsar astronomy scientifically. Pulsars signals are in general stronger at lower frequency (due to their spectral index). In some cases, this is inhibited by increased scattering in the ISM, but on the whole low frequency observations of pulsars can provide excellent timing precision. This is particularly true for sources with low dispersion measure, such as PSR J1744-1134, discussed later in this chapter (Lam et al., 2018a). Lower frequency observations are also more sensitive to variations in the ISM, which can be seen as DM variations. Therefore, beyond its direct applicability to NANOGrav timing, CHIME/Pulsar data will enable us to better understand the ISM. Understanding these variations is important to build models of DM variations and noise models for NANOGrav timing. Such noise models are required to tease out the red noise signal indicative of gravitational waves, so this is an important problem.

One of the exciting potentials but also challenges of CHIME/Pulsar data is the large relative bandwidth. With continuous bandwidth between 400–800 MHz, CHIME/Pulsar’s relative bandwidth is 0.66. This means that the process of sub-banding, partially compressing observations in frequency, fails to leverage the wideband nature of CHIME/Pulsar data. For this initial result, we present sub-

banded, narrow-band TOAs, but we look forward to incorporating wideband timing in the near future.

Our concerns about CHIME/Pulsar’s bandwidth are only a tiny microcosm of the bandwidth challenge we will have in the growing Ultra-Wideband (UWB) receiver era. The GBT UWB for example, will have a continuous frequency range between 0.7 and 4.2 GHz. Commissioning for this receiver is expected to begin in 2022. Pulsar timing is a major motivation for the UWB receiver, and in NANOGrav, we look forward to both using our telescope time more efficiently and improving our understanding of frequency dependent pulse properties. However, this is much broader bandwidth than is currently used in NANOGrav timing, and will require us to re-examine our analysis procedure. CHIME/Pulsar provides a test-bed for wideband analysis techniques.

One major change for future NANOGrav analysis is the end of observations from AO, following the facility’s collapse in late 2020. AO has been a huge part of NANOGrav from the beginning, and its loss creates new challenges for data collection. At present we have moved all NANOGrav pulsars to GBT, ceasing GBT 800 MHz observations for all pulsars with declination  $+10^\circ$ , and shifting our low-frequency observations to CHIME/Pulsar. This change increases the importance and the urgency attached to CHIME/Pulsar-NANOGrav data combination.

## **4.2 The Unique Challenges of CHIME/Pulsar Timing**

In determining timing solutions for CHIME/Pulsar data, we have several problems that are either not present or much less pernicious in the more conventional NANOGrav data.

The first is our comparatively poor understanding of CHIME/Pulsar as an instrument. With GBT, we now have more than 20 years of understanding the instrument’s beam and polarization properties, and with Arecibo’s Gregorian optics, we had about 25 years. Even our backends, PUPPI and GUPPI and now VEGAS have years of data, including carefully designed overlap observations with the previous generation of backends. NANOGrav data have never leveraged an entirely new observatory.

It has taken much time and effort to even start work on CHIME/Pulsar tim-

ing due to the long ramp-up to system stability for the CHIME/Pulsar instrument. The system has finally become stable and has reached acceptable sensitivity levels, allowing us to pursue more detailed analyses. That being said, we do not dig deeply into some of the more complex aspects of the instrument, like polarization, in this work. This is designed as a first effort at creating timing solutions that can be incorporated into a PTA with CHIME/Pulsar.

We also recognize the ongoing, dynamic nature of RFI excision at the CHIME site. Unlike other facilities used in pulsar timing, CHIME/Pulsar is not in a pristine location by any means. We do experience transient RFI, sometimes very pernicious transient RFI. Due to telecommunications constraints, many channels within our band are permanently contaminated, and the level of contamination is continually increasing. RFI excision, including time-dependent RFI excision, is a key component of our analysis strategy. As this is a starting point analysis, we use relatively simple RFI removal techniques, but more complex RFI removal for CHIME/Pulsar is underway.

CHIME/Pulsar data also differs from conventional NANOGrav data in its cadence. Even the high cadence timing program conducted with GBT has been intended only to obtain weekly observations of select pulsars. The default NANOGrav observing program generally observes pulsars once per month with each receiver (800 MHz and L-band at GBT, L-band and S-band<sup>1</sup> at Arecibo, and L-band and S-band at VLA). In contrast, CHIME/Pulsar observers nearly every pulsar every day. Therefore, a duration which would generate 2 wideband or 2 sets of sub-banded TOAs will have 30 wideband or 30 sets of sub-banded TOAs. This can make our analysis a bit odd, simply because it biases us towards this huge set of less sensitive TOAs. This is not necessarily a problem, but it is a new paradigm and requires thoughtfulness in combining the TOAs.

### 4.3 Creating CHIME/Pulsar TOAs

NANOGrav data follows a standard timing procedure, to allow multiple receivers and backends to be integrated into a single dataset consistently. To add CHIME/Pulsar data to this process, we had to first pre-process this data then create narrowband

---

<sup>1</sup>S-band is the frequency band from 2-4 GHz

TOAs.

### 4.3.1 RFI Removal

The first step in processing CHIME/Pulsar data for both narrow and wideband TOAs is to remove RFI. As discussed in Chapter 2, we accomplished this primarily using the PSRCHIVE tool `paz`. We use a standardized “kill” file, which lists frequency channels known to be contaminated with RFI at CHIME and the built-in `paz` median smoothed difference zapping algorithm.

### 4.3.2 Creating TOAs

After RFI removal, data files are “scrunched” using PSRCHIVE command `pam`. Scrunching refers to averaging data in frequency or time. Data can be scrunched to either a single time or frequency bin or a fixed number of bins. Folded observations of pulsars include many sub-integrations in time and, in the case of CHIME/Pulsar 1024 frequency channels. In principle, we could create a TOA for every sub-integration and for every frequency channel. However, as this would create thousands of TOAs per observation, each with low  $S/N$ , we do not. Instead, we scrunch the data in time and frequency, reducing its resolution. In time, we scrunch to a single point. In frequency, we scrunch to 32 frequency channels, each with a width of 12.48 MHz. This bandwidth is sufficiently narrow that the pulsar’s profile does not evolve substantially within a single scrunched file, but we do not have an overwhelming number of data points.

Following sub-banding, we create a standard profile to be used in generating narrowband TOAs. As discussed in Chapter 2, TOAs are generated by cross-correlating a standard profile with the observed profile in Fourier space. This standard profile is created by adding together and smoothing a representative sample of CHIME/Pulsar data.

We then use PSRCHIVE command `pat` to generate TOAs with the template and files provided. PSRCHIVE includes a number of different algorithms for calculating TOAs; we use the Fourier Domain Markov Chain Monte Carlo (FDM) model as is used in NANOGrav analysis. This model is based on Taylor (1992), but extends this algorithm to include an MCMC analysis. At this stage, we can also

add a variety of flags to the TOAs, to allow for further analysis, including  $S/N$  and information about the receiver and backend pair used to make the observation.

## 4.4 Creating the Timing Model

In this work, we create two separate and distinct timing models following the same procedures. The first is a CHIME/Pulsar only dataset and the second is a combined NANOGrav and CHIME/Pulsar dataset. The timing models are created independently for the two datasets, though the combined solution is used to inform some parameters in the CHIME/Pulsar-only solution, due to its much longer timespan.

The CHIME/Pulsar dataset includes only the data taken by CHIME with TOAs generated as described above. The NANOGrav and CHIME/Pulsar combined dataset includes both the CHIME/Pulsar data and also the preliminary 15 year NANOGrav dataset. The NANOGrav dataset for sources used in this work includes incorporates data from the GBT L-band and 800 MHz receivers and the AO L-wide and S-wide receivers. It also includes two sets of backends: the GASP and Green Bank Ultimate Pulsar Processing Instrument (GUPPI) backends at AO and the ASP and Puerto Rico Ultimate Pulsar Processing Instrument (PUPPI) backends at GBT. The GBT and AO TOAs are created with the NANOGrav TOA generation pipeline, `nanopipe`. Both models use as a starting point the NANOGrav 12.5 year narrowband results presented in Alam et al. (2021a).

Once we have created a set of CHIME/Pulsar TOAs, we can begin the process of fitting a timing model, following the procedure laid out in Chapter 2. Our TOAs can be analyzed with any pulsar timing software package. There are three major options: TEMPO, TEMPO2, and PINT. TEMPO (Nice et al., 2015) is a Fortran-based software, used since the early days of pulsar astronomy. It is also the primary software used in previous NANOGrav releases (NANOGrav Collaboration et al., 2015; Arzoumanian et al., 2018b; Alam et al., 2021a). TEMPO2 (Hobbs et al., 2006; Edwards et al., 2006; Hobbs et al., 2009) is a C-based software, also commonly used for pulsar timing. PINT (Luo et al., 2021) is a python-based software and the newest pulsar timing software. It was first used by NANOGrav in the 12.5 year dataset (Alam et al., 2021a) and will be the primary timing software for the NANOGrav 15-year dataset. With this in mind, we use PINT as our pulsar timing

software.

This decision allows us to leverage the NANOGrav computational resources, particularly the NANOGrav Jupyter Notebook Server. The notebook server provides pre-installed versions of pulsar timing software such as PINT and ENTERPRISE, as well as additional standard NANOGrav utilities. This allows our analysis to remain closely in step with the in-progress analysis effort to create the NANOGrav 15 year dataset. This is ideal as, although CHIME/Pulsar data will not appear in that release, we are using the 15 year dataset as the dataset to combine with CHIME/Pulsar. Additionally, these notebooks, developed by the NANOGrav Timing Working Group and Cyber-Infrastructure Working Group are efficient and user-friendly graphical interfaces for pulsar timing. They also ensure a reproducible analysis; as TEMPO and TEMPO2 are command line programs, reproducibility has traditionally been a challenge.

One of the challenges for NANOGrav data is including only “good” TOAs, but this analysis raises a question: What does it mean for something to be a “good” TOA, and how do we decide this in the least biased way possible? We apply an  $S/N$  cut, eliminating all TOAs with  $S/N < 8$  for narrow-band timing. This is particularly important for CHIME/Pulsar TOAs, where there are many TOAs that fail to reach this threshold. After these checks, there are still often outlier TOAs. In this work, we remove these TOAs manually, though in published NANOGrav data, this excision is automated (Arzoumanian et al., 2018b; Alam et al., 2021a). In light of the large number of TOAs in the dataset, we generally take an aggressive approach, preferring to lose a few good TOAs than to retain poor ones.

Once TOAs are generated and the dataset is cleaned, we follow NANOGrav convention for fitting timing models to sources. This means that for all combined datasets we fit at minimum the two spin parameters (period and period derivative) and five astrometric parameters (two proper motion parameters, two position parameters, and timing parallax) Alam et al. (2021a). Each of these parameters was introduced in more depth in Chapter 2.

For all sources in both the joint dataset and the CHIME/Pulsar only dataset, we use the TT(BIPM2019) clock standard and the DE440 ephemeris (Park et al., 2021).

Most of our sources are in binaries, and we follow the convention laid out

in the NANOGrav 12.5 year dataset for determining which binary parameters to include as well. For all binary sources, we fit for the Keplerian parameters (orbital period  $P_b$ , semi-major axis  $x$ , and either eccentricity  $e$ , longitude of periastron  $\omega$ , and epoch of periastron passage  $T_0$  or the epoch of the ascending node  $T_{\text{asc}}$  and the two Laplace-Lagrange parameters  $\varepsilon_1$  and  $\varepsilon_2$  (Alam et al., 2021a). As in the 12.5 year dataset, we select our binary model based on eccentricity and need for post-Keplerian parameters. We use the ELL1 model of Lange et al. (2001) with low eccentricity binaries. For systems with marginal Shapiro delay, we use the ELL1H model of Freire & Wex (2010). This model adds two additional parameters  $h_3$  and  $h_4$ . For systems with higher eccentricity, NANOGrav use the DD binary model of Damour & Deruelle (1986), though no such systems are included in this test set. Finally, for J1713+0747, we use the DDK binary model of Kopeikin (1995, 1996); van Straten et al. (2001), which is an extension of the DD binary model.

We use the Generalized Least Squares (GLS) fitter in PINT (Luo et al., 2021) to fit our timing models, to allow for the possibility of correlated noise in our data. GLS was first developed in Aitkin (1935), but first used in pulsar timing by Coles et al. (2011). In GLS fitting, we compensate for correlations by applying a linear transformation to the covariance matrix of the residuals to whiten the residuals (Coles et al., 2011).

The spin parameters, astrometric parameters, and binary parameters are described in more detail in Chapter 2 and have clear theoretical origin, from the pulsar’s spin, the motion of bodies in the solar system, or the orbital motion of a binary system. However, long-term, high precision pulsar timing requires us to add additional parameters to our pulsar model to account for the presence of multiple sources of data, variations in dispersion measure, frequency-dependent profile variations, and excess noise.

TOAs from different receivers are separated by arbitrary offsets; we fit arbitrary time “jumps” to bring all receivers into alignment. This is the same procedure used for combining receivers with the NANOGrav dataset, but is extended here to add time jumps for CHIME/Pulsar data. These jumps are not necessary in CHIME/Pulsar only datasets, as there is only one instrument. Although each pulsar beam can operate independently, all use the same analog front-end, and therefore are considered one instrument.

Although MSPs have consistent nominal DMs, they can also exhibit a variety of small variations about this nominal value. Previous work (Jones et al., 2017; Fonseca et al., 2021) has demonstrated the importance of including a model for short-timescale DM variations. Following in the footsteps of both Alam et al. (2021a) and Fonseca et al. (2021), we use the piecewise-constant DMX model to do this modeling. In creating a DMX model, we set a fixed window and fit for the difference between the observed and nominal DM during observations in that DMX window. The reported DM value then remains consistent at the nominal value, while we still account for temporal variations in DM. These variations can take a variety of forms, including stochastic variations and linear trends and can arise from a variety of physical mechanisms including planetary or pulsar motion and changes in the ISM. See Lam et al. (2016) for a more detailed discussion.

For both combined NANOGrav and CHIME/Pulsar data and for CHIME/Pulsar only data, we use the DMX bin used for that pulsar in previous NANOGrav observations (6.5 days for GBT and combined GBT and AO observations with GUPPI, 15 days for GBT observations with GASP, and 0.5 days for AO observations). Additionally, we use standard NANOGrav procedures for managing the effects of Solar wind on DM variation. Specifically, for pulsars where solar wind causes delay change of more than one microsecond in each bin, the bin is split. This primarily happens to pulsars near the plane of the ecliptic as they pass close to the Sun. Systematic studies of ideal DMX binning for CHIME/Pulsar data both in the NANOGrav dataset and in general are deferred to future work. As in the NANOGrav 12.5 year dataset (Alam et al., 2021a), we require that the ratio of the highest and lowest observing frequencies  $f_{\max}/f_{\min}$  be greater than 1.1 within each DMX time range. Though it may appear that the bin size decreases entirely at the introduction of the CHIME/Pulsar data, this is an illusion: daily CHIME/Pulsar observations mean that every 6.5 day span contains at least one and up to 7 observations, whereas many potential 6.5 day spans are skipped in monthly NANOGrav observations. This increases the density of DMX measurements without actually decreasing the bin width.

For all parameters besides the five required astrometric, two required spin, and five Keplerian binary parameters, we determine parameter significance with an F-test. As in Alam et al. (2021a), we require that  $p < 0.027$  for parameters to be

included. For binary parameters in the CHIME/Pulsar only analysis, we include all parameters included in the joint analysis, but fix values to the combined dataset value for parameters which are not significant or vary widely from the joint dataset.

Pulse profiles are frequency-dependent (Craft & Comella, 1968) and in long-timescale, high-precision pulsar timing experiments, we must account for these variation (Zhu et al., 2015a). In the case of the NANOGrav dataset and therefore in this work, we accomplish this using frequency dependent “FD” parameters. FD parameters are polynomial coefficients in log-frequency space and they account for frequency-dependent pulse shape and evolution (NANOGrav Collaboration et al., 2015; Alam et al., 2021a; Shapiro-Albert et al., 2021). We apply up to five such parameters, depending on the results of the F-test.

We calculate a noise model, using the PTA software ENTERPRISE (Ellis et al., 2019) and BAYESEPHM (Vallisneri et al., 2020). Again, we gain follow the framework of the 12.5 year dataset (Alam et al., 2021a). As discussed in Chapter 2, this analysis includes four total components: EQUAD, ECORR, EFAC, and (for some pulsars) red noise.

## **4.5 First Joint Timing Results from CHIME/Pulsar**

### **4.5.1 General Discussion**

The work presented here is merely the first stage of CHIME/Pulsar and NANOGrav data combination. In time, we will complete data combination for all NANOGrav sources observed with CHIME/Pulsar. However, for this work, we focus on providing analysis of eight representative test pulsars. The pulsars are well-understood in the NANOGrav dataset and represent a variety of pulsar properties, with the goal of creating a small but representative set.

These eight pulsars span the CHIME/Pulsar declination range, with PSR J1744-1134 at the southern extreme and PSR J1125+7819 at the northern extreme. They incorporate data from both GBT and AO observations, though most are GBT pulsars. PSR J1713+0747 and PSR B1937+21 have been observed with both observatories.

These sources also incorporate the variety of binary models employed in NANOGrav

analysis. PSRs J0645+5158, J1744-1134, and B1937+21 are isolated pulsars. PSRs J0740+6620, J1012+5307, and J1125+7819 employ the ELL1 model for binary modeling. PSR J2145-0750 uses the extended ELL1 model, and PSR J1713+0747 uses the DDK model.

The CHIME/Pulsar data in this analysis was obtained between MJD 58600 (2019-04-27) and MJD 59072 (2020-08-11). The beginning of this date range was chosen to correspond to the resolution of a phase jump issue in the CHIME/Pulsar system, and the end date was chosen to correspond to the end date of the NANOGrav 15 year dataset, itself set by the final day of AO operations. The NANOGrav data presented here is from the intermediate “v0.9” 15 year release, generated in March 2021. The start dates for this data vary depending on when the pulsar was added to NANOGrav, and the end date is approximately April 6, 2020 (MJD 58945), the final day of operations for the GBT GUPPI backend. (The official end date of the 15 year dataset was determined after this intermediate dataset was generated.)

These solutions were determined by an iterative process. We began by finding rough solutions for only CHIME/Pulsar data, working from the published NANOGrav 12.5 year results, with changes informed by the in-progress NANOGrav 15 year results. Once we had satisfactory solutions for CHIME/Pulsar-only sets of TOAs, we integrated the CHIME/Pulsar TOAs with the NANOGrav TOAs. We determined satisfactory pre-noise solutions for this combined dataset, then used NANOGrav computational resources to compute ENTERPRISE noise models for the combined dataset. We verified these combined solutions, then used the combined solutions as a basis for refining the CHIME/Pulsar only solutions.

The CHIME/Pulsar data is substantially shorter in duration than the NANOGrav dataset, which in turn causes its astrometric parameters to be less reliable. With this in mind, we fixed parallax and proper motion to the combined dataset values for the final CHIME/Pulsar-only dataset. We therefore fit only spin parameters, source position, and binary parameters with CHIME/Pulsar only data. We also use the combined fit FD parameters. FD parameters are highly covariant with other parameters in the model, particularly DMX, and the approximately one year data span for CHIME/Pulsar-only data is not sufficient for a reliable fit. This covariance between FD parameters and DMX is challenging to fully understand.

Figures 4.1 and following show the residuals from the timing model fits for both the CHIME/Pulsar only and combined datasets. For each source, we present first the CHIME/Pulsar only dataset, then the combined. Each figure includes the residuals as a function of time, the averaged residuals as a function of time, and the DM model as a function of time. Binaries also include the residuals as a function of phase.

In Table 4.1, we present a summary of our timing model fit parameters, providing the number of fit parameters per source and their types as well as the root-mean-square (RMS) and weighted root-mean-square (WRMS) values of the timing residuals for the combined set and for the CHIME-portion of the combined set. The RMS values for CHIME/Pulsar data are systematically higher than for the combined dataset; this is because CHIME/Pulsar is less sensitive than the GBT and AO systems used in the combined dataset.

In Table 4.2, we present values for each pulsar’s period, period derivative, dispersion measure, and binary period where applicable. This table is supplemented by complete parameter files for each source in Appendix A. Table 4.2 and the par files are the models for CHIME/Pulsar data only, as the combined NANOGrav-CHIME/Pulsar dataset uses the not yet finalized NANOGrav 15 year dataset.

In Table 4.3 and Table 4.4, we present the ratio of the parameter uncertainties for each parameter in the CHIME/Pulsar only uncertainties to the the combined NANOGrav and CHIME/Pulsar dataset. The uncertainties are much smaller for the combined dataset than the CHIME/Pulsar only dataset. This is unsurprising; based on the timespan difference, excluding instrumental differences, we would expect the combined uncertainties for spin parameters to be smaller by a factor of as much as  $\sqrt{15}$ , due to increased timespan. Increasing the time span increases our ability to precisely determine parameters, and the combined dataset includes between seven and fifteen years of data, while the CHIME/Pulsar only dataset includes approximately one year of data.

Additionally, not all binary parameters can be reliably determined from the CHIME/Pulsar dataset only. A timespan of only 472 days is minimal in the context of timing many binary pulsars, and we would not expect to be able to create complete, reliable solutions. Therefore, we compare our CHIME/Pulsar results for binary parameters to the combined results. If any of binary parameters differs from

**Table 4.1:** Summary of joint NANOGrav and CHIME/Pulsar timing model fits, including number of spin, astrometric, and binary parameters, number of DMX bins, and frequency dependent profile bins, jumps, and RMS residual values for both the total dataset and the CHIME components.

Source	Total $N_{\text{TOA}}$	CHIME $N_{\text{TOA}}$	Number of fit Parameters						RMS ( $\mu\text{s}$ )		CHIME/Pulsar RMS ( $\mu\text{s}$ )	
			$S^a$	$A^b$	$B^c$	DMX <sup>d</sup>	FD <sup>e</sup>	$J^f$	RMS <sup>g</sup>	WRMS <sup>h</sup>	CHIME RMS <sup>i</sup>	CHIME WRMS <sup>j</sup>
J0645+5158	14076	3332	2	5	–	203	2	2	1.009	0.219	1.024	0.292
J0740+6620	17303	3681	2	5	7	138	1	2	1.632	0.473	1.878	0.792
J1012+5307	32493	5885	2	5	6	264	4	2	1.998	0.461	2.728	2.097
J1125+7819	14193	5543	2	5	5	116	2	2	2.862	1.017	3.014	1.876
J1713+0747	60395	5655	2	5	9	492	4	4	1.467	0.198	3.471	2.225
J1744-1134	15660	1743	2	5	–	262	4	2	1.778	0.399	2.667	1.574
B1937+21	29449	4554	2	5	–	306	4	4	0.400	0.177	0.629	0.628
J2145-0750	22506	3667	2	5	6	291	2	2	3.645	0.776	5.959	3.777

<sup>a</sup> Spin parameters: spin frequency  $F0$  and spin frequency derivative  $F1$ .

<sup>b</sup> Astrometric parameters: ecliptic longitude, ecliptic latitude, two proper motion parameters, and parallax.

<sup>c</sup> Binary parameters. The specific parameters vary between pulsars.

<sup>d</sup> DMX bins.

<sup>e</sup> Frequency dependent profile parameters.

<sup>f</sup> Phase jumps between receivers

<sup>g</sup> Root-mean-square residual for joint dataset

<sup>h</sup> Weighted root-mean-square residual for joint dataset

<sup>i</sup> Root-mean-square residual for CHIME/Pulsar data

<sup>j</sup> Weighted root-mean-square residual for CHIME/Pulsar

**Table 4.2:** Timing parameters and their uncertainties (in parentheses) from combined NANOGrav and CHIME/Pulsar data

Source	$v_s$ (1/s)	$\dot{v}_s$ (1/s <sup>2</sup> )	$P_b$ d
J0645+5158	112.94972318752797(5)	$-6.2793(2) \times 10^{-17}$	–
J0740+6620	346.5319964546545(4)	$-1.46391(1) \times 10^{-15}$	4.76694461949(8)
J1012+5307	190.2678373290250(7)	$-6.2004(2) \times 10^{-16}$	0.604672713842(5)
J1125+7819	238.00405315443748(8)	$-3.9330(3) \times 10^{-16}$	15.3554459658(8)
J1713+0747	218.81184376641056(7)	$-4.08394(3) \times 10^{-16}$	67.8251299263(1)
J1744-1134	245.4261234309294(5)	$-5.3816(1) \times 10^{-16}$	–
B1937+21	641.9282301614627(3)	$-4.330871(7) \times 10^{-14}$	–
J2145-0750	62.2958887916192(2)	$-1.15624(6) \times 10^{-16}$	6.83890250968(6)

the combined dataset by more than 5%, we removed the parameter from the list of fitting parameters and repeated the CHIME/Pulsar-only analysis, keeping this parameter fixed at the combined dataset value. We used this justification to fix  $\dot{P}_b$ ,  $\epsilon_1$ , and  $\epsilon_2$  for PSR J1012+5307, to fix  $\dot{A}_1$  and  $h_3$  for PSR J2145-0750, and to fix  $\dot{P}_b$ ,  $m_2$ ,  $KIN$ , and  $K_\Omega$  for PSR J1713+0747.

However, it is important to note that the CHIME/Pulsar only dataset already has uncertainties competitive with the combined dataset after only about a year for PSR J1125+7819 and nearly competitive for PSR J0740+6620. This is a striking example of the power of daily cadence data in helping us better characterize binary systems. This particularly exciting in the broader context of CHIME/Pulsar timing, as binary parameters are what allow us to test GR and to measure the neutron star mass.

Another important component of our timing model are the noise parameters. Although we calculate noise parameters for all receivers simultaneously, we present only the CHIME/Pulsar noise model parameters in Table 4.5, for brevity’s sake. Other analyses within the CHIME ecosystem suggest that there remains room for improvement in our understanding of the instrument, so these noise parameters may yet be able to be improved by improving our understanding of the TOA uncertainties from CHIME/Pulsar. Encouragingly, though, the current EFAC parameters are close to 1, indicating relatively good understanding of our system.

Overall, our results are suggestive of good timing model fits and are very en-

**Table 4.3:** Ratio of timing parameter uncertainties between CHIME/Pulsar solutions and combined solutions for isolated pulsars, for all parameters fit separately in the two models. The large ratios for  $v_s$  and  $\dot{v}_s$  uncertainties reflect the much shorter dataspan in the CHIME/Pulsar only dataset.

Parameter	J0645+5158	J1744-1134	B1937+21
$v_s$	65.5	34.7	133.1
$\dot{v}_s$	609.3	424.5	526.8
$\lambda$	4.97	16.4	43.3
$\beta$	7.13	21.3	12.3

**Table 4.4:** Ratios of timing parameters uncertainties between CHIME/Pulsar solutions and combined solutions for binary pulsars, for all parameters fit in both models. Parameters marked “NG” are held constant at combined CHIME/Pulsar-NANOGrav values, parameters marked “–” are not included for this pulsar.

	J0740+6620	J1012+5307	J1125+7819	J1713+0747	J2145-0750
$v_s$	30.7	42.89	35.6	386	67.62
$\dot{v}_s$	490	769.3	512.6	4922	1415.6
$\lambda$	6.19	8.63	7.33	179	10.27
$\beta$	9.42	17.6	9.81	115	22.67
$P_b$	11.8	155	14.2	102	73.12
$\dot{P}_b$	NG	NG	–	NG	–
$A_1$	2.76	10.82788407	1.40	4.92	7.00
$\dot{A}_1$	–	–	NG	–	NG
$m_2$	3.05	–	–	NG	–
$\sin i$	4.97	–	–	–	–
$T_{\text{asc}}$	1.19	9.04	1.31	–	6.66
$\varepsilon_1$	2.50	NG	1.39	–	9.82
$\varepsilon_2$	2.19	NG	1.34	–	9.36
$h_3$	–	–	–	–	NG
$e$	–	–	–	41.5	–
$T_0$	–	–	–	16.6	–
$\Omega$	–	–	–	16.6	–
$K_{\text{in}}$	–	–	–	NG	–
$K_{\Omega}$	–	–	–	NG	–

**Table 4.5:** Noise parameters for CHIME/Pulsar, calculated with the joint CHIME/Pulsar-NANOGrav dataset.

Source	EFAC	EQUAD	ECORR	$A_{\text{red}}$	$\gamma_{\text{red}}$
J0645+5158	0.957	0.0832	0.0163	–	–
J0740+5520	0.964	0.0234	0.0123	–	–
J1012+5307	1.02	9.26	7.09	0.831	-1.49
J1125+7819	1.10	5.61	0.158	–	–
J1713+0747	1.10	5.12	0.0244	0.0146	-3.04
J1744-1134	1.07	0.313	0.469	0.00573	-4.55
B1937+21	0.815	0.447	0.570	0.147	-3.37
J2145-0750	1.09	7.47	2.66	0.76	-1.39

couraging. CHIME/Pulsar data integrates well with NANOGrav data, and the increase in cadence is visually striking. NANOGrav has had high-cadence programs in the past, but the daily cadence of CHIME/Pulsar data dwarfs the cadence of those programs. CHIME/Pulsar residuals are somewhat larger than GBT and AO residuals; this is not surprising as the telescope sensitivity is less than for CHIME/Pulsar than for GBT and AO. However, averaged RMS residuals for CHIME/Pulsar data remain around  $< 5\mu\text{s}$  and generally about  $1\mu\text{s}$ . We are continually improving our understanding of the CHIME/Pulsar system, and can expect improved timing residuals as we optimize our data processing and TOA generation for CHIME/Pulsar. As we do not attempt to propagate the timing solutions to search for gravitational waves, we cannot say definitively the effect of CHIME/Pulsar data on those results, but this is an interesting question for further study. The CHIME/Pulsar-only dataset is also very encouraging, demonstrating that CHIME/Pulsar is indeed a precision timing instrument.

## 4.5.2 Individual Pulsar Discussion

### PSR J0645+5158

PSR J0645+5158 is an isolated millisecond pulsar with a spin period of 8.85 ms, discovered with the Green Bank Telescope North Celestial Cap Survey (GBNCC) (Stovall et al., 2014). In the NANOGrav dataset, J0645+5158 has been observed

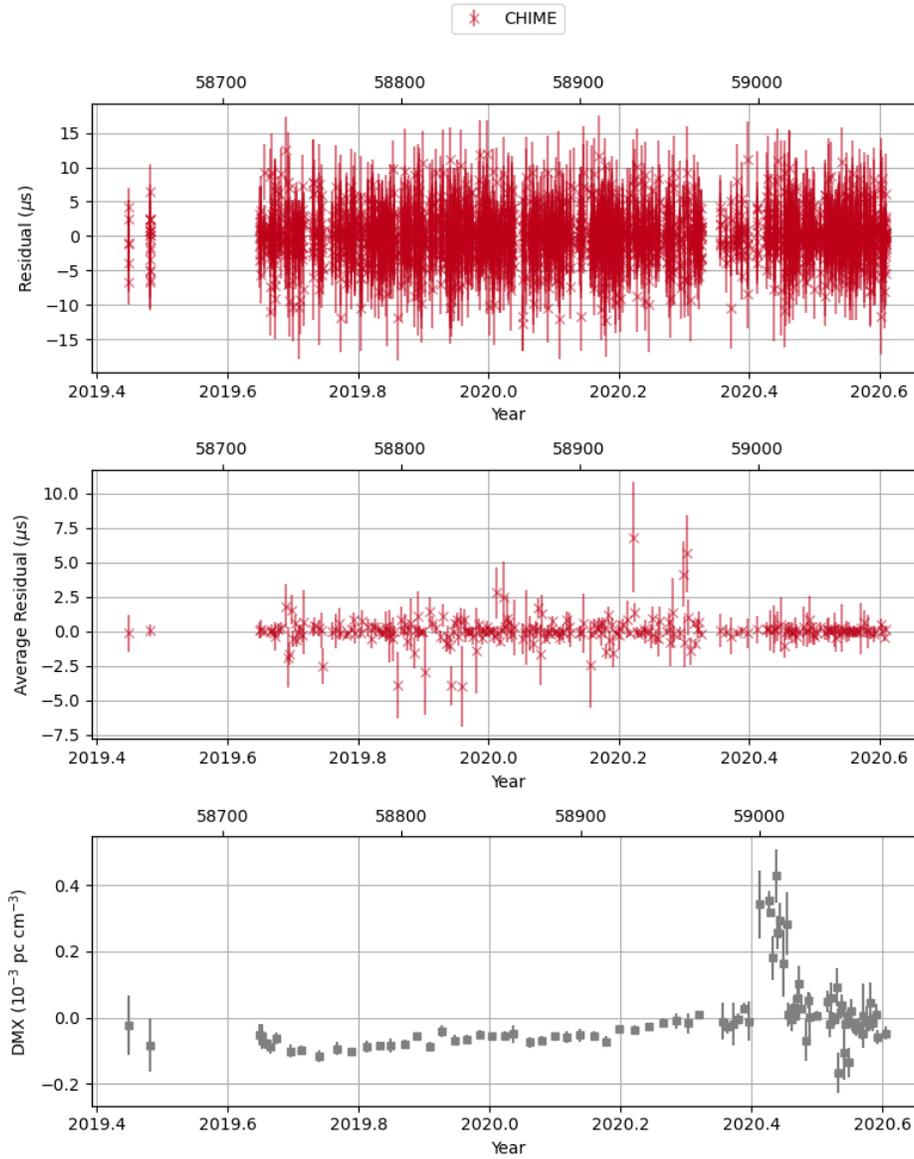
with the GBT since 2011. The overall RMS residual is  $1.009 \mu\text{s}$ , with a weighted RMS residual value of  $0.219 \mu\text{s}$ . The CHIME/Pulsar only residual is  $1.024 \mu\text{s}$  and the weighted rms residual is  $0.292 \mu\text{s}$ .

In the CHIME/Pulsar dataset, the DMX model shows a slight upward trend, and the most notable feature is the large excursion to higher DMX in approximately June 2020. In the combined dataset, the DMX fit for this source does not have an obvious overall trend; there are several periods of increasing DMX, followed by periods of decreasing DMX. The large upward excursion from June 2020 is also visible in the combined dataset. There are a few DMX points that may be outliers in this model or indicative of large but short-term changes that are poorly captured by the NANOGrav dataset's cadence.

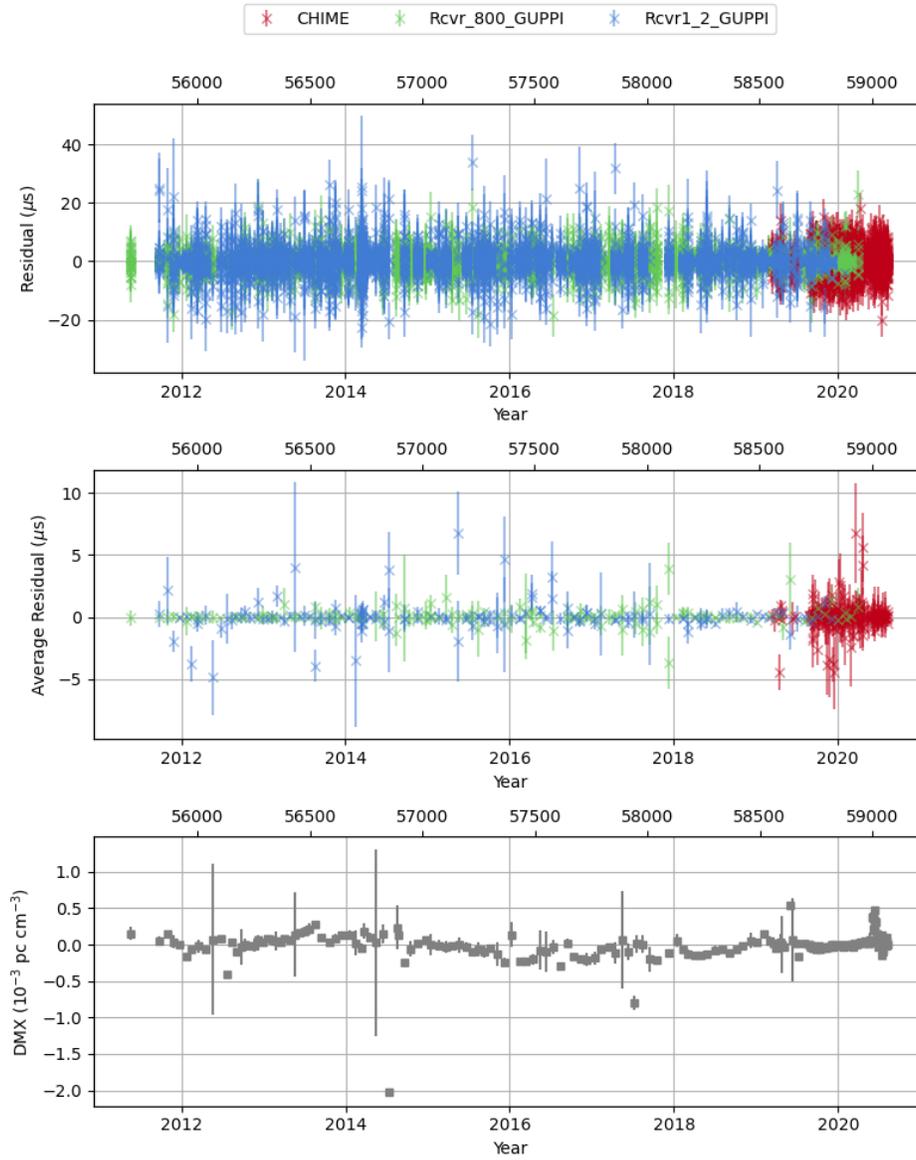
### **PSR J0740+6620**

PSR J0740+6620 is a binary pulsar with spin period of 2.89 ms, discovered by GBNCC (Stovall et al., 2014). In the NANOGrav dataset, this source has been observed with the GBT since late 2013. PSR J0740+6620 has been observed with high cadence at the GBT. The overall residual rms is  $1.632 \mu\text{s}$ , and the weighted RMS residual is  $0.473 \mu\text{s}$ . The CHIME/Pulsar residual is  $1.878 \mu\text{s}$  and the weighted RMS residual is  $0.792 \mu\text{s}$ .

This source has been observed extensively with CHIME/Pulsar and timing results including a mass measurement have been made by Fonseca et al. (2021). This is an important comparison dataset, as it is the most complete CHIME/Pulsar timing solution released to date. The CHIME/Pulsar TOA generation and all timing analyses for the data included in Fonseca et al. (2021) and in this work are entirely independent. This creates a natural venue for comparing results; such a comparison between the parameter values is presented in Table 4.6. Note that for spin frequency and position, the small uncertainties on the fit parameters mean that a small difference between the two solutions results in a many  $\sigma$  difference in results. The absolute difference between solutions for each of the three parameters is approximately an order of magnitude smaller than the difference between the NANOGrav 12.5 year result and the result of Fonseca et al. (2021). The discrepancy in spin parameters and also in position parameters arises from the use of



**Figure 4.1:** Timing residuals and DMX fit for PSR J0645+5158, from CHIME/Pulsar data only. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model.



**Figure 4.2:** Timing residuals and fit for DMX fit for PSR J0645+5158, from CHIME/Pulsar and NANOGrav data. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Blue data points are collected with the GBT L-band receiver, green points are collected with the GBT 800 MHz receiver, and red data points are collected with CHIME.

different reference epochs for the two analyses.

There is tension in the value for the companion mass; this may be due to differences in DMX modeling between the two analyses. We expect more careful DMX modeling for this version of the CHIME/Pulsar dataset and the resolution of the previously mentioned parameter covariance issue will bring the values into alignment. (More detailed analysis was conducted for Fonseca et al. (2021).) Additionally, the DMX values for the CHIME-only dataset demonstrate a systemic offset from the DMX values for the combined dataset in the same time range, likely due to the interplay between FD and DMX parameters. For timing solutions, the structure of the DMX fit is more important than the absolute value.

In the NANOGrav model, PSR J0740+6620's DMX model has an overall linear trend with additional stochastic variations. This is borne out by the CHIME/Pulsar only results; they are consistent with a continuation of the trend, but also allow us to clearly discern a small periodic variation on top of the broader linear trend.

### **PSR J1012+5307**

PSR J1012+5307 is a binary pulsar with a spin period 5.26 ms, discovered at Jodrell Bank Observatory (Nicastro et al., 1995). This source has been observed with the GBT since late 2004. The overall RMS residual is  $1.998 \mu\text{s}$  and the weighted RMS residual is  $0.461 \mu\text{s}$ . For CHIME/Pulsar data only, the overall RMS residual is  $2.728 \mu\text{s}$  and the weighted residual is  $2.097 \mu\text{s}$ . For this source, we use the ELL1 model to parameterize binary parameters.

This source demonstrates clearly one of the advantages of adding CHIME/Pulsar data to the NANOGrav dataset: the scatter in DMX fit noticeably decreases after the addition of CHIME/Pulsar data. Though we are still characterizing the specific reasons for this improvement, it is likely a function of the increased cadence. As DMX bin sizes are fixed, a 6.5 day bin in CHIME/Pulsar data contains 6–7 days of data, while a 6.5 day bin includes only one observation for the NANOGrav dataset.

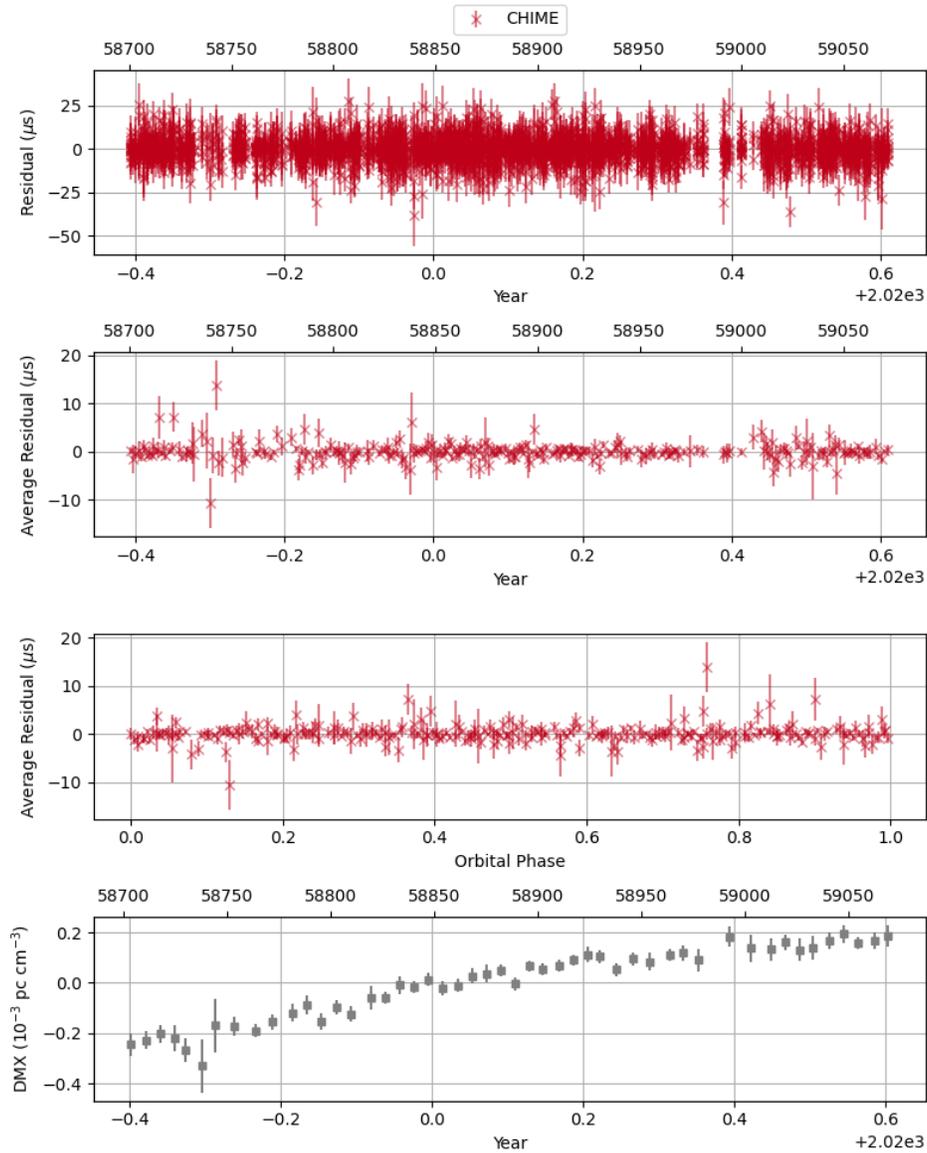
Additionally, prior to the addition of CHIME/Pulsar data, the DMX model appears to show an overall linear trend with sinusoidal variations. During the period where CHIME/Pulsar data is added, however, the DMX variation appears to flat-

**Table 4.6:** Parameter comparison between the combined NANOGrav and CHIME/Pulsar dataset and the results of Fonseca et al. (2021).

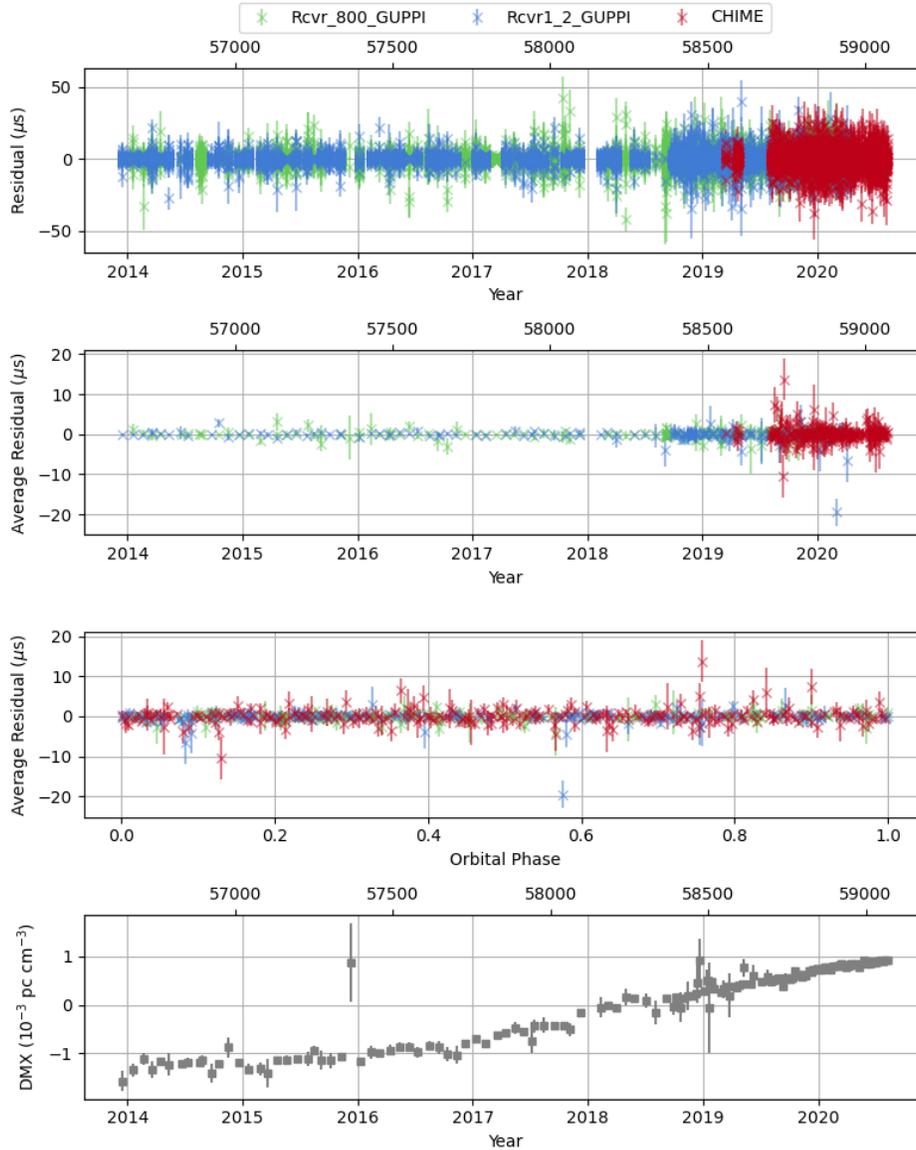
Parameter	Fonseca et al. (2021)	NANOGrav + CHIME/Pulsar	Difference ( $\sigma_{\max}$ )
Spin Parameters			
$\nu_s$ ( $s^{-1}$ )	346.5319964608338(3)	346.5319964546545(3)	17342.04 <sup>a</sup>
$\dot{\nu}_s$ ( $10^{-15}s^2$ )	-1.463874(11)	-1.463909(12)	2.76
Astrometry			
Ecliptic Longitude $\lambda$	103.759135333(12)	103.759135170(14)	11.9 <sup>b</sup>
Ecliptic Latitude $\beta$	44.102478368(13)	44.102477136(14)	88.0 <sup>b</sup>
Binary Parameters			
$P_b$ (days)	4.76694461933(8)	4.76694461949(8)	2.16
$x$ (lt-s)	3.97755608(10)	3.97755622(11)	1.30
$\epsilon_1$ ( $10^{-6}$ )	-5.68(3)	-5.66(3)	0.72
$\epsilon_2$ ( $10^{-6}$ )	-1.833(18)	-1.832(19)	0.031
$T_{asc}$	57804.731308893(17)	57857.1676997(14)	- <sup>b</sup>
$m_c$ ( $M_{\odot}$ )	0.251(5)	0.243(5)	1.44
$\sin i$	0.99909(12)	0.99928(12)	1.67
Configuration			
Reference Epoch	57807	57856	-
Terrestrial Clock Standard	TT(BIPM2019)	TT(BIPM2019)	-
Barycentric Timescale	TDB	TDB	-
Solar system ephemeris	DE438	DE440	-

<sup>a</sup> Due to the use of different reference epochs, the  $\nu_s$ ,  $\dot{\nu}_s$ , and position values are many  $\sigma$  different between the two works. This is expected given the change in reference epochs, and additionally, the real difference in spin period between the two values equates to roughly  $1 \times 10^{-11}$  s. The difference between this work and Fonseca et al. (2021) is approximately an order of magnitude smaller than the difference between the NANOGrav 12.5 year dataset Alam et al. (2021a) and Fonseca et al. (2021).

<sup>b</sup> Due to a slightly longer CHIME/Pulsar dataset in the results in this work, these values differ by about 50 MJD and therefore we present this comparison only for completeness.



**Figure 4.3:** CHIME/Pulsar Timing residuals and DMX fit for PSR J0740+6620. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model.



**Figure 4.4:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J0740+6620. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME.

ten entirely. This is borne out by the CHIME/Pulsar-only DM model in Figure 4.6. This may represent the first sign of CHIME/Pulsar measuring different ISM behaviour than is perceived by the broader NANOGrav dataset or it may simply represent a change in the ISM behaviour in the direction of PSR J1012+5307. Additionally, as in PSR J0740+6620, the CHIME-only dataset fails to capture the overall DMX level, again likely due to the interplay of FD parameters and DMX parameters. As in that case, the overall structure is more important than the value.

### **PSR J1125+7819**

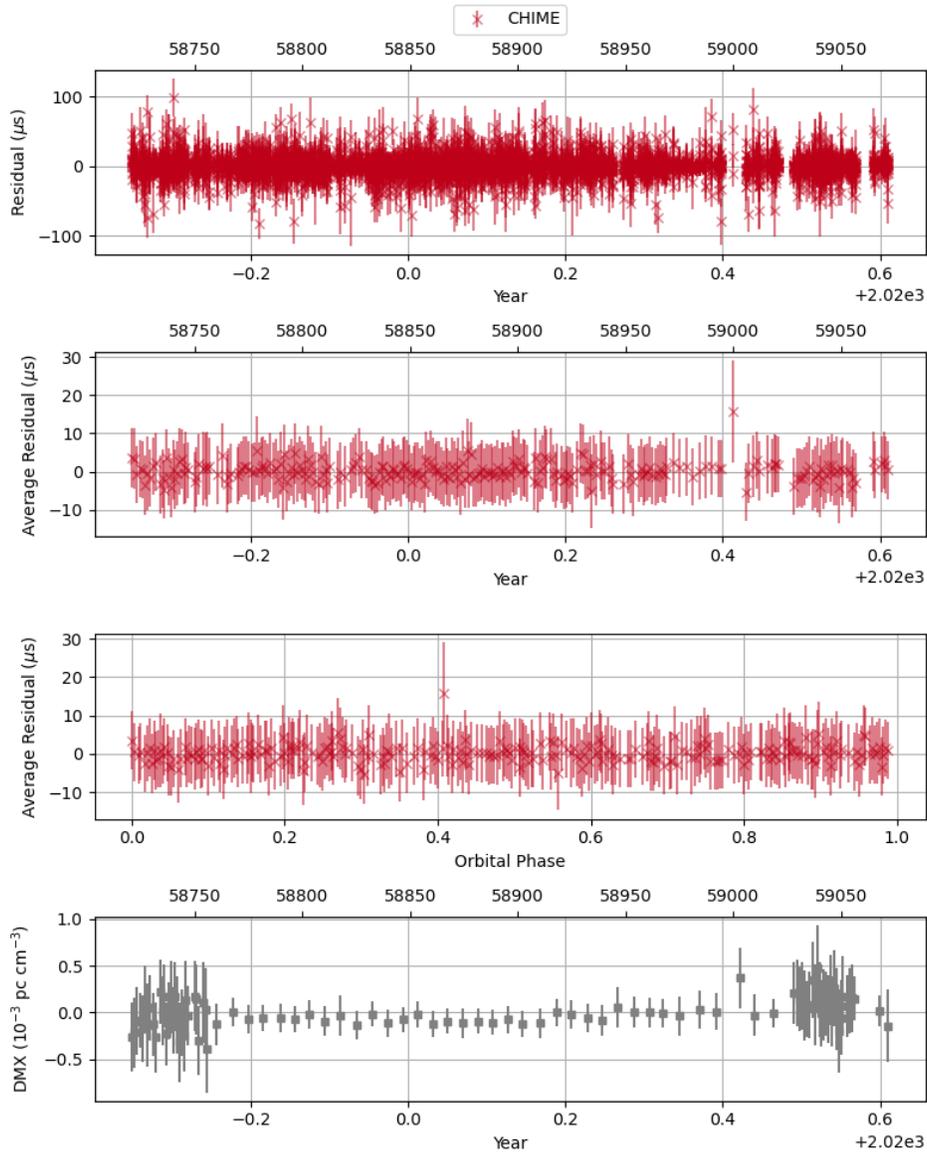
PSR J1125+7819 is a binary pulsar with a spin period 4.20 ms, discovered by GBNCC (Stovall et al., 2014). The NANOGrav observations for this source were conducted with GBT beginning in 2014. The overall RMS residual value is 2.862  $\mu\text{s}$  and the weighted RMS residual value is 1.017  $\mu\text{s}$  for the NANOGrav and CHIME/Pulsar combined dataset. For CHIME/Pulsar, the RMS residual value is 3.014  $\mu\text{s}$  and the weighted residual is 1.876  $\mu\text{s}$ . For this source, we again use the ELL1 model to parametrize binary parameters.

This source is notable in that the uncertainties in the binary parameters are equivalent for the CHIME/Pulsar only dataset and for the combined dataset. This is an exciting portent for the potential of CHIME/Pulsar data in the broader NANOGrav dataset, particularly for those sources that are newer additions to the dataset and do not yet have exquisite binary parameter measurements.

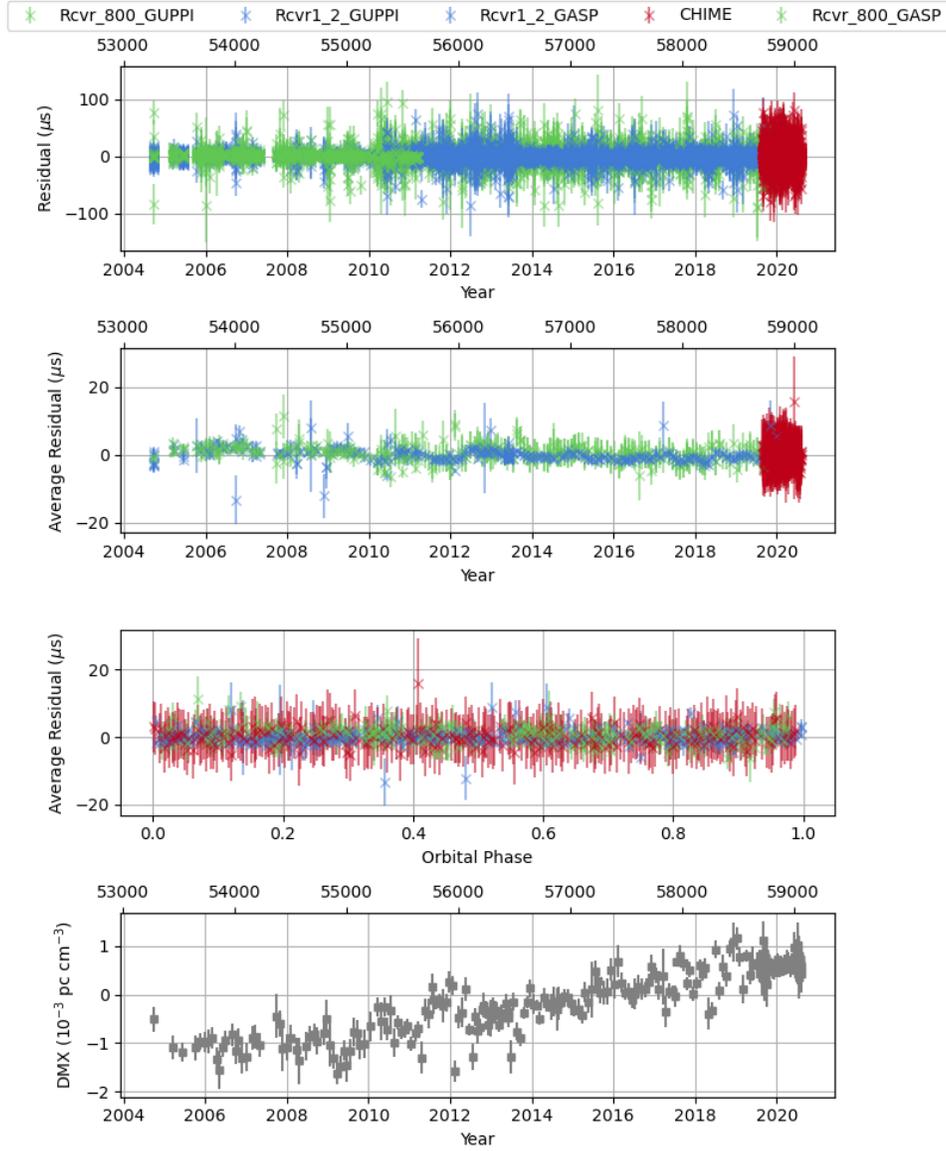
The DMX model for PSR J1125+7819 demonstrates primarily stochastic variations in both the NANOGrav and CHIME/Pulsar datasets. However, there is a noticeable decrease in the overall scatter after CHIME/Pulsar data are added. This may be due to the increased number of observations included in each bin and CHIME/Pulsar's very high sensitivity to DM.

### **PSR J1713+0747**

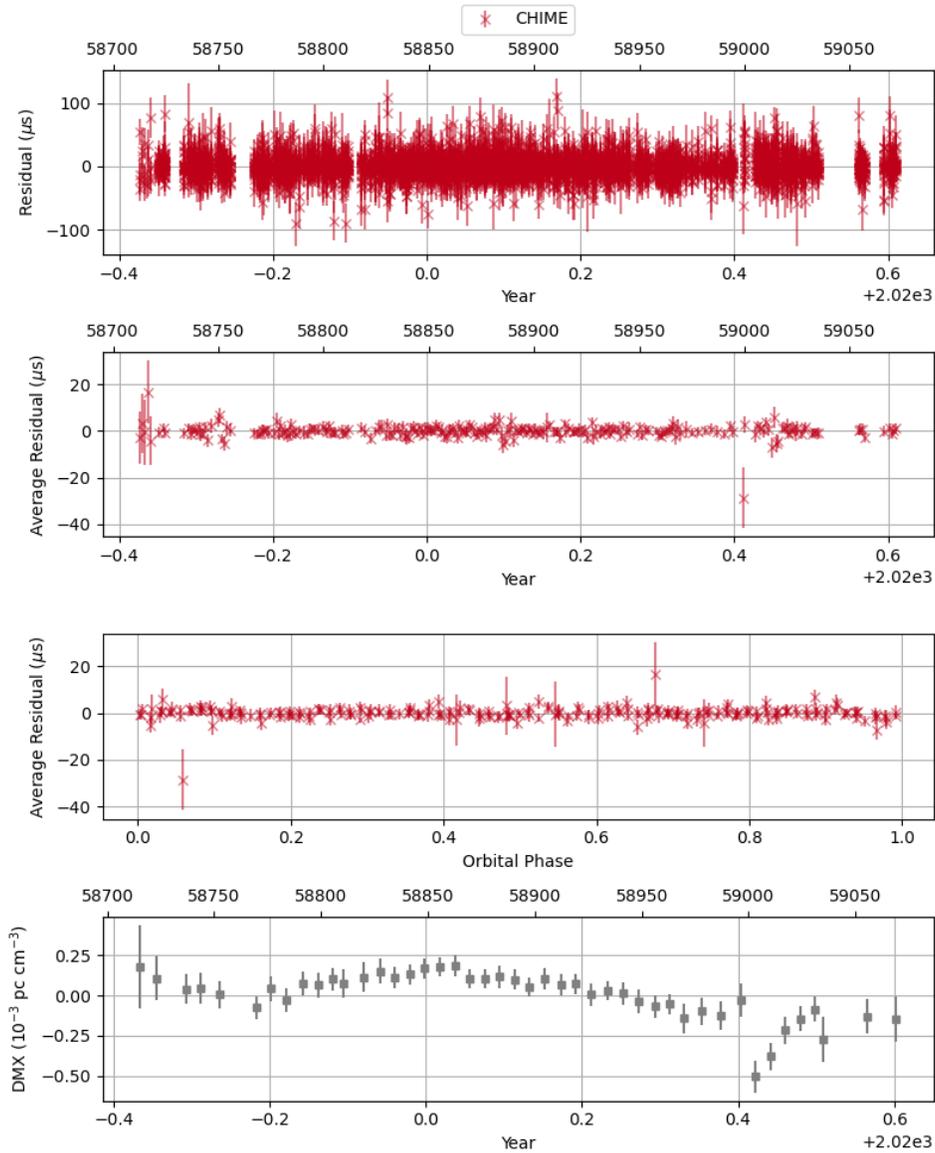
PSR J1713+0747 is a binary pulsar with a spin period of 4.57 ms, discovered with AO (Foster et al., 1993). It is one of the best studied MSPs, with more than 20 years of timing data (Zhu et al., 2015b, 2019). The RMS residual for the combined dataset is 1.467  $\mu\text{s}$  and the weighted RMS residual is 0.198  $\mu\text{s}$ . For the



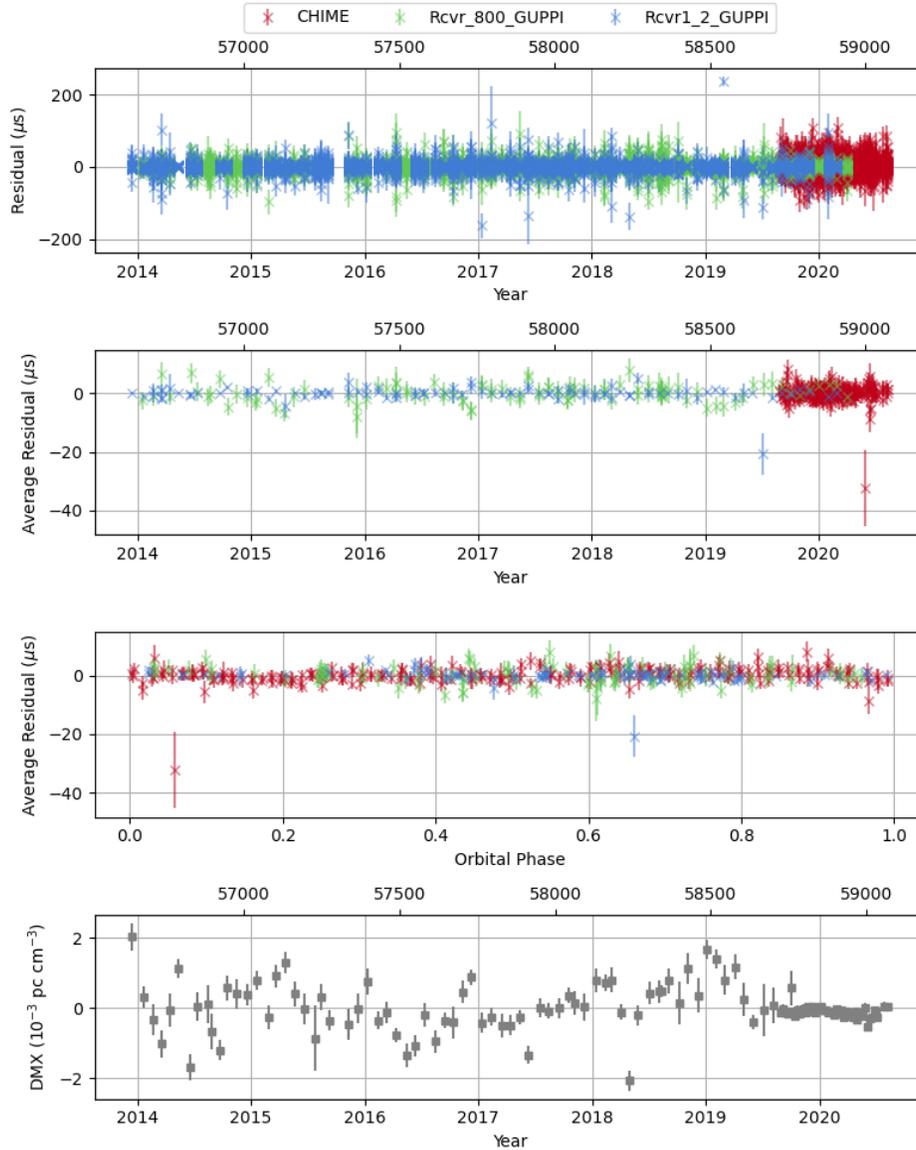
**Figure 4.5:** CHIME/Pulsar Timing residuals and DMX fit for PSR J1012+5307. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model.



**Figure 4.6:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1012+5307. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME.



**Figure 4.7:** CHIME/Pulsar Timing residuals and DMX fit for PSR J1125+7819. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model.



**Figure 4.8:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1125+7819. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME.

CHIME/Pulsar, the RMS residual is  $3.471 \mu\text{s}$  and the weighted RMS residual is  $2.225 \mu\text{s}$ .

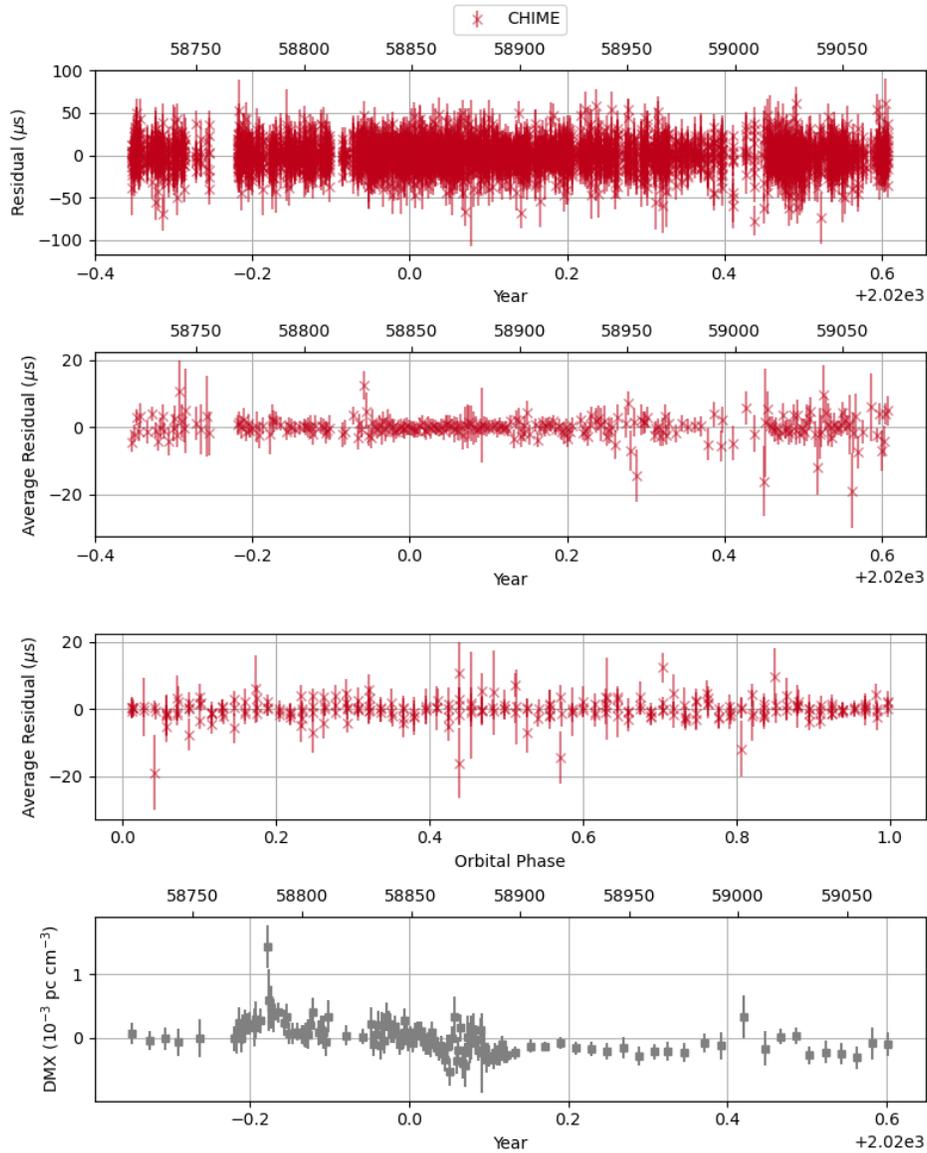
It has been a part of the NANOGrav dataset since 2005, and is observed with both AO and GBT. J1713+0747 is also the only pulsar in the test set to use the more complex DDK binary model. As one of NANOGrav’s best observed pulsars, PSR J1713+0747 benefits comparatively little from the addition of CHIME/Pulsar data. However, it is traditionally regarded as one of the best timing MSPs and therefore is an important source to include in this analysis. Poor results from CHIME/Pulsar would be concerning, but a failure to dramatically improve the dataset is unsurprising.

DMX modeling for PSR J1713+0747 is complicated by the presence of two chromatic timing events in timespan covered by the NANOGrav dataset. These events caused temporary, frequency-dependent changes in the timing parameters, corresponding to a drop in DM (Demorest et al., 2013; Keith et al., 2013; Lam et al., 2018b). The NANOGrav DMX model is altered to account for these events. Other than these known complications, the DMX model is fairly simple, with a slight downward linear trend and some additional structure.

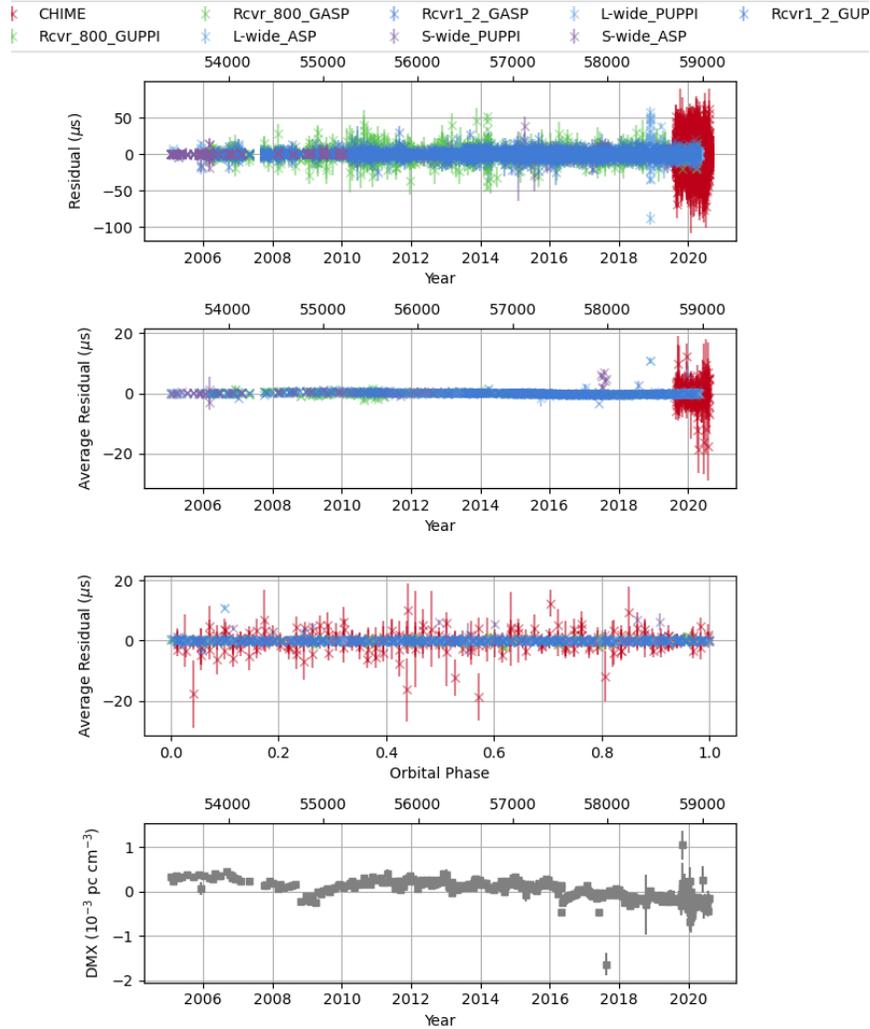
#### **PSR J1744-1134**

PSR J1744-1134 is an isolated MSP with a spin period of 4.07 ms. It has been a part of the NANOGrav dataset since 2005. The RMS residual value for the joint dataset is  $1.778 \mu\text{s}$  and the weighted RMS residual value is  $0.399 \mu\text{s}$ . The RMS and weighted RMS residual values for CHIME/Pulsar data are  $2.667 \mu\text{s}$  and  $1.574 \mu\text{s}$  respectively.

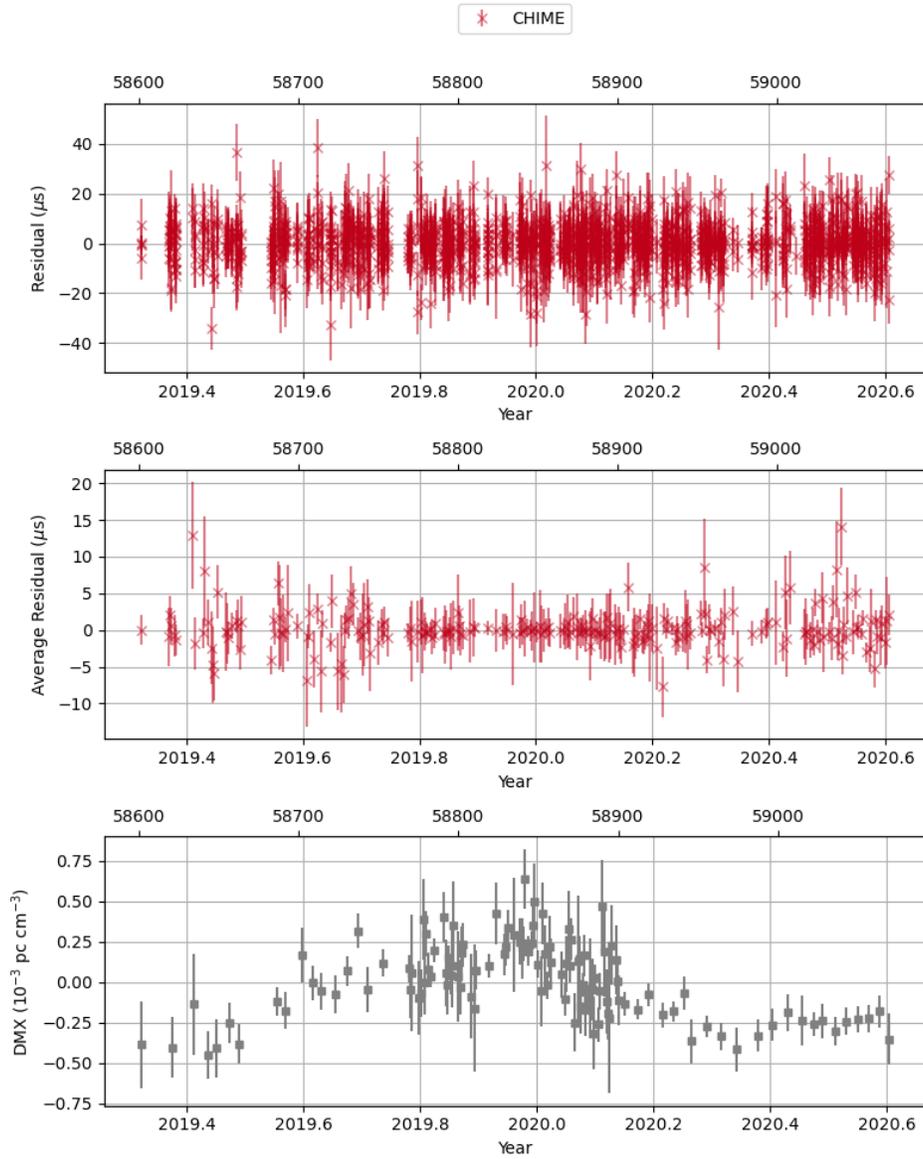
The DMX model for PSR J1744-1134 in the combined NANOGrav and CHIME/Pulsar dataset suggests possible long-timescale periodic variations. The most notable feature is the spike in the DMX model corresponding to the the addition of CHIME/Pulsar data. In the CHIME/Pulsar only dataset, this feature is apparent as a hump, with the same magnitude as the combined dataset’s spike. This is another example of our high cadence data with CHIME/Pulsar providing us a clearer model for DM behaviour.



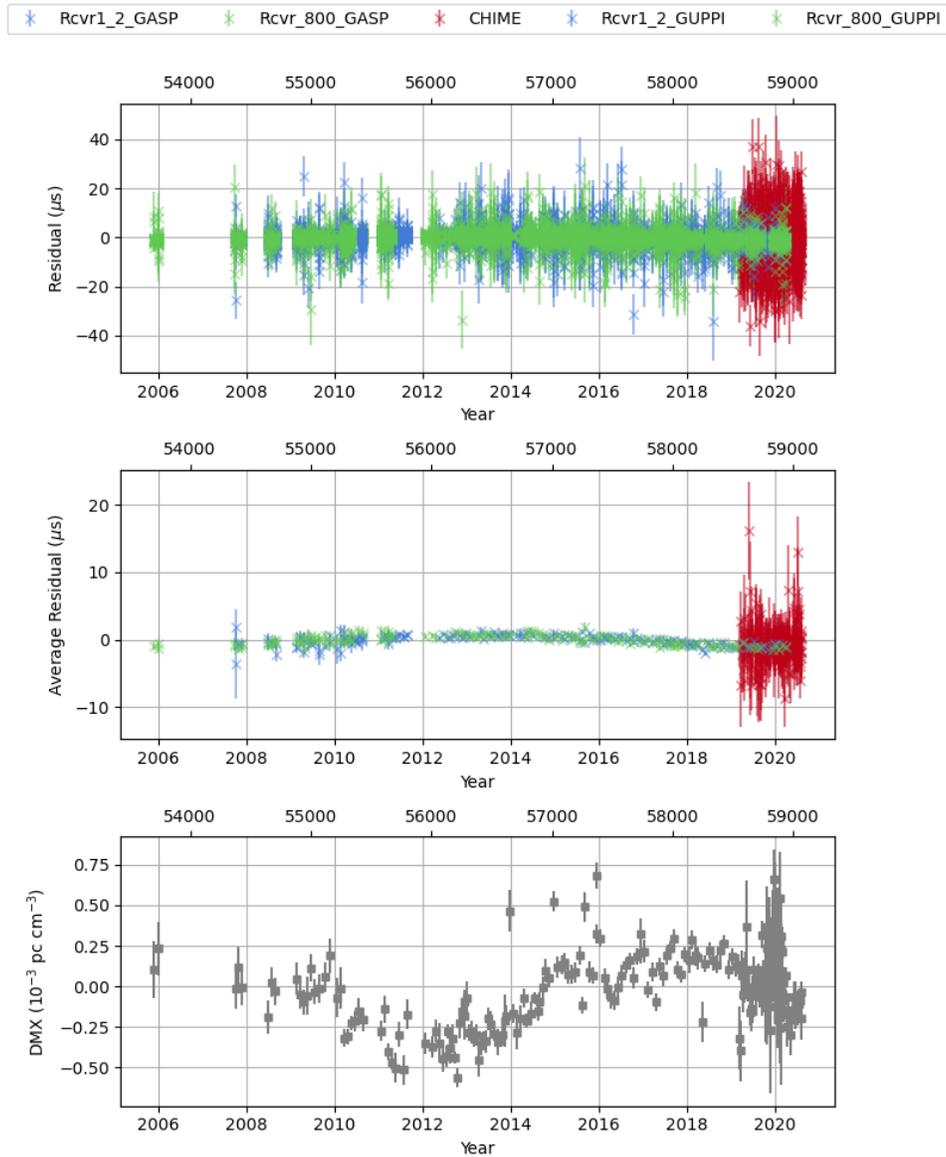
**Figure 4.9:** CHIME/Pulsar Timing residuals and DMX fit for PSR J1713+0747. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model.



**Figure 4.10:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1713+0747. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Light blue and purple data points were collected with the AO L-wide and S-wide receivers, respectively. Dark blue and green data points were collected with the GBT L-band and 800 MHz receivers, respectively. Red data points were collected with CHIME.



**Figure 4.11:** CHIME/Pulsar Timing residuals and DMX fit for PSR J1744-1134. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model.



**Figure 4.12:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J1744-1134. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME.

### **PSR B1937+21**

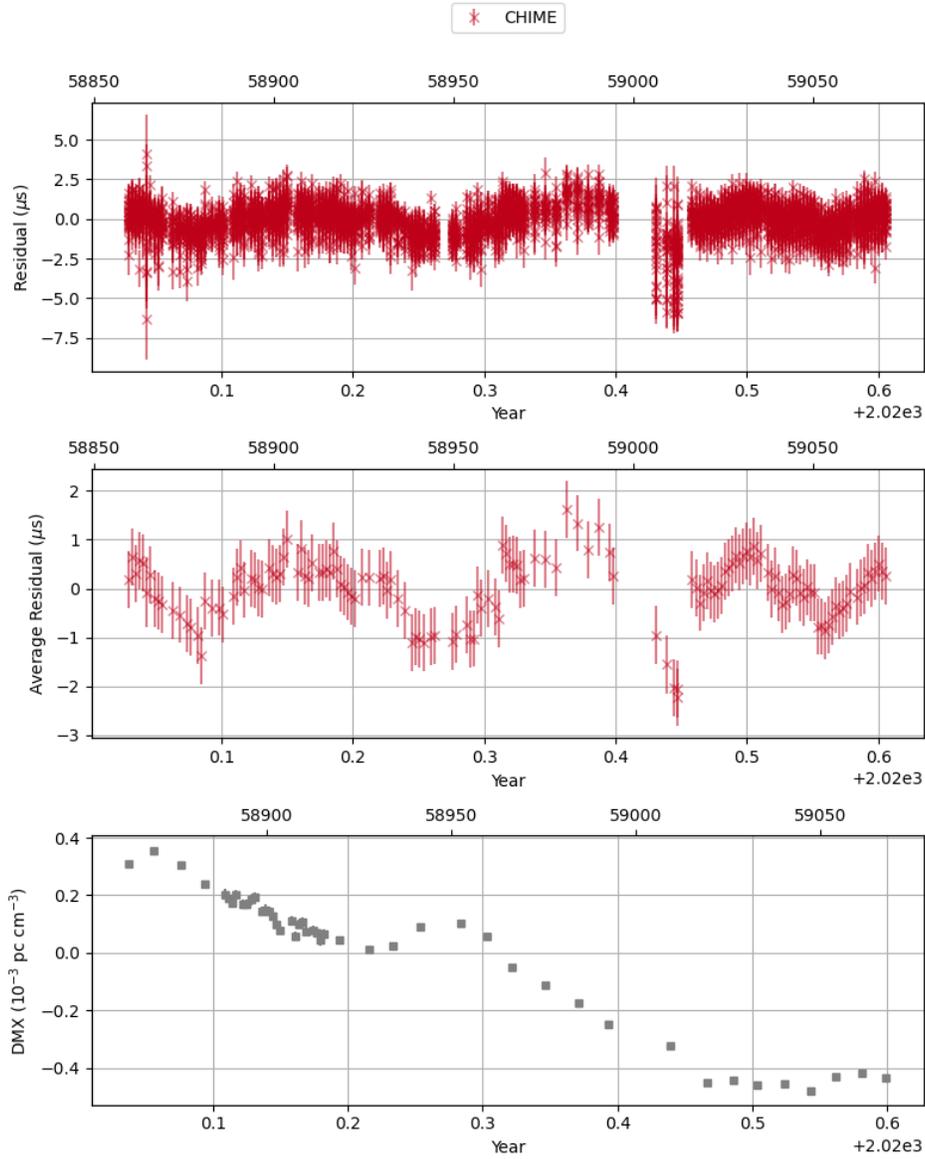
PSR B1937+21 is the fastest known MSP, with a spin period of only 1.56 ms. It was also the first MSP, discovered by Backer et al. (1982). It is a unique source in many ways. The most obvious is that the overall timing solution is excellent but influenced by significant red noise. This results in the obvious structure in the timing residuals for PSR B1937+21; Figure 4.15 demonstrates that this structure can be removed by properly fitting for red noise. The RMS residual for the combined dataset is  $0.400 \mu\text{s}$  and the weighted RMS residual is  $0.177 \mu\text{s}$ . For the CHIME/Pulsar dataset, the RMS residual and weighted RMS residual are  $0.629 \mu\text{s}$  and  $0.628 \mu\text{s}$  respectively.

The source has been part of the NANOGrav dataset since the outset and is observed with both AO and GBT.

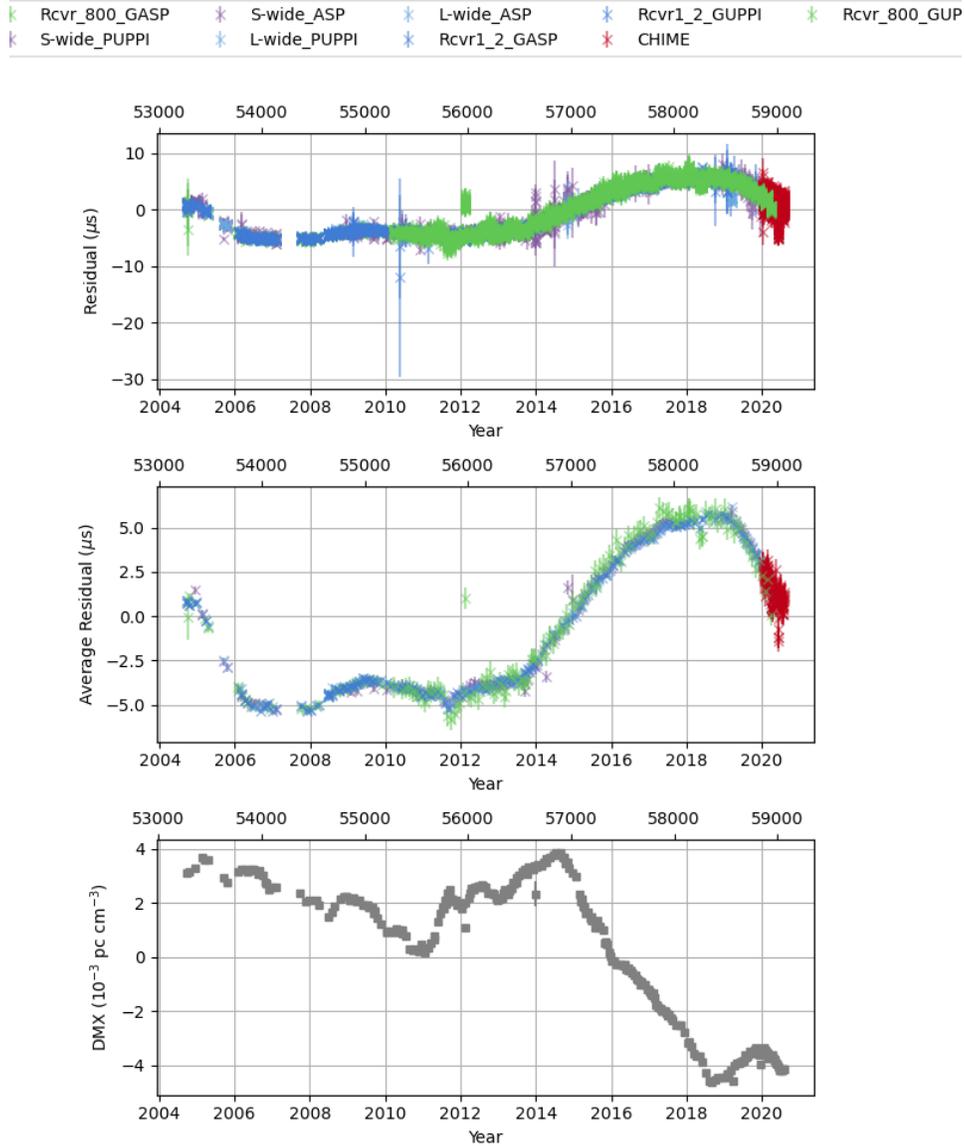
PSR B1937+21 experiences large-scale DM variations as parametrized by DMX, but the CHIME/Pulsar results are consistent with the overall trend in the NANOGrav dataset. DM variations for this source have previously been studied in depth, e.g. by Ilyasov et al. (2005); Ramachandran et al. (2006); Lam et al. (2016); Jones et al. (2017). As a result, the initial addition of CHIME/Pulsar data does not appear to make a large difference in the DMX modeling. However, the changes we see in DMX models for other sources with CHIME/Pulsar plus our prior knowledge that PSR B1937+21 has interesting DM structure suggest that systematic DM studies for this source with CHIME/Pulsar are likely to be profitable.

### **PSR J2145-0750**

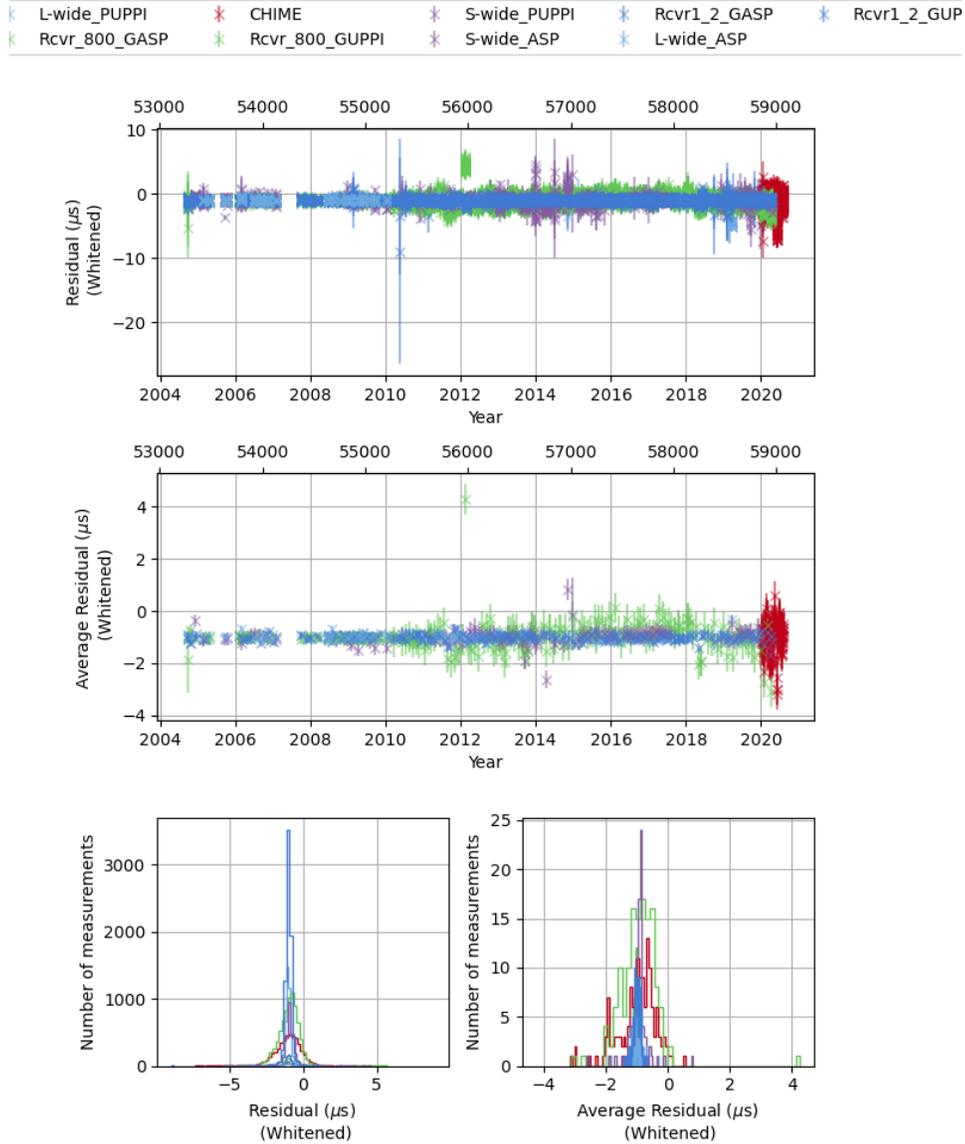
PSR J2145-0750 is a binary pulsar with a spin period of 16.1 ms, discovered with the Parkes Radio Telescope (Bailes et al., 1994). This is the MSP with the longest spin period in the NANOGrav dataset. It has been part of the NANOGrav dataset since 2004. The RMS residual value is  $3.645 \mu\text{s}$  and the weighted RMS residual value is  $0.776 \mu\text{s}$ . For CHIME/Pulsar, the RMS residual value and weighted RMS residual values are  $5.959 \mu\text{s}$  and  $3.777 \mu\text{s}$  respectively. The uncertainties in the binary parameters are higher than for similar pulsars; this may be attributable to the source's proximity to CHIME/Pulsar's lower declination limit or to our sampling of the orbit.



**Figure 4.13:** CHIME/Pulsar Timing residuals and DMX fit for PSR B1937+21. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. The large-scale structure in the residuals is due to red noise, not poor fit.



**Figure 4.14:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR B1937+21. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, and the third shows the DMX model. Light blue and purple data points were collected with the AO L-wide and S-wide receivers, respectively. Dark blue and green data points were collected with the GBT L-band and 800 MHz receivers, respectively. Red data points were collected with CHIME.



**Figure 4.15:** Whitenened NANOGGrav and CHIME/Pulsar timing residuals for PSR B1937+21. The first panel shows whitened timing residuals vs. time, the second shows epoch-averaged, whitened residuals vs. time, and the histograms show the distribution of whitened residuals. The colours correspond to receivers as discussed in Figure 4.14.

The DMX model for PSR J2145-0750 follows a general upward trend. Like the DMX model for PSR J1744-1134, the DMX model for PSR J2145-0750 includes a large spike shortly after the addition of CHIME/Pulsar data. Again, as in PSR J1744-1134, this corresponds to a real increase in the DMX model seen in the CHIME/Pulsar data.

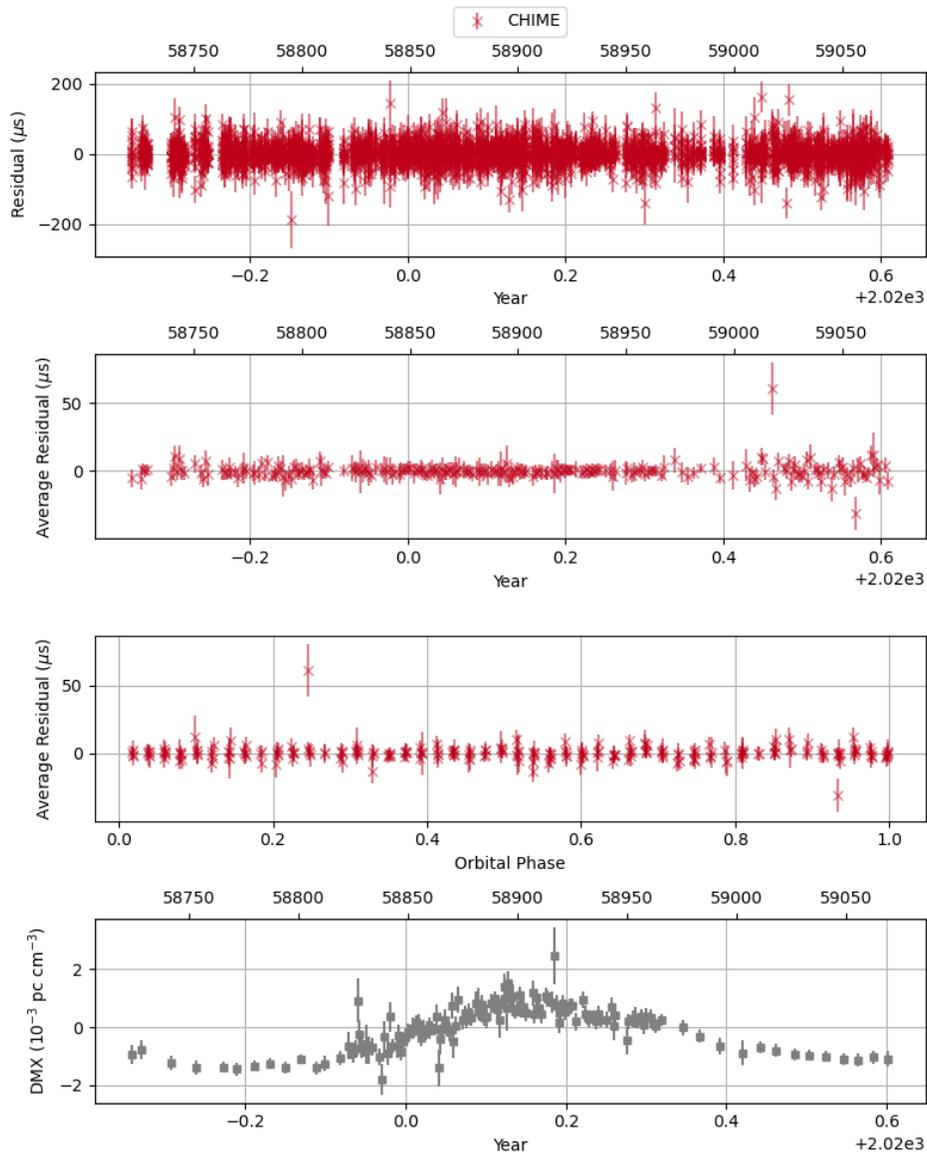
## 4.6 Future Directions

### 4.6.1 Improving Timing Solutions

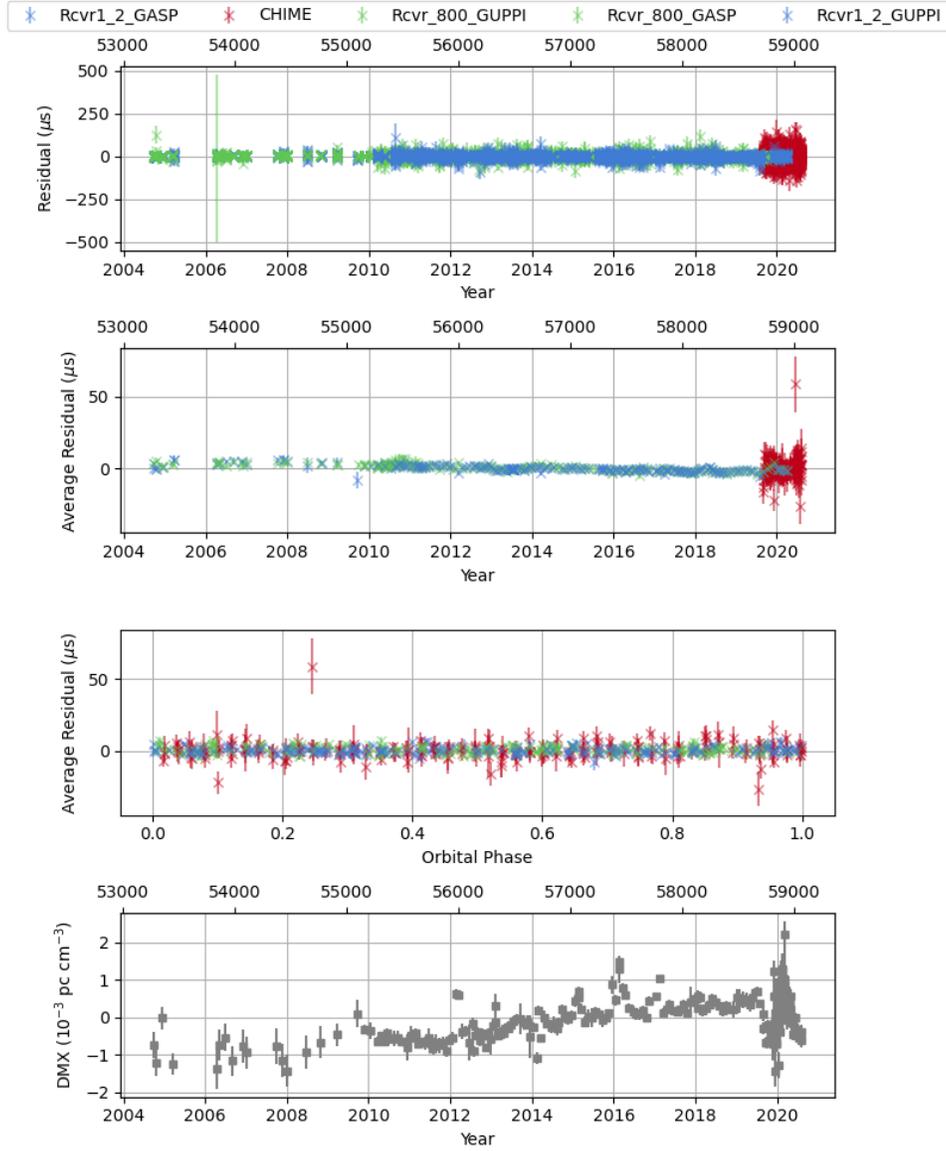
Though this is a very promising preliminary dataset, there are many things that can be improved within this analysis.

First, and most obviously, this work covers only a small sample of eight pulsars. There are approximately 70 pulsars in the upcoming NANOGrav 15 year dataset and a complete analysis should be conducted for each, with declination greater than  $-15^\circ$ , our usual cut-off for pulsar observation with CHIME/Pulsar. To enable us to conduct thorough studies for a larger number of sources, we will need to develop an automated pipeline to prepare NANOGrav-compliant TOAs. This can be accomplished by modifying the existing CHIME/Pulsar automated TOA generation pipeline to include components of `nanopipe`, the NANOGrav TOA generation pipeline. These pipelines are different not only in specifics but also in philosophy: `nanopipe` is intended to be run to create discrete releases while the CHIME/Pulsar TOA generation pipeline is intended to be run daily.

Second, we need to improve the rigor of our TOA excision. In final NANOGrav datasets, such as Alam et al. (2021a), all excised TOAs are given a flag to indicate the reason for their excision here, and an automated outlier analysis is performed to rigorously determine outliers in the TOAs. We have not used this automated pipeline in our current analysis; all of our TOA excision except for an initial  $S/N$  cut is manual. Outlier analysis is computationally expensive, so this is consistent with the process used in intermediate NANOGrav processing stages. However, we will want to demonstrate the feasibility of automated outlier analysis with CHIME/Pulsar TOAs prior to incorporating them into a NANOGrav data release. Additionally, DMX modeling is a new feature of PINT and continues to evolve.



**Figure 4.16:** CHIME/Pulsar Timing residuals and DMX fit for PSR J2145-0750. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model.



**Figure 4.17:** NANOGrav and CHIME/Pulsar Timing residuals and DMX fit for PSR J2145-0750. The first panel shows timing residuals vs. time, the second shows epoch-averaged timing residuals vs. time, the third shows the epoch-averaged timing residuals vs. binary orbital phase, and the fourth shows the DMX model. Blue data points were collected with the GBT L-band receiver, green points were collected with the GBT 800 MHz receiver, and red data points were collected with CHIME.

Third, we should optimize our DMX modeling for CHIME/Pulsar. In this analysis, we have simply used the DMX bin size used for NANOGrav analyses of each pulsar. However, Fonseca et al. (2021) have demonstrated that changes in DMX bin size can cause significant variations in timing solution with CHIME/Pulsar data. An immediate next step will be to systematically determine the ideal bin size for CHIME/Pulsar data when integrated into the NANOGrav dataset.

We should also implement wide-band timing for both CHIME/Pulsar data individually and as part of the NANOGrav set. This is beneficial for our increasingly wide bandwidth instruments, as it allows us to be more precise without ungainly datasets. In particular, wideband timing removes the need for FD parameters, incorporating frequency-dependence into the TOA generation process. It also dramatically decreases the number of TOAs in a timing model; this is increasingly important as our long-term pulsar timing datasets grow.

To create wideband TOAs for CHIME/Pulsar data, we will use `PulsePortraiture`, a python package designed for this purpose and used in Alam et al. (2021b). The philosophy of `PulsePortraiture`, introduced in Pennucci et al. (2014) and extended in Pennucci (2019), is to provide a minimal extension to the classic model of TOA creation to allow for frequency dependent templates and to fit DM. The details of the algorithm are discussed in those works. To create our profile models, we will use the spline interpolation method for creating profiles outlined in Pennucci (2019) for maximum adaptability.

Finally, improving our understanding of the polarimetry of CHIME/Pulsar will help us improve our CHIME/Pulsar TOAs and therefore our combined timing models. Such analyses are far beyond the scope of this work, but work by van Straten (2006) demonstrates that we can improve the uncertainty in our TOAs by accounting for the polarization properties of our receiver. Understanding CHIME/Pulsar polarization is a major challenge, and well beyond the scope of our initial CHIME/Pulsar and NANOGrav joint datasets. However, it is an extension with the potential to dramatically improve the efficacy of our CHIME/Pulsar observations.

### 4.6.2 Applying the Joint Dataset

The final area for extension is in applying the dataset to the ultimate NANOGrav goal: searching for low-frequency gravitational waves. We have not examined the implications of the combined dataset for gravitational wave detection, as it is beyond the scope of this work, but such simulations should be pursued. Ideally, we would add the complete set of CHIME/Pulsar solutions for this analysis, but we could also conduct a preliminary analysis by including CHIME/Pulsar data for a sub-set of the most important NANOGrav pulsars and running a gravitational wave search as in Arzoumanian et al. (2020).

In addition to applying the dataset to gravitational wave searches, the combined dataset will also be a powerful tool for improving our understanding of the effects of the ISM. As we move into an era of wide-band pulsar astronomy, frequency-dependent DM variations, which can be substantial, are likely to become an important topic of study (Cordes et al., 2016). CHIME/Pulsar is well positioned, with a large fractional bandwidth at moderately low frequency, to examine such DM variations and the ISM processes that cause them.

It is usually the case that a strong dataset breeds unexpected applications. In addition to those we already anticipate, we look forward to seeing the new science that the full CHIME/Pulsar-NANOGrav dataset will enable in the future.

## Chapter 5

# Detection of New Galactic Sources with CHIME/FRB

1

*The most damaging phrase in the language is: 'It's always been done that way.'* — Grace Hopper

### 5.1 Introduction

In the five decades since their initial discovery, radio pulsars (rapidly rotating, highly magnetized neutron stars) have proven to be excellent tools to study physics, including tests of general relativity. To date, astronomers have discovered more than 2,800 pulsars in the Milky Way and Magellanic Clouds (See the ATNF Pulsar Catalog, Manchester et al., 2005)<sup>2</sup>. Most pulsars emit consistently at radio wavelengths, but a subset emit only intermittently or experience nulling. Nulling pulsars experience sharp, sudden drops in pulse energy for a few periods, as first reported by Backer (1970). Intermittent pulsars are an intermediate category, displaying intermittent emission but behaving as a conventional pulsar during active periods (Lyne, 2009; Kramer et al., 2006a). In 2006, careful re-examination of

---

<sup>1</sup>A version of this chapter has been submitted to *The Astrophysical Journal*

<sup>2</sup><http://www.atnf.csiro.au/people/pulsar/psrcat>.

Parkes Multibeam Pulsar Survey data (PKMBS; Manchester et al., 2001) introduced Rotating Radio Transients (RRATs) to the field—sources with occasional pulsar-like emission and underlying radio periodicity (McLaughlin et al., 2006; McLaughlin et al., 2009). RRATs remain somewhat loosely defined. They are pulsar-like sources detected via single pulses instead of using periodicity search, and folded observations are not possible for these source: in the absence of single pulses, they do not appear to be emitting. These sources may exist as part of a continuum including nulling and intermittent pulsars (Burke-Spolaor, 2013). Conventional pulsars can be detected via single pulses but RRATs cannot be detected with a periodicity search.

Pulsar search strategies are well established, usually falling into two categories: untargeted searches and targeted searches. In untargeted searches, a specified region of sky is searched for pulsar emission. In targeted searches, a region of interest such as a known supernova remnant, globular cluster, or other likely pulsar environment is selected and thoroughly searched, as in, e.g., Ransom et al. (2005). Modern pulsar searches such as the Parkes High Time Resolution Universe (HTRU), the recently terminated Arecibo L-Band Feed Array Pulsar Survey (PALFA), and the Green Bank North Celestial Cap (GBNCC) survey are highly effective at discovering both “slow” and millisecond pulsars (Keith et al., 2010; Cordes et al., 2006; Stovall et al., 2014).

These surveys, however, are limited in the amount of time they can spend focusing on each pointing, as most such surveys make use of shared observatories with limited time allocations. This poses two relevant challenges. First, surveys are susceptible to missing intermittent pulsars or pulsars with a significant nulling fraction. Second, pulsar searching programs are sometimes unable to conduct sufficient follow-up observations to determine complete timing solutions.

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) is a transit telescope with several backends enabling a wide-range of science. Among those, the CHIME Fast Radio Burst (CHIME/FRB) and CHIME Pulsar Timing (CHIME/Pulsar) systems are uniquely situated to detect and time intermittent or otherwise unusual sources which produce bright single pulses. The sister systems were primarily designed to separately discover FRBs and to time known pulsars, but when combined, the search engine can identify bright single pulses from Galac-

tic sources each day and the pulsar instrument can collect search-mode observations to find initial timing solutions then daily fold-mode observations to improve timing solutions. CHIME is discussed in more detail in Chapter 3.

In this chapter, we discuss the discovery of seven new sources, including timing solutions where possible. Though detections of FRBs in pulsar searches are common (e.g. Lorimer et al., 2007; Champion et al., 2016; Patel et al., 2018), this work is among the first discovery of pulsars detected via single pulses in an FRB search.

In Section 5.2, we describe our method for detecting and characterizing new Galactic sources from single pulses, including discussion of the possibility of future detections. In Sections 5.3 and 5.4, we discuss the individual sources and their interpretation, with a focus on the future potential of this method.

## 5.2 Detection, timing, and analysis methods

### 5.2.1 Initial detection

As CHIME/FRB is a passive monitor, it will detect any emission from a source as it transits. Each of the seven sources presented in this work was a serendipitous discovery, not the result of a targeted search. CHIME/FRB scientific personnel acting as system monitors manually identified these sources as not corresponding to any known Galactic sources and flagged them for follow-up. This was not a systematic process and therefore it is not possible to estimate its completeness. However, we do realize the value for a thorough search through the CHIME/FRB database; in a subsequent work we will present a detailed method and the results for a systematic search of the CHIME/FRB database. The pulsars discovered in this alternate method have already been announced on the CHIME/FRB public galactic webpage.<sup>3</sup>

These sources have been primarily identified using metadata detections, without intensity data. However, after initial detection, we manually configure the CHIME/FRB system to produce a small set of saved intensity data for each source

---

<sup>3</sup>New Galactic pulsars discovered by CHIME/FRB can be viewed at <https://www.chime-frb.ca/galactic>.

to confirm that they are astrophysical and not RFI by examining a frequency vs. time “waterfall” plot of the CHIME/FRB intensity data. This intensity data are used only to confirm that sources are astrophysical and do not contribute to timing solutions.

### 5.2.2 CHIME/FRB metadata analysis

Intensity data from the CHIME/FRB system are saved only for sources with DMs consistent with Ambiguous or Extragalactic origin. This criterion excludes new pulsar and RRAT candidates. However, the CHIME/FRB system does preserve metadata for all events, including RFI and Galactic astrophysical events.

This metadata includes the source position, DM, and SNR, as well as the time of detection. This means it provides a basic daily monitor for new Galactic sources, recording any signals detected by the systems.

For this analysis, we include all bursts with  $\text{SNR} > 8$ , for the date range 2018-08-28 (MJD 58358) to 2020-05-01 (MJD 58970), inclusive.

Generally, we also require the event to be detected within the full-width-at-half-maximum (FWHM) at 600 MHz for a CHIME/FRB beam. In Section 5.3.2, we discuss a source which does not fall within the declination range of a formed beam at 600 MHz, and instead use the wider 400 MHz beam.

Figures 5.1 to Figure 5.7 provide examples of the information we can glean from CHIME/FRB metadata. These figures show the per-pulse SNR for each detection and the number of detections per day. They also include system-wide CHIME/FRB information, including the daily relative system sensitivity and the exposure times. Uncertainties in daily exposure times are the standard deviations between exposure times within the source’s uncertainty region; this is consistent with the uncertainty calculations used in Chapter 6 and previous CHIME/FRB publications, e.g., CHIME/FRB Collaboration (2019); Fonseca et al. (2020)

This information can be used to generate a simple estimate for burst rates. As all source locations are observed daily by CHIME/FRB, this metadata are the most complete record of detections. Though CHIME/Pulsar cadence is high by pulsar standards, the automated scheduler can lead to uneven coverage for sources at different sky locations; a source in a right ascension region dense with sources will

be observed less frequently than an isolated source. Therefore, using the exposure and detection data collected by CHIME/FRB is preferable for determining rates.

In addition to rate information, we can create initial TOAs and timing solution estimates using CHIME/FRB metadata. These TOAs, created using the arrival information included in the metadata, have standardized uncertainties determined not by the TOA uncertainty description discussed in Chapter 1 and 4 but by the instrumental constraints of CHIME/FRB. These metadata have uniform DM uncertainties of  $1.62 \times 2^i \text{ pc cm}^{-3}$ , and uniform TOA uncertainties,  $\sim 31.5 \times 2^i \text{ ms}$ , where  $i \in \{0, 1, 2, 3, 4\}$  is the tree index CHIME/FRB Collaboration (2018). Pulsars are detected in different trees based on their widths. Uncertainties are determined based on coarse graining steps in the tree dedispersion algorithm.<sup>4</sup>

With a set of single pulse TOAs from CHIME/FRB metadata, we can estimate the source’s spin period using brute force methods, finding the period between bursts which allows for the maximum integer number of periods between detections. This method is biased towards detecting multiples of the coarse-graining sampling time as the period. We remedy this by including a Monte Carlo resampling of TOAs, where each TOA is placed randomly within its error region prior to period estimation. The goodness-of-fit metric for all instances is averaged to create the global best spin period estimate. We confirmed the efficacy of this method by using it to recover known pulsar periods from CHIME/FRB metadata. This method was developed by Ziggy Pleunis, then Ph.D. student at McGill University.

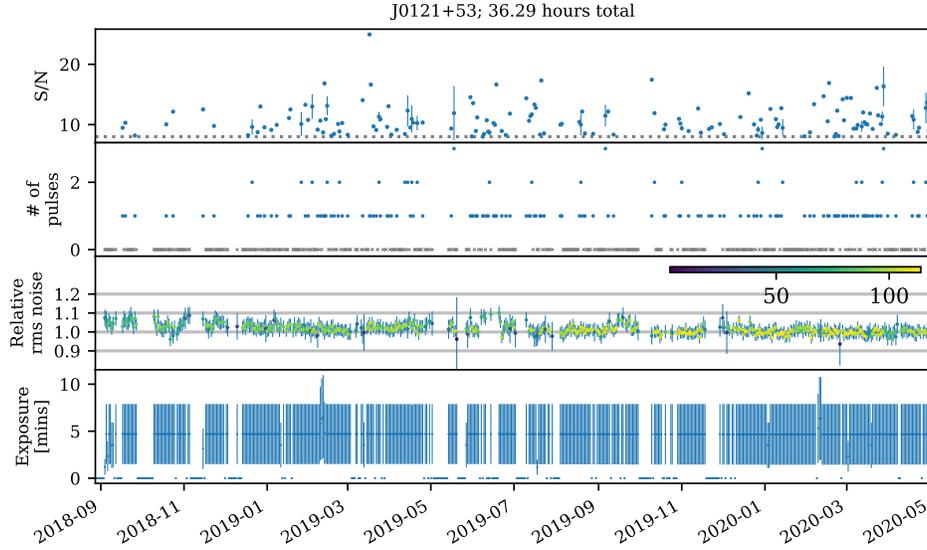
These initial TOAs and period estimation can be used with pulsar timing software such as TEMPO to create rudimentary timing solutions. These solutions are less precise than CHIME/Pulsar solutions due to the relatively large, fixed timing uncertainties. However, they are useful as starting points for more precise timing solutions.

### 5.2.3 CHIME/Pulsar data analysis

After candidate sources are detected with CHIME/FRB, they are added to the CHIME/Pulsar automated observing scheduler described in Amiri et al. (2021) to obtain more sensitive observations. CHIME/Pulsar initially observes sources

---

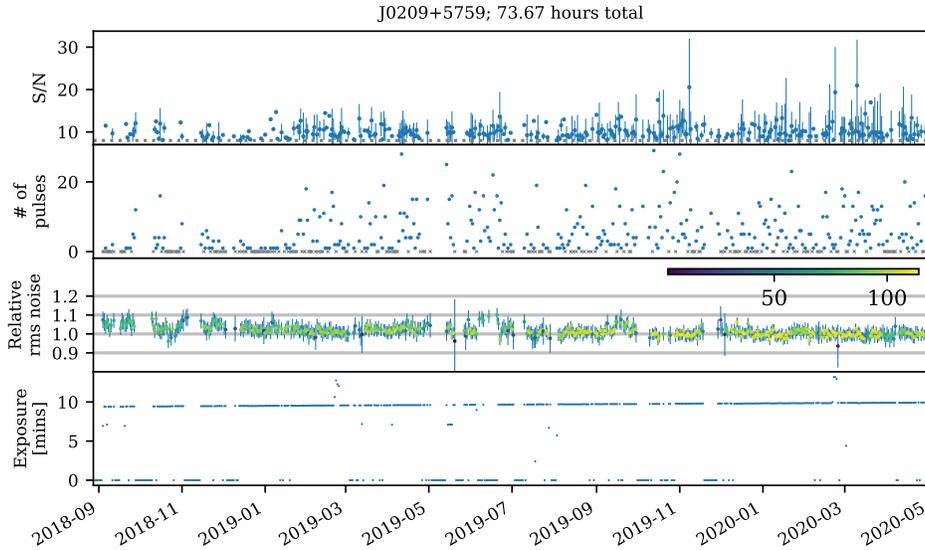
<sup>4</sup>Coarse-graining factors are configurable and were lowered in late 2020.



**Figure 5.1:** CHIME/FRB detections of new RRAT-like source PSR J0121+53, shown as an example, along with sensitivity to this location and exposure time at this location. The top panel shows the daily median detection S/N and standard deviation of the in-beam events with  $S/N > 8$ . The second panel shows the number of single pulses that were detected. If no pulse was detected on a day with non-zero exposure the marker is a gray cross. The third panel shows the daily sensitivity of the experiment as characterized by the relative rms noise of known pulsar detections. The color bar in the inset shows how many pulsars were used to obtain the average rms noise. The bottom panel shows the exposure to the source’s localization uncertainty region. These plots have been reproduced for the other six sources in Figures 5.2 to 5.7.

in search-mode. The exact duration of these observations varies, but ranges between 1,400 and 1,800 seconds. We examine this search mode data using PRESTO (Ransom, 2011). Once we have discovered a reliable spin period, we can collect fold-mode data with CHIME/Pulsar.

We use a combination of search-mode and fold-mode data from CHIME/Pulsar to constrain properties of these new sources, including flux density and width. Where possible, we also construct timing solutions.



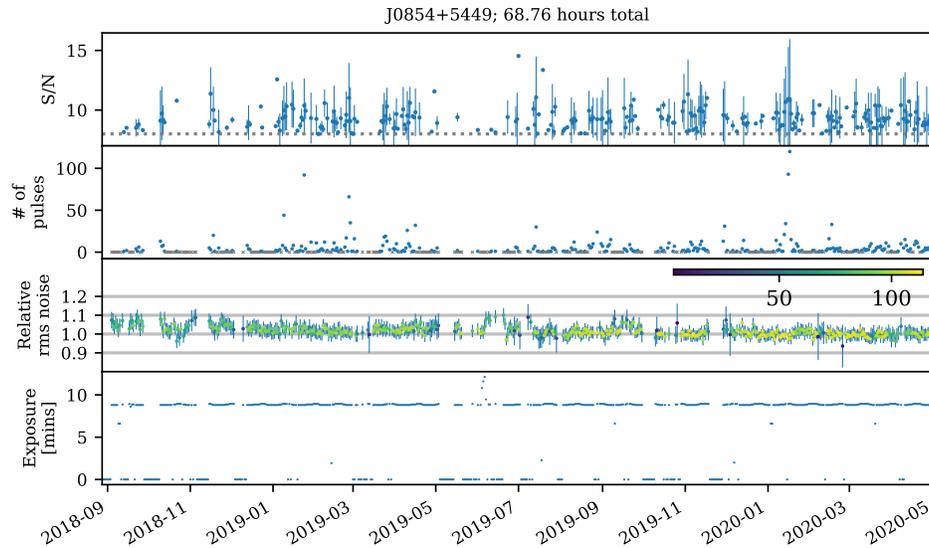
**Figure 5.2:** CHIME/FRB detections of new pulsar-like source PSR J0209+5759, which is detected consistently with CHIME/FRB. After an initial CHIME/FRB timing solution, this source was determined with CHIME/Pulsar data to be a nulling pulsar, with period 1.06 s. See Figure 5.1 for more details of the figure structure, and Section 5.3.3 for more discussion of the source.

### RFI Excision

We begin by removing a list of common known corrupted frequency channels in the CHIME frequency band with the PSRCHIVE (Hotan et al., 2004; van Straten et al., 2012) `paz` routine; this excises  $\sim 15\%$  of the band.

Before searching for single pulses, we use the PRESTO `rfifind` tool to determine individualized RFI masks (in addition to the known bad channels) for search-mode data. We also construct a common list of periodic RFI signals in CHIME/Pulsar data, colloquially known as a “birdies” file and apply it when appropriate.

To remove RFI from folded data, we again employ the PSRCHIVE `paz` routine, removing the known bad channels and allowing automated excision by using the built-in median difference filter algorithm. In some cases further RFI excision



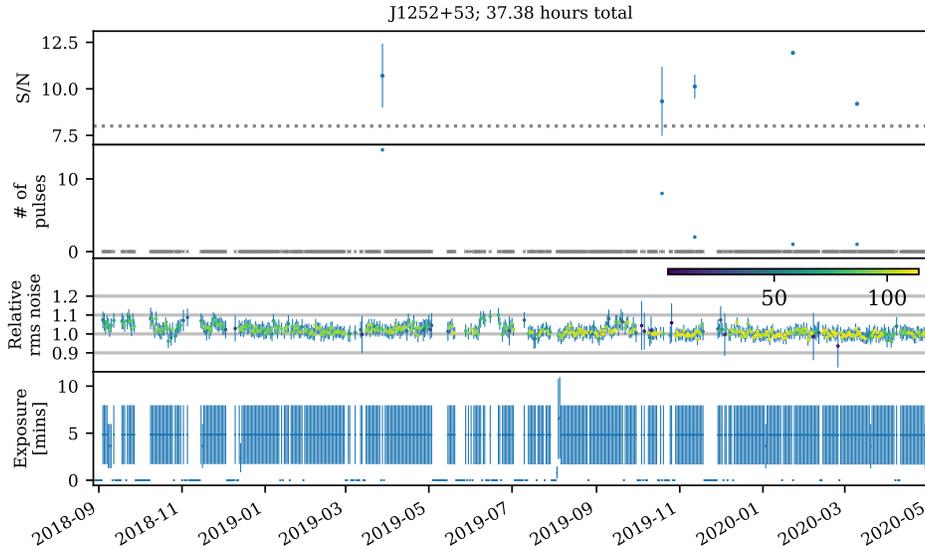
**Figure 5.3:** CHIME/FRB detections of new pulsar-like source PSR J0854+5449, which is detected consistently with CHIME/FRB. This pulsar has a period of 1.23 s, and was determined by CHIME/Pulsar data to be a conventional pulsar. See Figure 5.1 for more details of the figure structure and Section 5.3.3 for more discussion of the source.

was conducted using the interactive `pazi` tool.

### Single pulse detections

We analyze search-mode data from CHIME/Pulsar using PRESTO tools, searching for both single pulse detections and periodicity. Apparent single pulse detections are then examined in more detail using PSRCHIVE tools, the PRESTO waterfaller tool, or both to confirm detection. Once a source is confirmed, initial single pulse detections are also added to the public CHIME/FRB pulsar detection webpage.<sup>5</sup>

<sup>5</sup>New Galactic pulsars discovered by CHIME/FRB can be viewed at <https://www.chime-frb.ca/galactic>.

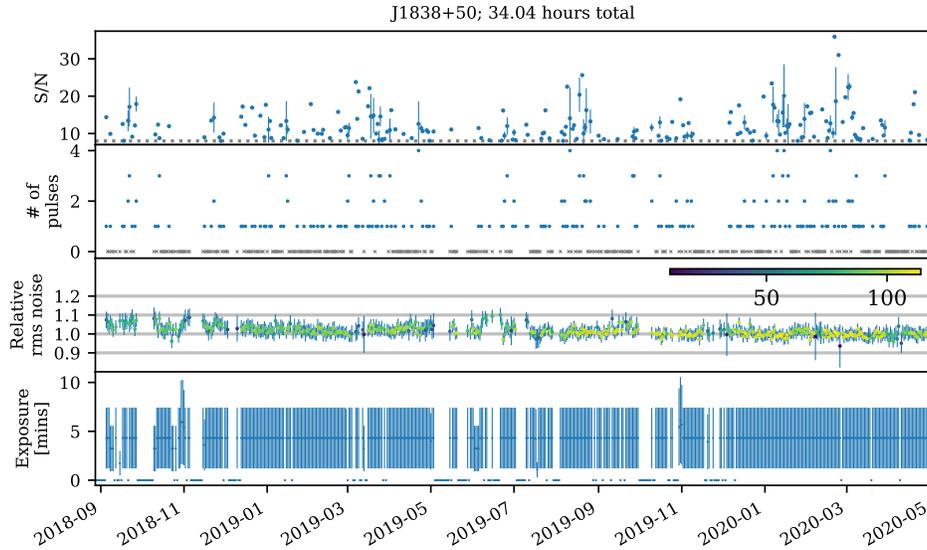


**Figure 5.4:** CHIME/FRB detections of new RRAT-like source PSR J1252+53. This source is detected only sparsely with CHIME/FRB and with CHIME/Pulsar. We were unable to determine a period estimate using CHIME/FRB data, but using CHIME/Pulsar data, we find an underlying period of 0.220 s. See Figure 5.1 for more details on the figure structure and Section 5.3.2 for more discussion of the source.

### Period Search

Where possible, we use PRESTO’s `accelsearch` function to search for the candidate’s period. Where `accelsearch` failed, we have also used fast folding algorithm (FFA) methods as in Morello et al. (2020) to search for periodicity. However, several of these sources are likely RRATs, making it challenging to robustly detect periodicity using standard pulsar tools.

In these cases, we also make use of brute force methods, in particular PRESTO utilities `single_pulse_search` and `rrat_period`, to determine periods. For comparison purposes, we also calculate spin periods using CHIME/FRB metadata as described in Section 5.2.2 where possible. The input to `rrat_period` is a list of times, at which the pulsar was detected within the span of one observation. It then computes intervals between those times and calculates the period which

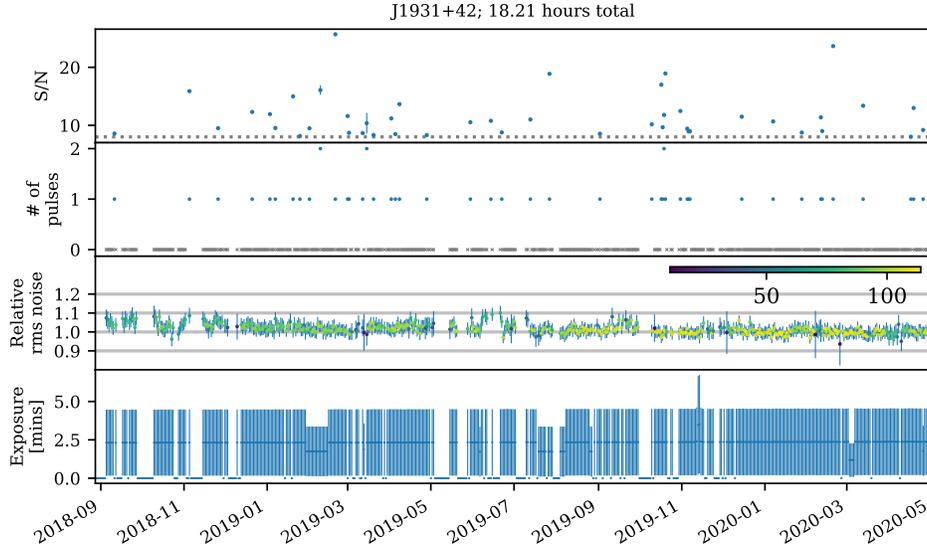


**Figure 5.5:** CHIME/FRB detections of new RRAT-like source PSR J1838+50. This source is detected regularly with CHIME/FRB and CHIME/Pulsar, allowing determination of a period with both sets of data. The period is measured to be 2.577 s. See Figure 5.1 for more details of the figure and Section 5.3.2 for more discussion for the source.

makes the number of periods in each interval as close to an integer as possible. However, our sources are highly intermittent and often `single_pulse_search` only detected two or three pulses within an observation.

As part of this analysis, UBC Ph.D. student Kathryn Crowter developed an extension to the existing PRESTO utility `rrat_period: rrat_period_multiday`. Both functions use brute force methods to search PRESTO single pulse search results for possible periodicity. `rrat_period_multiday` introduces a slight modification to allow for the inclusion of multiple observations in a single search. It is important to note that the function does not seek to account for gaps between observations: it merely enables us to combine information from multiple sets of intra-observation detection spacings. This tool has been used extensively for the sources in Section 5.3.2. These measured periods range from 0.220 s to 2.577 seconds.

As all observations are conducted within a narrow window in hour angle, we



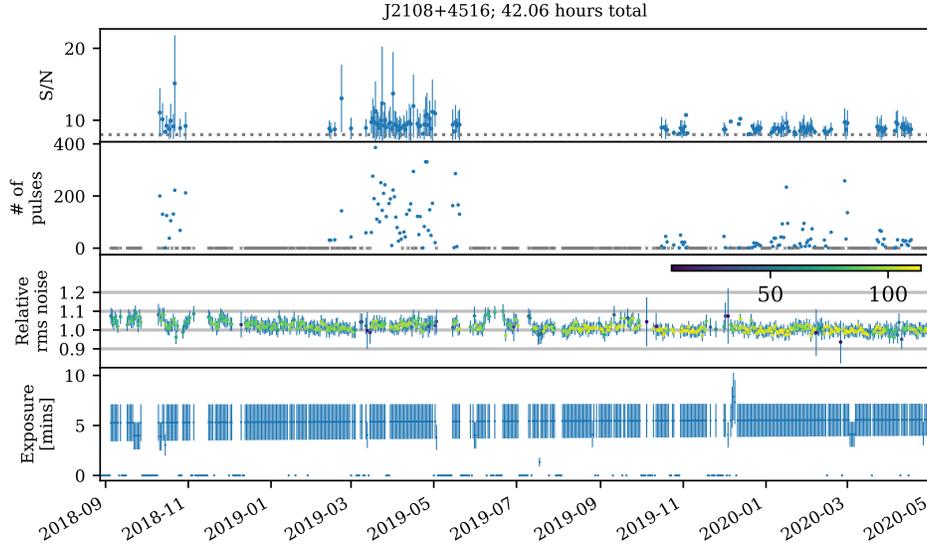
**Figure 5.6:** CHIME/FRB detections of new RRAT-like source PSR J1931+42, detected sparsely with CHIME/FRB. This is the only source presented in this work for which a period has not been determined. This is not unheard of for RRATs, but further observations may enable us to better characterize this source. See Figure 5.1 for more details on the figure structure and Section 5.3.1 for more information about the source.

cannot be certain that we have not captured a sidereal-day sampling alias with CHIME/Pulsar, i.e., the detected spin frequency  $f_{\text{det}} = f_0 + n/t_{\text{sidereal}}$  Hz, where  $f_0$  is the true frequency in Hz,  $n$  is a small integer and  $t_{\text{sidereal}}$  is the duration of a sidereal day in seconds. This degeneracy could be broken by additional observations with instruments other than CHIME.

### CHIME/Pulsar flux density estimation

To estimate band-averaged flux densities for each observation, this work employs the standard radiometer equation

$$S_{\text{peak}} = \frac{T_{\text{rec}} + T_{\text{sky}}}{G\sqrt{n_p\Delta V\Delta t}} \times \text{S/N}, \quad (5.1)$$



**Figure 5.7:** CHIME/FRB detections of new pulsar PSR J2108+4516. Its spin period is 0.577 s. This pulsar is notable for its intermittency, with months-long timespans where the pulsar is not seen by CHIME/FRB or CHIME/Pulsar. This behaviour is due to interaction within the binary system in which the source resides. See Figure 5.1 for more details about the figure and 5.3.3 for more information about the source.

where  $T_{\text{rec}}$  is the receiver temperature,  $G$  is the antenna gain,  $T_{\text{sky}}$  is the sky temperature,  $n_p = 2$  is the number of polarization streams summed,  $\Delta\nu$  is the effective bandwidth in Hz, and  $\Delta t$  is the integration time in seconds. Flux density calculations for this work were conducted by Bradley Meyers, a postdoctoral researcher at the University of British Columbia.

CHIME has an effective gain is  $G \approx 1.16 \text{ KJy}^{-1}$ , based on its illuminated collecting area of roughly  $80 \text{ m} \times 80 \text{ m} = 6400 \text{ m}^2$  and nominal 50% efficiency. In theory, this is also the antenna gain for the tied array beam formed by CHIME/Pulsar. In practice, imperfections our understanding of the CHIME beam and uncertainty in phase calibration leads to deviations from this value of  $G$  for CHIME/Pulsar. Therefore, a beam correction is applied in computing the effective gain to account for attenuation at large zenith angles and frequency variations. The complexity of the calculation is limited by considering only the gain for the mid-point of the ob-

servation, when the pulsar is centred in the primary beam) and computing the gain only at 600 MHz.

$T_{\text{rec}}$  is taken to be the nominal value for the temperature of receiver electronics and structures of CHIME, 30 K. Values for  $T_{\text{sky}}$  are calculated at the sources' nominal positions from the 408 MHz Haslam map (Haslam et al., 1982; Remazeilles et al., 2015), and scaled to 600 MHz assuming that the sky contribution scales as  $\nu^{-2.5}$  (e.g., de Oliveira-Costa et al., 2008).

Though the exact value for effective bandwidth varies from observation to observation, we treat the nominal value of  $\nu$  as 330 MHz, representing a masking fraction of about 15%. Recent analyses suggest that this may be an underestimate; telecommunications infrastructure continues to affect more and more of the CHIME band.

For mean flux densities, where appropriate, the equation  $S_{\text{mean}} = \delta S_{\text{peak}}$  is used, where the duty cycle  $\delta = A/S_{\text{peak}}$  and  $A$  is the area under the pulse profile.

Preliminary instrumental evidence suggests that the system temperature may be an underestimate, perhaps by a factor of two to three. Investigations into this apparent discrepancy are ongoing and combine efforts from CHIME, CHIME/FRB and CHIME/Pulsar. Therefore, the flux density values presented here should be considered lower limits only.

### **Pulse Widths**

We measure pulse widths by fitting a Gaussian profile to each single pulse or integrated profile (for persistent sources), using the `scipy.optimize.curve_fit` implementation of a non-linear least squares method. To ensure consistency between sources, this fit is completed after downsampling to 256 phase bands for all sources. Pulse widths are the full-width-at-half-maximum of the best fit Gaussian, i.e.  $W_{50} = 2\sqrt{2\ln 2}\sigma$ , where  $\sigma^2$  is the variance.

### **Timing Solutions**

When we are able to find a spin period, our next step is to attempt to build the best possible timing solution. We begin by attempting to fold search-mode data at the source's nominal period, using the PRESTO `prepfold` routine and Digital

Signal Processing Software for Pulsar Astronomy (DPSR; van Straten & Bailes, 2011).

Depending on the sources’ response to folding, our timing procedure bifurcates. Three can be folded using their nominal periods. For these sources, discussed in Section 5.3.3, we follow a standard pulsar timing procedure, folding all search-mode data using DPSR and excising RFI and generating TOAs with PSRCHIVE routines `paz` and `pat` respectively. We use a standard profile created with PSRCHIVE routine `paas` with a representative observation. We then determine timing solutions using TEMPO (Nice et al., 2015) and TEMPO2 (Hobbs et al., 2006). We also shift to collecting fold-mode data for persistent sources.

The other four sources are sufficiently RRAT-like that folding the data does not result in an improved detection. For these sources, discussed in Section 5.3.2, we use DPSR to create single pulse archives for each detection, generate a TOA for each single pulse, and use TEMPO and TEMPO2 to refine our initial single pulse period solution. Where available, we also make use of CHIME/FRB metadata solutions as starting points for a timing model.

Table 5.1 lists the period determination methods used for each source; methods used for a given source are designated by “Yes” and methods not used are designated by “No.” Due to the minimal data present for the four RRAT-like sources, we do not fit timing parameters besides the period. For the period, we fit with TEMPO, using single pulsar archives where folded archives are not available. Uncertainties in these periods are determined from this TEMPO fitting. However, these uncertainties should be regarded as substantial underestimates; these solutions are based on a small number of TOAs which are created from single pulse archives instead of folded archives. Where possible, we compare CHIME/Pulsar periods with CHIME/FRB metadata periods. Though both use data collected by CHIME, the CHIME/FRB and CHIME/Pulsar provide near independent determinations of spin period.

### 5.3 New sources detected with CHIME/FRB

For all sources, we report the source’s position, DM, rotation period, and burst rate in Table 5.3 and flux density in Table 5.4. Uncertainties are 68% confidence

**Table 5.1:** Period determination methods for each of the seven sources.

Pulsar Name	CHIME/FRB <sup>a</sup> metadata	Single Pulse TOAs	Folded TOAs	$N_{toas}$ <sup>b</sup>	Period <sup>c</sup> Determination	DM <sup>d</sup> Determination
PSR J0121+53	Yes	Yes	No	11	TEMPO	$S/N$ maximization
PSR J0209+5759	Yes	No	Yes	136	TEMPO	prepfold search
PSR J0854+5449	Yes	No	Yes	1,909 <sup>e</sup>	TEMPO	Subbanded timing
PSR J1252+53	No	Yes	No	10	TEMPO	$S/N$ maximization
PSR J1838+50	Yes	Yes	No	7	<code>rrat_period_multiday</code>	$S/N$ maximization
PSR J1931+42	No	No	No	–	–	$S/N$ maximization
PSR J2108+4516	Yes	No	Yes	124	TEMPO	PulsePortraiture

<sup>a</sup> See Section 5.2.2 for a discussion of CHIME/FRB metadata timing.

<sup>b</sup>  $N_{toas}$  is the number of TOAs used in the single single pulse TOA TEMPO analysis for RRAT-like sources or the folded TOA TEMPO analysis, depending on which is present for the source.

<sup>c</sup> See Section 5.2.3 for a description of `rrat_period_multiday`. See Section 5.2.3 for a discussion of TEMPO-based timing with single pulse and folded archives. <sup>d</sup> See Section 5.3 for an expanded discussion of DM determination.

<sup>e</sup> The TOA count for J0854+5449 is based on subbanded TOAs, instead of single TOAs per observation.

intervals.

Periods and flux densities for all sources are derived from CHIME/Pulsar data, as are the positions of PSR J0209+5759 and PSR J0854+5449. Period uncertainties are determined as outlined in Section 5.2.3, and flux densities are determined as outlined in Section 5.2.3. CHIME/Pulsar positions and their uncertainties are based on timing solutions. Burst rates and all other positions are determined from CHIME/FRB data, as discussed in 5.2.3.

DM uncertainties are determined uniquely for each source. For PSR J0854+5449, DM uncertainty is determined from the subbanded timing solution. For PSR J0209+5759, the DM uncertainty is determined based on `prepfold` DM search step size. For PSR J2108+4516, DM uncertainty is found with PulsePortraiture DM fitting (Penucci, 2019). For highly intermittent sources where folding is not possible, DM uncertainty is determined by finding the signal-to-noise maximizing DM during each detection and averaging.

In Figure 5.8, we present dynamic spectra or “waterfall” plots and single pulse profiles from the CHIME/Pulsar system for sources detected only in single pulses. For sources with complete pulsar timing solutions, we present timing residuals in Figure 5.9 and single observation profiles in Figure 5.10.

### 5.3.1 Source without spin period measurement: PSR J1931+42

PSR J1931+42 is the only source reported here for which we have not yet been able to determine a spin period. It was first detected by CHIME/FRB on 2018-08-03 (MJD 58333) with a DM of  $50.9 \text{ pc cm}^{-3}$ , and we have since recorded 59 pulses from the source with CHIME/FRB. We also conducted 40 CHIME/Pulsar search mode observations of PSR J1931+42 between MJD 58443 and MJD 58542.

This source has proved hard to detect, and harder still to detect with a sufficient number of pulses to enable a spin period determination. With CHIME/FRB, we have never seen the source more than twice in a single observation. We have observed two pulses on only seven days; all other days had only a single detection. With CHIME/Pulsar, we have detected PSR J1931+42 only four times (on MJDs 58523, 58534, 58537, and 58540), with two pulses detected on MJD 58523 and only one on the latter three days. The peak flux density of these pulses ranges

**Table 5.2:** Positions and dispersion measures for all new sources. Distances are determined by comparison with the NE2001 Galactic DM map (Cordes & Lazio, 2002).

Pulsar name	RA (hms)	Dec (dms)	DM ( $\text{pc cm}^{-3}$ )	$D_{\text{NE2001}}$ (kpc)
PSR J0121+53	$01^{\text{h}}21^{\text{m}} \pm 11^{\text{m}}$	$+53^{\circ}29' \pm 16'$	91.38(3)	$3.3^{+1.0}_{-0.7}$
PSR J0209+5759	$02^{\text{h}}09^{\text{m}}37.38^{\text{s}} \pm 0.03^{\text{s}}$	$+57^{\circ}59'45.35'' \pm 0.26''$	55.3(6)	$2.09^{+0.21}_{-0.23}$
PSR J0854+5449	$08^{\text{h}}54^{\text{m}}25.733^{\text{s}} \pm 0.002^{\text{s}}$	$+54^{\circ}49'28.81'' \pm 0.01''$	18.837(1)	$0.75^{+0.15}_{-0.14}$
PSR J1252+53	$12^{\text{h}}52^{\text{m}} \pm 13^{\text{m}}$	$+53^{\circ}42' \pm 17'$	20.70(3)	$1.00^{+0.33}_{-0.23}$
PSR J1838+50	$18^{\text{h}}38^{\text{m}} \pm 8^{\text{m}}$	$+50^{\circ}51' \pm 15'$	21.81(1)	$1.54^{+0.19}_{-0.17}$
PSR J1931+42	$19^{\text{h}}31^{\text{m}} \pm 7^{\text{m}}$	$+42^{\circ}30' \pm 5'$	50.90(2)	$3.13^{+0.44}_{-0.40}$
PSR J2108+4516	$21^{\text{h}}08^{\text{m}} \pm 7^{\text{m}}$	$+45^{\circ}16' \pm 4'$	82.4(3)	$3.37^{+0.37}_{-0.43}$

**Table 5.3:** Basic timing parameters for all new sources.

Pulsar name	Pulse Period <sup>a</sup> (s)	Pulse width <sup>b</sup> (ms)	Burst rate <sup>c</sup> (hr <sup>-1</sup> )
PSR J0121+53	2.7247846(4)	20(5)	2.4(2)
PSR J0209+5759	1.0639060415(1)	20(2)	21.4(5)
PSR J0854+5449	1.233032602667(5)	9.3(0.1)	21.5(6)
PSR J1252+53	0.22010358290(8)	20(3)	0.09(4)
PSR J1838+50	2.577223412(5)	13(2)	3.9(2)
PSR J1931+42	–	32(6)	8(1)
PSR J2108+4516	0.57722824(7)	15.1(1)	204(2)

<sup>a</sup>Based on CHIME/Pulsar single pulse detections unless otherwise stated in the text. Uncertainties are substantial underestimates for all sources using single pulse TOAs.

<sup>b</sup> The full-width-at-half-maximum of a Gaussian fit to the profile (or the median of all the values for each single pulse detected).

<sup>c</sup> Based on CHIME/FRB detections. The burst rate is calculated as the total number of detections with S/N > 8 divided by the total time CHIME/FRB has observed the candidate position, assuming with Poissonian uncertainty for burst numbers.

**Table 5.4:** Flux density measurements for all sources. The range of peak flux densities for the intermittent sources are given in the  $S_{600}^{\text{peak}}$  column, while for the more persistent sources we provide flux density estimates averaged over the entire pulse period in the  $S_{600}^{\text{mean}}$  column.

Pulsar name	$S_{600}^{\text{peak}}$ (mJy)	$S_{600}^{\text{mean}}$ (mJy)
PSR J0121+53	50–350	–
PSR J0209+5759	–	$\gtrsim 0.3$
PSR J0854+5449	–	$\gtrsim 0.5$
PSR J1252+53	100–230	–
PSR J1838+50	75–330	–
PSR J1931+42	25–150	–
PSR J2108+4516	–	$\gtrsim 3$

between 25–150 mJy at 600 MHz.

We can therefore provide only the weak constraint that the period must be less than  $\sim 98$  seconds, which is the smallest gap between detections. However, this is much too large a spacing to represent a single spin period, so the period remains unknown.

A single pulse from PSR J1931+42 is shown in Figure 5.8d.

### 5.3.2 Highly intermittent sources

Several of the sources reported have measured spin periods but display high intermittency. These sources may be RRATs or may be highly-intermittent pulsars, but do not have full timing solutions. Each period is initially determined via brute force methods using single pulse detections, then refined with pulsar timing software as discussed in 5.2.3.

Analysis of CHIME/FRB metadata suggests that many of our sources are detected at low  $S/N$ , near our FRB detection cutoff. It is possible that some portion of our source’s intermittency may not be intrinsic but derived from failure to detect lower luminosity emission from these sources. Additionally, our relatively large position uncertainties restrict our ability to optimally detect single pulses. These factors further muddy the RRAT vs. intermittent pulsar distinction, so we classify sources by their apparent intermittency and do not speculate as to their exact nature.

#### PSR J0121+53

PSR J0121+53 was first detected by CHIME/FRB on 2018-09-05 (MJD 58366) at a DM of  $87.4 \text{ pc cm}^{-3}$ . Single pulse detections with CHIME/Pulsar search-mode data indicate the DM is slightly higher:  $91.38 \text{ pc cm}^{-3}$ . Though this is a large discrepancy, such uncertainties are possible in initial detections due to CHIME/FRB’s tree dedispersion structure. For this reason, final DM calculations in CHIME/FRB publications such as CHIME/FRB Collaboration (2019) and Fonseca et al. (2020) do not publish initial DM determinations, but instead fit for a structure maximizing DM. The source is seen consistently in single pulse plots, but has not been successfully folded. This source has the highest DM of the seven sources presented here,

but since it is in the Galactic Plane it is at a distance of roughly 3.3 kpc (Cordes & Lazio, 2002). An example pulse is shown in Figure 5.8a. Though this pulse is single peaked, some pulses from this source have shown double-peaked structure, indicating complex morphological structure.

Analysis with the new multi-day RRAT period finder determines an underlying spin period of 2.725 s, consistent with the independently determined period from CHIME/FRB metadata, which found a spin period of 2.726 s. The final value is reported in Table 5.3 along with other parameters for this source.

### **PSR J1252+53**

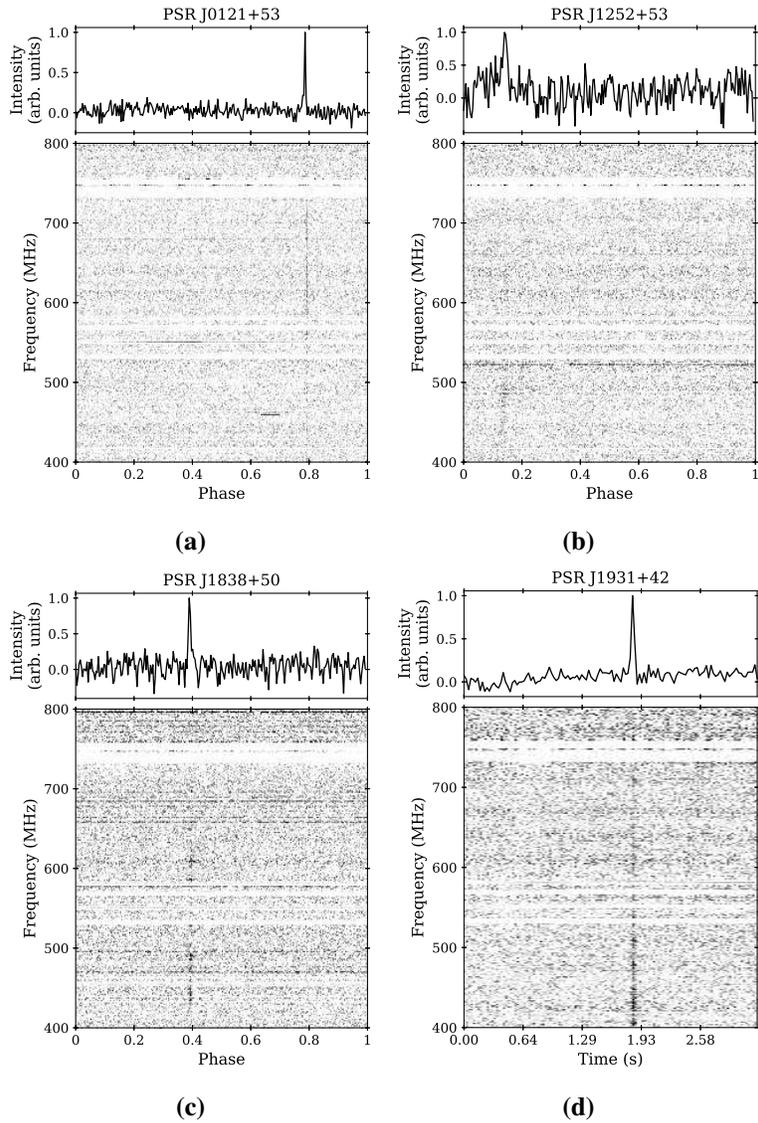
PSR J1252+53 was first detected by CHIME/FRB on 2018-10-23 (MJD 58414) at a DM of  $21.03 \text{ pc cm}^{-3}$ , and subsequently search mode observations were conducted using CHIME/Pulsar. Single pulse detections from CHIME/Pulsar filter-bank data suggest the correct DM is slightly lower than suggested by CHIME/FRB detections,  $20.70 \text{ pc cm}^{-3}$ . PSR J1252+53 appears only four times in over 100 near-daily observing sessions. Results from CHIME/Pulsar single pulse detections indicate that the underlying period is 220 ms; no CHIME/FRB metadata solution has been found.

Although detected, it is clear in Figure 5.8b that PSR J1252+53 has a lower S/N than our other six sources. It is also the source with the fewest CHIME/FRB and CHIME/Pulsar detections.

### **PSR J1838+50**

PSR J1838+50 was first detected by CHIME/FRB on 2018-07-27 (MJD 58326) with a DM of  $21.8 \text{ pc cm}^{-3}$ . PSR J1838+50 was followed up regularly using CHIME/Pulsar search mode observations between MJD 58463 and 58561 and again between MJD 59033 and 59123. CHIME/FRB metadata analysis and CHIME/Pulsar single pulse analysis independently find a period of 2.577 s. An example pulse is shown in Figure 5.8c.

Like PSR J0121+53, PSR J1838+50 is sufficiently intermittent that folding data for the duration of an observation does not improve the detection. Due to scheduling and data storage constraints, PSR J1838+50 has the fewest observations



**Figure 5.8:** Single pulses detected by CHIME/Pulsar from each highly intermittent source. The single-pulse profile versus pulse phase is in the inset panels and the frequency spectrum versus pulse phase in the main panels.

of the seven sources presented here. However, the CHIME/FRB burst rate suggests this source is moderately active. This raises the possibility that the solution could be substantially improved by continued observation.

### 5.3.3 Persistent Sources

The other sources presented here are all likely “slow” pulsars. CHIME/FRB is unlikely to detect millisecond pulsars due to its intrinsic time resolution (0.983 ms). These sources demonstrate some periods of inactivity or more limited activity, providing a possible explanation for the failure to detect them in previous pulsar searches.

For these sources, we are able to determine complete pulsar timing solutions. We present complete solutions and timing residuals for isolated pulsars PSR J0209+5759 and PSR J0854+5449 in Table 5.5 and Figure 5.9. We show pulse profiles for single observations for each source in Figure 5.10. We choose to present single observation profiles instead of average profiles as these sources are intermittent. Due to its greater complexity, the complete solution for PSR J2108+4516 is deferred to a future publication.

#### **PSR J0209+5759**

PSR J0209+5759 was initially detected by CHIME/FRB on 2018-09-06 (MJD 58367) with a DM of  $56.6 \text{ pc cm}^{-3}$ , but subsequent PRESTO analysis demonstrated the actual DM was slightly lower ( $55.3 \text{ pc cm}^{-3}$ ). This source was observed in search-mode by CHIME/Pulsar on 293 days between MJD 58411 to 58922. Since MJD 59100, we have been taking daily fold-mode observations of this source. PSR J0209+5759 displays a high degree of nulling, with 1,500–1,800 second observations generally including fewer than five single pulse detections. After folding, PSR J0209+5759 continues to display intermittent emission within observations. We did not detect a rotation measure for this source, using the method outlined in Ng et al. (2020). Analysis following the procedure outlined in Ng et al. (2020) finds a lower limit nulling fraction of 21%. However, we are able to determine a full timing solution for this source, with measured and derived parameters presented in Table 5.5. Timing residuals are shown in the first panel of Figure 5.9.

### **PSR J0854+5449**

PSR J0854+54 is a slow pulsar first detected by CHIME/FRB on 2018-07-25 (MJD 58324) with a DM of  $17.8 \text{ pc cm}^{-3}$ . CHIME/Pulsar subsequently conducted filterbank observations of this source from MJD 58592 to MJD 58868, when we moved to in fold-mode observations. Both the filterbank and fold-mode datasets are included in the final coherent timing solution presented here.

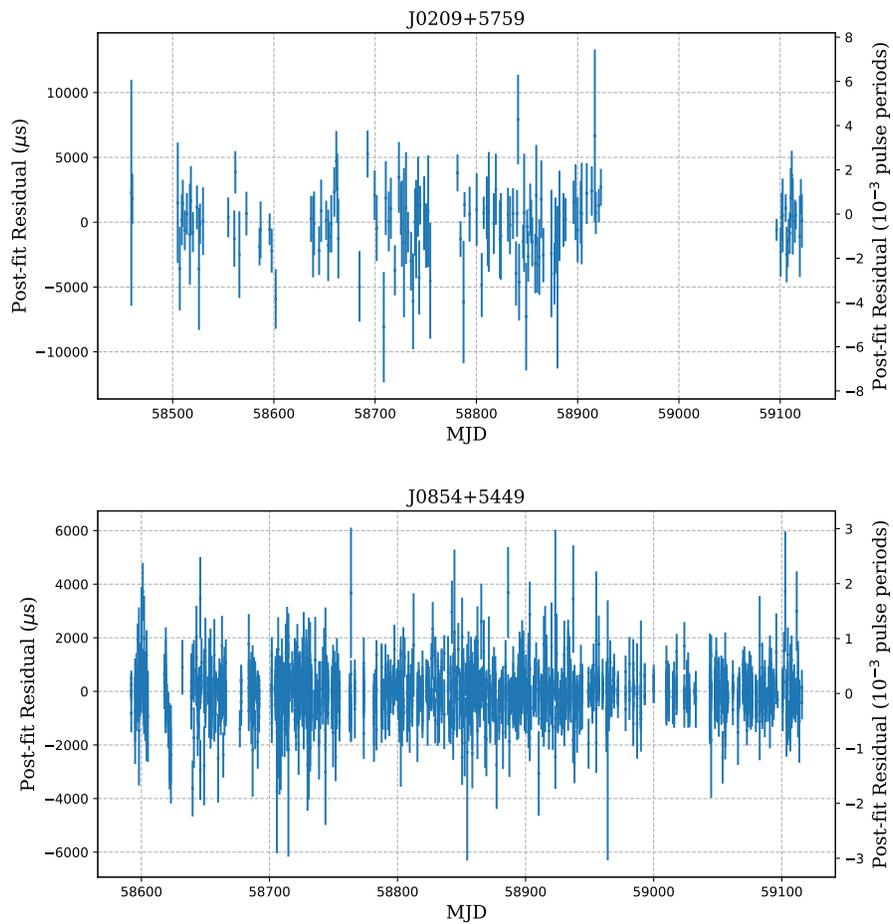
J0854+5449 has a rotation measure value of  $-8.8 \pm 2 \text{ rad m}^{-2}$ , determined using the procedure described in Ng et al. (2020). Measured and derived parameters are given in Table 5.5 and the timing residuals are shown in the second panel of Figure 5.9.

### **PSR J2108+4516**

PSR J2108+4516 was first detected by CHIME/FRB on 2018-10-11 (MJD 58402) and thereafter monitored with both filterbank and fold-mode observations by CHIME/Pulsar. This source is a complex binary system, consisting of a pulsar and a post-main sequence star. The system undergoes substantial DM and scattering variations, as well as lengthy periods of eclipse. It is the latter property that has likely prevented it from being detected before now. A detailed study of this source, including a full timing solution and discussion of the binary system’s properties, is being conducted by CHIME/Pulsar and CHIME/FRB team members. That work will be presented elsewhere (Andersen et al., in prep). A pulse profile for PSR J2108+4516 is shown in Figure 5.10.

### **5.3.4 Discussion of results**

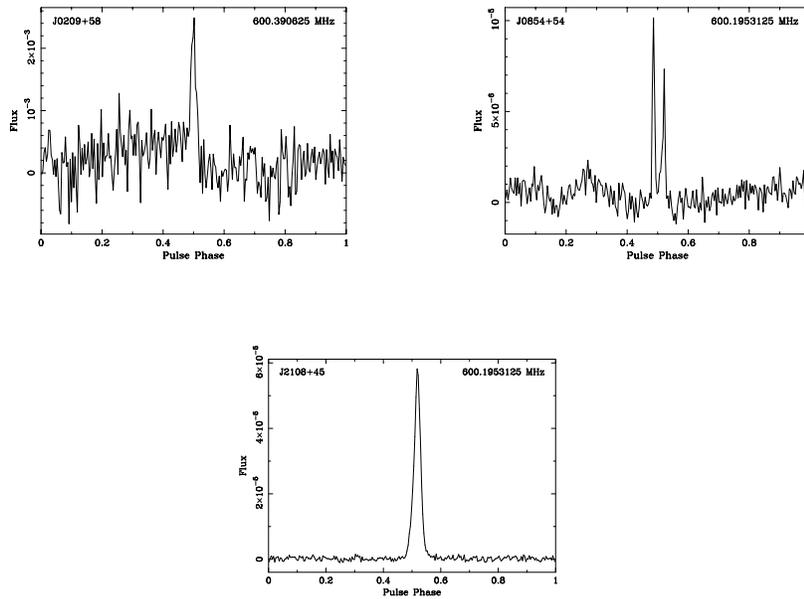
Although not primarily a pulsar discovery instrument, CHIME/FRB is a useful tool in detecting pulsars and RRATs using single pulse analyses. This method samples a distinctly different section of parameter space than conventional pulsar searching methodologies. As the initial detection method does not incorporate any folding, this method can detect only sources which are bright enough to be seen at least occasionally by CHIME/FRB in single pulses. However, CHIME/FRB passively observes the entire Northern sky each day; this suggests that if a source emits sufficiently brightly during the period of time when it is visible to the system, we



**Figure 5.9:** Best-fit timing residuals of PSRs J0209+5759 and J0854+5449 estimated from CHIME/Pulsar filterbank and fold-mode data. The smaller data set for J0209+5759 is due to fewer pulses emitted by the source, in a manner consistent with other known RRATs. Best-fit parameters of the timing model for these sources are shown in Table 5.5.

**Table 5.5:** Best-fit parameters and derived quantities for PSRs J0854+5449 and J0209+5759.

Global Parameters & Best-fit Metrics		
Pulsar name . . . . .	J0854+5449	J0209+5759
Reference epoch (MJD) . . . . .	58854.0	58790.0
Observing timespan (MJD) . . . . .	57592–59115	58459–59121
Number of Sub-bands . . . . .	32	1
Reduced $\chi^2$ . . . . .	1.04	1.03
RMS timing residual ( $\mu\text{s}$ ) . . . . .	338	1960
TOA uncertainty scale factor . . . . .	1.1	1.6
Ephemeris . . . . .	DE436	DE436
Clock Standard . . . . .	TT(BIPM2017)	TT(BIPM2017)
Timing Solutions		
Right ascension (J2000), $\alpha$ . . . . .	$08^{\text{h}}54^{\text{m}}25.7255(8)^{\text{s}}$	$02^{\text{h}}09^{\text{m}}37.304(16)^{\text{s}}$
Pulse frequency, $\nu$ ( $\text{s}^{-1}$ ) . . . . .	$0.8110085638790(6)$	$0.939932413620(12)$
Frequency derivative, $\dot{\nu}$ ( $\text{s}^{-2}$ ) . . . . .	$-2.0535(14) \times 10^{-16}$	$-1.22971(12) \times 10^{-14}$
Dispersion measure, DM ( $\text{pc cm}^{-3}$ ) . . . . .	$18.8368(15)$	$55.282\text{b}$
Derived Quantities		
DM distance from NE2001 (kpc) . . . . .	$0.75_{-0.14}^{+0.15}$	$2.09_{-0.23}^{+0.21}$
Galactic latitude, $l$ (deg) . . . . .	162.8	133.2
Galactic longitude, $b$ (deg) . . . . .	39.4	–3.3
Characteristic age, $\log_{10}(\tau_{\text{c}})$ (yr) . . . . .	7.79	6.15
Surface magnetic field, $\log_{10}(B_{\text{surf}})$ (G) . . . . .	11.8	12.5
Nulling Fraction . . . . .	–	0.21
Rotation Measure, $\text{rad m}^{-2}$ . . . . .	$-8.8(20)$	–



**Figure 5.10:** Single observations profiles for J0209+5759, J0854+5449, and J2108+4516. J0209+5759 has S/N 11.6, J0854+5449 has S/N 21.5, and J2108+4516 has S/N 196.

will detect that emission even if it is infrequent. Though CHIME/Pulsar observations require intentionally including the source in the automated scheduler, these observations can also be taken on daily or near-daily cadences. Very high cadence observations allow us to materially increase the chances that we observe an intermittent source when it is active.

The success of CHIME/FRB at detecting highly intermittent sources suggests that other FRB searches, e.g., ASKAP’s CRAFT survey (James et al., 2019) or MeerKAT’s MeerTRAP (Stappers, 2016) could be successfully applied to this task. CHIME/FRB has a particularly wide field of view and is fortunate to be guaranteed daily observations, but even FRB search instruments less optimized for this method could use metadata from Galactic detections to seed targeted pulsar searches.

### Comparison to previous results

One of the most salient questions for the implications of this work is a comparison to the population of known RRATs. At present, we cannot distinguish our new detection from the population of RRATs more broadly. Parameters of known RRATs were retrieved from the RRATalog,<sup>6</sup> and compared the distribution of DM, period, pulse width, and burst rate between the catalog and our sources. The 95% confidence intervals for the  $p$  values resulting from the bootstrapped Kruskal-Willis  $H$  tests this comparison between our sources and the known RRAT population are reported in Table 5.6. These calculations are discussed in more depth in Good et al. (2021).

**Table 5.6:** Confidence intervals of  $p$ -values from the bootstrapped Kruskal-Willis  $H$ -tests comparing the CHIME and known RRAT distributions.

Parameter	95% confidence interval	
	Lower	Upper
DM	0.009	0.878
Period	0.012	0.984
Burst rate	0.003	0.383
Pulse width		
..... Raw	0.004	0.135
.... Scaled	0.005	0.161

Though at present, our results are inconclusive, in the future CHIME will provide us with a uniquely well-honed RRAT observatory. One pitfall of such population level comparisons in the study of RRATS has been the differing selection functions amongst telescopes and surveys. CHIME enables us not only to detect many new RRATs but also to study these RRATs alongside known RRATs with CHIME/Pulsar. Such a survey is beyond the scope of this thesis but an exciting prospect for upcoming years.

### 5.3.5 Potential for Future Discoveries

In presenting a new procedure for discovering pulsars and RRATs, we must also be prepared to discuss how many pulsars and RRATs such a method can be expected

<sup>6</sup><http://astro.phys.wvu.edu/rratalog/>

to find. This is a more challenging task than it may initially appear, and a detailed discussion is deferred to Good et al. (2021) and Dong (2021).

One challenge in conducting this estimate was our relatively poor understanding of the single pulse characteristics of RRATs and pulsars. This is exacerbated by the relatively small known population of RRATs. We are therefore forced to extrapolate from a minimal number of datapoints and are left with a correspondingly ambiguous result. We estimate the number of potential sources, using existing population synthesis code, PsrPopPy (Bates et al., 2014) and drawing on the example of the Parkes Multibeam Survey. We estimate we will detect a maximum of 1500 sources with CHIME/FRB single pulse data (Good et al., 2021).

It is also important to note the the value reported here is based on the total number of pulsars and FRBs present in the universe: it does not attempt to distinguish between known and unknown objects. This means that our final calculation is not an estimate for the number of new pulsars discovered with CHIME/FRB but the total number of pulsars detected, including all previously known sources. Determining the exact number of pulsars seen with CHIME/FRB is not entirely straightforward, but the number is believed to be about 700. Thus, an estimate of up to 1500 detections suggests a maximum number of potential new sources around 800. In the coming months, we will be able to better characterize this number as we both detect more pulsars with CHIME and streamline our process for observing pulsars.

## 5.4 Conclusions & Future Work

The discovery of four new RRAT-like sources and three new pulsars represents only a starting point for detecting intermittent sources with CHIME/FRB and CHIME/Pulsar. Simulation work conducted with PsrPopPy altered for single bursts suggests that we could observe several hundred sources using CHIME/FRB single pulse triggers. Though these simulation results are preliminary and likely an overestimate, we do expect further detections with CHIME/FRB and are currently examining other candidate sources. Future work will improve these simulations in the future by using the CHIME/FRB injection system, used in the first CHIME/FRB catalog (CHIME/FRB Collaboration, 2021).

In this work, we focus only on reporting initial detections and solutions for these sources, but it is clear that there is more scope for inquiry within these seven sources. As we continue to detect more sources using CHIME/FRB, we will be able to better constrain the population of RRATs and pulsars detectable with single pulses.

The possibility of long-term monitoring of these sources and further sources we will likely detect with CHIME/FRB opens doors in studying intermittent pulsars and RRATs. In more systematic future work, we may be able to better understand the continuum between persistent and intermittent emission, including intermittent pulsars and RRATs. At present, it is difficult to compare our seven new sources to the existing population of RRATS, but further discovery of new RRATs and study of existing RRATs with CHIME/Pulsar will allow us to add new constraints to our understanding of the RRAT population.

## Chapter 6

# Follow-up Observations of CHIME/FRB Repeaters with Arecibo Observatory

*Besides learning to see, there is another art to be learned - not to see  
what is not.— Maria Mitchell*

### 6.1 Introduction

Fast radio bursts (FRBs) are luminous, millisecond transient signals detected at radio frequencies. First discovered in 2007, the published catalog was recently expanded by the addition of the CHIME/FRB catalog to a total of more than 500 known FRBs (CHIME/FRB Collaboration, 2021). Though initially hard to detect in large quantities, FRBs are apparently ubiquitous, with all-sky rates near 800 per sky per day above 5 Jy ms (CHIME/FRB Collaboration, 2021). Their large dispersion measures (DM) mark them as extragalactic, and some bursts appear to be located in extreme environments, e.g. (Michilli et al., 2018b).

The initial discovery of repeated bursts from FRB 20121102 with Arecibo Observatory (AO) (Spitler et al., 2014, 2016; Scholz et al., 2016), revolutionized the field as repeating FRBs allow follow-up observations. For FRB 20121102, these follow-up studies yielded a localization, a measurement of RM, and the discov-

ery revealed complex and interesting burst sub-structure. This structure, discussed in more detail in Chapter 1, shows the burst drifting down in frequency as time progresses and may hold clues to the emission mechanism of repeating FRBs and relevant propagation effects (Hessels et al., 2019; Metzger et al., 2019; Lyutikov, 2019; Cordes et al., 2017). Such structure has also been observed in other repeating FRBs detected by CHIME/FRB (CHIME/FRB Collaboration, 2019; Fonseca et al., 2020).

One of the most important questions regarding repeating FRBs is the nature of their repetition. Until recently, repeater bursts had been seen to be non-Poissonian, but had not been seen to be periodic (Oppermann et al., 2018). However, FRB 20180916B has been shown to have a  $16.35 \pm 0.18$  day periodicity detectable with CHIME, and FRB 20121102 has also been shown to be periodic (CHIME/FRB Collaboration, 2020b; Rajwade et al., 2020; Cruces et al., 2021). Like the initial discovery of FRB repetition, this has the potential to add exciting new angles to the study of FRBs; in particular, periodic repetition allows the possibility of predicting epochs of repeater activity.

The relationship between repeating and single burst FRBs remains unclear. All FRBs may be repeaters, but with a wide range of repetition rates, or some FRBs may repeat and some may not, indicating two different source classes (Ravi, 2019; Caleb et al., 2019a; Connor et al., 2020). Understanding the repetition behaviour of a large sample of FRBs and understanding in detail the behaviour of individual repeating FRBs are important steps towards understanding repeaters' place in the FRB ecosystem.

As we seek to understand repeating FRBs, their burst properties are of significant interest. Some of these bursts show complex downward-drifting time-frequency sub-structure as in FRB 20121102, FRB 20180814A, and other repeaters (Hessels et al., 2019; CHIME/FRB Collaboration et al., 2019a; CHIME/FRB Collaboration, 2019; Fonseca et al., 2020). High time resolution data enables us to examine such structures in depth, as in e.g. the detection of sub-30  $\mu$ s structure FRB 20121102 with AO (Michilli et al., 2018b). Additionally, developing a catalog of bursts from repeaters other than FRB 20121102 and FRB 20180916B will allow us to better determine what characteristics are common among repeaters and what are unique to particular sources.

Another open question about repeating FRBs is their luminosity distribution. Some have proposed a steep distribution reminiscent of emission from the Crab pulsar (Connor & Petroff, 2018; Lawrence et al., 2017). Such a distribution could explain the complicated repetition properties for many bursts, as lower sensitivity experiments like CHIME/FRB are insensitive to a potential dim sub-population of bursts. Recent results from observations of ASKAP’s FRB 170119 with the Green Bank Telescope revealed two repeat bursts a factor of about 600 fainter than the original ASKAP detection, lending strong support to the existence of low-luminosity bursts from repeating FRBs (Kumar et al., 2019). This conclusion is also supported by ongoing observations of FRB 20121102 with FAST<sup>1</sup>. Failure to detect dimmer bursts from known repeaters with a more sensitive telescope and greater exposure time would provide evidence against such an underlying population as an explanation the irregular repetition rates observed in many repeaters.

Examining these repeating FRBs in depth also provides the opportunity to study the repetition behaviour and morphology a broader range of repeating sources. One of the reasons FRB 20121102 and FRB 20180916B have both been studied repeatedly is their frequent repetition. The sources we propose to observe here are qualitatively less active than FRB 20180916B and FRB 20121102, providing us with the opportunity to obtain the same kind of detailed high time resolution and polarization data for what may be more “typical” repeating FRBs. Studying sources with a wide range of repetition rates can offer insights into the key open question of whether all FRBs repeat.

CHIME/FRB is a powerful FRB detection instrument, but it has comparatively low time resolution in intensity data (only 0.983 ms) and is not very sensitive to dim bursts (CHIME/FRB Collaboration, 2018). AO observations enable us to obtain data with greater sensitivity than CHIME/FRB intensity data and also provides polarization information and higher time resolution than is available in CHIME/FRB intensity data. This offers us the best opportunities examine short-timescale burst sub-structure and other properties like pulse luminosity function as well as to detect faint bursts.

Among the repeaters CHIME/FRB has detected to date, several are in AO’s

---

<sup>1</sup>Unpublished results shown at the FRB2020 virtual meeting in July 2020.

declination range. This work reports on follow-up observations of two of these of these low-declination repeating sources with Arecibo Observatory. More information about these bursts is reported in Table 6.1. There are also a number of single burst FRBs detected with CHIME/FRB with morphologies suggestive of possible repetition, particularly the distinctive downward-drifting frequency-time structure commonly seen in FRBs. These sources were also followed up with AO to search for repetition. More information about these bursts are reported in Table 6.2.

## 6.2 Motivation

We are beginning to have detailed pictures of the behaviour of our most popular repeating FRBs, namely FRB 20121102 and FRB 20180916B. The former has been studied extensively since it was found to repeat in 2016. Likewise, since FRB 20180916B was discovered to be the third repeater, it has been subject to extensive study. However, we do not find that the two sources behave in ways that are identical. First, they are in very different host environments. FRB 20121102 is in a high RM environment within a dwarf galaxy, while FRB 20180916B is in a Milky Way-like spiral galaxy, in a low RM environment. (See Chapter 1 for discussion of RM.) In terms of repetition, both demonstrate non-Poissonian grouped variation in their burst rates. Both have now been shown to have periodic activity, with 161 and 16 day periods between the start of active phases (Cruces et al., 2021; CHIME/FRB Collaboration, 2020b). And yet, that in itself is striking – the FRB 20121102 quiescent phase is approximately ten times as long as that of FRB 20180916B period. Is one typical or are they representative of a range of possibilities? We do not yet know. In one symmetry between the two sources, both demonstrate a variety of morphologies, including downward-drifting structure.

As we have many more observations for FRB 20121102 and FRB 20180916B than other repeaters, we are tempted to take our cues in possible repeater properties from these sources. Yet, given a population of only about 20 published repeaters and specific focus on only two, how are we to say whether these bursts are in any way representative of a broader population of repeating FRBs? Indeed, we have some reason to believe these FRBs may be non-standard. Specifically, these two bursts may be much more active than the average repeater. If true, we run a serious

risk of misunderstanding what exactly a repeater is like, in particular our sense of how probable it is that a source which has not yet had repeat bursts will at a later time. With this in mind, it is important to add further detailed observations of other repeating FRBs. This is one of the major motivations for this project.

Our sources were selected based on an additional motivation to mitigate an instrumental bias of CHIME. As a transit telescope, CHIME observes generally higher declinations for longer times with some declinations even being perpetually visible. On the other hand, CHIME/FRB observes declinations near its southern boundaries for only about 10 minutes per day, even including side-lobe detections. This is obvious in the spatial distribution of repeating FRBs discovered with CHIME, shown in Figure 6.1; detecting repetition from sources at low-declination can be challenging. Though we found in CHIME/FRB Collaboration (2021) that there does not appear to be a systematic bias in the location of repeaters, in-depth analysis of these low declination FRBs may be difficult. For this project we sought specifically to examine low declination FRBs, those south of  $+38^\circ$  declination. This corresponds to the northern limit of the AO declination range.

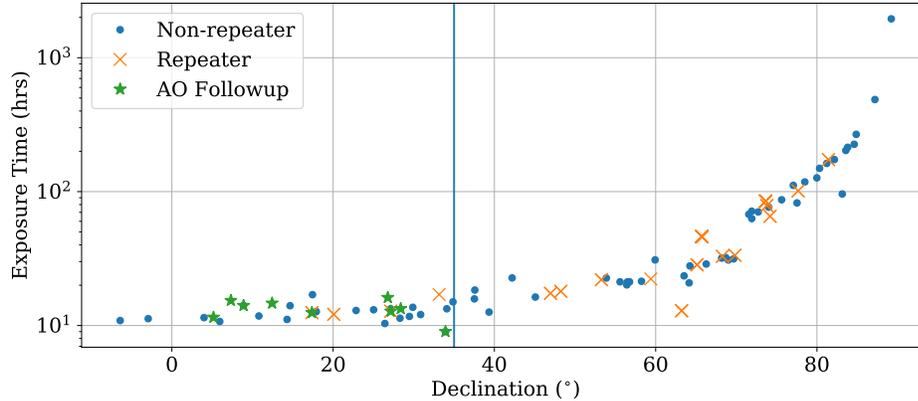
Another motivation for this project was the potential for low luminosity populations invisible to CHIME/FRB but visible to the more sensitive AO. The sensitivity gains from observing with AO may be hampered by the frequency dependent effects of scattering in the interstellar medium, but on the whole, we can detect bursts with AO that would be sub-threshold in CHIME/FRB.

## **6.3 Technical Details of Follow-up Program**

### **6.3.1 Source information**

The sources discussed in this chapter were chosen for further observation specifically for their combination of low-declination and either known repetition or repeater-like morphology. As morphology was a key factor in selecting these sources, particularly those without known repetition, we show dynamic spectra plots for each in Figures 6.2, 6.3, and 6.4.

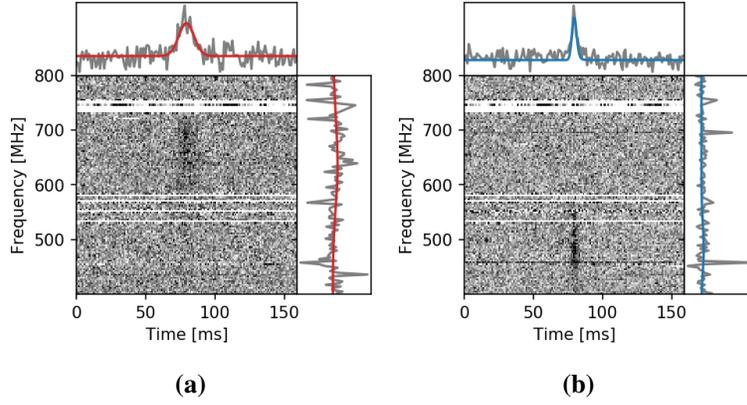
Our complete source list is shown in Tables 6.1 and 6.2. This project was conducted over the course of about one year, or two complete semesters of AO



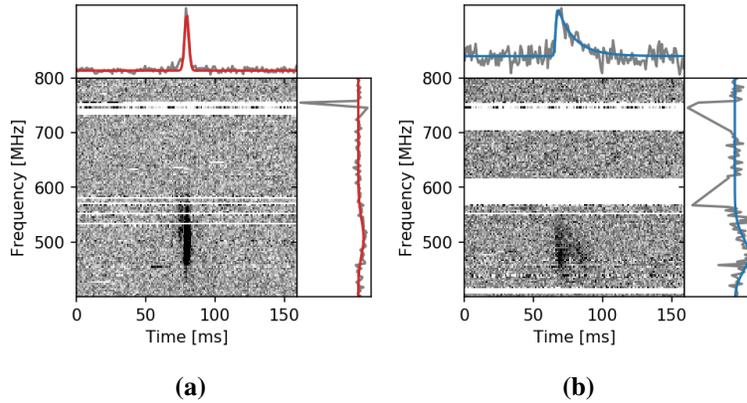
**Figure 6.1:** This figure plots declination (in degrees) against CHIME/FRB exposure time (shown on a log scale) in hours, for all bursts included in the first CHIME/FRB catalog. Arecibo Observatory’s maximum declination is approximately  $+38^\circ$  (shown as a blue vertical line). Sources followed-up in this work are marked as green stars; blue dots are apparently non-repeating FRBs. The additional repeating FRBs within AO’s declination range were observed after the proposal deadline for the data displayed here.

observing time. During the first semester, we observed FRB 20190116A, then the only known repeater in the relevant declination range, as well as several sources which are not known to be repeaters but had demonstrated the downward drifting frequency structure. Due to the difficulty in quantifying the meaning of a non-detection for unconfirmed repeaters and the slightly divergent science cases, we did not continue to observe the morphologically interesting one burst sources the following semester. In that semester, we observed 2 sources, FRB 20190116A and FRB 20190117A. In the the following semester, we intended to observe an additional set of sources, but due to 327 MHz issues in July 2020 and the auxiliary cable breakage in August 2020, we were only able to collect a handful of sessions, which are not presented here.

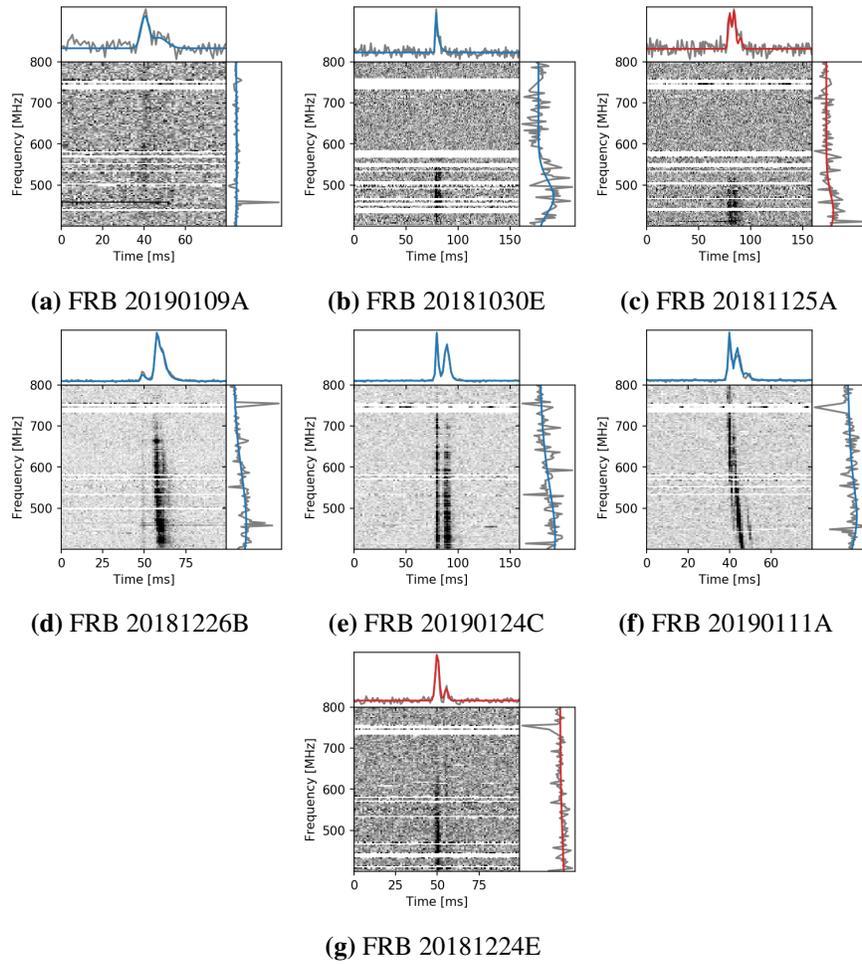
As CHIME/FRB is a transit telescope at roughly  $49.25^\circ$  N, exposure is very limited for bursts near the Southern declination limit of CHIME/FRB. In Tables 6.1 and 6.2, we include the total exposure time with CHIME/FRB, as well as with



**Figure 6.2:** Time vs. frequency dynamic spectra for the two known bursts from FRB 20190116A. The red or blue traces represent fits to the profile in time (top panel) and frequency (right panel). The colour indicates whether (blue) or not (red) the final burst fit includes scattering information. Neither is a clear demonstration of downward-drifting morphology; the difference in burst widths and bandwidths highlight the variability of repeat bursts.



**Figure 6.3:** Time vs. frequency dynamic spectra for two published bursts from FRB 20190117A. Six total bursts from this source have been detected by CHIME/FRB. This source, especially in the burst shown in panel (b) demonstrates clear downward-drifting structure.



**Figure 6.4:** Time-vs-frequency dynamic spectra for morphologically interesting FRBs at declination  $< +38^\circ$  detected with CHIME/FRB and followed-up with Arecibo Observatory. All plots are from CHIME/FRB data (CHIME/FRB Collaboration, 2021)

AO and the date range over which the source was observed with AO. All exposures for CHIME/FRB are based on the period of time included in the CHIME/FRB catalog, i.e., 2018-08-30 to 2019-07-01. AO observations for FRB 20190116A were conducted between MJD 58712 and MJD 59009. AO observations for FRB 20190117A were conducted between MJD 58888 and MJD 59010. AO observations for all apparently non-repeating sources were conducted between MJD 58695 and MJD 58848.

### 6.3.2 Search Strategy

The first step in determining a strategy for repeater follow-up was selecting an observatory and designing the observing program. A challenge we face in follow-up observation is the comparatively poor localization available with CHIME/FRB data. In general, our positional uncertainties are of the order  $10'$ , while many telescope beams are much smaller. This means we are often forced to conduct gridded observations to determine precise localizations. Alternatively, if localization is not the goal of our follow-up project, we can take a different approach. We can instead use a receiver with a similarly proportioned beam to observe a single patch of sky for a longer duration. In this case, we chose to use the Arecibo Observatory 327 MHz receiver, which is of a similar beam size as a CHIME/FRB formed beam. See comparison of receiver properties in Table 6.3.

The 327 MHz receiver operates at a complementary frequency range to that of CHIME/FRB, and a relevant one for FRB astronomy. Recent detections with the Green Bank Telescope and Sardinia Radio Telescope demonstrate that lower frequencies are fruitful for FRB survey (Chawla et al., 2020; Parent et al., 2020; Pilia et al., 2020).

One major advantage to AO when compared to CHIME/FRB is time and frequency resolution. With AO 327 and the PUPPI backend, we can set nearly arbitrarily high time and frequency resolution. With CHIME/FRB, our options are more limited. For this project we observed with  $10 \mu\text{s}$  resolution and downsampled to  $80 \mu\text{s}$  for initial analysis, providing about a factor of ten higher time resolution than CHIME/FRB intensity data. Had we successfully detected any bursts, we would have had exquisite resolution, a key consideration in ascertaining whether

**Table 6.1:** A summary of positions and dispersion measures for known repeaters discovered with CHIME/FRB and observed with AO. Reported DMs are structure optimizing. Methods for determining DM and position are explained in detail in (CHIME/FRB Collaboration, 2019) and (Fonseca et al., 2020).  $N_{\text{bursts}}$  refers only to CHIME/FRB bursts; there were no detections with AO. All exposure times for CHIME/FRB data are calculated between August 30, 2018 and July 1, 2019.

Source	RA (hh:mm)	Dec. (dd:mm)	DM ( $pc/cm^3$ )	$N_{\text{bursts}}$	CHIME/FRB exposure exposure (hrs)	AO exposure exposure (hrs)
FRB 190116A	$12:49 \pm 8'$	$27:09 \pm 14'$	441(2)	2	$12.8 \pm 7.69$	27.1
FRB 19017A	$22:07 \pm 8'$	$17:23 \pm 15'$	393.6(8)	6	$12.5 \pm 6.91$	20.1

**Table 6.2:** A summary of positions and dispersion measures for repeater-like FRBs discovered with CHIME/FRB and observed with AO. Reported DMs and locations are those presented in CHIME/FRB Catalog 1, and determined by the methods discussed therein. All exposure times for CHIME/FRB are calculated between August 30, 2018 and July 1, 2019.

Source	RA (hh:mm)	Dec. (dd:mm)	DM ( $pc/cm^3$ )	CHIME/FRB exposure (hrs)	AO exposure (hrs)
FRB 20190109A	07:12	05:09	324.6	$11.52 \pm 6.69$	4.78
FRB 20181030E	09:03	08:53	159.69	$14.1 \pm 4.87$	4.00
FRB 20181125A	09:51	34:06	272.19	$9.01 \pm 8.70$	2.90
FRB 20181226B	12:10	12:36	287.04	$14.68 \pm 4.79$	4.85
FRB 20190124C	14:29	28:23	303.64	$13.4 \pm 5.13$	2.12
FRB 20190111A	14:28	26:49	171.97	$16.2 \pm 5.13$	3.07
FRB 20181224E	15:57	07:19	581.85	$15.4 \pm 3.61$	2.69

**Table 6.3:** Basic properties of Arecibo 327 MHz Receiver and of CHIME/FRB.

	AO 327	CHIME/FRB
Beam Width	$4' \times 15'$	$40 - 20'$ <sup>a</sup>
Centre Frequency (MHz)	327	600
Bandwidth (MHz)	50	400
System Temperature (K)	115	$\sim 50$
Antenna Gain (K/Jy)	10	1.16

<sup>a</sup> Due to CHIME/FRB’s large fractional bandwidth, the formed beam width changes substantially over the band.

the downward-drifting spectral structure is truly only present in a small subset of bursts or whether it is simply lost in time resolution. Placing new constraints on the temporal limits of structure might exist within FRBs would have been valuable takeaway.

As an additional advantage, AO was relatively undersubscribed at the right ascensions we were seeking to observe (away from the Galactic plane). Our proposals were highly successful; we were awarded time for each (though rarely the full time request) and collected almost 200 hours of total telescope time.<sup>2</sup> We have compiled a large scale dataset, with many epochs of observation.

We conducted our observations in coherent dedispersion search mode, using the PUPPI backend. We used coherent search mode for this project as the DM of each source was known from CHIME/FRB detections.

Data from AO were analyzed using standard single pulse search methods in PRESTO. The process of searching for a repeat burst from a known FRB is much simpler than searching a data stream for an unknown FRB and identical to the process used in Chapter 5 to detect single pulses from RRATs. For starters, we used coherent dedispersion at the expected DM, using our foreknowledge to work more efficiently. Before data is downsampled to  $80 \mu\text{s}$ , we search it for RFI using PRESTO’s `rfifind` method. We then used a downsampling script developed by Dr. Scott Ransom to reduce the time and frequency resolution fo the data. We determine an optimal dedispersion plan with `DDplan.py` and follow it to look

<sup>2</sup>The sum of the exposure time for each source is less than the total time allocated, due to test pulsar observations, calibration scans, slew time, and a variety of instrumental and observing snafus.

for bursts both slightly above and slightly below the nominal DM of our source. Finally, we run `single_pulse_search` to find signals from single pulses. Anything which appeared in the `single_pulse_search` with  $S/N > 6$ , we then examined manually. We found no bursts from any of our sources during our follow-up program.

## 6.4 Results

During the course of this project, we did not observe repeat bursts from any of the sources we observed. It is worth noting that not only did we not observe any burst with AO, no further bursts from any of these sources were observed by CHIME/FRB during this period. This means that none of our seven non-repeating, morphologically interesting bursts have been observed to be repeaters at this time. Similar bursts from CHIME/FRB with such structure have been subsequently shown to be repeaters, but none of these bursts has yet been seen to repeat. This also suggests genuinely rare activity from known repeaters observed, particularly FRB 20190116A, which has not been observed to burst since its initial discovery in January 2019.

In the remainder of this chapter, we will set upper limits on the repetitions from each source, using both Poissonian and extended methods. We will also discuss possible explanations for the non-detection of further repeat bursts and the possible implications of this low-repetition rate within the context of the broader FRB repetition discussion.

### 6.4.1 Understanding Sensitivity Limits

In order to fully understand the implications of our non-detection, we must set a sensitivity limit for our survey. To do this, we follow the extension to the radiometer equation derived by Cordes & McLaughlin (2003) which says that the minimum detectable flux density for a radio transient with a given receiver is

$$S_{\min} = \frac{\beta S/N (T_{\text{rec}} + T_{\text{sky}})}{GW_i} \sqrt{\frac{W_b}{n_p \Delta \nu}}, \quad (6.1)$$

where  $\beta$  is an instrumental factor accounting for digitization loss, (approximately equal to 1),  $S/N$  is signal-to-noise ratio,  $T_{\text{rec}}$  is receiver temperature,  $T_{\text{sky}}$  is average sky temperature,  $G$  is antenna gain,  $W_i$  is intrinsic pulse width,  $W_b$  is broadened pulse width,  $n_p$  is number of polarizations and  $\Delta\nu$  is bandwidth.

The broadened width,  $W_b$  is the quadrature sum of the pulse's intrinsic width, the sampling time  $t_{\text{samp}}$ , the scattering time  $t_{\text{scatt}}$ , and the per-channel dispersive delay  $t_{\text{chan}}$

$$W_b = \sqrt{W_i^2 + t_{\text{samp}}^2 + t_{\text{chan}}^2 + t_{\text{scatt}}^2}. \quad (6.2)$$

The dispersive delay  $t_{\text{chan}}$  as derived in e.g. Lorimer & Kramer (2005) is

$$t_{\text{chan}} = 8.3\mu\text{s} \left( \frac{\Delta\nu_{\text{chan}}}{\text{MHz}} \right) \left( \frac{\nu}{\text{GHz}} \right)^{-3} \left( \frac{\text{DM}}{\text{pc cm}^{-3}} \right), \quad (6.3)$$

where  $\Delta\nu$  is the frequency channel bandwidth and  $\nu$  is the central observing frequency.

For known repeating bursts, FRB 20190116A and FRB 20190117A, we have taken the average of measured width parameters and scattering parameters; for the potential repeaters, we have used the measured CHIME/FRB parameters. For comparison purposes, we have also calculated the Cordes & McLaughlin (2003) sensitivity threshold for these sources at CHIME, using the same parameters. These results are shown in Table 6.4.

It is apparent that the AO sensitivity limits are better than the CHIME/FRB sensitivity limits, though not as dramatically as a simple radiometer equation calculation might suggest. This is largely due to the effects of scattering. Scattering timescales for CHIME/FRB detections are measured relative to 600 MHz. To determine the scattering timescales for our AO observations, we have used the standard scaling used for pulsar scattering (introduced in Chapter 2:  $\tau_s \propto \nu^{-4}$ . Scaling to 327 MHz, therefore, greatly increases the scattering timescales from those measured at 600 MHz, limiting the increase in sensitivity we are able to achieve by observing with AO.

The sensitivity limit for FRB 20181030E is dramatically higher than for the other bursts. This is due to the fact that FRB 20181030E is much narrower than other bursts. Though in its initial detection, it appeared morphologically complex,

**Table 6.4:** Sensitivity thresholds for AO and CHIME/FRB, calculated using the formalism of Cordes & McLaughlin (2003) and the parameters in Table 6.3. FRB 20190116A and FRB 20190117A (bolded) are known repeaters .

Source Name	AO 327 $S_{\min}$ (327 MHz)	CHIME/FRB $S_{\min}$ (600 MHz)
<b>FRB 20190116A</b>	$0.84 \pm 0.12$	$1.0 \pm 0.15$
<b>FRB 20190117A</b>	$0.51 \pm 0.068$	$0.88 \pm 0.089$
FRB 20190109A	$0.29 \pm 0.16$	$0.50 \pm 0.35$
FRB 20181030E	$2.5 \pm 0.11$	$3.2 \pm 0.10$
FRB 20181125A	$0.30 \pm 0.049$	$0.53 \pm 0.12$
FRB 20181226B	$0.48 \pm 0.045$	$0.73 \pm 0.084$
FRB 20190124C	$0.55 \pm 0.025$	$0.80 \pm 0.046$
FRB 20190111A	$0.56 \pm 0.018$	$0.93 \pm 0.032$
FRB 20181224E	$0.74 \pm 0.076$	$0.82 \pm 0.10$

subsequent analysis for the CHIME/FRB catalog has revealed that it is in fact a moderately narrow burst with structure likely better described by scattering than by sub-burst drifting.

#### 6.4.2 Constraints on Burst Rate

Though we embarked on the project with several potential scientific results, our non-detection means that the only repeater property we can discuss in depth based on these results is repeat rate.

As repeater science is a young field, there is not a definitive correct distribution to use for predicting FRB burst rates. With that in mind, we approach this problem from two directions. First, and simplest, we calculate limits on burst rate based on a Poisson distribution. For a Poisson process with burst rate  $r$ , the distribution of intervals  $\delta$  between bursts is exponential

$$P(\delta|r) = r e^{-\delta r}. \quad (6.4)$$

This is likely not the best approach to describe burst rates, as known repeaters show clear evidence of burst clustering and some have observed periodicity. However,

in the absence of reliable priors for periodicity, this remains a reasonable starting point for our analysis.

To determine our repetition rate, we follow the example of CHIME/FRB Collaboration (2019) and present both observed rates and scaled rates. All of our rates have Poisson confidence intervals, calculated as outlined in Kraft et al. (1991), and as implemented in `astropy.stats`. We calculate rates for CHIME/FRB-only and for CHIME/FRB-AO combined data. We do not calculate AO-only rates, as these would be trivially zero. The exposure time and uncertainty in exposure time for CHIME/FRB bursts is calculated as described in CHIME/FRB Collaboration (2021). Exposure time for AO is simply the length of time for each observation. These results are shown in Table 6.5 and Figure 6.5. As we would expect, the addition of significant AO exposure without a detection decreases the burst rates.

In CHIME/FRB Collaboration (2019), this scaled rate was determined by multiplying the observed rate by  $(S/S_0)^{1.5}$ , where  $S$  is our sensitivity limit in Jy and  $S_0 = 1$  Jy. This scaling, used in CHIME/FRB Collaboration (2019) achieves two purposes. First, scaling by  $S$  allows us to compare across the two telescopes. Second, the  $3/2$  power allows us to account for the way in which source counts scale with cosmological volume sampled assuming Euclidean spacetime. This scaling was used only for comparison purposes in CHIME/FRB Collaboration (2019), but we leverage it to allow us to combine data from both CHIME/FRB and AO. Instead of scaling the final rate, we instead scale the exposure time for each telescope, so that we can add the exposure times and calculate a final combined, scaled rate:

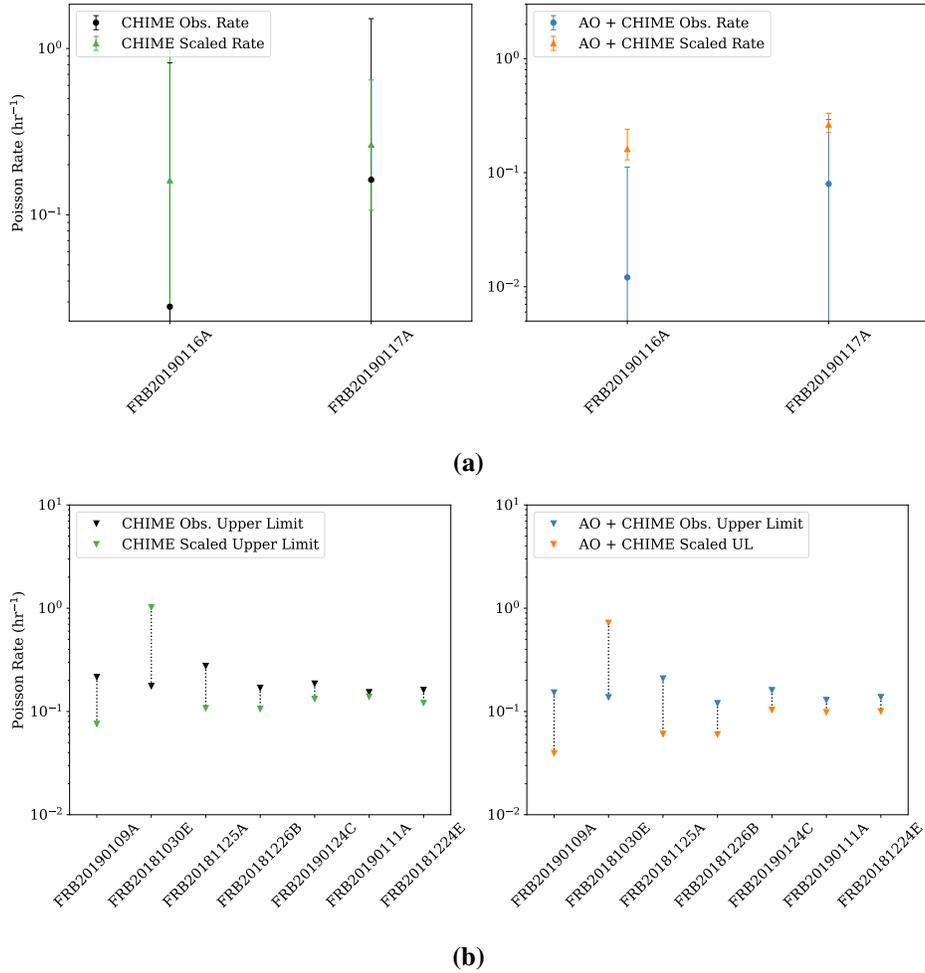
$$r_{\text{scaled}} = \frac{N_{\text{bursts}}}{t_{\text{CHIME}} (S_{\text{CHIME}}/S_0)^{-1.5} + t_{\text{AO}} (S_{\text{AO}}/S_0)^{-1.5}} \quad (6.5)$$

However, it has been well established that a pure Poisson distribution does not well describe FRB emission. One approach, used previously for FRB 20121102 analysis by Oppermann et al. (2018) and population level constraints by Caleb et al. (2019b) is to use the Weibull distribution in place of the Poisson distribution. The Weibull distribution is itself an extension of the Poisson distribution with the addition of a clustering parameter.

In this work, we follow closely the framework outlined by Oppermann et al. (2018), creating a Weibull distribution with parameters  $r$  (rate) and  $k$  (clustering),

**Table 6.5:** Poisson repetition rates for FRBs selected for followup with Arecibo. We report observed repeat rate, repeat rate scaled to account for sensitivity, and the observed CHIME/FRB and AO joint repeat rate, both observed and scaled. For bursts without known repetition, we present upper limits. Uncertainties are set by the uncertainty in CHIME/FRB exposure times, which is substantial.

Source	CHIME/FRB. Obs. (bursts/hr)	CHIME/FRB Scaled (bursts/hr)	CHIME/FRB + AO Obs. (bursts/hr)	CHIME/FRB + AO Scaled (bursts/hr)
<b>FRB 20190116A</b>	$0.16^{+0.13}_{-0.079}$	$0.16^{+0.95}_{-0.028}$	$0.050^{+0.15}_{-0.012}$	$0.043^{+0.121}_{-0.018}$
<b>FRB 20190117A</b>	$0.48^{+1.3}_{-0.32}$	$0.40^{+1.83}_{-0.16}$	$0.018^{+0.40}_{-8.8e-3}$	$0.077^{+0.144}_{-0.037}$
FRB 20190109A	< 0.22	< 0.076	< 0.15	< 0.039
FRB 20181030E	< 0.18	< 1.02	0.14	< 0.72
FRB 20181125A	< 0.28	< 0.11	< 0.14	< 0.061
FRB 20181226B	< 0.17	< 0.11	< 0.21	< 0.060
FRB 20190124C	< 0.19	< 0.13	< 0.12	< 0.10
FRB 20190111A	< 0.15	< 0.14	< 0.16	< 0.099
FRB 20181224E	< 0.16	< 0.12	< 0.13	< 0.010



**Figure 6.5:** Poisson repetition rates for known repeaters (panel a) and upper limits on repetition for non-repeating bursts (panel b). In both cases, the left-hand panel shows observed and sensitivity scaled CHIME/FRB only rates or limits, while the right-hand panel shows combined CHIME/FRB and AO rates and limits. Note that FRB 20181030E is much narrower than other bursts and therefore has a much higher sensitivity threshold. This causes the scaled rate to increase rather than decreasing

determining Bayesian probabilities for this data given those values and sampling over a grid of  $r$  and  $k$  values. The likelihood functions for Weibull-distributed repeater bursts are derived in detail in that work; we will briefly reproduce their results here. This method has been used extensively by the community since its introduction including in systematic efforts to understand FRB repetition such as Caleb et al. (2019a) and Connor & Petroff (2018).

If bursts from repeaters followed a Weibull distribution, the intervals between bursts  $\delta$  would be distributed as

$$W(\delta|k, r) = k\delta^{-1}[\delta r\Gamma(1 + 1/k)]^k e^{-[\delta r\Gamma(1 + 1/k)]^k}, \quad (6.6)$$

where  $r$  is the rate of bursts,  $\Gamma$  is  $\Gamma(x) = \int_0^\infty dt t^{x-1} e^{-t}$ , and  $k$  is a clustering parameter. In the limit that  $k$  equals one, the Weibull distribution reduces to the Poisson distribution. In cases where  $k < 1$ , the presence of one burst makes it more likely another burst will be emitted.

Under the assumption that gaps between observations are long relative to the length of observations, we can treat each individual observation as an independent event and therefore can determine our total probability by multiplying together different probabilities. Therefore, we focus on determining probabilities of potential observation outcome:  $N_{\text{bursts}} = 0$ ,  $N_{\text{bursts}} = 1$ , and  $N_{\text{bursts}} > 1$ .

The likelihood of  $N$  bursts with the time distribution  $t_1, \dots, t_N$  in an observation of length  $\Delta$  can be calculated by marginalizing over the time of the next burst ( $t_{N+1}$ ), resulting in a probability density  $p^3$

$$p(N, t_1 \dots t_N | k, r) = \int_{t_N}^{\infty} dt_{N+1} P(N | t_1, \dots, t_N, t_{N+1}) p(t_1, \dots, t_N, t_{N+1} | k, r). \quad (6.7)$$

The first term, the probability of observing  $N$  bursts is

$$P(N | t_1, \dots, t_N, t_{N+1}) = \theta(\Delta - t_N) \theta(t_{N+1} - \Delta), \quad (6.8)$$

where  $\theta(\cdot)$  is the Heaviside step function. This will be one only if exactly  $N$  bursts

---

<sup>3</sup>Throughout this section, we use  $p$  to represent probability density and  $P$  to represent probability.

occur in an observation of duration  $\Delta$ .

The second term, the probability density for arrival times of the subsequent bursts a Weibull distribution with certain parameters is

$$p(t_1, \dots, t_N, t_{N+1} | k, r) = p(t_1 | k, r) \prod_{i=1}^{N-1} W(t_{i+1} - t_i | k, r). \quad (6.9)$$

We must separately determine the probability of the first time interval, that between the observation start and the first burst. This probability density is proportional to the Weibull distribution and the length of the interval, thus

$$p(t_1 | k, r) = r \int_{t_1}^{\infty} d\delta_{01} W(\delta_{01} | k, r) = r \text{CCDF}(t_1 | k, r), \quad (6.10)$$

where  $\delta_{01}$  is the time interval between the last unobserved burst and the first observed burst and we define  $\text{CCDF} = \exp(-[\delta r \Gamma(1 + 1/k)]^k)$ .

All together, then, the probability density of a given burst distribution given a  $k$  and  $r$  is

$$p(N, t_1 \dots t_N | k, r) = r \text{CCDF}(t_1 | k, r) \text{CCDF}(\Delta - t_N | k, r) \prod_{i=1}^{N-1} W(t_{i+1} - t_i | k, r). \quad (6.11)$$

Equation 6.11 technically applies only to the case where two or more bursts are found within a single observation. In our dataset, this condition is met only by a single CHIME/FRB observation of FRB 20190116A – all other observations contain one or zero bursts.

Therefore, we predominantly employ the following special cases also derived by Oppermann et al. (2018)

$$P(N = 1 | k, r) = r \text{CCDF}(t_1 | k, r) \text{CCDF}(\Delta - t_N | k, r) \quad (6.12)$$

and

$$P(N = 0 | k, r) = \frac{\Gamma_i(1/k, (\Delta r \Gamma(1 + 1/k))^k)}{k \Gamma(1 + 1/k)}, \quad (6.13)$$

where  $\Gamma_i$  is the incomplete  $\Gamma$  function

$$\Gamma_i(x, z) = \int_z^{\infty} dt t^{x-1} e^{-t} \quad (6.14)$$

The ultimate quantity we are interested in is the likelihood of a given set of parameters; Bayes' Theorem dictates that

$$p(k, r | N, t_1 \dots t_N) \propto p(N, t_1 \dots t_N | k, r) p(k, r) \quad (6.15)$$

As in Oppermann et al. (2018), we employ a Jeffrey's prior, for which the likelihood is

$$p(k, r) \propto k^{-1} r^{-1}, \quad (6.16)$$

or

$$p(\log(k), \log(r)) \propto \text{constant}, \quad (6.17)$$

i.e. a uniform sampling in  $\log(k)$  and  $\log(r)$ . We therefore generate our values for  $r$  and  $k$  by creating a grid of values for  $\log(k)$  and  $\log(r)$ .

The implementation of Oppermann et al. (2018) was created with single dish observatories like AO and GBT in mind, and therefore it assumes observation durations and times of bursts within an observation are clearly defined. However, due to CHIME/FRB's complex beam structure, both quantities are somewhat ambiguous. The exposure time used in other calculations in this chapter and in previous CHIME/FRB publications is an attempt to quantify the amount of time a source is visible, but it does not account for the substantial daily time where sources are visible in the side-lobes. For the sake of sensible comparisons to Poisson rates calculated above, we take the approximations that observation durations are uniform across days and are the quotient of the total exposure time and the number of days deemed acceptable by the CHIME/FRB data quality monitoring team. Between August 30, 2018 and July 1, 2019, that is 237 days. The time an "observation" begins (i.e. the location of the burst becomes visible to CHIME/FRB) is frequency-dependent and complicated by side-lobe structure. Therefore, for simplicity's sake we take all bursts for FRB 20190117A and the first burst for FRB 20190116A to have occurred at the start of the observation. We maintain the spacing within the

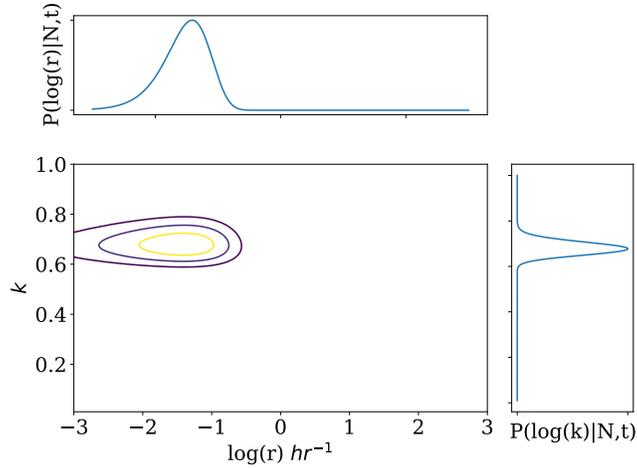
**Table 6.6:** Posterior means and for CHIME/FRB + AO repetition rates and clustering coefficients, based on a Weibull distribution of burst rates. The Weibull distribution reduces to the Poisson distribution when  $k = 1$ ; the low  $k$  values listed here suggest a high degree of clustering in bursts from these sources. Lower  $k$  values also decrease the predictive power of  $r$  estimates. Note that  $r$  is in units of  $\text{hr}^{-1}$ , whereas it is presented in units of  $\text{days}^{-1}$  in Oppermann et al. (2018). FRB 20190116A and FRB 20190117A (bolded) are known repeaters

Source	$\log(k)$	$\log(r)(\text{hr}^{-1})$
<b>FRB 20190116A</b>	$-0.168^{+0.014}_{-0.016}$	$-1.41^{+0.27}_{-0.34}$
<b>FRB 20190117A</b>	$-0.166^{+0.016}_{-0.016}$	$-1.04^{+0.156}_{-0.168}$
FRB 20190109A	$-0.164^{+0.016}_{-0.016}$	$-1.76^{+0.34}_{-0.45}$
FRB 20181030E	$-0.164^{+0.016}_{-0.016}$	$-0.520^{+0.336}_{-0.45}$
FRB 20181125A	$-0.164^{+0.016}_{-0.016}$	$-1.59^{+0.34}_{-0.45}$
FRB 20181226B	$-0.164^{+0.016}_{-0.016}$	$-1.59^{+0.34}_{-0.45}$
FRB 20190124C	$-0.164^{+0.016}_{-0.016}$	$-1.37^{+0.34}_{-0.44}$
FRB 20190111A	$-0.164^{+0.016}_{-0.016}$	$-1.39^{+0.34}_{-0.45}$
FRB 20181224E	$-0.164^{+0.016}_{-0.016}$	$-1.38^{+0.33}_{-0.45}$

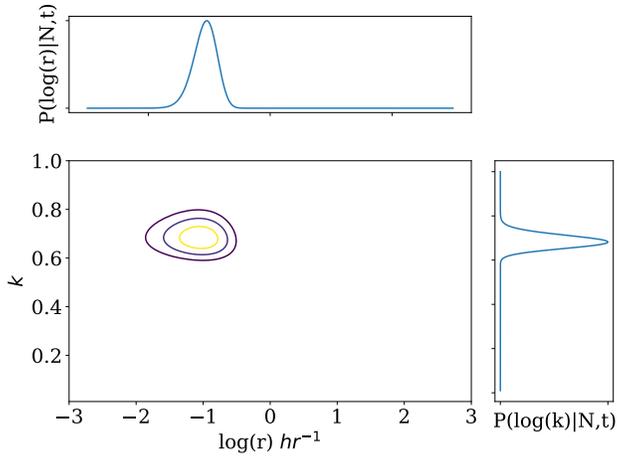
observation between the first and second FRB 20190116A bursts.

In Table 6.6, we present the posterior mean values for  $r$  and for  $k$  as calculated using this method. In Figure 6.6, we present a contour plot with the 68%, 95%, and 99% contours for  $k$  and  $r$ , for the two known repeating sources.

Based on the rate analyses, we can say that the repetition rates are generally low (between  $10^{-1}$  to  $10^{-2} \text{ hr}^{-1}$ ). We can also say that there is at least a moderate clustering effect. It is immediately obvious that the values determined by the Weibull analysis for  $k$  and  $r$  are nearly identical for all sources. This is likely a



(a)



(b)

**Figure 6.6:** Weibull distribution posterior probability distribution for clustering coefficient  $k$  and burst rate  $r$  for FRB 20190116A (Panel a) and FRB 20190117A (Panel b). Contours are at 68%, 95%, and 99% probability. The insets are marginalized posteriors for  $\log(r)$  and  $\log(k)$ . The Oppermann et al. (2018) implementation of the Weibull distribution finds a moderate amount of clustering and burst rates similar to the Poissonian confidence interval burst rates discussed earlier in the chapter. The preponderance of non-detections means that there is little distinction between maximum likelihood  $k$  and  $r$  values for the two sources.

consequence of the number of CHIME non-detections. The dominance of the non-detections is also suggested by the fact that the apparently non-repeating bursts have the same clustering values as the known repeating bursts, suggesting there may be little physical motivation for the clustering parameter chosen by the maximum likelihood analysis.

Each source’s list of observation includes between 231 and 236 non-detection observations made with CHIME/FRB. These CHIME/FRB non-detection are the dominant case (as opposed to CHIME/FRB detections or AO non-detections). Though each source has different total exposure times, once the exposure time is divided over the 237 days of observations, the differences become very small. Therefore, mathematically, we are looking at nearly identical datasets, particularly for the seven sources with only one detection.

We might hope to see a higher degree of clustering from this analysis, suggesting that we are simply missing clusters of observations by observing at the wrong times. However, only for FRB 20190116A are there any multi-burst observations. The Weibull implementation used here does not account for the time between observations, so only multi-burst observations appear as clusters. This means CHIME/FRB, with its short exposure times for these sources is unlikely to detect a cluster of detections; the FRB 20190116A pair of detections is an unusual case. In principle, the presence of this cluster should distinguish the value of  $k$  for FRB 20190116A. However, there is only one such cluster in over 250 observations, so each individual observation is only marginally impactful on the total parameters.

The excessive imbalance between detections and non-detections and the failure to account for gaps between observations emphasizes that the Oppermann et al. (2018) implementation of the Weibull distribution is likely insufficient for analyses that incorporate CHIME/FRB data or data from other all-sky FRB monitors. Though this structure has been commonly used in discussing FRB repetition, the time has come for a more sophisticated model. Developing that model is beyond the scope of this work, but it is needed.

The discovery of periodic activity from FRB 20121102 and FRB 20180916B suggests that a Weibull distribution may be an inadequate tool for understanding repetition, as it suggests that a simple repetition rate may be the wrong quantity to consider in discussing frequency of repetition. Our knowledge of repeater period-

icity would motivate a further extension of this method to account for the spacing between observations. However, we do not attempt such an extension in this work.

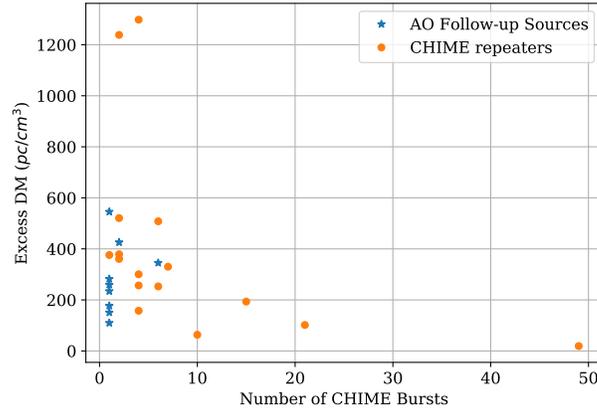
While we acknowledge the insufficiency of our tools in characterizing the behaviour of these sources, we also acknowledge that in the absence of additional detections, our options are limited. We cannot determine a periodicity with one epoch of detection, or even reliably with six separate detections, as in the case of FRB 20190117A. We cannot examine in high resolution bursts we do not detect. Unfortunately, with the recent demise of AO, we will not be able to continue this project further and are left with only rough constraints. We may be able to conduct a similar project in the future using FAST (the Five Hundred Metre Aperture Synthesis Telescope).

### **6.4.3 Comparison with other repeaters**

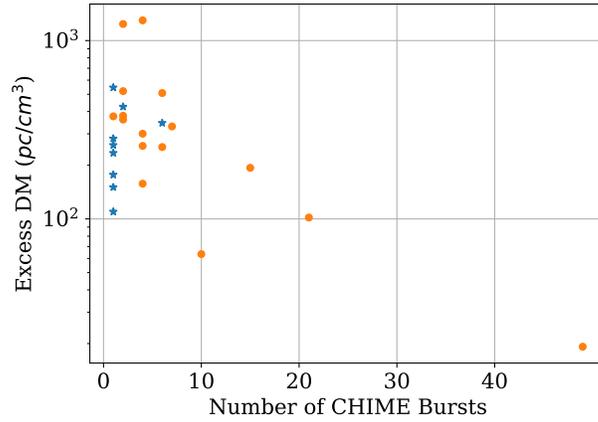
When noting the very low repetition rates suggested by both Poisson and Weibull analyses, the reader might be tempted to conclude that there was something unusual about these sources - perhaps they were particularly close to the fluence limits of our telescope or perhaps they are particularly distant. Of course, exposure time is a concern for a CHIME-only analysis, but the inclusion of AO data in this analysis mitigates that factor, bringing the amount of exposure for low declination bursts in line with higher declination CHIME bursts. In this section, we present a brief analysis of these sources in the context of published repeating FRBs. For brevity and for maximum relevance, we compare only CHIME detections and do not include detections from other telescopes such as ASKAP or GBT.

First, in Figure 6.7, we plot excess DM vs. number of detected repetitions. As is common in pulsar and FRB astronomy, we are using DM as a rough proxy for distance. Excess DM is calculated by subtracting the measured galactic DM from the total measured DM of the FRB. Here, we use the mean of the Yao et al. (2017) and Cordes & Lazio (2002) DM map values for the Galactic DM to these positions. It is important to note that this simple analysis does not account for contributions from the host DM.

In Figure 6.7a, we plot this comparison on linear axes. It is immediately apparent that the burst with the highest number of measured repetitions is also the clos-

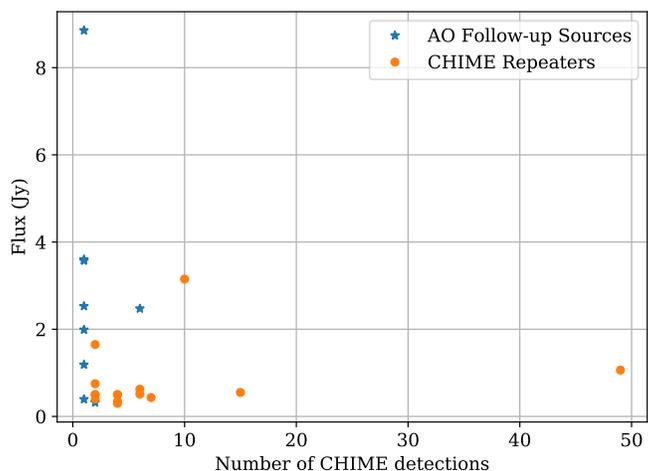


(a)



(b)

**Figure 6.7:** Excess DM ( $\text{pc cm}^{-3}$ ) plotted against number of bursts detected with CHIME/FRB for repeaters published in CHIME/FRB Collaboration (2019) and Fonseca et al. (2020), as well as FRB 20121102 and AO follow-up sources. AO follow-up sources are represented by blue stars; other repeaters are represented by orange circles. There is not a significant correlation between the number of bursts detected by CHIME/FRB and excess DM. Additionally, our AO follow-up sources do not appear to be more distant than observed CHIME/FRB repeaters.



**Figure 6.8:** Average CHIME/FRB peak flux measurement for known repeaters and for AO follow-up sources, plotted against number of bursts.

est, FRB 20180916B. However, the relationship between number of repetitions and excess DM is more ambiguous for other bursts. It is certainly the case that our repeaters of interest and our non-repeater, morphologically interesting bursts are not anomalously distant. They appear alongside the bulk of our published repeaters, suggesting they are not exceptional. In Figure 6.7b, we also plot excess DM on a log scale, again demonstrating no clear correlation between excess DM and number of repetitions (except for the extreme case of FRB 20180916B).

The second question we might ask is are these bursts anomalously dim and therefore unlikely to be observed again, i.e. could these be the very brightest bursts from this source and thus we are failing to see any further bursts? In Figure 6.8, we plot the number of peak flux for CHIME/FRB repeaters and target sources against the number of detections. Again, these bursts do not appear to be anomalously low in peak flux, so we do not have reason to consider them exceptional.

Overall, we assert that there is nothing exceptional about these bursts – they are not particularly distant nor are they particularly dim. This is important as it means we cannot write off the non-detection of further (or initial) repetition as being simply a function of the sources’ own peculiarities. We must confront it

directly.

## 6.5 Discussion

This work could be termed a “successful failure.” It was conceived with the idea that we could detect further bursts from repeater or repeater-like sources with higher resolution and potentially higher sensitivity. After up to a year of searching, we found no such bursts. But, the failure to detect such bursts is instructive and provides an important insight for the study of repeating FRBs broadly: low-burst rate repeating FRBs are a real component of the population, and increased exposure time does not ensure the discovery of further repeat bursts.

One potential stumbling point within this project was observing cadence. Our observing cadence was set essentially at random, based on the idea that we could not predict when FRBs would be active. However, our increasing understanding that FRBs go through active and inactive periods may lead to changes in how we conduct such follow-up observations. Ideally, we would detect this periodicity with a more passive monitor like CHIME/FRB and then be able to apply triggered observations as in Chawla et al. (2020). But even in the absence of such a program, we may be able to create more efficient programs. It may be better to propose focused blocks of observation, e.g. 1.5 hours per day every day for a week, repeated several times vs. the roughly weekly cadence of these proposals.

Another is our limited knowledge of FRB properties below 400 MHz. Our first sub-400 MHz FRB discoveries were made in the past year, and we still have only a handful. (See Parent et al. (2020); Chawla et al. (2020); Pilia et al. (2020); Pleunis et al. (2021b)). Of this handful, most are repeat detections of FRB 20180916B. Though obviously, FRBs can be seen at low frequency, it remains to be seen whether higher frequency knowledge is entirely applicable (particularly outside of FRB 20180916B).

A third potential challenge is the increasingly obvious tendency for repeating FRBs to be narrow-band. If narrow-band burst occurs outside of our bandwidth, we will never observe it, not matter how sensitive our telescope or how frequently we observe. As we begin to better understand the bandwidth distribution of repeating FRBs, we will be able to better plan for this challenge, though we may not be able

to mitigate it.

Though there may be a myriad of smaller lessons to be learned from these results, three are particularly cogent. First, not every search is successful even in our new statistical FRB science era. Second, while morphological studies of the first CHIME/FRB catalog have been highly suggestive, we should be cautious of over-interpreting about a year of data. Third, highly clustered and/or very low burst rate repeating FRBs are real, and something that should be considered in our population-level understanding of repetition. Not all repeating FRBs behave like FRB 20121102 and FRB 20180916B, and we should be cautious about building our understanding of the boundaries of repeating FRBs from a very small subset of sources.

Among published repeating FRBs, most have only 2-3 detected bursts. A small handful, most notably FRB 20121102 and FRB 20180916B, have large numbers of bursts. For some of these sources, we have been able to measure definite active and quiescent periods, which in turn make it easier to continue making repeat detections. Understandably, these easily detectable FRBs are the most studied sources. Yet, we must stop and ask ourselves what is a typical repeating FRB? And what are we treating as a typical repeating FRB? It seems entirely possible that the bulk of the population is highly clustered and/or has a low-repetition rate. If that's the case, we should be careful not to over-extrapolate from FRB 20121102 and FRB 20180916B data. These results are consistent with James et al. (2020), who found for a set of 27 bursts that no more than 60% could be repeaters consistent with a repetition rate of FRB 20121102. (However, these results, the results of James et al. (2020), and CHIME/FRB Collaboration (2021) would easily be consistent with an average repetition rate of, e.g. one burst per year.)

The phenomenal success of CHIME/FRB and of other dedicated FRB surveys such as ASKAP CRAFT have radically changed the face of FRB astronomy. Not long ago, a search that did not result in a burst would be commonplace; now it feels like an exception. This work provides a counter to the assumption that CHIME/FRB's sole impediment in detecting lower declination repeaters is exposure time and decreased sensitivity at lower declinations. Instrumental constraints are surely an impediment but not all-determining; adding these AO observation nearly triples the exposure time for FRB 20190116A but adds no bursts. If the

challenge is not purely instrumental, follow-up observations are hardly a guaranteed success. Targeted observations of a known FRB are also unable take advantage of the observation techniques that have led to our statistical era of FRBs, such as daily all-sky coverage. We are again dependent on luck of the observing schedule. Certainly, follow-up observations have a key place in FRB astronomy, but they are not certainties.

The continued non-detection of repetition from morphologically interesting sources is itself very intriguing, but hard to quantify. This may be evidence that all FRBs repeat, via a linkage between repeater like properties and bursts which appear to be one-off, it may be an indication that we have caught only the tail of brightest bursts (and gotten unlucky with Arecibo Observations), or it may mean nothing in particular. One reason why we only observed morphologically interesting one-off burst for a single semester is that quantifying and interpreting non-detections from FRBs that are not known to repeat is difficult. We can say that they do not repeat at a rate higher than some value, but that is a weak constraint in light of our general lack of information.

With the release of the CHIME/FRB catalog and accompanying meta-analyses, we are at a point where people are inclined to start drawing conclusions about the nature of the FRB population, including the properties of the repeater population. However, this work is an important reminder that we still have much to learn.

## Chapter 7

# Conclusion

*Science is a quest for Understanding.* — Jocelyn Bell Burnell

This thesis is ultimately a testament to the success of CHIME. We have achieved every application we dreamed of for CHIME/FRB and CHIME/Pulsar and almost all of them are presented here.

First, we are able to time pulsars with CHIME/Pulsar. This seemingly simple project is the culmination of more than five years of intense work, from developing the beamformer and calibration methods to installing the system to understanding its many quirks and finally to processing the data and timing pulsars. This work is truly the tip of the CHIME/Pulsar timing iceberg.

If Chapter 4 is the tip of the CHIME/Pulsar iceberg in the effort invested, it is also the tip of the iceberg in the science products we can expect from CHIME/Pulsar. In the immediate future, I will be continuing to refine and expand our CHIME/Pulsar and NANOGrav joint dataset. This will include moving from our set of eight test pulsars towards the full CHIME/Pulsar and NANOGrav combined dataset, encompassing all NANOGrav pulsars at declinations greater than  $-15^\circ$ . I will also be developing an automated pipeline and integrating it with `nanopipe` to efficiently manage this much larger dataset and to ensure reproducible results. I will also implement, first for test pulsars and later in pipeline format, wideband timing for CHIME/Pulsar. Our ultimate goal is that by the dataset subsequent to the 15 year dataset, CHIME/Pulsar will be able to be treated as simply one of NANOGrav's several observatories.

The potential applications of the CHIME/Pulsar-NANOGrav dataset are exciting. Not only will this provide a boost to efforts to detect low-frequency gravitational waves, it will also help us better understand the effects of the interstellar medium on pulsar timing and for its own sake. We can already see more detail in our DMX models based on even a cursory effort to build such models for CHIME/Pulsar data. In the future, we will determine the optimal DMX binning strategy for CHIME/Pulsar and begin to fully utilize its potential. Understanding dispersion measure variations caused by the ISM, especially frequency dependent dispersion measure variations, is an important part of understanding our dataset (Jones et al., 2017; Cordes et al., 2016). In addition to direct application to NANOGrav, we can expect CHIME/Pulsar to enable more in-depth pulsar studies, such as studies of rotation measure (Ng et al., 2020) and pulsar mass measurements such as Fonseca et al. (2021).

Additionally, recent events have demonstrated the usefulness of daily cadence observations for studying even pulsars we believe to be stable and well-understood. In mid-April 2021, a joint European PTA and FAST team reported a change in the pulse profile for PSR J1713+0747 (Xu et al., 2021). CHIME/Pulsar was able to almost immediately confirm the presence of this profile change and to further constrain the time range in which the change occurred (Meyers & CHIME/Pulsar Collaboration, 2021). We did not need to take any additional observations to make this confirmation; we simply had to process data we had already collected. The ability to conduct near-real-time monitoring of pulsars for profile changes or other interesting behaviour is another motivation for improving and increasingly automating our data analysis pipelines.

Chapter 5 demonstrated the exciting potential for transient monitoring with CHIME/Pulsar. This is one of the first cases where pulsars were discovered accidentally when searching for FRBs, instead of FRBs being discovered when searching for pulsars. The ability to search the entire sky every day for single pulses has the potential to greatly increase our catalog of RRATs, intermittent pulsars, and pulsars with high nulling fraction. This is particularly true following the work of Dong (2021), which allows us to scrape the FRB database and find all Galactic transients seen by CHIME/FRB.

The project also highlights the usefulness of daily or near daily monitoring with

CHIME/Pulsar for studying RRATs. We have long expected CHIME/Pulsar to be a useful instrument for studying RRATs, and Chapter 5 functions as an interesting pilot project for a large-scale RRAT survey with CHIME/Pulsar. Such a survey is beginning now. One of the major products of such a survey will be the ability to have a large set of RRATs detected with the same telescope selection effects. This will greatly simplify population-level studies of RRAT properties and may allow us to better understand the prevalence of RRATs in the the Milky Way.

Discovering RRATs with CHIME/FRB also highlights the similarities between FRBs and less luminous radio transients. It is possible, though by no means certain, that there exists some sort of continuum of repeating phenomenon from neutron stars, with MSPs and ordinary pulsars at one end, intermittent pulsars and RRATs in the middle, and repeating FRBs at the less repetitive end.

No discussion of the contributions of CHIME would be complete without reference to the phenomenal success of CHIME/FRB. In July 2018, when operations began, we knew of approximately 100 FRBs total and only one repeater. Now, the first CHIME/FRB catalog includes over 500 FRBs and there are more than 20 published repeaters (CHIME/FRB Collaboration, 2021). Unpublished data contains many more FRBs and repeaters. The past four years have been an extremely exciting (albeit, often hectic) time to be a part of the FRB field.

FRB repetition, as investigated in Chapter 6 is one of the most interesting questions in FRB astronomy at present. Though the association of an FRB with SGR 1935+2154 provides a strong argument in favour of magnetars as the source of FRBs, the possibility of multiple populations has not been definitively ruled out. The implications of the observed difference in burst morphology and bandwidth between repeaters and non-repeaters in the CHIME/FRB catalog has not yet been fully explored, and this will be an interesting topic in the coming years.

The inconclusiveness of the Weibull rate calculation for the FRBs examined in Chapter 6 is also an important opening for future study. While the Weibull rate is a good first step in developing a model for determining FRB repetition rates, it does reflect an expectation of a relatively small number of relatively long observations, with most observations including a detection. If all FRBs are repeating, but many have very low repetition rates, we may need another framework to allow us to differentiate between low-repeat rate bursts.

The results of Chapter 6 also provide an important opportunity for reflection. FRB science has had immense success in the past five years. It is still important, however, to design projects that push the boundaries of our understanding and may fail. Additionally, although we understand FRB 20121102 and FRB 20180916B in depth, this project is an important reminder that we do not understand the repeating FRB population well enough to know if those sources are in fact representative.

Since the start of commissioning operations for CHIME/Pulsar and CHIME/FRB in 2018, these instruments have achieved phenomenal success, and it has been an honour to be a part of the CHIME team. Both programs will continue to grow and develop over the coming years, changing the landscape of pulsar and FRB astronomy as they do.

# Bibliography

- Aggarwal, K., Law, C. J., Burke-Spolaor, S., et al. 2020, *Research Notes of the American Astronomical Society*, 4, 94, doi: 10.3847/2515-5172/ab9f33 → pages 15
- Aitkin, A. C. 1935, *Proceedings of the Royal Society of Edinburgh*, 55, 42 → pages 74
- Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021a, *ApJS*, 252, 4, doi: 10.3847/1538-4365/abc6a0 → pages 72, 73, 74, 75, 76, 87, 104
- . 2021b, *ApJS*, 252, 5, doi: 10.3847/1538-4365/abc6a1 → pages 107
- Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *Nature*, 300, 728, doi: 10.1038/300728a0 → pages 6
- Amiri, M., Bandura, K. M., Boyle, P. J., et al. 2021, *ApJS*, 255, 5, doi: 10.3847/1538-4365/abfdcb → pages 61, 65, 113
- Archibald, A. M., Gusinskaia, N. V., Hessels, J. W. T., et al. 2018, *Nature*, 559, 73, doi: 10.1038/s41586-018-0265-1 → pages 19
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, *ApJ*, 821, 13, doi: 10.3847/0004-637X/821/1/13 → pages 20
- Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018a, *ApJ*, 859, 47, doi: 10.3847/1538-4357/aabd3b → pages 20
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018b, *ApJS*, 235, 37, doi: 10.3847/1538-4365/aab5b0 → pages 72, 73
- Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, 905, L34, doi: 10.3847/2041-8213/abd401 → pages vi, xvi, 21, 22, 108

- Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2021, *ApJ*, 914, 121,  
doi: 10.3847/1538-4357/abfcd3 → pages 20
- Baade, W., & Zwicky, F. 1934, *Proceedings of the National Academy of Sciences*,  
20, 254 → pages 2
- Backer, D. C. 1970, *Nature*, 228, 42, doi: 10.1038/228042a0 → pages 9, 109
- Backer, D. C., & Hellings, R. W. 1986, *ARA&A*, 24, 537,  
doi: 10.1146/annurev.aa.24.090186.002541 → pages 41, 43
- Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982,  
*Nature*, 300, 615, doi: 10.1038/300615a0 → pages 6, 100
- Bailes, M., Harrison, P. A., Lorimer, D. R., et al. 1994, *ApJL*, 425, L41,  
doi: 10.1086/187306 → pages 100
- Bandura, K., Cliche, J. F., Dobbs, M. A., et al. 2016a, *Journal of Astronomical  
Instrumentation*, 5, 1641004 → pages 59
- Bandura, K., Bender, A. N., Cliche, J. F., et al. 2016b, *Journal of Astronomical  
Instrumentation*, 5, 1641005, doi: 10.1142/S2251171716410051 → pages 59
- Barr, E. D., Guillemot, L., Champion, D. J., et al. 2013, *MNRAS*, 429, 1633,  
doi: 10.1093/mnras/sts449 → pages 32
- Bassa, C. G., Tendulkar, S. P., Adams, E. A. K., et al. 2017, *ApJL*, 843, L8,  
doi: 10.3847/2041-8213/aa7a0c → pages 10
- Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, *MNRAS*, 439, 2893,  
doi: 10.1093/mnras/stu157 → pages 136
- Beloborodov, A. M. 2017, *ApJL*, 843, L26, doi: 10.3847/2041-8213/aa78f3 →  
pages 13
- Bhandari, S., Sadler, E. M., Prochaska, J. X., et al. 2020, *ApJL*, 895, L37,  
doi: 10.3847/2041-8213/ab672e → pages 11, 13
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics: Second Edition* → pages 44
- Bochenek, C. D., McKenna, D. L., Belov, K. V., et al. 2020a, *PASP*, 132, 034202,  
doi: 10.1088/1538-3873/ab63b3 → pages 52
- Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020b, *Nature*, 587, 59,  
doi: 10.1038/s41586-020-2872-x → pages 10, 13

- Booth, R. S., & Jonas, J. L. 2012, *African Skies*, 16, 101 → pages 52
- Burgay, M., D’Amico, N., Possenti, A., et al. 2003, *Nature*, 426, 531,  
doi: 10.1038/nature02124 → pages 19
- Burke-Spolaor, S. 2013, in *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, ed. J. van Leeuwen, Vol. 291, 95–100,  
doi: 10.1017/S1743921312023277 → pages 110
- Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, *A&A Rv*, 27, 5,  
doi: 10.1007/s00159-019-0115-7 → pages 20
- Caleb, M., Stappers, B., Rajwade, K., & Flynn, C. 2019a, *MNRAS*, 484, 5500 →  
pages 139, 156
- Caleb, M., Stappers, B. W., Rajwade, K., & Flynn, C. 2019b, *MNRAS*, 484, 5500,  
doi: 10.1093/mnras/stz386 → pages 153
- Caleb, M., Flynn, C., Bailes, M., et al. 2016, *MNRAS*, 458, 718,  
doi: 10.1093/mnras/stw109 → pages 26
- Caleb, M., Stappers, B. W., Abbott, T. D., et al. 2020, *MNRAS*, 496, 4565,  
doi: 10.1093/mnras/staa1791 → pages 15
- Champion, D. J., Petroff, E., Kramer, M., et al. 2016, *MNRAS*, 460, L30,  
doi: 10.1093/mnrasl/slw069 → pages 111
- Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, *Nature*, 466, 463,  
doi: 10.1038/nature09187 → pages 49
- Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, *Nature*,  
doi: 10.1038/nature20797 → pages 10
- Chawla, P., Andersen, B. C., Bhardwaj, M., et al. 2020, *ApJL*, 896, L41,  
doi: 10.3847/2041-8213/ab96bf → pages 15, 146, 165
- CHIME/FRB Collaboration. 2018, *ApJ*, 863, 48, doi: 10.3847/1538-4357/aad188  
→ pages 26, 58, 59, 61, 62, 63, 113, 140
- . 2019, *ApJL*, 885, L24, doi: 10.3847/2041-8213/ab4a80 → pages xiii, xxv,  
15, 17, 112, 127, 139, 147, 153, 163
- . 2020a, *Nature*, 587, 54, doi: 10.1038/s41586-020-2863-y → pages vi, xvi,  
10, 13, 14, 64

- . 2020b, *Nature*, 582, 351, doi: 10.1038/s41586-020-2398-2 → pages 11, 15, 139, 141
- . 2021, Submitted to *ApJS* → pages vii, xxiv, 11, 15, 60, 64, 136, 138, 142, 145, 153, 166, 170
- CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019a, *Nature*, 566, 235, doi: 10.1038/s41586-018-0864-x → pages 15, 64, 139
- . 2019b, *Nature*, 566, 230, doi: 10.1038/s41586-018-0867-7 → pages 64
- Coles, W., Hobbs, G., Champion, D. J., Manchester, R. N., & Verbiest, J. P. W. 2011, *MNRAS*, 418, 561, doi: 10.1111/j.1365-2966.2011.19505.x → pages 48, 74
- Connor, L., Miller, M. C., & Gardenier, D. W. 2020, *MNRAS*, 497, 3076, doi: 10.1093/mnras/staa2074 → pages 18, 139
- Connor, L., Pen, U.-L., & Oppermann, N. 2016, *MNRAS*, 458, L89, doi: 10.1093/mnras/slw026 → pages 17
- Connor, L., & Petroff, E. 2018, *ApJL*, 861, L1, doi: 10.3847/2041-8213/aacd02 → pages 140, 156
- Cordes, J. M., & Chatterjee, S. 2019, *ARA&A*, 57, 417, doi: 10.1146/annurev-astro-091918-104501 → pages 13
- Cordes, J. M., & Lazio, T. J. W. 2002, arXiv e-prints, astro. <https://arxiv.org/abs/astro-ph/0207156> → pages xiii, 10, 27, 63, 125, 128, 162
- Cordes, J. M., & McLaughlin, M. A. 2003, *ApJ*, 596, 1142, doi: 10.1086/378231 → pages xiv, 150, 151, 152
- Cordes, J. M., Shannon, R. M., & Stinebring, D. R. 2016, *ApJ*, 817, 16, doi: 10.3847/0004-637X/817/1/16 → pages 108, 169
- Cordes, J. M., Wasserman, I., Hessels, J. W. T., et al. 2017, *ApJ*, 842, 35 → pages 139
- Cordes, J. M., Freire, P., Lorimer, D. R., et al. 2006, *ApJ*, 637, 446 → pages 110
- Craft, H. D., & Comella, J. M. 1968, *Nature*, 220, 676, doi: 10.1038/220676a0 → pages 76
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *Nature Astronomy*, 4, 72, doi: 10.1038/s41550-019-0880-2 → pages 3

- Cruces, M., Spitler, L. G., Scholz, P., et al. 2021, MNRAS, 500, 448,  
doi: 10.1093/mnras/staa3223 → pages 11, 139, 141
- Cui, B. Y., Boyles, J., McLaughlin, M. A., & Palliyaguru, N. 2017, ApJ, 840, 5,  
doi: 10.3847/1538-4357/aa6aa9 → pages 9
- Damour, T., & Deruelle, N. 1986, Ann. Inst. Henri Poincaré Phys. Théor, 44, 263  
→ pages 47, 74
- Damour, T., & Schaefer, G. 1991, PhRvL, 66, 2549,  
doi: 10.1103/PhysRevLett.66.2549 → pages 19
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, MNRAS, 388,  
247, doi: 10.1111/j.1365-2966.2008.13376.x → pages 121
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, PASP, 129, 045001,  
doi: 10.1088/1538-3873/129/974/045001 → pages 52
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, ApJ, 762, 94,  
doi: 10.1088/0004-637X/762/2/94 → pages 20, 95
- Deneva, J. S., Stovall, K., McLaughlin, M. A., et al. 2013, ApJ, 775, 51,  
doi: 10.1088/0004-637X/775/1/51 → pages 32
- Deng, M., Campbell-Wilson, D., & for the CHIME Collaboration. 2017, ArXiv  
e-prints. <https://arxiv.org/abs/1708.08521> → pages 57
- Denman, N., Renard, A., Vanderlinde, K., et al. 2020, Journal of Astronomical  
Instrumentation, 9, 2050014, doi: 10.1142/S2251171720500142 → pages 59
- Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, MNRAS, 458, 3341,  
doi: 10.1093/mnras/stw483 → pages 21
- Detweiler, S. 1979, ApJ, 234, 1100, doi: 10.1086/157593 → pages 20
- Di, L., Zhang, X., Qian, L., et al. 2019, The Astronomer's Telegram, 13064, 1 →  
pages 15
- Dong, F. 2021, Finding new pulsars using CHIME/FRB single pulse events →  
pages 136, 169
- Duncan, R. C., & Thompson, C. 1992, ApJL, 392, L9, doi: 10.1086/186413 →  
pages 7

- Duncan, R. C., & Thompson, C. 1996, in American Institute of Physics Conference Series, Vol. 366, High Velocity Neutron Stars, ed. R. E. Rothschild & R. E. Lingenfelter, 111–117, doi: 10.1063/1.50235 → pages 7
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549, doi: 10.1111/j.1365-2966.2006.10870.x → pages xxviii, 40, 72
- Ellis, J. A., Siemens, X., & Creighton, J. D. E. 2012, ApJ, 756, 175, doi: 10.1088/0004-637X/756/2/175 → pages 20
- Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite. <http://ascl.net/1912.015> → pages xxvii, 76
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, ApJL, 875, L2, doi: 10.3847/2041-8213/ab0c96 → pages 52
- Everett, J. E., & Weisberg, J. M. 2001, ApJ, 553, 341, doi: 10.1086/320652 → pages 4
- Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, ApJL, 891, L6, doi: 10.3847/2041-8213/ab7208 → pages xiii, xxv, 112, 127, 139, 147, 163
- Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, ApJL, 915, L12, doi: 10.3847/2041-8213/ac03b8 → pages xii, 3, 65, 75, 83, 86, 87, 107, 169
- Foster, R. S., & Backer, D. C. 1990, ApJ, 361, 300, doi: 10.1086/169195 → pages 20
- Foster, R. S., Wolszczan, A., & Camilo, F. 1993, ApJL, 410, L91, doi: 10.1086/186887 → pages 90
- Freire, P. C. C., & Wex, N. 2010, MNRAS, 409, 199, doi: 10.1111/j.1365-2966.2010.17319.x → pages 74
- Gajjar, V., Siemion, A. P. V., Price, D. C., et al. 2018, ApJ, 863, 2, doi: 10.3847/1538-4357/aad005 → pages 14
- Gold, T. 1968, Nature, 218, 731, doi: 10.1038/218731a0 → pages 2
- Good, D. C., Andersen, B. C., Chawla, P., et al. 2021, Accepted by ApJ, arXiv:2012.02320. <https://arxiv.org/abs/2012.02320> → pages 23, 65, 135, 136
- Gourdji, K., Michilli, D., Spitler, L. G., et al. 2019, ApJL, 877, L19, doi: 10.3847/2041-8213/ab1f8a → pages 15

- Gregory, P. C., & Fahlman, G. G. 1980, *Nature*, 287, 805, doi: 10.1038/287805a0  
→ pages 7
- Hallinan, G., Ravi, V., Weinreb, S., et al. 2019, in *Bulletin of the American Astronomical Society*, Vol. 51, 255. <https://arxiv.org/abs/1907.07648> → pages 52
- Han, J. L., Wang, C., Wang, P. F., et al. 2021, *Research in Astronomy and Astrophysics*, 21, 107, doi: 10.1088/1674-4527/21/5/107 → pages 33
- Hankins, T. H., & Rickett, B. J. 1975, in *Methods in Computational Physics: advances in research and applications*, Vol. 14 (Elsevier), 55–129 → pages 28
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, 47, 1 → pages 121
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, 265, L39, doi: 10.1086/183954 → pages 20
- Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, *ApJL*, 876, L23, doi: 10.3847/2041-8213/ab13ae → pages 11, 15, 139
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709, doi: 10.1038/217709a0 → pages 2
- Hobbs, G., Jenet, F., Lee, K. J., et al. 2009, *MNRAS*, 394, 1945, doi: 10.1111/j.1365-2966.2009.14391.x → pages xxviii, 40, 72
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, 369, 655, doi: 10.1111/j.1365-2966.2006.10302.x → pages xxviii, 40, 72, 122
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, 21, 302, doi: 10.1071/AS04022 → pages 115
- Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, 195, L51 → pages 19, 20
- Ilyasov, Y. P., Imae, M., Hanado, Y., et al. 2005, *Astronomy Letters*, 31, 30, doi: 10.1134/1.1854794 → pages 100
- Jackson, J. D. 1962, *Classical Electrodynamics*, 808 → pages 5
- James, C. W., Bannister, K. W., Macquart, J. P., et al. 2019, *PASA*, 36, e009, doi: 10.1017/pasa.2019.1 → pages 26, 134
- James, C. W., Osłowski, S., Flynn, C., et al. 2020, *MNRAS*, 495, 2416, doi: 10.1093/mnras/staa1361 → pages 15, 166

- Jones, M. L., McLaughlin, M. A., Lam, M. T., et al. 2017, *ApJ*, 841, 125, doi: 10.3847/1538-4357/aa73df → pages 75, 100, 169
- Josephy, A., Chawla, P., Fonseca, E., et al. 2019, *ApJL*, 882, L18, doi: 10.3847/2041-8213/ab2c00 → pages 15
- Joshi, B. C., Arumugasamy, P., Bagchi, M., et al. 2018, *Journal of Astrophysics and Astronomy*, 39, 51, doi: 10.1007/s12036-018-9549-y → pages 21
- Karako-Argaman, C., Kaspi, V. M., Lynch, R. S., et al. 2015, *ApJ*, 809, 67, doi: 10.1088/0004-637X/809/1/67 → pages 9
- Kaspi, V. M., & Beloborodov, A. M. 2017, *ARA&A*, 55, 261, doi: 10.1146/annurev-astro-081915-023329 → pages 7
- Katz, J. I. 2017, *MNRAS*, 467, L96, doi: 10.1093/mnras/slx014 → pages 13
- Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, *MNRAS*, 415, 3065, doi: 10.1111/j.1365-2966.2011.18917.x → pages 9
- Keith, M., Jameson, A., Van Straten, W., et al. 2010, *MNRAS*, 409, 619 → pages 32, 110
- Keith, M. J., Coles, W., Shannon, R. M., et al. 2013, *MNRAS*, 429, 2161, doi: 10.1093/mnras/sts486 → pages 95
- Kerr, M., Reardon, D. J., Hobbs, G., et al. 2020, *PASA*, 37, e020, doi: 10.1017/pasa.2020.11 → pages 21
- Komesaroff, M. M. 1970, *Nature*, 225, 612, doi: 10.1038/225612a0 → pages 3
- Kopeikin, S. M. 1995, *ApJL*, 439, L5, doi: 10.1086/187731 → pages 47, 74
- . 1996, *ApJL*, 467, L93, doi: 10.1086/310201 → pages 47, 74
- Kraft, R. P., Burrows, D. N., & Nousek, J. A. 1991, *ApJ*, 374, 344, doi: 10.1086/170124 → pages 153
- Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006a, *Science*, 312, 549, doi: 10.1126/science.1124060 → pages 9, 109
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006b, *Science*, 314, 97, doi: 10.1126/science.1132305 → pages 19
- Kumar, P., Lu, W., & Bhattacharya, M. 2017, *MNRAS*, 468, 2726, doi: 10.1093/mnras/stx665 → pages 13

- Kumar, P., Shannon, R., Osłowski, S., et al. 2019, *ApJLetters*, 887, L30 → pages 140
- Lam, M. T. 2018, *ApJ*, 868, 33 → pages 67
- Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, *ApJ*, 821, 66,  
doi: 10.3847/0004-637X/821/1/66 → pages 75, 100
- Lam, M. T., McLaughlin, M. A., Cordes, J. M., Chatterjee, S., & Lazio, T. J. W. 2018a, *ApJ*, 861, 12, doi: 10.3847/1538-4357/aac48d → pages 68
- Lam, M. T., Ellis, J. A., Grillo, G., et al. 2018b, *ApJ*, 861, 132,  
doi: 10.3847/1538-4357/aac770 → pages 95
- Lange, C., Camilo, F., Wex, N., et al. 2001, *MNRAS*, 326, 274,  
doi: 10.1046/j.1365-8711.2001.04606.x → pages 46, 74
- Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, *Nature*, 220, 340 → pages 2
- Law, C. J., Abruzzo, M. W., Bassa, C. G., et al. 2017, *ApJ*, 850, 76,  
doi: 10.3847/1538-4357/aa9700 → pages 14
- Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2018, *ApJS*, 236, 8,  
doi: 10.3847/1538-4365/aab77b → pages 26
- Lawrence, E., Vander Wiel, S., Law, C., Burke Spolaor, S., & Bower, G. C. 2017, *AJ*, 154, 117, doi: 10.3847/1538-3881/aa844e → pages 17, 140
- Lazarus, P. 2012, *Proceedings of the International Astronomical Union*, 8, 35 → pages 32
- Lommen, A. N., & Demorest, P. 2013, *Classical and Quantum Gravity*, 30, 224001, doi: 10.1088/0264-9381/30/22/224001 → pages 37, 39, 48
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Science*, 318, 777, doi: 10.1126/science.1147532 → pages 26, 111
- Lorimer, D. R., & Kramer, M. 2005, *Handbook of Pulsar Astronomy*, Vol. 4 → pages 4, 6, 19, 26, 37, 39, 40, 41, 44, 53, 68, 151
- Lu, W., & Kumar, P. 2018, *MNRAS*, 477, 2470, doi: 10.1093/mnras/sty716 → pages 12
- Luo, J., Ransom, S., Demorest, P., et al. 2021, *ApJ*, 911, 45,  
doi: 10.3847/1538-4357/abe62f → pages xxviii, 40, 72, 74

- Lyne, A., & Manchester, R. 1988, MNRAS, 234, 477 → pages 3
- Lyne, A. G. 2009, Intermittent Pulsars, ed. W. Becker (Berlin, Heidelberg: Springer Berlin Heidelberg), 67–72, doi: 10.1007/978-3-540-76965-1\_4 → pages 9, 109
- Lyne, A. G., & Rickett, B. J. 1968, Nature, 219, 1339, doi: 10.1038/2191339a0 → pages 31
- Lyne, A. G., Burgay, M., Kramer, M., et al. 2004, Science, 303, 1153, doi: 10.1126/science.1094645 → pages 19
- Lyutikov, M. 2019, arXiv preprint arXiv:1901.03260 → pages 139
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993, doi: 10.1086/428488 → pages 109
- . 2016, VizieR Online Data Catalog, B/psr → pages xv, 8
- Manchester, R. N., & Taylor, J. H. 1977, Pulsars → pages 29
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, MNRAS, 328, 17, doi: 10.1046/j.1365-8711.2001.04751.x → pages 32, 110
- Marcote, B., Nimmo, K., Hessels, J. W. T., et al. 2020, Nature, 577, 190, doi: 10.1038/s41586-019-1866-z → pages 10, 13, 15
- Margalit, B., Beniamini, P., Sridhar, N., & Metzger, B. D. 2020, ApJL, 899, L27, doi: 10.3847/2041-8213/abac57 → pages 13
- Martinez, J. G., Stovall, K., Freire, P. C. C., et al. 2015, ApJ, 812, 143, doi: 10.1088/0004-637X/812/2/143 → pages 19
- Masui, K. W., Shaw, J. R., Ng, C., et al. 2019, ApJ, 879, 16, doi: 10.3847/1538-4357/ab229e → pages 50, 52, 55, 60
- Mazets, E. P., Golenetskij, S. V., & Guryan, Y. A. 1979a, Soviet Astronomy Letters, 5, 343 → pages 7
- Mazets, E. P., Golenskii, S. V., Ilinskii, V. N., Aptekar, R. L., & Guryan, I. A. 1979b, Nature, 282, 587, doi: 10.1038/282587a0 → pages 7
- McLaughlin, M., Lyne, A., Keane, E., et al. 2009, MNRAS, 400, 1431 → pages 110

- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, *Nature*, 439, 817, doi: 10.1038/nature04440 → pages 9, 110
- Metzger, B. D., Margalit, B., & Sironi, L. 2019, *MNRAS*, 485, 4091 → pages 139
- Meyers, B., & CHIME/Pulsar Collaboration. 2021, *The Astronomer's Telegram*, 14652, 1 → pages 169
- Michilli, D., Hessels, J. W. T., Donner, J. Y., et al. 2018a, *MNRAS*, 476, 2704, doi: 10.1093/mnras/sty368 → pages 14
- Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018b, *Nature*, 553, 182, doi: 10.1038/nature25149 → pages 138, 139
- Michilli, D., Masui, K. W., Mckinven, R., et al. 2021, *ApJ*, 910, 147, doi: 10.3847/1538-4357/abe626 → pages 63
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2021, arXiv e-prints, arXiv:2105.06979. <https://arxiv.org/abs/2105.06979> → pages 3
- Mingarelli, C. M. F., Levin, J., & Lazio, T. J. W. 2015, *ApJL*, 814, L20, doi: 10.1088/2041-8205/814/2/L20 → pages 13
- Morales, M. F. 2011, *PASP*, 123, 1265, doi: 10.1086/663092 → pages 55
- Morello, V., Barr, E. D., Stappers, B. W., Keane, E. F., & Lyne, A. G. 2020, *MNRAS*, 497, 4654, doi: 10.1093/mnras/staa2291 → pages 117
- Naidu, A., Joshi, B. C., Manoharan, P. K., & Krishnakumar, M. A. 2015, *Experimental Astronomy*, 39, 319, doi: 10.1007/s10686-015-9450-5 → pages 66
- Nakar, E. 2007, *PhR*, 442, 166, doi: 10.1016/j.physrep.2007.02.005 → pages 18
- NANOGrav Collaboration, Arzoumanian, Z., Brazier, A., et al. 2015, *ApJ*, 813, 65, doi: 10.1088/0004-637X/813/1/65 → pages 72, 76
- Newburgh, L. B., Bandura, K., Bucher, M. A., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9906, *Ground-based and Airborne Telescopes VI*, ed. H. J. Hall, R. Gilmozzi, & H. K. Marshall, 99065X, doi: 10.1117/12.2234286 → pages 52
- Ng, C., Vanderlinde, K., Paradise, A., et al. 2017, arXiv e-prints, arXiv:1702.04728. <https://arxiv.org/abs/1702.04728> → pages 61

- Ng, C., Wu, B., Ma, M., et al. 2020, *ApJ*, 903, 81,  
doi: 10.3847/1538-4357/abb94f → pages 9, 65, 130
- Ng, C., Pandhi, A., Naidu, A., et al. 2020, *MNRAS*, 496, 2836,  
doi: 10.1093/mnras/staa1658 → pages 65, 130, 131, 169
- Nicastro, L., Lyne, A. G., Lorimer, D. R., et al. 1995, *MNRAS*, 273, L68,  
doi: 10.1093/mnras/273.1.L68 → pages 86
- Nice, D., Demorest, P., Stairs, I., et al. 2015, *Tempo: Pulsar timing data analysis*.  
<http://ascl.net/1509.002> → pages xxviii, 40, 72, 122
- Nicholl, M., Williams, P. K. G., Berger, E., et al. 2017, *ApJ*, 843, 84,  
doi: 10.3847/1538-4357/aa794d → pages 13
- Olausen, S. A., & Kaspi, V. M. 2014, *ApJS*, 212, 6,  
doi: 10.1088/0067-0049/212/1/6 → pages xv, 7, 8
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *Physical Review*, 55, 374,  
doi: 10.1103/PhysRev.55.374 → pages 2
- Oppermann, N., Yu, H.-R., & Pen, U.-L. 2018, *MNRAS*, 475, 5109,  
doi: 10.1093/mnras/sty004 → pages xiv, xxiv, 17, 139, 153, 157, 158, 159,  
160, 161
- Pacini, F. 1967, *Nature*, 216, 567 → pages 2
- Paczynski, B. 1992, *AcA*, 42, 145 → pages 7
- Parent, E., Kaspi, V. M., Ransom, S. M., et al. 2018, *ApJ*, 861, 44,  
doi: 10.3847/1538-4357/aac5f0 → pages 35
- Parent, E., Chawla, P., Kaspi, V. M., et al. 2020, *ApJ*, 904, 92,  
doi: 10.3847/1538-4357/abddf6 → pages 146, 165
- Park, R. S., Folkner, W. M., Williams, J. G., & Boggs, D. H. 2021, *AJ*, 161, 105,  
doi: 10.3847/1538-3881/abd414 → pages 42, 73
- Patel, C., Agarwal, D., Bhardwaj, M., et al. 2018, *ApJ*, 869, 181,  
doi: 10.3847/1538-4357/aaee65 → pages 9, 111
- Pen, U.-L., & Connor, L. 2015, *ApJ*, 807, 179,  
doi: 10.1088/0004-637X/807/2/179 → pages 13
- Pennucci, T. T. 2019, *ApJ*, 871, 34, doi: 10.3847/1538-4357/aaf6ef → pages 107,  
124

- Pennucci, T. T., Demorest, P. B., & Ransom, S. M. 2014, *ApJ*, 790, 93,  
doi: 10.1088/0004-637X/790/2/93 → pages 107
- Perera, B. B. P., DeCesar, M. E., Demorest, P. B., et al. 2019, *MNRAS*, 490,  
4666, doi: 10.1093/mnras/stz2857 → pages 21
- Pilia, M., Burgay, M., Possenti, A., et al. 2020, *ApJL*, 896, L40,  
doi: 10.3847/2041-8213/ab96c0 → pages 15, 146, 165
- Piro, A. L. 2016, *ApJL*, 824, L32, doi: 10.3847/2041-8205/824/2/L32 → pages  
13
- Platts, E., Weltman, A., Walters, A., et al. 2019, *PhR*, 821, 1,  
doi: 10.1016/j.physrep.2019.06.003 → pages 13
- Pleunis, Z., Good, D. C., Kaspi, V. M., et al. 2021a, Submitted to *ApJ* → pages vi,  
xv, 11, 12, 15, 16, 18
- Pleunis, Z., Michilli, D., Bassa, C. G., et al. 2021b, *ApJL*, 911, L3,  
doi: 10.3847/2041-8213/abec72 → pages 15, 165
- Popov, S. B., & Postnov, K. A. 2010, in *Evolution of Cosmic Objects through  
their Physical Activity*, ed. H. A. Harutyunian, A. M. Mickaelian, & Y. Terzian,  
129–132. <https://arxiv.org/abs/0710.2006> → pages 13
- Radhakrishnan, V., & Cooke, D. J. 1969, *Astrophys. Lett.*, 3, 225 → pages 3
- Radhakrishnan, V., & Srinivasan, G. 1982, *Current Science*, 51, 1096 → pages 6
- Rajwade, K. M., Mickaliger, M. B., Stappers, B. W., et al. 2020, *MNRAS*, 495,  
3551, doi: 10.1093/mnras/staa1237 → pages 11, 139
- Ramachandran, R., Demorest, P., Backer, D. C., Cognard, I., & Lommen, A.  
2006, *ApJ*, 645, 303, doi: 10.1086/500634 → pages 100
- Rankin, J. M. 1983, *ApJ*, 274, 359 → pages 3
- Ransom, S. 2011, *PRESTO: Pulsar Exploration and Search TOolkit*.  
<http://ascl.net/1107.017> → pages xxviii, 32, 34, 114
- Ransom, S. M., Hessels, J. W. T., Stairs, I. H., et al. 2005, *Science*, 307, 892,  
doi: 10.1126/science.1108632 → pages 110
- Ravi, V. 2019, *ApJ*, 872, 88 → pages 139

- Remazeilles, M., Dickinson, C., Banday, A. J., Bigot-Sazy, M. A., & Ghosh, T. 2015, *MNRAS*, 451, 4311, doi: 10.1093/mnras/stv1274 → pages 121
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *ApJL*, 887, L21, doi: 10.3847/2041-8213/ab481c → pages 3
- Ryle, M., & Vonberg, D. D. 1946, *Nature*, 158, 339, doi: 10.1038/158339b0 → pages 52
- Scheuer, P. A. G. 1968, *Nature*, 218, 920, doi: 10.1038/218920a0 → pages 31
- Schilizzi, R. T., Dewdney, P. E. F., & Lazio, T. J. W. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7012, *Ground-based and Airborne Telescopes II*, ed. L. M. Stepp & R. Gilmozzi, 70121I, doi: 10.1117/12.786780 → pages 52
- Scholz, P., Spitler, L. G., Hessels, J. W. T., et al. 2016, *ApJ*, 833, 177, doi: 10.3847/1538-4357/833/2/177 → pages 17, 138
- Selina, R. J., Murphy, E. J., McKinnon, M., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10700, *Ground-based and Airborne Telescopes VII*, ed. H. K. Marshall & J. Spyromilio, 107001O, doi: 10.1117/12.2312089 → pages 52
- Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *ApJ*, 611, 623, doi: 10.1086/422185 → pages 20
- Shapiro, I. I. 1964, *PhRvL*, 13, 789, doi: 10.1103/PhysRevLett.13.789 → pages 19, 41
- Shapiro-Albert, B. J., Hazboun, J. S., McLaughlin, M. A., & Lam, M. T. 2021, *ApJ*, 909, 219, doi: 10.3847/1538-4357/abdc29 → pages 76
- Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014, *ApJ*, 781, 57, doi: 10.1088/0004-637X/781/2/57 → pages 50
- Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2015, *PhRvD*, 91, 083514, doi: 10.1103/PhysRevD.91.083514 → pages 50
- Spitler, L. G., Cordes, J. M., Hessels, J. W. T., et al. 2014, *ApJ*, 790, 101, doi: 10.1088/0004-637X/790/2/101 → pages 138
- Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, *Nature*, 531, 202, doi: 10.1038/nature17168 → pages 10, 17, 138

- Spitler, L. G., Herrmann, W., Bower, G. C., et al. 2018, *ApJ*, 863, 150,  
doi: 10.3847/1538-4357/aad332 → pages 15
- Staelin, D. H. 1969, *IEEE Proceedings*, 57, 724, doi: 10.1109/PROC.1969.7051  
→ pages 34
- Staelin, D. H., & Reifenstein, E. C. 1968, *Science*, 162, 1481,  
doi: 10.1126/science.162.3861.1481 → pages 2
- Stairs, I. H. 2003, *Living Reviews in Relativity*, 6, 5, doi: 10.12942/lrr-2003-5 →  
pages 19
- Stappers, B. 2016, in *MeerKAT Science: On the Pathway to the SKA*, 10 →  
pages 134
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, 791, 67,  
doi: 10.1088/0004-637X/791/1/67 → pages 32, 82, 83, 90, 110
- Swenson, G. W., J. 1969, *ARA&A*, 7, 353,  
doi: 10.1146/annurev.aa.07.090169.002033 → pages 50
- Taylor, J. H. 1974, *A&AS*, 15, 367 → pages 28, 62
- . 1992, *Philosophical Transactions of the Royal Society of London Series A*,  
341, 117, doi: 10.1098/rsta.1992.0088 → pages 39, 71
- Tegmark, M., & Zaldarriaga, M. 2009, *PhRvD*, 79, 083530,  
doi: 10.1103/PhysRevD.79.083530 → pages 55
- . 2010, *PhRvD*, 82, 103501, doi: 10.1103/PhysRevD.82.103501 → pages 55
- Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, *ApJL*, 834, L7,  
doi: 10.3847/2041-8213/834/2/L7 → pages 10, 13
- Thompson, A. R., Moran, J. M., & Swenson, George W., J. 2017, *Interferometry  
and Synthesis in Radio Astronomy*, 3rd Edition,  
doi: 10.1007/978-3-319-44431-4 → pages 50, 52
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, *PASA*, 30, e007,  
doi: 10.1017/pasa.2012.007 → pages 52
- Titus, N., Stappers, B. W., Morello, V., et al. 2019, *MNRAS*, 487, 4332,  
doi: 10.1093/mnras/stz1578 → pages 32
- Totani, T. 2013, *PASJ*, 65, L12, doi: 10.1093/pasj/65.5.L12 → pages 13

- Vallisneri, M., Taylor, S. R., Simon, J., et al. 2020, *ApJ*, 893, 112,  
doi: 10.3847/1538-4357/ab7b67 → pages 76
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, 556, A2,  
doi: 10.1051/0004-6361/201220873 → pages 52
- van Haasteren, R., & Levin, Y. 2013, *MNRAS*, 428, 1147,  
doi: 10.1093/mnras/sts097 → pages 48
- van Straten, W. 2006, *ApJ*, 642, 1004, doi: 10.1086/501001 → pages 107
- van Straten, W., & Bailes, M. 2011, *PASA*, 28, 1, doi: 10.1071/AS10021 →  
pages xxvi, 65, 122
- van Straten, W., Bailes, M., Britton, M., et al. 2001, *Nature*, 412, 158.  
<https://arxiv.org/abs/astro-ph/0108254> → pages 74
- van Straten, W., Demorest, P., & Osłowski, S. 2012, *Astronomical Research and  
Technology*, 9, 237. <https://arxiv.org/abs/1205.6276> → pages 115
- Voisin, G., Cognard, I., Freire, P. C. C., et al. 2020, *A&A*, 638, A24,  
doi: 10.1051/0004-6361/202038104 → pages 19
- Wex, N. 1995, arXiv e-prints, gr. <https://arxiv.org/abs/gr-qc/9511017> → pages 19
- Woosley, S. E., & Bloom, J. S. 2006, *ARA&A*, 44, 507,  
doi: 10.1146/annurev.astro.43.072103.150558 → pages 18
- Xu, H., Huang, Y. X., Burgay, M., et al. 2021, *The Astronomer's Telegram*,  
14642, 1 → pages 169
- Yadav, P. 2020, *Applying Convolutional Neural Networks to Classify Fast Radio  
Bursts Detected by The CHIME Telescope* → pages 64
- Yamasaki, S., Totani, T., & Kiuchi, K. 2018, *PASJ*, 70, 39,  
doi: 10.1093/pasj/psy029 → pages 13
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, 835, 29,  
doi: 10.3847/1538-4357/835/1/29 → pages 10, 27, 63, 162
- Ye, J., Wang, K., & Cai, Y.-F. 2017, *European Physical Journal C*, 77, 720,  
doi: 10.1140/epjc/s10052-017-5319-2 → pages 13
- Zhang, B. 2016, *ApJL*, 822, L14, doi: 10.3847/2041-8205/822/1/L14 → pages 13

Zhu, W. W., Stairs, I. H., Demorest, P. B., et al. 2015a, ApJ, 809, 41,  
doi: 10.1088/0004-637X/809/1/41 → pages 19, 76

Zhu, W. W., Desvignes, G., Wex, N., et al. 2019, MNRAS, 482, 3249,  
doi: 10.1093/mnras/sty2905 → pages 19, 90

Zhu, X.-J., Wen, L., Hobbs, G., et al. 2015b, MNRAS, 449, 1650,  
doi: 10.1093/mnras/stv381 → pages 20, 90

## Appendix A

# Parameter files from CHIME/Pulsar Timing Data

### PSR J0645+5158

PSR	J0645+5158		
EPHEM	DE440		
CLOCK	TT (BIPM2019)		
UNITS	TDB		
START	58647.8877313976238889		
FINISH	59071.7276610631348612		
INFO	-f		
TIMEEPH	FB90		
T2CMETHOD	IAU2000B		
DILATEFREQ	N		
DMDATA	N		
NTOA	3332		
CHI2	2726.032048222008		
ELONG	98.058551435927214	1	0.00000003643252872937
ELAT	28.852629254296826	1	0.00000008287939030866
PMELONG	2.191674678698748	0	0.00824870611624997
PMELAT	-7.274331319795477	0	0.012845094603894233

PX	0.8164997550737843	0	0.09199715964368278
ECL	IERS2010		
POSEPOCH	58860.0000000000000000		
F0	112.94972317954425365	1	3.2900323175063095062e-12
F1	-6.342235137860901675e-17	1	9.217956017998788782e-19
PEPOCH	58860.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW			0.0
SWM			0.0
DM	18.24845971652408006		
DM1			0.0
DMEPOCH	58860.0000000000000000		
DMX			6.5
FD1	1.352936117439847e-05	0	4.2800419807220464e-07
FD2	1.1267780078855536e-05	0	4.077461267871834e-07
TZRMJD	56823.7573902832733680		
TZRSITE			GB
TZRFRQ			803.59
EFAC	-f CHIME		0.9570633166331912
EQUAD	-f CHIME		0.08316171311706705
ECORR	-f CHIME		0.01629375827663391

## PSR J0740+6620

PSR	J0740+6620
EPHEM	DE440
CLOCK	TT (BIPM2019)
UNITS	TDB
START	58700.7798015187307291
FINISH	59071.7657914660322453
INFO	-f
TIMEEPH	FB90

T2CMETHOD			IAU2000B
DILATEFREQ			N
DMDATA			N
NTOA			3681
CHI2			0.0
ELONG	103.759132129331419	1	0.00000008493114544501
ELAT	44.102451737098178	1	0.00000013185365007810
PMELONG	-2.7442919344798518	0	0.01542709394755082
PMELAT	-32.48945417014345	0	0.022817515468766123
PX	0.9242700416517539	0	0.2052622168780123
ECL			IERS2010
POSEPOCH	58886.0000000000000000		
F0	346.53199632437922215	1	9.444075402418149887e-12
F1	-1.4671673742266940403e-15	1	5.47605161441997743e-18
PEPOCH	58886.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW			0.0
SWM			0.0
DM	14.962833906560818387		
DM1			0.0
DMEPOCH	58886.0000000000000000		
DMX			6.5
BINARY ELL1			
PB	4.766944620343574716	1	8.0411994905965373327e-10
PBDOT	1.2015470215405414e-12	0	2.295079423665566e-13
A1	3.9775561276430333	1	2.9591886660657727e-07
M2	0.2445258627098865	1	0.01677786022188767
SINI	0.9993591225354201	1	0.0005915975258265389
TASC	58886.8277376423531351	1	1.6788752378509608935e-08
EPS1	-5.758128330618990592e-06	1	7.511510995465730388e-08
EPS2	-1.7656292820956207832e-06	1	4.2355030487312342336e-08
FD1	-1.4279797378744484e-05	0	2.2285726289808528e-07

TZRMJD		57300.4995731594366435
TZRSITE		GB
TZRFRQ		1218.712
EFAC	-f CHIME	0.9640511384991127
EQUAD	-f CHIME	0.0233905378630033
ECORR	-f CHIME	0.012347374055675988

### PSR J1012+5307

PSR		J1012+5307
EPHEM		DE440
CLOCK		TT (BIPM2019)
UNITS		TDB
START		58720.8302808511292245
FINISH		59071.8704198927009608
INFO		-f
TIMEEPH		FB90
T2CMETHOD		IAU2000B
DILATEFREQ		N
DMDATA		N
NTOA		5885
CHI2		0.0
ELONG		133.361146264805399 1 0.00000030268409669900
ELAT		38.755255905494813 1 0.00000078582194657796
PMELONG		13.957262992173092 0 0.022794408506559182
PMELAT		-21.51902991829749 0 0.036709245849169014
PX		0.9462718696102586 0 0.23476878114972521
ECL		IERS2010
POSEPOCH		58896.0000000000000000
F0		190.26783718300371066 1 2.8147398184680409847e-11
F1		-6.201614708381611771e-16 1 1.6540851472227821461e-17
PEPOCH		58896.0000000000000000
CORRECT_TROPOSPHERE		Y

PLANET_SHAPIRO			Y
NE_SW			0.0
SWM			0.0
DM		9.03511782607564422	
DM1			0.0
DMEPOCH		58896.0000000000000000	
DMX			6.5
BINARY ELL1			
PB		0.6046727141126525569	1 8.4730167000108273297e-10
PBDOT		4.4947439850751304e-14	0 7.261213123483646e-15
A1		0.5818180809746374	1 7.416586121031073e-07
TASC		58895.8155912251422706	1 1.2186393422516063798e-07
EPS1		1.7085573335612043701e-06	0 2.5081368668595299374e-07
EPS2		4.731190565428794977e-07	0 2.6306669691444514114e-07
FD1		0.00011100616974179027	0 2.7307822832781177e-05
FD2		-0.00010223586912930505	0 2.78099853179052e-05
FD3		6.478894980876388e-05	0 2.1958073070385248e-05
FD4		-2.681422365432613e-05	0 1.1599588646379008e-05
TZRMJD		55578.4008005042764700	
TZRSITE			GB
TZRFRQ			800.459
EFAC	-f CHIME		1.0151061442811693
EQUAD	-f CHIME		9.261228696834493
ECORR	-f CHIME		7.092266055115262
RNAMP		0.8312814065852758	
RNIDX		-1.4902569725816166	

## PSR J1125+7819

PSR		J1125+7819
EPHEM		DE440
CLOCK		TT (BIPM2019)
UNITS		TDB

START	58712.8781154743049653		
FINISH	59071.9213352261506250		
INFO		-f	
TIMEEPH		FB90	
T2CMETHOD		IAU2000B	
DILATEFREQ		N	
DMDATA		N	
NTOA		5543	
CHI2	8641.657472275987		
ELONG	115.629347189573807	1	0.00000044155295951085
ELAT	62.452062143453752	1	0.00000032528323146292
PMELONG	16.33852629689881	0	0.05100165451374879
PMELAT	22.367067599866903	0	0.055834403822880324
PX	4.9774303026150255	0	1.9769059735765098
ECL		IERS2010	
POSEPOCH	58892.0000000000000000		
F0	238.00405311918742876	1	2.6992464590745881057e-11
F1	-3.8932539688626281008e-16	1	1.4006143016637585275e-17
PEPOCH	58892.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW		0.0	
SWM		0.0	
DM	11.216258838346784222		
DM1		0.0	
DMEPOCH	58892.0000000000000000		
DMX		6.5	
BINARY ELL1			
PB	15.355445981835102637	1	1.0876438166975216025e-08
A1	12.192421191118049	1	2.842965547764342e-07
TASC	58890.7457257015710823	1	5.7131729324027240712e-08
EPS1	-1.29224889310483844e-05	1	3.608294463439440656e-08

EPS2	8.199380806246728621e-07	1	3.552593195950699582e-08
FD1	-2.9262836960751683e-05	0	3.2556799005182287e-06
FD2	5.124743046389226e-05	0	2.7384536651118237e-06
TZRMJD	57334.5283861211098843		
TZRSITE			GB
TZRFRQ			1467.835
EFAC	-f CHIME		1.0985151019461477
EQUAD	-f CHIME		5.6070258131541735
ECORR	-f CHIME		0.1575919780097415

## PSR J1713+0747

PSR	J1713+0747		
EPHEM			DE440
CLOCK			TT (BIPM2019)
UNITS			TDB
START	58720.1252874834757176		
FINISH	59071.1647041045582639		
INFO			-f
TIMEEPH			FB90
T2CMETHOD			IAU2000B
DILATEFREQ			N
DMDATA			N
NTOA			4480
CHI2	53370.46325862117		
ELONG	256.668711800125891	1	0.00000039258590191528
ELAT	30.700351234959356	1	0.00000043586118615878
PMELONG	5.260880424961834	0	0.0016584099271442104
PMELAT	-3.4398761926087924	0	0.003459691494652847
PX	0.8876358558233133	0	0.026345581742783848
ECL			IERS2010
POSEPOCH	58896.0000000000000000		
F0	218.81184367239907854	1	2.7945257310905613892e-11

F1	-4.1642651402508896229e-16	1	1.4856063248080293587e-17
PEPOCH	58896.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW			0.0
SWM			0.0
DM	16.051165526247685746		
DM1			0.0
DMEPOCH	58896.0000000000000000		
DMX			6.5
BINARY_DDK			
PB	67.82513010247075732	1	2.0583617413766555197e-07
PBDOT	6.64515657593449e-13	0	2.944308088851705e-13
A1	32.34242482542219	1	7.75118610389174e-07
A1DOT			0.0
ECC	7.494962249143998e-05	1	3.915629110950788e-08
EDOT	1.709207862179033e-18	0	7.167627443972515e-18
T0	58915.7353371070894532	1	0.0056306848691458797826
OM	176.15793191368356531	1	0.029886661183250413237
OMDOT			0.0
M2	0.3124175003063606	0	0.015307360078344554
A0			0.0
B0			0.0
GAMMA			0.0
DR			0.0
DTH			0.0
KIN	70.16423730504314	0	0.724202444880325
KOM	285.4489332425267	0	1.8175884054422349
K96			Y
FD1	0.0005125022530566303	0	3.48335578011409e-06
FD2	-0.0005413440830725197	0	3.550500301437115e-06
FD3	0.00038120735612496746	0	2.2730362768176697e-06
FD4	-0.00018749446199034034	0	7.35300672373462e-07

FD5		4.815713565531543e-05	0	2.099743080220295e-06
TZRMJD		55670.1714780480012384		
TZRSITE			GB	
TZRFRQ		1442.413		
EFAC	-f CHIME	1.1043949694520103		
EQUAD	-f CHIME	5.119865976501277		
ECORR	-f CHIME	0.02444687948238013		
RNAMP		0.01461754978704134		
RNIDX		-3.044540033677644		

### PSR J1744-1134

PSR		J1744-1134		
EPHEM		DE440		
CLOCK		TT (BIPM2019)		
UNITS		TDB		
START		58601.4719896277822686		
FINISH		59070.1887776354781134		
INFO		-f		
TIMEEPH		FB90		
T2CMETHOD		IAU2000B		
DILATEFREQ		N		
DMDATA		N		
NTOA		1743		
CHI2		0.0		
ELONG		266.119454044584757	1	0.00000011616484561228
ELAT		11.805177249138746	1	0.00000082767210882404
PMELONG		19.043278983663264	0	0.006821054237798855
PMELAT		-8.906298555820596	0	0.03732239077582539
PX		2.5297108316510966	0	0.08568103161804638
ECL		IERS2010		
POSEPOCH		58836.0000000000000000		
F0		245.42612331684526611	1	1.8677634185216726496e-11

F1	-5.448299385233843407e-16	1	5.2516268419228153596e-18
PEPOCH	58836.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW			0.0
SWM			0.0
DM	3.1387901131626590129		
DM1			0.0
DMEPOCH	58836.0000000000000000		
DMX			6.5
FD1	9.166209325492177e-07	0	1.6193580054852453e-05
FD2	7.4816252778449685e-06	0	1.653172344360118e-05
FD3	-1.1410591445365322e-05	0	1.2595792526884228e-05
FD4	1.0364021984911366e-05	0	6.1552076852200395e-06
TZRMJD	55578.5384953116674421		
TZRSITE			GB
TZRFRQ			753.605
EFAC	-f CHIME		1.0672250470629603
EQUAD	-f CHIME		0.3133047926076498
ECORR	-f CHIME		0.46881567748142267
RNAMP			0.005725435344140706
RNIDX			-4.55180499581586

## PSR B1937+21

PSR		B1937+21
EPHEM		DE440
CLOCK		TT (BIPM2019)
UNITS		TDB
START	58859.8438089232807177	
FINISH	59070.2683512487421875	
INFO		-f
TIMEEPH		FB90

T2CMETHOD		IAU2000B		
DILATEFREQ			N	
DMDATA			N	
NTOA		4554		
CHI2	11310.961600873485			
ELONG	301.973243677646337	1	0.00000028836840546313	
ELAT	42.296751022061599	1	0.00000008929995049940	
PMELONG	-0.031033066367505368	0	0.004048294193910451	
PMELAT	-0.4105851193475993	0	0.006052064173815528	
PX	0.20050464309416707	0	0.048567205633431594	
ECL			IERS2010	
POSEPOCH	58965.0000000000000000			
F0	641.92821969918257535	1	4.3196051668004565527e-11	
F1	-4.3229483583083111257e-14	1	3.609801944316340895e-17	
PEPOCH	58965.0000000000000000			
CORRECT_TROPOSPHERE				Y
PLANET_SHAPIRO				Y
NE_SW			0.0	
SWM			0.0	
DM	70.99233767510194501			
DM1			0.0	
DMEPOCH	58965.0000000000000000			
DMX			6.5	
FD1	-0.00019459390677303807	0	1.204592917019061e-06	
FD2	0.0001976280066346023	0	1.2914431795783832e-06	
FD3	-0.00014003171355139445	0	1.0617903947621573e-06	
FD4	7.857345493546854e-05	0	4.855220197010139e-07	
FD5	-2.769471588620341e-05	0	4.0833817765975393e-07	
TZRMJD	55616.4178018924963079			
TZRSITE				GB
TZRFRQ			1742.942	
EFAC	-f CHIME		0.8148874622647135	
EQUAD	-f CHIME		0.44707391780059297	

ECORR	-f CHIME	0.5700843643399807
RNAMP		0.1466418720540249
RNIDX		-3.365582120674361

## PSR J2145-0750

PSR	J2145-0750		
EPHEM	DE440		
CLOCK	TT (BIPM2019)		
UNITS	TDB		
START	58721.3104730865673843		
FINISH	59071.3531848356217245		
INFO	-f		
TIMEEPH	FB90		
T2CMETHOD	IAU2000B		
DILATEFREQ	N		
DMDATA	N		
NTOA	3667		
CHI2	0.0		
ELONG	326.024583692137924	1	0.00000033550662816408
ELAT	5.313066727379374	1	0.00000803593463307988
PMELONG	-12.030784470696817	0	0.029077201730038374
PMELAT	-4.615346725440383	0	0.30772749022958135
PX	1.6226556997126	0	0.16700579972384064
ECL	IERS2010		
POSEPOCH	58896.0000000000000000		
F0	62.29588876435069184	1	1.454592934609826501e-11
F1	-9.8928729152025097776e-17	1	8.203816936452912655e-18
PEPOCH	58896.0000000000000000		
CORRECT_TROPOSPHERE			Y
PLANET_SHAPIRO			Y
NE_SW			0.0

SWM			0.0
DM		9.0017044656527312844	
DM1			0.0
DMEPOCH		58896.0000000000000000	
DMX			6.5
BINARY ELL1H			
PB		6.8389025103332941958	1 4.314620061355244852e-09
A1		10.164109887575357	1 7.385175574282122e-07
A1DOT		5.823370374237248e-15	0 9.221335651774318e-16
TASC		58892.7197792989014324	1 7.594221417322295073e-08
EPS1		-6.8806041120165291557e-06	1 1.4705437594857686985e-07
EPS2		-1.803448294534017451e-05	1 1.4639083274507585919e-07
H3		1.9235263667118798016e-07	0 5.70957065921761416e-08
NHARMS			3
FD1		-1.1001005007578426e-05	0 1.796558697359531e-06
FD2		3.6575836898993236e-05	0 1.3550970132126788e-06
TZRMJD		55616.5169794573625233	
TZRSITE			GB
TZRFRQ			1218.023
EFAC	-f CHIME		1.0910862229550096
EQUAD	-f CHIME		7.468863212023451
ECORR	-f CHIME		2.664968237984026
RNAMP		0.7626581397674012	
RNIDX		-1.3910128617802977	