

MASSIVE, BRIGHT, AND BRAND-NEW MILLISECOND PULSARS AS PROBES OF FUNDAMENTAL PHYSICS

Hannah Thankful Cromartie

Chapel Hill, North Carolina

B.S. Physics with Highest Honors, The University of North Carolina at Chapel Hill, 2014
M.S. Astronomy, University of Virginia, 2016

A Dissertation Presented to the
Graduate Faculty of the
University of Virginia
in Candidacy for the Degree of
Doctor of Philosophy

Department of Astronomy
University of Virginia
May 2020

Committee Members:
Dr. Scott Ransom
Dr. Phil Arras
Dr. Craig Sarazin
Dr. Kent Yagi

© Copyright by
Hannah Thankful Cromartie
All rights reserved
May 16, 2020

*Mom, I wish you were here to say:
“Sweetie, I don’t understand a word of this,
but I’m really proud of you.”*

Abstract

Spinning hundreds of times a second and with densities surpassed only by black holes, millisecond pulsars (MSPs) are among the most astonishing astrophysical objects in the Universe. Nearly 40 years of careful study have revealed their efficacy as probes of physical phenomena that could not otherwise be studied by scientists in Earth-based laboratories. This dissertation touches on several disparate aspects of pulsar science, all of which contribute (either directly or indirectly) to the goal of furthering our understanding of fundamental physics. All of the work detailed in this thesis was conducted under the umbrella of the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) collaboration, the goal of which is to detect the low-frequency stochastic background of gravitational waves from supermassive black hole binaries using an array of precisely timed MSPs called a Pulsar Timing Array (PTA).

Because a PTA's sensitivity to gravitational waves is linearly proportional to the number of MSPs in the array, the discovery of new MSPs with high timing precision is a critical component of NANOGrav's effort to detect gravitational waves. We report the discovery of the first six MSPs in *Fermi* unassociated Gamma-ray sources discovered with the Arecibo telescope (Cromartie et al. 2016), as well as the 2017 discovery of an additional thirteen with the Arecibo and Green Bank Telescopes (discussed later in the dissertation). Five of the six new Arecibo sources — a disproportionately

large fraction — are in highly accelerated binaries with low-mass companions. We therefore explore Arecibo’s predisposition to discovering interacting binaries quantitatively. Though the six original *Fermi* MSPs discovered with Arecibo are not appropriate for PTA inclusion, at least one of the thirteen 2017 discoveries is being provisionally included in the NANOGrav PTA.

Next, we discuss our measurement of the most massive neutron star observed to date. J0740+6620 is a $2.14_{-0.09}^{+0.10} M_{\odot}$ ($1\text{-}\sigma$ confidence interval) MSP in the NANOGrav dataset for which we obtained supplementary, orbital-phase-specific observations with the Green Bank Telescope in order to constrain its mass using the relativistic Shapiro delay (Cromartie et al. 2020). This discovery has significant implications for the extremely poorly understood neutron star equation of state, which describes the behavior of supranuclear-density matter deep within neutron star cores.

We then present the results of an additional Shapiro delay-powered endeavor, this time in order to constrain the mass of the bright Gamma-ray MSP J1231–1411. This source is of particular interest to the Neutron Star Interior Composition Explorer (NICER) mission. Forthcoming modeling of the source’s X-ray lightcurve promises to constrain its mass-to-radius ratio, which could further our understanding of the equation of state. An independent measurement of its mass via the radio Shapiro delay would improve the NICER team’s modeling of the MSP, and in turn, the project’s potential scientific payoff. We conducted a multi-wavelength analysis of timing data, including a new 22-hour campaign over orbital conjunction using the Green Bank Telescope. Both traditional χ^2 minimization fitting and Markov Chain Monte Carlo (MCMC)-based techniques indicate that this source is a low-mass MSP in a highly inclined binary orbit with a low-mass white dwarf. The MCMC trials, which prove the constraining power of our measurement of the white dwarf mass and orbital inclination, are informed by priors based on white dwarf evolutionary models. We

also conduct a single-photon MCMC fit to 12 years of *Fermi*-LAT Gamma-ray data, though the resulting constraints on MSP mass are not as stringent as the (provisional) constraints from radio data.

We conclude by sharing a potpourri of pulsar timing projects that are either in-progress, or that did not merit publication despite being of potential interest to the reader. The thirteen (unpublished) 2017 *Fermi* MSPs, a survey of highly accelerated “spider” MSP systems for inclusion in the NANOGrav array, and an additional three MSPs for which we conducted targeted Shapiro delay campaigns are all discussed in this final chapter.

Acknowledgements

Had I been asked ten years ago to place a bet on the subject of this doctoral dissertation, I would have put my money on “Journalism” or “French.” I owe my love of *astrophysics*, of all things, to two individuals. The first is Carl Sagan, because watching *Cosmos* as a senior in high school made me feel like an awestruck little girl looking through a telescope for the first time. The second is Dan Reichart, my undergraduate advisor at Carolina. Not only did his “Educational Research in Radio Astronomy” summer course at Green Bank convince me to pursue physics and astronomy, but it also sparked my love affair with *radio* astronomy that continues to this day. Dan, I owe you my deepest gratitude for your unmatched kindness and support.

Fernando Camilo turned me into a rabid pulsar hard-liner. If I hadn’t conducted my REU research with him at Arecibo — which quickly became my favorite place on Earth (and not just because I got to live in Jodie Foster’s cabin) — I doubt I’d have found a sub-field of astrophysics that I love quite so much. Thank you for everything.

I met my doctoral advisor, Scott Ransom, during that summer at Arecibo (you told me about `pushd` and `popd!`). Scott, I know that I haven’t always been an exemplary PhD student, and I thank you for sticking with me through troubled times. You’ve been a wonderful teacher and advocate, and are one of the reasons I’ll think fondly of graduate school — thank you.

I am also grateful to my scientific collaboration, the North American Nanohertz

Observatory for Gravitational Waves, and the National Radio Astronomy Observatory's Grote Reber Doctoral Fellowship program for financially supporting my graduate studies. Joining NANOGrav opened countless doors for me, and introduced me to colleagues and friends I hope to keep for the rest of my career (see Section 1.10).

The last six years have been the most trying of my life. I wouldn't have made it through without my friends, both in the department (Sean, Paul, Matt, Trey, Lauren, Molly, Robby, Allison M., Allison T., David, Luca, et al.), and outside of it (Gina, Dan, Katrina, Casey, Connor, Taylor, Zane, Ari, and Madeline, who still calls me her little astronaut). I cannot forget Sadie, Ernie, Samson, Gwydion, and Calvin, who are perfect dogs, and Ezreal, who sits on my head every morning to wake me up. Woof.

Riley, I have almost forgiven you for referring to me as "Grateful, or whatever the f— her name is" after we first met. You have been a lifesaver, both figuratively (moving back to Chapel Hill to help take care of my mom), and literally (pulling me out of the way of that Nazi's car when he almost killed us). You're a decent housemate. Thanks.

Dad, you've endured more than your fair share of hardship, and your strength has been an inspiration; what's more, you never failed to support me along the way. Thank you for everything, but especially the laminated Vim cheat sheet I still keep on my desk.

Mom, I miss you. I haven't come to terms with losing you, or with the pain and fear you endured so gracefully. I could only dream of being as tough as you always were. You used to tell me that I couldn't go on any dates until I got my PhD. Joke's on me... I won't be going on any dates even if they pass me!

To the rest of the Cromarties, I love you all. Not everyone is lucky to have such a supportive extended family. Thanks also to Jim and Caroline, who make me feel like

part of your clan.

Lastly, I would like to acknowledge that the University of Virginia was built on the unceded land of the Monacan peoples. In fact, Thomas Jefferson was lauded as a “father of American archaeology” after excavating a burial mound on his property where he observed Monacans grieving¹. UVA’s founder was a heinous man who enslaved his own children and countless others. We would do well to remember that this University’s grounds are steeped in innocent blood.

¹<https://www.monacannation.com/our-history.html>

Table of Contents

Abstract	v
Acknowledgements	ix
List of Figures	xviii
List of Tables	xix
1 An Introduction to Millisecond Pulsars, Pulsar Timing Arrays, and More	1
1.1 The Origins of Pulsar Astrophysics	2
1.2 Pulsars	3
1.2.1 Millisecond Pulsars	8
1.3 Instrumentation for Pulsar Observations	9
1.4 Searching for (Millisecond) Pulsars	10
1.4.1 RFI Mitigation	11
1.4.2 Dispersion	12
1.4.3 Fourier Transforms and Acceleration Searches	13
1.4.4 Inspecting Candidates	17
1.5 Pulsar Timing	18
1.6 Relativistic Shapiro Delay & the Neutron Star Equation of State	22
1.7 Pulsar Timing Arrays & NANOGrav	25
1.8 Pulsar Observations at High Energies	27
1.9 Overview	28
1.10 Collaborators	30
2 Six New Millisecond Pulsars in <i>Fermi</i> Gamma-Ray Sources	31
2.1 The First Six Arecibo <i>Fermi</i> MSPs	32
2.1.1 Introduction	32
2.1.2 Candidate Selection	34
2.1.3 Observations and Data Analysis	37
2.1.4 Results	39
2.1.5 Discussion	45

3	Relativistic Shapiro Delay Measurements of an Extremely Massive Millisecond Pulsar	55
3.1	The Most Massive Neutron Star to Date	56
3.1.1	Preface	56
3.1.2	Relativistic Shapiro Delay Measurements of an Extremely Massive Millisecond Pulsar	56
3.1.3	Methods	65
3.1.3.1	Green Bank Telescope Observations	65
3.1.3.2	Generation of TOAs and the Timing Model	66
3.1.3.3	Assessment of Timing Noise	68
3.1.3.4	Dispersion Measure Modeling	69
3.1.3.5	Simulations	70
3.1.3.6	Data Availability	71
3.1.3.7	Code Availability	71
4	Radio Shapiro Delay Observations of the Bright Gamma-Ray MSP J1231–1411	73
4.1	Abstract	73
4.2	Introduction	74
4.3	Observations & Initial Analysis	75
4.3.1	Green Bank Telescope, 2009-2010	78
4.3.2	Nançay Radio Telescope	78
4.3.3	<i>Fermi</i> LAT	79
4.3.4	Green Bank Telescope, 2019	79
4.4	Methods Overview	80
4.5	Analysis & Results	81
4.5.1	Noise Analysis for Radio Data	81
4.5.2	PINT & TEMPO2 χ^2 Analysis	82
4.6	PINT MCMC trials	86
4.6.1	Priors on m_c and $\sin(i)$	87
4.6.2	Results of Six MCMC Trials for Radio Observations	88
4.6.3	Single-Photon Gamma-Ray Timing	89
4.6.4	Discussion	104
4.7	Conclusions	108
4.8	Acknowledgements	108
5	Miscellanea	111
5.1	Introduction	111
5.2	Thirteen New Sources from <i>Fermi</i> Gamma-Ray Searches at the GBT and Arecibo	112
5.2.1	Observations	112
5.2.2	Results	113

5.2.3	Future Work	116
5.3	NANOGrav Timing	116
5.4	Other Shapiro Delay Projects	121
5.5	Supplemental Figures to Chapter 3	123
6	Summary & Closing Remarks	127
	References	131

List of Figures

1.1	$P - \dot{P}$ Diagram of Pulsar Evolution	5
1.2	Diagram of the Accepted Pulsar Model	7
1.3	Demonstration of Frequency-Dependent Smearing by the ISM	14
1.4	Results of a Discrete Fourier Transform	16
1.5	Doppler Smearing in the Fourier Domain	17
1.6	Example <code>prepfold</code> Plot for a New MSP Candidate	19
1.7	Solving for the Orbit of a Binary Pulsar	21
1.8	A Toy Model of Shapiro Delay	24
1.9	Results of a Bayesian Single-Photon Timing Analysis for B1937+21	29
2.1	Radio Sensitivity Comparison of <i>Fermi</i> -LAT-Directed MSP Searches	40
2.2	Pulse Profiles from the First Six Arecibo <i>Fermi</i> MSPs	42
2.3	New Spider MSPs from the First Arecibo <i>Fermi</i> Search	43
3.1	J0740+6620 Timing Residuals vs. Orbital Phase	60
3.2	J0740+6620 Timing Residuals and DMX vs. MJD	62
3.3	χ^2 Maps for J0740+6620: m_p , m_c , and i	63
4.1	J1231–1411 Discovery Plot	76
4.2	J1231–1411 NRT and <i>Fermi</i> Pulse Profiles	77
4.3	χ^2 Gridding Results for J1231–1411	84
4.4	MCMC Triangle Plot for Trial 1 (Strict TS99 UB)	90
4.5	3-Panel Posterior PDFs and Priors, Trial 1 (Strict TS99 UB)	91
4.6	MCMC Triangle Plot for Trial 2 (Wide TS99 UB)	92
4.7	3-Panel Posterior PDFs and Priors, Trial 2 (Wide TS99 UB)	93
4.8	MCMC Triangle Plot for Trial 3 (Very Wide UB)	94
4.9	3-Panel Posterior PDFs and Priors, Trial 3 (Very Wide UB)	95
4.10	MCMC Triangle Plot for Trial 4 (Strict Gaussian)	96
4.11	3-Panel Posterior PDFs and Priors, Trial 4 (Strict Gaussian)	97
4.12	Trial 4 Zoomed-In	98
4.13	MCMC Triangle Plot for Trial 5 (Mid-Width Gaussian)	99
4.14	3-Panel Posterior PDFs and Priors, Trial 5 (Mid-Width Gaussian)	100
4.15	MCMC Triangle Plot for Trial 6 (Wide Gaussian)	101

4.16	3-Panel Posterior PDFs and Priors, Trial 6 (Wide Gaussian)	102
4.17	MCMC Triangle Plot for <i>Fermi</i> Data	105
4.18	Post-Fit Pulse Profile for <i>Fermi</i> Data	106
5.1	Eclipses in the New MSP J1555–2908	115
5.2	Scintillation in the New MSP with a Massive Companion, J1304+12 .	117
5.3	Pulse Profiles from Four New <i>Fermi</i> MSPs Discovered with the GBT	117
5.4	Pulse Profiles for Six New <i>Fermi</i> MSPs Discovered at Arecibo	118
5.5	Timing Residuals for B1937+21 in the NANOGrav 12.5-year Dataset	120
5.6	J1125+7819 Orbital Coverage	122
5.7	Various EoS and Neutron Star Mass Constraints	124
5.8	Shapiro Delay Observation Scheme Comparison	125

List of Tables

2.1	Summary of Arecibo <i>Fermi</i> -LAT Source Search	35
2.2	Observing Parameters for Four <i>Fermi</i> -LAT Radio Surveys	38
2.3	Pulsars Discovered in Arecibo Searches of <i>Fermi</i> -LAT Sources	44
2.4	Arecibo Searches: Gamma-Ray Source Information from 3FGL Catalog	47
2.5	γ_{2m}^2 for PSR J2052+1218 as a Function of Integration Time	52
3.1	PSR J0740+6620 Best-Fit Parameters	67
4.1	Summary of White Noise Parameters Derived with <i>Enterprise</i>	83
4.2	Mass and Inclination Confidence Intervals from χ^2 Gridding with <i>TEMPO2</i>	83
4.3	PSR J1231–1411 Best-Fit Parameters from Linear Regression	85
4.4	1- and 2- σ Confidence Intervals on Shapiro Parameters for Six MCMC Trials	103
5.1	Summary of Arecibo and GBT <i>Fermi</i> -LAT Source Search	114

Chapter 1

An Introduction to Millisecond Pulsars, Pulsar Timing
Arrays, and More

1.1 The Origins of Pulsar Astrophysics

In late November 1967, Cambridge graduate student Jocelyn Bell (now Bell Burnell) noticed a peculiar, periodic “scruff” in the output of her strip chart recorder while collecting data for an 81-MHz extragalactic survey. Though spurious radio pulsations are common (frustratingly so; see Section 1.4), she chose to pursue the anomalous signal further. Months of poring over chart recordings led her to discover that the pulsations, which arrived once every sidereal day, were indeed astrophysical in origin. Bell Burnell’s tenaciousness and keen eye for periodicity gave way to one of the most consequential astrophysical discoveries of the 20th century: *pulsars*.

Bell Burnell’s doctoral advisor was first author (and Bell the second) on a paper detailing her discovery of a “rapidly pulsating radio source” (Hewish et al. 1968). The manuscript posited that LGM-1 (“little green men,” a playful name for what is now known as B1919+21) was nearby and most likely associated with white dwarfs or neutron stars. Though Walter Baade and Fritz Zwicky had conceived of neutron stars as early as 1934 (Baade & Zwicky 1934), Bell Burnell’s discovery marked the first direct detection of these small, dense astrophysical marvels. It was Hewish — not Bell Burnell — who was awarded the 1974 Nobel Prize in Physics for the discovery of pulsars¹.

In the fifty years since their discovery, pulsars have emerged as one of the most powerful space-based probes of astrophysical processes, and persist as sources of astrophysical intrigue in their own right. This dissertation runs the gamut of pulsar astronomy, addressing somewhat disparate topics: searching for new pulsars, timing pulsars at Gamma-ray frequencies, using an ensemble of pulsars to detect low-

¹Bell Burnell never publicly expressed anger over this injustice, later becoming a decorated astrophysicist.

frequency gravitational waves, observing general relativistic effects in pulsar binaries, and using pulsar observations to constrain the neutron star equation of state. Ideally, this text will impress upon its readers the remarkable *utility* and *versatility* that pulsars offer.

Much of the information contained in the following sections is mirrored in published works such as Lorimer & Kramer (2004) and Lyne & Graham-Smith (2006), both of which significantly expound upon the information presented here.

1.2 Pulsars

Pulsars are rapidly rotating, highly magnetized neutron stars that beam electromagnetic radiation from their magnetic poles. These exotic compact objects are the lighthouses of our Galaxy; a sufficiently nearby observer will notice a periodic brightening when the pulsar’s beam (often misaligned from its rotational axis) sweeps across their line of sight (e.g. Pacini 1967, 1968; Gold 1968; Richards & Comella 1969). In this way, pulsars do not *pulse* at all.

Core collapse supernovae (Types Ib, Ic, and II) are the progenitors of pulsars (Large et al. 1968). As a stellar core approaches the Chandrasekhar mass (a white dwarf’s maximum mass, $\sim 1.4 M_{\odot}$), electron degeneracy pressure eventually fails to counteract gravitational forces, and collapse begins. Electron capture by protons yields an abundance of neutrons and neutrinos, while conservation of angular momentum produces a neutron star with spin frequencies between ~ 1 and ~ 700 Hz. Conservation of magnetic flux as the core collapses yields a remnant with B -fields approximately 10^{10} times larger than the progenitor’s (typical pulsars have $B \sim 10^{12}$ G). Stellar progenitors of neutron stars are generally as light as $8 M_{\odot}$, though the upper limit — the mass at which a black hole would form in lieu of a neutron star — is a

contentious topic². The canonical pulsar, the characteristics of which are disputed in Chapter 3, is generally conceived of as a $1.4\text{-}M_{\odot}$ neutron star with a radius of ~ 10 km.

Pulsars are observed to spin down over time, though the rate of that spin-down varies significantly between populations (see Figure 1.1). Early models proposed that magnetic dipole radiation in a vacuum robs the pulsar of its rotational kinetic energy, releasing extremely low-energy radiation that causes the pulsar to slow, and in some circumstances, stimulate emission in a surrounding nebula. The rate of the resulting increase in P (i.e. dP/dt or \dot{P}) would be related to the “spin-down luminosity” (the neutron star’s total power output; \dot{E}) as:

$$\dot{E} = 4\pi^2 I \dot{P} P^{-3} \quad (1.1)$$

where the moment of inertia I can be simply calculated as $(2/5)MR^2$ for a uniform-density sphere (of order $\sim 10^{45}$ g cm²). Very little of this energy output (which can range from a fraction of L_{\odot} to more than $10^5 L_{\odot}$) is associated with the pulsar’s radio emission. In this simple case, a pulsar with aligned spin and magnetic axes could not spin down, however. More modern models attribute pulsar spin-down to the torque induced by electric currents established along field lines (see Contopoulos & Spitkovsky 2006; Spitkovsky 2006). Because neutron stars are not surrounded by a vacuum, it also becomes necessary to include a misalignment term describing the angle between the magnetic and spin axes.

Fifty years of careful study have yet to produce a comprehensive model of pulsar emission. The consensus model describes pulsars as spinning dipoles (see Goldreich & Julian 1969 and Figure 1.2). Their significant magnetic fields induce an electric

²This is due in part to the uncertainty in neutron star formation mechanisms. If they can merge to produce stable, ultra-massive neutron stars, this “mass gap” is not as easily understood.




Figure 1.1: The $P - \dot{P}$ diagram is sometimes likened to an HR diagram, as it conveys the evolutionary tracks of several pulsar populations. We are primarily concerned with radio pulsars (black dots), and especially binary pulsars (surrounded by black circles). Lines of constant magnetic field strength, luminosity, and characteristic age give a more detailed sense of these distinct populations. Young pulsars (pink and yellow; top left) and typical radio pulsars are spinning down rather quickly, and have high associated luminosities. Millisecond pulsars (bottom left; see Section 1.2.1) are older and have weak magnetic fields compared to young pulsars and magnetars (blue triangles; top right). *Figure courtesy of Scott Ransom*

field that pulls charge from the neutron star’s surface, creating a surrounding plasma called the *magnetosphere*. This plasma spins (co-rotates) with the pulsar, though a limit exists (the “light cylinder”) where the plasma velocity approaches c . Outside of the light cylinder, magnetic field lines are successful in reconnection. Inside this area (the polar caps as defined by the axis of rotation), magnetic field lines cannot close. Electrons emit curvature radiation during their escape as the magnetic force along these open lines accelerates the particles. Energetic photons produced during this process trigger a cascade of electron-positron pairs that produce still more photons, which may explain the observed radio emission. It is worthwhile to note that the cut-and-dry description of reconnection outside of the light cylinder has been challenged, and the situation is likely more complicated than this simplified model suggests (Contopoulos & Spitkovsky 2006; Philippov & Spitkovsky 2014). Additionally, the emission-by-cascade model is simplistic, and is only one of many proposed emission mechanisms (see Philippov et al. 2019).

The pulse profile width is negatively correlated with observing frequency for typical radio pulsars, implying that emission occurs at different heights. Millisecond pulsars have smaller light cylinders and magnetospheres, and therefore produce less variation in pulse width at a variety of observing frequencies.

The topics addressed in this thesis deal mainly with typical rotation-powered radio pulsars and recycled (millisecond) pulsars; however, these two groups do not represent the entire pulsar population. For example, it is thought that a small subset of newly born neutron stars may undergo an additional dynamo effect, transferring more energy to the creation of an extremely strong magnetic field (Martin et al. 2014). These aptly named “magnetars” — the existence of which was predicted in the early 1990s (Duncan & Thompson 1992) — are neutron stars with extreme magnetic fields of order $\sim 10^{15}$ G. Rapidly rotating radio transients (RRATs) are another subset




Figure 1.2: Pulsars are rapidly rotating neutron stars with (usually) misaligned spin and magnetic axes. The light cylinder marks the location where the co-rotating plasma velocity approaches c , and where magnetic field line reconnection is impossible. The pulsar's radio beam is formed by curvature radiation within this light cylinder. *Figure from Lorimer & Kramer (2004)*

of pulsars found more often in single-pulse searches due to their periodic bursting behavior. Accretion-powered pulsars (X-ray binary pulsars) obtain their rotational energy during gas accretion from a stellar companion. A significant population of pulsars can be found in globular clusters; however, this work pertains only to nearby pulsars in the Galactic plane.

The extremely steady spin period of pulsars — and especially millisecond pulsars — facilitates their use as astrophysical “clocks.” The careful accounting necessary to track each pulse from these clocks is known as pulsar timing (see Section 1.5). Pulsar timing has facilitated the use of pulsars as off-world physics laboratories, enabling validations of general relativity (Hulse & Taylor 1975; Archibald et al. 2019) and bolstering our understanding of nuclear physics (Demorest et al. 2010; Antoniadis et al. 2013; Cromartie et al. 2020). In short order, they will provide the first detection of low-frequency gravitational waves (see Section 1.7).

1.2.1 Millisecond Pulsars

Typical pulsars spin with periods of order \sim seconds, and represent a relatively tame³ contingent of the larger population. If a neutron star is born in a binary system with a stellar companion, the transfer of angular momentum to the pulsar by accretion from the companion can spin a slower, younger pulsar up to periods of \sim 1-10 ms. The eventual result is usually an old, rapidly rotating pulsar in a near-circular orbit with a white dwarf (Bhattacharya & van den Heuvel 1991). During the accretion and circularization process (“recycling”), the magnetic fields of these millisecond pulsars (MSPs) become orders of magnitude weaker. Though \sim 3000 pulsars are known, only \sim 300 of these are MSPs. The vast majority of MSPs are, unsurprisingly, in binary systems; the remaining “isolated” MSPs may have been kicked away from

³Author’s note: *boring*

their companions by a process that is not well understood.

MSPs are responsible for the most compelling science that pulsar timing enables. This is due in large part to their very slow spin-down rates and high timing precision (see Section 1.5), as well as the fact that typical pulsars more frequently exhibit anomalous behavior (such as “glitches” and other noise-inducing processes). Binary MSP systems provide unique opportunities to study general relativistic phenomena such as the strong equivalence principle (Archibald et al. 2019) and Shapiro delay (see Section 1.6).

1.3 Instrumentation for Pulsar Observations

Pulsars are somewhat anomalous in the broader context of radio astronomy, and require correspondingly unique observing considerations. They are weak (no more than a \sim few Jy, and often much less), so maximizing observing sensitivity is critical. Pulsars also have notably steep spectra (a recent paper reports the average spectral index of a >400 -pulsar sample to be -1.6 among those with simple power-law spectra; see Jankowski et al. 2018). In order to capture the short periodicity of MSPs, an instrument’s time resolution must be extremely high ($<100 \mu\text{s}$). Frequency-dependent dispersion as the signal travels through the interstellar medium (see Section 1.4.2) must be dealt with by splitting the observing bandwidth into a sufficiently large number of channels; the bandwidth must also be wide to maximize sensitivity. Lastly, sensitivity is strongly tied to telescope gain, making large-area dishes a requirement for studying most pulsars.

This work deals primarily with observations from the Robert C. Byrd Green Bank Telescope (GBT) in Green Bank, WV⁴, and the Arecibo radio telescope in Arecibo,

⁴<https://greenbankobservatory.org/science/telescopes/gbt/>

Puerto Rico⁵. Both instruments are ideal for studies of pulsars, and have been fundamental to pulsar astrophysics for much of the field’s history. The GBT is a 100-m diameter, fully steerable offset paraboloid dish, with its feed arm and focal point positioned to the side of the instrument to prevent aperture blockage. It is capable of observations in the range of ~ 300 MHz to 100 GHz, though pulsar observations primarily use the 430 MHz, 820 MHz, L-band, and S-band receivers. The main pulsar backends for the instrument are GUPPI (the Green Bank Ultimate Pulsar Processing Instrument, see DuPlain et al. 2008) and more recently, VEGAS (the Versatile GBT Astronomical Spectrometer, see Prestage et al. 2015). The Arecibo telescope, which has been operational for more than 50 years (and was the largest single-dish radio telescope until the construction of FAST, the Five-hundred-meter Aperture Spherical Telescope), is a 306-m stationary spherical reflector. Its frequency coverage spans the ~ 300 MHz - 10 GHz range, with pulsar observations conducted mainly with the 327-MHz, 430-MHz, L-band, and S-band receivers. Its pulsar backend, PUPPI (the Puerto Rican Ultimate Pulsar Processing Instrument), is a copy of the GBT’s GUPPI.

1.4 Searching for (Millisecond) Pulsars

The bright pulses of B1919+21 seen by Jocelyn Bell Burnell are not typical of pulsars; in fact, for the vast majority, individual pulses cannot be seen at all. Because individual pulses are extremely weak, it is necessary to integrate (or “fold”) many of them modulo the period of the pulsar in order to obtain a signal.

Pulsar searching is the art of picking an extremely weak periodic signal out of a sea of much louder instrumental noise, interference, and a number of other detractors. There are myriad motivations for searching for new pulsars. One driver is the desire

⁵<https://www.naic.edu/ao/>

to find exotic systems: additional neutron star-neutron star binaries (as in Hulse & Taylor 1975) or double pulsars (such as J0737–3039, the only such system known; see Burgay et al. 2003); pulsars in orbit with two or more other neutron stars or stellar objects (such as the “triple system;” see Ransom et al. 2014); or even an MSP-black hole binary, which would be a “holy grail” for tests of general relativity.

1.4.1 RFI Mitigation

The same remarkable periodicity that makes pulsars wonderful scientific probes leaves them vulnerable to mimicry by Earth-based signals. Radio frequency interference (RFI) refers to the (usually) human-made signals that contaminate astrophysical observations at radio frequencies. Common RFI sources include communications towers, aircraft, satellites, and local electronic devices (if you ever visit a radio telescope, please take the warnings to turn off your cell phone seriously!).

RFI from these sources is often periodic in nature; therefore, it is critical to excise it prior to searching a dataset. Pulsars are broadband emitters over common observing frequencies, while many RFI sources occur at specific frequencies. This is one easy way to distinguish their signals. Another giveaway is a pulsar’s vulnerability to dispersion by the interstellar medium (see Section 1.4.2). Signals of local origin experience no such frequency-dependent smearing, which is a helpful criterion to consider when judging pulsar candidates. Significant effort has been dedicated to the creation of automatic RFI-mitigating software such as the `rfifind` routine in the `PRESTO` software package⁶ (Ransom 2011), which roots out narrow-band, persistent RFI as well as broadband, short-duration RFI. Some statistical methods (such as spectral kurtosis thresholds and cyclostationary signal processing) have been tested for automatic RFI mitigation, which could prove valuable for large-scale radio telescope pipelines. For

⁶<https://www.cv.nrao.edu/~sransom/presto/>

searches utilizing PRESTO such as those described in Chapter 2, the first step in the data-reduction process is to run `rfifind`. This yields a mask that can be applied to narrow frequency or time ranges in order to avoid overwhelming a potential pulsar signal.

1.4.2 Dispersion

Radio emission from pulsars is subject to the deleterious effects of the interstellar medium (ISM) on its journey towards the Earth. This cold, ionized medium with electron number density n_e causes a time delay in pulse arrival according to frequency (see, e.g., Lorimer & Kramer 2004). The plasma frequency, f_p , and index of refraction of the plasma, μ , are:

$$f_p = \sqrt{\frac{e^2 n_e}{\pi m_e}} \quad (1.2)$$

$$\mu = \sqrt{1 - \left(\frac{f_p}{f}\right)^2} \quad (1.3)$$

where f is the observing frequency. When μ is multiplied by the speed of light (c), one obtains the group velocity of the wave (v_g). The time delay induced by dispersion is:

$$t = \left(\int_0^d \frac{dl}{v_g} \right) - \frac{d}{c} \quad (1.4)$$

where d is the distance between Earth and the pulsar. This expression reduces to:

$$t = \frac{e^2}{2\pi m_e c} \frac{\int_0^d n_e dl}{f^2}. \quad (1.5)$$

We define the numerator of the second portion of the expression to be the “dispersion measure” (DM):

$$\text{DM} = \int_0^d n_e dl, \quad (1.6)$$

which describes the integrated electron density along the line of sight (expressed in the rather bizarre units pc cm^{-3}).

Dispersion is a perpetual frustration in pulsar observing, partially because it is most significant at low frequencies where pulsars are generally strong emitters. A failure to correct for DM causes smearing in the pulse profile, rendering the signal extremely weak upon integration. It is therefore necessary to search over a variety of trial DMs when attempting to detect an unknown pulsar. These trial DMs are informed by models of Galactic electron density (e.g. Cordes & Lazio 2002). Pointings nearer to the Galactic plane will require searching to higher maximum DMs, while those farther off the plane will not. Tools such as `DDplan.py` have been created alongside packages like `PRESTO` in order to optimize one’s de-dispersion routine to maximize computing efficiency while ensuring that pulsars are not being lost between DM trials.

1.4.3 Fourier Transforms and Acceleration Searches

Pulsar observations are time series with hidden periodic signals. The most common approach to parsing them is a Fourier analysis, by which a time series is decomposed into periodic components and the extent to which each one contributes to the overall signal is represented by the frequency’s power. A power vs. frequency plot will often show strong RFI (e.g. the 60-Hz buzz emitted by power grids in the United States), which can be excised by hand with software such as `PRESTO`. Pulsar datasets are composed of discrete, evenly sampled points; therefore, it is necessary to use a




Figure 1.3: Pulse arrival time is inversely dependent on frequency squared. It is necessary to re-align recorded data according to a known DM prior to integrating in order to maximize the observed signal. In searching for new pulsars, many trial DMs must be applied in order to determine the DM that yields the highest-power signal upon integration. Note: the wrapping effect in this plot is an artifact of folding data modulo the pulsar's period. *From Lorimer & Kramer (2004)*

discrete — not continuous — Fourier transform (DFT) to extract their frequency information. For a sequence of N samples, the k^{th} Fourier component is:

$$\mathcal{F}_k = \sum_{j=0}^{N-1} \mathcal{T}_j \exp(-2\pi i j k / N). \quad (1.7)$$

The number of operations necessary to compute a DFT scales as N^2 . Algorithms such as the fast Fourier transform (FFT) contribute considerably to speeding up the transformation process.

Figure 1.4 demonstrates the results of an FFT. When searching for new pulsars, it is preferable to bolster signals in the frequency domain by adding power from subsequent harmonics to the fundamental frequency of interest through *harmonic summing*. This process would be unnecessary if pulsars had sinusoidal pulse profiles; however, their small duty cycles (narrow pulse widths) cause significant power to be pushed to higher harmonics. For example, a pulsar with a duty cycle of 5% will have its power distributed to approximately 20 harmonics above the fundamental.

This thesis deals mainly with MSPs, the vast majority of which are found in binary systems. Because the pulsar’s binary motion will induce a Doppler shift that causes the pulse period to change over time, which subsequently spreads signal into adjacent Fourier bins, it is necessary to perform an *acceleration search* (see Figure 1.5). The `accelsearch` routine in `PRESTO` implements a Fourier-domain correlation technique to “undo” the power-leaking effect. This is particularly important because it avoids performing more than one DFT per timeseries.

For especially highly accelerated binaries (such as those with extremely short orbits of \sim hours), acceleration searches are sometimes insufficient for detecting the pulsar. A *jerk search* has recently been implemented in `PRESTO` in order to extend searches past the constant-acceleration assumption (see Andersen & Ransom 2018).




Figure 1.4: The result of a DFT of a timeseries created during a search for new pulsars is shown. Inspecting this plot of power vs. frequency is important for identifying and masking RFI sources. These plots are not generally directly consulted for the identification of new candidates; rather, harmonic summing and folding of the raw dataset at the proposed period and DM is performed first. One source of RFI — the 60-Hz hum of AC circuits — can be identified in this plot.

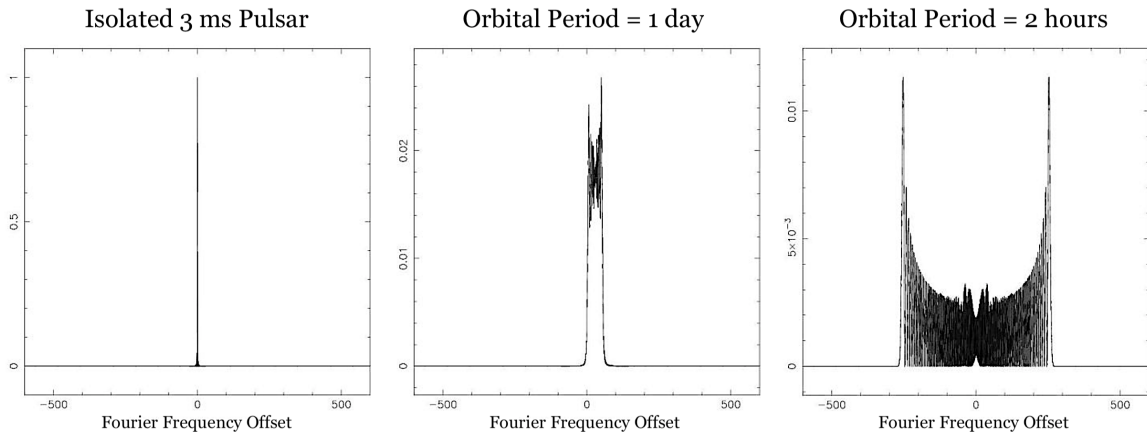


Figure 1.5: In binary MSP systems, Doppler smearing causes power to leak into adjacent Fourier bins. This does not occur in an isolated, 3-ms MSP (left; y-axis is normalized power), but when the MSP is placed in a binary with a 1-day (center) or 2-hour (right) orbital period, the signal dispersion becomes increasingly apparent. An acceleration search serves to pull signal back into the “true” frequency bin at various proposed acceleration values.

Though computationally expensive, this technique has already uncovered previously undetected MSPs.

1.4.4 Inspecting Candidates

The result of the aforementioned steps is a list of candidate signals at various DMs, periods, accelerations, and S/N measurements. A small amount of vetting can be performed automatically by eliminating signals at $DM = 0$ (as they are almost always RFI) and with extremely low S/N. The timeseries of all reasonable candidates are then folded with the given parameters, and inspection by eye ensues (see Figure 1.6). Good candidates are usually sharply peaked with small duty cycles and appear to be broadband emitters over the entire observing frequency range. The few resulting high-quality candidates are then subjected to a folding of their raw data, which is far more computationally intensive. Frequently, seemingly good candidates will expose themselves as RFI after the more comprehensive raw data folding.

Though it is sometimes clear that one has discovered a new pulsar, it is generally necessary to perform follow-up confirmation observations. Unfortunately, diffractive scintillation frequently causes sources to appear extremely bright on some days and to completely disappear on others. For these sources, confirmation can be a frustrating process. Following confirmation, if observing time is available, the process of timing the pulsar is the next step (see Section 1.5).

1.5 Pulsar Timing

Pulsar timing is the process of creating a model that accounts for *every single* pulse from a pulsar over long periods of time. Owing both to pulsars’ extremely stable spin periods and also to the fact that our models track *integer* numbers of pulses, pulsar timing has yielded some of the most precise measurements in astrophysics. The description of pulsars as “clock-like” is perhaps overused — though for a good reason. The extremely slow spin-down of MSPs, for example, means that our timing measurements rival the precision of atomic clocks.

The currency of pulsar timing is the “time of arrival” (TOA) measurement. For a standard pulsar (that is, one whose individual pulses cannot be detected), many hundreds or thousands of pulses are added together to create an integrated pulse profile. This profile is then matched with some nominal noise-free template, which is often derived from many observations of a pulsar at a given frequency. TOAs must be time-tagged using a clock (often a hydrogen maser) local to the observatory and GPS measurements. In order to ensure consistency between observatories and to nullify any relativistic effects induced by solar system bodies, TOAs are converted from the topocentric reference frame to a reference frame tied to the solar system barycenter.

If ν is the spin frequency of a pulsar, the pulse number N compared to some

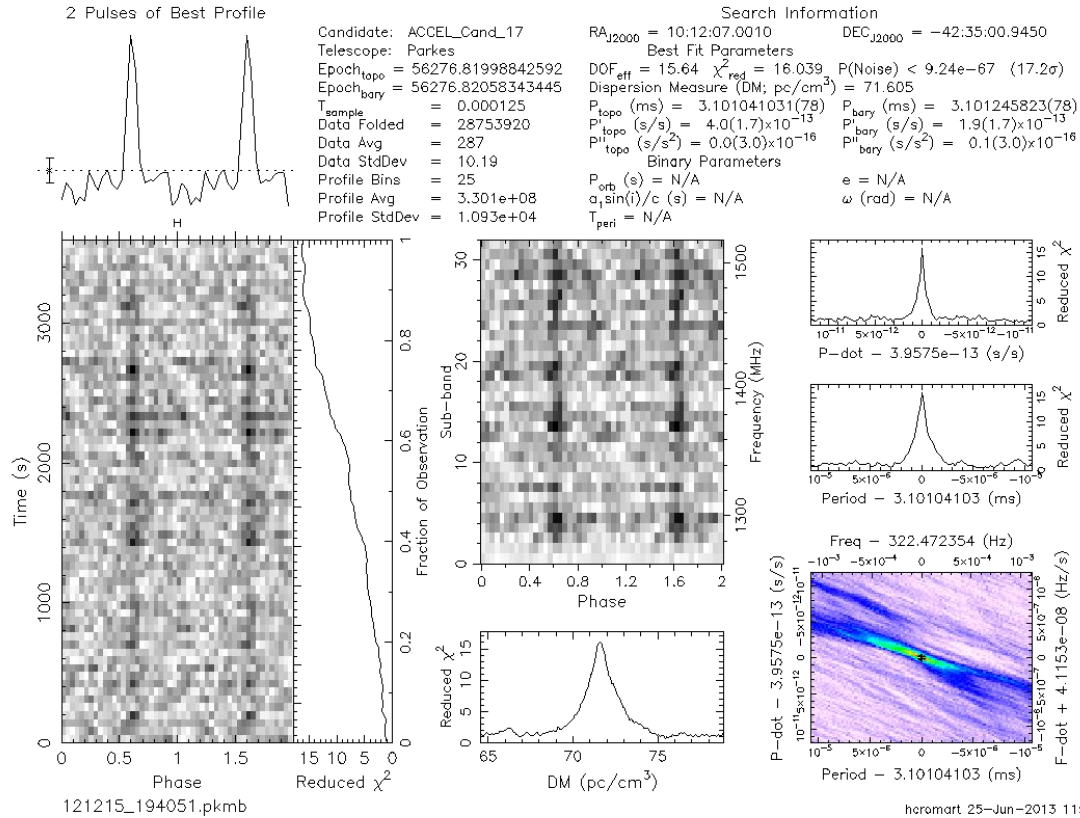


Figure 1.6: A single search can yield thousands of candidate pulsars. The highest S/N signals with nonzero DM are folded (in PRESTO, using a routine like `prepfold`) to produce a plot such as this one. A good candidate will present with several different features. The pulse profile (upper left) will be usually be sharp with a small duty cycle. The χ^2 of the signal (right of the time vs. phase plot) should build steadily over the course of the observation. A true pulsar, as opposed to RFI, will be broadband in nature (center). A peak should appear around the proposed DM (bottom center).

reference epoch t_0 can be derived from a Taylor expansion of the spin frequency:

$$N = N_0 + \nu_0(t - t_0) + \frac{1}{2}\dot{\nu}(t - t_0)^2 + \frac{1}{6}\ddot{\nu}(t - t_0)^3 + \dots \quad (1.8)$$

Assuming t_0 denotes the arrival time of a pulse, N is constrained to be an integer. This is the underpinning of all of pulsar timing. It is valuable to note that $\ddot{\nu}$ is extremely difficult to measure in pulsars that are not very appreciably spinning down.

The precision with which a TOA can be measured is dependent on the pulsar's spin period (P) and duty cycle (δ), the pulsar's flux density (S_{mean}) and the observing system's equivalent flux density (S_{sys}), and the integration time and observing bandwidth (t_{obs} and Δf , respectively):

$$\sigma_{TOA} \propto \frac{S_{\text{sys}}}{\sqrt{t_{\text{obs}}\Delta f}} \times \frac{P\delta^{3/2}}{S_{\text{mean}}}. \quad (1.9)$$

This highlights both the value of MSPs for timing (low P) and the importance of modern low-noise, wide-bandwidth receivers and backends. MSPs also yield more precise timing measurements than slow pulsars because one can integrate over many more pulses in a given amount of time, yielding a higher-quality pulse profile.

Timing models begin as a meager set of parameters derived from a pulsar's discovery observation (which hints at the spin period, position, and DM of the source). The addition of observations to one's dataset necessitates the use of new parameters such as parallax, spin-down rate, or proper motion (among many others). Timing binary pulsars is considerably more complicated, as it requires (at minimum) the addition of five Keplerian parameters to describe the system's orbital properties: P_b , the orbital period; projected semi-major axis $a \sin(i)$; orbital eccentricity e ; longitude of periastron ω ; and epoch of periastron passage T_0 . In order to time a binary pulsar, it




Figure 1.7: Binary motion induces a change in observed spin period over the course of an observation. In order to solve for the pulsar’s orbital parameters, it is therefore necessary to observe the system several times in quick succession (red sections) and fit for the orbit (black line) before continuing to time the pulsar. *Figure courtesy of Scott Ransom*

is therefore necessary to solve the system’s orbit with a series of closely spaced observations before proceeding with timing (see Figure 1.7). Section 1.6 will explore the addition of post-Keplerian parameters to one’s timing model to account for general relativistic effects.

The difference between a measured TOA and the time predicted by one’s timing model is referred to as a *timing residual*. The failure of a model to account for some parameter will induce structure in the resulting timing residuals (for example, neglecting to fit for parallax causes a growing sinusoidal structure over time). Fitting a timing model can be achieved by χ^2 minimization with a package such as **TEMPO** or

TEMPO2. More recently, a completely independent timing package called PINT (PINT is not TEMPO3; see the forthcoming paper Luo et al. 2020) has been developed by Jing Luo, Scott Ransom, Paul Ray, Paul Demorest, and many others. It serves as a check on timing performed with TEMPO, which has been the primary source of timing models for decades. PINT features both χ^2 reduction and Markov Chain Monte Carlo (MCMC) fitting capabilities (see Chapter 4).

1.6 Relativistic Shapiro Delay & the Neutron Star Equation of State

In a subset of binary MSPs, it becomes necessary to include post-Keplerian (general relativistic) parameters in one’s timing model. These include: $\dot{\omega}$, the rate of periastron advance; \dot{P}_b , the orbital decay rate; and relativistic γ , which includes gravitational redshift and transverse Doppler shift. The final two — the Shapiro delay range (r) and shape (s) parameters — are of particular relevance to this work. The measurement of two post-Keplerian parameters will break the degeneracy between the pulsar and companion masses, facilitating an independent measurement of each.

Relativistic Shapiro delay is a small delay in pulse TOAs (of order $\sim 10 \mu\text{s}$) occurring during orbital conjunction, induced by the curvature of spacetime near an MSP’s companion star (it occurs in other contexts; see Shapiro 1964). In general relativity, Shapiro delay can be characterized by two parameters: “shape” (s) and “range” (r):

$$r = T_{\odot} m_c \tag{1.10}$$

$$s = \sin(i) = x \left(\frac{P_b}{2\pi} \right)^{-2/3} T_\odot^{-1/3} M^{2/3} m_c^{-1} \quad (1.11)$$

where $T_\odot = GM_\odot/c^3$ (the Sun’s mass in units of time, $\sim 4.9 \mu\text{s}$), M is the system’s combined mass, and m_c is the companion’s mass. When combined with the Keplerian mass function, measurements of s and r also constrain the pulsar mass (m_p ; see Freire & Wex 2010):

$$f(m_p, m_c) = \frac{4\pi^2 (a \sin i)^3}{G P_b^2} = \frac{(m_c \sin i)^3}{M^2} \quad (1.12)$$

Shapiro delay manifests as a residual spike near orbital conjunction (see Figure 1.8). It is extremely sensitive to the inclination angle of the orbit (i), and is therefore only rarely measured.

Pulsar timing observations that probe relativistic Shapiro delay have yielded some of the best measurements of neutron star masses to date. They are critical for probing the neutron star interior equation of state (EoS), as the discovery of more and more massive neutron stars rules out each EoS that predicts collapse at a lower mass (for example, see Figure 3 in Watts et al. 2015). The EoS of cold supranuclear-density matter is poorly understood, partially because it is impossible to study in an Earth-based laboratory. Understanding it is one of the most puzzling and important problems in modern nuclear physics.

In 2010, Demorest et al. forced a reconsideration of the canonical $1.4\text{-}M_\odot$ neutron star with their Shapiro delay-enabled mass measurement of J1614–2230, the first two-solar-mass neutron star ever observed (Demorest et al. 2010). Then, in 2013, Antoniadis et al. used optical techniques and pulsar timing to yield a mass measurement of $2.01 \pm 0.04 M_\odot$ for the pulsar J0348+0432 (Antoniadis et al. 2013). Chapter 3 describes the recent measurement of the most massive ($2.14_{-0.09}^{+0.10} M_\odot$) neutron star to date (Cromartie et al. 2020). These three measurements have helped significantly con-

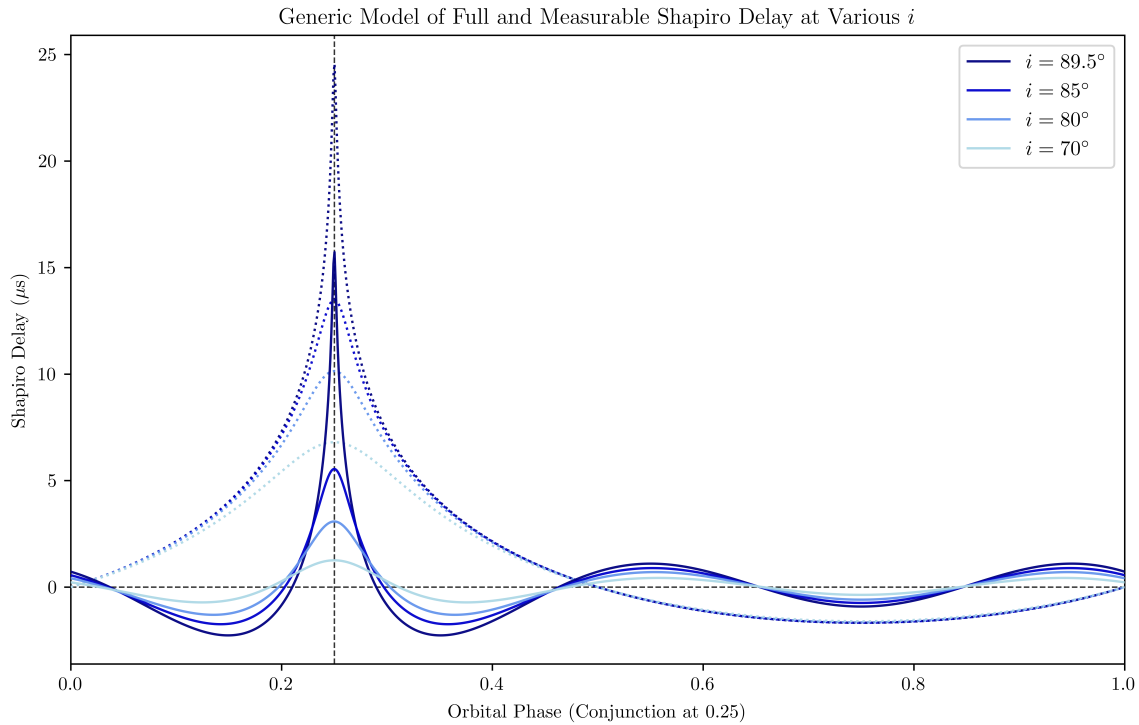


Figure 1.8: Shapiro delay occurs at orbital conjunction, which is defined as occurring when orbital phase = 0.25. The effect is extremely sensitive to inclination angle (i), and often induces a delay of no more than $\sim 10 \mu\text{s}$. The *measurable* Shapiro delay (that is, the signal to which we are actually sensitive without the *a priori* information about m_c and i) is shown with solid lines. The full extent of the Shapiro delay signal is denoted with dotted lines.

strain the EoS of supranuclear-density matter; for example, the “softer” non-nucleonic EoS (e.g. quark matter, hyperons, or Bose-Einstein meson condensates) have become increasingly difficult to justify in the presence of a 2.14-solar-mass neutron star (see Özel & Freire 2016).

1.7 Pulsar Timing Arrays & NANOGrav

As discussed in Section 1.5, MSPs can be timed with clock-like precision. Deviations in one’s timing model can highlight a number of intervening phenomena, such as DM variations and general relativistic effects. By measuring correlations in timing residuals between many pairs of MSPs, it becomes possible to determine if some portion of the measured delays are attributable to the changes in path length (between MSPs and the Earth, as well as between MSP pairs) induced by gravitational waves (Arzoumanian et al. 2018; Burke-Spolaor et al. 2019). These detectors are called Pulsar Timing Arrays (PTAs).

PTAs are similar, at least in principle, to experiments such as LIGO or the upcoming NASA-supported Laser Interferometer Space Antenna (LISA). Each line of sight to an MSP in a PTA is comparable to one of the arms of these interferometers. PTAs benefit from greater flexibility than these physical detectors; whereas LIGO and LISA would need to construct additional mirrors to improve their gravitational wave sensitivity, PTAs can (relatively) easily add new MSPs to their pulsar timing programs.

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) collaboration, of which the author of this dissertation is a member, conducts high-precision ($\leq \mu\text{s}$) timing of 70+ Galactic radio MSPs as part of its PTA. Adding to the 70+ NANOGrav PTA sources is critical to increasing the probability and signif-

ificance of a gravitational wave detection, as the S/N of a gravitational wave signal increases linearly with the number of MSPs in the array (Siemens et al. 2013). It is also necessary to include MSPs that are widely distributed across the sky, creating different detection baselines and increasing the significance of any quadrupolar correlation found in the timing residuals that is attributable to gravitational waves (Cornish & Sampson 2016). It is NANOGrav’s goal to continuously increase the current number of MSPs observed, and to time all of them to better than $1\text{-}\mu\text{s}$ accuracy. Though the collaboration’s first detection of gravitational waves is likely to be the isotropic stochastic background (resulting from an ensemble of supermassive black hole binaries), we hope to develop sensitivity to continuous wave (single) sources over time.

NANOGrav’s timing program includes \sim monthly observations from the GBT and Arecibo radio telescopes, and more recently, observations from the Karl G. Jansky Very Large Array (VLA) in New Mexico and the Canadian Hydrogen Intensity Mapping Experiment (CHIME) radio telescope at the Dominion Radio Astrophysical Observatory in Canada. Each MSP in the array is timed for 20-30 minutes, usually producing one or two TOAs per epoch.

The International Pulsar Timing Array (IPTA) comprises NANOGrav and several similar collaborations: the Parkes Pulsar Timing Array in Australia, the European Pulsar Timing Array, and more recently, South African, Indian, and Chinese PTA collaborations. The IPTA’s goal is to combine PTA observations from all of its constituent observatories in order to more swiftly and significantly detect the gravitational wave background. This process is not without complications, however. Though one would expect superior limits on the gravitational wave background amplitude from a combined data set, the act of combining data from so many disparate observatories has proven extremely difficult.

1.8 Pulsar Observations at High Energies

Because they are steep-spectrum astrophysical sources, MSPs have traditionally been searched for and timed at radio frequencies. Recently, however, instruments such as the *Fermi* Large Area Telescope (LAT) and the Neutron Star Interior Composition Explorer (NICER; see Chapter 4) have revolutionized our understanding of MSPs at Gamma and X-ray frequencies.

When the *Fermi* satellite launched in 2008 with the LAT in tow, astronomers were not counting on significant Gamma-ray emission from MSPs (Caraveo 2014). The *Fermi* Gamma-ray observation of PSR J0030+0451 — the first high-sigma detection of of an MSP in Gamma-rays — was the harbinger of a new era in high-energy pulsar astrophysics (Abdo et al. 2009).

Fermi's LAT yields a continually updated catalog of thousands of Gamma-ray point sources of unknown astrophysical origin. Poring over these “unassociated” sources enables the selection of candidates with MSP-like spectral and variability characteristics. Though gamma-emitting MSPs abound, most emit so few high-energy photons that it is nearly impossible to blindly search for them at high energies with traditional techniques. The small positional uncertainties of many *Fermi*-LAT sources facilitate single-pointing follow-up observations with large radio telescopes.

At high energies, photons emitted by a single pulsar can come as infrequently as once every few days. It is therefore near-impossible to derive a timing solution from high-energy photons alone. The best method for dealing with this scarcity is through single-photon timing. In this framework, one is given a list of photon events with a time of arrival, energy, and weight between 0 and 1 (which predicts the likelihood that a photon is actually from that pulsar). By creating a smoothed pulse profile that acts as a probability density function (PDF) as a function of phase (ϕ), we predict

ϕ for each event and take all TOAs and phases from our timing software. We then compute an overall likelihood, and repeat (thereby marginalizing over phase). The resulting peak is our returned value. The log likelihood is defined by Pletsch & Clark (2015):

$$\log \mathcal{L}(\mathbf{u}) = \sum_{j=1}^N \log [w_j F(\phi_j(t_j, \mathbf{u})) + (1 - w_j)] \quad (1.13)$$

where F is the pulse profile, ϕ is the rotational phase, w is the weight, \mathbf{u} is the matrix of timing parameters, and t is the TOA. The $1 - w_j$ term comes from the added complication that each photon may not come from the MSP itself.

This Bayesian single-photon timing routine is implemented in PINT (as `event_optimize`). The result of a timing run with `event_optimize` is shown in Figure 1.9.

1.9 Overview

This is a dissertation in four main parts:

- Chapter 2 describes the discovery of the first six *Fermi* MSPs discovered with the Arecibo telescope (Cromartie et al. 2016).
- Chapter 3 relates the discovery of the most massive neutron star measured to date, the NANOGrav MSP J0740+6620 (Cromartie et al. 2020).
- Chapter 4 describes the measurement of radio Shapiro delay in the bright Gamma-ray MSP J1231–1411, as well as its timing with the *Fermi*-LAT instrument over ~ 12 years (Cromartie et al. 2020b, in preparation).
- Chapter 5, *Miscellanea*, is more informal and personal than the preceding chapters. It details several minor and in-progress research projects that have been pursued alongside the rest of this thesis.

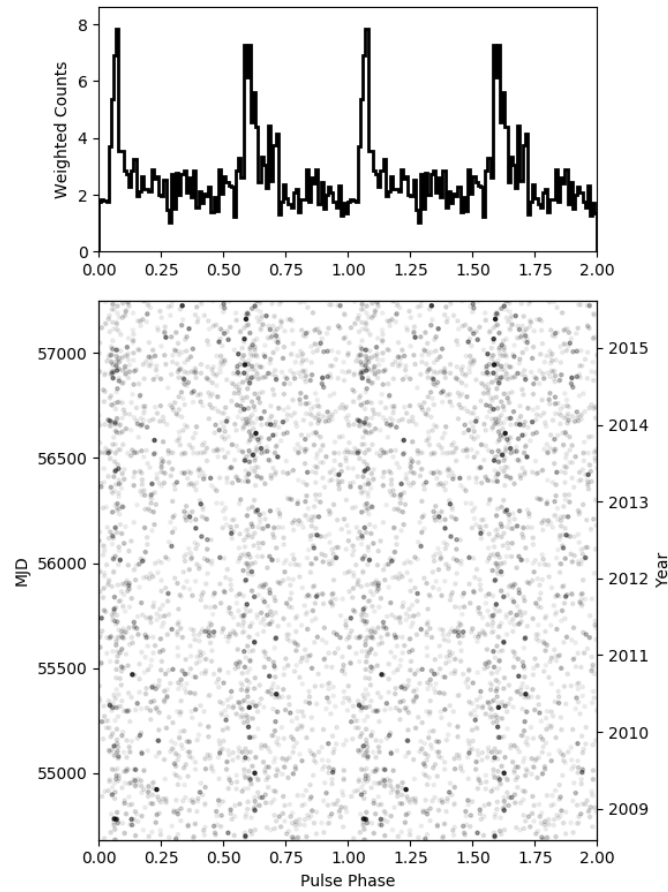


Figure 1.9: This figure presents *Fermi*-LAT Gamma-ray photons from the pulsar B1937+21 between 2008 and \sim 2016, timed using the Bayesian single-photon timing routine `event_optimize` in PINT. See Chapter 4 for a more thorough example of single-photon timing methodology (with a high S/N result).

- We offer summarizing and concluding remarks in Chapter 6.

1.10 Collaborators

In the modern era, high-quality science is rarely performed alone; therefore, it is critical to recognize the collaborators who have made this work possible. As a collaboration of more than 100 scientists, no part of NANOGrav can function in a vacuum. The author owes her thanks to every NANOGrav-associated scientist, and especially those in the Timing and Searching working groups. In no particular order, the following collaborators have been indispensable to the completion of this dissertation: Scott Ransom (NRAO), Fernando Camilo (MeerKAT), Julia Deneva (NRL), Elizabeth Ferrara (NASA/UMD), Paul Ray (NRL), Matthew Kerr (NRL), Emmanuel Fonseca (McGill), David Nice (Lafayette College), Megan DeCesar (NRL), Joe Swiggum (Lafayette College), Paul Demorest (NRAO), Tim Pennucci (Eötvös Loránd University), Ryan Lynch (GBT), and Dusty Madison (WVU).

Chapter 2

Six New Millisecond Pulsars in *Fermi* Gamma-Ray

Sources

2.1 The First Six Arecibo *Fermi* MSPs

The text in section 2.1 appeared in the *Astrophysical Journal* (Cromartie et al. 2016) as “*Six New Millisecond Pulsars from Arecibo Searches of Fermi Gamma-Ray Sources.*”

Abstract

We have discovered six radio millisecond pulsars (MSPs) in a search with the Arecibo telescope of 34 unidentified gamma-ray sources from the *Fermi* Large Area Telescope (LAT) 4-year point source catalog. Among the 34 sources, we also detected two MSPs previously discovered elsewhere. Each source was observed at a center frequency of 327 MHz, typically at three epochs with individual integration times of 15 minutes. The new MSP spin periods range from 1.99 to 4.66 ms. Five of the six pulsars are in interacting compact binaries (period ≤ 8.1 hr), while the sixth is a more typical neutron star-white dwarf binary with an 83-day orbital period. This is a higher proportion of interacting binaries than for equivalent *Fermi*-LAT searches elsewhere. The reason is that Arecibo’s large gain afforded us the opportunity to limit integration times to 15 minutes, which significantly increased our sensitivity to these highly accelerated systems. Seventeen of the remaining 26 gamma-ray sources are still categorized as strong MSP candidates, and will be re-searched.

2.1.1 Introduction

Of the 230 millisecond pulsars (MSPs) currently known in the Galactic disk¹, 30% have been discovered in previously unidentified sources of gamma rays detected by the

¹Public list of Galactic MSPs: <http://astro.phys.wvu.edu/GalacticMSPs/>

Fermi-LAT instrument² (Atwood et al. 2009). While only around 10% of all known pulsars rotate at millisecond rates, MSPs make up half of all pulsars observed to emit gamma rays (Caraveo 2014). The LAT source catalogs have been instrumental in the search for new MSPs, providing spectral data to aid in distinguishing possible MSPs from other gamma-ray-emitting objects, such as active galactic nuclei (AGNs). Once an MSP is discovered in a radio search and a phase-connected timing solution is available, the sparse gamma-ray photons are folded using the radio ephemeris in order to glimpse gamma-ray pulsations (e.g., Cognard et al. 2011). While it was possible to search for radio pulsars in gamma-ray sources prior to the *Fermi* era (e.g., Champion et al. 2005; Roberts 2013), the small positional uncertainty of LAT gamma-ray sources has enabled single-pointing radio searches. Overall, the search for MSPs in the Galactic disk has been made extremely efficient by employing *Fermi*-LAT data in selecting radio search targets.

Before 2013, no *Fermi* MSPs had been discovered using the 305-m Arecibo radio telescope in Puerto Rico. In contrast, the Green Bank (GBT), Parkes, Nançay, Giant Metrewave (GMRT), and Effelsberg telescopes had been used to discover dozens of new MSPs using the 1FGL and 2FGL catalogs (Abdo et al. 2010; Nolan et al. 2012) as guides. In this work, we present the first six MSPs discovered in unidentified LAT sources using the Arecibo telescope, along with preliminary orbital parameters gleaned from radio timing. We then quantitatively discuss the relative sensitivities — both in the flux density and acceleration regimes — between the *Fermi*-LAT MSP searches conducted at Arecibo and those done at the GBT and Parkes.

²For a list of all LAT pulsars, see <https://confluence.slac.stanford.edu/display/GLAMCOG/Public+List+of+LAT-Detected+Gamma-Ray+Pulsars>

2.1.2 Candidate Selection

All but two of the 34 candidates observed came directly from early versions of the third *Fermi*-LAT catalog (3FGL, also known as the 4-year catalog), which was later published by Acero et al. (2015). The two remaining sources were detected with the LAT but were below the significance threshold required for inclusion in the final catalog. While LAT source lists contain >1000 unidentified gamma-ray emitters, several constraints dramatically limit the number of sources appropriate for our searches.

Every source had to be located within the declination range of Arecibo ($-1^\circ < \delta < 39^\circ$). Justification for picking the 327-MHz receiver over L-band, for example, was two-fold. First, the target source error circles were required to fit within Arecibo's beam, allowing for single, long-duration pointings. The 327-MHz system, with a relatively large FWHM = $15'$, was the best choice. Second, pulsars have steep spectra, and therefore are brighter at such a relatively low frequency. Very few pulsars are known to be variable in gamma rays (Ray et al. 2012); thus, only non-variable LAT sources were selected. Also, each of the selected sources had a spectral energy distribution consistent with those of known gamma-ray pulsars, which typically have exponentially cut-off power-law spectra (Abdo et al. 2013). Because it is difficult to characterize LAT sources amid the Galaxy's diffuse gamma-ray background (see, e.g., Geringer-Sameth & Koushiappas 2012), and because the effects of dispersion, scattering, and synchrotron emission inhibit radio pulsar observations at low frequency along the Galactic plane, only sources with $|b| > 5^\circ$ were considered. After whittling the list down, we obtained the 34 sources in Table 2.1.

Table 2.1: Summary of Arecibo Searches of Unidentified *Fermi*-LAT Sources

Name	R.A. ^a (J2000.0)	Decl. ^a (J2000.0)	<i>l</i> (deg)	<i>b</i> (deg)	Integration Time (minutes)	DM _{max} ^b (pc cm ⁻³)
3FGL J0103.7+1323	01 ^h 03 ^m 46 ^s	13°23'33"	127.5	-49.4	15, 15, 15	72
3FGL J0134.5+2638	01 ^h 34 ^m 31 ^s	26°38'15"	134.7	-35.2	15, 15, 15	92
3FGL J0232.9+2606	02 ^h 32 ^m 56 ^s	26°06'13"	149.7	-31.4	15, 15, 15	130
3FGL J0251.1+2603	02 ^h 51 ^m 08 ^s	26°04'48"	153.9	-29.5	15, 15	124
3FGL J0318.1+0252	03 ^h 18 ^m 09 ^s	02°52'10"	178.4	-43.6	15, 15, 15	900
3FGL J0330.6+0437	03 ^h 30 ^m 40 ^s	04°37'32"	179.5	-40.1	15, 15, 15	680
3FGL J0342.3+3148c	03 ^h 42 ^m 18 ^s	31°48'33"	160.3	-18.4	15, 15, 15	100
3FGL J0421.6+1950	04 ^h 21 ^m 37 ^s	19°50'49"	175.9	-20.7	15, 15, 15	150
3FGL J0517.1+2628c	05 ^h 17 ^m 10 ^s	26°28'44"	178.6	-6.6	10, 15, 15, 15, 15	120
3FGL J0539.8+1434	05 ^h 39 ^m 48 ^s	14°33'53"	191.6	-8.6	5, 5	200
3FGL J1048.6+2338	10 ^h 48 ^m 41 ^s	23°38'29"	213.2	62.1	15, 15	66
3FGL J1049.7+1548	10 ^h 49 ^m 44 ^s	15°48'25"	228.5	59.6	10, 15	68
3FGL J1200.4+0202	12 ^h 00 ^m 27 ^s	02°02'31"	274.8	62.1	15, 15, 15	64
3FGL J1225.9+2953	12 ^h 25 ^m 59 ^s	29°53'25"	185.2	83.8	15, 15, 15	40
P7R4 J1250+3118 ^e	12 ^h 50 ^m 52 ^s	31°18'18"	124.6	85.8	15, 15, 15	40
3FGL J1309.0+0347	13 ^h 09 ^m 02 ^s	03°47'27"	313.9	66.3	15, 15, 10, 15, 10, 15	60
3FGL J1322.3+0839	13 ^h 22 ^m 20 ^s	08°39'27"	325.9	70.1	15, 15, 15	52
3FGL J1601.9+2306	16 ^h 01 ^m 57 ^s	23°06'39"	38.5	46.9	15, 15, 15	60
3FGL J1627.8+3217	16 ^h 27 ^m 52 ^s	32°17'56"	53.0	43.2	15, 15	70
3FGL J1704.1+1234	17 ^h 04 ^m 08 ^s	12°34'25"	32.5	29.4	15, 15, 15	116
3FGL J1720.7+0711	17 ^h 20 ^m 46 ^s	07°11'21"	29.0	23.4	15, 15, 15	156
3FGL J1805.9+0614	18 ^h 05 ^m 55 ^s	06°14'15"	33.4	13.0	15, 15	304
3FGL J1824.0+1017	18 ^h 24 ^m 05 ^s	10°17'27"	39.1	10.7	15, 15	356

Table 2.1 continued

Table 2.1: Summary of Arecibo Searches of Unidentified *Fermi*-LAT Sources (continued)

Name	R.A. ^a (J2000.0)	Decl. ^a (J2000.0)	<i>l</i> (deg)	<i>b</i> (deg)	Integration Time (minutes)	DM _{max} ^b (pc cm ⁻³)
3FGL J1827.7+1141	18 ^h 27 ^m 42 ^s	11°41′50″	40.8	10.5	15, 15	356
3FGL J1829.2+3229	18 ^h 29 ^m 08 ^s	32°30′42″	60.7	18.5	15, 15, 15	158
3FGL J1842.2+2742	18 ^h 42 ^m 15 ^s	27°42′09″	57.1	14.1	15, 15, 15	216
P7R4 J1909+2102^e	19 ^h 09 ^m 32 ^s	21°02′56″	53.7	5.6	15, 15	564
3FGL J1921.2+0136^d	19 ^h 21 ^m 14 ^s	01°36′26″	37.8	-5.9	5, 10, 15, 15	670
3FGL J2026.3+1430	20 ^h 26 ^m 21 ^s	14°30′53″	57.3	-13.4	15, 15, 15	226
3FGL J2042.1+0247^d	20 ^h 42 ^m 09 ^s	02°47′35″	49.0	-23.0	15, 15	140
3FGL J2052.7+1217	20 ^h 52 ^m 47 ^s	12°17′51″	59.1	-20.0	15, 15, 15	148
3FGL J2108.0+3654	21 ^h 08 ^m 02 ^s	36°55′19″	81.1	-7.2	15, 15, 15	360
3FGL J2212.5+0703	22 ^h 12 ^m 35 ^s	07°03′35″	68.7	-38.6	15, 30, 15, 15, 15, 5	84
3FGL J2352.0+1752	23 ^h 52 ^m 04 ^s	17°52′50″	103.5	-42.7	15, 15, 15, 15	74

NOTE — Boldfaced entries denote observations yielding MSP detections.

^a Arecibo telescope pointing position.

^b The maximum dispersion measure (DM) up to which we searched corresponds approximately to twice the maximum value predicted for each line of sight by the NE2001 electron density model (Cordes & Lazio 2002), with the exception of 3FGL J0318.1+0252 and J0330.6+0437, which were unintentionally searched to higher DMs.

^d Discovered at theGBT (S. Sanpa-Arsa et al. 2016, in preparation).

^e Source not included in 3FGL catalog.

2.1.3 Observations and Data Analysis

Observations of the 34 *Fermi*-LAT sources were conducted during 12 sessions between 2013 June and September using the Arecibo telescope. In order to combat the effects of scintillation, orbital acceleration, and eclipses (discussed further below), we aimed to observe each source three times for 15 minutes per pointing, though the exact number of pointings per source changed as data were analyzed.

Sessions in the months of June, July and early August were conducted in-person at the observatory, while later sessions occurred remotely. In either case, the standard CIMA³ telescope control software was used in conjunction with command-line control of the Puertorican Ultimate Pulsar Processing Instrument (PUPPI). The PUPPI backend (a replica of GUPPI⁴ at the GBT) was configured for the settings shown in Table 2.2. All data were taken in total-intensity, summed polarization mode. Once disks became full, they were shipped from Arecibo to Columbia University for data reduction. A summary of all observations is provided in Table 2.1.

Data were analyzed using the software package PRESTO (Ransom 2001). The data reduction process began with the detection and masking of significant radio frequency interference in the data. Dedispersion occurred up to a specified dispersion measure (DM), which we chose to be twice the maximum line-of-sight value given by the NE2001 model (Cordes & Lazio 2002). PRESTO can perform searches over spin period variations caused by orbital motion, searching over both period and period derivative. The extent of the acceleration search is specified by the `zmax` parameter, chosen to be 200 in our case. This means that linear pulsar spin frequency (f_0) drifts of up to 200 bins were searched in the highest harmonic, which in our analysis was the eighth (Ransom et al. 2002). If t_{int} is the total integration time and a_{max} is the maximum

³<http://www.naic.edu/~cima>

⁴<https://safe.nrao.edu/wiki/bin/view/CICADA/GUPPIUsersGuide>

Table 2.2: Observing Parameters for Four Radio Surveys of *Fermi*-LAT Sources

Parameter	AO 327-MHz Survey ^a	GBT 350-MHz Survey ^a	GBT 820-MHz Survey ^b	ParkeS 1390-MHz Survey ^c
Detection fraction ^d	8/34 (24%); 5/6	13/50 (26%); 3/13	3/25 (12%); 1/3	11/56 (20%); 2/11
Center frequency (MHz)	327	350	820	1390
Bandwidth, Δf (MHz)	68.75 ^e	100	200	256
Number of channels	2816 ^f	4096	2048	512
Sample time (μ s)	81.92	81.92	61.44	125
Receiver temperature ^g , T_{rec} (K)	62	20	18	25
Average sky temperature ^h , T_{sky} (K)	64	65	15	5
Telescope gain, G (K/Jy)	10	2	2	0.735
Effective threshold S/N ($\beta \cdot S/N_{\text{min}}$)	10	10	10	12
Integration time, t_{int} (minutes)	15	32	45	60

^a Hessels et al. (2011).^b Ransom et al. (2011).^c Kerr et al. (2012); Camilo et al. (2015).^d Number of *detected* MSPs divided by the total number of sources *observed*; number of black widow plus redback systems *discovered* divided by the total number of MSPs *discovered*.^e Recorded bandwidth: 100 MHz were sampled by PUPPI but we only recorded the section covering the receiver bandwidth.^f Number of recorded channels; 4096 channels were sampled across the entire 100-MHz bandwidth.^g Receiver temperature including spillover but excluding Galactic/CMB contribution. Values for Arecibo 327 MHz from NAIC⁵. GBT values are from page 11 of the proposer's guide⁶. See also Lynch et al. (2013) for 350 MHz. ParkeS values are based on those from the users' guide⁷.^h We calculated sky temperatures by scaling the Haslam et al. (1981) 408-MHz map to each survey's observing frequency using a spectral index of -2.6 (Lawson et al. 1987). For AO 327, we list the average T_{sky} for each of the 34 target locations. For GBT 350 and ParkeS, we calculated the average temperature at an evenly spaced grid of points encompassing the search regions. For GBT 820 we averaged the Ransom et al. (2011) values for their searches excluding the Galactic plane.

acceleration probed, $z_{\max} = a_{\max} t_{\text{int}}^2 f_0 / c$ (Ransom et al. 2001).

Sensitivity: Figure 2.1 shows the minimum flux density detectable by our Arecibo searches for a range of MSP spin periods and DMs, as determined by the radiometer equation for pulsars (Lorimer & Kramer 2004 Appendix 1.4):

$$S_{\min} = \beta \frac{S/N_{\min} (T_{\text{rec}} + T_{\text{sky}})}{G \sqrt{n_p t_{\text{int}} \Delta f}} \sqrt{\frac{W}{P - W}}, \quad (2.1)$$

where S_{\min} is the minimum detectable flux density, β is a normalization factor including corrections for, e.g., system digitization losses, S/N_{\min} is the threshold pulsar signal-to-noise ratio, T_{rec} is the receiver temperature (including contributions from the CMB and spillover), T_{sky} is the sky temperature, W is the effective pulse width (we assume the intrinsic pulse width to be $P/10$), P is the pulsar period, G is the telescope gain, n_p is the number of polarizations recorded (always 2 for the searches discussed here), t_{int} is the integration time, and Δf is the effective bandwidth. Relevant parameters for the Arecibo survey are shown in Table 2.2 under “AO 327”. Arecibo’s system equivalent flux density (SEFD) degrades for zenith angles exceeding 15 degrees⁸, but this had little impact for most of our searches. Overall our survey had an average SEFD = 13 Jy. The sensitivity of the Arecibo survey in the context of other *Fermi*-LAT searches is discussed in Section 2.1.5.

2.1.4 Results

In the 34 sources searched, we discovered six new MSPs (see Table 2.3). Pulse profiles from the discovery observations are shown in Figure 2.2. The rotation periods range between 1.99 and 4.66 ms, and their DMs span 17–65 pc cm⁻³. Orbital

⁸Detailed measurements for gain and system temperature of the 327-MHz Gregorian receiver at Arecibo can be found at <http://www.naic.edu/~astro/RXstatus/327/327greg.shtml>.




Figure 2.1: Sensitivity of four radio pulsar searches of *Fermi*-LAT sources as a function of spin period for a pulsar with 10% duty cycle. All surveys have been scaled to 327 MHz using a spectral index of -1.7 . The surveys presented here, in addition to this work, are: GBT 350 MHz (Hessels et al. 2011), GBT 820 MHz (Ransom et al. 2011), and Parkes 1390 MHz (Camilo et al. 2015). See Table 2.2 for details.

solutions have been obtained for all new discoveries from initial timing observations. However, phase-connected timing solutions (including precise positions, period derivatives, spin-down luminosities, and a study of the gamma-ray properties of the coincident gamma-ray sources) are not yet available for most of the MSPs, and will be presented elsewhere. A study of the redback PSR J1048+2339 is presented in Deneva et al. (2016).

Five of the new MSPs are neutron stars with short orbital periods. Three are “black widows”, in which much of the companion mass has been stripped away or accreted by the pulsar, leaving a (partially degenerate) companion with mass $\ll 0.1 M_{\odot}$. The remaining two short-orbit systems are “redbacks”, where the pulsar is frequently eclipsed by outflows from a non-degenerate companion with mass $\gg 0.1 M_{\odot}$. The final MSP is a more classical neutron star-white dwarf binary. For a diagram of orbital period vs. companion mass for such highly accelerated systems, see Roberts (2013).

Figure 2.3 shows the distribution of orbital periods vs. minimum companion masses for five of the new MSPs presented in this work. Minimum companion masses are calculated using Keplerian parameters derived from orbital timing solutions. Note that PSR J1909+21 is classified as a redback, despite its minimum companion mass being less than $0.1 M_{\odot}$. This is firstly because the $0.055 M_{\odot}$ value is the *minimum* companion mass, and it is larger than any known black widow minimum companion mass; secondly, its eclipses last for approximately half of the orbit, which is characteristic of a redback system with a dense circumstellar environment.

PSR J1805+06 is in a black widow system with an orbital period of 8.1 hr. The approximate position of 3FGL J1805.9+0614 was observed in 2009 at the Robert C. Byrd Green Bank Telescope (GBT) at 350 MHz (M. Roberts, private communication); however, only a quick search of the first five minutes of data was performed and the

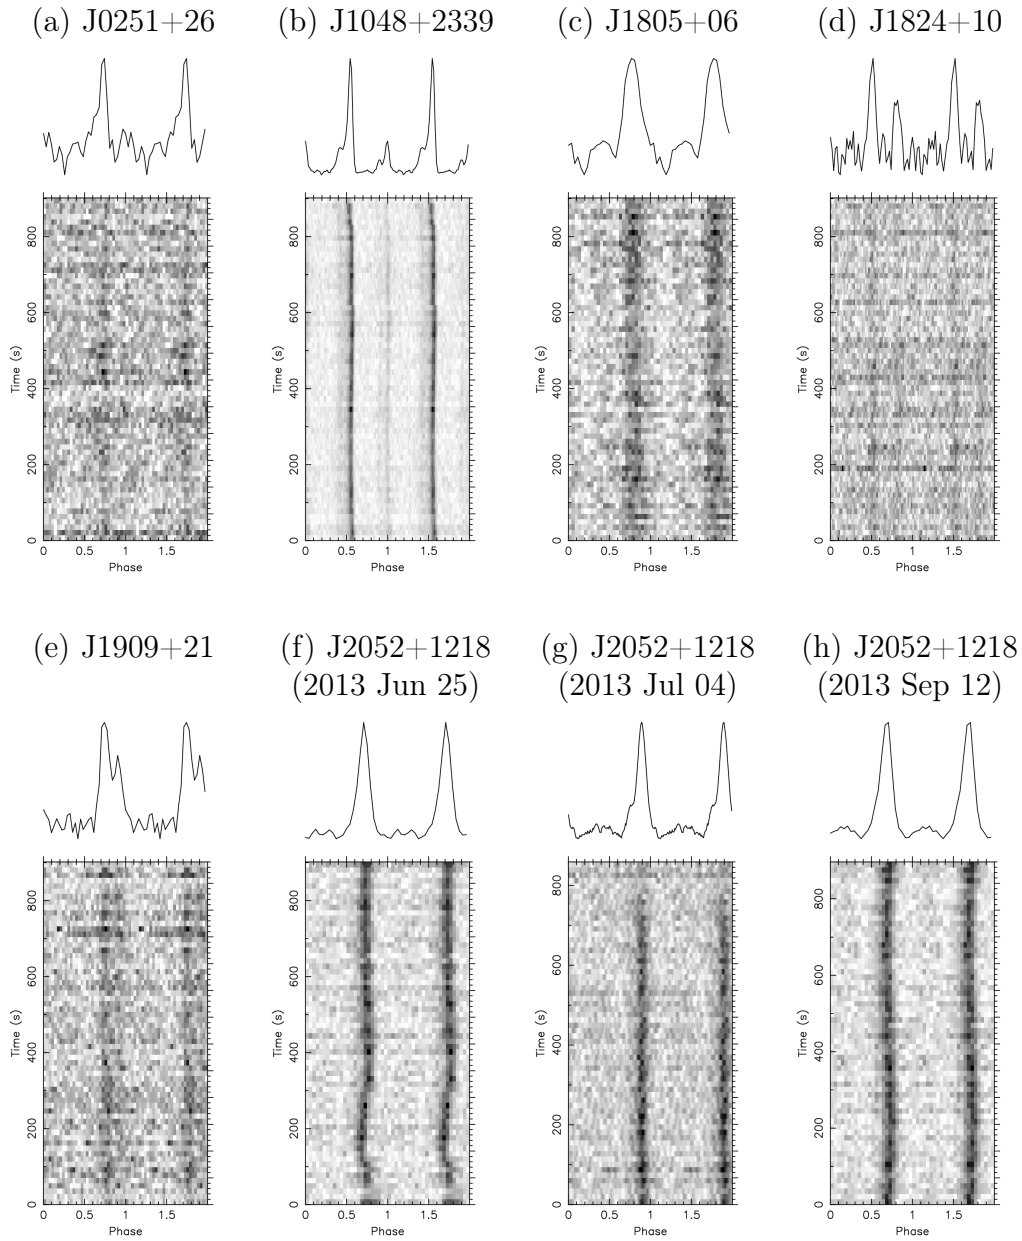


Figure 2.2: The best detections from the search observations of six new MSPs, folded modulo the period and period derivative returned by the software (two rotations are shown). The three search observations of the eclipsing PSR J2052+1218 show some effects related to likely eclipse egress (panel f) and ingress (g).

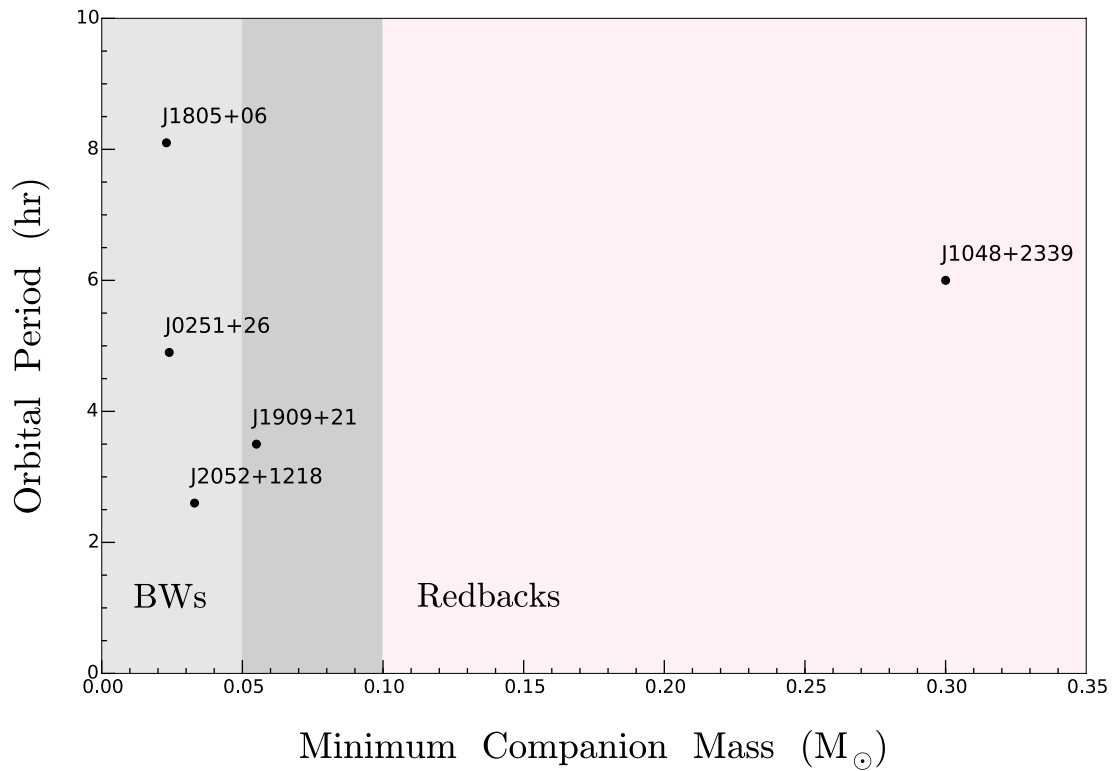


Figure 2.3: The new short-orbit MSPs from this work are presented in an orbital period vs. minimum companion mass plot. MSPs in the light grey area (leftmost block) are black widows, the one in the pink area (rightmost block) is a redback, and PSR J1909+21 is intermediate between the two, but classified as a redback (see Section 2.1.4).

Name ^a	P (ms)	DM (pc cm ⁻³)	Distance ^b (kpc)	P_{orbit} (hr)	Minimum Companion Mass ^c (M_{\odot})	Type ^d	Eclipses?	Discovery Flux Densities (mJy)
J0251+26	2.54	20	0.8	4.9	0.024	BW	Yes	0.3, 0.3
J1048+2339	4.66	17	0.7	6.0	0.30	RB	Yes	2.4
J1805+06	2.13	65	2.5	8.1	0.023	BW	No ^e	1.1, 1.5
J1824+10	4.07	60	2.5	1980.0	0.26	NSWD	No	0.09, 0.15
J1909+21	2.56	62	3.2	3.5	0.055	RB	Yes	0.6
J2052+1218	1.99	42	2.4	2.6	0.033	BW	Yes	1.3, 1.0, 1.2

^a Names with four digits of declination have been given only to MSPs with phase-connected timing solutions.

^b From the NE2001 model (Cordes & Lazio 2002).

^c Assuming a pulsar mass of 1.35 M_{\odot} (Özel et al. 2012).

^d BW = black widow, RB = redback, NSWD = neutron star-white dwarf.

^e PSR J1805+06 has not yet shown any eclipsing behavior; however, there is a gap in orbital coverage at phases 0.2–0.27, so eclipses cannot be ruled out.

pulsar was not detected. Searching the full data set following our discovery at Arecibo, the MSP is clearly detected. Ransom et al. also observed this source twice with the GBT at 820 MHz, but the MSP did not show up in a preliminary analysis of the first observation and the second dataset was not searched. Using the known DM and approximate period from our Arecibo detections reveals the pulsar in both GBT datasets.

PSR J2052+1218 is an intriguing system due to the pulsar’s very fast rotation (1.99 ms) and its short binary period (2.8 hr). Even after searching over acceleration, residual drifts in phase vs. time can be seen in this and other black widow and redback systems (see Figure 2.2, especially (b) and (h)).

We searched the sources containing PSRs J1921+01 and J2042+02 and detected the MSPs, unaware that they had already been discovered at the GBT. These will be published in a forthcoming paper detailing *Fermi*-LAT searches at the GBT (S. Sanpa-Arsa et al. 2016, in preparation).

The six new findings mark the first *Fermi* MSPs discovered using the Arecibo telescope and broke the 50-pulsar threshold for total LAT-guided radio MSP discoveries (which as of 2015 December stands at 69).

2.1.5 Discussion

Possible Candidates for Re-Observation: In searching 34 unidentified *Fermi*-LAT gamma-ray sources at Arecibo, we detected 8 MSPs, for a 24% success rate. This is in line with the success rate for LAT-guided radio surveys at the GBT and Parkes (Hessels et al. 2011; Ransom et al. 2011; Camilo et al. 2015 — see Table 2.2). While we find this to be a satisfying result, it is possible that some remaining sources in Table 2.4 could still be pulsars. Seventeen of the 26 sources currently without a known

pulsar counterpart are spectrally consistent with pulsars (denoted by a ranking of 1, 2, or 3 in the “Spectrum Notes” column) and have no known AGN association. Sources ranked 1 or 2 are very likely to be pulsars, while rank-3 sources lack definitive evidence to suggest they are not pulsars. An inability to make a detection does not preclude the presence of an MSP; rather, it may be due to a pulsar’s faintness, eclipses, scintillation, or extreme orbital or spin parameters. The large number of black widow and redback systems that have been discovered in *Fermi*-LAT sources make eclipses a distinct possibility for this collection of candidates. For example, both PSRs J1048+2339 and J1909+21 were only detected in the second of two search observations, owing to eclipses (Table 2.1). Additional observations of the 17 remaining “good” sources in Table 2.4 may result in the detection of new MSPs.

Uncertain Gamma-Ray Associations: Two new MSPs, PSRs J1048+2339 and J1909+21, may have been “lucky” discoveries within the error circle of a gamma-ray source, but not necessarily associated with that source. In 3FGL, J1048.6+2338 is listed as being possibly associated with a BL Lacertae-type blazar. Blazar associations are generally spatial, and accidental coincidence is a common cause for reclassification of non-variable sources. Until it is possible to fold the gamma-ray photons modulo the parameters obtained with a radio timing solution, it will remain unclear whether the *Fermi*-LAT source is an MSP or possibly a blazar. PSR J1909+21 is not associated with a nearby 3FGL source. We selected it for observation because in a preliminary source list internal to the LAT collaboration there appeared to be a promising source. As for all the MSPs we have discovered, we will know whether this one is associated with a LAT source once we have rotational ephemerides and can fold the gamma-ray photons.

Sensitivity in the Context of Other LAT Radio Surveys: If no break exists in the $\log N - \log S$ distribution of MSPs, one might expect that an increase in

Table 2.4: Arecibo Searches: Unidentified Gamma-Ray Source Information from 3FGL Catalog

3FGL name ^a	r_{95}^b (deg)	Class ^c	Sig ^d (σ)	Curve ^e (σ)	Var ^f	Spectrum ^g Notes	N_{obs}^h
J0103.7+1323	0.08	bcu	7.1	2.4	53	3 lh	3
J0134.5+2638	0.06	bcu	12.0	3.0	58	4 lh	3
J0232.9+2606	0.09	bcu	4.3	1.6	34	3 h	3
J0251.1+2603	0.11	psr	7.9	3.4	36	2 cp	2
J0318.1+0252	0.09	...	12.8	5.7	50	1 Pc	3
J0330.6+0437	0.11	...	7.4	2.4	62	2 cd	3
J0342.3+3148c	0.10	...	6.8	2.5	48	3 ld	3
J0421.6+1950	0.12	...	6.1	1.9	47	2 ld	3
J0517.1+2628c	0.12	...	7.0	1.6	50	3 ld	5
J0539.8+1434	0.08	fsrq	7.6	3.1	299	5 lVd	2
J1048.6+2338	0.12	bll	9.0	2.5	50	3 LD	2
J1049.7+1548	0.09	...	7.6	1.1	58	3 lh	2
J1200.4+0202	0.07	...	8.8	1.6	56	4 Lh	3
J1225.9+2953	0.05	...	17.4	5.2	58	1 Cp	3
P7R4 J1250+3118 ⁱ	3
J1309.0+0347	0.15	...	4.1	2.7	43	2 ?l	6
J1322.3+0839	0.12	bcu	7.8	0.4	74	3 ld	3
J1601.9+2306	0.11	...	7.8	4.3	47	2 P	3
J1627.8+3217	0.07	...	10.2	3.8	33	2 C	2
J1704.1+1234	0.07	...	9.4	0.5	47	4 LD	3
J1720.7+0711	0.09	...	9.6	1.6	44	3 cD	3
J1805.9+0614	0.09	psr	9.3	4.3	40	1 CP	2
J1824.0+1017	0.09	psr	6.7	3.5	48	2 lc	2

Table 2.4 continued

Table 2.4: Arecibo Searches: Unidentified Gamma-Ray Source Information from 3FGL Catalog (continued)

3FGL name ^a	r_{95}^b (deg)	Class ^c	Sig ^d (σ)	Curve ^e (σ)	Var ^f	Spectrum ^g Notes	N_{obs}^h
J1827.7+1141	0.10	...	6.5	3.8	39	2 lc	2
J1829.2+3229	0.15	...	5.7	3.5	49	2 ld	3
J1842.2+2742	0.08	...	8.3	2.4	39	2 c	3
P7R4 J1909+2102ⁱ	2
J1921.2+0136	0.10	psr	9.3	1.7	35	2 cD	4
J2026.3+1430	0.09	...	7.1	2.1	49	2 c	3
J2042.1+0247	0.14	PSR	7.5	4.8	28	1 CP	2
J2052.7+1217	0.11	psr	7.2	1.6	39	3 ID	3
J2108.0+3654	0.06	bcu	6.4	1.5	40	4 HI	3
J2212.5+0703	0.10	...	14.3	4.2	57	1 Cp	6
J2352.0+1752	0.07	bll	8.9	1.4	56	4 HI	4

NOTE — 3FGL source properties are from the Acero et al. (2015) catalog.

^a The six new Arecibo MSPs and two MSPs independently discovered at the GBT (S. Sampa-Arsa et al. 2016, in preparation) are shown in bold.

^b 3FGL source error circle radius at the 95% confidence level.

^c 3FGL pipeline classification scheme. PSR and psr are pulsars with and without LAT pulsations, respectively. The bll designation signifies a BL Lac object, bcu is an unclassified blazar, and fsrq is a flat-spectrum radio quasar.

^d 3FGL source significance.

^e Curvature significance for 3FGL source spectrum when fit to a log-parabolic model.

^f Variability index for source, where an index > 73 denotes variability at $> 99\%$ confidence level.

^g For a full description of *Spectrum Notes*, see Camilo et al. (2015). The first number in this scheme is a rating of how likely the source is to be a pulsar. A “1” means it is very likely, while sources with a “4” or “5” rating (or “3” with a possible AGN association) have been crossed off the list and will not be reobserved because they are unlikely to be pulsars. The source characteristics, on which the rating is based, are obtained from inspection of a source’s spectral energy distribution.

^h Number of times each 3FGL source was observed (from Table 2.1).

ⁱ Not included in the 3FGL catalog. P7R4 designators refer to unpublished source lists.

sensitivity would yield higher survey success rates. Instead, our discovery rates were comparable to those of other lower-sensitivity LAT-guided MSP surveys. However, we have based our target list on the 3FGL catalog, while previous surveys have been based largely on earlier catalogs, and newer *Fermi*-LAT catalogs include weaker, less well characterized sources that are more difficult to classify spectrally. In addition, some relatively bright MSPs, particularly those not subject to large accelerations, would have already been discovered in previous Arecibo “all-sky” surveys.

Figure 2.1 presents minimum detectable flux densities for four radio searches of *Fermi*-LAT sources, including our Arecibo work. Parameters for each of the surveys are provided in Table 2.2. Each of the sensitivity curves has been scaled to 327 MHz using an assumed MSP spectral index of $\alpha = -1.7$ (Stovall et al. 2014). As an example of their relative power, for spin period $P = 1.8$ ms the Arecibo searches are as sensitive at $DM = 100 \text{ pc cm}^{-3}$ as the GBT surveys are for $DM = 10 \text{ pc cm}^{-3}$. For identical low DMs, the Arecibo surveys are about twice as sensitive as the GBT searches. In other words, integration time at the GBT would have to be quadrupled to reach comparable raw sensitivities to Arecibo — but such an increase in integration time would have deleterious consequences for the detectability of compact binaries.

We list the radio flux densities for all discovery observations in Table 2.3 (these were obtained from an application of the radiometer equation and we estimate they have $\approx 25\%$ uncertainty). We see by comparison to Figure 2.1 that PSRs J0251+26 and J1824+10 could only have been discovered with Arecibo. Parkes could only have detected PSR J1048+2339. This only considers raw telescope sensitivity; it does not take into account sensitivity to high acceleration (discussed later), which further emphasizes the utility of large telescopes. For a discussion of selection effects related to interstellar scintillation and eclipses, see Camilo et al. (2015).

Why has the Arecibo survey turned up such a large proportion (5/6) of highly

accelerated interacting binaries, compared to the fractions found in other *Fermi*-LAT surveys? Though small-number statistics is a possible explanation, the result can likely be attributed to the Arecibo telescope’s very large gain, coupled with the relatively short integration times used, and the multiple-observation strategy used to search each good target. An integration time of just 15 minutes at Arecibo yields a minimum detectable flux density that is substantially lower than the longer integrations elsewhere (see Figure 2.1).

The population of *Fermi*-LAT MSPs contains a disproportionately large number of interacting binary systems for reasons that are currently poorly understood. For a time, it was thought that a tendency for intrabinary shocks to produce high-energy radiation could bias *Fermi*-LAT searches towards discovering these systems (e.g., Ray et al. 2012). More recent analyses, however, have found little evidence to support this claim⁹. The bias is likely due in part to previous surveys’ biases against finding binaries due to eclipses and acceleration.

The use of modern acceleration search techniques (as implemented within PRESTO in our case; Ransom 2001) was essential for the detection of the five compact MSP systems. Both Johnston & Kulkarni (1991) and Bagchi et al. (2013) have explored the detectability of binary pulsar systems, the latter having expanded the former’s work to include eccentric binaries. Johnston & Kulkarni (1991) provide a quantitative measure of the loss of power due to acceleration by way of an “efficiency factor”, γ_m . Squaring this value gives a ratio of the power in the m^{th} harmonic, which includes degradation due to acceleration a and jerk \dot{a} , to the power that would be present were the acceleration zero. Three such γ_m terms were reformulated in Bagchi et al. (2013). The first, γ_{1m}^2 , describes the ratio that would be found in a “standard” pulsar search in which acceleration is not searched over. The term γ_{2m}^2 describes the power that

⁹Johnson 2015: <http://www.iac.es/congreso/ns-ewass-2015/>

would be recovered in a constant acceleration search (like the ones we performed), and will be employed here. It is formulated as follows:

$$\gamma_{2m} = \frac{1}{t_{\text{int}}} \left| \int_0^{t_{\text{int}}} \exp \left[\frac{im\omega_p}{c} \left(\left(\int_0^t v_l dt \right) - \alpha_a t^2 - \alpha_v t \right) \right] dt \right|, \quad (2.2)$$

where t_{int} is the integration time of the observation, v_l is the pulsar's line-of-sight velocity, and ω_p is its angular spin frequency. A modern search algorithm yields values of acceleration α_a and velocity α_v that maximize γ_{2m} . Here, $\gamma_{2m}^2 = 1$ for a system with constant acceleration. The final term, γ_{3m}^2 , describes the power ratio recovered in a search over velocity, acceleration, and jerk. Such search algorithms are currently being developed, but have not yet been implemented.

Using software provided by Bagchi et al. (2013)¹⁰, we calculated values of γ_{2m}^2 for PSR J2052+1218 during a 15-minute integration, such as at Arecibo. We then recalculated these values using the integration times for surveys at the GBT and Parkes to compare the detectability of this fast-spinning, highly accelerated binary pulsar by the four different surveys. Results are given in Table 2.5.

As expected, the power recovered in successively higher harmonics decreases for each of the four surveys. The value of γ_{2m}^2 in the first harmonic is a reasonable proxy for binary detectability; that is, the higher the fraction of power that is recovered in an acceleration search, the more likely one is to detect the MSP in a given observation. For PSR J2052+1218, with a large and rapidly changing acceleration, 15-minute observations are significantly better at recovering power from a range of harmonics than longer integrations. Comparing γ_{2m}^2 in the first harmonic between the 15-minute Arecibo observation and the next-longest (32-minute GBT) observation shows that Arecibo yields a signal that is more than twice the strength of the GBT's (not consid-

¹⁰<http://psrpop.phys.wvu.edu/binary>

Table 2.5: Values of γ_{2m}^2 for PSR J2052+1218 as a Function of Integration Time

m (harmonic #)	15 min	32 min	45 min	60 min
1	0.745	0.315	0.188	0.120
4	0.358	0.138	0.099	0.052
8	0.253	0.087	0.051	0.032

NOTE — See Section 2.1.3 for a discussion of this comparison of relative sensitivity to a highly accelerated fast-spinning binary pulsar.

ering differences in telescope gain and system temperature). The difference becomes even more dramatic for successively longer observations. While longer integration times improve sensitivity, the effect is only proportional to the square root of the observation length, while telescope gain is a directly proportional parameter. One strategy to combat the effects of acceleration (useful for relatively bright MSPs) is to take a long observation and apply acceleration searches to small subsections of the data, as well as searching the entire observation.

The characteristics of the Arecibo telescope give it a two-fold advantage over similar instruments. First, it has a significantly better raw sensitivity than both the GBT and Parkes for similar *Fermi*-LAT source searches. We are therefore able to detect *fainter* systems, even in the absence of considerations relating to binary systems. Second, its large gain allows for short observations, which in turn increases its sensitivity to *highly accelerated* binaries, of which there are many among the *Fermi* source population. Short observations also allow us to split observing time over multiple epochs, rather than integrating for a long time at a single epoch, further increasing our ability to combat eclipses and scintillation. Though its declination range is limited, the Arecibo telescope’s raw sensitivity firmly establishes its indispensability as an MSP-finding resource.

Acknowledgements: H. Thankful Cromartie would like to thank the NSF, Fernando Camilo, the staff of the Arecibo Observatory, and its resident scientists for the opportunity to pursue this research, and for an experience that compelled her to continue in astrophysics.

The Arecibo Observatory is operated by SRI International under a cooperative agreement with the National Science Foundation (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana, and the Universities Space Research Association.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

The *Fermi* LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l’Energie Atomique and the Centre National de la Recherche Scientifique / Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden. Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d’Études Spatiales in France.

Chapter 3

Relativistic Shapiro Delay Measurements of an Extremely Massive Millisecond Pulsar

3.1 The Most Massive Neutron Star to Date

3.1.1 Preface

The text in Chapter 3 appears in Cromartie et al. (2020) (originally published in September 2019 in *Nature Astronomy*) as “*Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar.*” I have chosen to largely keep the traditional *Nature* formatting conventions, including an integrated abstract (bold text) and a methods section following the main text. See Section 5.5 for supplementary figures from the project that were not included in the *Nature Astronomy* manuscript.

3.1.2 Relativistic Shapiro Delay Measurements of an Extremely Massive Millisecond Pulsar

Despite its importance to our understanding of physics at supranuclear densities, the equation of state (EoS) of matter deep within neutron stars remains poorly understood. Millisecond pulsars (MSPs) are among the most useful astrophysical objects in the Universe for testing fundamental physics, and place some of the most stringent constraints on this high-density EoS. Pulsar timing — the process of accounting for every rotation of a pulsar over long time periods — can precisely measure a wide variety of physical phenomena, including those that allow the measurement of the masses of the components of a pulsar binary system (Lorimer & Kramer 2004). One of these, called relativistic Shapiro delay (Shapiro 1964), can yield precise masses for both an MSP and its companion; however, it is only easily observed in a small subset of high-precision,

highly inclined (nearly edge-on) binary pulsar systems. By combining data from the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) 12.5-year data set with recent orbital-phase-specific observations using the Green Bank Telescope, we have measured the mass of the MSP J0740+6620 to be $2.14_{-0.09}^{+0.10}$ solar masses (68.3% credibility interval; 95.4% credibility interval is $2.14_{-0.18}^{+0.20}$ solar masses). It is highly likely to be the most massive neutron star yet observed, and serves as a strong constraint on the neutron star interior EoS.

Relativistic Shapiro delay, which is observable when a pulsar passes behind its stellar companion during orbital conjunction, manifests as a small delay in pulse arrival times induced by the curvature of spacetime in the vicinity of the companion star. For a highly inclined MSP-white dwarf binary, the full delay is of order $\sim 10 \mu\text{s}$. The relativistic effect is characterized by two parameters, “shape” and “range.” In general relativity, shape (s) is the sine of the angle of inclination of the binary orbit (i), while range (r) is proportional to the mass of the companion, m_c . When combined with the Keplerian mass function, measurements of r and s also constrain the pulsar mass (m_p ; Freire & Wex 2010 provides a detailed overview and an alternate parameterization).

Precise neutron star mass measurements are an effective way to constrain the EoS of the ultra-dense matter in neutron star interiors. Although radio pulsar timing cannot directly determine neutron star radii, the existence of pulsars with masses exceeding the maximum mass allowed by a given model can straightforwardly rule out that EoS.

In 2010, Demorest et al. reported the discovery of a 2-solar-mass MSP, J1614–2230 (Demorest et al. 2010)¹. This Shapiro-delay-enabled measurement disproved the plau-

¹Though the originally reported mass was $1.97 \pm 0.04 M_\odot$, continued timing has led to a more

sibility of some hyperon, boson, and free quark models in nuclear-density environments. In 2013, Antoniadis et al. used optical techniques in combination with pulsar timing to yield a mass measurement of $2.01 \pm 0.04 M_{\odot}$ for the pulsar J0348+0432 (Antoniadis et al. 2013). These two observational results (along with others; see Freire et al. 2011) encouraged a reconsideration of the canonical $1.4 M_{\odot}$ neutron star. Gravitational wave astrophysics has also begun to provide EoS constraints; for example, the Laser Interferometer Gravitational-Wave Observatory (LIGO) detection of a double neutron star merger constrains permissible equations of state, suggesting that the upper limit on neutron star mass is $2.17 M_{\odot}$ (90% credibility; Margalit & Metzger 2017). Though the existence of extremely massive ($> 2.4 M_{\odot}$) neutron stars has been suggested through optical spectroscopic and photometric observations (e.g. Linares et al. 2018), radio timing can provide much more precise constraints on the existence of $\gtrsim 2 M_{\odot}$ neutron stars.

NANOGrav employs pulsar timing for an important general relativistic application: the detection of low-frequency gravitational waves primarily from supermassive black hole binaries. The collaboration’s observing program consists of high-cadence, multi-frequency radio observations of ~ 75 MSPs using the Green Bank and Arecibo telescopes (GBT and AO; see Arzoumanian et al. 2018 and the upcoming 12.5-year data release). Additionally, NANOGrav has begun using the Karl G. Jansky Very Large Array as the third observatory in its pulsar timing program. Using the Green Bank Telescope, NANOGrav regularly observes J1614–2230 and another high-mass radio MSP, J0740+6620.

PSR J0740+6620 (period = 2.89 ms) was discovered in the Green Bank Northern Celestial Cap 350-MHz survey (GBNCC) in 2012 Stovall et al. (2014). It is in a nearly circular (eccentricity = 5×10^{-6}), 4.77-day orbit (a recent GBNCC timing solution

precise mass measurement of $1.928 \pm 0.017 M_{\odot}$; see Fonseca et al. (2016).

is presented in Lynch et al. 2018). Recent optical and near-infrared observations revealed that its companion is likely the coolest white dwarf ever measured in orbit with an MSP (Beronya et al. 2019).

Here we present timing observations of the pulsar with the GBT taken between 2014 and 2019. We observed the pulsar regularly throughout this period as part of the NANOGrav timing program (Arzoumanian et al. 2018). This section of our data set includes ~ 70 epochs (occurring approximately monthly and at random orbital phases) during which the pulsar was observed at both 1.4 GHz and 820 MHz for ~ 20 minutes each. We were awarded additional time for two concentrated campaigns over superior conjunction (i.e. when the pulsar is behind its companion star), as probing the minima and maxima of the Shapiro delay signal is the best way to improve sensitivity to it (see the absorbed or “detectable” signal in the second panel of Figure 3.1).

After the second concentrated campaign consisting of two five-hour observations at orbital phases 0.15 and 0.25 (GBT 18B–372), the timing analysis (see details in Methods) yielded a pulsar mass of $2.14_{-0.09}^{+0.10} M_{\odot}$ at 68.3% credibility. The Methods section describes our rationale for choosing these two orbital phases, as well as the progression of mass measurements and precisions as more observations were added. Our final fits with and without Shapiro delay as a function of orbital phase are presented in Figure 3.1, and the top panel of Figure 3.2 shows timing residuals spanning the entire data set. Figure 3.3 shows a map of fitted probability distributions for the parameters m_c , m_p , and i .

Although our measured relative uncertainty is higher than, for example, the original relative error reported by Demorest et al. for J1614–2230 (5% vs. 2%), PSR J0740+6620 is a remarkably high-mass MSP. This measurement will help constrain high-density nuclear physics, as there are very few examples of $\gtrsim 2 M_{\odot}$ neutron stars. PSR J0740+6620 is 98% and 90% likely to be more massive than J1614–2230 and

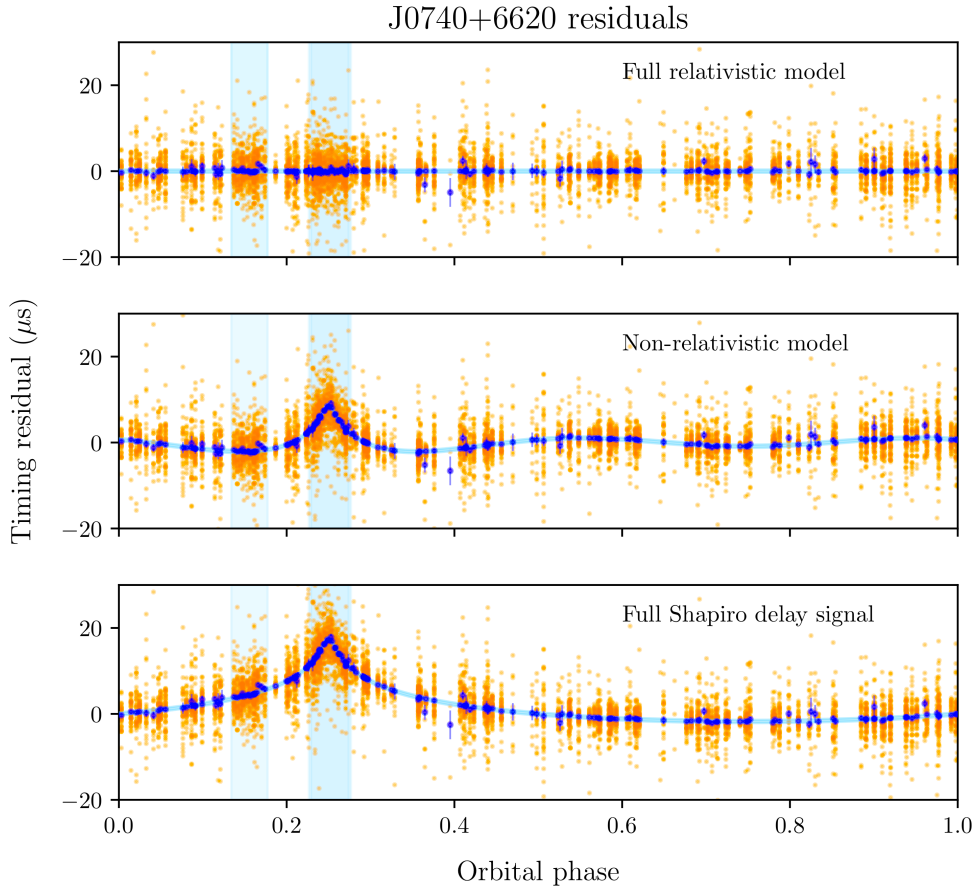


Figure 3.1: Timing residuals from all observations of J0740+6620 as a function of orbital phase, with superior conjunction at orbital phase = 0.25. Orange points are multi-frequency timing residuals, while dark blue points are averages of each group (i.e. timing epoch) of these points with $1\text{-}\sigma$ error bars. Averages were taken over a minimum of four data points to avoid showing misleading residuals from faint observations. Blue boxes indicate the orbital phases over which each of the three supplemental observations were taken (the box over conjunction is slightly darker because we made two superior conjunction observations). The top panel shows the full fit (including Shapiro delay parameters and all dispersion measure parameters — i.e. the full timing solution). The middle panel is the best fit with the measurable Shapiro delay signal added; this is the signal to which we are actually sensitive. The bottom panel is the “full” Shapiro delay signal. Both the second and third panels are calculated based on the orbital and system parameters determined from the full fit. The lighter blue line in the middle and bottom panels represents the theoretical measurable and full Shapiro delay, respectively (and marks a $0\text{-}\mu\text{s}$ residual in the top panel). The width of the line in each panel is equal to the root mean squared error of the averaged points.

J0348+0432, respectively, and is therefore likely to be the most massive well measured neutron star to date.

Taken together, these three massive MSPs serve as a strong validation of the existence of high-mass neutron stars. Due to the asymptotic nature of the relationship between maximum neutron star mass and nearly all EoS, even small increases in the measured mass of the most massive neutron stars force a reconsideration of the fundamental physics at play in their interiors (for example, see Figure 3 in Watts et al. 2015). Non-nucleonic solutions to the EoS problem, such as quark matter, hyperons, or Bose-Einstein meson condensates, yield softer equations of state (i.e. relatively compressible matter); however, more massive neutron stars necessitate stiffer EoS, which allow for higher maximum masses (see Özel & Freire 2016 for a review). The measurement of a 2.14-solar-mass neutron star is therefore in extreme tension with these non-nucleonic proposals, and underlines the necessity of untangling existing theoretical paradoxes. The most prominent of these may be the hyperon problem, which proposes that although the extreme densities inside neutron stars would favor the conversion of nucleons to hyperons, the presence of hyperons softens the EoS and excludes the possibility of high-mass neutron stars (for example, Bedaque & Steiner 2015). In addition, the mass measurement of J0740+6620 may have implications for the nature of neutron star mergers as detected by LIGO. Because several neutron stars with masses close to or greater than $\sim 2 M_{\odot}$ are now known, it may be the case that more mass-asymmetric neutron star mergers will occur than previously supposed.

Constraining the mass of J0740+6620 carries additional astrophysical benefits. Recent evidence from Antoniadis et al. (2016) suggests that the distribution of MSP masses may be bimodal, implying that many more neutron stars with masses greater than $\sim 1.6 M_{\odot}$ may exist than previously supposed (see also Özel & Freire 2016). Not only is it becoming clear that high-mass neutron stars make up a significant portion




Figure 3.2: Timing residuals and DMX for all epochs of J0740+6620 data. *Top panel:* Timing residuals from all epochs of J0740+6620 data, including both NANOGrav and superior conjunction-specific observations at all frequencies, are shown in orange (with $1\text{-}\sigma$ error bars). The superimposed blue points represent an average over each epoch ($\text{RMS} = 1.5 \mu\text{s}$; note that some days have two separately calculated averages from dual-frequency data). *Bottom panel:* Blue points indicate DMX values calculated for each epoch of data with $1\text{-}\sigma$ error bars. The DMX trend is somewhat simple (i.e., roughly quadratic); however, linear modeling is strongly disfavored. A single averaged epoch (one dark blue point) was removed from these plots, as its error bar was $\sim 8 \mu\text{s}$ due to a faint detection from which only one TOA could be extracted.




Figure 3.3: Map of fitted χ^2 distributions and corresponding probability density functions for m_p , m_c , and i . The left-hand heat map was generated by computing χ^2 values for different combinations of m_c and i ; the right-hand heat map was calculated by translating the $m_c - i$ probability density function to the $m_p - i$ phase space using the binary mass function. Darker blue regions correspond to lower χ^2 values. The three red circles correspond to 1, 2, and 3- σ significance cutoffs. Each of the three probability density functions (blue lines plotted on the tops and side of the heat maps) are projections of the χ^2 distributions. The solid red lines mark median values of each of the three parameters, while red dashed lines denote the upper and lower bounds of the 68.3% ($1-\sigma$) credibility interval.

of the population, but their existence also carries substantial implications for our understanding of MSP binary evolution. Because many fully recycled pulsars have been measured to have masses less than or equal to $1.4 M_{\odot}$, we know that recycling can be accomplished with only a small amount of mass transfer. We must therefore consider the possibility that some MSPs are not formed near the Chandrasekhar mass and increase to high masses through accretion; rather, they are born massive in the first place (see Tauris et al. 2011; Cognard et al. 2017b for earlier evidence of this phenomenon).

There exists a well known relationship between the mass of a pulsar’s white dwarf companion and the binary system’s orbital period (Rappaport et al. 1995; Tauris & Savonije 1999). For our measured orbital period of ~ 4.77 days, the predicted white dwarf companion masses (from equations 20 and 21 in Tauris & Savonije 1999) are $\sim 0.24 M_{\odot}$ for a mid-metallicity (Pop I+II) donor star and $\sim 0.25 M_{\odot}$ for a low-metallicity, Pop II star. Our measured mass of J0740+6620’s helium white dwarf companion is $0.260^{+0.008}_{-0.007} M_{\odot}$ (at 68.3% credibility). Given the stated uncertainties in convective mixing length, this discrepancy of 5–10% is not an indication that J0740+6620 is an exception to the orbital period vs. white dwarf mass relationship; however, it may indicate that this system was born in a relatively low-metallicity environment. There exist at least three other examples of MSP-helium white dwarf binaries with minimum companion masses greater than the Pop II masses predicted by Tauris & Savonije (J1125–6014, J1903–7051, and J1933–6211). Lastly, if J0740+6620 is measured to be at the high end of our mass credibility interval, it may provide evidence that the creation of a stable, high-mass neutron star is possible through the merger of two low-mass neutron stars (in a LIGO-like gravitational wave event).

Though it will require significant additional observing time to improve upon our

J0740+6620 measurement, high-cadence monitoring of the pulsar is a promising strategy. Daily observations with the Canadian Hydrogen Intensity Mapping Experiment (CHIME; see Ng 2018) telescope, in conjunction with the present data set, have the potential to determine the mass of J0740+6620 with 2-3% precision within a year. Additionally, the Neutron Star Interior Composition Explorer (NICER) is observing J0740+6620 at X-ray wavelengths (https://heasarc.gsfc.nasa.gov/docs/nicer/science_team_investigations). Modeling the thermal pulse profile of this MSP at X-ray energies will aid in constraining the mass and radius of J0740+6620. Continued collaboration with multifrequency observing programs will guarantee the steady improvement of this pulsar mass measurement in the long term.

3.1.3 Methods

3.1.3.1 Green Bank Telescope Observations

Both NANOGrav and targeted observations were conducted using the Green bank Ultimate Pulsar Processing Instrument (GUPPI; DuPlain et al. 2008). Observations at 1500 MHz were acquired with 800 MHz of bandwidth split into 512 frequency channels (which were summed to 64 channels before analysis), sampling every $0.64 \mu\text{s}$. At an observing frequency of 820 MHz, 200 MHz of bandwidth over 128 channels was acquired with an identical sampling rate (and later also summed to 64 channels). These dual-polarization observations at both frequencies were coherently dedispersed at the known DM of 15.0 pc cm^{-3} . Data were processed using NANOGrav pipelines for consistency with the existing four-year-long NANOGrav J0740+6620 data set (see Arzoumanian et al. 2015; Demorest 2018 for a thorough description of NANOGrav observing procedures and a description of NANOGrav’s main data processing pipeline, `nanopipe`).

3.1.3.2 Generation of TOAs and the Timing Model

The measurement and modeling of pulse times of arrival (TOAs) closely mirrors the procedure described by Arzoumanian et al. (2018). We provide a summary of the analysis procedure in this section.

During offline processing, total-intensity profile data were integrated over ~ 20 – 30 minute intervals to yield one or two TOAs per downsampled frequency interval for a normal NANOGrav observation, and ~ 10 minutes for the long scans near or during conjunction. We extracted TOAs from each of the 64 integrated channels over the entire observing bandwidth through cross correlation between the data and a smoothed profile template using the software package `PSRCHIVE` (source code in van Straten et al. 2011)².

We used standard pulsar-timing analysis tools, namely `TEMPO`³ and `TEMPO2`⁴ (source code in Hobbs & Edwards 2012) for modeling TOA variation in terms of many physical mechanisms. `TEMPO` and `TEMPO2`, while not fully independent timing packages, yield consistent results. For J0740+6620, fitted parameters include: celestial (ecliptic) coordinates; proper motion; spin frequency and its first derivative; and binary orbital parameters (see Table 3.1 which lists best-fit values for these parameters as determined with `TEMPO`).

We used the DE436⁵ Solar System ephemeris, maintained by the NASA Jet Propulsion Laboratory, for correction to the barycentric reference frame. The time standard used was BIPM2017. The overall RMS timing residual value for the timing model presented in this work is $1.5 \mu\text{s}$. The χ^2 of our fit is 7314.35 with 7334 degrees of freedom, yielding a reduced- χ^2 value of 0.997; note that the noise modeling (see

²<http://psrchive.sourceforge.net>

³<http://tempo.sourceforge.net>

⁴<https://www.atnf.csiro.au/research/pulsar/tempo2>

⁵<https://naif.jpl.nasa.gov/pub/naif/JUNO/kernels/spk/de436s.bsp.1b1>

Table 3.1: PSR J0740+6620 Best-Fit Parameters

Pulsar name.....	J0740+6620
Dates of Observations (MJD).....	56640 – 58462
Number of TOAs.....	7419
Measured Quantities	
Ecliptic longitude, λ (degrees).....	103.75913607(1)
Ecliptic latitude, β (degrees).....	44.10248468(2)
Epoch of position & period (MJD).....	57551.0
Proper motion in ecliptic longitude (mas yr ⁻¹).....	-2.75(3)
Proper motion in ecliptic latitude (mas yr ⁻¹).....	-32.43(4)
Parallax (mas).....	0.5(3)
Spin frequency, ν (Hz).....	346.5319964932129(6)
Spin frequency derivative, $\dot{\nu}$ (s ⁻²).....	$-1.46389(2) \times 10^{-15}$
Dispersion measure, DM (pc cm ⁻³)*.....	14.961787
Profile frequency dependency parameter, FD1.....	$-1.17(4) \times 10^{-5}$
Binary model.....	ELL1
Projected semi-major axis of orbit, x (lt-s).....	3.9775561(2)
Binary orbital period, P_b (days).....	4.7669446191(1)
Epoch of ascending node, TASC (MJD).....	57552.08324415(2)
EPS1 (first Laplace-Lagrange parameter), $e \sin \omega$	$-5.70(4) \times 10^{-6}$
EPS2 (second Laplace-Lagrange parameter), $e \cos \omega$	$-1.89(3) \times 10^{-6}$
Sine of inclination angle i	0.9990(2)
Companion mass, m_c (M _⊙).....	0.258(8)
Derived Parameters	
Orbital eccentricity, e	$5.10(3) \times 10^{-6}$
Longitude of periastron, ω (degrees).....	244.4(3)
Epoch of periastron, T_0 (MJD).....	57550.543(5)
Binary mass function (M _⊙).....	0.0029733870(4)
Pulsar mass (68.3% credibility interval, M _⊙).....	$2.14^{+0.10}_{-0.09}$
Pulsar mass (95.4% credibility interval, M _⊙).....	$2.14^{+0.20}_{-0.18}$
Companion mass (68.3% credibility interval, M _⊙)...	$0.260^{+0.008}_{-0.007}$
Companion mass (95.4% credibility interval, M _⊙)...	$0.260^{+0.016}_{-0.014}$
Inclination angle (68.3% credibility interval, degrees)	$87.38^{+0.20}_{-0.22}$
Inclination angle (95.4% credibility interval, degrees)	$87.38^{+0.39}_{-0.45}$

* Because this DM is an unfitted reference value, no error is reported. Values of DMX for each of the ~ 70 epochs are available upon request.

Assessment of Timing Noise) will always yield a χ^2 of ~ 1 .

We employed the ELL1 binary timing model (Lange et al. 2001) in describing the nearly-circular orbital dynamics of the J0740+6620 system. Parameters of the ELL1 binary model consist of the projected semi-major axis, orbital period, epoch of passage through the ascending orbital node, and two ‘‘Laplace-Lagrange parameters’’ (ϵ_1 and ϵ_2 ; the orbital eccentricity multiplied by the sine and cosine of periastron longitude, respectively; Lange et al. 2001) that quantify departures from perfectly circular orbits.

3.1.3.3 Assessment of Timing Noise

MSP rotation often exhibits a limit in achievable precision due to the presence of stochastic processes that act as noise to timing measurements. Examples of timing noise include systematic errors from cross-correlation template matching and ‘‘spin noise’’ due to irregular rotation of the neutron star. We use a noise model similar to those developed in the NANOGrav 9-year and 11-year data releases in order to quantify these noise terms in the J0740+6620 data set.

The noise model consists of white-noise components that combine to form additive Gaussian noise. For each of the two frontend receivers used in this work, we use three parameters to describe the white-noise contribution to timing noise: a scaling factor applied to all raw TOA uncertainties (‘‘EFAC’’); a term added in quadrature to the TOA uncertainties (‘‘EQUAD’’); and a noise term that quantifies TOA correlations purely across observing frequency (‘‘ECORR’’).

We used the **Enterprise**⁶ modeling suite for estimation of the white components of the noise model using a Markov chain Monte Carlo (MCMC)-based algorithm.

⁶<https://enterprise.readthedocs.io/en/latest>

Enterprise uses the `TEMPO(2)` fit as the maximum-likelihood fit for the timing parameters and the basis of the fit for the red noise parameters, should they be found to be significant. In our `TEMPO(2)` fits, we include an EFAC of 1.036 for L-band (1500-MHz) TOAs and 1.013 for 820-MHz TOAs. EQUAD for L-band is $0.00610 \mu\text{s}$, and $0.18310 \mu\text{s}$ for 820 MHz. ECORR values for L-band and 820-MHz TOAs are $0.00511 \mu\text{s}$ and $0.00871 \mu\text{s}$, respectively. Bayesian model selection via an Enterprise MCMC run disfavors the inclusion of red noise; therefore, the noise model includes only white noise components.

3.1.3.4 Dispersion Measure Modeling

The complexity of modeling DM variations arising from a dynamic interstellar medium has been discussed at length in previous works (see, for example, Lam et al. 2016; Jones et al. 2017). We have adopted the standard NANOGrav piecewise-constant model for DM trends wherein each epoch of data is fit with a constant “DMX” value; in other words, each of these parameters is a deviation from some nominal DM and is fixed over a single epoch. The observation that J0740+6620’s DM behavior is somewhat smooth over the duration of our data set (see Figure 3.2) led us to attempt alternatively modeling the entire data set by fitting for only the first and second derivatives of DM. In theory, this approach could be advantageous given the ability of DMX to absorb Shapiro delay signals (thanks to the similar duration of conjunction and a DMX epoch). While this strategy does reduce the formal parameter uncertainties from the fit, both an F-test and an Akaike information criterion test strongly favor the DMX model over the quadratic DM fit. This indicates the DM variation is not fully characterized by a quadratic model, and parameter values (including pulsar mass) derived from this model are likely to have systematic biases

not reflected in their formal uncertainties.

3.1.3.5 Simulations

Analysis of the NANOGrav 12.5-year data set without supplemental data yielded $m_p = 2.00 \pm 0.20 M_\odot$. After the initial 6-hour supplemental observation, we measured the mass of J0740+6620 to be $2.18 \pm 0.15 M_\odot$. We conducted simulations of future observations both to predict the constraining power of a concentrated Director’s Discretionary Time campaign as well as to determine how our mass measurement may improve with additional observations going forward. For these simulations, we first generated an arbitrary array of TOAs that mirror the desired observing cadence, starting date, etc. The TOAs were then fit (with pulsar timing software such as TEMPO or PINT⁷) using the known parameters for J0740+6620. Residuals from this fit were then subtracted from the original TOAs to create “perfect” TOAs, to which stochastic noise was then added. Two notable types of simulations were conducted. The first was an estimation of the improvement in our measurement of m_p given random orbital sampling (the “NANOGrav-only observation” scenario); this solidified our conclusion that the concentrated GBT campaigns were necessary. The second served to optimize our observing strategy during a targeted orbital phase campaign by trying various permutations of orbital phase, number of observing sessions, and observing session lengths. The results of this simulation informed our GBT Director’s Discretionary Time request for five hours over conjunction and five hours in one of the Shapiro “troughs” (we were awarded time in the first trough — around orbital phase 0.15 — in addition to conjunction). In order to ensure that obtaining data in this asymmetric fashion would not bias our mass measurement, we ran 10,000 simulations of a five-hour conjunction observation plus five hours in either the first or second

⁷<https://github.com/nanograv/PINT>

Shapiro trough. The averages of the 10,000 mass measurements obtained from each of these troughs were consistent within 1%, implying that our orbital sampling is not biasing our results (as one would expect, given that the Shapiro delay response curve is symmetric about superior conjunction).

3.1.3.6 Data Availability

PSR J0740+6620 TOAs from both the 12.5-year data set and from the two supplemental Green Bank Telescope observations will be available at <https://data.nanograv.org> upon publication of this manuscript.

3.1.3.7 Code Availability

All code mentioned in this work is open source and available at the links provided in the manuscript.

Chapter 4

Radio Shapiro Delay Observations of the Bright Gamma-Ray MSP J1231–1411

4.1 Abstract

We present a multiwavelength analysis of J1231–1411, the brightest known Gamma-ray millisecond pulsar. Not only is this source a primary science target for the Neutron Star Interior Composition Explorer telescope (NICER), but it has also been studied extensively with the Green Bank and Nançay radio telescopes, as well as the *Fermi* Large Area (Gamma-ray) Telescope. A new orbital-phase-targeted campaign with the Green Bank Telescope was designed to improve our constraint on the pulsar’s mass through the detection of relativistic Shapiro delay. This mass measurement could prove beneficial in constraining the poorly understood neutron star equation of state if it is taken in conjunction with NICER’s thermal X-ray lightcurve modeling of the source. Both a χ^2 minimization gridding analysis and several Markov-Chain Monte Carlo (MCMC)-enabled model fits (employing priors informed by white dwarf binary evolution models) of all available radio data suggest that this is a low-mass pulsar in a highly inclined binary orbit. A further MCMC analysis of 12 years of *Fermi*

Gamma-ray photons yields a far less constraining pulsar mass measurement; however, we are optimistic about the prospect of our radio Shapiro delay measurements proving useful for the NICER M/R constraint.

4.2 Introduction

Observations of neutron stars across the electromagnetic and gravitational wave spectra have played a critical role in the effort to elucidate the equation of state (EoS) of supranuclear-density matter. Recently, measurements of tidal deformability in neutron star mergers observed by the Laser Interferometer Gravitational-Wave Observatory (LIGO) experiment (e.g. Margalit & Metzger 2017) have set an upper limit on neutron star mass. Observations of relativistic Shapiro delay in radio millisecond pulsars (MSPs) have resulted in precise measurements of the heaviest neutron stars' masses (e.g. Demorest et al. 2010; Cromartie et al. 2020).

Yet another approach to constraining the neutron star interior EoS is X-ray light curve modeling facilitated by instruments such as NASA's Neutron Star Interior Composition Explorer (NICER) telescope (Gendreau & Arzoumanian 2017). Generally speaking, the X-ray light curve of an MSP depends on the compactness of the neutron star due to relativistic effects, such as light bending of photons emitted from the pulsar's surface (Bogdanov et al. 2007). The precise modeling of the observed light curve therefore permits extracting the mass-to-radius (M/R) ratio for the neutron star (an overview is presented in Watts et al. 2016).

In this letter, we present a multiwavelength analysis of the MSP J1231-1411 and a constraint on its mass in an effort to increase the significance of NICER's forthcoming light curve modeling for that source, and therefore, the neutron star interior EoS. Section 4.3 provides an overview of the GBT, Nançay Radio Telescope (NRT), and

Fermi Large Area Telescope (LAT) observations used in our analysis. Section 4.4 discusses the fundamentals of pulsar mass measurements using Shapiro delay, as well as the software used in our analysis. Section 4.5 discusses the noise modeling and analysis (both linear regression and MCMC) of all available radio data. We explain the reasoning behind our choice of priors on m_c and $\sin(i)$, and outline the six MCMC trials we conducted. Lastly, we provide an overview of our single-photon Gamma-ray analysis. These sections are followed by a discussion (Section 4.6.4) and conclusions (Section 4.7).

4.3 Observations & Initial Analysis

J1231–1411 was discovered with the Green Bank Telescope (GBT) during a radio follow-up campaign of 25 of the brightest unassociated *Fermi*-LAT Gamma-ray sources¹ (see Ransom et al. 2011 and the discovery plot in Figure 4.1). It is a relatively typical 3.68-ms MSP in orbit with a white dwarf companion; their binary orbital period is 1.86 days, and its projected semi-major axis is 2.043 lt-s. Since 2011, J1231–1411 has been intensively observed at radio wavelengths with the NRT, and more recently, the GBT. Although the timing precision of this source is better at 820 MHz than at L-band (i.e. ~ 1.4 GHz), strong diffractive scintillation sometimes renders this pulsar invisible at that frequency.

Ray et al. (2019) reported the existence of X-ray pulsations from J1231–1411, discovered by folding data from NASA’s Neutron Star Interior Composition Explorer (NICER) X-ray telescope with the *Fermi*-LAT timing solution (see Figure 4.2). Spectral modeling suggested that the observed pulsations were thermal in nature, originating from at least one hot spot on the neutron star’s surface.

¹In the *Fermi*-LAT Bright Gamma-Ray Source List; see Abdo et al. (2009).

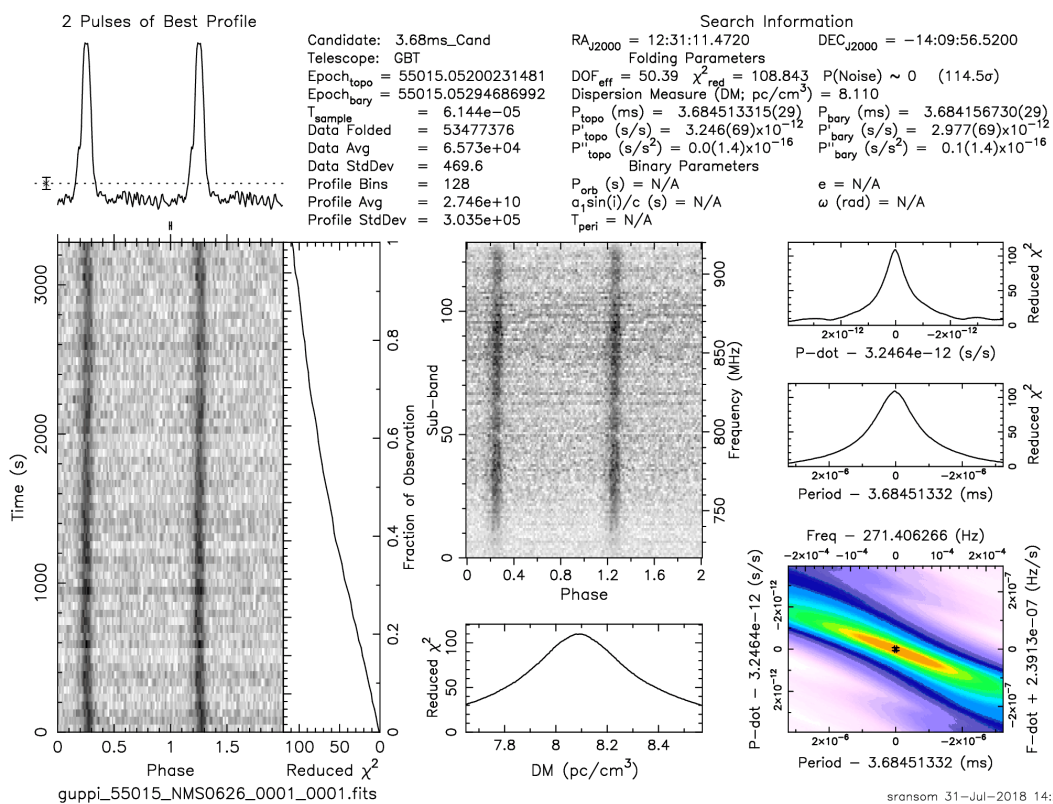


Figure 4.1: Discovery plot for J1231–1411 using the GBT at 820 MHz. At this frequency, the source is bright (~ 0.4 mJy); however, this observation was not subject to the scintillation that frequently dims this MSP.




Figure 4.2: X-ray and radio pulse profiles for J1231–1411 (two phases shown) from NICER observations (green) and NRT observations (orange; arbitrary amplitude). Detailed modeling of the profile shape of the X-ray emission will lead to constraints on M/R for the pulsar. *Figure 5 in Ray et al. (2019)*

4.3.1 Green Bank Telescope, 2009-2010

The GBT at the Green Bank Observatory in Green Bank, WV, is a 100-m fully steerable offset paraboloid radio telescope. PSR J1231–1411 was discovered at 820 MHz with the GBT in July 2009. A timing campaign for the source was conducted between December 2009 and August 2010 using the GBT 350-MHz receiver (see Ransom et al. 2011 for details). Observing epochs were frequent and evenly spaced throughout the entire timing campaign. These GBT data alone (unsurprisingly) do not suggest that Shapiro delay is detectable in J1231–1411.

4.3.2 Nançay Radio Telescope

The Nançay Radio Observatory (Station de Radioastronomie de Nançay) near Nançay, France, is home to “Le Grand Radiotélescope,” a transit telescope with a flat 200×40 -m primary mirror and spherical 300×35 -m secondary mirror. The telescope’s L-band receiver (centered at 1486 MHz with 512 MHz total bandwidth) is frequently used for pulsar timing campaigns. There were two long-term timing programs for J1231–1411 with the NRT. The first took place between November 2009 and August 2010 (concurrent with the GBT program), and the second was a significantly longer effort beginning in September 2011 and continuing through May 2018². This second series of observations was indispensable for obtaining an initial constraint on the inclination angle of the system, which provided a strong justification for our second GBT observing campaign. Taken alone, the NRT data supported a high orbital inclination angle, i , very close to $\sin(i) = 1$ for J1231–1411.

²See Cognard et al. (2017b) for details of the Nançay observational setup during this time period.

4.3.3 *Fermi* LAT

The *Fermi* LAT is a pair-conversion-based Gamma-ray telescope aboard NASA’s *Fermi* satellite, which was launched in June 2008 (Atwood et al. 2009). J1231–1411 is the brightest Gamma-ray MSP known (Abdo et al. 2013), and more than 10 years of *Fermi*-LAT data have facilitated its careful study at those wavelengths. A preliminary joint MCMC analysis of the NRT L-band data and *Fermi* Gamma-ray photons showed a slight preference for lighter MSP masses and high inclination angles (mean inclination $\approx 83.0^\circ$; 95th percentile $\approx 87.9^\circ$). An orbital inclination close to 90 degrees can allow the detection of Shapiro delay in systems with lower-mass companions or lower timing precision; therefore, our (quick-and-dirty) analysis made a compelling case for requesting further observing time with the GBT.

4.3.4 Green Bank Telescope, 2019

We present a new set of observations from the GBT taken in Spring 2019 (project code GBT19A-429; PI Cromartie) and analyzed in combination with archival radio timing data. The 22-hour project comprised five, two-hour observations scheduled at random orbital phases, and two, six-hour observations over superior conjunction (i.e. when the pulsar is behind its stellar companion). The latter pointings were critical in increasing our sensitivity to Shapiro delay. These observations were taken with the 820-MHz receiver and the Green Bank Ultimate Pulsar Processing Instrument backend (GUPPI; DuPlain et al. 2008) with 200 MHz of bandwidth split into 512 channels (later integrated into 128 channels). All observations were conducted using an online folding mode. Portions of several observations were removed, as diffractive scintillation sometimes rendered the MSP too dim.

Calibrated total-intensity data were integrated over 15-minute intervals, and the

resulting profiles were matched to a smoothed template using `PSRCHIVE` (source code in van Straten et al. 2011) to obtain time-of-arrival (TOA) measurements (see Arzoumanian et al. 2018 for a detailed explanation of this process).

4.4 Methods Overview

Pulsar timing — the process of constructing a model that accounts for every rotation of a pulsar over long time spans — has been responsible for the precise measurement of a wide variety of astrophysical phenomena. Among these is Shapiro delay, a general relativistic effect that is observable in a small subset of MSP binaries. When an MSP passes behind its stellar companion along our line of sight, the deformation of spacetime induced by the companion’s mass causes pulsations from the MSP to arrive later than would be expected otherwise. The delay is defined as occurring at orbital phase = 0.25 (superior conjunction), and is observable as a spike in timing residuals of order $\sim 10 \mu\text{s}$. By measuring the extent and shape of this delay, one obtains the post-Keplerian parameters range (r) and shape (s ; see Freire & Wex 2010):

$$r = T_{\odot} m_c \tag{4.1}$$

$$s = \sin(i) = x \left(\frac{P_b}{2\pi} \right)^{-2/3} T_{\odot}^{-1/3} M^{2/3} m_c^{-1} \tag{4.2}$$

where $T_{\odot} = GM_{\odot}/c^3$, M is the system’s combined mass, and m_c is the companion’s mass. When combined with the Keplerian mass function, measurements of s and r independently constrain the pulsar and companion masses (m_p and m_c):

$$f(m_p, m_c) = \frac{4\pi^2}{G} \frac{(a \sin i)^3}{P_b^2} = \frac{(m_c \sin i)^3}{M^2}. \tag{4.3}$$

The measurement of post-Keplerian orbital parameters is one of the best methods for parsing a binary system’s total mass into its constituent companion and pulsar masses.

The long timing baseline afforded to us by both the NRT L-band data and *Fermi*-LAT Gamma-ray photons permitted an initial measurement of Shapiro delay that suggested the system was highly inclined, and that the MSP had a relatively low mass. Though EoS constraints from neutron star mass measurements *alone* require the observation of more and more massive MSPs, the precise measurement of a wide variety of M/R ratios is critical in constraining the neutron star EoS. Therefore, NICER measurements of low-mass and high-mass MSPs alike can prove valuable.

This work employs a variety of pulsar timing tools to constrain the mass of J1231–1411 using both radio and Gamma-ray data. For radio observations, we use both the `TEMPO2` (see Hobbs & Edwards 2012) and `PINT` (Luo et al. 2020, submitted) pulsar timing software packages. Both packages feature traditional χ^2 fitting routines for pulse times of arrival (TOAs). Additionally, `PINT` features a Markov Chain Monte Carlo (MCMC) routine for fitting both radio and Gamma-ray data. The specifics of these techniques are described in Section 4.5.

4.5 Analysis & Results

4.5.1 Noise Analysis for Radio Data

Pulsar timing precision is limited by stochastic noise sources such as radiometer noise, radio frequency interference, and — in a subset of well timed MSPs — “spin noise” (irregularities in neutron star rotation). We contend with these noise sources by including two white noise parameters in our fits for each of the observing

frequency/backend combinations included in our analysis (see Table 4.1). The first is “EFAC”, a scaling error applied to TOA uncertainties; the second is “EQUAD”, a term added in quadrature to the TOA uncertainties. More details about the use of these parameters can be found in, e.g., Arzoumanian et al. (2018). The white noise parameter estimations were derived using the **Enterprise** modeling suite, which employs MCMC-based algorithms (Ellis et al. 2019). We also checked for the presence of red noise in the data; however, Bayesian model selection in **Enterprise** disfavors the inclusion of any red noise terms.

4.5.2 PINT & TEMPO2 χ^2 Analysis

The results of a normal χ^2 minimization fit using PINT and without including our Shapiro delay parameters (m_c and $\sin(i)$, the companion mass and the sine of the orbital inclination angle, i) are presented in Table 4.3. This fit incorporated all available radio data. We employed the binary timing model ELL1 (Lange et al. 2001) for this nearly circular MSP system. ELL1 includes the following orbital parameters: projected semi-major axis ($x \equiv a_p \sin i/c$); binary orbital period (P_b); epoch of ascending node (TASC); and two Laplace-Lagrange parameters, ϵ_1 and ϵ_2 , the orbital eccentricity multiplied by the sine and cosine of periastron longitude, respectively. These parameters are characterized with high precision; therefore, in subsequent analyses that include Shapiro delay parameters (beginning in Section 4.6), focus will be placed on a given trial’s preference for a certain combination of m_c and $\sin(i)$.

We also performed a traditional TEMPO2-based χ^2 gridding using all available GBT and Nançay data as a baseline comparison to the PINT-enabled MCMC analysis. This gridding routine (developed by Emmanuel Fonseca and presented in Fonseca et al. 2016) performs a χ^2 minimization using either TEMPO or TEMPO2³ for various

³We chose to use TEMPO2 because of its ability to handle clock corrections unique to the NRT.

Table 4.1: Summary of White Noise Parameters Derived with `Enterprise`

Backend & Receiver	EFAC (unitless)	EQUAD μs
GBT 350 MHz	1.242	0.04293
GBT 820 MHz (2009)	1.995	0.07266
NRT L-band (2009)	2.517	0.20958
GBT 820 MHz (2019)	0.904	3.43973
NRT L-band (2011)	1.165	0.74474

Table 4.2: Mass and Inclination Confidence Intervals from χ^2 Gridding with `TEMPO2`

Parameter	1- σ CI	2- σ CI
m_c (M_\odot)	$0.128^{+0.0573}_{-0.0364}$	$0.128^{+0.171}_{-0.0621}$
i (degrees)	$85.407^{+2.384}_{-4.274}$	$85.407^{+3.819}_{-12.342}$
m_p (M_\odot)	$0.751^{+0.547}_{-0.306}$	$0.751^{+1.523}_{-0.486}$

combinations of companion mass and inclination angle (see Figure 4.3). This yielded the constraints on m_c , m_p , and i presented in Table 4.2.

The χ^2 minimization favors a high inclination angle (~ 85 degrees), relatively low companion mass ($\sim 0.13 M_\odot$) and a correspondingly low pulsar mass. The 1- σ confidence interval for m_p is $0.751^{+0.547}_{-0.306}$ — a broad range that encompasses unreasonably low masses (well below $\sim 1 M_\odot$). The measured companion mass (1- σ confidence interval between 0.09 and $0.19 M_\odot$) gives us pause for reasons that are outlined in Section 4.6.1; however, the 2- σ confidence interval encloses a range of reasonable companion mass values.

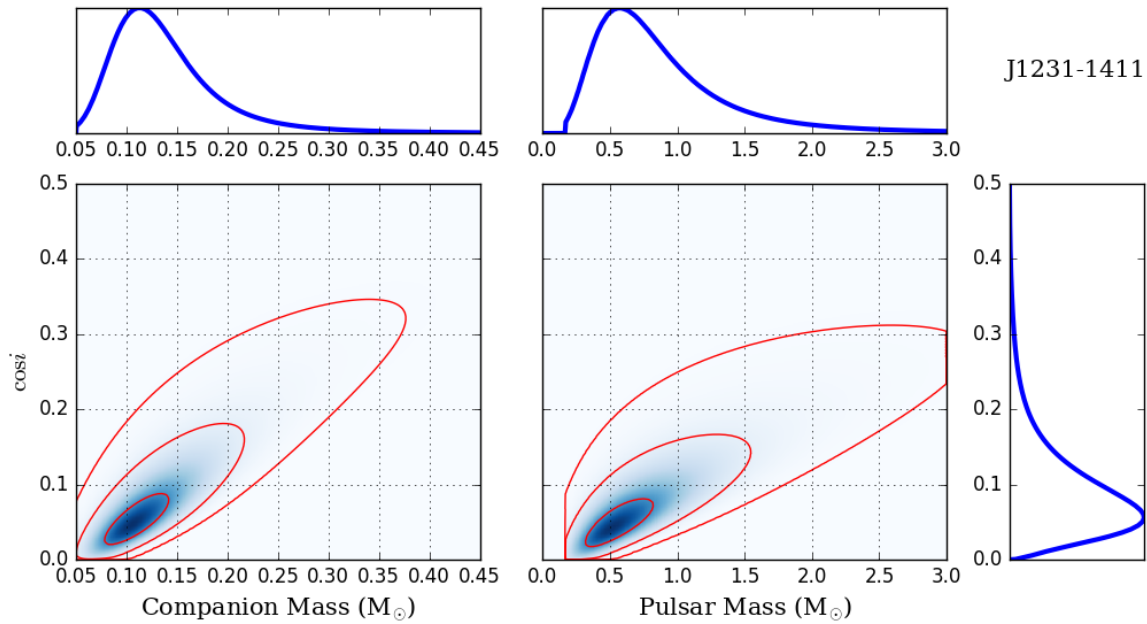


Figure 4.3: The left-hand heat-map shows the results of a χ^2 minimization routine over $\cos(i)$ and m_c . Darker blue areas indicate lower values of χ^2 . Red contours denote 1, 2, and 3- σ significance regions. The right-hand map is a translation of the $\cos(i)$ - m_c probability density function into m_p space using the binary mass function. Probability density functions denoted by blue lines on the top and side of the plot are projected from the χ^2 maps.

Table 4.3: PSR J1231–1411 Best-Fit Parameters from Linear Regression

Pulsar name.....	J0740+6620
Right ascension (degrees).....	12.51980922(2)
Declination (degrees).....	−14.1954561(5)
Proper motion in RA (mas yr ^{−1}).....	−61.6(2)
Proper motion in DEC (mas yr ^{−1}).....	6.9(3)
Spin frequency, ν (Hz).....	271.453019624396(6)
Spin frequency derivative, $\dot{\nu}$ (s ^{−2}).....	−1.66712(3) × 10 ^{−15}
Dispersion measure, DM (pc cm ^{−3})*	8.0877(4)
Binary model	ELL1
Projected semi-major axis of orbit, x (lt-s)	2.0426271(2)
Binary orbital period, P_b (days)	1.86014388501(8)
Epoch of ascending node, TASC (MJD)	55015.15346495(9)
EPS1 (first Laplace-Lagrange parameter), $e \sin \omega$..	−1.2(2) × 10 ^{−6}
EPS2 (second Laplace-Lagrange parameter), $e \cos \omega$	1.1(2) × 10 ^{−6}

4.6 PINT MCMC trials

Bayesian approaches to pulsar timing have only recently come into vogue; however, they offer a number of benefits over traditional pulsar timing techniques (see, for example, Vigeland & Vallisneri 2014). Pulsar timing models are complicated, necessitating the inclusion of parameters to describe astrometry, the behavior of the interstellar medium, the spin characteristics of the pulsar, and many other phenomena. The ability to include priors based on our physical understanding of timing model parameters is a powerful improvement on typical least-squares techniques. This is of particular interest to astronomers wishing to conduct pulsar timing fits using complementary data sets (for example, one could use astrometric measurements from Gaia to constrain a pulsar’s parallax or proper motion). Additionally, Bayesian inference provides more thorough estimates of parameter uncertainties; the presentation of error distributions that are not necessarily Gaussian in nature is useful, as is the ability to visualize the many correlations and degeneracies that occur between model parameters. The PINT software package employs `emcee`, an open-source Python implementation of an affine-invariant ensemble sampler for MCMC (Foreman-Mackey et al. 2013).

The wealth of data available for J1231–1411 has resulted in an extremely precise timing model for the source. Therefore, we are principally concerned with the estimation of the more poorly measured Shapiro delay parameters, especially in light of our 2019 GBT observations targeting superior conjunction. We present six different trials, each of which incorporates new priors on $\sin(i)$ and m_c .

4.6.1 Priors on m_c and $\sin(i)$

The evolutionary end-point of a low-mass white dwarf binary (that is, the system’s total mass and orbital velocity) can be directly predicted from the progenitor red giant star’s degenerate core mass and radius⁴ (Refsdal & Weigert 1971). In 1999, Tauris & Savonije used this fact to formulate a relationship between the binary orbital period of a pulsar system and the white dwarf companion’s mass (TS99; Tauris & Savonije 1999)⁵:

$$\frac{M_{WD}}{M_{\odot}} = \left(\frac{P_b}{b} \right)^{(1/a)} + c. \quad (4.4)$$

Values of a , b , and c are dependent on the chemical composition of the companion: Pop. I stars (young and metal-rich; $Z = 0.02$) have ($a = 4.50$, $b = 1.2 \times 10^5$, $c = 0.120$); the median Pop. I/II (medium- Z) values are (4.75, 1.1×10^5 , 0.115); and the Pop. II (old, metal-poor; $Z = 0.001$) values are (5.00, 1.0×10^5 , 0.110). This relationship holds for companions with masses between 0.18 and 0.45 M_{\odot} . A recent effort (Mata Sánchez et al. 2020) has been made to extrapolate this relationship to less massive white dwarfs using the relationships outlined in Istrate et al. (2016).

For J1231–1411’s 1.86-day orbit, the TS99 Pop. I, Pop. I/II, and Pop. II values are 0.205, 0.214, and 0.223 M_{\odot} , respectively. The priors on m_c for each MCMC trial (described later in this section) are based on results from TS99.

The most agnostic prior on inclination angle is that which describes a random distribution of inclinations; that is, a “flat in $\cos(i)$ ” distribution. The prior in $\sin(i)$ -space is therefore:

$$p(z) = \frac{z}{(1 - z^2)^{1/2}} \quad (4.5)$$

⁴The star’s metallicity (Z) is the biggest predictor of this mass-radius relationship.

⁵The authors have been involved in re-visiting the famous relationship more recently; see, e.g. Istrate et al. (2014, 2016)

for values of $z \equiv \sin(i)$ between 0 and 1. This prior is applied to all MCMC trials, for both radio and Gamma-ray data.

4.6.2 Results of Six MCMC Trials for Radio Observations

We present the results from six MCMC trials using all available radio data. Each of these trials was run with 256 walkers and 2000 steps, and all priors (with the exception of those on m_c and $\sin(i)$) were Gaussians centered on the best-fit values from a least-squares fit with widths equal to 10 times the reported uncertainty from that fit. Priors on m_c were all centered around the median TS99 value ($0.214 M_\odot$), but vary in width and shape.

- **Trial 1 (*Strict TS99 UB*):** A uniform bounded prior (top hat shape) between the TS99 Pop. I and Pop. II predicted values (0.205 to $0.223 M_\odot$)
- **Trial 2 (*Wide TS99 UB*):** A uniform bounded prior between 0.1868 and $0.2418 M_\odot$ (the TS99 Pop. I to Pop. II range extended by twice their difference on each side)
- **Trial 3 (*Very Wide UB*):** A uniform bounded prior spanning a large range of companion mass values (0.05 to $0.35 M_\odot$)
- **Trial 4 (*Strict Gaussian*):** A Gaussian centered on the TS99 Pop. I/II value with a width of $0.01 M_\odot$
- **Trial 5 (*Mid-Width Gaussian*):** Same as above, but Gaussian width is $0.1 M_\odot$
- **Trial 6 (*Wide Gaussian*):** Same as above, but Gaussian width is $1.0 M_\odot$

For each of the six trials, we present 1 and 2- σ confidence intervals for m_c , m_p , and i (see Table 4.4). So-called “waterfall” plots demonstrate the posterior probability distributions (PDFs) and covariances for each of the 18 parameters used in our fit. We also include parameter-specific PDFs plotted alongside the assigned prior for m_c and $\sin(i)$, as well as the m_p that can be derived from the results of these fits (see Figures 4.4-4.16). In order to facilitate an accurate comparison between trials, the $\sin(i)$ and m_c plots have identical axis ranges and identical numbers of bins.

For accurate results, it is critical that one ensures an MCMC run has been allowed to take a sufficient number of steps to converge completely. The most common procedure involves calculating the autocorrelation time (τ) for each chain in one’s fit⁶. Suggestions for chain lengths vary, though `emcee` recommends values $> 50\tau$. For each run, we have ensured that our integrated autocorrelation times and chain lengths meet that criterion (at the very least; many runs have factors of 1000 or more).

4.6.3 Single-Photon Gamma-Ray Timing

At radio frequencies, it is possible to fold a substantial number of pulses to create an integrated pulse profile from which TOAs can be extracted. Pulsar timing at high energies poses a unique problem: detections of high-energy photons are rare — sometimes, only a few photons from an MSP will arrive on a given day. A recent solution to this problem is so-called “single-photon timing,” an MCMC-based technique that assigns phases to individual photons and computes an overall likelihood in order to fit a timing model to high-frequency data (Pletsch & Clark 2015 provides a formulation of the likelihood and a more thorough description of this technique). This routine is implemented in PINT as `event_optimize`.

⁶A nice example can be found at <https://emcee.readthedocs.io/en/stable/tutorials/autocorr/>; and is based on Goodman & Weare (2010).




Figure 4.4: Triangle plot for the output of the strict TS99 UB trial (number 1). Top-to-bottom (and left-to-right), parameters are: RAJ, DECJ, PMRA, PMDEC, DM, PB, A1, M2, SINI, TASC, EPS1, EPS2, F0, F1, JUMP1, JUMP2, JUMP3, and JUMP4. “JUMP” parameters facilitate a DC offset between datasets, mainly to account for receiver and backend variability. Plots at the right end of each row show posterior PDFs (black histograms) and maximum likelihood values (blue lines). All other plots aid in visualizing covariances between parameters.

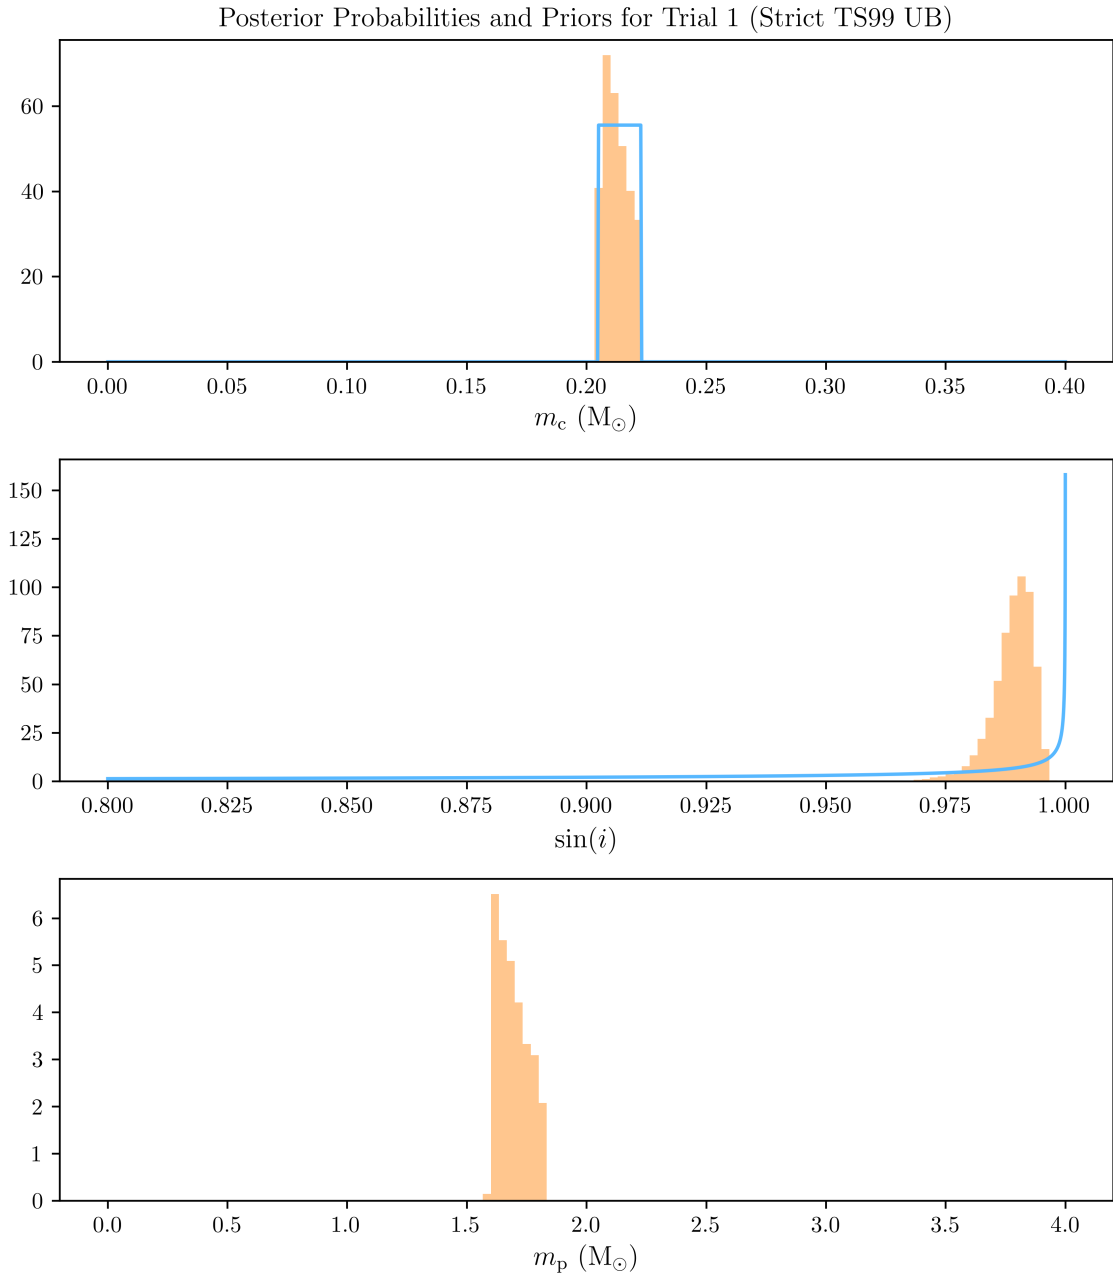


Figure 4.5: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; strict TS99 UB/trial 1) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot.

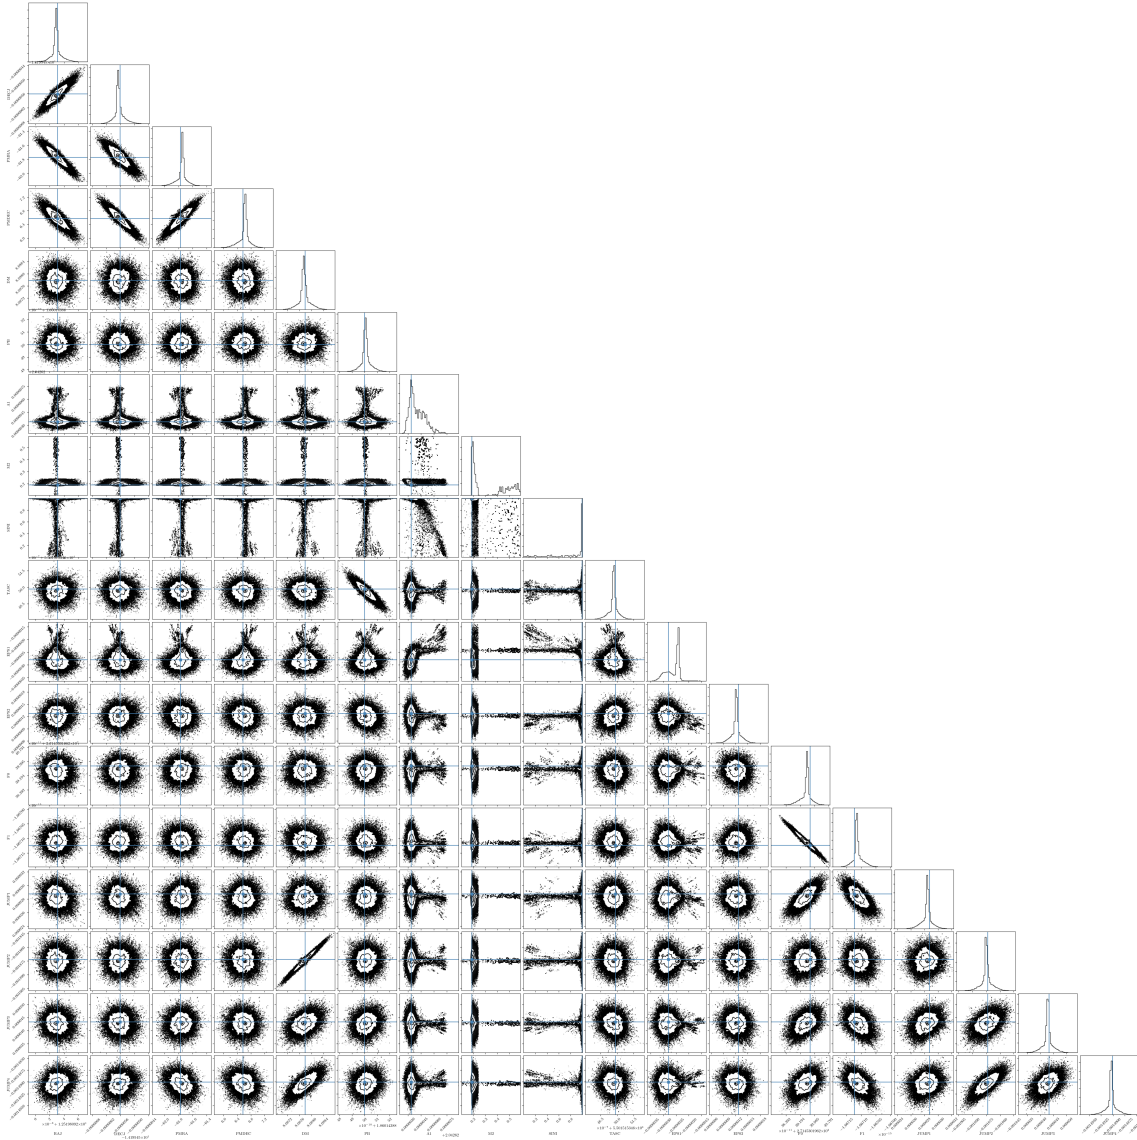


Figure 4.6: Triangle plot for the output of the wide TS99 UB trial (number 2). See Figure 4.4 caption for details.

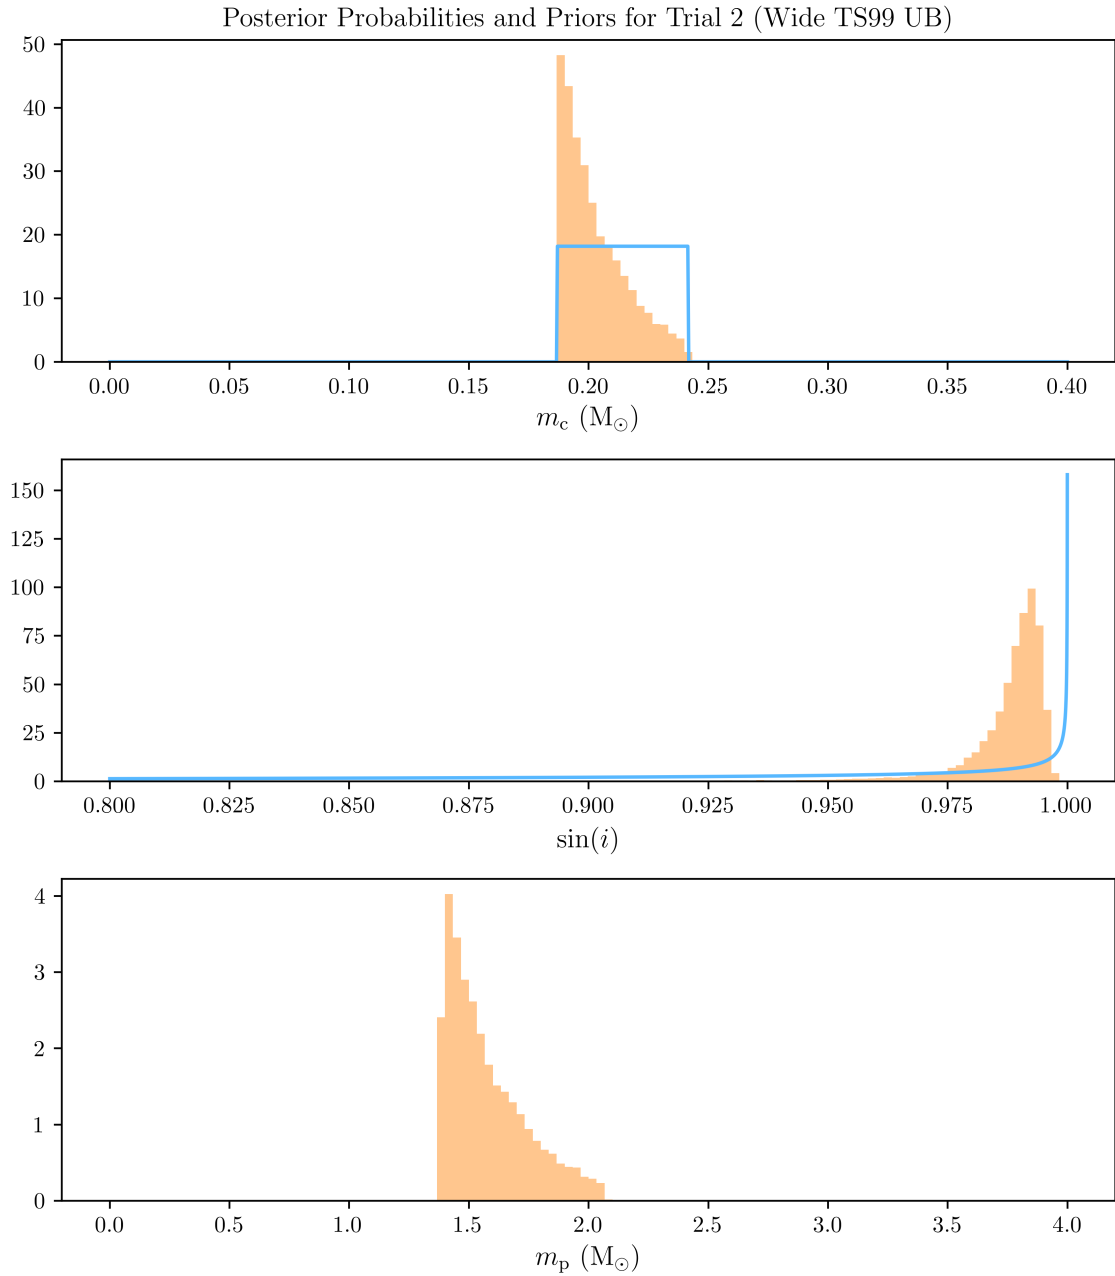


Figure 4.7: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; wide TS99 UB/trial 2) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot.




Figure 4.8: Triangle plot for the output of the very wide UB trial (number 3). See Figure 4.4 caption for details.

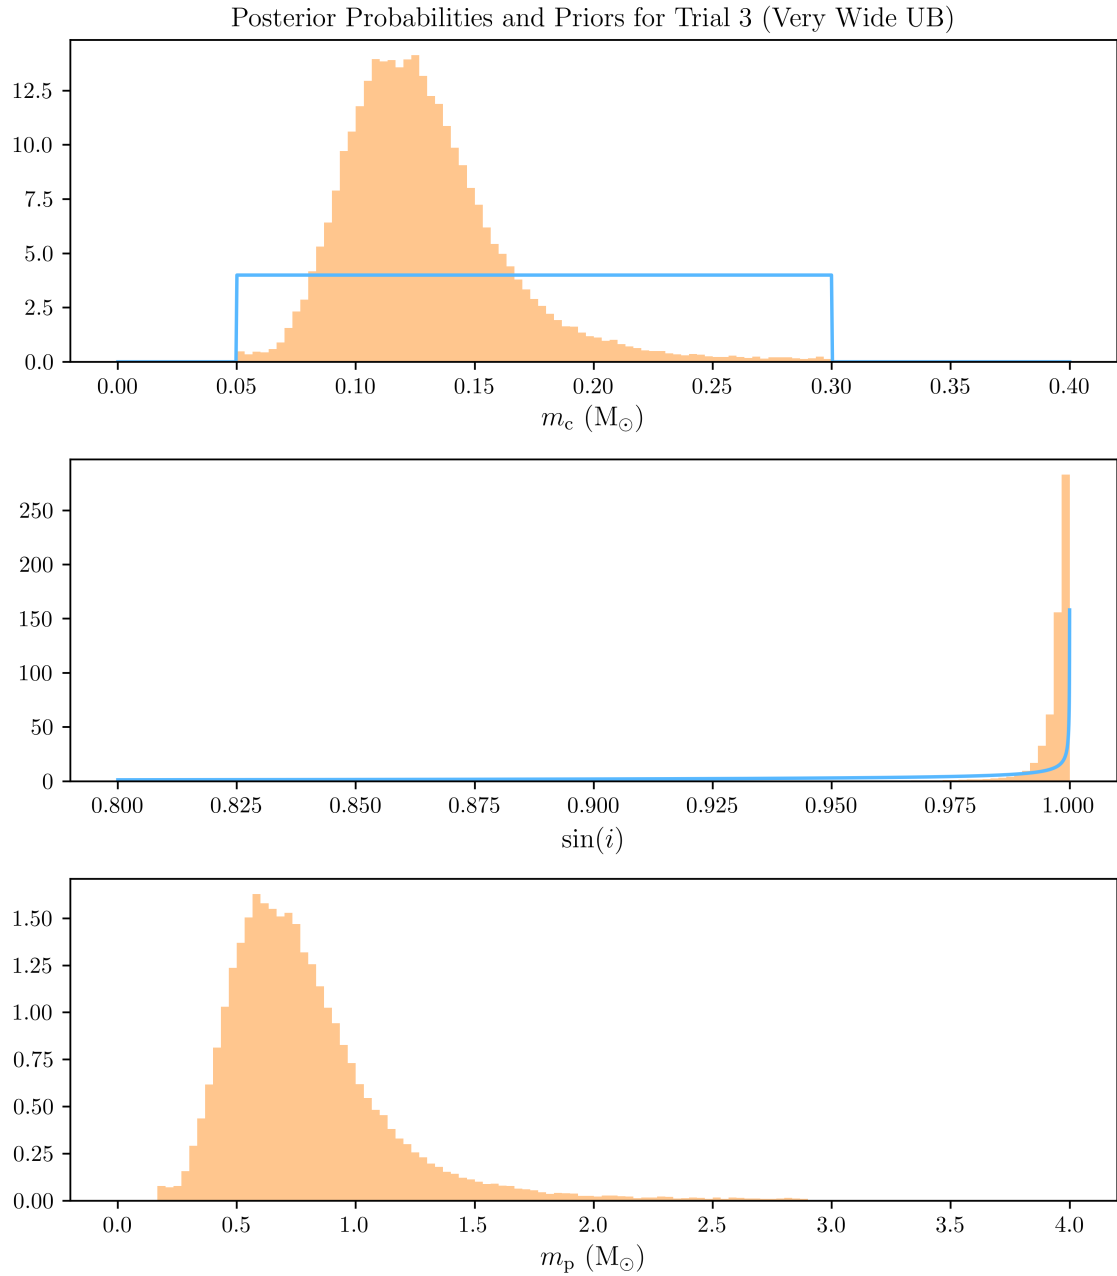


Figure 4.9: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; very wide UB/trial 3) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot.




Figure 4.10: Triangle plot for the output of the Strict Gaussian trial (number 4). See Figure 4.4 caption for details.

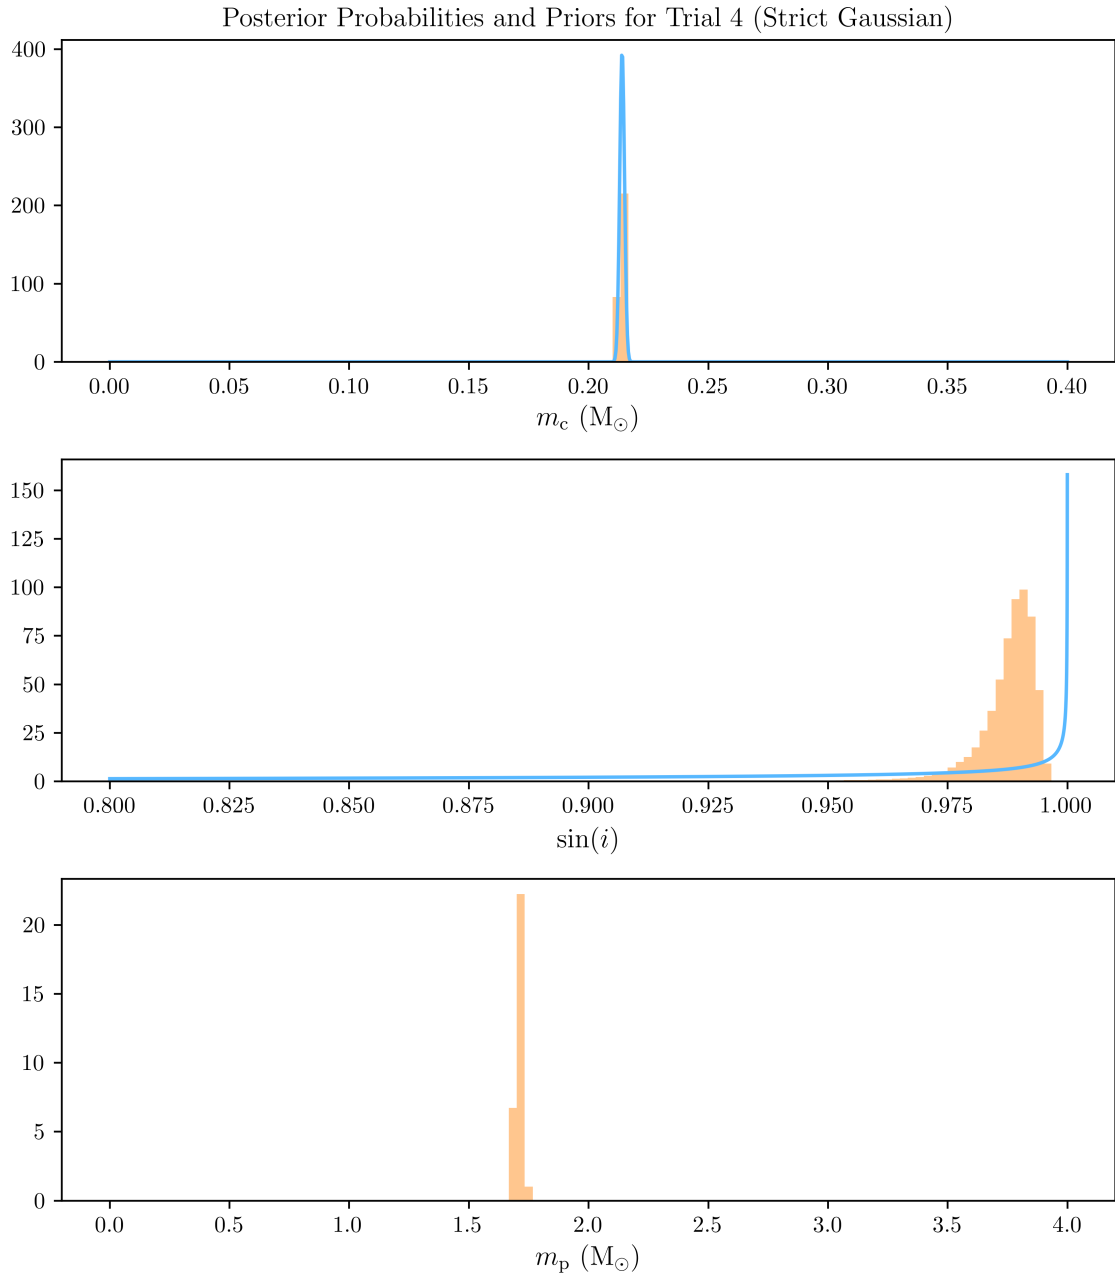


Figure 4.11: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; strict Gaussian/trial 4) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot. See also Figure 4.12.

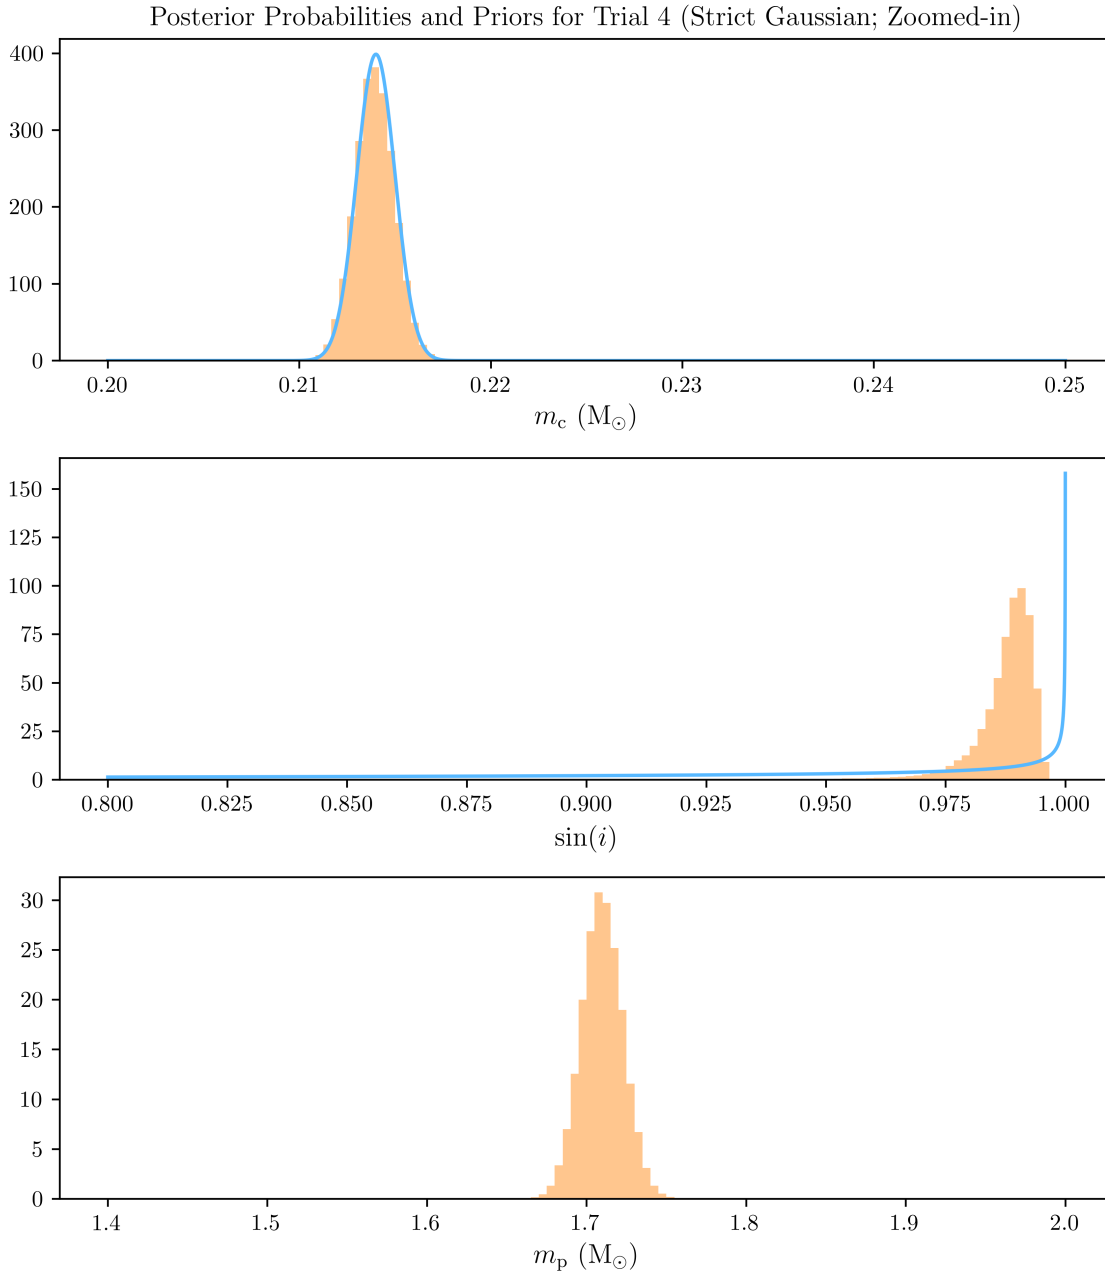


Figure 4.12: A “zoomed-in” version of Figure 4.11 to show detail in the distributions. All other 3-panel PDF plots are presented on the same scale for all three parameters.




Figure 4.13: Triangle plot for the output of the mid-width Gaussian trial (number 5). See Figure 4.4 caption for details.

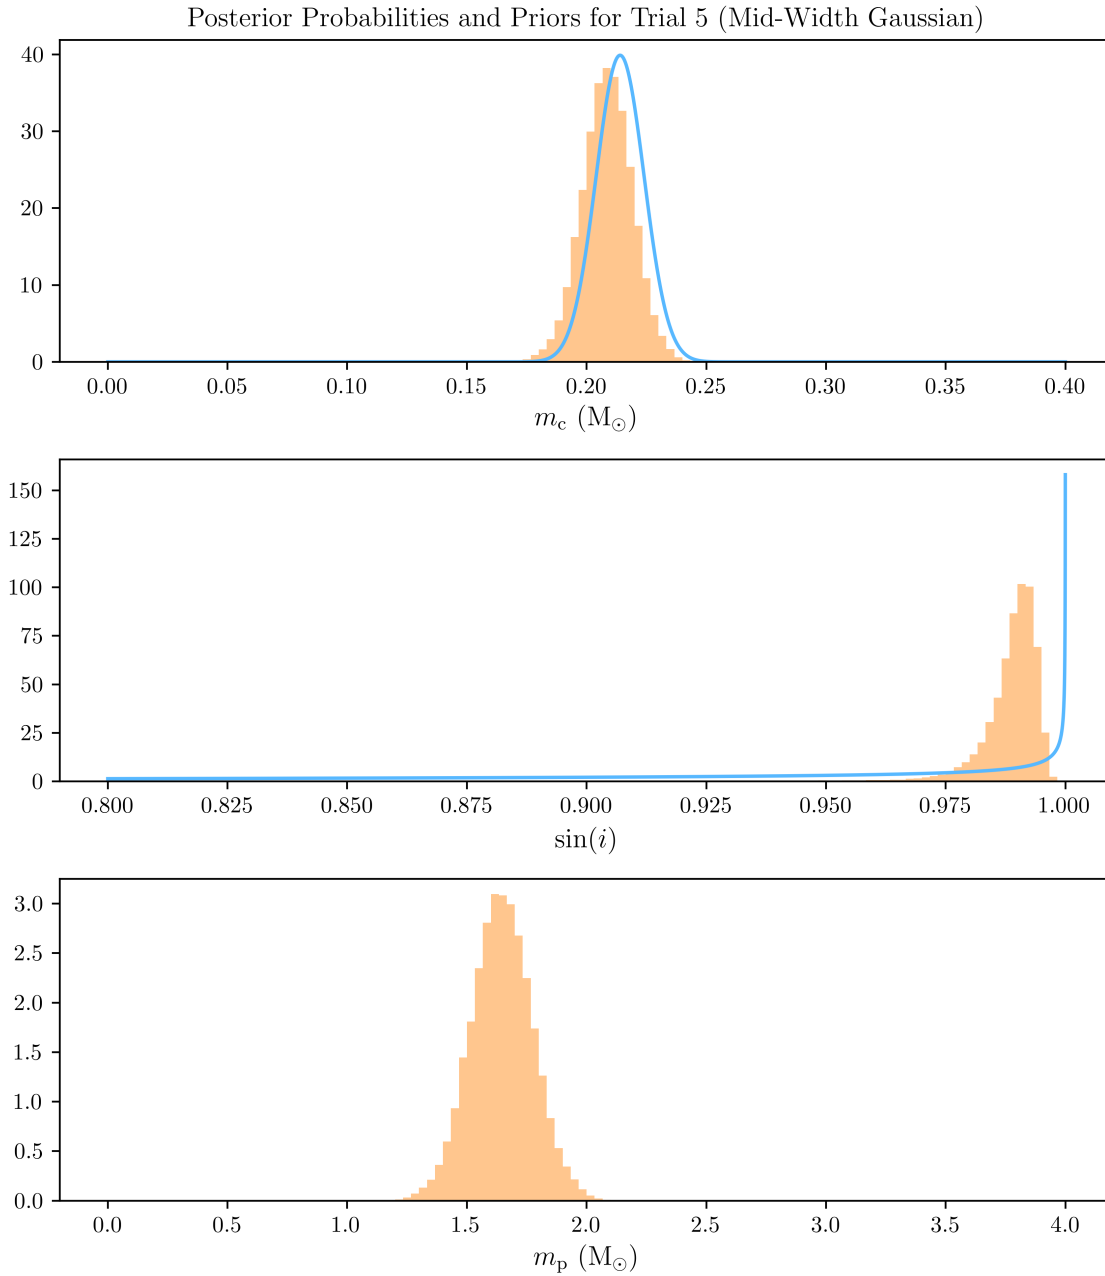


Figure 4.14: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; mid-width Gaussian/trial 5) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot.




Figure 4.15: Triangle plot for the output of the wide Gaussian trial (number 6). See Figure 4.4 caption for details.

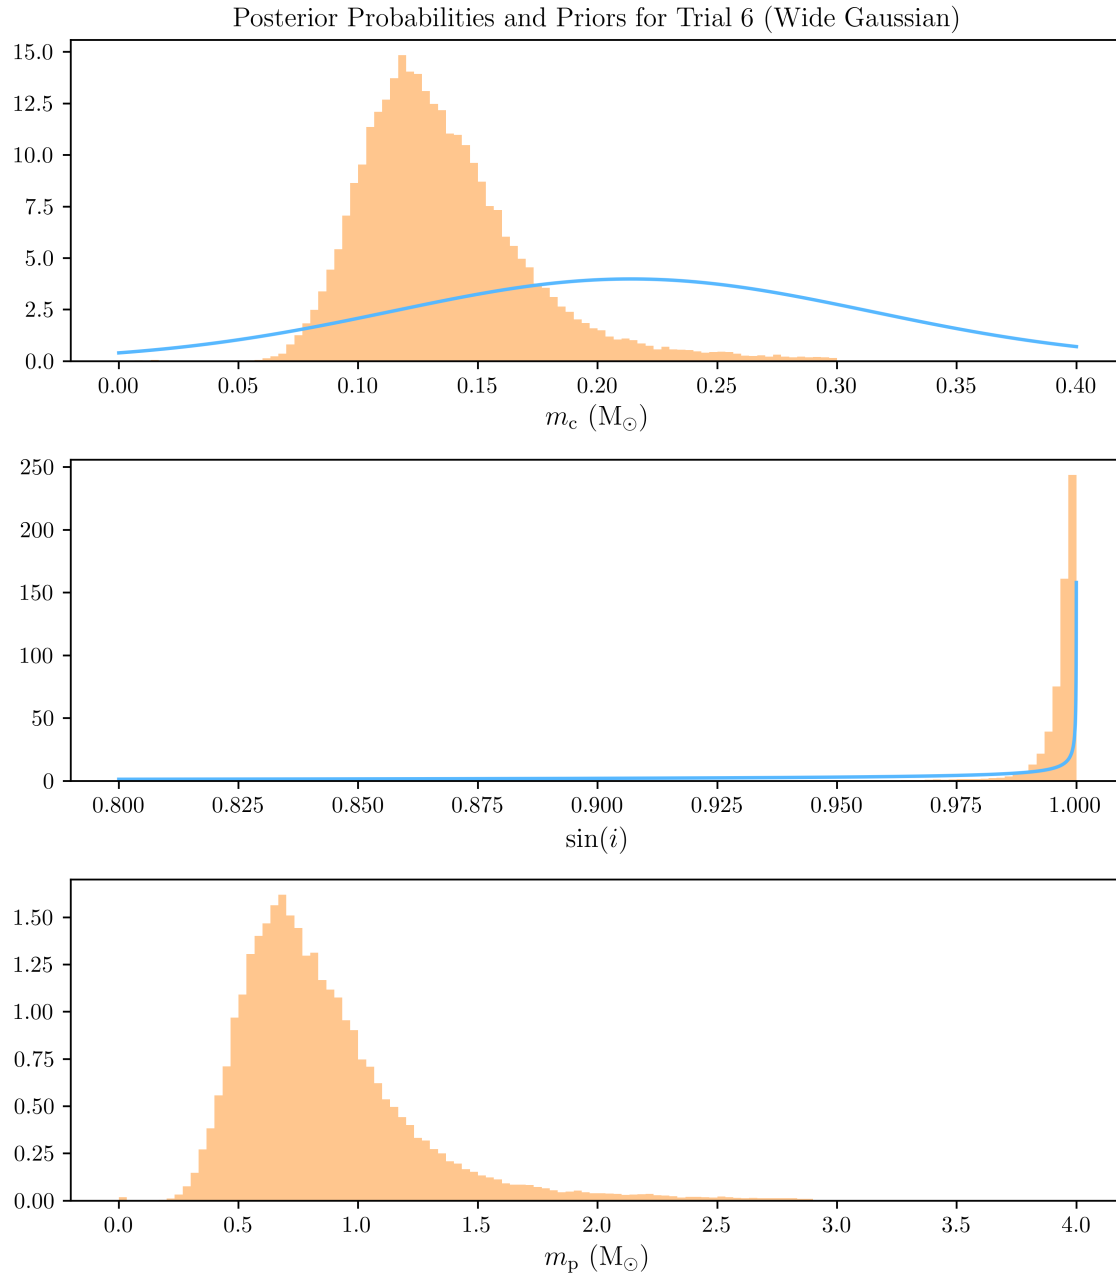


Figure 4.16: This 3-panel plot shows posterior PDFs (orange histograms) and priors (blue lines; wide Gaussian/trial 6) for m_c (top panel) and $\sin(i)$ (middle panel). The resulting posterior PDF for m_p is shown in the bottom panel. Some spurious chains have been removed here, but not in the triangle plot.

Table 4.4: 1- and 2- σ Confidence Intervals on Shapiro parameters for Six MCMC Trials with J1231–1411 Radio Data

Parameter	1 σ CI	2 σ CI
Strict TS99 UB (Trial 1) m_c	$0.221^{+0.318}_{-0.012}$	$0.221^{+0.374}_{-0.015}$
Strict TS99 UB (Trial 1) $\sin(i)$	$0.953^{+0.038}_{-0.620}$	$0.953^{+0.041}_{-0.881}$
Strict TS99 UB (Trial 1) m_p	$1.684^{+0.086}_{-0.061}$	$1.684^{+0.131}_{-0.081}$
Wide TS99 UB (Trial 2) m_c	$0.220^{+0.308}_{-0.028}$	$0.220^{+0.356}_{-0.033}$
Wide TS99 UB (Trial 2) $\sin(i)$	$0.976^{+0.016}_{-0.572}$	$0.976^{+0.019}_{-0.906}$
Wide TS99 UB (Trial 2) m_p	$1.528^{+0.221}_{-0.108}$	$1.528^{+0.449}_{-0.141}$
Very Wide UB (Trial 3) m_c	$0.133^{+0.355}_{-0.031}$	$0.133^{+0.456}_{-0.054}$
Very Wide UB (Trial 3) $\sin(i)$	$0.997^{+0.002}_{-0.201}$	$0.997^{+0.002}_{-0.824}$
Very Wide UB (Trial 3) m_p	$0.723^{+0.333}_{-0.223}$	$0.723^{+1.001}_{-0.384}$
Strict Gaussian (Trial 4) m_c	$0.214^{+0.001}_{-0.001}$	$0.214^{+0.002}_{-0.002}$
Strict Gaussian (Trial 4) $\sin(i)$	$0.989^{+0.004}_{-0.009}$	$0.989^{+0.006}_{-0.560}$
Strict Gaussian (Trial 4) m_p	$1.710^{+0.013}_{-0.013}$	$1.710^{+0.026}_{-0.025}$
Mid-Width Gaussian (Trial 5) m_c	$0.209^{+0.011}_{-0.011}$	$0.209^{+0.024}_{-0.022}$
Mid-Width Gaussian (Trial 5) $\sin(i)$	$0.990^{+0.004}_{-0.008}$	$0.990^{+0.006}_{-0.334}$
Mid-Width Gaussian (Trial 5) m_p	$1.646^{+0.127}_{-0.126}$	$1.646^{+0.256}_{-0.252}$
Wide Gaussian (Trial 6) m_c	$0.132^{+0.042}_{-0.027}$	$0.132^{+0.170}_{-0.049}$
Wide Gaussian (Trial 6) $\sin(i)$	$0.998^{+0.002}_{-0.006}$	$0.998^{+0.002}_{-0.257}$
Wide Gaussian (Trial 6) m_p	$0.777^{+0.367}_{-0.229}$	$0.777^{+1.107}_{-0.392}$

We present the results of an MCMC fit of all available J1231–1411 *Fermi*-LAT data through early 2020. This was conducted using PINT, and implements the same geometric prior on $\sin(i)$ that was used for the radio data. The prior on m_c was uniform between 0 and $2 M_\odot$.

Figure 4.17 presents a waterfall plot showing posterior probability distributions and covariances between all fitted parameters. Figure 4.18 shows the post-fit pulse profile resulting from the single-photon timing phase alignment process. 1- σ con-

fidence intervals for m_c , i , and m_p were $0.587^{+0.644}_{-0.283} M_\odot$, $63.981^{+4.845}_{-12.25}$ degrees, and $\sim 0.28 \pm 1.04 M_\odot$, respectively.

4.6.4 Discussion

Taken together, the linear least-squares gridding, six radio MCMC trials, and independent *Fermi* analysis do not point to a single, obvious conclusion. However, their joint consideration is still valuable in our effort to constrain J1231–1411’s mass.

All three analyses provide reasonable constraints on m_c . On the radio side, the least-squares grid and three PINT MCMC analyses with Gaussian priors (strict, mid-width, and wide Gaussian; or trials 4, 5, and 6) constrain m_c with a relative error ranging between $<1\%$ and $\sim 20\%$. In the wide Gaussian trial (number 6) especially, the posterior PDF makes it clear that m_c is indeed informative when given a comparably large amount of freedom. While the *Fermi* data do not provide a small fractional error for companion mass, the posterior PDFs indicate that m_c is similarly informative in those fits (see Figure 4.17).

The three analyses also constrain i with reasonable success; however, while the χ^2 minimization gridding and PINT MCMC trials prefer a very high inclination, the *Fermi* data do not agree ($63.981^{+4.845}_{-12.25}$ degrees). Only the strict TS99 UB trial (number 1) approaches such a low inclination angle ($72.36^{+2.18}_{-38.3}$ degrees).

It is not clear that the implementation of more strict TS99-based priors on m_c in our PINT MCMC fits was helpful, especially considering that the least constraining priors (very wide UB and wide Gaussian; trials 3 and 6) highlight the constraining power of m_c and suggest much lower companion masses than the TS99 median of $0.214 M_\odot$. This presents a potential astrophysical complication; namely, that a companion mass of ~ 0.13 is inconsistent at the $1\text{-}\sigma$ level with the TS99 relation, even for an




Figure 4.17: Triangle plot for the output of the *Fermi* MCMC fit with PINT. See Figure 4.4 caption for details.




Figure 4.18: The `event_optimize` routine in PINT uses single-photon MCMC-based timing for sparse Gamma-ray data. Pictured is the output from our J1231–1411 *Fermi* analysis.

extremely young, metal-rich Pop. I star (predicted mass = $0.205 M_{\odot}$). However, both the very wide UB and wide Gaussian trials agree with the TS99 prediction within their $2\text{-}\sigma$ confidence intervals. The situation would likely not be dire even if our $2\text{-}\sigma$ mass predictions were not in agreement with TS99. As previously mentioned, Mata Sánchez et al. (2020) and others have extrapolated the TS99 relationship to lower white dwarf mass values, as TS99 cannot account for the several low-mass, tightly orbiting binaries that are known to exist. It seems that a less strict TS99 prior is reasonable given these recent developments. If the collection of further data for J1231–1411 facilitates a more precise mass constraint on m_p , it will be instructive to compare its mass-orbital period relationship with both TS99 and more recent formulations.

Radio Shapiro delay constraints on m_p for the least informative m_c priors agree with our initial joint *Fermi*-NRT modeling: that J1231–1411 is likely a low-mass MSP with a high-inclination (edge-on) orbital plane. The NICER light curve data, which show an extreme flux ratio between the primary and secondary pulses (see Figure 4.2), are provisionally in agreement with this finding. Figure 1, Model IV in Bogdanov et al. (2008) is helpful in visualizing the situation at hand. If J1231–1411 is in a highly inclined orbit with its hot spots and spin axis closely aligned, it follows geometrically that the observer may only see a single, dominant lightcurve peak (see also Bogdanov et al. 2007; Bogdanov 2013). Additionally, in this configuration, the flux received from the neutron star’s two hot spots could be comparable for a high-mass MSP that causes significant bending in photon trajectories; however, a lower mass MSP (that lacks quite such extreme light-bending) with edge-on orbital geometry could easily have higher flux from one peak than the other.

4.7 Conclusions

We have conducted a multiwavelength analysis of timing data for the Gamma-ray-bright MSP J1231–1411 in order to constrain Shapiro delay and in turn, the pulsar’s mass. This study introduces 22 hours of new GBT observations for the source in combination with archival NRT and GBT data. We find that both a linear least-squares gridding and MCMC fit with an uninformative m_c prior prefer low companion and pulsar masses ($0.132^{+0.042}_{-0.027}$ and $0.777^{+0.367}_{-0.229}$ M_\odot , respectively) and high inclinations ($86.380^{+0.115}_{-0.340}$ and $85.407^{+2.384}_{-4.274}$ degrees, respectively). MCMC-enabled trials with PINT were informed by the TS99 white dwarf mass vs. orbital period relationship, and these results are consistent with predicted values of m_c within their $2\text{-}\sigma$ confidence intervals, even for the trials with less informative priors. Single-photon timing with PINT for 12 years of J1231–1411 timing data yielded a significantly less constraining pulsar mass measurement; however, it will be important to keep in mind potential tension between the radio and Gamma-ray analyses going forward. It is our hope that these constraints, though not on par with the precision of the best Shapiro delay observations, will serve as a helpful prior during NICER’s X-ray lightcurve modeling of J1231–1411. A more precise measurement of the M/R ratio will be a boon to the effort to constrain the neutron star EoS.

4.8 Acknowledgements

The *Fermi* LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the

United States, the Commissariat à l’Energie Atomique and the Centre National de la Recherche Scientifique / Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden. Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d’Études Spatiales in France.

The National Radio Astronomy Observatory and the Green Bank Observatory are facilities of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

The author would like to thank Dr. Emmanuel Fonseca for the gridding code used in Section 4.5.2, and for his help in implementing it during this analysis.

Chapter 5

Miscellanea

5.1 Introduction

The author of this dissertation took a rather winding road through her graduate research, having dabbled in many different millisecond pulsar astrophysics-related projects along the way. The result has been an unwieldy collection of side-projects, some of which are interesting but may not merit publication (for example, developing a method for the only-slightly lossy conversion of 600 TB of 8-bit pulsar data into 2-bit format). The observations presented in Chapters 2, 3, and 4 account for only a portion of the 300+ hours of observing using the GBT and Arecibo telescopes completed by the author while pursuing her doctorate. This chapter endeavors to provide quick summaries of several minor projects conducted contemporaneously with the research thus far presented in this thesis.

Section 5.2 details the discovery of 13 new MSPs in *Fermi* unassociated Gamma-ray sources that took place after the publication of Cromartie et al. (2016). Section 5.3 provides a brief summary of some of the author’s work as a member of the NANOGrav collaboration, including an observing project to study so-called “black widow” MSPs.

Section 5.4 touches on Shapiro delay observing campaigns carried out by the author for sources other than J0740+6620. Lastly, in lieu of an appendix, Section 5.5 provides supplementary figures to Chapter 3 that were not published in the Nature Astronomy manuscript.

5.2 Thirteen New Sources from *Fermi* Gamma-Ray Searches at the GBT and Arecibo

Since the publication of Cromartie et al. 2016 (see Section 2.1), the total number of MSPs discovered in *Fermi* unidentified point sources has steadily increased to more than 90. Incredibly, these objects now comprise about a third of all known MSPs in the Galactic plane! Among the new discoveries were 13 new MSPs found in a 2017 joint GBT-Arecibo *Fermi* search program (GLST091250¹; Ransom, Ray, Deneva, Cromartie et al.), the results of which are largely unpublished. Section 5.2 provides a brief overview of the discoveries, highlights their most interesting characteristics, and discusses current and future timing efforts.

5.2.1 Observations

The GLST091250 joint observing program featured pointings toward more than 100 *Fermi* unassociated Gamma-ray sources, the vast majority of which had not been previously searched. Sources were selected from a preliminary version of the 8-year *Fermi*-LAT unidentified point source catalog (Abdollahi et al. 2020). These thousands of sources were whittled down to 133 according to several criteria: 1. They were far from the Galactic plane ($|b| < 10$ degrees in most cases); 2. The LAT error ellipse was smaller than the 327-MHz beam at Arecibo or the 820-MHz beam at the

¹One source, J1845+02, was first observed during P2860 at Arecibo (PI Deneva).

GBT; and 3. They were not variable sources, and featured MSP-like, steep spectral cut-offs.

Observations at Arecibo employed two twin backends: PUPPI, mentioned in Section 2.1, and GUPPI, mentioned in Section 4.3. GBT observations with the 820-MHz receiver used 2000 channels, 64- μ s sampling, and 200 MHz of bandwidth. For the 327-MHz program at Arecibo, observations used 67 MHz of bandwidth, 2816 channels, and 81.92- μ s sampling. Each pointing lasted between \sim 14 and 30 minutes, which allowed sufficient exposure time without sacrificing efficiency or sensitivity to highly accelerated sources.

5.2.2 Results

Thus far, this search project has yielded 10 new MSPs in “quick look” searches, which search only the first \sim 10 minutes of each observation. Three more MSPs were very recently discovered in full searches of the data (employing the “jerk search” techniques developed by Andersen & Ransom 2018 — see Tabassum et al. 2020, in prep.); however, the quick-looks were useful for the immediate discovery of sources (MSP searches, as indicated in the Introduction, necessitate significant computing power).

Table 5.1 summarizes the properties of the first ten Arecibo and GBT discoveries as measured *at the time of discovery*. A forthcoming publication (Cromartie et al. 2020c, in prep.) will detail the results of a timing campaign for some of these sources. Pulse profiles for all 10 original sources are provided in Figures 5.3 and 5.4. Two of these sources, J1555–2908 and J1304+12, are of particular note.

J1555–2908: This MSP was discovered shortly after observations for GLST091250 began at the GBT, which afforded a small amount of follow-up on this source. Or-

Table 5.1: Summary of Arecibo and GBT Discoveries in *Fermi*-LAT Sources

Name	P (ms)	DM (pc cm ⁻³)	P_{orbit} (days)	Minimum Companion Mass (M_{\odot})
J0843+67	2.85	20.7	7.35	0.28
J1555–2908	1.79	75.9	0.23	0.05
J1727–1609	2.45	93.1	?	?
J1852–13	4.31	45.0	?	?
J1304+12	4.18	14.1	?	?
J1402+13	5.89	8.6	?	?
J1803+13	1.52	59.2	1.57	0.14
J1814+31	2.09	34.6	0.20	0.16
J1845+02	4.31	56.6	5.33	0.21
J1919+23	4.63	38.4	93.44	0.24

NOTE — Parameters reflect the information derived from discovery observations and a limited amount of follow-up; recent timing campaigns have improved our understanding of these sources.

bital parameters were able to be derived quickly thanks to NANOGrav exploratory observations that occurred because of the quality of the source, and our interest in including it in the pulsar timing program. J1555–2908 is a fast MSP ($P = 1.79$ ms) that was found to eclipse briefly over the course of its roughly 5.5-hour orbit (see Figure 5.1). Studying its eclipsing behavior is a valuable endeavor in its own right, as there is still much to learn about MSP eclipse physics. This source is relatively bright and sharply peaked, and could be a good candidate for PTAs.

Initial radio timing facilitated a search for the MSP in the *Fermi*-LAT data set, which was successful (details will be presented in Ray et al. 2020, in preparation). In fact, the ~ 10 years of LAT data with which the search was performed ended up informing the radio timing solution, which was relatively simplistic (e.g. the addition of several spin frequency derivative parameters). This source was found to be extremely energetic, with measured spin-down power (\dot{E}) of 3.1×10^{35} erg/second.




Figure 5.1: J1555-2908 shows brief eclipsing behavior during its 5.5-hour orbit. Opportunities to study clearly defined MSP eclipses do not come frequently; however, these observations can inform our understanding of MSP evolution.

J1304+12: A year-long supplementary timing program of the six Arecibo MSP discoveries commenced in January 2018 (P3227, PI Cromartie). Through monthly monitoring of the MSP, it became possible to obtain an orbital solution for the pulsar. The initial results of our observations were exciting: the 4.17-ms MSP is in a 10.33-day binary with a companion which, by measuring the semi-major axis $x = 16.35$ lt-s, we know to be at least $0.55 M_{\odot}$. Fast-rotating MSPs with such heavy companions are nearly unheard of, and could be informative in our understanding of binary evolution. We were excited to acquire an additional set of observations (a 17-hour targeted Shapiro delay campaign as described in Chapter 3). The observations (PI Cromartie) were successful; however, it has been next-to-impossible to obtain a phase-connected timing solution for this MSP. Once a longer observing baseline is available, the Shapiro observations may prove enlightening. One difficulty in obtaining the timing solution is the MSP’s wildly varying scintillation: the source will be sharply peaked one day, and barely detectable the next.

5.2.3 Future Work

As mentioned, all six of the Arecibo MSPs were timed as part of a 12-month campaign in 2018. The results of this work, including phase-connected timing solutions for these sources, will be documented in a forthcoming publication. Recent multi-band lightcurves have been obtained for J1555–2908; however, analysis has not been completed. The team also hopes to obtain spectra for further optical modeling.

5.3 NANOGrav Timing

Throughout her time in graduate school, the author has worked as a member of the North American Nanohertz Observatory for Gravitational Waves collabora-




Figure 5.2: J1304+12, despite having been subject to ~ 1.5 years of monthly observations, has been difficult to phase-connect. This is due in part to its scintillation behavior, which occasionally renders it barely visible.

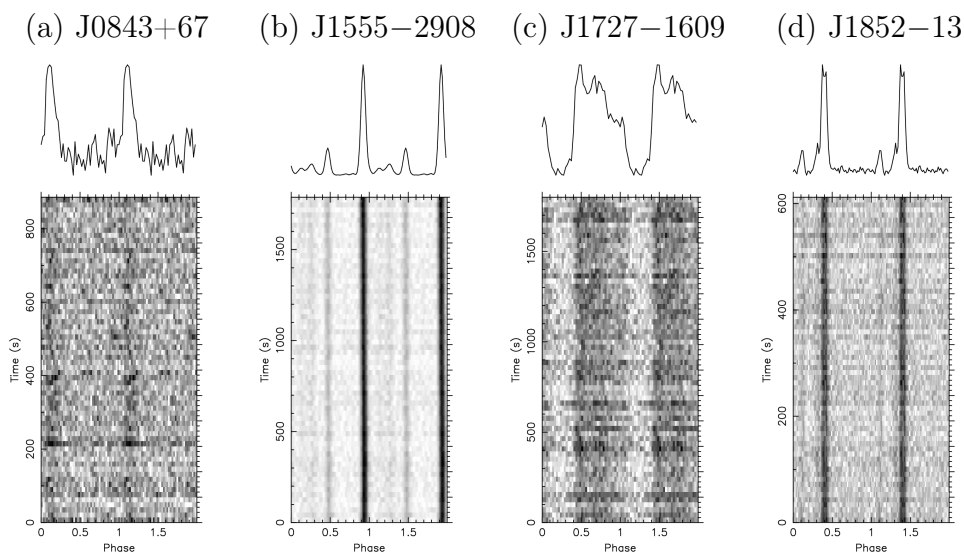


Figure 5.3: Pulse profiles from the four new *Fermi* MSPs discovered with the GBT. All observations presented here were taken at 820 MHz with the GBT.

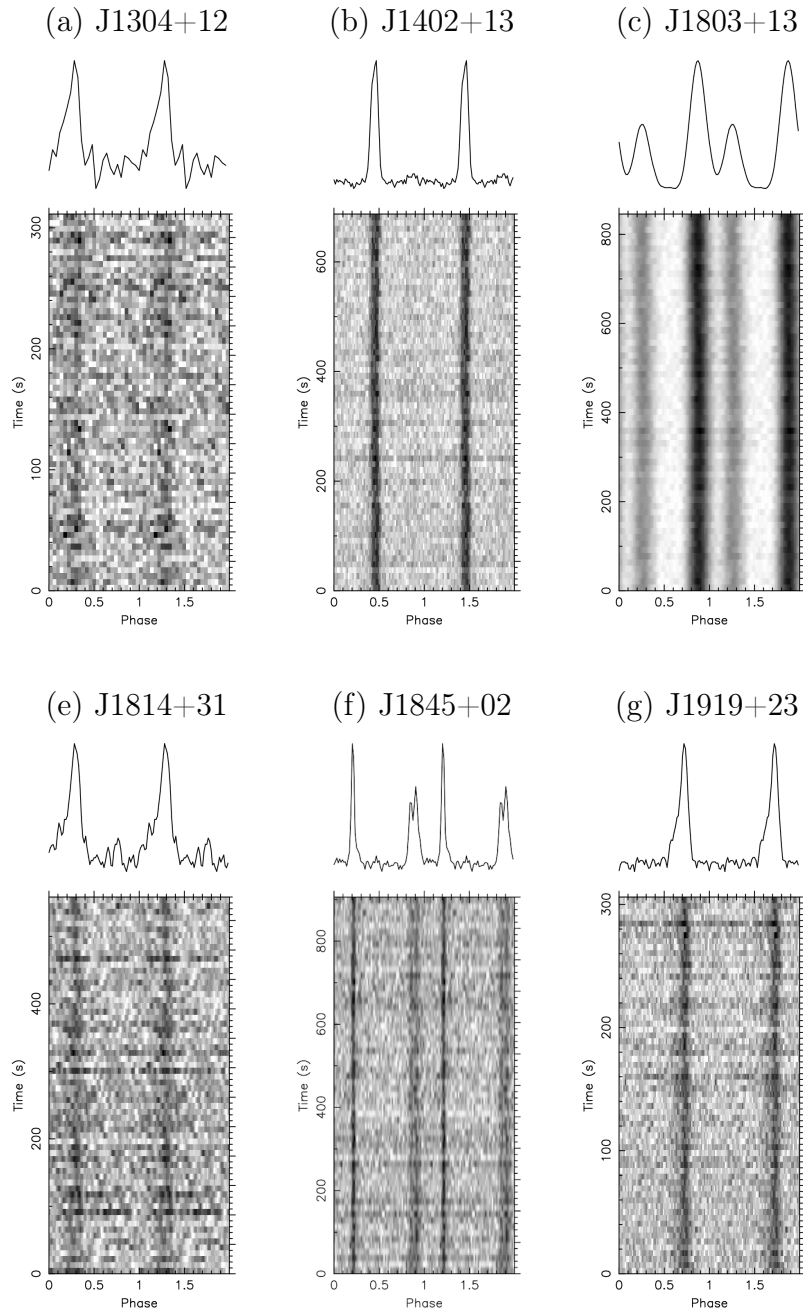


Figure 5.4: Pulse profiles from the six new *Fermi* MSPs discovered with Arecibo.

tion (NANOGrav; see Section 1.7), principally in the Timing and Searching working groups. In addition to contributing to working group activities, NANOGrav members observe in rotation using the GBT and Arecibo telescopes.

Primary contributions to the Timing working group include: assisting in the evaluation and development of MSP timing solutions for data releases (see the upcoming 12.5-year data release and Figure 5.5); serving as PI on a GBT project to detect Shapiro delay in two NANOGrav MSPs (see Chapter 3 and Section 5.4); and serving as PI on two observing campaigns to study “black widow” MSPs in the NANOGrav dataset.

The latter endeavor consisted of two PI observing projects — one at the GBT (GBT17A-230; 44 hours) and one at Arecibo (P3132; 12 hours) — to study six black widow “spider” pulsars. Black widows are MSPs that have obliterated their low-mass binary companions, first by accreting their mass and spinning up to millisecond rotational rates, and then by ablating them with high-energy emission. If the MSP is bright enough, it becomes possible to model these systems over long time periods and evaluate their fitness as PTA pulsars (the only real criterion that presently exists for determining the fitness of a black widow system is whether or not the system eclipses; the vast majority of black widows do). One paper (Bochenek et al. 2015) claims that modeling higher orbital derivatives (as is necessary for black widow timing) will only decrease sensitivity to gravitational waves by $\sim 5\%$. Because the sensitivity of a pulsar timing array to the stochastic gravitational wave background scales linearly with the number of pulsars in the array, including black widow pulsars could be valuable for NANOGrav science. The secondary science related to this project will also be fruitful (studying MSP eclipses and binary evolutionary history). As of this writing, analysis for this project is still underway.

The author has also pursued secondary science projects with NANOGrav col-

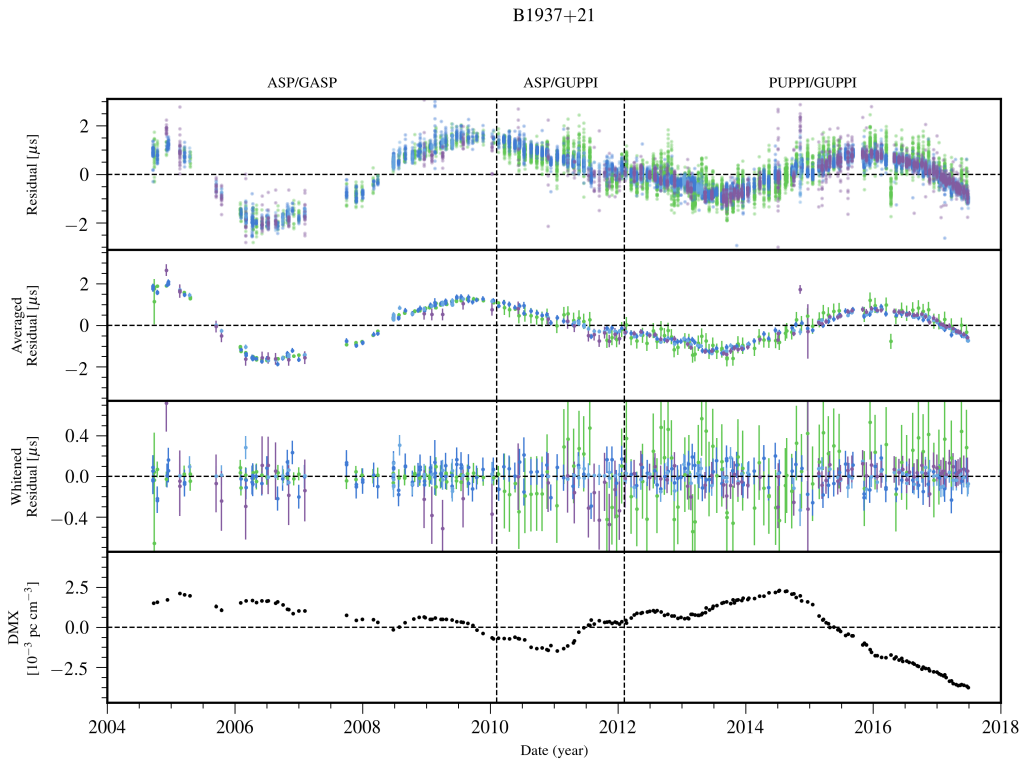


Figure 5.5: As a member of NANOGrav, the author of this dissertation contributes to the timing and analysis of the more than 70 MSPs that constitute the NANOGrav pulsar timing array. This is an example of one of the many plots she created for the upcoming 12.5-year dataset, which shows timing residuals (color-sorted by frequency; top panel), epoch-averaged residuals (second panel), whitened residuals (third panel) and epoch-specific DM measurements (DMX; bottom panel).

leagues, such as searching for fast radio burst signatures in short-duration gamma-ray burst remnants (Madison et al. 2019).

5.4 Other Shapiro Delay Projects

Four separate PI observing campaigns were conducted in order to detect Shapiro delay in four MSP sources. The first (GBT17B-306; 24 hours) consisted of both conjunction-phase and random-orbital-phase observations of two promising Shapiro delay candidates, J0721–2038 and J1400–14. The MSP J0721–2038 was first reported by Burgay et al. (2013) along with 13 other pulsars that were discovered in the Perseus Arm Pulsar Survey using the 64-m Parkes Radio Telescope. It has a relatively long spin period of 15.5 ms and a binary orbital period of 5.5 days. J1400–1431 was discovered at the GBT by high school students in the Pulsar Search Collaboratory in 2012 (see Swiggum et al. 2017), and has a spin period of 3.1 ms and orbital period of 9.6 days.

The second campaign (GBT18B-289; 30 hours) includes targeted conjunction observations of the aforementioned two sources, and adds in two NANOGrav MSPs: J0740 (the most massive neutron star ever observed; see Chapter 3), and J1125+7819 ($P = 4.2$ ms, $P_b = 15.4$). Both J0740+6620 and J1125+7819 were discovered in the Green Bank North Celestial Cap Survey and added to the NANOGrav pulsar timing array. GBT18B-289 included both targeted conjunction observations and randomly sampled pointings. Though J0740+6620 largely shielded our team from disappointment, we have not made a significant detection of Shapiro delay in the other three sources. For J1125+7819 especially, we suffered from bad luck in the (randomly scheduled) orbital phase samples dealt to us during the NANOGrav observing program (see Figure 5.6).

The third and fourth campaigns are the J0740+6620 DDT described in Chap-




Figure 5.6: Orbital coverage for NANOGrav MSP J1125+7819 at 820 MHz: sensitivity to Shapiro delay is strongly improved by sampling on both sides of the peak (at orbital phase = 0.25); we have been unfortunate in the random scheduling of observations.

ter 3 (GBT18B-372; 10 hours) and the J1304+12 Arecibo campaign (described in Section 5.2 (P3308; 17 hours).

We are currently obtaining \sim daily observations of J0740+6620 with CHIME as mentioned in Chapter 3, which are expected to significantly improve our constraint on the source’s mass. Additionally, we will continue to monitor our constraint on J1125+7819’s mass as NANOGrav gathers more data (and therefore improves its orbital phase sampling).

5.5 Supplemental Figures to Chapter 3

Here, we present two figures that were not included in the original Nature Astronomy publication detailing our mass measurement of the MSP J0740+6620.

Figure 5.7: It is instructive to visualize the mass of J0740+6620 in the context of other high-mass neutron stars, and to obtain a more comprehensive understanding of how each mass measurement can easily rule out EoS proposals.

Figure 5.8: This is a demonstration of the capability of our simulation techniques presented in Section 3.1.3.5. By comparing a host of different observing techniques (varying time requests, Shapiro delay curve sampling, etc.), we were able to effectively argue for the additional 10 hours of DDT time at the GBT. In the end, we saw that one five-hour observation in one of the Shapiro troughs in addition to one extra five-hour conjunction observation would most improve our mass measurement precision.

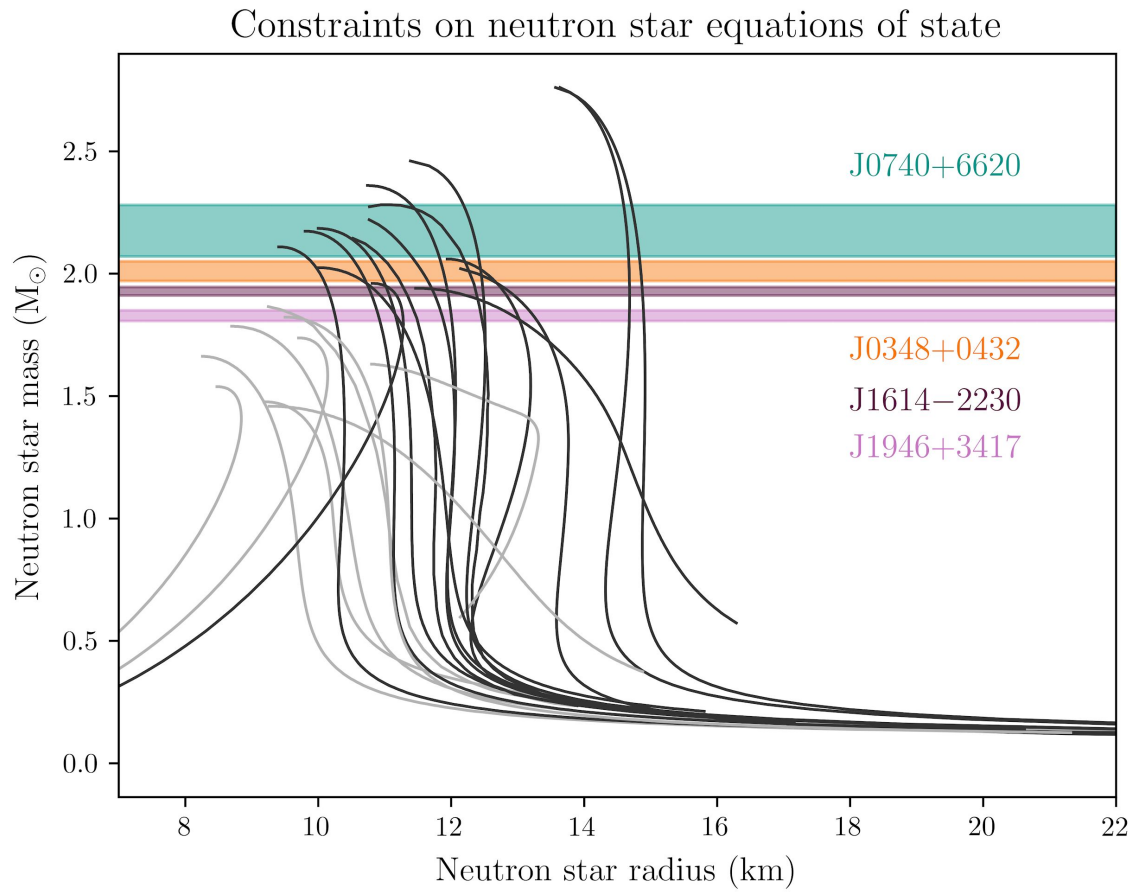


Figure 5.7: The discovery of a $2.14\text{-}M_{\odot}$ neutron star calls several EoS into question. Various EoS are represented by black and grey lines, while bars of color (J0740+6620 in teal) indicate the minimum mass for which an EoS must account (because a neutron star of that mass has been observed).




Figure 5.8: Employing the same simulation techniques described in Section 3.1.3.5, we carefully considered the most effective observing strategy to pursue for the 10-hour DDT request. The width of each Gaussian indicates the resulting 1- σ uncertainty obtained using a given observing scheme. Dotted lines indicate the bottom of the 3- σ confidence interval. Overall, we found that the proposal to observe for five hours in a Shapiro trough and five hours at conjunction would be the best use of DDT time.

Chapter 6

Summary & Closing Remarks

This dissertation explores a diversity of topics within the field of pulsar astrophysics; however, its goal is ultimately singular: to improve our understanding of fundamental physics using millisecond pulsars (MSPs) as tools. With extraordinarily steady spin frequencies and remarkable timing precision at radio, Gamma-ray, and X-ray frequencies, MSPs have repeatedly shown themselves to be among the Universe’s most powerful physical probes.

In Chapter 2, we reported the discovery of the first six MSPs in a search of 34 *Fermi* unassociated Gamma-ray sources with the Arecibo telescope. Before Cromartie et al. (2016), Arecibo was the only major radio observatory to have not contributed to the surge of new MSP discoveries in *Fermi* sources (which now account for $\sim 1/3$ of all known Galactic MSPs). Of the six new sources, five were so-called “spider” MSPs in highly accelerated interacting binary systems. Though this was initially surprising, a brief radio sensitivity comparison with similar *Fermi* unassociated source searches led us to the conclusion that Arecibo’s gain was responsible for the abundance of spiders. Sensitivity to binary systems decreases with increasing observing time (as Doppler smearing causes power to be lost to adjacent Fourier bins in accelerated systems);

therefore, the shorter integration times afforded to us by Arecibo’s substantial gain was advantageous in finding accelerated binaries.

Since 2016, dozens more *Fermi* MSPs have been discovered in searches of unassociated Gamma-ray sources, including 13 new sources from a joint search with the GBT and Arecibo telescopes. Searches for new MSPs are critical for several reasons, including improving the sensitivity of pulsar timing arrays to gravitational waves and discovering systems that could serve as powerful probes of general relativity and nuclear physics.

Chapter 3 details the observation of an unprecedentedly massive neutron star in the NANOGrav dataset (as presented in Cromartie et al. 2020). By combining several years of NANOGrav timing data with supplemental GBT observations (six hours over conjunction followed by an additional five over conjunction and five in one of the Shapiro “troughs”), we measure the mass of the MSP J0740+6620 to be $2.14^{+0.10}_{-0.09} M_{\odot}$. J0740+6620 is the first neutron star whose entire $1\text{-}\sigma$ mass credibility interval is greater than $2 M_{\odot}$.

The discovery of the first two 2-solar-mass neutron stars by Demorest et al. (2010) and Antoniadis et al. (2013) forced a fundamental reconsideration of the supposedly canonical $1.4 M_{\odot}$ neutron star. Similarly, our discovery has significant implications for our understanding of the neutron star equation of state. While we allude to our result ruling out certain non-nucleonic equation of state formulations (i.e. soft equations of state involving quark matter, hyperons, and Bose-Einstein meson condensates), many subsequent publications have explored the way in which a $2.14\text{-}M_{\odot}$ neutron star might affect various aspects of fundamental physics. For example, Zhou & Chen (2019) argue that the J0740+6620 result rules out the possibility of neutron star cores being composed purely of neutrons, as its mass is not compatible with supersoft high-density symmetry energy. The second-most-massive neutron star, however, does

not rule out that theory. Dietrich et al. (2020) conducted a multimessenger analysis using our result in combination with LIGO’s detection of GW170817 and NICER observations of J0030+0451 to constrain both the neutron star equation of state and the Hubble constant.

Going forward, we anticipate improvements in our measurement of the mass of J0740+6620 (and quantitatively verify that expectation) from \sim daily observations with the Canadian Hydrogen Intensity Mapping Experiment telescope (CHIME).

In the second-to-last chapter, we present the results of a multiwavelength pulsar timing analysis of J1231–1411, the brightest Gamma-ray MSP. Forthcoming X-ray lightcurve modeling with NICER will constrain the mass-to-radius ratio of the neutron star, thereby helping to constrain the neutron star equation of state. We were awarded 22 hours on the GBT to conduct orbital-phase-specific timing observations for relativistic Shapiro delay detection in hopes of further constraining the equation of state with an independent mass measurement. These data were analyzed along with archival data from both the Nanc ay and Green Bank radio telescopes. We performed least-squares regression gridding and MCMC-based analyses using priors informed by white dwarf orbital evolutionary theory to constrain the mass of J1231–1411. We also conducted a single-photon MCMC fit for 12 years of *Fermi* Large Area Telescope data. The radio data suggest that J1231–1411 is a low-mass MSP in a highly inclined binary orbit; however, the *Fermi* data are somewhat in tension with that conclusion.

Lastly, in Chapter 6, we summarize a host of minor and ongoing projects undertaken during the author’s time as a graduate student. These include new searches for MSPs in *Fermi* sources, attempts at detecting relativistic Shapiro delay in more MSPs (including another NANOGrav source), and a project to study highly accelerated black widow MSPs for inclusion in the NANOGrav pulsar timing array.

Radio pulsar astrophysics is on the brink of a technological and scientific revo-

lution. With next-generation instruments (such as the Square Kilometer Array and the Next-Generation VLA) poised to come online in the coming decades, we eagerly await the discovery of exotic MSP systems that could further our understanding of nuclear and relativistic physics. Until then, we will settle for NANOGrav's imminent detection of the nanohertz-frequency gravitational wave background.

References

- Abdo, A. A., Ackermann, M., Atwood, W. B., et al. 2009, *ApJ*, 699, 1171
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *The Astrophysical Journal Supplement Series*, 188, 405
- Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *The Astrophysical Journal Supplement Series*, 208, 17
- Abdollahi, S., Acero, F., Ackermann, M., et al. 2020, *The Astrophysical Journal Supplement Series*, 247, 33
- Acero, F., Ackermann, M., Ajello, M., et al. 2015, *The Astrophysical Journal Supplement Series*, 218, 23
- Andersen, B. C., & Ransom, S. M. 2018, *The Astrophysical Journal Letters*, 863, L13
- Antoniadis, J., Tauris, T. M., Ozel, F., et al. 2016, arXiv e-prints, arXiv:1605.01665
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Science*, 340, 448
- Archibald, A., Gusinskaia, N., Hessels, J., et al. 2019, in *American Astronomical Society Meeting Abstracts*, Vol. 233, *American Astronomical Society Meeting Abstracts #233*, 228.03
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, *The Astrophysical Journal*, 813, 65
- . 2018, *ApJS*, 235, 37
- Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *The Astrophysical Journal*, 697, 1071

- Baade, W., & Zwicky, F. 1934, Proceedings of the National Academy of Science, 20, 254
- Bagchi, M., Lorimer, D. R., & Wolfe, S. 2013, Monthly Notices of the Royal Astronomical Society, 432, 1303
- Bedaque, P. F., & Steiner, A. W. 2015, Phys. Rev. C, 92, 025803
- Beronya, D. M., Karpova, A. V., Kirichenko, A. Y., et al. 2019, Monthly Notices of the Royal Astronomical Society, 485, 3715
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, PhysR, 203, 1
- Bochenek, C., Ransom, S., & Demorest, P. 2015, The Astrophysical Journal Letters, 813, L4
- Bogdanov, S. 2013, The Astrophysical Journal, 762, 96
- Bogdanov, S., Grindlay, J. E., & Rybicki, G. B. 2008, The Astrophysical Journal, 689, 407
- Bogdanov, S., Rybicki, G. B., & Grindlay, J. E. 2007, The Astrophysical Journal, 670, 668
- Burgay, M., D'Amico, N., Possenti, A., et al. 2003, Nature, 426, 531
- Burgay, M., Keith, M. J., Lorimer, D. R., et al. 2013, Monthly Notices of the Royal Astronomical Society, 429, 579
- Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, A&A Rev., 27, 5
- Camilo, F., Kerr, M., Ray, P. S., et al. 2015, The Astrophysical Journal, 810, 85
- Caraveo, P. A. 2014, Annual Review of Astronomy and Astrophysics, 52, 211
- Champion, D. J., McLaughlin, M. A., & Lorimer, D. R. 2005, Monthly Notices of the Royal Astronomical Society, 364, 1011
- Cognard, I., Guillemot, L., Johnson, T. J., et al. 2011, The Astrophysical Journal, 732, 47

- Cognard, I., Freire, P. C. C., Guillemot, L., et al. 2017b, *The Astrophysical Journal*, 844, 128
- Contopoulos, I., & Spitkovsky, A. 2006, *The Astrophysical Journal*, 643, 1139
- Cordes, J. M., & Lazio, T. J. W. 2002, arXiv e-prints, astro
- Cornish, N. J., & Sampson, L. 2016, *Phys. Rev. D*, 93, 104047
- Cromartie, H. T., Camilo, F., Kerr, M., et al. 2016, *The Astrophysical Journal*, 819, 34
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *Nature Astronomy*, 4, 72
- Demorest, P. B. 2018, nanopipe: Calibration and data reduction pipeline for pulsar timing, ASCL, ascl:1803.004
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, 467, 1081
- Deneva, J. S., Ray, P. S., Camilo, F., et al. 2016, *The Astrophysical Journal*, 823, 105
- Dietrich, T., Coughlin, M. W., Pang, P. T. H., et al. 2020, arXiv e-prints, arXiv:2002.11355
- Duncan, R. C., & Thompson, C. 1992, *The Astrophysical Journal Letters*, 392, L9
- DuPlain, R., Ransom, S., Demorest, P., et al. 2008, *Society of Photo-Optical Instrumentation Engineers Conference Series*, Vol. 7019, *Launching GUPPI: the Green Bank Ultimate Pulsar Processing Instrument (SPIE)*, 70191D
- Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, ASCL, ascl:1912.015
- Fonseca, E., Pennucci, T. T., Ellis, J. A., et al. 2016, *The Astrophysical Journal*, 832, 167
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306

- Freire, P. C. C., & Wex, N. 2010, *Monthly Notices of the Royal Astronomical Society*, 409, 199
- Freire, P. C. C., Bassa, C. G., Wex, N., et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 412, 2763
- Gendreau, K., & Arzoumanian, Z. 2017, *Nature Astronomy*, 1, doi:10.1038/s41550-017-0301-3
- Geringer-Sameth, A., & Koushiappas, S. M. 2012, *Monthly Notices of the Royal Astronomical Society*, 421, 1813
- Gold, T. 1968, *Nature*, 218, 731
- Goldreich, P., & Julian, W. H. 1969, *The Astrophysical Journal*, 157, 869
- Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, 5, 65
- Haslam, C. G. T., Klein, U., Salter, C. J., et al. 1981, *Astronomy and Astrophysics*, 100, 209
- Hessels, J. W. T., Roberts, M. S. E., McLaughlin, M. A., et al. 2011, in *American Institute of Physics Conference Series*, Vol. 1357, American Institute of Physics Conference Series, ed. M. Burgay, N. D'Amico, P. Esposito, A. Pellizzoni, & A. Possenti, 40–43
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- Hobbs, G., & Edwards, R. 2012, *Tempo2: Pulsar Timing Package*, ASCL, ascl:1210.015
- Hulse, R. A., & Taylor, J. H. 1975, *The Astrophysical Journal Letters*, 195, L51
- Istrate, A. G., Marchant, P., Tauris, T. M., et al. 2016, *Astronomy and Astrophysics*, 595, A35
- Istrate, A. G., Tauris, T. M., & Langer, N. 2014, *Astronomy and Astrophysics*, 571, A45

- Jankowski, F., van Straten, W., Keane, E. F., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 473, 4436
- Johnston, H. M., & Kulkarni, S. R. 1991, *The Astrophysical Journal*, 368, 504
- Jones, M. L., McLaughlin, M. A., Lam, M. T., et al. 2017, *The Astrophysical Journal*, 841, 125
- Kerr, M., Camilo, F., Johnson, T. J., et al. 2012, *The Astrophysical Journal Letters*, 748, L2
- Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, *The Astrophysical Journal*, 821, 66
- Lange, C., Camilo, F., Wex, N., et al. 2001, *Monthly Notices of the Royal Astronomical Society*, 326, 274
- Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, *Nature*, 220, 340
- Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *Monthly Notices of the Royal Astronomical Society*, 225, 307
- Linares, M., Shahbaz, T., & Casares, J. 2018, *The Astrophysical Journal*, 859, 54
- Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy*, Vol. 4 (Cambridge University Press)
- Lynch, R. S., Boyles, J., Ransom, S. M., et al. 2013, *The Astrophysical Journal*, 763, 81
- Lynch, R. S., Swiggum, J. K., Kondratiev, V. I., et al. 2018, *The Astrophysical Journal*, 859, 93
- Lyne, A. G., & Graham-Smith, F. 2006, *Pulsar Astronomy* (Cambridge University Press)
- Madison, D. R., Agarwal, D., Aggarwal, K., et al. 2019, *The Astrophysical Journal*, 887, 252
- Margalit, B., & Metzger, B. D. 2017, *The Astrophysical Journal Letters*, 850, L19

- Martin, J., Rea, N., Torres, D. F., & Papitto, A. 2014, *Monthly Notices of the Royal Astronomical Society*, 444, 2910
- Mata Sánchez, D., Istrate, A. G., van Kerkwijk, M. H., Breton, R. P., & Kaplan, D. L. 2020, arXiv e-prints, arXiv:2004.02901
- Ng, C. 2018, in *IAU Symposium*, Vol. 337, *Pulsar Astrophysics the Next Fifty Years*, ed. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, 179–182
- Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, *The Astrophysical Journal Supplement Series*, 199, 31
- Özel, F., & Freire, P. 2016, *Annual Review of Astronomy and Astrophysics*, 54, 401
- Özel, F., Psaltis, D., Narayan, R., & Santos Villarreal, A. 2012, *The Astrophysical Journal*, 757, 55
- Pacini, F. 1967, *Nature*, 216, 567
- . 1968, *Nature*, 219, 145
- Philippov, A., Uzdensky, D. A., Spitkovsky, A., & Cerutti, B. 2019, *The Astrophysical Journal Letters*, 876, L6
- Philippov, A. A., & Spitkovsky, A. 2014, *The Astrophysical Journal Letters*, 785, L33
- Pletsch, H. J., & Clark, C. J. 2015, *The Astrophysical Journal*, 807, 18
- Prestage, R. M., Bloss, M., Brandt, J., et al. 2015, in *2015 URSI-USNC Radio Science Meeting*, 4
- Ransom, S. 2011, *PRESTO: Pulsar Exploration and Search Toolkit*, ASCL, ascl:1107.017
- Ransom, S. M. 2001, PhD thesis, Harvard University
- Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *The Astronomical Journal*, 124, 1788
- Ransom, S. M., Greenhill, L. J., Herrnstein, J. R., et al. 2001, *The Astrophysical Journal Letters*, 546, L25

- Ransom, S. M., Ray, P. S., Camilo, F., et al. 2011, *ApJ L*, 727, L16
- Ransom, S. M., Stairs, I. H., Archibald, A. M., et al. 2014, *Nature*, 505, 520
- Rappaport, S., Podsiadlowski, P., Joss, P. C., Di Stefano, R., & Han, Z. 1995, *Monthly Notices of the Royal Astronomical Society*, 273, 731
- Ray, P. S., Abdo, A. A., Parent, D., et al. 2012, arXiv e-prints, arXiv:1205.3089
- Ray, P. S., Guillot, S., Ransom, S. M., et al. 2019, *The Astrophysical Journal Letters*, 878, L22
- Refsdal, S., & Weigert, A. 1971, *Astronomy and Astrophysics*, 13, 367
- Richards, D. W., & Comella, J. M. 1969, *Nature*, 222, 551
- Roberts, M. S. E. 2013, in *IAU Symposium*, Vol. 291, *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, ed. J. van Leeuwen, 127–132
- Shapiro, I. I. 1964, *Physical Review Letters*, 13, 789
- Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, *Classical and Quantum Gravity*, 30, 224015
- Spitkovsky, A. 2006, *The Astrophysical Journal Letters*, 648, L51
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *The Astrophysical Journal*, 791, 67
- Swiggum, J. K., Kaplan, D. L., McLaughlin, M. A., et al. 2017, *The Astrophysical Journal*, 847, 25
- Tauris, T. M., Langer, N., & Kramer, M. 2011, *Monthly Notices of the Royal Astronomical Society*, 416, 2130
- Tauris, T. M., & Savonije, G. J. 1999, *Astronomy and Astrophysics*, 350, 928
- van Straten, W., Demorest, P., Khoo, J., et al. 2011, *PSRCHIVE: Development Library for the Analysis of Pulsar Astronomical Data*, ASCL, ascl:1105.014
- Vigeland, S. J., & Vallisneri, M. 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 1446

Watts, A., Espinoza, C. M., Xu, R., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 43

Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, *Reviews of Modern Physics*, 88, 021001

Zhou, Y., & Chen, L.-W. 2019, *The Astrophysical Journal*, 886, 52