

TESTING THE LIMITS OF PRECISION PULSAR
TIMING FOR GRAVITATIONAL WAVE
DETECTION

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Ross James Jennings
December 2021

© 2021 Ross James Jennings
ALL RIGHTS RESERVED

TESTING THE LIMITS OF PRECISION PULSAR TIMING FOR
GRAVITATIONAL WAVE DETECTION

Ross James Jennings, Ph.D.

Cornell University 2021

The NANOGrav collaboration and other pulsar timing arrays (PTAs) aim to use pulsar timing to detect low-frequency (nanohertz) gravitational waves. To do so, they have amassed sets of observations of millisecond pulsars spanning many years. Understanding the various processes that effect the time-of-arrival measurements made using these pulsar observations is extremely important for the goal of gravitational wave detection. Conversely, the pulsar observations are an extraordinarily rich source of information about pulsars and their astrophysical environments, including the interstellar medium along the line of sight to each pulsar, as well as gravitational waves.

In this work, we investigate several such processes. Using the *Gaia* second data release, we measure the parallaxes and proper motions of several pulsars, allowing us to determine their distances and velocities within the Galaxy. We determine the form of the pulsar timing signals which would result from hyperbolic gravitational encounters between pulsars and interstellar objects, or from asteroid belts in binary pulsar systems, and outline methods for searching for signals of these forms. We examine in detail the effect of pulse shape changes on time-of-arrival measurements and develop methods for mitigating this effect. Finally, we describe observations of an recent, unexpectedly large pulse shape change affecting a PTA pulsar, and speculate as to its origins and consequences for pulsar timing.

BIOGRAPHICAL SKETCH

Ross James Jennings was born in Highland Park, Illinois on February 20, 1993. He is the son of Kathryn “Kate” Jennings, a dedicated singer and former librarian who currently serves as the Deputy Treasurer of St. Croix County, Wisconsin, and Talbot “Tad” Jennings, a native Chicagoan and industrial-design school graduate who works as facilities manager for the River Falls Housing Authority. His younger sister, Zoë Jennings, is his only sibling. Zoë graduated from Coe College in Cedar Rapids, Iowa, in 2018, with a degree in theater. She currently resides in Minneapolis, Minnesota, where she works for the commercial real estate management company Cushman & Wakefield while pursuing her theater career.

Ross grew up on an inactive farm just outside the unincorporated town of New Centerville, Wisconsin, approximately midway between the towns of Baldwin and River Falls. The farm had once belonged to his maternal grandparents, Ross and Sue Pierson, both Iowa natives, who operated it until their retirement. Kate and Tad purchased the farm from them and moved there shortly after Ross was born. Ross attended the Baldwin–Woodville school district, which meant taking a bus to Baldwin for elementary school and high school, and to Woodville for middle school.

He was a voracious reader, primarily of science fiction and fantasy, where he benefitted from his mother’s experience as a children’s librarian. Besides the *Harry Potter* series, which he and many others of his generation grew up on, he remembers fondly the works of Diana Wynne Jones, Garth Nix, Ursula K. LeGuin, and Lloyd Alexander. He inherited a love of *Star Trek* from his mother, and his oldest friend, Forrester Smith, introduced him to Dungeons and Dragons and other tabletop role-playing games. He and his father built a

PC together in the mid-2000s, an experience which helped kick-start his journey into the world of electronics and computer hardware.

He began his undergraduate studies at Carleton College in the fall of 2011. There, he found himself surrounded by like-minded people, including a number of friends he made on the “quiet floor” where he lived as a freshman who would stay with him for all four years. He was drawn in by Carleton’s outstanding physics department, whose large contingent of students had a rare degree of camaraderie fostered by late nights working on problem sets in the Olin Hall labs, but many of his favorite classes were math classes, and he graduated in 2015 as a math and physics double major. His senior “comps” project in physics was advised by a pulsar astronomer, Professor Joel Weisberg. At the time, neither of them knew that Ross’ graduate studies would eventually lead him back to pulsars and to Joel’s longtime friend Jim Cordes. It was also at Carleton that Ross took up his hobby of swing dancing, which he continues today.

In the fall of 2013, Ross travelled to Hungary for the Budapest Semester in Mathematics, and in the summer of 2014, he participated in a Research Experience for Undergraduates with Professor David Gerdes at the University of Michigan, where he learned the Python programming language and helped to discover several new Kuiper Belt objects using data from the Dark Energy Survey.

He began his graduate career at Cornell in the fall of 2015, and spent some time working with Professor Lisa Kaltenegger in the field of extrasolar planets, before a class in astrostatistics and data mining lead him to Jim Cordes’ research group and the NANOGrav collaboration, which he joined as an associate member in 2018, and as a full member in 2019. An unexpected stumbling

block in his graduate career has been the Covid-19 pandemic, which began in the United States in early March of 2020, and kept him away from the Cornell campus and his fellow students for many months. Fortunately, he was able to continue his thesis work and apply for jobs remotely, and was able to return to campus in late May 2021. He is excited to begin the next stage of his career at West Virginia University in January 2022, where he will join the NANOGrav team lead by Professor Maura McLaughlin.

This document is dedicated to the memory of Phillip Hopkins, a true friend.

He is sorely missed.

ACKNOWLEDGEMENTS

This work would not have been possible without the support and guidance of Jim Cordes. Jim's laid-back personality and many years of experience advising graduate students have made him an excellent guide to the world of academic astronomy, and his advice has never led me astray. I also cannot fail to mention Shami Chatterjee, Jim's right-hand man, who has provided a valuable alternate perspective on many a topic during our group meetings, and has done nearly as much as Jim to keep me on task and help me finish the projects I start. There is also Adam Brazier, our Cyber-I guy, who is a friendly face at our group meetings when he has time to spare from his duties at the Center for Advanced Computing and an excellent traveling companion at NANOGrav meetings.

During my time in Jim's research group, I have had the privilege of working alongside several other graduate students. At the beginning, I shared my corner office with Gianfranco Grillo, a master's degree student, and Shen Wang, a visiting student from China. While neither was easy to talk to at first, sharing our small office and taking road trips to NANOGrav meetings Lancaster, Pennsylvania and Green Bank, West Virginia brought the three of us closer together. The first new PhD student to join our group was Akshay Suresh. Akshay is a force to be reckoned with, and while he has often pushed me outside of my comfort zone, I have just as often come to appreciate the perspectives he offered. I had the pleasure of serving on the AGN executive board with Stella Ocker, still Jim's newest graduate student, in 2019, and have come to consider her a good friend. Most recently, I have been delighted to welcome Thankful Cromartie, who I had met previously through NANOGrav meetings, to Jim's group as a postdoctoral fellow. Thankful came to Cornell during the Covid-19

pandemic, and it is only within the past few months that she has been able to join us for in-person group meetings.

I am also indebted to a number of other graduate students who have made my time at Cornell a better one. That includes my fellow physics graduate students Patrick Hollister, Thomas Oseroff, Ryan Porter, Eric Schwen, and Michelle Kelley, with whom I endured a semester of Jackson E & M, along with other classic classes including thermodynamics, quantum mechanics, and the graduate lab course. It also includes game night and unicorn czar Alex Grant, trivia master Cody Lamarche, mathematical artist Gabe Bonilla, Irishman and numerical relativist Éamonn O’Shea, and monk-to-be Dante Iozzo. Lisa Kaltenegger’s two students, Thea Kozakis and Jack Madden, shared an office just down the hall from mine for years, and I often turned to them when I was in need of the kind of guidance that comes only from fellow grad students. In addition to Stella, my friends on the AGN board, a group which overlaps heavily with the Tuesday night Big Red Barn trivia team, included Abhinav Jindal, Christopher Rooney, Catie Ball, and Chris O’Connor. Then there are my fellow swing dancers, especially Furkat Mukhtarov, Mari Tader, Colin Bundschu, Sierra Cook, and Maddie Reynolds, without whom I would have spent many more Monday nights alone.

Being part of the NANOGrav collaboration and attending NANOGrav meetings has been an immensely rewarding experience, and in the last few years I have benefitted greatly from my interactions with many NANOGravers. In addition to the aforementioned Thankful Cromartie, these include Dusty Madison and Michael Lam, two of Jim’s former students I barely missed overlapping with, but who are now both in faculty positions, and Joe Swiggum and Sarah Vigeland, my fellow Carls in NANOGrav. There are also numer-

ous others: graduate students Haley Wahl, Caitlin Witt, Deborah Good, and Brent Shapiro-Albert; postdocs Jeff Hazboun, Tim Pennucci, Nihan Pol, and Emmanuel Fonseca; and faculty members Scott Ransom, Xavi Siemens, David Kaplan, Anne Archibald, and Maura McLaughlin. There are also two research scientists: Matthew Kerr at the Naval Reserach Laboratory, and Paul Demorest at NRAO in Socorro, the architect of NANOGrav’s TOA generation pipeline. And I can’t leave out Nate Garver-Daniels, who almost single-handedly kept the computers running for years. Although Nate still drops by from time to time, his departure for a new and better job led to the hiring of “Nate replacement” Joe Glaser. While he can never replace Nate in our hearts, Joe G. seems an excellent sort of fellow, and I look forward to working with him at WVU.

I wouldn’t have made it this far without many great teachers along the way. This includes Alan Giambattista, Nick Taylor, and Natasha Holmes in the Cornell physics department, who taught me how to be a good TA, as well as several other Cornell faculty, including Jim Sethna, Liam McAllister, Tom Loredo, Jonathan Lunine, and Alex Hayes. It also includes many of my Carleton professors: Joel Weisberg, Cindy Blaha, Helen Wong, Bill Titus, Jay Tasson, Cherlon Ussery, Eric Egge, and Deanna Haunsperger, to name just a few. Some of my high school teachers should be included, too, particularly Marti Koller, Deb Challe, Jeremy Gerrits, and Tony Stitt (Yay, Mr. Stitt! We love you, Mr. Stitt! You’re the best, Mr. Stitt! Thank you, Mr. Stitt!).

Several members of Cornell’s administrative staff have also been instrumental throughout my graduate career. In particular, I’d like to thank Jessica Jones, Monica Carpenter, Jill Tarbell, Mary Mulvanerton, and Lynda Sovocool in the astronomy department, and Kacey Acquilano and Craig Wiggers in the physics department.

Finally, I'd like to thank my family. My parents, Tad and Kate Jennings, not only supported me throughout my graduate school career, but provided me with a place to stay rent-free during the height of the Covid-19 pandemic, between July 2020 and May 2021, and put up with my stressed-out state throughout that whole period. To my sister, Zoë Jennings, and my surviving grandparents, Sue Pierson and Rebecca Hunter: know that your support is appreciated. And to my grandfather and namesake, Ross Pierson, who passed away in March 2020 at the age of 98: know that you are missed.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	vi
Acknowledgements	vii
Table of Contents	xi
List of Tables	xiv
List of Figures	xv
1 Introduction	1
1.1 Pulsars	1
1.1.1 Interstellar Dispersion and Scattering	7
1.2 Gravitational Waves	14
1.3 Pulsar Timing Arrays	20
2 Binary Pulsar Distances and Velocities from <i>Gaia</i> Data Release 2	27
2.1 Introduction	28
2.2 Methods	31
2.2.1 Identification of Sources	31
2.2.2 Parallax Inversion and Lutz-Kelker Correction	34
2.2.3 Velocity Estimation	36
2.3 Results	37
2.3.1 Notes On Individual Objects	42
2.4 Discussion	44
2.5 Acknowledgements	49
3 Detecting Gravitational Scattering of Interstellar Objects Using Pulsar Timing	51
3.1 Introduction	51
3.2 Pulsar Timing Precision	55
3.3 TOA Perturbations	57
3.3.1 Timing Signature	62
3.3.2 Similar Signals	64
3.4 Detecting Pulsar Encounters	66
3.5 Detecting Earth Encounters	73
3.6 Frequency of Encounters	76
3.7 Summary and Conclusions	79
3.8 Acknowledgements	81
4 Pulsar Timing Signatures of Circumbinary Asteroid Belts	82
4.1 Introduction	82
4.2 Spectral Characteristics	85
4.2.1 Effects of Model Subtraction	89
4.3 Results	91

4.4	Discussion	94
4.4.1	Comparison with Other Low-Frequency Timing Signals	95
4.4.2	Stationarity	97
4.4.3	Survival of Asteroids	98
4.4.4	Distinguishing Individual Asteroids	99
4.4.5	Implications for Gravitaional-Wave Searches	102
4.5	Summary and Conclusions	103
4.6	Acknowledgements	104
4.7	Appendix: Stable Orbits in a Binary System	105
5	The Effect of Pulse Shape Changes on Pulsar Timing Precision	108
5.1	Introduction	109
5.2	Causes of Pulse Shape Changes	111
5.3	The Amplitude-Modulated Shot Noise Model	113
5.4	TOA Estimation	115
5.5	Errors Caused by Shape Differences	117
5.5.1	Component Amplitude and Phase Variations	119
5.5.2	Simulations with Gaussian Components	121
5.6	Tools for Assessing Single-Pulse Stochasticity	126
5.6.1	TOA and Mismatch Statistics	126
5.6.2	Profile Autocorrelation Functions	129
5.6.3	Correlations Between Frequency Channels	131
5.6.4	Principal Component Analysis	132
5.6.5	Application to the Vela Pulsar	137
5.7	State Switching	144
5.8	Instrumental Effects and RFI	145
5.9	Summary and Conclusions	146
5.10	Acknowledgements	149
5.11	Appendix: Periodic Extensions and Fourier-Domain Shifts	149
6	Improving Pulsar Timing by Correcting for Pulse Shape Changes	152
6.1	Introduction	152
6.2	Overview of Methods	154
6.3	PCA and Shift Invariance	158
6.4	Flexible Template Matching Algorithms	161
6.5	Application to Simulated Data	162
6.6	Application to the Vela Pulsar	165
6.7	Summary and Conclusions	167
6.8	Acknowledgements	168
7	Recent Observations of a Pulse Shape Change in PSR J1713+0747	169
7.1	Introduction	169
7.2	Observations	171
7.3	Interpretation	174

7.4	Conclusions	180
8	Future Directions	182
8.1	Searching for Interstellar Objects in Pulsar Timing Data	183
8.2	Distinguishing Asteroid Belts from Timing Noise	184
8.3	Surveying Profile Variability in PTA Data Sets	186
8.4	Characterizing and Mitigating Interstellar Scattering	187
A	The NANOGrav TOA Generation Process	190
A.1	Getting the Software and Data	190
A.2	Setup and Directory Structure	193
A.3	Per-Pulsar Makefiles	195
B	Transmission Functions for Polynomial Fits	202

LIST OF TABLES

2.1	Comparison of <i>Gaia</i> parallaxes with other distance measurements	39
2.2	Distances and velocities derived from <i>Gaia</i> parallaxes	40
4.1	Stability boundaries for some MSP binaries	91

LIST OF FIGURES

1.1	A $P - \dot{P}$ diagram.	4
1.2	Dispersion in PSR J2145–0750.	8
1.3	Dynamic spectrum and secondary spectrum of PSR B0833-45. . .	12
1.4	The Hellings and Downs curve.	24
2.1	Comparison of <i>Gaia</i> parallaxes with previous distance estimates	38
2.2	Comparison of <i>Gaia</i> parallaxes with DM distances	41
2.3	Comparison of <i>Gaia</i> proper motions with timing measurements	42
3.1	Geometry of a hyperbolic orbit	61
3.2	Timing perturbations caused by gravitational encounters	63
3.3	Signal amplitude as a function of velocity	69
3.4	Signal amplitude as a function of periapsis time	70
3.5	Signal amplitude as a function of impact parameter and velocity	71
3.6	Detectability of interstellar objects with PTAs	72
4.1	Power spectrum of TOA perturbations from an asteroid belt . .	85
4.2	Effect of power-law index on TOA residual spectrum	89
4.3	Spectra created by asteroid belts around particular pulsars . .	92
4.4	Example TOA residuals from an asteroid belt	95
5.1	TOA error in simulations with two identical components	122
5.2	TOA error in simulations with several identical components .	123
5.3	TOA error vs. signal-to-noise ratio	127
5.4	TOA error and mismatch for simulated pulses	128
5.5	Comparison of autocorrelation functions for simulated data .	130
5.6	Principal components in various simulated cases	133
5.7	Shape variations corresponding to principal component dot products	134
5.8	TOA error and mismatch for the Vela pulsar	138
5.9	Comparison of autocorrelation functions for Vela pulsar data .	139
5.10	Correlations between TOAs in nearby frequency channels for the Vela pulsar	140
5.11	Principal components for the Vela pulsar	141
5.12	The relationship between pulse shape and TOA estimates for the Vela pulsar	142
5.13	Average profiles for three millisecond pulsars	147
6.1	A comparison of TOA correction methods	163
6.2	The results of PCA-based TOA correction methods applied to the Vela pulsar	166
7.1	Average profiles before and after the event	175
7.2	Overview of the shape change event	176

7.3	Profile residuals around the shape change event	177
7.4	Time of arrival measurements around the shape change	178
7.5	The magnitude of the shape change over time	179

CHAPTER 1

INTRODUCTION

The work in this dissertation is concerned with pulsar timing arrays (PTAs). In particular, it was undertaken in an effort to understand the various signals and sources of noise which might interfere with a PTA's ability to detect gravitational waves. Some of those signals are astrophysically interesting in their own right: PTAs use pulsars as part of a galactic-scale gravitational wave detector, but pulsars themselves are examples of extreme astrophysical systems about which much remains to be learned. In the pages to come, I will describe the measurement of distances to several pulsars using *Gaia* optical astrometry of their binary companions; the potentially detectable signals that would be produced by encounters between pulsars and interstellar asteroids or planets, or by asteroid belts orbiting pulsars in binary systems; and the timing errors created by changes in the shape of the pulses emitted by pulsars, together with ways in which those errors might be mitigated. First, however, some explanation is in order.

1.1 Pulsars

The first pulsar was discovered in 1967 by Jocelyn Bell, then a graduate student at the University of Cambridge under the supervision of Antony Hewish. Bell and Hewish were studying interplanetary scintillation of quasars using an array of dipole antennas located at the Mullard Radio Astronomy Observatory near Cambridge when Bell noticed an unusual signal in a chart recording, consisting of repeating pulses that appeared to be coming from a consistent location on the sky (Hewish et al., 1968). The source of the signal, initially

dubbed “LGM-1”¹, was ultimately identified as a rapidly rotating neutron star, and given the name PSR B1919+21, derived from its right ascension and declination. Hewish and his collaborator Martin Ryle would go on to win the 1974 Nobel Prize in Physics for the discovery; to the enduring shame of the Nobel committee, Bell was left out of the award. However, she continued to be active in the field, and in 2018 was awarded the Special Breakthrough Prize in Fundamental Physics. In her more recent work, she is known by her married name, Jocelyn Bell Burnell. At the time of writing, Bell Burnell is a Visiting Professor of Astrophysics at the University of Oxford.

Neutron stars, first theorized to exist by Baade & Zwicky (1934), are extremely dense objects composed almost entirely of neutrons, with radii between 10 and 20 km, and typical masses between 1 and $2 M_{\odot}$. They form as a result of core-collapse supernovae, and retain much of the angular momentum of their progenitor stars. As a result, they spin extremely rapidly, completing approximately one rotation per second in typical cases. They are also extremely highly magnetized, with surface magnetic field strengths in excess of 10^{12} G being common. Before the discovery of B1919+21, the fact that some neutron stars produce detectable radio emission was entirely unknown, but the connection was made within two years (Gold, 1968; Pacini, 1968; Gold, 1969). At the same time, the identity of pulsars as neutron stars was confirmed by the discovery of pulsations originating from the Crab nebula (Staelin & Reifenstein, 1968; Comella et al., 1969; Richards & Comella, 1969), which is the remnant of a bright supernova recorded by Chinese astronomers in 1054 AD.

The Crab pulsar, PSR B0531+21, has a period of approximately 33 ms, much

¹LGM stood for “little green men”, a reference to the unlikely possibility that it represented a signal from extraterrestrial intelligence.

shorter than the 1.34 s period of B1919+21, which is more typical of pulsars as a whole. This turns out to be a consequence of the Crab’s relative youth. Pulsars gradually radiate away their rotational energy over time, causing them to spin down as they age. Most of the energy is radiated away not as radio waves, but as a particle wind driven by magnetic dipole radiation at the rotational frequency of the pulsar. It is this particle wind that powers the synchrotron radiation emitted by supernova remnants like the Crab nebula. The corresponding rate of increase in the period is very small, but nevertheless can be measured directly, by recording the precise times of arrival (TOAs) of pulses over a period of several years. This was one of the key predictions of Gold (1968). The period of the Crab pulsar increases at a rate $\dot{P} = 4.2 \times 10^{-13} \text{ s s}^{-1}$, while PSR B1919+21 has $\dot{P} = 1.3 \times 10^{-15} \text{ s s}^{-1}$.

Plotting \dot{P} against the period, P , for all known pulsars, gives the famous $P - \dot{P}$ diagram, which plays a role somewhat analogous to the Hertzsprung–Russell diagram for main-sequence stars, in that it allows several different classes of pulsars to be identified and their properties distinguished. A $P - \dot{P}$ diagram based on data from psr² (Manchester et al., 2005), a pulsar catalog maintained by the Australia Telescope National Facility (ATNF), is shown in Figure 1.1. The largest class of pulsars, the so-called canonical pulsars (CPs), a group which includes B1919+21, forms a cluster near the center of the $P - \dot{P}$ diagram, around $P = 1 \text{ s}$ and $\dot{P} = 10^{-15} \text{ s s}^{-1}$.

As they age, CPs move down and to the left in the $P - \dot{P}$ diagram, towards longer periods and lower period derivatives. This behavior is determined by

²<https://www.atnf.csiro.au/research/pulsar/psrcat/>

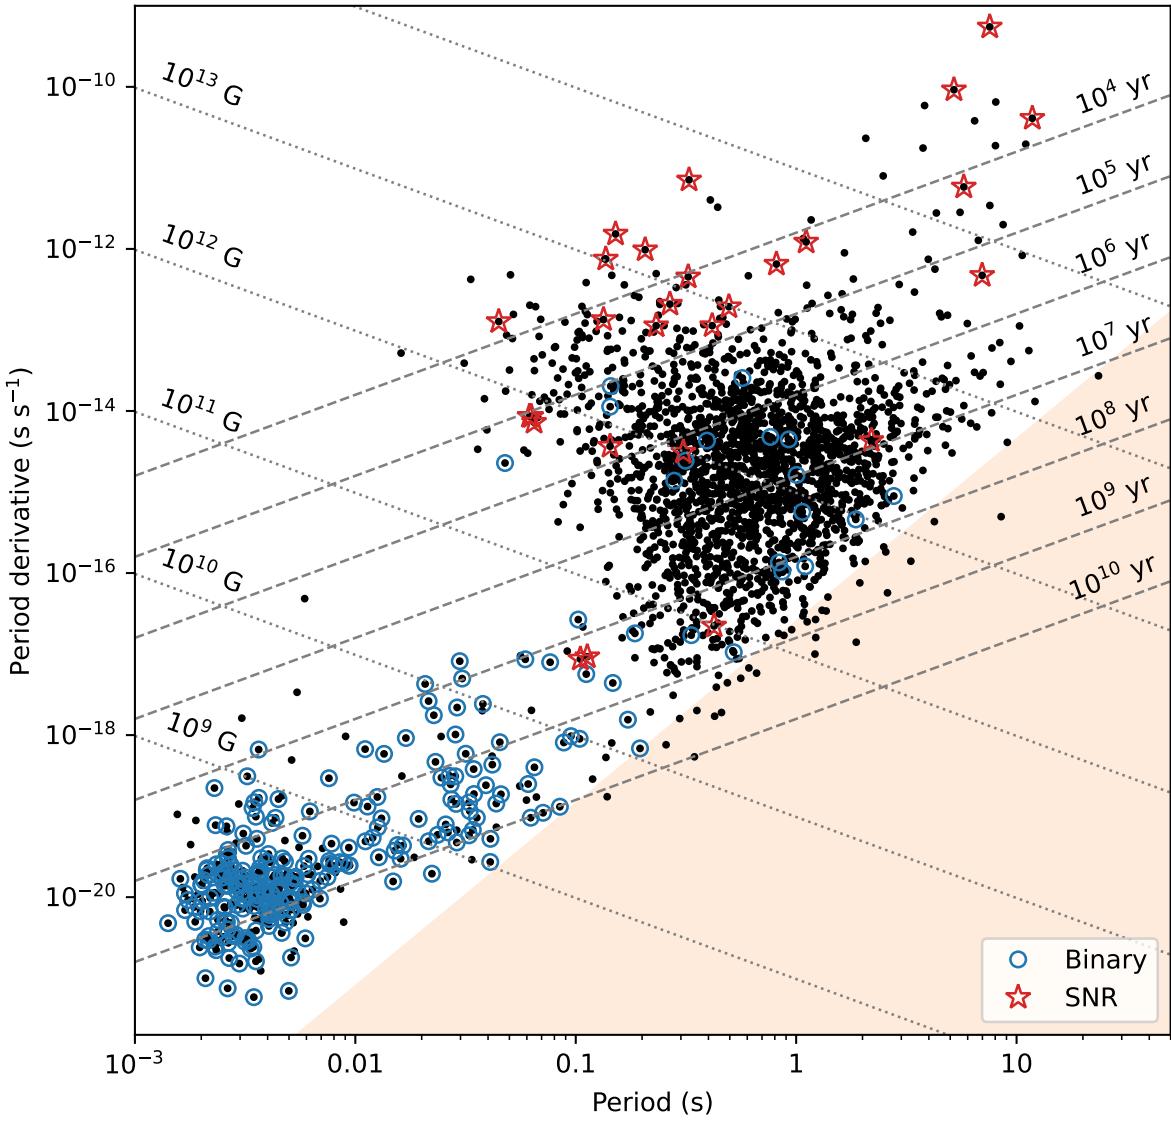


Figure 1.1: A $P - \dot{P}$ diagram, based on data from the ATNF catalog (Manchester et al., 2005). Lines of constant characteristic age (dashed) and surface magnetic field strength (dotted) are overplotted. Pulsars with known binary companions are shown with blue circles, and pulsars associated with known supernova remnants (SNRs) are shown with red stars. The region below the so-called “death line” is shaded orange. This plot was created with the help of the Python package `psrqpy` (Pitkin, 2018).

the relationship between P and \dot{P} , which is commonly given in the form

$$\dot{P} \propto P^{2-n}, \quad (1.1)$$

where n is the so-called braking index, which is equal to 3 if the spindown is entirely due to dipole radiation. As long as $n \neq 1$, equation (1.1) can be integrated to give

$$T = \frac{P}{(n-1)\dot{P}} \left[1 - \left(\frac{P_0}{P} \right)^{n-1} \right], \quad (1.2)$$

where P_0 is the initial period of the pulsar, and T is the amount of time elapsed since the pulsar was spinning with period P ; in other words, the age of the pulsar. Assuming that $P_0 \ll P$ and $n = 3$, the age of the pulsar is given by $T = P/(2\dot{P})$; this estimate is often referred to as the characteristic age of the pulsar. Contours of constant characteristic age are shown as dashed lines in Figure 1.1. Typical CPs have characteristic ages of a few million years, but those associated with supernova remnants (SNRs; red stars in Figure 1.1), including the Crab pulsar, are generally significantly younger, indicating that pulsars continue to produce radio emission after the associated SNRs have dispersed.

The total power radiated by the pulsar as it spins down is given by

$$\dot{E} = -I\Omega\dot{\Omega} = \frac{4\pi^2 I \dot{P}}{P^2}, \quad (1.3)$$

where I is the moment of inertia of the pulsar and $\Omega = 2\pi/P$ is its angular velocity. Combining this with equation (1.1) shows that the radiated power decreases over time; once it decreases past a certain point, the pulsar's radio emission shuts off entirely. From the perspective of a radio observer, the pulsar has died. This corresponds with the pulsar crossing a line in the $P - \dot{P}$ diagram known as the “death line”, into the orange shaded region in Figure 1.1 (the “graveyard”).

There is a second major class of pulsars in Figure 1.1, besides the canonical pulsars. These are the millisecond pulsars (MSPs), also called recycled pulsars. They have typical periods of a few milliseconds, and \dot{P} values around 10^{-20} , giving them characteristic ages of hundreds of millions to many billions of years. The first known MSP, PSR B1937+21, was discovered by Backer et al. (1982), and has a period of 1.56 ms, still one of the shortest known. The term “recycled pulsars” is a reference to their commonly accepted formation mechanism: MSPs are thought to arise when a “dead” pulsar in a binary system is spun up for a second time by accreting matter from its companion star. As a result, such a pulsar can attain an extremely short period without a correspondingly large period derivative. An important piece of evidence for this formation mechanism is the fact that a large number of MSPs are found in binary systems (blue circles in Figure 1.1), typically with white dwarf companions. Such a white dwarf is the leftover core of a red giant star, the loss of whose outer layers served to spin up the MSP. The particle wind generated by the MSP will gradually ablate away its companion, causing it to lose mass and eventually disintegrate entirely. This process may explain the existence of a small number of isolated MSPs.

Pulsars, especially MSPs, are extremely stable in their rotation, and produce a very predictable sequence of pulses. As a result, they can be thought of as naturally-occurring, high-precision clocks (with built-in radio transmitters, no less!). This is the basis of pulsar timing: one can construct a timing model and fit it to a series of TOA measurements to obtain precise estimates of many of the pulsar’s properties, including its position, proper motion, and binary orbital parameters. By timing MSPs, it is possible to construct a pulsar-based time scale that can compete with state-of-the-art ground-based atomic clocks (e.g.,

Hobbs et al., 2020). Pulsars also make an excellent tool for interplanetary and even interstellar navigation (Reichley et al., 1971; Downs, 1974; see also Shemar et al., 2016) and other applications requiring changes in the distance between the pulsar and the observer to be precisely determined. One such application is the direct detection of gravitational waves, which will be discussed further below.

Changes in distance can be determined to a precision limited only by the uncertainty in TOA measurements. In favorable cases, this can be less than 100 ns, corresponding to distance changes of less than 30 m. Importantly, however, without knowledge of the absolute phase, period, and period derivative at the point of emission, it is impossible to measure the absolute distance, or any linear or quadratic trend in the distance, to this high precision. Only signals with higher-frequency components can benefit from the full precision.

1.1.1 Interstellar Dispersion and Scattering

Before reaching a detector on Earth, the radio waves emitted by a pulsar must travel through the interstellar medium (ISM), a diffuse plasma consisting primarily of ionized hydrogen. This leads to two major complications affecting radio pulsar timing: dispersion and scattering of the signal by the ISM.

As a radio wave passes through the ISM, its electric and magnetic fields cause the electrons and ions of the ISM plasma to accelerate, creating a time-varying current density that gives rise to its own electromagnetic field, which interferes with the original wave. The electric field has a much more significant effect than the magnetic field, and the electrons, being less massive than the

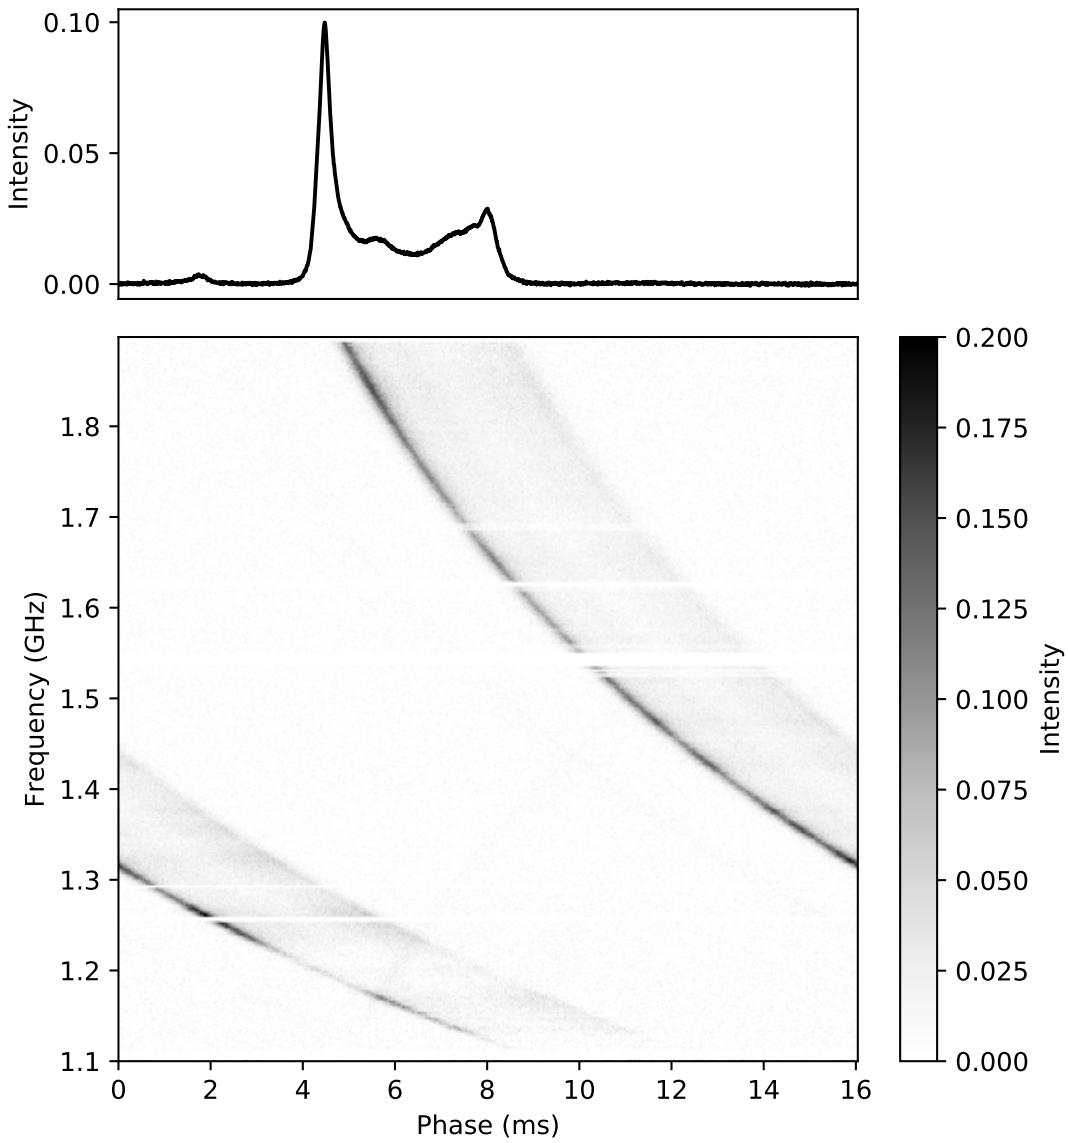


Figure 1.2: Interstellar dispersion in an observation of an MSP, PSR J2145–0750, made with the 1–2 GHz receiver on the 100-meter Green Bank Telescope (GBT). PSR J2145–0750 has a period of 16.05 ms and a dispersion measure of 9.00 pc cm^{-3} . Notably, the dispersion delay over this 800 MHz band changes by more than the period of the pulsar, causing the pulse to wrap in phase.

ions, accelerate to a much greater degree. The current density, \mathbf{j} , can therefore be approximated by assuming (a) that all the acceleration is due to the electric field associated with the radio wave, and (b) that only electrons move. The result, for a wave with electric field strength $\mathbf{E} = \mathbf{E}_0 e^{2\pi i(\nu t - \mathbf{k} \cdot \mathbf{x})}$, is

$$\mathbf{j} = -\frac{in_e e^2}{2\pi m_e \nu} \mathbf{E}_0 e^{2\pi i(\nu t - \mathbf{k} \cdot \mathbf{x})}, \quad (1.4)$$

where n_e is the density of free electrons in the ISM, e is the fundamental charge, and m_e is the electron mass. Applying Maxwell's equations in differential form, we find that, as in vacuum, the electric and magnetic fields are transverse to the direction of motion ($\mathbf{k} \cdot \mathbf{E} = \mathbf{k} \cdot \mathbf{B} = 0$), and the magnetic field satisfies $\mathbf{B} = \nu^{-1} \mathbf{k} \times \mathbf{E}$. Unlike in vacuum, electromagnetic waves propagating through the plasma satisfy the dispersion relation

$$\nu^2 = \nu_p^2 + c^2 k^2, \quad (1.5)$$

where the so-called plasma frequency, ν_p , is given by³

$$\nu_p = \frac{1}{2\pi} \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}}. \quad (1.6)$$

Electromagnetic waves with $\nu < \nu_p$ do not propagate, but instead are reflected by the plasma. Waves at higher frequencies travel with group velocity

$$v_g = \frac{d\nu}{dk} = c \sqrt{1 - \frac{\nu_p^2}{\nu^2}}, \quad (1.7)$$

which is slower at lower frequencies. This means that the low-frequency components of a pulse will arrive after the high-frequency components: at frequency ν , the pulse will be delayed by an amount of time

$$\Delta t = \frac{L}{v_g} - \frac{L}{c} = \frac{L}{c} \left[\left(1 - \frac{\nu_p^2}{\nu^2} \right)^{-1/2} - 1 \right]. \quad (1.8)$$

³Here I am working in SI units. To get the equivalent results in CGS units, replace ϵ_0 with $1/(4\pi)$.

Pulsars are typically observed at frequencies $\nu \gg \nu_p$, so this simplifies to

$$\Delta t = \frac{\nu_p^2 L}{2c\nu^2} = \frac{n_e e^2 L}{8\pi^2 \epsilon_0 m_e c \nu^2}. \quad (1.9)$$

That is, the delay is proportional the electron density, n_e , and the path length, L , and inversely proportional to ν^2 . The degree of dispersion is traditionally quantified using the dispersion measure, $\text{DM} = n_e L$, measured in units of pc cm^{-3} . In terms of the DM, equation (1.9) takes the form

$$\Delta t = \frac{k \text{DM}}{\nu^2} = 4.149 \text{ ms} \left(\frac{\text{DM}}{\text{pc cm}^{-3}} \right) \left(\frac{\nu}{\text{GHz}} \right)^{-2}, \quad (1.10)$$

where the dispersion constant, k , is given by

$$k = \frac{e^2}{8\pi^2 \epsilon_0 m_e c} = \frac{\hbar \alpha}{2\pi m_e}. \quad (1.11)$$

Here \hbar is the reduced Planck constant and $\alpha \approx 1/137$ is the fine-structure constant. In reality, the ISM is an inhomogeneous medium, and the signal passes through regions of differing electron density on its way from the pulsar to the receiver, so the DM is given by an integral of the line of sight to the pulsar:

$$\text{DM} = \int_0^L n_e d\ell. \quad (1.12)$$

Figure 1.2 shows the effect of dispersion in an observation of PSR J2145–0750, an MSP with a relatively long period and low DM, made with the 1–2 GHz receiver on the 100-meter Green Bank Telescope (GBT), between 1.1 and 1.9 GHz. J2145–0750 has a DM of 9.00 pc cm^{-3} , so the dispersion delay (relative to infinite frequency) is 10.34 ms at 1.9 GHz and 30.86 ms at 1.1 GHz. The difference between these, 20.51 ms, is more than the 16.05 ms period of the pulsar, so the pulse wraps in phase. Typical millisecond pulsars have even shorter periods and larger DMs, so this phase wrapping happens several times across the 800 MHz band.

Fortunately for pulsar timing experiments, dispersion has a predictable frequency dependence (given by equation 1.10), and the DM can be measured for any given pulsar, allowing the bulk of the effect to be corrected for. However, the DM is not perfectly constant over time: it varies, typically at the level of a few times 10^{-4} pc cm $^{-3}$. This variability results from the motion of density variations in the ISM across the line of sight, and, particularly for pulsars which are close to the ecliptic, the motion of the line of sight through the solar wind. To achieve the highest possible timing precision, these time-variable dispersion delays must be carefully accounted for.

Radio waves are not only dispersed as they travel through the ISM, but also scattered. Scattering is the result of diffraction and refraction by inhomogeneities in the ISM. It is analogous to the blurring effect of the atmosphere at visible wavelengths, and similarly depends on the apparent angular size of sources. At visible wavelengths, stars twinkle, while planets do not; this is because stars have a smaller angular size and so their light is more strongly scattered by the atmosphere. At radio frequencies, pulsars “twinkle”, while active galactic nuclei (AGN), which are the most common point sources, do not; this is similarly because pulsars have a smaller angular size and thus are more strongly scattered, in this case by the ISM.

For pulsars and other sources in the strong scattering regime, there are two qualitatively different kinds of scattering (Rickett, 1990): diffractive interstellar scintillation (DISS) and refractive interstellar scintillation (RISS). DISS produces amplitude modulations on time scales of seconds to minutes, which are correlated only across narrow frequency bands, while RISS causes the pulsar to gradually brighten and dim on much longer timescales (weeks). In analyz-

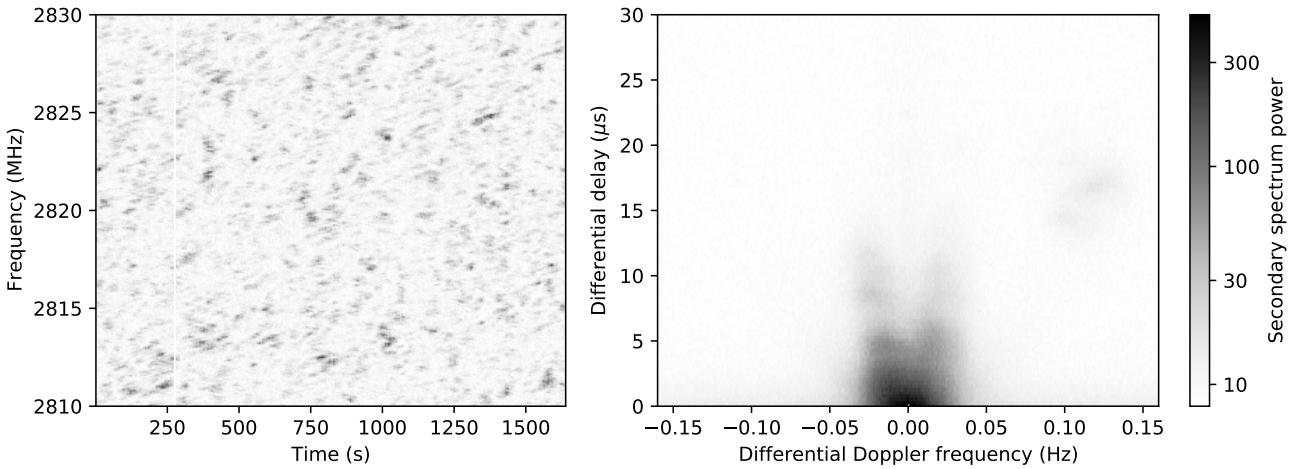


Figure 1.3: A portion of the dynamic spectrum (left) and the secondary spectrum (right) of an observation of the Vela pulsar, PSR B0833–45, made using the 10 cm receiver on the 64-meter Parkes telescope. This illustrates the effects of diffractive interstellar scintillation (DISS). The dynamic spectrum shown on the left is integrated over the on-pulse region, with each time point derived from all the pulses observed during a 2-second interval. The secondary spectrum on the right shows evidence of a parabolic arc, although it is relatively poorly defined compared to those seen in some other pulsars.

ing pulsar timing observations, DISS is most obvious, since the variations may well be visible within a single observing session, while RISS is more a matter of chance: on some days, you will be lucky, and find that the pulsar is brighter than expected because it is scintillated up, while on other days you might not detect the pulsar at all because it is scintillated down.

Figure 1.3 illustrates the effect of DISS on the Vela Pulsar, PSR B0833–45. The left panel shows the dynamic spectrum of the pulsar, averaged over the on-pulse region. The spectrum consists of a pattern of bright patches, called “scintles”, which have a characteristic width in both frequency and time. The power spectrum derived from the Fourier transform of the dynamic spectrum, called the secondary spectrum, is shown in the right panel of Figure 1.3. The secondary spectrum encodes information about the statistical properties of the

pattern of scintles. For reasons to be outlined below, the secondary spectrum power is typically concentrated in a parabolic arc.

DISS is typically modeled as taking place within a discrete, relatively thin screen, with turbulent fluctuations in electron density across its surface. In passing through such a screen, radio waves are diffracted through a small angle, θ . Those that reach the observer therefore end up taking a slightly longer path, and experience a delay proportional to θ^2 . In particular, if the pulsar is located at $z = 0$, the observer at $z = L$, and the screen at $z = sL$, a simple geometric argument shows that, to second order in θ , the new path length is given by

$$L' = L \left[1 + \frac{1}{2} s(1-s)\theta^2 \right]. \quad (1.13)$$

This causes the pulse to be delayed by an amount of time

$$\Delta t = \frac{L'}{c} - \frac{L}{c} = \frac{s(1-s)}{2c} L \theta^2. \quad (1.14)$$

The pulse seen by an observer is the superposition of many different images, each diffracted along a different path. In the simplest model, $\theta^2 = \theta_x^2 + \theta_y^2$, where the probability that a ray is scattered in the x direction by θ_x and in the y direction by θ_y is Gaussian, with width θ_{rms} in both directions. The distribution over θ^2 is then a χ^2 distribution with two degrees of freedom, and mean θ_{rms}^2 :

$$P(\theta^2) = \theta_{\text{rms}}^{-2} e^{-\theta^2/\theta_{\text{rms}}^2} d(\theta^2), \quad (1.15)$$

so the effect of scattering by the screen is to broaden the pulse by convolving it with the one-sided exponential function $h(t) = \tau^{-1} e^{-t/\tau}$, where the pulse broadening time, τ , is given by

$$\tau = \frac{s(1-s)}{2c} L \theta_{\text{rms}}^2. \quad (1.16)$$

Scattered images are also Doppler shifted by an amount

$$\Delta\nu \propto v_x\theta_x + v_y\theta_y, \quad (1.17)$$

where v_x and v_y are the x and y components, respectively, of the velocity of the intersection point between the screen and the line of sight. Because $\Delta\nu$ depends linearly on θ , while Δt depends on it quadratically, there is a quadratic relationship between the two quantities, and this is what gives rise to the quadratic arc in the secondary spectrum (Stinebring et al., 2001; Walker et al., 2004; Cordes et al., 2006).

1.2 Gravitational Waves

Gravitational waves (GWs) were Einstein's answer to the "spooky action at a distance" that plagued Newtonian gravity. In Einstein's general theory of relativity (GR), rather than being transmitted instantaneously, changes in gravitational forces are propagated by waves in the gravitational field which travel at the speed of light, c , in a manner directly analogous to the situation in the theory of electromagnetism. Because GR describes gravity through the geometry of spacetime, gravitational waves are quantified by the changes they produce in the metric tensor, $g_{\mu\nu}$, that defines this geometry.

Specifically, the dynamics of GWs are described by the Einstein field equations, which take the form

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (1.18)$$

The tensor $G_{\mu\nu}$ on the left-hand side of this equation, called the Einstein tensor, describes the curvature of spacetime. It depends only on the metric tensor, $g_{\mu\nu}$,

including through the Ricci tensor, $R_{\mu\nu}$, and its trace, the Ricci scalar, R . The stress-energy tensor, $T_{\mu\nu}$, on the right-hand side, describes the distribution of energy and momentum in spacetime, and plays the role of a source.

In general, the Einstein field equations constitute a system of 10 coupled non-linear differential equations for $g_{\mu\nu}$, which can only be solved numerically (and then only with great effort), except in a handful of special cases. However, in the limit where the gravitational field strength is small, they can be replaced by a linear system of equations:

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}. \quad (1.19)$$

Here the metric tensor, $g_{\mu\nu}$, has been expanded about its flat-space value, $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (1.20)$$

The symmetric tensor $h_{\mu\nu}$, referred to as the metric perturbation or GW strain, appears on the left-hand side of equation (1.19) in trace-reversed form: $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}h\eta_{\mu\nu}$, where the trace $h = \eta^{\mu\nu}h_{\mu\nu}$. The operator \square is the d'Alembertian, or wave operator, defined by

$$\square u = \eta^{\mu\nu}\partial_\mu\partial_\nu u = \left(-\frac{1}{c^2}\frac{\partial^2}{\partial t^2} + \nabla^2\right)u. \quad (1.21)$$

There is a significant amount of gauge freedom in equation (1.19). In particular, by choosing our coordinate system appropriately, we are free to enforce the Lorentz gauge condition $\partial^\nu h_{\mu\nu} = 0$. In free space, far from any sources, we can also demand that the trace $h = 0$, putting us in the so-called transverse-traceless (TT) gauge. Thanks to these gauge conditions, in our appropriately chosen coordinate system, only two of the ten components of $h_{\mu\nu}$ are nonzero.

For a wave propagating in the z direction, we have

$$h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (1.22)$$

The numbers h_+ and h_\times are the amplitudes of the two different GW polarizations, termed the “plus” and “cross” polarizations, respectively. The names come from the effect they have on test particles in the plane perpendicular to the wave’s motion: distances between particles are first stretched along one axis and compressed along the perpendicular axis, then compressed along the first axis and stretched along the second. For plus-polarized waves, this stretching and compression happens along the coordinate axes, while for cross-polarized waves, it happens along axes oriented at 45° from the coordinate axes.

Just as in the case of electromagnetic waves, gravitational waves far from a source can be described by a multipole expansion. However, unlike in the electromagnetic case, the dipole term in this expansion is always zero, a consequence of the conservation of momentum. This means that the leading term is the quadrupole term. In the TT gauge, the GW strain produced by a source with quadrupole moment tensor $I_{jk}(t)$ is given by

$$h_{jk} = \frac{2G}{c^4 r} \ddot{\mathcal{I}}_{jk}(t_r), \quad (1.23)$$

where $\mathcal{I}_{jk}(t) = I_{jk}(t) - \frac{1}{3} \delta_{jk} \delta^{\ell m} I_{\ell m}(t)$ is the trace-free quadrupole moment tensor, and $t_r = t - r/c$ is the retarded time, also used in the theory of electromagnetic waves. Equation (1.23) only gives the spacelike components of the strain; in the TT gauge, the remaining components are identically zero. The total GW

luminosity of such a source is given by

$$\dot{E}_{\text{GW}} = \frac{G}{5c^5} \left\langle \ddot{\mathcal{I}}_{jk} \ddot{\mathcal{I}}^{jk} \right\rangle, \quad (1.24)$$

where the angle brackets indicate an average taken over all directions.

The prototypical astrophysical source of gravitational waves is a binary system, that is, two massive objects in orbit around one another. Even in GR, as long as the objects are separated by a sufficient distance, their orbits are adequately described by Kepler's laws of planetary motion. Using the orbit solutions, it is possible to write down an expression for the quadrupole moment tensor, and thereby compute the GW strain. The results are worked out in detail for the general case of elliptical orbits in Peters & Mathews (1963). In the simplest case of a circular orbit with angular momentum in the z direction, one has

$$h_{ij}(t) = h_0 \begin{bmatrix} -\cos 2\pi f t_r & -\sin 2\pi f t_r & 0 \\ -\sin 2\pi f t_r & \cos 2\pi f t_r & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (1.25)$$

where the GW strain amplitude, h_0 , is given by

$$h_0 = \frac{4G^2 M \mu}{c^4 a r}. \quad (1.26)$$

Here $\mu = m_1 m_2 / (m_1 + m_2)$ is the reduced mass of the binary, a is the semi-major axis (equivalently for a circular orbit, the orbital separation), and f is the GW frequency, given by

$$f = \frac{1}{\pi} \sqrt{\frac{GM}{a^3}}, \quad (1.27)$$

where $M = m_1 + m_2$ is the total mass of the binary. Notably, the GW frequency is twice the Keplerian orbital frequency.

The GW luminosity of the binary can be determined using equation (1.24),

with the result

$$\dot{E}_{\text{GW}} = \frac{32G^4M^3\mu^2}{5c^5a^5}. \quad (1.28)$$

Since the gravitational binding energy of the binary is given by

$$E_B = -\frac{GM\mu}{2a}, \quad (1.29)$$

the semimajor axis, a , must shrink as energy is radiated away, and the GW frequency must increase correspondingly: the GW signal is “chirped”. Using equations (1.27), (1.28), and (1.29), it can be shown that f satisfies the differential equation

$$\dot{f} = \frac{96\pi}{5c^5}(\pi G\mathcal{M}_c)^{5/3}f^{11/3}. \quad (1.30)$$

The parameter $\mathcal{M}_c = (M^2\mu^3)^{1/5}$ is called the chirp mass of the binary. The GW strain amplitude (equation 1.26) and GW luminosity (equation 1.28) can also be expressed in terms of f and \mathcal{M}_c . In particular, we have

$$h_0 = \frac{4}{\pi c^4 r}(\pi G\mathcal{M}_c)^{5/3}f^{2/3}, \quad (1.31)$$

$$\dot{E}_{\text{GW}} = \frac{32}{5Gc^5}(\pi G\mathcal{M}_c)^{10/3}f^{10/3}. \quad (1.32)$$

For many years after they were initially predicted by Einstein, gravitational waves remained entirely theoretical. The first confirmation of the existence of gravitational waves came from pulsar timing: monitoring of PSR B1913+16, which is in a binary system with another neutron star of nearly equal mass, showed that its orbit was decaying in a manner consistent with energy loss due to gravitational radiation (Taylor & Weisberg, 1982; Weisberg et al., 2010; Weisberg & Huang, 2016). Astronomers Russell Hulse and Joe Taylor were awarded the 1993 Nobel Prize in Physics for their discovery of the B1913+16 system, which was the first double neutron star system known, and was immediately recognized as important for the study of gravitation.

The first direct detection of GWs came later, in September 2015, when the Advanced LIGO experiment, consisting of a pair of modified Michelson interferometers, detected the chirp created by a pair of merging stellar-mass black holes (Abbott et al., 2016). The signal, named GW150914 for the date it was observed, swept upwards in frequency from 35 to 250 Hz over the course of about 12 ms, and was eventually identified as resulting from the merger of two black holes with masses of 29_{-4}^{+4} and $36_{-4}^{+5} M_{\odot}$, at a distance of 410_{-180}^{+160} Mpc. Kip Thorne, Rainer Weiss, and Barry Barish, all of whom made major contributions to the LIGO scientific collaboration, received the 2017 Nobel Prize in Physics for the discovery.

Since the detection of GW150914, LIGO has continued to operate, and has been joined by two similar detectors, Virgo in Italy and KAGRA in Japan. LIGO detected two more binary black hole (BBH) mergers during its first observing run (O1), which ran from September 2015 to January 2016. During the second observing run (O2), which ran from November 2016 to August 2017, LIGO and Virgo made the first detection of GWs from the merger of two neutron stars (GW170817; Abbott et al., 2017). Unlike the previously detected black hole mergers, the neutron star merger had electromagnetic counterparts across much of the spectrum, including a short gamma-ray burst (GRB). The observation of GW170817 established that at least some short GRBs are created by neutron star mergers, and provided an opportunity to learn much more about the nature of these events. Seven more BBH mergers were also detected during O2 (Abbott et al., 2019). LIGO was shut down between August 2017 and April 2019 for a series of upgrades before its third observing run (O3), which was divided into two halves. Results from the first half, O3a, which ran from April to October 2019, were published in April 2021 (Abbott et al., 2021), and include

39 candidate BBH merger events. Results from the second half, O3b, which ran from November 2019 to March 2020 and detected 35 additional candidate BBH merger events, were submitted for publication in November 2021 (The LIGO Scientific Collaboration et al., 2021).

Gravitational-wave astronomy is now a very active area of research. In addition to ground-based interferometric detectors like those mentioned above, which are sensitive to GWs at frequencies of tens to hundreds of hertz, planning is underway for the Laser Interferometer Space Antenna (LISA; Amaro-Seoane et al., 2017). Planned for launch in the mid-2030s, LISA will be a space-based GW detector sensitive to GWs in the millihertz band. The research in this thesis, however, is concerned with another kind of GW detector entirely: the pulsar timing array.

1.3 Pulsar Timing Arrays

Pulsar timing arrays are experiments that aim to use pulsar timing to detect gravitational waves in the nanohertz frequency band. Compared to LIGO and other ground-based GW detectors, which search for transient chirped signals associated with the merger of stellar-mass objects like neutron stars and black holes, PTAs are sensitive to much more massive and longer-lived binaries. Specifically, the sources that PTAs are most likely to detect are supermassive black hole binaries (SMBHBs). Most galaxies are now thought to contain supermassive black holes at their centers, with masses ranging from millions to billions of solar masses. When galaxies merge, dynamical friction and other processes will eventually bring their central black holes close enough

together to form a binary capable of emitting gravitational waves. The most massive of these binaries, whose constituents have masses of $10^9 M_\odot$ or more, are the most promising targets for GW detection by PTAs.

The operating principle of PTAs is similar to that of earlier spacecraft radar ranging experiments (e.g., Estabrook & Wahlquist, 1975). As was first pointed out by Detweiler (1979), the Earth and a pulsar can be used as two ends of a GW antenna. Changes to the distance between them can be monitored by measuring the arrival times of the pulses. The change in distance caused by a passing GW then takes the form of a time delay

$$\Delta t = [F_+ H_+(t_e) + F_\times H_\times(t_e)] - [F_+ H_+(t_p) + F_\times H_\times(t_p)], \quad (1.33)$$

where $H(t)$ is the time antiderivative of the GW strain, $h(t)$, t_e is the time at which the GW reaches the Earth, and t_p is the “pulsar time”, i.e., the time at which the information that the GW has reached the pulsar arrives at Earth. The two terms in square brackets in equation (1.33) are known as the “earth term” and the “pulsar term”, respectively. The factors F_+ and F_\times are the “antenna pattern functions” (cf. Ellis et al., 2012), which depend on the propagation direction, \hat{k} , of the GW, as well as the direction \hat{p} from the earth to the pulsar. They are defined by

$$F_+ = \frac{(\hat{p} \cdot \hat{m})^2 - (\hat{p} \cdot \hat{n})^2}{2(1 + \hat{p} \cdot \hat{k})} \quad (1.34)$$

$$F_\times = \frac{(\hat{p} \cdot \hat{m})(\hat{p} \cdot \hat{n})}{1 + \hat{p} \cdot \hat{k}}. \quad (1.35)$$

Here \hat{m} and \hat{n} are the unit vectors, orthogonal to each other and to \hat{k} , that make up the frame which defines the + and \times polarizations.

PTAs use timing measurements from a large number of MSPs to improve of their GW measurements. MSPs are preferred over other kinds of pulsar

because of their shorter periods and greater spin stability, which means that a higher timing precision is achievable. Since a GW will affect the distances to every pulsar simultaneously, signals from multiple pulsars can be combined to achieve even higher sensitivity. Detection of the same signal in multiple pulsars also provides confirmation of its GW origin, and in principle allows the direction of propagation of the wave to be determined.

The gravitational wave strain produced by a SMBHB can be estimated using equation (1.31). Plugging in representative numbers for the most massive SMBHBs gives

$$h_0 = 5.5 \times 10^{-16} \left(\frac{r}{\text{Gpc}} \right)^{-1} \left(\frac{\mathcal{M}_c}{10^9 M_\odot} \right)^{5/3} \left(\frac{f}{10 \text{ nHz}} \right)^{2/3}. \quad (1.36)$$

Because pulsar timing can achieve a precision $\sigma_\tau \lesssim 100$ ns on individual TOAs, and the length T of a data set can exceed 10 yr, PTAs can detect signals at amplitudes similar to this. A rough estimate of the strain uncertainty, σ_h , achievable by timing a single pulsar is

$$\sigma_h \sim \frac{\sigma_\tau}{T} = 3.2 \times 10^{-16} \left(\frac{\sigma_\tau}{100 \text{ ns}} \right) \left(\frac{T}{10 \text{ yr}} \right)^{-1}. \quad (1.37)$$

A somewhat more careful analysis shows that, for a sinusoidal signal at frequency f , σ_h is approximately given by

$$\begin{aligned} \sigma_h &= 2\pi f \sigma_\tau \sqrt{\frac{\delta t}{T}} \\ &= 5.7 \times 10^{-16} \left(\frac{f}{10 \text{ nHz}} \right) \left(\frac{\sigma_\tau}{100 \text{ ns}} \right) \left(\frac{\delta t}{30 \text{ d}} \right)^{1/2} \left(\frac{T}{10 \text{ yr}} \right)^{-1/2}, \end{aligned} \quad (1.38)$$

where δt , the observing cadence, is the time between successive observations of the pulsar. While this result scales correctly with frequency, time, TOA uncertainty, and observing cadence, it turns out to be somewhat too optimistic compared to the sensitivity achieved with single pulsars in practice, largely because of the presence of additional sources of noise in the TOA measurements.

A wave with a frequency of 10 nHz has a period of approximately 3 yr, so a ten-year data set would only contain about three cycles of the wave. Even lower frequency waves, with periods equal to or longer than the length of the data set, eventually become difficult to distinguish from the linear or quadratic trends that would arise from a slightly incorrect estimate of the pulsar's period or period derivative. This is the main factor limiting the sensitivity of PTAs at the lowest frequencies. Similarly, the need to fit for the pulsar's precise position and parallax leads to a reduction in sensitivity at frequencies of approximately $1 \text{ yr}^{-1} = 31.7 \text{ nHz}$ and $2 \text{ yr}^{-1} = 63.4 \text{ nHz}$, respectively. A more detailed discussion of the factors that determine a PTA's sensitivity can be found in Hazboun et al. (2019).

Before they are capable of detecting any individual SMBHB source, PTAs are expected to be able to detect the stochastic gravitational-wave background (SGWB) resulting from the superposition of all of them. One of the defining characteristics of the SGWB is its frequency spectrum. This can be derived using equations (1.30) and (1.31). From equation (1.30), the amount of time that a given binary spends at frequencies between f and $f + df$ is

$$dt \propto f^{-11/3} df. \quad (1.39)$$

Meanwhile, the variance it contributes to the timing residuals from any given pulsar is proportional to

$$H_0^2 = \frac{h_0^2}{4\pi^2 f^2} \propto f^{-2/3}. \quad (1.40)$$

This means that, averaged over a large population of binaries, the spectrum, $S(f)$, of the timing residuals will be given by

$$S(f) df \propto H_0^2 dt \propto f^{-13/3} df. \quad (1.41)$$

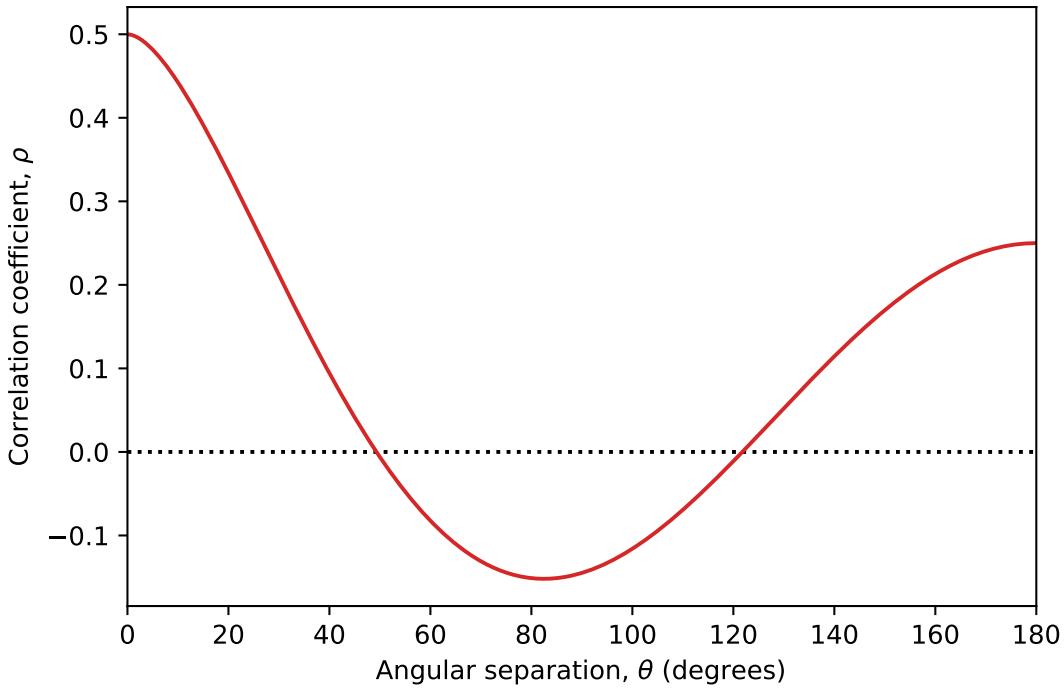


Figure 1.4: The Hellings and Downs curve (equation 1.42), which gives the correlation coefficient, ρ , between timing residuals created by an isotropic SGWB in pulsars separated by a given angle, θ , on the sky.

A stochastic process with a spectrum of the form $S(f) \propto f^{-\gamma}$ ($\gamma \gtrsim 1$), which has most of its power concentrated at low frequencies, is commonly referred to as a “red noise” process. The specific power-law index, $\gamma = 13/3$, seen in equation 1.41, is characteristic of the SGWB.

Another defining characteristic of the SGWB is the pattern of correlations it creates between timing residuals in different pulsars. A plane GW from a single source will cause the distances to pulsars in different parts of the sky to change in a quadrupolar pattern. As originally determined by Hellings & Downs (1983), for an isotropic background, this means that the timing residuals in a pair of pulsars separated by an angle θ will be correlated, with a correlation

coefficient ρ given by

$$\rho(\theta) = \frac{3}{2} \left(\frac{1 - \cos \theta}{2} \right) \left[\ln \left(\frac{1 - \cos \theta}{2} \right) - \frac{1}{6} \right] + \frac{1}{2}. \quad (1.42)$$

The function $\rho(\theta)$, also shown in Figure 1.4, is now referred to as the Hellings and Downs curve, and correlations of this form are regarded as a kind of “smoking gun” signature of a GW signal in pulsar timing residuals. No other known phenomena can produce correlations of this form, with the signal in pulsars on opposite sides of the sky being positively correlated, while those separated by 90 degrees are negatively correlated.

The PTA idea has led to the formation of several consortia of scientists around the world, who collaborate to produce data sets and search them for evidence of GWs. These include the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) in North America, the European Pulsar Timing Array (EPTA) in Europe, and the Parkes Pulsar Timing Array (PPTA) in Australia, which work together as part of the International Pulsar Timing Array (IPTA) project. Recently, they have been joined by the Indian Pular Timing Array (InPTA, because IPTA was already taken), as well as scientists associated with the MeerKAT radio telescope in South Africa and the FAST radio telescope in China. Efforts by the IPTA and its member collaborations have lead to increasingly stringent limits on the amplitude of the SGWB, typically quoted as a characteristic strain value, h_c , at a frequency $f = 1 \text{ yr}^{-1}$.

The analysis of the NANOGrav 11-year data set (Arzoumanian et al., 2018a), which included data from 45 pulsars, led to a 95% upper limit $h_c < 1.45 \times 10^{-15}$ (Arzoumanian et al., 2018b). Subsequently, the analysis of the NANOGrav 12.5-year data set (Alam et al., 2021a,b), which included data from two additional pulsars, revealed some evidence for a red noise process with a common am-

plitude and spectral slope across pulsars, at an amplitude $h_c \approx 1.92 \times 10^{-15}$, a value larger than the previous upper limit. This may represent the first hints of the SGWB. However, there is not (yet) significant evidence for Hellings and Downs correlations. The NANOGrav 15-year data set, currently being prepared, should reveal more about the nature of this common process.

The subsequent chapters in this dissertation describe efforts to increase the detection sensitivity of PTAs and to understand the additional signals which might appear in PTA data sets. In Chapter 2, we identify binary companions to pulsars in data from the *Gaia* spacecraft, and thereby obtain information about the distances to the corresponding pulsars. In Chapter 3, we investigate the signals resulting from hyperbolic encounters between pulsars and interstellar objects such as asteroids or rogue planets. Similarly, in Chapter 4, we describe the signals produced asteroid belts in orbit around pulsars, with particular emphasis on the role of a binary companion in limiting the inner extent of such a belt. Next, in Chapter 5, we quantify the ways in which pulse shape changes influence TOA estimates, and, in Chapter 6, describe how those influences may be mitigated. In Chapter 7, we describe recent observations of a shape change event in an important PTA pulsar, PSR J1713+0747, and speculate as to its origins and implications for PTA timing efforts. Finally, Chapter 8 gives some concluding remarks about possible future research directions.

CHAPTER 2

BINARY PULSAR DISTANCES AND VELOCITIES FROM GAIA DATA RELEASE 2

The second data release from the *Gaia* mission (*Gaia* DR2) includes, among its billion entries, astrometric parameters for binary companions to a number of known pulsars, including white dwarf companions to millisecond pulsars and the non-degenerate components of so-called “black widow” and “redback” systems. We find 22 such counterparts in DR2, of which 12 have statistically significant measurements of parallax. These DR2 optical proper motions and parallaxes provide new measurements of the distances and transverse velocities of the associated pulsars. For the most part, the results agree with existing radio interferometric and pulsar timing-based astrometry, as well as other distance estimates based on photometry or associations, and for some pulsars they provide the best known distance and velocity estimates. In particular, two of these pulsars have no previous distance measurement: PSR J1227–4853, for which *Gaia* measures a parallax of 0.62 ± 0.16 mas, and PSR J1431–4715, with a *Gaia* parallax of 0.64 ± 0.16 mas. Using the *Gaia* distance measurements, we find that dispersion measure-based distance estimates calculated using the NE2001 and YMW16 Galactic electron density models are on average slightly underestimated, which may be a selection effect due to the over-representation of pulsars at high Galactic latitudes in the present *Gaia* sample. While the *Gaia* DR2 results do not quite match the precision that can be achieved by dedicated pulsar timing or radio interferometry, taken together they constitute a small but important improvement to the pulsar distance scale, and the subset of millisecond pulsars with distances measured by *Gaia* may help improve the

sensitivity of pulsar timing arrays to nanohertz gravitational waves.

2.1 Introduction

The estimation of distances is a fundamental problem in astronomy, and since observers measure apparent quantities, distance estimates underpin much of astrophysics as well. For neutron stars, distance measurements (whether direct or indirect) enable a range of inferences, from constraints on the nuclear equation of state using radius measurements (e.g., Kaplan et al., 2002; Ho & Heinke, 2009; Shaw et al., 2018; Özel & Freire, 2016), to the physics of energy transport and the conversion of Poynting flux to particle flows in neutron star winds (e.g., Abdo et al., 2013; Spiewak et al., 2016). Precision astrometry of neutron stars allows the determination of birth sites and associations (e.g., Vlemmings et al., 2004; Kaplan et al., 2007, 2008; Tetzlaff et al., 2010; Kirsten et al., 2015) and high transverse velocity measurements constrain the physics of supernova core collapse (e.g., Chatterjee et al., 2005) as well as the evolution of close binary systems. Distances and transverse velocities are also important to determining both the underlying spin-down rates of pulsars and the intrinsic orbital decay rates of binaries, correcting for the kinematic Shklovskii effect (Shklovskii, 1970) and Galactic acceleration (e.g., Arzoumanian et al., 1999; Wex et al., 2000), which are important for the accurate determination of spin-down luminosities (e.g., Guillemot et al., 2016) and relativistic orbital decay (e.g., Damour & Taylor, 1991). However, distances to neutron stars are difficult to measure, and for most radio pulsars, they are only estimated indirectly through the observed pulse dispersion measure ($\text{DM} \equiv \int_0^d n_e ds$, the line-of-sight integral of the electron density). These DM-based distance estimates rely on models of

the Galactic electron density distribution such as that of Cordes & Lazio (2002a, henceforth NE2001) or that of Yao et al. (2017, henceforth YMW16). Precise, independent distance measurements are therefore vital, both in their own right and for their use in calibrating such models and thereby improving distance estimates for the rest of the radio pulsar population.

Independent distance estimates for radio pulsars come from a variety of sources, including limits based on H I absorption or associations with stellar clusters, spectroscopic parallaxes of binary companions, and other techniques, but the primary techniques rely on geometric parallax, which is equivalent to measuring the curvature of the pulse wavefront over the Earth’s orbit. Parallax can be measured via annual variations in pulse times of arrival (requiring timing precision $\lesssim 1 \mu\text{s}$ for distances of ~ 1 kpc; e.g., Toscano et al. 1999; Lorimer & Kramer 2012; Reardon et al. 2016; Desvignes et al. 2016; Guillemot et al. 2016; Matthews et al. 2016) or from astrometric imaging observations with sub-milliarcsecond precision, which can be accomplished with Very Long Baseline Interferometry (VLBI; e.g., Chatterjee et al. 2009; Deller et al. 2016). Both techniques (pulsar timing and interferometry) yield astrometry tied to the International Celestial Reference Frame (ICRF; Fey et al. 2015), as defined by extragalactic radio quasars (Ma et al., 1998). However, pulsar timing measurements depend on the position of the earth relative to the Solar System barycenter, and so can be affected by errors in the Solar System ephemeris (Fomalont et al., 1984; Wang et al., 2017). Interferometric measurements use a geocentric frame and so do not suffer from that issue.

The *Gaia* spacecraft (Gaia Collaboration et al., 2016), launched in 2013, is conducting an all-sky optical survey of more than 1 billion astrophysical

sources with magnitudes $\lesssim 21$. Like its predecessor *Hipparcos* (Perryman et al., 1997), *Gaia* solves simultaneously for the astrometric parameters of sources distributed across the sky. The results are connected to the ICRF through optical quasars (Mignard et al., 2018). The second *Gaia* data release (DR2, based on the first two years of data of a planned five year mission) contains the first astrometric solutions based entirely on *Gaia* data, including, for the majority of the catalog, the first measurements of parallax and proper motion (Gaia Collaboration et al., 2018).

Neutron stars are not the most obvious targets for optical astrometry, given their small sizes and lack of optical emission in all but a few cases (e.g., Mignani 2011; discussed further below). However, the companions of pulsars in binary systems offer opportune targets. We present results from a search through the *Gaia* DR2 catalog, where we have identified the companions to known binary pulsars by matching positions and proper motions, as well as the expected magnitude of the companions themselves where available, in order to weed out false positives. We describe our methods for source identification in Section 2.2, present our results in Section 2.3, including comparisons with previous distance measurements and estimates from Galactic electron density models, and conclude with a discussion of our results and future prospects in Section 2.4.

2.2 Methods

2.2.1 Identification of Sources

We identified probable companions to pulsars in the *Gaia* catalog based on their positions. Of the 2424 pulsars outside of globular clusters listed in the ATNF pulsar catalog (Manchester et al., 2005, 2016, v. 1.57), we selected the 188 that are in binary systems. We further selected those that have position uncertainties $< 1''$ in both coordinates. This is largely to exclude pulsars that cannot be reliably matched against *Gaia* astrometry, and effectively selects for those that have full phase-coherent timing solutions spanning at least a year. We selected binary pulsars because, with the exception of the Crab pulsar¹, no pulsar is intrinsically bright enough at optical wavelengths to be visible with *Gaia* (e.g., Mignani, 2011).

The resulting 155 sources were then cross-matched with *Gaia* DR2. Our cross-match used a radius of $1''$ after correcting the position of the pulsar to the epoch of *Gaia* (J2015.5) using the measured proper motion from the ATNF catalog. In most cases this is still far larger than the position uncertainty, with the majority of genuine counterparts agreeing to $< 0''.1$. In principle there could be even better agreement between pulsar timing positions and *Gaia*, but discrepancies between timing positions and the ICRF to which both *Gaia* and VLBI should be tied are present at the level of a few mas (Wang et al., 2017). Moreover, there are covariances between astrometric and other parameters in pulsar

¹The Crab pulsar is detected by *Gaia*, but the parallax is not detected at a statistically significant level. The proper motion of $\mu_\alpha = -11.8 \pm 0.2 \text{ mas yr}^{-1}$, $\mu_\delta = +2.6 \pm 0.2 \text{ mas yr}^{-1}$ is consistent with the proper motion measured using archival *Hubble Space Telescope* observations by Kaplan et al. (2008, $\mu_\alpha = -11.8 \pm 0.4 \pm 0.5 \text{ mas yr}^{-1}$, $\mu_\delta = +4.4 \pm 0.4 \pm 0.5 \text{ mas yr}^{-1}$).

binary fitting that make some proper motions less reliable than their formal uncertainties would indicate, and some pulsar positions are only known from optical observations (notably that of PSR J1417–4402; see below). Therefore we examined all sources with potential matches out to $1''$ by hand, consulting the literature for cases with known counterparts and examining images from the Pan-STARRS1 3π survey (Chambers et al., 2016a) for sources north of $\delta = -30^\circ$ and the SkyMapper Southern Sky Survey (Wolf et al., 2018) for sources south of $\delta = -30^\circ$ where there were ambiguous cases.

In almost every case, the potential association of a pulsar with a *Gaia* source was obvious and supported by previous detections in the literature (including requiring the *Gaia* match to have a similar brightness to that reported in the literature). However, in six cases—those of PSRs J1056–7117, J1125–6014, J1435–6100, J1543–5149, J1755–3716, and B1953+29—the putative matches appear from visual inspection to be unrelated foreground stars, inconsistent with the expected masses of the companions and distances to the systems (e.g., Mignani et al., 2014; Bates et al., 2015). For J1431–4715, whose white dwarf companion had not previously been detected (Bates et al. 2015 predict $V > 23$), we have verified that the companion is correct through spectroscopy (Kaiser et al., in prep.). For B1957+20 visual inspection showed that the *Gaia* source was an unrelated star $0.^{\prime}7$ away from the pulsar, despite the pulsar having a variable companion with a V-band magnitude which reaches a minimum of about 20 (Kulkarni et al., 1988; Callanan et al., 1995). In all cases we did a further test and estimated a rate of false detection based on the number of *Gaia* sources within $1'$ with magnitudes brighter than the putative companion. We found typical false positive rates of 10^{-5} , and as high as 10^{-3} for a few sources where either the proper motion was very high or the initial position measurement

was several decades ago.

Only in the cases of PSRs J0045–7319, J1417–4402, J1723–2837, and J2129–0429 were the position offsets greater than $0''.2$. However, all four of these have particularly bright companions with secure identifications. For PSR J0045–7319 the timing solution is over 20 years old and was complicated by the significant variations in spin and orbital properties encountered in the system (Kaspi et al., 1996). For PSR J1417–4402 the position is obtained from optical measurements of the companion star (Strader et al., 2015) and lacks the typical radio pulsar precision, but we can also be sure of the identification. For PSR J1723–2837 we can again be certain of the identification of a bright companion from the literature (Crawford et al., 2013) with radial velocity confirmation. Finally, for PSR J2129–0429, even though the companion position is offset from the pulsar position by $1''.3$, the companion was identified through radial velocity and photometric variations and is confirmed to be the same object, with the apparent offset between the optical position and the radio position already noted by Bellm et al. (2016).

We list all the sources with confident *Gaia* counterparts in Table 2.1, along with their dispersion measure distances and any other distance measurement found in the literature. Such distances include VLBI parallaxes, timing parallaxes, OB associations (for the Be-binary systems), spectroscopic parallaxes, and modeling of distorted companions in tight binaries.

2.2.2 Parallax Inversion and Lutz-Kelker Correction

In a field of objects with uniform density in space, more objects are located at larger distances from the observer than at smaller ones, because a spherical shell of a given thickness has a volume that scales as the square of its radius. As Lutz & Kelker (1973) observed, this has important implications for estimating the distance to an object from its trigonometric parallax. In effect, they argue that in such a case the prior distribution of the distance d to an object, $p(d)$, should be proportional to d^2 . In terms of the parallax π , this translates into $p(\pi) \propto \pi^{-4}$. In general, if the objects under consideration have space density $\rho(d, \theta, \phi)$, the prior on the distance d to an object known to lie in the direction (θ, ϕ) should have the form $p(d | \theta, \phi) \propto d^2 \rho(d, \theta, \phi)$.

As is typical, we assume the likelihood function for the true parallax π to have the form of a Gaussian centered on its measured value $\hat{\pi}$ with uncertainty σ_π , so that the posterior distribution for d takes the form

$$p(d | \hat{\pi}, \theta, \phi) \propto d^2 \rho(d, \theta, \phi) \exp\left[-\frac{(1/d - \hat{\pi})^2}{2\sigma_\pi^2}\right]. \quad (2.1)$$

In the case of constant space density, as originally observed by Lutz & Kelker, this has the unfortunate consequence that the posterior distribution $p(d | \pi)$ for d diverges as $d \rightarrow \infty$. When the uncertainty on the parallax σ_π is $\ll \pi$, the divergence occurs only for extremely large distances and can be ignored in practice. On the other hand, if $\rho(d, \theta, \phi)$ decreases rapidly enough as $d \rightarrow \infty$, the divergence disappears entirely.

With this in mind, we adopt a prior on the distance d to each pulsar based on a model of the Galactic pulsar population. This approach is similar to that advocated by Bailer-Jones et al. (2018) for estimating distances from *Gaia*

parallax data, but uses a pulsar-specific model for the space density ρ . Our model is based on that of Lorimer et al. (2006), but we adopt a larger vertical scale height of 500 pc for millisecond pulsars, following Verbiest et al. (2012). We use a DM-based upper bound on distance in a forthcoming analysis of the pulsar velocity distribution (Jennings et al., in prep.), but it is not included in the analysis here because its effect is insignificant. When the line of sight to the pulsar points away from the Galactic center, the effect of our prior can be to decrease the distance estimate, unlike the classical Lutz-Kelker correction based on an assumed uniform space density, which always increases the distance estimate.

Our prior does not account for all effects which could potentially bias parallax-based distance estimates. In particular, the fact that nearer objects are more likely to be detected due to their greater apparent brightness means that our Lutz-Kelker corrected distances may be systematically too large. Furthermore, the *Gaia* parallax estimates are subject to a systematic zero point offset, which has been estimated to average -0.029 mas but can vary as a function of source brightness, color, and position (Lindegren et al., 2018). We correct for this in the only way presently possible, by adding 0.029 mas to the central value $\hat{\pi}$ of the parallax likelihood function in calculating our distance estimates. In Table 2.2, in addition to giving Lutz-Kelker corrected distances d_{LK} , we also give distances d_{π} estimated by assuming a uniform prior on parallax. In both cases, we give the mode of the posterior distribution as a point estimate of d , and indicate the 16th and 84th percentiles of the distribution as errors. The distances d_{π} are compared with previous distance estimates in Figs. 2.1 and 2.2. The conclusions for d_{LK} , however, are essentially the same. For our sample of pulsars, the Lutz-Kelker correction increases the distance estimates by

no more than 30% (in most cases less than 25%), and in all cases the mode of the corrected distance distribution lies within the 68% credible region of the uncorrected distribution.

2.2.3 Velocity Estimation

In Table 2.2, we give three estimates of the component of the velocity of each pulsar transverse to the line of sight. All three estimates are calculated based entirely on *Gaia* parallax and proper motion data. The first and second, labelled $v_{\perp,\pi}$ and $v_{\perp,\text{LK}}$, are estimates of the pulsar velocity with respect to the solar system barycenter, making use of the distance estimates d_{π} and d_{LK} , respectively. The third, $v_{\perp,\text{DGR}}$, is an estimate of the pulsar velocity with respect to its standard of rest, taking into account the effects of differential Galactic rotation and the Sun’s peculiar velocity with respect to the local standard of rest. This estimate is also based on d_{LK} . In all three cases, the estimates are based on the marginal distribution of the perpendicular velocity, assuming the proper motions follow normal distributions with means and standard deviations as given by *Gaia* DR2, the distance follows the appropriate posterior distribution, and the distance and proper motion are independent. As with the distance estimates, we give the mode as a point estimate and the 16th and 84th percentiles as errors. Differential Galactic rotation is calculated based on a simple model of Galactic motion in which stars follow circular orbits around the center of the galaxy with a constant velocity Θ_0 . To maintain consistency with the Galactic pulsar population model of Lorimer et al. (2006), we set $\Theta_0 = 220 \text{ km s}^{-1}$ and the distance R_{\odot} from the Sun to the center of the Galaxy to 8.5 kpc, their 1985 IAU standard values (Kerr & Lynden-Bell, 1986). The components of the Sun’s

peculiar velocity are taken from Schönrich et al. (2010).

2.3 Results

The *Gaia* DR2 parallax and proper motion measurements are given in Table 2.2. Pulsars with more significant parallax measurements (less than 40% uncertainty in parallax) are given in the upper section of the table, and pulsars whose companions are readily identifiable in *Gaia* DR2 but which have less significant parallax measurements are given in the lower section. Previous distance estimates for the pulsars identified in Table 2.2 are given in Table 2.1, and these estimates are compared with the uncorrected *Gaia* distances (d_{π}) in Figure 2.1. Overall the agreement is good, and suggests that many of the alternate (non-astrometric) means of distance estimation for the pulsars have been reliable.

Of the pulsars with companions seen by *Gaia*, five also have pulsar timing or VLBI parallaxes, as indicated in Table 2.1. With the exception of PSR J0437–4715 (see below), all of these distance estimates are consistent with our *Gaia* DR2 results. Of these, PSR J1024–0719 does not have a significant parallax measurement in DR2. In the cases of PSRs J0437–4715 and J1012+5307, the most precise distance estimates have been achieved through orbital period derivative modeling (Reardon et al., 2016; Desvignes et al., 2016), and are significantly more precise than the *Gaia* distances. For PSR J1023+0038 the VLBI parallax is more precise than the *Gaia* value, while the reverse is true for PSR B1259–63.

In Figure 2.2 we compare the *Gaia* distances with the DM distances. We

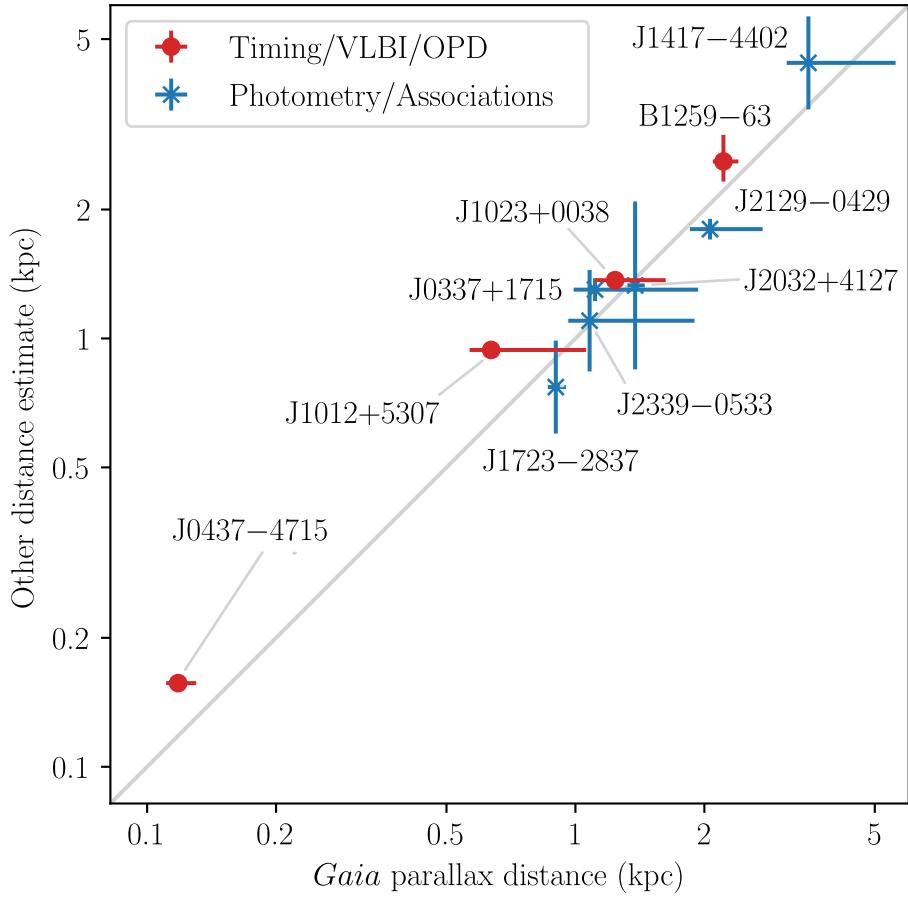


Figure 2.1: Comparison of *Gaia* parallax distances with the previous estimates of distance given in Table 2.1. Error bars represent $1-\sigma$ uncertainties or 68% credible intervals. For PSRs J1417–4402 and J1723–2837, which have distance determined assuming the companion overflows its Roche lobe, 25% error bars have been added for illustrative purposes. The *Gaia* distances are those determined without a Lutz-Kelker correction (d_π in Table 2.2).

Table 2.1: Comparison of *Gaia* parallaxes with other distance measurements

PSR	System Type ^a	DM	Distance			Other Method ^d	Refs.
			NE2001 ^b (pc cm ⁻³)	YMW16 ^c (kpc)	Other (kpc)		
Pulsars with significant <i>Gaia</i> parallax measurements							
J0337+1715	Triple	21.3	0.76	0.81	1.30 ± 0.08	Photometry	1
J0437-4715	He WD	20.4	0.14	0.16	0.15679 ± 0.00025	OPD	2
					0.1563 ± 0.0013	VLBI	3
J1012+5307	He WD	9.0	0.41	0.80	0.94 ± 0.03	OPD	4
J1023+0038	Transitional	14.3	0.62	1.11	$1.368^{+0.042}_{-0.039}$	VLBI	5
J1227-4853	Transitional	43.4	1.37	1.24	6
B1259-63	Be Binary	146.7	2.79	2.21	$2.59^{+0.37}_{-0.28}$	VLBI	7
					2.3 ± 0.4	Association	8
J1417-4402	Redback	55.0	1.60	2.16	4.4	RLO	9, 10
J1431-4715	He WD	59.4	1.57	1.82	11
J1723-2837	Redback	19.7	0.74	0.72	0.77	RLO	12
J2032+4127	Be Binary	114.7	3.65	4.62	1.33 ± 0.60	Association	13
J2129-0429	Redback	16.9	0.91	1.39	1.8 ± 0.1	Spectroscopy	14
J2339-0533	Redback	8.7	0.45	0.75	1.1 ± 0.3	Spectroscopy	15, 16
					0.7	Photometry	17
Pulsars with low significance <i>Gaia</i> parallax measurements							
J0045-7319	MS	105.4	> 36	58.7 ^e	60 ± 4	Association (SMC)	18
J0348+0432	He WD	40.5	2.08	2.26	2.1 ± 0.2	Spectroscopy	19
J1024-0719	MS	6.5	0.39	0.38	1.13 ± 0.18	Timing	20
					1.08 ± 0.04	Spectroscopy	21, 22
J1048+2339	Redback	16.7	0.70	2.00	23
J1311-3430	BW	20.5	1.41	2.43	24
J1628-3205	Redback	42.1	1.25	1.22	25
J1810+1744	Redback	39.7	2.00	2.36	26
J1816+4510	He WD	38.9	2.42	4.36	4.5 ± 1.7	Photometry	27
J1957+2516	Redback	44.1	3.07	2.66	28
J2215+5135	Redback	69.2	3.01	2.78	27

References—(1) Ransom et al. (2014); (2) Reardon et al. (2016); (3) Deller et al. (2008); (4) Desvignes et al. (2016); (5) Deller et al. (2012); (6) Roy et al. (2015); (7) Miller-Jones et al. (2018); (8) Negueruela et al. (2011); (9) Strader et al. (2015); (10) Camilo et al. (2016); (11) Bates et al. (2015); (12) Crawford et al. (2013); (13) Kiminki et al. (2015); (14) Bellm et al. (2016); (15) Romani & Shaw (2011); (16) Ray et al. (2020); (17) Kong et al. (2012); (18) Storm et al. (2004); (19) Antoniadis et al. (2013); (20) Guillemot et al. (2016); (21) Bassa et al. (2016); (22) Kaplan et al. (2016); (23) Deneva et al. (2016); (24) Romani et al. (2012); (25) Li et al. (2014); (26) Breton et al. (2013); (27) Kaplan et al. (2013); (28) Stovall et al. (2016).

^a We identify systems with helium-core white dwarf (WD) companions, Be binaries, redback systems, black widow (BW) systems, systems with main sequence (MS) companions, and transitional systems that alternate between accretion and rotation-powered. PSR J0337+1715 is a member of a triple system with two helium-core WDs.

^b DM-derived distance using the NE2001 model (Cordes & Lazio, 2002a).

^c DM-derived distance using the YMW16 model (Yao et al., 2017).

^d Methods used for DM-independent distance estimates: photometry and modeling of companion; direct VLBI astrometric parallax; combination of orbital period derivative (OPD) with kinematic models (Bell & Bailes, 1996); direct pulsar timing parallax; spectroscopic parallax; association with star cluster or galaxy; modeling of lightcurve assuming Roche lobe overflow (RLO).

^e Using YMW16 in the Magellanic Cloud mode.

Table 2.2: Distances and velocities derived from *Gaia* parallaxes

PSR	G mag.	π (mas)	μ_α (mas yr $^{-1}$)	μ_δ (mas yr $^{-1}$)	d_π (kpc)	d_{LK} (kpc)	$v_{\perp,\pi}$ (km s $^{-1}$)	$v_{\perp,\text{LK}}$ (km s $^{-1}$)	$v_{\perp,\text{DGR}}$ (km s $^{-1}$)
Pulsars with significant <i>Gaia</i> parallax measurements									
J0337+1715	18.08	0.73(24)	4.81(49)	-4.42(42)	1.11 $^{+0.82}_{-0.12}$	1.34 $^{+1.07}_{-0.17}$	34 $^{+26}_{-4}$	46 $^{+44}_{-6}$	48 $^{+44}_{-6}$
J0437-4715	20.41	8.33(67)	122.9(11)	-71.2(16)	0.118 $^{+0.012}_{-0.007}$	0.121 $^{+0.013}_{-0.008}$	79.5 $^{+8.2}_{-5.1}$	82.1 $^{+9.0}_{-5.3}$	78.8 $^{+9.0}_{-5.4}$
J1012+5307	19.63	1.33(41)	2.98(52)	-26.94(63)	0.64 $^{+0.42}_{-0.07}$	0.79 $^{+0.73}_{-0.09}$	82 $^{+55}_{-9}$	113 $^{+133}_{-12}$	113 $^{+132}_{-12}$
J1023+0038	16.27	0.73(14)	4.75(13)	-17.35(13)	1.24 $^{+0.39}_{-0.13}$	1.32 $^{+0.43}_{-0.14}$	106 $^{+33}_{-11}$	116 $^{+42}_{-13}$	116 $^{+46}_{-14}$
J1227-4853	18.08	0.62(16)	-18.73(20)	7.39(11)	1.37 $^{+0.69}_{-0.15}$	1.70 $^{+1.72}_{-0.17}$	131 $^{+66}_{-15}$	180 $^{+257}_{-15}$	144 $^{+187}_{-12}$
B1259-63	9.63	0.418(30)	-6.986(43)	-0.416(44)	2.21 $^{+0.19}_{-0.12}$	2.26 $^{+0.20}_{-0.13}$	73.5 $^{+6.1}_{-4.1}$	75.4 $^{+6.7}_{-4.3}$	36.8 $^{+1.3}_{-1.1}$
J1417-4402	15.79	0.221(71)	-4.70(10)	-5.10(87)	3.50 $^{+2.09}_{-0.38}$	4.06 $^{+2.24}_{-0.52}$	115 $^{+69}_{-13}$	144 $^{+86}_{-20}$	118 $^{+33}_{-15}$
J1431-4715	17.75	0.64(16)	-12.01(33)	-14.51(26)	1.35 $^{+0.69}_{-0.15}$	1.70 $^{+2.10}_{-0.15}$	121 $^{+59}_{-13}$	169 $^{+324}_{-10}$	170 $^{+263}_{-9}$
J1723-2837	15.55	1.077(54)	-11.713(82)	-23.990(62)	0.90 $^{+0.05}_{-0.04}$	0.91 $^{+0.05}_{-0.04}$	114 $^{+7}_{-5}$	115 $^{+7}_{-5}$	124 $^{+7}_{-5}$
J2032+4127	11.36	0.693(33)	-2.991(48)	-0.742(55)	1.38 $^{+0.07}_{-0.06}$	1.39 $^{+0.08}_{-0.06}$	20.1 $^{+1.1}_{-0.8}$	20.4 $^{+1.2}_{-0.9}$	25.7 $^{+1.2}_{-1.0}$
J2129-0429	16.84	0.424(88)	12.34(15)	10.19(14)	2.06 $^{+0.67}_{-0.21}$	2.18 $^{+0.72}_{-0.23}$	157 $^{+51}_{-16}$	172 $^{+61}_{-19}$	191 $^{+77}_{-23}$
J2339-0533	18.97	0.75(25)	4.15(48)	-10.31(30)	1.08 $^{+0.82}_{-0.12}$	1.25 $^{+0.84}_{-0.16}$	57 $^{+47}_{-6}$	72 $^{+56}_{-10}$	64 $^{+54}_{-9}$
Pulsars with low significance <i>Gaia</i> parallax measurements									
J0045-7319	16.22	0.040(60)	0.31(11)	-0.931(99)
J0348+0432	20.64	-2.0(10)	3.1(20)	-0.1(14)
J1024-0719	19.18	0.53(42)	-35.52(64)	-47.93(65)
J1048+2339	19.65	0.96(80)	-16.3(10)	-11.7(12)
J1311-3430 ^a	20.53
J1628-3205	19.52	1.20(56)	-6.4(10)	-19.81(82)
J1810+1744	20.08	1.05(69)	6.4(17)	-7.2(20)
J1816+4510	18.22	0.22(15)	-0.17(29)	-4.42(33)
J1957+2516	20.30	0.69(86)	-5.7(11)	-8.9(14)
J2215+5135	19.24	0.28(36)	0.31(54)	1.88(60)

NOTE—Quantities in parentheses represent the 1- σ uncertainties reported by *Gaia*. Along with astrometric parameters, we give the *Gaia* G-band mean magnitude of the associated source. Pulsars with detected companions (§ 2.2.1) but parallax measurements with errors greater than 40% of their values are shown below the line. For each pulsar, d_π is the distance estimate calculated assuming a uniform prior on parallax, and may be considered a maximum-likelihood estimate, while d_{LK} is the Lutz-Kelker corrected distance, calculated using the volumetric prior described in § 2.2.2. Both distance estimates take into account the zero-point offset of -0.029 mas given by Lindegren et al. (2018). Velocities are calculated using the *Gaia* proper motions given here. The transverse velocity $v_{\perp,\pi}$ of the pulsar is calculated using d_π as a distance estimate, while $v_{\perp,\text{LK}}$ is calculated using d_{LK} . The third velocity $v_{\perp,\text{DGR}}$ is an estimate of the pulsar’s velocity with respect to its standard of rest. It also uses d_{LK} as a distance estimate, and includes corrections due to differential Galactic rotation and the peculiar motion of the Sun. For all distances and velocities, we give the posterior mode as a point estimate, and the 16th and 84th percentiles of the posterior distribution as errors.

^a The companion to PSR J1311-3430 is detected in *Gaia* DR2, but only the two-parameter astrometric solution (position on the sky) is reported, because the five-parameter solution (including parallax and proper motion) failed the acceptance criteria detailed in Lindegren et al. (2018). A full solution is expected in future data releases, and so it is included here for the sake of completeness.

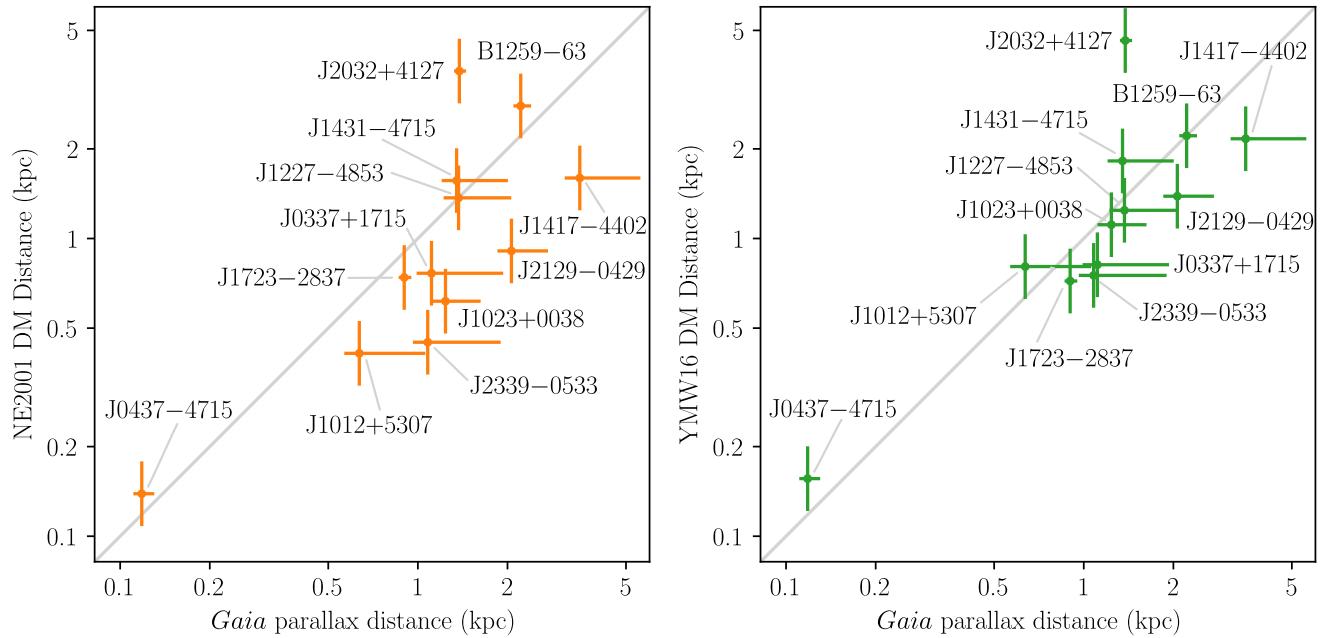


Figure 2.2: Comparison of *Gaia* parallax distances with DM distances calculated using the NE2001 (left; Cordes & Lazio, 2002a) and YMW16 (right; Yao et al., 2017) Galactic electron density models. As in Figure 2.1, the *Gaia* distances are those determined without a Lutz-Kelker correction (d_π in Table 2.2), and the error bars represent $1-\sigma$ uncertainties. The error bars on the DM distances are illustrative, and are estimated at 25%. PSR J0045–7319, which is in the Small Magellanic Cloud, is not shown.

see that both DM models do reasonably well in comparison with *Gaia*, with YMW16 showing somewhat better agreement for the sample of objects considered here. An assumed 25% uncertainty gets the majority of points consistent between the astrometric and DM distances. However, the distances predicted by both DM models appear to be somewhat underestimated compared to those measured by *Gaia*—the median ratio of DM distance to *Gaia* distance is 0.76 for NE2001 and 0.90 for YMW16. This agrees roughly with the findings of Gaensler et al. (2008) and Roberts (2011), who claimed that for pulsars out of the Galactic plane (which is true for most recycled pulsars), distances predicted by the NE2001 model were underestimated.

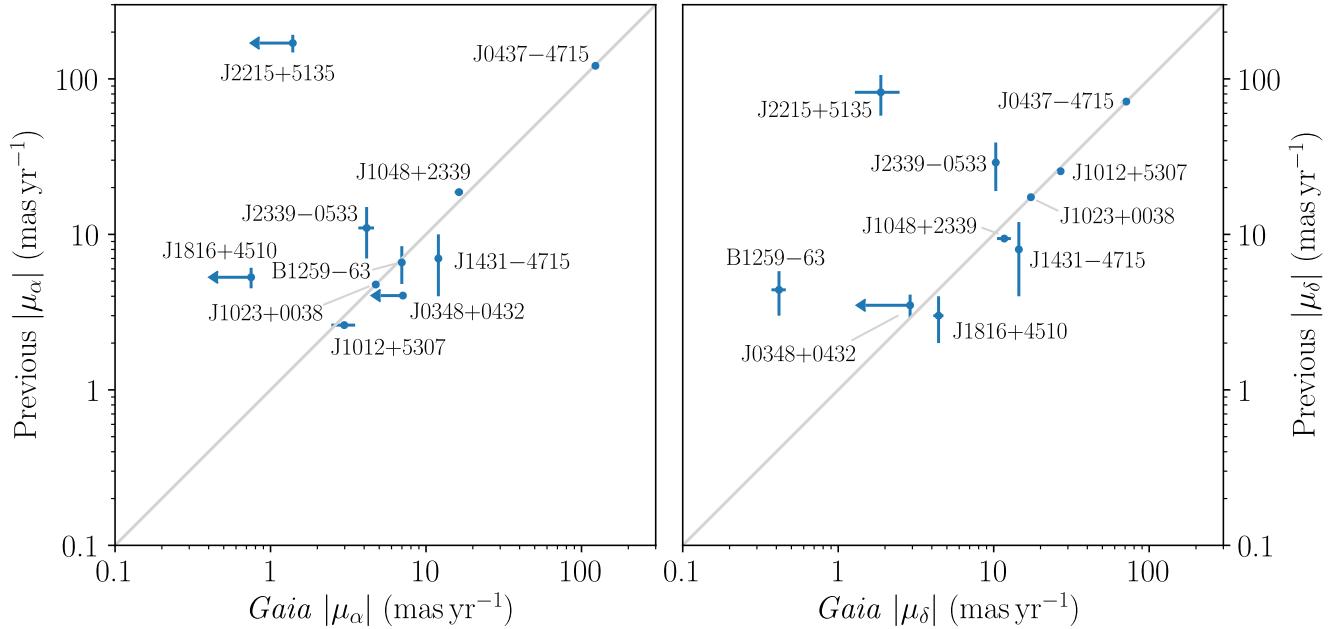


Figure 2.3: Comparison of *Gaia* proper motions with previous proper motion measurements from the ATNF pulsar catalog (Manchester et al., 2005, 2016). Proper motions in right ascension α include the geometric correction ($\mu_\alpha = \dot{\alpha} \cos \delta$). We plot the absolute values of the proper motions to allow for better comparison over a large dynamic range. Sources without 3σ detections of *Gaia* proper motions are plotted as 2σ upper limits.

We compare the proper motions from *Gaia* and the literature in Figure 2.3. Again the agreement for most sources is reasonably good, with the sources with the most precise prior measurements (PSRs J0437-4715, J1012+5307, J1023+0038) agreeing very well.

2.3.1 Notes On Individual Objects

J0337+1715 This pulsar is in a hierarchical triple system with two white dwarfs (Ransom et al., 2014). The orbital period of the outer white dwarf companion is 327 days, which is comparable to the 1-year parallactic period

and not too much smaller than the 22-month data span of *Gaia* DR2. As a result, the orbital motion of the inner white dwarf (which is the optically dominant one) with the outer white dwarf may have caused systematic errors in the proper motion and (especially) parallax measurements. However, we find reasonable agreement with the distance inferred from white dwarf photometry (Kaplan et al., 2014) and from VLBI astrometry (Deller et al., in prep.).

J0437–4715 This is the closest system in our sample by a factor of about 5 (and one of the closest known neutron stars). The most precise distance measurement comes from orbital period derivative modeling (Reardon et al., 2016). The *Gaia* distance is only about 77% of this value, and the results are formally inconsistent. The inconsistency may be due to unmodeled binary reflex motion, as the companion should trace out an ellipse with size ≈ 0.3 mas (based on the masses from Verbiest et al. 2008) every 5.5 d, comparable to the size of the parallax uncertainty. Most other systems (with the exceptions of the Be binaries discussed below, the outer companion of PSR J0377+1715, and PSR J1024–0719; see Bassa et al. 2016; Kaplan et al. 2016) have much smaller projected orbital separations. The discrepancy may be resolved with future *Gaia* data releases when the individual astrometric measurements are released and binary orbits are included in the fits.

B1259–63 The system orbital period is 1237 days, which less than twice the 22-month span of the data used to derive the *Gaia* astrometric parameters, and may have caused residual systematic errors in those parameters.

J1417–4402 Here we find that the *Gaia* distance constraint is much more consistent with the distance estimate from assuming that the companion fills

its Roche lobe (4.4 kpc, Strader et al., 2015; Camilo et al., 2016) than the dispersion-measure distance of 1.6 kpc.

J1816+4510 The *Gaia* proper motion μ_α disagrees significantly with the one from pulsar timing, a discrepancy to be investigated further.

J2032+4127 This is a long-period (46 ± 2 yr) binary system with a wide orbit. Reflex motion of the companion may have affected the *Gaia* astrometric fit, especially for proper motion.

J2215+5135 While there is no significant parallax for this source, the *Gaia* proper motion (magnitude 1.9 ± 0.8 mas yr $^{-1}$) is much smaller than the value listed in Abdo et al. (2013, magnitude 189 ± 23 mas yr $^{-1}$). We believe the latter to be an error: it would also imply an unlikely transverse velocity of 2700 km s $^{-1}$ at the DM distance.

2.4 Discussion

As discussed above, we focus on binary systems since they are, with the exception of the Crab pulsar, the only pulsar systems with optical emission detectable by *Gaia*. We find that the systems we have identified fall into 3 broad categories: recycled pulsars with white dwarf companions; recycled pulsars with non-degenerate companions (so-called “black widows” and “redbacks”; Roberts 2013); and young pulsars with massive B star companions.

Both of the first two classes form from similar progenitors, where millisecond pulsars (MSPs) are produced in binary systems when old NS are “recycled” by accreting matter from the companion star. After recycling, the evolution of such systems can take a number of paths, with some diverging and

leaving white dwarf companions and some converging and gradually destroying their companions through accretion and ablation (Tauris, 2011; Chen et al., 2013). Some of the converging systems are visible as redbacks, typically with distorted main-sequence-like companions of a few tenths of a solar mass, or black widows with companions of a few hundredths of a solar mass (Roberts, 2013). Some of the systems that we identified are still in the process of transitioning from accreting to rotation-powered (e.g., the transitional millisecond pulsar J1023+0038; Archibald et al., 2009), but resemble redbacks when not accreting. For the white dwarf systems, *Gaia* detections are biased in favor of the lower-mass (and hence larger) companions, especially as many of them are helium-core white dwarfs that can have hydrogen shell burning (stable or sporadic) which keeps them hotter and brighter for Gyr (e.g., Althaus & Benvenuto, 1998; Driebe et al., 1999; Istrate et al., 2016).

The young pulsars such as PSR B1259–63 follow a different evolutionary path, with wide orbits around massive (and luminous) stars. However, all of these systems share characteristics that made them suitable for *Gaia* detections: they are all both bright and nearby.

Our sample above contains a number of objects of individual interest. Some are used as laboratories for tests of alternative theories of gravity (PSR J0348+0432, Antoniadis et al. 2013; PSR J0337+1715, Ransom et al. 2014), while others (PSR J1023+0038, Shahbaz et al. 2018) help test theories of accretion, and the Be systems are testbeds for particle acceleration (PSR B1259–63, Dubus 2013).

The ensemble of distances to the sample of recycled pulsars is also of particular interest. A major effort is underway to regularly time MSPs in pul-

sar timing arrays (PTAs) to detect low-frequency (nHz) gravitational waves (e.g., Bizouard et al., 2013; Detweiler, 1979; Lommen, 2015). Precise measurements of the distances to MSPs can significantly improve the PTA sensitivity when searching for coherent signals from individual sources (Madison et al., 2016). The most relevant pulsars for PTA purposes are the most stable, particularly fully-recycled millisecond pulsars which typically have low-mass white dwarf companions. Our sample includes four such systems, as well as the partially-recycled pulsar PSR J0348+0432. *Gaia* provides the first independent distance estimate for PSR J1431–4715, while for PSR J1012+5307, it provides a distance estimate with precision comparable to that achievable with pulsar timing. However, the *Gaia* parallax distance estimates cannot compete with the precision achieved with VLBI for PSR J0437–4715, and the *Gaia* parallax measurement for the distant PSR J1816+4510 will likely remain insignificant until the end of *Gaia*’s nominal mission (see below).

Gaia DR2 is based on a $T = 22$ month data span. The final *Gaia* data release, planned for 2022, will be based on data collected during *Gaia*’s entire five-year nominal mission. Since errors in parallax measurements scale as $T^{-1/2}$ and errors in proper motion scale as $T^{-3/2}$, the additional 38 months of data should allow the errors in parallax to be reduced to approximately 60% of their current values, while the errors in proper motion should decrease by a factor of approximately 4.

While *Gaia* DR2 is essentially complete for sources with magnitude $G < 17$ (Gaia Collaboration et al., 2018), many pulsar companions are fainter than that limit, and it is likely that at least some new pulsar companions will be present in the final data release, which is expected to be complete to mag-

nitude $G \approx 20.7$ (Gaia Collaboration et al., 2016). As mentioned previously, redbacks and other pulsars with main sequence-like companions present the best targets for *Gaia*, but the 7 redback systems and 3 other pulsars with main sequence-like companions given in Table 2.1 already represent the majority of known examples of such systems. It is most likely that new pulsar companions detected in forthcoming *Gaia* data releases will represent both redbacks and millisecond pulsars with white dwarf companions, several of which also appear in the DR2 data discussed here.

In *Gaia* DR2, all sources are modeled as single stars (Gaia Collaboration et al., 2018). This raises the possibility that parallax and proper motion measurements for objects in binaries, like those considered here, may be corrupted by unmodeled orbital motion. A possible indicator of this corruption is the *Gaia* excess noise parameter (Lindegren et al., 2018, `astrometric_excess_noise` in the *Gaia* Archive), which is a measure of the size of effects not accounted for by the single-star model. Most companions to redback pulsars, and some Be stars, also show significant flux variability, but this is unlikely to have a systematic effect on the *Gaia* parallax estimates, since the period of the flux variations is not a simple fraction of Earth’s orbital period. For the majority of the sources in Table 2.2, corruption by orbital motion is unlikely to be an issue, as their orbital periods are small compared to the 22-month observing time of *Gaia* DR2, and any orbital motion will average out. For the nearest pulsar in our sample, PSR J0437–4715, the orbit of the white dwarf companion spans 0.3 mas on the sky, which is about half the uncertainty in its parallax. However, its orbital period is 5.74 days, meaning that it completes approximately 64 orbits in a year, and 116 in the 22-month *Gaia* DR2 observing span, so the contribution of this orbital motion to the parallax and proper motion measure-

ments should largely average out. Additionally, *Gaia* reports no excess noise for the source associated with this pulsar, which is perhaps surprising considering that its parallax gives a distance which is somewhat inconsistent with the highest-precision value derived from orbital period derivative modeling.

On the other hand, PSR B1259–63 has an orbital period of 1237 days, and the triple system PSR J0337+1715 has an outer orbital period of 327 days, both of which are comparable to the observing time, and so the *Gaia* parallaxes and proper motions of these pulsars may be affected by a systematic error due to orbital motion. That appears to be evident in the quality of the *Gaia* astrometric fit and the parallax uncertainty compared with what *Gaia* predicts for a source of that G magnitude². For PSR J0337+1715 the reduced χ^2 of the astrometric fit is 1.8, on the high end for our sources, and the measured parallax uncertainty of 0.25 mas agrees very well with the predicted uncertainty of 0.21 mas. For PSR B1259–63 the fit reduced χ^2 is even higher at 2.0, and the measured parallax uncertainty of 0.03 mas exceeds the predicted uncertainty of 0.01 mas. The only other source with a high reduced χ^2 and a large ratio of measured to predicted parallax uncertainty is PSR B2032+4127: like PSR B1259–63, this is a Be binary with a bright ($G = 11.4$) companion. For PSR J0337+1715, the excess noise reported by *Gaia* is approximately 0.3 mas, while for PSR B1259–63, no excess noise is reported. The only other pulsar with a significant parallax and nonzero excess noise parameter is PSR J1431–4715, for which the excess noise is again around 0.3 mas. So while there may be some corruption of the astrometric fits due to unmodeled orbital motion, the bigger effect may just be an overestimate of the attainable precision for the brightest sources. The full *Gaia* data-sets with longer timespan, better calibration, and binary model fits,

²Using the formula from <https://www.cosmos.esa.int/web/gaia/science-performance> scaled by 1.5 to account for the limited duration of DR2.

will help to resolve the question. Ultimately, the *Gaia* data represent a small but potentially important contribution to the pulsar distance scale, and will be incorporated in future versions of the pulsar population analysis and electron density models.

2.5 Acknowledgements

We thank C. Mingarelli for valuable discussions regarding pulsar cross-matching and the parallax zero-point offset, and W. Vlemmings for helpful comments. Additionally, we thank the anonymous referee for providing useful feedback. RJJ, DLK, SC, JMC are members of the NANOGrav Physics Frontiers Center, which is supported by the National Science Foundation award 1430284. ATD is the recipient of an Australian Research Council Future Fellowship (FT150100415).

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

The Pan-STARRS1 Surveys (PS1) and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching.

ing, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen’s University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation Grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation. The national facility capability for SkyMapper has been funded through ARC LIEF grant LE130100104 from the Australian Research Council. SkyMapper is owned and operated by The Australian National University’s Research School of Astronomy and Astrophysics.

CHAPTER 3

DETECTING GRAVITATIONAL SCATTERING OF INTERSTELLAR OBJECTS USING PULSAR TIMING

Gravitational scattering events, in which the path of an interstellar object is deflected by a pulsar or the solar system, give rise to reflex motion which can potentially be detected using pulsar timing. We determine the form of the timing signal expected from a gravitational scattering event, which is ramp-like and resembles the signal produced by a glitch or a gravitational wave burst with memory (BWM), and investigate the prospects for detecting such a signal using a pulsar timing array. The level of timing precision currently achieved for some millisecond pulsars makes it possible to detect objects as small as $10^{-10} M_{\odot}$, less than the mass of the dwarf planet Ceres, at impact parameters as large as 1 au. The signals produced by gravitational scattering could provide independent constraints on models of dark matter involving asteroid-mass objects or subhalos, and should be considered as potential false positives in searches for BWMs.

3.1 Introduction

If a compact astrophysical object were to pass near enough to a pulsar or to the solar system to interact gravitationally, it would slightly alter the motion of the pulsar relative to the solar system barycenter, potentially producing a detectable effect on the pulse times of arrival (TOAs). The North American Nanohertz Observatory for Gravitational Waves (NANOGrav; Arzoumanian

Published: Jennings, R. J., Cordes, J. M., & Chatterjee, S. 2020, ApJ, 889, 145.

et al. 2018a) and other pulsar timing arrays (PTAs; Desvignes et al. 2016; Hobbs 2013) have collected hundreds of pulsar-years of high-precision (sub-microsecond) pulsar timing data as part of their efforts to detect gravitational waves. The various regional PTAs collaborate to form the international pulsar timing array (IPTA; Hobbs et al. 2010), which occasionally releases combined data sets (e.g. Verbiest et al. 2016; Perera et al. 2019). The qualities of PTAs that make them well suited for detecting gravitational waves – namely, the precision of arrival time measurements and the long spans of time over which they may be collected – also make them highly sensitive to perturbations of this type.

An interstellar object (ISO) gives rise to a gravitational scattering signal in pulsar timing data if it passes close enough to a pulsar or to the solar system. One important category of ISOs consists of free-floating planets and smaller asteroid- and comet-like bodies. Almost certainly, a large number of these bodies exist in interstellar space, since the planet formation process is thought to result in the ejection of large numbers of rocky bodies, ranging in size from planetesimals to fully-formed planets (Charnoz & Morbidelli, 2003). Some such bodies have even been directly observed: in particular, the asteroid 1I/‘Oumuamua (Meech et al., 2017) and the comet 2I/Borisov (Guzik et al., 2019) have both been identified as interstellar in origin. To produce a detectable timing signal, an asteroid-like ISO would have to be a few hundred kilometers in diameter and pass within a few au of a pulsar. This is significantly larger than either ‘Oumuamua or Borisov, but smaller than the largest objects in the asteroid belt.

Some theories of dark matter predict that at least a fraction of it is composed of massive ISOs, such as primordial black holes (PBHs) or subhalos. As

a result, several searches for ISOs have been carried out with the goal of understanding the nature of dark matter. For the most part, these have involved microlensing surveys, which are sensitive to any object with mass along the line of sight to a star. The pioneering MACHO (Alcock et al., 2001), EROS (Tisserand et al., 2007), and OGLE (Wyrzykowski et al., 2011) surveys largely ruled out the possibility that the majority of dark matter could consist of objects between 10^{-8} and $100 M_{\odot}$, but did not significantly constrain objects smaller than $10^{-8} M_{\odot}$ (approximately 20 times the mass of Ceres). Subsequently, attention has focused on primordial black holes with masses between 10^{-11} and $10^{-8} M_{\odot}$ as dark matter candidates (Carr et al., 2016). The strongest limits on PBHs and other relatively compact objects in this mass range come from a recent microlensing survey of the Andromeda galaxy using the Hyper Suprime Cam (HSC) on the 8.2-meter Subaru telescope (Niikura et al., 2019). The implications of these constraints for the expected event rate are discussed further in Section 3.6.

A PBH or similar dark matter constituent would interact with a pulsar in a manner indistinguishable from a baryonic ISO of equal mass. If they constituted a significant fraction of dark matter by mass, PBHs would also be much more numerous than baryonic ISOs of the same mass. Because of this, pulsar timing has been proposed as a means of detecting PBHs. Some proposals have focused on the Shapiro delay signal caused by objects along the line of sight to the pulsar (Siegel et al., 2007; Clark et al., 2016), but others (Seto & Cooray, 2007; Kashiyama & Seto, 2012; Kashiyama & Oguri, 2018; Dror et al., 2019) have discussed the Doppler signal, which is more closely related to the timing signal described here.

Detectable encounters with ISOs are likely to be rare, but examining a large volume of data makes their observation more likely. In this respect, it is important to distinguish between encounters with a pulsar and encounters with the solar system. By analogy with the terminology used for gravitational wave signals, we call the former “pulsar-term” scattering events and the latter “Earth-term” events. If many pulsars are observed, the two scenarios give complementary sensitivities — since Earth-term events affect all pulsars, it is possible to detect weaker events by cross-correlating the signals, but objects are more likely to pass near one of the pulsars in a PTA simply because there are more pulsars. This means that Earth-term events are more detectable for large populations of very small objects, but pulsar-term events are more detectable for small populations of larger objects.

In what follows, we calculate the shape of the timing signal expected from a gravitational scattering event in detail, and use the results to assess the circumstances under which such events may be detectable. In Section 3.2, we discuss the factors which limit the precision of pulsar timing measurements. In Section 3.3, we derive the expected timing signal. In Section 3.4, we discuss the conditions under which pulsar-term events may be detected. Section 3.5 describes the corresponding conditions for Earth-term events. In Section 3.6, we review methods of estimating the number density of potential perturbers in the galaxy, and speculate as to the frequency of detectable events. Finally, in Section 3.7, we summarize our results and draw conclusions.

3.2 Pulsar Timing Precision

Contemporary pulsar timing methods rely on the fact that every pulsar has a characteristic average pulse shape which is stable on time scales of decades. Arrival times are determined by comparing observed pulse shapes with the template pulse shape using a Fourier-domain matched filtering algorithm (Taylor, 1992). Generally, averages with $N \gg 1$ pulses are used, with typical values of N being between 10^5 and 10^6 .

The minimum uncertainty in estimating an arrival time comes from noise in the pulse profile measurement introduced by the receiver (radiometer noise). Computing TOAs using a matched filtering algorithm minimizes this contribution to the arrival time error for a given level of radiometer noise. In the absence of other sources of error, the uncertainty in arrival times computed using matched filtering is given by (Cordes, 2013a; Lam et al., 2016):

$$\sigma_{\text{MF}} = \frac{\sqrt{PW_{\text{eff}}}}{S\sqrt{N_\phi}}. \quad (3.1)$$

Here P is the pulse period, S is the signal-to-noise ratio (SNR), N_ϕ is the number of samples in pulse phase, and W_{eff} is an effective pulse width, given by

$$W_{\text{eff}} = \left(\int_0^P U'(t)^2 dt \right)^{-1}, \quad (3.2)$$

where $U(t)$ is the pulse shape (normalized to unit maximum).

In practice, TOA errors are larger than equation (3.1) would predict, because other effects contribute. These include pulse jitter from the motion of emission regions in pulsar magnetospheres; dispersion and scintillation, caused by propagation of the signal through the ionized interstellar medium; and spin noise, caused by interactions between the neutron star crust and its magnetosphere

and superfluid interior. Of these, pulse jitter is uncorrelated in time, but scintillation, dispersion measure variations, and spin noise are generally correlated, producing gradual but random drifts in pulse times of arrival. Spin noise has a “red” power spectrum with most of the power concentrated at low frequencies. By contrast, radiometer and jitter noise have a “white” power spectrum, contributing approximately equal power at all frequencies.

Broadly speaking, pulsars can be divided into two categories: canonical pulsars (CPs), which are relatively young and tend to have periods of order one second and surface magnetic fields of order 10^{12} gauss; and millisecond pulsars (MSPs), which are much older and are thought to have been spun up by accreting matter from a companion. MSPs have significantly shorter periods (a few milliseconds) and weaker magnetic fields (of order 10^8 gauss) than canonical pulsars. Additionally, MSPs spin down much more slowly and have lower levels of spin noise (by a factor of around 10^6) than CPs (Shannon & Cordes, 2010; Lam et al., 2017; Parthasarathy et al., 2019). In part because of their short periods, pulses from MSPs can also be localized more precisely (cf. equation 3.1). For these reasons, pulsar timing arrays, which require very high-precision timing, almost exclusively time MSPs.

For many MSPs, TOA precision significantly better than 1 μ s is already routinely achieved. For example, 38 of the 45 pulsars included in the NANOGrav 11-year data set (Arzoumanian et al., 2018a) had a median TOA uncertainty in L-band (1–2 GHz) observations which was below 1 μ s, with the best-timed pulsar, PSR B1937+21, having a median uncertainty of a mere 12 ns. The overall median TOA uncertainty in the data set is 319 ns. A detailed breakdown of per-pulsar contributions to excess noise in a previous NANOGrav data release,

the 9-year data set, is given by Lam et al. (2017).

Because of its red spectrum, spin noise presents an additional challenge in searches for very low-frequency and quasi-static effects. While not all MSPs possess detectable levels of red spin noise, many do. We defer a full analysis of the effects of red spin noise on the detectability of scattering events to future work, but note that its effects are likely to be significant for events with large impact parameter, b , for which most of the signal power is concentrated at low frequencies (see Section 3.3 below).

The demands on timing precision for gravitational wave detection require attention to all processes that contribute more than ~ 10 ns RMS to timing residuals. Shannon et al. (2013) considered the effects of asteroid belts around isolated MSPs and in particular determined that the nonstationary timing residuals for B1937+21 are not inconsistent with an asteroid belt interpretation.

3.3 TOA Perturbations

The detectable effect of a perturbing object on measured TOAs is a result of the reflex motion of the pulsar or the solar system barycenter. For concreteness, we assume the perturbing object passes near the pulsar, with the understanding that the roles of the pulsar and the solar system barycenter can be interchanged if the perturber passes through the solar system. We also treat the perturber as a point mass, and assume that general relativistic effects and non-gravitational forces (such as the Yarkovsky effect, e.g. Rubincam 1998) can be ignored. A potential complicating factor is the fact that many MSPs are found in binary systems with white dwarfs, with typical separations between 0.01 and 0.6 au.

In this case, we are concerned with the reflex motion of the center of mass of the pulsar system. If the perturber comes closer than a few times the binary separation, three-body interactions, the effects of which are beyond the scope of this paper, may become important.

The position, \mathbf{r} , of the perturber relative to the center of mass of the pulsar-perturber system follows a hyperbolic trajectory, parameterized as

$$\mathbf{r}(H) = \frac{b(e - \cosh H)}{\sqrt{e^2 - 1}} \hat{\mathbf{x}} + b \sinh H \hat{\mathbf{y}} \quad (3.3)$$

(cf. Roy 1988, Section 4.7). Here b is the impact parameter, e is the eccentricity ($e > 1$ for a hyperbola), and H is the hyperbolic anomaly. The unit vector $\hat{\mathbf{x}}$ points from focus to periapse, and $\hat{\mathbf{y}}$ is the unit vector perpendicular to $\hat{\mathbf{x}}$ in the plane of motion, oriented such that $\hat{\mathbf{x}} \times \hat{\mathbf{y}}$ is in the direction of the orbital angular momentum (see Fig. 3.1).

The eccentricity, e , is related to the asymptotic velocity, v , of the perturber relative to the pulsar, by

$$e = \sqrt{1 + \left(\frac{bv^2}{GM}\right)^2}, \quad (3.4)$$

where G is the universal gravitational constant and M is the total mass of the pulsar and the perturber. All the perturbers considered here have masses much less than $1 M_\odot$, so, to a good approximation, M is equal to the mass of the pulsar (or, in the Earth-term case, the Sun). The asymptotic velocity, v , appearing here is also called the hyperbolic excess velocity. The hyperbolic anomaly, H , is related to time, t , by the hyperbolic Kepler equation:

$$t = t_0 + \frac{b(e \sinh H - H)}{v \sqrt{e^2 - 1}}. \quad (3.5)$$

Here t_0 is the time of periapsis, at which $H = 0$. Importantly, as long as $e > 1$, t is an increasing function of H , and equation (3.5) may be inverted to give H

as a function of t .

The pulsar's position, \mathbf{x} , relative to the solar system barycenter is given by

$$\mathbf{x}(t) = \bar{\mathbf{x}}(t) - \frac{m}{M} \mathbf{r}(t), \quad (3.6)$$

where

$$\bar{\mathbf{x}}(t) = \bar{\mathbf{x}}_0 + \bar{\mathbf{v}}t \quad (3.7)$$

is the unperturbed trajectory of the pulsar, which coincides with the trajectory of the pulsar-perturber center of mass when a perturber is present. Letting

$$\Delta\mathbf{x}(t) = \mathbf{x}(t) - \bar{\mathbf{x}}(t) \quad (3.8)$$

be the perturbation to the pulsar's position, and using (3.3) for \mathbf{r} , we find

$$\Delta\mathbf{x}(t) = -\frac{mb}{M} \left(\frac{e - \cosh H(t)}{\sqrt{e^2 - 1}} \hat{\mathbf{x}} + \sinh H(t) \hat{\mathbf{y}} \right). \quad (3.9)$$

Perturbations to the times of arrival of pulses are caused by changes in the length $d = |\mathbf{x}|$ of the path from the pulsar to the solar system barycenter. Let $d_0 = |\mathbf{x}_0|$ be the path length at $t = 0$ and $\hat{\mathbf{n}}$ be the unit vector pointing from the position of the pulsar at $t = 0$ to the solar system barycenter, so that $\bar{\mathbf{x}}_0 = -d_0 \hat{\mathbf{n}}$.

We then have

$$d = d_0 \sqrt{1 - \frac{2\Delta\mathbf{x} \cdot \hat{\mathbf{n}}}{d_0} + \frac{|\Delta\mathbf{x}|^2}{d_0^2}}. \quad (3.10)$$

For $d_0 \gg |\Delta\mathbf{x}|$, d can be expanded in powers of $\|\Delta\mathbf{x}\|/d_0$:

$$d \approx d_0 - \Delta\mathbf{x} \cdot \hat{\mathbf{n}} + \frac{\|\Delta\mathbf{x}\|^2 - (\Delta\mathbf{x} \cdot \hat{\mathbf{n}})^2}{2d_0}. \quad (3.11)$$

The successive terms in equation (3.11) have, respectively, constant, dipolar, and quadrupolar dependences on \mathbf{x} . Since the constant term contributes only to the absolute pulse phase, the TOA perturbations are given by

$$\Delta\tau = \frac{d - d_0}{c} \approx -\frac{\Delta\mathbf{x} \cdot \hat{\mathbf{n}}}{c} + \frac{\|\Delta\mathbf{x}\|^2 - (\Delta\mathbf{x} \cdot \hat{\mathbf{n}})^2}{2cd_0}. \quad (3.12)$$

Using equations (3.6)–(3.9), this becomes

$$\Delta\tau \approx -\frac{\bar{\mathbf{v}} \cdot \hat{\mathbf{n}}}{c} t + \frac{|\bar{\mathbf{v}} \times \hat{\mathbf{n}}|^2}{2cd_0} t^2 + \frac{m\mathbf{r} \cdot \hat{\mathbf{n}}}{Mc}, \quad (3.13)$$

where terms involving products of the small quantities $|\mathbf{r}|/d_0$ and $|\bar{\mathbf{v}}|/c$ have been dropped. The remaining terms represent, respectively, the integrated Doppler shift of the pulsar's spin period; the Shklovskii correction to the period derivative (Shklovskii, 1970); and the delay caused by the reflex motion of the pulsar. From this point forward, we will ignore the Doppler and Shklovskii terms, as the low-order polynomial effects they introduce are degenerate with the period and intrinsic spin-down rate of the pulsar, and focus only on the perturbation due to reflex motion:

$$\Delta\tau = \frac{m\mathbf{r} \cdot \hat{\mathbf{n}}}{Mc}. \quad (3.14)$$

The orientation of the interaction relative to the observer can be characterized by two angles (see Fig. 3.1) — the inclination, i , of the orbital plane with respect to the plane of the sky, and the argument of periape, ω , the angle between the line of nodes and the line of apsides. A third angle, the longitude of the ascending node, Ω , gives the orientation of the line of apsides relative to a reference direction, and is necessary to fully specify the orientation of the orbit with respect to a fixed coordinate system. However, it does not affect the distance between the observer and any point on the orbit, so it is irrelevant to the size of the timing perturbations under consideration.

We use the convention that ω is measured in the direction of motion and is positive when the angle from the line of apsides to the line of nodes, in the direction of motion, is acute. Furthermore, i is positive when the periape of the asteroid is closer to the observer than the center of mass. Subject to these

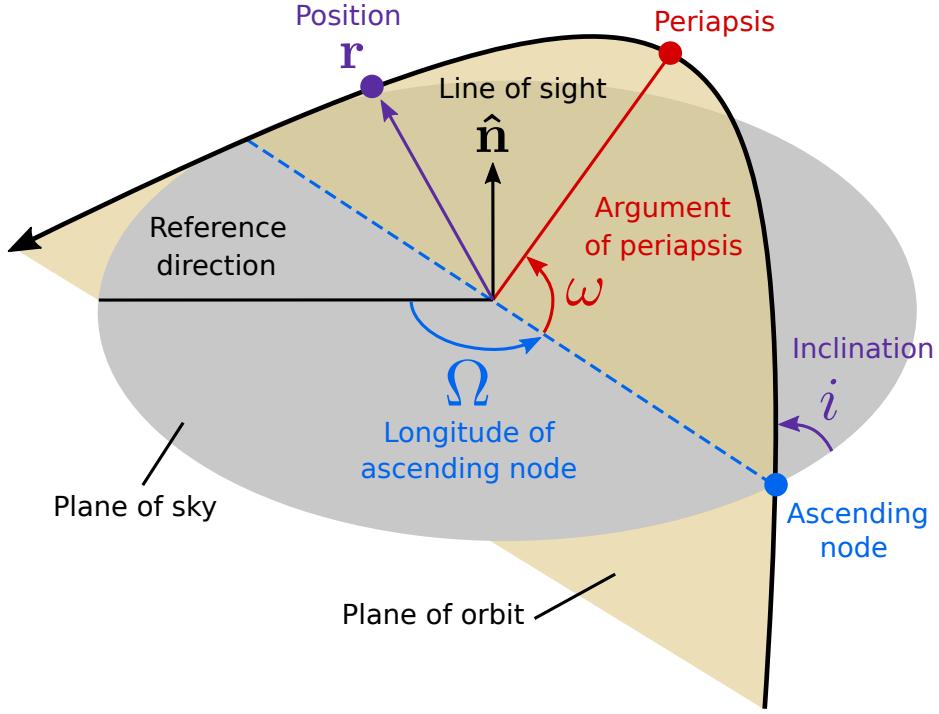


Figure 3.1: Geometry of a hyperbolic orbit, showing the three angles necessary to specify its orientation: the inclination, i , the argument of the periapsis, ω , and the longitude of the ascending node, Ω . Also shown are the position vector \mathbf{r} of the orbiting body (here, the perturbing ISO) and the line of sight direction $\hat{\mathbf{n}}$. The solid red line is the line of apsides, and the dashed blue line is the line of nodes.

conventions, the unit vector in the direction of the line of sight is given by

$$\hat{\mathbf{n}} = \sin i \sin \omega \hat{\mathbf{x}} + \sin i \cos \omega \hat{\mathbf{y}} + \cos i \hat{\mathbf{z}}. \quad (3.15)$$

The TOA perturbation then becomes

$$\Delta\tau = \frac{mb}{Mc} \sin i \left(\frac{e - \cosh H}{\sqrt{e^2 - 1}} \sin \omega + \sinh H \cos \omega \right). \quad (3.16)$$

Together with the relation between hyperbolic anomaly, H , and time, t , given by equation (3.5), this completely specifies the form of the timing perturbations produced by a close encounter.

Dror et al. (2019) also arrived at an expression, analogous to equation (3.16),

giving the expected timing signal resulting from a close encounter between a pulsar and a pointlike perturber. Our result differs in two key respects. First, while the Dror et al. result relies on the assumption that the orbit is highly unbound ($e \gg 1$), equation (3.16) has no such limitation, and is exact for all unbound orbits. Second, whereas Dror et al. expressed their result as a fractional change in frequency, our result is written in terms of TOA delays that are directly measurable. To compare our expression with the Dror et al. result, differentiate equation (3.16) with respect to t and expand to first order in the quantity

$$\frac{GM}{bv^2} = \frac{1}{\sqrt{e^2 - 1}}, \quad (3.17)$$

which is small when $e \gg 1$. This gives

$$\begin{aligned} \frac{d\Delta\tau}{dt} = \frac{\Delta\nu}{\nu} \approx \frac{Gm}{bvc} & \frac{(\cos\omega - \sinh H \sin\omega)}{\cosh H} \\ & + \frac{mv}{Mc} \sin i \cos\omega, \end{aligned} \quad (3.18)$$

which is equivalent (up to an additive constant that is degenerate with the unperturbed velocity of the pulsar) to equation (8) in Dror et al. (2019).

3.3.1 Timing Signature

Over the course of an encounter between a pulsar and a perturbing object, the slope of the TOA delay changes from its initial value to a different, stable final value. This is because momentum has been transferred from the perturber to the pulsar, and so the projection of the pulsar's velocity onto the line of sight has changed. The long-term change in the pulsar's velocity is given by

$$\Delta\mathbf{v}_p = \frac{2mv_0}{Me} \hat{\mathbf{x}}, \quad (3.19)$$

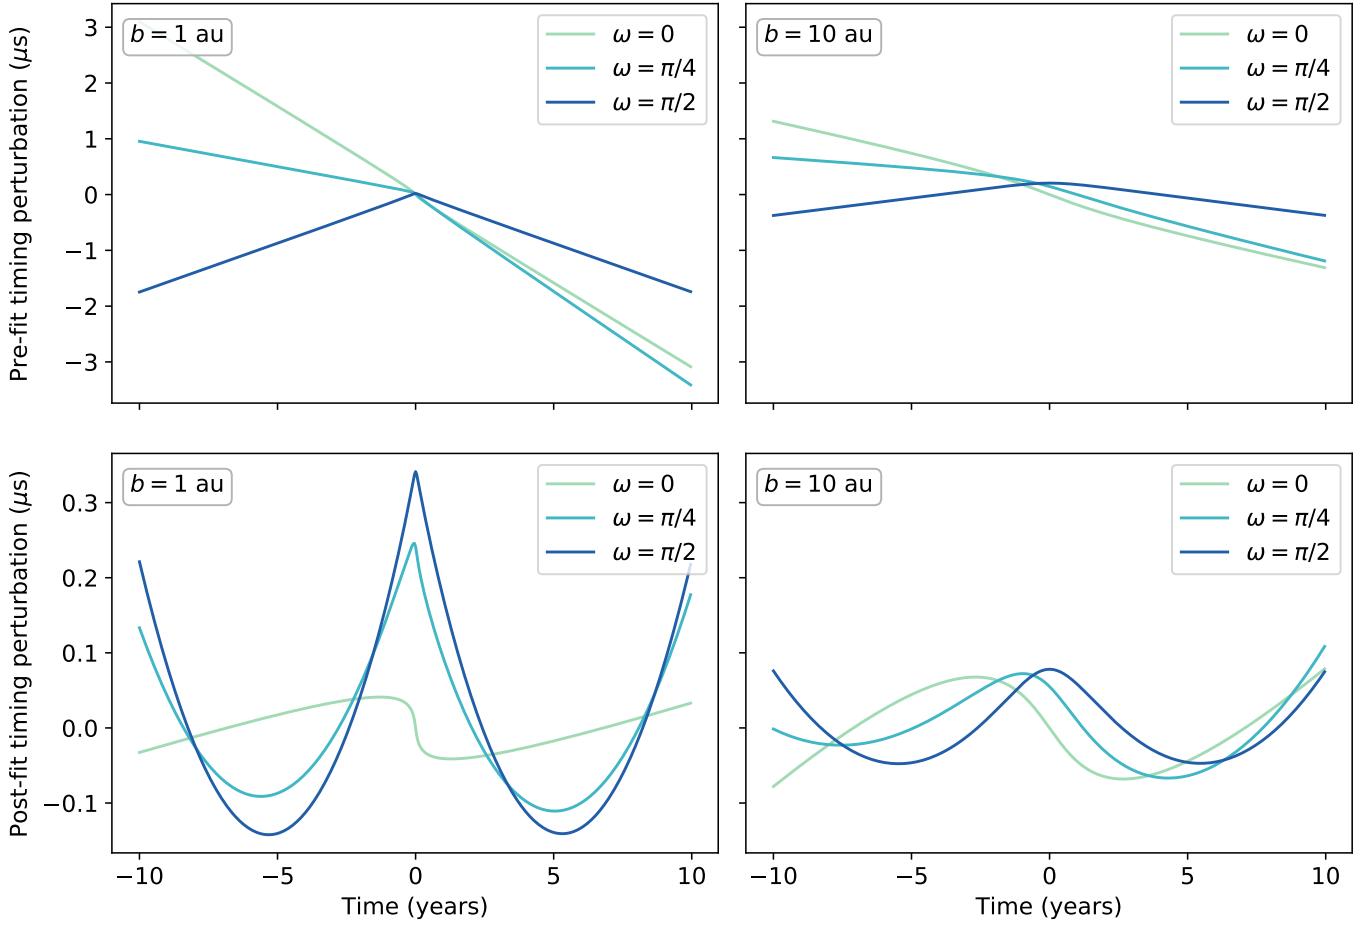


Figure 3.2: Timing perturbations caused by gravitational encounters with several different geometries. The upper two panels show the signal shape before fitting, while the lower two show the same signals after the model with the best-fit period and period derivative has been subtracted. All are shown for a putative 20-year data set, and correspond to an eccentricity $e = 2$, a pulsar mass $M = 1.4 M_{\odot}$, and a projected perturber mass $m \sin i = 10^{-10} M_{\odot}$. The left panels correspond to an impact parameter $b = 1$ au, while the right panels correspond to $b = 10$ au. The angle ω is the argument of the periapsis, with $\omega = \pi/2$ corresponding to a case in which the periapsis occurs along the line of sight between the Earth and the pulsar.

which corresponds to a change in the slope of the delay of

$$\Delta\dot{\tau} = -\frac{\Delta\mathbf{v}_p \cdot \hat{\mathbf{n}}}{c} = -\frac{2mv_0}{Mce} \sin i \sin \omega. \quad (3.20)$$

Notice that this can be either positive or negative, depending on the signs of i and ω . This permanent change in the velocity of the pulsar, and hence in the slope of the timing residuals, means that a long data set including an encounter will show a distinct change in apparent spin frequency.

Fig. 3.2 shows the shape of the timing perturbations for several particular geometries: the antisymmetric $\omega = 0$ and symmetric $\omega = \pi/2$ cases, as well as the intermediate $\omega = \pi/4$ case, for impact parameters of 1 au and 10 au. The signals are shown both before (upper panels) and after (lower panels) fitting and removing a quadratic model describing the period and period derivative of the pulsar (see Section 3.4 below). In the post-fit perturbation shown in the lower panels, the maximum amplitude of the signal is substantially smaller, and the change from one asymptotic slope to another is obscured. Nonetheless, the post-fit perturbation for a $10^{-10} M_\odot$ perturber can be hundreds of nanoseconds, comparable to the signal expected from the gravitational wave stochastic background. Comparing the 1 au (right) and 10 au (left) cases shows that the difference between favorable and unfavorable geometries is more significant for close encounters than it is for glancing encounters.

3.3.2 Similar Signals

The timing signal produced by a close encounter superficially resembles a glitch, an abrupt change in the period of a pulsar thought to be caused by an interaction between the crust and the superfluid interior. Glitches are mainly

observed in canonical pulsars, with only two cases seen so far in millisecond pulsars (Cognard & Backer, 2004; McKee et al., 2016). Glitch signals differ from those produced by close encounters in that they may be accompanied by changes in period derivative, sometimes display exponential recovery to the previous state, and have a preferred direction, spinning up the pulsar rather than slowing it down in all but a few anomalous cases (termed “anti-glitches”; Archibald et al. 2013). Exponential recovery has not been identified in either of the observed millisecond pulsar glitches, but this is not surprising since the frequency of recovery has been shown to decrease as a pulsar’s characteristic age increases (Lyne et al., 1995). Nevertheless, this means it is possible that one or both of the millisecond pulsar glitches could be misidentified gravitational scattering events. Assuming reasonable values for the hyperbolic excess velocity ($v = 100$ km/s) and interaction timescale ($t_i = 30$ days; see Section 3.4 below), the measured fractional changes in period of 9.5×10^{-12} (Cognard & Backer, 2004) and 2.5×10^{-12} (McKee et al., 2016) correspond to perturber masses of $2.0 \times 10^{-7} M_\odot$ ($0.067 M_\oplus$) and $5.2 \times 10^{-8} M_\odot$ ($0.017 M_\oplus$), respectively.

The same kind of ramp-like signal, involving a persistent change in period, could also be produced by a gravitational wave burst with memory (BWM; van Haasteren & Levin 2010; Cordes & Jenet 2012; Madison et al. 2017). An Earth-term BWM is distinguishable from an Earth-term gravitational scattering event in that the BWM produces quadrupolar correlations between pulsars, whereas the gravitational scattering event produces dipolar correlations. A pulsar-term BWM, however, is almost indistinguishable from a pulsar-term gravitational scattering event, except perhaps by its amplitude and time scale, and the detailed shape near the center of the event (when it is resolvable). Compared to BWMs, gravitational scattering events are much more likely to produce a

gradually varying (rather than cusp-like) signal, since the interaction timescale can easily be many years, or even longer than the observing span. On the other hand, the BWM signals produced by mergers of binary black holes with astrophysically realistic masses will always be cusp-like. Some BWM events may have electromagnetic counterparts, which would serve to distinguish them from gravitational scattering events. However, any searches for such counterparts will be complicated by the fact that BWM signals may not be detected until years after the merger events that produce them.

3.4 Detecting Pulsar Encounters

For the timing perturbation produced by an encounter between an ISO and a pulsar to be detectable, it must have an amplitude large enough to be distinguishable from noise. An important consideration here is that ISO encounter signals are always at least partially degenerate with terms involving the period and period derivative of the pulsar, and so it is appropriate to measure the amplitude only after those terms have been removed. The lower panels of Fig. 3.2 show the effect of removing the period and period derivative terms in a few sample cases, illustrating that this can reduce the signal amplitude significantly.

Searching for an ISO encounter signal in a time series of TOA residuals involves fitting a model of the form

$$\Delta\tau_i = ax_i + \varepsilon_i. \quad (3.21)$$

Here $\Delta\tau_i$ is the TOA residual at epoch i , a is the amplitude of the signal, x_i is its shape (normalized to unit amplitude), and ε_i is the noise. The least-squares

estimate of the amplitude parameter is

$$\hat{a} = \frac{x_i \Sigma_{ij}^{-1} \Delta \tau_j}{x_i \Sigma_{ij}^{-1} x_j}, \quad (3.22)$$

where Σ_{ij}^{-1} is the inverse covariance matrix of the noise, and there is an implied sum over each pair of repeated indices. The corresponding uncertainty is

$$\sigma_a = \left(x_i \Sigma_{ij}^{-1} x_j \right)^{-1/2}. \quad (3.23)$$

A signal is said to be detected if $z = \hat{a}/\sigma_a$ exceeds a given threshold, z_0 , which may be set to obtain a particular false positive probability, using the fact that z follows a standard normal distribution when the true amplitude, a , is zero. A reasonable default is $z_0 = 3$, but for searches over a large number of test shapes, larger values of z_0 may be necessary. It follows that for a signal to be detectable, its amplitude should satisfy $a \gtrsim z_0 \sigma_a$, or

$$\left(\Delta \hat{\tau}_i \Sigma_{ij}^{-1} \Delta \hat{\tau}_j \right)^{-1/2} \gtrsim z_0, \quad (3.24)$$

where $\Delta \hat{\tau}_i = a x_i$ is the model TOA residual at epoch i . When the timing noise is uncorrelated, which holds for radiometer and jitter noise, but not for scintillation or spin noise, we have $\Sigma_{ij}^{-1} = \sigma_\tau^{-2} \delta_{ij}$, where σ_τ is the timing precision, so equation (3.24) reduces to the simpler form

$$\Delta \hat{\tau}_{\text{rms}} \gtrsim \frac{z_0 \sigma_\tau}{\sqrt{N}}. \quad (3.25)$$

Here $\Delta \hat{\tau}_{\text{rms}} = (\Delta \hat{\tau}_i \Delta \hat{\tau}_i)^{1/2}$ is the RMS of the model TOA residuals, and N is the number of data points (epochs).

This RMS TOA residual is plotted as a function of hyperbolic excess velocity, v , in Fig. 3.3, and as a function of the position of the event within the data set in Fig. 3.4. Fig. 3.5 shows the RMS TOA residual in the two-dimensional

space of impact parameter, b , and hyperbolic excess velocity, v . An important consideration in determining the detectability of a given signal is the fact that the fit for the period and period derivative of the pulsar will remove some power from the signal. For instance, interactions which take place entirely outside the data set result in a permanent velocity change, but this is degenerate with the initial velocity of the pulsar.

As seen in Fig. 3.2, the shape of the signal, and thus the extent to which it is degenerate with the period and period derivative terms, depends on how the interaction time scale, $t_i \sim b/v$, compares to the data set length, T . In the limit where $t_i \ll T$, the signal produced by an encounter reduces to an instantaneous change in slope. This means that the signal is primarily distinguished by a sharp cusp, and the fit for the pulsar's period and period derivative does not influence its detectability, which depends only on the radial component of the the momentum exchanged during the encounter. In the opposing limit, where $t_i \gg T$, the signal is well approximated by its Taylor series expansion, and so the RMS residual falls off much faster as the event signature becomes degenerate with terms involving the period and period derivative of the pulsar.

For interstellar objects occurring as a Poisson process with constant number density, n , the probability of finding k objects within a volume V is

$$P(k) = \frac{(nV)^k e^{-nV}}{k!} \quad (3.26)$$

Assuming the same velocity, v , for all objects, the volume containing objects that will pass the pulsar at an impact parameter less than b within the span of the data set is $V_1 = \pi b^2 v T$. A PTA with N_{psr} pulsars then probes a total volume

$$V = \pi N_{\text{psr}} V_1 = \pi N_{\text{psr}} b^2 v T. \quad (3.27)$$

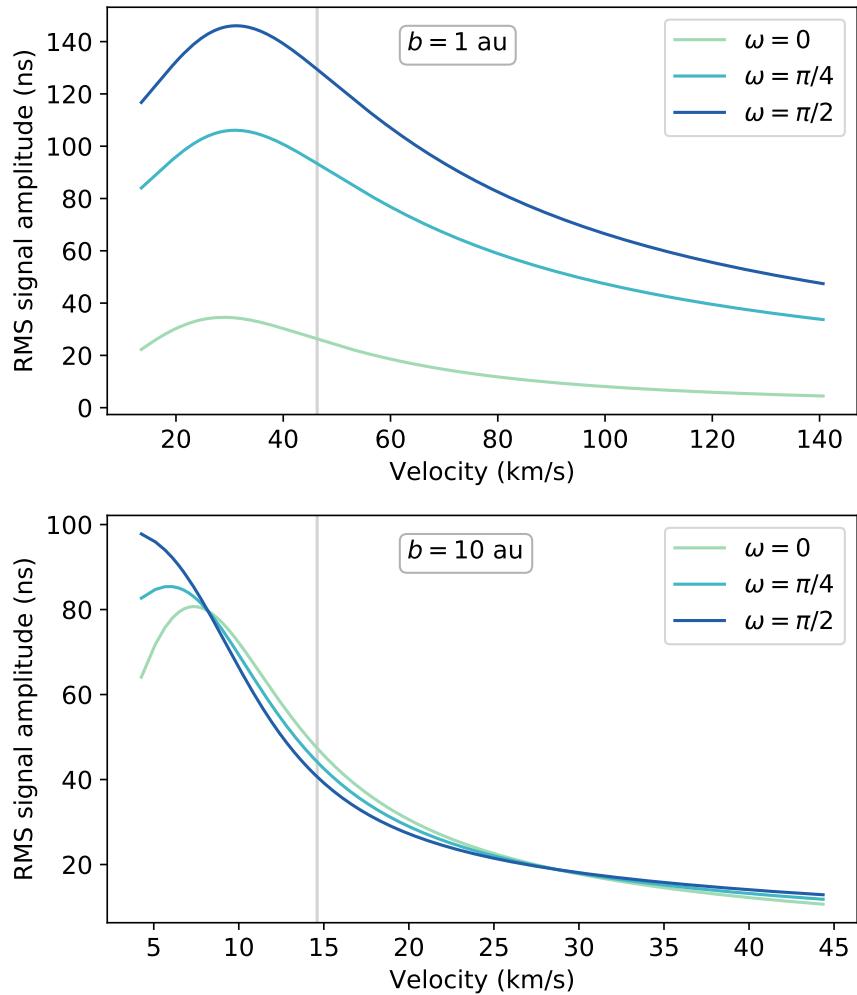


Figure 3.3: The RMS signal amplitude, after removing the best-fit period and period derivative, in a fiducial 20-year pulsar timing data set, as a function of hyperbolic excess velocity v . The top panel corresponds to encounters with an impact parameter $b = 1 \text{ au}$ and the bottom panel corresponds to $b = 10 \text{ au}$. The three lines in each panel correspond to different values of ω . The amplitudes are for $m \sin i = 10^{-10} M_{\odot}$ and scale linearly with the mass, m , of the perturber. As in Fig. 3.2, a fiducial pulsar mass $M = 1.4 M_{\odot}$ has been used. The light gray vertical line indicates the velocity also used in Fig. 3.4. In the $b = 1 \text{ au}$ case, the interaction timescale is much shorter, so the bulk of the signal comes from the difference in the pulsar's line-of-sight velocity before and after the encounter. This means that for $\omega = 0$, where the net transfer of momentum is perpendicular to the line of sight, the signal amplitude is much smaller.

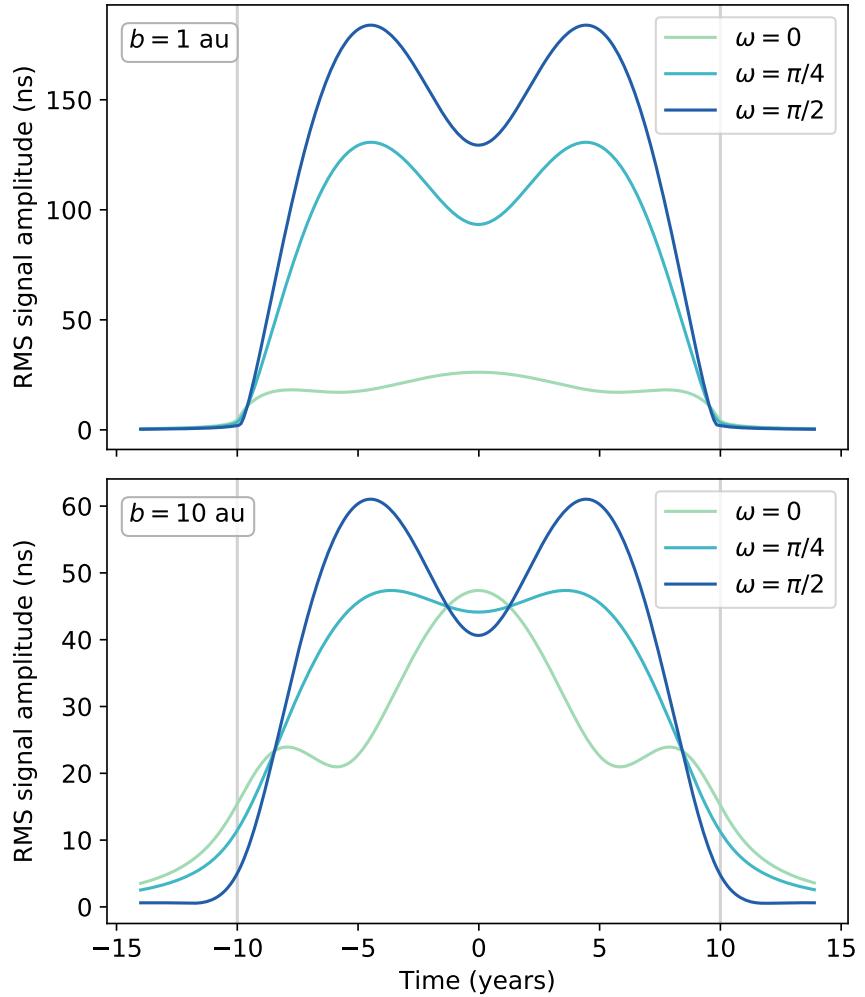


Figure 3.4: The RMS signal amplitude in a fiducial 20-year pulsar timing data set as a function of time of periapsis, measured from the center of the data set. The top panel corresponds to encounters with an impact parameter $b = 1$ au and the bottom panel corresponds to $b = 10$ au. The three curves in each panel correspond to different values of the argument of periapsis ω , and the light gray lines indicate the boundaries of the data set. As in Fig. 3.2, the curves are for $M = 1.4 M_{\odot}$, $m \sin i = 10^{-10} M_{\odot}$, and $e = 2$. The corresponding velocity in each case is indicated by the light gray vertical line in the appropriate panel of Fig. 3.3.

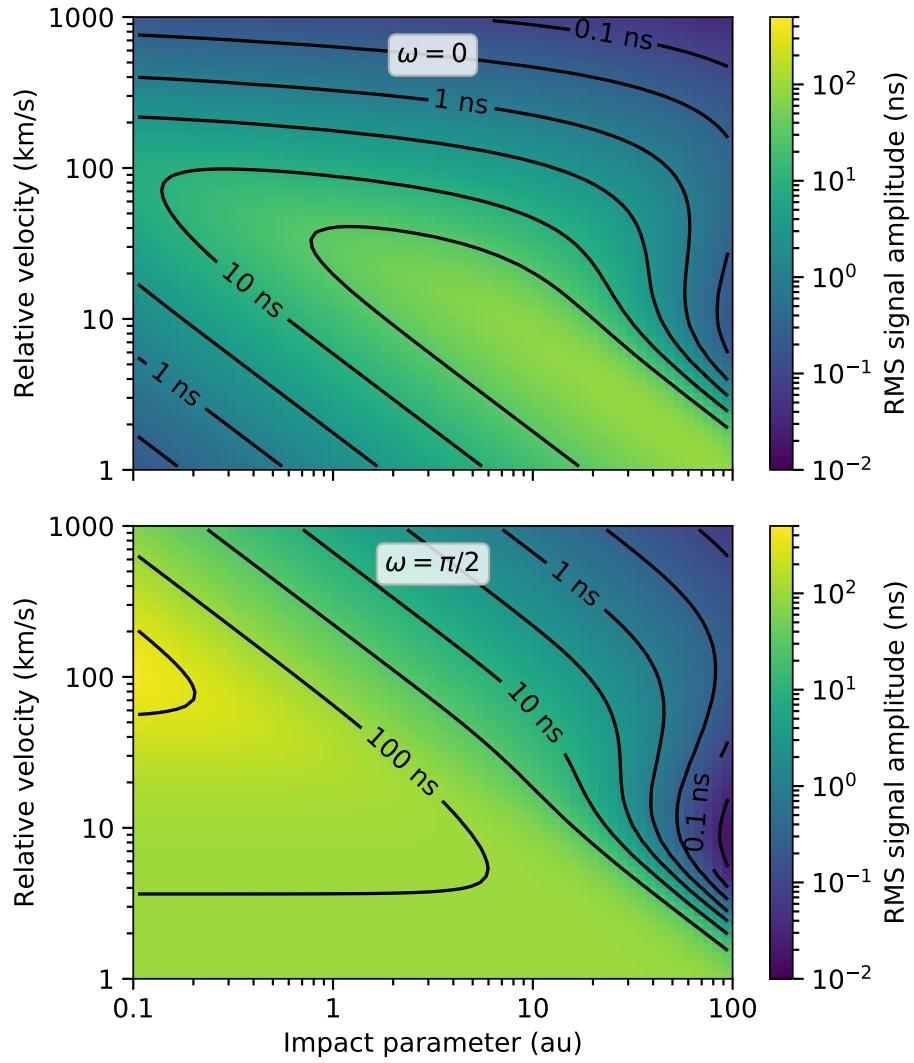


Figure 3.5: The RMS signal amplitude as a function of impact parameter and relative velocity, for a periapsis time of zero (corresponding to periapsis occurring at the center of the data set). As in previous figures, the amplitudes correspond to an object of mass $m = 10^{-10} M_{\odot}$ with zero inclination, and a fiducial 20-year data set. The upper panel corresponds to the case $\omega = 0$, where the line of apsides lies in the plane of the sky, whereas the lower panel corresponds to the case $\omega = \pi$, where the line of apsides lies along the line of sight.

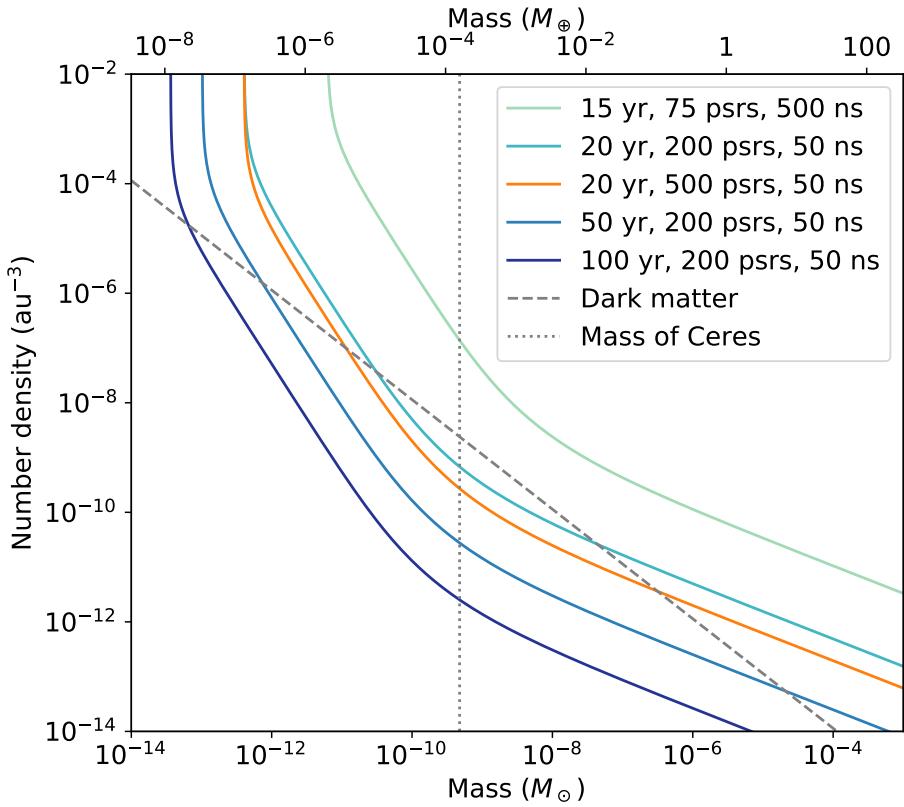


Figure 3.6: Detectability of interstellar populations of massive objects with a given number density and per-object mass in various PTA scenarios. For a population of objects to be detectable by a given PTA configuration, it should lie above the corresponding curve. The first case, with an observing baseline of 15 years and 75 pulsars and an RMS timing error of 500 ns, reflects capabilities similar to that of the current NANOGrav array. The remaining scenarios are increasingly optimistic future possibilities. In all cases, pulsars are assumed to be observed with a two-week cadence. The minimum detectable mass for a given number density is proportional to the timing precision. The dashed gray line indicates populations with a total mass density equal to the local density of dark matter, approximately $0.010 M_\odot \text{ pc}^{-3}$.

It follows that the probability that a total of k objects pass by the pulsars with impact parameter less than b within a time T is

$$P(k) = \frac{(\pi N_{\text{psr}} n b^2 v T)^k e^{-\pi N_{\text{psr}} n b^2 v T}}{k!}. \quad (3.28)$$

The probability that the smallest impact parameter encountered in the data set is b_0 is then

$$P(b_0) = 2\pi N_{\text{psr}} n b_0 v T e^{-\pi N_{\text{psr}} n b_0^2 v T} db_0. \quad (3.29)$$

The minimum impact parameter observed in a PTA data set will likely approximate the mean of this distribution, which is

$$\langle b_0 \rangle = \frac{1}{2\sqrt{N_{\text{psr}} n v T}}. \quad (3.30)$$

Fig. 3.6 plots the number density, n , against the minimum detectable mass, assuming that the signal is dominated by the event with the smallest impact parameter, and that the latter is given by equation (3.30).

3.5 Detecting Earth Encounters

If an ISO passes through the solar system, it will alter the motion of the solar system barycenter in much the same way it would alter the motion of a pulsar if it passed near one. Sufficiently massive objects passing near the inner solar system, even if not directly observable, would be detectable via their influence on the orbits of the planets. For objects which are smaller or pass further from the Sun, however, this influence becomes harder to discern. One of the most sensitive ways to find these smaller or more distant objects is by using pulsar timing to measure the reflex motion of the solar system barycenter. Since the passage of such an object through the solar system perturbs the TOAs

from every pulsar on the sky in a correlated way, it is possible to distinguish these flyby events from processes affecting only a single pulsar. With multiple pulsars, it is also possible to resolve signals of lower amplitude than would be detectable with a single pulsar, and to reconstruct the full geometry of the perturbing body's orbit.

Ignoring any perturbations to the ISO's orbit caused by planets, the expected TOA signal is still that described by equation (3.16), but the geometric interpretation is slightly different. To compare the signals observed in different pulsars, it is useful to write the TOA perturbation in terms of the ecliptic latitude, β , and longitude, λ , of the pulsar. This can be accomplished by replacing $\hat{\mathbf{n}}$ in equation (3.14) with the unit vector from the solar system barycenter to the pulsar, expressed in ecliptic coordinates λ and β :

$$\hat{\mathbf{n}} = \sin \beta \cos \lambda \hat{\mathbf{u}} + \sin \beta \sin \lambda \hat{\mathbf{v}} + \cos \beta \hat{\mathbf{w}}. \quad (3.31)$$

Here $\hat{\mathbf{u}}$ is the unit vector in the direction of the vernal equinox, $\hat{\mathbf{w}}$ is the unit vector in the direction of the north ecliptic pole, and $\hat{\mathbf{v}}$ satisfies $\hat{\mathbf{u}} \times \hat{\mathbf{v}} = \hat{\mathbf{w}}$. In terms of $\hat{\mathbf{u}}$, $\hat{\mathbf{v}}$, and $\hat{\mathbf{w}}$, the unit vectors $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ from Section 3.3 can be written

$$\begin{aligned} \hat{\mathbf{x}} = & (\cos \omega \cos \Omega - \sin \omega \sin \Omega \cos i) \hat{\mathbf{u}} \\ & + (\cos \omega \sin \Omega + \sin \omega \cos \Omega \cos i) \hat{\mathbf{v}} \end{aligned} \quad (3.32)$$

$$\begin{aligned} \hat{\mathbf{y}} = & (-\sin \omega \cos \Omega - \cos \omega \sin \Omega \cos i) \hat{\mathbf{u}} \\ & + (-\sin \omega \sin \Omega + \cos \omega \cos \Omega \cos i) \hat{\mathbf{v}} \\ & + \cos \omega \sin i \hat{\mathbf{w}}. \end{aligned} \quad (3.33)$$

The TOA perturbation for a pulsar in a direction $\hat{\mathbf{n}}$ is then

$$\Delta \tau = \frac{mb}{Mc} \left(\frac{e - \cosh H}{\sqrt{e^2 - 1}} \hat{\mathbf{x}} + \sinh H \hat{\mathbf{y}} \right) \cdot \hat{\mathbf{n}}. \quad (3.34)$$

Notably, the signal has a dipolar dependence on the line-of-sight direction, \hat{n} . This has a few important consequences, which are as follows: First, the signals from different pulsars differ only by a scale factor which depends in a predictable manner on the value of \hat{n} for each pulsar. This makes it possible to detect weaker signals than could be detected in a single pulsar by adding the signals from several pulsars coherently. Second, if the same event is detected in several pulsars, it should be possible to measure the orientation of the orbit completely, and thereby determine the direction the perturbing object came from. Finally, because Earth-term gravitational wave signals have a quadrupolar dependence on the line-of-sight direction, this makes Earth-term gravitational scattering signals distinguishable from Earth-term gravitational wave bursts with memory (BWMs).

Recently, errors in the solar system ephemeris have been recognized as a significant obstacle in current PTA searches for gravitational waves (Arzoumanian et al., 2018b). As a result, PTAs have developed techniques for marginalizing over uncertainties in the masses and trajectories of solar system bodies, and it has been recognized that PTA data sets can be used to make precision measurements of the position of the earth relative to the solar system barycenter. Such measurements were carried out by Caballero et al. (2018), who used the first IPTA data release (Verbiest et al., 2016) to constrain the masses and orbits of all the major planets and the largest five main-belt asteroids. Caballero et al. also conducted a search for unmodeled objects in closed orbits around the solar system barycenter, using methods described in Guo et al. (2018). They found no such objects larger than a few times $10^{-10} M_\odot$ interior to the orbit of Jupiter, with a somewhat weaker limit at larger semi-major axis values. Since their search was restricted to closed orbits, however, it did not produce direct

constraints on encounters of the sort considered here.

3.6 Frequency of Encounters

The total number of asteroid-mass ISOs is currently not well understood. A weak upper bound on their number density can be derived from the assertion that such objects cannot have a total mass density exceeding the local mass density of dark matter, which is about $0.01 M_\odot \text{ pc}^{-3}$ — recent estimates range from $0.008 M_\odot \text{ pc}^{-3}$ (Eilers et al., 2019) to $0.016 M_\odot \text{ pc}^{-3}$ (Buch et al., 2019). This shows that the number density of ISOs in the solar neighborhood with masses greater than $10^{-10} M_\odot$ (which is approximately 20% of the mass of Ceres) cannot exceed approximately $1.1 \times 10^{-8} \text{ au}^{-3}$.

The current best constraints on the abundance of these objects come from two directions. First, microlensing surveys, particularly the Subaru/HSC Andromeda survey (Niikura et al., 2019), have placed upper bounds on the frequency of large objects. Second, the detections of ‘Oumuamua by the Pan-STARRS1 survey (Meech et al., 2017) and C/2019 Q4 (Guzik et al., 2019) show that objects around 1 km in diameter are relatively common. These detections, along with the nondetection of other similar objects by solar system surveys, can be used to constrain the number density of asteroid- and comet-like bodies in interstellar space (Engelhardt et al., 2017; Do et al., 2018).

Niikura et al. (2019) report on a microlensing survey of 10^8 stars in the Andromeda galaxy, conducted using the Hyper Suprime Cam (HSC) on the 8.2-meter Subaru telescope. The survey, which focused on constraining the abundance of PBHs, produced a single candidate microlensing event. They

place an upper bound of approximately 2×10^{-3} on the fraction of dark matter in the galactic halo which consists of $10^{-10} M_\odot$ objects. This corresponds to a number density upper bound of about $2.2 \times 10^{-11} \text{ au}^{-3}$.

On the solar system side, Engelhardt et al. (2017), writing before the discovery of ‘Oumuamua, find the largest number densities of asteroid- and comet-like compatible with non-detections by the Pan-STARRS1 (Chambers et al., 2016b) and Catalina Sky Surveys (Christensen et al., 2012). For objects more than 1 km in diameter, their limits are $1.4 \times 10^{-4} \text{ au}^{-3}$ for comet-like ISOs, and $2.4 \times 10^{-2} \text{ au}^{-3}$ for asteroid-like ISOs. Incorporating ‘Oumuamua, Do et al. (2018) find a number density of 0.2 au^{-3} for objects at least 100 m in diameter. This can be extrapolated to objects of larger radii by assuming a power-law distribution of masses, in which the number density of objects with radii between r and $r + dr$ is given by

$$dn \propto r^{-(\alpha+1)} dr. \quad (3.35)$$

Dohnanyi (1969) made a theoretical argument for taking $\alpha = 2.5$ for material in collisional equilibrium in protoplanetary disks, and this appears to be broadly representative of small-body populations in the solar system (Jedicke et al., 2002). Attempts have also been made to determine the exponent α for interstellar objects empirically, based on direct measurements of interstellar dust by the *Ulysses* and *Galileo* spacecraft (Landgraf et al., 2000), and optical and radar detections of meteors identified as interstellar (Taylor et al., 1996; Baggaley, 2000; Meisel et al., 2002; Weryk & Brown, 2004; Hajduková, 2008, 2011; Siraj & Loeb, 2019a). On the basis of the Ulysses and Galileo data alone, Landgraf et al. (2000) estimated $\alpha = 3.3$. Most recently, Siraj & Loeb (2019b) arrived at $\alpha = 3.41 \pm 0.17$ in an analysis incorporating ‘Oumuamua and the bolide me-

teor CNEOS 2014-01-08 (Siraj & Loeb, 2019a), which is unique among meteors identified as interstellar because of its relatively large size, estimated at 0.45 m.

Caution is necessary in applying these results to the present context, both because the power-law distribution is being extrapolated well beyond the range of sizes for which it was derived, and because the velocity measurements required to establish a meteor as interstellar in origin are challenging. Hajduk (2001) gives a critique of the velocity measurement techniques used to establish the interstellar origin of radar meteors in the AMOR dataset (Taylor et al., 1996; Baggaley, 2000). Additionally, the velocity measurement used to establish the interstellar character of CNEOS 2014-01-08 relies on United States government sensors whose performance characteristics are not made public (Zuluaga, 2019; Devillepoix et al., 2019). Nevertheless, there is currently no better way to estimate the number density of asteroid-like ISOs.

Using the Dohnanyi (1969) scaling, the Do et al. (2018) result becomes $6.3 \times 10^{-4} \text{ au}^{-3}$ for objects larger than 1 km and $7.2 \times 10^{-11} \text{ au}^{-3}$ for objects more than 600 km in diameter, which for a density of 2.0 g/cm^3 corresponds to a mass of about $10^{-10} M_\odot$. Using the empirical scaling law from Siraj & Loeb (2019b) gives considerably smaller estimates of $7.8 \times 10^{-5} \text{ au}^{-3}$ for objects larger than 1 km and $2.6 \times 10^{-14} \text{ au}^{-3}$ for objects larger than 600 km.

All of these results are estimates for the density of asteroid-mass objects in the Galaxy as a whole; it is entirely possible that local overdensities of these objects may exist, for instance in globular clusters. The density of stars near the center of a globular cluster can be hundreds to thousands of times greater than it is in the solar neighborhood. If planetesimal-mass objects were similarly over-represented in globular clusters, it would be significantly more likely for

one to interact closely with a pulsar than the previous estimates would suggest. Additionally, it is possible that populations of distant, gravitationally bound asteroids exist around at least some pulsars. Stellar flybys of the pulsar could put some of these objects on marginally unbound orbits that pass very close to the central pulsar. Such an event would be comparatively easy to detect.

3.7 Summary and Conclusions

This paper describes the shape of the pulsar timing signal that would be produced by an interstellar asteroid or other massive object in the course of flying by a pulsar or the solar system, and evaluates the likelihood of detecting such a signal in current or future PTA data sets. We find that the signal produced by a scattering event of this form would be ramp-like, since the interaction would cause a persistent change to the velocity of the pulsar relative to the solar system barycenter that would create a corresponding persistent change in the slope of the timing residuals. This is similar to the shape of the signal produced by a pulsar glitch or a gravitational wave burst with memory. The persistent nature of the signal allows SNR to build up over the course of an observing span which is long compared to the duration of the interaction.

Interactions in which a $10^{-10} M_{\odot}$ ISO passes within a few au of a pulsar or the solar system barycenter should be strong enough to be detectable with current PTAs. However, estimates of the interstellar number density of objects of this mass suggest that such interactions should be rare. No such encounters have yet been detected unambiguously. It is possible, although unlikely, that one or both of the known millisecond pulsar glitches may be misidentified

encounters with ISOs.

Current PTAs are unlikely to detect interactions between pulsars and ISOs unless the number density of ISOs is enhanced in the immediate vicinity of one or more of the pulsars, as may be the case in globular clusters. On the other hand, future PTAs with 200 or more well-timed pulsars and observing baselines of at least 20 years may be able to place astrophysically interesting constraints on compact objects with masses between 10^{-12} and $10^{-10} M_{\odot}$ as constituents of dark matter. The sensitivity of a PTA to ISO encounters is primarily determined by the observing span, a fact which demonstrates the benefits of timing pulsars continuously for many decades.

Because ISO encounters produce signals very similar to those expected from gravitational wave bursts with memory, candidate BWM events detected by PTAs should be carefully scrutinized to be sure they are not mistakenly identified ISO encounter signals. For Earth-term events which are detected in multiple pulsars, the pattern of spatial correlations between pulsars can distinguish between the two types of signal, but pulsar-term events will be much more challenging to interpret. The absence of a sharp cusp in the timing signal can rule out a BWM explanation in favor an ISO encounter, but the reverse is not true, as close encounters between pulsars and ISOs can produce arbitrarily sharp cusps in the timing residuals. Unless an electromagnetic counterpart can be found in archival observations, it may be impossible to distinguish a pulsar-term BWM from an ISO encounter.

3.8 Acknowledgements

The authors are members of the NANOGrav Physics Frontiers Center, which receives support from the National Science Foundation (NSF) under award number 1430284.

CHAPTER 4

PULSAR TIMING SIGNATURES OF CIRCUMBINARY ASTEROID BELTS

The gravitational pull of a large number of asteroids perturbs a pulsar's motion to a degree that is detectable through precision timing of millisecond pulsars. The result is a low-frequency, correlated noise process, similar in form to the red timing noise known to affect canonical pulsars, or to the signal expected from a stochastic gravitational-wave background. Motivated by the observed fact that many millisecond pulsars are in binary systems, we describe the ways in which the presence of a binary companion to the pulsar would affect the signal produced by an asteroid belt. The primary effect of the companion is to destabilize the shortest-period orbits, cutting off the high-frequency component of the signal from the asteroid belt. We also discuss the implications of asteroid belts for gravitational-wave search efforts. Compared to the signal from a stochastic gravitational-wave background, asteroid belt noise has a similar frequency and amplitude, and is similarly independent of radio frequency, but is not correlated between different pulsars, which should allow the two kinds of signal to be distinguished.

4.1 Introduction

The high precision of pulsar timing makes it possible to detect orbital reflex motion created by small bodies orbiting pulsars, especially millisecond pulsars (MSPs). A striking demonstration of this is the discovery of pulsar planets. The smallest known pulsar planet, around PSR B1257+12, is less than twice as

Published: Jennings, R. J., Cordes, J. M., & Chatterjee, S. 2020, ApJ, 904, 191.

massive as the Moon (Wolszczan, 1994; Wolszczan et al., 2000), and a recent search for planets around MSPs (Behrens et al., 2020) was able to rule out planets more massive than the Moon at periods as long as 100 days. Since such small bodies can be detected individually, it is worth considering the effects of even smaller bodies, such as asteroids, which can combine to produce a detectable signal.

The orbital reflex motion created by a body in a circular orbit around a pulsar produces a sinusoidal variation in pulse times of arrival (TOAs). For a body of mass m in a circular orbit of radius a around a pulsar of mass M , the amplitude, $\Delta\tau$, of this signal is

$$\begin{aligned}\Delta\tau &= \frac{ma}{Mc} \sin i \\ &= 50 \text{ ns} \left(\frac{m}{10^{-10}M} \right) \left(\frac{a}{1 \text{ au}} \right) \sin i.\end{aligned}\tag{4.1}$$

Here c is the speed of light and i is the inclination of the orbital plane relative to the plane of the sky. If the ratio of the mass of the asteroid to the mass of the pulsar is 10^{-10} (around a typical pulsar, this corresponds to a mass similar to that of the large solar-system asteroid 4 Vesta) and it orbits at 1 au, the signal amplitude can be as large as 50 ns. This is large enough that a single asteroid could be detectable if it orbited one of the best-timed MSPs. The incoherent superposition of many such signals produced by an asteroid belt can be significantly stronger – a belt with a total mass of only $5 \times 10^{-4} M_{\oplus}$ can produce a signal in excess of 250 ns. For this reason, Shannon et al. (2013) suggested that an asteroid-belt signal could explain some of the low-frequency noise observed in TOAs from the bright MSP PSR B1937+21. Indeed, asteroid belts could be an important source of timing noise in MSPs in general, including in the large number of MSPs occurring in binary systems.

Asteroid belts around pulsars may be formed by supernova fallback material, or, in the case of MSPs, by material left over from the recycling process that spun the pulsar up. The latter is often thought to be the origin of the planets around PSR B1257+12. In addition to these, there are a number of other lines of evidence supporting the existence of debris disks or asteroid belts around at least some pulsars. The magnetars 4U 0142+61 (Wang et al., 2006) and 1E 2259+586 (Kaplan et al., 2009) exhibit infrared emission suggestive of debris disks. Furthermore, theoretical work related to the B1257+12 system (Miller & Hamilton, 2001; Bryden et al., 2006) indicates that asteroids larger than 1 km in radius can survive for 1 Gyr or more around an MSP.

Since many MSPs are found in binary systems with white dwarf companions, we are lead to consider additional effects that may arise in the binary context. Small objects such as asteroids orbiting in a binary star system have one of two types of stable orbits: satellite (S-type) orbits, circling one of the components; or planetary (P-type) orbits, encircling both components (Murray & Dermott, 1999). To maintain stability, S-type orbits must have much shorter periods (and smaller semi-major axes) than P-type orbits: there is an outermost stable S-type orbit and an innermost stable P-type orbit, which has a longer period. Combined with the scaling of TOA variations with orbital radius given by equation (6.11), this means that objects on P-type orbits will generally contribute more power to low-frequency TOA variations. In the remainder of this paper we will concern ourselves primarily with P-type asteroid belts.

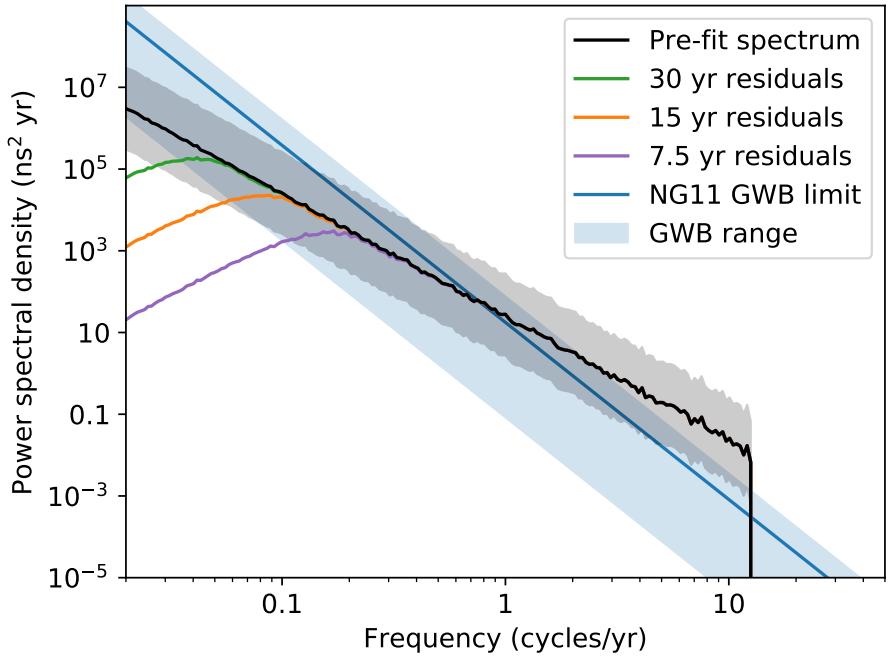


Figure 4.1: An example of the power spectrum of TOA perturbations created by a P-type asteroid belt around a pulsar and its companion. The gray shaded region corresponds to an order of magnitude change, in either direction, in the asteroid belt mass. The residual spectrum is shown for 7.5, 15, and 30 year dataset lengths, demonstrating the effects of model subtraction (see § 4.2.1). The spectrum expected from a GW background is shown for comparison, with the solid line corresponding to the upper limit on the strain amplitude, $A_{\text{GWB}} = 1.45 \times 10^{-15}$, set by Arzoumanian et al. (2018b), and the shaded region to A_{GWB} between 1×10^{-16} and 3×10^{-15} .

4.2 Spectral Characteristics

A single asteroid of mass m orbiting a pulsar of mass M on a circular orbit of semimajor axis a inclined at an angle i from the plane of the sky produces a sinusoidal timing delay with amplitude given by equation (4.1). The effects of additional asteroids on the timing residuals are additive as long as interactions between the asteroids can be neglected. When a large number of asteroids are present, they combine to produce a signal with an apparently continuous spectrum, as shown in Figure 4.1.

The semimajor axis, a , of a particular asteroid is related to its orbital frequency, f , by

$$a = \left(\frac{GM}{4\pi^2} \right)^{1/3} f^{-2/3}. \quad (4.2)$$

This relationship is exact in the two-body problem with M the sum of the masses. If one body is much less massive, as is the case for an asteroid orbiting a pulsar, it holds approximately with M the mass of the larger body. In a binary system, it holds approximately for both S-type and P-type orbits with appropriate choice of M . For S-type orbits, the appropriate M is the mass of the primary, and for P-type orbits, it is the total mass of the binary. The approximation becomes less accurate as the asteroid's orbit approaches the boundary of stability, but numerical evidence (Holman & Wiegert, 1999; Nagler, 2005) suggests that orbits typically become unstable before deviations from equation (4.2) are significant.

The spectrum of the TOA variations introduced by a large number of asteroids is related to the distribution of asteroid masses and orbital periods. Suppose that the number of asteroids with mass between m and $m + dm$ and orbital frequency between f and $f + df$ is

$$dN = n(m, f) dm df. \quad (4.3)$$

The number for which the amplitude is between τ and $\tau + d\tau$ and the orbital frequency is between f and $f + df$ is therefore

$$dN = n_*(\tau, f) d\tau df, \quad (4.4)$$

where

$$n_*(\tau, f) = K f^{2/3} n\left(K f^{2/3} \tau, f\right), \quad (4.5)$$

and

$$K = \left(\frac{4\pi^2 M^2}{G} \right)^{1/3} \frac{c}{\sin i}. \quad (4.6)$$

The power spectral density of the TOA signal, $S(f)$, is calculated by summing the contributions from each asteroid. A single asteroid contributes power $\tau^2/4$ to the power spectrum at frequency f , so the power spectral density (defined to contain only positive frequencies) is given by

$$S(f) = \frac{1}{4} \int_0^\infty \tau^2 n_*(\tau, f) d\tau. \quad (4.7)$$

If the distribution of asteroids has a power-law form in mass and frequency, with upper (lower) cutoffs m_+ (m_-) in mass and f_+ (f_-) in frequency, $n(m, f)$ is given by

$$n(m, f) = \frac{N\alpha\beta m^{\alpha-1} f^{\beta-1}}{(m_+^\alpha - m_-^\alpha)(f_+^\beta - f_-^\beta)}, \quad (4.8)$$

where $m_- < m < m_+$ and $f_- < f < f_+$, and 0 elsewhere. The average power spectral density at frequency f is therefore

$$S(f) = \left(\frac{G}{4\pi^2 M^2} \right)^{2/3} \frac{M_{\text{belt}} m_{\text{eff}}}{4c^2} \frac{\beta f^{\beta-7/3}}{(f_+^\beta - f_-^\beta)} \sin^2 i, \quad (4.9)$$

for $f_- < f < f_+$, and 0 otherwise. Here M_{belt} is the total mass of the belt, and m_{eff} is effective mass of a single asteroid, given by

$$m_{\text{eff}}(\alpha) = \frac{\langle m^2 \rangle}{\langle m \rangle} = \frac{(\alpha+1)(m_+^{\alpha+2} - m_-^{\alpha+2})}{(\alpha+2)(m_+^{\alpha+1} - m_-^{\alpha+1})}, \quad (4.10)$$

which always lies between the minimum and maximum masses, m_- and m_+ . When $m_+ \gg m_-$, m_{eff} is dominated by m_+ for $\alpha > -1$ and by m_- for $\alpha < -2$. For $-2 < \alpha < -1$, it behaves as a weighted geometric mean of the two.

An example spectrum is shown in Figure 4.1. It is an ensemble average calculated from 10^4 realizations of an asteroid belt containing 10^4 asteroids, distributed in mass and semimajor axis according to equation (4.8). The overall shape of the spectrum is described by equation (4.9). The high-frequency cutoff corresponds to the period of the innermost stable P-type orbit in the

PSR J1614–2230 system, determined using the formula of Holman & Wiegert (1999) (equation 4.31). The total mass of the belt is set to $2 \times 10^{-9} M_\odot$, approximately that of the solar system’s asteroid belt. The spectral indices in mass and frequency are $\alpha = -5/6$ and $\beta = -2/3$, respectively, and the total number of asteroids, N , is 10^4 . The upper and lower cutoffs in frequency are $f_- = 0.0173 \text{ yr}^{-1}$ and $f_+ = 12.5 \text{ yr}^{-1}$, and the cutoffs in mass are $m_- = 8.4 \times 10^{-15} M_\odot$ and $m_+ = 1.0 \times 10^{-10} M_\odot$. These are the same as the parameter values used for PSR J1614–2230 in Figure 4.3. More information on why these parameter values were selected can be found in Section 4.3.

Notably, the shape of the spectrum described by equation (4.9) depends only on the distribution of asteroids in orbital frequency. In particular, it has the form $S(f) \propto f^{-\gamma}$, where $\gamma = \frac{7}{3} - \beta$. The mass distribution enters only through $m_{\text{eff}(\alpha)}$, which scales the spectrum linearly. The effect of β on the shape is shown in Figure 4.2. In addition to the value used in Figure 4.1 ($\beta = -2/3$), spectra corresponding to two other values of β are shown: $\beta = 2/3$, which corresponds to a uniform disk surface density, and $\beta = -2$, which produces the same spectral slope, $\gamma = 13/3$, expected from the stochastic gravitational-wave (GW) background.

The total variance that asteroids contribute to the TOAs is

$$\sigma_\tau^2 = \int_{-\infty}^{\infty} S(f) df = 2 \int_0^{\infty} S(f) df. \quad (4.11)$$

For a power-law distribution (equation 4.8), this is

$$\sigma_\tau^2 = \left(\frac{G}{4\pi^2 M^2} \right)^{2/3} \frac{M_{\text{belt}} m_{\text{eff}}}{2c^2} \frac{\beta (f_+^{\beta-4/3} - f_-^{\beta-4/3})}{(\beta - \frac{4}{3})(f_+^\beta - f_-^\beta)} \sin^2 i. \quad (4.12)$$

Fixing the binary mass and asteroid frequency distribution to the values used

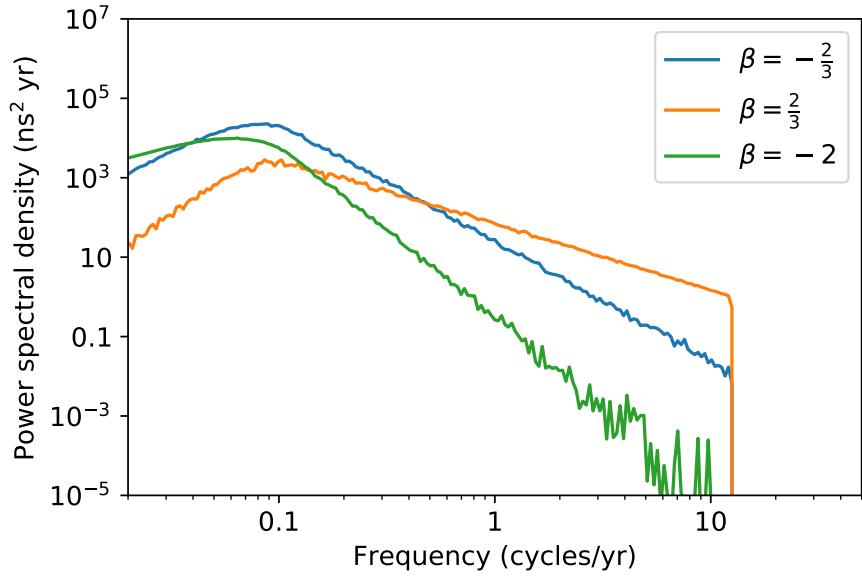


Figure 4.2: The effect of the power-law index, β , of the orbital frequency distribution of asteroids on the TOA residual spectrum. The β values shown correspond to a uniform distribution in semimajor axis ($\beta = -2/3$, $\gamma = 3$), a uniform disk surface density ($\beta = 2/3$, $\gamma = 5/3$), and a spectral slope equivalent to that expected from the stochastic gravitational-wave background ($\beta = -2$, $\gamma = 13/3$). All other parameters used are the same as those in Figure 4.1. The transmission function corresponding to a dataset length of 15 years has been applied in each case.

in Figure 4.1, this becomes

$$\sigma_\tau = 170 \text{ ns} \left(\frac{M_{\text{belt}}}{10^{-9} M_\odot} \right)^{1/2} \left(\frac{m_{\text{eff}}}{10^{-11} M_\odot} \right)^{1/2} \sin i. \quad (4.13)$$

For the asteroid mass distribution used in Figure 4.1, $M_{\text{belt}} = 2 \times 10^{-9} M_\odot$ and $m_{\text{eff}} \approx 1.5 \times 10^{-11} M_\odot$, so the standard deviation of the TOA perturbations in the case shown there is approximately 290 ns.

4.2.1 Effects of Model Subtraction

Any high-precision pulsar timing analysis requires fitting a timing model describing the pulsar's position, spindown, parallax, proper motion, and, if the

pulsar has a binary companion, its orbit. An additional signal not present in the timing model will often be partially degenerate with the timing model parameters, especially if the signal is spread over a wide range of frequencies. The spindown component, in particular, amounts to a polynomial trend, absorbing power at the lowest frequencies – those with periods comparable to, or longer than, the length of the data set. This sets a lower limit on the frequency of observable signals.

The amount of power removed by subtracting the timing model is quantified by the transmission function, $\mathcal{T}(f)$, the factor by which the power at frequency f is reduced (Blandford et al., 1984; Madison et al., 2013; Hazboun et al., 2019). For most pulsars, the spindown is described by a model in which the delay is a quadratic function of time. The transmission function for such a model behaves as

$$\mathcal{T}(f) \sim \left(\frac{f}{f_c}\right)^6, \quad f \lesssim f_c, \quad (4.14)$$

when f is significantly less than the critical frequency

$$f_c = \frac{1575^{1/6}}{\pi T_{\text{obs}}} \approx \frac{1.086}{T_{\text{obs}}}, \quad (4.15)$$

where T_{obs} is the length of the data set (see Appendix B). This causes the observed spectrum of the TOA residuals created by an asteroid belt to roll off at low frequencies, as illustrated in Figure 4.1 for various data set lengths. The same effect is applied in Figure 4.2. For a timing model including position and astrometric parameters, the transmission function would also include a dip around $f = 1 \text{ yr}^{-1}$, but this has a much less significant effect on the total variance, so we have not included it in Figures 4.1 and 4.2.

Table 4.1: Stability boundaries for some MSP binaries

PSR	P_B (days)	P_S (days)		P_P (days)	
		HW99	Q+20	HW99	Q+18
J0437–4715	5.741	1.65	1.67	16.7	18.2
J1614–2230	8.687	2.34	2.03	29.0	28.3
J1713+0747	67.825	18.7	15.8	217	220
J1741+1351	16.335	4.56	4.41	50.9	52.6
B1855+09	12.327	3.47	3.46	37.7	40.5
J1903+0327	95.174	7.40	5.58	664	634
J1909–3744	1.533	0.441	0.451	4.44	4.87
J1918–0642	10.913	3.07	3.07	33.4	34.9
J2043+1711	1.482	0.431	0.440	4.2	4.67

NOTE— P_S is the period of the outermost stable circular orbit around the pulsar, and P_P is the period of the innermost stable circular orbit around the binary. The orbital period, P_B , of each pulsar is included for reference. The estimates labeled HW99 are based on the results of Holman & Wiegert (1999) and those labeled Q+18 and Q+20 are based on those of Quarles et al. (2018) and Quarles et al. (2020), respectively. For further details, see § 4.7.

4.3 Results

The boundaries of stability for several MSP binary systems are given in Table 4.1. For each pulsar system, two estimates of the stability boundary are given, one using the interpolating polynomials given by Holman & Wiegert (1999) (equations 4.30 and 4.31), and the other based on the results of Quarles et al. (2018) and Quarles et al. (2020). We give further context on these estimates in Section 4.7. The systems in the table were selected because measurements of the pulsar and companion masses are available in the literature. All but one have almost perfectly circular orbits, a common characteristic of pulsar-white dwarf binaries. The exception, PSR J1903+0327, has a main-sequence companion on an eccentric orbit ($e = 0.437$). Simulated spectra of (P-type) asteroid belts around each system are shown in Figure 4.3, demonstrating the relation-

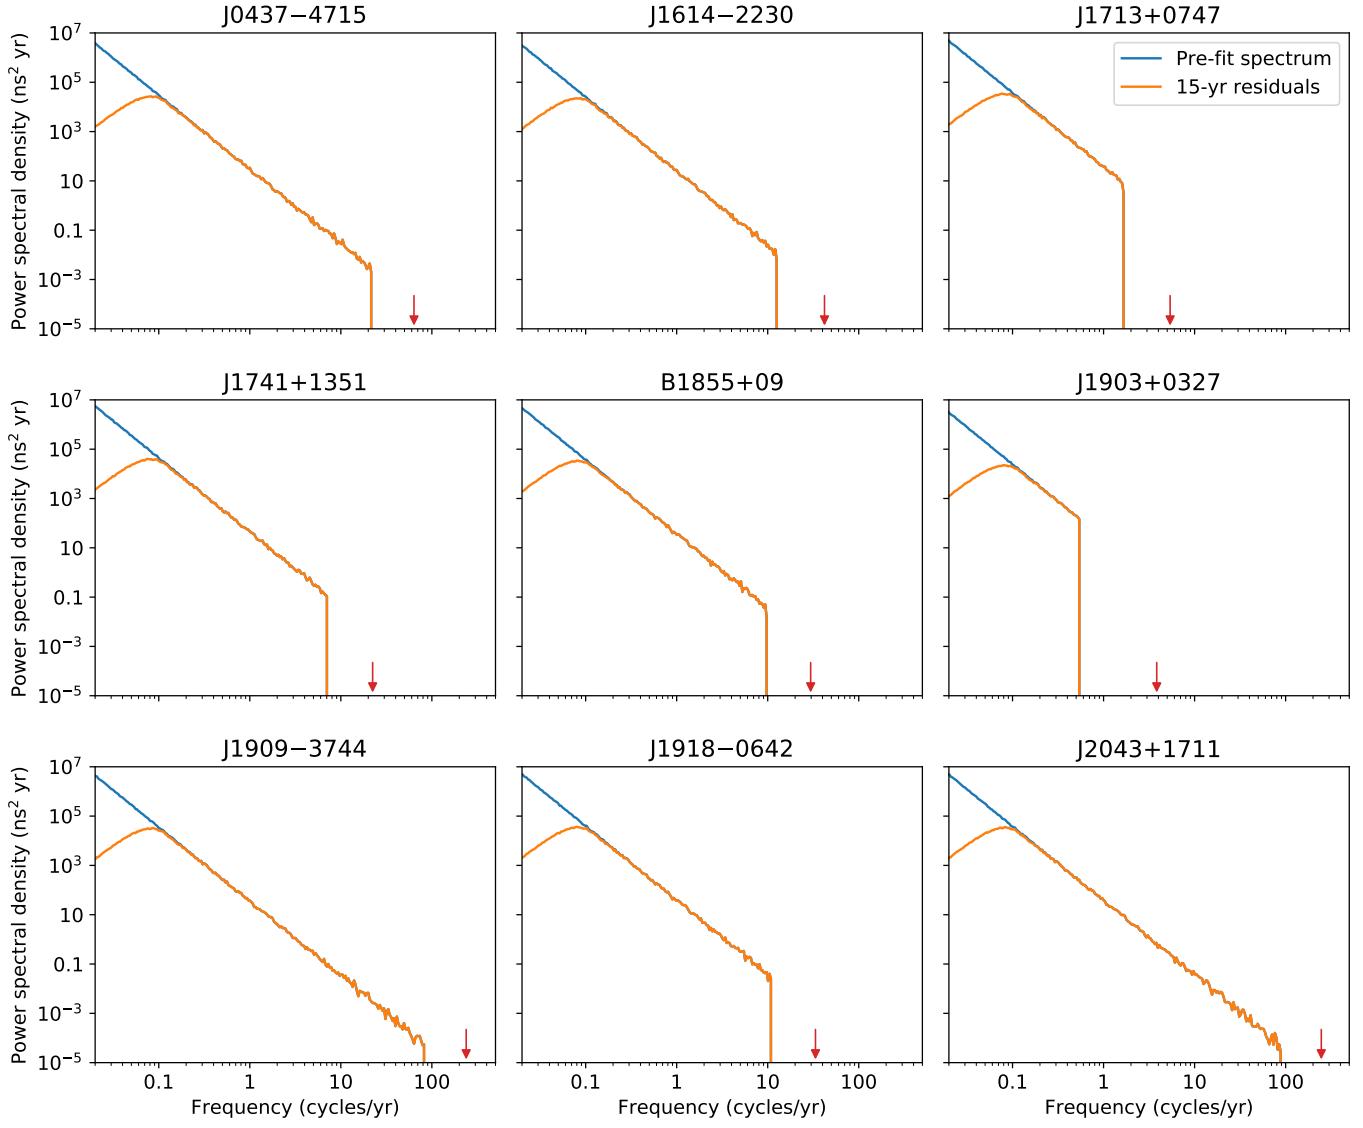


Figure 4.3: Simulated spectra of TOA perturbations for model P-type asteroid belts around each of the systems in Table 4.1. In each case, the spectrum is shown both before and after the subtraction of a quadratic spindown model for a nominal dataset length of 15 years, and the small red arrow indicates the pulsar’s orbital frequency. As in Figure 4.1, the total mass of the model asteroid belt in each case is chosen to match that of the solar system’s asteroid belt, and the distribution of asteroids in mass and orbital frequency has a power-law form (cf. equation 4.8), with indices $\alpha = -5/6$ in mass and $\beta = -2/3$ in frequency. More information on how parameters were selected can be found in Section 4.3.

ship between the binary orbital frequency and the frequency of the innermost stable P-type orbit.

While all the systems in Table 4.1 have orbital periods between 1 and 100 days, the orbital periods of known pulsar–white dwarf binary systems range from 2.46 hours in the case of PSR J0348+0432 (Antoniadis et al., 2013) to 3.37 years in the case of PSR B0820+02 (Hobbs et al., 2004). Even more extreme orbital periods are found in 47 Tuc R (Camilo et al., 2000), which has an ultra-light companion in a 96-minute orbit, and PSR J2032+4127 (Ho et al., 2017), which has a main-sequence companion in a highly eccentric 46-year orbit. The inner boundary of stability has a similar range. Using the minimum companion mass in each case, we find that it corresponds to a period of 6 hours for PSR J0348+0432 and 10 years for PSR B0820+02. For the longest-period objects, S-type asteroid belts may be more relevant than their P-type counterparts, but we will not consider them in detail here.

Each of the spectra seen in Figure 4.3 is an average of 10^4 realizations of an asteroid belt containing 10^4 asteroids with masses and orbital frequencies drawn at random from power-law distributions of the form given by equation (4.8). In each case the belt has a total mass of 2×10^{-9} solar masses (6.7×10^{-4} earth masses), comparable to that of the solar system’s asteroid belt. The maximum orbital frequency, f_+ , is chosen to match the frequency of the innermost stable P-type orbit in the particular binary system, as given in Table 4.1. The minimum orbital frequency, f_- , corresponds in each case to a semimajor axis of 20 au. This was chosen, somewhat arbitrarily, to make the longest orbital period longer than the 15-year nominal dataset length, since there are few constraints on how far a P-type asteroid belt extends away from

the binary system. Increasing the minimum orbital frequency narrows the range of frequencies over which the asteroid signals are distributed, and, if the total mass of the belt is held fixed, also increases the density of asteroids in orbital frequency space. Both of these effects tend to make the asteroid belt signal more recognizable.

As demonstrated in Figure 4.2, the shape of the spectrum is primarily determined by the orbital frequency distribution. In every case shown in Figure 4.3, the power-law index of this distribution is $\beta = -2/3$, which represents a uniform distribution in semimajor axis. This is the same as what is used in Figure 4.1.

The value $\alpha = -5/6$ is used for the power-law index of the mass distribution throughout Figure 4.3. This was chosen to match the value predicted by Dohnanyi (1969) for material in collisional equilibrium. As equations (4.9) and (4.10) demonstrate, the mass distribution affects the ensemble average spectrum only by determining the relationship between the mass cutoffs m_+ and m_- , the effective mass m_{eff} , and the total mass. However, in single realizations, the mass distribution is important because it determines how common asteroids with masses much larger than m_{eff} are. This has important implications for detecting individual asteroids (see Section 4.4.4).

4.4 Discussion

The results shown so far are based on a model asteroid belt whose total mass is equal to that of the solar system's asteroid belt. As shown in equation 4.9, the power spectral density of the TOA signal produced by the belt scales linearly

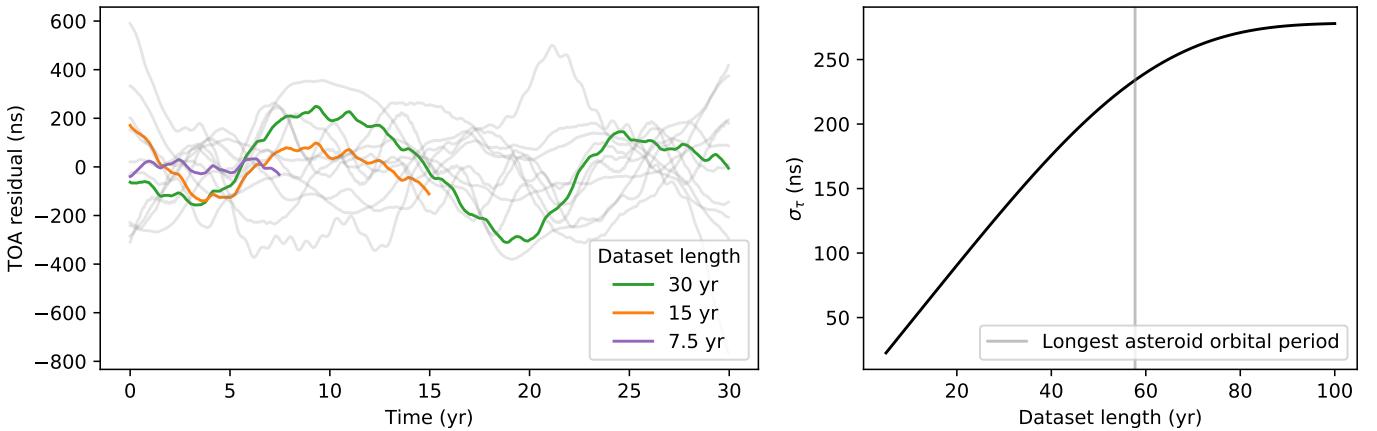


Figure 4.4: Example time series of TOA residuals created by an asteroid belt (after timing model subtraction), illustrating how the variance of the residuals increases with dataset length. The light gray lines correspond to different realizations with the same parameters used for Figure 4.1, with a quadratic trend removed. For one particular realization, highlighted in color, the first 7.5 and 15 years of the time series have been isolated, and the fitting and subtraction of the trend performed separately. This is quantified in the right panel, which shows how the standard deviation, σ_τ , of the residuals behaves as a function of the dataset length in the ensemble average. While σ_τ initially grows linearly with time, it eventually stops growing and stabilizes at approximately 290 ns, the value predicted by equation (4.12). This happens after the dataset length exceeds the longest asteroid orbital period.

with both the total mass of the belt and the mass of a typical asteroid.

4.4.1 Comparison with Other Low-Frequency Timing Signals

An asteroid belt is hardly the only phenomenon capable of producing red noise in pulsar timing residuals. In fact, red noise of a much larger amplitude than that considered here is common in canonical pulsars (those with periods of order one second and period derivatives of order 10^{-15} , as opposed to MSPs, which have periods of a few milliseconds and period derivatives of order 10^{-20}), and is generally understood to arise from stochastic variations

in the pulsar’s rotation rate, caused by some combination of magnetospheric torques (Kramer et al., 2006; Lyne et al., 2010) and instabilities arising from differential rotation between the neutron star’s surface and its superfluid interior (Jones, 1990). It is precisely the fact that MSPs are relatively free of this kind of timing noise that, combined with their shorter periods, makes them ideal for high-precision timing applications, including GW searches.

The amplitude of red timing noise in canonical pulsars has been observed to scale with the period derivative, \dot{P} (Cordes & Helfand, 1980; Arzoumanian et al., 1994), making it unsurprising that the effect should be smaller in MSPs. Indeed, Shannon & Cordes (2010) developed a scaling relation that is consistent with the lower levels of red noise observed in MSPs as well as the higher levels observed in canonical pulsars, indicating that they may have the same origin. Lam et al. (2017) followed this up with a fit including more MSPs, again finding similar behavior between canonical pulsars and MSPs. However, it remains possible that at least some of the observed red noise in MSPs may have another origin.

Perturbations caused by encounters with nearby stars, or with smaller objects such as interstellar planets or asteroids, could also give rise to low-frequency pulsar timing signals. However, encounters with stars typically take place too slowly for changes in the acceleration of the pulsar to be detectable (Phinney, 1993), and detectable encounters with smaller objects are rare events with a distinctive signature (Jennings et al., 2020a). The effects of both kinds of encounters can therefore be distinguished from the red noise produced by an asteroid belt.

4.4.2 Stationarity

The TOA perturbations generated by an asteroid belt can be compared to those produced by random walks in pulse phase, frequency, or frequency derivative, all of which are statistically non-stationary. However, the signal from an asteroid belt is in principle stationary when measured over a long enough span of time — there is always a longest-period asteroid. If the asteroids have a power-law distribution in orbital frequency (equation 4.8), so that the spectrum of the TOA perturbations also has a power law form, the variance of the residuals will grow with dataset length, but only up to a certain point, determined by the low-frequency (long-period) cutoff of the distribution. This can be seen in Figure 4.4. In the case shown there, where asteroids are distributed uniformly in semimajor axis, the variance initially grows linearly with time, but other power-law indices are possible depending on the asteroid distribution. As discussed in Section 4.2, if the power-law index of the orbital frequency distribution is β , the spectrum will have the form $S(f) \propto f^{-\gamma}$, where $\gamma = \frac{7}{3} - \beta$ (equation 4.9). In this case, the standard deviation of the TOA residuals will increase with time as $\sigma_{\text{TOA}} \propto T^\delta$, with $\delta = \frac{1}{2}(\gamma - 1) = \frac{2}{3} - \frac{1}{2}\beta$. A uniform distribution in semimajor axis gives $\gamma = 3$, while a uniform surface density gives $\gamma = 5/3$. For comparison, random walks in pulse phase, frequency, and frequency derivative correspond to $\gamma = 2, 4$, and 6 , respectively, and the expected spectrum of the stochastic gravitational-wave background corresponds to $\gamma = 13/3$.

4.4.3 Survival of Asteroids

Several factors other than the presence of a binary companion affect the stability of asteroid belts around pulsars. Many of these are especially relevant for millisecond pulsars, which are generally billions of years old. The seasonal Yarkovsky effect (e.g. Rubincam, 1998), a thermal drag mechanism acting on objects whose spin axes are tilted relative to their orbits, can cause small asteroids to migrate inward until they are evaporated. However, the migration timescale is proportional to asteroid radius, and 5 km asteroids can survive for at least 250 Myr at 1 au (Cordes & Shannon, 2008; Shannon et al., 2013).

An asteroid will be evaporated if it is heated enough for its equilibrium temperature to exceed the melting point of the materials that compose it. Asteroids in the vicinity of a pulsar are heated by a number of mechanisms in addition to thermal emission from the pulsar's surface, including particle and X-ray emission driven by the pulsar's spindown and Ohmic dissipation driven by currents between the asteroid and the pulsar's magnetosphere. A detailed analysis of these effects was carried out by Cordes & Shannon (2008), who reached the conclusion that asteroids do not begin to evaporate until they are within approximately 10^{10} cm (6×10^{-4} au) of the pulsar, which corresponds to a seven-minute orbital period, or a frequency of 7×10^5 cycles per year. It follows that asteroids on stable P-type orbits in pulsar binary systems, even those with the shortest orbital periods, are almost certainly safe from evaporation.

For asteroids that are massive enough or close enough together, their mutual gravitational interactions can also be a destabilizing influence. Examining this issue, Heng & Tremaine (2010) conclude that the timescale for destabilization in "cold" disks increases exponentially with the separation between the

asteroids, with a distance of around 10 times the Hill radius sufficing for a lifetime of 100 Myr.

In globular clusters, stellar encounters can limit the outer radius of an asteroid belt. An asteroid's orbit is likely to be disrupted if another star approaches the pulsar to within a few times the orbit's radius. In a cluster with stellar density n_* and velocity dispersion σ , the rate of encounters with periapsis distance less than r is given by

$$\Gamma_*(r) = \pi r^2 \sigma n_* \left(1 + \frac{2GM}{\sigma^2 r} \right) \quad (4.16)$$

(cf. Verbunt & Hut, 1987), where the factor in parentheses accounts for enhancement of the encounter cross-section due to gravitational focusing, and M is the total mass of the interacting bodies. For a typical globular cluster with $n_* = 10^4 \text{ pc}^{-3}$ and $\sigma = 10 \text{ km s}^{-1}$, encounters with $r < 10 \text{ au}$ happen approximately once every 250 Myr, assuming typical masses of 1.0 and $1.4 M_\odot$ for the star and pulsar, respectively. In the less dense environment of the galactic disk, where the stellar density is some four orders of magnitude smaller, encounters are much rarer, and even at 100 AU are unlikely to happen in the lifetime of the pulsar.

4.4.4 Distinguishing Individual Asteroids

One way to test the hypothesis that observed red noise in TOAs from a pulsar is caused by an asteroid belt is to separate out the signals caused by individual asteroids. This is possible only if the dataset is long enough to acquire sufficient resolving power in the frequency domain. If there are only a handful of asteroids, it may be possible to observe for long enough to resolve each indi-

vidual asteroid in frequency space, but for denser asteroid belts this becomes impractical. The average number, ΔN , of asteroids in a single frequency bin is related to the density,

$$\frac{dN}{df} = \frac{N\beta f^{\beta-1}}{f_+^\beta - f_-^\beta}, \quad (4.17)$$

of asteroids in frequency space by

$$\Delta N = \frac{dN}{df} \Delta f, \quad (4.18)$$

where Δf is the width of the frequency bin. Asteroids will be individually resolvable if $\Delta N \lesssim 1$. Because the minimum Δf achievable with a dataset of length T is approximately $1/T$, this means that asteroids will be individually resolvable in frequency only if

$$T \gtrsim \frac{dN}{df}. \quad (4.19)$$

Using the parameters from Section 4.2, this becomes

$$T \gtrsim 940 \text{ yr} \left(\frac{N}{10^4} \right) \left(\frac{f}{1 \text{ yr}^{-1}} \right)^{-5/3}. \quad (4.20)$$

This suggests that realistic observations will usually be in the opposite regime, in which $\Delta N \gg 1$. In such cases, the central limit theorem may be applied to individual frequency bins, so the probability distribution of the complex signal amplitude in each bin will be Gaussian, and the probability distribution of the power will be exponential, with mean $S(f)\Delta f$. Since the standard deviation of an exponential distribution is equal to its mean, the standard deviation of the power in the frequency bin centered at f will also be $S(f)\Delta f$.

An additional sinusoidal signal may be considered distinguishable from the asteroid belt signal if the probability of it arising by chance in the asteroid belt model is sufficiently small. For a signal of amplitude $\Delta\tau$, the ratio of the power

in the signal to power in asteroid belt noise is

$$S/N = \frac{(\Delta\tau)^2}{4S(f)\Delta f}, \quad (4.21)$$

where Δf is the frequency resolution. The signal may be said to be detected if the signal-to-noise ratio exceeds a threshold value, s . Under the null hypothesis $\Delta\tau = 0$, the probability that the threshold will be exceeded is

$$p = \exp(-s). \quad (4.22)$$

If N_f frequency bins are examined, the probability that at least one will exceed the threshold value by chance is

$$P = 1 - (1 - p)^{N_f} \approx N_f p. \quad (4.23)$$

It follows that, to achieve a false positive probability of P , the threshold should be set to

$$s = \ln\left(\frac{N_f}{P}\right). \quad (4.24)$$

The signal is then detectable if

$$\frac{(\Delta\tau)^2}{4S(f)\Delta f} \gtrsim \ln\left(\frac{N_f}{P}\right). \quad (4.25)$$

In the case that the signal is created by another asteroid, so that $\Delta\tau$ is given by equation (4.1), and the asteroids in the belt have power law distributions in mass and orbital frequency (equation 4.8), so that $S(f)$ is given by equation (4.9), the signal-to-noise ratio (equation 4.21) becomes

$$S/N = \frac{Nm^2}{M_{\text{belt}}m_{\text{eff}}} T \left(\frac{dN}{df} \right)^{-1}. \quad (4.26)$$

Here m_{eff} is given by equation (4.10), and dN/df is the density of asteroids in orbital frequency (equation 4.17). An additional asteroid of mass m is therefore distinguishable from the bulk of the belt if

$$m \gtrsim \sqrt{\frac{\beta f^{\beta-1}}{(f_+^\beta - f_-^\beta) T} \ln\left(\frac{N_f}{P}\right) \sqrt{M_{\text{belt}}m_{\text{eff}}}}. \quad (4.27)$$

Taking $N_f = 100$ and $P = 0.05$ (so that the signal-to-noise threshold is $s = 7.6$), and using the parameters from Section 4.2, this becomes $m \gtrsim m_{\text{thres}}$, where

$$m_{\text{thres}} = 6.7 \times 10^{-11} M_{\odot} \left(\frac{f}{1 \text{ yr}^{-1}} \right)^{-5/6} \left(\frac{T}{15 \text{ yr}} \right)^{-1/2}. \quad (4.28)$$

The expected number of asteroids with masses greater than the threshold, m_{thres} , is

$$N_{>} = \frac{N(m_{+}^{\alpha} - m_{\text{thres}}^{\alpha})}{m_{+}^{\alpha} - m_{-}^{\alpha}}. \quad (4.29)$$

For the parameters in Section 4.2, this is $N_{>} = 11$. In other words, in an asteroid belt like the one shown in Figure 4.1, approximately 11 of the 10^4 asteroids would be detectable individually in a 15-year dataset, assuming that the uncertainties in the TOA measurements are negligible compared to the signal from the rest of the asteroid belt.

4.4.5 Implications for Gravitaional-Wave Searches

In a number of ways, the pulsar timing signal produced by an asteroid belt closely resembles the signal expected from gravitational wave (GW) sources. Like the stochastic GW background, asteroid belts around pulsars should produce correlated noise in TOAs with frequencies of order 1 yr^{-1} and amplitudes of tens of nanoseconds; and, like continuous wave sources, individual large asteroids should produce approximately sinusoidal TOA perturbations. Both GW and asteroid-belt signals can be distinguished from the effects of dispersion and scattering produced by the interstellar medium in that they are achromatic, i.e., they do not depend on radio frequency. This means that unlike, for example, TOA variations caused by changes in dispersion measure (DM),

variations caused by the presence of an asteroid belt cannot be measured and corrected for by comparing signals at different radio frequencies.

However, there is at least one important way in which asteroid-belt and GW signals differ. In particular, the Earth-term component of any GW signal should be correlated across different pulsars, with a characteristic spatial pattern, originally described by Hellings & Downs (1983), that arises from the quadrupolar nature of GWs. Asteroid belts, on the other hand, belong to particular pulsars; even if all pulsars had asteroid belts with identical statistical properties, there would be no reason to expect the signals they produced to be correlated, since the masses, phases, and orbital frequencies of individual asteroids would differ from one pulsar to the next.

4.5 Summary and Conclusions

Asteroid belts in pulsar systems produce achromatic, low-frequency timing noise, which arises from the orbital reflex motion of the pulsar. The presence of a binary companion in a pulsar system has the effect of destabilizing orbits close to the pulsar. This provides a natural upper cutoff for the orbital frequencies of any asteroids, but does not exclude the possibility of an asteroid belt entirely, since sufficiently distant orbits remain stable.

The hypothesis that observed timing noise in one or more MSPs is produced by an asteroid belt may be tested by looking for evidence of stationarity, which should be present only if the frequency corresponding to the outer edge of the belt is observable; or by trying to isolate the signal from individual large asteroids. Completely resolving asteroids in orbital frequency is possible only

for very sparse asteroid belts or very long datasets, but it may be possible to detect individual large asteroids well before this point. Nevertheless, in the near term it is likely to remain challenging to determine whether particular instances of pulsar timing noise are caused by asteroid belts.

Asteroid belts are potentially important as noise sources in searches for low-frequency gravitational waves. The TOA signal produced by an asteroid belt is similar to that expected from a GW background in its frequency, amplitude, and achromatic nature, and may have a similar power-law spectrum. However, it differs in that it is not expected to be correlated between different pulsars. This makes it particularly important for GW searches to consider the correlation between signals in different pulsars.

4.6 Acknowledgements

The authors are members of the NANOGrav Physics Frontiers Center, which receives support from the National Science Foundation (NSF) under award number 1430284. This paper made use of data from the ATNF pulsar catalogue¹ (Manchester et al., 2005). We thank D. Lai for useful comments regarding orbital stability.

¹<https://www.atnf.csiro.au/research/pulsar/psrcat/>

4.7 Appendix: Stable Orbits in a Binary System

In any binary system, there is an approximate maximum radius for stable S-type orbits and an approximate minimum radius for stable P-type orbits. The boundary of the stable region is actually irregular and fractal in nature, with “teeth” corresponding to resonant orbits (Nagler, 2005; Shevchenko, 2015), but this approximation will suffice for our purposes.

Holman & Wiegert (1999) conducted a numerical investigation of the stability of circular S- and P-type orbits for various combinations of the eccentricity e and mass ratio μ of the central binary. They found that S-type orbits with semi-major axes less than a critical value a_S , and P-type orbits with semi-major axes greater than a critical value a_P , were stable for the duration of their simulations (10^4 orbital periods). They give the maximum semi-major axis for stable S-type orbits as a quadratic polynomial in e and μ :

$$\begin{aligned} a_S = & [(0.464 \pm 0.006) + (-0.380 \pm 0.010)\mu \\ & + (-0.631 \pm 0.034)e + (0.586 \pm 0.061)\mu e \\ & + (0.150 \pm 0.041)e^2 + (-0.198 \pm 0.074)\mu e^2]a_B. \end{aligned} \quad (4.30)$$

Here a_B is the binary separation. We adopt Holman & Wiegert’s convention for the mass ratio, letting μ denote the ratio $m_2/(m_1 + m_2)$, where m_1 is the mass of the primary (the star the asteroid orbits), and m_2 is the mass of the secondary. Similarly, they give the minimum semi-major axis for stable P-type

orbits as

$$\begin{aligned}
a_P = & [(1.60 \pm 0.04) + (5.10 \pm 0.05)e \\
& + (-2.22 \pm 0.11)e^2 + (4.12 \pm 0.09)\mu \\
& + (-4.27 \pm 0.17)e\mu + (-5.09 \pm 0.11)\mu^2 \\
& + (4.61 \pm 0.36)e^2\mu^2]a_B.
\end{aligned} \tag{4.31}$$

Here μ is taken to lie between 0 and 0.5, since the mass ratios μ and $1 - \mu$ are equivalent up to labelling of the component stars.

More recently, extensive simulations were carried out by Quarles et al. (2018) for P-type orbits and Quarles et al. (2020) for S-type orbits. Their results largely agree with those of Holman & Wiegert (1999), but cover the (e, μ) parameter space in greater detail, allowing for grid interpolation of the critical semi-major axes a_S and a_P .

A different, more analytical, approach to determining the boundaries of the region of stable orbits was taken by Szebehely (1980), who made use of the fact that the quantity

$$C_J = \frac{2GM_1}{r_1} + \frac{2GM_2}{r_2} + \omega^2 r^2 - v^2, \tag{4.32}$$

called the Jacobi constant, is conserved in the circular restricted three-body problem. Here M_1 and M_2 are the masses of the two primary bodies, ω is their orbital frequency, and r_1 and r_2 are the distances between each and the third, small body, while r is the small body's distance from the center of mass and v is its velocity in the synodic frame (co-rotating with the binary). For a particular value of C_J , the surfaces corresponding to $v = 0$ are called zero-velocity surfaces. They bound regions in space that a small body with that particular Jacobi constant cannot enter, since the square of its velocity must always be positive.

Szebehely (1980) made use of this by calculating the Jacobi constant for initial conditions that, in the appropriate two-body approximation, would correspond to a circular orbit. S-type orbits were considered stable if the zero-velocity surface prevented the small body from escaping to infinity. Similarly, P-type orbits were considered stable if the zero-velocity surface prevented the small body from approaching arbitrarily closely to either of the primaries. This definition of stability, termed Hill-type stability (Szebehely, 1978), differs qualitatively from that of Holman & Wiegert (1999): on the one hand, it is global and not limited by a finite integration time or numerical precision; but on the other hand, it does not take into account all modes of instability. S-type orbits that are Hill stable may eventually result in a collision with the primary, while P-type orbits that are Hill stable may eventually escape. Nevertheless, the results obtained by Szebehely (1980) for S-type orbits and by Szebehely & McKenzie (1981) for P-type orbits are in broad agreement with the results of Holman & Wiegert (1999).

We estimate the boundary of stability for P-type orbits using equation (4.31) throughout the paper. Because the results of Szebehely & McKenzie (1981) largely agree with those of Holman & Wiegert (1999), our results should not be sensitive to the particular criterion for stability adopted.

CHAPTER 5

THE EFFECT OF PULSE SHAPE CHANGES ON PULSAR TIMING PRECISION

Time-of-arrival (TOA) measurements of pulses from pulsars are conventionally made by a template matching algorithm that compares folded profiles to a long-term average pulse shape, but the shapes of pulses can and do vary, leading to errors in TOA estimation. In particular, all pulsars show stochastic variations in shape between successive pulses that only partially average out in any finite-length observation, a phenomenon which will only become more problematic for pulsar timing precision as more sensitive telescopes are built. Pulse shape changes also occur for other reasons, including pulsar nulling and mode changing, time-variable dispersion and scattering by the interstellar medium, and radio frequency interference. In this paper, we quantify the ways in which pulse shape changes affect TOA estimates, introducing a formula which relates shape changes to the TOA estimation errors they produce, and examining the behavior of TOA estimates in a model in which pulses are made up of several components that vary independently in amplitude and pulse phase. We also describe several techniques for identifying and characterizing such pulse shape variations, and apply them to data from the Vela pulsar, PSR B0833–45. The results described here are valuable in any application of pulsar timing, including but not limited to efforts by pulsar timing arrays to detect the low-frequency stochastic gravitational-wave background.

5.1 Introduction

Pulsar timing involves making precise measurements of times of arrival (TOAs) of pulses from pulsars. Such measurements can be used for a variety of scientific purposes, including characterizing energy loss from gravitational-wave emission in binary systems (e.g., Weisberg et al., 2010), testing other predictions of the general theory of relativity (Stairs, 2003), detecting small bodies orbiting pulsars (e.g., Wolszczan, 1994), and studying the structure of the ionized interstellar medium (e.g., Cordes & Lazio, 2002b). Additionally, current efforts by pulsar timing arrays (PTAs; Desvignes et al. 2016; Hobbs 2013), including the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; Arzoumanian et al. 2018a), seek to use precision pulsar timing to directly detect nanohertz gravitational radiation.

All applications of pulsar timing rely on the link between arrival times and the spin phase of a neutron star (NS) and the spin stability of the NS itself. While past successes have shown that pulsars make excellent clocks, they are not perfect in this respect, and understanding their limitations is crucial to all pulsar timing efforts. Such limitations arise from a number of sources, including rapid spin-up events called glitches (Espinoza et al., 2011), which are thought to be a result of differential rotation between a neutron star's crust and its superfluid interior (Glampedakis & Andersson, 2009); and torque fluctuations in the magnetosphere, which produce spin noise with power concentrated at low frequencies (Kramer et al., 2006; Lyne et al., 2010). Additionally, and most importantly for our purposes here, the radio emission from a pulsar is intrinsically variable, causing the shapes of individual pulses vary stochastically. This variability is likely the result of changes in the relativistic flows in

the magnetosphere. Fortunately, averages of many pulses converge to shapes that have proven to be stable on timescales of decades.

This stability is the basis of the template matching algorithms used to convert pulse shapes into TOA measurements (Taylor, 1992; Cordes, 2013b). For the most part, these algorithms rely on a simple model in which an observed pulse profile, constructed as the average of a large number of pulses, is treated as having a known, fixed shape, modified only by the addition of white noise. When the profile deviates from this template shape, TOA estimation errors result. Such deviations occur in practice for several reasons, many of which are outlined in Section 5.2 below. In particular, since a profile is the sum of only finitely many pulses, and the shapes of individual pulses vary stochastically, the shape of the profile also varies, albeit at a lower level. This single-pulse stochasticity limits the precision of TOA estimates for pulsars observed at very high signal-to-noise ratio (S/N).

In the remainder of this paper, we will describe the TOA estimation errors that arise from various kinds of pulse shape changes, and describe methods for identifying the presence of such shape changes and ascertaining the extent to which they affect TOA measurements. In Section 5.3, we describe the amplitude-modulated shot noise model underlying our current understanding of pulsar emission. In Section 5.4, we introduce the conventional algorithm used for TOA generation, and describe how the corresponding uncertainties may be estimated. Formulas for predicting the TOA estimation errors caused by particular shape changes are introduced and applied to simulated data in Section 5.5. A series of techniques for characterizing pulse shape changes, particularly those created by single-pulse stochasticity, and measuring the TOA

estimation errors they create, is outlined in Section 5.6. Sections 5.7 and 5.8 describe other phenomena which may cause shape changes, and outline how the previously-described methods may be applied to them. Chapter 6 will discuss methods for mitigating the effects of TOA estimation errors caused by pulse shape changes in certain cases.

5.2 Causes of Pulse Shape Changes

Observed pulse shapes deviate from the template pulse shape for several reasons. Here, we will primarily be concerned with shape changes caused by the intrinsic stochastic behavior of single pulses (Cordes & Downs, 1985; Cordes et al., 1990; Liu et al., 2012). These shape changes are generally broadband, i.e., strongly correlated between adjacent frequencies, and are thought to be caused by the motion of emitting regions within the pulsar’s magnetosphere. We describe several ways to characterize timing errors caused by shape changes of this form, with reference to simulations and data from observations of the Vela pulsar, PSR B0833–45, in Section 5.6.

Some pulsars also show other kinds of intrinsic pulse shape changes. Nulling pulsars (Backer, 1970a; Gajjar et al., 2012; Sheikh & MacDonald, 2021) occasionally cease to emit detectable pulses for short periods of time, ranging from a few pulse periods to several days. Mode-changing pulsars, including PSR B0329+54 and PSR B1237+25 (Backer, 1970b; Lyne, 1971; Bartel et al., 1982) alternate between two or more modes of emission on similar time scales. In at least one millisecond pulsar, PSR J1643–1224, evidence has been seen for a shape change that appears abruptly, decays over a period of months, and re-

sults in a permanently modified pulse shape (Shannon et al., 2016). A similarly abrupt shape change, which may have a similar origin, was recently seen in the millisecond pulsar J1713+0747 (Xu et al., 2021; Lam, 2021). Some recent observations of PSR J1713+0747 shortly before and after this event, made using the Green Bank Telescope (GBT) and the Canadian Hydrogen Intensity Mapping Experiment (CHIME), are given in Chapter 7. Approaches to characterizing these phenomena and their effects on timing will be described in Section 5.7.

Radio emission from pulsars is dispersed and scattered as it travels through the interstellar medium (ISM). This, too, causes pulse shape changes (see, e.g., Rickett, 1990). The most significant ISM propagation effect is dispersion, which creates a delay proportional to the inverse square of the observation frequency, ν , and to the column density of electrons along the line of sight, known in this context as the dispersion measure (DM). Another significant effect is interstellar scattering, a result of the propagation of radio waves along multiple paths through the nonuniform distribution of electrons in the ISM. This leads to a frequency-dependent broadening of pulse profiles, as well as to delays proportional to $\nu^{-4.4}$ (assuming density fluctuations in the ISM can be described by a Kolmogorov spectrum). The detailed analysis of these phenomena and their effects on pulsar timing is closely tied to the physics and astrophysics of the interstellar medium, and is beyond the scope of this paper.

Instrumental effects and radio-frequency interference (RFI) can also produce effective pulse shape changes. Because pulsar emission is often strongly polarized, imperfections in the polarization response of a receiver can change the apparent pulse shape unless they are carefully accounted for (Heiles et al., 2001; Wahl et al., 2021). Additionally, channel gain and timing mismatches in

time-interleaved analog-to-digital converter (ADC) systems can lead to image rejection artefacts that present themselves as frequency-reversed “ghost images” of the dispersed pulse (Alam et al. 2021a, §2.3.1; see also Kurosawa et al. 2001). The most significant effect, however, is generally RFI. This is often manifest as short impulses, which can alter the shape of an observed pulse if they coincide with it in time, or as long-term, narrow-band signals, which, if not removed from the data, can create a sinusoidal ripple effect most noticeable in the off-pulse region of a profile. Some of the techniques described here lend themselves well to identifying data which has been corrupted by RFI and understanding the extent to which it affects pulsar timing. This will be explored further in Section 5.8.

5.3 The Amplitude-Modulated Shot Noise Model

At very high time resolution, radio emission from pulsars consists of polarized shot noise (Cordes, 1976; Jenet et al., 2001; Cordes et al., 2004). That is, the electric field is given by a superposition of coherent shot pulses:

$$\mathbf{E}(t) = \Re \left[\sum_j \mathbf{x}_j s(t - t_j) e^{2\pi i \nu_0 t} \right]. \quad (5.1)$$

Here we have assumed that all shots have the same shape, $s(t)$, but this need not be the case. The Fourier transform of the shot shape determines the spectrum of the pulsar, so the fact that pulsars are detected at radio frequencies above 1 GHz means that the intrinsic width of a typical shot is less than 1 ns.

When a pulsar is observed with a quadrature heterodyned receiver system, the signal is first mixed to baseband by combining it with a local oscillator

and then band-pass filtered. This results in a time series of complex baseband voltages in each polarization channel. The polarization components of the baseband signal are given by

$$\mathbf{E}(t) = \sum_j \mathbf{x}_j s(t - t_j) * h(t), \quad (5.2)$$

where $h(t)$ is the impulse response function of the receiver, and $*$ represents convolution. The width of $h(t)$ is roughly the inverse of the receiver bandwidth, which in most cases is significantly larger than the intrinsic duration of individual shots and so dominates the width of the observed shots.

Shots can be assumed to occur as a Poisson process in time, with rate $\eta(t)$. The probability distribution $f_{\mathbf{x}}(\mathbf{x}; t)$ of the complex amplitude vector, \mathbf{x} , can have arbitrary form, but the fact that shots are polarized means that its components are correlated.

Both the polarization components of the baseband signal, as well as the total intensity, $I(t)$, can be described using an “envelope function”, $a(t)$:

$$\mathbf{E}(t) = a(t) \mathbf{m}(t), \quad (5.3)$$

$$I(t) = a(t)^2 |\mathbf{m}(t)|^2. \quad (5.4)$$

Here $\mathbf{m}(t)$ is a two-component stationary shot noise process. The two components are correlated with each other as a result of the pulsar’s polarization, and satisfy

$$\langle |\mathbf{m}(t)|^2 \rangle = 1. \quad (5.5)$$

The envelope function is determined by both the distribution of shot noise amplitudes and the rate at which they occur. In particular, we have

$$a(t)^2 = \eta(t) \int |\mathbf{x}|^2 f_{\mathbf{x}}(\mathbf{x}; t) d^2 x. \quad (5.6)$$

The envelope function is itself stochastic: not all pulses have the same mean shape. Indeed, the stochastic behavior of the envelope function will be our main concern in the remainder of this paper. Before we move on, however, it will be useful to consider how pulsars are typically observed.

In most pulsar observations, the baseband signal is channelized and folded profiles are produced as a function of pulse phase, ϕ , in each of several frequency channels. This process effectively smooths the data by dividing the period into a specified number, N_ϕ , of phase bins. If the expected number of shots contributing to each phase bin is sufficiently large, which is typically the case, the ability to distinguish individual shots is lost, and the averaged shot noise process, $\mathbf{m}(\phi)$, approaches complex-valued, white, Gaussian noise. In this case, the profiles are well described by the amplitude-modulated noise model of Rickett (1975):

$$p(\phi) = a(\phi)^2 M(\phi). \quad (5.7)$$

If no averaging is involved in constructing the profile, $M(\phi)$ will have exponential statistics, with 100% modulation. However, typically the process of forming profiles is accompanied by some averaging in frequency and/or time, so the modulation due to this “self-noise” is lessened, and the profile resembles the envelope function with only a small amount of added noise.

5.4 TOA Estimation

The conventional matched-filtering algorithm used for TOA estimation is laid out in detail in Taylor (1992, Appendix A). It can be characterized as the maximum likelihood estimator for the model in which the profile, $p(\phi)$, is described

as

$$p(\phi) = au(\phi - \tau) + n(\phi), \quad (5.8)$$

where $u(\phi)$ is the template (normalized to unit maximum), $n(\phi)$ is white noise with variance σ_n^2 , and the amplitude a and phase offset τ are the parameters. Both $p(\phi)$ and $u(\phi)$ should be thought of as periodic in ϕ with period 1. We are using a continuous notation for clarity here – in reality, values of $p(\phi)$ and $u(\phi)$ are known only at a finite number, N_ϕ , of phase bins (see Section 5.11 for details). Shifting $u(\phi)$ by a fractional number of phase bins can be performed using a standard Fourier-domain technique.

The likelihood for this model has the form $\mathcal{L}(a, \tau) \propto e^{-\frac{1}{2}\chi^2(a, \tau)}$, where

$$\chi^2(a, \tau) = \frac{N_\phi}{\sigma_n^2} \left(\int_0^1 p(\phi)^2 d\phi - 2a \int_0^1 p(\phi)u(\phi - \tau) d\phi + a^2 \int_0^1 u(\phi)^2 d\phi \right). \quad (5.9)$$

Completing the square in a gives

$$\chi^2(a, \tau) = \frac{N_\phi \langle u^2 \rangle}{\sigma_n^2} \left[(a - \hat{a}(\tau))^2 - \hat{a}(\tau)^2 \right] + \frac{N_\phi \langle p^2 \rangle}{\sigma_n^2}, \quad (5.10)$$

where $\langle u^2 \rangle = \int_0^1 u(\phi)^2 d\phi$, $\langle p^2 \rangle = \int_0^1 p(\phi)^2 d\phi$ and

$$\hat{a}(\tau) = \frac{1}{\langle u^2 \rangle} \int_0^1 p(\phi)u(\phi - \tau) d\phi. \quad (5.11)$$

It follows that the maximum likelihood (minimum χ^2) occurs at $\tau = \hat{\tau}$ and $a = \hat{a}(\hat{\tau})$, where

$$\hat{\tau} = \operatorname{argmax}_\tau \hat{a}(\tau)^2. \quad (5.12)$$

The quantity $\hat{a}(\tau)$ is, up to a constant factor, the cross-correlation between the profile and the template. Because a is positive (the profile is not inverted compared to the template), the maximum value of $\hat{a}(\tau)^2$ corresponds to the maximum (rather than the minimum) value of $\hat{a}(\tau)$. The time-of-arrival estimate, $\hat{\tau}_{\text{MF}}$, that maximizes the likelihood of model (5.8) can therefore be computed

by maximizing $\hat{a}(\tau)$ numerically. This procedure is identical to that described by Taylor (1992).

The uncertainties in a and τ may be determined by expanding χ^2 in a power series around its minimum. By calculating the appropriate derivatives, one can determine that

$$\chi^2(a, \tau) \approx \chi^2(\hat{a}, \hat{\tau}) + \frac{(a - \hat{a})^2}{\sigma_a^2} + \frac{(\tau - \hat{\tau})^2}{\sigma_\tau^2}, \quad (5.13)$$

where the variances, σ_a^2 and σ_τ^2 , are given by

$$\sigma_a^2 = \frac{\sigma_n^2}{N_\phi \langle u^2 \rangle}, \quad (5.14)$$

$$\sigma_\tau^2 = \frac{\sigma_n^2}{\hat{a} N_\phi} \left(- \int_0^1 p(\phi) u''(\phi - \hat{\tau}) d\phi \right)^{-1}. \quad (5.15)$$

If the profile is well described by the model (equation 5.8), one can substitute $p(\phi) = \hat{a}u(\phi - \hat{\tau})$ in equation (5.15) and integrate by parts to obtain

$$\sigma_\tau^2 = \frac{W_{\text{eff}}^2}{S^2 N_\phi}, \quad (5.16)$$

where

$$W_{\text{eff}} = \left(\int_0^1 u'(\phi)^2 d\phi \right)^{-1/2} \quad (5.17)$$

is a quantity with units of phase which can be interpreted as an effective pulse width, and $S = \hat{a}/\sigma_n$ is the signal-to-noise ratio of the profile.

5.5 Errors Caused by Shape Differences

Suppose that the profile, $p(\phi)$, differs from the shifted, scaled template by a small amount. That is,

$$p(\phi) = au(\phi - \tau) + r(\phi), \quad (5.18)$$

where the profile residual, $r(\phi)$, is much smaller in magnitude than $au(\phi - \tau)$. The presence of $r(\phi)$ means that the matched-filtering estimate of the phase shift, $\hat{\tau}$, will differ from the true phase shift by a small amount, $\delta\tau$. As long as $\delta\tau$ is small compared to the width of the template, the maximization in equation (5.12) can be carried out analytically, by expanding the condition $\hat{a}'(\tau + \delta\tau) = 0$ to first order in $\delta\tau$ and solving for $\delta\tau$. This gives

$$\delta\tau \approx \frac{\int_0^1 p(\phi)u'(\phi - \tau) d\phi}{\int_0^1 p(\phi)u''(\phi - \tau) d\phi}. \quad (5.19)$$

Making use of equation (5.18), we can rewrite this as

$$\delta\tau \approx \frac{a \int_0^1 u(\phi - \tau)u'(\phi - \tau) d\phi + \int_0^1 r(\phi)u'(\phi - \tau) d\phi}{a \int_0^1 u(\phi - \tau)u''(\phi - \tau) d\phi + \int_0^1 r(\phi)u''(\phi - \tau) d\phi}. \quad (5.20)$$

The first term in the numerator is the integral of a total derivative and so vanishes. Furthermore, since $r(\phi) \ll au(\phi - \tau)$, the second term in the denominator can be neglected compared to the first. The above expression for $\delta\tau$ can therefore be simplified to

$$\delta\tau \approx -\frac{\int_0^1 r(\phi)u'(\phi - \tau) d\phi}{a \int_0^1 u'(\phi - \tau)^2 d\phi}. \quad (5.21)$$

In other words, the error in the matched-filtering estimate of τ is proportional to the projection of $r(\phi)$ onto $u'(\phi - \tau)$.

Importantly, while equation (5.21) is useful for estimating $\delta\tau$ when τ is known, it cannot be used directly to make an improved estimate of τ when only $\hat{\tau}$ is known. In such a case, $r(\tau)$ may be estimated as

$$\hat{r}(\phi) = p(\phi) - \hat{a}(\hat{\tau})u(\phi - \hat{\tau}). \quad (5.22)$$

Solving for $p(\phi)$ and substituting the result into equation (5.19), replacing τ with $\hat{\tau}$ where appropriate, would appear to give the following estimate of $\delta\tau$:

$$\delta\hat{\tau} = -\frac{\int_0^1 \hat{r}(\phi)u'(\phi - \hat{\tau}) d\phi}{\hat{a}(\hat{\tau}) \int_0^1 u'(\phi - \hat{\tau})^2 d\phi}. \quad (5.23)$$

However, as a result of the way it is constructed, $\hat{r}(\phi)$ is approximately orthogonal to $u'(\phi - \hat{\tau})$, as may be determined by expanding the numerator of equation (5.23) to first order in $\delta\tau$. In particular, we have

$$\int_0^1 \hat{r}(\phi) u'(\phi - \hat{\tau}) d\phi = 0 + \mathcal{O}(\delta\tau^2), \quad (5.24)$$

so equation (5.23) becomes $\delta\hat{\tau} = 0 + \mathcal{O}(\delta\tau^2)$. This means that $\delta\hat{\tau}$ is not a useful estimator of $\delta\tau$.

5.5.1 Component Amplitude and Phase Variations

As a demonstration of how shape changes can give rise to TOA estimation errors, consider a model in which the pulse is made up of several components of fixed shape, each of which varies in amplitude and phase. The profile is then given by

$$p(\phi) = \frac{a}{N} \sum_{i=1}^N \sum_{j=1}^{N_c} (1 + b_{ij}) c_j(\phi - \psi_{ij}), \quad (5.25)$$

where N is the number of pulses being averaged, N_c is the number of components, $c_j(\phi)$ is the shape of component j , and b_{ij} and ψ_{ij} are the amplitude and phase offsets, respectively, of component j in pulse i . Without loss of generality, we can take the ensemble average amplitude and phase offsets for each component to be zero. Equation (5.25) can then be expanded to first order in ψ_j , giving

$$p(\phi) \approx \frac{a}{N} \sum_{i=1}^N \sum_{j=1}^{N_c} \left[(1 + b_{ij}) c_j(\phi) - (\psi_{ij} + b_{ij}\psi_{ij}) c'_j(\phi) \right]. \quad (5.26)$$

The template is the normalized ensemble average profile:

$$u(\phi) = \frac{\langle p(\phi) \rangle}{a} \approx \sum_{j=1}^{N_c} c_j(\phi). \quad (5.27)$$

The profile residual, $r(\phi) = p(\phi) - au(\phi)$, is therefore

$$r(\phi) \approx a \sum_{j=1}^{N_c} \left[\bar{b}_j c_j(\phi) + (\bar{\psi}_j + \bar{b}_j \bar{\psi}_j) c'_j(\phi) \right], \quad (5.28)$$

where barred quantities represent sample averages. The first term in equation (5.28) arises from the difference in amplitude between the profile and the template, and the second from the difference in phase. Substituting this into equation (5.21) gives

$$\delta\tau \approx \frac{\sum_{j=1}^{N_c} \left[(\bar{\psi}_j + \bar{b}_j \bar{\psi}_j) \Gamma_j + \bar{b}_j \Delta_j \right]}{\sum_{j=1}^{N_c} \Gamma_j}, \quad (5.29)$$

where we have made use of the symbols

$$\Gamma_j = \sum_{k=1}^{N_c} \int_0^1 c'_j(\phi) c'_k(\phi) d\phi \quad \text{and} \quad (5.30)$$

$$\Delta_j = \sum_{k=1}^{N_c} \int_0^1 c_j(\phi) c'_k(\phi) d\phi \quad (5.31)$$

for cross-correlations between the component shapes, and ignored all but the leading term in the denominator. When there is only one component, equation (5.29) reduces to

$$\hat{\tau} \approx \bar{\psi} + \bar{b}\bar{\psi}. \quad (5.32)$$

The variance of the TOA estimate is therefore

$$\sigma_{\tau}^2 = \langle \hat{\tau}^2 \rangle = \frac{1+m^2}{N} \langle \psi^2 \rangle, \quad (5.33)$$

where $m^2 = \langle b^2 \rangle$ is the square of the single-pulse modulation index. This shows that, for a single-component pulse, phase variation alone can produce TOA errors. Amplitude variation may enhance TOA errors created by phase variation, but cannot produce them on its own.

For an arbitrary number of components, we can also derive the variance of the TOA estimate from equation (5.29). This gives

$$\sigma_\tau^2 = \frac{\sum_{j=1}^{N_c} \left[(1 + m_j^2) \langle \psi_j^2 \rangle \Gamma_j^2 + m_j^2 \Delta_j^2 \right]}{N \left(\sum_{j=1}^{N_c} \Gamma_j \right)^2}, \quad (5.34)$$

where $m_j^2 = \langle b_j^2 \rangle$ is the square of the single-pulse modulation index for component j . Unlike in the single component case, here amplitude variation can create TOA errors on its own, even in the absence of phase variation: setting $\langle \psi_j \rangle^2 = 0$ in equation (5.34) leaves the term involving Δ_j^2 . Conceptually, this is because amplitude variations represent true shape changes only when there are multiple components with separately varying amplitudes. For single component pulses, amplitude variations simply scale the pulse, and can be absorbed into the phase factor, a , but for pulses with multiple components that vary separately, this is no longer the case.

5.5.2 Simulations with Gaussian Components

We performed a number of simulations in which shape variations were introduced into a simple pulse shape model consisting of Gaussian components. In each simulation, N_p profiles were generated by averaging N simulated pulses each, in the manner of equation (5.25). The shape of each component was taken to be Gaussian, with amplitude a_j and width w_j :

$$c_j(\phi) = a_j \exp \left[-\frac{(\phi - \mu_j)^2}{2w_j^2} \right]. \quad (5.35)$$

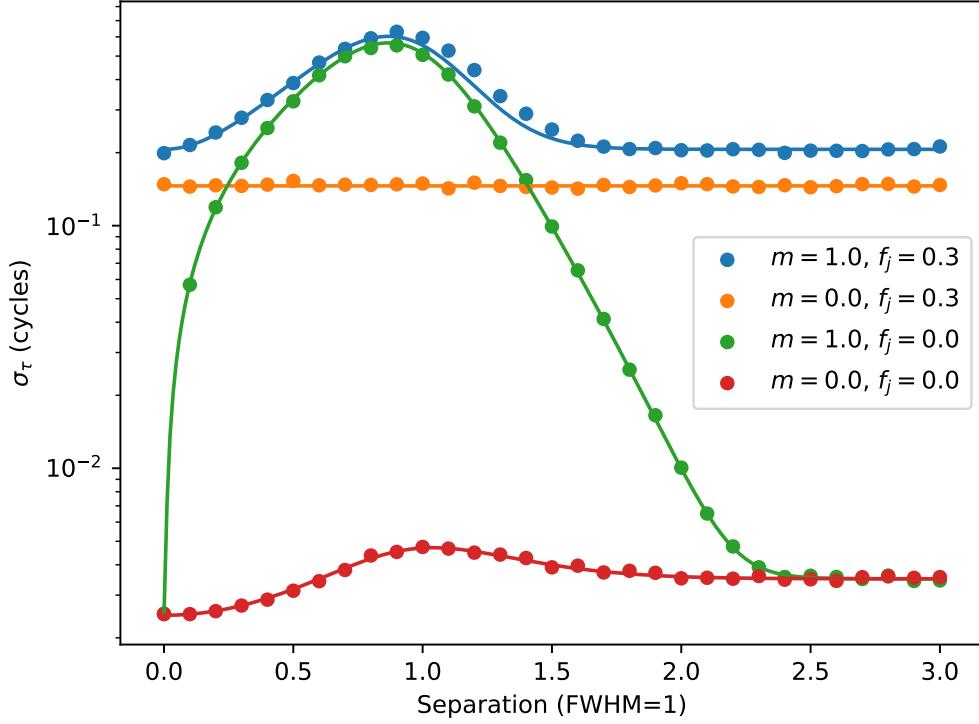


Figure 5.1: The standard deviation of TOA estimates in simulated profiles with two identical components with varying spacing. Crosses indicate results of simulated cases, and solid lines indicate predictions based on equations (5.16) and (5.40). Four cases are shown: one in which the pulses have amplitude and phase variations ($m = 1.0, f_j = 0.3$ for each component), one in which the pulses have only phase variations ($m = 0.0, f_j = 0.3$), one in which they have only amplitude variations ($m = 1.0, f_j = 0.3$), and a reference case where the pulses are copies of the template with additive white noise ($m = 0.0, f_j = 0.0$). Each point was calculated based on $N_p = 2048$ profiles, each the average of $N = 1000$ pulses. The average profiles had a signal-to-noise ratio $S = 1000$.

With this choice of component shape, we have

$$\Gamma_j = \sqrt{2\pi} a_j w_j \sum_{k=1}^{N_c} \frac{a_k w_k}{\left(w_j^2 + w_k^2\right)^{\frac{3}{2}}} \left[1 - \frac{(\mu_j - \mu_k)^2}{w_j^2 + w_k^2} \right] \exp \left[-\frac{(\mu_j - \mu_k)^2}{2(w_j^2 + w_k^2)} \right] \quad (5.36)$$

$$\Delta_j = \sqrt{2\pi} a_j w_j \sum_{k=1}^{N_c} \frac{a_k w_k (\mu_j - \mu_k)}{\left(w_j^2 + w_k^2\right)^{\frac{3}{2}}} \exp \left[-\frac{(\mu_j - \mu_k)^2}{2(w_j^2 + w_k^2)} \right]. \quad (5.37)$$

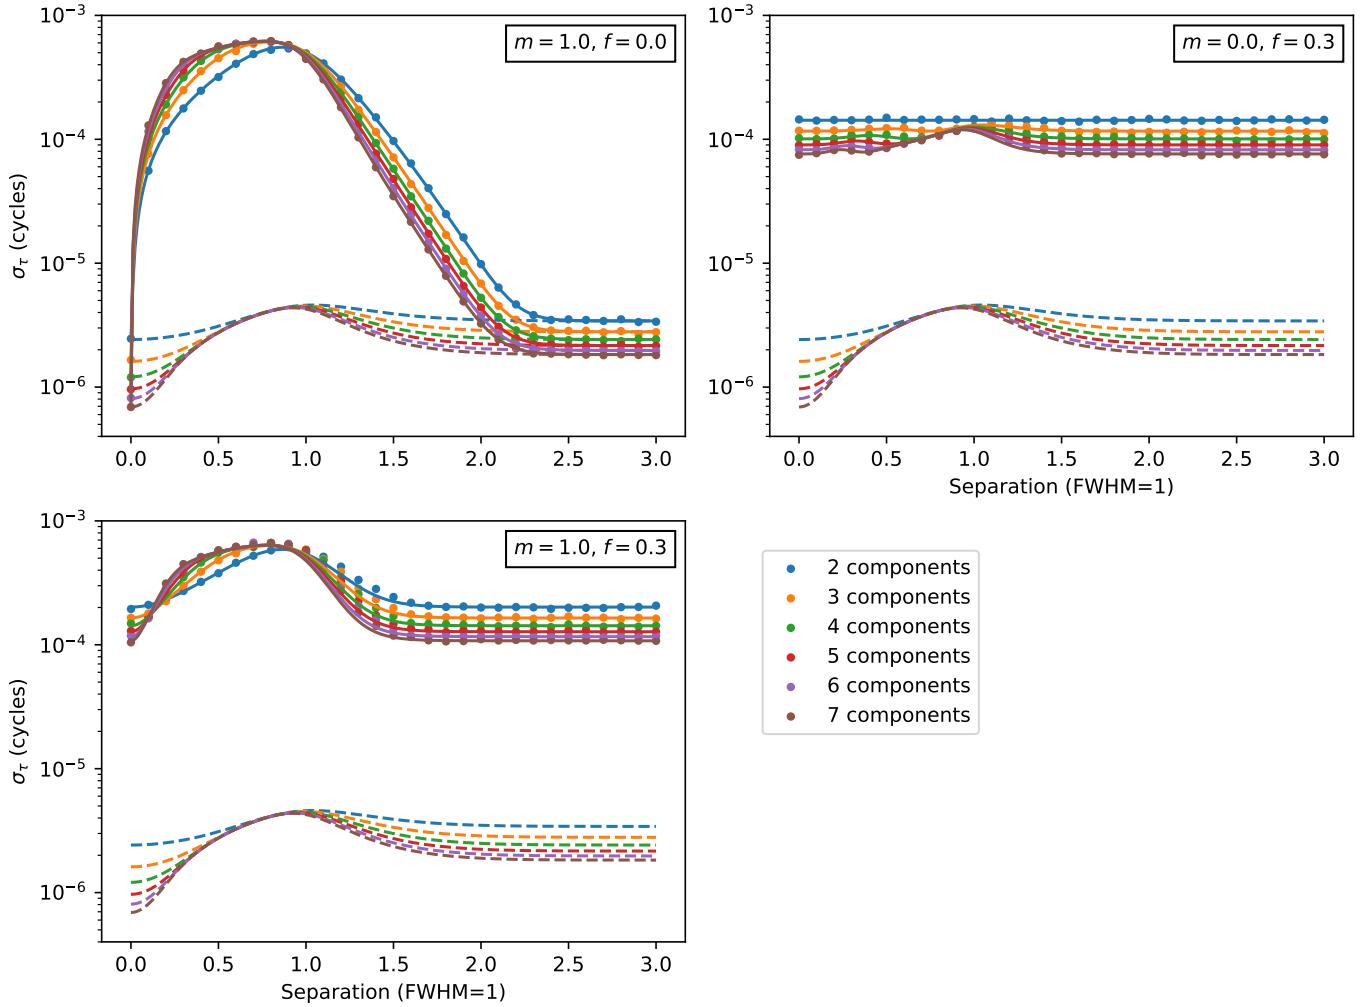


Figure 5.2: The standard deviation of TOA estimates in simulated profiles with several identical components with varying spacing. Results with amplitude variations only ($m = 1.0$) is shown in the upper left, results with phase variations only ($f_j = 0.3$) in the upper right, and results with both amplitude and phase variations ($m = 1.0, f_j = 0.3$) in the lower left. As in Figure 5.1, filled circles indicate simulation results, and solid lines indicate predictions based on equation (5.34). For comparison, the predictions for the reference case with no amplitude or phase variations (based on equation 5.16) are shown as dashed lines.

Amplitudes $(1 + b_{ij})$ were drawn from a log-normal distribution with modulation index m_j , and phase offsets (ψ_{ij}) were drawn from a normal distribution with standard deviation $\langle \psi_j^2 \rangle = f_j w_j$. The ratio $f_j = \langle \psi_j^2 \rangle / w_j$, where w_j , as here, is the single-pulse width of a particular component, is called the jitter parameter. Below, we will distinguish between different simulated cases by specifying the amplitudes and widths of each component, the separation between components, and the modulation index, m_j and jitter parameter, f_j , of each component.

As an initial comparision of simulations with theoretical results, consider the case where pulses are made up of two identical components. In this case, we have $\Gamma_1 = \Gamma_2 = \Gamma$ and $\Delta_1 = -\Delta_2 = \Delta$, where

$$\Gamma = \frac{\sqrt{\pi}a^2}{2w} \left[1 + \left(1 - \frac{\delta^2}{2w^2} \right) \exp\left(-\frac{\delta^2}{4w^2}\right) \right], \quad (5.38)$$

$$\Delta = \frac{\sqrt{\pi}a^2\delta}{2w} \exp\left(-\frac{\delta^2}{4w^2}\right), \quad (5.39)$$

and $\delta = \mu_1 - \mu_2$ is the separation between the components. Equation (5.34) then predicts that the variance of the TOA estimates will be given by

$$\begin{aligned} \sigma_\tau^2 &= \frac{(1 + m^2)\langle \psi^2 \rangle}{2N} + \frac{m^2\Delta^2}{2N\Gamma^2} \\ &= \frac{(1 + m^2)f^2w^2}{2N} + \frac{m^2\delta^2}{2N} \left[\exp\left(\frac{\delta^2}{4w^2}\right) + \left(1 - \frac{\delta^2}{2w^2}\right) \right]^{-2}. \end{aligned} \quad (5.40)$$

Figure 5.1 compares the predictions of equation (5.40) with simulation results for various values of δ in three different cases: one with only amplitude variations, another with only phase variations, and a third with both types of variation. With only phase variations, σ_τ is independent of δ , while in cases where phase variations are included, it reaches a maximum when $\delta \approx w$. For reference, a fourth case with neither amplitude nor phase variations, but only a small amount of additive white noise, is shown. In this case, σ_τ is predicted by

equation (5.16). It has some dependence on δ because W_{eff} is a function of the template shape (equation 5.17).

A more extensive set of simulation results is shown in Figure 5.2. Shown there are the results of simulations with various numbers of identical components (between 2 and 7), in each of the three primary cases from Figure 5.1, along with the predictions of equation (5.34). One notable trend visible in Figure 5.2 is that, for large separations ($\delta \gg w$), σ_τ decreases as the number of components, N_c , is increased. Indeed, it is approximately proportional to $N_c^{-1/2}$. This happens because, for large values of δ , the cross-correlation Δ_j rapidly approaches 0, as do all terms in equation (5.30) for Γ_j except the $j = k$ term, which is independent of δ . Using equation (5.36), we can see that, for the identical Gaussian components used in the simulations, Γ_j approaches $\sqrt{\pi}a^2/(2w)$ for each component. It follows that, in the limit of large δ , we have

$$\sigma_\tau^2 = \frac{(1 + m^2)f^2w^2}{NN_c}. \quad (5.41)$$

A similar phenomenon occurs when the components are not identical, complicated only by the fact that the components do not contribute equally to the variance. In general, for sufficiently well-separated components, we have

$$\sigma_\tau^2 = \frac{\sum_{j=1}^{N_c} a_j^4(1 + m_j^2)f_j^2}{N(\sum_{j=1}^{N_c} a_j^2w_j^{-1})^2}. \quad (5.42)$$

In other words, the TOA estimation errors associated with each component combine in a manner weighted by the combination $a_j^2w_j^{-1}$, and tending to decrease as the number of components increases. However, the trend of decreasing σ_τ with increasing N_c is not universal. As seen in Figure 5.2, in some cases with amplitude variations included, when $\delta \lesssim w$, σ_τ actually increases with increasing N_c at a similar rate. In combination with the fact that σ_τ also de-

pends significantly on the modulation index, m , and jitter parameter, f , this means that it is generally not possible to predict σ_τ from the number of pulse components without further information.

5.6 Tools for Assessing Single-Pulse Stochasticity

Several statistical tools can be used to identify the effects of single-pulse stochasticity in pulsar observations, and assess the degree to which these effects influence TOA estimates based on the observations. One can estimate the effect of single-pulse stochasticity on timing by comparing the standard deviation of TOA residuals to the signal-to-noise ratio and number of pulses averaged in the profiles used to generate them. The phase autocorrelation functions of single pulses and average profiles can be used to establish that their widths differ. Correlations between TOAs in adjacent frequency bands can be used to separate timing noise caused by single-pulse stochasticity from other forms of noise. And principal component analysis (PCA) allows the dominant modes of shape variation to be identified. Below, we use each of these tools to compare the results of simulations to observations of the Vela pulsar, PSR B0833-45. First, however, we will describe the particular simulations and observations under consideration.

5.6.1 TOA and Mismatch Statistics

The TOA error caused by single-pulse stochasticity (so-called “jitter noise”) dominates over error from radiometer noise only when the signal-to-noise ratio

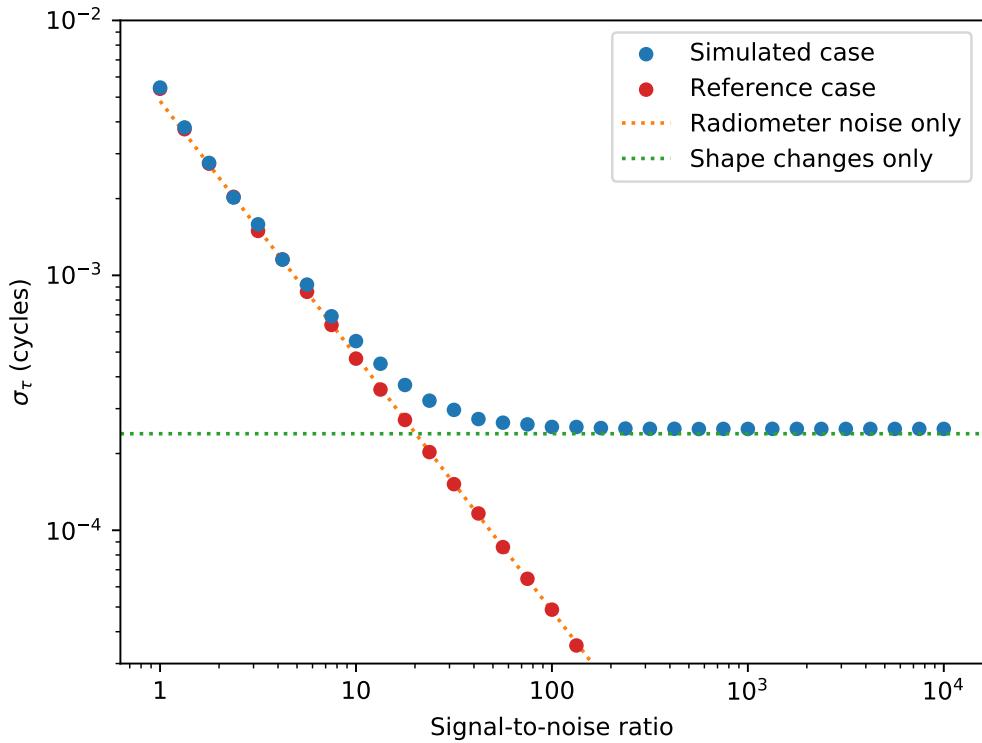


Figure 5.3: The standard deviation of TOA estimates in simulated pulses with amplitude and phase jitter and varying signal-to-noise ratio. Each point was calculated based on $N_p = 2048$ simulated profiles, each the average of $N = 1000$ simulated pulses. The dashed line shows the predicted TOA error based on equation (5.16) alone, ignoring single-pulse stochasticity, and the dotted line shows the predicted TOA error from single-pulse stochasticity, based on equation (5.34).

is large. One can therefore measure its magnitude by comparing the standard deviation of a set of TOA residuals to the signal-to-noise ratio, S , of the profiles used to produce them, as shown in Figure 5.3. At lower signal-to-noise ratios, the standard deviation of the TOA residuals, σ_τ , will be approximately proportional to $S^{-1/2}$, a consequence of equation (5.16), whereas at higher signal-to-noise ratios, σ_τ will be approximately independent of S , approaching a level determined by single-pulse stochasticity alone. Taking advantage of variability due to interstellar scintillation, Lam et al. (2016) used this technique to mea-

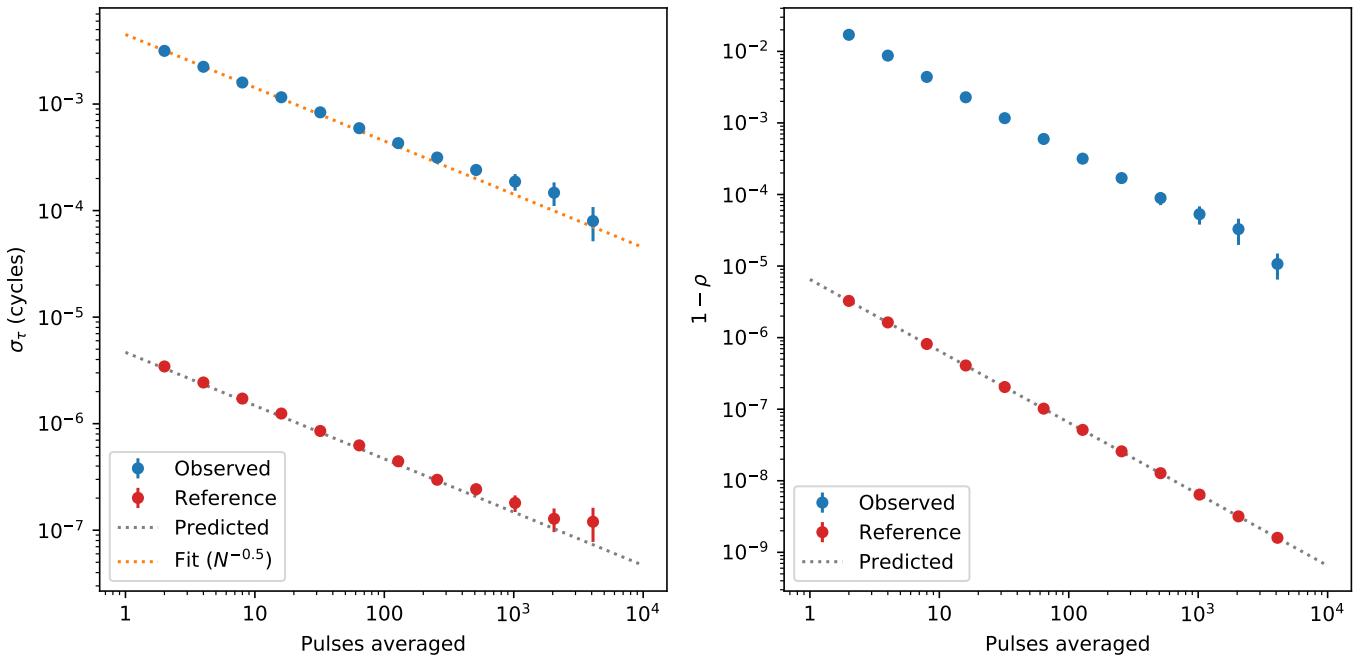


Figure 5.4: The standard deviation of TOA residuals (left) and profile-template mismatch (right) for average profiles constructed from various numbers of simulated pulses with amplitude and phase variations.

sure the magnitude of jitter noise for a number of millisecond pulsars timed by the NANOGrav collaboration. A later study (Lam et al., 2019) extended this analysis to measure the frequency-dependence of jitter noise in NANOGrav pulsars.

If the shapes of successive pulses are uncorrelated, the magnitude of jitter noise in an average profile should be proportional to $N^{-1/2}$, where N is the number of pulses averaged. Error from radiometer noise alone depends on N in the same way. As a result, the relative contribution of jitter noise to TOA errors is determined by the single-pulse signal-to-noise ratio alone – averaging a larger number of pulses does not change it. Figure 5.4 demonstrates this using simulated single-pulse data.

Figure 5.4 also shows the so-called “mismatch” between the template and

average profiles for several values of N . The mismatch, \mathcal{M} , between a template, $u(\phi)$, and a profile, $p(\phi)$, is given by $\mathcal{M} = 1 - \rho$, where

$$\rho = \frac{\int_0^1 p(\phi)u(\phi - \hat{\tau}) d\phi}{\sqrt{\int_0^1 p(\phi)^2 d\phi \int_0^1 u(\phi)^2 d\phi}} \quad (5.43)$$

is the correlation coefficient between the (shifted) template and the profile, and $\hat{\tau}$ is the matched-filtering TOA estimate. Defining the estimated residual, $\hat{r}(\phi)$, as in equation (5.22), we have

$$\mathcal{M} \approx \frac{\int_0^1 \hat{r}(\phi)^2 d\phi}{2 \int_0^1 u(\phi)^2 d\phi}. \quad (5.44)$$

Since $\hat{r}(\phi)$ is proportional to $N^{-1/2}$, the mismatch is proportional to N^{-1} . Comparing the TOA errors and mismatch values for average profiles created from different numbers of single pulses to expectations for profiles differing from the average only by additive white noise, as is done in Figure 5.8, demonstrates the relative size of the contribution from single-pulse stochasticity, as well as the fact that it depends on N in the same way as the contribution from additive white noise.

5.6.2 Profile Autocorrelation Functions

Another way to characterize single-pulse stochasticity is in terms of the phase autocorrelation function (ACF), $R(\tau)$, of a profile, $p(\phi)$, defined as

$$R(\tau) = \frac{1}{\langle p^2 \rangle} \int_0^1 p(\phi)p(\phi - \tau) dt. \quad (5.45)$$

Here, as in Section 5.4, we define $\langle p^2 \rangle = \int_0^1 p(\phi)^2 d\phi$. The ACF is particularly useful as a measure of pulse shape because it does not depend on how the profile is aligned in phase. Its width, which is a measure of the profile width,

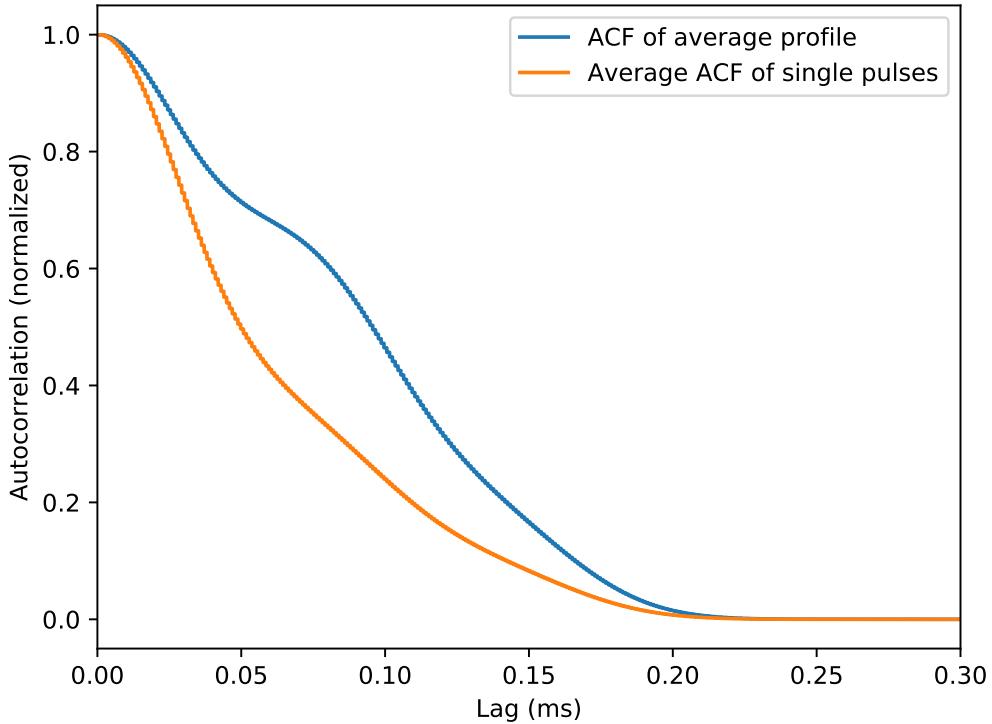


Figure 5.5: Comparison of single-pulse and average-profile autocorrelation functions for simulated data.

tends to be larger for average profiles than for single pulses. This is a consequence of single-pulse stochasticity: the average of many individual narrow pulses as slightly different phases is a broader pulse. This is illustrated using simulated data in figure 5.5.

For a single-component pulse with RMS width w which is subject to phase jitter with standard deviation fw , the RMS width of the single-pulse ACF will be $\sqrt{2}w$, while the RMS width of the average-profile ACF will be $\sqrt{2(1 + f^2)}w$. As a result, one can estimate the jitter parameter, f , as

$$f = \sqrt{\frac{W_{\text{avg}}^2}{W_{\text{sp}}^2} - 1}, \quad (5.46)$$

where W_{sp} and W_{avg} are the RMS widths of the single-pulse and average-profile

ACFs, respectively.

5.6.3 Correlations Between Frequency Channels

Because pulse-to-pulse shape changes are broadband, the associated TOA estimation errors are strongly correlated between nearby frequency channels. If TOAs are calculated for many frequency channels within the same observation, this characterization can serve as another way to separate TOA estimation errors from single-pulse stochasticity from other sources of noise, and thereby estimate their magnitude. Such estimates can be used to inform pulsar timing models. For example, Arzoumanian et al. (2014) and subsequent NANOGrav analyses (e.g. Alam et al., 2021a) employ a model of the form

$$C_{\nu\nu',tt'} = \delta_{tt'} \left[\delta_{\nu\nu'} \left(\mathcal{Q}^2 + \mathcal{E}^2 \sigma_r^2 \right) + \mathcal{J}^2 \right] \quad (5.47)$$

for the covariance between TOA residuals at times t and t' and frequencies ν and ν' . Here, σ_r is the estimated TOA uncertainty based on radiometer noise alone (equation 5.15), and the noise parameters \mathcal{Q} , \mathcal{E} , and \mathcal{J} , representing, respectively, additional white noise added in quadrature, a scaling factor, and a white noise component correlated in frequency but uncorrelated in time, are inferred alongside the timing model parameters. The inclusion of the correlated-noise parameter, \mathcal{J} , allows for TOA estimation errors of the sort described here to be taken into account. However, estimates of \mathcal{J} , which are based on TOAs alone, represent a limited view, and are best understood in combination with one or more of the other metrics presented here.

5.6.4 Principal Component Analysis

A more direct way of analyzing variations in the shapes of the profiles is to perform principal component analysis (PCA) on the profile residuals. PCA finds the set of orthogonal and statistically independent modes of variation which most effectively captures the variance in the data. Specifically, it results in a set of k principal components, $v_i(\phi)$, and associated variances, σ_i^2 , which allow the profile residuals to be represented as

$$r(\phi) = \sum_{i=1}^K x_i v_i(\phi), \quad (5.48)$$

where the coefficients x_i are statistically independent and have variance σ_i^2 :

$$\langle x_i x_j \rangle = \sigma_i^2 \delta_{ij} \quad (5.49)$$

and the principal components $v_i(\phi)$ are orthonormal:

$$\int_0^1 v_i(\phi) v_j(\phi) d\phi = \delta_{ij}. \quad (5.50)$$

The principal components are constructed from the data in a way that maximizes the variance “explained” by the first few: the first principal component represents the direction in which the data variance is largest, the second represents the direction, orthogonal to the first, in which the remaining variance is the largest, and so on.

One can find the principal components of a set of profile residuals, $r_i(\phi)$, by calculating the singular value decomposition (SVD) of the $N_\phi \times N$ matrix \mathbf{R} with the residuals as columns, i.e., with entries $R_{ij} = r_j(\phi_i)$, where ϕ_i is the phase value corresponding to phase bin i . The SVD gives the factorization

$$\mathbf{R} = \sqrt{N} \mathbf{U} \mathbf{S} \mathbf{V}^\top, \quad (5.51)$$

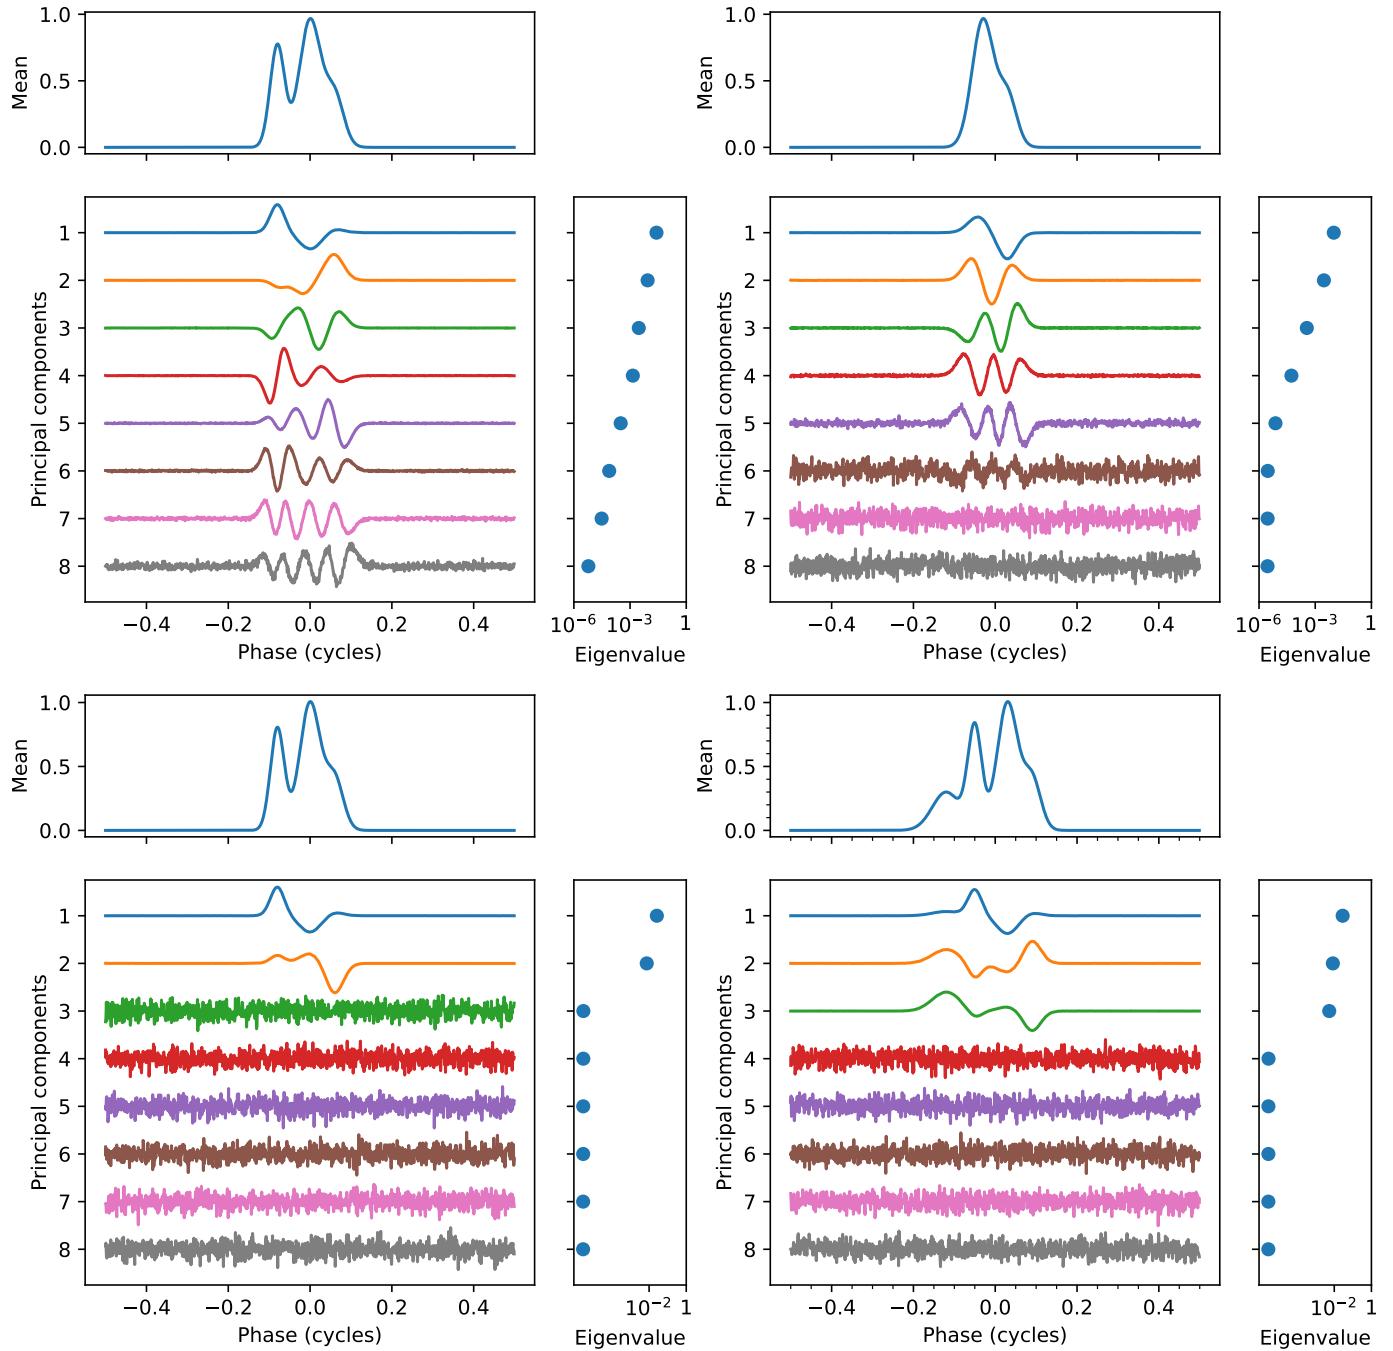


Figure 5.6: Principal components in various simulated cases. The upper two panels show cases with both amplitude and phase variations enabled, while the lower two show cases with only amplitude variations. The latter panels demonstrate that the number of principal components in this case is one less than the number of pulse components, while in the former, additional principal components can be seen.

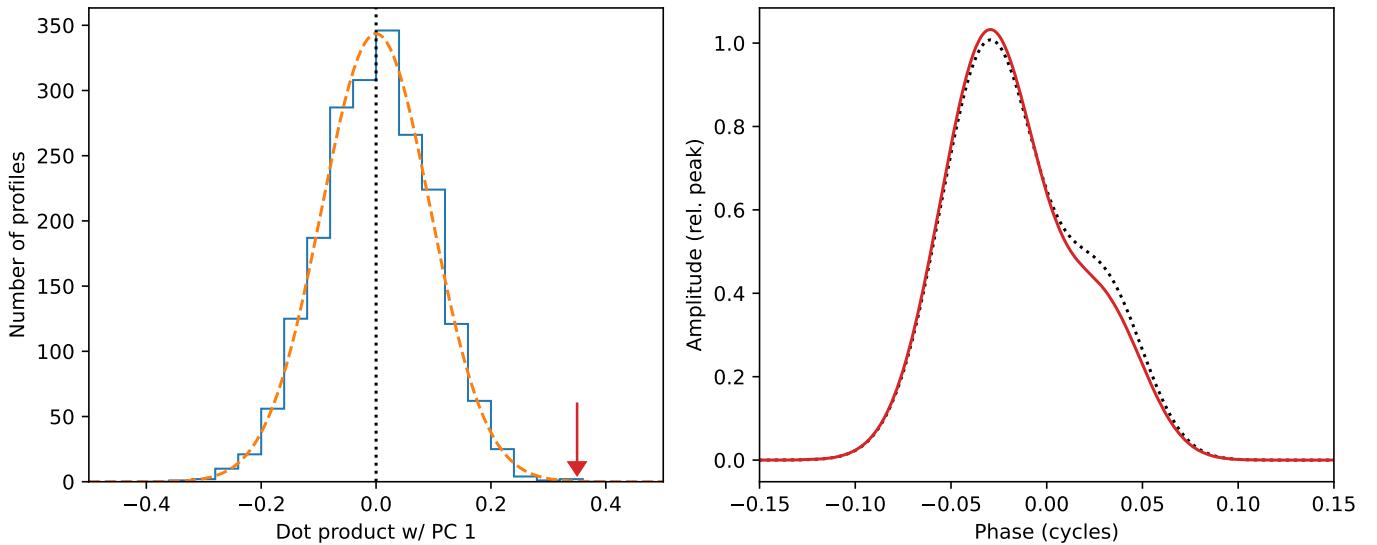


Figure 5.7: The correspondence between shape variations and principal component dot products. The left panel shows a histogram of the dot product x_1 with the first principal component for $N_p = 2048$ simulated profiles, each constructed from $N = 1000$ simulated pulses. The dashed orange line gives the expected distribution of dot product values, whose width is the square root of the eigenvalue, σ_1^2 . The right panel shows a profile (red) corresponding to a large value of the dot product, indicated by an arrow in the left panel. The template shape (dotted black line) is included for comparison. The data shown is for a set of simulated profiles with two Gaussian components, each with modulation index $m = 1$ and no phase jitter ($f = 0$). The leading component, centered at $\phi = -0.03$, has a FWHM of 0.06 cycles and a mean amplitude of 1.0, while the trailing component, at $\phi = 0.03$, has a FWHM of 0.05 cycles and a mean amplitude of 0.4.

where N is the number of residuals, $\mathbf{S} = \text{diag}(\sigma_i)$ is a diagonal matrix whose nonzero entries are the square roots of the variances, σ_i^2 , and \mathbf{U} and \mathbf{V} are orthogonal matrices, i.e., $\mathbf{U}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{V} = \mathbf{I}$. The total number, K , of nonzero variances, and thus of principal components, is equal to the rank of the matrix \mathbf{R} , which is almost certainly equal to whichever of N or N_ϕ is smaller. The columns of \mathbf{V} are the principal components, $v_i(\phi)$, i.e., its entries are $V_{ij} = v_j(\phi_i)$. Equivalently, one may also compute the eigenvalue decomposition of

the covariance matrix of the residuals,

$$\mathbf{C} = \frac{1}{N} \mathbf{R} \mathbf{R}^T, \quad (5.52)$$

which, using equation 5.51, is given by

$$\mathbf{C} = \mathbf{V} \mathbf{S}^2 \mathbf{V}^T. \quad (5.53)$$

This is a somewhat less numerically stable computation, but may be faster if the number of residuals, N , is significantly greater than the number of phase bins, N_ϕ . It also reveals an important property of the principal components. If the pulses from a pulsar are stationary at fixed pulse phase (which they are generally assumed to be, in the absence of mode changes), the covariance matrix in equation (5.52) is a stable property of the pulsar, and so the principal components are as well.

In simple cases, the shape and number of the significant principal components is directly related to the type of variation present in the data. Consider the case of simulated pulses with independent amplitude variations in each of several components, but no other shape variations. In this case, the significant principal components, together with the average pulse shape, form a basis for the space spanned by the components of the profile. Since the dimension of this space is equal to the number of profile components, the number of significant principal components is one less than the number of profile components.

For a single-component pulse with only phase variations, the profile is given by

$$p(\phi) = \frac{a}{N} \sum_{i=1}^N c(\phi - \psi_i), \quad (5.54)$$

which is a special case of equation (5.25). Expanding this in a power series in

the phase variations, ψ_i , shows that the residual is given by

$$r(\phi) = a\bar{\psi}c'(\phi) + \mathcal{O}(\psi_i^2), \quad (5.55)$$

where $\bar{\psi}$ is the mean of the phase variations. The dominant principal component in this case, then, should be proportional to $c'(\phi)$. This component is most dominant when $\bar{\psi}$ is much less than the width of the pulse; i.e., $f_j \ll 1$. For larger values of f_j , the higher-order terms in equation (5.55) become increasingly important, and contribute additional significant principal components related to higher derivatives of the component shape $c(\phi)$. These tend to have larger numbers of zero crossings as they become less significant, which is primarily results from the fact that each new principal component must be orthogonal to all of the previous ones.

Principal components for several simulated cases are shown in Figure 5.6, demonstrating the relationships described above: in the absence of phase variations, the number of significant principal components is one less than the number of independently varying components in the profile, and in the presence of phase variations, additional significant principal components are present, with the less-significant principal components having a larger number of zero crossings.

PCA is a very sensitive probe of subtle variations in large data sets. This means that isolating the characteristic principal components for a given pulsar requires a very clean data set, from which RFI has been thoroughly excised. It also means that PCA can be used as a tool for identifying low-level RFI in places where it might not otherwise be obvious, an idea which will be explored further below. An important tool in this sort of analysis is the set of principal component dot products. Given a set of residuals, $r_i(\phi)$, the principal compo-

nent dot products are the numbers

$$x_{ij} = \int_0^1 r_i(\phi) v_j(\phi) d\phi. \quad (5.56)$$

Since the principal components form an orthonormal basis for the space spanned by the residuals, these can be used to expand the residuals:

$$r_i(\phi) = \sum_{j=1}^k x_{ij} v_j(\phi). \quad (5.57)$$

In other words, x_{ij} quantifies the extent to which the principal component $v_j(\phi)$ contributes to the residual $r_i(\phi)$. By examining the dot products x_{ij} for a fixed principal component $v_j(\phi)$, one can find the profiles it contributes to the most, which is particularly useful if $v_j(\phi)$ represents an identifiable feature in the data.

5.6.5 Application to the Vela Pulsar

The Vela pulsar, PSR B0833–45, is one of the brightest radio pulsars. This makes it a particularly good test case for the study of single-pulse stochasticity, since it is possible to observe single pulses at a relatively high signal-to-noise ratio. Vela is a young pulsar, with a period $P = 89.3$ ms and a spindown rate $\dot{P} = 1.25 \times 10^{-13}$, giving it a characteristic age $P/(2\dot{P}) = 11\,300$ yr. Its period falls between those of millisecond pulsars and those of most canonical pulsars. As in the case of the Crab pulsar, this is attributable to its relative youth. Its significance as a particularly bright and young pulsar has made Vela one of the best-studied of all pulsars, and several authors have previously described the behavior of its single pulses. In particular, Krishnamohan & Downs (1983) discovered that there is a strong correlation between the intensity of single

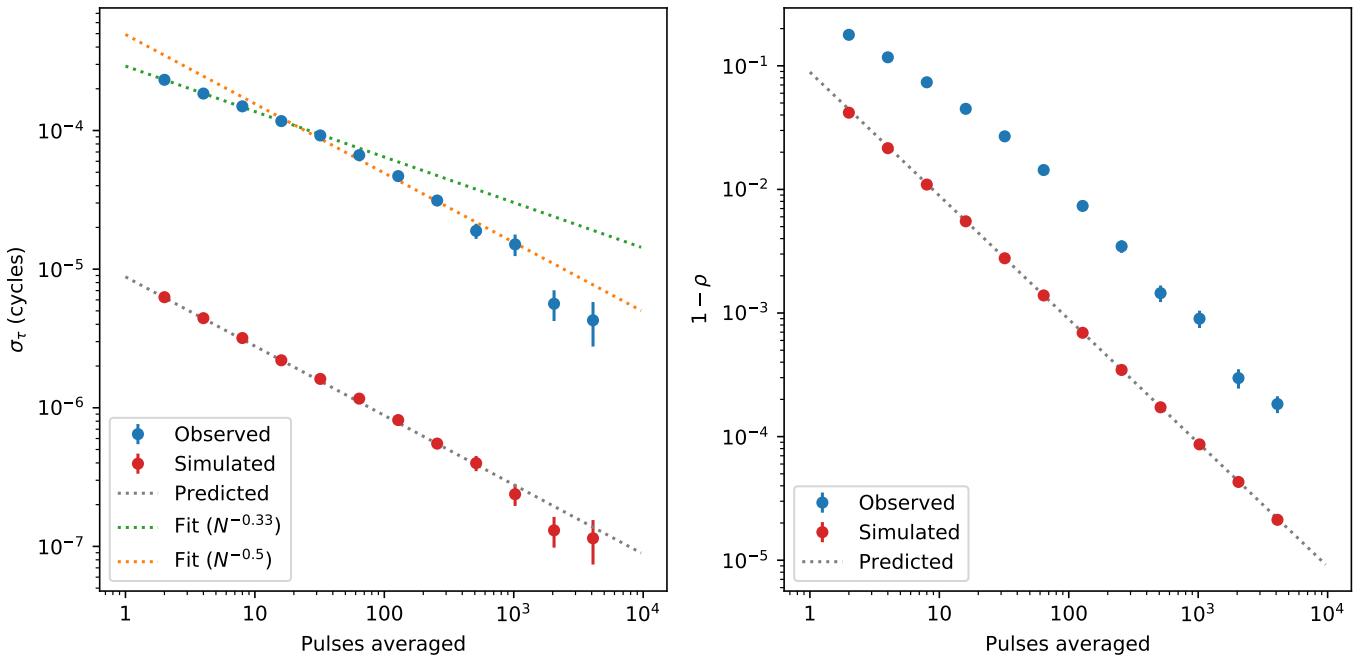


Figure 5.8: The standard deviation of TOA residuals (left) and profile-template mismatch (right) for average profiles constructed from various numbers of single pulses from the Vela pulsar. The results seen here are derived from $2^{14} = 16384$ consecutive single pulses, representing approximately 23 minutes of observing time. For each power-of-two value of N between 2 and 4096, profiles were constructed from groups of N consecutive pulses and the corresponding TOAs and mismatch values computed.

pulses from Vela and their arrival times, with brighter pulses arriving earlier. More recently, Kerr (2015) constructed a basis set of pulse shapes by grouping single pulses from Vela with similar shapes, showing that they could be used to partially correct for the TOA errors associated with shape variation.

Here, we characterize pulse shape changes in the Vela pulsar using a set of single pulses observed with the Parkes telescope on May 4, 2016. The 28-minute observation was made using the 10 cm receiver on the Parkes telescope at a center frequency of 2820 MHz, and processed using the CASPSR backend. Baseband voltages from two orthogonal, linearly polarized receivers were sampled at 800 million samples per second, and recorded in 204 largely con-

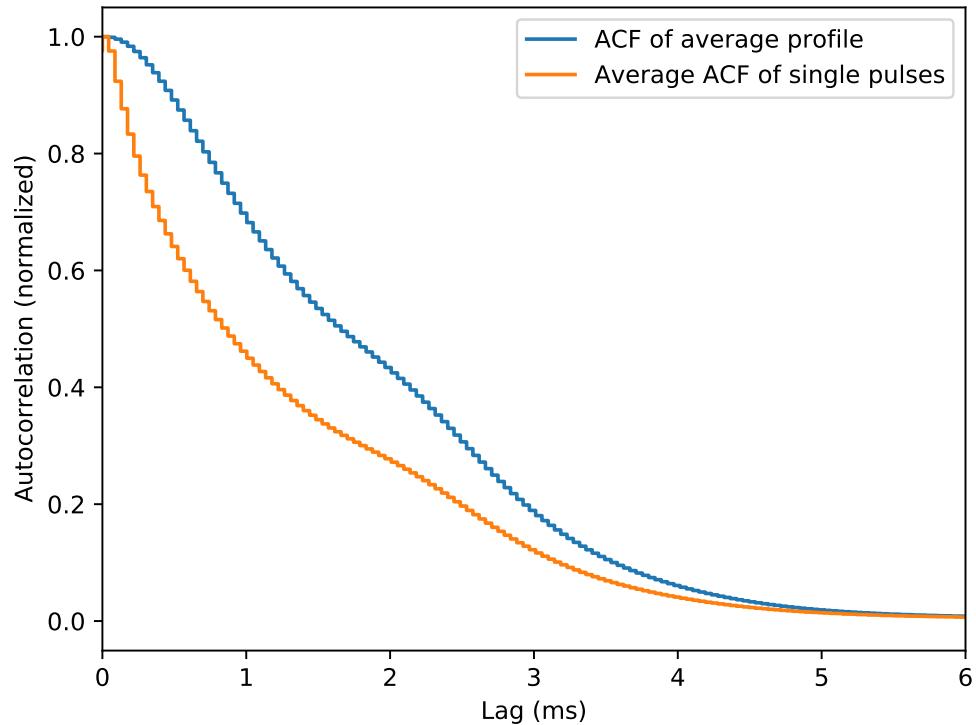


Figure 5.9: Comparison of single-pulse and average-profile autocorrelation functions for Vela pulsar data. A clear difference in width is visible, demonstrating the broadening effect of single-pulse stochasticity. Both autocorrelation functions were calculated from the same May 4, 2016 observation, consisting of 18254 single pulses. An autocovariance function was computed for each single pulse, and the results were averaged to produce the single-pulse autocovariance function, which was then corrected by subtracting an estimate of the contribution of radiometer noise to the covariance at zero lag and normalized to produce the single-pulse autocorrelation function seen here. The pulses were also combined to produce an average profile, from which the average-profile autocorrelation function was computed.

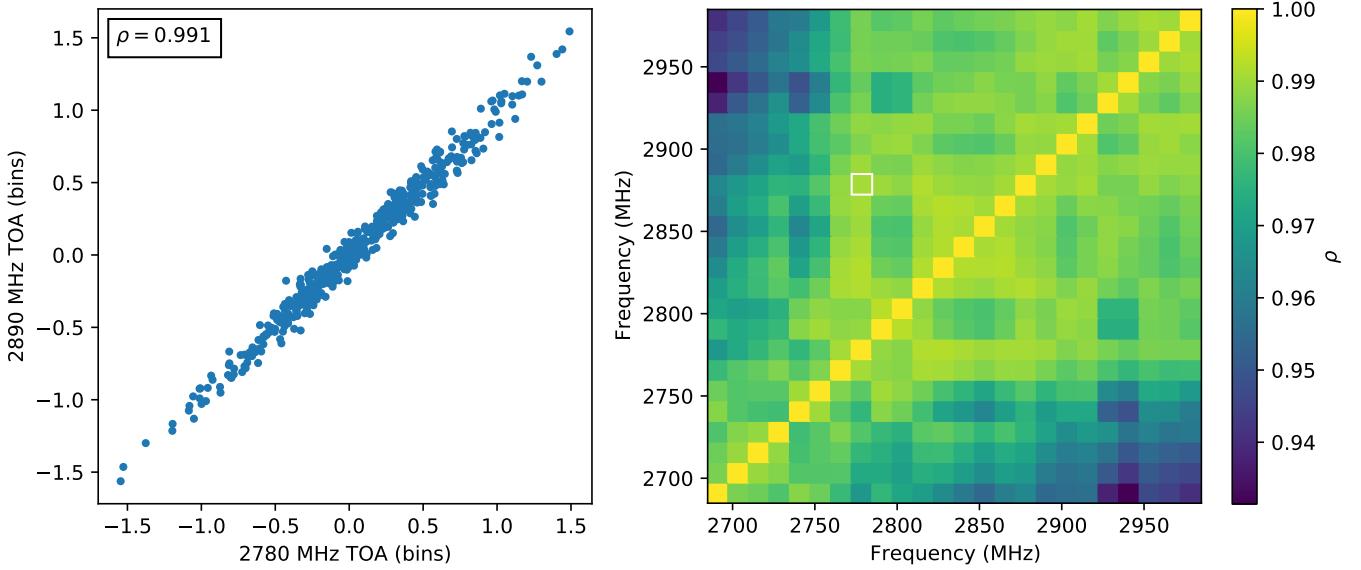


Figure 5.10: Correlations between TOAs in nearby frequency channels for the Vela pulsar. Here the observing band was divided into 24 channels, and a correlation coefficient was calculated between the TOA estimates each pair of channels. The TOAs for a single pair of channels are shown in the left panel. The right panel shows correlations between all channel pairs, with a small white square identifying the pair used in the left panel. Notably, even the weakest correlation here is still very strong, at 0.93, indicating that the single pulses decorrelate only marginally, even over 300 MHz.

tiguous 8-second blocks. We synthesized coherently dedispersed filterbanks in all four Stokes parameters from the baseband voltage recordings using DSPSR (van Straten & Bailes, 2011), producing spectra for each of the 18 254 single pulses in 512 frequency channels and 2048 phase bins. The spectra were then filtered to remove the bandpass edges, leaving 300 MHz of usable bandwidth between 2685 and 2985 MHz, and integrated over frequency and polarization to produce a set of frequency-averaged total-intensity profiles.

The standard deviation of the residuals and profile-template mismatch are plotted as a function of the number of pulses in each average in Figure 5.8. The results resemble the simulated data shown in Figure 5.4, with one notable

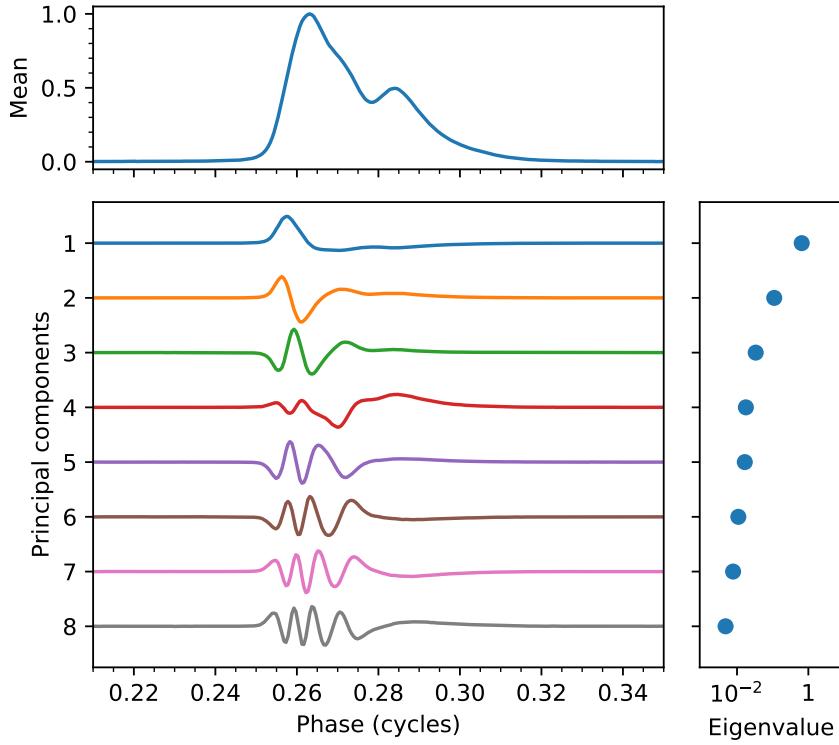


Figure 5.11: Principal components derived from single pulses from a May 4, 2016 observation of the Vela pulsar.

exception: the $N^{-1/2}$ power law fails to hold for low values of N , where fitting a broken power-law model to the data reveals that they are approximately proportional to $N^{-0.34}$. This is likely a result of small but measurable correlations between successive pulses.

The single-pulse and average-profile ACFs for the Vela data are shown in Figure 5.9. In comparing the single-pulse ACF to the average-profile ACF, it is necessary to account for the effect of additive white noise, which makes a contribution to $R(0)$ proportional to its variance. In Figure 5.5, this contribution has been estimated based on the variance of the noise in the off-pulse region, and subtracted prior to normalization. This allows for a more direct comparison between the two ACFs. Estimating f from the Vela pulsar data shown in

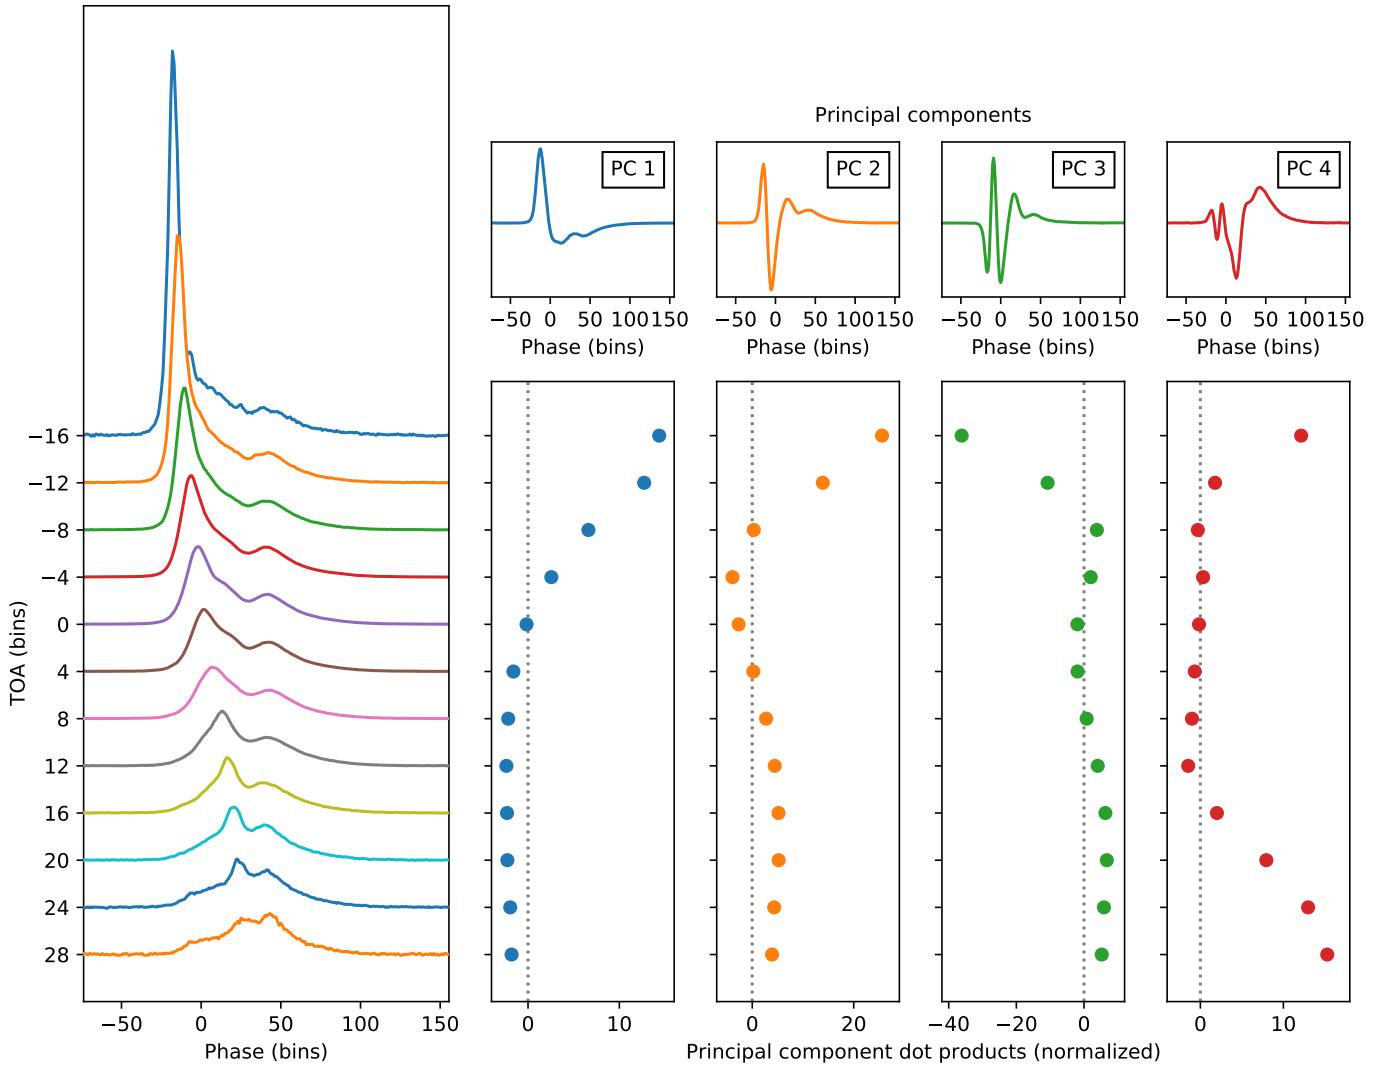


Figure 5.12: An illustration of the relationship between pulse shape and the matched-filtering estimate of the TOA for the Vela pulsar. A matched-filtering TOA estimate, $\hat{\tau}$, was calculated for each of the 18254 single pulses in the May 4, 2016 observation. The topocentric period was estimated by performing a linear fit to the $\hat{\tau}$ values, and the pulses were rotated to align them with respect to this estimated period. Pulses were then assigned to one of twelve groups according to the $\hat{\tau}$ value after rotation. Each represents a range of four phase bins, or $1/512$ of the pulse period. 283 pulses with outlying $\hat{\tau}$ values were not included in any group. The average pulse shape in each group is shown in the left panel above. The correlation between arrival time and intensity identified by Krishnamohan & Downs (1983) is clearly visible. The panels to the right show the dot products of the first four principal components (see Figure 5.11) with the group average pulses, which have been normalized by dividing them by the square roots of the corresponding variances.

Figure 5.9 using equation (5.46) gives $f = 0.28$. However, as will be seen below, a single-component model is not adequate to describe the pulse-to-pulse variation seen in this data.

As expected for errors caused by single-pulse stochasticity, the TOA estimation errors in the Vela data (estimated using differences between the matched-filtering TOA estimates and the predictions of the timing model) are highly correlated between different frequencies within the band. This is demonstrated in Figure 5.10, which shows the correlations between TOAs in different channels. To calculate these, the full 300 MHz band was split into 24 channels, each 12.5 MHz wide, and 507 average profiles, each consisting of 36 pulses, were constructed in each channel. A TOA was then estimated for each profile using the matched-filtering algorithm described in Section 5.4, and a Pearson product-moment correlation coefficient was calculated between each pair of channels. The results show that TOAs at the top and bottom of the band are correlated slightly less than those in adjacent channels, but only slightly: even the weakest correlation shown in Figure 5.10 is very large, at 0.93.

Principal components computed from an observation of the Vela pulsar are shown in Figure 5.11. They are not altogether dissimilar from the principal components seen in the more complex simulated cases in Figure 5.6, demonstrating that the simulations are a reasonable model. However, the presence of a large number of significant principal components indicates that the variability seen in pulses from the Vela pulsar cannot be explained by a model with only amplitude variations – some more complex shape change mechanism, such as phase variations, is required for a complete explanation.

5.7 State Switching

As previously mentioned in Section 5.2, some pulsars appear to alternate between two or more different modes of emission with distinct average pulse shapes. Other pulsars exhibit nulling, appearing to cease emitting entirely for short periods of time. Nulling may be thought of as an extreme form of mode changing, in which one or more of the modes shows little to no emission. Indeed, Esamdin et al. (2005) found that PSR B0826–34 showed faint emission, with a different average pulse shape, during its apparent nulls, confirming that the apparent nulling was actually mode changing. Reviews of the nulling and mode-changing phenomena may be found in Rankin (1986) and Wang et al. (2007), and more detailed studies of nulling in particular are given by Biggs (1992), Gajjar (2017), and Sheikh & MacDonald (2021).

The state transitions seen in nulling and mode-changing pulsars are not periodic, although they do exhibit characteristic time scales, with certain states sometimes being longer-lasting on average than others. Cordes (2013c) showed that transitions between the states may be modeled as a Markov process.

In general, timing of mode-changing pulsars may be carried out using a separate template for each mode, and the phase alignment between profiles for different modes may be determined using observations which extend across mode transitions. Nulling pulsars may not, of course, be timed during nulls, but in other respects can be timed in the same way as other pulsars.

PCA responds to mode changing in the same way that it responds to amplitude modulation: several principal components will be produced, which, together with the average pulse shape, span the same space as the profiles

in each emission mode. The dot products of the profile residuals with these principal components will show a pattern of abrupt changes at the transitions between modes. Such an analysis could in principle be used to identify mode changing where it is not obvious by eye, such as when the difference between modes is small.

Abrupt changes like those seen in PSR J1643–1224 (Shannon et al., 2016) and PSR J1713+0747 (Xu et al., 2021; Lam, 2021), which may be accompanied by a decay toward a final stable state, can be quantified using PCA in a similar way. Any gradual time evolution of the profile, while it might not be confined to a single principal component, will almost certainly be well described using a small number of principal components, as long as the profiles remain continuous functions of pulse phase.

5.8 Instrumental Effects and RFI

The shape changes described thus far have been astrophysical in origin, but instrumental effects and RFI can also give rise to apparent shape changes. Such effects only get in the way of making astrophysical measurements, but it is for exactly this reason that they must be identified where they occur and taken into account. Ordinarily, this is done by discarding the corrupted data entirely, but it is also possible in some cases to account for the effect of the RFI and recover some useful information from the data. Some of the techniques described above can be useful both for identifying data which has been corrupted by RFI and for understanding the extent to which low-level RFI influences timing.

As mentioned previously, PCA is a sensitive probe of low-level variation

in a dataset. As such, it can be used to quickly identify parts of a large dataset which are affected by RFI. Significant RFI will often appear in the first few principal components, even if it affects only a single profile in the entire dataset. The affected profile or profiles can then be determined by examining the dot products of profiles with the component in which the RFI appears. As a method for identifying RFI, this can be significantly faster than manual examination of the profiles, and lends itself well to automation, which is important for large datasets such as those assembled by PTAs.

One particularly common form of interference is baseline ripple – a sinusoidal signal usually most visible in the off-pulse region of a profile, and typically caused by interference related to AC transmission lines, occurring at harmonics of either 50 or 60 Hz, depending on the frequency used for power distribution in the region in which the telescope is located. Because of its predictable shape, it is possible to quantify the size of TOA errors produced by baseline ripple using equation (5.21). The profile residual produced by baseline ripple has the form

$$r(\phi) = \Re[z \exp(2\pi i f P \phi)], \quad (5.58)$$

where $z = x + iy$ is the ripple amplitude, P is the pulse period and f is the ripple frequency. Substituting this into equation (5.21) gives

$$\delta\tau \approx -\frac{\Re\left[z \int_0^1 u'(\phi - \tau) \exp(2\pi i f P \phi) d\phi\right]}{\int_0^1 u'(\phi)^2 d\phi}. \quad (5.59)$$

5.9 Summary and Conclusions

In this paper, we have quantified the ways in which pulse shape changes lead to errors in the estimation of TOAs from pulsars. Shape changes can occur for

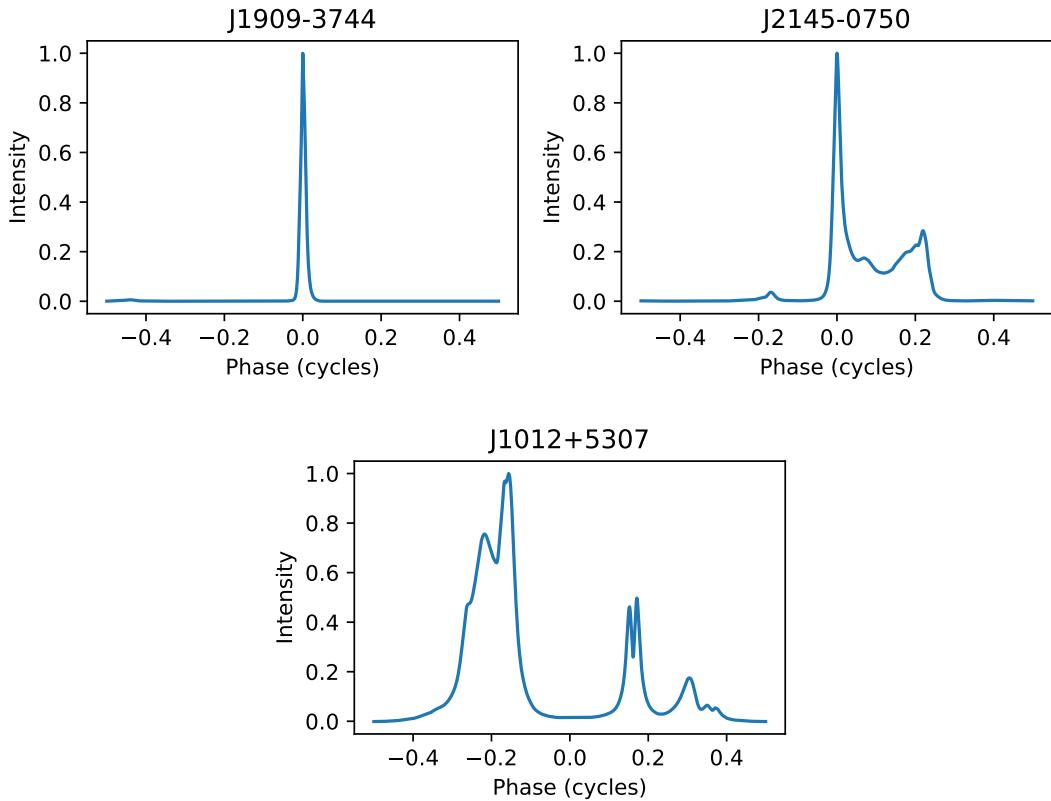


Figure 5.13: Average profiles for three millisecond pulsars (PSRs J1909–3744, J2145–0750, and J1012+5307), demonstrating the variety of pulse shapes exhibited by MSPs. All three profiles were taken from the NANOGrav 12.5-year data release, and are derived from observations made using the 1–2 GHz receiver and GUPPI backend on the Green Bank Telescope between 2010 and 2017.

a number of reasons. All pulsars show stochastic shape changes in the short term, between individual pulses, and others exhibit nulling or mode-changing behavior, or abrupt shape changes whose effect decays over a period of months. Effects related to dispersion and scattering by the ISM may also cause effective shape changes, as might RFI. Regardless of their origin, differences between observed profiles and the template used for TOA determination lead to TOA estimation errors whose magnitude can be predicted using equation (5.23), as long as the difference is small and of known form.

We have also characterized the phenomenology of single-pulse shape variations, both by examining the TOA estimation errors produced by a generative model that includes the effects of component amplitude and shape variations, and by applying various assessment tools to both simulated data and observations of the Vela pulsar. A general formula for the magnitude of TOA estimation errors created by component amplitude and phase variations (equation 5.34) is derived. We find that, while amplitude variations alone cannot cause TOA estimation errors in a single-component pulse, they can do so if there are multiple components that vary independently in amplitude, and that these effects are most pronounced when two or more components overlap substantially. If there are multiple components that are well-separated in phase, the associated TOA estimation errors combine in a manner weighted by the squared amplitude, and inversely by the width, of each component (equation 5.42). If the components are identical, the estimation errors decrease in proportion to the inverse square root of the number of components.

These results will have significant consequences for pulsar timing experiments, particularly pulsar timing array searches for gravitational waves. The MSPs timed by PTAs exhibit a diverse set of profile shapes (see Figure 5.13), and these are known to have significant consequences for timing as a result of equation (5.16), with sharp, narrow profiles like that of PSR J1909–3744 (seen in the left panel of Figure 5.13) yielding the best results. The ability to further characterize the timing performance of pulsars by examining the degree to which single-pulse stochasticity affects TOA estimates could be used to help determine whether a pulsar is suitable for inclusion in a PTA, or to determine whether noise parameter estimates are reasonable. Further, the ability to quantify the extent to which differences from the template shape influence TOA

estimates could be used to determine when data affected by shape changes represents a problem for timing.

5.10 Acknowledgements

The authors are members of the NANOGrav Physics Frontiers Center, which receives support from the National Science Foundation (NSF) under award number 1430284. This project made use of the Parkes 64-meter “Murriyang” radio telescope. The Parkes radio telescope is part of the Australia Telescope National Facility (grid.421683.a) which is funded by the Australian Government for operation as a National Facility managed by CSIRO. We acknowledge the Wiradjuri people as the traditional owners of the Observatory site. We would particularly like to thank M. Kerr for providing baseband voltage recordings of the Vela pulsar, PSR B0833–45, taken as part of the P816 observing project, as well as providing feedback on a draft version of this document.

5.11 Appendix: Periodic Extensions and Fourier-Domain Shifts

Throughout this paper, we describe profiles and templates as continuous functions of phase, but in practice, they are known only at a finite number, N_ϕ , of phase bins, corresponding to phase values $\phi_j = j/N_\phi$. Our use of a continuous notation here can be justified in terms of Fourier analysis. The discrete Fourier

transform (DFT) of a profile p sampled at the points ϕ_j is given by

$$\tilde{p}_k = \sum_{j=0}^{N_\phi-1} p(\phi_j) e^{-2\pi i k \phi_j}. \quad (5.60)$$

The DFT amplitudes are periodic in k with period N_ϕ : that is, $\tilde{p}_k = \tilde{p}_{k+N_\phi}$ for any integer k . Because $p(\phi_j)$ is real, the amplitudes also have a Hermitian symmetry, satisfying $\tilde{p}_{N_\phi-k} = \tilde{p}_k^*$. We can reconstruct the values $p(\phi_j)$ from any N_ϕ consecutive amplitudes using the inverse DFT:

$$p(\phi_j) = \frac{1}{N_\phi} \sum_{k=0}^{N_\phi-1} \tilde{p}_k e^{2\pi i k \phi_j}. \quad (5.61)$$

Although the individual terms are complex numbers, the sum on the right-hand side of equation (5.61) is a real number. We can see this more explicitly by writing it as a trigonometric polynomial in ϕ_j . When N_ϕ is odd ($N_\phi = 2n + 1$), equation (5.61) becomes

$$p(\phi_j) = \frac{1}{2n+1} \left[a_0 + 2 \sum_{k=1}^n a_k \cos(2\pi k \phi_j) - 2 \sum_{k=1}^n b_k \sin(2\pi k \phi_j) \right], \quad (5.62)$$

where a_k and b_k are the real and imaginary parts of \tilde{p}_k , respectively: i.e., $\tilde{p}_k = a_k + i b_k$. Similarly, when N_ϕ is even ($N_\phi = 2n$), we have

$$p(\phi_j) = \frac{1}{2n} \left[a_0 + 2 \sum_{k=1}^{n-1} a_k \cos(2\pi k \phi_j) + 2 \sum_{k=1}^{n-1} b_k \sin(2\pi k \phi_j) + a_n \cos(2\pi n \phi_j) \right], \quad (5.63)$$

where, as above, $\tilde{p}_k = a_k + i b_k$.

We can extend p to arbitrary real values of ϕ by replacing ϕ_j in equation (5.62) or (5.63) with an arbitrary value of ϕ . The result is a periodic function of ϕ with unit period. This is not equivalent to making the same substitution in equation (5.61): in passing from equation (5.61) to equations (5.62) and (5.63), we made use of the fact that $e^{2\pi i (N_\phi - k) \phi_j} = e^{-2\pi i k \phi_j}$, and, in the even case, the fact that $e^{\pi i N_\phi \phi_j} = e^{-\pi i N_\phi \phi_j}$, neither of which remains true if ϕ_j is replaced by an

arbitrary value of ϕ . Extending p using equation (5.62) or (5.63) is preferable because the resulting function will always be real-valued, which is not true if p is extended using equation (5.61).

The extension of p described above is a form of interpolation, and can also be described as the convolution of p with a kernel function, $K(N_\phi, \phi)$, which is as a discrete analog of the sinc function. In particular, we have

$$p(\phi) = \sum_{j=0}^{N_\phi-1} p(\phi_j) K(N_\phi, \phi - \phi_j), \quad (5.64)$$

where the kernel function is given by

$$K(N, \phi) = \frac{\sin(N\pi\phi)}{N \sin(\pi\phi)}, \quad (5.65)$$

when N is odd, and

$$K(N, \phi) = \frac{\sin(N\pi\phi)}{N \tan(\pi\phi)} \quad (5.66)$$

when N is even.

We use this extension throughout the paper to allow profiles to be shifted by a fractional number of phase bins. In practice, the shifted profile can be computed using a fast Fourier transform (FFT) algorithm: First, a “real” FFT is used to compute the DFT of the profile at non-negative frequencies. Then, the result is multiplied by the appropriate phase factor, determined by the Fourier shift theorem. Finally, the inverse transform, given by equation (5.62) or (5.63) is computed, also using an FFT. Additionally, $K(N, \phi)$ satisfies

$$\int_0^1 K(N, \phi) d\phi = \frac{1}{N}, \quad (5.67)$$

so the average value of the interpolated function is equal to the average of the sample values:

$$\int_0^1 p(\phi) d\phi = \frac{1}{N_\phi} \sum_{j=0}^{N_\phi-1} p(\phi_j). \quad (5.68)$$

CHAPTER 6

IMPROVING PULSAR TIMING BY CORRECTING FOR PULSE SHAPE CHANGES

6.1 Introduction

Pulsar timing depends on the link between measurable pulse times of arrival (TOAs) and the rotation of a neutron star, which is very stable as a result of the neutron star's large angular momentum. Monitoring the TOAs allows very small changes to the path length from the pulsar to the observatory to be measured. Such measurements have enabled the first detection of planets outside the Solar System (Wolszczan, 1994), the detection of energy loss from gravitational wave emission in a binary neutron star system (Weisberg et al., 2010), and the direct measurement of neutron star masses using Shapiro delay (e.g., Cromartie et al., 2020), among other significant results. Currently, pulsar timing arrays (PTAs) are attempting to use pulsar timing measurements to directly detect low-frequency gravitational waves from extragalactic sources.

The accuracy and precision of the underlying TOA measurements is crucial to all of these applications. One factor which limits both of these is the fact that the shapes of single pulses vary significantly. Conventional TOA measurement techniques rely on matched filtering against a template based on the average pulse shape, which is characteristic of each pulsar. This is useful because, when a large number of pulses are averaged to form a pulse profile, the result will converge toward the template shape. However, the presence of large shape variations in single pulses means that any finite average will differ from the template, if only by a small amount. These small differences in shape give rise

to TOA errors, an effect often referred to as jitter noise. In a previous chapter (Chapter 5), we quantified how shape changes give rise to jitter noise, and gave several methods for characterizing the effect at the profile level. In what follows, we will explore the possibility of mitigating the effect of jitter noise by taking into account more information about the profile shape than is used in the conventional matched filtering algorithm.

We are not the first to consider the prospect of mitigating at least some of the error caused by pulse shape variations using pulse shape information. The idea of using principal component analysis (PCA) as a correction method was introduced by Demorest (2007), and further developed by Osłowski et al. (2011), who used a variant of the principal component dot product method, to be described below, on 25 hours of data from the PSR J0437–4715. Lentati et al. (2015) devised a related scheme, using a basis of Hermite functions referred to as “shapelets”, as part of a broader program in which profile shapes were inferred simultaneously with the timing model. This “generative pulsar timing analysis” has the potential to be very computationally expensive for large data sets, and as a result has not, to date, been widely adopted. Kerr (2015), working with data from the Vela pulsar, PSR B0833–45, used a clustering technique to produce a a representative set of single pulses, and modeled profiles as linear combinations of these. In this paper, we introduce a new technique for mitigating errors caused by pulse shape changes, which extends the traditional template matching algorithm, and compare it to two other techniques: the principal component dot product method used by Demorest (2007) and Osłowski et al. (2011), and an approach based on the profile skewness function, a shift-invariant measure of pulse asymmetry.

6.2 Overview of Methods

As described in Chapter 5, the conventional matched filtering TOA estimate, $\hat{\tau}$, is calculated from a profile, $p(\phi)$, and a template $u(\phi)$, by maximizing the likelihood (equivalently, minimizing the χ^2 value) for a model of the form

$$p(\phi) = au(\phi - \tau) + n(\phi), \quad (6.1)$$

where $n(\phi)$ is white noise with variance σ_n^2 . The model parameters are a , the profile amplitude, and τ , the phase shift, which is related to the TOA in a physical time scale by

$$t_a = t_0 + P\tau, \quad (6.2)$$

where t_a is the time of arrival in the target time scale, t_0 is the time, closest to the center of the observation from which $p(\phi)$ was determined, that corresponds to a whole number of turns in the model used for folding, and P is the period of the pulsar (more precisely, the topocentric period at time t_0). In what follows, we will refer to τ as the TOA.

The likelihood for model (6.1) has the form $\mathcal{L}(a, \tau) = e^{-\frac{1}{2}\chi^2(a, \tau)}$, where

$$\chi^2(a, \tau) = \frac{N_\phi}{\sigma_n^2} \int_0^1 [p(\phi) - au(\phi - \tau)]^2 d\phi. \quad (6.3)$$

The maximum likelihood (minimum χ^2) occurs at $(\hat{a}, \hat{\tau})$, where

$$\hat{\tau} = \operatorname{argmax}_\tau \hat{a}(\tau), \quad (6.4)$$

$$\hat{a} = \hat{a}(\hat{\tau}), \quad (6.5)$$

$$\hat{a}(\tau) = \frac{1}{\langle u^2 \rangle} \int_0^1 p(\phi)u(\phi - \tau) d\phi. \quad (6.6)$$

By expanding the likelihood in a power series about its maximum value, the

uncertainties in the estimates \hat{a} and $\hat{\tau}$ can be determined:

$$\sigma_a^2 = \frac{\sigma_n^2}{N_\phi \langle u^2 \rangle} \quad (6.7)$$

$$\sigma_{\tau}^2 = \frac{\sigma_n^2 W_{\text{eff}}^2}{N_\phi \hat{a}^2}. \quad (6.8)$$

Here the effective pulse width, W_{eff} , is given by

$$W_{\text{eff}} = \left(\int_0^1 u'(\phi)^2 d\phi \right)^{-1/2}. \quad (6.9)$$

If the profile differs from the shifted and scaled template by a small amount, $r(\phi)$ (the profile residual), so that

$$p(\phi) = au(\phi - \tau) + r(\phi), \quad (6.10)$$

then, as shown in Chapter 5, the TOA estimation error, $\delta\tau = \hat{\tau} - \tau$, will be given by

$$\delta\tau \approx -\frac{\int_0^1 r(\phi) u'(\phi - \tau) d\phi}{a \int_0^1 u'(\phi - \tau)^2 d\phi}. \quad (6.11)$$

This equation is useful for predicting the amount of TOA error created by shape variations of known form, as explored in Chapter 5. However, it cannot be used to correct an individual TOA estimate if τ is not known, since approximating τ as $\hat{\tau}$ and a as \hat{a} will always give $\delta\tau \approx 0$. Instead, the starting point for any mitigation technique must be a measure of pulse shape, independent of τ , that correlates with the TOA estimation error, $\delta\tau$.

Suppose $x[p(\phi)]$ a measure of pulse shape satisfying $x[au(\phi)] = 0$ for any a , and that the profile residual, $r(\phi)$, is drawn from a consistent distribution with mean zero. We can then calculate the correlation coefficient, ρ , between x and $\delta\tau$:

$$\rho = \frac{\langle x \delta\tau \rangle}{\sqrt{\langle x^2 \rangle \langle \delta\tau^2 \rangle}}, \quad (6.12)$$

where angle brackets denote ensemble averages. If ρ is nonzero, we can exploit the correlation to construct a new, improved estimate of τ , $\hat{\tau}'$. The best linear, unbiased estimate is

$$\hat{\tau}' = \hat{\tau} - \rho \sqrt{\frac{\langle \delta\tau^2 \rangle}{\langle x^2 \rangle}} x, \quad (6.13)$$

which has mean squared error given by

$$\langle (\hat{\tau}' - \tau)^2 \rangle = (1 - \rho^2) \langle \delta\tau^2 \rangle. \quad (6.14)$$

Certain restrictions need to be placed on corrections of this sort to avoid overfitting, i.e., removing real TOA shifts when they are present. This means that the shape parameter, x , should be insensitive to true shifts, ideally satisfying

$$x[p(\phi)] = x[p(\phi - \delta\tau)] \quad (6.15)$$

for any shift, $\delta\tau$. One function of the profile which has this property is the autocorrelation function (ACF):

$$R(\phi) = \frac{1}{\int_0^1 p(\phi')^2 d\phi'} \int_0^1 p(\phi') p(\phi' + \phi) d\phi'. \quad (6.16)$$

Unfortunately, in addition to being insensitive to shifts, $R(\phi)$ is insensitive to profile asymmetries, giving the same result for the symmetrized profile $p_s(\phi) = \frac{1}{2}[p(\phi) + p(-\phi)]$ as for $p(\phi)$. Because the profile is often nearly symmetric, its derivative is often nearly antisymmetric, meaning that, as a consequence of equation (6.11), pulse asymmetries are some of the best predictors of TOA estimation errors. This means that $R(\phi)$ and statistics derived from it are not particularly useful as shape parameters. However, there is a generalization of $R(\phi)$ which does capture information about pulse asymmetry, while remaining insensitive to shifts. This is the profile skewness function:

$$K(\phi) = \frac{1}{\int_0^1 p(\phi')^3 d\phi'} \int_0^1 [p(\phi')^2 p(\phi' + \phi) - p(\phi') p(\phi' + \phi)^2] d\phi'. \quad (6.17)$$

The skewness function is the antisymmetric part of the cross-correlation function between the profile, $p(\phi)$, and its square, $p(\phi)^2$. It is also related to the three-point ACF, part of an infinite family of n -point ACFs:

$$R_n(\phi_1, \dots, \phi_{n-1}) = \frac{1}{\int_0^1 p(\phi')^n d\phi'} \int_0^1 p(\phi') \prod_{i=1}^{n-1} p(\phi' + \phi_i) d\phi'. \quad (6.18)$$

The n -point ACF is symmetric under arbitrary permutations of its $n - 1$ arguments, as well as substitutions which replace one argument, ϕ_j , with $-\phi_j$, while replacing each other argument, ϕ_i , with $\phi_i - \phi_j$. The ordinary ACF is the two-point ACF: $R(\phi) = R_2(\phi)$, while the skewness function is related to the three-point ACF by

$$K(\phi) = R_3(0, \phi) - R_3(\phi, \phi) = R_3(0, \phi) - R_3(0, -\phi). \quad (6.19)$$

Any statistic, x , which depends on the profile, $p(\phi)$, solely through one or more of the n -point autocorrelations, $R_n(\phi)$, is insensitive to shifts in the sense of equation (6.15). One example of such a statistic, which will be explored in further detail below, is the skewness statistic $S = K(\phi_0)$ obtained by evaluating the skewness function at a lag value, ϕ_0 , much smaller than the pulse width. S is a measure of pulse asymmetry which is insensitive to shifts. If the pulse shape is continuous, expanding equation (6.17) in a Taylor series in ϕ shows that $K(\phi) \propto \phi^3$ for sufficiently small values of ϕ . As a result, as long as the value of ϕ_0 is sufficiently small, it does not have a significant effect on the relevance of S as a predictor of TOA estimation error, since the S values corresponding to different values of ϕ_0 are related by a constant factor.

6.3 PCA and Shift Invariance

A more direct way to quantify the shape of pulses is to use principal component analysis (PCA). As described in Paper I, PCA finds a set of orthogonal and linearly independent modes that best capture the variance in a set of data. In this case, the data are the profile residuals, $r(\phi)$, which measure the difference between the profile shape and a template shape. PCA amounts to an eigenvalue decomposition of the covariance matrix, $\mathbf{C} = \langle \mathbf{r} \mathbf{r}^T \rangle$, of the residuals. The result is a set of K orthonormal principal components, $v_i(\phi)$, and associated eigenvalues, σ_i^2 , numbered in order of decreasing magnitude. If \mathbf{C} is non-singular, K is equal to the total number of phase bins, N_ϕ , and the principal components form an orthonormal basis for the space of residuals. This means that any particular residual can be written in the form

$$r(\phi) = \sum_{i=1}^K x_i v_i(\phi). \quad (6.20)$$

As a result of the way the principal components are constructed, the coefficients, x_i , are statistically independent, and satisfy $\langle x_i x_j \rangle = \sigma_i^2 \delta_{ij}$. Because the $v_i(\phi)$ are orthonormal, the coefficient values, x_i , can be determined from $r(\phi)$ by a dot product:

$$x_i = \int_0^1 r(\phi) v_i(\phi - \tau) d\phi. \quad (6.21)$$

Because of this, we will refer to them as principal component dot products. Note that this requires a preliminary estimate of the TOA, τ .

One can construct a shape parameter, x , for use in correcting TOA estimates, as a linear combination of the first k principal component dot products, x_i :

$$x = \sum_{i=1}^k c_i x_i = \sum_{i=1}^k c_i \int_0^1 r(\phi) v_i(\phi - \tau) d\phi. \quad (6.22)$$

If all of the principal components are used ($k = K$), this is a general linear function of the profile. Restricting to a small number of principal components has a regularizing effect, in that only the most significant modes of variation can contribute. This makes it less likely that the model will overfit, by reducing the number of degrees of freedom while retaining the ability to account for shape changes of almost any form. However, there is still a potential problem with using a shape parameter of the form given by equation (6.21): it is sensitive to shifts; i.e., equation (6.15) is not satisfied. Indeed, as long as the shift, $\delta\tau$, is much less than the width of the pulse, we have

$$p(\phi - \delta\tau) \approx p(\phi) - \delta\tau p'(\phi) \approx p(\phi) - a \delta\tau u'(\phi - \tau), \quad (6.23)$$

where we have used the assumption that the residual, $r(\phi)$, is much smaller in magnitude than $au(\phi - \tau)$, and not much more rapidly varying, so that the product $\delta\tau r'(\phi)$ can be considered negligibly small in comparison to $p'(\phi)$. This means that the change in each of the principal component dot products will be given by

$$\Delta x_i = -a \delta\tau \int_0^1 u'(\phi) v_i(\phi) d\phi. \quad (6.24)$$

This change will be zero if each principal component is orthogonal to the template derivative, $u'(\phi)$. Such a state of affairs can be arranged by projecting out the template derivative from the residual before computing the principal components. This amounts to using a modified residual

$$r^*(\phi) = p(\phi) - \hat{a}u(\phi - \hat{\tau}) - \hat{b}u'(\phi - \hat{\tau}) \quad (6.25)$$

in place of $r(\phi)$. Here the best-fit values, \hat{a} and \hat{b} , are given by

$$\hat{a} = \frac{\int_0^1 p(\phi)u(\phi - \hat{\tau}) d\phi}{\int_0^1 u(\phi)^2 d\phi}, \quad (6.26)$$

$$\hat{b} = \frac{\int_0^1 p(\phi)u'(\phi - \hat{\tau}) d\phi}{\int_0^1 u'(\phi)^2 d\phi}. \quad (6.27)$$

By construction, residuals of the form (6.25) are orthogonal to $u'(\phi - \hat{\tau})$; i.e., they satisfy

$$\int_0^1 r^*(\phi) u'(\phi - \hat{\tau}) d\phi = 0. \quad (6.28)$$

This means that the principal components, $v_i^*(\phi)$, derived from them are orthogonal to $u'(\phi)$ as well, and so (importantly) the dot products with these principal components are approximately shift invariant, at least for small values of the shift, $\delta\tau$.

This, then, is the principal component dot product method, a minor variant of which was applied by Osłowski et al. (2011) to PSR J0437–4715. To use it, one must first compute profile residuals of the form (6.25) for each of the profiles in a set of training data. One then computes the eigenvalues, σ_i^2 , and eigenvectors, $v_i^*(\phi)$, of the covariance matrix

$$C(\phi, \phi') = \sum_{i=1}^N r_i^*(\phi) r_i^*(\phi'), \quad (6.29)$$

and selects a number k of eigenvectors to use. Next, the coefficients, c_i , which define the shape parameter

$$x = \sum_{i=1}^k c_i x_i^*, \quad (6.30)$$

where

$$x_i^* = \int_0^1 r^*(\phi) v_i^*(\phi - \tau) d\phi, \quad (6.31)$$

can be selected so as to maximize the correlation between x and the TOA estimation error, $\delta\hat{\tau} = \hat{\tau} - \tau$. Here $\hat{\tau}$ is the standard template-matching estimate of the TOA. As long as a sufficient number of profiles are included in the training set, a timing model describing the topocentric phase and period of the pulsar can be fit to the data to produce an estimate of τ for each profile which is much better than the matched-filtering estimate, $\hat{\tau}$, for any single profile. Using these

estimates of τ , $\delta\hat{\tau}$ can be estimated for each profile. The values of c_i which maximize the correlation with $\delta\hat{\tau}$ are then

$$c_i = \frac{\sum_{j=1}^N x_{ij}^* \delta\hat{\tau}_j}{\sum_{j=1}^N (x_{ij}^*)^2}. \quad (6.32)$$

This is the “multiple regression” described by Osłowski et al. (2011). Finally, the principal components, $v_i^*(\phi)$, and the coefficients, c_i , can be applied to produce improved TOA estimates for new data. These are computed as

$$\hat{\tau}' = \hat{\tau} + \sum_{i=1}^k c_i \int_0^1 r^*(\phi) v_i^*(\phi - \tau) d\phi, \quad (6.33)$$

where $\hat{\tau}$ is the traditional matched-filtering estimate computed for the profile $p(\phi)$, and $r^*(\phi)$ is the residual for the same profile, computed using equation (6.25).

6.4 Flexible Template Matching Algorithms

The algorithm described above operates by correcting the estimate produced by the conventional template matching algorithm. It is also possible to modify the template matching algorithm itself using similar ideas. This can be done by extending the profile model (equation 6.1) to include principal components, obtaining

$$p(\phi) = au(\phi - \tau) + bu'(\phi - \tau) + \sum_{i=1}^k x_i^* v_i^*(\phi - \tau) + n(\phi). \quad (6.34)$$

Note that the shift, τ , is common between all terms, but that a term proportional to the derivative of the template, $u'(\phi)$, has been added in addition to the principal component terms. The corresponding coefficient, b , is taken to satisfy

$$b = \sum_{i=1}^k m_i x_i^* + \varepsilon, \quad (6.35)$$

where ε is normally distributed with variance σ_ε^2 . The coefficients m_i and the variance σ_ε^2 can be estimated from a training data set, as with the coefficients c_i in the previous section. Equation (6.34) can be understood as a generative model for $p(\phi)$, and therefore can be used to derive a likelihood function. The result may be expressed most easily by writing $p(\phi)$ in the basis of shifted principal components, as

$$p(\phi) = \hat{a}(\tau)u(\phi - \tau) + \hat{b}(\tau)u'(\phi - \tau) + \sum_{i=1}^{N_\phi} \hat{x}_i(\tau)v_i^*(\phi - \tau). \quad (6.36)$$

Here only the first k principal components are considered significant. The others are included only for the purpose of forming a complete basis, and may be replaced with any other vectors orthogonal to u , u' , and the first k principal components. The likelihood has the form $\mathcal{L}(\tau, x_i^*) = e^{-\frac{1}{2}\chi^2(\tau, x_i^*)}$, where

$$\begin{aligned} \chi^2(\tau, x_i^*) &= \frac{(\hat{a}(\tau) - a)^2}{\sigma_n^2} + \frac{(\hat{b}(\tau) - \sum_{i=1}^k m_i x_i^*)^2}{\sigma_\varepsilon^2 + \sigma_n^2} \\ &\quad + \sum_{i=1}^k \frac{(\hat{x}_i(\tau) - x_i^*)^2}{\sigma_i^2 + \sigma_n^2} + \sum_{i=k+1}^{N_\phi} \frac{\hat{x}_i(\tau)^2}{\sigma_n^2}. \end{aligned} \quad (6.37)$$

Maximizing this likelihood function allows τ to be estimated with an accuracy comparable to that of the principal component dot product method described in the previous section.

6.5 Application to Simulated Data

Figure 6.1 shows a comparison of several correction methods, applied to simulated data. The methods used were a simple shape-parameter method based on the skewness coefficient, the principal component dot product method described in Section 6.3, and the principal component profile reconstruction method described in Section 6.4.

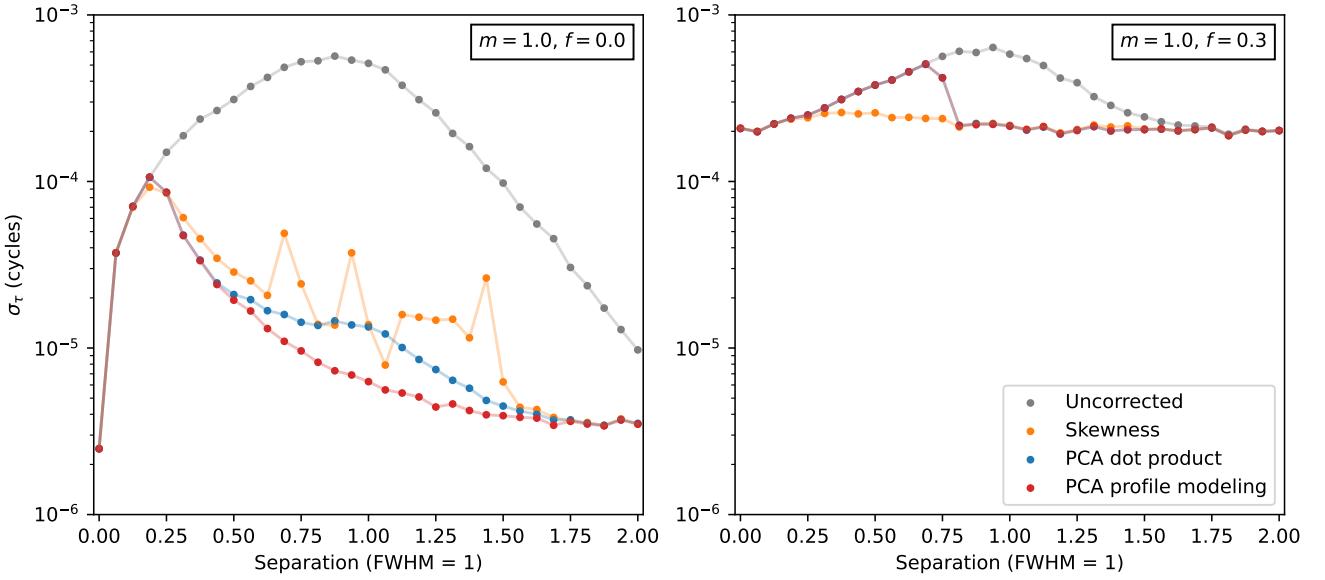


Figure 6.1: A comparison of TOA correction methods, as applied to simulated data. The data shown are for simulated pulses with two identical Gaussian components, having varying separation. The standard deviation of the TOA estimation error is shown for uncorrected data (gray) and for data corrected using the skewness method (orange) and the PCA dot product (blue) and profile modeling (red) methods. The left panel shows the results for pulses exhibiting component amplitude variations, but no phase variations, while the right panel shows the results for pulses that also exhibit phase variations. As shown, the two PCA-based methods exhibit similar performance, while the skewness method does somewhat worse in most of the cases shown at left.

The simulations are based on a model in which individual pulses are composed of several Gaussian components, each of which is subject to amplitude and phase variations. In brief, each simulated profile is given by a sum of the form

$$p(\phi) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{N_c} a_j (1 + b_{ij}) \exp \left[-\frac{(\phi - \mu_j - \psi_{ij})^2 \ln 2}{w_j^2} \right] + n(\phi), \quad (6.38)$$

where N is the number of pulses; N_c is the number of components; a_j , μ_j , and w_j are the average amplitude, central phase, and full width at half maximum (FWHM) of each component, respectively; and $n(\phi)$ represents white noise. The parameters b_{ij} and ψ_{ij} are responsible for the amplitude and phase

variations, with ψ_{ij} being drawn from a normal distribution with variance $\langle \psi_{ij}^2 \rangle = f_j^2 w_j^2$, and $1 + b_{ij}$ drawn from a log-normal distribution with unit mean and variance $\langle b_{ij}^2 \rangle = m_j^2$. The parameters f_j and m_j are known as the jitter parameter and modulation index, respectively.

The results shown in Figure 6.1 are for simulations in which the pulses had two components, each with amplitude $a = 1$, FWHM $w = 0.05$, and modulation index $m = 1$. The separation between the central phases of the two components varies between simulations, and is shown on the horizontal axis in units of the FWHM. The left panel shows results for simulations without phase variations ($f = 0$), while the right panel shows results for simulations with jitter parameter $f = 0.3$. For each of the principal component based methods, the number of significant principal components was set at $k = 8$.

In many of the cases shown, the correction methods perform well. They perform particularly well in cases where the majority of the TOA estimation is the result of component amplitude variations, rather than phase variations. In the absence of phase variations, the correction methods reduce the TOA estimation error by well over 90% in many cases. When they are present, the reduction in error using all three methods is much less, indicating that phase variations introduces a component of error that is uncorrectable, at least with these methods. It is also notable that, in the presence of phase variations, the skewness-based method outperforms the two PCA-based methods for some parameter values, when the reverse is usually the case.

6.6 Application to the Vela Pulsar

We also tested the correction methods by applying them to data from the Vela pulsar, PSR B0833–45. As discussed in Paper I, the Vela pulsar makes a particularly good test case for understanding single-pulse stochasticity because it is extremely bright, and as a result, single pulses can be detected with high signal-to-noise. Here we used a fold-mode observation made with the 10 cm receiver on the Parkes radio telescope on September 26, 2012. The observation, which was processed using the third-generation pulsar digital filterbank (PDFB3), is publicly available online through the CSIRO data access portal.

Data were recorded in 368 32-second subintegrations and one 28-second subintegration, for a total of 3.28 hours of observing time. The observation spans 400 MHz of bandwidth centered on 2820 MHz, nearly all of which is usable. In post-processing, we averaged the dedispersed data across the band to produce a pulse profile for each subintegration. We then calculated TOAs for each profile using the conventional template-matching algorithm and fit for the drift caused by the unknown topocentric period of the pulsar, which we then removed by rotating the profiles in the Fourier domain.

Using the first half of the data (first 185 profiles) as a training set, we then calculated the principal components, $v_i^*(\phi)$, and fit for the coefficients c_i and m_i described in Sections 6.3 and 6.4, respectively. Subsequently, we tested the two PCA-based TOA correction methods for varying numbers of significant principal components, k , by calculating corrected TOAs for the profiles in the second half of the data set. The results are shown in Figure 6.2. As seen there, the best results are achieved for $k = 18$, where the correction reduces

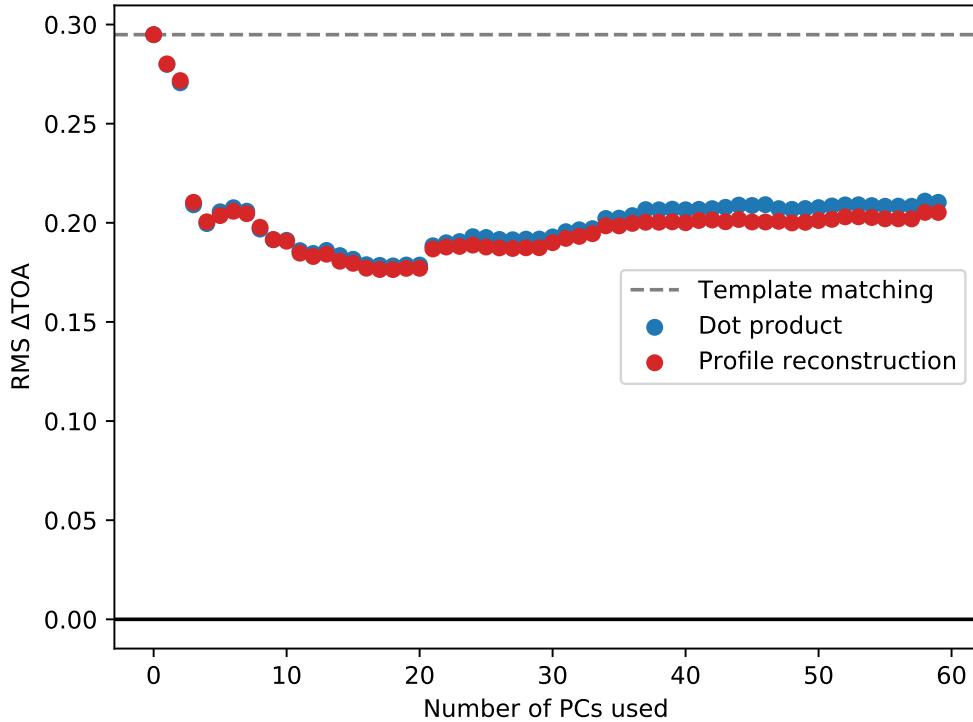


Figure 6.2: The results of two different PCA-based TOA correction methods, applied to data from the Vela pulsar. The results of the PCA dot product method are shown in blue, while those of the profile reconstruction method are shown in red. The number of principal components, k , used in correction, which is a free parameter in both methods, is shown on the horizontal axis.

the variance of the TOA estimates by 40%. The effectiveness of both methods decreases after this point, as the correction coefficients start to overfit. Most of the improvement is present already for $k = 3$. The skewness-based correction method was also tested, but did not produce an improvement compared to the uncorrected template-matching algorithm.

6.7 Summary and Conclusions

We have presented a comparison of three methods for correcting pulse time-of-arrival measurements for the effects of pulse shape changes, including a method based on principal component dot products that was previously investigated by Demorest (2007) and Osłowski et al. (2011), a method based on the profile skewness function, and an extension of the conventional template-matching algorithm which incorporates information from principal component analysis of the profile residuals. All three methods demonstrate some effectiveness, but the PCA-based methods are more effective than the skewness-based method in most cases. Clearly, it is possible to correct at least some of the TOA error caused by pulse shape changes by making use of the additional shape information present in pulsar observations.

Our conclusions could have important implications for pulsar timing in the regime where single-pulse stochasticity dominates the TOA error budget. In particular, these correction techniques could be used to achieve higher timing precision for the brightest pulsars timed by pulsar timing arrays. However, the correction techniques do have limits. In particular, some fraction of the estimation error introduced by component phase variations is not correctable with any of the methods considered here, and may not be correctable by any method. To the extent that single-pulse stochasticity represents a “noise floor” which limits the precision achievable by pulsar timing, TOA correction methods can only push that floor somewhat lower, rather than removing it entirely.

6.8 Acknowledgements

The authors are members of the NANOGrav Physics Frontiers Center, which receives support from the National Science Foundation (NSF) under award number 1430284. This project made use of data from the Parkes 64-meter “Murriyang” radio telescope. The Parkes radio telescope is part of the Australia Telescope National Facility (grid.421683.a) which is funded by the Australian Government for operation as a National Facility managed by CSIRO. We acknowledge the Wiradjuri people as the traditional owners of the Observatory site. We thank M. Kerr for providing baseband voltage recordings of the Vela pulsar, PSR B0833–45, taken as part of the P816 observing project. Additional publically-available fold-mode data were accessed using the CSIRO Data Access Portal (<https://data.csiro.au/>).

CHAPTER 7

RECENT OBSERVATIONS OF A PULSE SHAPE CHANGE IN PSR

J1713+0747

7.1 Introduction

PSR J1713+0747 is a bright millisecond pulsar (MSP) that forms a sensitive component of pulsar timing arrays (PTAs). It is currently observed by all member collaborations of the International Pulsar Timing Array (IPTA; Perera et al., 2019), namely, the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; Alam et al., 2021a; Alam et al., 2021b), the European Pulsar Timing Array (EPTA; Desvignes et al., 2016), the Parkes Pulsar Timing Array (PPTA; Kerr et al., 2020; Reardon et al., 2021), and the Indian Pulsar Timing Array (InPTA; Joshi et al., 2018), as well as by participating scientists from China and South Africa. The high timing precision achievable with J1713+0747 also makes it useful for tests of general relativity (e.g. Zhu et al., 2019).

PTA observations of J1713+0747 show that it has undergone two unusual chromatic timing events in the last 15 years. The first of these, which took place around MJD 54750 (October 11, 2008), was identified by Demorest et al. (2013) as part of the first NANOGrav data release. The second event took place around MJD 57510 (May 2, 2016), and was identified by Lam et al. (2018) using the NANOGrav 12.5-year data set (Alam et al., 2021a). During both events, time-of-arrival (TOA) measurements show an abrupt jump towards earlier times, taking place within a few days, followed by a gradual, approximately exponential recovery taking place over several months. The effect is larger at lower frequencies, and can be modeled by a delay proportional to ν^{-2} , where

ν is the radio frequency. This is consistent with a change in dispersion measure (DM), i.e., the column density of free electrons along the line of sight, of approximately $\Delta \text{DM} \approx -6 \times 10^{-4} \text{ pc cm}^{-3}$ for each event. Lam et al. (2018) attribute the events to lensing of the radio emission by some structure in the ionized interstellar medium. Follow-up by Lin et al. (2021), also using the NANOGrav 12.5-year data, showed that the data are consistent with a model in which the lensing is produced by folds in an underdense sheet.

Lin et al. (2021) also identified low-amplitude achromatic shape changes associated with each of the events. A shape change associated with the second event, but not the first, had previously been identified by Goncharov et al. (2021) in the PPTA second data release. In each case, the profile changes by a few percent of its peak amplitude, and the change appears and decays on the same timescale as the associated chromatic timing event. Lensing can produce additional images that can interfere with the main pulse to produce shape changes, but these are expected to show some dependence on radio frequency, so the achromatic nature of the changes is puzzling in the context of a lensing model.

Xu et al. (2021) recently announced the detection of a much larger pulse shape change in J1713+0737, also associated with chromatic pulse shape variations. They presented profiles observed with the Effelsberg and Nançay radio telescopes, as well as the Kunming 40-meter telescope and the Five-hundred-meter Aperture Spherical Telescope (FAST). FAST also observed an associated change in the linearly polarized emission from the pulsar. The shape change was subsequently confirmed using other telescopes, including the Giant Metre-wave Radio Telescope (GMRT; Singha et al., 2021) and the Canadian Hydrogen

Intensity Mapping Experiment (CHIME; Meyers & CHIME/Pulsar Collaboration, 2021). As mentioned above, the shape change is much larger than those associated with the two previous chromatic timing events, in some cases changing the intensity by more than 50% of the peak pulse amplitude. The daily cadence of CHIME data can be used to narrow down the period in which the event occurred to 24 hours, between MJD 59320 and MJD 59321 (April 16 and 17, 2021). Below, we report on recent observations of J1713+0747 made by NANOGrav and the CHIME/Pulsar Collaboration using the Green Bank Telescope (GBT), the Very Large Array (VLA), and CHIME, both before and after the shape change event, and speculate as to its origin.

7.2 Observations

The observations considered below include NANOGrav observations made using the GBT and the VLA, as well as CHIME observations. The GBT observations were made using two receivers: the 820 MHz receiver, which has a bandwidth of 200 MHz, and the 1–2 GHz receiver, which has 800 MHz of bandwidth centered on 1500 MHz. They were processed using the VEGAS backend (Roshi et al., 2011; Ford et al., 2014) in pulsar mode. The VLA observations were made using the 1.5 GHz receiver, which has 1024 MHz of bandwidth centered on 1500 MHz, and the 3 GHz receiver, which has 2048 MHz of bandwidth centered on 3000 MHz; and processed using the YUPPI backend. The CHIME observations were made using CHIME’s single receiver, which has 400 MHz of bandwidth centered on 600 MHz, and processed using the CHIME/Pulsar backend.

The GBT observations were made as part of the main NANOGrav observing program, and have integration times of about 30 minutes and roughly monthly cadence at both 820 and 1500 MHz. Seven observations at 820 MHz and thirteen at 1–2 GHz were made between January and September 2021, with five at each frequency taking place after the shape change event. The VLA observations made before the shape change event represent about 6 hours of observing time spread across 18 different 20-minute sessions between March 2015 and April 2020. Those taken afterwards were part of a dedicated campaign to observe the pulsar in its altered state, and represent four different epochs during a single week in late May and early June 2021 (MJDs 59361–59368), with four consecutive 10-minute integrations taken at each epoch. CHIME took data on J1713+0747 whenever the it passed through the CHIME beam while the telescope was operational. The pulsar was observed for approximately 15 minutes each day in 2021, with only a few interruptions. Included here are 218 observations taken between January 2 and September 16, 2021 (MJDs 59216–59473).

All of these observations will be included in a future NANOGrav data set. However, due to the recentness of the event, they are not part of NANOGrav’s published 12.5-year data set or the 15-year data set currently being prepared¹.

The fold-mode data were first dedispersed, and then post-processed by removing portions of the band edges in which the bandpass filter rolled off significantly, as well as channels that were identified as containing significant RFI. The presence of RFI was determined by manually inspecting the average pro-

¹The 15-year data set includes data from two sets of pulsar backends: the Astronomical Signal Processor (ASP) at Arecibo and Green Bank Astronomical Signal Processor (GASP) at the GBT (Demorest, 2007), and the later Green Bank Ultimate Pulsar Processing Instrument (GUPPI) at the GBT (DuPlain et al., 2008) and its sibling instruments PUPPI at Arecibo and YUPPI at the VLA. GUPPI was decommissioned in April 2020 due to a hardware failure, and observations at Arecibo ceased after a cable failure in August 2020. The Arecibo telescope ultimately collapsed in November 2020, so no further observations were made using PUPPI.

file for each receiver as a function of frequency and pulse phase. Finally, the data were averaged in frequency within each band to produce a set of profiles, which served as the basis for later analysis. The average profiles before and after the event are shown in Figure 7.1. The full set of profiles for the GBT and CHIME data observed within the calendar year 2021 is shown in Figure 7.2.

A time-of-arrival (TOA) estimate was produced for each profile by matched filtering with a template profile (cf. Taylor, 1992). A separate template profile was created for each receiver, using the average of all the data from before the shape change. The TOA estimates from the GBT data showed evidence of sinusoidal variation at the 67.8-day orbital period of the J1713+0747 system, so a simple timing model consisting of a single sinusoid with linear amplitude modulation was fit to the TOAs before the shape change. The timing model was extrapolated beyond the shape change, and its predictions were subtracted from the TOA estimates to form TOA residuals. The profiles were then aligned by rotating them by the amount predicted by the timing model using a Fourier-domain algorithm. Then, the best-fit scaled and shifted copy of the template was subtracted from each profile to produce a profile residual. The profile residuals for the GBT and CHIME data from 2021 are shown in Figure 7.3, and the TOA residuals are shown in Figure 7.4.

Finally, the dominant mode of variation of each set of profile residuals was determined using principal component analysis (PCA). In each case, a single principal component accounted for the vast majority of the variance in the data. The amount that each profile deviated from the template shape was quantified by taking the dot product of the profile residual with the most significant principal component, and dividing by the maximum value of the principal

component. The results are shown in Figure 7.5.

7.3 Interpretation

The average profiles seen in Figure 7.1 illustrate the form of the shape change at different radio frequencies. All bands are affected to a similar degree. The profile changes gradually as a function of frequency both before and after the shape change event, which is typical for pulsar emission. While the specific form of this frequency dependence is different after the event, its magnitude is similar.

Five main emission components can be identified: in addition to the bright central peak, there are two leading components and two trailing components, one of which is evident in the original profile only as a slight change in slope on the trailing edge of the main peak, but appears much more distinct after the shape change. After the shape change, the main peak is narrower and perhaps fainter than previously, while the leading “shoulder” component is brighter, and the trailing shoulder is broader. The second trailing component also appears to be brighter than previously, when compared to the other components. Notably, the overall amplitude of emission from the pulsar is expected to vary between observing epochs as a result of scintillation, so it is much more difficult to measure absolute changes in the brightness of particular components than it is to measure relative changes in amplitude between the components. Observations do seem to indicate that the pulsar is less bright overall after the event, but this may be due to scintillation, at least in part: the dimming was much more significant in the initial observations after the shape change than it

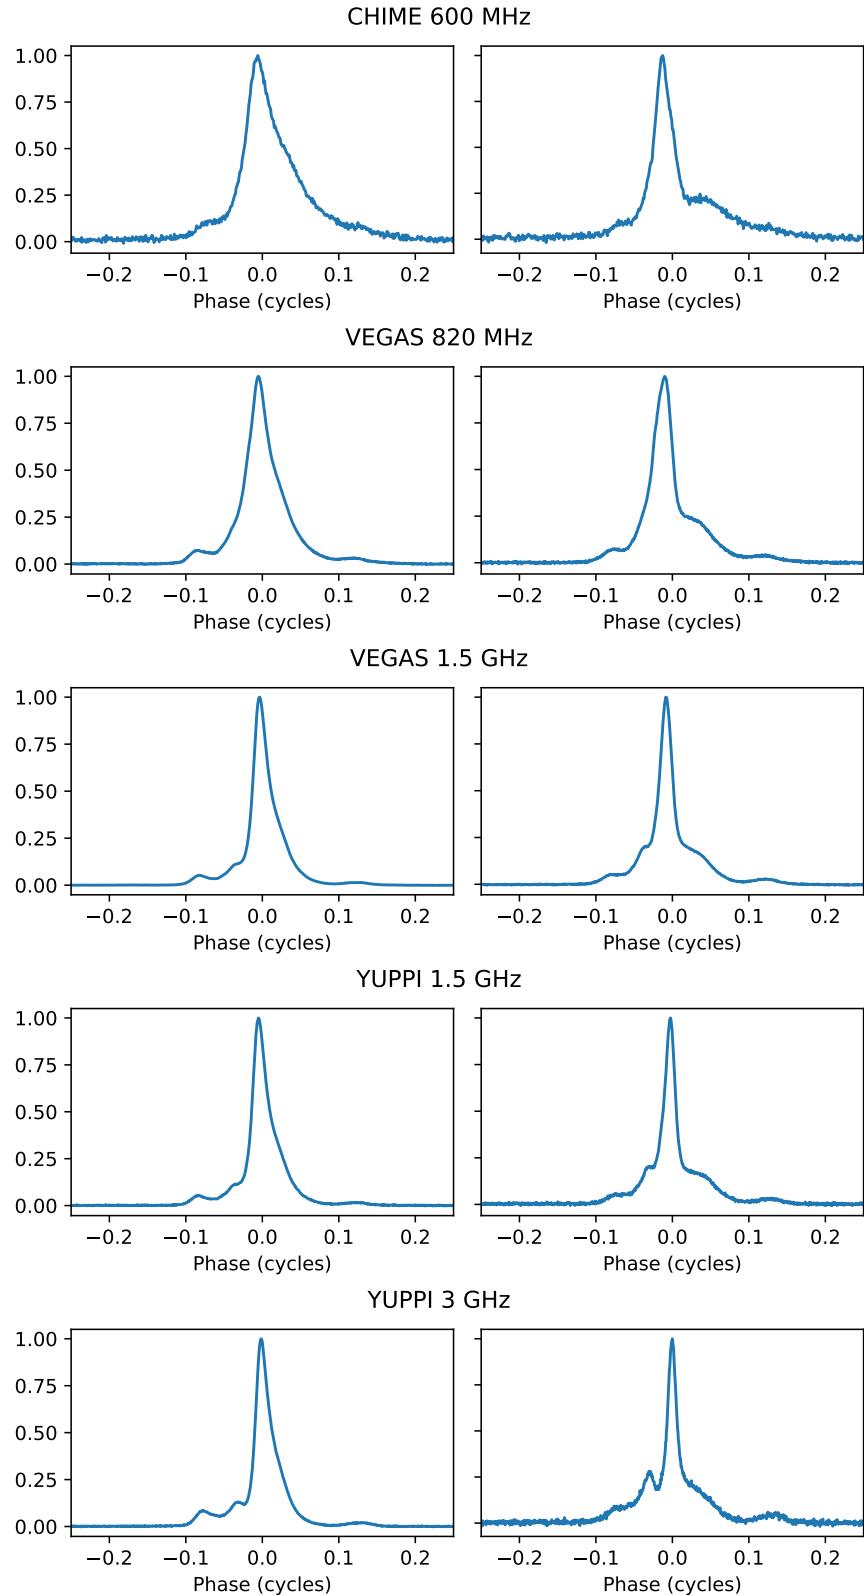


Figure 7.1: Average profiles before and after the event, demonstrating that the shape change occurred across a broad range of frequencies, and that the new profile shows a gradual frequency dependence which differs from that of the old profile.

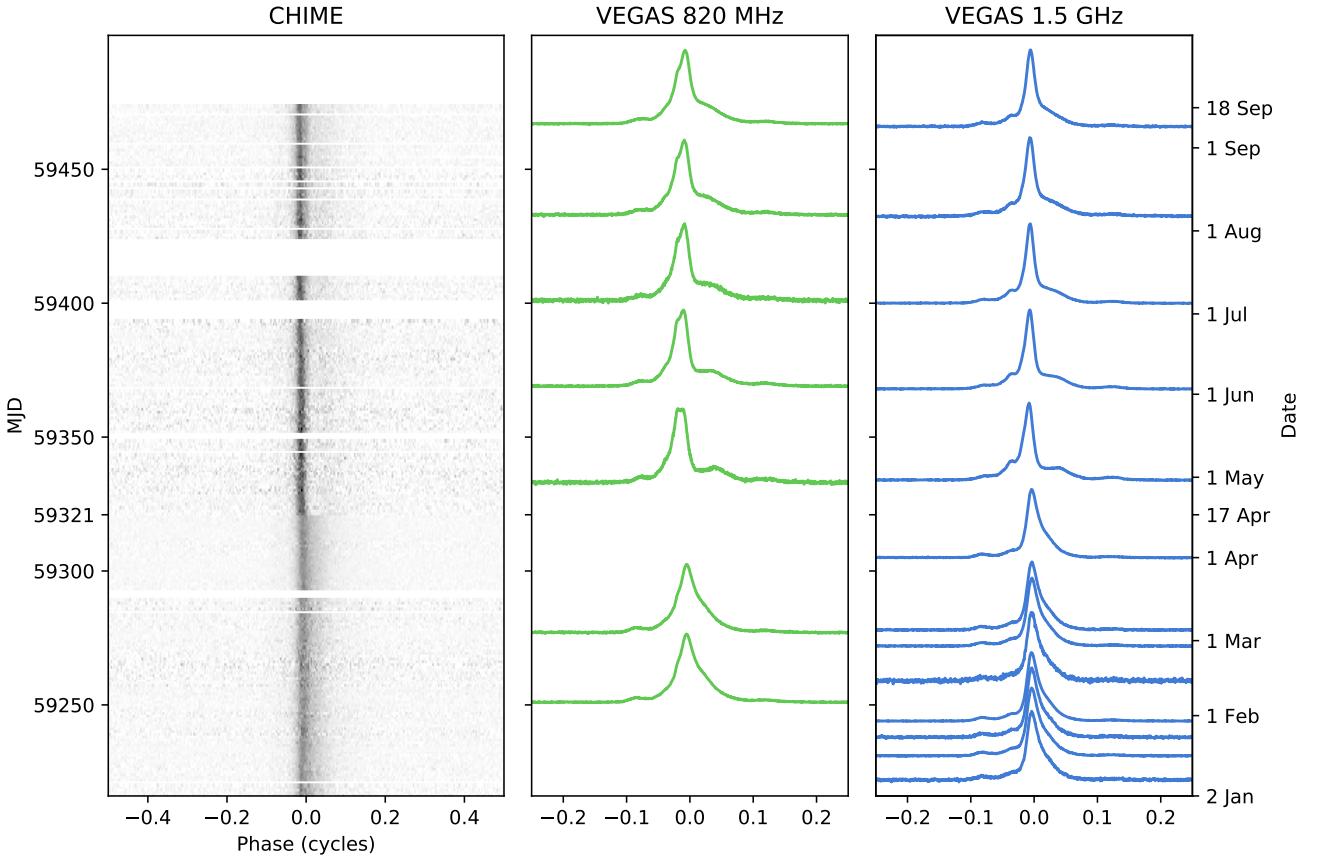


Figure 7.2: Overview of the shape change event as seen by NANOGrav and CHIME. Profiles observed at CHIME are seen in the left panel, while those observed at GBT, using the VEGAS backend, are shown in the right two panels.

is in subsequent observations.

Figure 7.2 shows that the profile shape continues to change gradually over time, even several months after the initial shape change event. In general, the pulse is changing in a direction that indicates it may be reverting to its original shape. At 1.5 GHz, the leading shoulder component is brightest, and the trailing shoulder most distinct, immediately after the shape change, with both effects tending to diminish over time. At 820 MHz, the leading shoulder appears to have briefly become approximately as bright as the main peak, and gradually declines in brightness over time, while, as at 1.5 GHz, the trailing

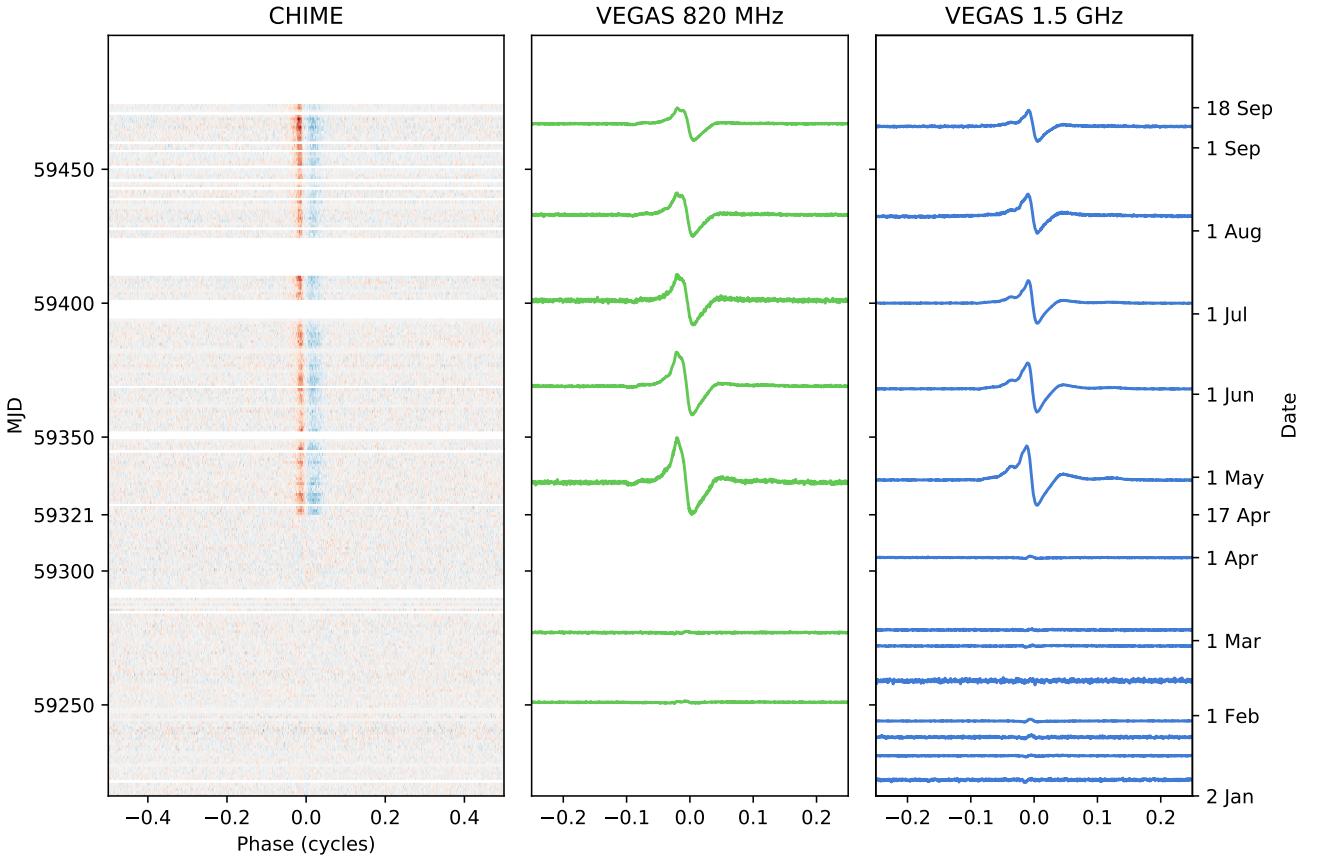


Figure 7.3: Profile residuals around the shape change event. As in Figure 7.2, CHIME observations are shown in the left panel, and GBT (VEGAS) observations are shown in the right two panels.

shoulder blends more and more with the main pulse. The tendency of the profile to revert to its original state is even clearer in Figure 7.3, where the profile residuals are normalized by dividing by the peak amplitude of the corresponding profile. The residuals can be seen to gradually decrease over time, while approximately maintaining their shape.

The TOAs and principal component dot product values shown in Figures 7.4 and 7.5 allow this to be quantified. In general, the change is more significant at lower frequencies, as can be seen by comparing the GBT 820 MHz data to the 1.5 GHz data. Going by the comparison between these

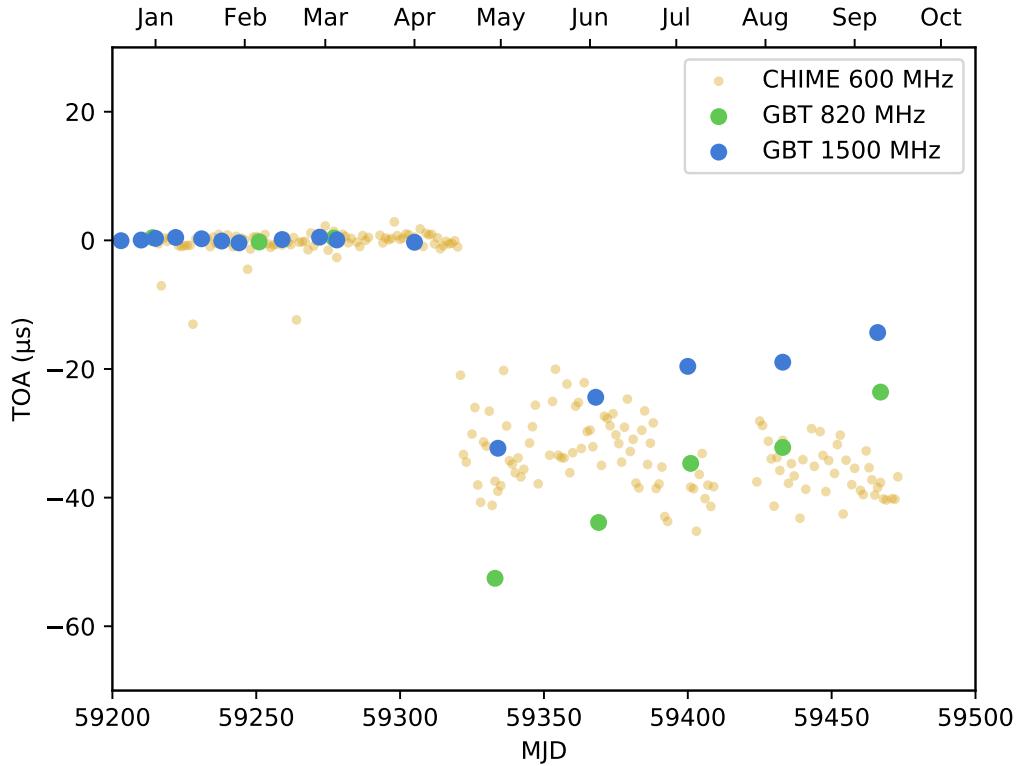


Figure 7.4: Time of arrival measurements around the shape change, after removing the best-fit timing model.

two bands alone, the corresponding change in DM would be approximately $-5 \times 10^{-3} \text{ pc cm}^{-3}$, nearly ten times larger than the change in DM associated with each of the two previous chromatic timing events, but in the same direction. In many ways, the most recent event resembles a larger version of those previous events. Not only does the change in DM have the same sign, but it appears to be decaying back toward its original value at a similar rate – some six months after the event, the chromatic TOA variations have been reduced by slightly more than half. This is approximately twice as long as it took for a similar fractional reduction in each of the two previous events.

However, the CHIME data present a somewhat different picture. Despite

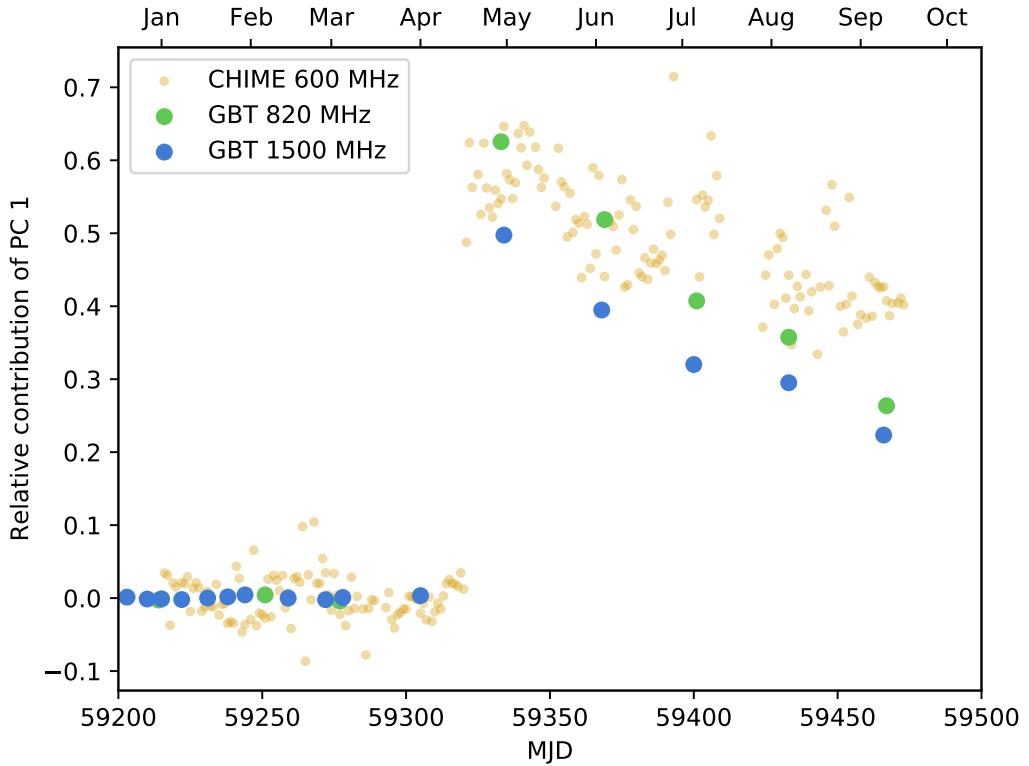


Figure 7.5: The magnitude of the shape change, as quantified by the dot product with the first principal component in each data set, over time. The dot product has been normalized by dividing it by the maximum absolute value of the associated principal component, so that the vertical axis may be thought of as the maximum fractional contribution of the principal component to the pulse shape.

having been taken at a lower frequency even than the GBT 820 MHz data, it shows a smaller TOA jump at the time of the shape change. The results shown in Figure 7.4 are averaged over the band, and, moreover, have not been barycentered, so the apparent DM is affected by the Earth’s orbital motion: the Doppler effect causes the observed frequency to differ from the frequency in the ISM. When these factors are accounted for, however, the effect appears to be larger for the higher frequencies within the CHIME band, indicating that the effect is not purely dispersive, and may be greatest at a frequency of about 800 MHz.

7.4 Conclusions

The recent shape change observed in PSR J1713+0747 may have originated either in the pulsar’s magnetosphere or as a result of propagation through discrete structures in the ISM. In several ways, it resembles a larger version of the two previously observed chromatic timing events (Demorest, 2007; Lam et al., 2018; Lin et al., 2021), which have previously been attributed to lensing of the pulsar emission by underdense regions in the ISM. As in these events, the DM, as measured between 820 and 1500 MHz, has decreased abruptly, and appears to be recovering toward its original value on a timescale of several months. Although frequency-dependent, the effect is not purely dispersive, and is accompanied by pulse shape changes that are nearly achromatic. There is some evidence that the same may be true of the two previous events.

Chromatic (i.e., radio-frequency dependent) changes in TOAs are usually thought of as a signature of ISM propagation effects in pulsar timing. However, it is difficult to produce complex shape changes in this way. Lensing may produce multiple images of the pulse that interfere with each other to produce an altered profile shape, but systematic changes in the widths and relative amplitudes of pulse components, such as those seen in the recent J1713+0747 event, would be much more likely if the event had a magnetospheric origin.

The new event also bears some resemblance to the profile shape change seen in PSR J1643–1224 in early 2015 (Shannon et al., 2016), which is thought to have had a magnetospheric origin. In that case, however, the frequency dependence was inverted, with the shape changes being stronger at higher frequencies. An event of similar form may also have occurred in PSR J1640+2224 in mid-

2012. Events like these may represent an entirely new phenomenon that will complicate millisecond pulsar timing, including pulsar timing array searches for gravitational waves. However, they have a very different signature than that expected from gravitational waves, and are unlikely to produce a false positive detection of the latter. Depending on what their origin turns out to be, these “chromatic timing events” may provide additional insight into the nature of millisecond pulsars or the structure of the interstellar medium.

CHAPTER 8

FUTURE DIRECTIONS

In this final section, I give some potential future directions for research which are enabled by the results presented previously. Here I can scarcely avoid mentioning the primary goal of NANOGrav and other PTAs, which is to detect and characterize sources of gravitational waves in the nanohertz frequency band. It is this gravitational-wave search that serves as motivation for the majority of the research discussed in this and previous chapters. Perhaps the most important future application of this research will be the eventual detection of nanohertz GWs, and the characterization of the population of supermassive black hole binaries that emit them, which has important implications for our understanding of cosmology and galaxy formation. However, since the GW analysis itself is not the primary topic here, I also outline here several examples of the kinds of non-GW science that can be carried out using these methods, both using PTA data sets and in pulsar timing more broadly construed.

With the prospect of detection on the horizon, the ability to understand PTA data sets and the processes that affect them in detail is becoming more important than ever. Conversely, these data sets are extraordinarily rich sources of information on pulsars and their astrophysical environments. In Chapters 3 and 4, I described the form of the signals expected from two related causes: hyperbolic encounters between pulsars and interstellar objects, and asteroid belts around pulsars. In Sections 8.1 and 8.2 below, I describe how these signals may be searched for in pulsar timing data. Chapters 5, 6, and 7 focused on how pulsar timing could be improved by accounting for changes in pulse shape.

Sections 8.3 and 8.4 describe two related prospects: conducting a thorough survey of profile variability in PTA data sets, and correcting for the effects of interstellar scattering.

8.1 Searching for Interstellar Objects in Pulsar Timing Data

Chapter 3 describes the form of the signal expected from a hyperbolic gravitational encounter between a pulsar and any other interstellar object (ISO) with a significant amount of mass, such as a rogue planet, large interstellar asteroid, or primordial black hole (PBH). To date, no direct search for signatures of this form has been carried out. Carrying out such a search, even if no events were detected, would allow limits to be placed on the mass and number density of such interstellar objects, which would contribute to the process of eliminating potential dark matter candidates.

Even though the MACHO, EROS, and OGLE surveys have largely ruled out dark matter candidates with masses above $10^{-8} M_{\odot}$, it remains possible that a significant fraction of dark matter is composed of PBHs with masses between 10^{-13} and $10^{-8} M_{\odot}$, and searches for hyperbolic encounters with pulsars may be able to detect such objects. Indeed, the strongest constraints on PBHs in this mass range currently come from the prospect that they may be captured by neutron stars, a process which would destroy the neutron star (Capela et al., 2013), and is essentially the most extreme example of a neutron star–PBH encounter. Lensing surveys of stars in nearby galaxies have also begun to put constraints on PBHs in this mass range (Niikura et al., 2019), but searches in pulsar timing data would serve as an independent check on those results.

The timing signatures of ISO encounters are largely ramp-like, producing a change in apparent pulse frequency due to the recoil motion of the pulsar created by the velocity change imparted to the ISO during the encounter. This means that some ISO encounters will appear similar to pulsar glitches or gravitational wave bursts with memory (BWMs). Unlike glitches, however, the changes in frequency associated with an ISO encounter have no preferred sense. Examining known glitches, especially the few known examples of inverted glitches, in more detail, could reveal whether they could be interpreted as ISO encounters for physically reasonable masses and velocities. Similarly, ISO encounters should be considered as an alternative explanation for any candidate BWMs, especially if they are seen in only a single pulsar.

A large number of pulsars are found in globular clusters (GCs). PTAs generally exclude these pulsars for reasons having to do with their unusual astrophysical environment. However, because of this environment, which contains a much higher-than-average density of stars and other massive objects, GC pulsars may be the most likely to experience ISO encounters. This means that a search for ISO encounters with well-timed GC pulsars could be particularly informative.

8.2 Distinguishing Asteroid Belts from Timing Noise

Chapter 4 establishes that the signal produced by an asteroid belt in a pulsar system takes the form of a red noise process, with an upper cutoff frequency determined by the orbit of the pulsar around its companion, if such a companion is present. Many pulsars, including MSPs, exhibit red noise of this

form. In canonical pulsars, where this noise occurs much more commonly and with a larger amplitude, it is thought to be caused by torque fluctuations in the magnetosphere. However, the much lower-amplitude process occurring in MSPs may be of a different nature, and, as demonstrated by the famous “pulsar planets” around PSR B1257+12, some MSPs are known to have small objects in orbit around them.

For MSPs with measurable red noise, it could be informative to determine the total mass and distribution of asteroids that would be necessary to produce the observed timing signal. Additional techniques could also be used to assess whether the signal could possibly originate with an asteroid belt, including searches for the signals associated with the largest individual objects, or for a high-frequency cutoff at the expected location in a binary pulsar system. Non-uniform asteroid belts could also produce spectra which deviate significantly from the power-law form expected for most other sources of red noise. One of the best targets for such a detailed study is the original millisecond pulsar, PSR B1937+21, already discussed in this context by Shannon et al. (2013), which possesses an unusually large amount of red noise for an MSP.

The red noise produced by asteroid belts is of particular consequence for PTA GW searches because the signal produced by the stochastic GW background is also expected to take the form of a red noise process. The possibility that additional red noise in the timing residuals of some pulsars is the result of asteroid belts makes it particularly important to establish that a candidate GW signal is correlated across several pulsars in sense described by Hellings & Downs (1983). If a specific pulsar were conclusively determined to possess an asteroid belt, it should almost certainly be excluded from subsequent GW

analyses because of the possible degeneracy between the asteroid belt and GW signals.

8.3 Surveying Profile Variability in PTA Data Sets

Pulsar timing array data sets, including those produced by NANOGrav and the IPTA, tend to be scrutinized mainly at the level of TOA measurements, rather than at the level of profiles. To a certain degree, this is understandable, since the TOA measurements represent a significantly smaller data volume than the raw or pre-processed profiles, while in principle retaining everything necessary to fit timing models and conduct searches for GWs and other phenomena such as the ISO encounters or asteroid belts discussed above. On the other hand, this means that information about the shape of the profiles and how closely they match the template profile used for TOA generation often goes largely ignored. We propose conducting an analysis of profile variability on all timescales, using the techniques described in Chapter 5, on a future NANOGrav or IPTA data set, making direct use of the associated timing solutions.

In a typical pulsar timing analysis, a significant number of TOAs must be discarded as outliers because of contamination by RFI. Examining the profiles used in generating the TOAs means that the determination of which points represent outliers can be made on the basis of something other than the TOAs themselves, drastically reducing the possibility of misidentifying a real but unexpected timing effect as the result of RFI contamination. However, this approach to outlier identification is uncommon, largely because of the labor involved in examining a large number of profiles by eye. For this reason, tech-

niques like those described in Chapter 5, which allow the degree to which a given profile deviates from the template to be quantified, and the effect of this deviation on timing to be estimated, could prove extremely helpful in pulsar timing analysis.

As PTAs come closer to detecting a GW background, PTA data sets are being scrutinized more and more closely. A few recent analyses, including Brook et al. (2018), have examined profiles across large PTA data sets. However, an analysis of single-pulse stochastic behavior over a large PTA dataset is still lacking. Furthermore, tying an analysis of profile shape changes directly to process of fitting timing models could potentially yield large benefits, including identification of events such as that described in Chapter 7 as well as profile-based outlier identification.

8.4 Characterizing and Mitigating Interstellar Scattering

Chapter 6 describes several techniques for mitigating the effect of TOA estimation errors due to profile shape changes. An additional phenomenon which limits pulsar timing is time-variable interstellar scattering, which is the result of multipath propagation of radio signals through the ISM. As described in the introduction to this thesis, diffractive interstellar scattering (DISS) causes the emission from a pulsar to break up into distinct patches, called scintles, in the frequency-time plane. It also leads to pulse broadening, a phenomenon which becomes increasingly important at low radio frequencies, and results in a delay typically proportional to ν^{-4} or $\nu^{-4.4}$, where ν is the radio frequency.

To properly characterize DISS, it is necessary to resolve scintles in the dy-

namic spectrum of a pulsar, and to mitigate it, one must be able to measure the impulse response function, $h_{\text{ISM}}(t)$, of the ISM. Current pulsar timing data sets are limited in both respects. The characteristic bandwidth of scintles is often less than 1 MHz, and for some pulsars it can be much smaller. This means that fully resolving scintles across bandwidths of 800 MHz or more requires high frequency resolution, which, in typical digital filterbank systems, means constructing the spectrum from larger blocks of samples. This is a manifestation of the fundamental uncertainty relation $\Delta\nu\Delta t \gtrsim 1$. Thus, frequency resolution must be traded off against pulse phase resolution, and, for pulsar observations, the latter is limited roughly by the inverse of the pulse width – if the pulse phase resolution is too coarse, the pulse will not be resolved. In addition, typical pulsar observations record only intensity-like quantities, which are sensitive only to the magnitude of $h_{\text{ISM}}(t)$. Recovering any information about the function’s phase requires baseband voltage data, which comes in volumes that are sometimes prohibitively large.

A promising new technique which has the ability to work around both of these limitations is cyclic spectroscopy, which makes use of the quasi-periodic (cyclostationary) nature of pulsar signals to describe them using a two-dimensional function known as the cyclic spectrum. The cyclic spectrum can be computed with a frequency resolution which is not limited by its pulse phase resolution because it combines information across multiple pulses coherently. Additionally, it is sensitive to the phase of $h_{\text{ISM}}(t)$, allowing the full impulse response function to be retrieved, at least in principle, from a cyclic spectrum observation, which occupies much less space than the corresponding baseband data. A dedicated cyclic spectroscopy backend is currently under development at the Green Bank Telescope. With or without cyclic spectroscopy,

more work on the timing effects of scattering is clearly needed if we are to fully understand the limitations of pulsar timing and how they might be improved.

APPENDIX A

THE NANOGRAV TOA GENERATION PROCESS

This appendix is intended to describe the process of data reduction and TOA generation used by NANOGrav, including the generation of intermediate data files, in sufficient detail to allow it to be reproduced exactly.

A.1 Getting the Software and Data

The data reduction and TOA generation process makes use of several pieces of software. The full list of requirements is:

- PSRCHIVE (<http://psrchive.sourceforge.net/>). This is the main tool used to view and manipulate data in PSRFITS format. It consists primarily of command-line utilities, written in C, but also has a Python interface and a scripting language, psrsh. Traditionally, PSRCHIVE has been installed by building it from source (see <http://psrchive.sourceforge.net/installation.shtml> for full instructions). However, successfully configuring the build tools can be challenging, especially if its dependencies are installed in non-standard locations. A much easier way to install it is using conda. Once a conda environment has been set up, PSRCHIVE can be installed simply by running

```
conda install -c conda-forge psrchive
```

Published: Jennings, R. J. 2021, NANOGrav Memorandum 007. <http://nanograv.org/assets/files/memos/NANOGrav-Memo-007.pdf>

This should automatically handle downloading and installing dependencies.

- `psrtools` (<https://github.com/demorest/psrtools>). This is a set of two command-line utilities (`autotoa` and `normalize_rms`) that depend on PSRCHIVE. Like PSRCHIVE, they are written in C, and can either be built from source or installed with conda:

```
conda install -c demorest psrtools
```

- `nanopipe` (<https://github.com/demorest/nanopipe>). This is a collection of Python and psrsh scripts used to coordinate the various steps of the NANOGrav calibration and TOA generation pipeline. It can be installed by cloning the repository from GitHub:

```
git clone  
  https://github.com/demorest/nanopipe.git  
cd nanopipe  
pip install .
```

- `toagen` (<https://gitlab.nanograv.org/nano-time/toagen>). This is a special-purpose repository which holds the full directory structure used for the calibration and TOA generation process, including configuration files which are not part of nanopipe. Currently access is limited to NANOGrav members. It is hosted on NANOGrav's internal GitLab server. If you are a NANOGrav member but have not used GitLab previously, you will have to log in to `gitlab.nanograv.org` with your NANOGrav username and password and add an SSH key (user icon in upper right > Settings > SSH keys) before you can access the repository. You can then clone it by running:

```
git clone  
git@gitlab.nanograv.org:nano-time/toagen.git
```

toagen is not an installable program or library, but rather contains a set of configuration files and Makefiles that are used in the TOA generation process, along with the current template profiles and .tim files for all NANOGrav pulsars. It does not contain the actual profile .fits files, as these are much to large to be version-controlled (see below).

- Profile data. The working copy of NANOGrav’s profile data is stored on servers at WVU and is accessible from Bowser (bowser.phys.wvu.edu) or the notebook server (notebook.nanograv.org). VEGAS, GUPPI, and PUPPI data are stored under

```
/nanograv/timing/data/
```

The data files are grouped into directories by pulsar name (without the “B” or “J” prefix), backend, and year, and are always in a directory called rawdata at the lowest level. For example, the first part of an Arecibo obervation of PSR J1713+0747 taken on March 29, 2013 can be found at

```
/nanograv/timing/data/1713+0747/puppi/2013/  
rawdata/puppi_56380_1713+0747_0542_0001.fits
```

Calibration scans are found in the same directories, and can easily be identified by the presence of cal in the filenames. Flux calibrator obser-vations are found in directories corresponding to the name of the flux calibrator.¹

¹A quasar, usually B1442+101 = J1445+099, but at Arecibo sometimes J1413+151, and at GBT 3C190 is used for one 2011 observation. There are multiple directories with slight variations on the name, and some with suffixes, but the most recent calibrator observations all seem to go into 1442+101.

Some data from other telescopes used by NANOGrav are also available on the WVU servers, although not in the same place as the GBT and Arecibo data. YUPPI (VLA) data are stored under

```
/hyrule/data/users/pdemores/VLA/
```

and are grouped by semester and proposal number. Some CHIME data, grouped by pulsar, are also available, under

```
/hyrule/data/CHIME/NANOGrav/
```

As of September 29, 2021, only observations of PSR J1713+0747 from January to September of 2021 are available, but additional data should be copied over soon.

A.2 Setup and Directory Structure

The basic processing scripts are contained in `nanopipe`. As established in the instructions for `nanopipe` (https://github.com/demorest/nanopipe/blob/master/doc/basic_instructions.txt), the work takes place in a base directory which has subdirectories for each pulsar being processed. Within `toagen`, there is a separate base directory for each backend (ASP, GASP, PUPPI, GUPPI, YUPPI).

Flux calibrator observations are stored in a separate directory, called `fluxcal`, inside the base directory for each backend. Before processing the rest of the observations, one needs to create `.fcal` files from the flux calibrator observations. This is done using the PSRCHIVE tool `fluxcal`:

```
fluxcal -f -e fcal
```

The results can be checked using the PSRCHIVE calibrator viewer, pacv.

The base directory for each backend also contains several psrsh scripts (copied from nanopipe’s config directory): the basic list is

```
zap_minmax
zap_and_tscrunch
update_be_delay.psrsh
process_fluxcal
img_$be
```

Here \$be is the backend name.

There should also be a configuration file called `make_psr_make.config.py` in the base directory for each backend, containing, at minimum, these lines:

```
ver_id = "$project_id"
tscrunch_arg = "-J ../zap_and_tscrunch"
```

Here \$project_id is a string identifying the particular version of the processing being carried out.

This basic setup is customized in a backend-dependent way in the real NANOGrav TOA generation process, in ways that are captured in the toagen repository. Once the scripts and configuration files are in place, a makefile can be generated for each pulsar by running (from within the corresponding subdirectory)

```
make_psr_make $basedir > Make.psr
```

For \$basedir one should substitute the full path of the base directory. The

makefile can be used to produce TOAs by running

```
make -f Make.psr toas
```

(again from within the pulsar subdirectory). This command can be altered to run several processes in parallel by adding the option `-j $n_procs`, where `$n_procs` is the number of processes, or by replacing `toas` with another makefile target from this list:

```
rf
calibration
zap
scrunch
templates
toas
```

to only run part of the analysis.

A.3 Per-Pulsar Makefiles

The main process is laid out in the per-pulsar makefiles (`Make.psr`), which are generated by the `make_psr_make` script in `nanopipe`. The processing carried out by `Make.psr` breaks down into the following seven steps, most of which correspond directly to makefile targets:

1. Target `rf`: Combine the raw `.fits` files into `.rf` files. The raw data files must be copied in from where they are stored on WVU’s servers. Several corrections are applied at this stage, namely:

- (a) Epoch error due to polyco REF_MJD precision in PSRFITS is corrected. This accounts for the effects of a bug in PSRCHIVE, discovered during the 9-year analysis, which meant that polyco reference MJDs were not stored using enough numerical precision.
- (b) Backend delays (stored in the BE_DELAY PSRFITS header item) are set appropriately. These measure extra signal latency within each backend, usually just digital filter latency, and are determined from the backend design (not measured experimentally). For GUPPI and PUPPI, the delay is a function of the number of channels and channel bandwidth in the observation.
- (c) ADC ghost image corrections (Alam et al. 2021a, §2.3.1; cf. Kurosawa et al. 2001) are applied.
- (d) Par files are installed into the FITS header. This also means realigning the profiles according to the par file being installed. Usually the “predictive” par file from the previous data set is used here, if it is available.
- (e) The DM value in the FITS header is set appropriately. This DM value is taken from the par file used above, and is used for de-dispersion prior to frequency averaging.
- (f) Specified bands are zapped and time-frequency zapping (Offringa et al., 2010) is applied.
- (g) The source and receiver names are fixed to make them uniform across the data set.

If this process fails, the file is noted in the list of cal failures. Finally, the database of calibrator archives is built up.

Commands:

```
psradd -T -j 'fix refmjd' -J
  ./update_be_delay.psrsh -J ./img_$be -j
  'install par $psr.basic.par' -j "e dm=$dm" -J
  ./zap_and_tscrunch_with_list_$tel -j "e
  rcvr:name='fix_receiver_name $obsj.fits'" -j
  "e name='get_proper_name $obsj.fits'" -o
  $obs.rf $obs1.fits $obs2.fits

psradd -T -j 'fix refmjd' -J
  ./update_be_delay.psrsh -J ./img_$be -J
  ./zap_and_tscrunch_with_list_$tel -j "e
  rcvr:name='fix_receiver_name $calj.fits'" -j
  "e name='get_proper_name $calj.fits'" -o
  $cal.cf $cal1.fits $cal2.fits
```

These commands (and the others quoted below) are taken from the Make.ps1 file, but reformatted to make them independently runnable and as generic as possible. The psrsh scripts

```
update_be_delay.psrsh
img_puppi
img_guppi
zap_and_tscrunch_with_list_gb
zap_and_tscrunch_with_list_ao
```

can be found in nanopipe's config directory (however, the ones found in nanopipe use zap median, as was done for the 12.5-year and earlier data sets, rather than zap tfzap, used in toagen and the 15-year data set). The

Python scripts

```
fix_receiver_name  
get_proper_name
```

can be found in nanopipe's scripts directory.

2. Target calibration: Calibrate the data (.rf files) to produce .calib files.

This performs a very basic flux/polarization calibration (cal gain and phase correction) using pac. Matching with the calibrator observations is done on a per-channel basis, and the pulse phase of the cal transition is set to 0.5 turns. Calibrator Stokes parameters are derived from the flux-cal data. If this process fails, the file is noted in the list of cal failures. Then the data goes through another round of RFI zapping, removing a specified fraction of the band edge, as well as more specified frequency bands, but these seem to mostly be duplicates. Any zapping options in the variable \$XPAZ are applied. Finally, if the file \$obs.calib does not exist, the file \$obs.calibP is removed.

Commands:

```
pac -w  
pac -a -j "config  
    SquareWave::transition_phase=0.5" -x -e  
    x.calib -d database.txt $obs.rf  
    paz -m -E2.0 $zapchannels $(XPAZ) $obs.x.calib
```

\$zapchannels is a space-separated list of -F options specifying channels to zap on a per-receiver basis, configured in make_psr_make.

3. Target zap: Apply “min-max” zapping to produce .zap files.

Command:

```
psrsh -e zap ./zap_minmax $obs.x.calib
```

The psrsh script `zap_minmax` can be found in `nanopipe`'s `config` directory. (But, just as for `zap_and_tscrunch_with_list_$tel` above, this is an older version of `zap_minmax` that uses `zap median` rather than `zap tfzap`. The `zap tfzap` version can be found in `toagen`.)

4. Target scrunch (1/2): Scale the weights so that the offpulse RMS of the weighted data is 1, producing `.norm` files.

Command:

```
normalize_rms -w $obs.x.zap
```

5. Target scrunch (2/2): Frequency average to 64 channels, and time average to a number of subintegrations set by the observation length (the length divided by 1800 s, rounded down, plus one), to produce `.ff` files.

Command:

```
pam -e ff -f$factor --setnsub 'psredit -Q -c
length $obs.x.norm | awk '{print
int($2/1800.0) + 1}'' $obs.x.norm
```

`$factor` varies depending on the backend. It is 8 for receivers with 512 channels (GUPPI and PUPPI L-band, PUPPI S-band) and 2 for GUPPI 820 MHz, which has 128 channels.

6. Target templates: Iteratively determine the template and TOAs using `autotoa` (an implementation of the method outlined in Chapter 2 of Demorest 2007), with a list of all `.ff` files for this pulsar-frontend-backend combination as input. Use a Gaussian of width 0.1 as the initial guess

for the template, and perform at most 3 iterations. Then rotate the final profile by half a turn, and apply UD8 wavelet smoothing to it.

Commands:

```
for f in {$obs.ff}; do echo $f >> $rcvr.fflist
autotoa -g0.1 -i3 -S $rcvr.sum $rcvr.fflist
pam -r0.5 -m $rcvr.sum
psrsmooth -W -t UD8 $rcvr.sum
```

7. Target toas: Use pat to calculate the final narrowband TOAs from the .ff files and template and write them to a .tim file in tempo2/IPTA format, adding appropriate flags. Calculate errors using MCMC. Then use the same .ff files to calculate wideband TOAs based on the 12.5-year wideband templates, and write them to a .tim file with a similar format and flags.

Commands:

```
echo 'MODE 1' > $rcvr.nb.tim
pat -A FDM -e err=num -C chan -C subint -C snr
-C wt -C flux -C fluxe -f "tempo2 IPTA" -X
"-proc 15y -pta NANOGrav -ver TEMP $(XFLAG)"
-s $rcvr.sum.sm -M $rcvr.fflist >>
$rcvr.nb.tim
pptoas.py --print-flux --quiet
--flags=proc,15y,pta,NANOGrav,ver,TEMP$(XFLAG)
-m $rcvr.12y.x.avg_port.spl -d $rcvr.fflist
-o $rcvr.wb.tim
```

Acknowledgements

The `toagen` repository was assembled by the members of the NANOGrav git-flowers team, including but not limited to Joe Swiggum, Paul Demorest, Anne Archibald, Deborah Good, Adam Brazier, Nate Garver-Daniels, and Joe Glaser. I thank Paul Demorest, who developed the original version of the TOA generation process and wrote several of the software tools used to implement it, for answering a series of questions that made the current form of this document possible. I also thank Michael Lam and Haley Wahl for providing feedback on early drafts.

APPENDIX B

TRANSMISSION FUNCTIONS FOR POLYNOMIAL FITS

Fitting a polynomial model to a time series and subtracting it removes variance from the data at the lowest frequencies. How much variance is removed is quantified by the transmission function, $\mathcal{T}(f)$, which is the ratio of the variance remaining at frequency f to the initial variance at the same frequency. In the pulsar timing array context, this formalism was introduced Blandford et al. (1984), and further developed by Cordes (2013a) and Madison et al. (2013). Here we will compute the transmission function for polynomial fits of various orders.

Let us consider the idealized case of data sampled continuously over a finite observing time, T . For convenience, we will set $t = 0$ at the midpoint of the observation, so that the data is represented by a function, $x(t)$, supported on the interval $[-T/2, T/2]$. Appropriately scaled versions of the Legendre polynomials $P_n(x)$ form a complete orthogonal basis for the space of (appropriately well-behaved) functions of this form. In particular, we have

$$\int_{-T/2}^{-T/2} P_m\left(\frac{2t}{T}\right) P_n\left(\frac{2t}{T}\right) dt = \frac{T\delta_{mn}}{2n+1}. \quad (\text{B.1})$$

It follows that we can expand $x(t)$ in a series of the form

$$x(t) = \sum_{k=0}^{\infty} C_k P_k\left(\frac{2t}{T}\right), \quad (\text{B.2})$$

where the coefficients C_k are given by

$$C_k = \frac{2k+1}{T} \int_{-T/2}^{-T/2} P_k\left(\frac{2t}{T}\right) y(t) dt. \quad (\text{B.3})$$

Published: Jennings, R. J. 2021, NANOGrav Memorandum 006. <http://nanograv.org/assets/files/memos/NANOGrav-Memo-006.pdf>

The sum of the first n terms in this series is the polynomial of degree $n - 1$ most closely approximating $x(t)$ in the least-squares sense – in other words, the result of least-squares fitting a polynomial of degree $n - 1$ to $x(t)$.

Because fitting and subtracting a polynomial is linear in $x(t)$, we can consider its effect on each frequency separately, so we need only consider the case where $x(t) = e^{2\pi i f t}$. In this case, the coefficients $C_k(f)$ are given by

$$C_k(f) = \frac{2k+1}{T} \int_{-T/2}^{-T/2} P_k\left(\frac{2t}{T}\right) e^{2\pi i f t} dt. \quad (\text{B.4})$$

The total variance of the original signal, $x(t) = e^{2\pi i f t}$, is given by

$$\sigma_0^2(f) = \frac{1}{T} \int_{-T/2}^{T/2} |e^{2\pi i f t}|^2 dt = 1, \quad (\text{B.5})$$

whereas the variance of the approximating polynomial

$$\hat{x}(t) = \sum_{k=0}^n C_k(f) P_k\left(\frac{2t}{T}\right) \quad (\text{B.6})$$

is given by

$$\sigma^2(f) = \frac{1}{T} \int_{-T/2}^{T/2} |\hat{x}(t)|^2 dt = \sum_{k=0}^n \frac{|C_k(f)|^2}{2k+1} \quad (\text{B.7})$$

It follows that the transmission function, $\mathcal{T}_n(f)$, for fitting and subtracting a polynomial of degree n is

$$\mathcal{T}_n(f) = \frac{\sigma_0^2(f) - \sigma^2(f)}{\sigma_0^2(f)} = 1 - \sum_{k=0}^n \frac{|C_k(f)|^2}{2k+1}. \quad (\text{B.8})$$

Equation (B.4) can be simplified by replacing t and f with the dimensionless variables $u = 2t/T$ and $v = \pi f T$, in which case it becomes

$$C_k(v) = \left(k + \frac{1}{2}\right) \int_{-1}^1 P_k(u) e^{iuv} du, \quad (\text{B.9})$$

In particular, we have

$$C_0(v) = \frac{1}{2} \int_{-1}^1 e^{iuv} du = \frac{\sin v}{v}, \quad (\text{B.10})$$

$$C_1(v) = \frac{3}{2} \int_{-1}^1 ue^{iuv} du = \frac{3i(\sin v - v \cos v)}{v^2}, \quad (\text{B.11})$$

$$C_2(v) = \frac{5}{4} \int_{-1}^1 (3u^2 - 1) e^{iuv} du = \frac{5[(v^2 - 3) \sin v + 3v \cos v]}{v^3}. \quad (\text{B.12})$$

Explicit expressions for $\mathcal{T}_n(v)$ for $n = 0, 1, 2$ may be obtained by substituting these results into equation (B.8). In particular,

$$\mathcal{T}_0(v) = 1 - \frac{\sin^2 v}{v^2}, \quad (\text{B.13})$$

$$\mathcal{T}_1(v) = \mathcal{T}_0(v) - \frac{3(\sin v - v \cos v)^2}{v^4}, \quad (\text{B.14})$$

$$\mathcal{T}_2(v) = \mathcal{T}_1(v) - \frac{5[(v^2 - 3) \sin v + 3v \cos v]^2}{v^6}. \quad (\text{B.15})$$

Qualitatively, each of these transmission functions rises from $\mathcal{T} = 0$ at $v = 0$ to oscillate just under $\mathcal{T} = 1$ at large values of v . We can find the asymptotic behavior of the transmission function at low frequencies by expanding each $C_k(v)$ around $v = 0$. This gives

$$C_0(v) = 1 - \frac{v^2}{6} + \frac{v^4}{120} - \frac{v^6}{5040} + \mathcal{O}(v^8), \quad (\text{B.16})$$

$$C_1(v) = v - \frac{v^3}{10} + \frac{v^5}{280} + \mathcal{O}(v^7), \quad (\text{B.17})$$

$$C_2(v) = -\frac{v^2}{3} + \frac{v^4}{42} + \mathcal{O}(v^6). \quad (\text{B.18})$$

Substituting these series expansions into equation (B.8), we obtain

$$\mathcal{T}_0(v) = \frac{v^2}{3} - \frac{2v^4}{45} + \frac{v^6}{315} + \mathcal{O}(v^8), \quad (\text{B.19})$$

$$\mathcal{T}_1(v) = \frac{v^4}{45} - \frac{4v^6}{1575} + \mathcal{O}(v^8), \quad (\text{B.20})$$

$$\mathcal{T}_2(v) = \frac{v^6}{1575} + \mathcal{O}(v^8). \quad (\text{B.21})$$

This means that, near $f = 0$, we have

$$\mathcal{T}_0(f) \sim \frac{\pi^2 T^2}{3} f^2, \quad (\text{B.22})$$

$$\mathcal{T}_1(f) \sim \frac{\pi^4 T^4}{45} f^4, \quad (\text{B.23})$$

$$\mathcal{T}_2(f) \sim \frac{\pi^6 T^6}{1575} f^6. \quad (\text{B.24})$$

One could continue this process to find the asymptotic behavior of $\mathcal{T}_n(f)$ for higher values of n , but it turns out there's a way to solve the problem for arbitrary n . First, note that $C_k(v) = \mathcal{O}(v^k)$ for each value of k considered so far. This is true in general: expanding the right-hand side of equation (B.9) in a power series in v gives

$$C_k(v) = \left(k + \frac{1}{2}\right) \sum_{n=0}^{\infty} \frac{(iv)^n}{n!} \int_{-1}^1 u^n P_k(u) du, \quad (\text{B.25})$$

but since $P_k(u)$ is orthogonal to all polynomials of degree less than k , all terms with $n < k$ are zero.

Since we can write

$$\sum_{k=0}^{\infty} C_k(v) P_k(u) = e^{iuv}, \quad (\text{B.26})$$

it follows that

$$\begin{aligned} \sum_{k=0}^n C_k(v) P_k(u) &= e^{iuv} - \sum_{k=n+1}^{\infty} C_k(v) P_k(u) \\ &= e^{iuv} + \mathcal{O}(v^{n+1}) \\ &= \sum_{k=0}^n \frac{(iuv)^k}{k!} + \mathcal{O}(v^{n+1}). \end{aligned} \quad (\text{B.27})$$

Up to terms of order v^{n+1} , each side of this equation is a polynomial of degree n in u . Equating the leading terms gives

$$\frac{1}{2^n} \binom{2n}{n} u^n C_n(v) = \frac{i^n u^n v^n}{n!} + \mathcal{O}(v^{n+1}), \quad (\text{B.28})$$

where we have used a well-known expression for the leading coefficient of $P_n(u)$. It follows that

$$C_n(v) = \frac{(iv)^n}{(2n-1)!!} + \mathcal{O}(v^{n+1}). \quad (\text{B.29})$$

Computing the total variance on both sides of equation (B.26) gives

$$\sum_{k=0}^{\infty} \frac{|C_k(v)|^2}{2k+1} = 1. \quad (\text{B.30})$$

Combining this with equation (B.8) gives

$$\mathcal{T}_n(v) = \sum_{k=n+1}^{\infty} \frac{|C_k(v)|^2}{2k+1} = \frac{|C_{n+1}(v)|^2}{2n+3} + \mathcal{O}(v^{n+2}). \quad (\text{B.31})$$

It follows that, near $v = 0$,

$$\mathcal{T}_n(v) = \frac{v^{2n+2}}{(2n+3)(2n+1)!!^2} + \mathcal{O}(v^{n+2}), \quad (\text{B.32})$$

or, in terms of f ,

$$\mathcal{T}_n(f) \sim \frac{(\pi T)^{2n+2}}{(2n+3)(2n+1)!!^2} f^{2n+2}. \quad (\text{B.33})$$

Substituting $n = 0, 1, 2$ into equation (B.33) reproduces the results of equations (B.22), (B.23), and (B.24). For $n = 3, 4$ we find

$$\mathcal{T}_3(f) \sim \frac{\pi^8 T^8}{99\,225} f^8, \quad (\text{B.34})$$

$$\mathcal{T}_4(f) \sim \frac{\pi^{10} T^{10}}{9\,823\,275} f^{10}. \quad (\text{B.35})$$

It is also possible to relax the assumption that the data are sampled continuously and uniformly, at least for $\mathcal{T}_0(f)$ and $\mathcal{T}_1(f)$. Suppose that the data are instead sampled at N discrete times t_i . For convenience, we can assume that the mean of the sample times is zero:

$$\frac{1}{N} \sum_{i=1}^N t_i = 0. \quad (\text{B.36})$$

Let σ^2 , γ , and κ be the variance, skewness, and kurtosis of the sample times, i.e.,

$$\frac{1}{N} \sum_{i=1}^N t_i^2 = \sigma^2, \quad (\text{B.37})$$

$$\frac{1}{N} \sum_{i=1}^N t_i^3 = \gamma \sigma^3, \quad (\text{B.38})$$

$$\frac{1}{N} \sum_{i=1}^N t_i^4 = \kappa \sigma^4. \quad (\text{B.39})$$

We can define an inner product between functions sampled at the times t_i by

$$\langle x(t), y(t) \rangle = \frac{1}{N} \sum_{i=1}^N x(t_i) y(t_i). \quad (\text{B.40})$$

The three polynomials

$$p_0(t) = 1, \quad (\text{B.41})$$

$$p_1(t) = t, \quad (\text{B.42})$$

$$p_2(t) = t^2 - \gamma \sigma t - \sigma^2 \quad (\text{B.43})$$

are orthogonal with respect to this inner product; that is, $\langle p_m(t), p_n(t) \rangle = 0$ whenever $m \neq n$. They also satisfy

$$\langle p_0(t)^2 \rangle = 1, \quad (\text{B.44})$$

$$\langle p_1(t)^2 \rangle = \sigma^2, \quad (\text{B.45})$$

$$\langle p_2(t)^2 \rangle = (\kappa - \gamma^2 - 1) \sigma^4. \quad (\text{B.46})$$

Just as in the continuously sampled case, we can expand an arbitrary function $x(t)$ in terms of these polynomials, writing

$$x(t) = c_0 p_0(t) + c_1 p_1(t) + c_2 p_2(t) + r(t), \quad (\text{B.47})$$

where the remainder $r(t)$ is orthogonal to each of the previous terms, and the coefficients are given by $c_k = \langle x(t), p_k(t) \rangle$. As before, to compute the transmission function, we must consider $x(t) = e^{2\pi i f t}$. In this case, the coefficients are

given by

$$c_0(f) = \frac{1}{N} \sum_{j=1}^N e^{2\pi i f t_j}, \quad (\text{B.48})$$

$$c_1(f) = \frac{1}{N} \sum_{j=1}^N t_j e^{2\pi i f t_j}, \quad (\text{B.49})$$

$$c_2(f) = \frac{1}{N} \sum_{j=1}^N (t_j^2 - \gamma \sigma t_j - \sigma^2) e^{2\pi i f t_j}. \quad (\text{B.50})$$

The full transmission functions $\mathcal{T}_0(f)$, $\mathcal{T}_1(f)$, and $\mathcal{T}_2(f)$ can be written

$$\mathcal{T}_0(f) = 1 - |c_0(f)|^2, \quad (\text{B.51})$$

$$\mathcal{T}_1(f) = 1 - |c_0(f)|^2 - \frac{|c_1(f)|^2}{\sigma^2}, \quad (\text{B.52})$$

$$\mathcal{T}_2(f) = 1 - |c_0(f)|^2 - \frac{|c_1(f)|^2}{\sigma^2} - \frac{|c_2(f)|^2}{(\kappa - \gamma^2 - 1)\sigma^4}, \quad (\text{B.53})$$

but cannot be expressed in any simpler form. However, in the low frequency limit, the expressions do simplify somewhat. In particular, expanding $e^{2\pi i f t}$ in a power series around $f = 0$, we have

$$c_0(t) = \frac{1}{N} \sum_{j=1}^N \left(1 + 2\pi i f t_j - 2\pi^2 f^2 t_j^2 - \frac{4\pi^3}{3} i f^3 t_j^3 + \frac{2\pi^4}{3} f^4 t_j^4 + \mathcal{O}(f^5) \right) \quad (\text{B.54})$$

$$= 1 - 2\pi^2 \sigma^2 f^2 - \frac{4\pi^3}{3} i \gamma \sigma^3 f^3 + \frac{2\pi^4}{3} \kappa \sigma^4 f^4 + \mathcal{O}(f^5),$$

$$c_1(t) = \frac{1}{N} \sum_{j=1}^N \left(t_j + 2\pi i f t_j^2 - 2\pi^2 f^2 t_j^3 - \frac{4\pi^3}{3} i f^3 t_j^4 + \mathcal{O}(f^4) \right) \quad (\text{B.55})$$

$$= 2\pi i \sigma^2 f - 2\pi i \gamma \sigma^3 f^2 - \frac{4\pi^3}{3} i \kappa \sigma^4 f^3 + \mathcal{O}(f^4).$$

Substituting these results into equations (B.51) and (B.52) gives

$$\mathcal{T}_0(f) \sim 4\pi^2 \sigma^2 f^2, \quad (\text{B.56})$$

$$\mathcal{T}_1(f) \sim 4\pi^4 (\kappa - \gamma^2 - 1) \sigma^4 f^4. \quad (\text{B.57})$$

Using the variance $\sigma^2 = T^2/12$, skewness $\gamma = 0$, and kurtosis $\kappa = 9/5$ for a uniform distribution in these expressions reproduces equations (B.22) and (B.23).

A similar result is possible for $\mathcal{T}_2(f)$, showing that it is asymptotically proportional to f^6 , but with a much more complicated coefficient.

BIBLIOGRAPHY

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Physical Review Letters*, 116, 061102, doi: 10.1103/PhysRevLett.116.061102
- . 2017, *The Astrophysical Journal, Letters*, 848, L13, doi: 10.3847/2041-8213/aa920c
- . 2019, *Physical Review X*, 9, 031040, doi: 10.1103/PhysRevX.9.031040
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, *Physical Review X*, 11, 021053, doi: 10.1103/PhysRevX.11.021053
- Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *The Astrophysical Journal, Supplement*, 208, 17, doi: 10.1088/0067-0049/208/2/17
- Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021a, *The Astrophysical Journal, Supplement*, 252, 4, doi: 10.3847/1538-4365/abc6a0
- . 2021b, *The Astrophysical Journal, Supplement*, 252, 5, doi: 10.3847/1538-4365/abc6a1
- Alcock, C., Allsman, R. A., Alves, D. R., et al. 2001, *The Astrophysical Journal*, 550, L169, doi: 10.1086/319636
- Althaus, L. G., & Benvenuto, O. G. 1998, *Monthly Notices of the RAS*, 296, 206, doi: 10.1046/j.1365-8711.1998.01332.x
- Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv e-prints, arXiv:1702.00786. <https://arxiv.org/abs/1702.00786>
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Science*, 340, 448, doi: 10.1126/science.1233232

- Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, *Science*, 324, 1411, doi: 10.1126/science.1172740
- Archibald, R. F., Kaspi, V. M., Ng, C. Y., et al. 2013, *Nature*, 497, 591, doi: 10.1038/nature12159
- Arzoumanian, Z., Cordes, J. M., & Wasserman, I. 1999, *The Astrophysical Journal*, 520, 696, doi: 10.1086/307482
- Arzoumanian, Z., Nice, D. J., Taylor, J. H., & Thorsett, S. E. 1994, *The Astrophysical Journal*, 422, 671, doi: 10.1086/173760
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2014, *The Astrophysical Journal*, 794, 141, doi: 10.1088/0004-637X/794/2/141
- . 2018a, *The Astrophysical Journal, Supplement*, 235, 37, doi: 10.3847/1538-4365/aab5b0
- Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018b, *The Astrophysical Journal*, 859, 47, doi: 10.3847/1538-4357/aabd3b
- Baade, W., & Zwicky, F. 1934, *Proceedings of the National Academy of Science*, 20, 259, doi: 10.1073/pnas.20.5.259
- Backer, D. C. 1970a, *Nature*, 228, 42, doi: 10.1038/228042a0
- . 1970b, *Nature*, 228, 1297, doi: 10.1038/2281297a0
- Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, *Nature*, 300, 615, doi: 10.1038/300615a0
- Baggaley, W. J. 2000, *Journal of Geophysics Research*, 105, 10353, doi: 10.1029/1999JA900383

- Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Mantelet, G., & Andrae, R. 2018, ArXiv e-prints, arXiv:1804.10121. <https://arxiv.org/abs/1804.10121>
- Bartel, N., Morris, D., Sieber, W., & Hankins, T. H. 1982, The Astrophysical Journal, 258, 776, doi: 10.1086/160125
- Bassa, C. G., Janssen, G. H., Stappers, B. W., et al. 2016, Monthly Notices of the RAS, 460, 2207, doi: 10.1093/mnras/stw1134
- Bates, S. D., Thornton, D., Bailes, M., et al. 2015, Monthly Notices of the RAS, 446, 4019, doi: 10.1093/mnras/stu2350
- Behrens, E. A., Ransom, S. M., Madison, D. R., et al. 2020, The Astrophysical Journal, Letters, 893, L8, doi: 10.3847/2041-8213/ab8121
- Bell, J. F., & Bailes, M. 1996, The Astrophysical Journal, Letters, 456, L33, doi: 10.1086/309862
- Bellm, E. C., Kaplan, D. L., Breton, R. P., et al. 2016, The Astrophysical Journal, 816, 74, doi: 10.3847/0004-637X/816/2/74
- Biggs, J. D. 1992, The Astrophysical Journal, 394, 574, doi: 10.1086/171608
- Bizouard, M. A., Jenet, F., Price, R., & Will, C. M. 2013, Classical and Quantum Gravity, 30, 220301. <http://stacks.iop.org/0264-9381/30/i=22/a=220301>
- Blandford, R., Narayan, R., & Romani, R. W. 1984, Journal of Astrophysics and Astronomy, 5, 369, doi: 10.1007/BF02714466
- Breton, R. P., van Kerkwijk, M. H., Roberts, M. S. E., et al. 2013, The Astrophysical Journal, 769, 108, doi: 10.1088/0004-637X/769/2/108
- Brook, P. R., Karastergiou, A., McLaughlin, M. A., et al. 2018, The Astrophysical Journal, 868, 122, doi: 10.3847/1538-4357/aae9e3

- Bryden, G., Beichman, C. A., Rieke, G. H., et al. 2006, *The Astrophysical Journal*, 646, 1038, doi: 10.1086/504976
- Buch, J., Leung, S. C., & Fan, J. 2019, *Journal of Cosmology and Astro-Particle Physics*, 2019, 026, doi: 10.1088/1475-7516/2019/04/026
- Caballero, R. N., Guo, Y. J., Lee, K. J., et al. 2018, *Monthly Notices of the RAS*, 481, 5501, doi: 10.1093/mnras/sty2632
- Callanan, P. J., van Paradijs, J., & Rengelink, R. 1995, *The Astrophysical Journal*, 439, 928, doi: 10.1086/175229
- Camilo, F., Lorimer, D. R., Freire, P., Lyne, A. G., & Manchester, R. N. 2000, *The Astrophysical Journal*, 535, 975, doi: 10.1086/308859
- Camilo, F., Reynolds, J. E., Ransom, S. M., et al. 2016, *The Astrophysical Journal*, 820, 6, doi: 10.3847/0004-637X/820/1/6
- Capela, F., Pshirkov, M., & Tinyakov, P. 2013, *Physical Review D*, 87, 123524, doi: 10.1103/PhysRevD.87.123524
- Carr, B., Kühnel, F., & Sandstad, M. 2016, *Physical Review D*, 94, 083504, doi: 10.1103/PhysRevD.94.083504
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016a, *ArXiv e-prints*. <https://arxiv.org/abs/1612.05560>
- . 2016b, *arXiv e-prints*, arXiv:1612.05560. <https://arxiv.org/abs/1612.05560>
- Charnoz, S., & Morbidelli, A. 2003, *Icarus*, 166, 141, doi: 10.1016/S0019-1035(03)00213-6

- Chatterjee, S., Vlemmings, W. H. T., Brisken, W. F., et al. 2005, *The Astrophysical Journal, Letters*, 630, L61, doi: 10.1086/491701
- Chatterjee, S., Brisken, W. F., Vlemmings, W. H. T., et al. 2009, *The Astrophysical Journal*, 698, 250, doi: 10.1088/0004-637X/698/1/250
- Chen, H.-L., Chen, X., Tauris, T. M., & Han, Z. 2013, *The Astrophysical Journal*, 775, 27, doi: 10.1088/0004-637X/775/1/27
- Christensen, E., Larson, S., Boattini, A., et al. 2012, in *AAS/Division for Planetary Sciences Meeting Abstracts #44*, AAS/Division for Planetary Sciences Meeting Abstracts, 210.13
- Clark, H. A., Lewis, G. F., & Scott, P. 2016, *Monthly Notices of the RAS*, 456, 1394, doi: 10.1093/mnras/stv2743
- Cognard, I., & Backer, D. C. 2004, *The Astrophysical Journal, Letters*, 612, L125, doi: 10.1086/424692
- Comella, J. M., Craft, H. D., Lovelace, R. V. E., & Sutton, J. M. 1969, *Nature*, 221, 453, doi: 10.1038/221453a0
- Cordes, J. M. 1976, *The Astrophysical Journal*, 210, 780, doi: 10.1086/154887
- . 2013a, *Classical and Quantum Gravity*, 30, 224002, doi: 10.1088/0264-9381/30/22/224002
- . 2013b, *Classical and Quantum Gravity*, 30, 224002, doi: 10.1088/0264-9381/30/22/224002
- . 2013c, *The Astrophysical Journal*, 775, 47, doi: 10.1088/0004-637X/775/1/47

- Cordes, J. M., Bhat, N. D. R., Hankins, T. H., McLaughlin, M. A., & Kern, J. 2004, *The Astrophysical Journal*, 612, 375, doi: 10.1086/422495
- Cordes, J. M., & Downs, G. S. 1985, *The Astrophysical Journal, Supplement*, 59, 343, doi: 10.1086/191076
- Cordes, J. M., & Helfand, D. J. 1980, *The Astrophysical Journal*, 239, 640, doi: 10.1086/158150
- Cordes, J. M., & Jenet, F. A. 2012, *The Astrophysical Journal*, 752, 54, doi: 10.1088/0004-637X/752/1/54
- Cordes, J. M., & Lazio, T. J. W. 2002a, ArXiv Astrophysics e-prints
- . 2002b, arXiv e-prints, astro. <https://arxiv.org/abs/astro-ph/0207156>
- Cordes, J. M., Rickett, B. J., Stinebring, D. R., & Coles, W. A. 2006, *The Astrophysical Journal*, 637, 346, doi: 10.1086/498332
- Cordes, J. M., & Shannon, R. M. 2008, *The Astrophysical Journal*, 682, 1152, doi: 10.1086/589425
- Cordes, J. M., Wolszczan, A., Dewey, R. J., Blaskiewicz, M., & Stinebring, D. R. 1990, *The Astrophysical Journal*, 349, 245, doi: 10.1086/168310
- Crawford, F., Lyne, A. G., Stairs, I. H., et al. 2013, *The Astrophysical Journal*, 776, 20, doi: 10.1088/0004-637X/776/1/20
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *Nature Astronomy*, 4, 72, doi: 10.1038/s41550-019-0880-2
- Damour, T., & Taylor, J. H. 1991, *The Astrophysical Journal*, 366, 501, doi: 10.1086/169585

- Deller, A. T., Verbiest, J. P. W., Tingay, S. J., & Bailes, M. 2008, *The Astrophysical Journal, Letters*, 685, L67, doi: 10.1086/592401
- Deller, A. T., Archibald, A. M., Brisken, W. F., et al. 2012, *The Astrophysical Journal, Letters*, 756, L25, doi: 10.1088/2041-8205/756/2/L25
- Deller, A. T., Vigeland, S. J., Kaplan, D. L., et al. 2016, *The Astrophysical Journal*, 828, 8, doi: 10.3847/0004-637X/828/1/8
- Demorest, P. B. 2007, PhD thesis, University of California, Berkeley. <https://www.cv.nrao.edu/~pdemores/thesis.pdf>
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, *The Astrophysical Journal*, 762, 94, doi: 10.1088/0004-637X/762/2/94
- Deneva, J. S., Ray, P. S., Camilo, F., et al. 2016, *The Astrophysical Journal*, 823, 105, doi: 10.3847/0004-637X/823/2/105
- Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, *Monthly Notices of the RAS*, 458, 3341, doi: 10.1093/mnras/stw483
- Detweiler, S. 1979, *The Astrophysical Journal*, 234, 1100, doi: 10.1086/157593
- Devillepoix, H. A. R., Bland, P. A., Sansom, E. K., et al. 2019, *Monthly Notices of the RAS*, 483, 5166, doi: 10.1093/mnras/sty3442
- Do, A., Tucker, M. A., & Tonry, J. 2018, *The Astrophysical Journal*, 855, L10, doi: 10.3847/2041-8213/aaae67
- Dohnanyi, J. S. 1969, *Journal of Geophysics Research*, 74, 2531, doi: 10.1029/JB074i010p02531

Downs, G. 1974, Interplanetary Navigation Using Pulsating Radio Sources, Tech. Rep. 32-1594, NASA Technical Reports Server. <https://ntrs.nasa.gov/citations/19740026037>

Driebe, T., Blöcker, T., Schönberner, D., & Herwig, F. 1999, *Astronomy and Astrophysics*, 350, 89

Dror, J. A., Ramani, H., Trickle, T., & Zurek, K. M. 2019, *Physical Review D*, 100, 023003, doi: 10.1103/PhysRevD.100.023003

Dubus, G. 2013, *Astronomy and Astrophysics Reviews*, 21, 64, doi: 10.1007/s00159-013-0064-5

DuPlain, R., Ransom, S., Demorest, P., et al. 2008, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7019, Advanced Software and Control for Astronomy II, ed. A. Bridger & N. M. Radziwill, 70191D, doi: 10.1117/12.790003

Eilers, A.-C., Hogg, D. W., Rix, H.-W., & Ness, M. K. 2019, *The Astrophysical Journal*, 871, 120, doi: 10.3847/1538-4357/aaf648

Ellis, J. A., Jenet, F. A., & McLaughlin, M. A. 2012, *The Astrophysical Journal*, 753, 96, doi: 10.1088/0004-637X/753/2/96

Engelhardt, T., Jedicke, R., Vereš, P., et al. 2017, *The Astronomical Journal*, 153, 133, doi: 10.3847/1538-3881/aa5c8a

Esamdin, A., Lyne, A. G., Graham-Smith, F., et al. 2005, *Monthly Notices of the RAS*, 356, 59, doi: 10.1111/j.1365-2966.2004.08444.x

Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, *Monthly Notices of the RAS*, 414, 1679, doi: 10.1111/j.1365-2966.2011.18503.x

Estabrook, F. B., & Wahlquist, H. D. 1975, General Relativity and Gravitation, 6, 439, doi: 10.1007/BF00762449

Fey, A. L., Gordon, D., Jacobs, C. S., et al. 2015, The Astronomical Journal, 150, 58, doi: 10.1088/0004-6256/150/2/58

Fomalont, E. B., Goss, W. M., Lyne, A. G., & Manchester, R. N. 1984, Monthly Notices of the RAS, 210, 113, doi: 10.1093/mnras/210.1.113

Ford, J. M., Prestage, R. M., & Bloss, M. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9152, Software and Cyberinfrastructure for Astronomy III, ed. G. Chiozzi & N. M. Radziwill, 915218, doi: 10.1117/12.2056883

Gaensler, B. M., Madsen, G. J., Chatterjee, S., & Mao, S. A. 2008, Publications of the Astron. Soc. of Australia, 25, 184, doi: 10.1071/AS08004

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, ArXiv e-prints. <https://arxiv.org/abs/1804.09365>

Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, Astronomy and Astrophysics, 595, A1, doi: 10.1051/0004-6361/201629272

Gajjar, V. 2017, arXiv e-prints, arXiv:1706.05407. <https://arxiv.org/abs/1706.05407>

Gajjar, V., Joshi, B. C., & Kramer, M. 2012, Monthly Notices of the RAS, 424, 1197, doi: 10.1111/j.1365-2966.2012.21296.x

Glampedakis, K., & Andersson, N. 2009, Physical Review Letters, 102, 141101, doi: 10.1103/PhysRevLett.102.141101

Gold, T. 1968, Nature, 218, 731, doi: 10.1038/218731a0

- . 1969, *Nature*, 221, 25, doi: 10.1038/221025a0
- Goncharov, B., Reardon, D. J., Shannon, R. M., et al. 2021, *Monthly Notices of the RAS*, 502, 478, doi: 10.1093/mnras/staa3411
- Guillemot, L., Smith, D. A., Laffon, H., et al. 2016, *Astronomy and Astrophysics*, 587, A109, doi: 10.1051/0004-6361/201527847
- Guo, Y. J., Lee, K. J., & Caballero, R. N. 2018, *Monthly Notices of the RAS*, 475, 3644, doi: 10.1093/mnras/stx3326
- Guzik, P., Drahus, M., Rusek, K., et al. 2019, arXiv e-prints, arXiv:1909.05851. <https://arxiv.org/abs/1909.05851>
- Hajduk, A. 2001, in *ESA Special Publication*, Vol. 495, *Meteoroids 2001 Conference*, ed. B. Warmbein, 557–559
- Hajduková, Mária, J. 2011, *Publications of the ASJ*, 63, 481, doi: 10.1093/pasj/63.3.481
- Hajduková, M. 2008, *Earth Moon and Planets*, 102, 67, doi: 10.1007/s11038-007-9171-5
- Hazboun, J. S., Romano, J. D., & Smith, T. L. 2019, *Physical Review D*, 100, 104028, doi: 10.1103/PhysRevD.100.104028
- Heiles, C., Perillat, P., Nolan, M., et al. 2001, *Publications of the ASP*, 113, 1274, doi: 10.1086/323289
- Hellings, R. W., & Downs, G. S. 1983, *The Astrophysical Journal, Letters*, 265, L39, doi: 10.1086/183954
- Heng, K., & Tremaine, S. 2010, *Monthly Notices of the RAS*, 401, 867, doi: 10.1111/j.1365-2966.2009.15739.x

- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709, doi: 10.1038/217709a0
- Ho, W. C. G., & Heinke, C. O. 2009, *Nature*, 462, 71, doi: 10.1038/nature08525
- Ho, W. C. G., Ng, C. Y., Lyne, A. G., et al. 2017, *Monthly Notices of the RAS*, 464, 1211, doi: 10.1093/mnras/stw2420
- Hobbs, G. 2013, *Classical and Quantum Gravity*, 30, 224007, doi: 10.1088/0264-9381/30/22/224007
- Hobbs, G., Lyne, A. G., Kramer, M., Martin, C. E., & Jordan, C. 2004, *Monthly Notices of the RAS*, 353, 1311, doi: 10.1111/j.1365-2966.2004.08157.x
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *Classical and Quantum Gravity*, 27, 084013, doi: 10.1088/0264-9381/27/8/084013
- Hobbs, G., Guo, L., Caballero, R. N., et al. 2020, *Monthly Notices of the RAS*, 491, 5951, doi: 10.1093/mnras/stz3071
- Holman, M. J., & Wiegert, P. A. 1999, *The Astronomical Journal*, 117, 621, doi: 10.1086/300695
- Istrate, A. G., Marchant, P., Tauris, T. M., et al. 2016, *Astronomy and Astrophysics*, 595, A35, doi: 10.1051/0004-6361/201628874
- Jedicke, R., Larsen, J., & Spahr, T. 2002, *Observational Selection Effects in Asteroid Surveys*, 71–87
- Jenet, F. A., Anderson, S. B., & Prince, T. A. 2001, *The Astrophysical Journal*, 558, 302, doi: 10.1086/322469

- Jennings, R. J. 2021a, Transmission Functions for Polynomial Fits, Tech. Rep. 006, NANOGrav Memorandum Series. <http://nanograv.org/assets/files/memos/NANOGrav-Memo-006.pdf>
- . 2021b, The NANOGrav TOA Generation Process, Tech. Rep. 007, NANOGrav Memorandum Series. <http://nanograv.org/assets/files/memos/NANOGrav-Memo-007.pdf>
- Jennings, R. J., Cordes, J. M., & Chatterjee, S. 2020a, *The Astrophysical Journal*, 889, 145, doi: 10.3847/1538-4357/ab64df
- . 2020b, *The Astrophysical Journal*, 904, 191, doi: 10.3847/1538-4357/abc178
- Jennings, R. J., Kaplan, D. L., Chatterjee, S., Cordes, J. M., & Deller, A. T. 2018, *The Astrophysical Journal*, 864, 26, doi: 10.3847/1538-4357/aad084
- Jones, P. B. 1990, *Monthly Notices of the RAS*, 246, 364
- Joshi, B. C., Arumugasamy, P., Bagchi, M., et al. 2018, *Journal of Astrophysics and Astronomy*, 39, 51, doi: 10.1007/s12036-018-9549-y
- Kaplan, D. L., Bhalerao, V. B., van Kerkwijk, M. H., et al. 2013, *The Astrophysical Journal*, 765, 158, doi: 10.1088/0004-637X/765/2/158
- Kaplan, D. L., Chakrabarty, D., Wang, Z., & Wachter, S. 2009, *The Astrophysical Journal*, 700, 149, doi: 10.1088/0004-637X/700/1/149
- Kaplan, D. L., Chatterjee, S., Gaensler, B. M., & Anderson, J. 2008, *The Astrophysical Journal*, 677, 1201, doi: 10.1086/529026
- Kaplan, D. L., van Kerkwijk, M. H., & Anderson, J. 2002, *The Astrophysical Journal*, 571, 447, doi: 10.1086/339879

- . 2007, *The Astrophysical Journal*, 660, 1428, doi: 10.1086/513309
- Kaplan, D. L., van Kerkwijk, M. H., Koester, D., et al. 2014, *The Astrophysical Journal, Letters*, 783, L23, doi: 10.1088/2041-8205/783/1/L23
- Kaplan, D. L., Kupfer, T., Nice, D. J., et al. 2016, *The Astrophysical Journal*, 826, 86, doi: 10.3847/0004-637X/826/1/86
- Kashiyama, K., & Oguri, M. 2018, arXiv e-prints, arXiv:1801.07847. <https://arxiv.org/abs/1801.07847>
- Kashiyama, K., & Seto, N. 2012, *Monthly Notices of the RAS*, 426, 1369, doi: 10.1111/j.1365-2966.2012.21935.x
- Kaspi, V. M., Bailes, M., Manchester, R. N., Stappers, B. W., & Bell, J. F. 1996, *Nature*, 381, 584, doi: 10.1038/381584a0
- Kerr, F. J., & Lynden-Bell, D. 1986, *Monthly Notices of the RAS*, 221, 1023, doi: 10.1093/mnras/221.4.1023
- Kerr, M. 2015, *Monthly Notices of the RAS*, 452, 607, doi: 10.1093/mnras/stv1296
- Kerr, M., Reardon, D. J., Hobbs, G., et al. 2020, *Publications of the Astron. Soc. of Australia*, 37, e020, doi: 10.1017/pasa.2020.11
- Kiminki, D. C., Kobulnicky, H. A., Vargas Álvarez, C. A., Alexander, M. J., & Lundquist, M. J. 2015, *The Astrophysical Journal*, 811, doi: 10.1088/0004-637X/811/2/85
- Kirsten, F., Vlemmings, W., Campbell, R. M., Kramer, M., & Chatterjee, S. 2015, *Astronomy and Astrophysics*, 577, A111, doi: 10.1051/0004-6361/201425562

- Kong, A. K. H., Huang, R. H. H., Cheng, K. S., et al. 2012, *The Astrophysical Journal, Letters*, 747, L3, doi: 10.1088/2041-8205/747/1/L3
- Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006, *Science*, 312, 549, doi: 10.1126/science.1124060
- Krishnamohan, S., & Downs, G. S. 1983, *The Astrophysical Journal*, 265, 372, doi: 10.1086/160682
- Kulkarni, S. R., Djorgovski, S., & Fruchter, A. S. 1988, *Nature*, 334, 504, doi: 10.1038/334504a0
- Kurosawa, N., Kobayashi, H., Maruyama, K., Sugawara, H., & Kobayashi, K. 2001, *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 48, 261, doi: 10.1109/81.915383
- Lam, M. T. 2021, *Research Notes of the American Astronomical Society*, 5, 167, doi: 10.3847/2515-5172/ac1670
- Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, *The Astrophysical Journal*, 819, 155, doi: 10.3847/0004-637X/819/2/155
- . 2017, *The Astrophysical Journal*, 834, 35, doi: 10.3847/1538-4357/834/1/35
- Lam, M. T., Ellis, J. A., Grillo, G., et al. 2018, *The Astrophysical Journal*, 861, 132, doi: 10.3847/1538-4357/aac770
- Lam, M. T., McLaughlin, M. A., Arzoumanian, Z., et al. 2019, *The Astrophysical Journal*, 872, 193, doi: 10.3847/1538-4357/ab01cd
- Landgraf, M., Baggaley, W. J., Grün, E., Krüger, H., & Linkert, G. 2000, *Journal of Geophysics Research*, 105, 10343, doi: 10.1029/1999JA900359

- Lentati, L., Alexander, P., & Hobson, M. P. 2015, *Monthly Notices of the RAS*, 447, 2159, doi: 10.1093/mnras/stu2611
- Li, M., Halpern, J. P., & Thorstensen, J. R. 2014, *The Astrophysical Journal*, 795, 115, doi: 10.1088/0004-637X/795/2/115
- Lin, F. X., Lin, H.-H., Luo, J., et al. 2021, *Monthly Notices of the RAS*, 508, 1115, doi: 10.1093/mnras/stab2529
- Lindgren, L., Hernandez, J., Bombrun, A., et al. 2018, ArXiv e-prints, arXiv:1804.09366. <https://arxiv.org/abs/1804.09366>
- Liu, K., Keane, E. F., Lee, K. J., et al. 2012, *Monthly Notices of the RAS*, 420, 361, doi: 10.1111/j.1365-2966.2011.20041.x
- Lommen, A. N. 2015, *Reports on Progress in Physics*, 78, 124901, doi: 10.1088/0034-4885/78/12/124901
- Lorimer, D. R., & Kramer, M. 2012, *Handbook of Pulsar Astronomy* (Cambridge, UK: Cambridge University Press)
- Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, *Monthly Notices of the RAS*, 372, 777, doi: 10.1111/j.1365-2966.2006.10887.x
- Lutz, T. E., & Kelker, D. H. 1973, *Publications of the ASP*, 85, 573, doi: 10.1086/129506
- Lyne, A., Hobbs, G., Kramer, M., Stairs, I., & Stappers, B. 2010, *Science*, 329, 408, doi: 10.1126/science.1186683
- Lyne, A. G. 1971, *Monthly Notices of the RAS*, 153, 27P, doi: 10.1093/mnras/153.1.27P

- Lyne, A. G., Pritchard, R. S., & Shemar, S. L. 1995, *Journal of Astrophysics and Astronomy*, 16, 179, doi: 10.1007/BF02714833
- Ma, C., Arias, E. F., Eubanks, T. M., et al. 1998, *The Astronomical Journal*, 116, 516, doi: 10.1086/300408
- Madison, D. R., Chatterjee, S., & Cordes, J. M. 2013, *The Astrophysical Journal*, 777, 104, doi: 10.1088/0004-637X/777/2/104
- Madison, D. R., Chernoff, D. F., & Cordes, J. M. 2017, *Physical Review D*, 96, 123016, doi: 10.1103/PhysRevD.96.123016
- Madison, D. R., Zhu, X. J., Hobbs, G., et al. 2016, *Monthly Notices of the RAS*, 455, 3662, doi: 10.1093/mnras/stv2534
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *The Astronomical Journal*, 129, 1993, doi: 10.1086/428488
- . 2016, *VizieR Online Data Catalog*, 1
- Matthews, A. M., Nice, D. J., Fonseca, E., et al. 2016, *The Astrophysical Journal*, 818, 92, doi: 10.3847/0004-637X/818/1/92
- McKee, J. W., Janssen, G. H., Stappers, B. W., et al. 2016, *Monthly Notices of the RAS*, 461, 2809, doi: 10.1093/mnras/stw1442
- Meech, K. J., Weryk, R., Micheli, M., et al. 2017, *Nature*, 552, 378, doi: 10.1038/nature25020
- Meisel, D. D., Janches, D., & Mathews, J. D. 2002, *The Astrophysical Journal*, 579, 895, doi: 10.1086/342919
- Meyers, B., & CHIME/Pulsar Collaboration. 2021, *The Astronomer's Telegram*, 14652, 1

Mignani, R. P. 2011, *Advances in Space Research*, 47, 1281, doi: 10.1016/j.asr.2009.12.011

Mignani, R. P., Corongiu, A., Pallanca, C., et al. 2014, *Monthly Notices of the RAS*, 443, 2223, doi: 10.1093/mnras/stu1300

Mignard, F., Klioner, S., Lindegren, L., et al. 2018, ArXiv e-prints. <https://arxiv.org/abs/1804.09377>

Miller, M. C., & Hamilton, D. P. 2001, *The Astrophysical Journal*, 550, 863, doi: 10.1086/319813

Miller-Jones, J. C. A., Deller, A. T., Shannon, R. M., et al. 2018, ArXiv e-prints. <https://arxiv.org/abs/1804.08402>

Murray, C. D., & Dermott, S. F. 1999, *Solar system dynamics*

Nagler, J. 2005, *Physical Review E*, 71, 026227, doi: 10.1103/PhysRevE.71.026227

Negueruela, I., Ribó, M., Herrero, A., et al. 2011, *The Astrophysical Journal, Letters*, 732, L11, doi: 10.1088/2041-8205/732/1/L11

Niikura, H., Takada, M., Yasuda, N., et al. 2019, *Nature Astronomy*, 3, 524, doi: 10.1038/s41550-019-0723-1

Offringa, A. R., de Bruyn, A. G., Biehl, M., et al. 2010, *Monthly Notices of the RAS*, 405, 155, doi: 10.1111/j.1365-2966.2010.16471.x

Osłowski, S., van Straten, W., Hobbs, G. B., Bailes, M., & Demorest, P. 2011, *Monthly Notices of the RAS*, 418, 1258, doi: 10.1111/j.1365-2966.2011.19578.x

Özel, F., & Freire, P. 2016, *Annual Review of Astron. and Astrophys.*, 54, 401, doi: 10.1146/annurev-astro-081915-023322

Pacini, F. 1968, *Nature*, 219, 145, doi: 10.1038/219145a0

Parthasarathy, A., Shannon, R. M., Johnston, S., et al. 2019, *Monthly Notices of the RAS*, 2056, doi: 10.1093/mnras/stz2383

Perera, B. B. P., DeCesar, M. E., Demorest, P. B., et al. 2019, *Monthly Notices of the RAS*, 490, 4666, doi: 10.1093/mnras/stz2857

Perryman, M. A. C., Lindegren, L., Kovalevsky, J., et al. 1997, *Astronomy and Astrophysics*, 323, L49

Peters, P. C., & Mathews, J. 1963, *Physical Review*, 131, 435, doi: 10.1103/PhysRev.131.435

Phinney, E. S. 1993, in *Astronomical Society of the Pacific Conference Series*, Vol. 50, *Structure and Dynamics of Globular Clusters*, ed. S. G. Djorgovski & G. Meylan, 141

Pitkin, M. 2018, *Journal of Open Source Software*, 3, 538, doi: 10.21105/joss.00538

Quarles, B., Li, G., Kostov, V., & Haghishipour, N. 2020, *The Astronomical Journal*, 159, 80, doi: 10.3847/1538-3881/ab64fa

Quarles, B., Satyal, S., Kostov, V., Kaib, N., & Haghishipour, N. 2018, *The Astrophysical Journal*, 856, 150, doi: 10.3847/1538-4357/aab264

Rankin, J. M. 1986, *The Astrophysical Journal*, 301, 901, doi: 10.1086/163955

Ransom, S. M., Stairs, I. H., Archibald, A. M., et al. 2014, *Nature*, 505, 520, doi: 10.1038/nature12917

- Ray, P. S., Polisensky, E., Parkinson, P. S., et al. 2020, Research Notes of the American Astronomical Society, 4, 37, doi: 10.3847/2515-5172/ab7eb5
- Reardon, D. J., Hobbs, G., Coles, W., et al. 2016, Monthly Notices of the RAS, 455, 1751, doi: 10.1093/mnras/stv2395
- Reardon, D. J., Shannon, R. M., Cameron, A. D., et al. 2021, Monthly Notices of the RAS, 507, 2137, doi: 10.1093/mnras/stab1990
- Reichley, P., Downs, G., & Morris, G. 1971, Jet Propulsion Lab, 1, 80
- Richards, D. W., & Comella, J. M. 1969, Nature, 222, 551, doi: 10.1038/222551a0
- Rickett, B. J. 1975, The Astrophysical Journal, 197, 185, doi: 10.1086/153501
- . 1990, Annual Review of Astron. and Astrophys., 28, 561, doi: 10.1146/annurev.aa.28.090190.003021
- Roberts, M. S. E. 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, 127–130, doi: 10.1063/1.3615095
- 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, doi: 10.1017/S174392131202337X
- Romani, R. W., Filippenko, A. V., Silverman, J. M., et al. 2012, The Astrophysical Journal, Letters, 760, L36, doi: 10.1088/2041-8205/760/2/L36
- Romani, R. W., & Shaw, M. S. 2011, The Astrophysical Journal, Letters, 743, L26, doi: 10.1088/2041-8205/743/2/L26

Roshi, D. A., Bloss, M., Brandt, P., et al. 2011, in 2011 XXXth URSI General Assembly and Scientific Symposium, 1–4, doi: 10.1109/URSIGASS.2011.6051280

Roy, A. E. 1988, Orbital Motion, 3rd edn. (Philadelphia: IOP)

Roy, J., Ray, P. S., Bhattacharyya, B., et al. 2015, *The Astrophysical Journal, Letters*, 800, L12, doi: 10.1088/2041-8205/800/1/L12

Rubincam, D. P. 1998, *Journal of Geophysics Research*, 103, 1725, doi: 10.1029/97JE03034

Schönrich, R., Binney, J., & Dehnen, W. 2010, *Monthly Notices of the RAS*, 403, 1829, doi: 10.1111/j.1365-2966.2010.16253.x

Seto, N., & Cooray, A. 2007, *The Astrophysical Journal*, 659, L33, doi: 10.1086/516570

Shahbaz, T., Dallilar, Y., Garner, A., et al. 2018, *Monthly Notices of the RAS*, 477, 566, doi: 10.1093/mnras/sty562

Shannon, R. M., & Cordes, J. M. 2010, *The Astrophysical Journal*, 725, 1607, doi: 10.1088/0004-637X/725/2/1607

Shannon, R. M., Cordes, J. M., Metcalfe, T. S., et al. 2013, *The Astrophysical Journal*, 766, 5, doi: 10.1088/0004-637X/766/1/5

Shannon, R. M., Lentati, L. T., Kerr, M., et al. 2016, *The Astrophysical Journal, Letters*, 828, L1, doi: 10.3847/2041-8205/828/1/L1

Shaw, A. W., Heinke, C. O., Steiner, A. W., et al. 2018, *Monthly Notices of the RAS*, 476, 4713, doi: 10.1093/mnras/sty582

Sheikh, S. Z., & MacDonald, M. G. 2021, *Monthly Notices of the RAS*, 502, 4669, doi: 10.1093/mnras/stab282

Shemar, S., Fraser, G., Heil, L., et al. 2016, *Experimental Astronomy*, 42, 101, doi: 10.1007/s10686-016-9496-z

Shevchenko, I. I. 2015, *The Astrophysical Journal*, 799, 8, doi: 10.1088/0004-637X/799/1/8

Shklovskii, I. S. 1970, *Soviet Astronomy*, 13, 562

Siegel, E. R., Hertzberg, M. P., & Fry, J. N. 2007, *Monthly Notices of the RAS*, 382, 879, doi: 10.1111/j.1365-2966.2007.12435.x

Singha, J., Surnis, M. P., Joshi, B. C., et al. 2021, *Monthly Notices of the RAS*, 507, L57, doi: 10.1093/mnrasl/slab098

Siraj, A., & Loeb, A. 2019a, arXiv e-prints, arXiv:1904.07224. <https://arxiv.org/abs/1904.07224>

—. 2019b, arXiv e-prints, arXiv:1906.03270. <https://arxiv.org/abs/1906.03270>

Spiewak, R., Kaplan, D. L., Archibald, A., et al. 2016, *The Astrophysical Journal*, 822, 37, doi: 10.3847/0004-637X/822/1/37

Staelin, D. H., & Reifenstein, Edward C., I. 1968, *Science*, 162, 1481, doi: 10.1126/science.162.3861.1481

Stairs, I. H. 2003, *Living Reviews in Relativity*, 6, 5, doi: 10.12942/lrr-2003-5

Stinebring, D. R., McLaughlin, M. A., Cordes, J. M., et al. 2001, *The Astrophysical Journal Letters*, 549, L97, doi: 10.1086/319133

- Storm, J., Carney, B. W., Gieren, W. P., et al. 2004, *Astronomy and Astrophysics*, 415, 531, doi: 10.1051/0004-6361:20034634
- Stovall, K., Allen, B., Bogdanov, S., et al. 2016, *The Astrophysical Journal*, 833, 192, doi: 10.3847/1538-4357/833/2/192
- Strader, J., Chomiuk, L., Cheung, C. C., et al. 2015, *The Astrophysical Journal Letters*, 804, L12, doi: 10.1088/2041-8205/804/1/L12
- Szebehely, V. 1978, *Proceedings of the National Academy of Science*, 75, 5743, doi: 10.1073/pnas.75.12.5743
- . 1980, *Celestial Mechanics*, 22, 7, doi: 10.1007/BF01228750
- Szebehely, V., & McKenzie, R. 1981, *Celestial Mechanics*, 23, 3, doi: 10.1007/BF01228541
- Tauris, T. M. 2011, in *Astronomical Society of the Pacific Conference Series*, Vol. 447, *Evolution of Compact Binaries*, ed. L. Schmidtobreick, M. R. Schreiber, & C. Tappert, 285. <https://arxiv.org/abs/1106.0897>
- Taylor, A. D., Baggaley, W. J., & Steel, D. I. 1996, *Nature*, 380, 323, doi: 10.1038/380323a0
- Taylor, J. H. 1992, *Philosophical Transactions of the Royal Society of London Series A*, 341, 117, doi: 10.1098/rsta.1992.0088
- Taylor, J. H., & Weisberg, J. M. 1982, *The Astrophysical Journal*, 253, 908, doi: 10.1086/159690
- Tetzlaff, N., Neuhäuser, R., Hohle, M. M., & Maciejewski, G. 2010, *Monthly Notices of the RAS*, 402, 2369, doi: 10.1111/j.1365-2966.2009.16093.x

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2021, arXiv e-prints, arXiv:2111.03606. <https://arxiv.org/abs/2111.03606>

Tisserand, P., Le Guillou, L., Afonso, C., et al. 2007, *Astronomy and Astrophysics*, 469, 387, doi: 10.1051/0004-6361:20066017

Toscano, M., Britton, M. C., Manchester, R. N., et al. 1999, *The Astrophysical Journal, Letters*, 523, L171, doi: 10.1086/312276

van Haasteren, R., & Levin, Y. 2010, *Monthly Notices of the RAS*, 401, 2372, doi: 10.1111/j.1365-2966.2009.15885.x

van Straten, W., & Bailes, M. 2011, *Publications of the Astron. Soc. of Australia*, 28, 1, doi: 10.1071/AS10021

Verbiest, J. P. W., Weisberg, J. M., Chael, A. A., Lee, K. J., & Lorimer, D. R. 2012, *The Astrophysical Journal*, 755, 39, doi: 10.1088/0004-637X/755/1/39

Verbiest, J. P. W., Bailes, M., van Straten, W., et al. 2008, *The Astrophysical Journal*, 679, 675, doi: 10.1086/529576

Verbiest, J. P. W., Lentati, L., Hobbs, G., et al. 2016, *Monthly Notices of the RAS*, 458, 1267, doi: 10.1093/mnras/stw347

Verbunt, F., & Hut, P. 1987, in *IAU Symposium*, Vol. 125, *The Origin and Evolution of Neutron Stars*, ed. D. J. Helfand & J. H. Huang, 187

Vlemmings, W. H. T., Cordes, J. M., & Chatterjee, S. 2004, *The Astrophysical Journal*, 610, 402, doi: 10.1086/421439

Wahl, H. M., McLaughlin, M., Gentile, P. A., et al. 2021, arXiv e-prints, arXiv:2104.05723. <https://arxiv.org/abs/2104.05723>

- Walker, M. A., Melrose, D. B., Stinebring, D. R., & Zhang, C. M. 2004, *Monthly Notices of the RAS*, 354, 43, doi: 10.1111/j.1365-2966.2004.08159.x
- Wang, J. B., Coles, W. A., Hobbs, G., et al. 2017, *Monthly Notices of the RAS*, 469, 425, doi: 10.1093/mnras/stx837
- Wang, N., Manchester, R. N., & Johnston, S. 2007, *Monthly Notices of the RAS*, 377, 1383, doi: 10.1111/j.1365-2966.2007.11703.x
- Wang, Z., Chakrabarty, D., & Kaplan, D. L. 2006, *Nature*, 440, 772, doi: 10.1038/nature04669
- Weisberg, J. M., & Huang, Y. 2016, *The Astrophysical Journal*, 829, 55, doi: 10.3847/0004-637X/829/1/55
- Weisberg, J. M., Nice, D. J., & Taylor, J. H. 2010, *The Astrophysical Journal*, 722, 1030, doi: 10.1088/0004-637X/722/2/1030
- Weryk, R. J., & Brown, P. 2004, *Earth Moon and Planets*, 95, 221, doi: 10.1007/s11038-005-9034-x
- Wex, N., Kalogera, V., & Kramer, M. 2000, *The Astrophysical Journal*, 528, 401, doi: 10.1086/308148
- Wolf, C., Onken, C. A., Luvaal, L. C., et al. 2018, *Publications of the Astron. Soc. of Australia*, 35, e010, doi: 10.1017/pasa.2018.5
- Wolszczan, A. 1994, *Science*, 264, 538, doi: 10.1126/science.264.5158.538
- Wolszczan, A., Hoffman, I. M., Konacki, M., Anderson, S. B., & Xilouris, K. M. 2000, *The Astrophysical Journal*, 540, L41, doi: 10.1086/312872
- Wyrzykowski, L., Skowron, J., Kozłowski, S., et al. 2011, *Monthly Notices of the RAS*, 416, 2949, doi: 10.1111/j.1365-2966.2011.19243.x

Xu, H., Huang, Y. X., Burgay, M., et al. 2021, The Astronomer's Telegram, 14642, 1

Yao, J. M., Manchester, R. N., & Wang, N. 2017, The Astrophysical Journal, 835, 29, doi: 10.3847/1538-4357/835/1/29

Zhu, W. W., Desvignes, G., Wex, N., et al. 2019, Monthly Notices of the RAS, 482, 3249, doi: 10.1093/mnras/sty2905

Zuluaga, J. I. 2019, Research Notes of the American Astronomical Society, 3, 68, doi: 10.3847/2515-5172/ab1de3