# The detectability of millisecond pulsars in eccentric binary systems

by

Erik Madsen

BSc, University of British Columbia, 2011

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

**Master of Science**

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Astronomy)

The University Of British Columbia
(Vancouver)

August 2013

# Abstract

With new, highly sensitive telescopes, increased computational power, and improved search algorithms, the present century has seen a great increase in the discovery of pulsars in globular clusters. These are typically fast-spinning 'millisecond' pulsars, and more often than not they are members of binary systems. Unlike in the Galactic field, millisecond pulsars in globular clusters are often found in eccentric systems because of disruptions and exchanges due to the high stellar density of the cluster environment.

A long-standing problem is that of characterizing our sensitivity to pulsars in binary systems, particularly those with non-zero eccentricity. A pulsar's orbital motion modulates its observed pulse period, making its detection through standard Fourier analysis difficult or impossible. A common technique to mitigate this problem is the 'acceleration search', which corrects for uniform line-of-sight acceleration, but not higher-order variations. This is often a valid approximation, and many pulsars have been found this way. However, it is not clear where such a search breaks down.

This is a problem with a many-dimensional phase space that includes all of the binary parameters, the pulsar parameters, and the various search inputs. Past studies have approached the problem analytically, and have made valuable insights; however, until recently they have been restricted to circular orbits, and have not accounted for pulsar brightness or signal digitization.

Here I approach the problem empirically. I simulate 1.8 million pulsars in a variety of orbital configurations and explore the frequency of pulsar recovery across various dimensions of the phase space. I find in particular that, at very short orbital periods, high eccentricities make binary systems easier to detect.

# Table of Contents

# List of Tables

# List of Figures

# Acknowledgments

I would like to thank Dr. Ingrid Stairs for her patience and advice, and for always being an impressive source of knowledge with regards to all things pulsar. By now it comes as no surprise to me that even in a year as busy and exciting as this one has been for her, she finds a way to make time for her students.

Thanks to Dr. Jeremy Heyl and Dr. Scott Ransom for taking the time to read through this thesis even though I withheld it until the last possible moment.

I would also like to acknowledge the support of CANARIE for providing funding for the computing system used in this work.

Finally, thanks to Chelsea for keeping me sane and being patient with me despite my less-than-ideal organizational skills. I never would have made it half this far without her love and support.

# Chapter 1

# Introduction

## 1.1 What is a pulsar?

The first pulsar was discovered serendipitously in 1967 (Hewish et al., 1968). A very fast, steady series of radio-frequency pulses was observed from a particular point on the sky, evidently originating in deep space. Following this discovery, many more pulsars were found in quick succession, and before long it was agreed upon that they were highly magnetized, rapidly rotating neutron stars (Gold, 1969; Pacini, 1968); radio emission beamed from the stars' magnetic poles was sweeping across our line of sight like a lighthouse beam, and this was the reason it was observed as regular radio pulses. This remains our basic description of pulsars nearly half a century later.

Neutron stars were proposed by Baade & Zwicky (1934) as theoretical compact products of supernovae, and the first direct evidence for their existence was in the discovery of pulsars. Pulsars are a class of neutron star, though not every neutron star can be observed as a pulsar. It is very likely that there are many neutron stars with beamed radio emission that simply never points towards Earth, any many without such emission at all.

A very good description of the modern pulsar model is laid out in Lorimer & Kramer (2005), and much of the information in this section can be found there and in references within—where necessary, I provide additional references. Neutron stars are formed in core-collapse supernovae. When an evolved massive star is no

longer able to create energy through nuclear fusion in its core (i.e., the core nuclei are iron), the hydrostatic equilibrium between radiative expansion and gravitational collapse breaks down, and an enormous amount of mass rapidly falls inward. Because multiple electrons cannot exist in identical states, electrons are forced to occupy higher and higher energy states as the volume is compressed. This creates an outward *electron degeneracy pressure*, which is the pressure that withholds a white dwarf star from collapse. If there is sufficient stellar mass to compress the core beyond this state, then electrons and protons form neutrons via inverse beta decay, and an analagous *neutron degeneracy pressure* withholds further collapse. If the compression can proceed no further, we are left with a neutron star, with a radius of about 10 km. Further collapse would produce a black hole or perhaps another exotic object such as a quark star, though the latter remains speculative.

While the physics of neutron star interiors remains a subject of debate, much has been observationally established about their masses and exteriors. The Chandrasekhar limit of 1.4 $M_\odot$ has often been used as a canonical neutron star mass, though it has become clear that one 'typical mass' is not sufficient to describe the current distribution of measurements, with the mass distribution increasingly appearing multi-modal and dependant upon formation scenario (e.g., Özel et al., 2012; Schwab, Podsiadlowski & Rappaport, 2010). Their masses do tend to be scattered in a small range about this value, nonetheless; thus far they have been measured between 1.25 and 2 $M_\odot$ (Demorest et al., 2010; Kramer et al., 2006). The Chandrasekhar limit is the mass beyond which a white dwarf star should theoretically no longer be able to withstand gravitational pressure, so it is not surprising that neutron stars are not found well below this mass (though a neutron star's measured mass will be less than that of a white dwarf that collapses to form it, due to its higher gravitational binding energy, as discussed by, e.g., Podsiadlowski et al. (2005)). The analagous upper-limit mass for neutron stars is not well-determined, but it might be as high as about 3 $M_\odot$, depending upon the correct neutron star equation of state.

Neutron stars spin fast and have strong magnetic fields. Angular momentum conservation and magnetic field compression during core collapse can basically account for these properties. Because of the magnetic field, there is a corotating magnetosphere, extending no further than the *light cylinder* (Figure 1.1), the dis-

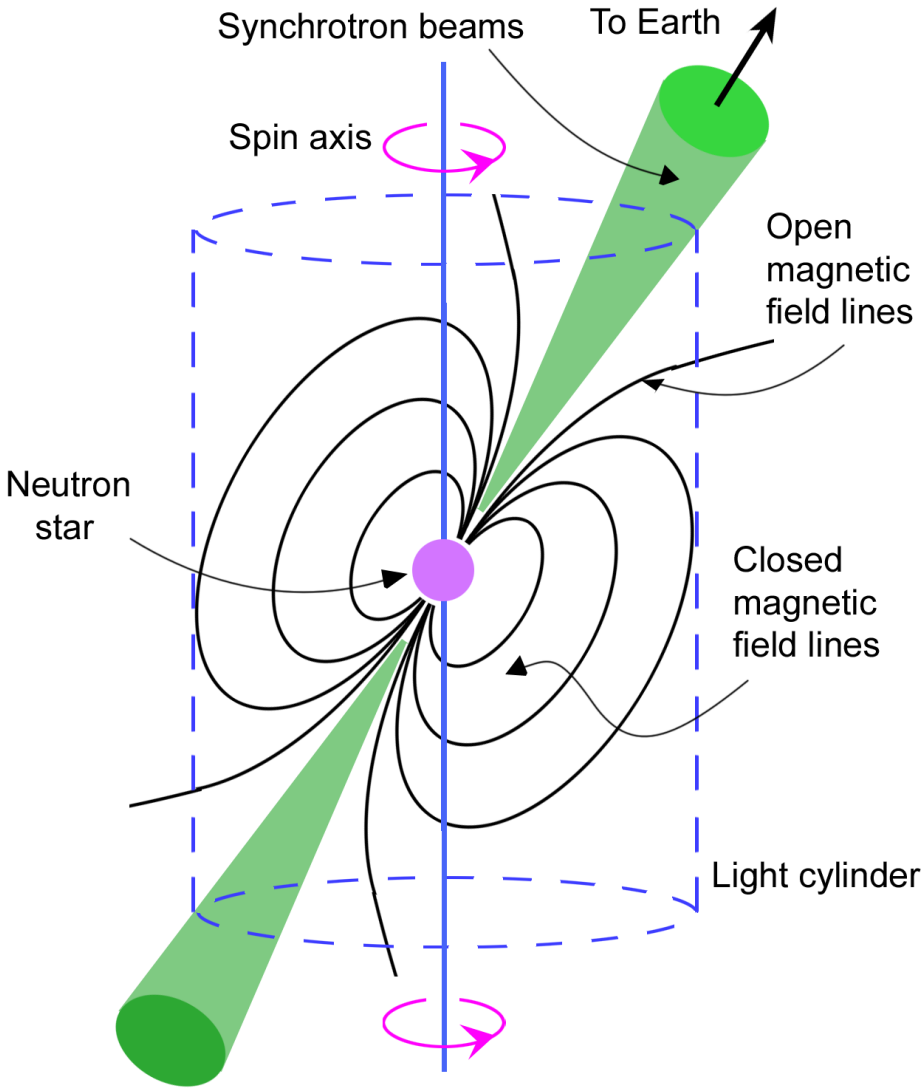


**Figure 1.1:** The basic pulsar model. Magnetic field lines that extend beyond the light cylinder cannot close on the neutron star's opposite pole, as corotation would require them to move faster than the speed of light. Particles streaming along these "open" field lines are thought to play a part in the formation of the radio beams.

tance beyond which anything corotating with the star would exceed the speed of light; a faster-spinning star has a tighter light cylinder. Field lines that originate nearer the star's equator turn back and close before reaching this distance. Field lines that originate closer to the poles, however, extend beyond the light cylinder and do not close on the star's opposite pole. It is in these polar cap regions that a process involving the streaming of charged particles is thought to be responsible for the beamed radio emission observed in pulsars.

We are able to observe a pulsar's beamed emission as a series of short pulses because of a misalignment between the magnetic and spin axes; this causes the beams aligned with the magnetic axis to be swept about, and if at least one of the beams points towards the Earth once per sweep, we see a series of pulses, much like the brief flashes of a lighthouse beam. Probably some neutron stars have their spin and magnetic axes aligned, but even if radio beams are emitted, we would not see these as pulsars.

## 1.2 Millisecond pulsars

A pulsar's emission is powered by its rotational energy, and this leads to a gradual slowing of the pulsar's spin rate. The youngest pulsars have lower spin periods ($P$) and higher period derivatives ($\dot{P}$), while older pulsars spin more slowly (periods from hundreds of milliseconds up to a few seconds) and their rotational speeds decay more slowly. This is shown on the $P$-$\dot{P}$ diagram of pulsars shown in Figure 1.2, where diagonal lines show characteristic ages, defined as $\tau_c = P/2\dot{P}$. The main cluster of points on this plot represents these "typical" pulsars, but there is another distinct cluster at much lower spin periods and period derivatives. These are the millisecond pulsars (MSPs), whose positions on the diagram do not correspond to their ages, as they are old pulsars that have had their spin rates increased through mass-accretion; looking at the plot, it is immediately striking how many of the pulsars in this region are members of binary systems. It has been widely thought for some time that these binary systems were at one time much like the accreting systems known as low-mass X-ray binaries (LMXBs) (e.g., Kiziltan & Thorsett, 2010; Phinney & Kulkarni, 1994). Recently, direct links between MSPs and LMXBs have even been established, with the 3.93-ms globular cluster pulsar J1824−2452I

(Papitto et al., 2013, in press) apparently switching between rotation-powered and accretion-powered emission, and the 1.69-ms field pulsar J1023+0038 (Archibald et al., 2009) showing evidence for a recent accretion disk. The angular acceleration that occurs during this accretion is often called *spinning up*, and MSPs are sometimes referred to as *recycled* pulsars because of this process. The *spin-up line*, drawn as a thick black line in Figure 1.2, represents the minimum period allowed by Eddington-limited accretion.

Tidal forces in an accreting binary system also tend to circularize the orbit, so that binary systems with MSPs typically have very low eccentricities ($e$). An exception in the Galactic field is PSR J1903+0327, an MSP in a binary system with a main-sequence star that has a high eccentricity of 0.44 (Champion et al., 2008). Initial possible explanations included the system having been formed in and ejected from a globular cluster; alternatively, Freire et al. (2011) and Portegies Zwart et al. (2011) both suggest that the system previously had a third body, that which spun up the pulsar, which was later ejected due to chaotic three-body interactions. Such a scenario could have resulted in the ejection of the MSP from the system instead of the donor star, and may thus explain the presence of isolated MSPs in the Galaxy. Whatever the explanation for its own properties, it is clear that J1903+0327 is not a typical system.

It is in globular clusters, however, that a population of eccentric MSP systems exists; in these clusters, very high stellar densities greatly increase the probability for close encounters that disturb binaries.

## 1.3 Globular clusters and their MSPs

Globular clusters are gravitationally bound systems of hundreds of thousands of stars. They orbit in the Galaxy's halo, formed early in the history of the Universe, and contain some of the oldest stars ever observed. It has long been thought that all of the stars in a globular cluster formed at the same time, and while star-formation clearly ceased early in the clusters' formation, the discovery of multiple distinct main sequences with different helium abundances in some clusters has shown that this idea is not quite true (e.g., D'Antona et al., 2005; Milone et al., 2012). It appears as though there can be a few generations of star-formation, with the chemical

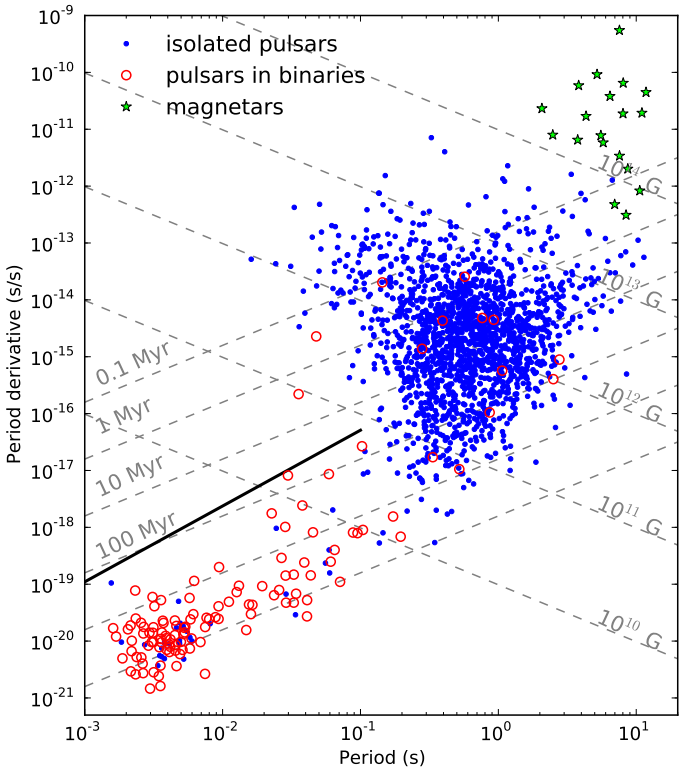


**Figure 1.2:** The pulsar $P$-$\dot{P}$ diagram, featuring all known pulsars except those in globular clusters, because the Doppler effect due to their intra-cluster acceleration makes it difficult or impossible to accurately establish intrinsic $P$ and $\dot{P}$ values. A point is placed based on the pulsar's spin period and its derivative. Dashed lines show constant characteristic age and constant magnetic field strength. Pulsars known to be members of binary systems are shown as red open circles and magnetars are shown as green stars. Note that the millisecond pulsars, in the lower-left region of this diagram, are typically members of binary systems. Pulsar radio beams are thought to "shut off" with slow rotation speed and low magnetic field strength, hence the lack of points in the lower right portion of the diagram. The thick black line is the spin-up line, below which we expect all of the recycled pulsars. The points in this diagram are not corrected for Shklovskii or Galactic-acceleration contributions to $\dot{P}$. Based on Figure 1.13 from Lorimer & Kramer (2005); data from the ATNF pulsar database at http://www.atnf.csiro.au/people/pulsar/psrcat/ (Manchester et al., 2005).

content of each generation slightly different from the last due to contamination from stellar winds and supernovae. Richer et al. (2013) have even measured dynamical differences between stellar generations in 47 Tucanae. Even with this interesting complication, the nearly uniform age of a globular cluster's population provides valuable insights into stellar evolutionary models. The most massive stars have long since evolved beyond the main sequence, becoming white dwarfs, neutron stars, or black holes. Of more than 150 known globular clusters in orbit about the Milky Way (Harris, 1996), 144 pulsars are currently known to reside in 28 of them.[1]

Most of the pulsars found in globular clusters are MSPs, which is perhaps not very surprising; if a cluster's neutron stars were formed billions of years ago when the early-type stars were running out of fuel for nuclear fusion, enough time has passed that we would expect "normal" pulsars to have spun down and ceased radio emission, and those in binary systems have had time to accrete matter and be spun up to MSP speeds. Adding to this, the high stellar densities in globular clusters can lead to binary systems losing or exchanging companions—such exchanges are thought to be responsible for a number of the pulsar systems that are observed in clusters. It is also the likely reason that MSP systems in globular clusters, unlike those in the Galactic disk, frequently have high eccentricities, with 8 currently known that have $P < 20$ ms and $e > 0.2$.

Exchange scenarios are also thought to be responsible for the high density (per unit mass) of LMXBs in globular clusters compared to the Galactic disk, with primordial neutron stars capturing main sequence companions that subsequently evolve off of the main sequence and accrete onto the neutron stars (Clark, 1975). Thus, despite the very old age of cluster populations, pulsars are still being actively recycled.

## 1.4 Why are unusual MSP binaries interesting?

There is much scientific interest in finding neutron stars in globular clusters, particularly in eccentric systems. Eccentric systems are scientifically valuable because the orbit's shape introduces a reference direction into the system: the longitude of

[1] http://www.naic.edu/~pfreire/GCpsr.html

periastron. With this reference direction present and easily measured through pulsar timing, we also gain the ability to measure whether it changes with time, and at what rate. While the orbital trajectories of an ideal, classical two-point-mass system are static, various physical phenomena, including general relativistic effects, can cause the longitude of periastron to precess. General-relativistic precession is a function of the total mass of the two bodies; it is one of five major 'post-Keplerian' (PK) parameters that can in principle be measured in a relativistic system (e.g., Damour & Deruelle, 1985, 1986), a measurement of any two of which allows both masses to be determined. Each further PK parameter then provides an independent measurement of the two masses and acts as a test of general relativity or other theories of gravity.

A famous example of the use of PK parameters is the double-pulsar system PSR J0737−3039A/B, the only case so far in which both neutron stars could be observed as pulsars: an MSP with period 22.7 ms and a slow pulsar with period 2.77 s. In this system, all five PK parameters (and a sixth contraint in the mass ratio) were used to constrain the two masses, and all contraints proved mutually consistent (Kramer et al., 2006). The double-pulsar system has a short orbital period (2.5 hours) and an eccentricity of 0.088, which is a significant eccentricity for a recycled system, and is probably due to the supernova explosion that created the second pulsar (Burgay et al., 2003). Because of the possibility of binary exchange scenarios, there are more ways in which a double-pulsar system can form in a globular cluster than in the field, and an exchange is probably the *only* way to form a double-MSP binary. An eccentric MSP-MSP binary would provide an excellent astrophysical laboratory for high-precision measurements of orbital precession and other general relativistic effects.

Another scientifically valuable system would be a pulsar in orbit with a black hole. A close orbit with such a companion could provide very precise tests of gravity (thanks to the high mass) and black hole physics. Such systems have yet to be discovered, but globular clusters may be the best places to find them, and they might have very high ($\gtrsim$ 0.9–0.99) eccentricities (Sigurdsson, 2003).

## 1.5 Pulsar surveys of globular clusters

Significant efforts have been made to find globular cluster pulsars, with increasing success in the last decade thanks to improvements in computing and search algorithms, upgrades to the 305-metre Arecibo telescope in Puerto Rico, and the construction of new high-sensitivity instruments, especially the 100-metre NRAO Green Bank Telescope (GBT) in West Virginia. While a typical survey for field pulsars consists of many adjacent telescope pointings that are each only a few minutes in duration, globular cluster pulsar surveys operate quite differently; because of its small angular size on the sky, a globular cluster typically fits entirely within the beam of a radio telescope. This is convenient, as it is necessary to point a telescope at a single cluster for several hours in order to get sufficient signal-to-noise from such distant sources. Here I give a brief overview of some of the most recent searches and their discoveries.

In 2001 and 2002, Hessels et al. (2007) surveyed the 22 globular clusters observable with Arecibo and within a distance of 70 kpc from the Sun. They took data using the Wideband Arecibo Pulsar Processor (WAPP) backend in the L-band frequency range, generally with a sampling time of 64 $\mu$s. 11 new MSPs were found in M3, M5, M13, M71, and NGC 6749, all but one of which are in binary systems.

Nearly half of the pulsars currently known in globular clusters have been found since 2005 with the GBT. Many of these are in the cluster Terzan 5, with 21 MSP discoveries reported by Ransom et al. (2005), and the discovery of the fastest-spinning pulsar yet known, PSR J1748$-$2446ad, with a spin period of only 1.40 ms (Hessels et al., 2006). These were found in data taken with the Pulsar Spigot backend (Kaplan et al., 2005). This system acquired data with 3-level sampling. The data were then correlated and accumulated and each cumulative sample was saved with 16-bit precision with a time resolution of 81.92 $\mu$s. To date, 34 pulsars have been found in Terzan 5 (e.g., Lynch et al., 2013), and 18 of these are in binary systems, 7 of which have high eccentricities.

Bégin (2006) and collaborators searched for pulsars in the globular clusters M28, NGC 6440 and NGC 6441, and one observation each in NGC 6522 and NGC 6624, using data taken on the GBT, again mostly with the Spigot backend.

23 new MSPs were discovered in these searches: 11 in M28 (7 in binaries, 2 with high eccentricities), 5 in NGC 6440 (3 in binaries, 1 with high eccentricity), 3 in NGC 6441 (1 in a binary), 2 in NGC 6522, and 3 in NGC 6624. The eccentric binaries that were found have periods from 8–30 days, while many of the non-eccentric binaries have periods less than 1 day. The eight new pulsars in NGC 6440 and NGC 6441 were followed up in Freire et al. (2008), those in NGC 6624 are discussed in Lynch et al. (2012), those in M28 will be reported in Bégin et al (in prep.), and the best information currently available for the new pulsars in NGC 6522 is in the online catalogue of globular-cluster pulsars referred to in Section 1.3.

All of the above surveys used the PRESTO[2] software package to reduce and search the telescope data, using a method known as *acceleration searching*. This allows the detection of a periodic signal whose origin is accelerating due to orbital motion, and will be described in more detail in Chapter 2. Without such a technique, many of these pulsars' fast orbits would keep them from being detected.

## 1.6 Investigating search biases

A major challenge of these surveys is in characterizing our sensitivity to various configurations of binary systems. In short, pulsars are detected as periodic signals in time-series data, and the varying Doppler shift caused by a pulsar's orbital motion dilutes the signal's power across a range of frequencies so that it is not strongly peaked at the pulsar's rotation frequency. This effect gets worse as the fraction of the pulsar's orbit contained in the data becomes larger, and so it's a problem that particularly affects our ability to find short-period binaries.

Various attempts have been made to characterize our sensitivity to pulsars in binary systems (e.g., Johnston & Kulkarni, 1991; Jouteux et al., 2002), all taking an analytical approach. Until very recently (Bagchi, Lorimer & Wolfe, 2013), all of these assumed circular orbits in order to make the problem tractable. They have also needed to come up with a metric by which to judge whether or not a system is "detected."

I take an empirical "brute force" approach to this problem; I have written code to generate large quantities of simulated survey data based on the GBT Spigot

[2] http://www.cv.nrao.edu/~sransom/presto/

searches, and I run this data through standard search algorithms. Upon determining which systems are found and which are not, the results can be histogrammed to see which parameters or combinations of parameters lead to a significant drop in pulsar-recovery.

In Chapter 2, I briefly discuss pulsar-observing. I then explain the acceleration search method and the software package PRESTO, which implements it. At the end of the chapter, I discuss previous efforts to characterize our sensitivity to pulsars in binary systems, including a brief explanation of their methods and their shortcomings. In Chapter 3, I describe the software I wrote to simulate large quantities of time series data and how this was run. In Chapter 4, I show the results obtained with this software and discuss them in the context of previous results as well as theoretical ideas concerning cluster pulsars. I end with a brief discussion of potential future directions for this work in Chapter 5.

# Chapter 2

# Background

## 2.1 Pulsar observations

### 2.1.1 Pulse profiles and TOAs

When one observes a known pulsar, data are acquired at a telescope for some duration, and the signal is dedispersed and *folded* at the pulsar's rotation period. This means cutting the time series up into segments that each contain one rotation of the neutron star and stacking these to obtain a high signal-to-noise pulse profile: an "image" of one pass of the beam across our line of sight. While individual pulses vary dramatically in profile shape (in cases where they are bright enough to be seen above the noise level), a pulsar's *average* profile is very steady in most cases; clearly there are short-timescale variations that cancel out when the signal is averaged. Thus, a pulsar's "standard profile" may be constructed by stacking many pulses to reduce the noise to a low level.

A well-constructed standard profile can then be matched to future observations that have been folded in order to measure the precise time at which the pulsar passed through a particular point in its rotation—the peak of its pulse profile, for instance. This surprisingly powerful piece of data is known as a *time of arrival* (TOA), and I will discuss its application shortly.

### 2.1.2 Dispersion measure

The vast space between stars is populated by free electrons that act as a dispersing medium. This is the ionized component of the *interstellar medium* (ISM) and its effect is a frequency-dependent delay in the TOAs of the pulses as seen at Earth. Since any observation has some finite bandwidth, it is always necessary to correct for this dispersion, otherwise the broad-band profiles are smeared out. The conceptually simple way of doing this is to split the full band into narrow subbands, each of which has a relatively negligible level of dispersion smearing. These bands can then be rotated into alignment before being summed to obtain the full signal. This is known as incoherent dedispersion. A more sophisticated method, coherent dedispersion (Hankins & Rickett, 1975), removes dispersion effects completely by convolving the signal (whose amplitude and phase are both recorded) with the inverse of the ISM transfer function.

The amount by which the signal needs to be corrected can be characterized by a single number, the *dispersion measure* (DM). The DM is a measure of the column density of electrons along the line of sight to the pulsar. This turns out to be a useful property: using a model of the electron distribution in the Galaxy (Cordes & Lazio, 2002), a pulsar's DM can be used to estimate its distance. This is not extremely accurate in practice, but it is often the only distance estimate available. Conversely, if we have independent measurements of the distances to a number of pulsars, their DM values can be used to help construct a model of the Galactic electron distribution.

### 2.1.3 Pulsar timing

While this work does not use the technique of pulsar timing itself, a brief discussion on the subject is warranted since the technique will be mentioned and pulsar-timing software is used for some important calculations.

Although pulsar timing involves very complicated many-parameter models, the principles behind it are quite simple. A pulsar spins (and spins down) with predictable regularity, and yet there are apparent variations in its spin-rate as seen by an Earthbound observer. There are many distinct sources for these variations both local and distant, mostly due to the Doppler effect, and we see a superposition of

all of them. The goal of pulsar timing, then, is to account for all of these variations in a model whose parameters (when fit to data) contain a wealth of information about the behaviour of the observed system.

The data that are needed for this are pulse TOAs. These are precise times at which the pulsar is known to be at some chosen fiducial point in its rotation, for instance the peak of the pulse profile, and are obtained by taking several minutes of data and folding to get a high signal-to-noise pulse profile. There is no need (nor is it possible) to have such a measurement for every single rotation; we just need to "check in" occasionally to see if the rotation phase is where the model predicts that it should be, and to make minor adjustments if necessary. As long as the time between two TOAs is short enough that the current model is wrong by only a small fraction of the rotation period, *phase connection* is maintained and we can be confident that every rotation that took place between TOAs has been correctly accounted for. If simply adjusting parameters cannot account for the data, additional parameters may be needed. Pulsar timing becomes a game of looking for patterns in fit residuals and recognizing which parameters need to be added to remove those patterns.

Distant sources of pulse-period variations include a periodically varying Doppler effect from motion within a binary system and delays due to the signal traversing strong gravitational fields, while local causes come from the Earth's rotation and its motion within the Solar System, which need to be corrected using accurate Solar System ephemerides and a good model of the Earth's rotation.

A commonly-used software package for pulsar timing is TEMPO[1] (and its successor, TEMPO2[2]). It takes TOAs and parameter estimates as inputs and performs a fit using those parameters. TEMPO also has a 'prediction mode' that uses a set of parameters (including binary system elements) to predict the observed spin frequency and the phase at a particular point in time and at a particular telescope or the Solar System barycentre. This is done using a polynomial expansion, and the output is a set of polynomial coefficients, or *polycos*. This is the TEMPO function that is used as part of the data-simulation routine in this work.

[1] http://tempo.sourceforge.net
[2] http://tempo2.sourceforge.net

## 2.2 PRESTO and acceleration searching

PRESTO is a software package designed primarily to detect pulsar signals in time series data. It performs many functions, both closely and peripherally related to this goal, including detecting and removing radio-frequency interference (RFI); dedispersing and barycentring data; acceleration, single-pulse, and sideband pulsar-searching; pulsar candidate sifting and folding; and, through various utilities, visualizing, manipulating, and analyzing data. Ransom (2001) provides an overview of PRESTO, while Ransom, Eikenberry & Middleditch (2002) describe the acceleration search in particular, which is the primary method employed in globular cluster pulsar surveys with the GBT, and is therefore the method that is tested in this work.

A typical pulsar search using PRESTO follows a series of steps. Raw telescope data contain information over time but also over frequency, the range of the latter depending on the observing bandwidth. Ideally any signals present above the noise level are cosmic in origin, but of course there is always radio-frequency interference (RFI). The PRESTO script `rfifind` masks out the most obvious RFI, which it identifies as a signal that is not persistent either in time or across the frequency band—either of these indicates that the signal does not come from a pulsar. PRESTO currently handles raw data from various machines at the GBT, Arecibo, Parkes, and Jodrell Bank, as well as the pulsar standard SIGPROC format and simple floating-point time series.

Before searching for pulsar signals, the data are barycentred—the arrival times of the data corrected to the Solar System barycentre—and the frequency dimension is collapsed by de-dispersing over a range of DM values one wishes to search, generating a one-dimensional time series for each step in DM. It is beneficial to first create a *topocentric* (non-barycentred) time series at $\mathrm{DM} = 0$, however, as another check for RFI. Even if a signal is found in this time series that appears very pulsar-like, it must be terrestrial in origin or else its signal would be dispersed and undetectable in a $\mathrm{DM} = 0$ time series.

Once a set of de-dispersed, barycentred time series is made, a Fourier transform is carried out on each of them. One of the strengths of PRESTO is its ability to efficiently calculate the fast Fourier transform (FFT) of very long high-resolution time series that may be much too large to work with in their entirety in computer

memory.

Having thus prepared the data, the `accelsearch` script carries out the actual search, known as an *acceleration search*. Because a pulsar signal is typically a repeated narrow spike, much of the power in Fourier space is in harmonics rather than the fundamental frequency. Harmonic-summing is performed in order to recover this power, and the user can specify whether to sum 1, 2, 4, 8, or 16 harmonics. A signal from a source that is accelerated due to motion in a binary system will have harmonics that drift across Fourier bins (linearly if the acceleration is constant), with higher harmonics drifting proportionally further than lower harmonics. For a particular search, one specifies the maximum number of Fourier bins the highest harmonic is allowed to drift; if this is set to zero, a standard Fourier-domain search is conducted rather than an acceleration search. When this value is set higher, pulsars undergoing stronger linear accelerations may be detected. However, the time it takes to perform a search increases linearly with the number of harmonics summed and with the number of bins we allow the highest harmonic to drift.

Once the acceleration search is complete, we are left with a list of pulsar candidates and their properties. Additional functions in PRESTO are used to sift through these and to match candidates at neighbouring DM values that may be the same source.

## 2.3 Detectability studies

It has been evident for many years that we are not equally sensitive to all regions of parameter space when it comes to searching for pulsars in binary systems, but characterizing our sensitivity has presented an ongoing challenge. Here I briefly cover a few significant efforts that have been made.

An early approach to the problem can be found in Johnston & Kulkarni (1991). Motivated by early globular cluster searches and the expectation for many cluster pulsars to be in short-period binary systems, they approach the problem analytically, defining a stationary pulsed signal as a sum of Fourier components whose power spectrum is straightforward to calculate. They consider a time-varying distance to the source that can be expanded in a Taylor series, providing a velocity term, an acceleration term, and higher derivatives. They show that the acceleration

and higher-order terms produce time-dependent phase errors that cause a reduction in the amplitude of the signal's harmonics in the power spectrum. The extent to which these amplitudes are reduced is taken as a measure of how difficult these systems are to detect.

To quantify the reduction in detectability, they define the function

$$\gamma(\alpha_1, \alpha_2, T) = \frac{1}{T}\left|\int_0^T \exp\left[\frac{im\omega_p}{c} \times \left(v_0 t + \frac{a_0 t^2}{2!} + \frac{j_0 t^3}{3!} + \ldots - \alpha_1 t^2 - \alpha_2 t\right)\right] dt\right|, \tag{2.1}$$

where $T$ is the observation length, $m$ is the harmonic number of the power spectrum peak under consideration, and $\omega_p$ is the angular frequency of the pulsar. $c$ and $i$ are the speed of light and $\sqrt{-1}$ respectively, and $v_0$, $a_0$, and $j_0$ are the line-of-sight velocity, acceleration, and jerk (first derivative of acceleration) respectively at time $t = 0$. $\gamma$ takes values between 0 and 1.

$\gamma^2$ is the ratio of the peak height of the $m$th harmonic in the power spectrum when the signal is accelerated compared to when it has a constant velocity. $\alpha_1$ and $\alpha_2$ are free parameters. The authors test the effectiveness of a standard Fourier analysis (no acceleration search) by setting $\alpha_1 = 0$, and $\alpha_2$ is chosen to maximize $\gamma$. $\alpha_2$ can be thought of as a mean velocity over the integration time $T$. For testing the effectiveness of an acceleration search, both $\alpha_1$ and $\alpha_2$ are varied to maximize $\gamma$. The authors wrote code that maximizes $\gamma$ for circular orbits with chosen binary parameters such as inclination angle, orbital period, and companion mass, as well as integration time $T$, and they run this over a grid of orbital and spin period values. They take $\gamma = 0.5$ as their detection threshold, acknowledging that this is arbitrary, as a strong signal reduced to a fraction of its strength may still be much easier to detect than an undiminished weak signal. Regardless, they are able to demonstrate that a pulsar signal is reduced most severely for fast-spinning pulsars in tight binaries.

Jouteux et al. (2002) define their own set of $\gamma$ functions for various search methods, including the standard Fourier analysis and the acceleration search. These functions carry a meaning similar to the $\gamma$ of Equation 2.1, in that they represent the loss of signal strength due to source acceleration, with a larger value of

$\gamma$ (closer to 1 than 0) implying a better-recovered signal. Their main comparison is between the acceleration search method and their proposed 'partial coherence recovery technique' (PCRT). A method implemented in PRESTO known as a 'side-band search' is based on the same principles, which involve taking advantage of the binary periodicity. So, unlike the acceleration search, which breaks down if the time series contains whole orbits or even an appreciable fraction of an orbit, the PCRT method works best if the time series covers several orbits.

While interesting insights are made in both of these studies, there are notable limitations owing both to computing power and the analytical expressions used. Neither can test for the effect of pulse amplitude compared to the noise level of the time series, and both are restricted to circular orbits for the purpose of tractability. In Johnston & Kulkarni (1991), the amount of additional calculation required to handle eccentric orbits would have been unfeasible with the computing power available. Jouteux et al. (2002) describe the Fourier response of their signal using simple Bessel functions, which would not be possible had they used non-circular orbits. Recently, the non-circular-orbit problem was addressed by Bagchi, Lorimer & Wolfe (2013), who extend the methods of Johnston & Kulkarni (1991) to include orbital eccentricities.

Bagchi, Lorimer & Wolfe (2013) define $\gamma_{1m}$ and $\gamma_{2m}$, identical to Equation 2.1 with and without $\alpha_1 = 0$, respectively, and $\gamma_{3m}$, which includes an additional $t^3$ term, testing the idea of a 'jerk search', which corrects not just a linear acceleration of the source, but also the first derivative of that acceleration (this is rarely used in practice as it greatly increases the processing time). The subscript $m$ is simply a notational clarification, denoting the harmonic under examination. The extension to eccentric systems significantly increases the calculation required to determine the line-of-sight terms in these expressions; in particular, Kepler's equation must be solved iteratively to find the true anomaly at each point in time. Furthermore, a weighted-average $\gamma_{nm}$ is calculated over the integration time in each case, accounting for the varying amounts of time spent at different points in an eccentric orbit. This extension to eccentric systems means that more parameters need to be set, namely the eccentricity and argument of periastron, though tools are provided by the authors to test the signal degradation for any given set of parameters.

The results of the new study are consistent with those of Johnston & Kulkarni

(1991), but the inclusion of arbitrary eccentricties makes them more robust. In each case, the importance of a higher-order search, such as an acceleration search compared to a straightforward Fourier analysis, is greater when searching for faster-spinning pulsars.

I have approached the problem of characterizing our sensitivity over the parameter space empirically, rather than analytically. In keeping with the recent survey practices, particularly those of the GBT, I use PRESTO's `accelsearch` function to search large quantities of artificial survey data made up of noise and randomly-generated binary-system pulsar signals. While past studies use $\gamma^2$ as their best analytical stand-in for the fraction of pulsars recovered, here I am able to use the recovered fraction directly, and the amplitude of individual pulses is a variable parameter. In the next chapter, I discuss my methods for generating the artificial data and searching it with PRESTO.

# Chapter 3

# Methods

In this chapter I outline the basic ingredients of my simulated survey data, how they are generated, and how they are searched using the PRESTO software.

## 3.1 Parameters of the artificial pulsars

I generate pulsars in binary systems by drawing parameters from a set of chosen distributions. These distributions are listed in Table 3.1 with those that are non-uniform illustrated in Figure 3.1. The inclination angle $i$ is actually drawn from a distribution that is uniform in $\cos i$, as this is the actual distribution of orbits that are oriented completely at random. For the companion mass, most of the distribution is uniform between 0.05 and 3 $M_\odot$, since pulsar companions tend to be white dwarfs

**Table 3.1:** The distributions from which pulsar parameters are selected. Those that are not uniform are illustrated and explained in Figure 3.1.

| **Parameter** | **Range** | **Uniform?** |
|---|---|---|
| Period ($P$) | 1–20 ms | yes |
| Eccentricity ($e$) | 0–1 | yes |
| Binary period ($P_b$) | 10–1440 minutes | yes |
| Longitude of periastron ($\omega$) | 0–360 degrees | yes |
| Inclination angle ($i$) | 0–90 degrees | no |
| Companion mass ($m_2$) | 0.5–10 $M_\odot$ | no |
| Pulse amplitude | 0.01–0.1 | no |

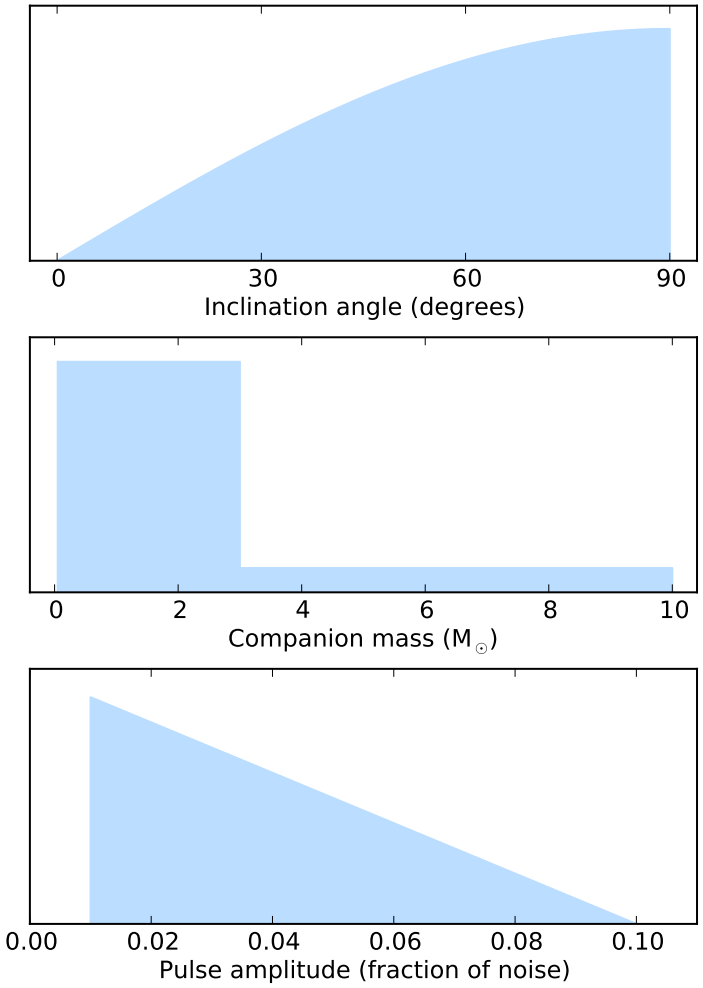


**Figure 3.1:** Distributions for three of the parameters randomly chosen when generating pulsars. These are the only three that are not uniformly distributed. The inclination angle $i$ is actually uniform in $\cos i$, as is the case for any completely random distribution of orbit orientations. The companion mass is split between two uniform distributions with 80% between 0.05 and 3 $M_\odot$ and 20% between 3 and 10 $M_\odot$. The pulse amplitude distribution decreases from 0.01 to 0.1 because the detection of lower-amplitude pulsars is more interesting.

or neutron stars, but 20% of the distribution is uniform between 3 and 10 $M_\odot$, for companions that are non-degenerate stars or black holes. This distribution was simply chosen out of a desire for better statistics in the lower-mass range.

The orbital periods are distributed between 10 minutes and 1 day. The shortest pulsar orbital period currently known is that of PSR J1311−3430, at 94 minutes (Pletsch et al., 2012). Because short-period compact objects spiral inwards due to the emission of gravitational radiation, neutron stars should exist in extremely short-period binaries, although it is not clear to what point they would continue to behave as pulsars. A lower limit of 10 minutes is taken as plausible while sufficiently below the current observed limit.

The fastest-spinning pulsar is PSR J1748−2446ad with a spin period of 1.40 ms (Hessels et al., 2006), so 1 ms is taken as a suitable lower limit for the spin period distribution. The upper limit of 20 ms is a period sometimes taken as that below which pulsars may be considered *millisecond* pulsars (though recycled pulsars slower than this are often still referred to as MSPs).

Once the parameters for a particular pulsar are selected, they are saved in a TEMPO-style parameter file, or *parfile*. Use of these parfiles allows me to take advantage of TEMPO's prediction mode for performing orbital calculations.

## 3.2 The pulse profile

A simulated pulse profile is generated as a closely-spaced superposition of 1 to 5 narrow Gaussian curves, in order to mimic the structure observed in real pulse profiles; the peaks are scaled so that the maximum point of the superposition is at unit height. The profile is then a function with values between 0 (the baseline) and 1 (the peak) whose height can be returned as a function of phase, which is also a value between 0 and 1. Each profile is associated with a TEMPO parfile, containing the pulsar's rotational and binary parameters.

The placement of the Gaussian curves proceeds as follows. An overall 'width scale' $W$ is chosen from a truncated normal distribution with mean 1.0 and standard deviation 0.2, constrained between 0.1 and 1.9. The width (Gaussian $\sigma$) of each curve is chosen from a normal distribution with mean $0.01W$ and standard deviation $0.001W$. The position in phase of a particular curve is chosen from a normal

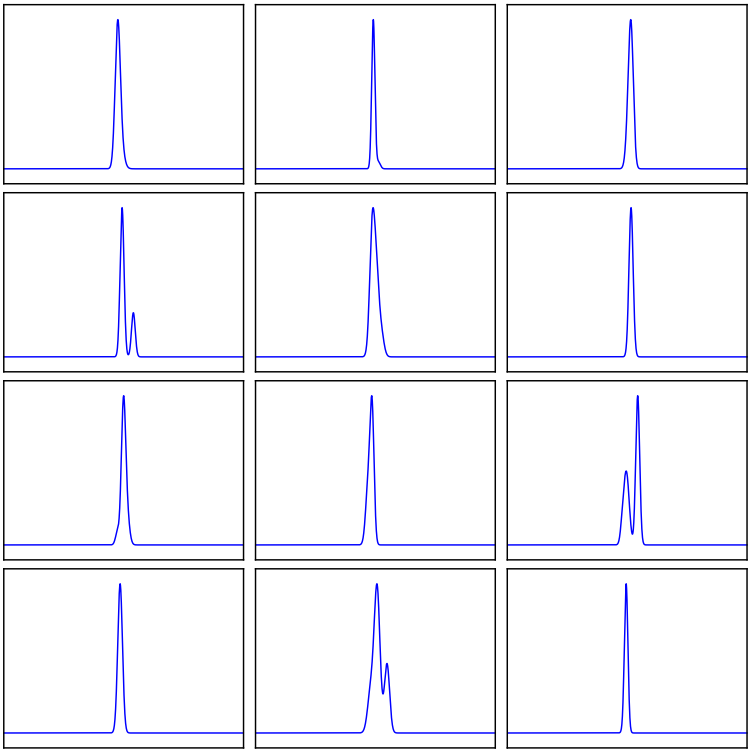

**Figure 3.2:** A set of twelve randomly-generated pulse profiles constructed as described in the text. Profiles such as these are made for every pulsar generated. The horizontal axis in each frame is rotation phase and runs from 0 to 1. The vertical axis is dimensionless pulse amplitude, with the baseline at 0 and the peak of the pulse at 1.

distribution with mean 0.5 and standard deviation $0.02W$. The first curve placed is given a height of 1. The height of each subsequent curve is set by multiplying the previously placed curve's height by the square of a uniform random number between 0 and 1. Once all of the Gaussian curves are set, they are scaled so that the maximum point of their superposition (i.e., the pulse profile) has a height of 1. All of this simply comes from tuning numbers to obtain what look to be reasonably consistent and realistic pulse profiles with some random variation. A set of sample profiles is shown in Figure 3.2.

## 3.3 The time series

Artificial time series are created with a time resolution of 81.92 $\mu$s, the same as the resolution in the GBT globular cluster surveys. They are 1000 s (12.2 million data points) long. Typically, in real surveys, long (several-hour) globular cluster observations are made, but the time series are split up into overlapping segments of about this length in order to find pulsars in tight binaries. By generating these shorter time series directly, I am able to spend my computation time on a much larger set of artificial pulsars than if I generated long time series and split them up. There is certainly an orbital-phase dependence to the difficulty of detecting a pulsar in an eccentric orbit, as will be seen in Chapter 4, and the reader should bear in mind that in a real survey with long observations, a pulsar in a short orbit that is highly accelerated along the line of sight in one $\sim$1000 s time series may be accelerated much less in another from the same long observation. To ensure that orbital phase is realistically sampled, at the start of a time series each pulsar is at a uniformly random mean anomaly between 0 and 360°. This means that a pulsar in an eccentric orbit is more likely to be near apastron than periastron, due to the relative durations spent at each point in the orbit.

The effects of DM are ignored in these time series, or to put it another way, all of the pulsars are assumed to have the same single DM to which the telescope data would have been de-dispersed. This is not far from the real situation when observing globular clusters, since all of the pulsars in the beam are at almost the same distance from us, and therefore have almost the same DM. Thus, I can add noiseless pulsar signals directly to an array of white noise.

If I wished to construct a noisy time series containing a set of signals belonging to pulsars that are *inertial* relative to the reference frame (the Solar System barycentre), I could simply add the appropriately-scaled repeated pulse profile to a time series of noise. At each time step, the code would check precisely how much time had passed, find the rotational phase of the pulsar at that moment (trivial for an inertial pulsar), and get the value of the pulse signal at that phase. I might want to take a typical pulsar spin-down into account, but over a single observation its effect on the spin period would be negligible.

However, matters become much more complicated when the pulsar orbits an-

other body. The apparent spin period is modulated about the intrinsic period according to a varying Doppler shift, which can be determined by the pulsar's velocity along the line of sight to the Solar System at any given time. The tricky step is calculating the instantaneous line-of-sight velocity, whose periodic behaviour depends upon the orbit's shape and orientation, and at a far more precise level, on general relativistic effects too, which cause the orbit to deviate from a simple Keplerian ellipse. All of these factors can be accounted for by using TEMPO's prediction mode.

The particular binary system model that is used is the DD (Damour & Deruelle, 1986) theory-independent relativistic model, although I calculate and include all of the post-Keplerian parameters with the appropriate general relativistic equations, effectively making it equivalent to the DDGR variation of the DD model (Taylor & Weisberg, 1989), in which general relativity is assumed to be correct. I did this because of an error (subsequently fixed) in the DDGR code that sometimes caused the program to crash when making predictions for systems with very high eccentricities.

Given a set of parameters for a pulsar system (including a reference epoch for the parameters), a start time for the range of validity of the expansion, and the length of the range of validity, a set of polynomial coefficients, or polycos, is output that can be used to calculate the rotation phase $\phi$ as

$$\phi = \phi_{ref} + 60 f_0 \Delta T + c_1 + c_2 \Delta T + c_3 \Delta T^2 + \ldots \tag{3.1}$$

at a particular time $T$, where $\Delta T = (T - T_{mid}) \times 1440$. Here $f_0$ (the reference rotation frequency) is in units of hertz and $T$ and $T_{mid}$ (the midpoint of the chosen range of validity) are in days, so the factors of 60 and 1440 scale everything to minutes. This is an arbitrary choice of units, of course, as $\phi$ itself is dimensionless. The coefficients $c_n$ are the polycos themselves and are determined by the same kind of many-parameter model TEMPO uses in pulsar timing—in this case, the parameters are simply input, not fit.

I use 12 polycos per segment, where a segment is the range of validity of the polycos, which is taken to be the whole time series if the pulsar has a long enough orbit, but which I never allow to be greater than one-fifth the length of the orbit.

When necessary, multiple sets of polycos are generated throughout a time series. To keep memory-usage at a manageable level, time series are made in chunks of size $2^{22}$ (4.2 million) points, each chunk appended to a file on disk before generating the next.

Pulsar signals are superposed onto white noise, with the pulse amplitude defined relative to the noise level. This amplitude is proportional to the signal-to-noise ratio (SNR) that appears in the pulsar radiometer equation (Dewey et al., 1985; Hessels et al., 2007),

$$S_{min} = \frac{(\mathrm{SNR})\beta T_{sys}}{G\sqrt{n\Delta\nu t_{obs}}} \left( \frac{w}{P_0 - w} \right)^{1/2}, \tag{3.2}$$

which is used to estimate the flux density detection threshold for a survey. Typically the SNR here is the cutoff below which a candidate is not considered. The $\beta$ term is a correction factor to account for the signal digitization, and is taken to be 1.2 for the 3-level sampling of the Arecibo and GBT surveys (Hessels et al., 2007). $T_{sys}$ is the system temperature and $G$ is the telescope gain which converts between temperature and flux density. These are 20 K and 2.0 K/Jy respectively for the GBT survey, and 40 K and 10.5 K/Jy for Arecibo. The number of orthogonal polarizations summed is $n$, which is 2 for both surveys. The observing bandwidth $\Delta\nu$ is 600 MHz for the GBT (Ransom et al., 2005) and 100 MHz for Arecibo, and the observation length is $t_{obs}$. The pulse period and pulse width are $P_0$ and $w$. Equation 3.2 produces a value with dimensions of flux density.

To convert the amplitudes of my simulated pulses into flux density, I begin with Equation 7.1 from Lorimer & Kramer (2005),

$$\mathrm{SNR} = \frac{1}{\sigma_p \sqrt{w_{eq}}} \sum_{i=1}^{n_{bins}} (p_i - \bar{p}), \tag{3.3}$$

where $\sigma_p$ is the noise level in the folded pulse profile, and thus $\sigma_p = \sigma_r \sqrt{t_{obs}/P_0}$, where $\sigma_r$ is the noise level in the raw time series. The equivalent width $w_{eq}$ is the width of a top hat function with the same integrated area and peak height as the pulse profile, in bins. Thus, if $w_p = w/P_0$ is the width as a fraction of the full rotation period ($w_p = 0.03$ is the average value in my simulation), $w_{eq} = w_p \times$

$P_0/t_{samp}$, where $t_{samp}$ is the sampling rate (81.92 $\mu$s). The last term is simply a sum of the amplitudes of the bins across the profile. If we imagine again the top hat function with the pulse's peak height $h$ and width $w_p$, we can replace this term with $\frac{w_p P_0}{t_{samp}} h$. Combining all of these we get

$$\mathrm{SNR} = \sqrt{\frac{t_{obs} w_p}{t_{samp}}} \frac{h}{\sigma_r}, \tag{3.4}$$

noticing that the last term $h/\sigma_r$ is equivalent to the peak pulse amplitude as a fraction of the time series noise level, as I have defined my amplitudes.

Pulse-to-pulse fluctuations are not modeled; individual pulses are equivalent in shape to the average summed pulse profile. The rotation phase at each time step is efficiently calculated using the Numerical Python[1] polynomial evaluation. $\Delta T$ in Equation 3.1 is input as a precise array of times, and $\phi$ is returned as an array of phases, which is then used as input to obtain an array of signal strengths between 0 and 1. This array is scaled as needed and added to the time series array, and this process is repeated for each pulsar. The final array is saved in the PRESTO time series format (a simple array of 4-byte numbers), and a corresponding time series information file is generated, containing the start time of the data, the number of bins, the time resolution, and a flag indicating that the data are barycentred.

## 3.4 Creating and searching a large quantity of data

I generate these data on a parallel-processing system, using 36 CPUs: there are 3 nodes that each consist of 12 CPUs. For every run of the program, each CPU generates 10 time series 1000 s in length, with each time series randomly containing up to 20 pulsars. A time series has a corresponding directory containing the TEMPO parfiles for all of its pulsars. These time series are then searched using PRESTO's `accelsearch` function. As explained in Section 2.2, the number of harmonics to be summed and the number of Fourier bins the highest harmonic can drift must both be input. Here the first 8 harmonics are summed and the maximum drift is set at 300 bins. Frequencies between 1 and 10,000 Hz are searched.

[1] http://www.numpym.org/

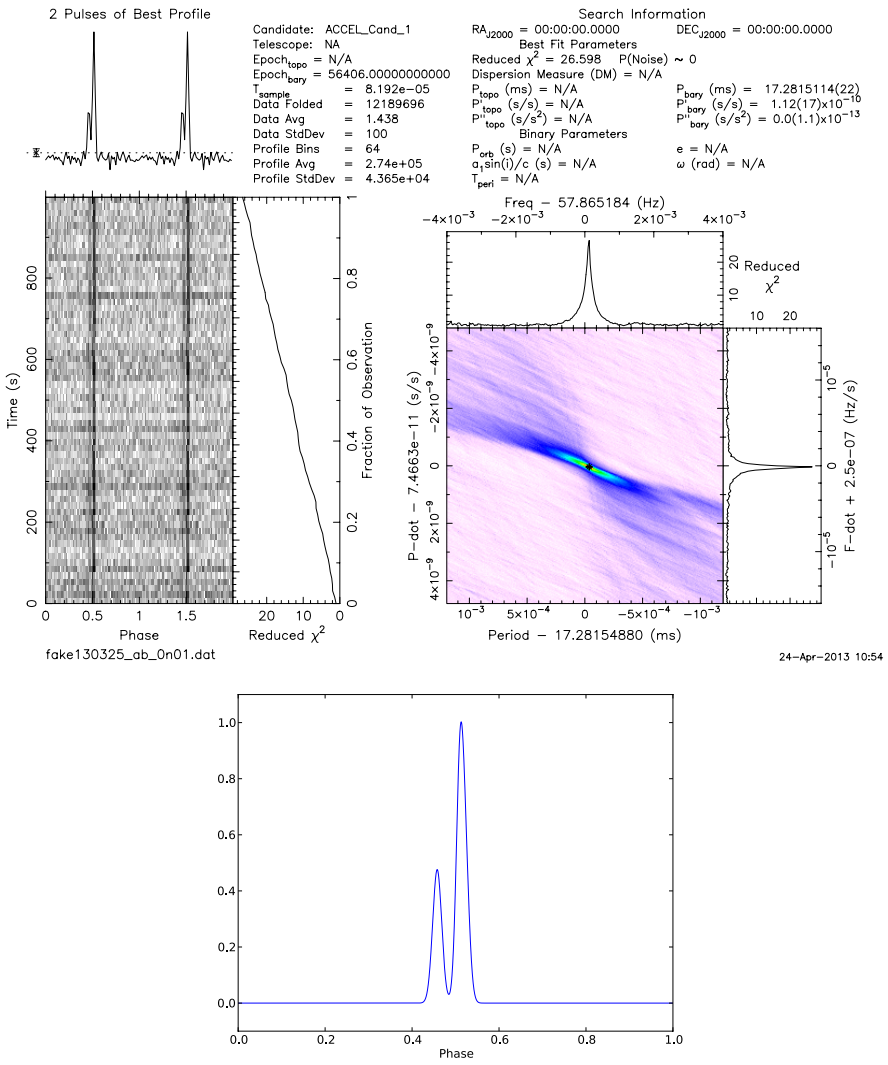


**Figure 3.3:** An example of a PRESTO candidate (upper plots and text) and the profile of the artificial pulsar that is detected here (lower plot). The uppermost PRESTO plot is the summed profile and the plot below that shows the persistence of the signal over time, both plotted over two rotations in phase. The plot to the right is the significance of the signal in period-vs-period-derivative space. The amplitude of this pulsar is 8.5% of the noise level, which is quite high in the amplitude distribution.

The product of an acceleration search is a list of pulsar candidates. This list is parsed and used to fold the time series using the frequency and frequency derivative of each candidate. We are thus left with a set of folded PRESTO candidate files, or *pfd* files, which PRESTO can display graphically. The graphical candidate for one of the artificial pulsars is shown in Figure 3.3 along with the pulsar's generated profile. Slight differences between the input and candidate profiles arise from the noise and the coarser binning of the candidate. While not indicated in the figure, the amplitude of this pulsar is 8.5% the noise level, which is quite high in the distribution (Table 3.1), so that a pulse every $\sim$17 ms in a 1000 s time series leads to a folded signal $\sim$20$\times$ the height of the noise in the upper-left plot. A similar pulsar with an amplitude of 1% the noise level (the low end of the distribution) would have a folded profile at 2–3$\times$ the noise level.

## 3.5 Matching candidates

At this stage, I have many time series, for each time series I have a set of parfiles containing the information about the pulsars contained within, and I have a set of candidates found by PRESTO. The goal is to figure out which of the pulsars were actually found—i.e., to determine which of the pulsars are among the candidates.

In a real survey, a human expert would look through the candidate plots and mark those that look like pulsars for further investigation. It is possible in this scenario for a candidate that represents a real pulsar to be disregarded as random noise or RFI. In this work, I consider a pulsar that has been recognized by PRESTO as a candidate as 'detected' and ignore the possibility of human failure beyond this point, since it would be very difficult to properly account for this in any kind of realistic way. The enormous number of candidates produced also prohibits my looking at each one individually.

There are certain to be candidates that, while present due to an injected pulsar signal, show up not at the actual pulse period but rather at some harmonic ratio of the period. In an effort to reduce false positives, I do not attempt to identify these; instead, I assume that these pulsars will also show up as candidates at their true periods (modulated by binary motion).

The matching process proceeds as follows for each time series.

- A list of the parfiles associated with the time series is created. For each entry, polycos are generated to determine the apparent spin period at the start of the time series, which in general is not the same as the spin period listed in the parfile because of the pulsar's line-of-sight velocity in a binary system. Each entry contains the pulsar label and this apparent period.

- A list of the candidates associated with the time series is created. Each entry contains the candidate label and the spin period of the candidate.

- For each candidate, a list of pulsars for which $|1 - \frac{P'_P}{P_C}| < 10^{-5}$ is created, where $P'_P$ and $P_C$ are the apparent pulsar period and the candidate period, respectively. This is usually either an empty list or a list containing 1 item.

- Taking the best match from each candidate's list (the match with the smallest value of $|1 - \frac{P'_P}{P_C}|$), make a list of the pulsars that were detected.

This is done for each time series, and a database of all of the pulsars is created, along with many of the parameters contained in their parfiles and whether or not they were detected. In total, I generated 1,782,261 pulsars, of which 1,039,258 were detected. I will explain these results in detail in Chapter 4.

# Chapter 4

# Results and Discussion

## 4.1 Detections across parameter space

Of 1.8 million pulsars generated, about 60% were detected in the acceleration searches described in Chapter 3. These results are shown as histograms of the fraction of pulsars recovered, displaying the detection distribution over one or two parameters while marginalizing over the remaining parameters. The colour-mapping in the 2-dimensional histograms is chosen so that regions with more than 50% of pulsars detected are in blue and regions with less than 50% detected are in red. For ease of communication, I will often refer to blue and red regions as "detected" and "undetected" respectively, even though the colours represent a continuous range between 0 and 1.

One of the motivating goals of this simulation is to observe the effects of non-zero eccentricities on the detectability of pulsars in short-period binary systems. As seen in Figure 4.1, the detected pulsars actually push down to shorter orbital periods at higher eccentricities. This is because, while an eccentric system is very highly accelerated near periastron when compared to the acceleration of a non-eccentric system with the same orbital period, it moves through this portion of its orbit very quickly and spends most of its time undergoing relatively little acceleration. This phase-dependence of eccentric orbits is evident in Figure 4.2, which shows the fraction of pulsars detected as a function of mean anomaly and eccentricity. The detectability of pulsars at periastron (mean anomaly 0) begins to drop

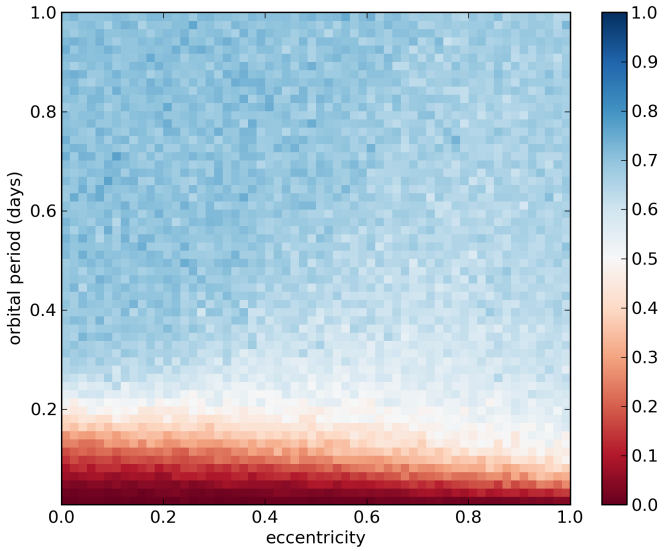


**Figure 4.1:** Pulsar detections across eccentricity and orbital period. Because of the relatively long time spent on the slow, distant portion of its orbit, a pulsar in a high-eccentricity system can be found at somewhat shorter orbital periods than one in a circular orbit. In this and all subsequent 2-dimensional histograms, the red-blue colour gradient represents the fraction of pulsars recovered.

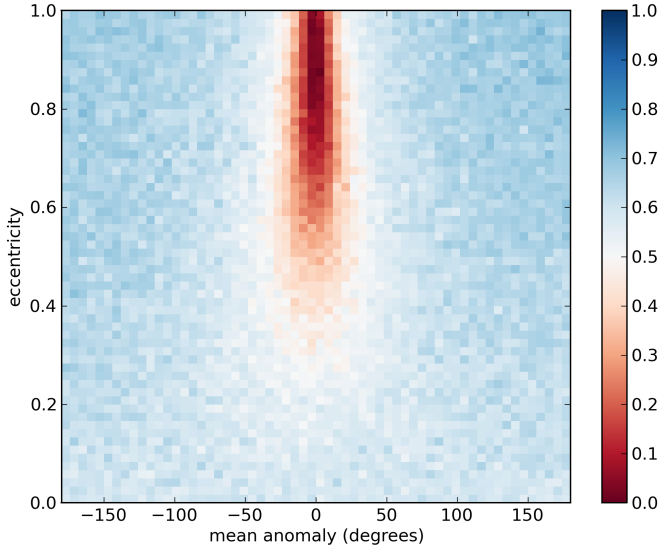


**Figure 4.2:** Orbital phase dependence of detectability. Pulsars with higher-eccentricity orbits are not detected at periastron (mean anomaly 0) due to very high accelerations there but are easier to detect at other mean anomalies than less eccentric systems.

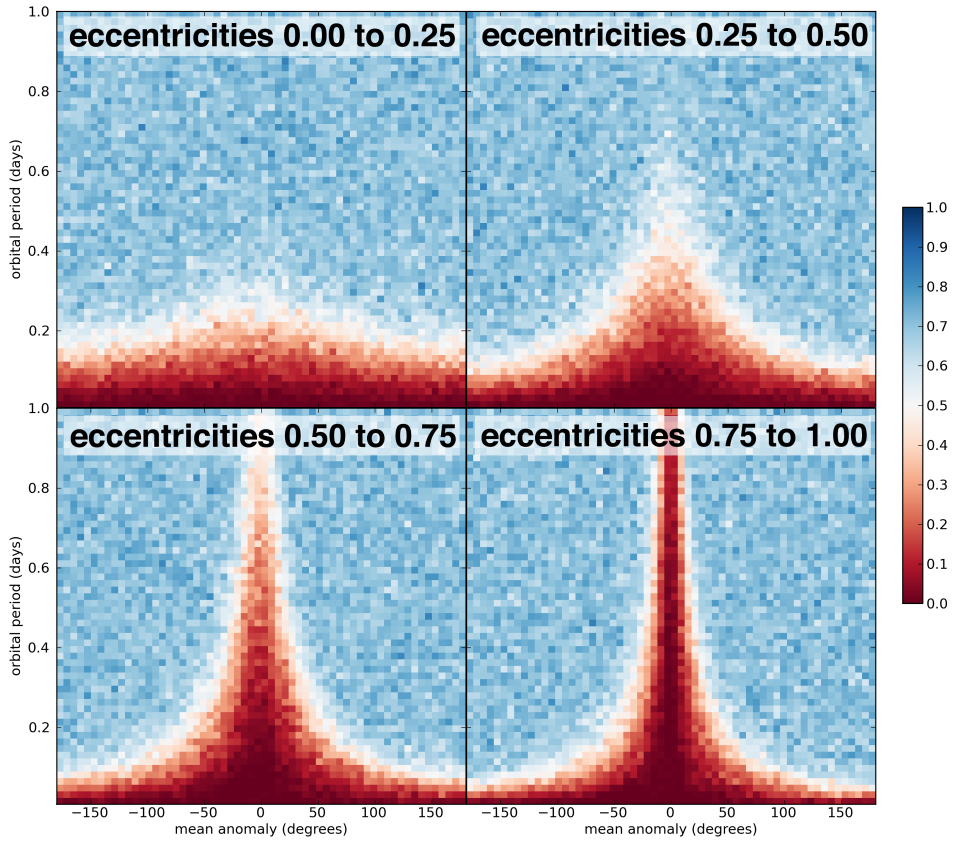


**Figure 4.3:** Another look at the orbital phase dependence. Each of the four panels integrates one quarter of the total eccentricity range in order to show a progression from low to high eccentricities. While pulsars near periastron (mean anomaly 0) become increasingly difficult to detect at higher eccentricities, those at other orbital phases are detected at ever-shorter orbital periods. An eccentric binary system spends very little time at periastron, so this explains the slight improvement in detections seen at high eccentricities in Figure 4.1.

noticeably at $e \approx 0.3$, reaching almost zero by $e \approx 0.7$. To understand the high-eccentricity behaviour in Figure 4.1, we can look at the three parameters ($e$, $P_b$, and mean anomaly) of Figures 4.1 and 4.2 simultaneously.

Figure 4.3 plots detections as a function of mean anomaly and orbital period integrated over four eccentricity ranges. The phase-eccentricity relationship seen in Figure 4.2 is apparent as the increasing spike of the undetected region at periastron.

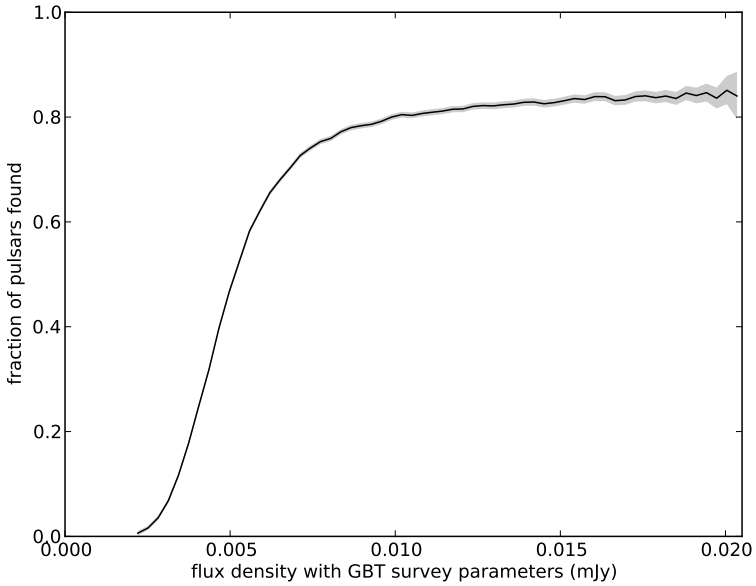


**Figure 4.4:** Pulsar detections as a function of GBT survey flux density. The grey region surrounding the black line represents the statistical uncertainty in the histogram due to the sample size. The sample size decreases linearly towards higher amplitudes.

However, as the central spike increases, the wings of this region are pushed down to lower orbital periods, meaning that shorter-period binaries are detected there. This improvement over most of the orbit leads to the overall increase in the detectability of short-orbit pulsars with increasing eccentricity.

The amplitude of the pulses are scaled to units of flux density using Equations 3.2 and 3.4 with the GBT survey parameters. This is plotted in Figure 4.4, where it is seen that the lowest amplitude pulses are not found at all; there is a steep increase in detectability at low flux densities, but by $\sim$0.01 mJy, brightness makes little difference. As discussed in Section 3.1, brighter pulsars are generated less frequently in this simulation, so the high-amplitude region of Figure 4.4 has greater uncertainty.

When eccentricity is introduced, a symmetry is broken: the orbit has a well-defined periastron, and we may ask how the position of periastron affects detectability. A histogram across longitude of periastron $\omega$ and eccentricity $e$ is

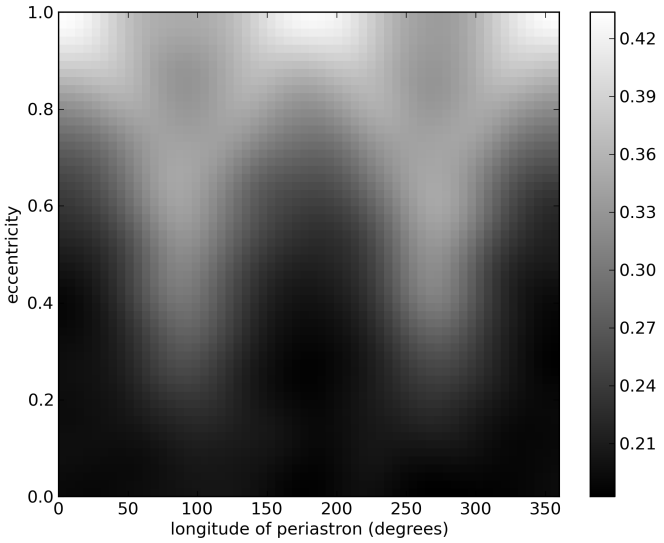


**Figure 4.5:** Detections across longitude of periastron $\omega$ and eccentricity $e$. To increase the contrast, this is only for orbital periods $< 0.2$ days (above which the effect of $\omega$ is weaker), and the greyscale is mapped only onto the range of values present in the plot. Gaussian smoothing is also applied, with $\sigma$ of 3 bins. The change in behaviour at $e \approx 0.7$ is discussed in the text.

shown in Figure 4.5. The effect of $\omega$ is quite small, but it is more noticeable at short orbital periods, so I only include pulsars with $P_b < 0.2$ days. To increase the contrast, I also use a greyscale colour map, only spanning the values present in the histogram, and apply Gaussian smoothing to reduce noise. If we look back at Figure 4.1, it is no surprise that pulsars at the bottom of the histogram in Figure 4.5 are poorly detected in this $P_b$ range. The interesting feature of Figure 4.5 is the change in the pattern of $\omega$-dependence at higher eccentricities.

Up to $e \approx 0.7$, pulsars with $\omega$ at 90° or 270° are easiest to detect—these are orbits whose major axis is oriented along the line of sight. Above this eccentricity, however, the easiest detections shift symmetrically to values above and below 90° and 270° until, at the highest eccentricities, the situation is completely reversed: pulsars in systems with $\omega$ at 90° or 270° are the hardest to detect, while those with $\omega$ at 0 or 180° are easiest. These latter systems have their *minor* axis oriented along the line of sight, and at such high eccentricties, the minor axis is very small relative to the major axis, so that highly eccentric systems with $\omega$ at 0 or 180° move

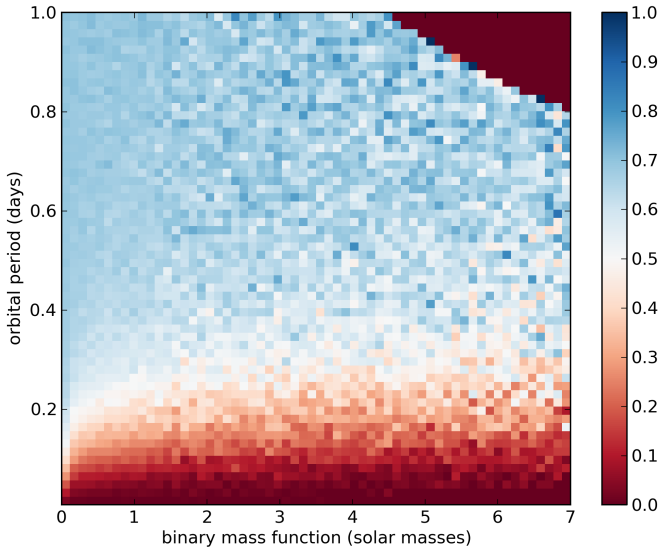


**Figure 4.6:** Detections across binary mass function and orbital period. The upper-right region of the plot is beyond the range of the parameters used to generate these pulsars.

and accelerate almost entirely transverse to the line of sight, resulting in very little radial Doppler-shifting behaviour. At lower eccentricities, the greatest acceleration changes, near periastron, are largely along the line of sight.

Detection levels also depend on the mass of the pulsar's companion and the inclination angle of the orbit. There is a degeneracy in these parameters' effects on the pulsar's Doppler shift, and so they are often combined into the binary mass function $f(m_1,m_2) = (m_2 \sin i)^3 (m_1+m_2)^{-2}$, where $m_1$ and $m_2$ are the pulsar mass and the companion mass, respectively. The pulsar mass $m_1$ is taken to be 1.4 $M_\odot$ here. A larger mass function (due to a more massive companion or a greater inclination angle) leads to greater difficulty detecting pulsars with short orbital periods, as seen in Figure 4.6. For a given orbital period, a pulsar with a more massive companion experiences greater acceleration throughout its orbit, and thus the variations in acceleration with time are more severe. In addition to the acceleration variations, pulsars with the largest mass functions in this simulation would at times be undergoing line-of-sight accelerations that shift their Fourier signals across more bins than the 300 allowed for in the acceleration search.

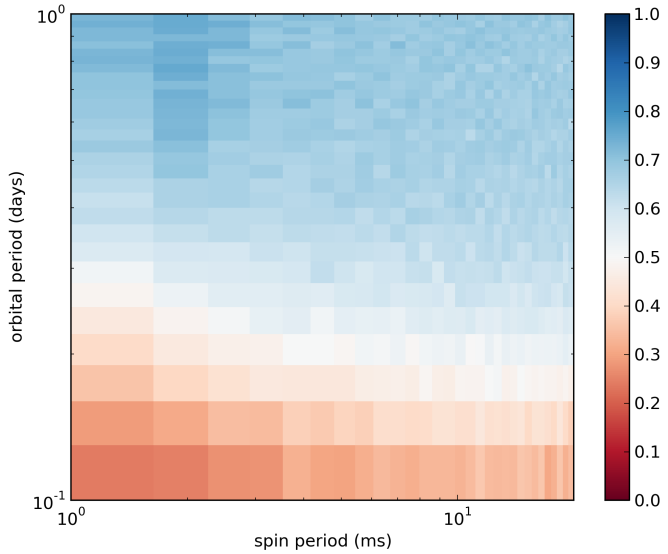


**Figure 4.7:** Detections across spin and orbital period. This is a qualitatively similar plot to those of Johnston & Kulkarni (1991) and Bagchi, Lorimer & Wolfe (2013), who plot their analytical $\gamma$ detectability functions against these two parameters.

## 4.2 Comparisons with recent analytical work

Here I make some qualitative comparisons with the recent work of Bagchi, Lorimer & Wolfe (2013).

Plotting detections against spin and orbital period, as in the log-log histogram of Figure 4.7, we see results similar to those of Bagchi, Lorimer & Wolfe (and Johnston & Kulkarni 1991) when they assume the use of an acceleration search. In their plots of the detectability parameter $\gamma$ across spin and orbital period, detections are reduced most severely for faster-spinning pulsars in shorter orbits, which is the behaviour seen here. The poorly-detected region below orbital period 0.1 days is cut off in Figure 4.7 so that the axes begin at the same values as the spin-orbit $\gamma$ plots of these previous works.

In Figure 4.8 I divide this histogram into four eccentricity ranges. As Bagchi, Lorimer & Wolfe observe, eccentricity has little effect except at high values, where it results in a larger region of phase space becoming easier to detect. This improvement in the orbital period parameter is also seen in Figure 4.1.

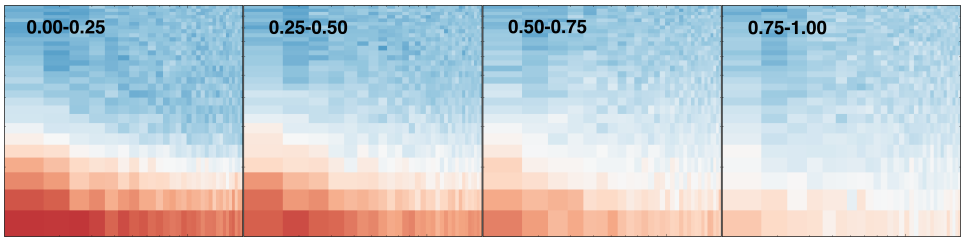


**Figure 4.8:** Improvement in detections on the spin-orbit histogram with increasing eccentricity. The integrated eccentricity ranges are shown at the top of each panel. The axes and colour scheme are the same as in Figure 4.7. As noted by Bagchi, Lorimer & Wolfe, $e$ has little effect except at high values.

Bagchi, Lorimer & Wolfe find that increases in inclination angle or companion mass make a larger portion of phase space difficult to detect, and we see this here in Figure 4.6 with the binary mass function.

While these observations are made primarily by examining the behaviour of their $\gamma_{1m}$ parameter, which assumes no acceleration search, the increase in companion mass can be seen in their paper to worsen the $\gamma_{2m}$ (acceleration search) phase space, and observationally, an increase in inclination angle is like an increase in the companion mass. Bagchi, Lorimer & Wolfe do not test how the $\gamma_{2m}$ phase space changes with variations in eccentricity. They do note that an increase (between 0 and 90°) in the longitude of periastron $\omega$ makes a larger portion of the $\gamma_{1m}$ phase space difficult to detect with $e = 0.5$, which is opposite to the behaviour seen at $e = 0.5$ in Figure 4.5 in this work, where pulsar detection improves as $\omega$ increases from 0 to 90°. This is a qualitative difference between a standard Fourier search and an acceleration search. The former is negatively affected by line-of-sight acceleration, even if that acceleration is varying quite slowly, while the latter is, to a large degree, immune to the effects of a slowly-varying acceleration. It is a rapid *change* in acceleration, i.e., a large line-of-sight *jerk*, that negatively affects acceleration search results.

Using expressions for the line-of-sight acceleration and jerk from Bagchi, Lorimer & Wolfe (their Equations 26 and 27, respectively), I perform a crude test to demonstrate this point. Over a grid of $e$ and $\omega$ values, I find the mean acceleration and jerk across the complete orbit (in mean anomaly, so that it is a time average), with the results shown in Figure 4.9. For each step in $e$, I divide the acceleration or jerk

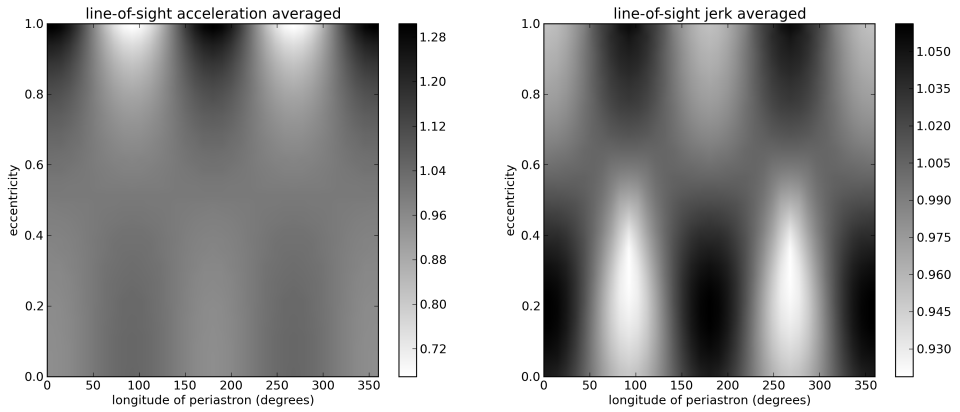


**Figure 4.9:** Line-of-sight acceleration (left) and jerk (right) averaged over the time of a complete orbit on a grid of $e$ and $\omega$ values. Details on scaling are given in the text. Whiter values are lower and considered better for detection, via standard Fourier search in the left plot and via acceleration search in the right plot. Note the similarity of the righthand plot to Figure 4.5.

by its average value across $\omega$ at that step so that the much larger values at high $e$ do not obscure the behaviour at low $e$. Thus, both plots in Figure 4.9 show the relative variations in acceleration or jerk over $\omega$, but not the variations over $e$. The greyscale is chosen so that whiter regions have lower values and are thus deemed easier to detect; the attempted detection is assumed to be via a standard Fourier search in the left plot, and via an acceleration search in the right plot. While there is very little variation across the $e = 0.5$ region of the left plot, the maxima are at $90^\circ$ and $270^\circ$, so this agrees with the increased difficulty of detection at $90^\circ$ seen by Bagchi, Lorimer & Wolfe The plot on the right is similar in appearance to Figure 4.5, supporting my observations.

## 4.3 Conclusions

Binary pulsar detectability has long been a well-known problem with a daunting range of parameters. Past studies have taken an analytical approach to characterizing our coverage of the expansive phase space, and here I have supplemented these studies with an empirical approach, one tailored in particular to recent globular cluster surveys, and covering a broad range of parameters including orbital

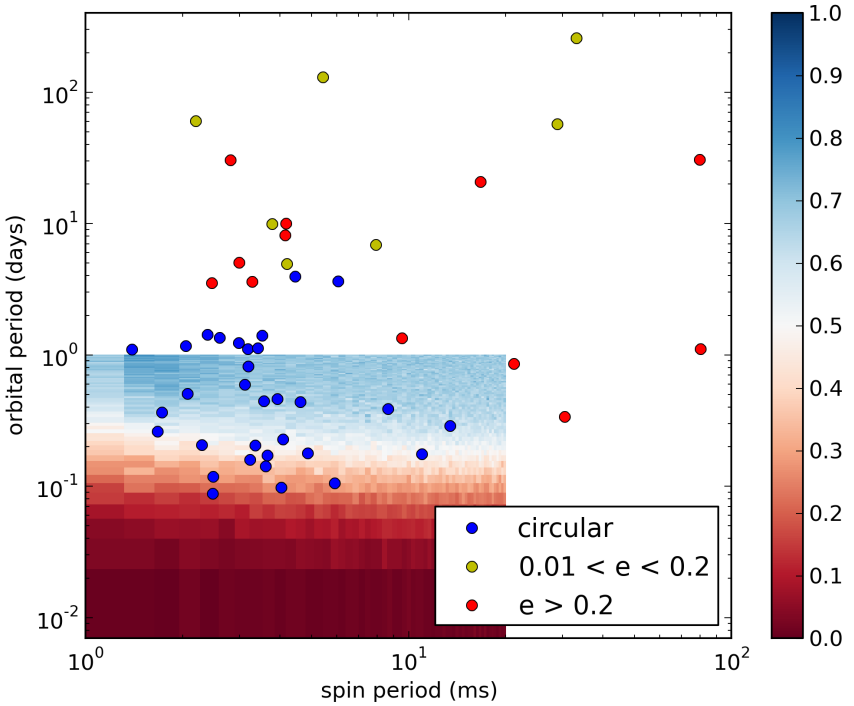


**Figure 4.10:** Globular cluster pulsars discovered with the Arecibo and Green Bank telescopes, plotted over my detection levels for comparison.

eccentricity. The results are consistent with those of Johnston & Kulkarni (1991) and Bagchi, Lorimer & Wolfe (2013), the latter being the first analytical study to account for orbital eccentricity. Like that study, I find that pulsars in short-period binary systems are actually easier to detect when they orbit at very high eccentricities.

Comparing my spin-orbital-period detection plot with binary globular cluster pulsars that have been found in Green Bank and Arecibo telescope surveys (Figure 4.10), it does not appear as though discoveries are pushing down to the orbital-period detection limits, at least for MSPs in eccentric orbits. While it is difficult to draw strong conclusions based on the relatively little information here, this may reflect a lack of very eccentric orbits with orbital periods in this range. For instance, if the eccentricity is the result of a binary exchange scenario, it is expected that the semimajor axis of the resulting binary system will scale as the ratio of the mass of

the neutron star to that of its replaced companion (Heggie, Hut & McMillan, 1996; Ivanova et al., 2008). Since neutron star companions tend to have masses lower than the neutron star, this will often lead to an increase in semimajor axis following the exchange, and thus an increase in orbital period unless the new companion is much more massive than the old one.

Faucher-Giguère & Loeb (2011) expect that a high-density environment with exchanges resulting in black-hole-MSP binaries would lead to highly eccentric orbits with periods as fast as $\sim$5 minutes, slightly lower even than the minimum orbital periods I examine here. However, it is not known how many stellar-mass black holes are retained in globular clusters, if any.

# Chapter 5

# Future Work

The data generated in this simulation take quite a long time to create and search, even with a 36-CPU system. With sufficient time, the present work could be expanded in several ways. It would be interesting to run the acceleration search with some variation in both the number of harmonics summed and the maximum drift in Fourier bins in order to see the particular effects these search parameters have on the behaviour across various orbital parameters. Even testing a few values of each could increase the total computing time by a factor of 10 or more. It would be possible to make improvements in speed elsewhere to partially compensate for this, such as including more pulsars per time series. It is likely that some improvements in computational efficiency could also be made. Testing multiple values of time series length and time resolution might also be of interest, but would again be computationally time-consuming. Such extensions could help guide future surveys in terms of how they acquire and search their data.

In spite of some limitations, this work represents a first attempt to explore this problem's parameter space empirically, It has provided results that are intuitively sensible, yet difficult to arrive at analytically. While at present this serves primarily as a description of a phase space that modern surveys have worked within, with some fairly minor expansions and sufficient computing time, these results could be used to help direct the course of future surveys.

# Bibliography


Archibald A. M. et al., 2009, Science, 324, 1411, arXiv:0905.3397 → pages 5

Baade W., Zwicky F., 1934, Phys. Rev., 45, 138 → pages 1

Bagchi M., Lorimer D. R., Wolfe S., 2013, MNRAS, 432, 1303 → pages 10, 18, 37, 38, 39, 40

Bégin S., 2006, Master's thesis, University of British Columbia → pages 9

Burgay M. et al., 2003, Nature, 426, 531 → pages 8

Champion D. J. et al., 2008, Science, 320, 1309 → pages 5

Clark G. W., 1975, ApJ, 199, L143 → pages 7

Cordes J. M., Lazio T. J. W., 2002, ArXiv:astro-ph/0207156 → pages 13

Damour T., Deruelle N., 1985, Ann. Inst. H. Poincaré (Physique Théorique), 43, 107 → pages 8

Damour T., Deruelle N., 1986, Ann. Inst. H. Poincaré (Physique Théorique), 44, 263 → pages 8, 25

D'Antona F., Bellazzini M., Caloi V., Pecci F. F., Galleti S., Rood R. T., 2005, ApJ, 631, 868 → pages 5

Demorest P. B., Pennucci T., Ransom S. M., Roberts M. S. E., Hessels J. W. T., 2010, Nature, 467, 1081 → pages 2

Dewey R. J., Taylor J. H., Weisberg J. M., Stokes G. H., 1985, ApJ, 294, L25 → pages 26

Faucher-Giguère C.-A., Loeb A., 2011, MNRAS, 415, 3951 → pages 41

Freire P. C. C. et al., 2011, MNRAS, 412, 2763 → pages 5

Freire P. C. C., Ransom S. M., Bégin S., Stairs I. H., Hessels J. W. T., Frey L. H., Camilo F., 2008, ApJ, 675, 670 → pages 10

Gold T., 1969, Nature, 221, 25 → pages 1

Hankins T. H., Rickett B. J., 1975, in Methods in Computational Physics Volume 14 — Radio Astronomy, Academic Press, New York, pp. 55–129 → pages 13

Harris W. E., 1996, AJ, 112, 1487, updated version at http://www.physics.mcmaster.ca/Globular.html → pages 7

Heggie D. C., Hut P., McMillan S. L. W., 1996, ApJ, 467, 359 → pages 41

Hessels J. W. T., Ransom S. M., Stairs I. H., Freire P. C. C., Kaspi V. M., Camilo F., 2006, Science, 311, 1901 → pages 9, 22

Hessels J. W. T., Ransom S. M., Stairs I. H., Kaspi V. M., Freire P. C. C., 2007, ApJ, 670, 363 → pages 9, 26

Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, Nature, 217, 709 → pages 1

Ivanova N., Heinke C. O., Rasio F. A., Belczynski K., Fregeau J. M., 2008, MNRAS, 386, 553 → pages 41

Johnston H. M., Kulkarni S. R., 1991, ApJ, 368, 504 → pages 10, 16, 18, 37, 40

Jouteux S., Ramachandran R., Stappers B. W., Jonker P. G., van der Klis M., 2002, A&A, 384, 532 → pages 10, 17, 18

Kaplan D. L. et al., 2005, PASP, 117, 643 → pages 9

Kiziltan B., Thorsett S. E., 2010, ApJ, 715, 335 → pages 4

Kramer M. et al., 2006, Science, 314, 97 → pages 2, 8

Lorimer D. R., Kramer M., 2005, Handbook of Pulsar Astronomy. Cambridge University Press → pages 1, 6, 26

Lynch R. S. et al., 2013, in IAU Symposium, Vol. 291, IAU Symposium, pp. 41–46 → pages 9

Lynch R. S., Freire P. C. C., Ransom S. M., Jacoby B. A., 2012, ApJ, 745, 109 → pages 10

Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, AJ, 129, 1993 → pages 6

Milone A. P., Marino A. F., Piotto G., Bedin L. R., Anderson J., Aparicio A., Cassisi S., Rich R. M., 2012, ApJ, 745, 27 → pages 5

Özel F., Psaltis D., Narayan R., Santos Villarreal A., 2012, ApJ, 757, 55 → pages 2

Pacini F., 1968, Nature, 219, 145 → pages 1

Phinney E. S., Kulkarni S. R., 1994, Ann. Rev. Astr. Ap., 32, 591 → pages 4

Pletsch H. J. et al., 2012, Science, 338, 1314 → pages 22

Podsiadlowski P., Dewi J. D. M., Lesaffre P., Miller J. C., Newton W. G., Stone J. R., 2005, MNRAS, 361, 1243 → pages 2

Portegies Zwart S., van den Heuvel E. P. J., van Leeuwen J., Nelemans G., 2011, ApJ, 734, 55 → pages 5

Ransom S. M., 2001, PhD thesis, Harvard University → pages 15

Ransom S. M., Eikenberry S. S., Middleditch J., 2002, AJ, 124, 1788 → pages 15

Ransom S. M., Hessels J. W. T., Stairs I. H., Freire P. C. C., Camilo F., Kaspi V. M., Kaplan D. L., 2005, Science, 307, 892 → pages 9, 26

Richer H. B., Heyl J., Anderson J., Kalirai J. S., Shara M. M., Dotter A., Fahlman G. G., Rich R. M., 2013, ApJ, 771, L15 → pages 7

Schwab J., Podsiadlowski P., Rappaport S., 2010, ApJ, 719, 722 → pages 2

Sigurdsson S., 2003, in Astronomical Society of the Pacific Conference Series, Vol. 302, Radio Pulsars, Bailes M., Nice D. J., Thorsett S. E., eds., p. 391 → pages 8

Taylor J. H., Weisberg J. M., 1989, ApJ, 345, 434 → pages 25