Pulsar Searches: From Radio to Gamma-Rays

Thesis by

Adam M. Chandler

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

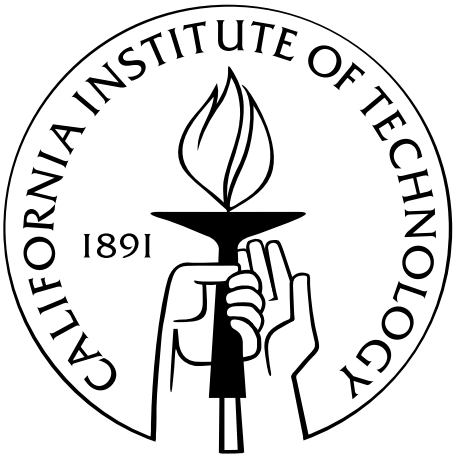

California Institute of Technology

Pasadena, California

2003

(Defended December 5, 2002)

Acknowledgements

Physics isn't done in a vacuum.

Well, *some* physics may be done in a vacuum, but this particular thesis would be a pale shadow of what you see here, had I not received help from a great many individuals and organizations. Most of the people mentioned in these acknowledgements merit paragraphs all to themselves. In the interest of brevity, however, many have been undeservedly relegated to impersonal lists. This laconic approach belies my true level of gratitude.

First, I must thank my advisor, Tom Prince. Whenever I became mired in the details of a problem, he always managed to keep me moving in the right direction. Tom kept my focus on science, never letting me lose sight of the big picture.

If Tom was my strategic advisor, Stuart Anderson was my tactical advisor. Stuart always made time to help with any questions I could come up with, on virtually any topic. On the rare occasions when Stuart was unable to address an issue outright, he always had a good suggestion on how to proceed.

Thanks are also due to my other Caltech collaborators: Bryan Jacoby, David Kaplan, Danny Koh, Shri Kulkarni, Dick Lamb, and Bob Rutledge; as well as the faculty members who served on my candidacy and thesis defense committees: Peter Goldreich, Ken Libbrecht, Nick Scoville, and Kip Thorne.

My time at Caltech would have been far less pleasant without fellow graduate students Bob Beach, Matt Evans, Rick Jenet, Steffen Jensen, and Matthew Sumner. And all those awful classes would have been even more awful without Elliot Lipeles, Anna Shapiro, and Werner Sun.

Whether they realize it or not, the following people contributed tangibly to this thesis or to my graduate school experience: Don Backer, Walter Brisken, Minerva Calderon, Fernando Camilo, Alok Chatterjee, Heather Cox, Donna Driscoll, Mihai Gavrila, Duncan Lorimer, Daryl Macomb, John Mattox, Maura McLaughlin,

Maura McNicholas, Bill Metcalf, Brian Moore, David Nice, Scott Ransom, Paul Ray, Connie Ryan, Ingrid Stairs, Linqing Wen, and Alex Wolszczan.

I thank the Laser Interferometer Gravitational Wave Observatory project for their hospitality (and my lovely mountain view). For their invaluable assistance, I am indebted to the friendly people at the Caltech Center for Advanced Computing Research, the Arecibo Observatory, and the Green Bank Telescope. I am eternally grateful to those individuals responsible for the Astrophysics Data System, and the astro-ph preprint server. I can only imagine how painful it was to do astrophysics research before these services existed. This thesis also benefited from extensive use of Merriam-Webster's online dictionary. The graphics in this thesis were all produced with IDL or PGPLOT, and the thesis itself was laid out in LATEX. Some of the plots in the introduction used archived data from the European Pulsar Network.

I thank the institute proofreader for meticulously checking this thesis for typography, consistency, and style. Of course, any remaining errors are entirely the proofreader's fault.

On a more personal note, I thank my parents, Clare Chandler and Tom Chandler, for bringing me into this world and letting me find my own way in it.

I am extremely fortunate to have been able to share this experience with Mona Carroll. This is not the proper forum to gush about (nor can words express) how much she means to me, so I will simply thank her for her companionship, support, and patience, which I too often take for granted.

Most of all, I am grateful to Janine Shertzer. This thesis exists only because she took a chance hiring an unmotivated slacker in whom she saw some potential. If Janine had not shown me that I could have an "exciting career doing physics research," I might still be floundering now. I am who I am today because of Janine Shertzer. I cannot thank her enough.

Abstract

We report the results of four different pulsar searches, covering radio, X-ray, and gamma-ray wavelengths. These searches targeted pulsars in virtually all of their guises: young and old, long-period and short-period, accretion-powered and rotation-powered. Ten new pulsars were discovered.

There are very few known gamma-ray pulsars, all of which were found by folding gamma-ray data with a pulse period known from other wavelengths. Some emission models indicate that there may be a large number of gamma-ray pulsars that are undetectable at lower energies. We searched several of the brightest unidentified gamma-ray sources for pulsations. This was the first attempt to identify gamma-ray pulsars by a direct search of gamma-ray data. No new identifications resulted and we report upper limits.

Even more rare than gamma-ray pulsars are accreting millisecond pulsars. We searched for coherent pulsations from Aql X-1, a low-mass X-ray binary suspected of harboring such an object. No pulsations were detected, and we argue that the quiescent emission of this system has a thermal origin (i.e., it is not due to low-level accretion).

The two radio searches included here were both designed to detect millisecond pulsars. First, we report the results of a large area survey from Arecibo. Five new slow pulsars were discovered, including an apparent orthogonal rotator and an extremely unusual bursting radio pulsar. No short-period pulsars were discovered and we place some of the first useful observational constraints on the limiting spin period of a neutron star.

We also performed pointed searches of several globular clusters using the new Green Bank Telescope. Three new binary millisecond pulsars were found in M62. These were the first new objects found with the GBT, and they bring the total pulsar population in M62 to six. We also discovered two isolated pulsars, one each in NGC 6544 and NGC 6624.

Many of the methods we developed will be relevant to future searches. Perhaps the most significant contribution is a dynamic power spectrum-based technique that finally allows sensitive searches for binary pulsars whose orbital periods are of the same order as the observation time.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Pulsars are celestial sources of very regularly pulsed radiation (Figure 1.1). First discovered at radio wavelengths, pulsars have since been seen in virtually every electromagnetic band — radio, infrared, optical, ultraviolet, X-rays, and gamma-rays (Figure 1.2). The first pulsar was discovered over 30 years ago (Hewish et al., 1968), and new ones are still being found today, driven by improvements in hardware (telescopes and computers) and software (search algorithms). These continuing advances have allowed search sensitivities to improve, and have opened up new windows onto exotic pulsars, including high-energy gamma-ray pulsars, very-short-period pulsars, and pulsars in tight binary systems. These are the subject of this thesis.

The purpose of this introductory chapter is threefold. The first goal is to provide a brief introduction to pulsars for a potentially mixed audience, including a discussion of pulsar formation and evolution, some of the science that can be learned from pulsars, and the techniques used to find them. The second goal is to motivate the various searches described in the subsequent chapters of this thesis. As these four pulsar search projects were all quite different from each other, the final aim in this introduction is to tie these somewhat disparate chapters together.

We will attempt to address each of these three objectives concurrently by considering the standard model for pulsar evolution. The four projects described in this thesis span the electromagnetic spectrum, including searches at radio, X-ray, and gamma-ray wavelengths. As we shall see in section 1.1 below, these different wavelength searches targeted pulsars at virtually every stage of their life cycle.

Since the bulk of the work that went into this thesis involved search algorithm

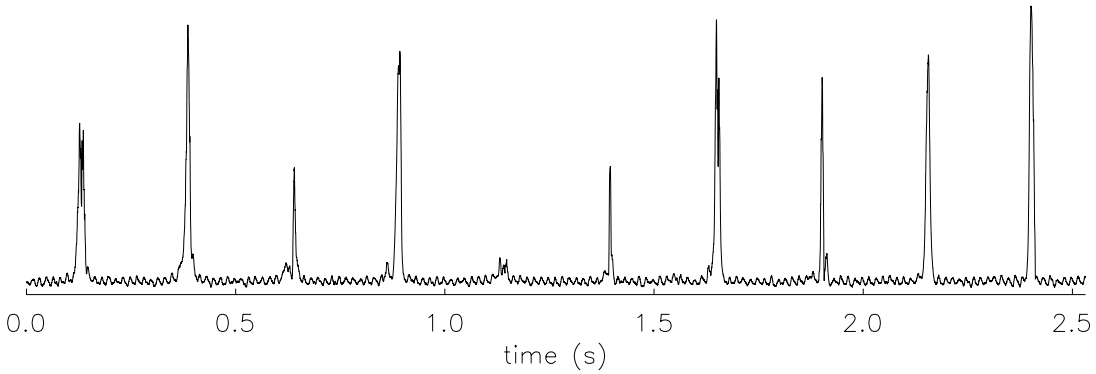

Figure 1.1: 430 MHz radio pulses from pulsar B0950+08.

development, computer coding, and analysis of large data sets, we also include, in section 1.2, a brief overview of pulsar search techniques. We highlight some of the similarities and differences between searches at different wavelengths, and discuss an important complication — pulsars whose spin periods change over the course of an observation. The last section of this chapter (sec. 1.3) briefly outlines the organization of the remainder of the thesis.

1.1 Introduction to Pulsars and Pulsar Evolution

1.1.1 Pulsar Overview

The first pulsar was discovered in 1967 by Anthony Hewish and Jocelyn Bell. They were studying interplanetary scintillation of radio galaxies and quasars when they serendipitously noticed an extraterrestrial source emitting regular radio pulses with a period of about 1.3 seconds. Within several months of the first discovery, Hewish et al. (1968) had already determined that the source was Galactic (extra-solar), and due to its extremely short period, it was likely to be a compact object — a white dwarf or neutron star. The pulsations were thought to be due to either radial oscillations, a binary orbit, or rotation of the compact object. At the time, neutron stars were purely theoretical objects. Their existence had been proposed over 30 years earlier by

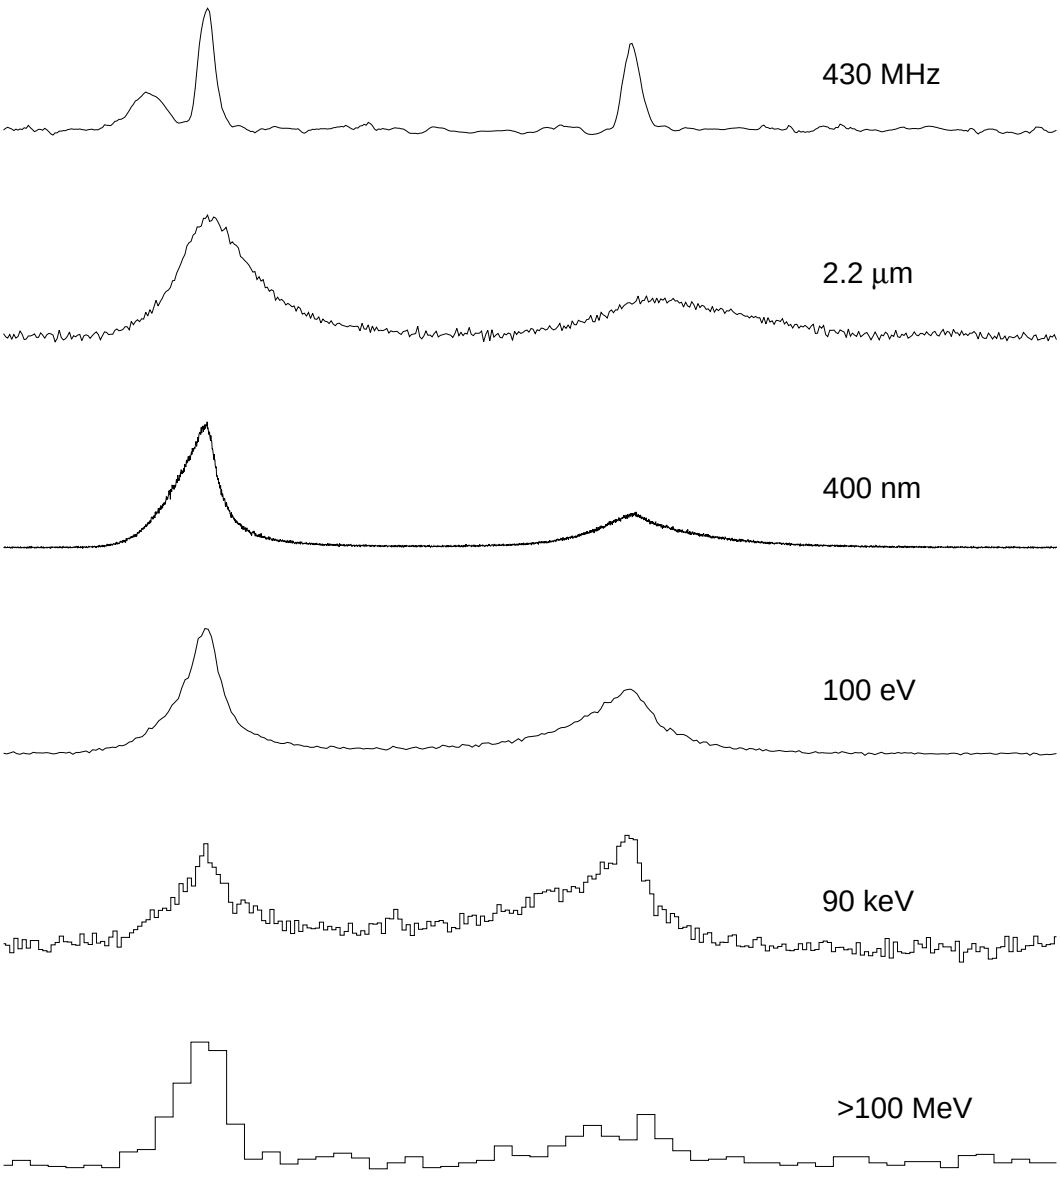

Figure 1.2: Pulse profiles for the Crab pulsar (PSR B0531+21) at radio, infrared, optical, ultraviolet, X-ray, and gamma-ray wavelengths (Moffett & Hankins, 1996; Lundgren, 1994; Percival et al., 1993; Harnden & Seward, 1984; Ulmer et al., 1994).

Baade & Zwicky (1934), shortly after the discovery of the neutron (Chadwick, 1932), but neutron stars had received very little attention from the astronomical community. As even shorter-period pulsars were discovered in the Vela and Crab nebulae (Large et al., 1968; Staelin & Reifenstein, 1968), the white dwarf, oscillation, and binary models had to be abandoned. Interestingly, just before the first pulsar discovery, Pacini (1967) had postulated that a spinning, magnetized neutron star could be the power source in the crab nebula, an idea independently advanced by Gold (1968) just after Hewish's discovery. All subsequent evidence has supported the model that pulsars are indeed rotating, highly magnetized neutron stars.

We pause here to consider how utterly outlandish these objects are. Neutron stars have a radius of approximately $10 - 15\,\mathrm{km}$, roughly equivalent in cross section to a small city, or in volume to the island of Oahu. Within that small volume, neutron stars contain more mass than the sun, for an average density of $\sim 10^{14}\,\mathrm{g\,cm^{-3}}$. This is about 10 trillion times the density of lead, or roughly the density of an atomic nucleus. Typical pulsars have magnetic fields of $10^{12}\,\mathrm{G}$, over a trillion times stronger than the earth's field, and 100,000 times stronger than the largest fields yet produced in terrestrial laboratories (incidentally, these laboratory fields last only a few microseconds before completely destroying the magnet producing them). These extremely massive, ultra-dense, ultra-magnetized objects have been observed spinning as rapidly as 642 revolutions per second, and moving through the Galaxy at speeds exceeding $1000\,\mathrm{km\,s^{-1}}$. Neutron stars have become an invaluable resource for studying physics under extreme conditions, not attainable on Earth. Not surprisingly, their study has had a profound influence on a number of diverse scientific fields, from nuclear physics to gravitation.

A pulsar can be thought of as a cosmic lighthouse — a rotating source of beamed radiation. However, the exact mechanism by which pulsars generate these beams is still not fully understood. A complete treatment of emission models is far beyond the scope of this introduction, but we will briefly describe the basic picture. The rotating magnetic field of a pulsar no doubt gives rise to an electric field, which in turn induces a charge distribution on and around the pulsar. For a pulsar (or any

object) rotating with angular frequency Ω, there is an imaginary co-rotating cylinder, of radius $r_{\rm lc} = c/\Omega$, whose surface moves at the speed of light. Only magnetic field lines entirely within this "light cylinder" can form a closed loop joining the two magnetic poles. A stable, co-rotating, charge-separated plasma exists within the region of closed field lines. Since charged particles can escape along the open field lines, vacuum regions develop in which the electric field can accelerate charges to extremely high energies. These charges can generate high-energy photons via synchrotron or curvature radiation or inverse compton scattering of thermal photons from the neutron star surface. Spontaneous pair production by these high-energy photons produces more charged particles, resulting in a cascade effect. An intense beam of radiation and relativistic particles will therefore be produced. If this beam is misaligned with the rotation axis and intercepts the earth, this object should be an observable pulsar.

Although many properties of pulsars can be explained with this simple model, it is far from complete. For example, it cannot account for the ludicrously high brightness temperatures of radio pulsar beams (which require coherent emission). Also, it remains unclear whether pulsar gamma-ray emission is produced close to the neutron star surface (at the magnetic "polar caps"), or farther out near the light cylinder (in the "outer gap" between the last closed field lines and the surface on which the induced charge density is zero). For further information on pulsar emission, we refer the reader to the wealth of literature on the subject. Most of this traces its roots back to the seminal work by Goldreich & Julian (1969). More recent reviews can be found in, for example, Jenet (2001), Thompson et al. (1997), and references therein.

Ultimately, the energy source for a pulsar's emission is its huge reservoir of rotational kinetic energy. Pulsar spin rates must therefore decrease over time, as this energy powers low frequency magnetic dipole radiation and the emission beam of the pulsar. The spin period P and its first time derivative $\dot{P}$ can tell us a lot about a pulsar, including its approximate magnetic field strength, age, and evolutionary state. In our discussion of pulsar evolution, we will find it useful to refer to Figure 1.3, which shows P vs. $\dot{P}$ for the majority of the known pulsar population (i.e., those

with measured $\dot{P} > 0$).

The rotational energy of a pulsar with moment of inertia I and angular rotation frequency Ω is $E = \frac{1}{2}I\Omega^2$. If we model a pulsar as a rotating magnetic dipole surrounded by vacuum, we expect it to lose rotational energy at a rate proportional to the fourth power of the rotation frequency

$$\dot{E} = I\Omega\dot{\Omega} = -\frac{B^2R^6\Omega^4\sin^2\alpha}{6c^3}, \tag{1.1}$$

where B is the field strength at the magnetic pole on the neutron star surface, R is the neutron star radius, c is the speed of light, and α is the angle between the rotation axis and the magnetic moment (Jackson, 1975; Shapiro & Teukolsky, 1983). More generally, we expect

$$\dot{\Omega} \propto \Omega^n, \tag{1.2}$$

where n is the braking index (for a pure dipole in vacuum, $n = 3$). If we differentiate Equation 1.2, we see that $n = \Omega\ddot{\Omega}/\dot{\Omega}^2$, which allows us to determine the braking index if the second derivative of the spin period can be measured. This is notoriously difficult in practice due to timing noise and glitches, and braking indices have been measured for only a handful of pulsars (Lyne et al., 1988; Kaspi et al., 1994; Boyd et al., 1995; Lyne et al., 1996). The median measured value, $n = 2.4$, indicates that the vacuum dipole braking model is not the whole story. Nevertheless, it is common practice to assume $n = 3$ when deriving approximate pulsar parameters based on P and $\dot{P}$.

If a pulsar was born at time $t = 0$ spinning with angular frequency $\Omega_0 \gg \Omega$, integration of Equation 1.2 from $t = 0$ to $t = \tau$ yields

$$\tau = \frac{\Omega}{(1-n)\dot{\Omega}}\left[1 - \left(\frac{\Omega}{\Omega_0}\right)^{n-1}\right] \approx \frac{P}{(n-1)\dot{P}}. \tag{1.3}$$

We set $n = 3$ to define the "characteristic age" of a pulsar

$$\tau_c = \frac{P}{2\dot{P}}. \tag{1.4}$$

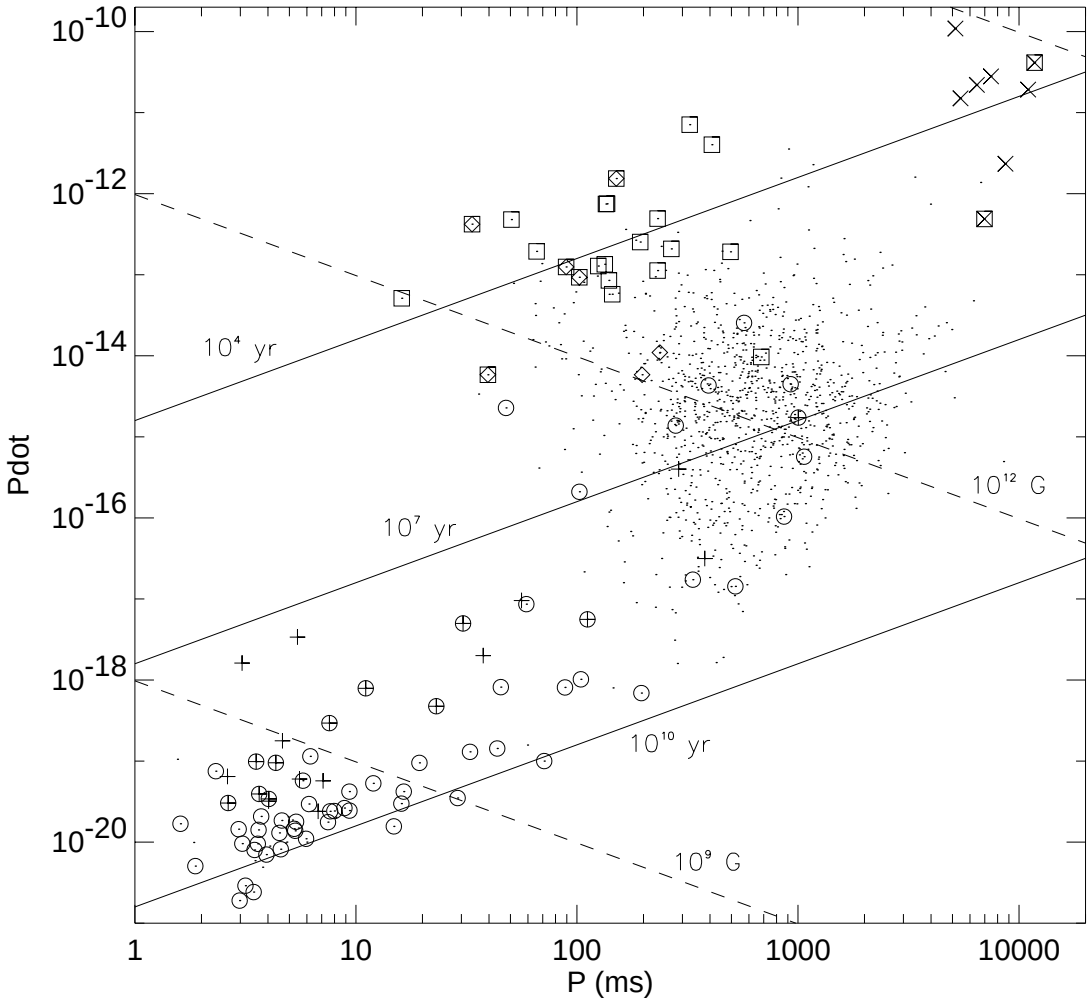

Figure 1.3: $P - \dot{P}$ diagram. Each plotted point represents the spin period P and the first time derivative of the spin period $\dot{P}$ for an individual pulsar. The 1251 known pulsars with measured $\dot{P} > 0$ are shown. The solid lines in the plot are lines of constant characteristic age: from top to bottom, $\tau_c = 10^4, 10^7, 10^{10}$ yr. The dashed lines are lines of constant estimated surface magnetic field: from upper right to lower left, $B = 10^{15}, 10^{12}, 10^9$ G. A surrounding square indicates a pulsar with a probable supernova remnant association. Diamonds indicate pulsars for which pulsed gamma-rays have been definitively detected. Soft gamma-ray repeaters and likely anomalous X-ray pulsars are denoted by Xs. Circles indicate pulsars in binary systems, and crosses indicate pulsars in globular clusters.

Lines of constant τ_c are plotted in Figure 1.3 for $\tau_c = 10^4,\ 10^7,$ and $10^{10}\,\mathrm{yr}$. Also included in the $P - \dot{P}$ diagram are lines of constant surface dipole magnetic field strength, estimated from Equation 1.1 with $I = 10^{45}\,\mathrm{g\,cm^2}$, $R = 10\,\mathrm{km}$, and $\alpha = 30°$:

$$B \approx 3.2{\times}10^{19}\,\sqrt{P\dot{P}}\,\mathrm{G}. \tag{1.5}$$

The dashed lines in the figure indicate $B = 10^9,\ 10^{12},$ and $10^{15}\,\mathrm{G}$.

The following subsections give a very brief description of pulsar formation and evolution. Much of the material presented here has been culled from the following sources, to which readers interested in a more complete account are referred: Manchester & Taylor (1977); Shapiro & Teukolsky (1983); Lyne & Graham-Smith (1998); Blandford et al. (1993); Phinney & Kulkarni (1994).

1.1.2 Pulsar Birth and Death

Normal stars are supported against gravitational collapse by thermal pressure from nuclear fusion. When a star's nuclear fuel runs out, and there is nothing to hold it up against its own gravity, it must collapse. For a relatively light star (up to a few times the mass of the sun, $M_\odot$), this collapse will cease when the core's density reaches $\sim 10^6\,\mathrm{g\,cm^{-3}}$. At this point, the core is supported against further collapse by a purely quantum mechanical pressure. As an electron is forced into a smaller and smaller volume, the uncertainty principle dictates that its momentum (and therefore kinetic energy) must increase. Electrons in close proximity to each other are prohibited from occupying the same quantum state by the Pauli exclusion principle, and as the lowest available energy levels fill up, electrons are forced to even higher energy states. These fast moving electrons generate a temperature-independent "electron degeneracy pressure," balancing the compressing effects of gravity. As the central region of the star contracts to this state, the liberated gravitational energy heats the outer layers, which expand and dissociate from the dense core. What remains is an incandescent cinder — a "white dwarf" — about the size of the earth, having roughly the mass of the sun.

Even electron degeneracy pressure is not strong enough to oppose the gravitational collapse of more massive stars ($M \gtrsim 8 M_\odot$). As the core of a massive star collapses, protons and electrons merge to form neutrons via inverse beta decay. If the star is not overwhelmingly massive (total mass $M \lesssim 20 M_\odot$), neutron degeneracy pressure may abruptly halt the collapse of the core. Imploding layers above the suddenly stiff core recoil outward in a supernova explosion (Woosley & Weaver, 1986). The neutron star that remains will be spinning rapidly due to the conservation of angular momentum as the core contracted. And magnetic flux conservation may result in a large magnetic field.

Evidence that short-period pulsars are born in supernovae can be seen in Figure 1.3. In the upper middle part of the diagram, we indeed find a large number of young pulsars ($\tau_c \sim 10^3 - 10^5\,\mathrm{yr}$) associated with supernova remnants (SNRs). As these pulsars age, and their spin rates decrease, they move rightward and downward in the diagram. After $\sim 10^6\,\mathrm{yr}$, we expect that these pulsars will join the main "island" in the plot with $P \sim 1\,\mathrm{s}$ and $\dot{P} \sim 10^{-15}$. The majority of the known pulsars ($\sim 85\%$) reside in this region of $P - \dot{P}$ phase space. SNRs fade from detectability after $\lesssim 10^5\,\mathrm{yr}$, and so we do not expect to see SNR associations with older pulsars.

The energy output of a young pulsar (recall, $\dot{E} \propto \dot{P}/P^3$) may be four orders of magnitude larger than that of an older pulsar. Not surprisingly, many of the closest energetic young pulsars are detected at high energies. The diamonds in Figure 1.3 show the seven pulsars for which high-confidence, high-energy gamma-ray pulsations have been detected. Note that five of these are also associated with SNRs.

The pulsars we have been discussing have typical magnetic field strengths of order $\sim 10^{12} - 10^{13}\,\mathrm{G}$. A distinct population of apparently young pulsars with much larger magnetic fields is visible in the upper right of Figure 1.3. These are the so-called anomalous X-ray pulsars (AXPs) and soft gamma-ray repeaters (SGRs), denoted by Xs in the figure. SGRs are so named because they emit occasional bursts of low-energy gamma-rays, though in quiescence SGRs can be seen in X-rays (Kouveliotou et al., 1998, 1999). The AXPs are considered anomalous because their inferred rotational energy loss is orders of magnitude too small to explain their observed luminosities

(Mereghetti, 2001). Accretion from an extremely low mass companion, from a supernova fall back disk, or from the interstellar medium is a possible explanation for the X-ray luminosity. If this is the case, the $\dot{P}$ measurements would be suspect due to torques on the pulsar from the infalling accreted matter, and their magnetic fields may not be so unusual. But the intriguing possibility exists that these objects are "magnetars," pulsars with extremely large magnetic fields, whose bright, high-energy emission is powered by magnetic field decay (Thompson & Duncan, 1995, 1996). Photon splitting dominates pair production in the magnetospheres of these objects, quenching any radio emission, but high-energy emission is not hindered (Baring & Harding, 1997). The birth rate of these high-field objects may be within an order of magnitude of the birth rate of normal, lower-field neutron stars (Kouveliotou et al., 1994; Regimbau & de Freitas Pacheco, 2001), indicating that these objects represent an important alternative track for neutron star birth and evolution.

These young, high-energy pulsars (of both types) are the subject of Chapter 2 of this thesis. All the known pulsars were discovered at radio or X-ray wavelengths. The few known gamma-ray pulsars were all identified by folding gamma-ray data with a pulse period known from observations in radio or X-rays. In Chapter 2 we describe the first attempt to search for new pulsars in gamma-ray data directly. Even a single discovery would have meant a significant increase in the known gamma-ray pulsar population. Unfortunately no new pulsars were found, but the techniques we developed will be applicable to searches with the upcoming *Gamma-Ray Large Area Space Telescope*.

Young pulsars are also relevant to the radio pulsar survey of Chapter 4. Although this search was not designed specifically to look for young pulsars, we did discover PSR J0627+07, an interesting radio pulsar with a possible SNR association. This survey also discovered four other "garden variety," or middle-aged pulsars.

As a pulsar of advanced age continues to slow down, it eventually sputters and dies. The radio pulsation mechanism shuts off, probably due to weakening of the induced electric field in the magnetosphere. Depending on its magnetic field, the active lifetime of a typical pulsar can be anywhere from $\sim 10^7$ yr to perhaps a few

times 10^8 yr.

1.1.3 Life After Death

We now turn our attention to the last distinct population in Figure 1.3 — the fast spinning, low magnetic field pulsars in the lower left of the diagram. These are collectively called "recycled pulsars" for reasons which will soon become apparent. The fastest pulsars, with $P \lesssim 30$ ms, are commonly referred to as millisecond pulsars (MSPs).

All pulsars by their very nature represent life after death, since they are born in the throes of a supernova explosion. As we saw in section 1.1.2, after radiating away most of their rotational energy, pulsars turn off and become undetectable. But if an extinct pulsar has (or acquires) a binary companion, it may be able to cheat death yet again.

If we consider the potential energy in a reference frame rotating with the binary (Fig. 1.4), the stars will sit at the bottom of two potential wells. Along the line connecting the two, there will be a saddle point joining the two separate "Roche lobes" in which the stars sit. If matter from one star crosses this equilibrium point, it may fall toward the other star. A steady flow of such material pouring in with non-zero angular momentum will form a disk around the recipient star. Matter that reaches the stellar surface (through viscous energy losses and momentum transfer in the disk) will surrender its remaining angular momentum to the accreting star. Depending on a number of factors (e.g., the mass accretion rate, the magnetic field of the star, the total time spent accreting), this process can potentially spin the recipient star up to very short periods. (For more detailed treatments of accretion physics and binary evolution see, e.g., King 1993; Bhattacharya & van den Heuvel 1991; Rappaport & Joss 1977).

The outflow from the mass donor could be from a stellar wind, or from true Roche lobe overflow. The latter is usually the result of normal stellar evolution. As an evolving star starts to run out of nuclear fuel, its outer envelope expands. In

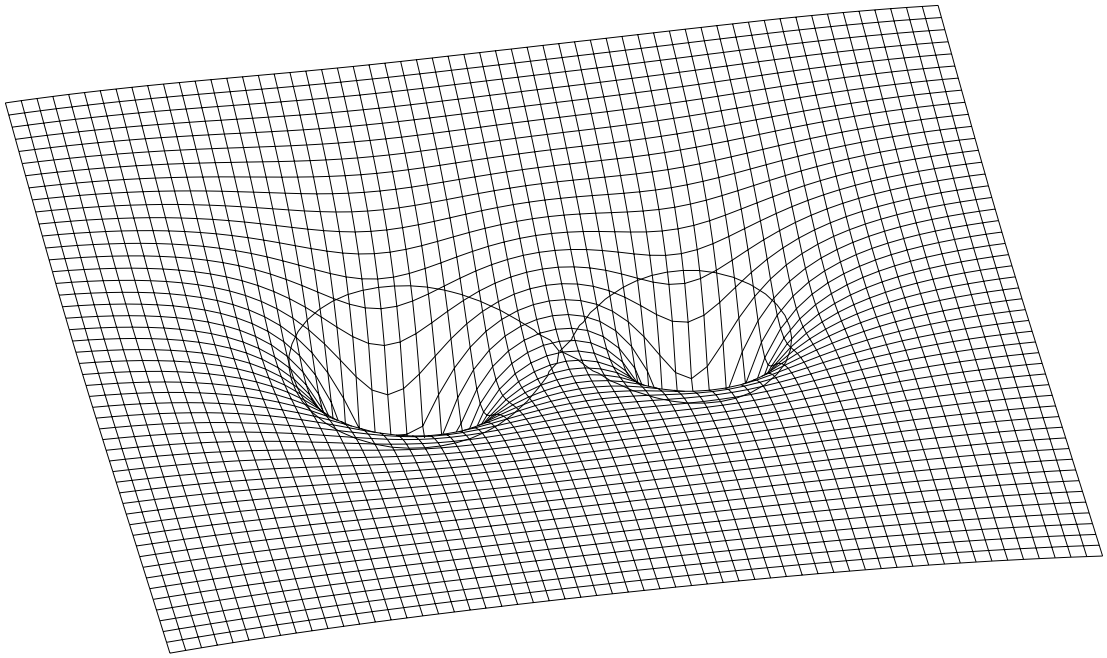

Figure 1.4: Roche potential in a binary system. The contour overlaid on the surface is the equipotential passing through the saddle point, defining the Roche lobes of the two stars. Above this potential, material can flow from one star to the other.

a compact binary, the expanded giant may fill its Roche lobe and begin funneling matter to its companion. Roche lobe overflow from a normal, unevolved star is also possible if the binary separation is reduced by, e.g., gravitational radiation or close encounters with nearby stars.

If the accretor is a compact object (a white dwarf, neutron star, or black hole), there is an enormous amount of gravitational potential energy available from the infalling matter. Much of this energy is converted to heat in the accretion disk and at the stellar surface (if it has a surface, i.e., not for black holes). Due to the high temperatures involved, accreting compact objects are often visible as bright X-ray sources.

When matter accretes onto a strongly magnetized neutron star, hot, ionized gas may be funneled along the field lines to the magnetic poles. These rotating hot spots may be visible from Earth as X-ray pulsations. The inner edge of the accretion disk

is held off from the pulsar's surface at a distance where the magnetic energy density is comparable to the ram pressure of the accreted material. The keplerian orbital period at this radius determines the minimum spin period to which the pulsar can be spun back up. The majority of observed accreting pulsars have large magnetic fields ($\sim 10^{12}\,\mathrm{G}$) and therefore large limiting spin periods $\sim 1\,\mathrm{s}$ (Bildsten et al., 1997). Clearly, we cannot directly relate these accreting X-ray pulsars to the MSPs in Figure 1.3. (Of course, accreting pulsars themselves are not represented in Figure 1.3 since they do not have steady $\dot{P} > 0$.)

The observed accreting pulsars also tend to have rather high-mass stellar companions, whose short lifetimes indicate that these systems have not been accreting very long. A neutron star with a lighter, longer-lived companion (i.e., in a low mass X-ray binary or "LMXB") could accrete for much longer. Long-term accretion may bury the magnetic field of a neutron star (e.g., Bhattacharya & Srinivasan 1995), reducing the field strength to MSP levels ($10^8 - 10^9\,\mathrm{G}$). The weaker magnetic field also allows the accretion disk to get closer to the neutron star surface, allowing spin up to very short periods.

Immediately after the discovery of the first MSP (Backer et al., 1982), LMXBs were recognized as their likely progenitors (Alpar et al., 1982). Over 15 years passed with no definitive proof of this claim, although there were indications of incoherent millisecond phenomena in LMXBs (i.e., kiloHertz QPOs and burst oscillations; van der Klis 2000). The discovery of 2.5 ms coherent pulsations in the transiently accreting LMXB SAX J1808.4-3658 (Wijnands & van der Klis, 1998; Chakrabarty & Morgan, 1998) finally provided the missing link between accreting pulsars and MSPs. Chapter 3 of this thesis describes a search for coherent millisecond pulsations from the LMXB Aql X-1. No pulsations were detected, a result that has implications for the origin of the quiescent luminosity in this system and others like it. And again, the search technique we developed should have wider application. Since the publication of Chapter 3 (Chandler & Rutledge, 2000), two additional accreting MSPs have been found (Markwardt et al., 2002; Galloway et al., 2002).

Eventually the accretion phase ends. This can happen, e.g., when the entire enve-

lope of the donor star has been exhausted, when the donor's Roche lobe expands due to widening of the binary as angular momentum is transferred to the pulsar, or when a radio pulsar phase begins and radiation pressure clears out the accretion disk. If sufficient spin-up has occurred, the result will be a binary, recycled, rotation-powered pulsar. Solitary recycled pulsars are observed, but they most likely formed in binary systems and have since lost or destroyed their companions (see Chapter 5). MSPs have relatively weak magnetic fields (though still gargantuan by terrestrial standards), and therefore spin down very slowly. They are also far less susceptible to the glitches and timing noise commonly seen in younger pulsars, making them exquisite clocks. As such, MSPs have been used to measure neutron star and companion masses, to test gravitational theories, and to search for a low frequency cosmic gravitational wave background, to name just a few applications (Thorsett & Chakrabarty, 1999; Weisberg & Taylor, 1984; Stinebring et al., 1990).

Figure 1.3 shows that the recycled pulsar population is comprised of two groups: those in the Galactic plane, and those in the globular cluster system (though apart from their locations, these pulsars are not intrinsically different). These two types of MSPs were the chief quarry of the searches described in Chapters 4 and 5, respectively.

Due to their relative proximity to us and their large scale height, the Galactic plane MSPs appear rather isotropically distributed on the sky (Fig. 1.5). Chapter 4 describes a sensitive, large-area, high Galactic latitude survey designed to detect MSPs. This was also the most sensitive large-area survey to date for sub-millisecond pulsars, and places some of the first observational constraints on the lower limit of pulsar spin periods.

Globular clusters are old, roughly spherical formations of $\sim 10^4 - 10^6$ stars, located primarily around the central bulge of the Galaxy. With such a large number of stars in such a small space, encounters between stars are inevitable. Even originally solitary neutron stars can acquire companions by tidally capturing a passing star or exchanging into an existing binary. Not surprisingly, both LMXBs and MSPs occur in globular clusters at a much higher rate (per unit mass) than in the Galactic plane. Currently, roughly half of the known MSPs are in globular clusters. In Chapter 5,

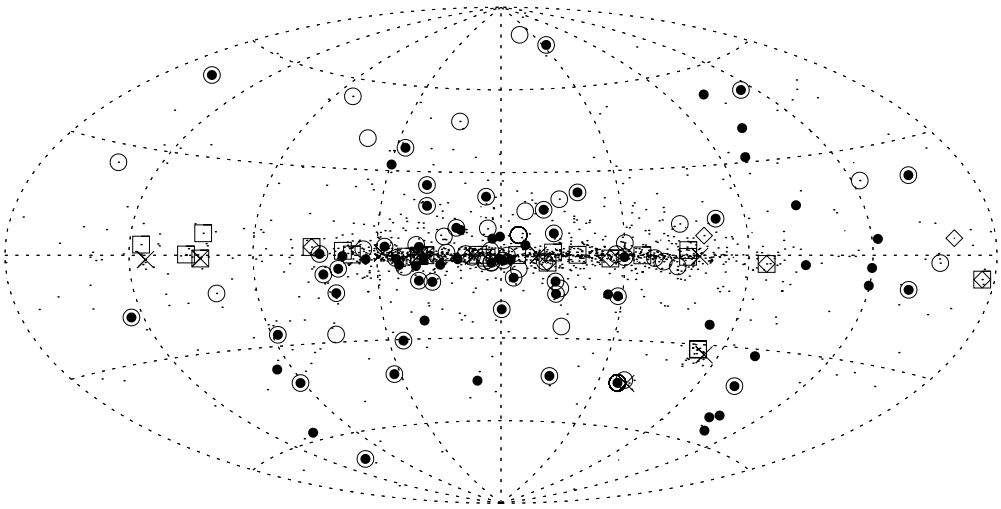

Figure 1.5: Distribution of pulsars on the sky. The locations of the known pulsars are plotted in Galactic (II) coordinates in an Aitoff projection. Pulsars in globular clusters are excluded. Plotting symbols are the same as in Figure 1.3, with the addition of small filled circles for MSPs ($P \leq 25\,\mathrm{ms}$). The majority of pulsars are found within a few degrees of the Galactic plane. This includes all the youngest pulsars (the SGRs, AXPs, gamma-ray pulsars, and SNR associations) except those found in the Magellanic clouds. The MSPs are much more isotropically distributed.

we report the discovery of three new binary MSPs in the globular cluster known as M62. These were the first objects discovered using the new Green Bank Telescope, and they elevate M62 to third place among clusters with the most known pulsars. Two additional globular cluster pulsar discoveries are reported in the appendix.

1.2 Pulsar Searches

Before we can search for pulsations from a given object at a given wavelength, we must measure its intensity as a function of time. The methods by which raw data are collected and used to generate this "time series" differ widely, depending on the wavelength of interest. In sections 1.2.1 and 1.2.2 below, we point out some of these differences. We will see that for the most part, they stem from the wave/particle nature of light, and propagation effects through the various media between the pulsar

and the telescope. Once a time series has been prepared, the remaining steps in a pulsation search are largely independent of wavelength; these steps are outlined in section 1.2.3. In section 1.2.4, we discuss the effect of a time dependent pulse period, one of the principal difficulties encountered in pulsar searches.

1.2.1 X-Ray and Gamma-Ray Data

After travelling relatively unimpeded through vast reaches of space, gamma-rays are unable to penetrate the earth's dense atmosphere. The same is true for X-rays, although they also suffer some absorption in the interstellar medium. Detection of high-energy radiation therefore requires high altitude balloon-, rocket-, or satellite-borne instruments. This makes them relatively expensive to deploy, severely limits their collecting area, and makes them prone to unrecoverable failures (ranging from inconvenient to catastrophic).

The interaction of high-energy (X-ray and gamma-ray) radiation with matter is dominated by quantum effects, e.g., photoelectric effect, compton scattering, bremsstrahlung, and pair production/annihilation. At high energies, we therefore deal with the propagation and detection of individual photons. For reviews of high-energy astronomical instrumentation, see Hoffman et al. (1999) and Ramsey et al. (1994). Our gamma-ray search in Chapter 2 used data from the Energetic Gamma Ray Experiment Telescope aboard the *Compton Gamma Ray Observatory*. Fundamentally, EGRET detects an incoming high-energy gamma-ray by converting it into an electron-positron pair. The energies and tracks of the resulting pair are measured and used to determine the energy of the original photon and the direction from which it came (Kanbach et al., 1989). Our X-ray search in Chapter 3 used data from the Proportional Counter Array aboard the *Rossi X-Ray Timing Explorer*, ultimately detecting the photoelectric absorption of an incident X-ray photon in the gas filled detector, with an output signal approximately proportional to the incident photon energy (Swank et al., 1996). The PCA has essentially no spatial resolution, though the field of view is collimated at $\sim 1^{\circ}$.

While radio search observations typically generate gigantic continuous data sets (sec. 1.2.2), discrete photon data sets are of much smaller size. Reducing a raw photon list to a searchable time series is computationally trivial. First, each photon's spacecraft time of arrival (TOA) is adjusted to a solar system barycentric TOA. If additional corrections are necessary (qv sec. 1.2.4), they can also be applied before building the time series. An empty, discrete time series array is then created, and bins corresponding to corrected photon arrival times are increased by 1. Weighting the photons with *a priori* knowledge of the spectral or spatial distribution of the source photons can increase signal-to-noise (e.g., sec. 2.2.2).

Also unlike radio searches, high-energy searches are practically immune to interference. Cosmic ray interactions in the detector can mimic high-energy photon events. Specialized systems attempt to reject these events on board, before they are telemetered to the ground, but not all spurious events are successfully screened. However, Galactic cosmic ray detections are better characterized as background or noise, rather than interference, since they rarely overwhelm a strong source and are unlikely to present like a periodic signal. Other effects can lead to detectable periodicity at the spacecraft orbital period — Earth albedo, South Atlantic Anomaly passage, and Earth occultations of the source, for instance. Additional sources of high-energy periodic "interference" are usually instrumental in origin (e.g., the 402 second *ROSAT* wobble) and are easily excised.

1.2.2 Radio Data

Radio waves can penetrate the atmosphere to reach ground-based telescopes. This allows the construction of sensitive telescopes with very large collecting areas, including single dishes (e.g., the 305 m diameter Arecibo telescope in Puerto Rico), and groups of individual dishes working in concert (e.g., the Very Large Array in New Mexico, or the planned Square Kilometer Array). Focused radio waves are detected as a continuous electric field oscillation in an antenna, not as discrete photons. For detailed descriptions of radio telescopes, see Burke & Graham-Smith (1997) and Kraus (1986).

Radio pulsars are very broad band and have a fairly steep radio spectrum (flux density $S \propto \nu^{-\alpha}$, where ν is the radio frequency and $\alpha \sim 1.6$ is the spectral index). On their own, these two facts indicate that a pulsar's signal-to-noise ratio is maximized by observing with a large bandwidth about a low central radio frequency. However, the situation is complicated by propagation effects in the interstellar medium, specifically dispersion and scattering.

Dispersion occurs when different frequency components of a wave travel at different speeds in a given medium. For an electromagnetic wave propagating in a plasma with a free electron number density n_e, the angular frequency ω and wave number k are related by the dispersion relation $\omega^2 = c^2k^2 + \omega_p^2$, where ω_p is the plasma frequency, $\omega_p^2 = 4\pi n_e e^2/m_e$, c is the speed of light in vacuum, e is the electron charge, and m_e is the electron mass. The group velocity is given by

$$v_g = \frac{\partial\omega}{\partial k} \approx c\left(1 - \frac{2\pi n_e e^2}{m_e\omega^2}\right) \tag{1.6}$$

(valid for $\omega \gg \omega_p$). Higher frequencies travel faster than lower frequencies. A pulse travels a distance ℓ in time

$$t = \int_0^\ell \frac{dx}{v_g} = \frac{\ell}{c} + \frac{2\pi e^2}{cm_e\omega^2}\int_0^\ell n_e dx. \tag{1.7}$$

The integral on the right hand side of Equation 1.7 defines the "dispersion measure" (DM), usually expressed in units of $\mathrm{cm}^{-3}\,\mathrm{pc}$. In a bandwidth $B_{\mathrm{MHz}} \ll \nu_{\mathrm{MHz}}$ (both expressed in MHz), the total delay between the highest and lowest frequency components is

$$\tau_{\mathrm{DM}} = \left(\frac{202}{\nu_{\mathrm{MHz}}}\right)^3 \mathrm{DM}\, B_{\mathrm{MHz}}\,\mathrm{ms}. \tag{1.8}$$

In the upper panel of Figure 1.6 we see intensity as a function of pulse phase and radio frequency for PSR B1257+12. The frequency dependent dispersive delay is obvious. If we construct a time series by simply recording the total power in the band as a function of time, search sensitivity will be seriously degraded (bottom panel, dotted line).

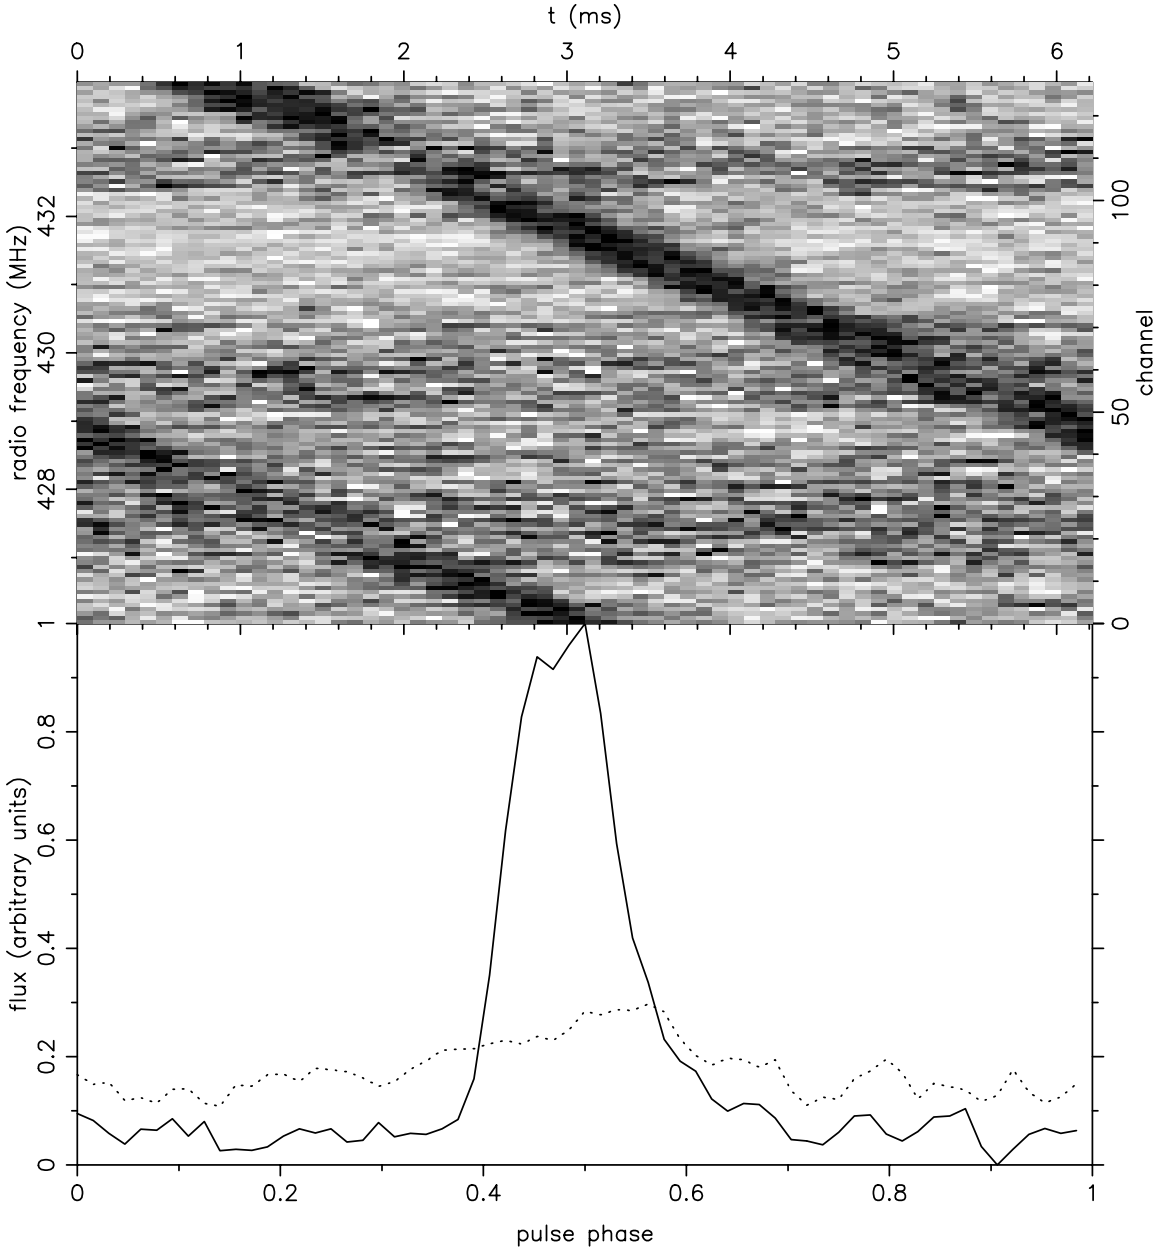

Figure 1.6: Radio pulse dispersion. Upper panel shows a radio frequency resolved pulse profile of PSR B1257+12. A 5.2 s data set was folded modulo the 6.218 ms pulse period to increase the signal-to-noise ratio. The pulse is clearly dispersed, arriving first at high frequencies, later at lower frequencies (DM = 10.186 $\mathrm{cm}^{-3}\,\mathrm{pc}$). Bottom panel shows the total power in the band as a function of pulse phase (dotted line), i.e., with no correction for dispersion. Solid line shows the pulse profile with the frequency channels appropriately shifted before adding.

Dispersion is commonly mitigated by dividing the pass band into $N_{\rm chan}$ smaller bandwidth channels, and recording a separate time series for each. This is accomplished by the use of a filterbank or autocorrelation spectrometer. A single time series can then be constructed by delaying the higher frequency channels before adding them together (Fig. 1.6, bottom panel, solid line). In this manner, the overall dispersive pulse smearing is reduced by a factor $1/N_{\rm chan}$. Of course, the size of the raw data set is increased by a factor of $N_{\rm chan}$, and if the proper DM is unknown *a priori*, the raw data must be dedispersed for a number of different DM trials. Each resultant time series is then searched separately. For this reason, an unknown DM greatly increases the computational complexity of a blind radio pulsar search.

Multipath propagation or scattering in the turbulent interstellar medium can also broaden a radio pulse. Scattering causes a pulse to develop an asymmetric tail, which is wider at lower frequencies (Fig. 1.7). By empirically fitting the observed broadening of 196 pulsars, Cordes et al. (1991) determined that

$$\tau_{\rm scat} \approx \left(\frac{1000}{\nu_{\rm MHz}}\right)^{4.4} 10^{-7.231+0.9255\log_{10}{\rm DM}+0.814613(\log_{10}{\rm DM})^2}. \qquad (1.9)$$

The only way to overcome the effects of scattering is to observe at high frequencies. Thus, a radio pulsar search should use the lowest frequency and largest bandwidth possible, while keeping the effects of dispersion and scattering at a reasonable level.

Another all too common problem encountered in radio pulsar searches is interference. Weak astronomical signals are easily drowned out by bright impulsive interference, and periodic interference can often be mistaken for a pulsar, sometimes even mimicking the effects of interstellar dispersion. Radio frequency interference can be of natural origin (e.g., lightning) or of human origin (e.g., TV, cell phones, radar, spark plugs, and even an occasional electric blanket). If interference does initially produce a candidate pulsar signal in a search, further analysis will usually reveal its true nature; for instance, it may be narrow band, not dispersed, or detected at widely separated sky positions. However, a single detection of a new pulsar is usually considered suspect until it can be verified by further observations.

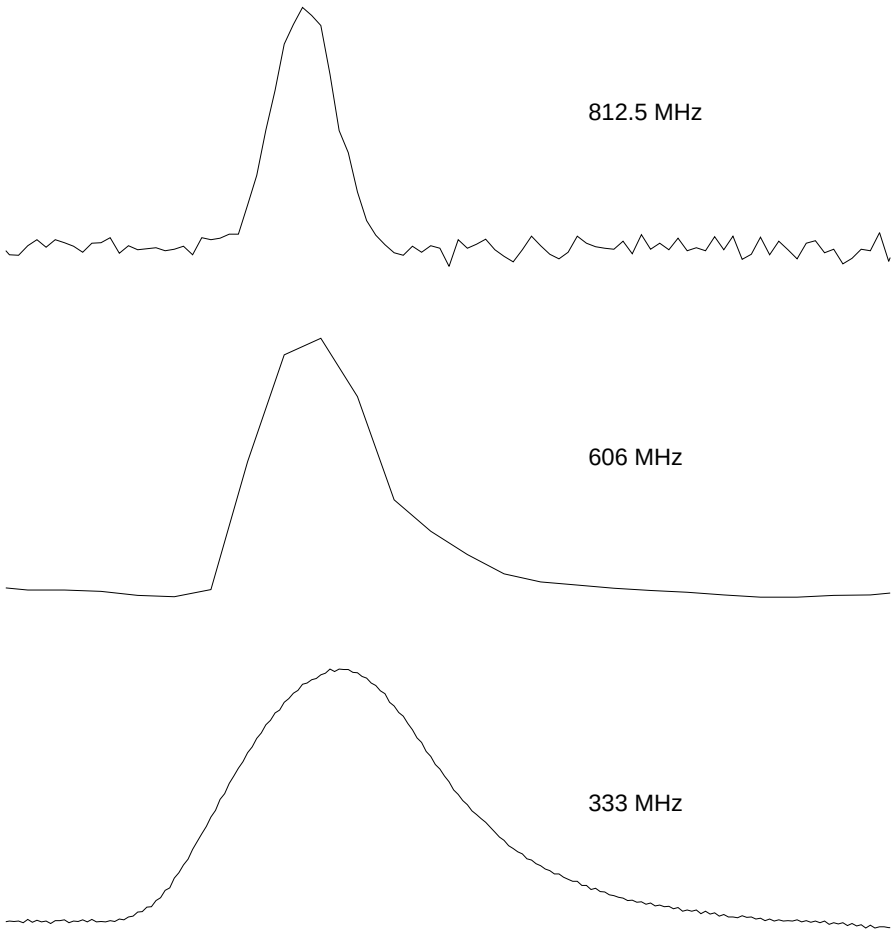

Figure 1.7: Pulse broadening due to scattering. A blowup of the Crab pulsar interpulse is shown at three different radio frequencies, 812.5 MHz, 606 MHz, and 333 MHz. The longer scattering tail at lower frequencies is clearly visible (Moffett & Hankins, 1996; Lundgren, 1994).

1.2.3 Detecting Periodicity

The earliest pulsar discoveries were made using chart recordings. However, even with large telescopes and modern low-noise receivers and amplifiers, very few pulsars are bright enough to have detectable single pulses. Large signal-to-noise ratios can be achieved by adding many pulses together (see, e.g., Fig. 1.6). When the pulse period is unknown, the most efficient method for searching a large range of periods is to use a fast Fourier transform (FFT). In this section, we provide only a brief recipe for FFT-based searches. For reviews and more complete analyses, see Burns & Clark (1969), Lyne & Graham-Smith (1998), Bhattacharya (1998), Middleditch (1976), van der Klis (1989), and Ransom et al. (2002). Also, details specific to the particular searches in this thesis are of course provided in the individual chapters.

The discrete Fourier transform (DFT) of a time series s is given by

$$a_k = \sum_{j=0}^{N-1} s_j \, e^{2\pi i j k/N}, \tag{1.10}$$

where N is the total number of points in the time series and $i = \sqrt{-1}$. The complex Fourier amplitude a_k corresponds to spin frequency $f_k = k/T_{\rm obs}$, where $T_{\rm obs}$ is the total duration of the time series. The index k runs from 0 to $N-1$, though for real data $a_{N-k} = a_k^*$, so only half of the Fourier amplitudes are independent. For most applications, the Fourier amplitudes are absolute squared to produce an estimate of the (real) power spectrum $P_k = |a_k|^2$. A bright, coherent, periodic signal will appear as a series of spikes in the power spectrum at frequencies corresponding to the various harmonics of the signal (or a single spike at the fundamental frequency in the case of a sinusoidal waveform). Figure 2.2 is a good example. Calculating the DFT of Equation 1.10 requires roughly $\sim N^2$ operations. The FFT yields the exact same result, but much more efficiently, requiring only $\sim N \log_2 N$ operations (for a derivation, see, e.g., Press et al. 1995). Still, the FFT stage is usually the most computationally demanding part of a pulsar search.

To find pulsar candidates, the power spectrum is searched for bins containing

power significantly above the local noise level. Although X-ray pulsars tend to have broad, nearly sinusoidal pulse profiles, gamma-ray and radio pulsars typically have narrower and/or more complicated pulse shapes (i.e., with rich harmonic content). Sensitivity to these pulsars can be improved by incoherently adding together the power at harmonically related frequencies and again looking for significant summed peaks.

The FFT, harmonic sum, and candidate search must be repeated for each individual time series in the search. These separate iterations might correspond to different sky positions, different DM trials, or different spin evolution trials (sec. 1.2.4 below). A modern pulsar search can easily use up many thousands of CPU hours on fast parallel supercomputers. Since the computational cost of the FFT search stage is so high, candidate thresholds may be set somewhat low, to minimize missed detections. Of course, this will result in a large number of false positives, so the FFT candidates may be subjected to further automated analysis (e.g., see sec. 2.2.2 and sec. 4.2.3) before human inspection of the final candidates. Ideally, strong candidates that are not eliminated as interference can then be verified in follow-up or archival data.

1.2.4 Detecting a Time-Dependent Period

If a pulsar's period changes appreciably over the course of an observation, it may not be detected in a standard FFT search, since its spectral power will be spread over a number of frequency bins. An apparent change in period may be due to a number of causes including intrinsic spin down, and acceleration in a binary system.

Intrinsic $\dot{P}$

As we saw in section 1.1, pulsar spin rates slowly decrease with time. To maintain good sensitivity to a realistic pulsar signal, we would like the change in spin frequency f over the course of an observation to be less than the independent Fourier frequency

step in the power spectrum ($1/T_{\rm obs}$):

$$\dot{f}\,T_{\rm obs} < \frac{1}{T_{\rm obs}}. \tag{1.11}$$

In terms of period and period derivative, we require

$$T_{\rm obs} < \frac{P}{\sqrt{\dot{P}}}. \tag{1.12}$$

For MSPs, this means that coherent observations are limited to $T_{\rm obs} \lesssim 1$ month, while for young pulsars, the limit is $T_{\rm obs} \lesssim 8$ hr. The vast majority of pulsar surveys use integration times much shorter than this (e.g., Chapter 4). Globular cluster observations may exceed 8 hr, but young pulsars are not expected in old clusters (and MSP sensitivity is not reduced). Note also that acceleration of a pulsar in the gravitational potential of a globular cluster produces an apparent $\dot{P}$ that is of the same order as its intrinsic spin down, i.e., not a problem for detection.

Under rare circumstances, intrinsic spin down can become a problem. This is in fact the case in Chapter 2, where $\gtrsim 2$ week observations were searched for young pulsars. In order to detect pulsations, we attempted to remove the effects of $\dot{P}$ before calculating the FFT. Without any prior knowledge of the exact value to use, we had to correct each time series for a number of different $\dot{P}$ trials and repeat the Fourier analysis for each.

Binary Motion

For a pulsar moving in a binary system, doppler shifts in the observed pulse period can have a far greater effect on searches than slow intrinsic spin down. A binary pulsar's instantaneous apparent $\dot{P}$ can be as much as 10 orders of magnitude larger than its intrinsic $\dot{P}$. Without some sort of correction, an ordinary FFT search would have little chance of detecting such a signal.

An extensive discussion of binary pulsar search techniques can be found in section 5.2.3. To summarize, existing methods for detecting binary pulsars are effective

only when the orbital period P_b is much longer or much shorter than the observation time $T_{\rm obs}$. When $P_b \gg T_{\rm obs}$, the time dependence of the observed spin period is approximately linear, and constant $\dot{P}$ trials can be used (very much like the intrinsic $\dot{P}$ searches described above). When $P_b \ll T_{\rm obs}$, phase modulation of the pulsar signal may produce a detectable comb of sidebands centered on the rest spin frequency in the power spectrum.

Traditionally, when P_b and $T_{\rm obs}$ are comparable, pulsar searches have had very poor sensitivity. In section 5.2.4, we describe a powerful new search method, based on dynamic power spectra, that nicely fills in the $P_b \sim T_{\rm obs}$ gap left by the other methods. This new technique has already successfully discovered two new binary MSPs, and promises to be an important addition to the pulsar searcher's toolkit.

1.3 Thesis Organization

The individual chapters speak for themselves, having been written as stand-alone papers. Chapters 2 and 3 have already been published (Chandler et al., 2001; Chandler & Rutledge, 2000), while Chapters 4 and 5 will appear shortly. Each chapter begins with its own abstract, which we will not attempt to distill even further in the present context.

Despite the title of this thesis, the chapters are actually arranged from high to low energies, i.e., from gamma-rays to radio. This was the order in which the projects were undertaken, but it also has an evolutionary significance; we begin with a search targeted at young pulsars, followed by a search for a pulsar during its pupal accretion stage, finishing with two searches designed to detect recycled MSPs.

Chapter 2

A Search for Radio-Quiet Gamma-Ray Pulsars

Abstract

Most Galactic point sources of gamma-rays remain unidentified. The few (extra-solar) sources that have been identified are all young, rotation-powered pulsars, all but one of which were identified using radio ephemerides. The radio-quiet Geminga pulsar was identified only after pulsations were discovered in a coincident X-ray source. Observational evidence indicates that many of the unidentified Galactic sources are likely to be pulsars, and some theoretical models predict a potentially large population of radio-quiet gamma-ray pulsars. We present a new method for performing sensitive gamma-ray pulsar searches. We used this method to search several of the strongest EGRET sources for pulsations. This was a blind search for new pulsars, covering a frequency and frequency derivative phase space large enough to detect Crab-like pulsars as well as lower frequency, high magnetic field "magnetars." No new pulsars were discovered and we report upper limits constraining the characteristics of any signals contained in the data sets searched.

(Originally appeared as Chandler et al., 2001, ApJ, 556, 59.)

2.1 Introduction

The nature of the unidentified Galactic plane gamma-ray sources has been a long-standing problem, dating back to the first dedicated gamma-ray astronomy satellites of the 1970s (*SAS-2* and *COS-B*). Despite vast technical improvements over earlier missions, the *Compton Gamma Ray Observatory* (*CGRO*) has only compounded the problem; although a small number of identifications have resulted from *CGRO* observations, many new unidentified sources have been discovered. Less than 10% of the currently cataloged Galactic plane gamma-ray sources have been identified. The principal difficulty is the large uncertainty in the gamma-ray positions. A typical gamma-ray source error box may contain hundreds of possible counterparts at other wavelengths. Nevertheless, all the identifications of persistent gamma-ray sources have been established through such multi-wavelength correlations. In this paper, we describe an unsuccessful attempt to identify several of the strongest high-energy gamma-ray sources via direct analysis of the gamma-ray data. Specifically, we attempted to identify them as pulsars.

Of the four instruments on board the *CGRO*, the Energetic Gamma Ray Experiment Telescope (EGRET) is sensitive to the highest-energy range, about 20 MeV to 30 GeV. Details of EGRET's design and capabilities can be found in Kanbach et al. (1989) and Thompson et al. (1993). Generally speaking, it simultaneously offers good imaging (source localization to within $\sim 0^\circ.5$), spectral resolution ($E/\Delta E \sim 5$), and temporal resolution (~ 50 μsec).

Shortly after the April 1991 launch of the *CGRO*, an EGRET all-sky survey was begun. During this first phase of its mission, EGRET verified ~ 15 previously known *COS-B* sources and discovered dozens of new point gamma-ray sources (Fichtel et al., 1994). As the amount of accumulated data has increased and the gamma-ray background model has improved, several updated source catalogs have been compiled (Thompson et al., 1995, 1996). The latest EGRET catalog, based on over four years of data, includes 271 sources, the majority of which are unidentified (Hartman et al., 1999). The sources can be divided into two groups — those that are in the Galactic

plane and those that are not. All of the identified sources in the latter group (having Galactic latitude $|b| > 10°$) are extragalactic. The 80 sources with $|b| < 10°$ include 5 pulsars, 1 solar flare, and 74 unidentified sources. Of these, a few are probably blazars seen through the Galactic plane and as many as 5 are probably spurious, leaving $\gtrsim 60$ unidentified Galactic point sources of high-energy gamma-rays.

Four of the five aforementioned pulsars had been seen previously at radio wavelengths: the Crab pulsar, the Vela pulsar, PSR B1706-44, and PSR B1055-52. These EGRET sources were identified by epoch folding the gamma-ray data with the known pulsar periods (Nolan et al., 1993; Kanbach et al., 1994; Thompson et al., 1992; Fierro et al., 1993). Additionally, pulsations from PSR B1951+32 have been seen in the EGRET data (≥ 300 MeV), though this pulsar is not detected as a statistically significant point source (Ramanamurthy et al., 1995). The remaining EGRET pulsar, Geminga, was not previously detected as a radio pulsar. Its identification came only after pulsations were discovered in a coincident *ROSAT* X-ray source (Halpern & Holt, 1992) and then verified in the EGRET data (Bertsch et al., 1992). Although there have been reports of extremely weak pulsations from Geminga at low radio frequencies (~ 100 MHz) (e.g., Malofeev & Malov, 1997), Geminga can still be regarded as a "radio-quiet gamma-ray pulsar."

Efforts to identify additional EGRET sources with radio pulsars have been largely unsuccessful. Many positional coincidences exist, but few of them are expected to be real associations. The most conclusive evidence for an association would be the detection of pulsations in the gamma-ray data. Thompson et al. (1994), Fierro et al. (1995), and Nel et al. (1996) epoch folded EGRET data with ephemerides of over 350 radio pulsars, with no significant detections, though marginal evidence for pulsed gamma emission has been found for PSR B0656+14, PSR B1046-58, and PSR J0218+4232 (Ramanamurthy et al., 1996; Kaspi et al., 2000; Kuiper et al., 2000). Deep radio searches of EGRET source error boxes (Nice & Sayer, 1997; Lundgren et al., 1995) have resulted in only one new pulsar detection, but the pulsar proved to be unrelated to the target gamma-ray source.

Despite these failures, there are a number of reasons to believe that many of the

unidentified Galactic EGRET sources are indeed young, rotation-powered pulsars. First of all, the only definitive identifications of persistent Galactic sources are all pulsars. It is reasonable to assume that even if other types of Galactic sources exist, we probably have not identified all of the detectable pulsars.

Kaaret & Cottam (1996) argued that a significant fraction of the Galactic plane sources in the second EGRET catalog were located in OB associations, which are likely to house young pulsars. Using the known distances of the OB associations, they estimated the luminosities of the coincident gamma-ray sources and found them to be consistent with the known EGRET pulsars. They ultimately estimated that ~ 20 of the 25 sources in the second EGRET catalog with $|b| < 5^\circ$ are pulsars. That these pulsars have not been detected in radio wavelengths can be attributed to narrow beaming of radio pulses and the high interstellar dispersion expected in star forming regions.

Sturner & Dermer (1995) found a significant correlation between the high confidence point sources in the first EGRET catalog and supernova remnants (SNRs). SNR associations have been suggested for as many as 7 of the 32 low-latitude unidentified sources in the second EGRET catalog (Sturner et al., 1996; Esposito et al., 1996). It is plausible to assert that the gamma-rays are due to either a young pulsar, born in the supernova explosion, or cosmic rays accelerated in the expanding supernova shockwave. At present, neither cause can be ruled out. Several authors have argued for the presence of pulsars in a few specific cases (Bhattacharya et al., 1997; Brazier et al., 1998; Roberts & Romani, 1998; Brazier et al., 1996), but evidence for shock front cosmic ray production also exists for several of the SNR associations. Until future missions resolve spatial structure in the gamma-ray sources or improve their spectral characterization, conclusive identifications of shock-powered SNRs will be very difficult to establish. On the other hand, detection of pulsations in the EGRET data has the potential to resolve the issue on a source by source basis.

McLaughlin et al. (1996) characterized the time variability of the sources in the second EGRET catalog (using data from phases 1, 2, and 3). In their classification scheme, the known EGRET pulsars were shown to be non-variable or marginally

variable. Although they argue for the existence of a genuine population of variable low-latitude sources, they also find a significant number of non-variable low-latitude sources, which they conclude are most likely pulsars. The gamma-ray pulsar population model of McLaughlin & Cordes (2000) predicts that EGRET should see 20 (low-latitude) pulsars, a result that is consistent with the variability arguments, since 17 of the unidentified low-latitude sources are non-variable. Wallace et al. (2000) revisited the question of EGRET source variability. Looking at shorter time scales ($\sim$ days), they found strong evidence of short-term variability for only four unidentified cataloged sources. This result is consistent with the hypothesis that many of the unidentified sources could be pulsars.

Thus, despite indications that many of the unidentified EGRET sources are pulsars, previous attempts to identify them with radio pulsars have failed. This raises the possibility that many of the unidentified sources could be radio-quiet pulsars like Geminga. Romani & Yadigaroglu (1995) have proposed an outer gap model of pulsar gamma-ray production which predicts that pulsed gamma radiation is beamed into a larger solid angle than the radio emission. This model is shown to account for observed properties of individual pulsars (e.g., pulse profiles, relative phase of radio and gamma-ray pulses) and the overall observed population of radio-only and radio/gamma-ray pulsars. Their model predicts a large number of detectable radio-quiet gamma-ray pulsars, 2.5$\times$ the number of detected radio-loud gamma-ray pulsars. In a more recent analysis (Yadigaroglu & Romani, 1997), they showed that radio-quiet pulsars can account for essentially all of the strongest unidentified low-latitude EGRET sources.

Another possible mechanism for producing radio-quiet gamma-ray pulses has recently been suggested. Even before the identifications of the soft gamma repeaters SGR 1806-20 and SGR 1900+14 as magnetars (Kouveliotou et al., 1998, 1999; Thompson & Duncan, 1995, 1996), Baring & Harding (1997) proposed the existence of a class of high magnetic field, radio-quiet high-energy pulsars. According to the model, intense magnetic fields can inhibit pair production, suppressing a pulsar's radio emission. A highly magnetized pulsar "can still emit gamma-rays prolifically." While

there has been a report of weak low frequency radio pulsations from SGR 1900+14 (Shitov, 1999), this is not inconsistent with the model (Baring & Harding, 1998).

The EGRET instrument long outlived its intended two to four year lifetime. After a gyroscope failure in December 1999, NASA decided to de-orbit the *CGRO* in June 2000. The next generation high-energy gamma-ray satellite *GLAST* will launch no sooner than 2005. Thus, no new pertinent gamma-ray data are expected in the near future. Searches of EGRET error boxes at other wavelengths continue, but it now seems that further application of conventional approaches is unlikely to result in new identifications. In this paper we test the hypothesis that many of the unidentified EGRET sources are radio-quiet pulsars, by searching for pulsations in the existing gamma-ray data directly (Mattox et al., 1996). Previous attempts to find pulsations have involved epoch folding with known signal parameters. We instead performed a blind search for unknown pulsars on several of the strongest EGRET point sources.

As we will explain in more detail below, this is a difficult problem. The size and number of fast Fourier transforms (FFTs) that must be calculated make this a computationally intensive problem — one that would have been infeasible 10 years ago. Due to the small photon flux of the sources, signal-to-noise ratios are very low. Special techniques must be employed to keep detection sensitivity as high as possible. Even then, we require a somewhat fortuitous pulse waveform. Despite the difficulties, even one detection would add significantly to our knowledge of gamma-ray pulsars and pulsar emission mechanisms. Unfortunately, no candidate signal survived all of our detection criteria.

In section 2.2 below, we describe our analysis method. We indicate the specific EGRET sources we searched for pulsations and the range of pulsar frequencies and frequency derivatives to which the search was sensitive. In section 2.3, we describe the results of our pulsation search. We discuss the determination of detection significance and the connection between signal waveform and detectability, which we then use to place limits on the characteristics of any signals contained in the EGRET data searched. Finally, in section 2.4, we discuss these results and their implications.

2.2 Analysis

2.2.1 Principles

Finding a pulsar signal in the EGRET data is not a trivial problem. Even the strongest unidentified EGRET sources have count rates ($\geq 100\,\mathrm{MeV}$) of ~ 1 source photon per hour, over a typical background of ~ 5 photons per hour (in the Galactic plane). To achieve an adequate signal-to-noise ratio to allow an FFT search, long data sets — on the order of weeks — are required. The data sets used in the present analysis spanned 14 to 38 days each. The independent frequency spacing of the FFT is given by $1/T$, where T is the total duration of the time series data. The long EGRET data sets therefore result in sub-microHertz spectral resolution, requiring large FFTs to cover a given frequency range.

Pulsar rotation frequencies are generally observed to decrease slowly over time. This is attributed to the conversion of rotational kinetic energy into the radiation emitted by the pulsar, hence the term "rotation-powered pulsar." Such intrinsic spin down is not usually a problem for pulsar searches at radio and X-ray wavelengths, where shorter data sets ($\sim$ minutes to hours) are often used. Over a 14-day observation, however, a pulsar's spin frequency can change by as much as several hundred microHertz. The signal from such a pulsar will be spread out over hundreds of frequency bins in the FFT, rendering it undetectable.

Figure 2.1 shows a plot of the 472 pulsars in the Princeton catalog (Taylor et al., 1993) that are observed to be spinning down. The pulsars are plotted in a phase space of spin frequency f and its first-order time derivative $\dot{f}$. If we Taylor expand a pulsar's frequency evolution $f(t)$, we see that over the course of an observation of duration T, the pulsar frequency will visit the range $\Delta f = \dot{f}T + \ddot{f}T^2/2 + \cdots$ (a dot denotes a derivative with respect to time, in this case, evaluated at the beginning of the observation). In the discrete Fourier spectrum, this is equivalent to a drift of $\Delta B = T\Delta f = \dot{f}T^2 + \ddot{f}T^3/2 + \cdots$ independent frequency bins. Using $T = 14$ days, we find that a frequency derivative of $|\dot{f}| = 1/T^2 = 6.8 \times 10^{-13}\ \mathrm{Hz\,s^{-1}}$ causes a drift equivalent to one power spectrum bin. As is evident in Figure 2.1, many

pulsars are observed with larger first-order frequency derivatives than this. Likewise, a second derivative of $|\ddot{f}| = 2/T^3 = 1.1 \times 10^{-18}$ Hz s^{-2} will cause a similar one-bin drift. Fortunately, this is more than two orders of magnitude larger than observed second frequency derivatives. We are therefore justified in adopting a linear model for the frequency evolution.

To counteract this spin down effect, we attempt to remove the frequency drift of a pulsar signal in the time domain before calculating the FFT. In a blind search, little is known *a priori* about a pulsar's frequency evolution (except that its behavior is expected to be similar to that of the known pulsars). A number of frequency drift trials must therefore be performed for each source, whereby the data are corrected for an assumed frequency derivative, and an associated FFT is calculated and analyzed.

The known EGRET pulsars are highlighted (with triangles) in Figure 2.1. The gamma-ray pulsars tend to have high frequencies and large (negative) values of $\dot{f}$. This is not surprising since we expect pulsars in this region of phase space to be the most energetic; i.e., to have the largest rotational energy loss $\dot{E} = 4\pi^2 I f \dot{f}$, where I is the pulsar's moment of inertia. This fact is rather unfortunate for our purposes, since the phase space that our search must cover is therefore large. For this project, we chose to cover frequencies up to 40 Hz and frequency derivatives large enough to include the Crab pulsar. This search phase space is represented by the darker shaded region in the figure. As described below (section 2.2.2), our frequency drift search was carried out over trial values of $\dot{f}/f$, not simply $\dot{f}$. The $\dot{f}$ limit of the dark shaded region in the figure corresponds to

$$-\frac{\dot{f}}{f} \leq \frac{3.7 \times 10^{-10} \text{ Hz s}^{-1}}{30 \text{ Hz}} \approx 1.2 \times 10^{-11} \text{ s}^{-1}. \tag{2.1}$$

For the actual FFT calculation, we require a Nyquist frequency of $f_{\text{Nyq}} \geq 160$ Hz, so that we may calculate sums of up to four harmonics without aliasing. Given the required length of the data sets, this means we must calculate FFTs of $N = 2 \cdot 160 \cdot T = 2^{29} - 2^{30} \sim 10^9$ points. Each source requires $\sim 3.7 \times 10^{-10} \cdot T^2 \sim 10^3$ frequency derivative trials. Such a search has only recently become computationally

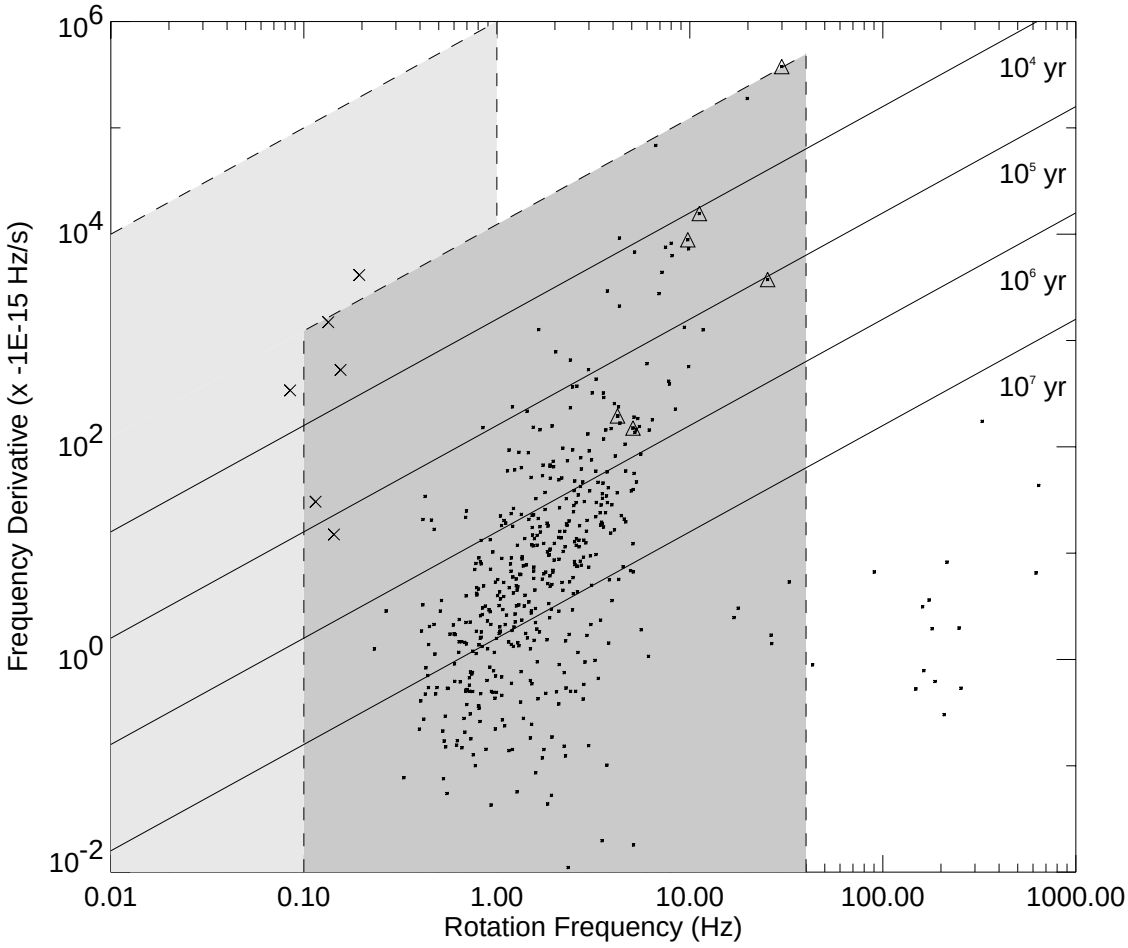

Figure 2.1: Gamma-ray pulsar search phase space. The $f - \dot{f}$ phase space relevant to our search is highlighted. The 472 pulsars in the Princeton catalog (Taylor et al., 1993) which have positive period derivatives are denoted by black dots. The known EGRET pulsars are indicated by triangles. They are, in order of decreasing rotation frequency, Crab, PSR 1951+32, Vela, PSR 1706-44, Geminga, and PSR 1055-52. The Xs denote, in order of decreasing $|\dot{f}|$, SGR 1900+14, SGR 1806-20, 1E 1048-5937, 1E 1841-045, 4U 0142+615, and 1E 2259+586. The solid lines plotted on the figure are lines of constant age for pulsars born at short periods, assuming a vacuum dipole braking law. The dark shaded region corresponds to the Crab-like pulsar search and the lighter region denotes the magnetar search phase space described in the text.

feasible.

2.2.2 Search Methodology

We began by selecting a number of unidentified EGRET sources for analysis. For each source, we prepared a data set, essentially a list of photon arrival times. Before beginning the search, we determined a list of frequency derivative trials to cover our target phase space.

For a given frequency derivative trial, the photon arrival times were corrected for the assumed pulsar spin down rate. Using these adjusted arrival times, we calculated an FFT, which we used to construct a normalized power spectral estimate. This power spectrum was then searched for significant candidates, i.e., frequency bins containing statistically significant excess power. Sums of two and four harmonics were also searched in this way. This process was repeated for each frequency derivative trial.

Candidates from the FFT stage of the analysis were subjected to further scrutiny. In the vicinity of each candidate, we sampled the $f - \dot{f}$ phase space more finely, again looking for significant powers. If any candidates survived this (more stringent) cut, we put them though several final verification procedures. Each of these steps is described in detail below.

Source Choice and Data Preparation

The unidentified EGRET sources chosen for analysis in this project were selected simply on the basis of strength. The most appropriate measure of source strength in this case is the quantity N_s^2/N_t, where N_s is the number of source photons and N_t is the total number of photons in a data set (i.e., source plus background). This is the square of the signal-to-noise ratio for a data set. As we show in detail in section 2.3, the expected spectral power from a periodic source is proportional to N_s^2/N_t and the significance of a detection is exponential in N_s^2/N_t. So the general idea was to choose an EGRET viewing period and select a set of photon events so as to maximize N_s^2/N_t. In some cases, adjacent (or nearly adjacent) observations were concatenated.

The sources so chosen and the viewing periods (VPs) analyzed are listed in Table 2.1, ranked in order of N_s^2/N_t. Also indicated in the table are possible supernova remnant associations and the source positions, though these were not source selection criteria.

Having chosen an EGRET viewing period, we selected photons based on three criteria — one spectral, two spatial. Due to the diffuse gamma-ray background, source signal-to-noise is very low for photon energies below 100 MeV. Our analysis included only photons above this threshold. The imaging capabilities of the EGRET instrument were characterized during its pre-launch calibration (Thompson et al., 1993). From the measured point spread function (PSF), it was determined that approximately 67% of the photons from a source will be observed to come from within a cone of energy-dependent half angle

$$\theta_{67} = 5^\circ.85 \left(\frac{E_\gamma}{100\,\mathrm{MeV}} \right)^{-0.534} \tag{2.2}$$

(E_γ is the photon energy). In-flight analyses of EGRET's angular dispersion characteristics have been shown to be in very good agreement with the pre-launch PSF (Esposito et al., 1999). Only photons measured to have arrived from within θ_{67} of the assumed source direction were included for analysis. Tests with the known EGRET pulsars indicated that this cut was robust and optimal — small changes in this cutoff angle had little effect and large changes either way reduced the recovered spectral power. Finally, we selected only photons that were more than four times this angle from the earth limb.

The next step in the data preparation was to correct the spacecraft photon arrival times to solar system barycentric arrival times, using the known position and orientation of the instrument and the assumed direction of the source. For the source directions, we used the maximum likelihood positions from the first or second EGRET catalog. The only exception is 2EG J2020+4026 for which we used the position of a coincident *ROSAT* source (Brazier et al., 1996). The actual positions used for each source are listed in Table 2.1. The errors on the position estimates (semimajor axis of 95% confidence error ellipse) are typically $\lesssim 0^\circ.5$ (see Thompson et al., 1995).

Table 2.1: EGRET point sources searched

2EG Source	SNR	l II	b II	EGRET VP analyzed	Duration (days)	N_s^2/N_t	$[N_s^2/N_t]_{\rm eff}$	$\dot{f}$ trials	$\alpha\mathcal{F}^2$	pulse FWHM[a]	pulsed fraction[b]
J2020+4026	γ Cygni	78.1	2.3	203	21	76	73	1210	< 0.94	$> 23\%$	$< 49\%$
J0618+2234	IC 443	189.2	3.0	$.2 - 1.0$	38	46	32	3947	< 2.14	$> 12\%$	$< 75\%$
J2019+3719[c]		75.46	0.60	203	18	34	27	859	< 2.53	$> 10\%$	$< 81\%$
J1835+5919		88.8	25.1	212	14	28	28	540	< 2.44	$> 10\%$	$< 80\%$
J1021-5835[c]		284.45	-1.20	14	14	27	27	524	< 2.53	$> 10\%$	$< 81\%$

[a] FWHM of a single Gaussian peak, expressed as a percentage of the pulse period; assumes 100% pulsed fraction.

[b] For typical radio pulsar waveform (see text).

[c] The third EGRET catalog (Hartman et al., 1999) split 2EG J2019+3719 and 2EG J1021-5835 into multiple sources. The upper limits reported assume the validity of the second EGRET catalog (Thompson et al., 1995) (see section 2.2.2).

Our analysis was carried out before the publication of the third EGRET source catalog (Hartman et al., 1999). The updated positions in the new catalog adversely affect only two of our sources. The highly significant second EGRET catalog source 2EG J2019+3719 was split into two lower-confidence detections in the third catalog (3EG J2016+3657 and 3EG J2021+3716). Although there is significant overlap between our searched error box and the error boxes of the two new sources, neither would individually be strong enough to produce pulsations detectable in our blind search. A similar situation exists for 2EG J1021-5835 (split into 3EG J1013-5915, 3EG J1014-5705, and 3EG J1027-5817). If the third catalog is correct, we could not have detected pulsations from these weaker individual sources. We still report our search results for 2EG J2019+3719 and 2EG J1021-5835 in Tables 2.1 and 2.3, though the upper limits apply only if the second catalog is assumed to be correct.

We now consider the effect of the position uncertainties on our search. A position error of ϵ, measured in radians, leads to a Doppler shift from the earth's orbital motion of

$$\delta f = \frac{v_\oplus}{c} \epsilon f \sin\theta = 10^{-5} \frac{\epsilon}{10^{-2}} \frac{f}{10 \text{ Hz}} \sin\theta \text{ Hz}, \tag{2.3}$$

where $v_\oplus/c = 10^{-4}$ is the magnitude of the earth's orbital velocity in units of the speed of light, and θ is the angle between the earth's orbital velocity and the EGRET source direction. A pure shift of frequency will have no effect on the *detection* of pulsations, though the frequency estimate will be incorrect. In the event of a detection, the correct rest frequency of the pulsar and its position can hopefully be refined using other EGRET observations.

Differentiation of Equation 2.3 leads to

$$\delta \dot{f} = \frac{v_\oplus}{c} \epsilon \dot{\theta} f \cos\theta + \mathcal{O}\left[\dot{\theta}^2\right], \tag{2.4}$$

where $\dot{\theta} = 2 \times 10^{-7}$ rad s^{-1} is the earth's orbital angular velocity. Thus

$$\delta \dot{f} = 2 \times 10^{-12} \frac{\epsilon}{10^{-2}} \frac{f}{10 \text{ Hz}} \cos\theta \text{ Hz s}^{-1}. \tag{2.5}$$

This is much less than the limiting $\dot{f}$ of our search. Therefore, as for frequency, $\delta\dot{f}$ does not compromise the detection sensitivity, but does prevent a precise determination of the intrinsic pulsar $\dot{f}$ unless more observations are used to estimate the pulsar position and rotation parameters simultaneously.

Differentiation of Equation 2.4 leads to

$$\delta\ddot{f} = \dot{\theta}^2 \delta f + \mathcal{O}\left[\dot{\theta}^3\right] \tag{2.6}$$

$$\delta\ddot{f} = 4 \times 10^{-19} \frac{\epsilon}{10^{-2}} \frac{f}{10 \text{ Hz}} \sin\theta \text{ Hz s}^{-2}. \tag{2.7}$$

Previously, we noted that an $\ddot{f}$ of approximately $2/T^3 \sim 10^{-18}$ Hz s^{-2} would require a search over trial $\ddot{f}$ values, like the $\dot{f}$ search. For our search phase space, including frequencies up to 40 Hz, we see that $\delta\ddot{f}$ is of this same order. Since we did not include frequency second derivative trials in our search, we must carefully consider the effect $\delta\ddot{f}$ has on our search sensitivity.

The effect of the position error is to cause a source frequency to drift over time. In the FFT, the signal power will be spread over several spectral bins. If we require that this spreading not exceed some fixed number of bins, over the course of an entire observation, we see from Equation 2.7 that we can cover the full source error box in a restricted frequency range, or cover the full 40 Hz frequency range over a smaller region of the error box. Alternatively, we can claim sensitivity to the entire error box over the entire range of search frequencies by reducing the effective source strength. After we correct for $\dot{f}$, the apparent source frequency diverges from its initial value according to $\Delta f(t) = \delta\ddot{f}t^2/2$. Equivalently, in units of independent Fourier bins, $\Delta b(t) = \delta\ddot{f}Tt^2/2$. The fundamental has drifted over a range equivalent to one spectral bin after a time $t_1 = [2/(\delta\ddot{f}T)]^{1/2}$. Since the source strength parameter N_s^2/N_t scales linearly with time, only a reduced effective signal strength

$$\left[\frac{N_s^2}{N_t}\right]_{\text{eff}} = \frac{t_1}{T}\frac{N_s^2}{N_t} = \left(\frac{2}{\delta\ddot{f}T^3}\right)^{1/2} \frac{N_s^2}{N_t} \tag{2.8}$$

has been confined to an acceptable range. The effective strength $[N_s^2/N_t]_{\text{eff}}$ for each

source is listed in Table 2.1. These values were used to determine the sensitivity limits of our search. Note that not all of our target sources were affected by position error. And for those that were, the sensitivity is actually non-uniform over the source error box and searched frequency range. The limits reported are for the worst-case scenario: $f = 40$ Hz and extreme position error (source located on its 95% error contour).

Frequency Derivative Trials

Ultimately, a data set consisted of a list of each photon's barycentric time of arrival (TOA). We assume the data contain a signal whose frequency evolves in time according to $f(t) = f_0 + \dot{f}t$. In order to remove the frequency drift, we introduce a new time $\tilde{t}$, which is a function of the original time t, such that the frequency as a function of $\tilde{t}$ is constant. Equivalently, we require the integrated phase to be linear in $\tilde{t}$

$$\phi(t) - \phi_0 = \int_0^t f(t')\, dt' = f_0 \tilde{t}. \tag{2.9}$$

The i^{th} TOA t_i (measured from $t_0 = 0$) will therefore be corrected to

$$\tilde{t}_i = t_i + \frac{1}{2}\frac{\dot{f}}{f_0} t_i^2. \tag{2.10}$$

A time series is constructed by dividing the total duration of the data T into N bins. The entire array is initially set to zero, then for each TOA, the value in the corresponding time bin is increased by one. Before being corrected, the i^{th} TOA would fall into bin $b_i = t_i/\Delta t$, where $\Delta t = T/N$, the resolution of the time series. We can then express the frequency derivative correction in terms of bin number

$$\tilde{b}_i = b_i - \frac{b_i^2}{a^2}, \tag{2.11}$$

where the parameter a is given by

$$a = \sqrt{\frac{2f_0}{\Delta t|\dot{f}|}} \tag{2.12}$$

(note that $\dot{f} < 0$).

The spacing of the $\dot{f}$ trials is chosen so that the fourth harmonic of a Crab-like pulsar will drift by no more than two power spectrum bins over the course of the entire observation. This results in a maximum trial spacing of $\Delta\dot{f}_{trial} = 1/T^2$. We use a fiducial $f_0 = 30$ Hz in Equation 2.12 and search up to $\dot{f}_0 = -3.7 \times 10^{-10}$ Hz s^{-1}. The number of $\dot{f}$ trials required for a given source is then

$$N_{\dot{f}} = \frac{|\dot{f}_0|}{\Delta\dot{f}_{trial}} = 3.7 \times 10^{-10} T^2. \tag{2.13}$$

The resulting numbers are listed in Table 2.1.

FFT and Initial Candidates

Having prepared a TOA list and generated a list of trial frequency derivatives for a given source, we proceeded with the first stage of the actual search. For each $\dot{f}$ trial, we constructed a time series as per section 2.2.2. We used this time series to calculate an FFT, from which we calculated the power spectrum. We then searched this power spectrum for significant peaks (single harmonics and harmonic sums). As a whole, this step was the most computationally intensive part of the analysis, involving the calculation of almost 10^4 billion-point FFTs, including all the sources searched. The gigapoint FFTs were calculated in-core using the 512-processor Intel Touchstone Delta supercomputer.

If a single power spectrum bin contained power that exceeded a predetermined threshold, then the corresponding f and $\dot{f}$ values were saved for further analysis. The FFT is most sensitive to source frequencies which are exactly equal to the discrete Fourier frequencies (integer multiples of $1/T$). When the source frequency is not equal to a Fourier frequency, the signal power can be spread over several power

spectrum bins, and the single bin peak power can be reduced by almost 60%. Even in this worst case, however, 80% of the power will still be in the two bins closest to the source frequency. For this reason, and also to allow for some frequency drift due to $\dot{f}$ error, we also saved candidates that showed significant power in a sum over neighboring spectral bins. Since gamma-ray and radio pulsar signals tend to have rich harmonic content, we computed two- and four-harmonic power sums, again saving the best candidates. Neighboring bins were also included in the harmonic sums.

The end result of this stage of the search was a list of $(f, \dot{f})$ candidates corresponding to the top one-, two-, and four-harmonic powers. The minimum power thresholds were set very low, resulting in $\sim 10^3$ candidates of each type from each source. Obviously, virtually all of these candidates are expected to be due to random noise fluctuations.

Post-FFT Analysis

Due to the small number of photons in a data set ($N_t \sim 10^3$), small sections of power spectra can be calculated with minimal computational effort. The raw (unnormalized) power in bin k is given by

$$P_k^{\mathrm{raw}} = \left| \sum_{j=0}^{N-1} x_j e^{-2\pi i j k/N} \right|^2 , \tag{2.14}$$

where $\boldsymbol{x}$ is the spin-down-corrected time series described above and $N = 2^{30}$ is the total number of points in the time series. Since only N_t of the N time series points are non-zero, we can reduce the sum over N to a sum over N_t. If we globally normalize the power spectrum to a mean value of 1, Equation 2.14 becomes

$$P_k = \frac{1}{N_t - N_t^2/N} \left| \sum_{j=0}^{N_t-1} e^{-2\pi i \tilde{b}_j k/N} \right|^2 , \tag{2.15}$$

where $\tilde{b}_j$ is the corrected bin number of the j^{th} photon. This requires only $\sim N_t \sim 10^3$ floating point calculations per frequency bin (and very little memory). It is therefore

a simple matter to calculate sections of power spectra in the vicinity of the FFT candidates on a workstation computer. We take advantage of this fact to refine our candidate list.

In the neighborhood of each candidate, we searched a local $f - \dot{f}$ phase space with higher resolution. Compared with the FFT search, we reduce the spacing of the $\dot{f}$ trials by a factor of 8, and by oversampling the power spectrum, we reduce the spacing of the frequency trials by the same factor. The result is a significant decrease in the spin-down induced frequency drift and the power loss due to the discrete Fourier sampling, obviating the neighboring bin sums. For a real signal, roughly the same signal power is recovered, but in fewer power spectrum bins, dramatically increasing the detection significance. An example illustrating the effectiveness of this method is given below in section 2.2.4.

Final Verification Procedures

At this stage, we attempted to confirm or refute the strongest candidate detections by means of several final tests. One such test was to look for spectral power in higher harmonics. Since the search only included sums of up to four harmonics, this method was used primarily to check for power in the 5th through 8th harmonics of four-harmonic candidates. Note that this is a statistically independent test. Since the candidates resulted from a search involving many trials (over a trillion different four-harmonic sums were calculated in the FFT stage of the analysis), a large four-harmonic power was required for a candidate to survive the search stage. Checking the next four harmonics involves only one trial for each of the final candidates (of which there were $\sim 10^2$), so a little excess power can be highly significant. Of course, candidates were not rejected simply on the basis of this test since the typically broad, multicomponent waveforms of gamma-ray pulsars may not exhibit significant power in the higher harmonics.

For each candidate, we also folded a pulse profile. This is essentially a histogram of photon counts as a function of pulse phase. The pulse phase of a photon is equal to its arrival time modulo the putative pulse period. Pulse profiles (also known as

light curves) can be usefully analyzed "by eye" for comparison with known pulsars or possible emission models. They can also be used to calculate a statistical significance by determining the probability that a given profile could have resulted from a flat distribution.

Another powerful verification technique involved photon weighting. In the initial search, photons were selected from an EGRET observation as described in section 2.2.2. All photons that made the cut in Equation 2.2 were given equal weight in the construction of the time series and calculation of the power spectrum. In reality, the EGRET point spread function (PSF) is not shaped like a step; the larger a photon's angular separation from a source, the less likely it is to have originated from that source. Using the derived EGRET (source-free) background model and the measured PSF, we applied weights to the photons in a data set and recalculated a candidate's spectral power. Testing this method on real and simulated EGRET pulsars showed that signal power could be increased by $15-70\%$. The smaller gains came from using the exact same photons as were originally used. The larger gains came from using more photons — either extending the angular cutoff out to where the PSF is essentially zero, or extending to lower energies (< 100 MeV), or both. At the power levels of our candidates, even a 20% gain can increase a candidate's significance by several orders of magnitude.

The above confirmation procedures can be extremely useful in weeding out noise fluctuations and increasing the confidence of real signals. But perhaps the most conclusive verification of a candidate would be to detect the same pulsation in a different observation of the same source. In practice, this can be very difficult. The measured frequency and frequency derivatives are subject to small but non-trivial error. When we also take the position uncertainty into account and propagate forward or backward to the second observation, the $f-\dot{f}$ phase space we must search is considerable. Of course this phase space is small compared to the original search space, but it is not insignificant. If the two observations have comparable source count rates and are not widely separated in time, then verification may be possible. But even a slight reduction in source strength and a separation of several months can

make the source signal impossible to distinguish from noise in the second observation.

2.2.3 Magnetar Search

In addition to the search for Crab-like pulsars that we have been describing so far, we also ran a separate search over a phase space more suited to detecting magnetars. Although they are not in the Princeton catalog, we have indicated six known or suspected magnetars in Figure 2.1. The soft gamma repeaters SGR 1806-20 and SGR 1900+14 and the anomalous X-ray pulsars 1E 1841-045, 1E 1048-5937, 4U 0142+615, and 1E 2259+586 are marked with Xs in the plot (Baring & Harding, 1998 and references therein). (Note however that in some cases the $\dot{f}$ errors are larger than the plotted symbols.) They are clearly distinguished from the "normal" pulsars by their low rotation frequencies and comparatively large spin down rates. This combination implies a large surface magnetic dipole field and gives the magnetars their name. For this search, we chose to cover the lighter shaded region in Figure 2.1: frequencies from 0.01 to 1.0 Hz and $-\dot{f}/f < 1 \times 10^{-9}$ s^{-1}. This region includes magnetic fields up to approximately 2×10^{15} G at $f = 1$ Hz and 2×10^{17} G at $f = 0.01$ Hz. The search was technically very similar to the Crab-like search, with several exceptions which we now describe.

There are indications that the long-term frequency evolution of magnetars may not be well described by a simple low order expansion, i.e., by the lowest order frequency derivatives (Kouveliotou et al., 1999; Melatos, 1999). For this reason, we chose to shorten our two longest data sets, as compared with the Crab-like search, and to include $\ddot{f}$ trials in our frequency drift corrections (up to $|\ddot{f}| < 3 \times 10^{-18}$ Hz/s^2, both positive and negative values). Since the frequency range we wanted to cover was smaller, the FFTs could be smaller, but the number of required spin down corrections was much higher. Treating each order derivative separately, we decided to limit the effect of each so that the fourth harmonic of any signal in the search phase space would drift by no more than the equivalent of one power spectrum bin. This necessitated a trial $\dot{f}$ spacing of $\Delta\dot{f}_{trial} = 1/2T^2$ and a trial $\ddot{f}$ spacing of $\Delta\ddot{f}_{trial} = 1/T^3$. For a

14-day data set, this amounts to over 3×10^4 spin down trials (per source).

To reduce this number, we decided to utilize incoherently stacked power spectra (e.g., see Brady & Creighton, 1998). The continuous data set is divided into a number of smaller sections and a power spectrum is calculated for each. The separate spectra are then added together. If we divide the data into $\mathcal{S}$ sections, the independent Fourier spacing is increased by a factor of $\mathcal{S}$. We can therefore tolerate more frequency drift and the total number of drift trials goes down by a factor of $\mathcal{S}^2$ (one factor of $\mathcal{S}$ for each of the two search parameters $\dot{f}$ and $\ddot{f}$). For this search, using $\mathcal{S} = 4$ stacks for all sources brought the number of trials down to an acceptable value, with only about a 20% reduction in sensitivity as compared with the coherent (single stack) method.

We calculated the power spectra with an oversampling factor of 2. Neighboring bin sums were not used for the two- and four-harmonic candidates. The single harmonic candidates included significant single bin powers and neighboring bin sums. These sums were, of course, of *independent* bins, not truly neighboring bins, which are not independent in an oversampled spectrum. Good candidates were put through the same kinds of final analyses as described above for the Crab-like search. Note that the $\ddot{f}$ search absorbs the effect of the position error, so we need not consider reduced source strengths when calculating the sensitivity of the magnetar search.

2.2.4 Testing Our Method on Geminga

Our method easily detects the known strong EGRET pulsars (Crab, Vela, and Geminga). The FFT stage alone detects each of them with better than $S \sim 10^{-50}$ significance (S is essentially the probability that the observed spectral power was produced by random noise fluctuations). Figure 2.2 shows a section of the Geminga power spectrum calculated from Viewing Period 1.0 data prepared as described above. Only spectral bins with power $P > 7$ are plotted. The three highest peaks correspond to the first, second, and fourth harmonics of Geminga's 4.217 675 1 Hz rotation frequency. Most of the power is in the second harmonic. This is understandable since Geminga's pulse profile, shown below in the figure, has two peaks separated by about 180° in phase.

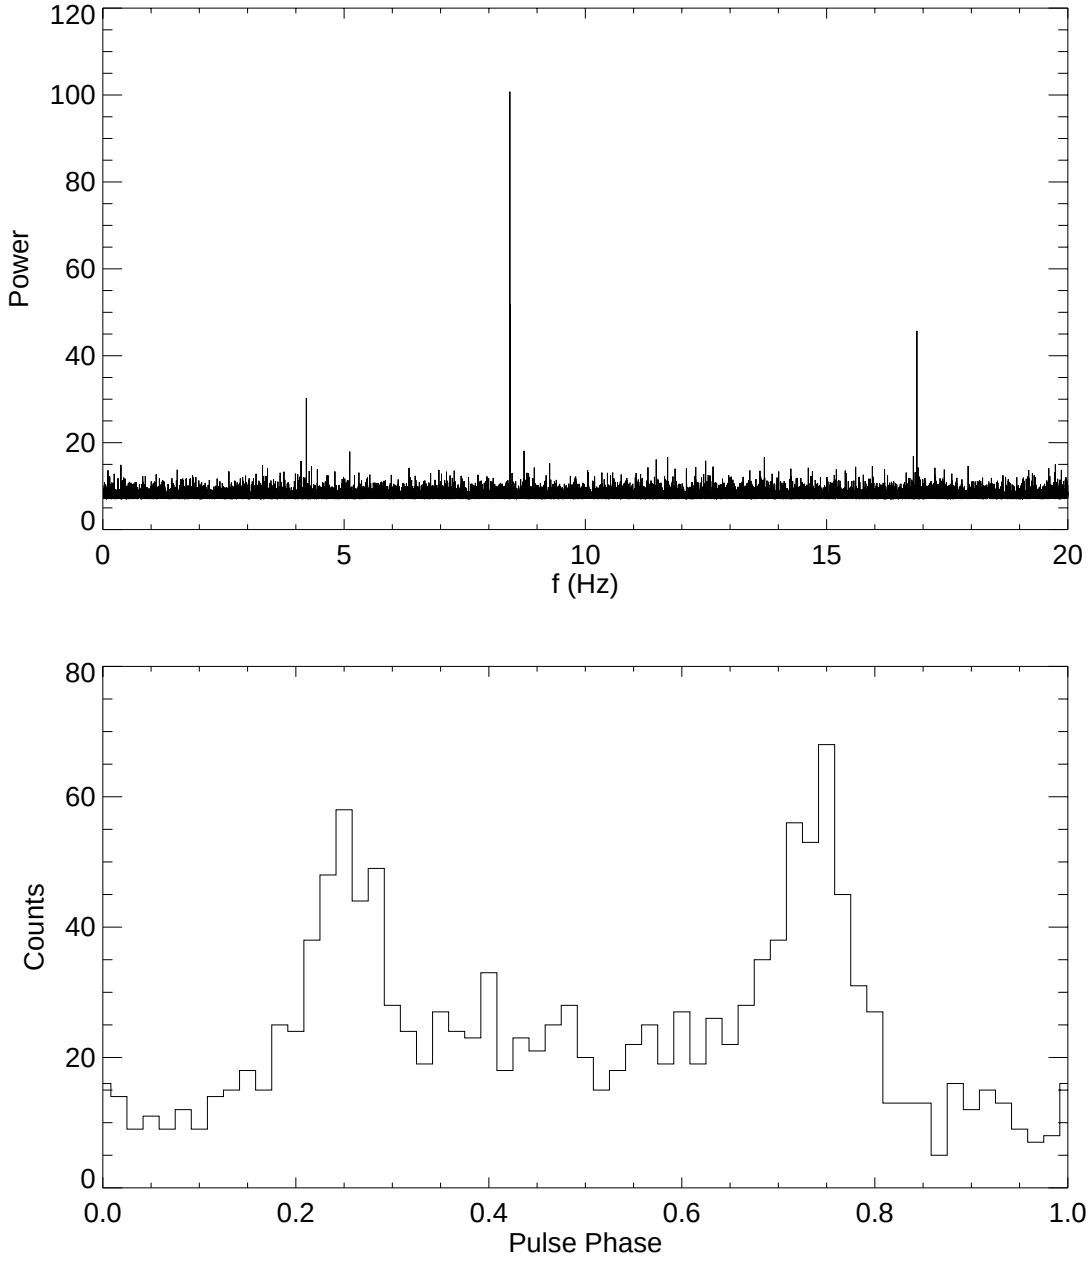

Figure 2.2: Section of the power spectrum from the EGRET VP 1.0 observation of Geminga, along with the corresponding pulse profile.

We can simulate a weaker source by selecting a subset of the data. If we include only every n^{th} photon, the source strength N_s^2/N_t scales down by a factor of $1/n$. We now illustrate the main steps of our search method using a data set consisting of every 6^{th} VP 1.0 Geminga photon. The data set contained 232 photons, 126 of which are expected to have actually come from Geminga, yielding $N_s^2/N_t = 68$. This is comparable to our strongest unidentified source.

The pulsar shows up in the FFT candidate list with a four-harmonic power of 47.51. Note that this is a sum of power from 12 spectral bins, 3 consecutive bins at each of the 4 harmonics. We did not run an entire set of frequency derivative trials; Geminga's small $\dot{f}$ is actually covered by the first $\dot{f}$ trial. Even so, it was not the highest candidate in the list. Had we run all 540 trials, we would have expected to see about 20 noise candidates with at least this much power. Still, this power was well above the threshold we used in the actual searches (our four-harmonic (twelve-bin) power cutoff was 44, which should produce $\sim 10^3$ noise candidates per source). So at this stage, the pulsar signal has made our candidate list, but is certainly not significant.

When we refine the search space in the vicinity of the FFT candidates, the pulsar signal emerges above the noise. After the local oversampling analysis, the weakened Geminga has a four-harmonic, four-bin power of 40.88. This was by far the best candidate remaining. If we consider only the full Geminga search, the candidate now has a significance of $S \sim 0.09$. If we take all the searches into account, the significance is only $S \sim 0.9$, but this is well within the range of candidates that were subjected to the final verification procedures. In fact, this would have been our second best search candidate at this point.

Although the power in the next 4 harmonics is not significant, the profile alone looks good enough to warrant attempting a search in other observations. But we can convincingly verify the detection in this data set by PSF-weighting the photons. With the weighted data, the four-harmonic spectral power increases to 55.56, with significance $S \sim 10^{-6}$, taking all the search trials into account. The main steps of this analysis are summarized in Table 2.2.

Table 2.2: Test results for weakened Geminga data (see text).

	After FFT	After Local Analysis	After PSF Weighting
4-Harmonic Power	47.51 (12 bins)	40.88 (4 bins)	55.56 (4 bins)
Significance	~ 1.0	~ 0.9	$\sim 10^{-6}$

2.3 Results

No candidate from any source passed all of our detection criteria. In section 2.3.1 below, we describe the statistics used to determine the significance of a power spectrum candidate. We then discuss the selection of our final candidate list and the application of our final verification procedures. Section 2.3.2 describes how source strength and waveform affect spectral power. This information is used in section 2.3.3 to place limits on the pulsed fraction and duty cycle for each source. In section 2.3.4 we describe our best candidate and how it was debunked.

2.3.1 Significance

In the absence of a signal, the power P in a given spectral bin follows a χ^2 distribution. More precisely, since our individual spectra were normalized to a mean value of one, $2P$ is χ^2 distributed with 2 degrees of freedom. If P_m is a sum of powers from m independent frequency bins, then $2P_m$ is χ^2 distributed with $2m$ degrees of freedom. The probability p that a power sum P_m will exceed some given threshold P_0 is

$$p(P_m > P_0) = e^{-P_0} \sum_{j=0}^{m-1} \frac{P_0^j}{j!}. \tag{2.16}$$

Our searches of four-harmonic power sums provided our best sensitivity, and we consider only four-harmonic candidates from now on.

The above probability applies to an individual statistical trial. To determine the significance of a candidate, we must account for all the trials in the search — all the

fundamental frequencies (up to 40 Hz) in all of the power spectra (one for each $\dot{f}$ trial) for all of the EGRET sources searched. Corresponding to each fundamental frequency bin, there are 9 different four-harmonic sums, and after the local oversampling analysis, we include two additional factors of 8, accounting for the increases in frequency and frequency derivative resolution. Thus, by the final verification stage, we had effectively searched

$$N_{\rm trials} \approx 4 \times 10^{14} \quad \text{(Crab-like search space)} \tag{2.17}$$

four-harmonic trials. For the magnetar search, we had

$$N_{\rm trials} \approx 5 \times 10^{11} \quad \text{(Magnetar search space).} \tag{2.18}$$

Note that these four-harmonic sums were not all truly statistically independent, so we are overestimating $N_{\rm trials}$.

The significance of a candidate with power P_0 (ie., the probability that its power was produced by noise) is therefore

$$S \approx 1 - [1 - p(P_m > P_0)]^{N_{\rm trials}} \approx N_{\rm trials} \cdot p(P_m > P_0), \tag{2.19}$$

where the last approximation holds for large P_0. For the Crab-like searches $m = 4$, while for the magnetar searches $m = 4\mathcal{S} = 16$.

Unfortunately, even our best search candidates had significance $S \sim 1$. All four-harmonic, four-bin (Crab search) candidates with $P_0 \geq 36$ were subjected to further testing. These "candidates" were not at all statistically significant; rather they represented only the expected tail of the noise power distribution. Sixty-four independent $(f, \dot{f})$ candidates met this criterion. Most appeared in the candidate list several times, with various nearby values of f or $\dot{f}$, for a total of 582 final candidates. For each candidate, we calculated the power in the next four harmonics and analyzed the pulse profile. One candidate stood out after these analyses, and we describe it in some detail below (section 2.3.4). All the candidates were then re-analyzed using

PSF-weighted data, and in all cases the powers dropped, which is unequivocally fatal. For the magnetar search, the cutoff power was $P_0 = 53$ (again, corresponding to the highest expected noise powers, $S \sim 1$), resulting in a total of 224 candidates (56 independent), all of which were similarly eliminated.

2.3.2 Relating Spectral Power to Source Properties

For a given point source, the selected data set contains a total of N_t photons, N_s of which were emitted by the source. We assume that the source has some DC component and that only a fraction $\mathcal{F}$ of these source photons actually contribute to the pulsation. Using the arrival times of these photons, we calculate a power spectrum and normalize it as described above. The expected spectral power in a single harmonic is

$$\langle P \rangle \approx 1 + \alpha \mathcal{F}^2 \frac{N_s^2}{N_t} \tag{2.20}$$

(Buccheri et al., 1987). The parameter α ranges from 0 to 1 as the waveform's duty cycle decreases from 100% (no pulsation) to 0% (δ-function). For a sinusoidal signal, $\alpha = 1/4$ (see for example van der Klis, 1989). If we sum the powers from h harmonics, we expect

$$\langle P_h \rangle \approx h + \alpha_h \mathcal{F}^2 \frac{N_s^2}{N_t}. \tag{2.21}$$

Note that for a sinusoid, $\alpha_h = 1/4$ for all values of h since there is no signal power in higher harmonics. For a pure δ-function waveform, in which all the photons are emitted with the same phase, $\alpha_h = h$. In all cases, $0 \leq \alpha_h \leq h$.

The source strength N_s^2/N_t is estimated for each source *a priori* and the minimum detectable power threshold is set by the specifics of the search (e.g., for our Crab search, we can say that none of the data sets contain a pulsar signal with a four-harmonic power exceeding $P_{\rm max} = 36$). We can therefore place constraints on the pulsed fraction and the duty cycle of any pulsar signal contained in the data (provided it lies within our search phase space).

2.3.3 Upper Limits

If we wish to place upper limits on the source parameters with better than 50% confidence, we cannot simply solve Equation 2.21 for $\alpha\mathcal{F}^2$. The spectral power produced by a periodic source will vary depending on particular samplings of the TOAs. Conversely, a measured spectral power can conceivably result from a wide range of intrinsic source parameters. We must allow for this intrinsic statistical variation when calculating high confidence limits on the source parameters. We must also consider the fact that our search covers only a discrete grid of frequencies and frequency derivatives and generally will not recover the total signal power available in a data set. For these reasons, we determined our upper limits using Monte Carlo simulations. We calculated these limits at a confidence level of 95%.

Rather than simply report upper limits on the somewhat arcane combination $\alpha\mathcal{F}^2$, we chose a fiducial value for each parameter to constrain the other. Although it was certainly possible for us to have detected pulsations with realistic gamma-ray pulsar waveforms, we are unable to place useful upper limits on such signals. For most of our sources, we can only rule out pulsations at the 95% confidence level for waveforms that would have been easier to detect. The upper limits we report for the pulsed fraction are for a waveform typical of radio pulsars. Specifically, we assumed a single narrow Gaussian peak, with a full width at half maximum (FWHM) equal to 2.87% of the pulse period. This is the median radio pulse width quoted in the Princeton catalog (Taylor et al., 1993). To limit the duty cycle, we assumed a pulsed fraction of $\mathcal{F} = 100\%$, and we report lower limits on the FWHM of a single-peaked Gaussian waveform. The Crab-like search results are shown in Table 2.1; the magnetar search results are in Table 2.3.

2.3.4 Best Candidate

The source 2EG J1835+5919 produced our best candidate. Its largest four-harmonic, twelve-bin power (after the FFT stage) was 55.98 and its largest four-harmonic, four-bin power (after the local oversampling stage) was 38.67, neither of which was

Table 2.3: Magnetar search results.

2EG Source	N_s^2/N_t	$\dot{f}$ trials	$\ddot{f}$ trials	total trials	$\alpha\mathcal{F}^2$	pulse FWHM[a]	pulsed fraction[b]
J2020+4026	73	1161	6	6966	< 0.75	$> 27\%$	$< 44\%$
J0618+2234[c]	32	728	3	2184	< 1.71	$> 15\%$	$< 67\%$
J2019+3719[d]	34	1161	6	6966	< 1.61	$> 16\%$	$< 65\%$
J1835+5919	28	730	3	2190	< 1.96	$> 13\%$	$< 71\%$
J1021-5835[d]	27	708	3	2124	< 2.03	$> 13\%$	$< 73\%$

[a] FWHM of a single Gaussian peak, expressed as a percentage of the pulse period; assumes 100% pulsed fraction.
[b] For typical radio pulsar waveform (see text).
[c] VP 1.0 only.
[d] The third EGRET catalog (Hartman et al., 1999) split 2EG J2019+3719 and 2EG J1021-5835 into multiple sources. The upper limits reported assume the validity of the second EGRET catalog (Thompson et al., 1995) (see section 2.2.2).

particularly significant, though the candidate was strong enough to make our final list. In fact, it appeared in the list a total of 27 times, with frequencies in the range $f = [32.967\,559\,8, 32.967\,560\,2]$ Hz and frequency derivatives in the range $\dot{f} = [-1.348\,3 \times 10^{-10}, -1.342\,7 \times 10^{-10}]$ Hz s^{-1}.

It was the power in the next four harmonics that first distinguished this candidate. Its best entry yielded a 5^{th}-8^{th} harmonic power of 19.86, which has significance better than $S \sim 0.002$. (Here, we have treated all 582 final candidates as independent, which they are not, so this estimate of the significance is conservative.) The pulse profile, shown in Figure 2.3, also looked very promising. The signal not only has rich harmonic content, but as we can see from the profile, those harmonics have similar phases; i.e., they have conspired to produce a single narrow peak.

We attempted a search for this signal in EGRET Viewing Periods 201-203, but did not find anything significant. This is not at all conclusive, however. As described above, even for a real pulsar, this verification method is not expected to have a high success rate. The final test was the most decisive. PSF-weighting the photons reduced the four-harmonic power from 38.67 to 24.01 (the power in the next four harmonics dropped from 19.86 down to 10.28), completely eliminating this as a pulsar candidate.

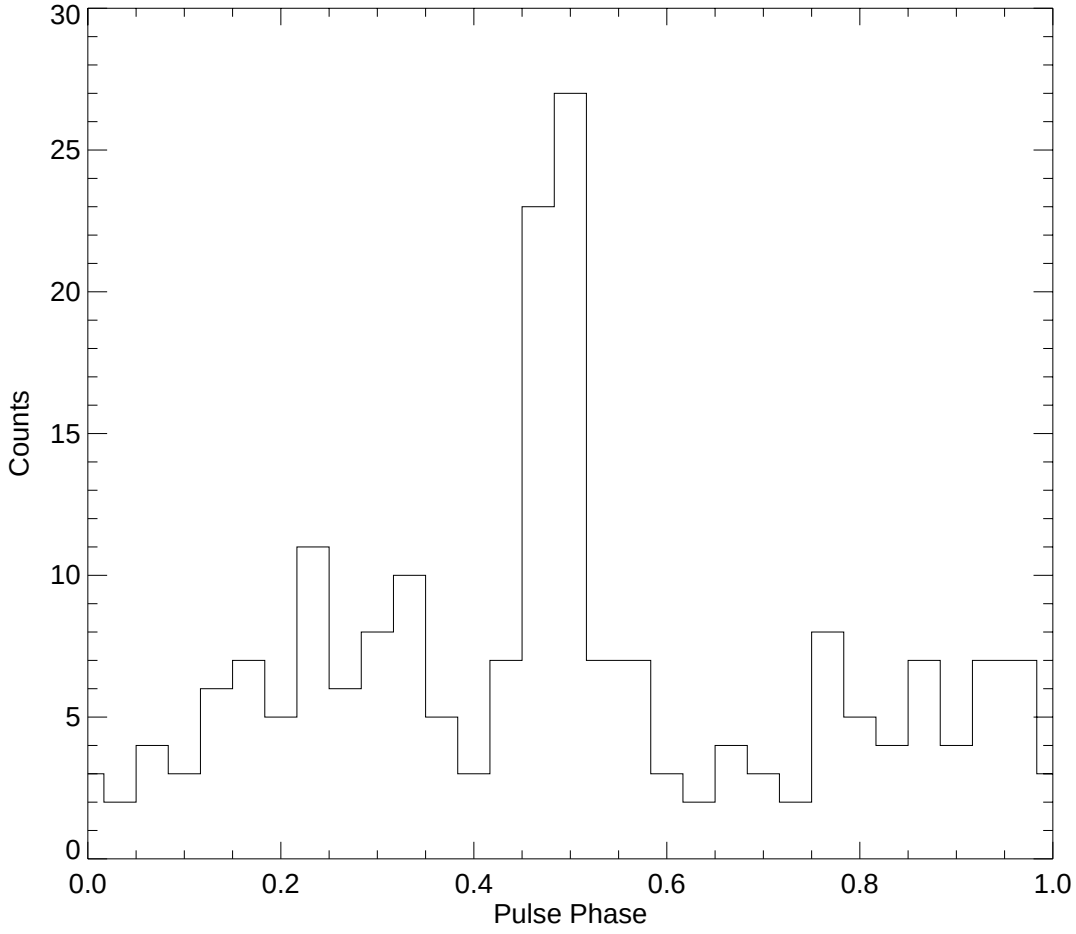

Figure 2.3: Best gamma-ray search candidate. The pulse profile shown is from 2EG J1835+5919, VP 212, with $f = 32.967\,560\,053$ Hz and $\dot{f} = -1.346\,3 \times 10^{-10}$ Hz s^{-1}. It was eliminated after PSF-weighting drastically degraded its significance.

2.4 Discussion

We knew from the start that our probability of detecting a new pulsar was not terribly high. With such a small number of known gamma-ray pulsars, however, even a single new identification would have been an important discovery. Unfortunately our search detected no new gamma-ray pulsars, and we report here only upper limits.

At first glance, our upper limits on the pulsed fraction for each source (for the Crab-like search; Table 2.1), ranging from 49% to 81%, may not seem very stringent. Analysis of the Crab, Vela, and Geminga pulsars, however, reveals that all of these sources are essentially 100% pulsed in gamma-rays. (It is noteworthy that the outer gap model of pulsar gamma-ray emission also predicts a pulsed fraction of 100%.) It is therefore not unreasonable to expect most (if not all) of the photons from a source to contribute to any pulsation.

The waveform limits are a bit more troublesome. Our search was reasonably sensitive to signals with sharply peaked waveforms, such as those produced by typical radio pulsars. The pulse shapes of the known strong gamma-ray pulsars, however, exhibit dual narrow peaks, separated in phase by about 140°–180°, with some emission (a "bridge") between the two peaks (Fierro, 1996). The Geminga profile shown in Figure 2.2 is a good example. For our purposes, these waveforms can be characterized by the parameter α (see Eq. 2.21). By analyzing the power in the first four harmonics of the Crab, Vela, and Geminga pulsars we determine $\alpha_{\rm C} = 0.93$, $\alpha_{\rm V} = 0.64$, and $\alpha_{\rm G} = 0.41$, respectively. These values are to be compared with the "$\alpha\mathcal{F}^2$" column in Table 2.1, with $\mathcal{F} = 100\%$.

Thus, with very nearly 95% confidence, we would have detected a pulsar with a Crab-like waveform in 2EG J2020+4026 (100% pulsed). We cannot make such a strong statement about the Vela or Geminga waveforms, nor can we make any such claim for the weaker EGRET sources we searched. That is not to say that we had no chance of detecting pulsars with Vela-like or Geminga-like waveforms. We simply cannot claim with a high degree of confidence that we would have found them.

Direct timing of the EGRET data has not resulted in any new identifications.

It now seems likely that further identifications of gamma-ray sources will come only from ongoing searches in other wavelengths and from future gamma-ray missions. In particular, the new generation of X-ray satellites (*CXO* and *XMM*) may resolve new counterparts to EGRET sources. Looking further ahead, the next generation gamma-ray satellite *GLAST* promises to improve on EGRET's source localization and sensitivity by at least an order of magnitude. The improved source localization will make multi-wavelength correlations easier and more reliable, while the increased effective area will greatly facilitate blind pulsar searches. McLaughlin & Cordes (2000) predict that 140 new gamma-ray pulsars will be detectable in blind searches with *GLAST* data, and the techniques described in this paper will be directly applicable to those searches.

A. M. C. would like to thank Stuart B. Anderson for helpful discussions. Access to the Intel Touchstone Delta computer system was provided by the Caltech Center for Advanced Computing Research. Basic research in X-ray astronomy at the Naval Research Laboratory is supported by the Office of Naval Research. This work was supported by NASA grants NAG5-2833, NAGW-4761, and NAG5-3384.

Chapter 3

No Persistent Pulsations in Aql X-1 as it Fades into Quiescence

Abstract

We searched for coherent X-ray pulsations from Aql X-1 in a series of *RXTE* observations taken shortly after a recent outburst. During the course of these observations, Aql X-1 passes through an apparent "propeller" phase as its luminosity fades to its quiescent value. No pulsations were detected, and we place upper limits (ranging from 0.52% to 9.0%) on the fractional RMS amplitude of any periodic signal contained in the various data sets searched. This result has implications for the geometry of the system, if the quiescent luminosity is due to continued low-level accretion. Alternatively, our result supports the idea that the quiescent luminosity may be due to thermal emission.

(Originally appeared as Chandler & Rutledge, 2000, ApJ, 545, 1000.)

3.1 Introduction

The origin of the quiescent emission observed from low-magnetic field transient neutron stars (NSs) has posed something of a puzzle. When first discovered observationally (in the transient Cen X-4; van Paradijs et al., 1987), the X-ray emission, fit with a blackbody spectrum, was too faint to have originated over the entire surface of a 10 km NS. It was suggested that the emission was due to accretion over a fraction ($\sim$ few%) of the NS surface. As additional quiescent NS X-ray spectra were observed and fit with a blackbody spectrum (Verbunt et al., 1994; Asai et al., 1996a,b), the small implied emission areas seemed to confirm this view.

The cause of accretion anisotropy over the surface of a low B-field NS has been purported to be the "propeller effect" (Illarionov & Sunyaev, 1975; Stella et al., 1986), in which at low accretion rates, the magnetic field increases in importance relative to the gravitational field in determining the accretion geometry, perhaps expelling much of the accretion from the system when the effect of the magnetic field is comparable to or stronger than that of the gravitational field at the radius where the keplerian frequency is equal to the spin frequency of the NS. When this occurs, some fraction of the accretion may follow the magnetic field lines to the magnetic poles which, if these are offset from the rotational poles, could conceivably produce X-ray pulsations. The efficiency of such a propeller has been estimated recently (Menou et al., 1999). However, some observations of decreasing torque with increasing luminosity may indicate deviations from this standard magnetic accretion scenario in low mass X-ray binaries (Chakrabarty et al., 1997). If, after an outburst, accretion continues into quiescence then pulsations might be expected, particularly if the NS can be shown to have gone through a "propeller phase," in which the accretion geometry in the vicinity of the NS is dominated by the magnetic field, rather than by the gravitational field. No pulsations from these objects in quiescence have yet been reported, although they have been searched for (Verbunt et al., 1994; Asai et al., 1996b).

Thus, until recently, the view of the quiescent luminosity of low-B transient NSs

has been as due to continued accretion over a fraction of the surface of the NS, perhaps caused by modified accretion geometry due to the effects of the magnetic field, which may therefore give rise to pulsations in the X-ray intensity.

A different view of the origin of the quiescent luminosity has recently been described (Brown et al., 1998), in which nuclear reactions at the base of the NS crust keep the NS core heated to temperatures ($\approx 10^8$ K) sufficient to explain a large fraction, if not all, of the quiescent luminosity of these objects (10^{32-33} erg s^{-1}); in this manner, the quiescent luminosity can be produced in the absence of active accretion. The discrepancy in the emission area can be explained as due to an incorrect blackbody assumption – specifically, if accretion is shut off, metals stratify in the NS atmosphere (Alcock & Illarionov, 1980; Romani, 1987), producing a pure H photosphere, in which the free-free opacity, which decreases with increasing frequency, dominates, permitting higher E photons to originate at greater depth, thus at higher temperatures (Rajagopal & Romani, 1996; Zavlin et al., 1996). This produces an emergent spectrum which is spectrally harder than a simple blackbody, and – when described as a simple blackbody – requires a higher T, and subsequently lower emission area, to parameterize it. Spectral fits with H atmosphere spectra show that the observed X-ray spectra from the three such objects for which data exists (Cen X-4, Aql X-1, and 4U 1608-522) imply emission areas consistent with a 10 km NS (Rutledge et al., 1999). Thus, both their quiescent luminosities and their emergent X-ray spectra can be explained in this manner, without invoking active accretion during quiescence. An attractive feature of this explanation is that it accounts for the similar luminosities of these three objects, without requiring their quiescent mass accretion rates (a product of the propeller efficiency and mass accretion rate through the accretion disk) to be serendipitously similar.

It remains an open question if accretion (at a level of $\dot{M} \lesssim 10^{-12} M_\odot$ yr^{-1}) occurs at all, contributing some fraction of the quiescent X-ray luminosity. There are few observational means by which this may be investigated. One such possibility is to search for metal lines in the quiescent X-ray spectrum, which, if present, would imply that metals are being fed into the atmosphere at a rate greater than gravity can

stratify it (Bildsten et al., 1992; see Zavlin et al., 1996 for X-ray spectra of metallic atmospheres). Such spectra can be investigated by the latest generation of X-ray spectroscopy missions (*AXAF* and *XMM*). A second method is to search for intensity variability during quiescence. In particular, as described above, if pulsational variability at a particular frequency were to be found in quiescence, this would support the scenario of active accretion along the magnetic field lines to the NS surface. Alternatively, if the NS magnetic field were strong enough to significantly affect the photospheric opacity ($> 10^{10}$ G; cf. Zavlin et al., 1995), this would produce apparent "hot spots" near the magnetic poles of the NS which, if offset from the rotation axis, could give rise to pulsations; however, it is thought that type-I X-ray bursting sources have magnetic fields below this value. Clearly discovery of pulsations or constraints on the pulsed fraction during X-ray quiescence (following the onset of a "propeller") would provide useful limits for models of quiescent NS emission.

The system Aql X-1 lends itself to such a search for pulsations. During the X-ray decline of a recent outburst, the intensity at first decreased by one order of magnitude over 17 days, then abruptly (beginning at $L_x \approx 10^{36}$ erg s^{-1}) by 3 orders of magnitude over 3 days, which was coupled with the sudden spectral hardening of the X-ray intensity (Campana et al., 1998; Zhang et al., 1998a). The X-ray intensity stopped dropping at a luminosity $L_x \sim 10^{33}$ erg s^{-1}, comparable to the quiescent luminosity observed previously (Verbunt et al., 1994). This was interpreted as evidence for a propeller-phase, though this remains to be confirmed through repeated observations of Aql X-1 in outburst. If indeed this behavior does indicate a propeller phase in Aql X-1, it offers the opportunity to search for pulsations during the period when the NS magnetic field substantially alters the accretion geometry.

Here we have searched for, and not found, pulsations in Aql X-1 during and immediately following its supposed propeller phase. In section 3.2, we describe our search method and the results of our analysis. In section 3.3 we discuss these results and our conclusions.

3.2 Analysis

3.2.1 Search Methodology

We used high time resolution data taken with *RXTE*/PCA (Swank et al., 1996). The PCA detectors have a total geometric area of 6500 cm^2, and a nominal energy range of 2-60 keV. We selected for analysis four observations, based on the results of Zhang et al. (1998b). These were listed by Zhang et al. as observations 8, 10, 11, plus a later observation, 12, which was not analyzed by Zhang due to the faintness of the source – which made their spectral analysis infeasible, but can still be useful in our search for pulsations. Based on Zhang's results, the beginning of the propeller-phase appears to be between observations 9 and 10, and so our timing analyses of observations 10-12 are expected to reveal pulsations created by the magnetic-field modified accretion geometry (if any), while observation 8 is included simply for comparison purposes while the source was moderately brighter. We list details of the observations in Table 3.1. Data were obtained from the XTE archive at HEASARC.

For our data analysis, we used the standard FTOOLS/XTE v4.0. Using standard-2 formatted data, we extracted the X-ray spectrum from each observation and subtracted background counts (estimated using pcabackest tool), and examined the counts spectra. Based on the spectrum of the faintest observation (number 12), which becomes background dominated above $\sim 12\,\mathrm{keV}$, we selected for our analysis PHA bins 0-20, corresponding in energy roughly to 2-12 keV. We then extracted individual photon events from data in E125 mode (which has time resolution of $\sim 125\mu s$), and corrected the photon arrival times (time of arrival; TOAs) to the solar system

Table 3.1: *RXTE* Aql X-1 analyzed observations

Observation	Start Time / End Time (UT)	Start Time (JD)	Start Time (MET)
20098-03-08-00	1997-03-01 21:33 1997-03-02 00:21	2450509.398	99869572
20098-03-10-00	1997-03-05 22:55 1997-03-06 01:28	2450513.456	100220227
20098-03-11-00	1997-03-08 21:31 1997-03-09 01:11	2450516.394	100474022
20098-03-12-00	1997-03-10 21:25 1997-03-10 23:55	2450518.389	100646431

barycenter.

Zhang et al. (1998b) reported strong evidence for a 549 Hz pulsation frequency, observed during $\sim$ 10 seconds following a type-I X-ray burst from Aql X-1. (It is possible that this may be the second harmonic of the NS spin frequency.) Similar pulsations have been observed from type-I bursts of 5 other sources (1728-34, Strohmayer et al., 1996; 1636-53, Zhang et al., 1996; 1731-260, Smith et al., 1997; a Galactic center source, possibly MXB 1743-29, Strohmayer et al., 1997; and 1702-43, Markwardt et al., 1999). Rather than restrict our efforts to a small range of frequencies around the previously measured value, we chose to do a more general search for pulsations over the broad frequency range $0.5 - 1024$ Hz. This range includes all previously reported rotation frequencies for type-I X-ray bursters. We were able to do this more general search without a significant loss of sensitivity for reasons which we explain below, in section 3.2.4.

The Aql X-1 system consists of a NS and a companion star, both orbiting the system's center of mass with a period of about 19 hours (see section 3.2.2 below). As a result of this orbital motion, the NS is accelerating along our line of sight. Any pulsations will therefore be observed with a doppler shifted frequency which is not constant in time. In a standard fast Fourier transform (FFT) analysis, the spectral power resulting from such a signal may be spread out over many frequency bins, drastically decreasing the probability of detection. For example, consider a 550 Hz pulsar in a 19 hour circular orbit with a maximum projected velocity of $100\,\mathrm{km\,s^{-1}}$. If we calculate a coherent power spectrum using $T = 2048$ seconds of data, the power from this pulsar will be spread over as many as 70 independent frequency bins (for the worst case orbital phase, the observed pulsar frequency covers a range $\Delta f = 0.034\,\mathrm{Hz}$, giving $T\Delta f = 70$ bins). To counteract this effect, we attempt to remove the acceleration of the pulsar signal in the time domain before performing the FFT. Because we do not have complete knowledge of Aql X-1's orbital parameters, we must cover the possible orbital parameter space with a number of acceleration trials and repeat the Fourier analysis for each.

For our purposes, the orbit (assumed circular) can be characterized by three in-

dependent parameters — the orbital period $P_{\rm orb}$, the projected circular velocity v, and the orbital phase x_0 of the NS at the start of the observation in question. The projected circular velocity is the magnitude of the NS's circular velocity projected along our line of sight, $v = v_{\rm NS} \sin i$, where i is the angle between the line of sight and the normal to the plane of the orbit. Note that we could equally well have chosen the projected orbital radius ($a \sin i = v P_{\rm orb}/2\pi$) in lieu of this parameter. The orbital phase x_0 is measured in cycles, and is therefore in the range $[0, 1]$, with $x_0 = 0$ corresponding to longitude = 0 (where the observed pulsar frequency would be measured at its minimum). The signal from a pulsar in such an orbit would be observed with a doppler shifted frequency

$$f(t) = f_0 \left[1 - \frac{v}{c} \cos \left(\frac{2\pi}{P_{\rm orb}} t + 2\pi x_0 \right) \right], \tag{3.1}$$

where f_0 is the rest frequency of the pulsar, c is the speed of light, and t is the elapsed time since phase x_0.

Pulsar orbital acceleration searches are typically carried out (see, for example, Anderson et al., 1990) by correcting data with an assumed constant acceleration $f(t) = f_0 + \dot{f} t$. This is applicable when the data cover only a small part of the pulsar's orbit, or the pulsar is very strong. In the first case, the short span of orbit is well approximated by a constant acceleration, while in the second case, some spreading of the feature in the power spectrum can be tolerated without the signal disappearing into the noise. For the Aql X-1 searches reported here, we assumed that neither of these conditions was satisfied. Our acceleration searches therefore fully correct for an assumed circular orbit $(P_{\rm orb}, v, x_0)$. As compared with the constant acceleration method, this circular acceleration method may require more trial accelerations to cover a given orbital phase space, but the detection significance is greatly increased. This is because the circular method recovers significantly more signal power in a single power spectrum bin, more than making up for the increased number of trials (the detection significance is exponential in recovered power, but only linear in the number of trials).

After preparing an energy selected, barycentered photon TOA list, we proceeded with the search method as follows. First, we assumed particular values for the orbital parameters (i.e., a trial acceleration) from within the search phase space (of orbital period, velocity, and initial phase; see Table 3.2). We corrected the TOAs for the assumed acceleration by introducing a corrected time $\tilde{t}$, which is a function of the original time t, such that the frequency as a function of $\tilde{t}$ is constant. Equivalently, we require the integrated phase to be linear in $\tilde{t}$:

$$x(t) - x_0 = \int_0^t f(t')dt' = f_0\tilde{t}. \tag{3.2}$$

Integrating Equation (3.1) for a given acceleration trial $(P_{\rm orb}, v, x_0)$, we see that the i^{th} TOA t_i is corrected to:

$$\tilde{t_i} = t_i + \frac{\beta}{\omega}\left[\sin(2\pi x_0) - \sin(\omega t_i + 2\pi x_0)\right], \tag{3.3}$$

where $\beta = v/c$ and $\omega = 2\pi/P_{\rm orb}$.

These corrected TOAs were used to construct a time series which was then FFTed and used to produce an estimate of the power density spectrum (PDS). The PDS was searched for candidates (frequency bins containing statistically significant excess power). The process was repeated for each acceleration trial. The spacing of acceleration trials in (orbital period, velocity, and phase)-parameter space was chosen to allow sensitivity to as weak a signal as possible while keeping the computational requirements reasonable.

3.2.2 Determination of Searched Parameter Space

Similar, but inconsistent, values for the orbital period of Aql X-1 have been published by two groups. Based on observations of Aql X-1 during outburst, Chevalier & Ilovaisky (1991) measured the orbital period to be 18.97±0.02 hr. Shahbaz et al. (1998) determined the orbital period in quiescence to be 19.30±0.05 hr. In a more recent analysis, Chevalier & Ilovaisky (1998) report the period to be 18.9479±0.0002 hr,

again measured during outburst, but they also report a quiescent period within 0.02% of this value. To be safe, we chose to cover a range of orbital periods that encompasses a 3σ range in all of these measurements (as well as the periods in between):

$$18.91 < P_{\rm orb} < 19.45 \text{ hr}. \tag{3.4}$$

Fortunately, we were able to cover this entire range with one trial value of $P_{\rm orb}$ (see section 3.2.3 below).

The projected circular velocity v has not been measured. We can determine an upper limit on v for our search by assuming a value for the NS mass, $m_{\rm NS} = 1.4\,M_\odot$ ($M_\odot$ = one solar mass), and choosing a maximum companion mass to which we will be sensitive. We can then use our knowledge of the orbital period to calculate the orbital velocity of the NS. Recently, the true optical counterpart of Aql X-1 has been identified as a late K type star (Chevalier et al., 1999). Again, to be safe we decided to cover a range of velocities corresponding to a companion star mass as high as $1.0\,M_\odot$ (for all possible orbital inclination angles). This results in a search range of

$$0 \leq v \leq 130\,\text{km}\,\text{s}^{-1}. \tag{3.5}$$

Note that our search was also sensitive to larger companion masses in a restricted range of inclination angles; e.g. our search covered companion masses up to $2.0\,M_\odot$ for $0^\circ \leq i \leq 40^\circ$.

For the starting orbital phase for each observation, rather than rely on a particular ephemeris, we simply chose to search the full range

$$0 \leq x_0 \leq 1 \text{ cycle}. \tag{3.6}$$

Our search phase space is summarized in Table 3.2.

Table 3.2: Aql X-1 searched parameter space

Parameter	Parameter Space
$P_{\rm orb}$	$18.91 - 19.45$ hr
$v_{\rm NS} \sin i$	$0 - 130\,{\rm km\,s^{-1}}$
Initial Orbital Phase	$0 - 1$ cycles
Pulsar Spin Frequency	$0.5 - 1024$ Hz

3.2.3 Searches Performed

The *RXTE* observations included in this analysis were roughly $9 - 13$ kiloseconds in duration. The TOA data in each observation are broken up into two or three cohesive sections, separated by gaps due to earth occultations. Each continuous section is at least 2048 seconds long, and we chose to study these continuous blocks individually. Observation 10, with one occultation drop-out, is therefore divided into two sections, which we call observations 10 a and 10 b. Observations 11 and 12 are similarly divided (into 11 a, 11 b, 11 c, 12 a, and 12 b). We analyzed only the first section of observation 8, for a total of eight separate data sets. The photons from a given data set were binned into a 2^{23} point time series, with a time resolution of $\Delta t = 244\,\mu$sec. Over the course of the observations, the source count rate decreases from about 670 counts s^{-1} at the beginning of observation 8 to less than 10 counts s^{-1} by the end of observation 12, with a background rate of about $31 - 37$ counts s^{-1} (see Table 3.3).

Our search method was described above in section 3.2.1. We now describe in detail the method used to determine the grid of trial values for the orbital parameters. Since we must discretize a continuous phase space, there will likely be some offset between a signal's actual parameters and our nearest trial values. The effect of this offset will be some spreading out of the signal power in the power spectrum, since the time dependence of the signal's frequency will not have been completely removed. In choosing the orbital trials, the general idea was to space them just finely enough to keep the remaining frequency drift within tolerable limits (to be quantified below).

We began analytically, and then made some empirical adjustments. Considering each of the three orbital parameters ($P_{\rm orb}, v, x_0$) separately, we took partial derivatives of Equation (3.1), to determine the frequency drift that results from offsets in the individual parameters. For example, an error in the velocity ($\partial v = c\partial\beta$) causes a drift of

$$\partial f = f_0 \cos\left(\omega t + 2\pi x_0\right) \partial\beta. \tag{3.7}$$

We are interested only in the extrema of the frequency range covered by the signal over the course of the observation time T, and for simplicity, we chose to make our velocity trial values independent of the other parameters. We therefore replace the cosine factor in Equation (3.7) by its maximum change over the course of an observation ($T = 2048$ s), and we replace f_0 by our maximum search frequency ($f_0 \rightarrow 1024$ Hz). The maximum frequency drift caused by a velocity offset is therefore

$$\partial f_{\rm max} = 1024\ 0.19\ \partial\beta\ ({\rm Hz}). \tag{3.8}$$

A "tolerable" drift must be no greater than the independent frequency resolution of the power spectrum, F. Since we are considering each orbital parameter separately, for these initial calculations we restrict the error-induced drift from each to $F/2$. This is somewhat arbitrary, and the actual spacings used were decided on with the help of simulations. Note that the allowable spacing of the trials can be twice the maximum allowable offset. We now have for the velocity trial spacing

$$\Delta v_{\rm trial} \lesssim \frac{cF}{1024\ 0.19} = 0.77\ \left(\frac{F}{5\times10^{-4}\ {\rm Hz}}\right)\ {\rm km\,s^{-1}}. \tag{3.9}$$

For a coherent FFT of 2048 seconds of data, the independent Fourier resolution is $F = 1/T = 4.88\times10^{-4}$ Hz. The search strategy that we settled on actually used a larger frequency resolution (see below); for now we will leave the expressions for the trial spacings in terms of F explicitly.

Again, for simplicity, we chose to keep the spacing of the $P_{\rm orb}$ and x_0 trials independent of $P_{\rm orb}$ and x_0, but we did allow for velocity dependence. For the orbital

period, we ultimately find

$$\Delta P_{\rm trial} \lesssim 2800 \left(\frac{F}{5\times10^{-4}\ {\rm Hz}}\right) \left(\frac{v}{100\,{\rm km\,s^{-1}}}\right)^{-1} \text{ seconds.} \tag{3.10}$$

Even for the smallest Fourier resolution ($F = 4.88\times10^{-4}$ Hz) and the largest velocity in our search range ($v = 130\,{\rm km\,s^{-1}}$), the orbital period trial spacing turns out to be larger than our target search range of (19.45 hr − 18.91 hr) 3600 s hr^{-1} = 1944 s. In other words, a single trial value was sufficient to cover our entire search range in $P_{\rm orb}$. Thus, *we did not actually search over trial values of the orbital period.* For the initial orbital phase trial spacing, we find

$$\Delta x_{\rm trial} \lesssim 1.2\times10^{-3} \left(\frac{F}{5\times10^{-4}\ {\rm Hz}}\right) \left(\frac{v}{100\,{\rm km\,s^{-1}}}\right)^{-1} \text{ cycles.} \tag{3.11}$$

If we were to use these spacings with coherent 2048 second FFTs, each of the eight data sets would require over 95 000 trial accelerations. This would require $\gtrsim$ 100 days of CPU time on a Sun Ultra 10 workstation. To reduce the computational requirements, we chose to utilize incoherently stacked power spectra. The original data of duration T are divided into $\mathcal{S}$ sections of duration $T/\mathcal{S}$. Each section is FFTed individually and the $\mathcal{S}$ individual power spectra are (incoherently) added together. The independent Fourier step size has been increased by a factor of $\mathcal{S}$ ($F = 1/(T/\mathcal{S}) = \mathcal{S}/T$) and the trial parameter spacings increase by the same factor. The total number of trials is therefore reduced by a factor of $\mathcal{S}^2$ (since we are searching over two parameters). The reduction in the number of trials comes at the cost of reduced sensitivity, so the number of stacks $\mathcal{S}$ should be kept as small as possible, just barely bringing the number of trials to a feasible level. For our analyses, we chose to use $\mathcal{S} = 4$.

With the above analytical calculations as guides, we used simulations to determine the trial spacings actually used in the search. Our simulations upheld the decision to use a single $P_{\rm orb}$ trial. Since there are indications that the orbital period is more likely to lie towards the lower end of our search range (18.91 − 19.45 hr), we chose to

use $P_{\rm orb} = 19.00$ hr in our trials. We decided to use a velocity trial spacing of

$$\Delta v_{\rm trial} = 2.0\,{\rm km\,s^{-1}}. \tag{3.12}$$

This is about 1.5× finer than the spacing calculated above (recall that we are using $\mathcal{S} = 4$ stacks). For the orbital phase, we decided on a trial spacing of

$$\Delta x_{\rm trial} = 0.011\ \left(\frac{v}{100\,{\rm km\,s^{-1}}}\right)^{-1}\ {\rm cycles}, \tag{3.13}$$

which is about 2.3× coarser than the spacing indicated in the analytical calculation. For small velocities, we used a minimum of 16 x_0 trials (except for the single $v = 0$ trial).

Using these spacings with 2× oversampled power spectra (exactly like our actual searches), we determined that at least 98% of the time, we were able to recover at least 77% of a simulated signal's power in a single spectral bin, even in the least favorable regions of our search phase space. We simulated signals whose orbital parameters were offset from our search trials and whose pulsation frequencies were offset from our discrete Fourier frequencies. The reduction in recovered signal power is due to a combination of these factors. (Note that it is mere coincidence that a simple FFT recovers, on average, 77% of a signal's power based solely on offset from the discrete Fourier frequencies. In a 2× oversampled spectrum, the *minimum* power recovered, based solely on frequency offset, is 81%.) Our peak detected power was not always in the frequency bin closest to the rest frequency of the simulated pulsar. This has no effect on the detection of pulsations; the true frequency and orbit of the pulsar can be refined after the initial detection.

In total, we covered the orbital phase space with 4069 acceleration trials. The same accelerations were used for each of the eight data sets.

3.2.4 Estimation of Detection Sensitivities

To characterize the sensitivity of our search, we wish to place quantitative limits on the minimum signal strength required for a significant detection. Since the noise statistics of our power spectra are well understood, it is a simple matter to determine the detection threshold, the minimum spectral power $P_{\rm det}$ required for a detection. This information can be used to determine the minimum required signal strength. For a review of detection thresholds, detection sensitivities, and upper limits, see Vaughan et al. (1994).

Our stacked power spectra were constructed by summing four individual spectra, each normalized to a mean power of one (note that this convention differs from the commonly used Leahy normalization (Leahy et al., 1983), for which the mean noise power is two). In the absence of a signal, the power P in a given spectral bin follows a χ^2 distribution. Specifically, $2P$ is χ^2 distributed with $2\mathcal{S} = 8$ degrees of freedom. The probability that the power P in a single PDS bin will exceed a given value P_0 is therefore

$$p(P > P_0) = e^{-P_0}\left(1 + P_0 + \frac{1}{2}P_0^2 + \frac{1}{6}P_0^3\right). \tag{3.14}$$

Considering each frequency bin of each PDS to be one search "trial," our entire search consisted of $N_{\rm trials} = 3.4 \times 10^{10}$ such trials. This is the product of the number of observations searched (8), the number of acceleration trials for each observation (4069), and the number of frequencies searched in each power spectrum (2^{20}). The latter number is equal to the frequency range searched (1024 Hz) divided by the independent frequency resolution ($F = (\mathcal{S}/2048)$ Hz) times the oversampling factor (2). Of course, neighboring bins in an oversampled spectrum are not independent, and neighboring acceleration trials may not produce truly independent spectra. Thus $N_{\rm trials}$ is an upper limit to the number of statistically independent trials in the search. Since we are overestimating the number of independent trials, our detection threshold will be conservative.

The statistical significance S of a measured power P_0 is equal to the probability that the power was produced by a random noise fluctuation. For small p (large

powers), the significance is

$$S = N_{\rm trials}\, p(P > P_0). \tag{3.15}$$

To achieve our target significance of 10^{-4}, we require a PDS power of at least

$$P_{\rm det} = 43. \tag{3.16}$$

For a given data set containing N_t total counts, N_s of which were emitted by the source, we can relate the source strength to the expected spectral power by

$$\langle P \rangle = 4 + \frac{1}{4}\frac{N_s^2}{N_t}\mathcal{F}^2 \tag{3.17}$$

(Buccheri et al., 1987; van der Klis, 1989), where $\mathcal{F}$ is the pulsed fraction, i.e. the fraction of source counts that actually contribute to the pulsation. The first term in Equation (3.17) is the expected noise power ($\langle P_N \rangle = 4$ because we have summed four power spectra each normalized to unity), while the second term is the expected signal power. Here, we have assumed that the signal waveform is sinusoidal. N_s and N_t are determined for each data set before beginning the pulsation search. Therefore, the sensitivity of each search is determined as a limit on the pulsed fraction $\mathcal{F}$.

We now calculate the detection sensitivities for the eight data sets searched. For a given observation, the detection sensitivity is expressed as the minimum pulsed fraction required to produce a spectral power exceeding $P_{\rm det}$ with high confidence. Instead of simply solving Equation (3.17) for $\mathcal{F}$ using the detection threshold power $P_{\rm det} = 43$, we must allow for statistical variation of the spectral power produced by a source of a given strength, as well as variation in the recovered power due to the discrete nature of the search trials.

To determine the detection sensitivity, we must consider the probability distribution of power in a spectral bin containing a signal plus noise. If, in the absence of noise, the signal produces a power $P_{\rm sig}$, then the power P in the presence of noise will

be distributed according to

$$p_n(P > P_0; P_{\rm sig}) = \exp[-(P_0 + P_{\rm sig})] \sum_{m=0}^{\infty} \sum_{k=0}^{m+n-1} P_0^k P_{\rm sig}^m / (k!m!) \qquad (3.18)$$

(Groth, 1975), where n is the number of independent bins that were summed to produce P (i.e., in our case $n = \mathcal{S} = 4$). Thus, $p_n(P > P_0; P_{\rm sig})$ is the probability that the power will exceed P_0 in the spectral bin containing the signal. It is important to note that noise power and signal power are not simply additive; P will not always exceed $P_{\rm sig}$ (see Vaughan et al. (1994) for a discussion of this point). For a given N_s and N_t, we would like to find the pulsed fraction $\mathcal{F}$ that is 95% likely to produce power $P > P_{\rm det}$.

To ensure that our reported sensitivity is as conservative as possible, we will not naively invert Equation 3.18 for $P_{\rm sig}$ with $P_0 = P_{\rm det}$. Instead we consider the worst-case scenario covered by our search parameter space. Due to the discrete binning of photons in the construction of the time series, the recovered signal power in the FFT falls off with increasing frequency. A signal near our maximum search frequency — 1024 Hz, which is half the Nyquist frequency of our power spectra — will produce 81% of the power that would be recovered from a low frequency signal of the same intrinsic strength. Also, due to our discrete grid of acceleration trials and frequency trials, we are only 98% likely to recover more than 77% of a signal's available power. Thus, in our worst case, we are 98% likely to recover at least $(0.77)(0.81)P_{\rm sig} = 0.62P_{\rm sig}$.

We account for this reduction in recovered signal power approximately by solving Equation 3.18 in the form

$$p_n(P > P_{\rm det}; 0.62P_{\rm sig}) = 0.97. \qquad (3.19)$$

The probability of our recovered signal power exceeding $0.62P_{\rm sig}$ *and* our total spectral power P exceeding $P_{\rm det}$ is then approximately $(0.98)(0.97) \approx 0.95$, giving our desired 95% confidence.

Solving Equation 3.19 (numerically) results in $P_{\rm sig} = 93.6$. For a given observa-

tion, the minimum required pulsed fraction is then given by $\mathcal{F} = (4N_t P_{\rm sig}/N_s^2)^{1/2}$ and converted to RMS by dividing by $2\sqrt{2}$. Thus, our 95% confidence detection sensitivities for observations 8 a, 10 a, 10 b, 11 a, 11 b, 11 c, 12 a, and 12 b are given by 0.60%, 1.2%, 1.2%, 2.7%, 2.9%, 3.0%, 5.6%, and 10%, respectively. We verified these numbers with Monte Carlo simulations.

Our search included spin frequencies up to 1024 Hz. Had we restricted our search to a smaller range of frequencies, corresponding to the previously reported pulsations, our detection sensitivities would not have been drastically different. To achieve the same detection significance, with far fewer trials, we require only a slightly reduced spectral power. And since the pulsed fraction limit depends essentially on the square root of the required power, the detection sensitivities are not terribly sensitive to the number of frequency trials in the search. Given Fox's best estimate for the pulsation frequency $549.76^{+0.05}_{-0.03}$ Hz (D. Fox 1999, private communication), if we had confined our search to a 5σ range (about 549.76 Hz and 274.88 Hz), our detection sensitivity in observation 10 a would have improved to 1.0%, while in observation 12 b the limit would have been 8.8%.

3.2.5 Verification of Procedures

Simulated data were used to test our search codes and to verify the sensitivity of the search. As an additional check, we applied our acceleration method to SAX J1808.4-3658 — the accreting millisecond pulsar (Wijnands & van der Klis, 1998; Chakrabarty & Morgan, 1998). Since our search was for a similar pulsar, with $\sim$ few hundred Hz frequency in a $\sim$ few hour orbit about a low-mass companion, SAX J1808.4-3658 provides a useful test of our acceleration method.

We analyzed an *RXTE*/PCA observation (observation #30411-01-01-04). We selected energy PHA bins corresponding to 0.4-17 keV from data of the same type we use in our search (E125). We barycenter corrected the arrival times with the position given by Roche et al. (1998), and used 2048 seconds of data beginning at Mission Elapsed Time (MET, seconds since 1994.0) 135395082.

The pulsations in these data are easily detected — the ~ 401 Hz signal is obvious in an unaccelerated PDS, with significant power detected in each of about 30 adjacent frequency bins, using 2048 seconds of data (Figure 3.1, dotted line). The highest single bin power represents about 22% of the total signal power in the extended feature.

Using the known orbital parameters (Chakrabarty & Morgan, 1998) to remove the acceleration from the data, the spreading of the signal power is reduced to about 6 independent Fourier bins around the correct rest frequency of the pulsar, 400.975 210 6(8) Hz, with $\sim 85\%$ of the total power in a single spectral bin (Figure 3.1, solid line). Oversampling reveals that the spreading that remains is consistent with the expected sinc-function response of the (discrete) FFT to a signal with the pulsar's rest frequency (van der Klis, 1989). Thus, we confirm that our technique successfully removes the effect of the doppler shift due to the pulsar's orbital motion.

3.2.6 Search Results

No candidate pulsations from any of the Aql X-1 observations exceeded our predetermined detection threshold. Using our largest detected search power, we can calculate upper limits on the strength of any pulsar signal contained in the various data sets searched (Vaughan et al., 1994), provided its orbit and spin frequency are covered by our search phase space. These upper limits are more restrictive than the *a priori* detection sensitivities.

The upper limit calculation is essentially the same as the detection sensitivity calculation, with the detection threshold power $P_{\rm det}$ replaced by the maximum observed power $P_{\rm max} = 30.96$. The results are shown in Table 3.3. For example, with 95% confidence, we can say that there is no sinusoidal pulsar signal in observation 10 a with a fractional RMS amplitude greater than 1.0%. The limits we can place on the fractional RMS become less stringent as the source flux decreases, increasing to an upper limit of 9.0% during observation 12 b, when Aql X-1 was observed to be faintest ($\sim$1 mCrab; 13 000 PCA counts $\rm s^{-1}$=1 Crab).

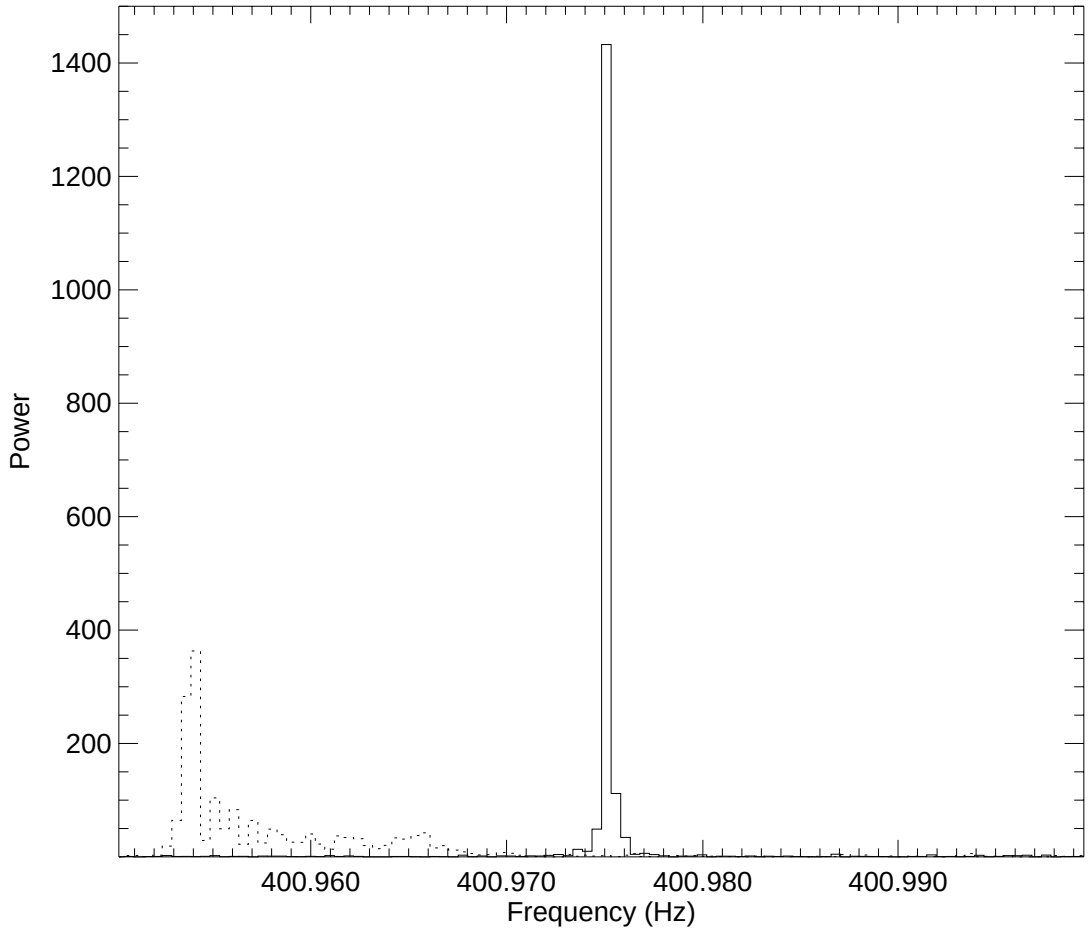

Figure 3.1: SAX J1808.4-3658. The dotted line shows a section of the unaccelerated power spectrum. The spreading of the signal power due to the orbital doppler shift is evident. The solid line shows the power spectrum after removal of the doppler shift using our acceleration technique. The signal power has not been confined to a single bin because the pulsar frequency lies between two discrete Fourier frequencies. During this observation, SAX J1808.4-3658 passes through longitude = 0. The peak in the uncorrected spectrum is therefore maximally offset from the pulsar's rest frequency, $\Delta f_{\max} = \beta f_0 = 0.022\,\mathrm{Hz}$.

Table 3.3: Aql X-1 observation parameters and search results

Observation	MET Analyzed	Source Count Rate[a]	Background Count Rate[a,b]	Fractional RMS Amplitude[c]
8 a	99869572-98871620	670.6±0.8	37.3	$< 0.52\%$
10 a	100220227-100222275	188.2±0.6	35.8	$< 1.0\%$
10 b	100226227-100228275	177.4±0.6	35.8	$< 1.1\%$
11 a	100474022-100476070	52.6±0.5	35.7	$< 2.3\%$
11 b	100479522-100481570	48.6±0.5	35.7	$< 2.5\%$
11 c	100485082-100487130	45.5±0.5	35.7	$< 2.6\%$
12 a	100646431-100648479	19.0±0.5	31.2	$< 4.9\%$
12 b	100652431-100654479	9.2±0.5	31.2	$< 9.0\%$

[a] counts s^{-1} (13 000 counts s^{-1} = 1 Crab)
[b] Systematic uncertainty ~ 0.5 counts s^{-1}
[c] Upper limits are 95% confidence

3.3 Discussion

We find no evidence of pulsations from Aql X-1 as it fades into quiescence. The upper limits on the pulsed fraction for the eight data sets searched range from 0.52% to 9.0% RMS. If, indeed, during the period identified as the "propeller phase" (Campana et al., 1998; Zhang et al., 1998a) the NS magnetic field significantly modifies the accretion flow geometry in the vicinity of the NS, the apparent absence of pulsations does not support the hypothesis that the quiescent emission is due to continued accretion. Our results cannot completely rule out such accretion, however, since it is possible that the geometry of the system may not lead to detectable pulsations; for example, the magnetic axis may be very nearly aligned with the rotation axis, or the rotation axis may point directly towards the earth. The non-detection of pulsations coupled with the observation of a thermal spectrum during quiescence (Rutledge et al., 1999) favors the interpretation that the quiescent luminosity is not due to accretion, but rather to a hot NS core (Brown et al., 1998).

We thank Stuart Anderson and Thomas Prince for helpful discussions. We also thank Lars Bildsten, Sergio Campana, and an anonymous referee for valuable comments on the manuscript. This work was supported by NASA Grant #NAG5-3239.

Chapter 4

Discovery of 5 New Pulsars in Caltech-Arecibo Drift Survey

Abstract

We report the discovery of five new pulsars in the latest Caltech-Arecibo drift survey. Nineteen previously known pulsars, including two millisecond pulsars (MSPs), were re-detected in our 840 square degree search area. Among the new discoveries are PSR J0627+07, an apparent orthogonal rotator, and PSR J1938+22, an unusual bursting radio pulsar. Our search resulted in no new fast pulsars, despite having millijansky sensitivity well below 1 ms. This is the most sensitive large area survey to date for sub-millisecond pulsars, and we place a 95% confidence upper limit on the sub-MSP surface density of one per 280 square degrees, to a limiting flux density of ~ 2 mJy for 0.5 ms pulsars at 430 MHz.

(Chandler, A. M., Anderson, S. B., Kulkarni, S. R., and Prince, T. A.,
to be submitted to the Astrophysical Journal)

4.1 Introduction

The earliest pulsar surveys were sensitive to pulsars with periods $P \gtrsim 30$ ms (e.g., Large & Vaughan, 1971; Davies et al., 1972; Davies et al., 1973; Hulse & Taylor, 1974). It quickly became clear that the population of slow pulsars was highly concentrated along the Galactic plane, and many subsequent large-area pulsar surveys have been restricted to low Galactic latitudes (e.g., Camilo et al., 1997). Improvements in computer power and data recording speeds have allowed faster sampling rates, leading to the discovery of millisecond pulsars (MSPs) (Backer et al., 1982). Compared to the long-period pulsars, MSPs have much weaker magnetic fields and therefore tend to be less luminous and have much longer active lifetimes. MSP space velocities appear to be smaller on average than long-period pulsar velocities (Toscano et al., 1999), but with such long lifetimes, they may travel far out of the Galactic plane. MSPs are therefore much more isotropically distributed on the sky than their relatively short-lived, longer-period counterparts (Johnston & Bailes, 1991).

A complete census of the local MSP population therefore requires high Galactic latitude surveys. These same surveys should also detect faint, nearby slow pulsars which may help fill in the low end of the pulsar luminosity distribution, crucial for modelling the total Galactic pulsar population.

In a modest ~ 150 deg^2 high latitude survey at Arecibo (Wolszczan, 1991; Wolszczan & Frail, 1992), two very interesting recycled pulsars were discovered — PSR B1534+12, a relativistic double neutron star binary, and PSR B1257+12, a pulsar with Earth-mass planets (the first planets discovered outside our solar system). A number of groups then undertook to survey the entire Arecibo sky for MSPs.

In this paper, we report the results of a large-area survey covering about 6% of the sky visible to the 305 m Arecibo telescope. This facility recently completed a lengthy upgrade period during which the high frequency feeds were largely inaccessible, and the ability to point the telescope was severely limited. However, with the 430 MHz line feed still operational, this was an ideal opportunity to conduct pulsar "drift" surveys of the Arecibo sky, in which the azimuth and zenith of the telescope are held

fixed while the sky simply drifts through the telescope beam at constant declination. Only infrequent pointing changes are required to adjust the declination.

For the purpose of these drift surveys, the sky visible to the Arecibo telescope was divided into one-degree declination strips and portioned out to five groups. To date, these surveys have detected 65 new pulsars, including 9 new millisecond pulsars (Foster et al., 1995; Camilo et al., 1996a; Ray et al., 1996; Xilouris et al., 2000; McLaughlin et al., 2000; Lommen et al., 2000). Here we report results from the latest Caltech survey, in which we have discovered 5 new slow pulsars and no new millisecond pulsars. The current survey is notable for its significant sensitivity ($\sim 1\,\mathrm{mJy}$) to pulsars with periods $< 1\,\mathrm{ms}$. Although no sub-millisecond pulsars were detected, this result places some of the first useful observational constraints on the low end of the pulsar period distribution. This has important implications for accretion scenarios and neutron star structure.

In section 4.2 below, we describe the details of the search including sky coverage, sensitivity, and analysis methods. In section 4.3 we report the results of our search and list the key parameters of the 5 newly discovered pulsars. Finally, in section 4.4 we provide more detail on two of the new pulsars, and discuss the overall implications of the survey.

4.2 Observations and Analysis

Our search observations were conducted between November 1996 and May 1998. The data were taken in an 8 MHz bandwidth about a central frequency of 430 MHz. The Penn State Pulsar Machine (PSPM) (Cadwell, 1997) was used to divide this full band into 128 separate frequency channels. Two orthogonal polarization signals were 4-bit digitized, after detection, and summed together. The PSPM search mode sampling rate was 80 μsec. Representing nearly 24 days of telescope time, the entire data set totaled about 2.5 TB.

Our analysis essentially followed standard radio pulsar search methods (Burns & Clark, 1969; Lyne & Graham-Smith, 1998). The data were divided into sections

("beams") corresponding to the time required for a source to transit the telescope beam, a Fourier transform was calculated for each of a number of dedispersion trials, and the resulting power spectra were searched for fundamental spin frequencies and harmonic sums that showed significant excess power. The top pulsar candidates were then re-observed at Arecibo for verification.

4.2.1 Sky Coverage

The Arecibo sky includes declinations from $-1°39'$ to $+38°21'$. The Caltech group was allotted one-degree declination strips centered on $\delta = 2°.5$, $7°.5$, $12°.5$, $17°.5$, $22°.5$, $27°.5$, and $32°.5$. The present survey covered the region shown in Figure 4.1. At 430 MHz, Arecibo's half power beam width (FWHM) is $10'$, so in principle a single (one degree) declination strip can be covered by six individual scans, separated by $\Delta\delta = 10'$.

In some regions of our search area, the separation between scans was less than $10'$, so some points on the sky passed through the central half power beam in more than one scan. This improves sensitivity to those declination regions that would otherwise be covered with only $\sim 50\%$ of peak sensitivity, and can facilitate the verification of pulsar candidates when a candidate is visible in multiple scans.

In total, this survey covered 840 independent square degrees, or 2.0% of the entire sky. About half of this area was covered by more than one scan.

In a previous search of 960 deg^2 taken from within these same declination strips, Ray et al. (1996) discovered 12 new pulsars. Since the AO upgrade drift scans were taken rather randomly, the current search overlaps the previous search area by about 300 deg^2 or 35%. As we show in the next section, however, our sensitivity was significantly better, particularly for millisecond pulsars.

4.2.2 Search Sensitivity

The sensitivity of a radio pulsar search is generally expressed as the minimum flux density (in mJy $= 10^{-29}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$) that the search can detect. This is a function

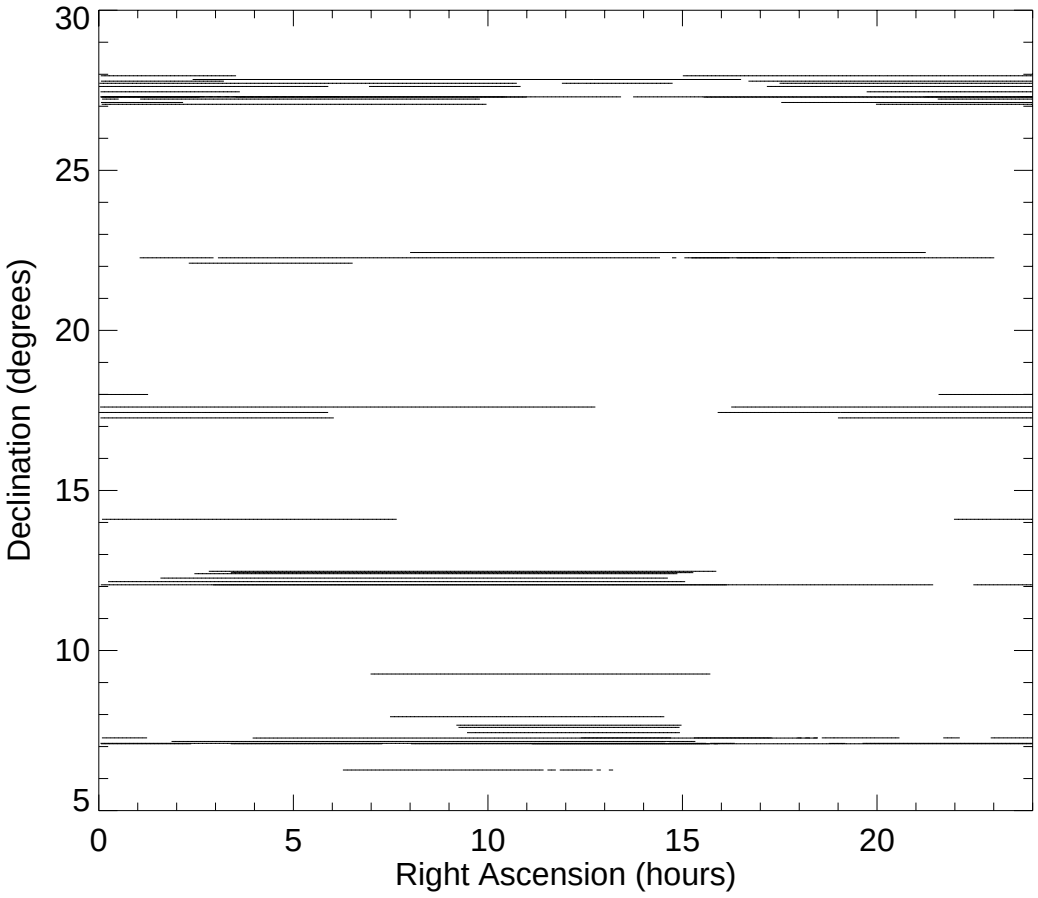

Figure 4.1: Sky coverage for this survey

of many factors, including the physical parameters of the telescope, receiver, and backend, as well as the properties of the pulsar itself, where it is located on the sky, and the length of the observation. A simple expression for the sensitivity is given by

$$S_{\rm min} = \alpha\beta \frac{T_{\rm sys}}{G(N_p B \tau_{\rm int})^{1/2}} \left(\frac{w_{\rm eff}}{P - w_{\rm eff}} \right)^{1/2} \tag{4.1}$$

(Dewey et al., 1985; Bhattacharya, 1998), where α is the threshold signal-to-noise ratio; $T_{\rm sys}$ is the sum of the receiver noise temperature, the sky temperature (a function of galactic latitude and longitude), and the spillover temperature (a function of zenith angle, due to increasing ground illumination with increasing zenith); G is the gain of the telescope in $\mathrm{K\,Jy^{-1}}$; N_p is the number of polarizations summed; B is the total bandwidth of the observation; $\tau_{\rm int}$ is the integration time; and P is the pulse period. The factor β accounts for various losses in the observing system and data processing. Based purely on the radiometer equation, β would be 1, but losses from, e.g., digitization and imperfect dedispersion degrade the theoretical sensitivity, resulting in $\beta > 1$. In practice, β is determined empirically by observing known calibration sources.

The effective pulse width $w_{\rm eff}$ is defined by

$$w_{\rm eff}^2 = w_0^2 + \tau_{\rm samp}^2 + \tau_{\rm DM}^2 + \tau_{\rm scat}^2, \tag{4.2}$$

where w_0 is the intrinsic pulse width; $\tau_{\rm samp}$ is the effective broadening of the pulse due to the anti-aliasing filter, and is of the same order as the sampling time

$$\tau_{\rm samp} \approx 2 \times 80\,\mu\mathrm{sec} = 0.160\,\mathrm{ms}; \tag{4.3}$$

$\tau_{\rm DM}$ is the dispersion smearing in a single filterbank channel

$$\tau_{\rm DM} = 8.3 \times 10^6 \frac{B_{\rm MHz}\,\mathrm{DM}}{\nu_{\rm MHz}^3 N_{\rm chan}} = 0.00652\,\mathrm{DM\,ms}, \tag{4.4}$$

where $\nu_{\rm MHz}$ is the central observing frequency in MHz and $N_{\rm chan}$ is the number of

filterbank channels; and $\tau_{\rm scat}$ is the broadening due to scattering

$$\tau_{\rm scat} \approx \left(\frac{1000}{\nu_{\rm MHz}}\right)^{4.4} 10^{-7.231+0.9255\log_{10}{\rm DM}+0.814613(\log_{10}{\rm DM})^2}\,{\rm ms} \qquad (4.5)$$

(Cordes et al., 1991).

For the current search, we adopt the values $T_{\rm sys} = 70\,{\rm K}$ and $G = 16\,{\rm K\,Jy^{-1}}$, appropriate for a high Galactic latitude source near zenith. Based on our observing hardware and analysis methods, we use $\alpha = 8$, $\beta = 1.5$, $N_p = 2$, $B = 8\,{\rm MHz}$, and $\tau_{\rm int} = 42\,{\rm s}$, and we assume an intrinsic pulse width $w_0 = 5\%$. Our search sensitivity as a function of pulse period is shown in Figure 4.2 for four values of the dispersion measure DM $= 0,\ 30,\ 100$, and $280\,{\rm cm^{-3}\,pc}$. The latter value is the largest DM trial used in our search. For comparison with other pulsar searches, we have kept with tradition and reported our search sensitivity assuming optimal conditions. For a pulsar in the Galactic plane, observed at a more extreme zenith angle, offset from the nominal declination of a scan by the half power beam radius, the minimum detectable flux can be an order of magnitude higher.

As is evident from Figure 4.2, our sensitivity to long-period pulsars was slightly better than that of the earlier generation of Arecibo drift surveys. The primary reason for this is the increased integration time used in the current survey. A point source at declination δ drifts through the center of the $10'$ half power telescope beam in $40/\cos\delta$ seconds. The optimal integration time turns out to be just slightly ($\sim 1\,{\rm sec}$) longer than this. Obviously, this hasn't changed since the previous drift surveys were conducted. The difference is due to the sampling rates used in the two types of surveys. Efficient fast Fourier transform (FFT) algorithms require that the number of data points be a power of 2. With a sampling rate of 250 μsec, the older Arecibo correlator data leads to 33 or 66 sec integrations (2^{17} or 2^{18} points). The PSPM search sample rate of 80 μsec allows for 42 sec integrations, much closer to optimal. Of course, data can always be resampled at a more advantageous rate, but this introduces correlations that change the statistics of calculated power spectra.

In contrast, our sensitivity to fast pulsars was quite a bit better than the ear-

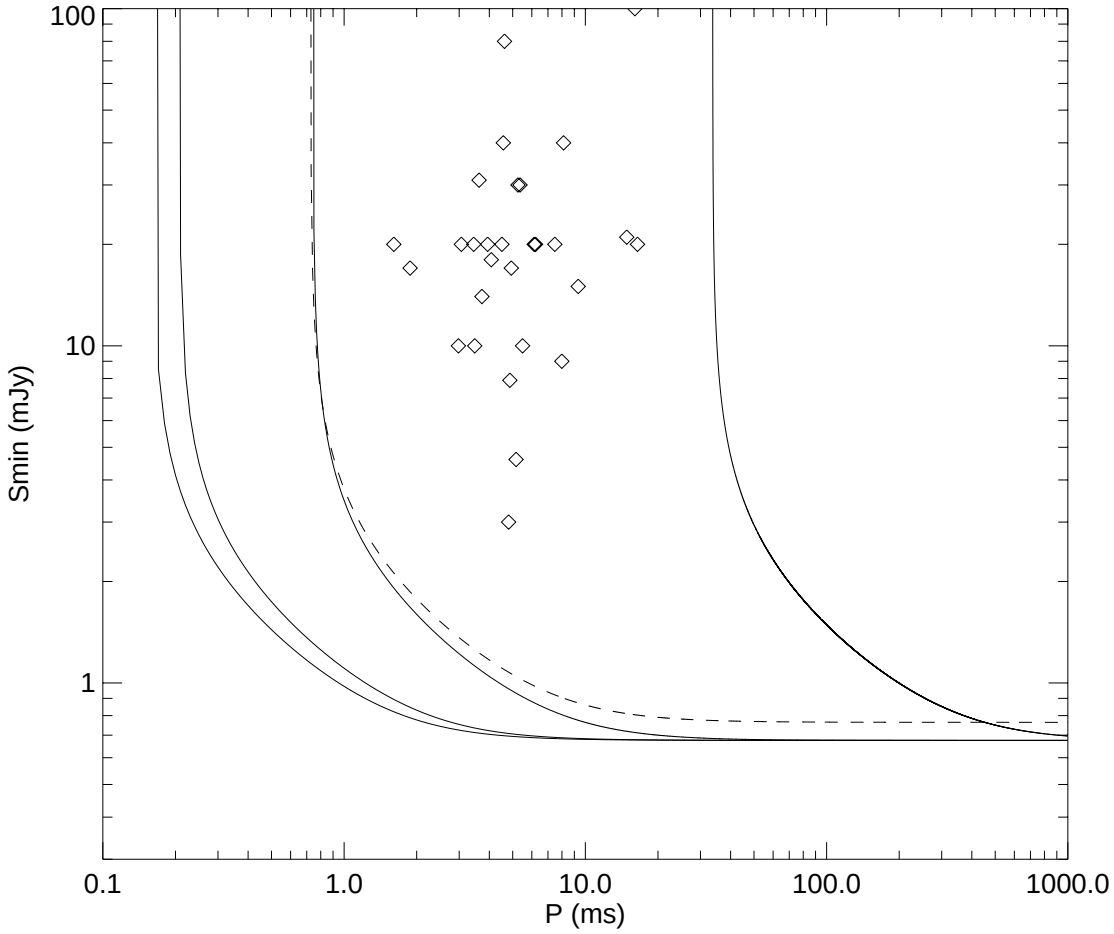

Figure 4.2: Search sensitivity. The solid lines show this survey's minimum detectable flux density as a function of pulsar period for DM $= 0, 20, 100,$ and $280\,\mathrm{cm}^{-3}\,\mathrm{pc}$ (left to right). The dashed line shows a DM $= 20\,\mathrm{cm}^{-3}\,\mathrm{pc}$ curve typical of the previous generation of Arecibo surveys. The diamonds indicate the known millisecond pulsars ($P < 25\,\mathrm{ms}$) with measured flux densities at 400 MHz. Pulsars in globular clusters are not plotted.

lier surveys. The previous surveys all divided 8 MHz of telescope bandwidth into 32 channels. For the current search, the same 8 MHz was divided into 128 channels and sampled ~ 3 times faster. The faster sample rate increases the number of unaliased harmonics visible in the power spectrum for high spin frequency pulsars. The use of smaller individual frequency channels reduces the (unrecoverable) single-channel dispersive pulse smearing. The result is that the present survey is not only significantly more sensitive to millisecond pulsars, but also has true millijansky sensitivity to sub-millisecond pulsars.

For sources at large zenith angles, our sensitivity is also better than pre-1994 searches due to the installation of the ground shield, which reduces the spillover temperature. This is true regardless of pulsar spin period.

We note that our search sensitivity applies equally to isolated and binary pulsars. For most plausible binary orbits, the doppler shifted period of an accelerating pulsar will not change appreciably during our 42 sec integration time. Only for extremely tight orbits ($P_b \lesssim 1/2\,\mathrm{hr}$) with massive companions ($M_c \gtrsim 1.0 M_\odot$), will our sensitivity suffer significantly. Drift surveys from Arecibo are able to achieve mJy sensitivity, despite the short integration time, due to the telescope's large collecting area. To reach this sensitivity with other (smaller) telescopes requires longer pointed observations, sacrificing some sensitivity to binary pulsars. Special techniques can be employed to remove the acceleration of the pulsar signal, but the computational costs can quickly become excessive for a large area survey.

4.2.3 Search Method

We analyzed the data 2^{19} samples at a time, since $2^{19}\,80\,\mu\mathrm{sec} = 41.94\,\mathrm{s}$. For the purposes of our analysis, we defined a data segment of this length to be one "beam." Each beam that we analyzed overlapped the previous one by 50%, so that a source would spend at least 75% of its transit in an analyzed beam. Our entire survey included $\sim 10^5$ such beams.

Each beam was dedispersed for 210 different dispersion measures (DMs), including

trial values up to $280\,\mathrm{cm}^{-3}\,\mathrm{pc}$. For each DM trial, a $2^{19}-$point FFT was calculated, from which a normalized power spectral estimate was formed. Spectral features due to known radio frequency interference (RFI) were removed, as were harmonics of the 60 Hz AC signal. If a spectral bin contained power exceeding a predetermined threshold, the corresponding spin frequency and DM were saved as a candidate pulsar. Harmonic folds of 2, 3, 4, 8, and 16 harmonics were also calculated, and candidates exhibiting significant total power were similarly saved. If a given beam produced multiple degenerate candidates, whose spin periods were very similar (or were harmonics of each other) for various DM trials, then duplicates were removed. Only the most statistically significant candidate with that spin period was saved. For each candidate, a pulse profile was folded and saved for each individual filterbank channel. A single, dedispersed profile was also saved for each candidate. The candidate signal-to-noise thresholds at this stage were intentionally set relatively low, so that each beam was expected to yield approximately one spurious noise candidate. Of course the data also contain a plethora of RFI signals, so after the FFT stage of the analysis, we were left with $\gtrsim 2 \times 10^5$ candidates.

To weed out the noise and RFI signals, the FFT candidates were put through an additional screening pass. The one-dimensional dedispersed profiles and the two-dimensional dispersed profiles were each correlated with a number of templates of varying DM and pulse width, basically to determine if the candidate was "pulsar-like." Also during this stage, if a candidate's DM was indistinguishable from zero, it was eliminated, since there were so many RFI signals that were not dispersed. Again, the thresholds in this stage were somewhat liberal, so if a candidate exhibited a hint of actually being dispersed and having a reasonable waveform, then it survived to the next stage of the analysis.

Candidates that survived the profile selection stage (of which there were ~ 5000), were scrutinized by a human to determine the final list of 50 candidates. These final candidates were then confirmed or confuted by taking new follow up data at Arecibo. During much of the time in which our search data were taken, the telescope pointing was inaccurate by $\gtrsim 10'$. For this reason, multiple pointings had to be performed

when attempting to verify candidates.

The screening described above strongly selected against low-DM, long-period pulsars. In an effort to recover some sensitivity to these apparently non-dispersed signals, we attempted to look for positional coincidences — candidates from different scans with similar period, right ascension (RA), and declination. One detection resulted from this analysis (PSR J1404+12, see Table 4.3). This method was also useful for attempting to verify normal (dispersed) candidates, without having to go back to the telescope. The first new pulsar discovered in this survey, PSR J0627+07, was initially confirmed in other search scans.

4.3 Results

Our search re-detected 19 previously known pulsars, including two MSPs. These are listed in Table 4.1. Three known (slow, bright) pulsars (Table 4.2) ostensibly in our search area were not detected. This could be due to nulling or scintillation, but is likely due to the pointing error mentioned in section 4.2.3 above. Since the recorded sky positions were potentially off by more than a full beam width, the missing pulsars may not actually have been in our search area at all.

The detected pulsars often appear in our search data significantly offset from their catalogued positions. Similar position errors have been seen by the other groups analyzing Arecibo upgrade data. Figure 4.3 shows the measured position offsets for the 19 previously known pulsars detected in this survey (three of them appear twice, as they were detected in multiple scans), along with 33 offsets measured by the Princeton group (D. Nice, private communication). (Occasionally, strong pulsars are observed with position errors exceeding $20'$; this could well be due to detection in the sidelobes of the power pattern.) The three circles of increasing radius represent the approximate area covered by a single telescope pointing and ($\sim$hexagonal) grids of 7 and 19 pointings. Based on the figure, when following up on a candidate corresponding to a real pulsar, we expect that a single pointing will be approximately 40% likely to cover the true position of the pulsar. Likewise, 7 pointings will be sufficient for about

Table 4.1: Previously known pulsars detected in this survey

Name	R.A. (J2000)	Dec (J2000)	P (ms)	DM cm^{-3} pc
B0525+21	05:28:52.3	22:00:00.2	3745.52	51
B0531+21	05:34:32.0	22:00:52.1	33.4033	57
B0611+22	06:14:17.1	22:29:58.2	334.925	97
B0834+06	08:37:05.6	06:10:14.1	1273.77	13
J0947+27	09:47:22.0	27:41:60.0	851.050	29
B0950+08	09:53:09.3	07:55:35.6	253.065	3
B1534+12	15:37:10.0	11:55:56.1	37.9044	12
J1640+2224	16:40:16.7	22:24:09.0	3.16332	18
J1741+2758	17:41:53.5	27:58:09.0	1360.74	30
J1811+0702	18:11:20.7	07:02:23.0	461.712	54
B1859+07	19:01:39.0	07:16:34.7	643.998	253
J1903+2225	19:03:53.0	22:25:12.3	651.185	109
B1915+22	19:17:43.8	22:24:29.0	425.906	201
B1930+22	19:32:22.7	22:20:56.6	144.455	219
B2028+22	20:30:40.4	22:28:21.7	630.512	72
J2033+17	20:33:21.0	17:35:60.0	5.94896	25
J2043+2740	20:43:43.5	27:40:56.0	96.1305	21
B2053+21	20:55:39.1	22:09:27.6	815.181	36
B2110+27	21:13:04.4	27:54:02.7	1202.85	25

Table 4.2: Previously known pulsars not detected in this survey

Name	R.A. (J2000)	Dec (J2000)	P (ms)	DM cm^{-3} pc
B1530+27	15:32:10.3	27:45:50.6	1124.84	15
B1903+07	19:05:53.8	07:09:21.1	648.039	269
J1910+0714	19:10:19.0	07:14:11.8	2712.42	124

Table 4.3: New pulsars detected in this survey

Name	R.A.[a] (J2000)	Dec[a] (J2000)	P (ms)	DM ($\mathrm{cm^{-3}\,pc}$)	S_{430}[b] (mJy)
J0517+22	05:17:10	22:16:00	222.344	20	7
J0627+07	06:27:44	07:06:00	475.831	138	6
J1404+12	14:04:40	12:03:00	2650.27	25	3
J1935+12	19:35:15	12:03:00	1939.90	190	1
J1938+22	19:38:10	22:16:00	166.118	95	1[c]

[a]Position uncertainty for PSR J0627+07 is $\sim 1.5'$; all other positions are uncertain by $\sim 5'$.
[b]Flux densities are uncalibrated and should be considered order of magnitude estimates.
[c]Time average flux density, including high and low states (see section 4.4.2).

73% of candidates, and roughly 91% of candidates will be verified with 19 pointings.

With a generous but finite allotment of telescope time, we were able to follow up on 50 candidates. The majority (31) received 7 pointings, but we were able to search 19 or more sky positions for 15 of our top candidates. If a candidate came from a search scan containing a known pulsar, we used the measured RA offset of the pulsar to correct the nominal position of the candidate. For example, the newly discovered pulsar PSR J1935+12 was originally detected in a scan which also contained the known pulsar PSR B1534+12. The candidate RA was adjusted by $14'$, the measured offset of PSR B1534+12, and the new pulsar was confirmed at the corrected position.

In Table 4.3 we list the main characteristics of the five new pulsars detected in this survey. Their average pulse profiles are shown in Figure 4.4. Despite our excellent sensitivity, we did not find any new MSPs or sub-millisecond pulsars.

We detected two (previously known) MSPs in our 840 $\mathrm{deg^2}$ search area, for a surface density of 1 MSP per 420 $\mathrm{deg^2}$. This is somewhat lower than the combined results of previous high latitude Arecibo surveys, which have found MSPs at a rate of roughly 1 per 300 $\mathrm{deg^2}$, to a sensitivity limit of several mJy (Camilo et al., 1996b). The lower detection rate of the current survey is rather surprising, considering our improved sensitivity in comparison to the earlier searches.

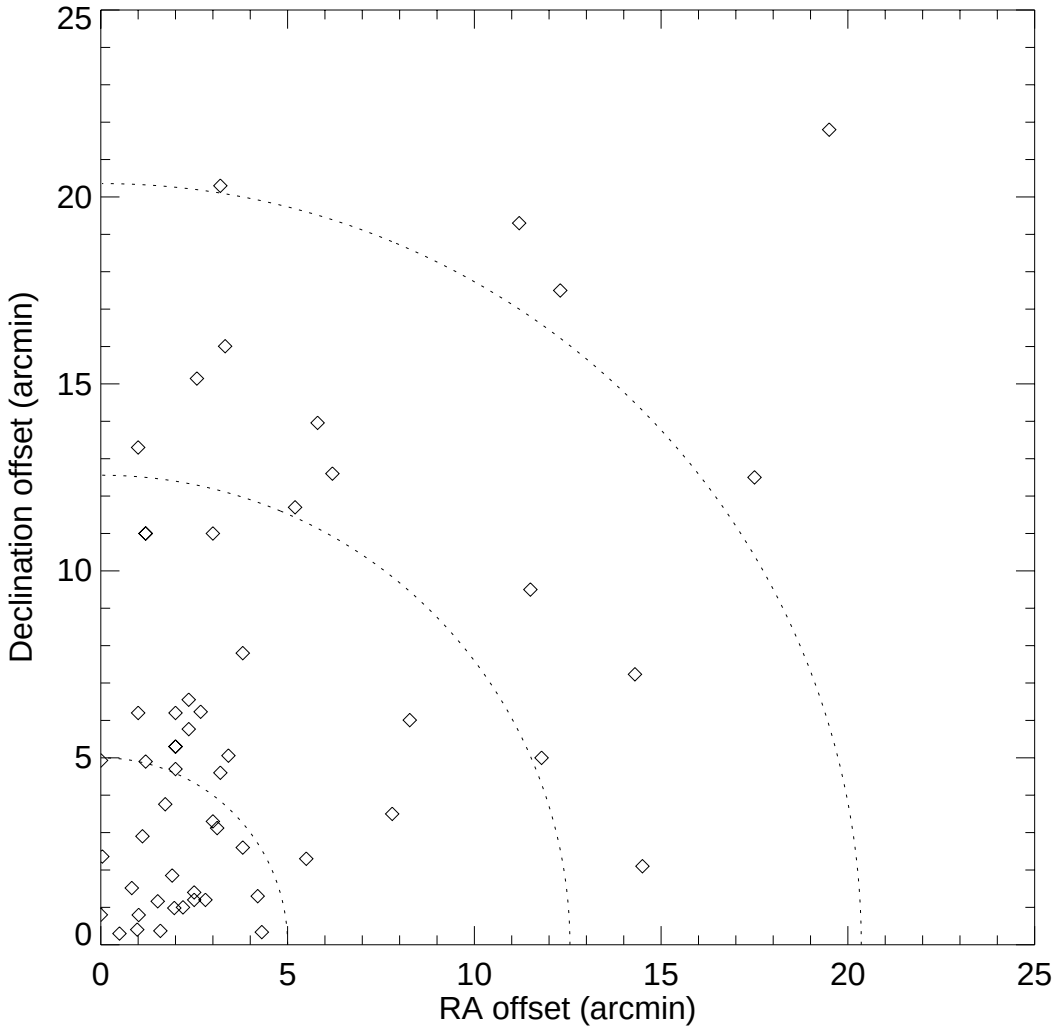

Figure 4.3: Measured position errors of previously known pulsars. The circular dotted lines represent the approximate area that can be covered by 1, 7, and 19 telescope pointings, including respectively 40%, 73%, and 91% of the plotted pulsars.

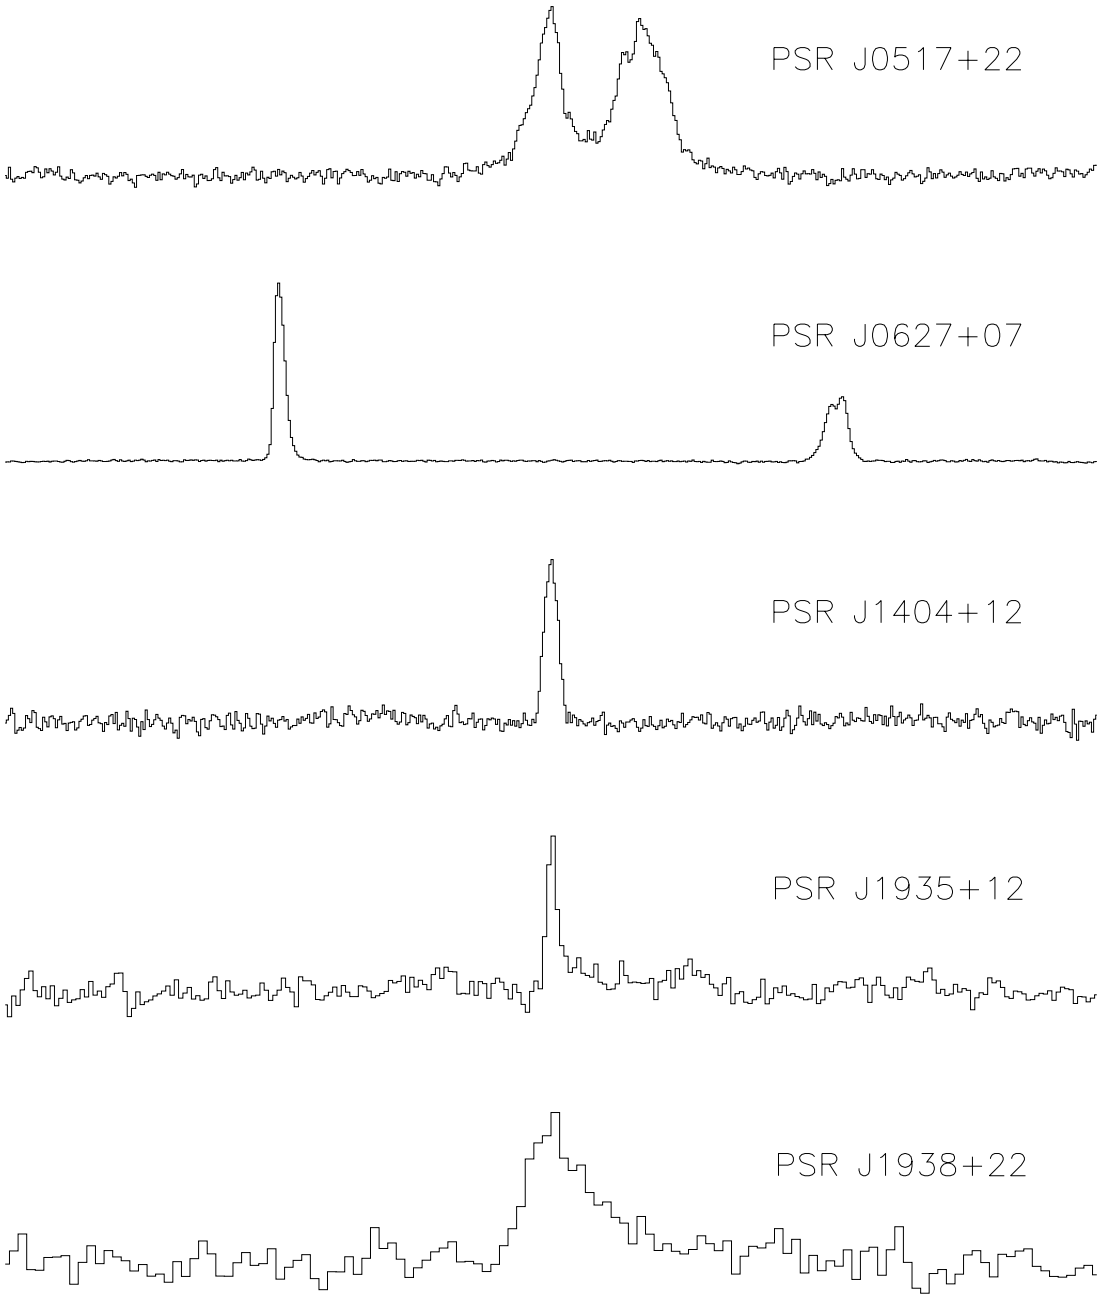

Figure 4.4: Folded pulse profiles of the 5 new pulsars discovered in this survey. The vertical (flux) scales are arbitrary.

4.4 Discussion

4.4.1 PSR J0627+07

PSR J0627+07 is remarkable in several ways. It has the largest DM of any pulsar yet discovered in the Galactic anticenter region. The Taylor & Cordes (1993) free electron model places it at a distance of $\sim 7.8\,\mathrm{kpc}$, which seems unlikely, given the pulsar's large apparent flux density of $\sim 6\,\mathrm{mJy}$. We may be seeing this pulsar through an ionized cloud.

The waveform of this pulsar is also quite interesting. As is evident from Figure 4.4, at 430 MHz PSR J0627+07 shows a clear interpulse which lags the main pulse by 183°. The same morphology is seen at 1400 and 2380 MHz. The apparent wavelength independence of the pulse-interpulse separation argues against conal emission from a single magnetic pole. Apparently, PSR J0627+07 is a rare example of an orthogonal rotator, where the magnetic axis is inclined to the rotation axis by 90°, and the main pulse and interpulse originate in opposite magnetic poles. It is also interesting to note that the main pulse appears to be quite narrow. According to Lyne & Manchester (1988), the opening angle of the emission beam from an orthogonal rotator at 400 MHz should satisfy $\rho \geq 13^\circ P^{-1/3}$. For PSR J0627+07, the width of the main pulse (full width at 10% of maximum) is approximately $w_{10} = 7.6^\circ$, much narrower than the predicted minimum beam width of 16.7°. Firm conclusions will await the determination of the impact parameter (the minimum angle between the magnetic axis and the line of sight).

PSR J0627+07 is located within about 3° of the center of the large, nearby Monoceros supernova remnant. Using the estimated distance ($0.8 - 1.6\,\mathrm{kpc}$) and age ($30 - 100\,\mathrm{kyr}$) of this remnant (Kirshner et al., 1978; Davies et al., 1978; Graham et al., 1982; Leahy et al., 1986), we find that a pulsar proper motion of about $400 < v_\perp < 2700\,\mathrm{km\,s^{-1}}$ would be required for the pulsar to have been born in the supernova. Pulsar velocities as high as $v \sim 1600\,\mathrm{km\,s^{-1}}$ have been measured (Cordes & Chernoff, 1998), so an association is certainly possible, provided the characteristic age of the pulsar turns out to be compatible. Timing of this pulsar is underway at

Arecibo. A full year of timing observations will allow precise fitting of the pulsar sky position (and therefore removal of the effects of earth motion) and precise determination of the intrinsic spin down rate, $\dot{P}$, from which the characteristic age can be determined.

4.4.2 PSR J1938+22

PSR J1938+22 is noteworthy because of its marked intensity fluctuations. Specifically, it appears to have distinct states of low and high emission. The pulsar will be nearly undetectable for $\sim 30 - 200$ s at a time and then abruptly turn on for $\sim 1 - 7$ s bursts. The pulse energy has a bimodal distribution, with the average high-state energy exceeding that of the low state by a factor of 34. Based on our somewhat limited data (apart from the discovery scan, we have only a single 336 s follow up observation), the pulsar appears to be "on" only 3.5% of the time. Three high-state episodes occurred during the 336 s observation, the longest of which is shown in Figure 4.5. There is no obvious periodicity associated with the onset of the bursts, nor is there any clear evidence of subpulse structure.

The rather large DM of $95 \pm 5\,\mathrm{cm}^{-3}\,\mathrm{pc}$ coupled with the abruptness of the phase transitions would argue against a scintillation origin. The high and low states are therefore probably due to the intrinsic emission of the pulsar. Conventional nulling also seems unlikely, since a null fraction of 96.5% would be unprecedented (typical null fractions are $\lesssim 1\%$ (Biggs, 1992), though fractions as high as 70% have been observed (Durdin et al., 1979)). At 166 ms, this would be the shortest period pulsar to show any significant nulling. Also, "null" pulse fluxes are typically much less than 1% of the mean flux of the detected pulses, whereas for PSR J1938+22, the ratio of mean low-state to high-state flux is $\sim 3\%$ (this is not at all obvious from Figure 4.5; the low-state pulsation is only detected in integrations of $\gtrsim 80$ s). This pulsar clearly warrants further observation and analysis.

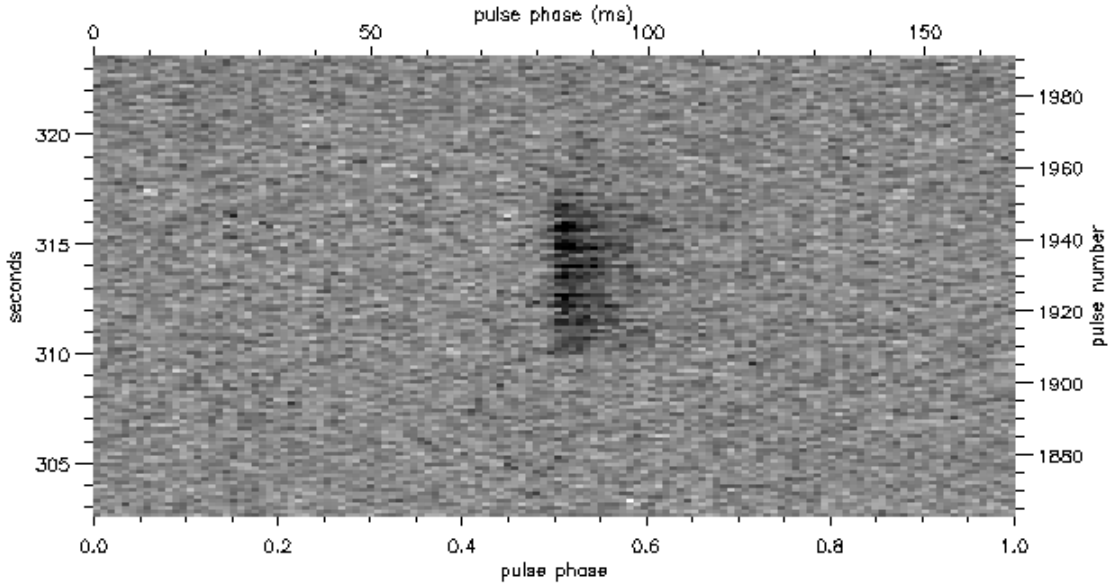

Figure 4.5: Single pulses from PSR J1938+22. Shown is the longest of the three bursts contained in the 336 s follow up observation. The pulsar is bright only $\sim 3.5\%$ of the time.

4.4.3 Sub-millisecond pulsars

The original MSP, PSR B1937+21 (Backer et al., 1982), with a spin period of 1.558 ms, still stands today as the fastest known pulsar. Over the last 20 years, the number of known pulsars with periods less than 25 ms has grown to at least 120; their period distribution is shown in Figure 4.6. It is tempting to infer from the figure that the distribution of MSP periods peaks around $3 - 4$ ms and has a hard minimum of about 1.5 ms. This conclusion may be supported by observations of low-mass X-ray binaries (LMXBs), long considered to be the progenitors of millisecond radio pulsars (Alpar et al., 1982). The X-ray evidence by itself seems to indicate a preponderance of neutron stars (NSs) spinning with ~ 3 ms periods, with a minimum period of perhaps 1.7 ms (van der Klis, 2000). However, these conclusions may be premature. The X-ray observations represent a small sample, and their interpretation is still a subject of debate, while the fastest observed radio pulsar periods are suspiciously close to the detection limits of the surveys in which they were discovered. In a detailed study

of the radio MSP population, accounting for the selection effects of these surveys, Cordes & Chernoff (1997) determined that faster spin periods are certainly possible. Their model indicates a 95% confidence lower limit of 1.0 ms and a 99% lower limit of 0.65 ms.

Due to the formidable storage and computational requirements, radio pulsar searches with useful sensitivity to sub-millisecond pulsars have only recently begun. Crawford et al. (2000) and Edwards et al. (2001) performed pointed searches towards unidentified FIRST and NVSS sources and globular clusters. The largest area survey to date is that of D'Amico (2000), covering nearly 10000 deg^2 (almost 25% of the sky) at 408 MHz with $\gtrsim$ 10 mJy sensitivity to sub-MSPs. Arecibo (PSPM) search sensitivities are about an order of magnitude better than this and will therefore provide much more stringent limits on the sub-MSP population. The only published PSPM search results belong to Cadwell (1997), who analyzed just over 220 deg^2. No sub-MSPs have been found in any of these surveys.

The current search covered 840 deg^2 or 2.0% of the sky. No sub-MSPs were discovered to a limiting sensitivity of about 2 mJy for DM = 20 cm^{-3} pc and P = 0.5 ms. (The dashed lines in Figure 4.6 indicate the periods to which our search had 1 mJy and 2 mJy sensitivity for DM = 20 cm^{-3} pc.) Using simple Poisson statistics (assuming an isotropic distribution on the sky), our result indicates a 95% confidence upper limit on the surface density of sub-MSPs of 1 per 280 deg^2 down to 2 mJy. If we include the area searched by Cadwell (1997), the upper limit becomes 1 per 350 deg^2. As more groups complete and publish their PSPM drift surveys, this upper limit will continue to improve.

Solid observational constraints on the lower limit of the pulsar spin period distribution have important consequences. This limit may be set by intrinsic properties of the NSs themselves, or by the process responsible for spinning them up to short periods. We discuss some of these processes in the next few paragraphs. Although our result cannot differentiate between these possible explanations, it does firm up their observational basis.

The maximum possible NS spin rate is determined by the nuclear equation of state

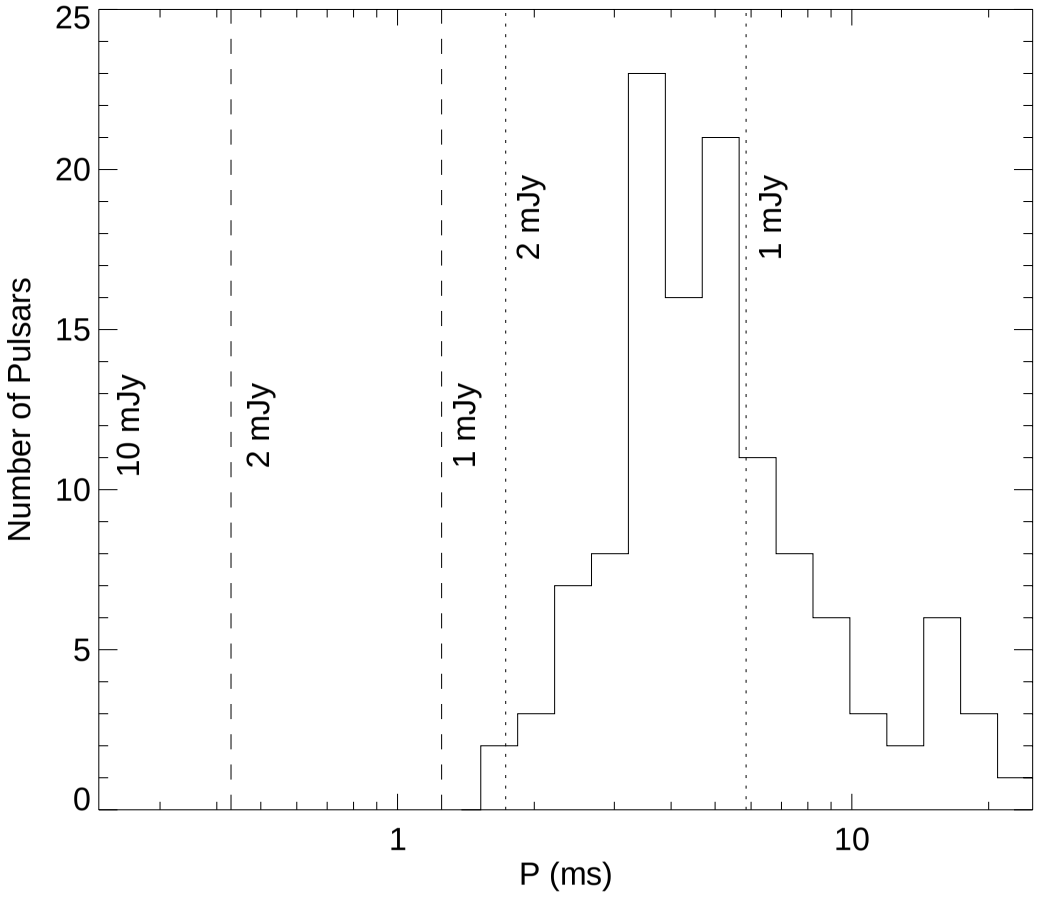

Figure 4.6: Millisecond pulsar period distribution. The dotted lines show the periods at which the previous generation of large-area Arecibo surveys had 1 mJy and 2 mJy sensitivity, for DM = 20 $\mathrm{cm}^{-3}\,\mathrm{pc}$. The dashed lines show the same for the current survey. Our search was sensitive to the entire plotted range of periods with at least 10 mJy sensitivity.

(EOS), which is still poorly understood. Cook et al. (1994) calculated the maximum spin rate for a NS, using 14 different proposed EOSs. The minimum spin periods they found ranged from 1.4 ms for the hardest EOSs (those with large radii for a given mass) down to 0.53 ms for soft EOSs. Some of the softest EOSs are already ruled out by NS mass measurements, while harder EOSs are ruled out by PSR J1937+21, whose phenomenal rotational stability indicates that it is not perilously close to the mass shedding limit. Our lack of sub-MSP detections provides some observational evidence in favor of a hard EOS, but clearly more sky must be searched before firm conclusions can be drawn.

In actuality, pulsars may never reach the maximum break-up spin rate allowed by the EOS, if the spin-up process is somehow interrupted. Several mechanisms for this have been proposed. If Figure 4.6 is truly representative of the pulsar period distribution, as our result would indicate, this lends support to these models.

Temperature gradients in the crust of an accreting NS can lead to density variations caused by differing electron capture rates. If these density variations are non-axisymmetric, the resulting mass quadrupole will radiate away angular momentum in the form of gravitational waves (Bildsten, 1998; Ushomirsky et al., 2000). With plausible model parameters, this mechanism can explain the clustering of LMXB periods around 3 ms, as the gravitational radiation balances the accretion torque.

R-modes (fluid modes with axial parity) in the NS interior may similarly halt spin-up (Andersson, 1998; Bildsten & Ushomirsky, 2000). Beyond some critical spin frequency (where the growth rate of the r-mode is equal to the viscous damping rate), the amplitude of the mode can become large enough to radiate away a significant amount of angular momentum via gravitational radiation. Calculated values for the critical spin rate have fluctuated by about an order of magnitude in the last few years, but the latest treatment (Levin & Ushomirsky, 2001) puts the onset of the instability in the interesting 3 ms range, though this maximum spin rate depends sensitively on the (unknown) thickness of the NS crust.

Burderi et al. (2001) recently proposed another possible explanation for the apparent paucity of sub-MSPs. A downward fluctuation in the mass transfer rate ($\dot{M}$)

in an accreting NS system can result in the onset of a radio pulsar phase. If the NS is spinning rapidly enough, radiation pressure might then expel material that would otherwise be accreted even after $\dot{M}$ rebounds. An attractive feature of this scenario is that it naturally explains the relatively small observed MSP masses (Thorsett & Chakrabarty, 1999) by halting the accretion itself, not just the spin up. If a system is compact enough ($P_b \lesssim 1\,\mathrm{hr}$), accretion can resume after a radio phase, and spin up to sub-millisecond periods should be possible. The authors attribute the fact that these sub-MSPs have not been detected to selection effects. Although the current survey did not seriously select against pulsars in tight binaries, other factors such as extremely low luminosity, long eclipses, and long duty cycles might make them difficult to detect.

This survey represents only the beginning of sub-MSP searches. In the near term, we can expect the completion and publication of the remainder of the Arecibo upgrade surveys. Future improvements in observing hardware (e.g., a possible Arecibo multi-beam receiver, the Square Kilometer Array) and improvements in computing power and algorithms will significantly drive down the upper limits or will produce sub-MSP detections. Further null results will continue to shed light on the nuclear equation of state and accretion scenarios, while a single sub-MSP detection would necessitate re-evaluation of current NS theories.

We thank Joshua Shapiro and Yuki Takahashi for help with the analysis, and Bryan Jacoby and Ramesh Bhat for help with observations. We are also grateful to David Nice, Walter Brisken, Maura McLaughlin, Duncan Lorimer, and Alex Wolszczan for providing information on the status of their groups' surveys. The Arecibo Observatory is part of the National Astronomy and Ionosphere Center, which is operated by Cornell University under a cooperative agreement with the National Science Foundation. Access to the Hewlett-Packard X-Class computer, located at the California Institute of Technology, was provided by the Center for Advanced Computing Research. This work was supported by NSF grants 0005-1-000024 and 00040-1-000262.

Chapter 5

Discovery of 3 New Binary Millisecond Pulsars in the Globular Cluster M62

Abstract

We report the discovery of three new binary millisecond pulsars in the globular cluster M62 = NGC 6266. These pulsars are the first new objects discovered with the 100 m Green Bank Telescope. Two of the three pulsars were found using a new search method sensitive to binaries whose orbital periods are of the same order as the observation time. With six pulsars, M62 is now the third ranking globular cluster in terms of total known pulsar population. Timing of these new pulsars should provide important information about the host cluster, including useful constraints on its central density and initial mass function.

(Chandler, A. M., Anderson, S. B., Backer, D. B., Jacoby, B. A.,
Kulkarni, S. R., and Prince, T. A., to be submitted to
the Astrophysical Journal)

5.1 Introduction

The discovery of the first millisecond pulsar (MSP) (Backer et al., 1982) quickly led to the suggestion (Alpar et al., 1982) that these objects are formed in low-mass X-ray binaries (LMXBs) when a neutron star gains angular momentum in the process of accreting material from a binary companion. Since the specific incidence of LMXBs in globular clusters is significantly higher than in the Galactic disk, clusters were recognized early on as promising targets for MSP searches. The first cluster MSP discovery came several years later in M28 (Lyne et al., 1987). Over the last 15 years, the known population has grown to at least 67 pulsars in 21 clusters (Kulkarni & Anderson, 1996; Camilo et al., 2000; Possenti et al., 2001; Lyne et al., 2000; Ransom, 2001; S. Ransom, private communication; this paper; Appendix A), now comprising roughly half of the known MSPs.

It is more than just their sheer numbers that make cluster pulsars interesting. Exchange interactions in the dense cluster cores can lead to the formation of exotic pulsar systems such as triples (Backer et al., 1993; Freire et al., 2001), extremely tight binaries (e.g., Camilo et al., 2000; Ransom et al., 2001; D'Amico et al., 2001), double neutron star systems (Prince et al., 1991), and perhaps even neutron star-black hole binaries (not yet discovered). When multiple pulsars are found in a single cluster, they can help elucidate the internal dynamics of the cluster as a whole. Precision timing of cluster pulsars can reveal the effects of acceleration in the gravitational potential of the cluster, and these measurements can be used to place useful limits on the cluster's central density and mass-to-light ratio. Additionally, pulsar positions can be used to make statistical neutron star mass measurements, and constrain the initial mass function of the cluster stars.

Detecting these pulsars is not easy. Pulsars in globular clusters are typically several times more distant than field pulsars, and their flux densities are therefore roughly an order of magnitude lower. To recoup this lost signal strength, integration times must increase by two orders of magnitude, while maintaining a fast sample rate. Large fast Fourier transforms (FFTs) are therefore required (perhaps as large as 10^9

points). To correct for a (usually) unknown degree of interstellar dispersion, and to maintain sensitivity to Doppler-shifted binary pulsars, many hundreds of these large FFTs must generally be calculated in a full search. The potential payoffs usually justify such a computational effort, and these searches have become more tractable in recent years.

All MSP searches benefit from advances in computing power, search algorithms, and data recording speeds. But improvements in telescopes and observing hardware can also drive new discoveries. Installation of the new low noise, wide-band multibeam receivers and high-resolution filterbanks at the Parkes Observatory in Australia has led to the discovery of at least 21 cluster pulsars in the last few years (Camilo et al., 2000; Possenti et al., 2001). Similarly, recent Arecibo upgrades have greatly improved pulsar search sensitivities, particularly at high frequencies. At least 3 new pulsars have been detected in L-band globular cluster searches from Arecibo (S. Ransom, private communication). The recently completed Robert C. Byrd Green Bank Telescope (GBT), however, has the potential to bring about a much more significant increase in the known population of globular cluster pulsars. With a 100-meter reflecting surface, the GBT is the largest fully-steerable telescope in the world, able to see declinations from $+90^\circ$ down to $\sim -45^\circ$. This covers 85% of the celestial sphere, including enough of the southern hemisphere to cover the Galactic center region, encompassing more than 80% of the cataloged globular clusters in the Galaxy (Harris, 1996). Our group recently installed the Berkeley-Caltech Pulsar Machine (BCPM) at the GBT for community use. This flexible digital filterbank is the GBT's first pulsar back end, and has allowed several groups to begin searching for new cluster pulsars under the GBT's Early Science program.

In this paper we report the discovery of 3 new binary MSPs in the globular cluster NGC 6266 = M62. These are the first new objects discovered using the GBT. A preliminary announcement of these results was made by Jacoby et al. (2002). Two of these new pulsars were discovered using a novel technique sensitive to pulsars whose orbital periods are approximately equal to the observation length, traditionally a poorly sampled regime of orbital periods. With the three previously known pulsars

in M62 (D'Amico et al., 2001; Possenti et al., 2001), there are now six known MSPs in this cluster, making it a prime candidate for dynamical studies. M62 is now the third ranking globular cluster in terms of the number of known pulsars; only 47 Tuc and M15 have more, with 20 and 8 respectively (Camilo et al., 2000; Anderson, 1993).

In section 5.2 below, we describe our search observations and analysis methods. We also describe our new binary pulsar search technique in great detail. In section 5.3, we present the results of our search, including the parameters of the newly discovered pulsars in M62. We discuss the importance of the new discoveries and future prospects for M62 and globular cluster searches in general.

5.2 Observations and Analysis

5.2.1 Data Collection

Installation of the Berkeley-Caltech Pulsar Machine (BCPM) at Green Bank was completed in early August, 2001. The BCPM is a copy of the 96-channel, 170-MHz bandwidth, analog/digital filter bank that was designed and built by a team from Berkeley, NRL, and Caltech with previous installations at Nançay, Effelsberg and Arecibo. A partial technical description is contained in Backer et al. (1997).

The BCPM accepts dual polarization IFs in the 300-500 MHz range. These are divided into 6 analog channels using Mixer/Filter/LO units with independent gain and leveling control. The signals are then sampled and divided digitally 16 ways, for a total of 96 channels in each polarization. The maximum bandwidth per channel is 1.8 MHz, though smaller bandwidths are selectable. The dual polarization data are then brought together, decimated down to the desired sample rate, and (optionally) summed. The mean is then removed and the remainder is quantized into 4 bits. The 4-bit data are formatted and passed to the host workstation (a Sun Ultra-1) over an EDT interface. Data are then passed to disk, and from disk to tape. Monitoring, control, and data acquisition in the host workstation use a version of the Penn State Pulsar Machine (Cadwell, 1997) software, which has been integrated into the GBT

observing system.

Our original search data were recorded on August 16, 2001, with confirmation data taken December 4-9, 2001. For our search and follow up observations, we used the BCPM fed by the L-band receiver, with the central frequency set to 1400 MHz. To limit the effects of dispersion, we used 0.7 MHz channels, for a total bandwidth of 67.2 MHz. We selected a sample rate of 100 μsec, with the two polarizations summed in hardware. Thus, our 4-hour search observation included just over 2^{27} time samples. Much of our analysis was actually carried out with the data resampled at 200 μsec resolution, since the rather large dispersion measure of the M62 pulsars (DM = 114.4 cm^{-3} pc) causes a single-channel dispersive delay of 240 μsec. The half power beam radius of the telescope (at 1400 MHz) is approximately 4.3′ (HWHM), which is nearly 24 times the cluster's core radius, and 3.5 times the half-mass radius. This beam should easily cover all dynamically relaxed pulsar systems in the cluster, as well as some that have been ejected from the core, but remain bound to the cluster (like, for example, PSR B2127+11C, located approximately 13.5 core radii from the center of M15; Anderson, 1993).

5.2.2 Search for Isolated Pulsars

We began with a standard search for isolated pulsars (Burns & Clark, 1969; Lyne & Graham-Smith, 1998). Of course, MSPs are almost certainly formed in binary systems, but solitary MSPs are not uncommon. Their existence can be explained by the obliteration of eclipsing companions, tidal disruption of the companion, or ionization of the binary system by near collisions (particularly for cluster pulsars) (van den Heuvel & van Paradijs, 1988; Biggs et al., 1994).

The first step in the data analysis was to shift and add the filterbank channels to create a single dedispersed time series. Since there were already known pulsars in M62 (D'Amico et al., 2001), we were able to use the cluster's previously determined dispersion measure (the column density of free electrons along the line of sight to the cluster) DM = 114.4 cm^{-3} pc. This saved us considerable effort, compared to a truly

blind search involving many trial DMs.

We calculated the FFT of this time series to generate an estimate of the normalized power spectrum. Frequency bins containing statistically significant power were saved as possible pulsar candidates. To take advantage of the short duty cycles typical of radio pulses, incoherent power sums of up to 16 harmonics were also calculated and searched for significant candidates.

The top candidates were then analyzed by human. Visualization techniques included time-domain and frequency-domain analyses. For each candidate, we folded a single dedispersed pulse profile, as well as a two-dimensional folded filterbank profile (to verify that the candidate was broad-band and exhibited the expected amount of dispersion). From the dedispersed time series, we also calculated a two-dimensional time-resolved pulse profile (D'Amico et al., 2001) and a time-resolved power spectrum, in which the data are divided into a number of shorter segments, and each is folded or Fourier transformed individually. These time-resolved techniques are useful for spotting binary motion, though this simple FFT search is expected to be sensitive only to rather long-period binaries. This is discussed in much greater detail in the next section. We note that one of our newly discovered pulsars was found using this straightforward FFT method, despite the fact that it is in a ~ 1 day binary.

Figure 5.1 shows the sensitivity of this search to isolated pulsars. Following standard conventions (e.g., Dewey et al., 1985; Bhattacharya, 1998; sec. 4.2.2), the sensitivity is approximated as

$$S_{\rm min} = \alpha\beta \frac{T_{\rm sys}}{G(N_p B \tau_{\rm int})^{1/2}} \left(\frac{w_{\rm eff}}{P - w_{\rm eff}} \right)^{1/2} . \qquad (5.1)$$

$S_{\rm min}$ is the minimum detectable flux density. The factors on the right hand side of Equation 5.1 are characteristic of a particular observing system and search method. The parameter α is the threshold signal-to-noise ratio (8.0 was used in Fig. 5.1); β is a factor which accounts for various losses in the observing system ($\beta \sim 1.5$); $T_{\rm sys}$ is the system temperature, the sum of the receiver temperature and the sky temperature (36 K used in Fig. 5.1); G is the gain of the telescope ($1.85\,{\rm K\,Jy^{-1}}$); N_p is the number

of polarizations summed (2); B is the bandwidth of the observation (67.2 MHz); $\tau_{\rm int}$ is the integration time (13421.8 s); P is the rotation period of the pulsar; and finally $w_{\rm eff}$ is the effective pulse width.

The effective width accounts for the broadening of the intrinsic pulse width w_0 due to the finite sampling time ($\tau_{\rm samp} \approx 200\,\mu$sec), and the effects of interstellar dispersion and scattering. The dispersion contribution is due to the dispersive delay within a single filterbank channel

$$\tau_{\rm DM} = 8.3 \times 10^6\ b_{\rm MHz}\ \nu_{\rm MHz}^{-3}\ {\rm DM} = 0.24\,{\rm ms}, \tag{5.2}$$

where $b_{\rm MHz}$ is the channel bandwidth, $\nu_{\rm MHz}$ is the central observing frequency, both expressed in MHz, and DM is the dispersion measure in ${\rm cm}^{-3}\,{\rm pc}$. The dispersive broadening dominates the effect of scattering

$$\tau_{\rm scat} \approx \left(\frac{1000}{\nu_{\rm MHz}}\right)^{4.4} 10^{-7.231+0.9255\log_{10}{\rm DM}+0.814613(\log_{10}{\rm DM})^2} = 0.003\,{\rm ms} \tag{5.3}$$

(Cordes et al., 1991). The effective width is approximated by summing these contributions in quadrature

$$w_{\rm eff}^2 = w_0^2 + \tau_{\rm samp}^2 + \tau_{\rm DM}^2 + \tau_{\rm scat}^2. \tag{5.4}$$

For Figure 5.1, we assumed an intrinsic pulse width equal to 10% of the period.

The dashed line in Figure 5.1 shows the approximate sensitivity of the previous Parkes search of M62 (D'Amico et al., 2001). Due primarily to the much larger reflecting area of the GBT, our search sensitivity is significantly better ($\sim 50\%$ lower flux limit at all periods). We note, however, that our sensitivity is severely limited by the number of frequency channels available in the BCPM. Once the GBT "spigot" comes on-line, an additional factor of 2 improvement in sensitivity should be realized.

5.2.3 Binary Pulsar Search

More than half of the known pulsars in globular clusters are in binary (or triple) systems. When a pulsar moves in an orbit with one or more companions, the observed,

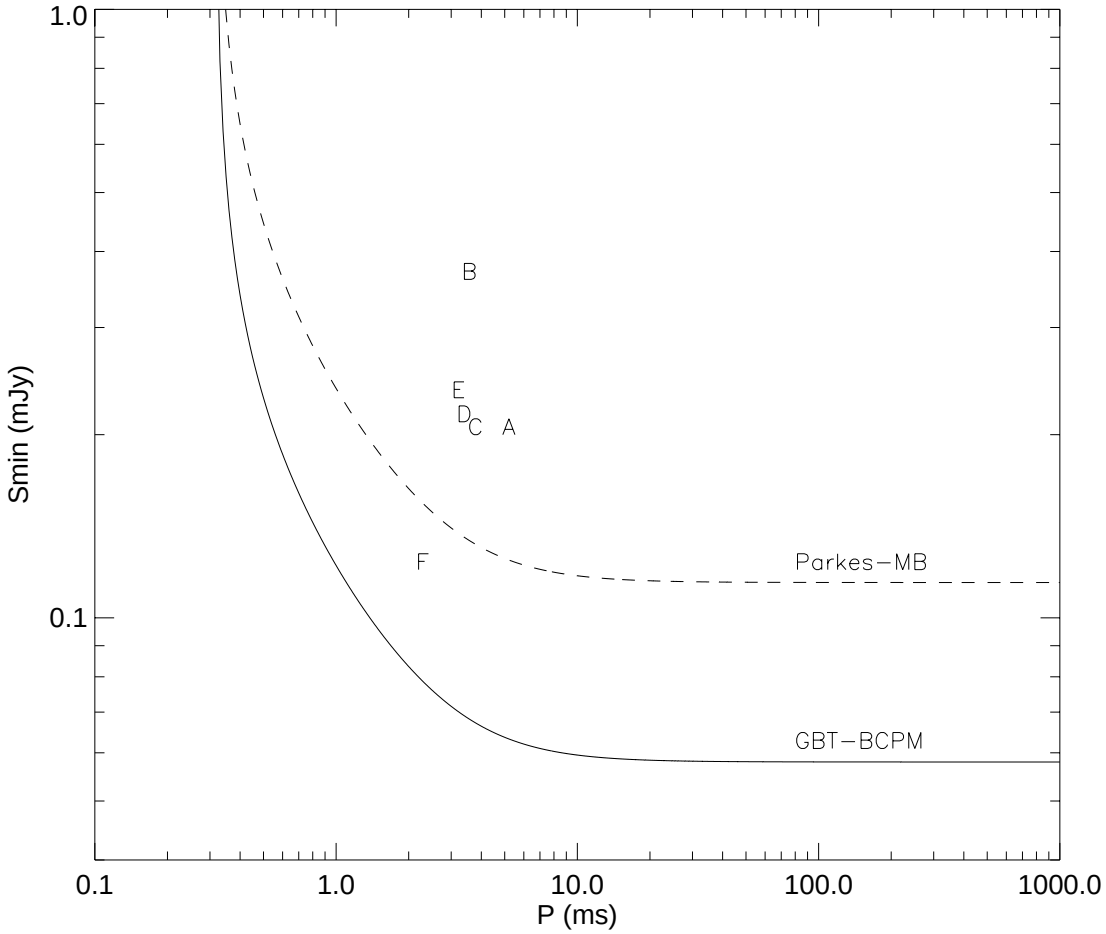

Figure 5.1: M62 search sensitivity for isolated pulsars. Solid line shows the approximate sensitivity of the current search; dashed line shows that of the previous search from Parkes. The periods and approximate flux densities of the six (binary) pulsars in M62 are represented by the letters A through F on the plot.

Doppler-shifted pulse frequency will change with time. In a simple FFT search, the acceleration of the pulsar causes its power to spread out over a range of spectral bins, making its signature difficult (if not impossible) to detect. A number of techniques have been developed to mitigate this effect, with various advantages, disadvantages, and realms of applicability. These binary search techniques increase the computational complexity of a search, but improvements in computer power have made their use much more practicable. The ratio of binary to solitary detection rates in recent years is nearly 3 to 1, suggesting that selection effects have artificially deflated the cluster pulsar binary fraction, which may be much greater than 50%.

Clearly, sensitivity to binary pulsars in a globular cluster search is vital. In this section, we begin by briefly discussing the existing binary search methods that we utilized. These include "acceleration" techniques which have proven their mettle in previous surveys, as well as the sideband or "cepstrum" method (i.e., spectrum of spectrum method) which shows great promise, but has not as yet yielded new pulsars. We did not discover any new pulsars using these methods in the current search. Next, we present a new, powerful binary search technique that we have developed to complement the other approaches. We describe how the new method works and show that in its intended regime of applicability, it compares quite favorably to the other binary search methods, both in terms of sensitivity and computational effort. Perhaps the most important proof of its efficacy is the fact that two of our newly discovered M62 pulsars were found using this new method.

To determine the most appropriate binary search technique for a given system, we must compare the orbital period, P_b, to the observation time, $T_{\rm obs}$. The following discussion is broken down into the three regimes $P_b \gg T_{\rm obs}$, $P_b \ll T_{\rm obs}$, and $P_b \sim T_{\rm obs}$.

$P_b \gg T_{\rm obs}$

If a pulsar spinning with rest frequency f_0 moves in a binary system, the observed, Doppler-shifted frequency is given by

$$f'(t) = f_0[1 + v(t)/c], \tag{5.5}$$

where c is the speed of light and v is the radial velocity of the pulsar (the velocity projected along the line of sight). When an observation covers only a very small part of an orbit, the time dependence of the radial velocity of the pulsar can be well approximated by a constant acceleration $\Delta v(t) \approx at$. This is equivalent to keeping only the lowest order derivative in a Taylor expansion of the frequency evolution, i.e., tantamount to assuming a constant frequency derivative $\dot{f}$.

In a blind search, one corrects the time series for an assumed acceleration before calculating the FFT, harmonic folds, etc., repeating the entire process for each acceleration trial. A full search typically involves $\sim 10^2 - 10^3$ iterations. Acceleration searches of this sort have proven to be quite successful (e.g., Anderson, 1993; Camilo et al., 2000). The acceleration correction can also be carried out in the frequency domain, using essentially a matched filtering approach (Middleditch et al., 1993; Ransom, 2001). This technique is computationally similar to the time domain method, usually involving roughly the same number of operations to carry out. But it may be more efficient than the time-domain method (for a reasonable degree of signal spreading) when the full FFT is distributed over multiple processors, since only one large FFT is required.

Acceleration searches are most effective for pulsars in very wide orbits. For such long-period systems (depending to a lesser extent on orbital phase), the acceleration method is capable of gathering nearly all of the original signal power back into a single bin of the power spectrum (subject to the sinc function response of the discrete Fourier transform). As the orbital period approaches the observation time, the minimum detectable pulsar flux increases, since only a small fraction of the observation is well fit by the assumed linear frequency dependence. We will return to this issue in section 5.2.4 below.

$P_b \ll T_{\rm obs}$

When an observation contains many orbits, the phase modulation of the pulsar signal results in a family of sidebands around the rest frequency of the pulsar in the Fourier spectrum (Ransom, 2001). The number of sidebands depends on the orbital semima-

jor axis and the spin frequency, $N_{\rm sb} \approx 4\pi x f_0$ ($x = a \sin i/c$ is the light travel time of the projected semimajor axis). The known binary pulsars with $P_b < 10\,{\rm d}$ have spin periods in the range $[1.607, 1004]\,{\rm ms}$ and orbital semimajor axes in the range $[0.010, 20.0]\,{\rm ls}$. For these systems, the number of sidebands $N_{\rm sb}$ covers the range $\sim [4, 13\,000]$. For all but the brightest pulsars, the spectral power distributed over these sidebands becomes buried in noise.

Since these side bands are evenly spaced at the orbital frequency, a convenient method of incoherently summing the power back up is to step through the power spectrum, calculating small FFTs. In general the number of sidebands and their location in the spectrum are unknown, so overlapping FFTs of several sizes must be tried. The family of sidebands acts as a pulse train which, when transformed back to the time domain, can produce detectable power at the orbital period (and at "harmonic" multiples of the period). Once a detection has been made, the orbital parameters can be estimated from the original complex Fourier spectrum.

This is the best method currently available for detecting ultra-short-period binary pulsars. While this method has successfully redetected previously known pulsars, its application has not yet resulted in any *new* detections, but it has only been in use for a short while. The shortest known radio pulsar orbital period currently stands at 96 minutes (Camilo et al., 2000), close to the limits of what can be detected using traditional acceleration searches. Two of the three currently known accreting MSPs are in ~ 40 minute orbits (Markwardt et al., 2002; Galloway et al., 2002), and neutron star binaries with orbital periods as short as 11 minutes have been observed (Stella et al., 1987). The modulation sideband technique finally allows pulsation searches with meaningful sensitivity to such systems.

The sideband search method is most effective when an observation encompasses more than a few orbits. Strictly speaking, the method can work as long as at least one orbit is covered by the data, but the sensitivity degrades rapidly as $T_{\rm obs}$ approaches P_b (when $T_{\rm obs} = P_b$, the spacing of the sidebands is equal to the Fourier step size of the spectrum).

$P_b \sim T_{\rm obs}$

Acceleration and sideband searches cover complementary regimes of $P_b/T_{\rm obs}$, but they leave a sizeable gap around $P_b/T_{\rm obs} \sim 1$. Previous searches have therefore had very little sensitivity to pulsars in this region. Traditionally, the best way to handle such systems would be to divide an observation into shorter segments and perform acceleration searches on each of them separately. This effectively eliminates the usual $S_{\rm min} \sim {T_{\rm obs}}^{-1/2}$ advantage of accumulating longer data sets. Alternatively, if the observation time can be increased, then the sideband search method can be applied. However, increasing the observation time significantly (so that $T_{\rm obs} \gtrsim 2P_b$) is not always possible. At Arecibo, for instance, transit times are limited to $\lesssim 3\,{\rm hr}$, while at other sites, sources are usually up for only $8 - 12\,{\rm hr}$. Of course, if a source is circumpolar or if the observatory is in space or (one day) on the moon, then longer observations may be a viable solution. In the next section, we present a third possibility. We describe a new pulsar search technique that we have developed to fill the $P_b/T_{\rm obs}$ gap left by the other two binary search methods.

5.2.4 The Dynamic Power Spectrum Method

Our new search method began as a candidate visualization technique, already alluded to in section 5.2.2 above. It involves the use of a dynamic power spectrum (DPS). The general idea is to divide an observation into a number of shorter segments, and calculate a power spectrum for each (optionally, harmonic folded). The power spectra are then stacked on top of one another in a two-dimensional array. In this spin frequency vs. time plane, a bright pulsar is easily visible as a contiguous curve of excess power tracing out $f(t)$. Our search method relies on the frequency-local nature of these pulsar signatures. Even if none of the segments contain globally significant power, we can detect pulsars by finding a pattern of *locally* significant powers, i.e., barely excess powers in several different time segments within a small range of spin frequencies. In principle, this method is similar to the segmented acceleration search mentioned in the last paragraph, but with a formalized way of lowering the individual

detection thresholds using a built-in coincidence requirement.

The concept of a dynamic power spectrum is not new — applications of DPS techniques have ranged from speech and musical analysis to the study of burst oscillations in LMXBs (Bracewell, 2000; van Straaten et al., 2001). Rudimentary DPS methods have been used in pulsar searches (in M15 by members of our group and in Terzan 5 by Lyne et al., 2000) without success. These earlier applications were very limited in scope, and involved analyzing spectra by eye. In this section, we present a comprehensive method to automate the search process and optimize searches for a wide range of pulsar parameters. The DPS method described here requires far less computation than acceleration searches, and nicely fills in the $P_b/T_{\rm obs}$ gap.

We begin by considering the ideal number of segments $\mathcal{S}$ into which an observation should be divided. This number should be large enough to keep the signal power from spreading out over too many spectral bins, but should also be as small as possible to avoid excessive reduction in the single-segment signal-to-noise ratio (SNR $\sim \mathcal{S}^{-1/2}$). The independent Fourier step size of the DPS will be $F = \mathcal{S}/T_{\rm obs}$. We will require that the fundamental frequency drift of a pulsar, over the course of a single segment, be less than F.

For a pulsar in a circular orbit, the projected radial velocity is given by

$$v(t) = -v_0 \sin i \cos\left(\frac{2\pi}{P_b}t + \phi_0\right), \tag{5.6}$$

where v_0 is the orbital velocity of the pulsar, i is the orbital inclination (the angle between the plane of the orbit and the plane of the sky), and ϕ_0 is the orbital phase at $t = 0$. Here we have defined ϕ to be zero at the descending node, where the pulsar is maximally moving away from the earth (i.e., where the pulse frequency is observed to be at its minimum). We can substitute this expression into Equation 5.5 and take a derivative with respect to time to obtain

$$\dot{f}(t) = \frac{2\pi f_0 \beta}{P_b} \sin\left(\frac{2\pi}{P_b}t + \phi_0\right), \tag{5.7}$$

where we have used $\beta = v_0 \sin i/c$. In the worst-case scenario, we have $|\dot{f}_{\max}| = 2\pi f_0 \beta / P_b$. We require that the frequency range visited by this pulsar over the course of a single segment be less than the Fourier step size, i.e., $\Delta f_{\max} = |\dot{f}_{\max}|(T_{\rm obs}/\mathcal{S}) < \mathcal{S}/T_{\rm obs}$. Solving for the minimum $\mathcal{S}$, we obtain

$$\mathcal{S} = \left(\frac{2\pi f_0 \beta {T_{\rm obs}}^2}{P_b} \right)^{1/2}. \tag{5.8}$$

For typical pulsar and search parameters, the ideal number of segments ranges from one to several hundred. For certain systems, if $T_{\rm obs}/P_b \gtrsim 5$, the number of segments may become unwieldy, but in such cases, the sideband search method is more appropriate.

In a blind search, the pulsar parameters are not known *a priori*, so several values of $\mathcal{S}$ must be attempted, with each resulting DPS searched independently. A fairly exhaustive pulsar search might employ 6 orders, $\mathcal{S} = 4 - 128$ (in powers of two). For non-overlapping segments, $\mathcal{S}$ is restricted to powers of two so that inefficient non-power-of-two FFTs need not be used. Note that if the segment spectra are harmonically folded (generally a good idea), then the frequency drift of the higher harmonics will be proportionately larger than the drift of the fundamental. In such cases optimal detection may occur with a larger number of segments than Equation 5.8 would indicate (although SNR losses associated with using more segments may very well negate any benefit from increasing $\mathcal{S}$).

Calculating a DPS involves roughly the same number of operations ($N_{\rm op}$) as calculating a single full-length FFT. In a full search we calculate a DPS for N_s different values of $\mathcal{S}$, for a total of roughly

$$N_{\rm op} \sim \sum_{\mathcal{S}=\mathcal{S}_1}^{\mathcal{S}_{N_s}} \left[5N \log_2 \frac{N}{\mathcal{S}} + N_{\rm search} \right] \approx 10 N_s N \log_2 N, \tag{5.9}$$

where N is the total number of samples in the full time series and $N_{\rm search}$ is the number of operations required to detect a pulsar signal, once the DPS has been calculated. As we shall see below, coincidentally $N_{\rm search} \sim N \log_2 N$. With $N_s = 6$, a full DPS

search requires about a factor of $10 - 100$ fewer operations than a full acceleration search. We also note that the DPS segments can almost always be processed with in-core, single-processor FFTs.

We now consider the problem of actually detecting the presence of a weak pulsar, having calculated a DPS for a given observation. In many ways, the optimal way to detect the $f(t)$ pattern of a pulsar is by eye. Since this would be rather tedious for a full DPS search, we have developed a hierarchical search algorithm which we describe below. The method we describe here works reasonably well in terms of false positive rate, false negative rate, and computational complexity, though other algorithms are certainly possible.

We begin with a first pass over the entire DPS, selecting the bins containing the top powers. For this first cut, the threshold is set fairly low, so that ~ 100 "candidates" survive to the next stage of the analysis. We do not expect any of these first cut candidates to be statistically significant, when considering all of the DPS bins as independent trials (of which there might typically be $\sim 10^8$). Essentially all of these hits are expected to be due to noise. With $N/2$ points in the DPS, this first cut requires $\sim N$ operations.

Given a first cut candidate in frequency bin b and segment s, we set about to look for locally significant powers in the other segments, in the vicinity of b. If a pulsar has a maximum line of sight acceleration magnitude $a_{\max} = c|\dot{f}_{\max}|/f_0$, then the largest possible change in spin frequency from one segment to the next is

$$\Delta f \leq |\dot{f}_{\max}|\Delta t = f_0 \, \frac{a_{\max}}{c} \, \frac{T_{\rm obs}}{\mathcal{S}}. \tag{5.10}$$

Therefore, in segment $s+1$ or $s-1$, we should look within a range $b \pm \Delta b$, where

$$\Delta b = b \, \frac{a_{\max}}{c} \, \frac{T_{\rm obs}}{\mathcal{S}}. \tag{5.11}$$

In segments $s+2$ and $s-2$, we search $b \pm 2\Delta b$, etc. We thus search within an hourglass figure drawn through the first cut candidate point (b, s), whose opening angle is

determined by $a_{\max}$. A candidate is assigned a statistical score (i.e., significance) based on the powers found and the number of bins searched. As we move up or down through the segments, if we encounter a (locally) significant power, we redraw the hourglass through the current point and continue outward until we reach the first or last segment. This "bootstrapping" technique effectively rewards candidates whose powers follow a contiguous path, by searching fewer points.

To maintain the best sensitivity to a variety of signals, we perform the hourglass search for a number of different $a_{\max}$ values. The largest $a_{\max}$ used should cover any signals detected using smaller $a_{\max}$ trials, but the significance of a detection will be best when the hourglass searched is just large enough to encompass the pulsar's path through the $f - t$ plane. A typical search might involve $a_{\max} = 5$, 10, 25, 50, 75, 100, & $200\,\mathrm{m\,s^{-2}}$, and higher values are easily accommodated. With ~ 100 first cut candidates, each $a_{\max}$ search requires $\sim 100\,\Delta b\,\mathcal{S}^2$ operations. For typical search parameters, the total number of operations required for all 7 $a_{\max}$ values listed above happens to be of the same order as $N \log_2 N$.

As with any pulsar search, the last stage of the process is human inspection of the top candidates.

DPS Examples

Figure 5.2 shows a section of an example DPS. Here we see a 3 hr observation of a simulated 3.0 ms pulsar in a 2.8 hr orbit about a $0.1 M_\odot$ companion (viewed at $i = 60°$). The sinusoidal signal had a single-pulse signal to noise ratio of 0.01. This signal is easily detected, despite the fact that none of the powers are globally significant. The largest power in the feature is consistent with the expected tail of the noise distribution, given the total number of bins in the DPS (in this case, $N = 2^{24}$).

According to Equation 5.8, the ideal number of segments for this pulsar is $\mathcal{S} = 48.25$. In this case, $\mathcal{S} = 32$ yielded a more significant detection than $\mathcal{S} = 64$. For this pulsar, $a_{\max}$ is $17.95\,\mathrm{m\,s^{-2}}$. The dashed lines in Figure 5.2 indicate the initial search hourglass corresponding to $a_{\max} = 25\,\mathrm{m\,s^{-2}}$, centered on the bin containing the largest power.

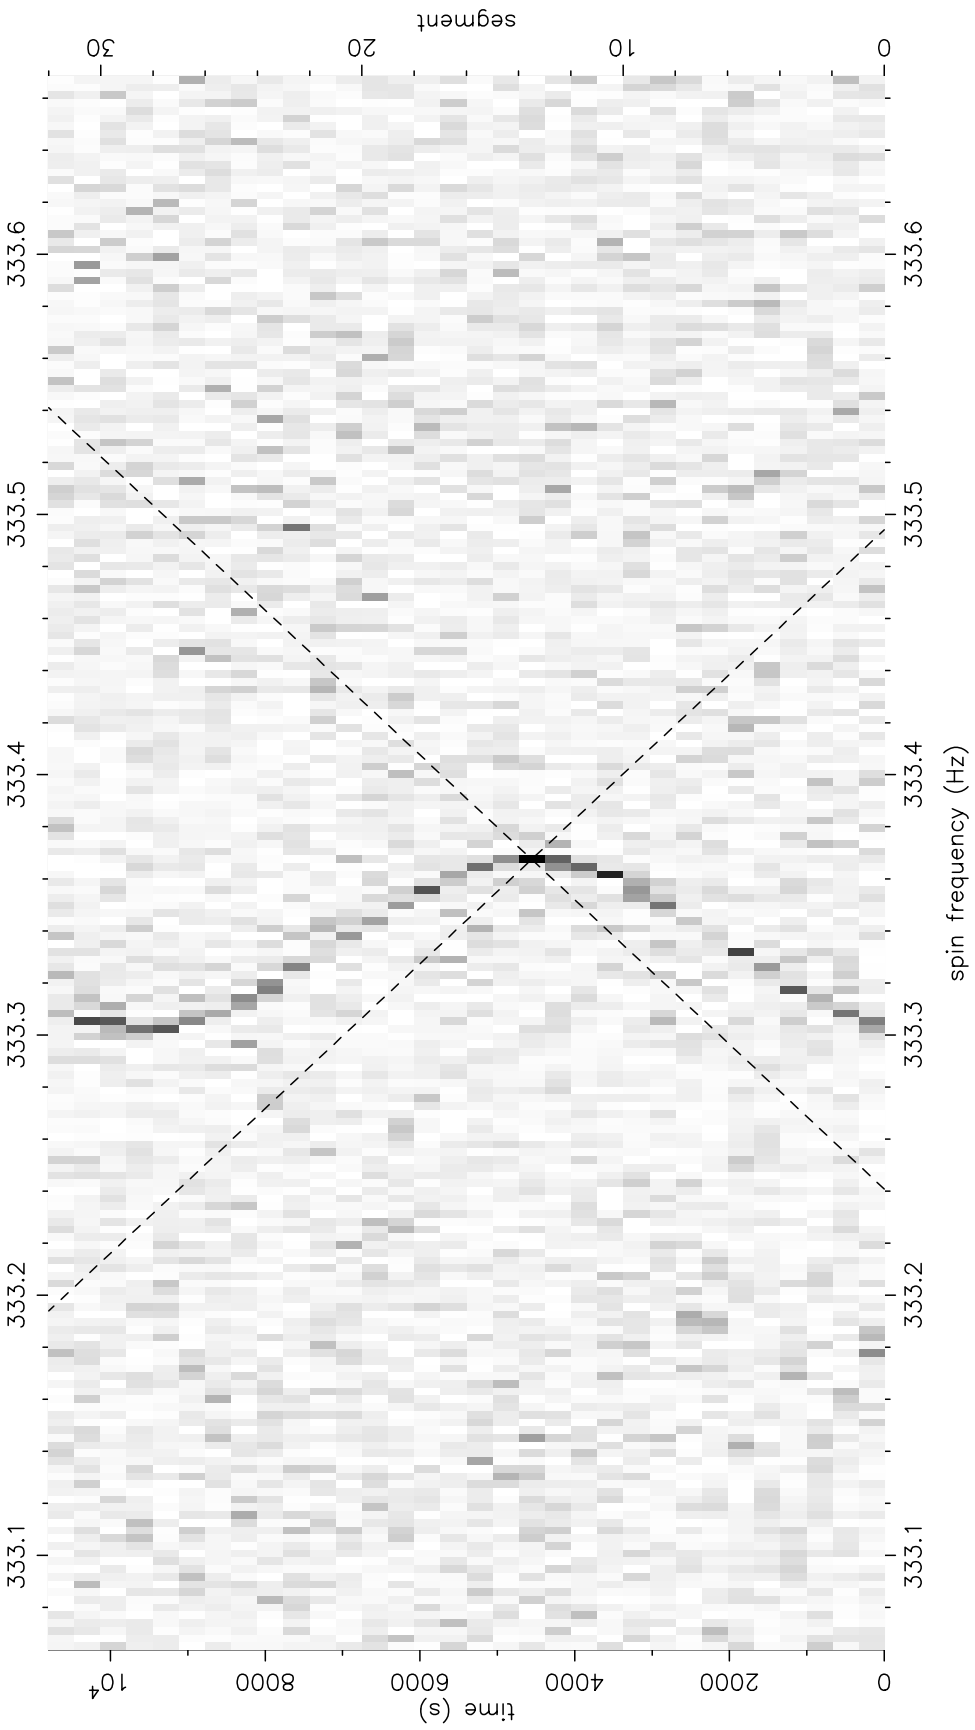

Figure 5.2: Example dynamic power spectrum. Section of the DPS of a simulated 3.0 ms pulsar in a 2.8 hr cirular orbit with a $0.1M_{\odot}$ companion is shown. The observation time was 3.0 hr. Dashed lines show the $a_{\rm max} = 25\,{\rm m\,s^{-2}}$ local search hourglass centered on the first cut hit in segment 13.

To this point our discussion has focused on circular orbits, but the method works for elliptical systems as well. Figure 5.3 shows a non-sinusoidal DPS detection of PSR B2127+11C, the relativistic binary in the globular cluster M15. Four harmonics have been summed in each segment's spectrum (so neighboring bins are not independent), and because the pulsar is so bright, the DPS is shown in a log scale. This pulsar is readily detected with small $a_{\max}$ trials, but if it were significantly less luminous, it would be best detected with a rather large $a_{\max}$ ($a_{\max} = 200\,\mathrm{m\,s^{-2}}$ is overlaid on the figure). An acceleration search would have serious trouble detecting such a signal.

Comparison with Other Binary Search Methods

We now consider a quantitative comparison between the DPS method, the acceleration method, and the sideband search method. In Figure 5.4, we plot the results of simulations showing the minimum detectable instantaneous (i.e., single-pulse) SNR as a function of observation time. All of these simulations were calculated with a 2.0 ms pulsar in a 1.0 hr circular orbit with a mass function $f = 2.89\times10^{-4}$, equivalent to a companion mass of $0.1M_\odot$ viewed at an inclination of $i = 60°$. The points plotted are averaged over the starting orbital phase of the observation. This plot is similar to Figure 4.6 of Ransom (2001), comparing the sensitivities of only the acceleration and sideband search methods.

Also included for reference in the plot is a standard "8-sigma" theoretical sensitivity curve (dotted line). This shows the single-pulse, single-harmonic SNR at which a stationary (or uniformly moving) pulsar would be detected with a single-trial significance equivalent to an 8σ Gaussian event. At the risk of being pedantic, we now explicitly derive the equation for this curve to aid the reader in understanding Figure 5.4. In a power spectrum normalized to unity, the probability that the power in a particular bin will exceed some threshold P is equal to $\exp(-P)$. We equate this to the probability of a Gaussian random variable exceeding 8σ, and solve for P:

$$P = -\ln\left[\int_8^\infty \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx\right] = 35.01. \tag{5.12}$$

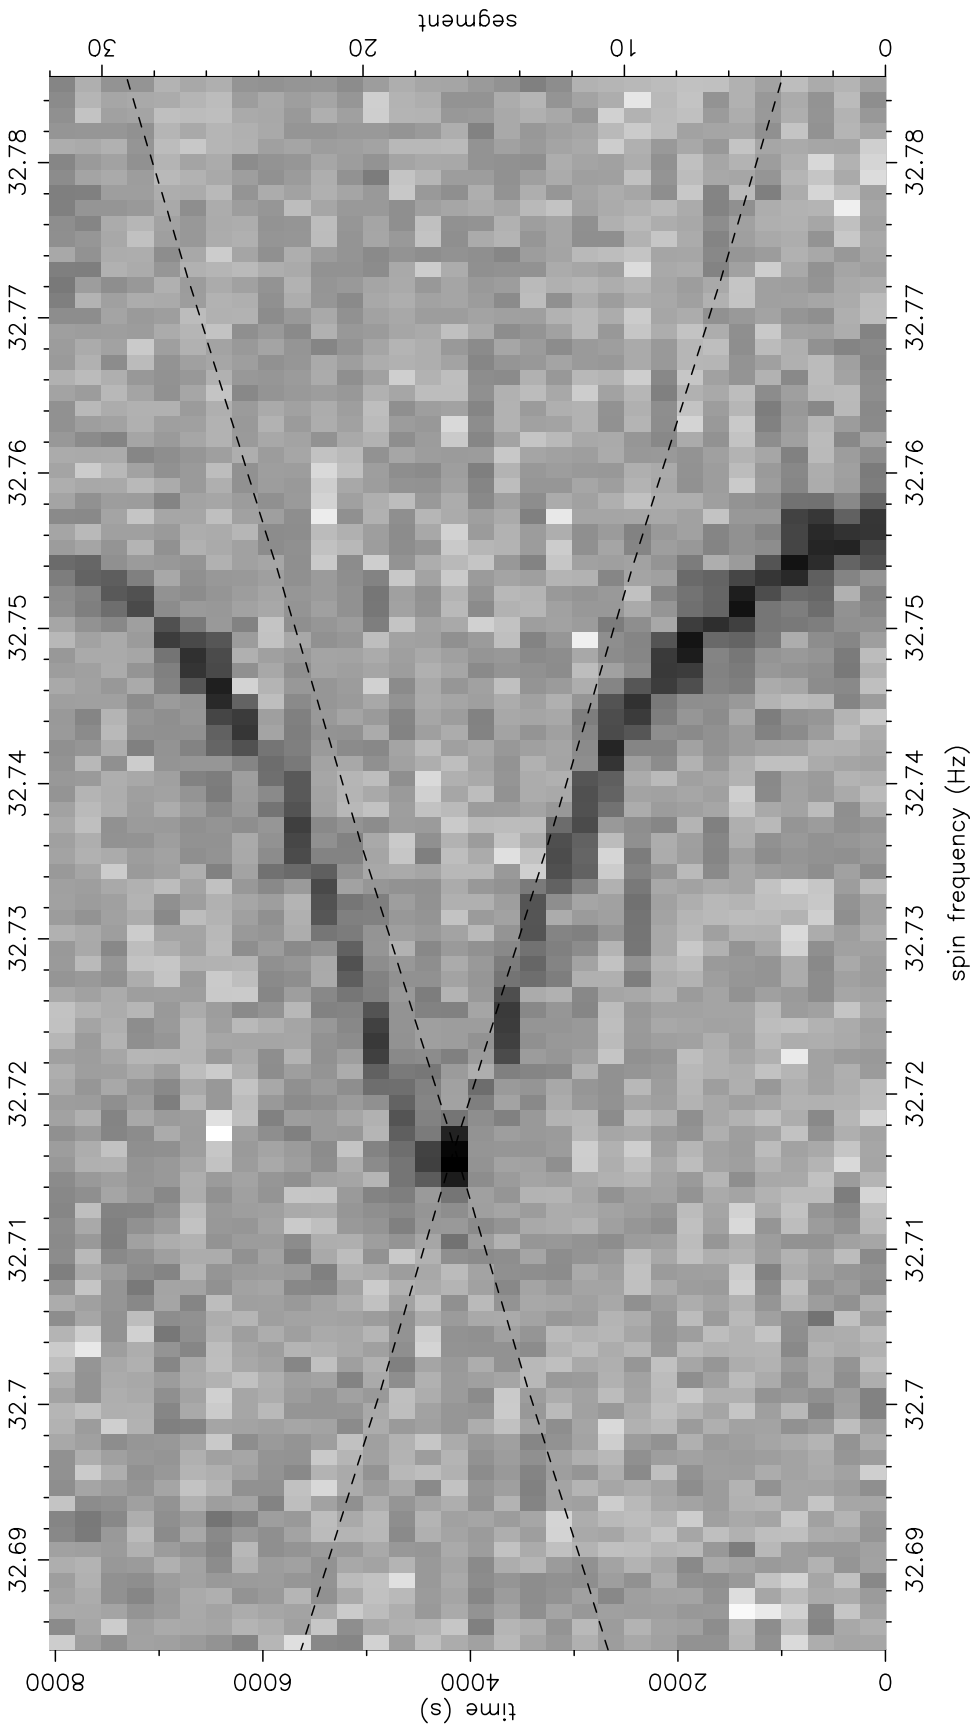

Figure 5.3: Elliptical DPS example. PSR B2127+11C in M15 is seen passing through periastron. The data are from an 8000 s Arecibo observation from January 1999.

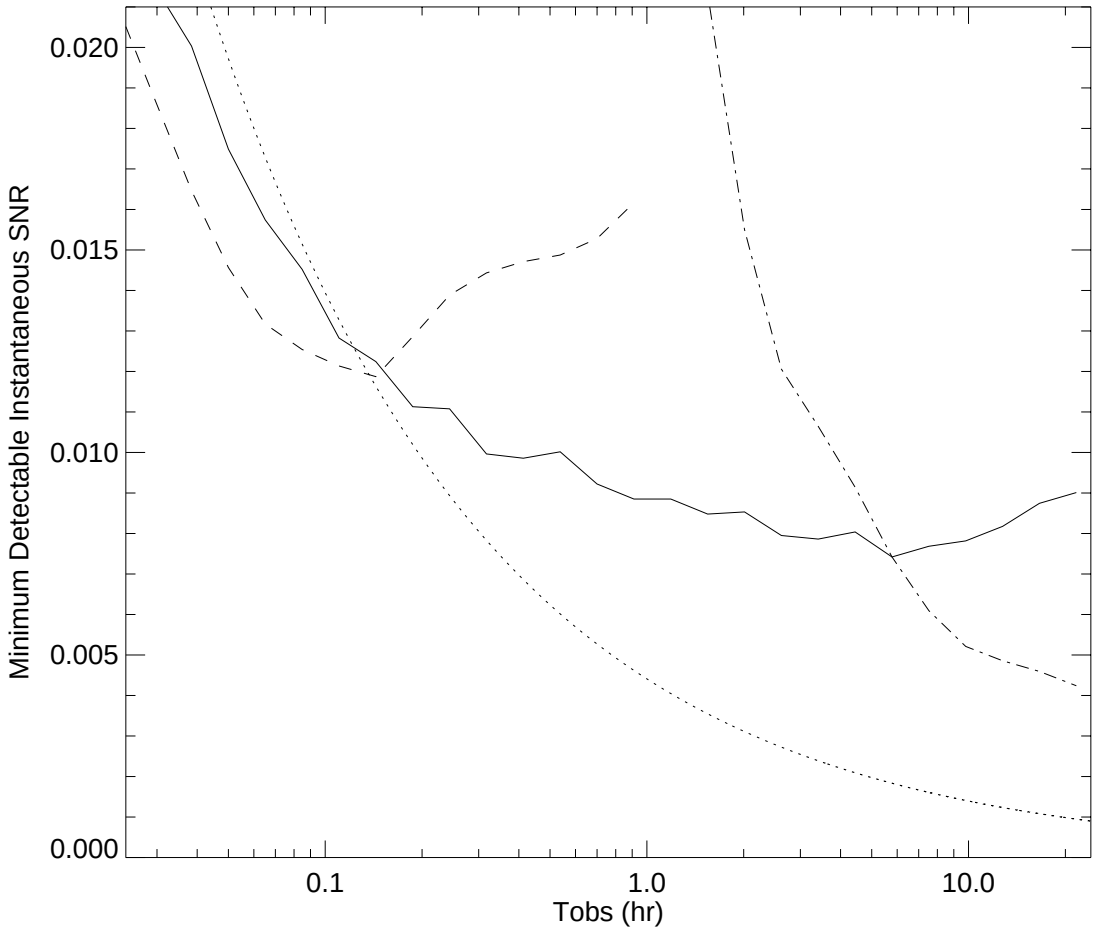

Figure 5.4: Comparison of binary search sensitivities. The curves show the minimum detectable instantaneous signal-to-noise ratio as a function of observation time, for a simulated 2.0 ms pulsar in a 1.0 hr circular orbit with a $0.1 M_{\odot}$ companion. The solid curve represents the DPS sensitivity, the dashed curve is for acceleration searches, and the dot-dashed curve is for the sideband search method. Plotted is the SNR of the fundamental (i.e., lowest harmonic) of our simulated Gaussian waveform (10% FWHM pulse). The simulated acceleration and DPS searches used 4-harmonic sums (see text). For reference, the dotted curve shows the theoretical 8-sigma single-harmonic coherent detection limit.

This is equivalent to an overall (i.e., integrated) SNR of $\sqrt{35.01} = 5.917$. For a given $T_{\rm obs}$, this overall SNR is due to a sum of $(T_{\rm obs}/P_{\rm spin})$ individual pulses, so the 8σ single-pulse SNR is

$$\mathrm{SNR}_{8\sigma} = 5.917 \left(\frac{P_{\rm spin}}{T_{\rm obs}} \right)^{1/2} . \tag{5.13}$$

Thus, longer observations are sensitive to weaker signals, with a coherent detection limit proportional to ${T_{\rm obs}}^{-1/2}$.

Our simulations did not actually use a simple 8-sigma threshold. Rather, we plot the single-pulse fundamental SNR that is 95% likely to result in a detection with better than 1% significance (by significance, we mean the probability that a candidate was produced by random noise, accounting for all the statistical trials involved). Practically speaking, however, our threshold is within 15% of the 8-sigma curve over the entire range of $T_{\rm obs}$ shown in Figure 5.4.

Another important difference between our simulations and those of Ransom (2001) is that we used a more realistic radio pulsar waveform — Gaussian with a 10% full width at half maximum (FWHM), rather than sinusoidal. This benefits the acceleration searches and the DPS searches since these methods can easily take advantage of incoherent harmonic summing. This may be possible with the sideband search method, but would be extremely cumbersome, involving incoherently summing secondary spectra, ideally of different lengths. For a fair comparison between the three methods (i.e., since the sideband method used only a single harmonic), we plot the SNR of only the fundamental harmonic of our simulated pulsar signal. Because of harmonic summing, at the short-$T_{\rm obs}$ end of the figure the acceleration and DPS sensitivities are below the theoretical *single-harmonic* threshold.

We have checked that our acceleration search simulation results are consistent with analytical predictions (Johnston & Kulkarni, 1991; Evans et al., in preparation). As has been pointed out previously (Johnston & Kulkarni, 1991; Ransom, 2001; Jouteux et al., 2002), acceleration searches are optimal when $T_{\rm obs} \approx P_b/7$, a fact verified by our simulations. From Figure 5.4, we see that once $T_{\rm obs} \gtrsim P_b/6$, the DPS method is more sensitive than the acceleration method (even with its optimal observation

length).

We also note that the acceleration method is most effective when $\dot{f}$ is most nearly constant. This fact leads to two difficulties. First, many pulsars in tight binaries are eclipsed by their companions for part of their orbit. These eclipses occur at precisely the orbital phase where the acceleration method is most effective, greatly reducing the probability of detection (when the pulsar passes in front of its companion, it can still be detected). The DPS method on the other hand is most likely to hit at orbital phases where the pulse frequency is most nearly constant (i.e., 90° away from eclipse). This is evident in Figures 5.2 and 5.3. Second, $\dot{f}$ is most nearly constant when $\dot{f}$ is at its extrema. This means that in acceleration searches, pulsars are most likely to be detected when they are experiencing their maximum acceleration. Because acceleration searches are so computationally expensive, they generally cover only a very limited range of accelerations; $|a_{\max}| = 30\,\mathrm{m\,s^{-2}}$ is typical. When a pulsar's maximum acceleration exceeds the search cutoff, it is less likely to be detected. The simulated pulsar of Figure 5.4 has a maximum acceleration of $70.8\,\mathrm{m\,s^{-2}}$, so our simulated acceleration search used $|a_{\max}| = 75\,\mathrm{m\,s^{-2}}$. And indeed, the most significant detection of the pulsar almost always occurs near $|a_{\mathrm{trial}}| \approx 70\,\mathrm{m\,s^{-2}}$. The sensitivity of a standard acceleration search would be considerably worse. Finally, we note that large $a_{\max}$ values can be searched with the DPS method at a far lower cost than moderate $a_{\max}$ acceleration searches.

It is clear from Figure 5.4 that if long, continuous observations are possible, then the sideband search method has the best sensitivity. But to beat the DPS method, sideband searches require $T_{\mathrm{obs}} \gtrsim 5P_b$. Thus, for globular cluster searches, the acceleration method is best suited to detecting binary systems with $P_b \sim$ days, the DPS method should be the method of choice for pulsars with $P_b \sim$ hours, while the sideband method is ideal for anything shorter. The DPS method is very similar in computational complexity to the sideband search method (although no large, parallel or out-of-core FFT is required), and is orders of magnitude more efficient than the acceleration method. With the addition of the DPS method, pulsar searches can now cover the complete range of T_{obs}/P_b with respectable sensitivity.

Table 5.1: Pulsars in M62[a]

Name	Period (ms)	P_b (days)	$a \sin i$ (ls)	$M_C^{\min}$ ($M_\odot$)
J1701-3006A	5.241	3.80	3.48	0.19
J1701-3006B	3.593	0.14	0.25	0.12
J1701-3006C	3.806	0.21	0.19	0.07
J1701-3006D	3.418	1.12	0.98	0.12
J1701-3006E	3.234	0.16	0.07	0.03
J1701-3006F	2.295	0.20	0.06	0.02

[a]Statistics for pulsars A, B, and C are from Possenti et al. (2001)

5.3 Results and Discussion

5.3.1 Search Results

Our M62 search resulted in the discovery of three new binary millisecond pulsars, PSR J1701-3006D, E, and F. Their initial spin and orbital parameters are listed in Table 5.1, along with the parameters of the three previously known pulsars in M62 (D'Amico et al., 2001; Possenti et al., 2001).

Pulsar D was discovered in a straightforward FFT search, i.e., without using any special binary search techniques. This was aided by the fact that the pulsar is relatively bright ($S_{1400} \sim 0.2\,\mathrm{mJy}$) and passes through its descending node during the August 2001 observation. Pulsar A, with its fairly long orbital period, was also detectable without any binary correction. For our acceleration searches, we divided the data into eight 1677.7 s segments. Each segment was searched separately with trial values up to $|a_{\max}| = 60\,\mathrm{m\,s^{-2}}$. The acceleration search did not uncover any new pulsars, but pulsars A, B, C, and D were re-detected. The sideband search also revealed no new pulsars.

Pulsar E was discovered in the August data using the DPS method. It was not detected in the acceleration search of the August data, but was seen in an acceleration search with the follow-up December data. Pulsar F was discovered with the DPS method in the December data. It is visible in a DPS of all five of the December observations, but is not detected by any other method, nor is it visible at all in the

August data. Pulsars E and F therefore establish the validity of the DPS method as a pulsar search tool. (In fact, the DPS method also detected the other 4 pulsars in M62.) In Figures 5.5, 5.6, and 5.7, we show the dedispersed pulse profile and DPS for the discovery observation of each of the three new pulsars. There is strong evidence that pulsar E is eclipsed by its companion's wind for $\sim 15\%$ of its orbit (at 1400 MHz).

We have begun regular observations of M62 at Green Bank, with the goal of determining the timing solutions for the six pulsars. After a year of such observations, we will be able to precisely determine the rest spin period P, period derivative $\dot{P}$, sky position, and the five keplerian orbital parameters (orbital period, projected semimajor axis, eccentricity, longitude of periastron, and epoch of ascending node passage) for each pulsar.

5.3.2 Cluster Dynamics

Until recently, 47 Tuc and M15 were the only two globular clusters known to contain more than 2 pulsars. With the recent flurry of new discoveries, there are now 4 clusters containing 5 or more known pulsars. These are listed in Table 5.2, along with some of their salient parameters. Some pulsar population statistics for these clusters are listed in Table 5.3. When several pulsars are known within a single cluster, their timing characteristics can be used to determine several important properties of the cluster itself. After the pioneering work of Phinney (1992) and Anderson (1993) with M15, similar dynamical studies have been applied to 47 Tuc (Freire et al., 2001) and NGC 6752 (D'Amico et al., 2002). We expect that M62 will provide another laboratory in which these ideas can be applied.

Cluster pulsar $\dot{P}$ measurements are contaminated by acceleration in the gravitational potential of the cluster. The effect of this motion on $\dot{P}$ can be of the same order as, or larger than the intrinsic magnetic spindown of the pulsar. This may make it difficult to determine key derived parameters of the pulsar, such as characteristic age and magnetic field, but it allows us to study the dynamics of the cluster as a whole.

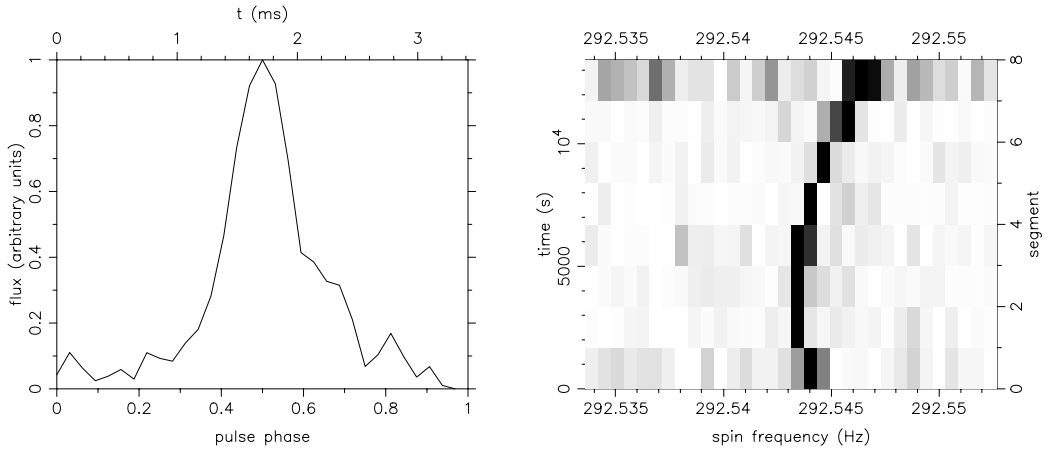

Figure 5.5: Pulse profile and DPS for PSR J1701-3006D.

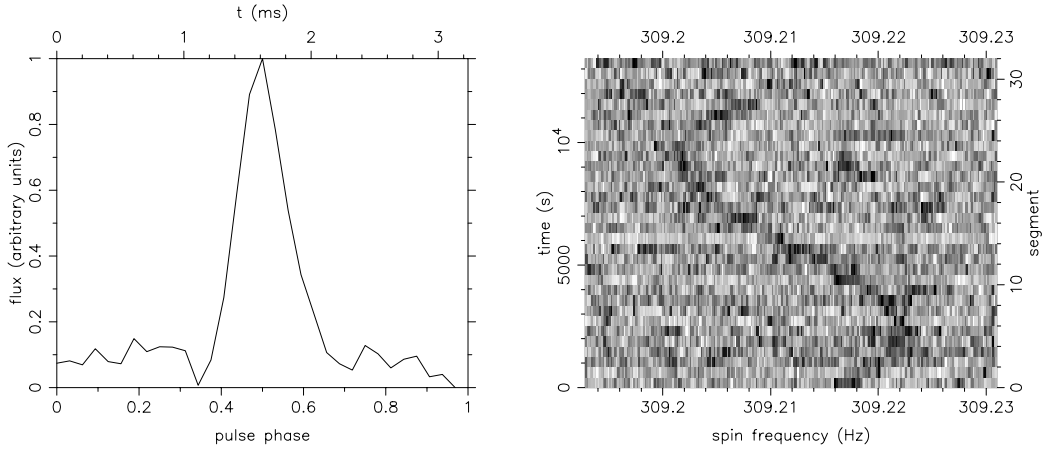

Figure 5.6: Pulse profile and DPS for PSR J1701-3006E.

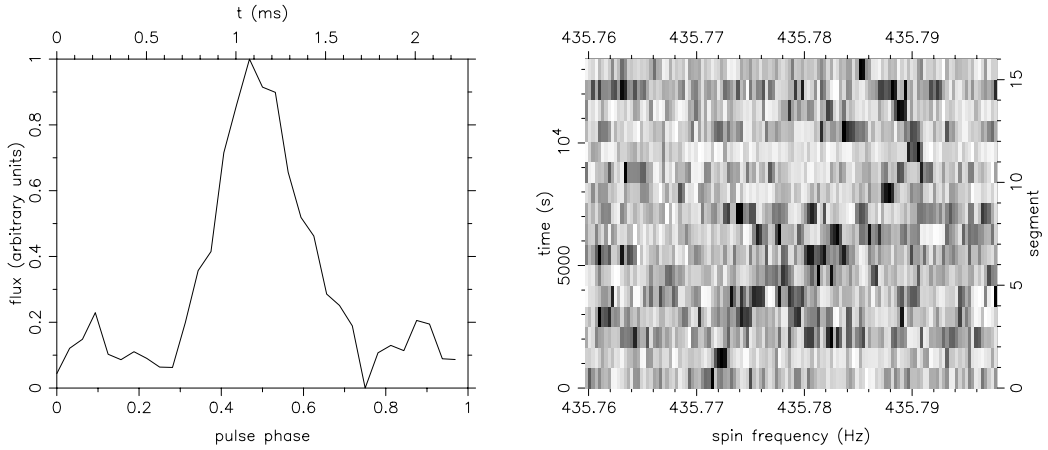

Figure 5.7: Pulse profile and DPS for PSR J1701-3006F.

Table 5.2: Globular clusters with the largest known pulsar populations. $D_\odot$ is the cluster's approximate distance from the sun. r_c is the core radius; a 'c' indicates a core collapsed cluster. Note that M62's classification as core collapsed is somewhat uncertain (Djorgovski, 1993; Djorgovski & King, 1986), and some have argued that the core of 47 Tuc may have been collapsed in the past (de Marchi et al., 1996). M_V is the integrated absolute visual magnitude. ρ_0 is the central luminosity density, and ρ_c is the central mass density derived from pulsar dynamical studies (Harris, 1996; Camilo et al., 2000; Freire et al., 2001; Anderson, 1993; D'Amico et al., 2002).

Cluster	Total Pulsars	$D_\odot$ (kpc)	r_c ($''$)	M_V (mag)	$\log_{10}\rho_0$ ($L_{\odot,V}\,\mathrm{pc}^{-3}$)	ρ_c ($M_\odot\,\mathrm{pc}^{-3}$)
47 Tuc	20	5.0	23	-9.42	4.77	$\gtrsim 4.0\times10^5$
M15	8	10.3	2.2c	-9.17	5.38	$\gtrsim 2.7\times10^6$
M62	6	6.9	11c	-9.19	5.14	
NGC 6752	5	4.0	10c	-7.73	4.91	$\gtrsim 7.1\times10^5$

Table 5.3: Pulsar population statistics of the top globular clusters.

Cluster	Isolated Pulsars	Binary, $P_b < 0.3$ d	Binary, $P_b > 0.3$ d	P_{min} (ms)	P_{median} (ms)	P_{max} (ms)
47 Tuc	7	7	6	2.101	3.590	7.589
M15	7	0	1	4.027	18.636	110.665
M62	0	4	2	2.295	3.506	5.241
NGC 6752	4	0	1	3.266	5.277	9.035

If a negative $\dot{P}$ is observed for a given pulsar, we can conclude that it lies in the more distant half of the cluster and is accelerating towards the earth. Assuming that the acceleration dominates the magnetic braking contribution to $\dot{P}$, we can use the measured acceleration and position to place a lower limit on the cluster's central density and mass-to-light ratio (Phinney, 1992). Using this derived mass-to-light ratio, the observed surface brightness profile, and the pulsar positions, we can determine the maximum line of sight acceleration for the remaining pulsars. This allows us to constrain their intrinsic spindown rates and quantities derived therefrom. We note that 13 of the 25 pulsars previously studied in 47 Tuc, M15, and NGC 6752 exhibit negative period derivatives, and we therefore reasonably expect approximately half of the pulsars in M62 to do the same.

In a dynamically relaxed cluster, the number density of objects of a given mass should follow a Boltzmann distribution, with heavier objects more tightly distributed about the core (Spitzer, 1987). Comparison of the radial distribution of pulsars and the observed distribution of turnoff main sequence stars in M62 should allow a statistical determination of the pulsar masses. Inspection of these distributions should also allow the determination of the dominant mass species in the central region of the cluster, and help constrain the cluster's initial mass function. This will tell us how many neutron stars were formed in M62 and might provide clues as to their retention rate and recycling efficiency. Of course, any pulsars whose characteristic ages do not exceed the core's two body gravitational relaxation time ($\sim 4{\times}10^7$ yr; Harris, 1996) or pulsars that have obviously been ejected from the core must be excluded from this analysis.

Dynamical analyses such as these have been applied to the 7 core pulsars in M15 (Phinney, 1992; Anderson, 1993), the 15 pulsars in 47 Tuc with coherent timing solutions (Freire et al., 2001), and the 3 core pulsars in NGC 6752 (D'Amico et al., 2002) with great success. Analysis of the pulsar positions and accelerations in these clusters has yielded results that are not attainable by other means, and we can expect the application of these techniques to M62 to be just as fruitful.

5.3.3 Cluster Pulsar Demographics

Once there is a large enough sample of populous clusters, we may be able to form clear links between cluster properties and pulsar demographics. With the current census, however, few definitive conclusions can be drawn.

Selection effects have undoubtedly skewed the list of clusters appearing in Table 5.2. Some deserving clusters are certainly absent because their correct dispersion measure has not been found and/or binary selection effects have prevented the full revelation of their pulsar populations. With the advent of new search techniques such as the sideband and DPS methods, binary selection effects should become less of a factor. One might also argue that 47 Tuc appears in first place because of a reverse selection effect — since it is visible at Parkes for long periods when practically nothing else is, it has been observed far more than any other cluster.

The only clear similarity between the clusters listed in Table 5.2 is that they all appear to be fairly dense. Indeed, they are all within the top 80^{th} percentile among the Galactic globular population. However, if we were to include the fifth ranking cluster (the very low-density M13, with 4 MSPs; Anderson, 1993, S. Ransom, private communication), this conclusion would fail to hold. In fact, the recycling rate is probably not a very sensitive function of central density (Phinney, 1996).

A glance at Table 5.3 shows that the pulsar populations in these top four clusters are distinctly different. There are, however, some similarities. Both 47 Tuc and M62 contain roughly equal numbers of so-called "normal" and "short-period" binaries. The "normal" binaries are characterized by orbital periods $P_b \sim 0.4 - 4\,\mathrm{d}$ and companion masses $M_c \sim 0.2 M_\odot$, and are very similar to a number of systems found in the Galactic disk (see, e.g., Phinney & Kulkarni, 1994). The "short-period" binaries (a.k.a. "eclipsing" binaries, although not all of them exhibit eclipses) have $P_b \sim 1.5 - 5.5\,\mathrm{hr}$ and $M_c \sim 0.03 M_\odot$ (Camilo et al., 2000). M62 A and D clearly fall into the former category, while M62 C, E, and F are members of the latter group. M62 B, 47 Tuc V, and 47 Tuc W (along with Ter 5 A; Lyne et al., 1990) have very short orbital periods, but somewhat more massive companions ($\sim 0.1 M_\odot$), and may be

indicative of an emerging class of their own.

These pulsar systems are probably formed when an NS exchanges into a primordial main sequence (MS) binary, supplanting the lighter star. The longer-period systems are presumably the result of stable mass transfer, once the $\sim 1 M_{\odot}$ companion evolves off the MS and overflows its Roche lobe. Accretion from the giant's expanded envelope spins up the NS (Joss & Rappaport, 1983; Paczynski, 1983; Webbink et al., 1983). The short-period systems cannot be produced by direct exchange of an NS into an extremely tight primordial binary, because the liberated binding energy would most likely eject the new NS-MS pair from the weak gravitational field of the cluster. Rather, the short-period systems may have resulted from the exchange of an NS into a binary system with a more massive ($\sim 1-3 M_{\odot}$) MS companion (Rasio et al., 2000). In this case, when the secondary enters the giant stage, a common envelope phase can occur, making the orbit more compact. Further orbital decay due to gravitational radiation may be followed by another episode of mass transfer once the white dwarf companion again overflows its Roche lobe, resulting in the very low mass, short-period systems observed. Clearly the initial exchanges that produce such systems must occur within the first $\sim 10^9$ yr of the cluster's life, when massive MS stars still exist. This fact indicates that 47 Tuc and M62 may have had similar early dynamical histories.

But 47 Tuc also contains a roughly equal number of isolated pulsars, whereas M62 has none, despite the fact that there are no systematic selection effects against detecting solitary pulsars. In fact, M62 is currently the only cluster with more than two pulsars, not containing an isolated pulsar. Based on the 47 Tuc population, we might expect ~ 3 isolated pulsars in M62. Detecting zero would be only a $\lesssim 2\sigma$ chance occurrence. In other words, we cannot confidently reject the hypothesis that the two populations are drawn from the same underlying distribution.

If the disparity is real, however, it is not easily explained. The large emerging population of short-period binaries (Deich et al., 1993; Camilo et al., 2000; Ransom et al., 2001; D'Amico et al., 2001) argues against companion ablation on short timescales, possibly indicating that these companions are no longer degenerate (van den Heuvel & van Paradijs, 1988) and may never completely evaporate (Robinson et al., 1995).

Companion destruction scenarios might therefore require systems even more compact than those currently observed, but it is unclear why these may have formed in 47 Tuc and not in M62. Complete tidal disruption of the mass donor is also an unlikely explanation for the isolated pulsars in 47 Tuc. In this case, the "companion" is disrupted during tidal capture and forms an accretion disk around the NS (Krolik, 1984). Significant spin-up is not expected from such systems, however (Verbunt et al., 1987), and this process is unlikely to form true MSPs. It appears that the most likely formation channel for the isolated pulsars in 47 Tuc involves collisions resulting from three-body encounters (Verbunt et al., 1987; Sigurdsson & Phinney, 1995). The lack of isolated pulsars in M62 might indicate that these catastrophic binary-single star encounters are less common (despite M62's larger density), or preferentially eject the NSs from the cluster for some unknown reason.

Similarities are also apparent between M15 and NGC 6752. These post-core-collapse (PCC) clusters each contain a single binary pulsar system, both of which seem to have been ejected from the central regions of their respective clusters. This is indicative of a high rate of encounters in these clusters' cores. This fact is also supported by the observation that the remaining pulsars in these clusters are all isolated, indicating that they have each likely experienced at least one exchange or collision in the past.

The differences between these two clusters are striking, however. The derived central mass-to-light ratio of NGC 6752 is $3-4$ times larger than that of M15, suggesting a much flatter IMF, and a much larger concentration of compact objects in the core of NGC 6752 than in M15. PSR B2127+11C in M15 is a highly eccentric double NS binary, likely the result of an interaction in which the pulsar exchanged into a binary containing its current companion (Prince et al., 1991). The recoil associated with this exchange ejected the new system to the cluster outskirts. Despite the fact that PSR J1911-5958A in NGC 6752 is a more normal MSP binary ($P = 3.3\,\mathrm{ms}$, eccentricity $e < 10^{-5}$, companion mass $\sim 0.2 M_{\odot}$), its dynamical history may have been quite unusual. Colpi et al. (2002) have suggested that a scattering event involving a massive black hole binary at the cluster's center ejected the system to its

current location, over 3 half mass radii from the cluster center. In any case, these two binary systems are markedly different.

The isolated pulsar populations in M15 and NGC 6752 are also very different. NGC 6752 contains only bona fide MSPs, while the distribution of periods in M15 is much broader. The long-period pulsars may be the result of destructive tidal capture (Tavani, 1992), but the shorter-period pulsars are most likely the result of spin up in a binary system, followed by ionization or a destructive stellar collision.

Clearly, a larger sample of highly populated clusters will be required before the study of cluster pulsar demographics can move beyond speculation. There are hints of similarities and differences among the current top clusters, but their explanation is still elusive. For example, it is unclear why M62 should be so similar to 47 Tuc and so different from the other PCC clusters in Table 5.2.

This vexing situation may get worse before it gets better, but one certainty is that the list of highly populated clusters will indeed keep growing, perhaps by leaps and bounds. The recent resurgence of cluster pulsar discoveries will undoubtedly continue, aided by new and improved observing hardware and search techniques.

We are grateful to the NRAO-Green Bank staff for all of their support. We also thank Scott Ransom and Ingrid Stairs for freely providing information on unpublished search results. The National Radio Astronomy Observatory is a facility of the National Science Foundation, operated under cooperative agreement by Associated Universities, Inc. Access to the Hewlett-Packard V2500 computer, located at the California Institute of Technology, was provided by the Center for Advanced Computing Research. This work was supported by NSF grants 0005-1-000024 and 00040-1-000262.

Appendix A

PSR J1807-2459B and PSR B1820-30C

In this brief appendix, we announce the discovery of two additional pulsars in our GBT globular cluster search campaign (see Chapter 5). These are PSR J1807-2459B and PSR B1820-30C in the extremely dense, core-collapsed clusters NGC 6544 and NGC 6624, respectively. Both pulsars appear to be isolated. Initial values for their basic parameters are given in Table A.1, along with the properties of the other known pulsars in the two clusters.

We have not yet attempted to confirm either of these new pulsars in follow-up data, but we regard the detections with very high confidence since they both exhibit the proper amount of dispersion for their respective clusters. Pulse profiles, both dedispersed and dispersed, are shown in Figures A.1 and A.2. Assuming that PSR B1820-30C is real and is truly associated with NGC 6624, it is of particular interest. Slow pulsars are quite rare in globular clusters — there are only 6 known

Table A.1: Pulsars in NGC 6544 and NGC 6624.

Name	Cluster	Period (ms)	DM ($\mathrm{cm}^{-3} - \mathrm{pc}$)	P_b (days)	$a \sin i$ (ls)	Ref.
J1807-2459A	NGC 6544	3.059	134	0.071	0.012	A,B
J1807-2459B	NGC 6544	4.186	134	-	-	-
B1820-30A	NGC 6624	5.440	87	-	-	C
B1820-30B	NGC 6624	378.6	87	-	-	C
B1820-30C	NGC 6624	405.9	87	-	-	-

A Ransom et al. 2001
B D'Amico et al. 2001
C Biggs et al. 1994

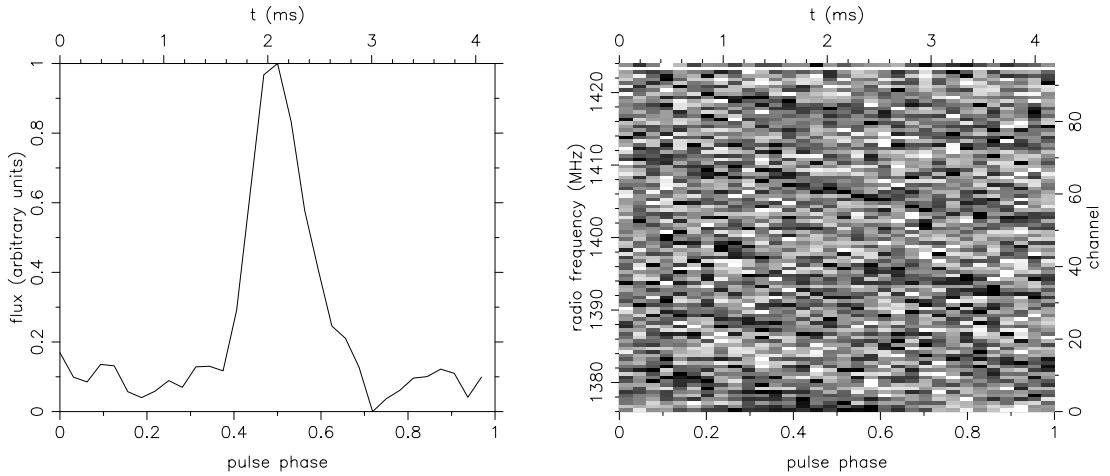

Figure A.1: Dedispersed and dispersed pulse profiles for PSR J1807-2459B.

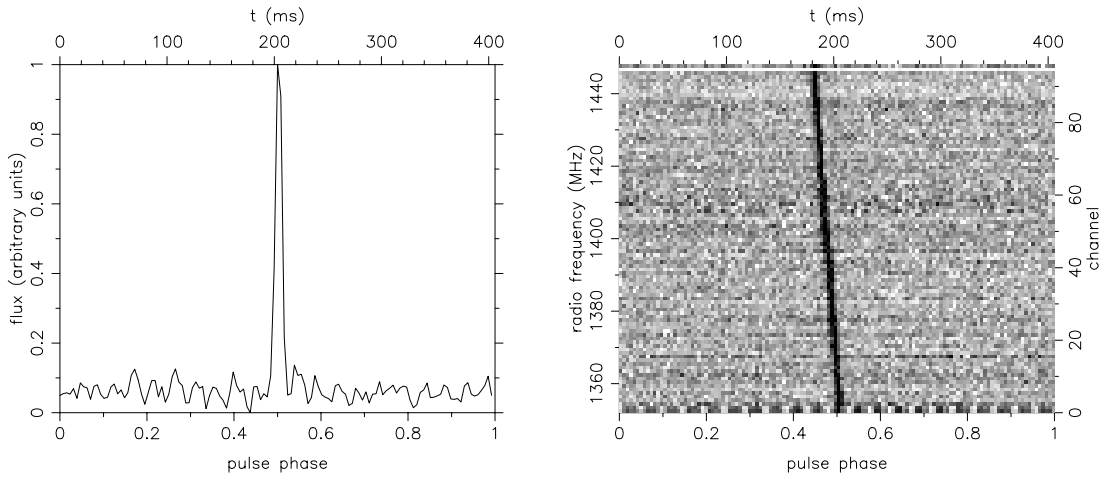

Figure A.2: Dedispersed and dispersed pulse profiles for PSR B1820-30C.

cluster pulsars with $P \geq 100\,\mathrm{ms}$ (including PSR B1820-30C) — and NGC 6624 now contains two.

Our searches of these clusters are not complete. In particular, sideband searches in these two clusters may hold some promise. NGC 6544 contains the radio pulsar with the second shortest orbital period, and NGC 6624 contains the shortest period neutron star binary (the LMXB 4U 1820-30 at 11.4 minutes; Stella et al. 1987).

Bibliography

Alcock, C. & Illarionov, A. 1980, ApJ, 235, 534

Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, Nature, 300, 728

Anderson, S. B. 1993, PhD thesis, California Institute of Technology

Anderson, S. B., Gorham, P. W., Kulkarni, S. R., Prince, T. A., & Wolszczan, A. 1990, Nature, 346, 42

Andersson, N. 1998, ApJ, 502, 708+

Asai, K., Dotani, T., Kunieda, H., & Kawai, N. 1996a, Publ. Astron. Soc. Japan, 48, L27

Asai, K., Dotani, T., Mitsuda, K., Hoshi, R., Vaughan, B., Tanaka, Y., & Inoue, H. 1996b, Publ. Astron. Soc. Japan, 48, 257

Baade, W. & Zwicky, F. 1934, Proc. Nat. Acad Sci. U.S.A., 20, 254

Backer, D. C., Dexter, M. R., Zepka, A., Ng, D., Werthimer, D. J., Ray, P. S., & Foster, R. S. 1997, PASP, 109, 61

Backer, D. C., Foster, R. S., & Sallmen, S. 1993, Nature, 365, 817+

Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, Nature, 300, 615

Baring, M. G. & Harding, A. K. 1997, in AIP Conf. Proc. 410: Proceedings of the Fourth Compton Symposium, ed. C. D. Dermer, M. S. Strickman, & J. D. Kurfess (American Institute of Physics), 638+

Baring, M. G. & Harding, A. K. 1998, ApJ, 507, L55

Bertsch, D. L., Brazier, K. T. S., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Kwok, P. W., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., von Montigny, C., Michelson, P. F., Nolan, P. L., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Sreekumar, P., & Thompson, D. J. 1992, Nature, 357, 306

Bhattacharya, D. 1998, in NATO ASI C Proc. 515: The Many Faces of Neutron Stars., 103–+

Bhattacharya, D., Akyuz, A., Case, G., Dixon, D., & Zych, A. 1997, in AIP Conf. Proc. 410: Proceedings of the Fourth Compton Symposium, ed. C. D. Dermer, M. S. Strickman, & J. D. Kurfess (American Institute of Physics), 1137+

Bhattacharya, D. & Srinivasan, G. 1995, in X-Ray Binaries, ed. W. Lewin, J. van Paradijs, & E. van den Heuvel, Vol. 1 (Cambridge University Press), 495

Bhattacharya, D. & van den Heuvel, E. P. J. 1991, Physics Reports, 203, 1

Biggs, J. D. 1992, ApJ, 394, 574

Biggs, J. D., Bailes, M., Lyne, A. G., Goss, W. M., & Fruchter, A. S. 1994, MNRAS, 267, 125+

Bildsten, L. 1998, ApJ, 501, L89

Bildsten, L., Chakrabarty, D., Chiu, J., Finger, M. H., Koh, D. T., Nelson, R. W., Prince, T. A., Rubin, B. C., Scott, D. M., Stollberg, M., Vaughan, B. A., Wilson, C. A., & Wilson, R. B. 1997, ApJS, 113, 367

Bildsten, L., Salpeter, E. E., & Wasserman, I. 1992, ApJ, 384, 143

Bildsten, L. & Ushomirsky, G. 2000, ApJ, 529, L33

Blandford, R. D., Hewish, A., Lyne, A. G., & Mestel, L., eds. 1993, Pulsars as Physics Laboratories (Oxford University Press; Philosophical Transactions of the Royal Society of London; 341)

Boyd, P. T., van Citters, G. W., Dolan, J. F., Wolinski, K. G., Percival, J. W., Bless, R. C., Elliot, J. L., Nelson, M. J., & Taylor, M. J. 1995, ApJ, 448, 365

Bracewell, R. N. 2000, The Fourier transform and its applications (Boston : McGraw Hill (McGraw-Hill series in electrical and computer engineering. Circuits and systems))

Brady, P. R. & Creighton, T. 1998, gr-qc/9812014

Brazier, K. T. S., Kanbach, G., Carraminana, A., Guichard, J., & Merck, M. 1996, MNRAS, 281, 1033

Brazier, K. T. S., Reimer, O., Kanbach, G., & Carraminana, A. 1998, MNRAS, 295, 819

Brown, E. F., Bildsten, L., & Rutledge, R. E. 1998, ApJ, 504, l95

Buccheri, R., Sacco, B., & Ozel, M. E. 1987, A&A, 175, 353+

Burderi, L., Possenti, A., D'Antona, F., Di Salvo, T., Burgay, M., Stella, L., Menna, M. T., Iaria, R., Campana, S., & D'Amico, N. 2001, ApJ, 560, L71

Burke, B. F. & Graham-Smith, F. 1997, An Introduction to Radio Astronomy (Cambridge University Press)

Burns, W. R. & Clark, B. G. 1969, A&A, 2, 280+

Cadwell, B. J. 1997, Ph.D. Thesis, Pennsylvania State University

Camilo, F., Lorimer, D. R., Freire, P., Lyne, A. G., & Manchester, R. N. 2000, ApJ, 535, 975

Camilo, F., Lyne, A. G., Bell, J. F., Sheppard, D., Manchester, R. N., D'Amico, N., & Kaspi, V. 1997, in American Astronomical Society Meeting, Vol. 191, 11113+

Camilo, F., Nice, D. J., Shrauner, J. A., & Taylor, J. H. 1996a, ApJ, 469, 819+

Camilo, F., Nice, D. J., & Taylor, J. H. 1996b, ApJ, 461, 812+

Campana, S., Stella, L., Mereghetti, S., Colpi, M., Tavani, M., Ricci, D., Fiume, D. D., & Belloni, T. 1998, ApJ, 499, L65

Chadwick, J. 1932, Proc. R. Soc. London A, 136, 692

Chakrabarty, D., Bildsten, L., Finger, M. H., Grunsfeld, J. M., Koh, D. T., Nelson, R. W., Prince, T. A., Vaughan, B. A., & Wilson, R. B. 1997, ApJ, 481, L101

Chakrabarty, D. & Morgan, E. H. 1998, Nature, 394, 346

Chandler, A. M., Koh, D. T., Lamb, R. C., Macomb, D. J., Mattox, J. R., Prince, T. A., & Ray, P. S. 2001, ApJ, 556, 59

Chandler, A. M. & Rutledge, R. E. 2000, ApJ, 545, 1000

Chevalier, C. & Ilovaisky, S. A. 1991, A&A, 251, L11

—. 1998, I.A.U. Circ., 6806

Chevalier, C., Ilovaisky, S. A., Leisy, P., & Patat, F. 1999, A&A, 347, L51

Colpi, M., Possenti, A., & Gualandris, A. 2002, ApJ, 570, L85

Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1994, ApJ, 424, 823

Cordes, J. M. & Chernoff, D. F. 1997, ApJ, 482, 971+

—. 1998, ApJ, 505, 315

Cordes, J. M., Ryan, M., Weisberg, J. M., Frail, D. A., & Spangler, S. R. 1991, Nature, 354, 121

Crawford, F., Kaspi, V. M., & Bell, J. F. 2000, AJ, 119, 2376

D'Amico, N. 2000, in ASP Conf. Ser. 202: IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond, 27+

D'Amico, N., Lyne, A. G., Manchester, R. N., Possenti, A., & Camilo, F. 2001, ApJ, 548, L171

D'Amico, N., Possenti, A., Fici, L., Manchester, R. N., Lyne, A. G., Camilo, F., & Sarkissian, J. 2002, ApJ, 570, L89

Davies, J. G., Lyne, A. G., & Seiradakis, J. H. 1972, Nature, 240, 229+

—. 1973, Nature Phys. Sci., 244, 84+

Davies, R. D., Elliott, K. H., Goudis, C., Meaburn, J., & Tebbutt, N. J. 1978, A&AS, 31, 271

de Marchi, G., Paresce, F., Stratta, M. G., Gilliland, R. L., & Bohlin, R. C. 1996, ApJ, 468, L51

Deich, W. T. S., Middleditch, J., Anderson, S. B., Kulkarni, S. R., Prince, T. A., & Wolszczan, A. 1993, ApJ, 410, L95

Dermer, C. D., Strickman, M. S., & Kurfess, J. D., eds. 1997, AIP Conf. Proc. 410: Proceedings of the Fourth Compton Symposium (American Institute of Physics)

Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, ApJ, 294, L25

Djorgovski, S. & King, I. R. 1986, ApJ, 305, L61

Djorgovski, S. G. 1993, in ASP Conf. Ser. 50: Structure and Dynamics of Globular Clusters, 373+

Durdin, J. M., Large, M. I., Little, A. G., Manchester, R. N., Lyne, A. G., & Taylor, J. H. 1979, MNRAS, 186, 39

Edwards, R. T., van Straten, W., & Bailes, M. 2001, ApJ, 560, 365

Esposito, J. A., Bertsch, D. L., Chen, A. W., Dingus, B. L., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Lin, Y. C., Mayer-Hasselwander, H. A., McDonald, L. M., Michelson, P. F., von Montigny, C., Mukherjee, R., Nolan, P. L., Reimer, O. L., Schneid, E., Sreekumar, P., Thompson, D. J., Tompkins, W. F., & Willis, T. D. 1999, ApJS, 123, 203

Esposito, J. A., Hunter, S. D., Kanbach, G., & Sreekumar, P. 1996, ApJ, 461, 820+

Fichtel, C. E., Bertsch, D. L., Chiang, J., Dingus, B. L., Esposito, J. A., Fierro, J. M., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Kwok, P. W., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., McDonald, L., Michelson, P. F., von Montigny, C., Nolan, P. L., Pinkau, K., Radecke, H.-D., Rothermel, H., Sommer, M., Sreekumar, P., Schneid, E. J., Thompson, D. J., & Willis, T. 1994, ApJS, 94, 551

Fierro, J. M. 1996, PhD thesis, Stanford University.

Fierro, J. M., Arzoumanian, Z., Bailes, M., Bell, J. F., Bertsch, D. L., Brazier, K. T. S., Chiang, J., D'Amico, N., Dingus, B. L., Esposito, J. A., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Johnston, S., Kanbach, G., Kaspi, V. M., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Manchester, R. N., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nel, H. I., Nice, D., Nolan, P. L., Schneid, E. J., Shriver, S. K., Sreekumar, P., Taylor, J. H., Thompson, D. J., & Willis, T. D. 1995, ApJ, 447, 807+

Fierro, J. M., Bertsch, D. L., Brazier, K. T. S., Chiang, J., D'Amico, N., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Johnston, S., Kanbach, G., Kaspi, V. M., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Manchester, R. N., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nolan, P. L., Schneid, E. J., & Thompson, D. J. 1993, ApJ, 413, L27

Foster, R. S., Cadwell, B. J., Wolszczan, A., & Anderson, S. B. 1995, ApJ, 454, 826+

Freire, P. C., Camilo, F., Lorimer, D. R., Lyne, A. G., Manchester, R. N., & D'Amico, N. 2001, MNRAS, 326, 901

Galloway, D. K., Chakrabarty, D., Morgan, E. H., & Remillard, R. A. 2002, ApJ, 576, L137

Gold, T. 1968, Nature, 218, 731

Goldreich, P. & Julian, W. H. 1969, ApJ, 157, 869

Graham, D. A., Haslam, C. G. T., Salter, C. J., & Wilson, W. E. 1982, A&A, 109, 145

Groth, E. J. 1975, ApJS, 29, 285

Halpern, J. P. & Holt, S. S. 1992, Nature, 357, 222

Harnden, F. R. & Seward, F. D. 1984, ApJ, 283, 279

Harris, W. E. 1996, AJ, 112, 1487+

Hartman, R. C., Bertsch, D. L., Bloom, S. D., Chen, A. W., Deines-Jones, P., Esposito, J. A., Fichtel, C. E., Friedlander, D. P., Hunter, S. D., McDonald, L. M., Sreekumar, P., Thompson, D. J., Jones, B. B., Lin, Y. C., Michelson, P. F., Nolan, P. L., Tompkins, W. F., Kanbach, G., Mayer-Hasselwander, H. A., Mücke, A., Pohl, M., Reimer, O., Kniffen, D. A., Schneid, E. J., von Montigny, C., Mukherjee, R., & Dingus, B. L. 1999, ApJS, 123, 79

Hewish, A., Bell, J., Pilkington, J., Scott, P., & Collins, R. 1968, Nature, 217, 709

Hoffman, C. M., Sinnis, C., Fleury, P., & Punch, M. 1999, Reviews of Modern Physics, 71, 897

Hulse, R. A. & Taylor, J. H. 1974, ApJ, 191, L59

Illarionov, A. F. & Sunyaev, R. A. 1975, A&A, 39, 185

Jackson, J. D. 1975, Classical Electrodynamics, 2nd edn. (John Wiley & Sons)

Jacoby, B. A., Chandler, A. M., Backer, D. C., Anderson, S. B., & Kulkarni, S. R. 2002, I.A.U. Circ., 7783, 1+

Jenet, F. A. 2001, PhD thesis, California Institute of Technology

Johnston, H. M. & Kulkarni, S. R. 1991, ApJ, 368, 504

Johnston, S. & Bailes, M. 1991, MNRAS, 252, 277

Joss, P. C. & Rappaport, S. A. 1983, Nature, 304, 419

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

Kaaret, P. & Cottam, J. 1996, ApJ, 462, L35

Kanbach, G., Arzoumanian, Z., Bertsch, D. L., Brazier, K. T. S., Chiang, J., Fichtel, C. E., Fierro, J. M., Hartman, R. C., Hunter, S. D., Kniffen, D. A., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nel, H. I., Nice, D., Nolan, P. L., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Sreekumar, P., Taylor, J. H., & Thompson, D. J. 1994, A&A, 289, 855

Kanbach, G., Bertsch, D. L., Favale, A., Fichtel, C. E., Hartman, R. C., Hofstadter, R., Hughes, E. B., Hunter, S. D., Hughlock, B. W., Kniffen, D. A., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., von Montigny, C., Nolan, P. L., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Thompson, D. J., & Walker, A. H. 1989, in Proceedings of the Gamma Ray Observatory Science Workshop, ed. W. N. Johnson (NASA Goddard Space Flight Center), 2-1

Kaspi, V. M., Lackey, J. R., Mattox, J., Manchester, R. N., Bailes, M., & Pace, R. 2000, ApJ, 528, 445

Kaspi, V. M., Manchester, R. N., Siegman, B., Johnston, S., & Lyne, A. G. 1994, ApJ, 422, L83

King, A. R. 1993, in Galactic High-Energy Astrophysics. High-Accuracy Timing and Positional Astronomy. Lectures Held at the Astrophysics School IV Organized by the European Astrophysics Doctoral Network (EADN) in Graz, Austria, 19-31 August 1991, Edited by Jan van Paradijs and Hans M. Maitzen. Springer-Verlag Berlin Heidelberg New York. Also Lecture Notes in Physics, volume 418, p.161, 161–+

Kirshner, R. P., Tull, T. R., & Parker, R. A. R. 1978, A&AS, 31, 261

Kouveliotou, C., Dieters, S., Strohmayer, T., van Paradijs, J., Fishman, G. J., Meegan, C. A., Hurley, K., Kommers, J., Smith, I., Frail, D., & Murakami, T. 1998, Nature, 393, 235

Kouveliotou, C., Fishman, G. J., Meegan, C. A., Paciesas, W. S., van Paradijs, J., Norris, J. P., Preece, R. D., Briggs, M. S., Horack, J. M., Pendleton, G. H., & Green, D. A. 1994, Nature, 368, 125

Kouveliotou, C., Strohmayer, T., Hurley, K., van Paradijs, J., Finger, M. H., Dieters, S., Woods, P., Thompson, C., & Duncan, R. C. 1999, ApJ, 510, L115

Kramer, M., Wex, N., & Wielebinski, R., eds. 2000, Pulsar Astronomy - 2000 and beyond, IAU Symposium No. 177 (Sheridan)

Kraus, J. D. 1986, Radio astronomy (Powell, Ohio: Cygnus-Quasar Books, 1986)

Krolik, J. H. 1984, ApJ, 282, 452

Kuiper, L., Hermsen, W., Verbunt, F., Thompson, D. J., Stairs, I. H., Lyne, A. G., Strickman, M. S., & Cusumano, G. 2000, A&A, 359, 615

Kulkarni, S. R. & Anderson, S. B. 1996, in IAU Symp. 174: Dynamical Evolution of Star Clusters: Confrontation of Theory and Observations, Vol. 174, 181+

Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, Nature, 220, 340

Large, M. L. & Vaughan, A. E. 1971, MNRAS, 151, 277+

Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Kahn, S., Sutherland, P. G., & Grindlay, J. E. 1983, ApJ, 266, 160

Leahy, D. A., Naranan, S., & Singh, K. P. 1986, MNRAS, 220, 501

Levin, Y. & Ushomirsky, G. 2001, MNRAS, 324, 917

Lommen, A. N., Zepka, A., Backer, D. C., McLaughlin, M., Cordes, J. M., Arzoumanian, Z., & Xilouris, K. 2000, ApJ, 545, 1007

Lundgren, S. C. 1994, Ph.D. Thesis, Cornell University

Lundgren, S. C., Zepka, A. F., & Cordes, J. M. 1995, ApJ, 453, 419

Lyne, A. G., Brinklow, A., Middleditch, J., Kulkarni, S. R., & Backer, D. C. 1987, Nature, 328, 399

Lyne, A. G. & Graham-Smith, F. 1998, Pulsar astronomy (Cambridge, U.K.; New York : Cambridge University Press, (Cambridge astrophysics series ; 31) ISBN 0521594138)

Lyne, A. G., Johnston, S., Manchester, R. N., Staveley-Smith, L., & D'Amico, N. 1990, Nature, 347, 650

Lyne, A. G. & Manchester, R. N. 1988, MNRAS, 234, 477

Lyne, A. G., Mankelow, S. H., Bell, J. F., & Manchester, R. N. 2000, MNRAS, 316, 491

Lyne, A. G., Pritchard, R. S., Graham-Smith, F., & Camilo, F. 1996, Nature, 381, 497

Lyne, A. G., Pritchard, R. S., & Smith, F. G. 1988, MNRAS, 233, 667

Malofeev, V. M. & Malov, O. I. 1997, Nature, 389, 697

Manchester, R. N. & Taylor, J. H. 1977, Pulsars (San Francisco: Freeman)

Markwardt, C. B., Strohmayer, T. E., & Swank, J. H. 1999, ApJ, 512, L125

Markwardt, C. B., Swank, J. H., Strohmayer, T. E., Zand, J. J. M. i., & Marshall, F. E. 2002, ApJ, 575, L21

Mattox, J. R., Koh, D. T., Lamb, R. C., Macomb, D. J., Prince, T. A., & Ray, P. S. 1996, A&AS, 120, C95

McLaughlin, M., Arzoumanian, Z. A., & Cordes, J. M. 2000, in Pulsar Astronomy - 2000 and beyond, ed. M. Kramer, N. Wex, & R. Wielebinski, IAU Symposium No. 177 (Sheridan), 41

McLaughlin, M. A. & Cordes, J. M. 2000, ApJ, 538, 818

McLaughlin, M. A., Mattox, J. R., Cordes, J. M., & Thompson, D. J. 1996, ApJ, 473, 763+

Melatos, A. 1999, ApJ, 519, L77

Menou, K., Esin, A. A., Narayan, R., Garcia, M. R., Lasota, J. P., & McClintock, J. E. 1999, ApJ, 520, 276

Mereghetti, S. 2001, in The Neutron Star–Black Hole Connection, ed. C. Kouveliotou, J. van Paradijs, & J. Ventura, NATO ASI C No. 567 (Dordrecht: Kluwer), 351

Middleditch, J. 1976, PhD thesis, University of California, Berkeley

Middleditch, J., Deich, W., & Kulkarni, S. 1993, in Isolated Pulsars, ed. K. A. van Riper, R. Epstein, & C. Ho (Cambridge University Press), 372–379

Moffett, D. A. & Hankins, T. H. 1996, ApJ, 468, 779+

Nel, H. I., Arzoumanian, Z., Bailes, M., Brazier, K. T. S., D'Amico, N., Esposito, J. A., Fichtel, C. E., Fierro, J. M., Hunter, S. D., Johnston, S., Kanbach, G., Kaspi, V. M., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Manchester, R. N., Mattox, J. R., Mayer-Hasselwander, H. A., Merck, M., Michelson, P. F., Nice, D. J., Nolan, P. L., Ramanamurthy, P. V., Taylor, J. H., Thompson, D. J., & Westbrook, C. 1996, ApJ, 465, 898+

Nice, D. J. & Sayer, R. W. 1997, ApJ, 476, 261

Nolan, P. L., Arzoumanian, Z., Bertsch, D. L., Chiang, J., Fichtel, C. E., Fierro, J. M., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Kwok, P. W., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny,

C., Nel, H. I., Nice, D., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Sreekumar, P., Taylor, J. H., & Thompson, D. J. 1993, ApJ, 409, 697

Pacini, F. 1967, Nature, 216, 567

Paczynski, B. 1983, Nature, 304, 421

Percival, J. W., Biggs, J. D., Dolan, J. F., Robinson, E. L., Taylor, M. J., Bless, R. C., Elliot, J. L., Nelson, M. J., Ramseyer, T. F., van Citters, G. W., & Zhang, E. 1993, ApJ, 407, 276

Phinney, E. S. 1992, Philos. Trans. R. Soc. London A, 341, 39

Phinney, E. S. 1996, in ASP Conf. Ser. 90: The Origins, Evolutions, and Destinies of Binary Stars in Clusters, 163+

Phinney, E. S. & Kulkarni, S. R. 1994, ARA&A, 32, 591

Possenti, A., D'Amico, N., Manchester, R., Sarkissian, J., Lyne, A., & Camilo, F. 2001, astro-ph/0108343

Press, W., Flannery, B., Teukolsky, S., & Vetterling, W. 1995, Numerical Recipes in C (Cambridge University Press)

Prince, T. A., Anderson, S. B., Kulkarni, S. R., & Wolszczan, A. 1991, ApJ, 374, L41

Rajagopal, M. & Romani, R. W. 1996, ApJ, 461, 327

Ramanamurthy, P. V., Bertsch, D. L., Dingus, B. L., Esposito, J. A., Fichtel, C. E., Fierro, J. M., Hunter, S. D., Kanbach, G., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Mattox, J. R., Mayer-Hasselwander, H. A., Merck, M., Michelson, P. F., von Montigny, C., Mukherjee, R., Nolan, P. L., & Thompson, D. J. 1995, ApJ, 447, L109

Ramanamurthy, P. V., Fichtel, C. E., Kniffen, D. A., Sreekumar, P., & Thompson, D. J. 1996, ApJ, 458, 755

Ramsey, B. D., Austin, R. A., & Decher, R. 1994, Space Science Reviews, 69, 139

Ransom, S. M. 2001, PhD thesis, Harvard University

Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, AJ, 124, 1788

Ransom, S. M., Greenhill, L. J., Herrnstein, J. R., Manchester, R. N., Camilo, F., Eikenberry, S. S., & Lyne, A. G. 2001, ApJ, 546, L25

Rappaport, S. & Joss, P. C. 1977, Nature, 266, 683

Rasio, F. A., Pfahl, E. D., & Rappaport, S. 2000, ApJ, 532, L47

Ray, P. S., Thorsett, S. E., Jenet, F. A., van Kerkwijk, M. H., Kulkarni, S. R., Prince, T. A., Sandhu, J. S., & Nice, D. J. 1996, ApJ, 470, 1103+

Regimbau, T. & de Freitas Pacheco, J. A. 2001, A&A, 374, 182

Roberts, M. S. E. & Romani, R. W. 1998, ApJ, 496, 827+

Robinson, C., Lyne, A. G., Manchester, R. N., Bailes, M., D'Amico, N., & Johnston, S. 1995, MNRAS, 274, 547

Roche, P., Chakrabarty, D., Morales-Rueda, L., Hynes, R., Slivan, S. M., Simpson, C., & Hewett, P. 1998, I.A.U. Circ., 6885

Romani, R. W. 1987, ApJ, 313, 718

Romani, R. W. & Yadigaroglu, I. A. 1995, ApJ, 438, 314

Rutledge, R. E., Bildsten, L., Brown, E. F., Pavlov, G. G., & Zavlin, V. E. 1999, ApJ, 514, 945

Shahbaz, T., Thorstensen, J. R., Charles, P. A., & Sherman, N. D. 1998, MNRAS, 296, 1004

Shapiro, S. L. & Teukolsky, S. A. 1983, Black Holes, White Dwarfs and Neutron Stars. The Physics of Compact Objects (New York: Wiley–Interscience)

Shitov, Y. P. 1999, I.A.U. Circ., 7110, 2+

Sigurdsson, S. & Phinney, E. S. 1995, ApJS, 99, 609

Smith, D. A., Morgan, E. H., & Bradt, H. 1997, ApJ, 479, L137

Spitzer, L. 1987, Dynamical evolution of globular clusters (Princeton, NJ, Princeton University Press, 1987, 191 p.)

Staelin, D. H. & Reifenstein, E. C. 1968, Science, 162, 1481

Stella, L., White, N. E., & Priedhorsky, W. 1987, ApJ, 312, L17

Stella, L., White, N. E., & Rosner, R. 1986, ApJ, 308, 669

Stinebring, D. R., Ryba, M. F., Taylor, J. H., & Romani, R. W. 1990, Physical Review Letters, 65, 285

Strohmayer, T. E., Jahoda, K., Giles, A. B., & Lee, U. 1997, ApJ, 486, 355

Strohmayer, T. E., Zhang, W., Swank, J. H., Smale, A., Titarchuk, L., Day, C., & Lee, U. 1996, ApJ, 469, L9

Sturner, S. J. & Dermer, C. D. 1995, A&A, 293, L17

Sturner, S. J., Dermer, C. D., & Mattox, J. R. 1996, A&AS, 120, C445

Swank, J. H., Jahoda, K., Zhang, W., & Giles, A. B. 1996, in The Lives of the Neutron Stars, ed. M. A. Alpar, U. Kiziloglu, & J. van Paradijs, NATO ASI Ser C., 450) (Boston: Kluwer)

Tavani, M. 1992, in X-Ray Binaries and Recycled Pulsars, ed. E. P. J. van den Heuvel & S. A. Rappaport, 387

Taylor, J. H. & Cordes, J. M. 1993, ApJ, 411, 674

Taylor, J. H., Manchester, R. N., & Lyne, A. G. 1993, ApJS, 88, 529

Thompson, C. & Duncan, R. C. 1995, MNRAS, 275, 255

—. 1996, ApJ, 473, 322+

Thompson, D. J., Arzoumanian, Z., Bertsch, D. L., Brazier, K. T. S., Chiang, J., D'Amico, N., Dingus, B. L., Esposito, J. A., Fierro, J. M., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Johnston, S., Kanbach, G., Kaspi, V. M., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Manchester, R. N., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nel, H. I., Nice, D. J., Nolan, P. L., Ramanamurthy, P. V., Shemar, S. L., Schneid, E. J., Sreekumar, P., & Taylor, J. H. 1994, ApJ, 436, 229

Thompson, D. J., Arzoumanian, Z., Bertsch, D. L., Brazier, K. T. S., D'Amico, N., Fichtel, C. E., Fierro, J. M., Hartman, R. C., Hunter, S. D., Johnston, S., Kanbach, G., Kaspi, V. M., Kniffen, D. A., Lin, Y. C., Lyne, A. G., Manchester, R. N., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nel, H. I., Nice, D., Nolan, P. L., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Sreekumar, P., & Taylor, J. H. 1992, Nature, 359, 615

Thompson, D. J., Bertsch, D. L., Dingus, B. L., Esposito, J. A., Etienne, A., Fichtel, C. E., Friedlander, D. P., Hartman, R. C., Hunter, S. D., Kendig, D. J., Mattox, J. R., McDonald, L. M., Mukherjee, R., Ramanamurthy, P. V., Sreekumar, P., von Montigny, C., Fierro, J. M., Jones, B. B., Lin, Y. C., Michelson, P. F., Nolan, P. L., Tompkins, W., Willis, T. D., Kanbach, G., Mayer-Hasselwander, H. A., Merck, M., Pohl, M., Kniffen, D. A., & Schneid, E. J. 1996, ApJS, 107, 227

Thompson, D. J., Bertsch, D. L., Dingus, B. L., Esposito, J. A., Etienne, A., Fichtel, C. E., Friedlander, D. P., Hartman, R. C., Hunter, S. D., Kendig, D. J., Mattox, J. R., McDonald, L. M., von Montigny, C., Mukherjee, R., Ramanamurthy, P. V., Sreekumar, P., Fierro, J. M., Lin, Y. C., Michelson, P. F., Nolan, P. L., Shriver, S. K., Willis, T. D., Kanbach, G., Mayer-Hasselwander, H. A., Merck, M., Radecke, H.-D., Kniffen, D. A., & Schneid, E. J. 1995, ApJS, 101, 259

Thompson, D. J., Bertsch, D. L., Fichtel, C. E., Hartman, R. C., Hofstadter, R., Hughes, E. B., Hunter, S. D., Hughlock, B. W., Kanbach, G., Kniffen, D. A., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., von Montigny, C., Nolan, P. L.,

Nel, H. I., Pinkau, K., Rothermel, H., Schneid, E. J., Sommer, M., Sreekumar, P., Tieger, D., & Walker, A. H. 1993, ApJS, 86, 629

Thompson, D. J., Harding, A. K., Hermsen, W., & Ulmer, M. P. 1997, in AIP Conf. Proc. 410: Proceedings of the Fourth Compton Symposium, ed. C. D. Dermer, M. S. Strickman, & J. D. Kurfess (American Institute of Physics), 39+

Thorsett, S. E. & Chakrabarty, D. 1999, ApJ, 512, 288

Toscano, M., Sandhu, J. S., Bailes, M., Manchester, R. N., Britton, M. C., Kulkarni, S. R., Anderson, S. B., & Stappers, B. W. 1999, MNRAS, 307, 925

Ulmer, M. P., Lomatch, S., Matz, S. M., Grabelsky, D. A., Purcell, W. R., Grove, J. E., Johnson, W. N., Kinzer, R. L., Kurfess, J. D., Strickman, M. S., Cameron, R. A., & Jung, G. V. 1994, ApJ, 432, 228

Ushomirsky, G., Cutler, C., & Bildsten, L. 2000, MNRAS, 319, 902

van den Heuvel, E. P. J. & van Paradijs, J. 1988, Nature, 334, 227+

van der Klis, M. 1989, in Timing Neutron Stars, ed. H. Ögelman & E. van den Heuvel, Lecture Notes in Physics No. 330 (Kluwer Academic Publishers), 27–69

van der Klis, M. 2000, ARA&A, 38, 717

van Paradijs, J., Verbunt, F., Shafer, R. A., & Arnaud, K. A. 1987, A&A, 182, 47

van Straaten, S., van der Klis, M., Kuulkers, E., & Méndez, M. 2001, ApJ, 551, 907

Vaughan, B. A., van der Klis, M., Wood, K. S., Norris, J. P., Hertz, P., Michelson, P. F., van Paradijs, J., Lewin, W. H. G., Mitsuda, K., & Penninx, W. 1994, ApJ, 435, 362

Verbunt, F., Belloni, T., Johnston, H. M., van der Klis, M., & Lewin, W. H. G. 1994, A&A, 285, 903

Verbunt, F., van den Heuvel, E. P. J., van Paradijs, J., & Rappaport, S. A. 1987, Nature, 329, 312

Wallace, P. M., Griffis, N. J., Bertsch, D. L., Hartman, R. C., Thompson, D. J., Kniffen, D. A., & Bloom, S. D. 2000, ApJ, in press, see astro-ph/0004252

Webbink, R. F., Rappaport, S., & Savonije, G. J. 1983, ApJ, 270, 678

Weisberg, J. M. & Taylor, J. H. 1984, Physical Review Letters, 52, 1348

Wijnands, R. & van der Klis, M. 1998, Nature, 394, 344

Wolszczan, A. 1991, Nature, 350, 688

Wolszczan, A. & Frail, D. A. 1992, Nature, 355, 145

Woosley, S. E. & Weaver, T. A. 1986, ARA&A, 24, 205

Xilouris, K. M., Fruchter, A., Lorimer, D. R., Eder, J., & Vazquez, A. 2000, in Pulsar Astronomy - 2000 and beyond, ed. M. Kramer, N. Wex, & R. Wielebinski, IAU Symposium No. 177 (Sheridan), 21

Yadigaroglu, I. A. & Romani, R. W. 1997, ApJ, 476, 347+

Zavlin, V. E., Pavlov, G. G., & Shibanov, Y. A. 1996, A&A, 315, 141

Zavlin, V. E., Pavlov, G. G., Shibanov, Y. A., & Ventura, J. 1995, A&A, 297, 441

Zhang, S. N., Yu, W., & Zhang, W. 1998a, ApJ, 494, L71

Zhang, W., Jahoda, K., Kelley, R. L., Strohmayer, T. E., Swank, J. H., & Zhang, S. N. 1998b, ApJ, 495, L9

Zhang et al., S. 1996, I.A.U. Circ., # 6462