

New Search Techniques for Binary Pulsars

A thesis presented

by

Scott Mitchell Ransom

to

The Department of Astronomy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Astronomy

Harvard University

Cambridge, Massachusetts

4 September 2001

© 2001, by Scott Mitchell Ransom

All rights reserved.

New Search Techniques for Binary Pulsars

Dr. Giovanni G. Fazio

Scott Mitchell Ransom

Abstract

We have developed two new methods for finding binary pulsars that search complementary portions of parameter space. Both techniques examine small regions of the Fourier transform of a time series and identify the distinctive but very low signal-to-noise patterns produced by binary pulsars. The first technique is a Fourier-domain version of traditional “acceleration” searches which can coherently detect a binary pulsar if the orbital period is much longer than the observation time ($P_{orb} \gtrsim 10T_{obs}$). Based on Fourier-domain matched filtering, it provides sensitivities similar to time-domain versions, but with great computational advantages (i.e. greatly decreased memory usage, increased speed, and an inherently parallel nature) when the time series is very long ($\gtrsim 10^8$ points). The second technique is a computationally “cheap” but incoherent search for binary pulsars where $P_{orb} \lesssim \frac{2}{3}T_{obs}$. This search detects periodic sidebands created by orbital phase modulation of a binary pulsar’s signal using a two-stage Fourier analysis. Sideband searches have the potential to detect very interesting exotic pulsar systems in ultra-compact binaries that would be impossible to detect using traditional techniques. We have implemented both searches along with many other advanced Fourier “tricks” into a suite of pulsar search software called PRESTO. Using this software we have discovered a very interesting 3 ms binary pulsar with a 102 m orbital period and a $\sim 10 M_{Jupiter}$ companion in the globular cluster NGC 6544, a most-likely isolated 4.71 ms pulsar in the globular cluster Terzan 5, and have confirmed the new 65.6 ms X-ray pulsar discovered in the supernova remnant 3C 58.

Contents

Abstract	iii
Acknowledgements	vii
1 Introduction	1
1.1 Overview	1
1.2 Why Globular Clusters?	3
1.3 Chapter Summaries	5
2 Fourier Techniques for Long Time Series	7
2.1 Introduction	8
2.2 Discrete Fourier Transform	9
2.2.1 Introduction to the DFT	9
2.2.2 Computation of Very Long FFTs	10
2.3 Fourier Transforms of Real Data	13
2.3.1 Fourier Response to Noise	13
2.3.2 Fourier Response to Periodic Signals	15
2.3.3 Periodic Signals With Noise	21
2.3.4 Photon Counting Data	25
2.4 Improving the DFT Response to Arbitrary Signals	27
2.4.1 Fourier Interpolation	27
2.4.2 General DFT Response Correction	31
2.5 Signal Property Estimation	37
2.5.1 Power, Phase and Signal Amplitude	37
2.5.2 Signal Location and Duration in Time	39
2.5.3 The Pulsation Frequency and Frequency Derivative	41
2.6 Conclusions	41
3 A Binary Millisecond Pulsar in NGC 6544	43
3.1 Introduction	43
3.2 Observations and Data Reduction	44
3.3 Results	46

3.4	Discussion	50
4	A New Binary Pulsar Search Technique	52
4.1	Introduction	52
4.2	Phase Modulation due to Orbital Motion	53
4.2.1	Circular Orbits	54
4.2.2	Small Amplitude Limit	56
4.2.3	Elliptical Orbits	56
4.3	Pulsars with Deterministic Periodic Phase Modulation	59
4.4	Sideband Searches	62
4.4.1	Two-Staged Fourier Analysis	63
4.4.2	Search Considerations	64
4.5	Discussion	70
5	Results and Future Prospects	74
5.1	PRESTO	74
5.2	Globular Cluster Searches	77
5.2.1	Terzan 5	77
5.2.2	Liller 1	78
5.2.3	47 Tucanae	78
5.3	Other Searches	82
5.3.1	3C 58	82
5.3.2	Supernova Remnants with <i>RXTE</i>	82
5.4	Future Possibilities	85
A	Derivation of Fourier Interpolation	86
B	Derivation of Centroid	87
C	Derivation of Purity	89
D	Centroid, Purity, and Fourier Response	91
E	Derivation of Frequency Uncertainty	93
F	Derivation of \dot{f} Uncertainty	97
	References	100

Those afraid of the universe as it really is, those who pretend to nonexistent knowledge and envision a Cosmos centered on human beings will prefer the fleeting comforts of superstition. They avoid rather than confront the world. But those with the courage to explore the weave and structure of the Cosmos, even where it differs profoundly from their wishes and prejudices, will penetrate its deepest mysteries.

— CARL SAGAN, *Cosmos*

For Ian and Ellie.

Acknowledgements

It is amazing to think that this long strange journey is almost over. Many people deserve thanks for helping me along the way, but I will mention only a select few.

To COL Bruce Oldaker, who encouraged me to apply for a “Graduate Fellowship of the Hertz Foundation,” when I was much more interested in a Rhodes or Marshall Scholarship, I cannot thank you enough. And to the Hertz Foundation, thank you for allowing me to pursue a dream I have had since I was a boy. Your fellowships are not nearly as well known, nor as appreciated, as they should be — I will attempt to change that.

To Brian Schmidt, thank you for encouraging me to try to understand the Middleditchian Fourier stuff present in those early drafts of the SN 1987A pulsar paper. Little did I know just what all that effort would turn into.

To Giovanni, thank you for starting me out with the *SSPM* and then encouraging me when I branched off onto the path that eventually led to this thesis. It has truly been an honor to work with you.

To Lincoln Greenhill, Vicky Kaspi, Josh Grindlay, and Giovanni (my TAC), a hearty thanks for going to bat for me over the NGC 6544 issue and encouraging me to finish my dissertation sooner rather than later. Special thanks go to Josh, Ramesh, and Giovanni for the numerous letters of reference.

To Jim Cordes, who travelled from Cornell for my defense and brought with him many comments that have improved this thesis significantly, I thank you for your troubles and hope that we can continue to work together in the future.

Lincoln and Jim Herrnstein deserve thanks for travelling to and almost destroying Parkes in an effort to take pulsar search data for me while I was stationed in Ft. Sill, Oklahoma as an Artillery instructor. Similar thanks go to Fernando Camilo, Nichi D’Amico, and Paulo Freire for providing lots of 47 Tuc data to play with over the last two years.

For those thankless individuals who developed astro-ph¹, ADS², and HEASARC³ just-in-time to allow me to use them from the intellectual wastelands of U.S. Army bases across this great country of ours, I thank you. And to the many MAJs, LTCs, and COLs who simply looked at me as if I was crazy when they found me reading *ApJ Letters* when I should have been firing artillery, I salute you.

Many thanks go to Lars Hernquist for giving me the opportunity to administer

¹<http://xxx.lanl.gov/archive/astro-ph>

²<http://adsabs.harvard.edu>

³<http://heasarc.gsfc.nasa.gov>

the theory group “Beowulf” and for allowing me the use of so *many* of its CPU-hours. And in a similar vein, thanks go to Rick Prairie for providing access to the machines at the Ohio Supercomputing Center and many lively discussions.

To my parents, thank you for allowing me to watch (over and over again) a short series on PBS that would essentially change my life. And to Carl Sagan, who wrote *COSMOS*, thank you for your efforts to teach us about the universe.

And finally, thank you Karen, for allowing me to spend *entirely* too much of the last seven years in front of my computer screen, when I should have been spending it with you. You are a gem.

Chapter 1

Introduction

1.1 Overview

Discovering new pulsars today is hard. Almost all of the bright pulsars have already been discovered, leaving only the non-luminous, distant, or exotic left to find. This thesis is about finding those pulsars. In particular, this thesis is primarily about finding pulsars that are non-luminous, distant, *and* exotic — the binary millisecond pulsars (MSPs) in globular clusters.

Radio pulsars would be difficult to detect if the only problem was their intrinsically low radio flux densities. Unfortunately, other problems plague pulsar searchers as well. The frequency dependent dispersive effects of the interstellar medium (ISM) require searches over not just pulsar period and location on the celestial sphere, but also over a range of Dispersion Measures (DMs) that are proportional to the integrated electron density along the line of sight to the pulsar. The orbital motion of binary pulsars causes the virtually constant spin period of a pulsar to be observed as changing, possibly drastically, over short intervals of time. And terrestrial radio interference, now so common in our modern world, often overwhelms feeble pulsar signals or perhaps more nefariously, *masquerades* as them.

Yet in spite of these difficulties pulsars are being found at a breathtaking rate. In the last few years, the Parkes Multibeam Pulsar Survey (Lyne et al. 2000a; Manchester et al. 2001) has nearly doubled the number of known pulsars (to ~ 1300) by observing the Galactic plane for relatively long integrations (~ 35 m) at 20 cm (1.4 GHz). Observations taken at 1.4 GHz, as opposed to the more traditional frequency of 400 MHz, greatly reduce the pulse smearing caused by uncorrected DM ($\propto \nu^{-3}$) and the pulse broadening effects of interstellar scattering ($\propto \nu^{-4}$)¹.

¹For a thorough discussion of interstellar propagation and instrumental effects on pulsar observations see Cordes (2001).

High-bandwidth and very-low-noise receivers help to compensate for the lower pulsar flux densities observed at these wavelengths².

Searches for pulsars in globular clusters tend to be even more difficult for two basic reasons. They are relatively distant and their pulsar populations are dominated by MSPs — many of which are in close binary systems (see e.g. Kulkarni & Anderson 1996; Fruchter & Goss 2000; Camilo et al. 2000).

The large distances of globular clusters imply the need for long observation times in order to detect the intrinsically low luminosity MSPs³. Such distances also mean that high frequency resolution search hardware (i.e. large numbers of filterbank channels or correlator lags) is required to contend with the higher expected DMs.

Searching for MSPs is difficult in general since their fast spin periods require high time resolution data (especially to see pulse structure from higher harmonics) and pulse smearing due to dispersion or scattering has progressively greater effects as spin periods decrease. The binary nature of most cluster pulsar MSPs makes matters even worse. The orbital motions of these systems "smear" the signal power over many bins of the Fourier power spectrum, greatly decreasing search sensitivities (see e.g. Johnston & Kulkarni 1991).

Traditional "acceleration" searches (e.g. Middleditch & Kristian 1984; Anderson et al. 1990; Wood et al. 1991) attempt to correct for these orbital motions by approximating the frequency evolution as linear in time and searching over this additional parameter. Such techniques have successfully uncovered such remarkable systems as the relativistic double neutron star (NS) system in M15 (Anderson 1992) and the extremely low mass systems with orbital periods shorter than 2 hours in 47 Tucanae and NGC 6544 (Camilo et al. 2000; Ransom et al. 2001b; D'Amico et al. 2001). Unfortunately, this technique only works when the time series length is much shorter than the orbital period. For short orbital period systems, this drastically limits the duration of data that can be analyzed coherently and therefore the resulting sensitivity.

The complex nature of globular cluster pulsar searches has made high-performance computing a requirement from the beginning. When Lyne et al. (1987) discovered the first globular cluster pulsar (a 3 ms pulsar in M28) the analysis took 5 h on a

²Radio pulsar flux densities go as $S_\nu \propto \nu^\alpha$ where $-1 \lesssim \alpha \lesssim -4$ (Lyne & Graham-Smith 1998).

³Another method of detecting weak MSPs is appropriate for certain globular clusters that display significant scintillation. Since diffractive interstellar scintillation modulates a pulsar's intensity with an exponential probability distribution, repeated observations of a cluster can occasionally catch a very weak pulsar when it is scintillated to a significantly higher-than-average flux density. Such an observation strategy has been used with great success in 47 Tuc (Camilo et al. 2000) and NGC 6752 (Possenti et al. 2001).

TABLE 1.1
Globular Cluster Search Raw Data

Telescope	Instrument	N_{chan}	BW (MHz)	Sample Bits	T_{samp} (μs)	Rate (MB/s)	8 h Obs (GB)
GBT	BCPM	96	170	4	50	0.92	25.7
Parkes	Multibeam	512	256	1	80	0.76	21.5
Arecibo	WAPP	256	100	16	64	7.63	214.6
GBT	Spigot Card	2048	800	8	82	23.8	670.6

Cray XMP supercomputer. As computer speeds and memory sizes increased, so did the complexity of the search algorithms and the size of the data sets, such that even today, searches of globular clusters with *known* DMs can take many times the duration of the observation to conduct — even using state-of-the-art computer hardware. For clusters where searches over DM and acceleration are performed many thousands of CPU hours can be spent searching an observation.

To give an idea of the computational requirements for typical cluster searches it is useful to simply look at the *sizes* of the raw data sets that are generated during an observation. Table 1.1 shows some of the basic properties of “typical” search data using current (or soon-to-be-available) instruments. For 8h globular cluster observations, which we argue are useful for detecting isolated pulsars and for “sideband” searches (see Chapter 4), raw data sets can run to over a half terabyte in size! For those of us attempting full-scale period, DM, and acceleration searches on such observations, we need all the CPU cycles (and disk space) we can handle.

1.2 Why Globular Clusters?

Searching globular clusters for pulsars obviously takes a lot of effort as well as prodigious quantities of telescope and computer time. This begs the question: “Why globular clusters?”

Globular clusters are known to be rich sources of radio pulsars — approximately 55 at the time of this writing. Mechanisms which are not yet fully understood caused clusters to manufacture large numbers of low-mass X-ray binaries (LMXBs; thought to be precursors of MSPs), well in excess of what one might initially expect based on simple mass and/or number arguments compared with the Galaxy. The high cross-sections for dynamical interactions between primordial binaries and other stars in the dense cluster environment undoubtedly has much to do with this fact (Davies & Hansen 1998; Rasio et al. 2000). Various arguments suggest that anywhere from a

few hundred to a few thousand active pulsars await discovery in the globular cluster system (Fruchter & Goss 1990; Anderson 1992; Fruchter & Goss 2000; Camilo et al. 2000).

As for the science produced by cluster pulsars, just as for Galactic systems, most of it comes from long-term high-precision pulse timing measurements. Timing of the 8 known pulsars in M15 allowed measurements of the radial and luminosity distributions of the pulsars, estimates of the central and average surface density of the cluster, and the possible identification of a gradient in the Galactic free-electron distribution (Anderson 1992). Timing observations of some of the more than 20 MSPs in 47 Tuc (Freire et al. 2001a) have allowed similar measurements, as well as the identification of X-ray and (in one case) optical counterparts (Grindlay et al. 2001; Edmonds et al. 2001), the possible detection of a triple system with two pulsars, the measurement of the cluster proper motion, and the detection of ionized intracluster gas (Freire et al. 2001b). These and other groups have also made many interesting measurements of cluster binary pulsar systems which have provided tests of accretion and ablation models, masses of neutron stars and white dwarfs, and insights into the evolutionary history of X-ray binaries (see Kulkarni & Anderson 1996). Future discoveries have the potential to offer tests of General Relativity of the kind made famous by the Hulse-Taylor binary pulsar.

We have been spurred to action in this field by four basic facts.

1. Radio telescopes and pulsar back-ends have undergone revolutionary improvements in sensitivity in the past few years are just now beginning to come on-line and produce results. Examples include the Arecibo upgrade, the new Green Bank Telescope (GBT), and the Multibeam receiver system at Parkes. Each of these systems will undoubtedly discover many hundreds of pulsars in the coming years (the unfinished Parkes Multibeam survey has already discovered more than 600 new pulsars as mentioned in §1.1).
2. The costs of the large numbers of CPU cycles required to analyze globular cluster data has dropped drastically in recent years. It is now possible to buy “Beowulf” style⁴ machines with fraction-of-a-teraflop capabilities for only \sim \$100K. These machines make large-scale searches of globular clusters possible for only a fraction of the cost required just a few years ago.
3. Search algorithm improvements, on-the-other-hand, have *not* kept pace with the advances in computers and pulsar search hardware. In fact, we would argue that search algorithm development has remained virtually *stagnant* for

⁴A collection of usually identically configured PC-style computers connected together on a fast local network.

the past 20 years. New search algorithms optimized for the “Beowulf” style machines mentioned above would be especially useful, and in fact are some of the main topics of this thesis.

4. Very Large Array (VLA) images of globular clusters taken over the years have identified a number of radio sources with characteristics indicative of radio pulsars, but where no pulsations have been detected (Hamilton et al. 1985; Fruchter & Goss 1990; Kulkarni et al. 1990; Johnston et al. 1991; Fruchter & Goss 2000). One possible reason why no pulsations have been detected is that these sources are located in high-DM clusters and contain MSPs in potentially very compact binaries. The selection effects described above have made them virtually *impossible* to detect using conventional search techniques.

This thesis addresses each of these items. In particular, a large fraction of this work is dedicated to the discussion of new algorithms and Fourier techniques that allow large scale searches for binary pulsars to be carried out on individual workstations.

1.3 Chapter Summaries

The thesis is divided into three central chapters that have been — or soon will be — published as separate papers, followed by a chapter of additional results and future prospects, and then various appendices.

- **Chapter 2** discusses a variety of useful Fourier analysis techniques that can be used to greatly increase the amount of information gleaned from a Fourier transform of a long astrophysical time series. Many of these techniques have been known and used by some of the wizards of Fourier analysis for some time (e.g. Middleditch 1976), but remain unknown and unused by most researchers. This chapter provides derivations of the techniques and discusses their use in the context of pulsar data analysis. One of the most useful techniques presented is a Fourier-domain matched filtering search optimized for use on workstation-class machines.
- **Chapter 3** presents the new 3 ms binary pulsar PSR J1807–2459 discovered in the globular cluster NGC 6544 using a Fourier-domain “acceleration” search discussed in Chapter 2. This fascinating system has a minimum companion mass of only $\sim 10 M_{\text{Jupiter}}$ and an orbital period less than 2 hours.
- **Chapter 4** introduces the phase modulation “sideband search”. We calculate the general Fourier response of a binary pulsar in a circular orbit and show how the new search can identify this response in an efficient manner. We compare

the sideband search with traditional “acceleration” searches and provide a detailed discussion of how real searches would be conducted.

- **Chapter 5** presents results from some of the searches we have conducted using the algorithms developed for this thesis. The newly discovered pulsar Terzan 5 D is presented as well as the new Crab-like X-ray pulsar in supernova remnant 3C 58. Future searches and prospects are mentioned.

Chapter 2

Fourier Techniques for Very Long Astrophysical Time Series Analysis

Scott M. Ransom, Stephen S. Eikenberry, & John Middleditch 2001, to be submitted to *The Astronomical Journal*

Abstract

In this paper, we present an assortment of both standard and advanced Fourier techniques that are useful in the analysis of astrophysical time series of very long duration — where the observation time is much greater than the time resolution of the individual data points. We begin by reviewing the operational characteristics of Fourier transforms (FTs) of time series data, including power spectral statistics, discussing some of the differences between analyses of binned data, sampled data, and event data, and briefly discuss algorithms for calculating discrete Fourier transforms (DFTs) of very long time series. We then discuss the response of DFTs to periodic signals, and present techniques to recover Fourier amplitude “lost” during simple traditional analyses if the periodicities change frequency during the observation. These techniques include Fourier interpolation which allows us to correct the response for signals that occur between Fourier frequency bins. We then present techniques for estimating additional signal properties such as the signal’s centroid and duration in time, the first and second derivatives of the frequency, the pulsed fraction, and an overall estimate of the significance of a detection. Finally, we present a recipe for a basic but thorough Fourier analysis of a time series for well-behaved pulsations.

2.1 Introduction

The analysis of time series data is an important tool in many areas of astrophysics, including research involving white dwarfs, black holes, and neutron stars. In the study of neutron stars, time series analysis has had particular importance for pulsar research due to the high coherence of pulsar periodicities. While many techniques can be used to investigate the properties of these periodic signals, the computational efficiency of the Fast Fourier Transform (FFT) makes Fourier analysis the preferred approach for many applications (see e.g. Burns & Clark 1969). The literature contains literally hundreds of descriptions of various aspects of Fourier analysis, most of which deal with signal detection using the power spectrum.

Because of this concentration on power spectra, many researchers discard a wealth of information provided by the Fourier phases. Techniques which use this phase information exist that can provide insight into many useful signal properties. While many of these techniques have been known for some time (see e.g. Middleditch 1976), few have appeared in textbooks or refereed journals, and fewer still have been presented with any sort of derivation or insight into their assumptions and/or limitations.

A second, more practical problem with most astronomical Fourier analysis is its concentration on short time series. We define “short” to mean either that the full time series of binned or sampled data can fit into the core memory of your computer ($N \lesssim 10^7$ points), or for data consisting of events (such as photon arrival times in X-ray astronomy), that the time resolution (dt) of each event makes up a non-negligible fraction of the total time duration (T) of the data ($T/dt \lesssim 10^7$). Billion point FFTs have been used successfully in the past (e.g. Anderson 1992; Mattox et al. 1996; Middleditch et al. 2000), but each required the use of state-of-the-art supercomputing facilities. Today, such analyses are possible using clusters of workstations or even individual desktop machines. Many projects, such as pulsar searches of globular clusters, astero- or helioseismological observations, and gravitational wave experiments, require extremely large Fourier transforms in order to make the highest sensitivity (i.e. *coherent*) searches for pulsations, and to extract the maximum amount of information from signals found in these searches.

It is our goal in this paper to present some useful Fourier analysis techniques, that for various reasons are used only rarely when working with long time series. Most of our examples pertain to pulsar searches of very long time series, but the methods can be used in the Fourier analysis of virtually any coherent periodicity. This paper will briefly discuss the properties of the DFT, its response to periodic signals and noise, and methods for its computation. We will discuss methods for interpolating Fourier amplitudes, estimating a signal’s duration and centroid in

time, accurately determining its frequency and frequency derivative, correcting for changing pulsation frequency during an observation, and estimating the phase- or amplitude modulation of a signal. Such techniques allow the detection of signals whose frequencies or amplitudes change significantly during an observation — for instance, due to orbital motion or pulsar spin-down.

2.2 Discrete Fourier Transform

In order to present more advanced Fourier techniques later, we first review some fundamentals of the DFT and its most common implementation, the FFT. Since thorough discussions of the Fourier transform in both its continuous and discrete variants exist in the literature (e.g. Bracewell 1999), we will mention only a few topics of particular relevance to astrophysical data analysis, closely following Middleditch (1976).

2.2.1 Introduction to the DFT

The k^{th} element of the Discrete Fourier Transform (DFT), of a uniformly-spaced time series n_j ($j = 0, 1, \dots, N - 1$) is defined as

$$A_k = \sum_{j=0}^{N-1} n_j e^{-2\pi i j k / N}, \quad (2.1)$$

where $i = \sqrt{-1}$ and k is the Fourier frequency or wavenumber ($k = 0, 1, \dots, N - 1$). For a time spacing dt between successive data elements, the frequency of the k^{th} element is $f_k = \frac{k}{N dt} = \frac{k}{T}$, where T is the total time duration of the sequence being transformed. This frequency spacing of $\frac{1}{T}$, for evenly sampled, un-padded and non-windowed data, is often called the Independent Fourier Spacing (IFS)¹, and defines the finest frequency resolution available while maintaining completely independent Fourier amplitudes. Fourier frequency $\frac{N}{2T}$ is known as the Nyquist frequency.

If we view the DFT summation in the complex plane, we see that it is a simple vector addition with each element rotated by $-\frac{2\pi k}{N}$ from the previous element. If the n_j have a constant value, the sum will form k regular $\frac{N}{k}$ -sided polygons with each polygon returning near to the origin, and with the last one returning exactly to the

¹The IFS is important when trying to determine the overall significance of a candidate from a search. The number of IFSs searched corresponds to the number of independent trials searched and should therefore be included in calculations that try to determine if a candidate was produced by noise. See Vaughan et al. (1994).

origin. Therefore, the DFT of a constant data string will be identically zero for all $k > 0$, and equal to the sum of the data elements for $k = 0$ (the “DC” frequency element).

For most astrophysical observations, the data points are real-valued. This property adds an important feature to Fourier transforms of these time series — they are symmetric about the Nyquist frequency, such that $A_{N-k} = A_k^*$, where A_k^* represents the complex conjugate of A_k . This symmetry allows us to calculate the full DFT of a time series by computing amplitudes at only half of the normal Fourier frequencies, thereby speeding up computation of the DFT by nearly a factor of two (e.g., Press et al. 1992).

When deriving properties and techniques based on DFTs, it is often both computationally and intuitively easier to work with a time-normalized data series, where $T = 1$ and u represents the fraction of the observation complete at any given instant (such that $0 \leq u \leq 1$). In this case, instead of working with frequencies f , or integral wavenumbers k , we define r , which represents any real wavenumber (including non-integers). If the number of samples, N , from our data source, $n(u)$, is large, we can compute a continuous FT

$$A_r = N \int_0^1 n(u) e^{-2\pi i r u} du, \quad (2.2)$$

which is almost identical to our DFT when $r = k$, and produces very high accuracy estimates of the Fourier amplitudes at any frequency such that $0 \ll r \ll \frac{N}{2}$. We will use this approximation in many of our derivations.

2.2.2 Computation of Very Long FFTs

The Fast Fourier Transform is a family of well-known computer algorithms which quickly calculate a DFT in only $O(N \log_2(N))$ operations as opposed to the $O(N^2)$ operations of a brute-force DFT. FFTs, their computation, and their origins have been described in numerous articles and books over the last few decades (see Bracewell 1999, for an introduction). Therefore, we will describe only a few special versions which have become generally known only recently and which are useful in the analysis of extremely long time series.

High energy pulsar searches using photon counting systems (infrared, optical, x-ray, and gamma-ray detectors) or pointed radio telescope searches (for example, searching globular clusters) often utilize very high sampling rates (i.e. 10 – 50 kHz) and very long integration times of hours, days, or even weeks. These observations result in time series with hundreds of millions or even billions of data points. The subsequent FFTs are impossible to perform using conventional FFT algorithms unless the full time series fits into the core memory of the computer being used.

Utilizing special “out-of-core” FFT methods, we can transform such huge arrays using distributed memory parallel systems or with a single workstation and manageable numbers of passes through the data stored on external media. Most of these methods are based on the “four-step” FFT and/or special external media array permutation methods (Fraser 1976; Bailey 1990).

The “Four-Step” FFT

Closely following the derivation described in Sweet & Wilson (1995), we can think of our one-dimensional time series as a “C-like” or row-ordered, two-dimensional matrix of size $N = N_r N_c$, where N_c is the number of columns (i.e. the length of a row) and N_r the number of rows (i.e. the length of a column). Using this data decomposition, the FFT is computed by 1. FFTing the columns of the matrix, 2. multiplying the data by complex “twiddle factors”, 3. FFTing the rows of the matrix, and 4. matrix transposing the result. If we define indices $x = 0, 1, \dots, N_c - 1$, $y = 0, 1, \dots, N_r - 1$, $l = 0, 1, \dots, N_r - 1$, and $m = 0, 1, \dots, N_c - 1$, we can write our signal and its DFT amplitudes as

$$n(x, y) = n_j, \quad j = N_c y + x \quad (2.3)$$

and

$$A(l, m) = A_k, \quad k = N_r m + l. \quad (2.4)$$

Substituting into the definition of the DFT (eqn. 2.1) and simplifying we get

$$A(l, m) = \sum_{x=0}^{N_c-1} \left[e^{-2\pi i x l / N} \sum_{y=0}^{N_r-1} n(x, y) e^{-2\pi i y l / N_r} \right] e^{-2\pi i x m / N_c}. \quad (2.5)$$

Notice that the bracketed terms are really N_c DFTs of length N_r — FFTs of the matrix columns multiplied by the “twiddle factor” $e^{-2\pi i x l / N}$. We denote these column FFTs as

$$A_c(x, l) = e^{-2\pi i x l / N} \sum_{y=0}^{N_r-1} n(x, y) e^{-2\pi i y l / N_r}. \quad (2.6)$$

We now see that the outer summation is N_r DFTs of length N_c — FFTs of the matrix rows composed of the $A_c(x, l)$ terms. We denote the result of these transforms as

$$A_r(m, l) = \sum_{x=0}^{N_c-1} A_c(x, l) e^{-2\pi i x m / N_c}. \quad (2.7)$$

To recover the full FFT in its normal order, we simply transpose $A_r(m, l)$

$$A(l, m) = A_r^T(m, l). \quad (2.8)$$

The “four-step” algorithm only needs small portions of the data in memory at any one time. Unfortunately, the short FFTs are strided in memory² by instead of being stored contiguously. This results in a significant inefficiency in today’s cache-based processors and requires inter-node communications when distributing the short FFTs over many processors on parallel computer systems. These shortcomings can be overcome with the “six-step” algorithm (Bailey 1990).

The “Six-Step” FFT

If the initial data set is transposed from a $N_c \times N_r$ matrix into a $N_r \times N_c$ matrix, the strided column FFTs of length N_r become contiguous row FFTs. This memory locality facilitates the use of processor cache systems, greatly increasing memory response times, and allows independent calculation of the row FFTs by different processors in parallel systems. A second transpose operation after the row FFTs have been calculated counteracts the effects of the first transpose operation and makes the next set of row FFTs contiguous in memory as well. The FFT is finished with a final transpose operation which brings the data into normal order.

These transpose operations are relatively efficient on distributed memory parallel systems with fast inter-node communications, and add little time to the overall FFT. In fact, for applications with very large N , the time required to move the data from external media to computer memory and vice versa tends to dominate the FFT time. The “six-step” algorithm has become the “standard” FFT algorithm for distributed memory systems where the serial nature and large number of short FFTs exploit parallel computation strengths.

Unfortunately, if the data does not all fit into the core memory of a workstation or parallel machine, the transposes become extremely slow operations since data must be read to and written from much slower external media. Fraser (1976) devised optimized methods for dealing with such data permutations on external media including the “two-pass” FFT algorithm.

The “Two-Pass” Out-of-Core FFT

Fraser (1976) and Bailey (1990) both describe how a very large data set may be Fourier transformed with only two read-write passes through externally stored data if a scratch area on the external media the same size as the input data set exists. The method uses the “four-step” algorithm with out-of-core transposes. These transpose algorithms allow blocked external media access and perform most of the

²“Strided” means that the data is stored in a sequence of non-contiguous memory locations spaced by a constant amount of memory known as the “stride”.

transposition work in core memory. While external media access speeds and transfer rates are orders of magnitude slower than core-memory systems, the “two-pass” algorithm allows huge arrays to be transformed in manageable times.

2.3 Fourier Transforms of Real Data

2.3.1 Fourier Response to Noise

If our data n_j are composed of some constant value c_j plus random real-valued noise fluctuations d_j with zero mean, the transform terms become

$$(c_j + d_j) e^{-2\pi ijk/N} = c_j e^{-2\pi ijk/N} + d_j e^{-2\pi ijk/N}. \quad (2.9)$$

The linear nature of the Fourier transform allows us to treat the DFT of the d_j independently from the constant length steps, c_j . Since the complex phase factor for a given j and k is fixed, the direction of each element in the sum is nearly fixed. However, since the sign of the d_j may be either positive or negative, the vector direction of the j^{th} element may be reversed. Thus, the DFT of the d_j creates a kind of random walk in the complex plane.

The statistical properties of this random walk for DFTs of pure noise have been well studied (see, e.g. Blackman & Tukey 1959), and result in power spectra distributed according to an exponential distribution (a χ^2 distribution with 2 degrees of freedom) with average and standard deviation equal to $N \langle d_j^2 \rangle$. If we normalize the powers by dividing by $N \langle d_j^2 \rangle$, the probability for a power $P = |A_k|^2$ in a single bin to equal or exceed a power P' by chance is³

$$Prob(P \geq P') = e^{-P'}. \quad (2.10)$$

Similarly, if we sum m properly normalized powers, the probability for the summed power P_m to exceed a power P' is given by

$$Prob(P_m \geq P') = \sum_{j=0}^{m-1} \frac{(P')^j}{j!} e^{-P'}, \quad (2.11)$$

which is the probability for a χ^2 distribution of $2m$ degrees of freedom to exceed $2P'$. Such an *incoherent* (since no phase information is used) summation of powers is often useful when searching for signals suspected of having power in many harmonics (see §2.3.2).

³This is different than the probability for an actual signal to produce a power $P > P'$ in the presence of noise (see §2.3.3). This difference is important in setting upper limits on the amplitudes of periodic signals as discussed in Vaughan et al. (1994).

Proper normalization of the powers is essential for an accurate estimate of a signal's statistical significance or lack thereof. We often cannot normalize our power spectrum by simply dividing by $N \langle d_j^2 \rangle$, since frequency-dependent noise may be present throughout our power spectrum — perhaps as a result of instrumental, atmospheric, or astrophysical processes. Typically, these processes produce noise which increases in strength towards the low-frequency part of the spectrum and is correspondingly called *red noise*.

Techniques to flatten or remove this “coloured” noise component from the power spectrum are described by Israel & Stella (1996), and usually involve dividing short portions of the power spectrum by the locally determined average power level, P_{local} , such that

$$P_{k,norm} = \frac{|A_k|^2}{N \langle d_j^2 \rangle} \simeq \frac{|A_k|^2}{P_{local}} = \frac{P_k}{P_{local}}. \quad (2.12)$$

As long as the number of averaged powers is small enough such that the power spectrum is roughly constant over the range in question, a *white noise* like power spectrum is produced with average and standard deviation of approximately one and an exponential distribution (eqn. 2.10).

Since strong narrow-band signals near some frequency of interest will skew a local power average upwards (and correspondingly decrease the calculated significance of a signal detection), it is important to exclude such powers from the calculation of P_{local} . A simple and effective way to accomplish this is by normalizing with a corrected local median power level instead of the local power average. An exponential distribution with unity mean and standard deviation has a median of $\ln(2)$. Therefore, dividing a section of raw powers by $\frac{1}{\ln(2)}$ times the local median value is theoretically equivalent to normalizing with the local mean, but has the advantage of being insensitive to high-power outliers in the spectrum.

More advanced algorithms for the removal of “coloured” noise and power normalization do exist. A simple example involves fitting polynomial models to portions of the power spectrum and then dividing by them out. These methods work well for Fourier frequencies near zero where the assumption of roughly equivalent power levels for the local powers may be unwarranted.

For the special case where the noise in our data is purely Poissonian (i.e. for binned photons in an optical or x-ray observation), we have $\langle d_j^2 \rangle = \langle n_j \rangle$. In this case, our properly normalized power for the k^{th} DFT element is

$$P_{k,norm} = \frac{|A_k|^2}{N \langle d_j^2 \rangle} = \frac{|A_k|^2}{n_{ph}}, \quad (2.13)$$

where $n_{ph} = N \langle n_j \rangle$ is the sum of the n_j (or the total number of photons for a photon-counting system), which is also equal to the “DC” frequency value of the FT

(see §2.2.1). However, the same processes that caused the “coloured” noise discussed above can significantly alter this situation and require a power normalization based on local powers (see eqn. 2.12).

2.3.2 Fourier Response to Periodic Signals

One of the more useful properties of the FT for astronomical purposes is its response to periodic signals. Since all real periodic signals can be expanded into a series of sinusoids it is important to understand the FT response to a simple sine wave.

Sinusoidal Signals

If we now let our n_j represent a sampled cosinusoid of amplitude a , phase ϕ , and frequency $f_r = \frac{r}{T}$ (where wavenumber r is an integer and f_r an “integral frequency”), we can write

$$n_j = a \cos(2\pi f_r j dt + \phi) \quad (2.14a)$$

$$= a \cos(2\pi jr/N + \phi) \quad (2.14b)$$

$$= \frac{a}{2} (e^{2\pi jr/N+i\phi} + e^{-2\pi jr/N-i\phi}). \quad (2.14c)$$

From this expression, we see that the k^{th} element of the DFT is given by

$$A_k = \frac{a}{2} \sum_{j=0}^{N-1} e^{-2\pi ijk/N} (e^{2\pi jr/N+i\phi} + e^{-2\pi jr/N-i\phi}) \quad (2.15a)$$

$$= \frac{a}{2} \sum_{j=0}^{N-1} e^{2\pi ij(k-r)/N+i\phi} + e^{-2\pi ij(k+r)/N-i\phi}, \quad (2.15b)$$

and represents the summation of two vectors in the complex plane. For $k \neq r$, the first term traces out $|k - r|$ complete “rotations” (pseudo-polygons which start and end at the origin) in the complex plane (since $k - r$ is an integer), giving a net contribution of zero to the k^{th} DFT element. The second term traces out $k + r$ complete rotations and once again contributes nothing to the k^{th} DFT element (since k and r are both positive and therefore $k + r \neq 0$).

When $k = r$, however, the consecutive terms in the summation add coherently (i.e. in phase and therefore without rotation), since the rotation caused by the DFT exponential exactly cancels that from the signal exponential (the Fourier transform “derotates” the signal). As a result, each element of the sum is a step in the complex plane of magnitude $\frac{a}{2}$ in a direction parallel to the one set by the arbitrary initial phase of the signal, ϕ . For a cosinusoidal signal with integral frequency f_r , the DFT

will be uniformly zero, except in the r^{th} frequency bin, where the response is

$$A_r = \frac{Na}{2} e^{i\phi}. \quad (2.16)$$

The Fourier response is more complicated for sinusoids with non-integral frequencies (i.e. wavenumber r is a non-negative real number). The k^{th} DFT element is still given by eqn. 2.15b, but not all of the signal ends up in a single DFT bin A_k . When $k = [r]$ (where $[r]$ is the nearest integer to r), the first term in eqn. 2.15b traces out a fraction $(k - r)$ of a complete rotation in the complex plane, while the second term traces out $k + [r]$ complete rotations, plus an additional fractional rotation.

When N is large, these complete and fractional “rotations” can be treated as circles and arcs respectively. Therefore, the first term of eqn. 2.15b results in a semi-circular arc of length $\frac{Na}{2}$ along the arc, while the second term produces a semi-circular arc of length $\frac{a}{2}(N \bmod k + [r])$ along the arc. The DFT response is simply a vector drawn from the origin to the end of the arc (see Fig. 2.1). Since virtually all astrophysical applications involve $r \gg 1$, where the first term dominates the response, we will ignore the second term in the rest of our analysis.

The final response is a chord subtending $2\pi(k - r)$ radians of a circle of radius $\frac{Na}{4\pi(k-r)}$. The equation for the length of a chord is

$$C = 2 \frac{s}{\Theta} \sin\left(\frac{\Theta}{2}\right), \quad (2.17)$$

where s is the arc length and Θ is the angle subtended by the chord. The curvature of the arc away from the signal’s starting phase ϕ results in a phase change of $e^{-i\pi(k-r)}$. Therefore, the DFT response and power are

$$A_k = e^{i\phi} e^{-i\pi(k-r)} 2 \frac{Na/2}{2\pi(k-r)} \sin\left[\frac{2\pi(k-r)}{2}\right] \quad (2.18a)$$

$$= \frac{Na}{2} e^{i\phi} e^{-i\pi(k-r)} \frac{\sin[\pi(k-r)]}{\pi(k-r)} \quad (2.18b)$$

$$= A_o e^{-i\pi(k-r)} \text{sinc}[\pi(k-r)] \quad (2.18c)$$

$$P_k = |A_k|^2 = P_o \text{sinc}^2[\pi(k-r)], \quad (2.18d)$$

where $A_o = \frac{Na}{2} e^{i\phi}$ is the DFT response for an integral frequency signal (eqn. 2.16), P_o is the corresponding Fourier power, and the sinc function is defined as $\text{sinc}(x) = \frac{\sin(x)}{x}$. This result is easily confirmed by a direct integration of eqn. 2.2 where $n(u)$ is equal to eqn. 2.14c with $\frac{j}{N} \rightarrow u$.

The sinc factor in eqn. 2.18c produces a loss of sensitivity for the standard FFT to most real-world signals (where r is not an integer). This effect, often

Figure 2.1 Fourier responses plotted as a series of vector additions in the complex plane. The outer circles in each plot show the Fourier amplitude of a signal where all power is recovered by the vector addition (i.e. calculation of the DFT at the correct signal frequency r and frequency derivative \dot{r} for signals with linear changes in frequency over time). The end-points of the vector additions are the Fourier amplitudes. For plots (a) and (b), a fully recovered signal would start at $0 + 0i$ and end at $1 + 0i$. (a) and (b) show the effects on Fourier amplitude and phase when a signal's intrinsic frequency (r in bins or wavenumber) or frequency derivative (\dot{r} in bins/observation) differs from the computed values. For plot (b), the average Fourier frequencies in each case were correct, and only the frequency derivatives were in error. (c) shows the response of PSR J1807-2459 during its discovery observation (Ransom et al. 2001b) with and without corrections for pulsar acceleration (\dot{r}) and interpolation in Fourier frequency (r). The vectors were calculated using the method shown in §2.4.2. The fact that the corrected vector does not quite reach the circle implies that higher order effects of the orbital motion remain un-corrected (see Fig. 2.6). (d) shows corrected and uncorrected responses of 10,000 randomly selected photons from a 2.4 day observation of the Crab pulsar with ROSAT.

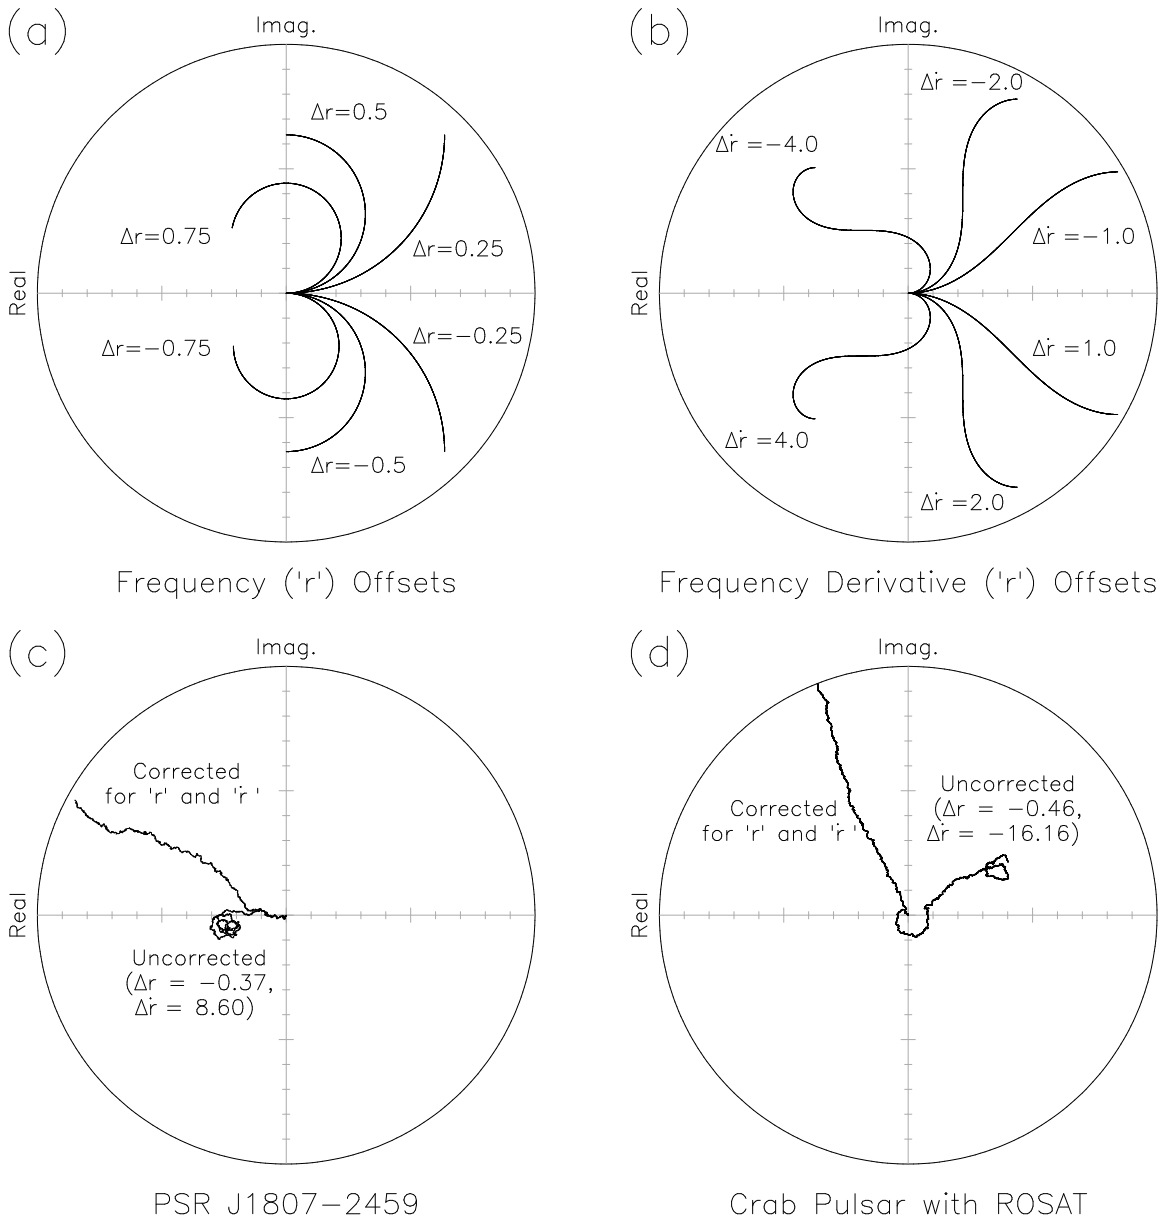


Figure 2.1 (Continued).

called “scalloping” (e.g. Middleditch et al. 1993), is shown in Fig. 2.2, and causes a worst-case (when $|k - r| = \frac{1}{2}$) amplitude reduction of $|A_k| = \frac{2}{\pi} |A_o|$ — a nearly 60% loss of signal power. On average, scalloping results in a $\sim 23\%$ loss of signal power (van der Klis 1989). It is important to remember, though, that this loss in sensitivity is due to the finite frequency resolution of the FFT algorithm rather than an intrinsic feature of the data itself. In §2.4.1 we discuss various methods to reduce or even eliminate this loss of sensitivity.

Non-Sinusoidal Signals

Many real-world periodic signals are not sinusoidal. Fortunately, we can expand all real-valued pulsations as a series of m sinusoidal components

$$n_j = \sum_{h=1}^m a_h \cos(2\pi jhr/N + \phi_h) \quad (2.19a)$$

$$= \sum_{h=1}^m \frac{a_h}{2} (e^{2\pi i jhr/N + i\phi_h} + e^{-2\pi i jhr/N - i\phi_h}) \quad (2.19b)$$

where $h = 1, 2, \dots, m$ specifies the harmonic number (with $h = 1$ known as the “fundamental”), and a_h and ϕ_h represent the amplitude and phase of each component respectively. Due to the linear nature of the FT, we can treat the harmonics as independent sinusoidal signals. Each of these sinusoids produces a Fourier response equivalent to eqn. 2.18c, except that A_o becomes $A_h = \frac{Na_h}{2} e^{i\phi_h}$.

For nearly sinusoidal pulsations only the first few terms of eqn. 2.19a contain significant amplitudes, a_h . This results in a similarly small number of significant peaks in the corresponding power spectrum of the data. Low duty-cycle pulsations (i.e. those with a pulse that is short compared to the pulse period), such as most radio pulsars, on the other hand, have dozens of significant terms in their expansions and therefore harmonics in their power spectra.

A useful pulsation model, particularly for radio and x-ray pulsars, can be constructed based on a modified von Mises distribution (MVMD)

$$f(\kappa, t) = a \frac{e^{\kappa \cos(2\pi f_r t + \phi)} - e^{-\kappa}}{I_0(\kappa) - e^{-\kappa}}, \quad (2.20)$$

where $0 \leq t \leq T$ is the instantaneous time, I_0 is the modified Bessel function of zeroth order, and the shape parameter, κ , determines the width of the function (e.g. Mardia & Zemroch 1975). In the limit as $\kappa \rightarrow 0$, the MVMD becomes a sinusoid, while as $\kappa \rightarrow \infty$, it becomes a Gaussian (see Fig. 2.3). The integral of the MVMD over a single pulse period is simply a , all of which is pulsed (i.e. the pulsed fraction is one). The full-width at half-maximum (FWHM) as a fraction of a pulse, is

$$\text{FWHM}_{\text{MVMD}} = \pi^{-1} \arccos \{ \ln [\cosh(\kappa)] \}, \quad (2.21)$$




Figure 2.2 The thin solid black lines plot the sinc function amplitude responses of raw (or integer) FFT bins. The point where these lines cross (at an offset of $1/2$ and an amplitude of $1 - \frac{2}{\pi} \sim 0.637$ of the full response) is the worst case response for an uncorrected FFT (see eqn. 2.18c). The thin dashed line is the response of an “interbin” as calculated using eqn. 2.31. The thick black line shows the overall “scalloping” when interbinning is used. Worst case responses with interbinning occur at offsets of $\pm (1 - \frac{\pi}{4}) \sim 0.215$ with amplitudes of ~ 0.926 of the full response.

and the maximum value is

$$\max_{\text{MVMD}} = \frac{2a \cosh(\kappa)}{I_0(\kappa) - e^{-\kappa}}. \quad (2.22)$$

The FT of the MVMD can be computed in a particularly convenient form for harmonic analysis. According to Abramowitz & Stegun (1972, eqn. 9.6.34), we can expand the exponential in the MVMD as

$$e^{\kappa \cos(x)} = I_0(\kappa) + 2 \sum_{h=1}^{\infty} I_h(\kappa) \cos(hx), \quad (2.23)$$

where I_h is the modified Bessel function of order h . When combined with the rest of the MVMD definition we have

$$f(\kappa, t) = a + \frac{2a \sum_{h=1}^{\infty} I_h(\kappa) \cos(2\pi h f_r t + h\phi)}{I_0(\kappa) - e^{-\kappa}}. \quad (2.24)$$

This expression is simply a ‘‘DC’’ term (since the integral over a pulse equals a) plus a series of independent cosinusoidal harmonics. After Fourier transforming (i.e. substituting into eqn. 2.2 with $f_r t \rightarrow ru$), we are left with a series of harmonics of amplitude $\frac{AN I_h(\kappa)}{I_0(\kappa) - e^{-\kappa}}$ and phase $h\phi$ at Fourier frequency hr . It is important to note that as $\kappa \rightarrow \infty$ and the pulse becomes narrower, the Fourier amplitudes of the low order harmonics are twice that of a sinusoid with the same pulsed fraction (see the dashed line in Fig. 2.4). This fact, along with the large number of harmonics that low duty-cycle pulsations generate, can significantly increase search sensitivities to such pulsations (see Fig. 2.5).

Fig. 2.4 shows the approximate number of significant harmonics (meaning that a harmonic’s amplitude is greater than one-half the amplitude of the fundamental) generated by an MVMD pulsation, as well as a histogram of the duty-cycles of over 600 radio pulsars (the majority of which are from Taylor et al. 1995). Most radio pulsars have duty-cycles $\lesssim 5\%$ corresponding to $\gtrsim 10$ significant harmonics — assuming a sufficient data sample rate.

2.3.3 Periodic Signals With Noise

When a periodic signal is present in a noisy time series, a sum of m powers P_m , containing some amount of signal power P_s , is no longer described by a χ^2 distribution with $2m$ degrees of freedom (see §2.3.1). Groth (1975) calculated the expectation value and variance of P_m as $\langle P_m \rangle = m + P_s$ and $\langle P_m^2 - \langle P_m \rangle^2 \rangle = m + 2P_s$ respectfully. He also derived the exact probability density function for P_m which can be integrated to give the probability that P_m is greater than or equal to some power




Figure 2.3 Sample pulse profiles from the modified von Mises distribution (MVMD) as described in §2.3.2. FWHM is the fractional full-width at half-maximum and κ is the MVMD shape parameter. High values of κ result in Gaussian profiles, while as $\kappa \rightarrow 0$, the pulse shape becomes more and more sinusoidal. The integral of a full pulse is equal to one unit, all of which is pulsed (i.e. the pulsed fraction is one).




Figure 2.4 The thick solid line represents the approximate number of harmonics from an MVMD signal (see §2.3.2 and Fig. 2.3) that produce Fourier amplitudes greater than one-half the amplitude of the fundamental. The thin solid line is the $\frac{1}{\text{FWHM}}$ rule-of-thumb that is often used to estimate the number of significant harmonics a signal will generate. The thin dashed line plots the ratio of the fundamental amplitude for an MVMD signal to a sinusoidal amplitude of the same pulsed intensity. The grey histogram shows the distribution of pulse widths (FWHM) for over 600 radio pulsars.

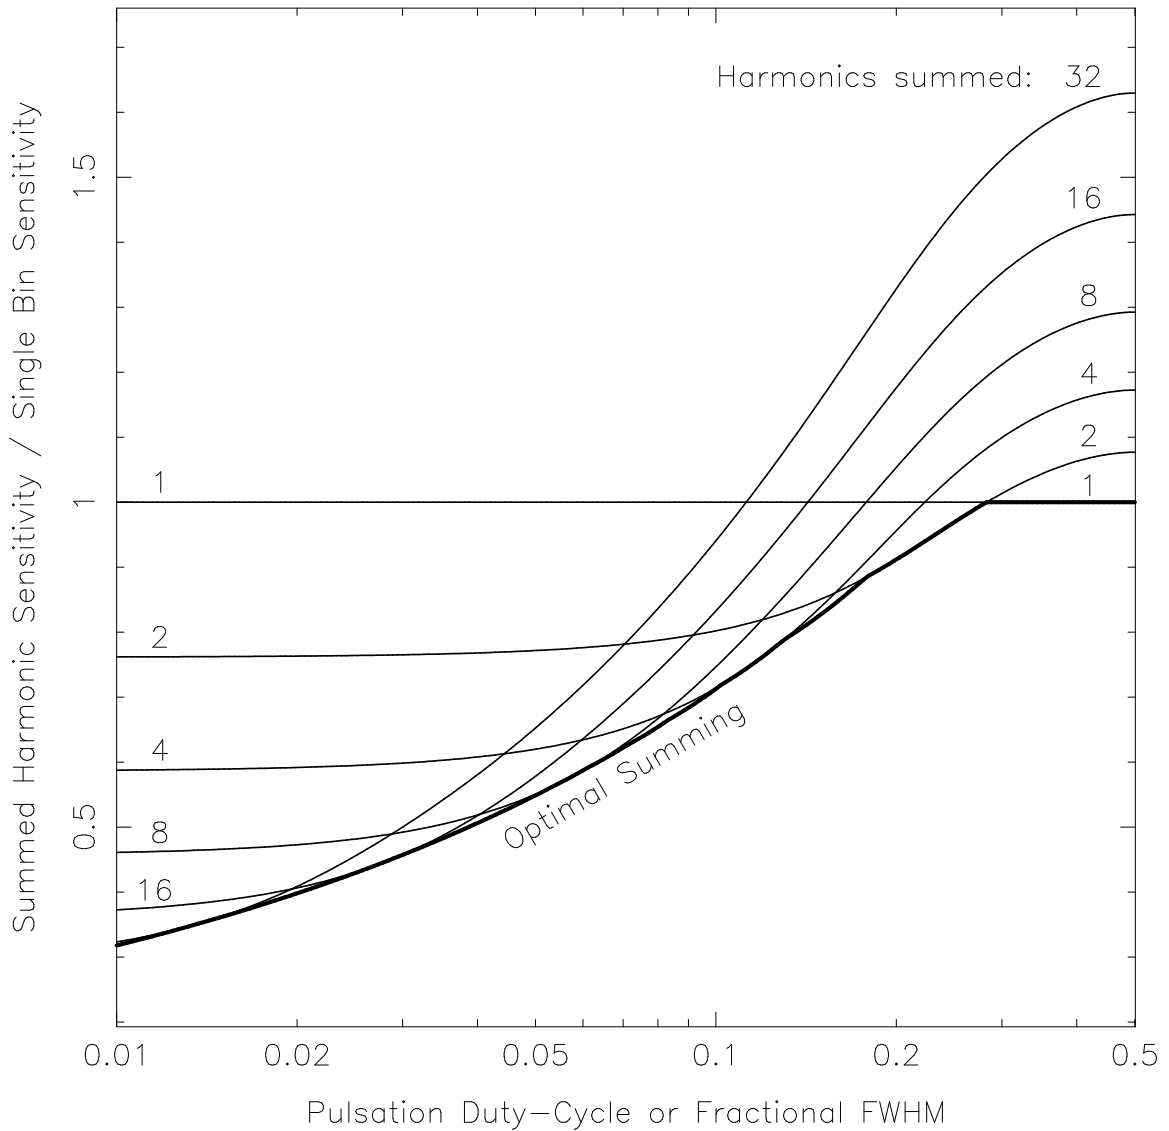


Figure 2.5 The thin lines represent sensitivities to MVMD signals (see §2.3.2 and Fig. 2.3) for the incoherent summing of 1, 2, 4, 8, 16, or 32 harmonics, as compared to searches made without harmonic summing (§2.3.1). Lower numbers represent better sensitivity (i.e. fainter signals are detectable). The best possible sensitivity using incoherent harmonic summing is shown by the thick black line. It should be noted that incoherent summing produces worse sensitivities than not summing if the duty cycles of the pulsations are large. This results from the fact that such pulsations have only a small number of significant harmonics (see Fig. 2.4), so that summing tends to add only noise rather than signal.

P' ,

$$Prob((P_m; P_s) \geq P') = e^{-(P'+P_s)} \sum_{k=0}^{\infty} \sum_{j=0}^{k+m-1} \frac{(P')^j P_s^k}{j!k!}. \quad (2.25)$$

When $P_s = 0$ this equation reduces to Eqn. 2.11.

The fact that the probability density function for a signal plus noise is different from a χ^2 distribution with $2m$ degrees of freedom is very important when trying to determine the sensitivity of a search for pulsations or an upper limit to the amplitude of a periodic signal present in a time series. Vaughan et al. (1994) describe a procedure⁴ for correctly determining search sensitivities and upper limits using the equations of Groth (1975).

2.3.4 Photon Counting Data

Since many of today's astronomical time series come from photon counting experiments, it is important to raise some of the issues particular to Fourier analysis of such data. If we can assume purely Poissonian statistics, a power spectrum of pure noise is flat, and can be normalized simply by dividing by the total number of photons in the data (the zeroth or "DC" frequency bin from the DFT — see §2.3.1). In addition to this difference in power spectrum normalization, the other points worth noting come from the fact that photon counting data is based on the measurement of events rather than the instantaneous sampling of a continuous process.

One important issue that is beyond the scope of this paper is dead-time correction. Dead-time effects modify a detector's sensitivity to photons for some time after the detection of an earlier photon. These effects can cause complicated non-linear and frequency-dependent systematics during Fourier analysis. We refer the reader to Zhang et al. (1995) and references therein, for a thorough discussion of this topic.

Binned vs. Sampled Data

Many high-energy telescopes and detectors produce time series of binned photons rather than the sampled data produced by radio telescopes. Since binning essentially averages a periodic signal's instantaneous rate over the binning time (dt), it modifies the Fourier response to the signal. Binning removes phase information from the data and causes the Fourier response to sinusoidal pulsations to become

⁴Note that Vaughan et al. (1994) use a power normalization that is a factor of two higher than ours.

frequency-dependent — resulting in decreased sensitivity at high frequencies (e.g. Middleditch 1976; Leahy et al. 1983).

The frequency-dependent loss in Fourier amplitude due to binning is $\text{sinc}(\pi f_r dt)$ or $\text{sinc}(\pi r/N)$. The binned data Fourier response to a sinusoid is therefore eqn. 2.18c times this factor. This decrease in sensitivity corresponds to a loss in signal power of about $\sim 19.8\%$ at half the Nyquist frequency and $\sim 59.5\%$ at the Nyquist frequency itself.

For Poissonian noise (i.e. from a photon-counting experiment that does not introduce count-rate dependent systematics), which is independent of the sampling interval, the Fourier response is flat over all frequencies. This is in contrast to a sinusoidal signal passing through the same system which suffers the frequency-dependent attenuation described above (Middleditch 1976). Such behavior is important when trying to estimate limits or amplitudes for pulsations in a time series (Vaughan et al. 1994).

Low Count Rate Data

The Fourier analysis of gamma-ray or x-ray observations often places us in a very unique regime — very long integration times ($\gtrsim 10^4$ s) with very low numbers of counts ($\lesssim 10^3$ photons). In addition, due to visibility constraints based on the orbits of the telescopes, large fractions of the time between the first and last photons may be devoid of counts.

Fourier analysis of such data can overwhelm present computational resources. For example, a 10^6 s observation (about 11.6 days) with photon time-of-arrival (TOA) resolution of 10^{-4} s would require a 10 Gpt FFT for a full-resolution analysis. Such FFTs, while possible, are extremely difficult to compute unless very special and dedicated hardware resources are available. If this data contains only a small number of photons, however, we can exactly compute the DFT over any frequency range and to any frequency resolution using a brute-force implementation of the FT.

If we treat each TOA as a sample of amplitude one, an exact DFT amplitude at arbitrary Fourier frequency r becomes

$$A_r = \sum_{j=1}^{n_{ph}} e^{-2\pi i r (t_j - t_o)/T}, \quad (2.26)$$

where n_{ph} is the number of photons, t_j is the TOA of the j^{th} photon, t_o is the time of the start of the observation, and T is the total duration of the observation. Very quick harmonic-summing searches of an observation are possible using this technique, with the added benefit that “scalping” (see §2.3.2) is non-existent.

Since eqn. 2.26 only involves a summation over the number of photons, it can be computed quickly if n_{ph} is relatively small. Great increases in computation speed can

be had if we search a regular grid of Fourier frequencies. Trigonometric recurrences such as

$$\cos(\theta + \delta) = \cos(\theta) - [\alpha \cos(\theta) + \beta \sin(\theta)] \quad (2.27a)$$

$$\sin(\theta + \delta) = \sin(\theta) - [\alpha \sin(\theta) - \beta \cos(\theta)], \quad (2.27b)$$

where $\alpha = 2 \sin^2(\frac{\delta}{2})$ and $\beta = \sin(\delta)$, allow extremely efficient calculation of the complex exponential for each TOA (Press et al. 1992). This technique allows one to calculate billions of Fourier frequencies from a few hundred photons using only modest computational resources.

2.4 Improving the DFT Response to Arbitrary Signals

2.4.1 Fourier Interpolation

The potential for up to a 30% decrease in signal-to-noise ($S/N \propto \sqrt{P}$) due to an essentially arbitrary difference between the signal frequency and the integer frequency of the nearest Fourier bin is clearly a drawback in the use of the DFT (see §2.3.2). However, if we could calculate the Fourier Transform (FT) at a higher frequency resolution than the $1/T$ spacing that results from the FFT, we could significantly reduce or eliminate scalloping and effectively flatten the Fourier response as a function of frequency.

One possibility for increasing the frequency resolution is to simply calculate the DFT by a brute force summation at frequencies between the integer frequencies. Such a technique is possible in special situations (see §2.3.4), but for most applications, the computational costs would be unacceptably high. Another well-known possibility is to “pad” the end of the time series with a large number of points with values equivalent to the mean of the data⁵. The padding adds no power to the data but it does increase the Fourier resolution since T has been artificially increased by the padding. While this technique is simple and effective for short time series, the difficulties involved in performing very long FFTs (§2.2.2) makes this technique difficult when dealing with long time series.

Yet another way to calculate a higher-resolution Fourier response is to use the complex amplitudes produced by the standard FFT to interpolate responses at non-integer frequencies — a process known as “fine-binning” or “Fourier interpolation”

⁵Padding with the data mean is preferable to zero-padding since zero-padding introduces low frequency power into the Fourier response.

(e.g. Middleditch et al. 1993). Similar techniques allow the full recovery of a signal’s theoretical coherent response provided that the signal’s behavior during the observation is either known or can be guessed.

The purpose of Fourier interpolation is to calculate a complex Fourier amplitude at an arbitrary frequency $f_r = \frac{r}{T}$, where r is any real number, such that the result is sufficiently close to the exact calculation,

$$A_r = \sum_{j=0}^{N-1} n_j e^{-2\pi i j r / N}. \quad (2.28)$$

We can rewrite this expression as

$$A_r = \sum_{k=0}^{N-1} A_k e^{-i\pi(r-k)} \text{sinc}[\pi(r-k)]. \quad (2.29)$$

where the A_k are the complex FFT amplitudes at the integer frequencies l (see Appendix A and §2.4.2 for a derivation and discussion of this result).

The sinc function in eqn. 2.29, provides the key to computing an accurate interpolated amplitude using a relatively small number of operations. Since $\text{sinc}[\pi(r-k)] \rightarrow 0$ as $\pi(r-k) \rightarrow \pm\infty$, the expansion of A_r in terms of the A_k is dominated by the local Fourier amplitudes (i.e. where $k \sim r$). We can therefore approximate A_r as

$$A_r \simeq \sum_{k=[r]-m/2}^{[r]+m/2} A_k e^{-i\pi(r-k)} \text{sinc}[\pi(r-k)], \quad (2.30)$$

where $[r]$ is the nearest integer to r , and m is the number of neighboring FFT bins used in the interpolation. Note that the interpolation is simply a correlation of the local FFT spectrum around the desired frequency element with a “template” response — in this case the theoretical response of a DFT to a sinusoid as described by eqn. 2.18c.

The upper panel in Fig. 2.6 shows both the raw FFT power spectrum (denoted by grey dots) and the interpolated power spectrum (the line connecting the dots) for a radio observation of the short-period binary pulsar PSR J1807–2459. The spectra cover a narrow frequency range near the pulsar’s rotational frequency and were calculated using $m = 32$ and a frequency step size of $\Delta r = \frac{1}{16}$ (compared to the raw FFT frequency step size of $\Delta r = 1$). Note that the interpolated spectrum reveals much more information regarding the shape of the frequency profile — including the true amplitude and location of the maximum in Fourier power. We will see in §2.5 how this can be used to deduce further information regarding the signal properties.

Figure 2.6 An 18σ (single trial) detection from the discovery observation of the 1.7 hour binary PSR J1807–2459 in the globular cluster NGC6544 using a Fourier-domain “acceleration” search. Contour intervals correspond to 30, 60, 90, 120, and 150 times the average local power level. The intrinsic pulsar period and $\dot{f} = 0$ (which corresponds to an un-accelerated FFT of the data) are marked by the solid gray lines. The dots correspond to the “raw” or un-interpolated powers from the original FFT of the observation. The gray ellipse is the predicted “path” of the pulsar in the f – \dot{f} plane given the known binary parameters. During the 28.9 minute observation, the pulsar moved from ~ 11 o’clock to ~ 3 o’clock on the ellipse. The peak’s slight offset from the ellipse and the presence of “shoulders” indicate that the constant \dot{f} assumption of the acceleration search could not fully correct for the orbital motion during this observation. The top and right-hand panels show cuts through the peak in the f and \dot{f} directions respectively. The line in the top panel with gray dots shows the Fourier interpolated $\dot{f} = 0$ power spectrum (calculated as per §2.4.1).




Figure 2.6 (Continued).

A computationally less expensive version of Fourier interpolation is “interbinning,” where we approximate the FT response at half-integer frequencies using only the nearest 2 integer frequency bins. By using the Fourier interpolation equation (eqn. 2.30) with $m = 2$, ignoring an overall phase shift, and boosting the response such that the best-case response (at half-integer frequencies) is equivalent to the full response, we get

$$A_{k+\frac{1}{2}} \simeq \frac{\pi}{4} (A_k - A_{k+1}). \quad (2.31)$$

This particular formulation of interbinning was reported by van der Klis (1989), and its response is shown in Fig. 2.2. Middleditch et al. (1993) have contributed a correction to this formula for use when the data are padded at the end.

Interbinning is extremely useful since such a computationally inexpensive calculation reduces the maximum loss of signal-to-noise from $1 - \frac{2}{\pi}$ or $\sim 36\%$ at a frequency offset of $1/2$ bin to $\sim 7.4\%$ at an offset of $\pm (1 - \frac{\pi}{4})$ bins. This large but cheaply-obtained reduction in scalloping can be extremely beneficial when searching large numbers of FFT bins and interbins.

It is important to note that interbins as defined above have three different properties than integer FFT bins. First, they have different noise properties, which makes calculation of the significance of interbin powers much more difficult. Second, each interbin is correlated with the integer bins it was created from, meaning that interbins are not independent Fourier trials (see §2.2.1 for a discussion of the IFS). And finally, interbins do not recover the correct phase of a sinusoid at the interbin frequency. In general, since interbins are most commonly used during searches to simply identify signals in the power spectrum that would otherwise have been lost due to scalloping, these weaknesses do not degrade the usefulness of their calculation. When a signal is identified, a full-scale interpolation of the Fourier amplitudes around the signal using eqn. 2.30 allows accurate estimates to be made of the signal’s significance and other properties (see §2.5).

2.4.2 General DFT Response Correction

Fourier interpolation serves as a specific example of a much more general technique — the ability to completely recover the fully coherent response for virtually any signal. For Fourier interpolation, we can exactly calculate the response of any Fourier frequency based purely on the properties of the FT. To correct the Fourier transform’s response to a particular signal, we must know not only the properties of the FT, but the properties of the signal we are looking for as well. For the cases we will discuss, this ability comes in one of two forms: matched filtering in the Fourier domain using only the “local” Fourier amplitudes near the Fourier frequencies of interest (which we call the “correlation technique”), or the straightening of the

curved Fourier vector addition in the complex plane (which we will call “vector bending”).

Correcting for Constant Frequency Derivative

In order to illustrate these two methods, we demonstrate how to correct for a signal whose response is reduced due to a constant frequency derivative \dot{f} (or in Fourier frequency bins, $\dot{r} = \dot{f}T^2$). The DFT operates, as we noted in §2.3.2, by “derotating” the vector addition of the data in the complex plane by changing the phases of each of the vector elements — causing a straight line to form for a sinusoidal signal with integral frequency. In the presence of a frequency derivative however, the signal frequency may change by one or more frequency bins over the course of the observation. The complex phase corrections provided by the DFT will fail to completely derotate the data, and pulsation power will be “smeared” across several nearby frequency bins — causing a decrease in the measured DFT response (e.g. Johnston & Kulkarni 1991). Fig. 2.1 illustrates this effect in the complex plane.

As with a frequency error, an uncorrected frequency derivative causes the vector addition to form an arc, although in this case quasi-parabolic rather than circular. The decreased DFT response equals the distance from the origin to the end of the arc. This distance is significantly shorter than that of a coherently detected signal which equals the length along the arc.

Signals with constant or nearly constant \dot{f} are quite common in pulsar astronomy — especially when dealing with time series of very long duration. In such long observations, even the very small spin-downs typical of pulsars can cause a signal to drift across numerous Fourier bins. The Doppler effects of binary pulsar orbits cause similar frequency drifting when the observation time is much shorter than the orbital period.

The “standard” method to correct for a constant frequency derivative is to “stretch” the original time series to compensate for the known or assumed \dot{f} . This process involves re-sampling the data ensemble n_j using a transformation similar to

$$t' = t + \frac{2\dot{f}}{f_o}t^2, \quad (2.32)$$

where t is the time used when sampling the original data, \dot{f} is the frequency derivative, and f_o is the initial frequency of the signal. Additional details and variations on the theme can be found in Middleditch & Kristian (1984); Anderson et al. (1990); Johnston & Kulkarni (1991) and Wood et al. (1991).

By stretching the data using the appropriate transform and then FFTing the corrected time series, we can recover the fully coherent response. Such techniques have been used with significant success in searches of relatively short time-series (e.g.

Camilo et al. 2000). However, this technique runs into significant difficulties when trying large numbers of transformations using long time series where computation of the FFT is non-trivial. Both techniques that we will mention allow full corrections to be made to a signal without requiring multiple FFTs of the full original data set.

Correlation Technique

The correlation technique is the more powerful of the two methods, and uses matched filtering in the Fourier domain to “sweep-up” signal spread over a number of frequency bins into a single bin. In astrophysical applications, we usually have some sort of “pure” signal (like a harmonic from a millisecond pulsar), whose frequency changes as a function of time due to some other process (such as orbital motion or pulsar spin-down). In the Fourier domain, these processes cause the perfect sinc-like response of a harmonic to be spread over numerous local Fourier bins — in effect, the sinc response is convolved with a finite impulse response (FIR) filter (where finite in this case refers to a small portion, say m bins, of the frequency range analyzed rather than a short period of time). If we can predict the complex form (and phase) of that FIR filter, we can recover the coherent response (i.e. the perfect sinc-function) by correlating the appropriate Fourier bins with a “frequency-reversed” and complex-conjugated template that matches the filter.

In mathematical terms, consider a signal with a normalized Fourier response of A_{k-r_o} , where $k - r_o$ is simply the frequency offset of bin k from some reference frequency r_o , which goes to zero as $|k - r_o|$ approaches some number of bins $m/2$. For Fourier interpolation as described in §2.4.1, this response is equal to eqn. 2.18c without the A_o factor (i.e. normalized to an amplitude of one for a coherent response). The complex-valued Fourier response of such a signal at frequency r_o can be calculated with the sum

$$A_{r_o} \simeq \sum_{k=[r_o]-m/2}^{[r_o]+m/2} A_k A_{r_o-k}^* \quad (2.33)$$

If r_o is initially unknown (i.e. we are searching for a signal with the response shape as defined by the template but at an unknown frequency) we simply compute this summation at a range of frequencies r .

Calculating eqn. 2.33 over a range of evenly spaced frequencies is equivalent to correlating the raw FFT amplitudes with the template and is therefore most efficiently computed using short FFTs and the convolution theorem. With FFTs of length M , such that $m \ll M \ll \frac{N}{2}$, we can search a very-long FFT of length $\frac{N}{2}$ for any signal whose A_{k-r_o} we can compute, using overlap-and-save or overlap-and-add techniques (see e.g. Press et al. 1992). Such calculations have advantages over

standard time-domain “stretching” techniques in that they are memory local and can be easily parallelized — important properties when dealing with very-long time-series and modern distributed memory computer architectures.

Moving to our example of a signal with a constant frequency derivative, a single harmonic of the signal has the form

$$n(u) = a \cos \left[2\pi \left(r_o u + \frac{\dot{r}}{2} u^2 \right) + \phi \right] \quad (2.34a)$$

$$= \frac{a}{2} \left[e^{2\pi i (r_o u + \frac{\dot{r}}{2} u^2)} e^{i\phi} + e^{-2\pi i (r_o u + \frac{\dot{r}}{2} u^2)} e^{-i\phi} \right], \quad (2.34b)$$

where we use the same notation as §2.2.1. Neglecting the second term as in §2.3.2 and Fourier transforming at some “center” or average frequency $r'_c = r_o + \dot{r}/2$, we get

$$A_{r'_c} = \frac{aN}{2} e^{i\phi} \int_0^1 e^{i\pi(\dot{r}u^2 + 2q_r u)} du, \quad (2.35)$$

where $q_r = r_c - r'_c$, and the real “center” frequency of the signal is $r_c = r_o + \dot{r}/2$. This integral can be evaluated in closed form

$$\int_0^1 e^{i\pi(\dot{r}u^2 + 2q_r u)} du = \frac{1}{\sqrt{2\dot{r}}} e^{-i\pi\frac{q_r^2}{\dot{r}}} \{S(Z_r) - S(Y_r) + i[C(Y_r) - C(Z_r)]\} \quad (2.36)$$

where $Y_r = \sqrt{\frac{2}{\dot{r}}} q_r$, $Z_r = \sqrt{\frac{2}{\dot{r}}} (q_r + \dot{r})$, and $C(x)$ and $S(x)$ are the Fresnel integrals

$$C(x) = \int_0^x \cos\left(\frac{\pi}{2}t^2\right) dt, \quad S(x) = \int_0^x \sin\left(\frac{\pi}{2}t^2\right) dt. \quad (2.37)$$

The Fourier transform response then becomes

$$A_{r'_c} = \frac{aN}{2\sqrt{2\dot{r}}} e^{i\left(\phi - \pi\frac{q_r^2}{\dot{r}}\right)} \{S(Z_r) - S(Y_r) + i[C(Y_r) - C(Z_r)]\}. \quad (2.38)$$

Using the correlation technique, the coherent response can be recovered by convolving local Fourier amplitudes with the “frequency-reversed” and complex-conjugated template as defined by eqn. 2.38. This response, at average Fourier frequency r_c and Fourier frequency derivative \dot{r} can therefore be written as

$$A_{r_c, \dot{r}} = \sum_{k=[r]-m/2}^{k=[r]+m/2} A_k \frac{1}{\sqrt{2\dot{r}}} \left(e^{i\pi\frac{q_k^2}{\dot{r}}} \{S(Z_k) - S(Y_k) - i[C(Y_k) - C(Z_k)]\} \right). \quad (2.39)$$

Eqn. 2.39 takes into account the fact that the signal has been “spread” relatively evenly into the \dot{r} closest frequency bins to r_c , while an additional small amount of signal has “leaked” into bins further away — much like the non-zero wings of the

sinc-like response to a constant frequency signal. As a rule of thumb, the correct Fourier amplitude will be well-approximated if m is chosen such that $\dot{r} < m \lesssim 2\dot{r}$.

Large scale searches of pulsations with constant frequency derivatives have been conducted using the correlation technique. A successful example is a search for pulsars in globular cluster NGC 6544 using data taken with the Parkes radio telescope (Ransom et al. 2001b). The search was conducted using an FFT of 13865600 points over Fourier \dot{r} values from -100 to 100 with a step-size of $\Delta\dot{r} = 2$ and included the calculation of amplitudes at half-bin frequency intervals. The search resulted in the detection of the 3.06 ms PSR J1807-2459 in a low-mass binary with orbital period 1.7 h, the second shortest radio pulsar orbital period known. A detailed view of the pulsar's fundamental harmonic is shown in Fig. 2.6. The plot was calculated using the correlation technique with spacings of $\Delta r = 0.0625$ and $\Delta\dot{r} = 0.25$. The generation of such a piece of the f - \dot{f} plane takes only a fraction of a second on a rather modest workstation.

Vector Bending

Vector bending is one of the simplest and most straightforward methods to correct a Fourier response that has been smeared over several local frequency bins. As we described in §2.2.1, the DFT can be thought of as the vector addition of N complex numbers. This addition produces a straight line in the complex plane for a coherently detected sinusoid. For a sinusoid with a non-integral or time-varying pulsation frequency, the standard DFT addition produces a curved shape (see Fig. 2.1). Since the amplitude of the Fourier response is the distance between the origin and the end-point of the vector addition, any curvature in the vector addition implies non-optimal signal detection.

The precise shape of the response curve in the complex plane depends on the mismatch of the signal's (possibly time-dependent) pulsation frequency and the frequency used in forming the DFT addition (i.e. the closest FFT bin). Regardless of the shape, though, for short enough segments of the curve, the segments differ little from straight lines. We can therefore approximate the shape of the curve as a sum of G linear segments, each of which contains $\frac{N}{G}$ points from the full-resolution vector addition. In terms of the r^{th} DFT amplitude, we can write this as

$$A_r = \sum_{g=0}^{G-1} B_{r,g} = \sum_{g=0}^{G-1} \sum_{h=0}^{\frac{N}{G}-1} n_j e^{-2\pi i j r / N}, \quad (2.40)$$

where $j = g\frac{N}{G} + h$. This is equivalent to calculating and then summing G independent DFTs (the $B_{r,g}$), each of which suffers virtually no loss in sensitivity when the curvature over a segment is small.

Using these vector addition segments or sub-vectors, we can correct for the loss of sensitivity due to curvature by simply straightening the vector addition. If we can predict the true pulsation frequency of a signal as a function of time, we can predict how much curvature will accumulate in each sub-vector and then remove it by rotating the segment appropriately.

For our example of a signal with constant frequency derivative, the instantaneous phase (ignoring the intrinsic phase of the signal) is equal to

$$\Phi_{true}(u) = 2\pi \int r(u) du \quad (2.41a)$$

$$= 2\pi \int (r_o + \dot{r}u) du \quad (2.41b)$$

$$= 2\pi r_o u + \pi \dot{r} u^2, \quad (2.41c)$$

where $r_o = f_o T$ is the initial pulsation frequency of the signal, and $\dot{r} = \dot{f} T^2$ is the frequency derivative. The process of taking a DFT removes an instantaneous phase equivalent to $\Phi_{DFT}(u) = 2\pi r u$ from the signal (see eqn. 2.2). So the instantaneous phase error is equal to

$$\Phi_{error}(u) = \Phi_{true}(u) - \Phi_{DFT}(u) \quad (2.42a)$$

$$= 2\pi(r_o - r)u + \pi \dot{r} u^2. \quad (2.42b)$$

Therefore, to correct a particular signal using vector bending, we first calculate the $B_{r,g}$ using eqn. 2.40 for a particular Fourier frequency r (such as the frequency of a known pulsar). Now we attempt to un-bend the full Fourier response by summing the $B_{r,g}$ after correcting for the phase errors $\Phi_{error}(u_g)$ as defined by eqn. 2.42a. The corrected response is equal to

$$A_{r,\dot{r}} = \sum_{g=0}^{G-1} B_{r,g} e^{-i\Phi_{error}(u_g)}. \quad (2.43)$$

A choice of $G \sim 10^3$ will essentially eliminate the loss of response for reasonable frequency offsets and frequency derivatives (i.e. less than a few 10s of Fourier bins).

While impractical for large-scale searches due to the fact that the $B_{r,g}$ must be recomputed every few r , vector bending offers significant computational advantages in certain situations. In particular, X-ray observations often consist of short (< 1 hour) “on-source” segments separated by hours, days, or even weeks of “off-source” time (see also §2.3.4). An FFT of the entire time series might be prohibitively expensive. However, if we can determine an “initial guess” frequency (e.g. by FFTing one segment of the observation or from an ephemeris), we can quickly calculate the $B_{k,g}$ at this frequency from the “on-source” intervals alone. We can reconstruct the

$f-\dot{f}$ plane around our frequency of interest without ever creating the full “filled-in” or padded time series, let alone calculating a potentially huge FFT. Such techniques have allowed us to perform frequency analysis of very-long stretches of data from ROSAT observations of PSR B0540-59 (Eikenberry et al. 1998). Similarly, Fig. 2.1 shows the use of the method for a 2.4 day observation of the Crab pulsar.

2.5 Signal Property Estimation

Besides correcting for losses of sensitivity, Fourier interpolation and the other response correcting techniques mentioned in §2.4.2 allow us to determine other useful properties of a detected signal. Using detailed amplitude and phase information from a signal’s Fourier harmonics, we can estimate properties such as the statistical significance of the signal, the location and duration of the signal in the time series (the “centroid” and “purity” respectively), the precise pulsation frequency, and estimates of the measurement errors for Fourier power and phase.

The first step when estimating signal properties in the Fourier domain is to isolate the true peak of the Fourier response in power. This is easily accomplished by using the matched filtering techniques to generate an oversampled grid of amplitudes near and around the signal candidate (see Figs. 2.7 and 2.6). Simple optimization algorithms such as the downhill simplex method can then be used to refine the peak location (e.g. Press et al. 1992). Once the peak has been located, estimates of the first and second derivatives of power and phase with respect to Fourier frequency, obtained using Fourier interpolation, can be used to calculate various useful signal properties (Middleditch et al. 1993).

2.5.1 Power, Phase and Signal Amplitude

When the peak of the Fourier response has been located as a function of Fourier frequency and the other search parameters, the measured power is defined as

$$P_{meas} = \frac{|A_{r,\dots}|^2}{P_{norm}}, \quad (2.44)$$

where P_{norm} is the expected noise power and is usually described by one of $N \langle d_j^2 \rangle$, P_{local} , or n_{ph} as discussed in §2.3.1. Groth (1975, see §2.3.3) showed that since the measured power is a random variable due to the presence of noise, its variance is $2P_{signal} + 1$, where P_{signal} is the power caused by the signal. Since we do not *a priori* know the true signal power, a good estimate for the variance of the measured power is simply $2P_{meas} - 1$ since $\langle P_{meas} \rangle = P_{signal} + 1$.




Figure 2.7 The theoretical response for a boxcar-windowed signal with a constant frequency derivative. Contours plotted are 10%, 30%, 50%, 70%, and 90% of peak (i.e. fully corrected) power. The top panel shows the familiar sinc response of the signal along the $\dot{f} = 0$ line. The right panel shows a similar cut along the $\Delta r = 0$ line (i.e. the calculated average frequency is the true average signal frequency). The relatively uniform spread of signal power over the local power spectrum is apparent for all values of frequency derivative.

Using P_{meas} , as well as the knowledge that a sinusoid of amplitude a in a noisy time series produces a power with an expectation value of $\langle P_{meas} \rangle = \frac{a^2 N^2}{4P_{norm}} + 1$ (see §2.3.2), we can estimate the signal amplitude as

$$\langle a \rangle = \frac{2}{N} \sqrt{P_{norm} (P_{meas} - 1)}. \quad (2.45)$$

For binned data containing a signal with Fourier frequency r , the measured power should be multiplied by $1/\text{sinc}^2(\pi r/N)$ to correct for the loss in sensitivity due to binning (see §2.3.4). Vaughan et al. (1994) provide detailed instructions on how to estimate upper limits on pulsation amplitudes as well as estimates of the overall sensitivity of a search.

The statistical significance of a signal is also determined by P_{meas} . The probability that noise caused a particular power to equal or exceed P_{meas} is given by $e^{-P_{meas}}$ (eqn. 2.10 with $P' = P_{meas}$). But for a search over N_{IFS} independent Fourier powers, the probability that at least one of the noise powers exceeds P_{meas} is given by

$$Prob(P_{noise} \geq P_{meas}) = 1 - (1 - e^{-P_{meas}})^{N_{IFS}}. \quad (2.46)$$

Vaughan et al. (1994) show how to use this information to set detection thresholds that minimize the number of spurious candidate signals and give high confidence that signals with powers above the thresholds are real.

Using the real and imaginary parts of the peak Fourier response, we can also calculate the phase of the sinusoidal signal as

$$\phi_{meas} = \arctan \left[\frac{\text{Im}(A_{r,\dots})}{\text{Re}(A_{r,\dots})} \right] \quad (2.47)$$

radians. Using similar arguments as for the measured power, the variance of the measured phase is approximately $1/(2P_{meas} - 1)$ radians.

2.5.2 Signal Location and Duration in Time

Astronomical observations of pulsations effectively consist of a *window* of on-source time where pulsations are present, and the rest of the time when they are not. For most of this paper we have assumed that a signal is present throughout the observation as evidenced by the limits of integration for eqn. 2.2, which in time-normalized units go from $0 \leq u \leq 1$, or equivalently, from $0 \leq t \leq T$. In effect, pulsations such as that defined in eqn. 2.14a are multiplied by a square window function defined as 1 during the observation and 0 at all other times. This window function is simply a property of the DFT and is due to the finite duration of our observation.

It is possible, though, and for various reasons often likely, that a signal we are observing turns on and off or varies in intensity during an observation. The behavior of the signal itself effectively defines a new window function, $W(u)$. By measuring the moments of this window function with respect to time, we can determine approximately where in our data a signal is located and for how long.

The approximate location of a signal in a time-normalized time series is described by its centroid, $\hat{C} = \langle u \rangle = \frac{\langle t \rangle}{T}$ which is proportional to the first moment of the window function with respect to time. More specifically,

$$\hat{C} = \frac{\int_{-\infty}^{\infty} u W(u) du}{\int_{-\infty}^{\infty} W(u) du}. \quad (2.48)$$

Middleditch & Cordova (1982) wrote this in terms of the measured Fourier response as

$$\hat{C} = -\frac{1}{2\pi} \frac{\partial \phi(r_o)}{\partial r}, \quad (2.49)$$

where $\phi(r_o)$ is the phase measured at the peak of the Fourier response (i.e. ϕ_{meas}) and r is the Fourier frequency (see Appendix B for a derivation). Signals present throughout an observation have $\hat{C} = 1/2$ while those present in only the first or second halves of the observation have $\hat{C} = 1/4$ or $\hat{C} = 3/4$ respectively.

The second moment of the window function with respect to time is related to the moment of inertia of a function and can therefore be used to estimate the root-mean-squared (rms) dispersion of the pulsations in time about the centroid. Using this information, Middleditch & Cordova (1982) defined a parameter called the ‘‘purity’’ (and symbolized by α) as

$$\alpha = \frac{1}{\pi} \sqrt{-\frac{3}{2P(r_o)} \frac{\partial^2 P(r_o)}{\partial r^2}}, \quad (2.50)$$

where $P(r_o)$ is the measured power at the the peak of the Fourier response (i.e. P_{meas} , see Appendix C for a derivation). The scaling in eqn. 2.50 is chosen such that the rms dispersion of the signal about the centroid for a window function $W(u)$, is equivalent to that of a rectangular window function of duration α (in units of the time series length) centered on the centroid. A signal present throughout the data would have $\alpha = 1$, while one present in only half the data (in a continuous section) would have $\alpha = 1/2$. Signals present only at the start and end of an observation but absent in the middle have $\alpha > 1$. Purity can also help to identify sidelobes caused by a periodic modulation of a signal as these Fourier amplitudes have $\alpha = \sqrt{3}$.

Since the location and duration of a signal in an time series affects the Fourier response, it is important to understand how eqn. 2.18c changes if a signal is present

during only part of an observation. In Appendix D, we show that when close to the peak of a signal's Fourier response,

$$A_r \simeq A_o e^{-2\pi i \hat{C}(r-r_o)} \text{sinc}[\pi\alpha(r-r_o)] \quad (2.51)$$

where $A_o = \frac{Na}{2} e^{i\phi_o}$, ϕ_o is the intrinsic phase of the signal, and r_o is the true signal frequency (in FFT bins). This equation demonstrates that for centroids different from $1/2$, the phase shift between consecutive FFT bins differs from the π radians shown in eqn. 2.18c. Similarly, for purity values different from 1, neighboring FFT bins show more or less correlation with each other (i.e. the central peak of the sinc function changes its width).

2.5.3 The Pulsation Frequency and Frequency Derivative

The true pulsation frequency of the signal is located at the point where $\frac{\partial P}{\partial r} = 0$. Furthermore, given the response in eqn. 2.51, we can show (Appendix E) that the uncertainty in this measurement (in Fourier bins) is given by

$$\sigma_r = \frac{3}{\pi\alpha\sqrt{6P_{meas}}}. \quad (2.52)$$

This uncertainty is considerably smaller than the often-quoted “frequency error” for the FFT of one bin width, which is simply the frequency resolution returned by the FFT algorithm.

If the correlation method is used to isolate a peak in the $f-\dot{f}$ plane as shown in Figures 2.7 and 2.6, we can calculate the uncertainty in the measured \dot{r} value by using similar arguments and methods as for the frequency uncertainty (see Appendix F for a derivation). The uncertainty in the \dot{r} (in Fourier bins) is approximately

$$\sigma_{\dot{r}} = \frac{3\sqrt{10}}{\pi\sqrt{P_{meas}}}. \quad (2.53)$$

2.6 Conclusions

In this paper we have described techniques that allow sophisticated and fully coherent Fourier analysis of very long time series. Most of these techniques use the wealth of information provided by the Fourier phases — information discarded during “standard” analyses based on raw power spectra.

Significant gains in sensitivity and efficiency are possible when using Fourier phase information during the search for periodic signals (using the Fourier domain matched filtering techniques described in §2.4.2) and when characterizing signals that are known to be present in the data (using the parameters described in §2.5).

The methods of Fourier domain matched filtering allow efficient, memory local, and inherently parallel analysis of extremely long time series with only modest computational resources. Billion point out-of-core FFTs followed by fully coherent matched filtering pulsation searches are possible on “standard” workstations. More traditional time domain based techniques (such as acceleration searches performed by stretching or compressing the time series followed by large in-core FFTs) on similarly sized time series require specialized high-performance computing resources, assuming they can be performed at all.

As astronomical instruments become more sophisticated and specialized, time series of ever increasing duration and time resolution will appear. The Fourier domain techniques described in this paper should prove to be essential tools in their analysis.

Acknowledgments

We would like to thank G. Fazio for supporting our research and encouraging us in this work. Additional thanks go to B. Schmidt, F. Seward, R. Narayan, and J. Grindlay for their encouragement and comments. Many of the computations for this paper were performed on equipment purchased with NSF grant PHY 9507695. S. Eikenberry is supported by a NSF CAREER Grant.

Chapter 3

A Binary Millisecond Pulsar in Globular Cluster NGC 6544

Scott M. Ransom, Lincoln J. Greenhill, James R. Herrnstein, Richard N. Manchester, Fernando Camilo, Stephen S. Eikenberry, & Andrew G. Lyne 2001, *The Astrophysical Journal*, **546**, L25

Abstract

We report the detection of a new 3.06 ms binary pulsar in the globular cluster NGC6544 using a Fourier-domain “acceleration” search. With an implied companion mass of $\sim 0.01 M_{\odot}$ and an orbital period of only $P_b \sim 1.7$ hours, it displays very similar orbital properties to many pulsars which are eclipsed by their companion winds. The orbital period is the second shortest of known binary pulsars after 47 Tuc R. The measured flux density of 1.3 ± 0.4 mJy at 1332 MHz indicates that the pulsar is almost certainly the known steep-spectrum point source near the core of NGC6544.

3.1 Introduction

Globular clusters are rich sources of radio pulsars. Since the discovery of the first pulsar in a cluster by Lyne et al. (1987), a series of deep searches have accumulated almost 50 pulsars in globular clusters, with the majority of these being binary millisecond pulsars (MSPs). See Lyne (1995) and Kulkarni & Anderson (1996) for reviews.

Fruchter & Goss (1990; 2000, hereafter FG00) used the Very Large Array (VLA) to detect steep-spectrum radio sources, which are likely pulsars, in numerous

globular clusters. They discovered a relatively bright point-like source within $5''$ of the optical center of cluster NGC6544 (1.2 ± 0.1 mJy at 20 cm and 6.5 ± 0.8 mJy at 90 cm). Biggs & Lyne (1996) searched NGC6544 at 50 cm as part of a survey of globular clusters and found no pulsars down to a limiting flux density of ~ 4 mJy.

It has been argued that Doppler smearing from orbital motion could explain the lack of detection of pulsations from the pulsar-like sources found in VLA surveys (FG00). The sensitivities of observations whose durations exceed even a few percent of the orbital period may be drastically reduced for traditional analyses (Johnston & Kulkarni 1991). Various schemes have been developed to correct, at least partially, for this loss of sensitivity when the Doppler effect of the orbital motion can be approximated as a constant frequency derivative (e.g. Middleditch & Kristian 1984; Anderson et al. 1990; Wood et al. 1991). These “acceleration” searches have met with significant success, including the discovery of nine new binary MSPs in the globular cluster 47 Tucanae (Camilo et al. 2000).

We have discovered a binary MSP in NGC6544 using Fourier-domain techniques to correct for pulsar orbital motion. This pulsar is almost certainly the radio source reported by FG00 near the center of the cluster. Results from our single observation indicate that this object shares many of the properties displayed by the burgeoning class of eclipsing binary MSPs. These systems are characterized by short orbital periods ($P_b \sim 1 - 10$ hours) and very low-mass companions ($m_2 \sin i \lesssim 0.05 M_\odot$) (Rasio et al. 2000; Nice et al. 2000). D’Amico et al. (2001) independently discovered this pulsar and report the results of follow-up observations.

3.2 Observations and Data Reduction

We observed NGC6544, a dense core-collapsed globular cluster, on 1998 February 7 with the multibeam data acquisition system on the Parkes radio telescope in Australia. Signals from each of two orthogonal linear polarizations were measured and summed from 256 contiguous 0.25 MHz wide channels centered at 1332 MHz. Each channel was one-bit sampled at 8 kHz and written to magnetic tape. The full observation comprised 28.9 minutes of data. Staveley-Smith et al. (1996) and Lyne et al. (2000a) discuss the observing system in detail.

Since no pulsars were known in this cluster before our observation, we searched a wide range of possible dispersion measures (DMs) for candidate signals. We dedispersed the data into time series of 13865600 points using 600 trial DMs from 0 to 600 pc cm^{-3} , in increments of 1 pc cm^{-3} . This stepsize allowed a worst-case dispersion smearing over our bandwidth of ~ 0.1 ms. After barycentering the data using the JPL DE200 ephemeris (Standish 1982), we performed a Fast Fourier Transform (FFT) on each time series and saved the resultant spectra to disk.

Traditional time-domain acceleration searches are performed by stretching or compressing a time series to compensate for a constant frequency derivative and then Fourier transforming the resulting series. We have developed a Fourier-domain acceleration search (S. Ransom, S. Eikenberry, & J. Middleditch 2001, in preparation) to compute the coherent Fourier response over portions of the frequency–frequency derivative ($f-\dot{f}$) plane using only local Fourier amplitudes from an FFT of the whole dataset. Similar techniques were described in Middleditch et al. (1993).

The method correlates a predicted Fourier response with a subset of the complex Fourier amplitudes in the initial full-length FFT. We effectively apply a matched filter to the data which “sweeps up” all of the signal power that the orbital motion spreads over nearby frequency bins.

Fourier-domain acceleration searches offer several significant advantages over their time-domain counterparts:

1. Fourier methods do not require stretching or compressing the time series. Time-domain stretching or compressing is usually performed using linear interpolation, which changes the statistics of the data. Portions of the data are effectively two-bin averaged which decreases sensitivity to high frequency signals.
2. Time-domain techniques require a full-length FFT for each trial acceleration. For very long integrations (which are common in cluster searches) that do not fit into the computer’s core memory, FFTs take orders-of-magnitude more time to compute than usual (e.g. Bailey 1990). The Fourier technique requires only a single full-length FFT for each observation (or trial DM).
3. Correlations using only localized Fourier amplitudes are fast and memory efficient. The correlations can always be performed in core memory and are calculated efficiently using short FFTs and pre-computed response templates. The memory locality also allows efficient parallelization of Fourier-domain searches.
4. The Fourier-domain method allows the calculation of only independent $f-\dot{f}$ trials. With time-domain techniques, as described by Camilo et al. (2000), any choice of acceleration stepsize results in either under- or over-sampling accelerations for the vast majority of frequencies searched. This amounts to loss of sensitivity or wasted CPU cycles respectively. A small amount of over-sampling can be incorporated to eliminate “scalloping” in both the f and \dot{f} directions (see e.g. van der Klis 1989).

We performed Fourier-domain acceleration searches on each of the long FFTs looking for signals that drifted by up to 100 Fourier frequency bins during the

observation. If we define z to be the number of Fourier bins that the pulsar frequency (f_o) drifts over the course of an integration, the corresponding average acceleration of the pulsar is $a = zc/(T^2 f_o)$, where T is the total integration time. For this observation, a frequency drift of 100 Fourier bins corresponds to an acceleration of 500 m s^{-2} for a 10 Hz signal or 5 m s^{-2} for a 1000 Hz signal. Since most known binary pulsars show maximal accelerations of only a few m s^{-2} , we were sensitive to all but the most exotic binaries or pulsars with spin periods (P_{spin}) much less than $\sim 2 \text{ ms}$. Pulse profiles of strong candidates were folded at a series of trial DMs, frequencies, and frequency derivatives around the acceleration search values in order to maximize the signal-to-noise.

3.3 Results

We detected a strong pulsar candidate (power/average local power ~ 165 corresponding to a single trial significance of 18σ) at a DM of 134 pc cm^{-3} . The average frequency was 326.85895 Hz and the signal had drifted by 8.6 Fourier bins ($a = 2.6 \text{ m s}^{-2}$) during the observation (see Figure 2.6). Search techniques that do not account for acceleration would have detected the signal at about ~ 34 times the average local power level, corresponding to an 8σ detection.

Optimization of the candidate by pulse-folding produced a narrow pulsar-like profile at a localized maximum in DM of 134 pc cm^{-3} indicating that the signal was almost certainly not due to terrestrial interference. The Taylor & Cordes (1993) free-electron model gives an estimated distance of 3.3 kpc ($\sim 25\%$ error) for this DM in the direction of NGC6544, in fair agreement with the published cluster distance of $\sim 2.6 \text{ kpc}$ obtained by fitting the mean V magnitudes of horizontal branch stars (Harris 1996).

We estimated a pulsar flux density at 1332 MHz of $1.3 \pm 0.4 \text{ mJy}$ by comparing the integrated pulsed flux and observed noise to the predicted system noise level given a system sensitivity and total system temperature of $T_{tot} \sim T_{sys} + T_{sky}$. The multibeam system has a cold-sky $T_{sys} \sim 21 \text{ K}$ and a sensitivity of 1.36 Jy K^{-1} . Haslam et al. (1982) measured a sky temperature at 408 MHz in the direction of NGC6544 of $T_{sky} \sim 260 \text{ K}$. Assuming a spectral index of -2.6 , which is typical for this region (Lawson et al. 1987), we obtain $T_{sky} \sim 12 \text{ K}$ and $T_{tot} \sim 33 \text{ K}$ at 1332 MHz . FG00 reported a 20 cm flux density of $1.2 \pm 0.1 \text{ mJy}$ and a spectral index of -1.1 between 90 cm and 20 cm which corresponds to a flux density of $\sim 1.4 \pm 0.1 \text{ mJy}$ at 1332 MHz . Our measured flux density agrees with that of FG00 within the errors, suggesting that the pulsar may be the previously unidentified point source near the core of NGC6544. Since the DM-based distance estimate for the pulsar matches published values for NGC6544, cluster membership is very likely.

TABLE 3.1
Parameters for PSR J1807–2459

Parameter	Value
Right Ascension ^a (J2000)	18 ^h 07 ^m 20 ^s .36(2)
Declination ^a (J2000)	−24° 59′ 52″.6(4)
Dispersion Measure (pc cm ^{−3})	134(2)
Flux Density at 1332 MHz (mJy)	1.3(4)
Pulse FWHM at 1332 MHz (%)	11(3)
Pulsar Period (s)	0.003059447(2)
Epoch (MJD)	50851.9
Orbital Period (days)	0.070(3)
Projected Semi-Major Axis (lt-s)	0.0116(8)
Eccentricity ^b	0.0
Epoch of Periastron (MJD)	50851.887(1)
Mass Function (M _⊙)	3.4(5) × 10 ^{−7}
Companion Mass Limit ^c (M _⊙)	≥ 0.009

^a A. Fruchter 2000, private communication.

Fruchter & Goss discovered an error in the 20 cm position reported in FG00. The corrected position shown above will appear in an erratum.

^b Parameter assumed to be exactly zero. See §3.3.

^c Assuming a pulsar mass of $m_1 = 1.4 M_{\odot}$.

NOTE — Numbers in parentheses represent 2σ uncertainties in the last digit quoted.

The relatively large signal acceleration indicates that the pulsar is in a short period binary with a low-mass companion. These systems, almost without exception, have eccentricities very near to zero due to tidal circularization (see e.g. Phinney & Kulkarni 1994 and Rasio & Heggie 1995).

In order to study the assumption that the pulsar is a member of a low-mass binary with an eccentricity of zero, we split the observation into 16 equal-length parts and folded the data in each at the nominal pulsar frequency. We determined phase offsets for each of the pulse arrival times and performed a Levenberg-Marquardt least-squares fit of a circular orbit (i.e. sinusoid) to the phase offsets. The results are tabulated in Table 3.1 and the residuals and error ellipses of the fit are shown in Figures 3.1 and 3.2 respectively.

We found no other candidates in our search of NGC6544 to a limiting flux density of ~ 1.1 mJy for 1 ms period signals and ~ 0.5 mJy for periods $\gtrsim 10$ ms.




Figure 3.1 Orbital fit for PSR J1807–2459 assuming an eccentricity of zero. The error-bars plotted are $\pm 2\sigma$ uncertainties in the pulse arrival times. An eclipse, if present, should occur near True Anomaly = 90° , which corresponds to superior conjunction. The apparent lack of an eclipse could be due to the relatively high observing frequency (1332 MHz). The inset plot shows the observed pulse profile after folding with the best-fit orbital ephemeris and pulsar frequency. The error-bar indicates $\pm 2\sigma$ uncertainties in the profile values.




Figure 3.2 Error ellipses (from inner to outer of 1σ , 2σ , and 3σ) for the orbital period and semi-major axis of PSR J1807–2459. The large covariance is due to the fact that the observation covered only $\sim 30\%$ of the orbital period.

3.4 Discussion

There are currently eight binary pulsar systems that display similar orbital characteristics to those of PSR J1807–2459: PSR B1957+20 (Fruchter et al. 1988) and PSR J2051–0827 (Stappers et al. 1996) in the Galactic disk, PSR J1910+0004 in NGC6760 (Deich et al. 1993), and pulsars 47 Tuc I, J, O, P, and R (Camilo et al. 2000). Each of these systems displays a pulse period of a few milliseconds, an orbital period of 1-10 hours, and a very low-mass companion ($m_2 \sin i \lesssim 0.05 M_\odot$). Some of these systems display radio eclipses, particularly at longer wavelengths (i.e. 50-90 cm).

Our current 1332-MHz data show no evidence for an eclipse (see Figure 3.1) at or near superior conjunction, but we cannot rule out the possibility of eclipses at longer wavelengths. Pulsars J2051–0827 and 47 Tuc J show eclipses at radio wavelengths between 50-90 cm, but are detected at all orbital phases at 20 cm, scintillation permitting (Stappers et al. 1996; Camilo et al. 2000). PSR J1910+0004 in NGC6760 was detected at all orbital phases at 20 cm, but Deich et al. (1993) could not rule out short duration eclipses or eclipses at longer wavelengths.

Most formation theories for short-period binary MSPs in globular clusters involve dynamical interactions with primordial binaries which have been shown to exist in significant numbers (Hut et al. 1992). Rasio et al. (2000) describe a recent variant of the standard “spin-up” model (see Phinney & Kulkarni 1994, for a review) using an exchange interaction between a neutron star (NS) and a hard primordial binary containing at least one relatively massive main sequence (MS) star ($\sim 1-3 M_\odot$). If the exchange interaction results in a binary containing the NS and the massive MS star, the system enters a common-envelope (CE) phase once the MS star evolves and fills its Roche lobe. At the end of the CE phase, the NS emerges with a very low-mass companion in a very short period circular orbit. Systems with orbital periods $P_b \lesssim 8$ hours undergo evolution due to gravitational radiation and enter a second phase of mass transfer that spins the NS up to millisecond periods.

The scenario described by Rasio et al. (2000) is possible only if the exchange interactions occurred at a time when massive MS stars were still present in the cluster. The NSs and primordial binaries must also have undergone mass segregation and ended up near the cluster core before the MS stars evolved. The timescale for mass segregation is the half-mass relaxation time, which for NGC6544 is ~ 200 Myr (Harris 1996). This is significantly less than the MS lifetime of a 1-3 M_\odot star and implies that this scenario may have produced PSR J1807–2459.

Alternatives to “spin-up” models exist in which the NS is created either by accretion induced collapse of a white dwarf (WD) (e.g. Bailyn & Grindlay 1990), or the coalescence of a pair of massive WDs (Chen & Leonard 1993). Pairs of

primordial binaries interact to produce hard binaries which contain a massive WD or WDs that these models require. The NSs that result from these models are born spinning rapidly ($P_{spin} \sim 2 - 10$ ms), have relatively weak ($10^9 - 10^{10}$ G) magnetic fields characteristic of MSPs, and may have low-mass companions in very short period orbits. Interestingly, these systems could be the progenitors of the low-mass X-ray binaries (LMXBs) rather than their descendants (Chen et al. 1993).

Future observations of PSR J1807–2459 will lead to improved measurements of rotational and orbital parameters, an estimate of or upper limit to the orbital eccentricity, and perhaps information about the circumstellar medium if eclipses are detected. Rasio et al. (2000) predict that many similar systems must exist in the globular cluster system — some with orbital periods as short as ~ 15 min. If recent successes such as the 20 cm searches of 47 Tucanae (Camilo et al. 2000) are any indication, this prediction will soon be put to the test. The ever-increasing speed of computers and improved algorithms for binary pulsar detection (e.g. Ransom 2000) will allow the analysis of extremely long observations (i.e. days or weeks) that should reveal even the weakest of pulsars in the tightest of binaries.

Acknowledgements

We would like to thank Vicky Kaspi, Josh Grindlay, Giovanni Fazio, and Nichi D’Amico for useful discussions. FC is supported by NASA grant NAG 5-3229. SE is supported by a NSF CAREER grant.

Chapter 4

A New Search Technique for Short Orbital Period Binary Pulsars

Scott M. Ransom, James M. Cordes, & Stephen S. Eikenberry 2001, to be submitted to *The Astrophysical Journal*

Abstract

We describe a new and efficient technique to search for short period binary pulsars in observations longer than the orbital period. The orbital motion of the pulsar during long observations effectively modulates the phase of the pulsar signal causing sidebands to appear around the pulsar spin frequency and its harmonics in the Fourier transform. For the majority of binary radio pulsars or LMXBs, large numbers of sidebands are present allowing efficient searches using Fourier transforms of short portions of the original power spectrum. Analysis of the complex amplitudes and phases of the sidebands can provide enough information to solve for the Keplerian orbital parameters. This technique is particularly applicable to radio pulsar searches in globular clusters and searches for coherent X-ray pulsations from LMXBs.

4.1 Introduction

Short orbital-period binary pulsars have proven to be sensitive laboratories that allow tests of a variety of physical processes including theories of gravitation. Unfortunately, except for those systems containing the brightest pulsars, they are notoriously difficult to detect. When the observation time T_{obs} is greater than a fraction of the orbital period P_{orb} , Doppler effects due to the orbital motion of the pulsar cause drastic reductions in the sensitivities of pulsar searches (Johnston &

Kulkarni 1991).

Numerous groups (e.g. Middleditch & Kristian 1984; Anderson et al. 1990; Wood et al. 1991) have developed variations of so called “acceleration” searches in order to mitigate the loss in sensitivity caused by orbital motion. These searches can almost completely recover this lost sensitivity in certain situations by taking advantage of the fact that orbital motion causes an approximately linear drift of the pulsar spin frequency f_{spin} when $T_{obs} \lesssim P_{orb}/10$. Recent acceleration searches of globular clusters have discovered numerous binary systems, some with orbital periods as short as ~ 90 minutes (Camilo et al. 2000; Ransom et al. 2001b; D’Amico et al. 2001).

Acceleration searches have a severe flaw, however. By their very nature they place an upper limit on T_{obs} for any given orbital period. Since sensitivities improve as $T_{obs}^{1/2}$, acceleration searches are limited to discovering only the very brightest binary pulsars with orbital periods less than a few hours — which are some of the most interesting scientifically. For a 2 ms pulsar in a 1 hour orbit with a low mass white dwarf (WD) companion, the optimal observation time T_{best} for an acceleration search (using eqn. 22 from Johnston & Kulkarni 1991) is only $\sim 300 - 600$ seconds. If this pulsar were located in a globular cluster, where observations of 5 – 10 hours in duration are possible, acceleration searches would be a factor of $\sqrt{T_{obs}/T_{best}} \sim 8$ times less sensitive than a coherent search of a full observation.

We have developed an efficient search technique for short orbital period binary pulsars that complements the use of acceleration searches. Since it requires $T_{obs} \gtrsim 1.5P_{orb}$ and its sensitivity improves as the ratio $r_{orb} = T_{obs}/P_{orb}$ increases, it allows the search of very interesting portions of orbital parameter space for the first time. When many orbits are present in an observation its sensitivity approaches that of a fully coherent search.

4.2 Phase Modulation due to Orbital Motion

During long observations of a binary pulsar where T_{obs} is longer than P_{orb} , the orbital motion causes a modulation of the phase of the pulsar signal. Differences in light travel time across the projected orbit advance or delay the pulse phase in a periodic fashion.

4.2.1 Circular Orbits

If the orbit is circular (where the eccentricity e and the angle of periapsis ω are both zero), the phase delays are sinusoidal. The fundamental harmonic¹ of the resulting signal as sampled at a telescope can be described by

$$n_j = n(u) = a \cos [2\pi r_{spin}u + \phi_{spin} + \Phi_{orb} \cos(2\pi r_{orb}u + \phi_{orb})], \quad (4.1)$$

where a and ϕ_{spin} are the amplitude and phase of the pulsation, Φ_{orb} and ϕ_{orb} are the amplitude and phase of the modulation due to the orbit, r is a Fourier frequency (i.e. $r = fT_{obs} = T_{obs}/P$), and $u = \frac{j}{N} = \frac{t_j - t_o}{T_{obs}}$ is the “normalized” time during the observation of the j^{th} sample of N such that $0 \leq u \leq 1$. We can write Φ_{orb} and ϕ_{orb} in terms of the five Keplerian orbital parameters (P_{orb} , e , w , the semi-major axis $x_{orb} = a_1 \sin(i)/c$, and the time of periapsis T_o) as

$$\Phi_{orb} = 2\pi x_{orb} f_{spin} = \frac{2\pi x_{orb}}{P_{spin}} \quad (4.2)$$

and

$$\phi_{orb} = 2\pi \frac{t_o - T_o}{P_{orb}} + \omega + \frac{\pi}{2}, \quad (4.3)$$

where the units of both are radians.

The Discrete Fourier Transform (DFT) of a simple (non-modulated) cosinusoidal signal present throughout an observation can be approximated as (see e.g. Ransom et al. 2001a)

$$A_r = \sum_{j=0}^{N-1} n_j e^{-2\pi i j r / N} \quad (4.4a)$$

$$\simeq N \int_0^1 a \cos(2\pi r_{spin}u + \phi_{spin}) e^{-2\pi i r u} du \quad (4.4b)$$

$$= \frac{aN}{2} \int_0^1 (e^{2\pi i (r_{spin} - r)u} e^{i\phi_{spin}} + e^{-2\pi i (r_{spin} + r)u} e^{-i\phi_{spin}}) du \quad (4.4c)$$

$$\simeq \frac{aN}{2} e^{i\phi_{spin}} \int_0^1 e^{2\pi i (r_{spin} - r)u} du \quad (4.4d)$$

$$= A_{coherent} e^{i\pi (r_{spin} - r)} \frac{\sin[\pi (r_{spin} - r)]}{\pi (r_{spin} - r)}, \quad (4.4e)$$

since the $r_{spin} + r$ component of the second term in eqn. 4.4c causes the vector addition to loop around the origin in the complex plane instead of adding coherently

¹The following analysis can be applied to any harmonic h of a pulsar’s signal by substituting hr_{spin} for r_{spin} and the appropriate harmonic phase for ϕ_{spin} .

in phase. The result is the traditional “sinc function”-like response of a square-windowed sinusoid, where $A_{coherent} = \frac{aN}{2} e^{i\phi_{spin}}$ represents the coherent Fourier response at the spin frequency r_{spin} . Similarly, we can approximate the DFT of our phase modulated signal from eqn. 4.1 as

$$A_r = \frac{aN}{2} \int_0^1 e^{2\pi i r_{spin} u + i\phi_{spin} + \Phi_{orb} \cos(2\pi r_{orb} u + \phi_{orb})} e^{-2\pi i r u} du \quad (4.5a)$$

$$= A_{coherent} \int_0^1 e^{i\Phi_{orb} \cos(2\pi r_{orb} u + \phi_{orb})} e^{2\pi i (r_{spin} - r) u} du. \quad (4.5b)$$

Using the Jacobi-Anger expansion (e.g. Arfken 1985)

$$e^{iz \cos(\theta)} = \sum_{m=-\infty}^{\infty} i^m J_m(z) e^{im\theta} = \sum_{m=-\infty}^{\infty} J_m(z) e^{im(\theta + \frac{\pi}{2})}, \quad (4.6)$$

where $J_m(z)$ is an integer order Bessel function of the first kind, we can expand the integrand in eqn. 4.5b to give

$$A_r = A_{coherent} \int_0^1 \sum_{m=-\infty}^{\infty} J_m(\Phi_{orb}) e^{im(2\pi r_{orb} u + \phi_{orb} + \frac{\pi}{2})} e^{2\pi i (r_{spin} - r) u} du \quad (4.7a)$$

$$= A_{coherent} \sum_{m=-\infty}^{\infty} J_m(\Phi_{orb}) e^{im(\phi_{orb} + \frac{\pi}{2})} \int_0^1 e^{2\pi i (r_{spin} + m r_{orb} - r) u} du. \quad (4.7b)$$

Since the integrand in eqn. 4.7b is identical in form to that in eqn. 4.4d, the Fourier transform of a phase modulated signal is equivalent to the Fourier transform of a series of cosinusoids at frequencies centered on the pulsar spin frequency r_{spin} , but separated from r_{spin} by $m r_{orb}$ Fourier bins. The response is therefore

$$A_r = A_{coherent} \sum_{m=-\infty}^{\infty} J_m(\Phi_{orb}) e^{im(\phi_{orb} + \frac{\pi}{2})} e^{i\pi(r_{spin} + m r_{orb} - r)} \frac{\sin[\pi(r_{spin} + m r_{orb} - r)]}{\pi(r_{spin} + m r_{orb} - r)}. \quad (4.8)$$

When r_{orb} is an integer (i.e. the observation covers an integer number of complete orbits), the Fourier response shown in eqn. 4.8 becomes particularly simple at Fourier frequencies $r_{spin} + s r_{orb}$, where s is an integer describing the sideband in question. In this case, the summation collapses to a single term when $m = s$, giving

$$A_{r_{spin} + s r_{orb}} = A_{coherent} J_s(\Phi_{orb}) e^{is(\phi_{orb} + \frac{\pi}{2})} \quad (4.9a)$$

$$= \frac{aN}{2} J_s(\Phi_{orb}) e^{is(\phi_{orb} + \frac{\pi}{2}) + i\phi_{spin}}. \quad (4.9b)$$

In effect, the phase modulated signal produces sidebands composed of a series of sidebands split in frequency from r_{spin} by $s r_{orb}$ Fourier bins, with amplitudes

proportional to $J_s(\Phi_{orb})$, and phases of $s(\phi_{orb} + \frac{\pi}{2}) + \phi_{spin}$ radians². Figure 4.1 shows the Fourier response for a “typical” pulsar-WD system in a circular 50 m orbit during an 8 hour observation.

While the number of sidebands implied by eqn. 4.8 is infinite, properties of the Bessel functions — which define the magnitude of the Fourier amplitudes of the sidebands — give a rather sharp cutoff at $\sim 2\Phi_{orb}$ sidebands (or $\sim \Phi_{orb}$ pairs of sidebands). The maximum value of $J_s(\Phi_{orb})$ also occurs near $s = \Phi_{orb}$ at a value of $\sim 0.675\Phi_{orb}^{-1/3}$, which gives the phase-modulation response its distinctive “horned” shape when Φ_{orb} is large (see Fig. 4.1). The average magnitude of the phase modulation sidebands falls off more quickly as $\sim 0.611\Phi_{orb}^{-1/2}$ (see Fig. 4.2).

For pulsars with narrow pulse shapes and many Fourier harmonics, phase modulation produces sidebands split by r_{orb} Fourier bins and centered around each spin harmonic. For harmonic number h , the phase modulation amplitude is $\Phi_{orb} = 2\pi x_{orb} h f_{spin}$. Therefore, there are h times more sidebands than around the fundamental, but the average amplitudes are smaller by a factor of $\sim h^{-1/2}$.

4.2.2 Small Amplitude Limit

When $\Phi_{orb} \ll 1$, a phase modulated signal will have only three significant peaks in its Fourier response. These peaks are located at r_{spin} and $r_{spin} \pm r_{orb}$ and have magnitudes proportional to $J_0(\Phi_{orb}) \sim 1 - \Phi_{orb}^2/4 \sim 1$ and $J_{\pm 1}(\Phi_{orb}) \sim \pm \Phi_{orb}/2$ respectively.

Middleditch et al. (1981) performed a detailed analysis of phase modulated optical pulsations in the small amplitude limit from the 7.7 second X-ray pulsar 4U 1626-67 in order to solve for the orbital parameters of the system. The separation of the sidebands provided a measurement of the ~ 42 minute orbital period and the magnitudes and phases of the three significant Fourier peaks allowed them to solve for Φ_{orb} (and therefore x_{orb}) and ϕ_{orb} directly.

4.2.3 Elliptical Orbits

Konacki & Maciejewski (1996) derived a Fourier expansion for elliptical orbital motion with the useful property that the magnitudes of higher order harmonics (which correspond to the modulation amplitude Φ_{orb}) decrease monotonically. In the limit of circular orbits the expansion reduces to the single phase-modulating sinusoid found in eqn. 4.1.

²Note that $J_{-n}(z) = (-1)^n J_n(z)$, and also that $J_n(z)$ can be negative. If the amplitude determined by the Bessel function is negative, the *measured* phase of the Fourier response will differ by π radians since the *measured* amplitude is always positive.

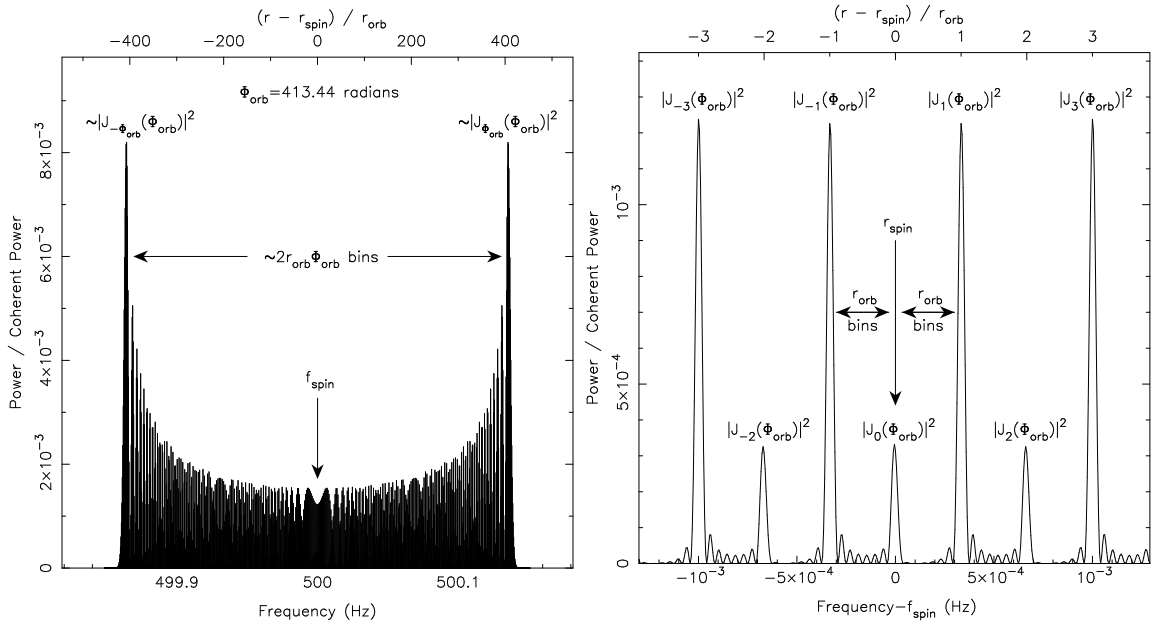


Figure 4.1 These two figures show the Fourier response of the fundamental spin harmonic of a phase modulated 2 ms binary pulsar. The properties of the response are typical for circular orbits with $\Phi_{orb} = 2\pi x_{orb} f_{spin} \gg 1$. The response was calculated for an 8 hour observation of a NS-WD binary with a 50 m orbital period, projected semi-major axis of 0.132 lt-s, orbital inclination of 60° , and a companion mass of $\sim 0.2 M_\odot$. The figure on the right shows the central portion of the full response on the left. The periodic nature of the sidebands with peak-to-peak spacing (in Fourier bins) of $r_{orb} = T_{obs}/P_{orb}$ is obvious.




Figure 4.2 Results of the measurement of the maximum and average values of $|J_s(\Phi_{orb})|$ where s are all the integers from $0 \rightarrow [\Phi_{orb}]$. The Bessel functions correspond to the magnitude of the s^{th} sideband pair of the Fourier response of a phase modulated sinusoid as a fraction of the unmodulated or fully coherent amplitude (see §4.2.1). Approximate asymptotic relations for large values of Φ_{orb} are given above the measured values.

Each harmonic of the orbital modulation generates sidebands around the spin harmonic of the form discussed in §4.2.1. The sidebands created by harmonic h of the orbital expansion are separated from each other by hr_{orb} Fourier bins. Since $\Phi_{orb,h}$ monotonically decreases for each higher harmonic, the number of sidebands generated by each harmonic decreases while their spacing and average amplitudes increase. The most important aspect of this superposition of sidebands is the fact that the overall response contains sidebands spaced by r_{orb} Fourier bins — just as for the limiting case of circular orbits.

4.3 Pulsars with Deterministic Periodic Phase Modulation

Figures 4.3 and 4.4 show P_{spin} , P_{orb} , and Φ_{orb} for 52 binary pulsars from the literature, all with orbital periods less than 10 days. A few conclusions can be reached immediately. First, millisecond pulsars (MSPs) are the most common binary pulsar. Second, there is a relatively flat distribution of P_{orb} which cuts off quite dramatically near $P_{orb} \sim 0.08$ d (~ 2 h). And third, a large fraction of the systems are located in globular clusters.

The fact that large numbers of relatively compact binary pulsar systems should contain MSPs is not surprising in light of the standard “recycling” model for MSP creation (Verbunt 1993; Phinney & Kulkarni 1994). This model explains not only the preponderance of pulsars with millisecond spin periods, but also the fact that many MSPs will have low-mass ($\sim 0.1 - 0.4 M_{\odot}$) companion stars in compact circular orbits. These recycled systems make up the majority of the pulsars shown in Figure 4.3 and have typical values of Φ_{orb} of 10^2 to a few times 10^3 radians.

The cutoff in P_{orb} at approximately two hours may be at least partly due to selection effects. The rapid orbital motion of these systems prevents their detection using conventional search techniques — including acceleration searches — for all but the brightest systems (Johnston & Kulkarni 1991). Since selection effects have almost certainly been at least partially responsible for the cutoff, it is difficult to determine how much the lack of known systems is due to a real lack of such systems in general. It is certain, though, that neutron star (NS) systems with much shorter orbital periods *do* exist, as they are observed as ultra-compact LMXBs such as 4U 1820–303 ($P_{orb} \sim 11$ m), 4U 1850–087 ($P_{orb} \sim 21$ m), 4U 1627–673 ($P_{orb} \sim 42$ m), and 4U 1916–053 ($P_{orb} \sim 50$ m) (Chou et al. 2001; Liu et al. 2001).

The large fraction of known systems inhabiting globular clusters can be explained by three basic facts. First, long observations, due to their targeted nature, allow fainter systems to be found in clusters (e.g. Anderson 1992). Second, repeated




Figure 4.3 Orbital versus spin periods for 52 currently known binary pulsars with $P_{orb} < 10$ d. The symbols denote the type of system (circles for NS-WD systems, squares for NS-NS systems, and a star for LMXB J1808.4-3658) as well as its location in the galaxy or a globular cluster (open or filled respectively). Lines of constant phase modulation amplitude $\Phi_{orb} = 2\pi x_{orb} f_{spin}$ are overlaid assuming orbital inclinations of 60° , pulsar masses of $1.4 M_\odot$, and companion masses of 0.1, 1, and $10 M_\odot$ for dotted, dashed, and solid lines respectively.

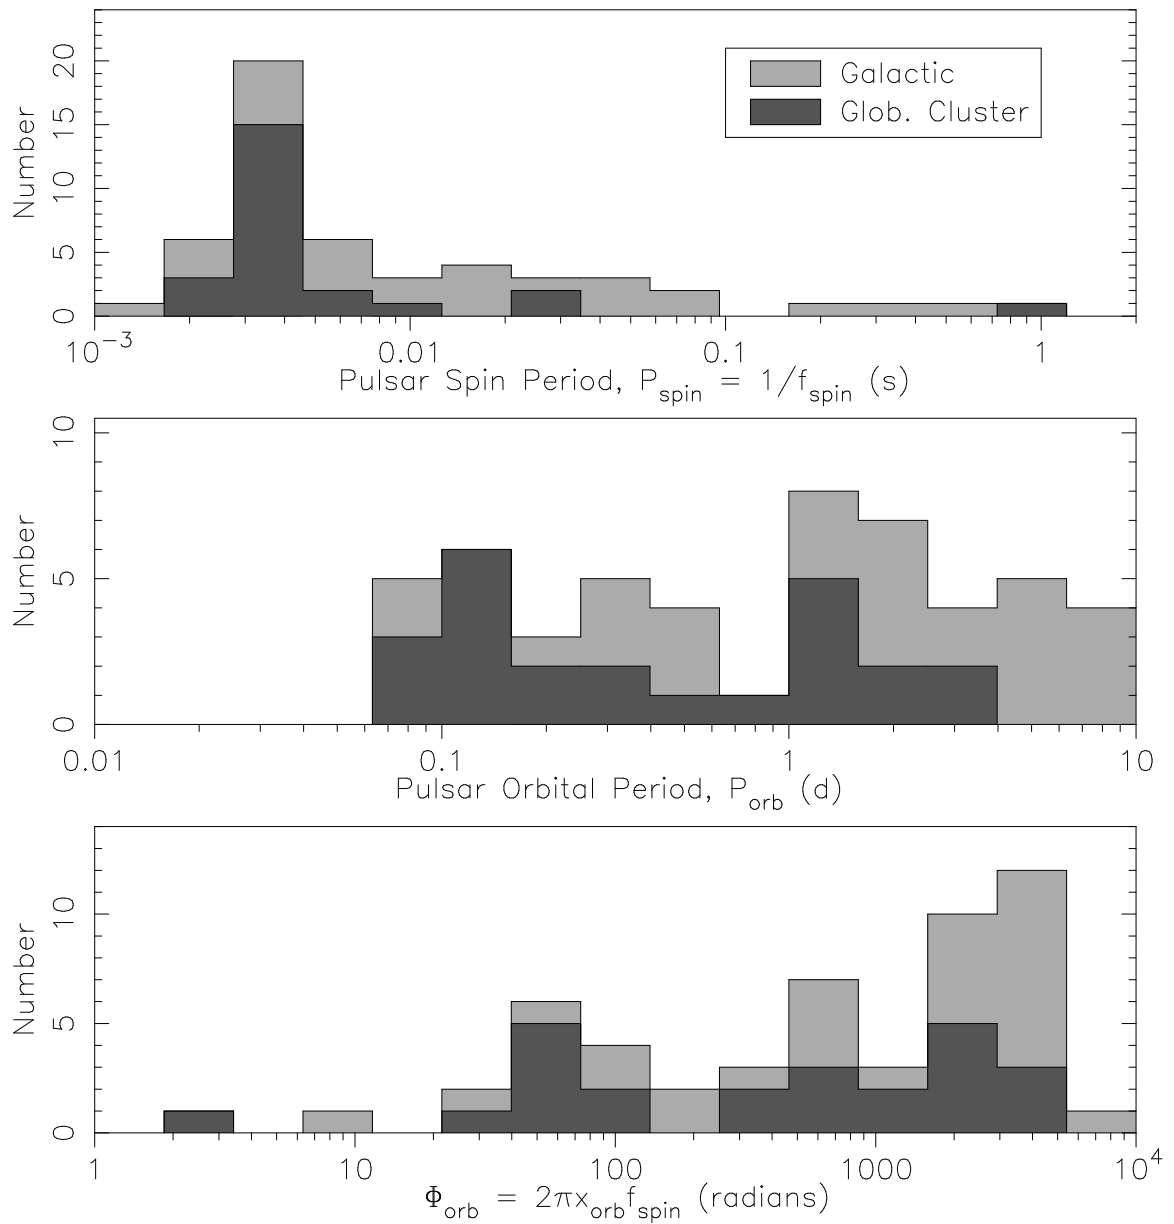


Figure 4.4 Histograms of spin period, orbital period, and phase modulation amplitude Φ_{orb} for the 52 binary pulsar systems described in Fig. 4.3.

observations of certain clusters have allowed many very weak pulsars to be detected with the help of amplification from scintillation (Camilo et al. 2000; Possenti et al. 2001). Finally, large numbers of binary pulsars are expected to exist in clusters due to dynamical interactions within the clusters (e.g. Kulkarni & Anderson 1996; Davies & Hansen 1998; Rasio et al. 2000).

Recently, Rasio et al. (2000) have conducted extensive population synthesis studies for binary pulsars in dense globular clusters like 47 Tucanae. Their initial results imply that exchange interactions among primordial binaries can produce large numbers of NSs in binary systems. In fact, their results predict that large populations of binary pulsars with low mass companions ($m_c \sim 0.05 M_\odot$) should exist in many globular clusters with orbital periods as short as ~ 15 m.

The evidence seems to show that a large population of compact binary pulsars should exist — particularly in globular clusters — which have so far eluded detection. With long observations, these pulsars should be detectable if search algorithms are able to identify the distinctive sidebands generated by the systems. The rest of this paper discusses just such an algorithm.

4.4 Sideband Searches

The fact that power from an orbitally modulated sinusoid is split into many sidebands allows us to increase the signal-to-noise of the detected response by incoherently summing the sidebands. In fact, according to one of the addition theorems of Bessel functions

$$\sum_{m=-\infty}^{\infty} J_m(z)^2 = 1, \quad (4.10)$$

theoretically *all* of the signal power can be recovered in this manner. Practically, though, a nearly complete recovery of signal power is very difficult to achieve — especially if $\Phi_{orb} \gg 1$. Such a summation would require interpolation of the $\sim 2\Phi_{orb}$ significant sidebands at the precise Fourier frequencies of their peaks $r_{spin} \pm sr_{orb}$, a computationally daunting task when Φ_{orb} , r_{spin} , and r_{orb} are unknown.

Even if we could recover all of the power using a sum of the sidebands, the significance of the measurement would be less than the significance of a non-modulated sinusoid with the same Fourier power. This loss in significance is due to the fact that noise is co-added along with signal. The exact loss in significance can easily be calculated since the probability for a sum of m noise powers to exceed some power P_{sum} is the probability for a χ^2 distribution of $2m$ degrees of freedom to exceed $2P_{sum}$ (see e.g. Groth 1975, and references therein).

In the case where only a few sidebands are suspected to be present (i.e. $\Phi_{orb} \lesssim 5$), a brute-force search using incoherent sideband summing may be computationally

feasible and worthwhile. Unfortunately, except for certain binary X-ray pulsars with long spin periods, the number of astrophysical systems with $\Phi_{orb} \lesssim 5$ is probably small compared to those with $\Phi_{orb} \gg 1$ (see §4.3).

4.4.1 Two-Staged Fourier Analysis

For systems with $\Phi_{orb} \gg 1$, the large number of regularly spaced sidebands allows a very efficient — although somewhat less sensitive — detection scheme. *Since the sidebands are separated from each other by r_{orb} Fourier bins, they appear as a short series of pulsations in the power spectrum of the original time series.* By stepping through the full-length power spectrum and taking short Fast Fourier Transforms (FFTs) of the powers, we incoherently sum any sidebands present and efficiently detect their periodic nature. As an added benefit, the measured “frequency” of these “sideband pulsations” gives a direct and accurate measurement of the *orbital* period.

We define the following notation to represent the two distinct stages of the Fourier analysis for detecting a phase modulated signal. As defined in eqn. 4.1, the initial time series n_j contains N points. After Fourier transforming the n_j (usually using an FFT of length N), the complex responses at Fourier frequencies r are represented by A_r as defined in eqn. 4.4a. The powers and phases are then simply P_r and ϕ_r respectively. The short FFTs of the P_r used to detect the sideband pulsations are of length M (where $M \ll N$) and generate complex responses A'_r , powers P'_r , and phases ϕ'_r .

In order for a pulsed signal to undergo enough modulation to produce periodic sidebands, the observation must contain more than one complete orbit (i.e. $r_{orb} > 1$). One might initially think that two complete orbits would be necessary to create a Nyquist sampled series of sidebands in the original power spectrum. However, since the sideband pulsations are not bandwidth limited, periodic signals with “wavelengths” less than two Fourier bins (i.e. $r_{orb} < 2$) will still appear in the A'_r . Instead of appearing at Fourier frequency $r = M/r_{orb}$ as for signals with $r_{orb} > 2$ (see §4.4.2), the fundamental harmonic of the response will appear aliased around the Nyquist frequency ($r_{Nyq} = M/2$) at $r_{alias} = M - M/r_{orb}$. While $r_{orb} > 1$ is the requirement to create detectable periodic sidebands, simulations have shown that this technique really requires $r_{orb} \gtrsim 1.5$ to detect most binary systems.

As the number of orbits present in the data increases, so does the spacing between the sidebands. Since each sideband has a traditional sinc function shape in A'_r (see §4.2.1), the full-width at half-maximum (FWHM) of the sideband is approximately one Fourier bin. Therefore, an increase in spacing effectively decreases the duty-cycle of the sideband pulsations (i.e. $\sim 1/r_{orb}$). Signals with small duty-cycles are easier to detect for two reasons. First, significant higher harmonics begin to appear in P'_r

which can be incoherently summed in order to improve the signal-to-noise ratio. A rule-of-thumb for the number of higher harmonics present in P'_r is approximately one-half the inverse of the duty-cycle ($\sim r_{orb}/2$). Second, as the duty-cycle decreases, the Fourier amplitudes of each harmonic increase in magnitude up to a maximum of twice that of a sinusoid with an identical pulsed fraction (see Ransom et al. 2001a, for a more detailed discussion of pulse duty-cycle effects). This effect makes each individual harmonic easier to detect as r_{orb} increases.

4.4.2 Search Considerations

As the number of orbits present in a time series increases, the signal-to-noise ratio of a detectable pulsar spin harmonic goes approximately as $\sim (r_{orb} - 1)^{-\alpha}$ where $0.3 \lesssim \alpha \lesssim 0.4$ (see Figures 4.6 and 4.7). This leads to the motto of those conducting sideband searches: “The longer the observation the better.” When making very long observations it is important to keep the on-source time (i.e. the window function) as continuous and uniform as possible. Significant gaps in the data — especially periodic or recurring gaps of the kind common in X-ray observations — create side-lobes around Fourier peaks of constant-frequency pulsations that a sideband search will detect. It is a good idea before a sideband search to take the power spectrum of the power spectrum of the known window function in order to identify spurious orbital periods that a sideband search might uncover.

Since phase modulation sidebands require coherent pulsations (i.e. f_{spin} is constant in time) in order to form, it may be necessary to correct the time series (or equivalently the Fourier transform) such that the underlying pulsations are coherent. As an example, if the telescope’s motion with respect to the solar system barycenter is not taken into account during long observations, interstellar pulsations — and therefore any phase modulation sidebands as well — will be “smeared” across numerous Fourier bins due to Doppler effects. By stretching or compressing the time series to account for the Earth’s motion before Fourier transforming (or equivalently, by applying the appropriate Fourier domain matched filter *after* the Fourier transform, see Ransom et al. 2001a) the smeared signal can be made coherent for the purposes of the sideband search.

If the time-dependence of f_{spin} is unknown *a priori*, due to accretion in an LMXB, the spin-down of a young pulsar, or timing (spin) noise, for example, one could in principle attempt a series of trial corrections and a sideband search for each trial. Such a search methodology is similar in computational complexity to traditional “acceleration” searches.

While §4.4.1 provided basic principles for detecting compact binary pulsars using the two-staged Fourier analysis, in the rest of this section we provide a more detailed

discussion of how a search might be conducted. Figure 4.5 shows two examples of sideband searches. The left column shows data from an 8 h observation of the globular cluster 47 Tucanae taken with the Parkes radio telescope at 20 cm on 2000 November 17. The signal is the fundamental harmonic of the 2.1 ms binary pulsar 47 Tuc J ($f_{spin} = 476.05$ Hz, $P_{orb} = 2.896$ h, $r_{orb} = 2.762$, $\Phi_{orb} = 120.87$ radians, and $m_c \sim 0.03 M_\odot$) with a signal-to-noise ratio of $\eta = a/\sigma_{n_j} \sim 0.0054$. The right column shows data from a simulated 8 h observation identical to the real 47 Tuc observation described above. The signal is the fundamental harmonic of the same 2 ms binary pulsar used in Figure 4.1 ($f_{spin} = 500$ Hz, $P_{orb} = 50$ m, $r_{orb} = 8$, $\Phi_{orb} = 413.44$ radians, and $m_c = 0.2 M_\odot$) with a signal-to-noise ratio one half that of the 47 Tuc J observation ($\eta \sim 0.0027$).

Once a long observation has been prepared as discussed above, we Fourier transform the data and compute the power spectrum P_r . We then begin a sideband search which is composed of two distinct parts: initial detection of a binary pulsar and the determination of a detected pulsar’s orbital elements.

Initial Detection

To maximize the signal-to-noise of a detection in P'_r , we want to match the length of a short FFT M with the width of a pulsar’s sideband response ($\sim 2r_{orb}\Phi_{orb}$) in P_r . Unfortunately, during blind searches an observed pulsar’s orbital parameters are unknown *a priori*, suggesting the need to search a range of short FFT lengths M and to overlap consecutive FFTs. We must therefore choose the number and lengths of the short FFTs, the fraction of each FFT to overlap, and the number of harmonics to sum in each FFT. These choices are influenced by the length of the observation and the nature of the systems that we are attempting to find.

Figure 4.4 shows that all currently known binary pulsars with $P_{orb} < 10$ d have modulation amplitudes of $\Phi_{orb} \lesssim 6500$ radians. This is due to the fact that the short spin periods of MSPs tend to offset the small semi-major axes (x_{orb}) of their typical binary systems. Conversely, binary systems with wider orbits tend to have longer spin periods. The vast majority of short orbital period systems lie in the range of $\Phi_{orb} \sim 20 - 5000$ radians. Multiplying this range by twice the minimum r_{orb} of 1.5 gives a rule-of-thumb range of $M \sim 60 - 15000$ where the high end should be increased based on expected values of r_{orb} . For 8 – 12 h globular cluster searches a reasonable range of values would be $64 \lesssim M \lesssim 65536$ in powers-of-two increments, for a total of 11 different values for M .

The choice of the overlap fraction is slightly more subjective. Since the sidebands can be very weak, even for a relatively bright pulsar, it pays to try and maximize the signal-to-noise ratio as much as possible. As sideband searches are quite efficient, we believe it is worth overlapping by 75% if possible. At an absolute minimum,

Figure 4.5 Two examples of key parts of a phase modulation sideband search. The left column shows the analysis of an 8 h observation of the globular cluster 47 Tucanae taken with the Parkes radio telescope of the binary millisecond pulsar 47TucJ. The right column shows analysis of simulated data containing a weaker millisecond pulsar. Plots *a* and *e* show portions of the full-length power spectrum centered on the fundamental spin harmonic of the pulsars. The grey regions are the sections of the power spectrum that were FFTd in order to create the power spectra shown in plots *b* and *f*. Plots *c* and *g* show histograms of the orbital phase as measured using the raw Fourier amplitudes in the original FFTs. Once the three Keplerian elements for circular orbits were determined, a Fourier domain orbital template was calculated and correlated with the raw Fourier amplitudes (plots *d* and *h*). This process recovered all of the power from the pulsar’s fundamental spin harmonic that had initially (plots *a* and *e*) been spread over many Fourier bins. The dotted lines in each plot show the known values of the independent variable. A detailed description of each plot is given in §4.4.2.

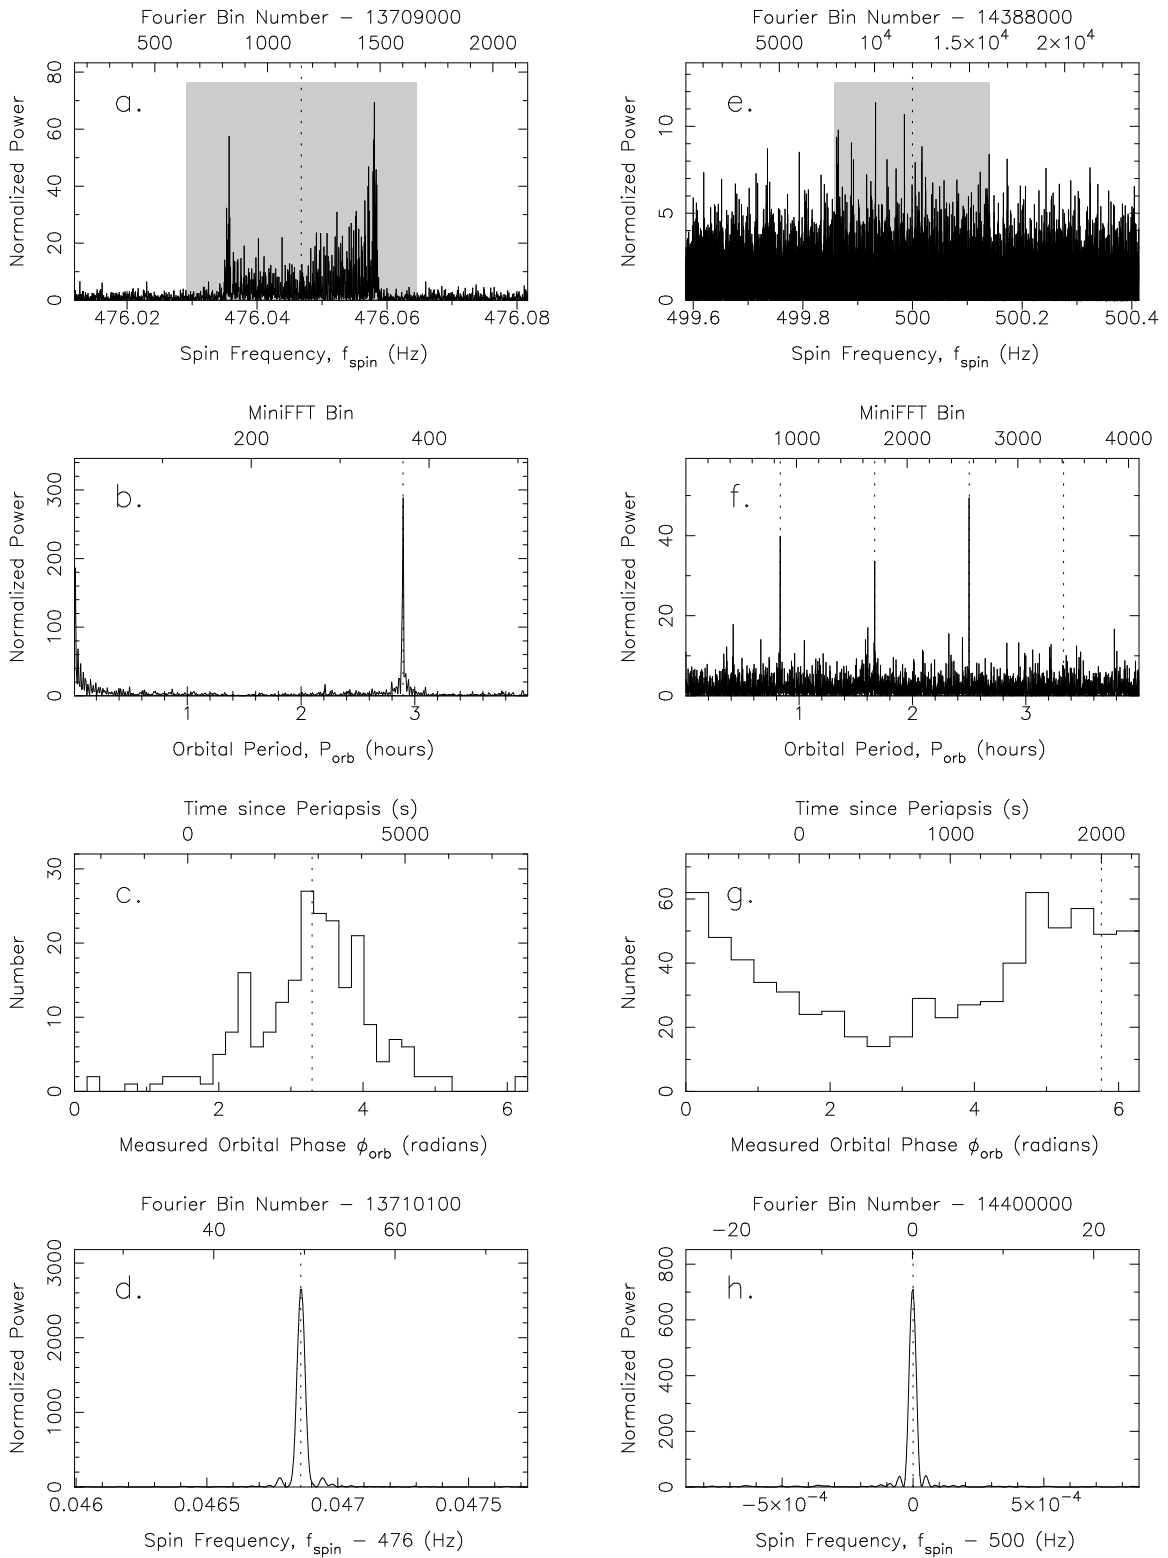


Figure 4.5 (Continued).

one should overlap the short FFTs by 50% in order to avoid excessive loss of signal-to-noise when a series of sidebands straddles consecutive FFTs.

In Figure 4.5 plots *a* and *e*, the grey shaded regions show the sections of the P_r that were Fourier transformed for the sideband search. The short FFTs of length $M = 1024$ and $M = 8192$ were centered on the known pulsar spin frequency and resulted in the P'_r displayed in plots *b* and *f* respectively. The short FFT lengths were chosen as the shortest powers-of-two greater than $2r_{orb}\Phi_{orb}$. For 47 Tuc J (plot *a*), the “horned” structure of the phase modulation sidebands described in §4.2.1 is easily recognizable. This shape is not apparent in plot *e*, and in fact, traditional searches would detect nothing unusual about this section of the power spectrum.

Normalization of the P'_r results in the usual exponential distribution with mean and standard deviation of one for a transform of pure noise. If one first normalizes the P_r as shown in plots *a* and *e* (see Ransom et al. 2001a, for several normalization techniques), the resulting P'_r can be normalized simply by dividing by M .

Plots *b* and *f* of Figure 4.5 show the highly significant detections of the binary pulsars in P'_r . For 47 Tuc J, only one harmonic is visible since $r_{orb} < 2$, but the single-trial significance is approximately $\sim 23\sigma$. For the simulated 50 m binary pulsar, 3 significant harmonics are detected (predicted harmonic locations are marked with a dotted line in plot *f*). When summed, the detection has a single-trial significance of $\sim 15\sigma$.

Orbit Determination

Once a binary pulsar has been detected, examination of the P_r and ϕ_r allow the estimation of the Keplerian orbital parameters and the pulsar spin period P_{spin} . Brute-force matched filtering searches centered on the estimates are then used to refine the parameters and recover the fully coherent response of the pulsar.

The simplest parameter to determine is P_{orb} . In a Discrete Fourier Transform we denote the Fourier frequency r as the number of integer oscillations in the transformed time series. In phase modulation the spacing between sidebands is r_{orb} bins, implying that the most significant peak in P'_r should occur at Fourier frequency M/r_{orb} . Conversely, if we measure the Fourier frequency r_{meas} in P'_r with the most power, we can solve for r_{orb} using $r_{orb} = M/r_{meas}$. By using Fourier interpolation or zero-padding to oversample A'_r and P'_r , we can determine r_{meas} to an accuracy of

$$\sigma_r = \frac{3}{\pi\alpha\sqrt{6P_{meas}}}, \quad (4.11)$$

where P_{meas} is the normalized power at the peak and α is the signal “purity” — a property that is proportional to the root-mean-squared (RMS) dispersion of the pulsations in time about the centroid — which is equal to one for pulsations

present throughout the data (see Middleditch et al. 1993; Ransom et al. 2001a). Remembering that $P_{orb} = T_{obs}/r_{orb}$ and solving gives

$$P_{orb} = \frac{T_{obs}r_{meas}}{M} \pm \frac{3T_{obs}}{\pi\alpha M\sqrt{6P_{meas}}}. \quad (4.12)$$

For aliased signals, the orbital period is

$$P_{orb,aliased} = \frac{T_{obs}(M - r_{meas})}{M} \pm \frac{3T_{obs}}{\pi\alpha M\sqrt{6P_{meas}}}. \quad (4.13)$$

For pulsars with $\Phi_{orb} \gg 1$ observed such that $r_{orb} > 2$, it is often possible to measure P_{orb} to a part in 10^4 , corresponding a ~ 1 s error for many of the pulsars shown in Figure 4.4.

Once P_{orb} is known it becomes possible to measure the projected semi-major axis x_{orb} by determining the total width of the sidebands in P_r . For bright pulsars like 47 Tuc J, the number of Fourier bins between the ‘‘horns’’ ($\Delta_r \sim 2\Phi_{orb}r_{orb}$) as shown in Figure 4.5a can be measured directly and converted into x_{orb} using

$$x_{orb} \simeq \frac{\Delta_r}{4\pi f_{spin}r_{orb}}, \quad (4.14)$$

where f_{spin} can be estimated as the frequency midway between the phase modulation ‘‘horns’’.

For weaker pulsars, where the sideband edges are not obvious (as in Figure 4.5e), Δ_r can be estimated using numerous short FFTs. By taking a series of short FFTs of various lengths around the detection region of P_r and measuring the power, ‘‘centroid’’ and ‘‘purity’’³ at the frequency corresponding to P_{orb} , one can map the extent of the hidden sidebands and compute an estimate for x_{orb} .

For pulsars in circular orbits, which constitute the majority of the systems in Figures 4.3 and 4.4, the only remaining orbital parameter is the time of periapsis passage T_o . Defining the time since periapsis passage as $\Delta_{T_o} = t_o - T_o$ and using eqn. 4.3 we get

$$\Delta_{T_o} = \frac{P_{orb}}{2\pi} \left(\phi_{orb} - \frac{\pi}{2} \right), \quad (4.15)$$

where ϕ_{orb} can be measured from the phases of the sidebands in A_r using the following technique.

From eqn. 4.9b we know that the sidebands have phases of $\phi_s = s \left(\phi_{orb} + \frac{\pi}{2} \right) + \phi_{spin}$ radians. Since some of the measured phases will be ‘‘flipped’’ by π radians

³The measured values of ‘‘centroid’’ and ‘‘purity’’ are estimates of a signal’s location and duration in a time series as determined from the derivatives of the Fourier phase and power at the peak of the signal’s response (see Middleditch et al. 1993; Ransom et al. 2001a).

due to the fact that the Bessel function amplitudes will be negative, we can use our knowledge of r_{orb} and f_{spin} to calculate the $J_s(\Phi_{orb})$ directly and determine which phases need to be “flipped”. After correcting these phases, we can estimate ϕ_{orb} simply by subtracting the phases of neighboring sidebands using $\phi_{orb} = \phi_{s+1} - \phi_s - \pi/2$.

For weak pulsars each measurement of ϕ_{orb} has a large uncertainty. Fortunately, we can make $\sim 2\Phi_{orb}$ measurements of ϕ_{orb} by using each pair of sidebands and then determine ϕ_{orb} and its uncertainty statistically. Figure 4.5 plots *c* and *g* show histograms of the measurement of ϕ_{orb} using this technique.

It is important to mention that for weak pulsars it may be difficult to determine the location of the sidebands in order to measure their phases. The location of the sideband peaks can be calculated by measuring the phase ϕ_{meas} of the fundamental harmonic discovered in A'_r . The number of Fourier bins from the first bin used for the short FFT to the peak of the first sideband in A_r is simply $\phi_{meas}r_{orb}/2\pi$. Subsequent peaks are located at intervals of r_{orb} Fourier bins.

Once estimates have been made for the Keplerian orbital parameters, a set of complex sideband templates are computed over the most likely ranges of the parameter values given uncertainties in each. These templates are then correlated with a small region of A_r around where the initial detection took place. When a template is used that matches the pulsar sidebands buried in A_r , the total power for that pulsar spin harmonic is recovered (Ransom et al. 2001a). Figure 4.5 plots *d* and *h* show the results of just such a matched filtering operation. The Doppler effects from the orbital motion have been completely removed from the data and the resulting Fourier response is that of an isolated pulsar.

For binary pulsars in eccentric orbits, the techniques discussed here could in principle be applied to each set of sidebands from the orbital Fourier expansion (see §4.2.3). Such an analysis would be much more complicated than that described here and would require a high signal-to-noise detection.

4.5 Discussion

We have described a new search technique for binary pulsars that can identify sidebands created by orbital modulation of a pulsar signal when $T_{obs} > P_{orb}$. These “sideband searches” allow the detection of short orbital period binary pulsars that would be completely undetectable using conventional search techniques.

Figures 4.6 and 4.7 show how the sensitivity of sideband searches compares to that of acceleration searches for a wide range of observation times T_{obs} . Acceleration search sensitivities are near-optimal when $T_{obs} \lesssim 0.1P_{orb}$ (Johnston & Kulkarni 1991), but degrade rapidly at longer integration times when the constant frequency

derivative approximation breaks down. Sideband searches, on the other hand, require $T_{obs} \gtrsim 1.5P_{orb}$ in order to approach the sensitivity of optimal duration acceleration searches, but for longer T_{obs} sensitivity improves as $\sim (r_{orb} - 1)^{-\alpha}$ where $0.3 \lesssim \alpha \lesssim 0.4$. *For targeted searches of duration $\sim 8-12$ h (e.g. globular cluster observations), sideband searches are $\sim 2-10$ times more sensitive to compact binary pulsars ($P_{orb} \lesssim 4$ h) than optimal duration acceleration searches.* In general, one should think of acceleration and sideband searches as complementary to each other — sideband searches allow the detection of ultra-compact binary pulsars, while acceleration searches maximize the detectability of isolated and longer-period binary pulsars.

The fact that sideband searches target a different portion of orbital parameter space than acceleration searches and yet require significantly less computation time, provides good reason to include them in future pulsar searches where $T_{obs} \gtrsim 30$ m. In fact, the Parkes Multibeam Pulsar Survey ($T_{obs} \sim 35$ m, Lyne et al. 2000a; Manchester et al. 2001) has included a sideband search in an on-going re-analysis of their survey data at the cost of only a marginal increase in computer time. While the probability of detecting a binary pulsar with $P_{orb} \lesssim 23$ m is almost certainly quite low, the discovery of a single such system would provide a wealth of scientific opportunity.

Acknowledgments

We would like to thank G. Fazio for supporting our research and encouraging us in this work. Many of the computations for this paper were performed on equipment purchased with NSF grant PHY 9507695. S. Eikenberry is supported by a NSF CAREER Grant.

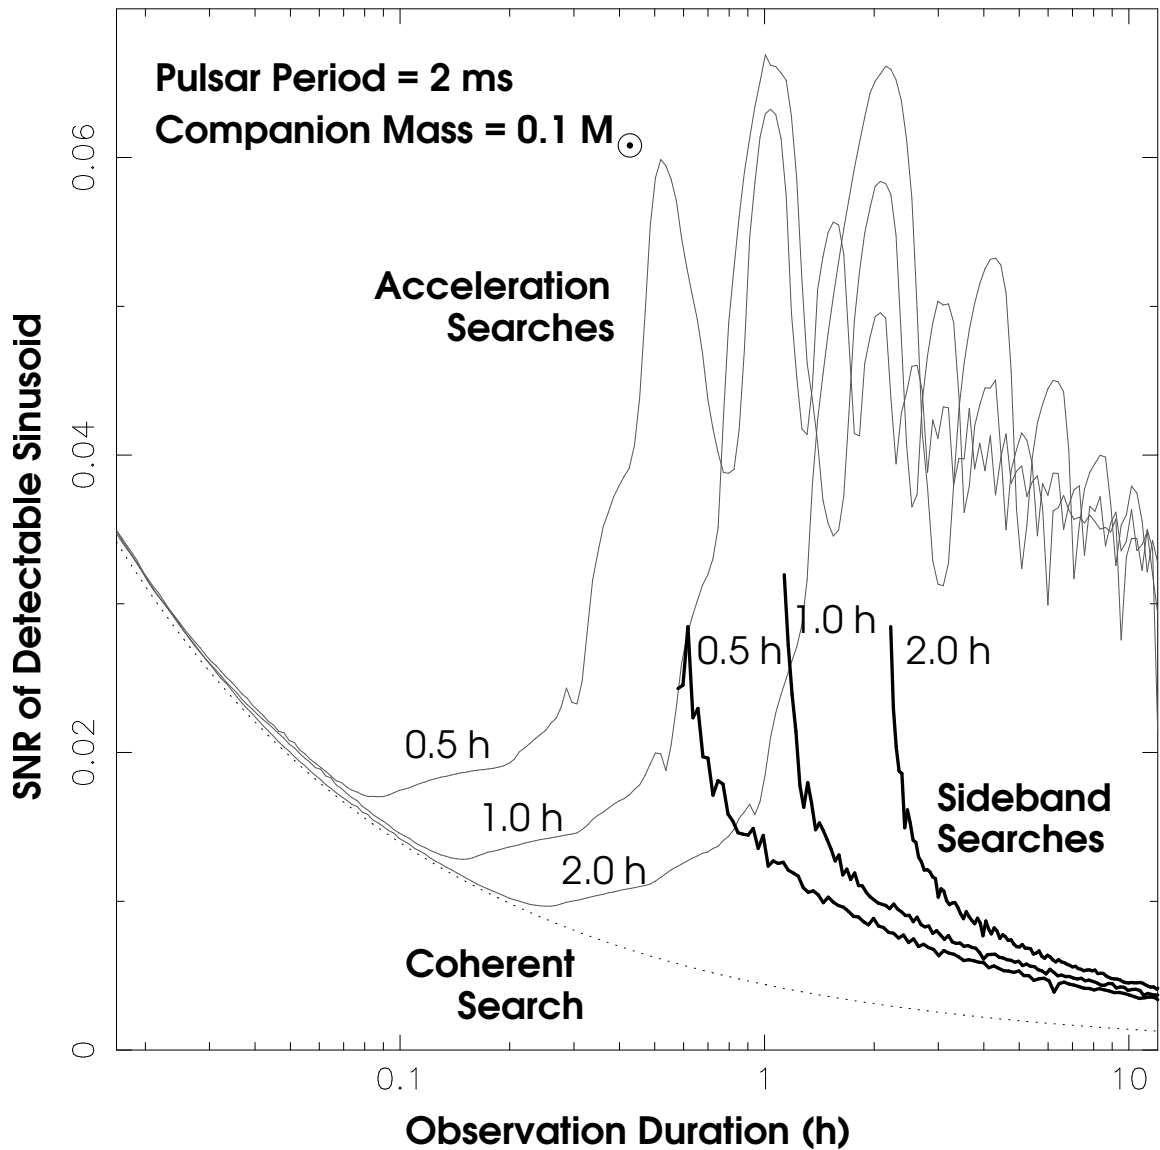


Figure 4.6 The results of simulations showing the detectable signal-to-noise ratios (SNR) for the fundamental harmonic (i.e. a sinusoid) of a 2 ms binary pulsar with a $0.1 M_{\odot}$ companion and orbital periods of 0.5, 1, and 2 h. The detection threshold was 8σ . Optimal detection of the sinusoid is shown as the dotted line labeled “Coherent Search” which is proportional to $T_{obs}^{-1/2}$. The minima in the “Acceleration Search” curves denote optimal observation durations for each orbital period P_{orb} ($T_{best} \sim 0.1P_{orb}$). Sideband searches are more sensitive to binary pulsars than optimal duration acceleration searches when $T_{obs} \gtrsim 2P_{orb}$.

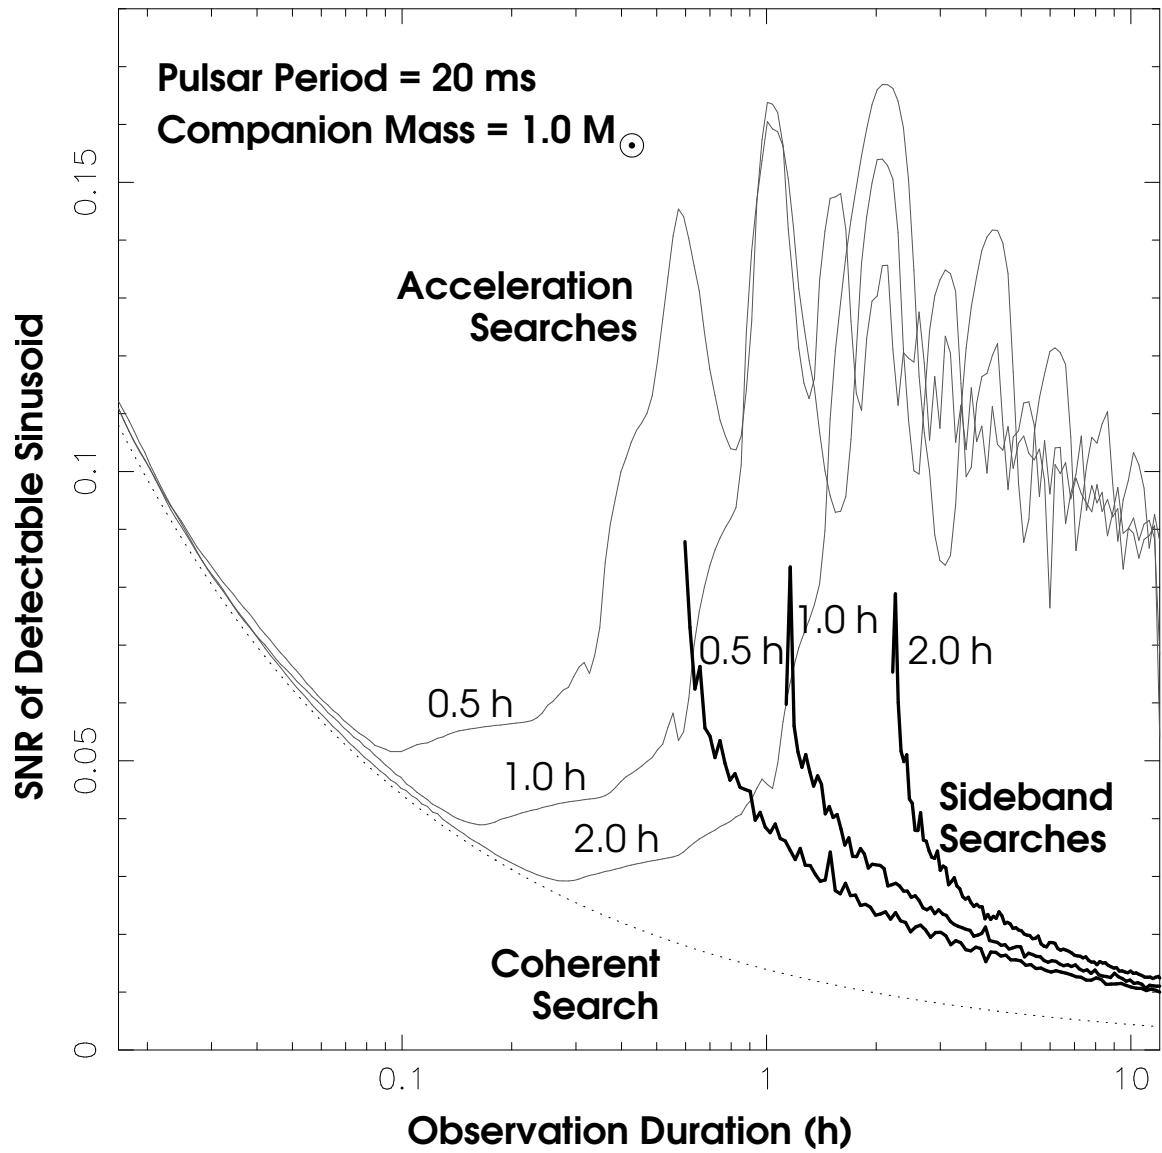


Figure 4.7 Results from simulations similar to those shown in Figure 4.6 but for a 20 ms binary pulsar with a $1.0 M_{\odot}$ companion.

Chapter 5

Results and Future Prospects

Unfortunate as it may sound, the vast majority of time spent working towards this thesis was spent in the design, coding, debugging, and testing of a software package called PRESTO that implements the various techniques discussed in Chapters 2 and 4. What has resulted, though, is a high-quality suite of pulsar search software that has been used to discover new pulsars. With a suite of working and debugged code, searching new or archival data sets becomes an almost trivial undertaking.

5.1 PRESTO

PRESTO is a large suite of pulsar search software developed from scratch in ANSI C. Written with portability, ease-of-use, and memory efficiency in mind, it can currently handle data from the Berkeley-Caltech Pulsar Machine (BCPM; located at the GBT), the Wideband Arecibo Pulsar Processor (WAPP), the Parkes Multibeam system, or any individual time series composed of single precision floating point data. This last format is useful for analyzing X-ray, optical, or dedispersed radio data from any source.

The software is composed of numerous routines designed for the three primary stages of pulsar search analysis:

1. Data Preparation
 - (a) `rfifind` searches the raw data in both the time and frequency domains for interference or other problems. Each channel is analyzed for short time intervals throughout the observation. The result is a list of recommended portions of the data to mask.
 - (b) `prepsubband` clips, masks, dedisperses, barycenters, and pads raw data into numerous time series over a range of trial DMs.

- (c) `realfft` performs either an in-core or out-of-core FFT based on the length of the time series to be transformed. The input data is not required to be a power-of-two in length.
- (d) `zapbirds` removes known sources of interference (or strong pulsars) from a Fourier transform by replacing them with the local average Fourier amplitudes.

2. Searching

- (a) `accelsearch` performs a Fourier domain acceleration search with Fourier interpolation and harmonic summing on an FFT. The code is written such that it uses a minimum of memory (~ 200 MB for a search over 200 independent accelerations no matter what the length of the FFT) and utilizes the processor cache as much as possible.
- (b) `search_bin` performs a phase modulation sideband search on an FFT including harmonic summing and Fourier interpolation.

3. Candidate Optimization

- (a) `prepfold` folds known pulsars or candidates from `accelsearch` over a range of DMs, periods, and period derivatives around the “best-guess” values and returns the optimized pulse profile. All plots of new and candidate pulsars in this chapter were generated using `prepfold`.
- (b) `bincand` attempts to find an orbital solution for a sideband search candidate uncovered with `search_bin`. Brute force matched filtering is used to search the orbital parameters around the “best-guess” values as discussed in chapter 4.

Many other utilities for displaying, plotting, and working with data and have been written in C or as Python¹ scripts. Since the software is coded in a modular form and is relatively mature, adding new capabilities, such as the ability to handle a new pulsar back-end, can be accomplished in only a few hours of programming. Similarly, with a large library of routines that implement the “building-block” functionality of Fourier-domain searches as discussed in chapter 2 (such as complex correlations and template generation), new searches or analysis techniques can be coded quickly and efficiently.

¹A free and aesthetically pleasing scripting language available from <http://www.python.org>

TABLE 5.1
Parkes Globular Cluster Observations

Cluster	ν_{center} (MHz)	BW (MHz)	N_{chan}	T_{obs} (h)	T_{samp} (μs)	N (10^6)	DMs (pc/cm ³)	Yr	Notes
47 Tuc	660	32	256	9.6	125	276	24.5	98	RFI
47 Tuc	1390	256	512	10.1	80	455	24.5	00	PM
47 Tuc	1390	256	512	8.0	80	240	24–25	00	
47 Tuc	1390	256	512	2.8	80	128	24–25	00	
Liller 1	1390	256	512	8.0	500	59	520–820	00	
NGC 2808	660	32	256	9.2	125	264	0–600	98	RFI
NGC 6544	1332	64	256	0.5	125	16.7	0–600	98	PSR
Terzan 5	1332	64	256	8.3	125	240	235–243	98	
Terzan 5	1332	64	256	3.0	125	86.4	230–246	98	
Terzan 5	1390	256	512	2.0	250	28.8	235–245	99	
Terzan 5	1390	256	512	8.3	125	240	232–244	00	PSR
Terzan 5	1390	256	512	3.3	125	100	N/A	01	RFI

Table 5.1 Parameters of globular cluster search data taken at the Parkes radio telescope and searched for this work. Under the notes column “RFI” indicates that extensive radio-frequency interference degraded the quality of the data significantly, “PSR” means that a new pulsar was discovered, and “PM” indicates that only a phase modulation search has been conducted so far. The new pulsar in NGC 6544 is discussed in Chapter 3.

5.2 Globular Cluster Searches

Using PRESTO, we have searched a number of long observations of globular clusters taken with the Parkes radio telescope. Table 5.1 shows the basic parameters of the data we have searched.

5.2.1 Terzan 5

In order to investigate very intriguing radio point sources discovered in VLA observations of Terzan 5 by Fruchter & Goss (1990); Fruchter & Goss (2000), we observed the cluster twice in 1998 for ~ 8 and ~ 3 hours and then again (with Nichi D’Amico and collaborators) for 8 hours in 2000. The data from 2000 were taken using an improved back-end for the Multibeam system with twice the number of channels and four times the bandwidth of the system used in 1998. This data was also analyzed with more sophisticated versions of the Fourier-domain acceleration and sideband searches.

The sideband search easily detected Terzan 5 A, the 11 ms eclipsing binary pulsar ($P_{orb} \sim 1.8$ h), in each of the long observations. These detections occurred in spite of the fact that the pulsar was visible for less than about 1/3 of each observation due to eclipses and “nulling” (Nice & Thorsett 1992). For new compact binary MSPs ($P_{orb} \lesssim 4$ h), our limiting search sensitivity was $\gtrsim 0.25$ mJy at 20 cm. The relatively bright (~ 1.4 mJy at 20 cm) and steep-spectrum point source “N” as observed by Fruchter & Goss (2000) should have easily been detected if it really is a compact binary pulsar (assuming that the pulsar signal is “on” throughout the observation). Acceleration searches easily detected both Terzan 5 A and Terzan 5 C (Lyne et al. 2000b) in each observation.

Full-length acceleration searches of the 8 h observation from 2000 uncovered a new 4.71 ms pulsar (see Figure 5.1) and a candidate pulsar at 2.2 ms (see Figure 5.3). A re-analysis of the 1998 data using the improved search algorithms immediately revealed the 4.71 ms pulsar (Terzan 5 D; Figure 5.2), but could not confirm the weaker 2.2 ms candidate. A follow-up observation was taken in 2001 in order to try and confirm the 2.2 ms candidate, but strong interference prevented a sensitive search of the data.

Terzan 5 D is most likely an isolated pulsar, although during the 1998 observation the best-fit period derivative is non-zero at a marginally significant level. While it is possible that this measurement indicates the pulsar is a member of a long-period binary, the most likely explanation is that it is due to the lower significance detection. Future measurements using more sensitive systems (i.e. the GBT) will provide a definitive answer and quite possibly a timing solution. A timing solution would also allow us to compare the location of the pulsar with the point sources observed by

Fruchter & Goss (2000).

5.2.2 Liller 1

During the Parkes run in 2000, we also made an 8 h observation of the rich globular cluster Liller 1. Liller 1 has been predicted to contain a wealth of pulsars (e.g. Kulkarni & Anderson 1996), yet none have been detected. The most likely reasons are the high predicted DM ($500\text{--}1000\text{ pc cm}^{-3}$) and extensive scattering due to its location near the Galactic plane and close to the Galactic center ($\ell = 354^\circ.8$, $b = -0^\circ.2$).

This cluster had been observed with the VLA by Fruchter & Goss (2000) who identified a radio source near its core. More recent observations have identified this radio source with the bursting LMXB known as the ‘‘Rapid Burster’’ (Fruchter & Goss 2000; Moore et al. 2000).

Our searches found no convincing pulsations despite the use of ~ 5000 CPU-hours of computer time. If millisecond pulsations are present in the cluster, the expected scatter broadening at 20 cm of $\sim 5\text{--}20$ ms would almost certainly have made them undetectable. Future searches at higher frequencies (i.e. 5 GHz) should allow more sensitivity to millisecond periods.

5.2.3 47 Tucanae

During the observing run at Parkes in 1998, we took a 9.3 h observation of 47 Tuc at 660 MHz. Unfortunately, the interference in this band was terrible and we were unable to make a very sensitive search of the data. Similar problems plagued the data taken of NGC 2808.

By 2000, Fernando Camilo and collaborators had been making many long observations of 47 Tuc using the improved 20 cm Multibeam system. As they were not analyzing the data in sections longer than about 17 m (see Camilo et al. 2000), they offered to provide some long integrations in order to test PRESTO.

Sideband searches of the long observations easily detected the pulsar 47 Tuc J with its orbital period of ~ 3 h. 47 Tuc I, O, R, and W should also be detectable in sideband searches of 47 Tuc as long as favorable scintillation conditions are present. The search limit for new ultra-compact binary pulsars was ~ 0.25 mJy assuming a constant flux density throughout the observation.

Fourier-domain acceleration searches of this data easily detected many of the known pulsars in 47 Tuc and uncovered numerous good candidates that remain unconfirmed. We currently have 6 more long ($\gtrsim 8$ h) observations of 47 Tuc awaiting analysis.

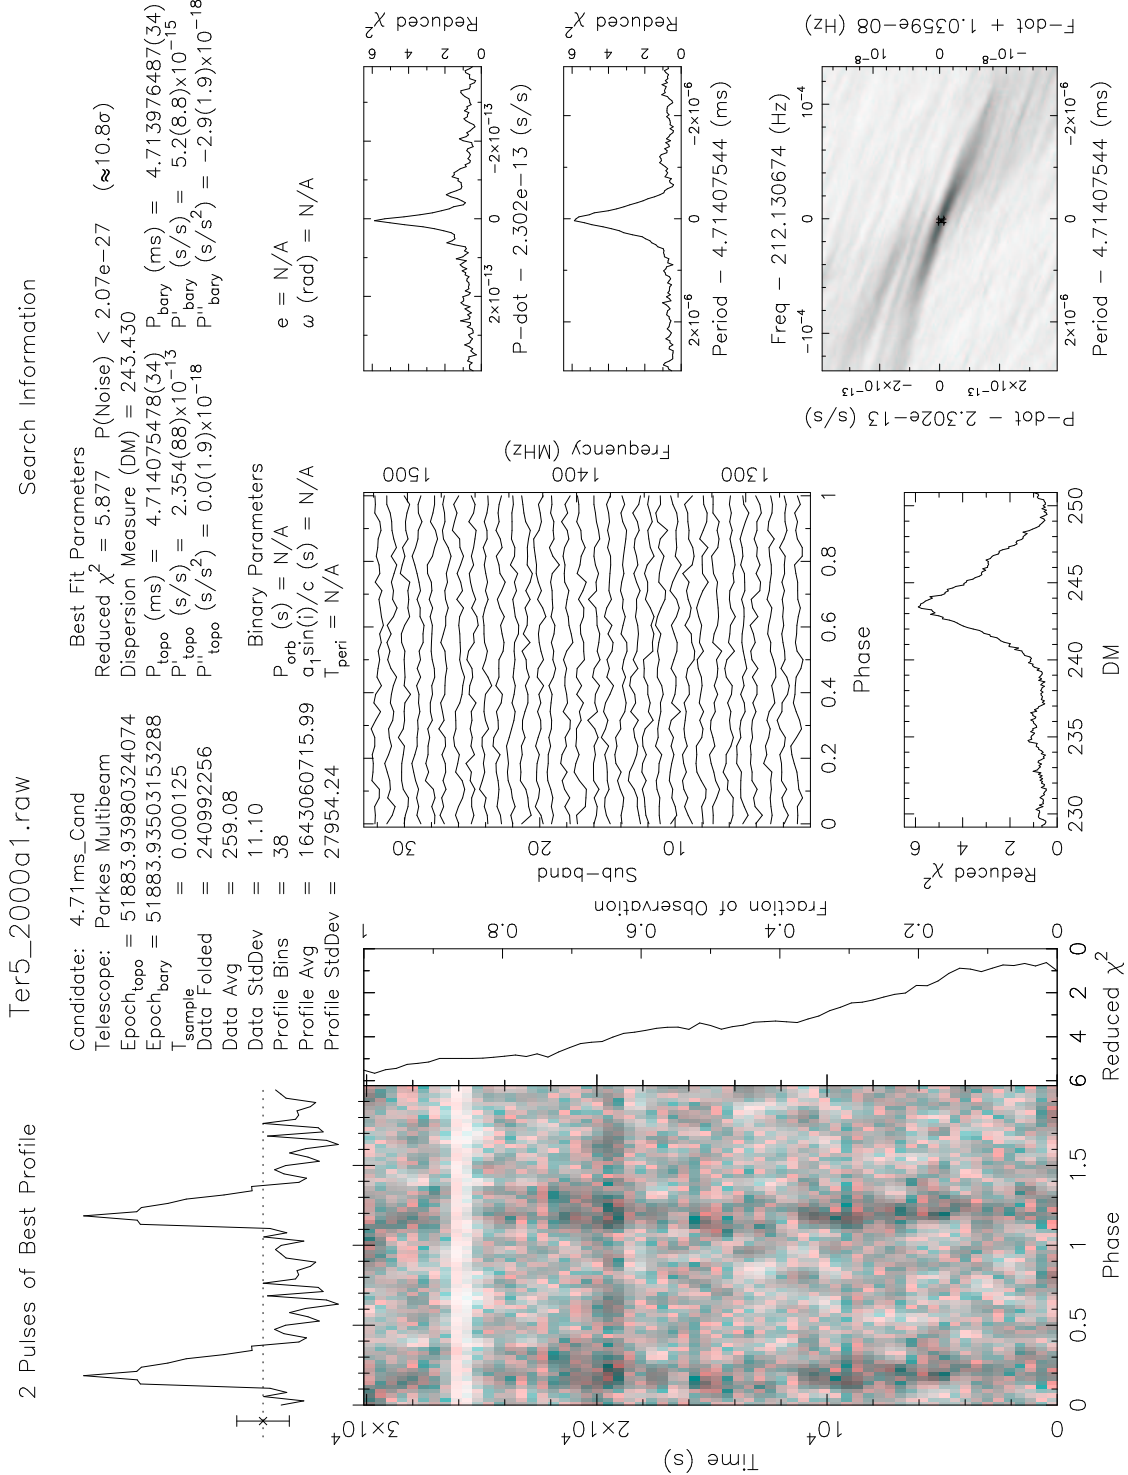


Figure 5.1 The discovery observation of pulsar Terzan 5 D taken 2000 December 5. The pulsar was discovered using a Fourier-domain acceleration search with harmonic summing as described in Chapter 2, although accelerations were not needed since the measured \dot{f} from the single observation is consistent with zero.

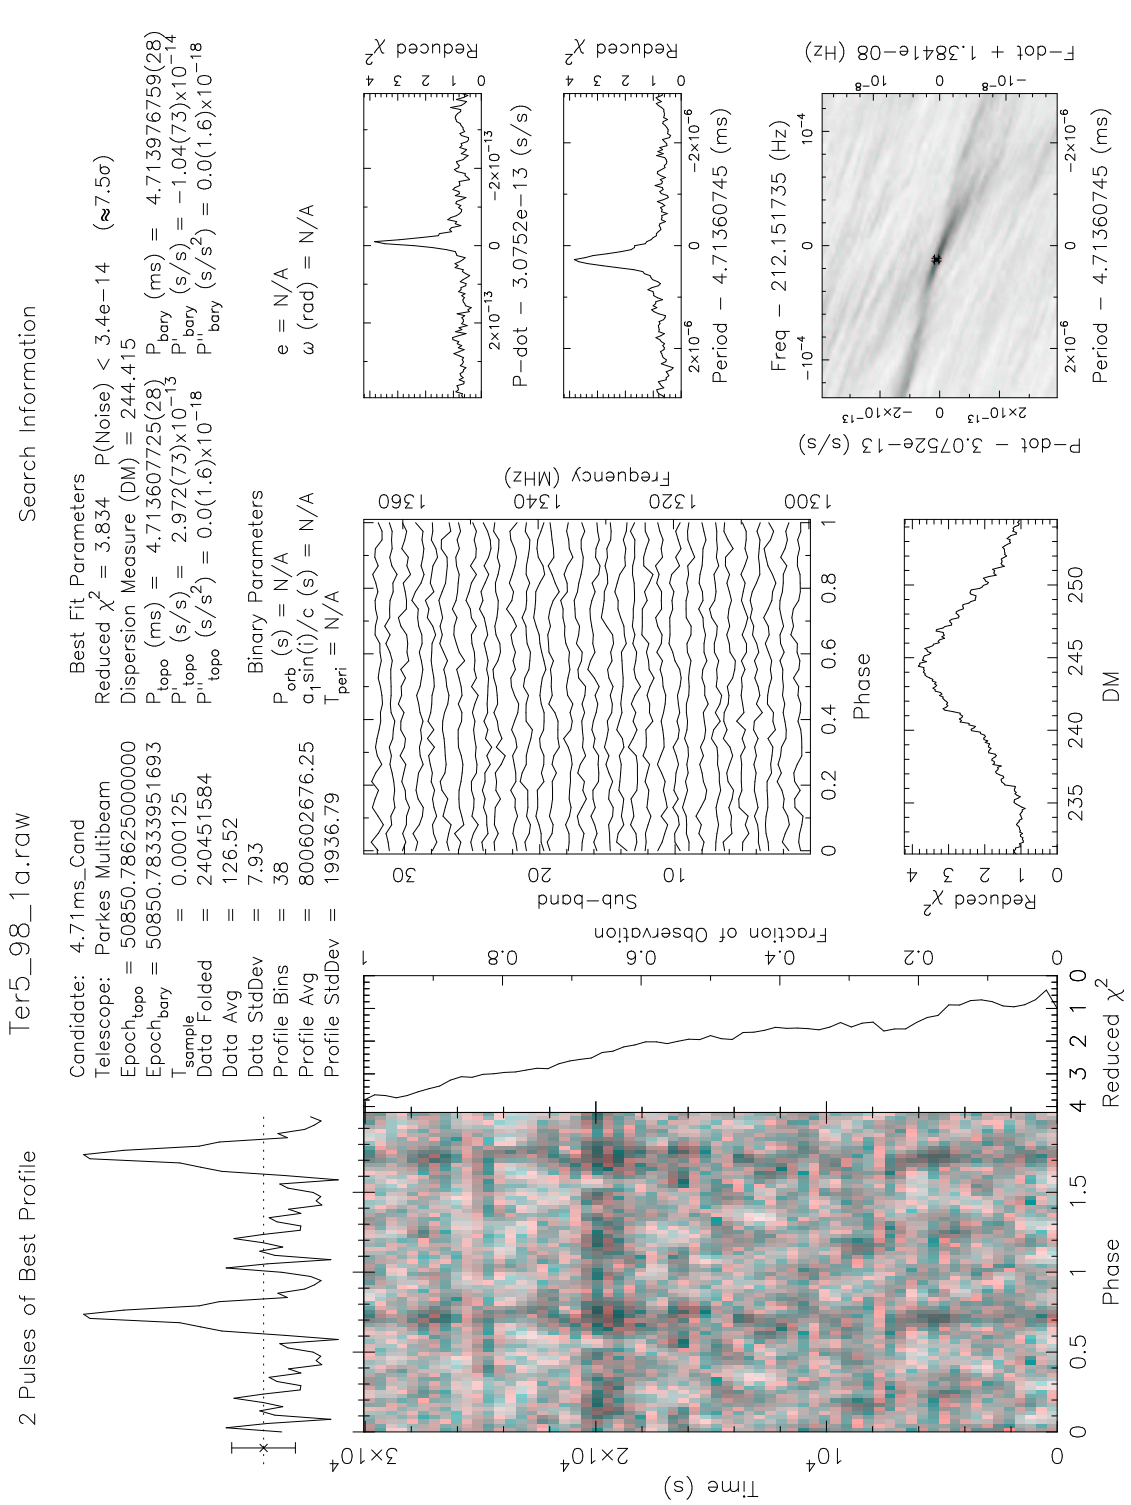


Figure 5.2 Pulsar Terzan 5 D as confirmed in an 8 h Parkes observation taken 1998 February 6 using a less sensitive system with only 64 MHz of bandwidth.

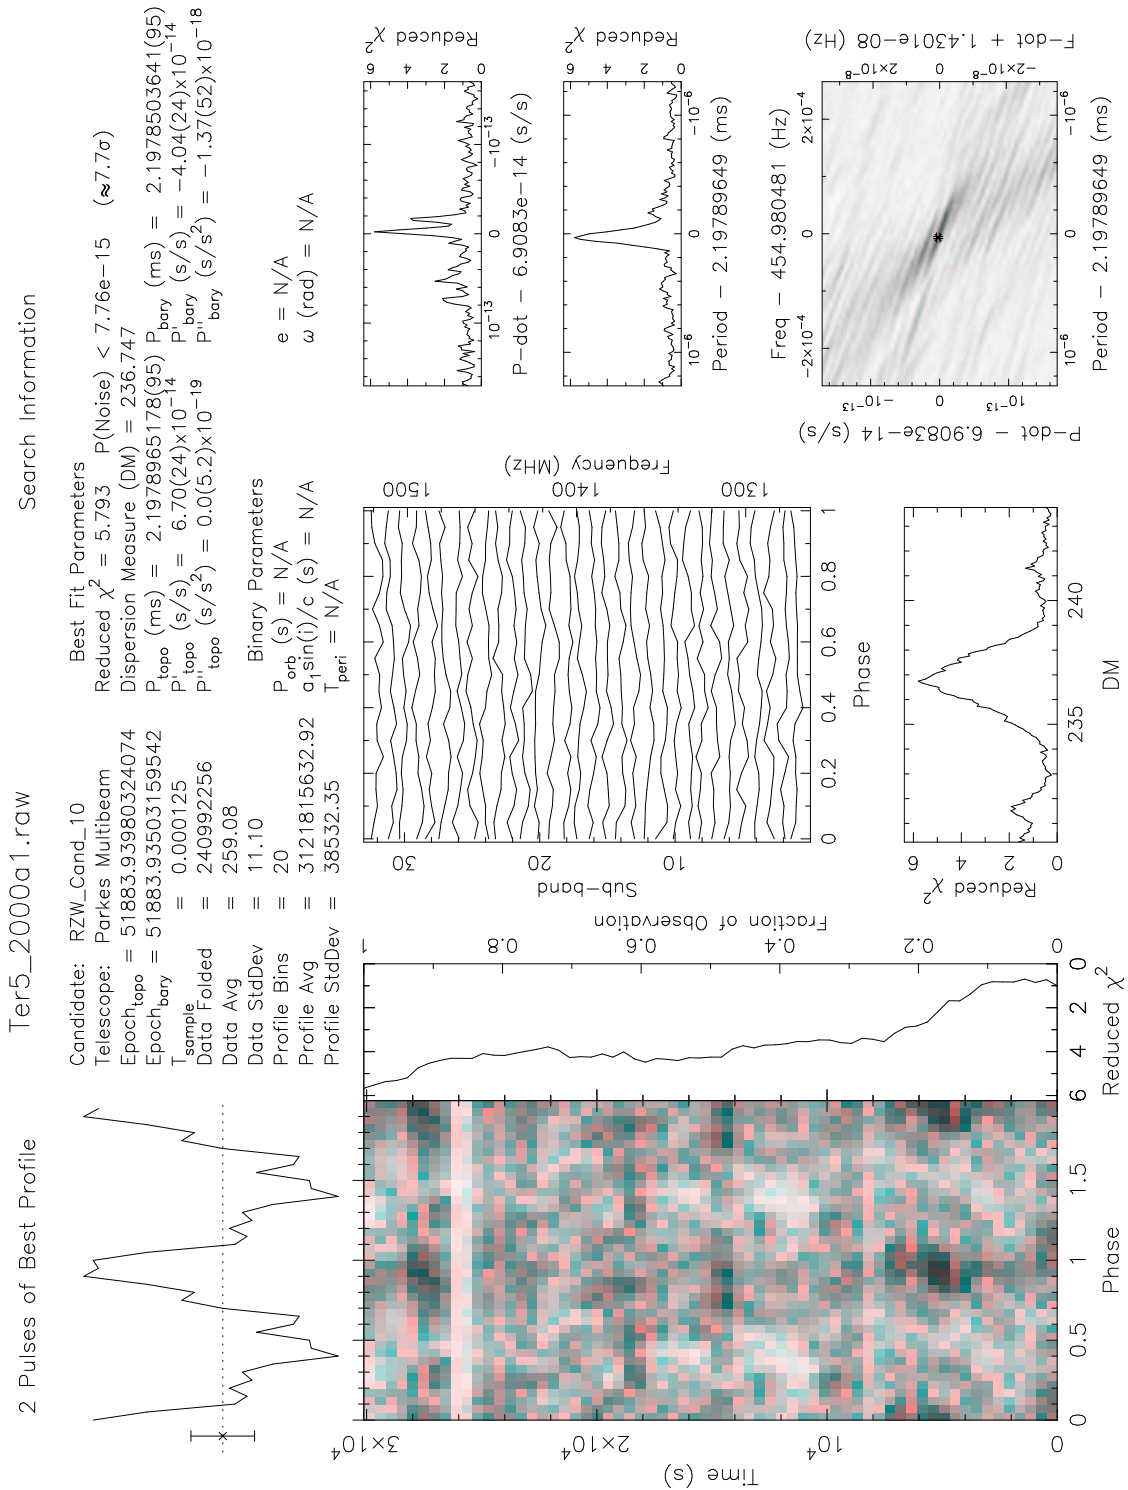


Figure 5.3 A good pulsar candidate (unconfirmed) from the same Parkes observation of Terzan 5 described in Figure 5.1. This candidate was detected with a highly significant f implying that, if real, it is the member of a relatively long period binary system (i.e. $P_{orb} \gtrsim 1$ d).

5.3 Other Searches

With working code, it is very easy to conduct searches of other long time series. Many of these “other” searches were made of archival X-ray data from HEASARC² or new observations (such as 3C 58 and the CasA SNR) taken with the *Chandra X-ray Observatory*. A few of these searches are described below.

5.3.1 3C 58

Steve Murray and collaborators (Murray et al. 2001, in preparation) detected 65.68 ms pulsations from the X-ray point-source in supernova remnant (SNR) 3C 58 at $\sim 6.7\sigma$ significance using the High Resolution Camera (HRC) on the *Chandra X-ray Observatory*. Working with Murray, Bryan Gaensler and I immediately searched a 20 ks observation of 3C 58 from the *Rossi X-ray Timing Explorer (RXTE)* archives in an attempt to confirm the new pulsar. The data had been previously searched for pulsations by other groups using various search codes (F. Seward, B. Gaensler, private communication).

We conducted a blind search of ~ 2 – 10 keV data from the observation using the Fourier-domain acceleration search with harmonic summing from PRESTO. The only candidate was a $\sim 6.4\sigma$ detection of a 65.66 ms signal using a sum of eight harmonics — the same signal found with the *Chandra* data. Folding and optimizing the candidate with various energy cuts produced the $\sim 11\sigma$ profile shown in Figure 5.4.

Using the detections from both the *Chandra* and *RXTE* observations, we were able to determine a spin-down rate for the new pulsar (PSR J0205+6449) of $\dot{P} = 1.94 \times 10^{-13}$. Assuming a magnetic dipole model for PSR J0205+6449, we infer a surface magnetic field of 3.6×10^{12} G. The characteristic age for this pulsar is 5400 yr, indicating either that 3C 58 was not the supernova of 1181 CE, or that the pulsar’s initial spin period was ~ 60 ms.

5.3.2 Supernova Remnants with *RXTE*

Encouraged by the success with 3C 58, Gaensler et al. (in preparation) undertook searches of 9 other plerionic/composite SNRs using archival data from *RXTE* in an attempt to discover new X-ray pulsars. While we have discovered no new pulsars, the firm upper limits on pulsed fraction (see Table 5.2) may enable us to say something about X-ray pulsar luminosities and/or beaming.

²Available at <http://heasarc.gsfc.nasa.gov/w3browse/w3blaunch.html>

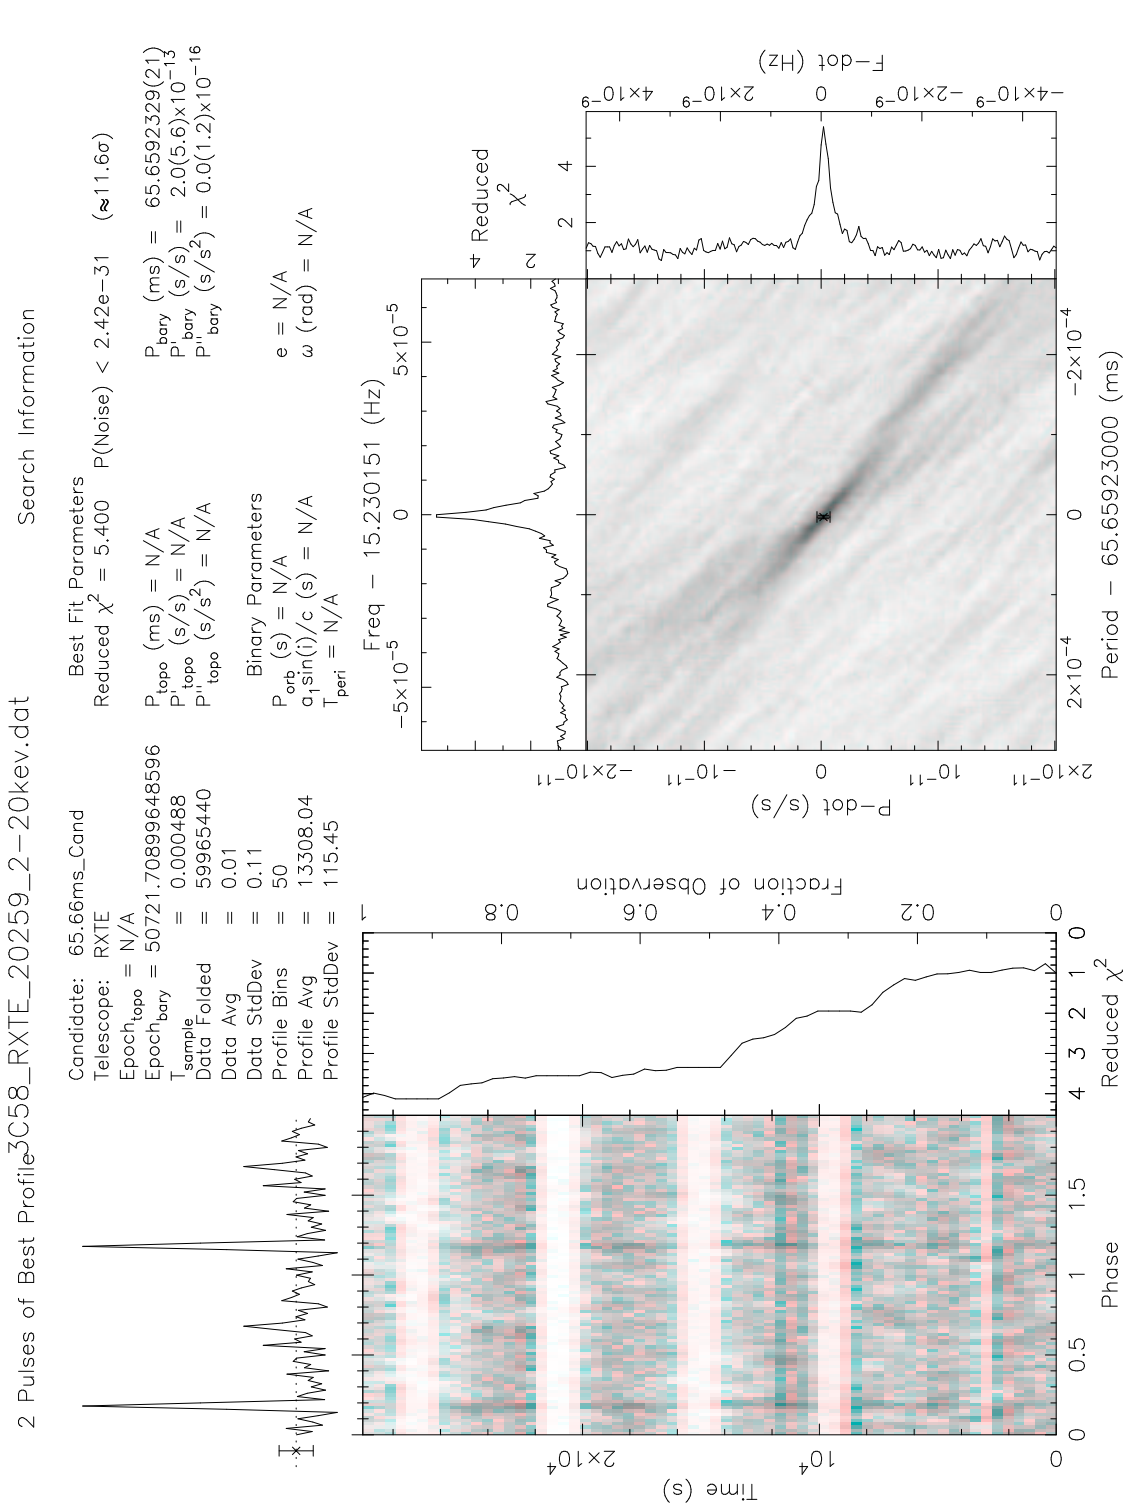


Figure 5.4 The new pulsar PSR J0205+6449 in the supernova remnant 3C 58 as seen in a 20 ks archival *RXTE* observation.

TABLE 5.2
RXTE SNR Pulsation Search

Supernova Remnant	T_{src} (10^3 s)	$N_{photons}$ (10^3)	N_{bins} (10^6)	N_{trials} (10^7)	$\log_2(T_{samp})$ (s)	Highest P_{norm}	Pulsed Fraction
G18.9–1.1	49	1636	48.5	6.8	-9	22	< 1.1%
G20.0–0.2	32	862	31	4.3	-9	22	< 1.5%
G21.5–0.9	20	1408	37	5.2	-10	21	< 1.1%
G24.7+0.6	45	1493	180	25	-11	22	< 1.1%
G27.8+0.6	15	467	50	7.1	-11	25	< 1.9%
G63.7+1.1	91	2814	192	27	-10	22	< 0.8%
G292.0+1.8	94	2716	63.5	8.9	-9	23	< 0.8%
G326.3–1.8	43	731	32	4.5	-9	20	< 1.6%
γ Cygni	197	5611	405	5.7	-10	24	< 0.6%
3C 58	20	445	15	2.1	-9	20 ^a	$\sim 0.6\%$ ^b

^a The maximum normalized power due to PSR J0205+6449 was ~ 13.5

^b This is the observed *RXTE* pulsed fraction for PSR J0205+6449.

Table 5.2 The values for “Highest P_{norm} ” are the highest individual powers (normalized with $N_{photons}$) found in the FFT. The N_{trials} column is the approximate number of independent Fourier trials that each search covered. All upper limits were calculated assuming a sinusoidal (i.e. worst case) profile, allowing for drifts of up to 10 Fourier bins during the observation. Frequencies up to 250, 500, or 1000 Hz were searched for sampling times $\log_2(T_{samp})$ of -9, -10, and -11 respectively. Upper limits were calculated according to the method of Vaughan et al. (1994). The pulsed fraction given corresponds to frequencies $\lesssim 100$ Hz. Higher frequencies have progressively worse limits.

5.4 Future Possibilities

Within the next few months, my collaborators (Vicky Kaspi, Ingrid Stairs, Jason Hessels) and I will make some of the first pulsar searches with the new GBT and the upgraded Arecibo telescope. These observations will include numerous 8 h integrations of globular clusters from the GBT and 1–2 h integrations using Arecibo. A large “Beowulf” cluster of fast workstations at McGill University will provide the muscle required by the searches.

Beyond radio data, there are many long observations of LMXBs, SNRs, and other interesting objects that could benefit from a sophisticated analysis for both binary and isolated pulsars. The application of sideband searches to long observations of LMXBs is especially intriguing.

With PRESTO debugged and working, lots of CPU-cycles available, and a torrent of data available or on the way, I am looking forward to many hours spent searching for new pulsars.

Appendix A

Derivation of Fourier Interpolation

Following the derivation found in Middleditch et al. (1993), we begin with the definition of the k^{th} DFT element

$$A_k = \sum_{j=0}^{N-1} n_j e^{-2\pi i j k / N}, \quad (\text{A.1})$$

which we then rewrite by substituting the inverse DFT for the n_j

$$A_k = \sum_{j=0}^{N-1} \left(\frac{1}{N} \sum_{l=0}^{N-1} A_l e^{2\pi i j l / N} \right) e^{-2\pi i j k / N} \quad (\text{A.2a})$$

$$= \frac{1}{N} \sum_{l=0}^{N-1} A_l \sum_{j=0}^{N-1} e^{-2\pi i j (k-l) / N}. \quad (\text{A.2b})$$

The last summation can be computed exactly using the identity

$$\sum_{j=0}^{N-1} e^{i\alpha j} = e^{\frac{i\alpha}{2}(N-1)} \frac{\sin\left(\frac{N\alpha}{2}\right)}{\sin\left(\frac{\alpha}{2}\right)}, \quad (\text{A.3})$$

such that when $N \gg 1$ we have

$$\sum_{j=0}^{N-1} e^{-2\pi i j (k-l) / N} = e^{-i\pi(k-l)(1-\frac{1}{N})} \frac{\sin[\pi(k-l)]}{\sin[\pi(k-l)/N]} \quad (\text{A.4a})$$

$$\simeq e^{-i\pi(k-l)} \frac{\sin[\pi(k-l)]}{\pi(k-l)/N} \quad (\text{A.4b})$$

$$\simeq N e^{-i\pi(k-l)} \text{sinc}[\pi(k-l)]. \quad (\text{A.4c})$$

Substituting this expression into eqn. A.2b and changing the integer frequency k into a continuous real-valued frequency r , we arrive at eqn. 2.29

$$A_r = \sum_{l=0}^{N-1} A_l e^{-i\pi(r-l)} \text{sinc}[\pi(r-l)]. \quad (\text{A.5})$$

Appendix B

Derivation of Centroid

The centroid is a measure of the approximate location of a signal in a time series as estimated by the first moment of the signal with respect to time (see §2.5.2). We can think of a sinusoidal signal in our data as being always present but modulated in intensity by some window function $W(u)$, where $u = t/T$ is the normalized time and T is the length of the observation. A “normal” observation of a pulsar of constant intensity would therefore have $W(u) = 1$ when $0 < u < 1$ and $W(u) = 0$ at all other times (i.e. a square window). Our signal is therefore described by

$$s(u) = a \cos(2\pi r_o u + \phi_o) W(u), \quad (\text{B.1})$$

where a is the amplitude, $r_o = f_o T$ is the Fourier frequency, and ϕ_o is the phase of the sinusoid at time $u = 0$.

Since the centroid of a function is proportional to the first moment of the function with respect to time, we can easily calculate the centroid using the Moment Theorem of Fourier transforms. Bracewell (1999) does this and defines the centroid as

$$\langle u \rangle = -\frac{A'(0)}{2\pi i A(0)}, \quad (\text{B.2})$$

where $A(0)$ and $A'(0)$ are the Fourier transform and its first derivative with respect to r measured at $r = 0$.

Eqn. B.2 is not directly applicable for our sinusoidal signal since the information about the window function in eqn. B.1 has been shifted to the frequency of the sinusoid in accordance with the Modulation Theorem of Fourier transforms. Accordingly, we can apply the Modulation Theorem to eqn. B.2 which gives us

$$\langle u \rangle_{s(u)} = -\frac{A'(r_o)}{2\pi i A(r_o)}. \quad (\text{B.3})$$

Finally, we can write $A(r)$ in phasor form as $A(r) = a(r)e^{i\phi(r)}$ where $a(r)$ and $\phi(r)$ represent the Fourier amplitude and phase as functions of the Fourier frequency

r . The derivative with respect to Fourier frequency can be written

$$A'(r) = \frac{\partial a(r)}{\partial r} e^{i\phi(r)} + a(r) i \frac{\partial \phi(r)}{\partial r} e^{i\phi(r)}. \quad (\text{B.4})$$

At the frequency of our signal, the amplitude is $A(r_o) = a(r_o) e^{i\phi(r_o)}$ and the Fourier response is at its peak, making $\frac{\partial a(r_o)}{\partial r} = 0$. Therefore

$$A'(r_o) = A(r_o) i \frac{\partial \phi(r_o)}{\partial r}. \quad (\text{B.5})$$

Substituting eqn. B.5 into eqn. B.3 we arrive at

$$\hat{C} = \langle u \rangle_{s(u)} = -\frac{1}{2\pi} \frac{\partial \phi(r_o)}{\partial r}, \quad (\text{B.6})$$

which is equivalent to eqn. 2.49.

Appendix C

Derivation of Purity

The “purity” of a signal (see §2.5.2) is a measure of the rms dispersion of the pulsations in time with respect to the centroid and is directly proportional to the variance in time of the window function of the sinusoid from eqn. B.1. The time variance is defined as

$$\langle (u - \langle u \rangle)^2 \rangle = \frac{\int_{-\infty}^{\infty} (u - \langle u \rangle)^2 W(u) du}{\int_{-\infty}^{\infty} W(u) du}, \quad (\text{C.1})$$

but can be written as

$$\langle (u - \langle u \rangle)^2 \rangle = \langle u^2 \rangle - \langle u \rangle^2 \quad (\text{C.2a})$$

$$= -\frac{A''(0)}{4\pi^2 A(0)} + \frac{[A'(0)]^2}{4\pi^2 [A(0)]^2}, \quad (\text{C.2b})$$

using the Moment Theorem for Fourier transforms (see e.g. Bracewell 1999). Since our signal is sinusoidal (see Appendix B), application of the Modulation Theorem gives

$$\langle (u - \langle u \rangle)^2 \rangle_{s(u)} = -\frac{A''(r_o)}{4\pi^2 A(r_o)} + \frac{[A'(r_o)]^2}{4\pi^2 [A(r_o)]^2} \quad (\text{C.3a})$$

$$= -\frac{A''(r_o)}{4\pi^2 A(r_o)} - \hat{C}^2. \quad (\text{C.3b})$$

Using $A(r) = a(r)e^{i\phi(r)}$ and remembering that $\frac{\partial a(r_o)}{\partial r} = 0$, the second derivative of the Fourier amplitude is

$$A''(r_o) = iA(r_o)\frac{\partial^2 \phi(r_o)}{\partial r^2} - A(r_o)\left(\frac{\partial \phi(r_o)}{\partial r}\right)^2 + \frac{A(r_o)}{a(r_o)}\frac{\partial^2 a(r_o)}{\partial r^2}. \quad (\text{C.4})$$

From Appendix B, we know that $\frac{\partial \phi(r_o)}{\partial r} = -2\pi\hat{C}$ which makes the first term of eqn. C.4 equal to zero since $\frac{\partial^2 \phi(r_o)}{\partial r^2} = 0$ and the second term equal to $-4\pi^2 A(r_o)\hat{C}^2$.

The second derivative of power at the peak response can be written as

$$\frac{\partial^2 P(r_o)}{\partial r^2} = \frac{\partial^2}{\partial r^2} (A^*(r_o)A(r_o)) = 2a(r_o) \frac{\partial^2 a(r_o)}{\partial r^2}, \quad (\text{C.5})$$

making the third term equal to $\frac{A(r_o)}{2(a(r_o))^2} \frac{\partial^2 P(r_o)}{\partial r^2}$. Substituting into eqn. C.3a and simplifying gives

$$\langle (u - \langle u \rangle)^2 \rangle_{s(u)} = -\frac{1}{8\pi^2 P(r_o)} \frac{\partial^2 P(r_o)}{\partial r^2}. \quad (\text{C.6})$$

If normalize the variance using the value obtained for a signal present throughout the observation (i.e. a square window where $W(u) = 1$ from $0 \leq u \leq 1$ and zero elsewhere — which we will call a *pure* signal) where

$$\langle (u - \langle u \rangle)^2 \rangle_{\text{pure}} = \langle u^2 \rangle - \langle u \rangle^2 = \frac{1}{3} - \left(\frac{1}{2}\right)^2 = \frac{1}{12}, \quad (\text{C.7})$$

and then taking the square-root, we are left with

$$\alpha = \sqrt{\frac{\langle (u - \langle u \rangle)^2 \rangle_{s(u)}}{\langle (u - \langle u \rangle)^2 \rangle_{\text{pure}}}} = \frac{1}{\pi} \sqrt{-\frac{3}{2P(r_o)} \frac{\partial^2 P(r_o)}{\partial r^2}}, \quad (\text{C.8})$$

which is equivalent to eqn. 2.50.

Appendix D

Centroid, Purity, and Fourier Response

In order to consider the effects of centroid and purity on the Fourier response to a sinusoidal signal as described by eqn. B.1, we initially assume a Fourier response equal to eqn. 2.18c of

$$A_r = A_o e^{-i\pi(r-r_o)} \text{sinc} [\pi(r - r_o)] \quad (\text{D.1})$$

where $A_o = \frac{Na}{2} e^{i\phi_o}$, ϕ_o is the intrinsic phase of the signal, and r_o is the true signal frequency (in FFT bins). This response is correct only for signals with a square window function (i.e. $W(u) = 1$ from $0 \leq u \leq 1$ and zero elsewhere).

From eqn. D.1 we see that a change in Fourier frequency of a single Fourier bin causes a change in the measured Fourier phase of π radians. This phase change is also visible from the centroid equation for a pure signal with $\hat{C} = \frac{1}{2}$, where $d\phi(r) = -\pi$ for every $dr = 1$. Rewriting the centroid equation and integrating, we see that the Fourier phase near the peak response goes as

$$\phi(r) = -2\pi r \hat{C} + c. \quad (\text{D.2})$$

When $r = r_o$, $\phi = \phi_o$, allowing us to solve for the constant of integration, $c = \phi_o + 2\pi r_o \hat{C}$. Substituting into eqn. D.2, we see that the phase of the Fourier response is equal to

$$\phi(r_o) = \phi_o - 2\pi \hat{C}(r - r_o). \quad (\text{D.3})$$

Therefore, for signals that have centroids different from $\frac{1}{2}$, the phase change across a single Fourier bin is different from the usual π radians.

The purity, as described in §2.5 and Appendix C, is the effective duration of a square window which reproduces the measured rms dispersion of the signal in time about the centroid. Since the Fourier response to a square window goes as $\text{sinc}(\pi fT)$, where $fT = r$, we can see that replacing the window of length T with

one of effective duration αT causes the Fourier response to go as $\text{sinc}(\pi\alpha fT)$. This fact is also approximately true for more complicated window functions as long as $|r - r_o| \ll 1$. The Fourier response to a windowed sinusoid is therefore

$$A_r \simeq A_o e^{-2\pi i \hat{C}(r-r_o)} \text{sinc} [\pi\alpha(r - r_o)]. \quad (\text{D.4})$$

Numerical simulations show that this approximation is valid for purity values $\alpha \lesssim 1.5$.

Appendix E

Derivation of Frequency Uncertainty

For a given Fourier frequency offset $\Delta_r = r - r_o$, where r_o is the Fourier frequency of the signal, the magnitude of the Fourier response and the power go as (see Appendix D)

$$|A(r)| = |A_o| \text{sinc}(\pi\alpha\Delta_r) \quad (\text{E.1a})$$

$$P(r) = P_{meas} \text{sinc}^2(\pi\alpha\Delta_r) \quad (\text{E.1b})$$

where α is the signal purity. We can expand the sinc function in order to approximate the expression for the power near the peak of the response as

$$P(r) = P_{meas} \left[\frac{\sin(\pi\alpha\Delta_r)}{\pi\alpha\Delta_r} \right]^2 \quad (\text{E.2a})$$

$$\simeq P_{meas} \left[1 - \frac{(\pi\alpha\Delta_r)^2}{3} \right]. \quad (\text{E.2b})$$

Taking the derivative of power with respect to r and solving for Δ_r we obtain

$$\Delta_r = -\frac{3}{2\pi^2\alpha^2 P_{meas}} \frac{\partial P}{\partial r}. \quad (\text{E.3})$$

As expected, when the Fourier frequency equals the true frequency of the pulsations (i.e. $\Delta_r = 0$), the Fourier response peaks and $\frac{\partial P(r)}{\partial r} = 0$.

In order to estimate the uncertainty in Δ_r , we apply standard propagation of errors to arrive at

$$\sigma_r = \frac{3}{2\pi^2\alpha^2 P_{meas}} \sigma_{P'(r)}, \quad (\text{E.4})$$

where we have replaced $\frac{\partial P(r)}{\partial r}$ with $P'(r)$ to simplify the notation. The derivative of the power at the true frequency can be approximated using finite differences as

$$P'(r) = \frac{P(r_o + \Delta_r) - P(r_o - \Delta_r)}{2\Delta_r}. \quad (\text{E.5})$$

The uncertainty in $P'(r)$ can also be approximated using finite differences and error propagation. Since $P(r_o + \Delta_r)$ and $P(r_o - \Delta_r)$ are highly correlated when $\Delta_r \ll 1$, their uncertainties are also correlated giving

$$\sigma_{P'(r)} \simeq \frac{1}{2\Delta_r} (\sigma_{P(r_o+\Delta_r)} + \sigma_{P(r_o-\Delta_r)}) \quad (\text{E.6a})$$

$$\simeq \frac{\sigma_{P(r_o+\Delta_r)}}{\Delta_r}. \quad (\text{E.6b})$$

Now, we turn to the question of the uncertainty in $P(r_o + \Delta_r)$, closely following Middleditch (1976). The amplitude of the Fourier response at the true signal frequency can be written as

$$\sqrt{P_{meas}} = \sum_{j=0}^{N-1} y_j \cos(\phi_j) \quad (\text{E.7})$$

where $y_j = \frac{n_j}{\sqrt{P_{norm}}}$ are the points in our time series as defined in eqn. 2.1, but scaled using the appropriate P_{norm} such that the measured power, P_{meas} , is properly normalized (see §2.3.1). Similarly, the ϕ_j represent the pulsation phases at times $u = \frac{j}{N}$, but rotated by the measured Fourier phase, ϕ_{meas} , such that the result of the vector addition lies along the real axis in the complex plane (i.e. the final complex phase is zero). In effect, this transform isolates components of the data that are parallel to the final Fourier response.

At a small frequency offset Δ_r from the true frequency, we can expand the power in a similar fashion as

$$\sqrt{P(r_o + \Delta_r)} = \sum_{j=0}^{N-1} y_j \cos(\phi_j + \delta_{\phi_j}), \quad (\text{E.8})$$

where δ_{ϕ_j} are the ‘‘phase errors’’ introduced by the frequency offset. The phase errors add curvature to the vector addition and are defined as

$$\delta_{\phi_j} = 2\pi\Delta_r \left(\frac{j}{N} - \frac{1}{2} \right) = 2\pi\Delta_r \left(u - \frac{1}{2} \right), \quad (\text{E.9})$$

where u is the normalized time, $u = \frac{t}{T} = \frac{j}{N}$. The $\frac{1}{2}$ term in eqn. E.9 removes the accumulated phase error over the course of the observation (i.e. $\int_0^1 2\pi\Delta_r u du = \pi\Delta_r$) and makes the vector summation of eqn. E.7 finish on the real axis. Expanding the cosine in eqn. E.8 gives

$$\sqrt{P(r_o + \Delta_r)} = \sum_{j=0}^{N-1} y_j \cos(\phi_j) \cos(\delta_{\phi_j}) - \sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}). \quad (\text{E.10})$$

Considering the uncertainties in the separate terms of eqn. E.10, since the cosine term is an even function of δ_{ϕ_j} , it is symmetric about r_o , and therefore does not contribute to the uncertainty in the $P'(r)$ measurement as defined by eqn. E.5. For the sine term, given that the $\cos(\phi_j)$ “derotates” the signal onto the real axis by definition, we see that

$$\sum_{j=0}^{N-1} y_j \sin(\phi_j) = 0. \quad (\text{E.11})$$

Furthermore, since the ϕ_j and δ_{ϕ_j} are uncorrelated, the average value of this term will be zero,

$$\left\langle \sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}) \right\rangle = 0, \quad (\text{E.12})$$

and has no systematic effect on $P(r_o + \Delta_r)$. However, we can calculate the fluctuations introduced by this term

$$\left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}) \right)^2 \right\rangle = \left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j) \right)^2 \right\rangle \langle \sin^2(\delta_{\phi_j}) \rangle. \quad (\text{E.13})$$

where the cross-terms average to zero since ϕ_j and δ_{ϕ_j} are uncorrelated. Due to the normalization of the y_j , the sum component averages to $\frac{1}{2}$. The δ_{ϕ_j} component has an average of

$$\langle \sin^2(\delta_{\phi_j}) \rangle = \int_0^1 \sin^2(\delta_{\phi_j}) \, du \quad (\text{E.14a})$$

$$\simeq \int_0^1 \delta_{\phi_j}^2 \, du \quad (\text{E.14b})$$

$$\simeq \int_0^1 4\pi^2 \Delta_r^2 \left(u^2 - u + \frac{1}{4} \right) \, du \quad (\text{E.14c})$$

$$\simeq \frac{\pi^2 \Delta_r^2}{3}. \quad (\text{E.14d})$$

Therefore eqn. E.13 is equal to $\frac{\pi^2 \Delta_r^2}{6}$, and the variance of $\sqrt{P(r_o + \Delta_r)}$ will be

$$\sigma_{\sqrt{P(r_o + \Delta_r)}}^2 = \sqrt{P(r_o + \Delta_r)} \frac{\pi^2 \Delta_r^2}{6}. \quad (\text{E.15})$$

To get the standard deviation, we take the square root of this expression, but then squaring to get the uncertainty in $P(r_o + \Delta_r)$ doubles the uncertainty to give

$$\sigma_{P(r_o + \Delta_r)} = 2\pi \Delta_r \alpha \sqrt{\frac{P_{meas}}{6}} \quad (\text{E.16})$$

where we have used the approximation $P(r_o + \Delta_r) \simeq P_{meas}$. Then, substituting into equation E.6b, we find

$$\sigma_{P'(r)} = 2\pi\alpha\sqrt{\frac{P_{meas}}{6}}. \quad (\text{E.17})$$

And finally, substituting this expression into eqn. E.4, we have

$$\sigma_r = \frac{3}{2\pi^2\alpha^2 P_{meas}} \left(2\pi\alpha\sqrt{\frac{P_{meas}}{6}} \right) \quad (\text{E.18a})$$

$$= \frac{3}{\pi\alpha\sqrt{6P_{meas}}}. \quad (\text{E.18b})$$

Appendix F

Derivation of \dot{f} Uncertainty

If a peak in the f - \dot{f} plane has been isolated using techniques similar to those shown in §2.4.2, we can calculate the error in the measurement of the true Fourier frequency derivative $\dot{r}_o = \dot{f}_o T^2$ in a manner similar to that for the frequency uncertainty as described in Appendix E. Signals with non-zero frequency derivatives have Fourier peaks which are located off the $\dot{r} = 0$ line in the f - \dot{f} plane, but the *shapes* of those peaks are independent of \dot{r} and in fact depend only on the window function of the signal (see Appendix D). The shape of the response in power as a function of $\Delta_{\dot{r}} = \dot{r} - \dot{r}_o$ at the “correct” Fourier frequency $r = r_o$ is described by eqn. 2.38, and can be written as

$$A(r_o, \dot{r}) = A_o e^{-i\pi \frac{\Delta_{\dot{r}}}{4}} \sqrt{\frac{2}{\Delta_{\dot{r}}}} \left[S \left(\sqrt{\frac{\Delta_{\dot{r}}}{2}} \right) - iC \left(\sqrt{\frac{\Delta_{\dot{r}}}{2}} \right) \right], \quad (\text{F.1})$$

when q_r is defined as

$$q_r = r_c - r'_c = \left(r_o + \frac{\dot{r}_o}{2} \right) - \left(r_o + \frac{\dot{r}_o + \Delta_{\dot{r}}}{2} \right) = -\frac{\Delta_{\dot{r}}}{2}. \quad (\text{F.2})$$

This definition of q_r keeps the magnitude of the Fourier response symmetric about q_r no matter what the value of $\Delta_{\dot{r}}$. The power as a function of $\Delta_{\dot{r}}$ is therefore

$$P(r_o, \dot{r}) = P_o \frac{2}{\Delta_{\dot{r}}} \left\{ \left[S \left(\sqrt{\frac{\Delta_{\dot{r}}}{2}} \right) \right]^2 + \left[C \left(\sqrt{\frac{\Delta_{\dot{r}}}{2}} \right) \right]^2 \right\}. \quad (\text{F.3})$$

If we expand the Fresnel integrals about 0 as

$$C(x) \simeq x - \frac{\pi^2}{40} x^5 \dots, \quad \text{and} \quad S(x) \simeq \frac{\pi}{6} x^3 \dots, \quad (\text{F.4})$$

and then substitute, the power becomes

$$P(r_o, \dot{r}) \simeq P_o \left(1 - \frac{\pi^2 \Delta_{\dot{r}}^2}{180} \right). \quad (\text{F.5})$$

Taking the derivative of $P(r_o, \dot{r})$ with respect to \dot{r} and solving for $\Delta_{\dot{r}}$ gives

$$\Delta_{\dot{r}} = -\frac{90}{\pi^2 P_o} \frac{\partial P(r_o, \dot{r})}{\partial \dot{r}}. \quad (\text{F.6})$$

From propagation of errors,

$$\sigma_{\dot{r}} = -\frac{90}{\pi^2 P_o} \sigma_{P'(r_o, \dot{r})}, \quad (\text{F.7})$$

where $P'(r_o, \dot{r}) = \frac{\partial P(r_o, \dot{r})}{\partial \dot{r}}$. Similarly, after using a finite difference estimate of the power derivative following Appendix E, we find

$$\sigma_{P'(r_o, \dot{r})} = \frac{\sigma_{P(r_o, \Delta_{\dot{r}})}}{\Delta_{\dot{r}}}, \quad (\text{F.8})$$

where $P(r_o, \Delta_{\dot{r}})$ represents the power as measured at $\dot{r} = \dot{r}_o + \Delta_{\dot{r}}$.

Closely following Appendix E, we represent $\sqrt{P(r_o, \Delta_{\dot{r}})}$ as a sum of the parallel components of the properly normalized time series points, y_j , as

$$\sqrt{P(r_o, \Delta_{\dot{r}})} = \sum_{j=0}^{N-1} y_j \cos(\phi_j + \delta_{\phi_j}). \quad (\text{F.9})$$

The δ_{ϕ_j} are the ‘‘phase errors’’ introduced when $\Delta_{\dot{r}} \neq 0$, and are defined as

$$\delta_{\phi_j} = \pi \Delta_{\dot{r}} \left(u^2 - u + \frac{1}{6} \right), \quad (\text{F.10})$$

where the u term comes from keeping the response symmetric about r_o (i.e. $q_r = -\frac{\Delta_{\dot{r}}}{2}$) and the $\frac{1}{6}$ removes the accumulated phase error over the course of the observation (i.e. $\int_0^1 \pi \Delta_{\dot{r}} (u^2 - u) du = -\frac{\pi \Delta_{\dot{r}}}{6}$) and makes the vector summation of eqn. F.9 finish on the real axis.

Expanding the cosine term of eqn. F.9 gives

$$\sqrt{P(r_o, \Delta_{\dot{r}})} = \sum_{j=0}^{N-1} y_j \cos(\phi_j) \cos(\delta_{\phi_j}) - \sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}). \quad (\text{F.11})$$

The first term shortens both $P(r_o, \Delta_{\dot{r}})$ and $P(r_o, -\Delta_{\dot{r}})$ by the same amount and does not affect the derivative of power. For the sine term, since ϕ_j and δ_{ϕ_j} are uncorrelated, the average value is zero (see Appendix E), but its fluctuations are important. To calculate the fluctuations, we square the terms and get

$$\left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}) \right)^2 \right\rangle = \left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j) \right)^2 \right\rangle \langle \sin^2(\delta_{\phi_j}) \rangle. \quad (\text{F.12})$$

Due to the normalization of the y_j , $\left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j) \right)^2 \right\rangle$ averages to $\frac{1}{2}$ as before, and we can directly calculate $\langle \sin^2(\delta_{\phi_j}) \rangle$ as

$$\left\langle \sum_{j=0}^{N-1} \sin^2(\delta_{\phi_j}) \right\rangle = \int_0^1 \sin^2(\delta_{\phi_j}) \, du \quad (\text{F.13a})$$

$$\simeq \int_0^1 \delta_{\phi_j}^2 \, du \quad (\text{F.13b})$$

$$= \pi^2 \Delta_{\dot{r}}^2 \int_0^1 \left(u^2 - u + \frac{1}{6} \right)^2 \, du \quad (\text{F.13c})$$

$$= \frac{\pi^2 \Delta_{\dot{r}}^2}{180}. \quad (\text{F.13d})$$

The fluctuations from the sine term are therefore $\frac{\pi^2 \Delta_{\dot{r}}^2}{360}$, and since squaring eqn. F.9 doubles the errors, the standard deviation of $P(r_o, \Delta_{\dot{r}})$ is

$$\sigma_{P(r_o, \Delta_{\dot{r}})} = \sqrt{P(r_o, \Delta_{\dot{r}})} \frac{2\pi \Delta_{\dot{r}}}{\sqrt{360}} \simeq \sqrt{P_{meas}} \frac{\pi \Delta_{\dot{r}}}{\sqrt{90}}. \quad (\text{F.14})$$

Substituting into equations F.8 and then F.7 as in Appendix E gives us the uncertainty in the frequency derivative

$$\sigma_{\dot{r}} = \frac{3\sqrt{10}}{\pi \sqrt{P_{meas}}}. \quad (\text{F.15})$$

References

- Abramowitz, M. & Stegun, I. A., eds. 1972, Handbook of Mathematical Functions (New York: Dover)
- Anderson, S. B. 1992, PhD thesis, California Institute of Technology
- Anderson, S. B., Gorham, P. W., Kulkarni, S. R., Prince, T. A., & Wolszczan, A. 1990, Nature, 346, 42
- Arfken, G. 1985, Mathematical Methods for Physicists, 3rd edition (San Diego: Academic Press)
- Bailey, D. H. 1990, Journal of Supercomputing, 4, 23
- Bailyn, C. D. & Grindlay, J. E. 1990, ApJ, 353, 159
- Biggs, J. D. & Lyne, A. G. 1996, MNRAS, 282, 691
- Blackman, R. B. & Tukey, J. W. 1959, The Measurement of Power Spectra (New York: Dover)
- Bracewell, R. 1999, The Fourier Transform and its Applications, 3rd edition (New York: McGraw-Hill)
- Burns, W. R. & Clark, B. G. 1969, A&A, 2, 280
- Camilo, F., Lorimer, D. R., Freire, P., Lyne, A. G., & Manchester, R. N. 2000, ApJ, 535, 975
- Chen, K. & Leonard, P. J. T. 1993, ApJ, 411, L75
- Chen, K., Middleditch, J., & Ruderman, M. 1993, ApJ, 408, L17
- Chou, Y., Grindlay, J. E., & Bloser, P. F. 2001, ApJ, 549, 1135
- Cordes, J. M. 2001, "Pulsar Observations I. Propagation Effects, Searching, Distance Estimates, Scintillations and VLBI", unpublished, available from <http://www.astro.cornell.edu/~cordes/documents/>

- D'Amico, N., Lyne, A. G., Manchester, R. N., Possenti, A., & Camilo, F. 2001, *ApJ*, 548, L171
- Davies, M. B. & Hansen, B. M. S. 1998, *MNRAS*, 301, 15
- Deich, W. T. S., Middleditch, J., Anderson, S. B., Kulkarni, S. R., Prince, T. A., & Wolszczan, A. 1993, *ApJ*, 410, L95
- Edmonds, P. D., Gilliland, R. L., Heinke, C. O., Grindlay, J. E., & Camilo, F. 2001, *astro-ph/0107096*
- Eikenberry, S. S., Fazio, G. G., & Ransom, S. M. 1998, *ApJ*, 492, 754
- Fraser, D. 1976, *Journal of the Association of Computing Machinery*, 23, 298
- Freire, P. C., Camilo, F., Lorimer, D. R., Lyne, A. G., Manchester, R. N., & D'Amico, N. 2001a, *astro-ph/0103372*
- Freire, P. C., Kramer, M., Lyne, A. G., Camilo, F., Manchester, R. N., & D'Amico, N. 2001b, *astro-ph/0107206*
- Fruchter, A. S. & Goss, W. M. 1990, *ApJ*, 365, L63
- Fruchter, A. S. & Goss, W. M. 2000, *ApJ*, 536, 865
- Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, *Nature*, 333, 237
- Grindlay, J. E., Heinke, C., Edmonds, P. D., & Murray, S. S. 2001, *Science*, 292, 2290
- Groth, E. J. 1975, *ApJS*, 29, 285
- Hamilton, T. T., Helfand, D. J., & Becker, R. H. 1985, *AJ*, 90, 606
- Harris, W. E. 1996, *AJ*, 112, 1487
- Haslam, C. G. T., Stoffel, H., Salter, C. J., & Wilson, W. E. 1982, *A&AS*, 47, 1
- Hut, P., McMillan, S., Goodman, J., Mateo, M., Phinney, E. S., Pryor, C., Richer, H. B., Verbunt, F., & Weinberg, M. 1992, *PASP*, 104, 981
- Israel, G. L. & Stella, L. 1996, *ApJ*, 468, 369
- Johnston, H. M. & Kulkarni, S. R. 1991, *ApJ*, 368, 504
- Johnston, H. M., Kulkarni, S. R., & Goss, W. M. 1991, *ApJ*, 382, L89

- Konacki, M. & Maciejewski, A. J. 1996, *Icarus*, 122, 347
- Kulkarni, S. R. & Anderson, S. B. 1996, in *Dynamical Evolution of Star Clusters – Confrontation of Theory and Observations: IAU Symposium 174*, Vol. 174, 181
- Kulkarni, S. R., Goss, W. M., Wolszczan, A., & Middleditch, J. 1990, *ApJ*, 363, L5
- Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *MNRAS*, 225, 307
- Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Sutherland, P. G., Kahn, S., & Grindlay, J. E. 1983, *ApJ*, 266, 160
- Liu, Q. Z., van Paradijs, J., & van den Heuvel, E. P. J. 2001, *A&A*, 368, 1021
- Lyne, A. G. 1995, in *Millisecond Pulsars - A Decade of Surprise*, ed. A. S. Fruchter, M. Tavani, & D. C. Backer, Vol. 72 (Astronomical Society of the Pacific), 35–45
- Lyne, A. G., Brinklow, A., Middleditch, J., Kulkarni, S. R., Backer, D. C., & Clifton, T. R. 1987, *Nature*, 328, 399
- Lyne, A. G., Camilo, F., Manchester, R. N., Bell, J. F., Kaspi, V. M., D’Amico, N., McKay, N. P. F., Crawford, F., Morris, D. J., Sheppard, D. C., & Stairs, I. H. 2000a, *MNRAS*, 312, 698
- Lyne, A. G. & Graham-Smith, F., eds. 1998, *Pulsar Astronomy*, 2nd edition (Cambridge: Cambridge University Press)
- Lyne, A. G., Mankelov, S. H., Bell, J. F., & Manchester, R. N. 2000b, *MNRAS*, 316, 491
- Manchester, R. N., Lyne, A. G., Camilo, F., Bell, J. F., Kaspi, V. M., D’Amico, N., McKay, N. P. F., Crawford, F., Stairs, I. H., Possenti, A., Kramer, M., & Sheppard, D. C. 2001, *MNRAS*, in press (astro-ph/0106522)
- Mardia, K. V. & Zemroch, P. J. 1975, *Applied Statistics*, 24, 268
- Mattox, J. R., Koh, D. T., Lamb, R. C., Macomb, D. J., Prince, T. A., & Ray, P. S. 1996, *A&AS*, 120, 95
- Middleditch, J. 1976, PhD thesis, California Univ., Berkeley

- Middleditch, J. & Cordova, F. A. 1982, *ApJ*, 255, 585
- Middleditch, J., Deich, W., & Kulkarni, S. 1993, in *Isolated Pulsars*, ed. K. A. Van Riper, R. Epstein, & C. Ho (Cambridge University Press), 372–379
- Middleditch, J. & Kristian, J. 1984, *ApJ*, 279, 157
- Middleditch, J., Kristian, J. A., Kunkel, W. E., Hill, K. M., Watson, R. D., Lucinio, R., Imamura, J. N., Steiman-Cameron, T. Y., Shearer, A., Butler, R., Redfern, M., & Danks, A. C. 2000, *New Astronomy*, 5, 243
- Middleditch, J., Mason, K. O., Nelson, J. E., & White, N. E. 1981, *ApJ*, 244, 1001
- Moore, C. B., Rutledge, R. E., Fox, D. W., Guerriero, R. A., Lewin, W. H. G., Fender, R., & van Paradijs, J. 2000, *ApJ*, 532, 1181
- Nice, D. J., Arzoumanian, Z., & Thorsett, S. E. 2000, in *ASP Conf. Ser. 202: Pulsar Astronomy - 2000 and Beyond*, 67
- Nice, D. J. & Thorsett, S. E. 1992, *ApJ*, 397, 249
- Phinney, E. S. & Kulkarni, S. R. 1994, *ARA&A*, 32, 591
- Possenti, A., D’Amico, N., Manchester, R. N., Sarkissian, J., Lyne, A. G., & Camilo, F. 2001, [astro-ph/0108343](#)
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes: The Art of Scientific Computing*, 2nd edition (Cambridge: Cambridge University Press)
- Ransom, S. M. 2000, in *ASP Conf. Ser. 202: Pulsar Astronomy - 2000 and Beyond*, 43
- Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2001a, submitted to *ApJ*.
- Ransom, S. M., Greenhill, L. J., Herrnstein, J. R., Manchester, R. N., Camilo, F., Eikenberry, S. S., & Lyne, A. G. 2001b, *ApJ*, 546, L25
- Rasio, F. A. & Heggie, D. C. 1995, *ApJ*, 445, L133
- Rasio, F. A., Pfahl, E. D., & Rappaport, S. 2000, *ApJ*, 532, L47
- Standish, E. M. 1982, *A&A*, 114, 297

- Stappers, B. W., Bailes, M., Lyne, A. G., Manchester, R. N., D'Amico, N., Tauris, T. M., Lorimer, D. R., Johnston, S., & Sandhu, J. S. 1996, *ApJ*, 465, L119
- Staveley-Smith, L., Wilson, W. E., Bird, T. S., Disney, M. J., Ekers, R. D., Freeman, K. C., Haynes, R. F., Sinclair, M. W., Vaile, R. A., Webster, R. L., & Wright, A. E. 1996, *Proc. Astron. Soc. Aust.*, 13, 243
- Sweet, R. & Wilson, J. 1995, Development of Out-of-Core Fast Fourier Transform Software for the Connection Machine, Tech. rep., Center for Computational Mathematics, University of Colorado at Denver
- Taylor, J. H. & Cordes, J. M. 1993, *ApJ*, 411, 674
- Taylor, J. H., Manchester, R. N., Lyne, A. G., & Camilo, F. 1995, unpublished (available at <ftp://pulsar.princeton.edu/pub/catalog>)
- van der Klis, M. 1989, in *Timing Neutron Stars*, (NATO ASI Series), ed. H. Ögelman & E. P. J. van den Heuvel (Dordrecht: Kluwer), 27–69
- 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. 1993, *ARA&A*, 31, 93
- Wood, K. S., Hertz, P., Norris, J. P., Vaughan, B. A., Michelson, P. F., Mitsuda, K., Lewin, W. H. G., van Paradijs, J., Penninx, W., & van der Klis, M. 1991, *ApJ*, 379, 295
- Zhang, W., Jahoda, K., Swank, J. H., Morgan, E. H., & Giles, A. B. 1995, *ApJ*, 449, 930