The Astronomical Journal, 124:1788–1809, 2002 September

FOURIER TECHNIQUES FOR VERY LONG ASTROPHYSICAL TIME-SERIES ANALYSIS

Scott M. Ransom,[1,2] Stephen S. Eikenberry,[3] and John Middleditch[4]
Received 2002 April 23; accepted 2002 June 11

ABSTRACT

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 of time-series data, including power-spectral statistics, discussing some of the differences between analyses of binned data, sampled data, and event data, and we 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.

Key words: methods: data analysis — pulsars: general

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 because of 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 that use this phase information exist and 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 one's 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 nonnegligible fraction of the total time duration (T) of the data ($T/dt \lesssim 10^7$). Gigapoint 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 discrete Fourier transform (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, as a result of orbital motion or pulsar spin-down.

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

[1] Department of Physics, McGill University, Montreal, QC H3A 2T8, Canada; ransom@physics.mcgill.ca.
[2] Center for Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139.
[3] Department of Astronomy, 212 Space Sciences Building, Cornell University, Ithaca, NY 14853.
[4] CCS-3, MS B256, Los Alamos National Laboratory, Los Alamos, NM 87545.

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).

2.1. *Introduction to the DFT*

The kth element of the discrete Fourier transform of a uniformly spaced time series n_j ($j = 0, 1, \ldots, N-1$) is defined as

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

where $i = (-1)^{1/2}$ and k is the Fourier frequency or wave number ($k = 0, 1, \ldots, N-1$). For a time spacing dt between successive data elements, the frequency of the kth element is $f_k = k/(N\,dt) = k/T$, where T is the total time duration of the sequence being transformed. This frequency spacing of $1/T$, for evenly sampled, unpadded, and nonwindowed data, is often called the independent Fourier spacing (IFS)[5] and defines the finest frequency resolution available while maintaining completely independent Fourier amplitudes. Fourier frequency $N/2T$ 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 $-2\pi k/N$ from the previous element. If the n_j have a constant value, the sum will form k regular N/k-sided polygons with each polygon returning near to the origin and with the last one returning exactly to the 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 2 (see, 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 $0 \le u \le 1$). In this case, instead of working with frequencies f, or integral wavenumbers k, we define r, which represents any real wavenumber (including nonintegers). If the number of samples, N, from our data source, $n(u)$, is large, we can compute a continuous Fourier transform (FT)

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

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 $0 \ll r \ll N/2$. We will use this approximation in many of our derivations.

2.2. *Computation of Very Long FFTs*

The fast Fourier transform is a family of well-known computer algorithms that 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 that have become generally known only recently and that 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 (e.g., 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).

2.2.1. *The Four-Step FFT*

Closely following the derivation described by 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 $x = 0, 1, \ldots, N_c - 1$, $y = 0, 1, \ldots, N_r - 1$, $l = 0, 1, \ldots, N_r - 1$, and $m = 0, 1, \ldots, N_c - 1$, we can write our signal and its DFT amplitudes as

$$n(x, y) = n_j , \quad j = N_c y + x , \quad (3)$$

$$A(l, m) = A_k , \quad k = N_r m + l . \quad (4)$$

Substituting into the definition of the DFT (eq. [1]) and simplifying, we obtain

$$A(l, m) = \sum_{x=0}^{N_c-1} \left[e^{-2\pi ixl/N} \sum_{y=0}^{N_r-1} n(x, y) e^{-2\pi iyl/N_r} \right] e^{-2\pi ixm/N_c} . \quad (5)$$

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

$$A_c(x, l) = e^{-2\pi ixl/N} \sum_{y=0}^{N_r-1} n(x, y) e^{-2\pi iyl/N_r} . \quad (6)$$

[5] The 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.)

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

$$A_r(m,l) = \sum_{x=0}^{N_c-1} A_c(x,l) e^{-2\pi i x m/N_c} . \quad (7)$$

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

$$A(l,m) = A_r^T(m,l) . \quad (8)$$

The four-step algorithm only needs small portions of the data in memory at any one time. Unfortunately, the short FFTs are strided in memory[6] instead of being stored contiguously. This results in a significant inefficiency in today's cache-based processors and requires internode 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).

2.2.2. *The Six-Step FFT*

If the initial data set is transposed from an $N_c \times N_r$ matrix into an $N_r \times N_c$ 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 that brings the data into normal order.

These transpose operations are relatively efficient on distributed memory parallel systems with fast internode communications, and they 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 do 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.

2.2.3. *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 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.

3. FOURIER TRANSFORMS OF REAL DATA

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

$$(c_j + d_j) e^{-2\pi i j k/N} = c_j e^{-2\pi i j k/N} + d_j e^{-2\pi i j k/N} . \quad (9)$$

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 jth 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 is[7]

$$\Pr(P \geq P') = e^{-P'} . \quad (10)$$

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

$$\Pr(P_m \geq P') = \sum_{j=0}^{m-1} \frac{(P')^j}{j!} e^{-P'} , \quad (11)$$

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 § 3.2.2).

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 that increases in strength toward the low-frequency part of the spectrum and is correspondingly called *red noise*.

Techniques to flatten or remove this "colored" 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

[6] "Strided" means that the data is stored in a sequence of noncontiguous memory locations spaced by a constant amount of memory known as the "stride."

[7] This is different from the probability for an actual signal to produce a power $P > P'$ in the presence of noise (see § 3.3). This difference is important in setting upper limits on the amplitudes of periodic signals, as discussed in Vaughan et al. (1994).

power level $P_{\rm local}$, such that

$$P_{k,\rm norm} = \frac{|A_k|^2}{N\langle d_j^2\rangle} \simeq \frac{|A_k|^2}{P_{\rm local}} = \frac{P_k}{P_{\rm local}} \ . \tag{12}$$

As long as the number of averaged powers is small enough 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 1 and an exponential distribution (eq. [10]).

Since strong narrowband signals near some frequency of interest will skew a local power average upward (and correspondingly decrease the calculated significance of a signal detection), it is important to exclude such powers from the calculation of $P_{\rm local}$. 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 $1/(\ln 2)$ times the local median value is theoretically equivalent to normalizing with the local mean, but it has the advantage of being insensitive to high-power outliers in the spectrum.

More advanced algorithms for the removal of "colored" noise and power normalization do exist. A simple example involves fitting polynomial models to portions of the power spectrum and then dividing the power spectrum by the models. 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 kth DFT element is

$$P_{k,\rm norm} = \frac{|A_k|^2}{N\langle d_j^2\rangle} = \frac{|A_k|^2}{n_{\rm ph}} \ , \tag{13}$$

where $n_{\rm ph} = N\langle n_j\rangle$ 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 § 2.1). However, the same processes that caused the "colored" noise discussed above can significantly alter this situation and require a power normalization based on local powers (see eq. [12]).

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.

3.2.1. *Sinusoidal Signals*

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

$$n_j = a\cos(2\pi f_r j\, dt + \phi) \tag{14a}$$

$$= a\cos(2\pi jr/N + \phi) \tag{14b}$$

$$= \tfrac{1}{2}a(e^{2\pi ijr/N + i\phi} + e^{-2\pi ijr/N - i\phi}) \ . \tag{14c}$$

From this expression, we see that the kth element of the DFT is given by

$$A_k = \frac{a}{2}\sum_{j=0}^{N-1} e^{-2\pi ijk/N}(e^{2\pi ijr/N + i\phi} + e^{-2\pi ijr/N - i\phi}) \tag{15a}$$

$$= \frac{a}{2}\sum_{j=0}^{N-1} e^{2\pi ij(k-r)/N + i\phi} + e^{-2\pi ij(k+r)/N - i\phi}, \tag{15b}$$

and represents the summation of two vectors in the complex plane. For $k \neq r$, the first term traces out $|k - r|$ complete "rotations" (pseudopolygons that start and end at the origin) in the complex plane (since $k - r$ is an integer), giving a net contribution of zero to the kth DFT element. The second term traces out $k + r$ complete rotations and once again contributes nothing to the kth 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 $a/2$ in a direction parallel to that 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 rth frequency bin, where the response is

$$A_r = \tfrac{1}{2}Nae^{i\phi} \ . \tag{16}$$

The Fourier response is more complicated for sinusoids with nonintegral frequencies (i.e., wavenumber r is a nonnegative real number). The kth DFT element is still given by equation (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 equation (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 equation (15b) results in a semicircular arc of length $Na/2$ *along the arc*, while the second term produces a semicircular arc of length $a(N \bmod k + [r])/2$ along the arc. The DFT response is simply a vector drawn from the origin to the end of the arc (see Fig. 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 $2\pi(k - r)$ rad of a circle of radius $Na/4\pi(k - r)$. The equation for the length of a chord is

$$C = 2\frac{s}{\Theta}\sin\left(\frac{\Theta}{2}\right) \ , \tag{17}$$

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,

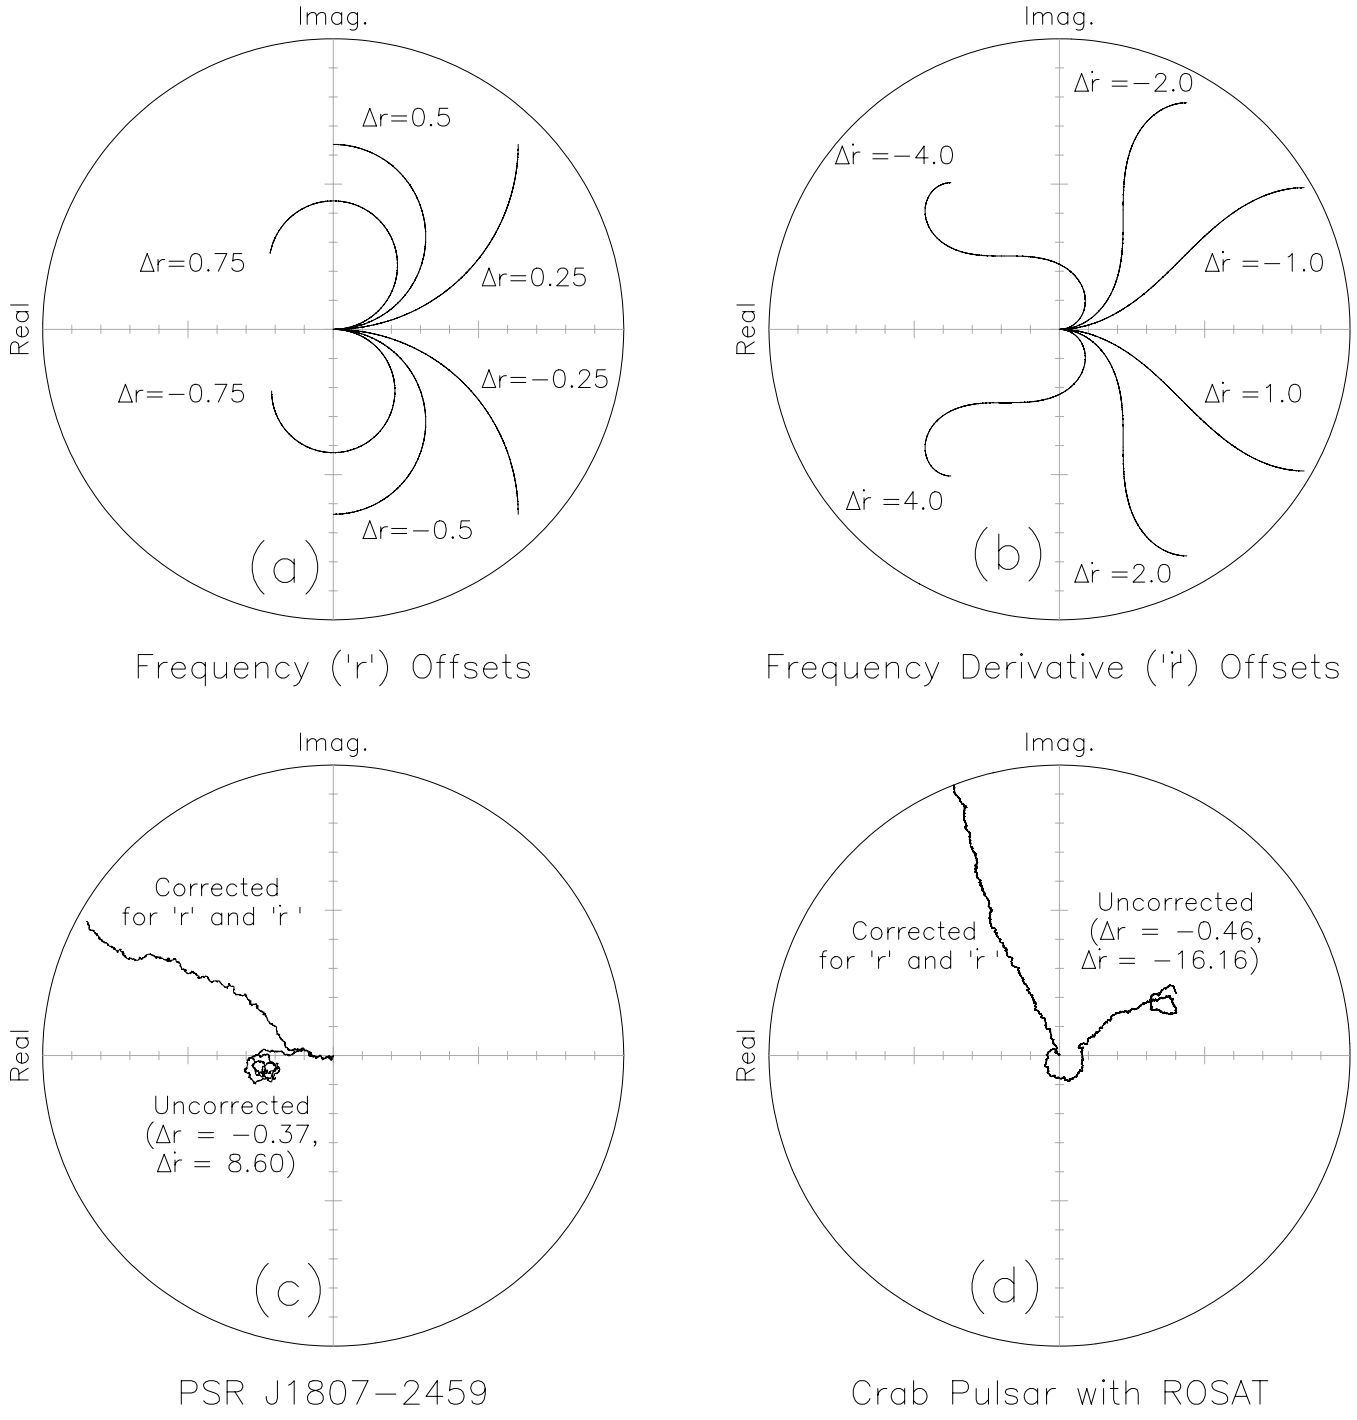

FIG. 1.—Fourier responses plotted as a series of vector additions in the complex plane. The outer circles in each plot show the Fourier amplitude of a signal where all power is recovered by the vector addition (i.e., calculation of the DFT at the correct signal frequency r and frequency derivative $\dot{r}$ for signals with linear changes in frequency over time). The endpoints of the vector additions are the Fourier amplitudes. For (*a*) and (*b*) a fully recovered signal would start at $0 + 0i$ and end at $1 + 0i$. (*a*) and (*b*) Effects on Fourier amplitude and phase when a signal's intrinsic frequency (r in bins or wavenumber) or frequency derivative ($\dot{r}$ in bins per observation) differs from the computed values. For (*b*) the average Fourier frequencies in each case were correct, and only the frequency derivatives were in error. (*c*) Response of PSR J1807−2459 during its discovery observation (Ransom et al. 2001) with and without corrections for pulsar acceleration ($\dot{r}$) and interpolation in Fourier frequency (r). The vectors were calculated using the method shown in § 4.2.3. The fact that the corrected vector does not quite reach the circle implies that higher order effects of the orbital motion remain uncorrected (see Fig. 6). (*d*) Corrected and uncorrected responses of 10,000 randomly selected photons from a 2.4 day *ROSAT* observation of the Crab pulsar.

the DFT response and power are

$$A_k = e^{i\phi}e^{-i\pi(k-r)}2\frac{Na/2}{2\pi(k-r)}\sin\left[\frac{2\pi(k-r)}{2}\right] \quad (18a)$$

$$= \frac{Na}{2}e^{i\phi}e^{-i\pi(k-r)}\frac{\sin[\pi(k-r)]}{\pi(k-r)} \quad (18b)$$

$$= A_0 e^{-i\pi(k-r)}\mathrm{sinc}\,[\pi(k-r)]\,, \quad (18c)$$

$$P_k = |A_k|^2 = P_0\,\mathrm{sinc}^2\,[\pi(k-r)]\,, \quad (18d)$$

where $A_0 = Na/2\ \mathrm{e}^{i\phi}$ is the DFT response for an integral frequency signal (eq. [16]), P_0 is the corresponding Fourier power, and the sinc function is defined as $\mathrm{sinc}\ x = \sin(x)/x$. This result is easily confirmed by a direct integration of equation (2), where $n(u)$ is equal to equation (14c) with $j/N \rightarrow u$.

The sinc factor in equation (18c) produces a loss of sensitivity for the standard FFT to most real-world signals (where r is not an integer). This effect, often called "scalloping" (e.g., Middleditch, Deich, & Kulkarni 1993), is shown in Figure 2, and it causes a worst-case (when $|k-r| = \frac{1}{2}$) amplitude reduction of $|A_k| = 2|A_0|/\pi$—nearly a 60% loss of signal power. On average, scalloping results in a ~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 themselves. In § 4.1 we discuss various methods to reduce or even eliminate this loss of sensitivity.

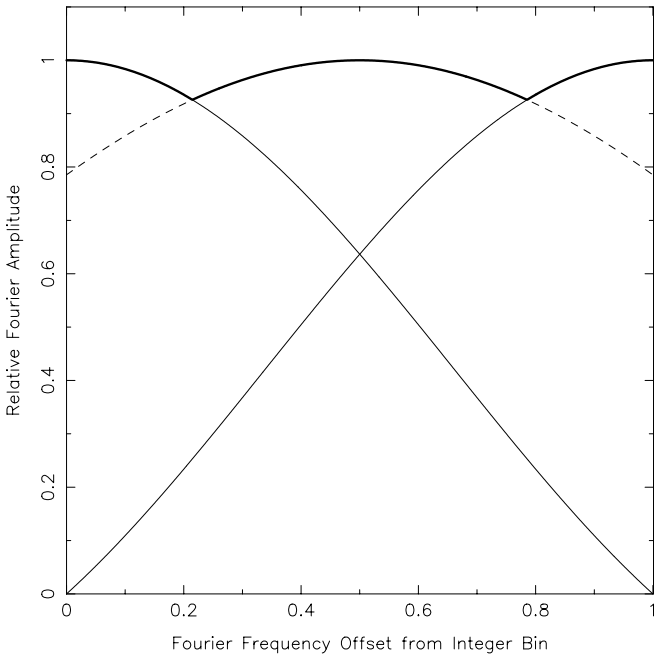

FIG. 2.—The sinc function amplitude responses of raw (or integer) FFT bins (*thin solid lines*). The point where these lines cross (at an offset of 1/2 and an amplitude of $1 - 2/\pi \sim 0.637$ of the full response) is the worst-case response for an uncorrected FFT (see eq. [18c]). The thin dashed line is the response of an "interbin" as calculated using eq. (31). The thick black line shows the overall "scalloping" when interbinning is used. Worst-case responses with interbinning occur at offsets of $\pm(1 - \pi/4) \sim 0.215$ with amplitudes of ~0.926 of the full response.

3.2.2. *Nonsinusoidal Signals*

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

$$n_j = \sum_{h=1}^{m} a_h \cos(2\pi jhr/N + \phi_h) \quad (19a)$$

$$= \sum_{h=1}^{m} \frac{a_h}{2}(e^{2\pi ijhr/N + i\phi_h} + e^{-2\pi ijhr/N - i\phi_h})\,, \quad (19b)$$

where $h = 1, 2, \ldots, m$ 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. Because of 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 equation (18c), except that A_0 becomes $A_h = Na_h/2\ \mathrm{e}^{i\phi_h}$).

For nearly sinusoidal pulsations only the first few terms of equation (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 with 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):

$$f(\kappa, t) = a\frac{e^{\kappa\cos(2\pi f_r t + \phi)} - e^{-\kappa}}{I_0(\kappa) - e^{-\kappa}}\,, \quad (20)$$

where $0 \le t \le T$ 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 $\kappa \rightarrow 0$, the MVMD becomes a sinusoid, while as $\kappa \rightarrow \infty$ it becomes a Gaussian (see Fig. 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 1). The full width at half-maximum (FWHM), as a fraction of a pulse, is

$$\mathrm{FWHM}_{\mathrm{MVMD}} = \pi^{-1}\arccos[\ln(\cosh\kappa)]\,, \quad (21)$$

and the maximum value is

$$\max_{\mathrm{MVMD}} = \frac{2a\cosh\kappa}{I_0(\kappa) - e^{-\kappa}}\,. \quad (22)$$

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

$$e^{\kappa\cos x} = I_0(\kappa) + 2\sum_{h=1}^{\infty} I_h(\kappa)\cos(hx)\,, \quad (23)$$

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

$$f(\kappa, t) = a + \frac{2a\sum_{h=1}^{\infty} I_h(\kappa)\cos(2\pi hf_r t + h\phi)}{I_0(\kappa) - e^{-\kappa}}\,. \quad (24)$$

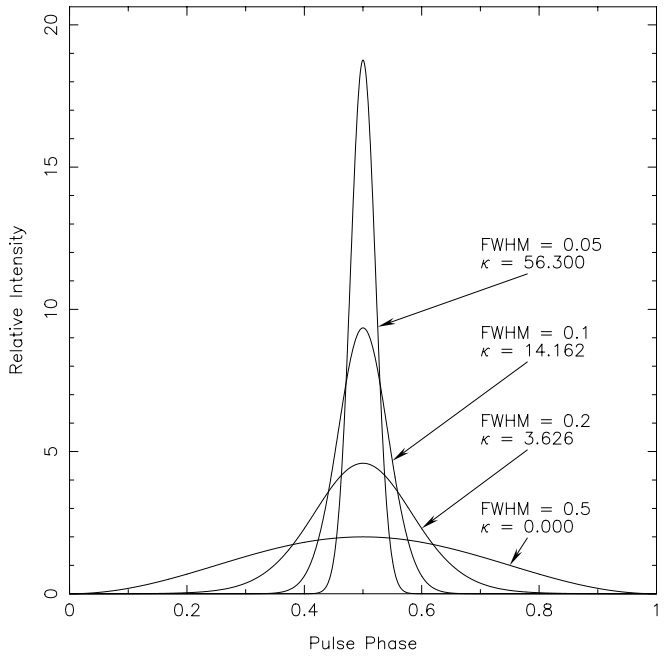

FIG. 3.—Sample pulse profiles from the modified von Mises distribution (MVMD) as described in § 3.2.2. FWHM is the fractional full width at half-maximum and κ is the MVMD shape parameter. High values of κ result in Gaussian profiles, while as $\kappa \rightarrow 0$, the pulse shape becomes more and more sinusoidal. The integral of a full pulse is equal to one unit, all of which is pulsed (i.e., the pulsed fraction is 1).

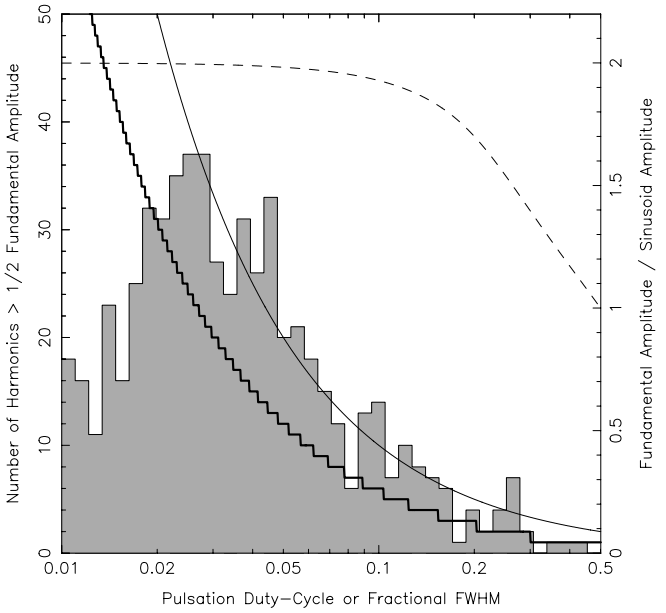

FIG. 4.—Approximate number of harmonics from an MVMD signal (*thick solid line*; see § 3.2.2 and Fig. 3) that produce Fourier amplitudes greater than one-half the amplitude of the fundamental. The thin solid line is the 1/FWHM rule of thumb that is often used to estimate the number of significant harmonics 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 gray histogram shows the distribution of pulse widths (FWHM) for over 600 radio pulsars.

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 eq. [2] with $f_r t \rightarrow ru$), we are left with a series of harmonics of amplitude $aN\, I_h(\kappa)/[I_0(\kappa) - e^{-\kappa}]$ 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 Fig. 4, *dashed line*). 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. 5).

Figure 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[8]). Most radio pulsars have duty cycles of $\lesssim 5\%$, corresponding to $\gtrsim 10$ significant harmonics—assuming a sufficient data sample rate.

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 § 3.1). Groth (1975) calculated the expectation value and variance of P_m as $\langle P_m \rangle = m + P_s$ and $\langle P_m^2 - \langle P_m \rangle^2 \rangle = m + 2P_s$, respectively. 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

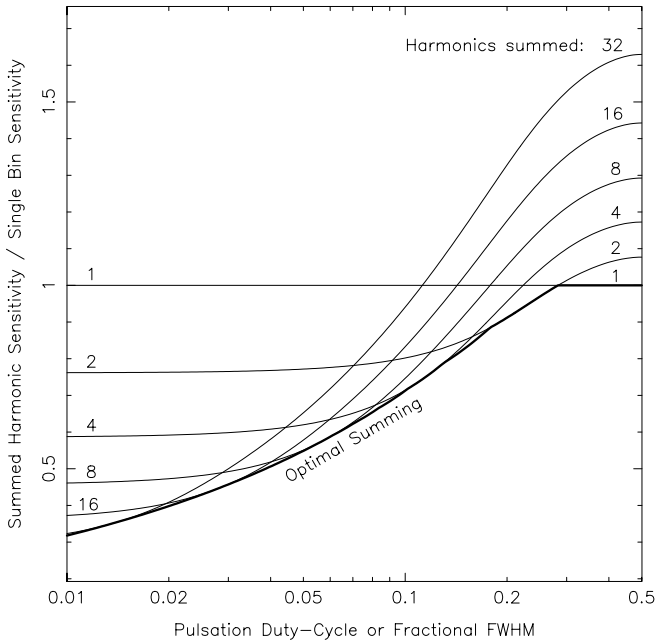

FIG. 5.—Sensitivities to MVMD signals (*thick solid line*; see § 3.2.2 and Fig. 3) for the incoherent summing of 1, 2, 4, 8, 16, or 32 harmonics, as compared to searches made without harmonic summing (§ 3.1). Lower numbers represent better sensitivity (i.e., fainter signals are detectable). The best 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. 4), so that summing tends to add only noise rather than signal.

[8] Available at http://pulsar.princeton.edu/pulsar/catalog.shtml.

equal to some power P',

$$\Pr\left[(P_m; P_s) \geq P'\right] = e^{-(P'+P_s)} \sum_{k=0}^{\infty} \sum_{j=0}^{k+m-1} \frac{(P')^j P_s^k}{j!k!} \ . \quad (25)$$

When $P_s = 0$, this equation reduces to equation (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 procedure[9] for correctly determining search sensitivities and upper limits using the equations of Groth (1975).

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 § 3.1). In addition to this difference in power-spectrum normalization, the other points worth noting comes 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, which 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 nonlinear 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.

3.4.1. *Binned versus 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 frequency dependent—resulting in decreased sensitivity at high frequencies (see, e.g., Middleditch 1976; Leahy et al. 1983).

The frequency-dependent loss in Fourier amplitude due to binning is sinc ($\pi f_r\, dt$), or sinc ($\pi r/N$). The binned-data Fourier response to a sinusoid is therefore equation (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).

3.4.2. *Low Count Rate Data*

The Fourier analysis of gamma-ray or X-ray observations often places us in a unique regime—very long integration times ($\gtrsim 10^4$ s) with very low numbers of counts ($\lesssim 10^3$ photons). In addition, because of 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 gigapoint 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 these data contain 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 1, an exact DFT amplitude at arbitrary Fourier frequency r becomes

$$A_r = \sum_{j=1}^{n_{\rm ph}} e^{-2\pi i r(t_j - t_0)/T} \ , \quad (26)$$

where $n_{\rm ph}$ is the number of photons, t_j is the TOA of the jth photon, t_0 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 "scalloping" (see § 3.2.1) is nonexistent.

Since equation (26) only involves a summation over the number of photons, it can be computed quickly if $n_{\rm ph}$ is relatively small. Great increases in computation speed can be had if we search a regular grid of Fourier frequencies. Trigonometric recurrences, such as

$$\cos(\theta + \delta) = \cos\theta - (\alpha\cos\theta + \beta\sin\theta) \ , \quad (27a)$$

$$\sin(\theta + \delta) = \sin\theta - (\alpha\sin\theta - \beta\cos\theta) \ , \quad (27b)$$

where $\alpha = 2\sin^2(\delta/2)$ 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.

4. IMPROVING THE DFT RESPONSE TO ARBITRARY SIGNALS

4.1. *Fourier Interpolation*

The potential for up to a 30% decrease in signal-to-noise ratio ($\mathrm{S/N} \propto \sqrt{P}$) 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 § 3.2.1). However, if we could calculate the 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.

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

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 § 3.4.2), 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 data.[10] 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) 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 noninteger frequencies—a process known as "fine-binning" or "Fourier interpolation" (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 $f_r = r/T$, where r is any real number, such that the result is sufficiently close to the exact calculation,

$$A_r = \sum_{j=0}^{N-1} n_j e^{-2\pi i j r/N} \ . \tag{28}$$

We can rewrite this expression as

$$A_r = \sum_{k=0}^{N-1} A_k e^{-i\pi(r-k)} \operatorname{sinc}\left[\pi(r-k)\right] \ , \tag{29}$$

where the A_k are the complex FFT amplitudes at the integer frequencies l (see Appendix A and § 4.2.2 for a derivation and discussion of this result).

The sinc function in equation (29) provides the key to computing an accurate interpolated amplitude using a relatively small number of operations. Since $\operatorname{sinc}\left[\pi(r-k)\right] \to 0$ as $\pi(r-k) \to \pm\infty$, 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

$$A_r \simeq \sum_{k=[r]-m/2}^{[r]+m/2} A_k e^{-i\pi(r-k)} \operatorname{sinc}\left[\pi(r-k)\right] \ , \tag{30}$$

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 equation (18c).

The top panel in Figure 6 shows the raw FFT power spectrum (*gray dots*) and the interpolated power spectrum (*gray line connecting the dots*) for a radio observation of the short-period binary pulsar PSR J1807−2459. The black line without dots, which traces a very significant peak in power, was calculated using an extension of the interpolation techniques that accommodates a signal with a nonzero but constant frequency derivative (see § 4.2.2). The power spectra cover a narrow frequency range near the pulsar's rotational frequency and were calculated using $m = 32$ and a frequency step size of $\Delta r = 1/16$ (compared with the raw FFT frequency step size of $\Delta r = 1$). The true signal amplitude, phase, and location (i.e., frequency) can be reconstructed using these techniques. We will see in § 5 how this information can be used to deduce further properties of the signal.

A computationally less expensive version of Fourier interpolation is "interbinning," where we approximate the FT response at half-integer frequencies using only the nearest two integer frequency bins. By using the Fourier interpolation equation (eq. [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 obtain

$$A_{k+(1/2)} \simeq \tfrac{1}{4}\pi(A_k - A_{k+1}) \ . \tag{31}$$

This particular formulation of interbinning was reported by van der Klis (1989), and its response is shown in Figure 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 ratio from $1 - 2/\pi$, or ~36% at a frequency offset of $\frac{1}{2}$ bin, to ~7.4% at an offset of $\pm(1 - \pi/4)$ 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 § 2.1 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 as a result of 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 equation (30) allows accurate estimates to be made of the signal's significance and other properties (see § 5).

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

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

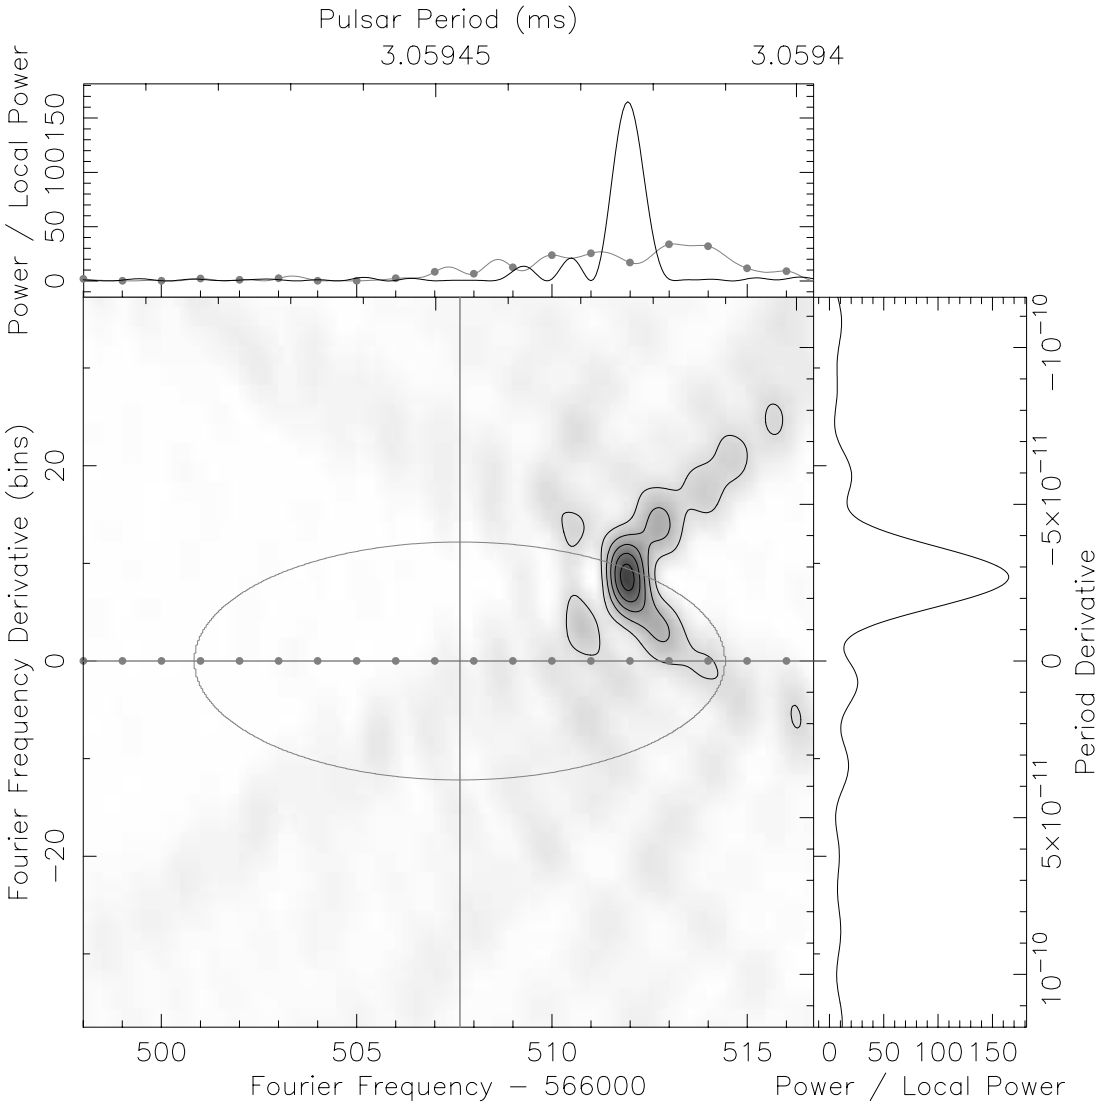

FIG. 6.—An 18 σ (single trial) detection from the discovery observation of the 1.7 hr binary PSR J1807−2459 in the globular cluster NGC 6544 using a Fourier-domain "acceleration" search. Contour intervals correspond to 30, 60, 90, 120, and 150 times the average local power level. The intrinsic pulsar period and $\dot{f} = 0$ (which corresponds to an unaccelerated FFT of the data) are marked by the solid gray lines. The dots correspond to the "raw" or uninterpolated 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 § 4.1).

vector addition in the complex plane (which we will call "vector bending").

4.2.1. *Correcting for Constant Frequency Derivative*

In order to illustrate these two methods, we demonstrate how to correct for a signal whose response is reduced because of a constant frequency derivative $\dot{f}$ (or in Fourier frequency bins, $\dot{r} = \dot{f}T^2$). The DFT operates, as we noted in § 3.2.1, 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 (see, e.g., Johnston & Kulkarni 1991). Figure 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 resampling the data ensemble n_j using a transformation similar to

$$t' = t + 2(f/\dot{f}_0)t^2 \ , \tag{32}$$

where t is the time used when sampling the original data, $\dot{f}$ is the frequency derivative, and f_0 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 nontrivial. 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.

4.2.2. *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 sinclike 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_0}, where $k - r_0$ is simply the frequency offset of bin k from some reference frequency r_0, which goes to zero as $|k - r_0|$ approaches some number of bins $m/2$. For Fourier interpolation as described in § 4.1, this response is equal to equation (18c) without the A_0 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_0 can be calculated with the sum

$$A_{r_0} \simeq \sum_{k=[r_0]-m/2}^{[r_0]+m/2} A_k A^*_{r_0-k} \ . \tag{33}$$

If r_0 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 equation (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 $m \ll M \ll N/2$, we can search a very long FFT of length $N/2$ for any signal whose A_{k-r_0} 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

$$n(u) = a\cos[2\pi(r_0 u + \tfrac{1}{2}\dot{r}u^2) + \phi] \tag{34a}$$

$$= \tfrac{1}{2}a[e^{2\pi i(r_0 u + \dot{r}u^2/2)}e^{i\phi} + e^{-2\pi i(r_0 u + \dot{r}u^2/2)}e^{-i\phi}] \ , \tag{34b}$$

where we use the same notation as § 2.1. Neglecting the second term as in § 3.2.1 and Fourier-transforming at some "center" or average frequency $r'_c = r + \dot{r}/2$, we obtain

$$A_{r'_c} = \frac{aN}{2}e^{i\phi}\int_0^1 e^{i\pi(\dot{r}u^2 + 2q_r u)}\,du \ , \tag{35}$$

where $q_r = r_c - r'_c$ and the real "center" frequency of the signal is $r_c = r_0 + \dot{r}/2$. This integral can be evaluated in closed form:

$$\int_0^1 e^{i\pi(\dot{r}u^2 + 2q_r u)}\,du$$

$$= \frac{1}{\sqrt{2\dot{r}}}e^{-i\pi q_r^2/\dot{r}}\{S(Z_r) - S(Y_r) + i[C(Y_r) - C(Z_r)]\} \ , \tag{36}$$

where $Y_r = (2/\dot{r})^{1/2}q_r$, $Z_r = (2/\dot{r})^{1/2}(q_r + \dot{r})$, and $C(x)$ and $S(x)$ are the Fresnel integrals

$$C(x) = \int_0^x \cos\left(\frac{\pi}{2}t^2\right)dt \ , \ S(x) = \int_0^x \sin\left(\frac{\pi}{2}t^2\right)dt \ . \tag{37}$$

The Fourier transform response then becomes

$$A_{r'_c} = \frac{aN}{2\sqrt{2\dot{r}}}e^{i(\phi - \pi q_r^2/\dot{r})}\{S(Z_r) - S(Y_r) + i[C(Y_r) - C(Z_r)]\} \ . \tag{38}$$

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 equation (38). This response, at average Fourier frequency r_c and Fourier frequency derivative $\dot{r}$, can therefore be written as

$$A_{r_c,\dot{r}} = \sum_{k=[r]-m/2}^{k=[r]+m/2} A_k \frac{1}{\sqrt{2\dot{r}}}$$

$$\times \left(e^{i\pi q_k^2/\dot{r}}\{S(Z_k) - S(Y_k) - i[C(Y_k) - C(Z_k)]\}\right) \ . \tag{39}$$

Equation (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 farther away—much like the nonzero wings of the sinclike 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. 2001). The search was conducted using an FFT of 13,865,600 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 hr, the second shortest radio pulsar orbital period known. A detailed view of the pulsar's fundamental harmonic is shown in Figure 6. The plot was calculated using the correlation technique with spacings of $\Delta r = 1/16$ and $\Delta\dot{r} = 1/4$. The generation of such a piece of the f-$\dot{f}$ plane takes only a fraction of a second on a rather modest workstation.

4.2.3. *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 § 2.1, 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 nonintegral or time-varying pulsation frequency, the standard DFT addition produces a curved shape (see Fig. 1). Since the amplitude of the Fourier response is the distance between the origin and the endpoint of the vector addition, any curvature in the vector addition implies nonoptimal 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 N/G points from the full-resolution vector addition. In terms of the rth DFT amplitude, we can write this as

$$A_r = \sum_{g=0}^{G-1} B_{r,g} = \sum_{g=0}^{G-1} \sum_{h=0}^{N/G-1} n_j e^{-2\pi i j r/N} , \tag{40}$$

where $j = gN/G + h$. 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 subvectors, 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 subvector 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

$$\Phi_{\rm true}(u) = 2\pi \int r(u)du \tag{41a}$$

$$= 2\pi \int (r_0 + \dot{r}u)du \tag{41b}$$

$$= 2\pi r_0 u + \pi \dot{r} u^2 , \tag{41c}$$

where $r_0 = f_0 T$ 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 $\Phi_{\rm DFT}(u) = 2\pi r u$ from the signal (see eq. [2]). So the instantaneous phase error is equal to

$$\Phi_{\rm error}(u) = \Phi_{\rm true}(u) - \Phi_{\rm DFT}(u) \tag{42a}$$

$$= 2\pi(r_0 - r)u + \pi \dot{r} u^2 . \tag{42b}$$

Therefore, to correct a particular signal using vector bending, we first calculate the $B_{r,g}$ using equation (40) for a particular Fourier frequency r (such as the frequency of a known pulsar). Now we attempt to unbend the full Fourier response by summing the $B_{r,g}$ after correcting for the phase errors $\Phi_{\rm error}(u_g)$ as defined by equation (42a). The corrected response is equal to

$$A_{r,\dot{r}} = \sum_{g=0}^{G-1} B_{r,g} e^{-i\Phi_{\rm error}(u_g)} . \tag{43}$$

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 tens of Fourier bins).

While impractical for large-scale searches because of 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 hr) on-source segments separated by hours, days, or even weeks of off-source time (see also § 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, Fazio, & Ransom 1998). Similarly, Figure 1 shows the use of the method for a 2.4 day observation of the Crab pulsar.

5. SIGNAL PROPERTY ESTIMATION

Besides correcting for losses of sensitivity, Fourier interpolation and the other response-correcting techniques mentioned in § 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, as well as 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. 6 and 7). 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).

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

$$P_{\rm meas} = \left|A_{r,\ldots}\right|^2 / P_{\rm norm} \ , \tag{44}$$

where $P_{\rm norm}$ is the expected noise power and is usually described by one of $N\langle d_j^2\rangle$, $P_{\rm local}$, or $n_{\rm ph}$, as discussed in § 3.1. Groth (1975; see § 3.3) showed that since the measured power is a random variable because of the presence of noise, its variance is $2P_{\rm signal} + 1$, where $P_{\rm signal}$ 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 $2P_{\rm meas} - 1$, since $\langle P_{\rm meas}\rangle = P_{\rm signal} + 1$.

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

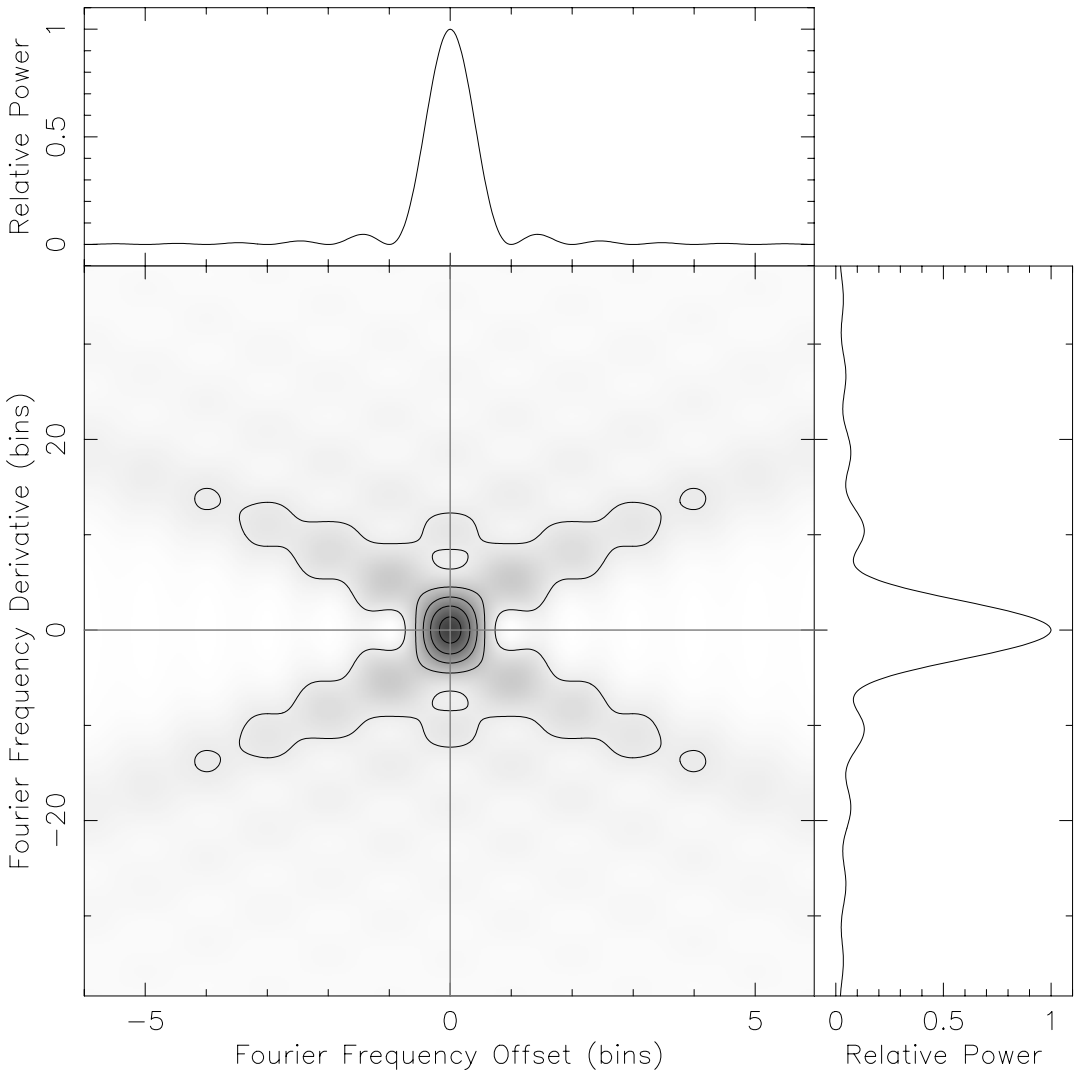

FIG. 7.—Theoretical response for a boxcar-windowed signal with a constant frequency derivative. Contours plotted are 10%, 30%, 50%, 70%, and 90% of peak (i.e., fully corrected) power. The top panel shows the familiar sinc response of the signal along the $\dot{f} = 0$ line. The right panel shows a similar 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.

§ 3.2.1), we can estimate the signal amplitude as

$$\langle a \rangle = 2N^{-1}\sqrt{P_{\rm norm}(P_{\rm meas}-1)}\ . \tag{45}$$

For binned data containing a signal with Fourier frequency r, the measured power should be multiplied by $1/\mathrm{sinc}^2(\pi r/N)$ to correct for the loss in sensitivity due to binning (see § 3.4.1). Vaughan et al. (1994) provide detailed instructions on how to estimate upper limits on pulsation amplitudes, as well as estimate the overall sensitivity of a search.

The statistical significance of a signal is also determined by $P_{\rm meas}$. The probability that noise caused a particular power to equal or exceed $P_{\rm meas}$ is given by $e^{-P_{\rm meas}}$ (eq. [10] with $P' = P_{\rm meas}$). But for a search over $N_{\rm IFS}$ independent Fourier powers, the probability that at least one of the noise powers exceeds $P_{\rm meas}$ is given by

$$\Pr\left(P_{\rm noise} \geq P_{\rm meas}\right) = 1 - \left(1 - e^{-P_{\rm meas}}\right)^{N_{\rm IFS}}\ . \tag{46}$$

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

$$\phi_{\rm meas} = \arctan\left[\frac{\mathrm{Im}\,(A_{r,\ldots})}{\mathrm{Re}\,(A_{r,\ldots})}\right]\ \mathrm{rad}\ . \tag{47}$$

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

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 equation (2), which in time-normalized units have the range $0 \leq u \leq 1$, or equivalently, $0 \leq t \leq T$. In effect, pulsations such as that defined in equation (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, $\hat{C} = \langle u \rangle = \langle t \rangle / T$, which is proportional to the first moment of the window function with respect to time. More specifically,

$$\hat{C} = \frac{\int_{-\infty}^{\infty} uW(u)du}{\int_{-\infty}^{\infty} W(u)du}\ . \tag{48}$$

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

$$\hat{C} = -\frac{1}{2\pi}\frac{\partial\phi(r_0)}{\partial r} \pm \frac{1}{\sqrt{24P(r_0)}}\ , \tag{49}$$

where $P(r_0)$ and $\phi(r_0)$ is the phase measured at the peak of the Fourier response (i.e., $\phi_{\rm meas}$) and r is the Fourier frequency (see Appendix D for a derivation). Signals present throughout an observation have $\hat{C} = \frac{1}{2}$, while those present in only the first or second half of the observation have $\hat{C} = \frac{1}{4}$ or $\hat{C} = \frac{3}{4}$, 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 rms dispersion of the pulsations in time about the centroid. Using this information, Middleditch & Córdova (1982) defined a parameter called the "purity" (and symbolized by α) as

$$\alpha = \frac{1}{\pi}\sqrt{-\frac{3}{2P(r_0)}\frac{\partial^2 P(r_0)}{\partial r^2}} \pm \frac{1}{\alpha\sqrt{10P(r_0)}}\ , \tag{50}$$

where $P(r_0)$ is the measured power at the peak of the Fourier response (i.e., $P_{\rm meas}$, see Appendix E for a derivation). The scaling in equation (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 $\alpha = \frac{1}{2}$. 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 equation (18c) changes if a signal is present during only part of an observation. In Appendix F we show that, when close to the peak of a signal's Fourier response,

$$A_r \simeq A_0 e^{-2\pi i\hat{C}(r-r_0)}\,\mathrm{sinc}\left[\pi\alpha(r-r_0)\right]\ , \tag{51}$$

where $A_0 = Na/2\ e^{i\phi_0}$, ϕ_0 is the intrinsic phase of the signal, and r_0 is the true signal frequency (in FFT bins). This equation demonstrates that for centroids different from $\frac{1}{2}$, the phase shift between consecutive FFT bins differs from the π radians shown in equation (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).

5.3. *Pulsation Frequency and Frequency Derivative*

The true pulsation frequency of the signal is located at the point where $\partial P/\partial r = 0$. Furthermore, given the response in equation (51), we can show (Appendix B) that the uncertainty in this measurement (in Fourier bins) is given by

$$\sigma_r = \frac{3}{\pi\alpha\sqrt{6P_{\rm meas}}}\ . \tag{52}$$

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 6 and 7, we can calculate the uncertainty in the measured $\dot{r}$-value by using similar arguments and methods as for the frequency uncertainty (see Appendix C for a derivation). The uncertainty in the $\dot{r}$ (in Fourier bins) is approximately

$$\sigma_{\dot{r}} = \frac{3\sqrt{10}}{\pi\alpha^2\sqrt{P_{\rm meas}}} \ . \tag{53}$$

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 § 4.2) and when characterizing signals that are known to be present in the data (using the parameters described in § 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. Gigapoint 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.

We would like to thank G. Fazio for supporting our research and encouraging us in this work. Additional thanks go to J. Grindlay, B. Schmidt, F. Seward, R. Narayan, V. Kaspi, and L. Greenhill for their encouragement and comments. Many of the computations for this paper were performed on equipment purchased with NSF grant PHY 95-07695. S. M. R. acknowledges the support of a Tomlinson Fellowship awarded by McGill University. S. S. E. is supported in part by an NSF CAREER Grant.

APPENDIX A

DERIVATION OF FOURIER INTERPOLATION

Following the derivation found in Middleditch et al. (1993), we begin with the definition of the kth DFT element

$$A_k = \sum_{j=0}^{N-1} n_j e^{-2\pi ijk/N} \ , \tag{A1}$$

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

$$A_k = \sum_{j=0}^{N-1} \left(\frac{1}{N} \sum_{l=0}^{N-1} A_l e^{2\pi ijl/N} \right) e^{-2\pi ijk/N} \tag{A2a}$$

$$= \frac{1}{N} \sum_{l=0}^{N-1} A_l \sum_{j=0}^{N-1} e^{-2\pi ij(k-l)/N} \ . \tag{A2b}$$

The last summation can be computed exactly using the identity

$$\sum_{j=0}^{N-1} e^{i\alpha j} = e^{i\alpha(N-1)/2} \frac{\sin(N\alpha/2)}{\sin(\alpha/2)} \ , \tag{A3}$$

such that, when $N \gg 1$, we have

$$\sum_{j=0}^{N-1} e^{-2\pi ij(k-l)/N} = e^{-i\pi(k-l)(1-N^{-1})} \frac{\sin[\pi(k-l)]}{\sin[\pi(k-l)/N]} \tag{A4a}$$

$$\simeq e^{-i\pi(k-l)} \frac{\sin[\pi(k-l)]}{\pi(k-l)/N} \tag{A4b}$$

$$\simeq N e^{-i\pi(k-l)} \, \mathrm{sinc}\,[\pi(k-l)] \ . \tag{A4c}$$

Substituting this expression into equation (A2b) and changing the integer frequency k into a continuous real-valued frequency

r, we arrive at equation (29):

$$A_r = \sum_{l=0}^{N-1} A_l e^{-i\pi(r-l)} \operatorname{sinc}\left[\pi(r-l)\right] . \tag{A5}$$

APPENDIX B

DERIVATION OF FREQUENCY UNCERTAINTY

For a given Fourier frequency offset $\Delta_r = r - r_0$, where r_0 is the Fourier frequency of the signal, the magnitude of the Fourier response and the power vary as

$$|A(r)| = |A_0| \operatorname{sinc}\left(\pi\alpha\,\Delta_r\right) , \tag{B1a}$$

$$P(r) = P_{\rm meas} \operatorname{sinc}^2\left(\pi\alpha\,\Delta_r\right) , \tag{B1b}$$

(see Appendix F), 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

$$P(r) = P_{\rm meas}\left[\frac{\sin(\pi\alpha\,\Delta_r)}{\pi\alpha\,\Delta_r}\right]^2 \simeq P_{\rm meas}\left[1 - \frac{(\pi\alpha\,\Delta_r)^2}{3}\right] . \tag{B2}$$

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

$$\Delta_r = -\frac{3}{2\pi^2\alpha^2 P_{\rm meas}}\frac{\partial P}{\partial r} . \tag{B3}$$

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

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

$$\sigma_r = \frac{3}{2\pi^2\alpha^2 P_{\rm meas}}\sigma_{P'(r)} , \tag{B4}$$

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

$$P'(r) \simeq \frac{P(r_0 + \Delta_r) - P(r_0 - \Delta_r)}{2\,\Delta_r} = \frac{P^+ - P^-}{2\,\Delta_r} , \tag{B5}$$

where we have simply renamed $P(r_0 + \Delta_r)$ and $P(r_0 - \Delta_r)$. The uncertainty in $P'(r)$ can also be approximated using finite differences and error propagation. Since P^+ and P^- are highly correlated when $\Delta_r \ll 1$, their uncertainties are also correlated, giving

$$\sigma_{P'(r)} \simeq \frac{1}{2\,\Delta_r}\left(\sigma_{P^+} + \sigma_{P^-}\right) \simeq \frac{\sigma_{P^+}}{\Delta_r} . \tag{B6}$$

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

$$\sqrt{P_{\rm meas}} = \sum_{j=0}^{N-1} y_j \cos(\phi_j) , \tag{B7}$$

where $y_j = n_j/P_{\rm norm}^{1/2}$ are the points in our time series as defined in equation (1), but scaled using the appropriate $P_{\rm norm}$ such that the measured power, $P_{\rm meas}$, is properly normalized (see § 3.1). Similarly, the ϕ_j represent the pulsation phases at times $u = j/N$, but rotated by the measured Fourier phase $\phi_{\rm meas}$, 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

$$\sqrt{P^+} = \sum_{j=0}^{N-1} y_j \cos(\phi_j + \delta_{\phi_j}) , \tag{B8}$$

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

$$\delta_{\phi_j} = 2\pi\alpha\,\Delta_r\left(\frac{j}{N} - \frac{1}{2}\right) = 2\pi\alpha\,\Delta_r\left(u - \frac{1}{2}\right)\,, \tag{B9}$$

where u is the normalized time, $u = t/T = j/N$. The $1/2$ term in equation (B9) removes the accumulated phase error over the course of the observation (i.e., $\int_0^1 2\pi\alpha\Delta_r u\,du = \pi\alpha\,\Delta_r$) and makes the vector summation of equation (B7) finish on the real axis. Expanding the cosine in equation (B8) gives

$$\sqrt{P^+} = \sum_{j=0}^{N-1} y_j \cos(\phi_j)\cos(\delta_{\phi_j}) - \sum_{j=0}^{N-1} y_j \sin(\phi_j)\sin(\delta_{\phi_j})\,. \tag{B10}$$

Considering the uncertainties in the separate terms of equation (B10), since the cosine term is an even function of δ_{ϕ_j}, it is symmetric about r_0 and therefore does not contribute to the uncertainty in the $P'(r)$ measurement as defined by equation (B5). For the sine term, given that the cos (ϕ_j) derotates the signal onto the real axis by definition, we see that

$$\sum_{j=0}^{N-1} y_j \sin(\phi_j) = 0\,. \tag{B11}$$

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

$$\left\langle \sum_{j=0}^{N-1} y_j \sin(\phi_j)\sin(\delta_{\phi_j}) \right\rangle = 0\,, \tag{B12}$$

and has no systematic effect on P^+. However, we can calculate the fluctuations introduced by this term

$$\left\langle \left(\sum_{j=0}^{N-1} y_j \sin(\phi_j)\sin(\delta_{\phi_j}) \right)^2 \right\rangle = \left\langle \left[\sum_{j=0}^{N-1} y_j \sin(\phi_j) \right]^2 \right\rangle \langle \sin^2(\delta_{\phi_j}) \rangle\,, \tag{B13}$$

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

$$\langle \sin^2 \delta_{\phi_j} \rangle = \int_0^1 \sin^2(\delta_{\phi_j})\,du \tag{B14a}$$

$$\simeq \int_0^1 \delta_{\phi_j}^2\,du \tag{B14b}$$

$$\simeq \int_0^1 (2\pi\alpha\,\Delta_r)^2 (u^2 - u + \tfrac{1}{4})du \tag{B14c}$$

$$\simeq \tfrac{1}{3}(\pi\alpha\,\Delta_r)^2\,. \tag{B14d}$$

Therefore equation (B13) is equal to $(\pi\alpha\,\Delta_r)^2/6$, and the variance of $\sqrt{P^+}$ will be

$$\sigma^2_{\sqrt{P^+}} = \sqrt{P^+}(\pi\alpha\,\Delta_r)^2/6\,. \tag{B15}$$

To get the standard deviation, we take the square root of this expression. Propagating errors to get the uncertainty in P^+ adds a factor of $2\sqrt{P^+}$ to give

$$\sigma_{P(r_0+\Delta_r)} = 2\pi\alpha\,\Delta_r\sqrt{P_{\rm meas}/6}\,, \tag{B16}$$

where we have used the approximation $P^+ \simeq P_{\rm meas}$. Finally, substituting into equations (B6) and (B4), we have

$$\sigma_r = \frac{3}{\pi\alpha\sqrt{6P_{\rm meas}}}\,. \tag{B17}$$

A much simpler derivation of equation (B17) is possible if we realize that properly normalized powers times 2 are distributed according to a χ^2 distribution with 2 degrees of freedom (see § 3.1). For such a distribution a 1 σ error corresponds to a change in the measured χ^2 of $\frac{1}{2}$. In the case of powers the ± 1 σ errors can therefore be found by starting with the expansion of the power around $P(r_0) = P_{\rm meas}$ as given in equation (B2),

$$P_{\rm meas}[1 - \tfrac{1}{3}(\pi\alpha\,\Delta_r)^2] = P_{\rm meas} - \tfrac{1}{2}\,, \tag{B18}$$

and then solving for the Δ_r that corresponds to $P_{\rm meas} - \frac{1}{2}$. This yields

$$\Delta_r = \sigma_r = \frac{3}{\pi\alpha\sqrt{6P_{\rm meas}}}\,. \tag{B19}$$

APPENDIX C

DERIVATION OF $\dot{f}$ UNCERTAINTY

If a peak in the f-$\dot{f}$ plane has been isolated using techniques similar to those shown in § 4.2.2, we can calculate the error in the measurement of the true Fourier frequency derivative $\dot{r}_0 = \dot{f}_0 T^2$ in a manner similar to that for the frequency uncertainty as described in Appendix B. Signals with nonzero frequency derivatives have Fourier peaks that 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 F). The shape of the response in power as a function of $\Delta_{\dot{r}} = \dot{r} - \dot{r}_0$ at the "correct" Fourier frequency $r = r_0$ is described by equation (38) and can be written as

$$A(r_0, \dot{r}) = A_0 e^{-i\pi\Delta_{\dot{r}}/4} \sqrt{\frac{2}{\alpha^2 \Delta_{\dot{r}}}} \left[S\left(\sqrt{\frac{\alpha^2 \Delta_{\dot{r}}}{2}} \right) - iC\left(\sqrt{\frac{\alpha^2 \Delta_{\dot{r}}}{2}} \right) \right] , \tag{C1}$$

when q_r is defined as

$$q_r = r_c - r_c' = \left(r_0 + \frac{\dot{r}_0}{2} \right) - \left(r_0 + \frac{\dot{r}_0 + \alpha^2 \Delta_{\dot{r}}}{2} \right) = -\frac{\alpha \Delta_{\dot{r}}}{2} . \tag{C2}$$

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

$$P(r_0, \dot{r}) = P_0 \frac{2}{\alpha^2 \Delta_{\dot{r}}} \left\{ \left[S\left(\sqrt{\frac{\alpha^2 \Delta_{\dot{r}}}{2}} \right) \right]^2 + \left[C\left(\sqrt{\frac{\alpha^2 \Delta_{\dot{r}}}{2}} \right) \right]^2 \right\} . \tag{C3}$$

If we expand the Fresnel integrals about 0 as

$$C(x) \simeq x - \frac{\pi^2}{40} x^5 \ldots , \quad S(x) \simeq \frac{\pi}{6} x^3 \ldots \tag{C4}$$

and then substitute, the power becomes

$$P(r_0, \dot{r}) \simeq P_0 \left[1 - \frac{(\pi \alpha^2 \Delta_{\dot{r}})^2}{180} \right] . \tag{C5}$$

Taking the derivative of $P(r_0, \dot{r})$ with respect to $\dot{r}$ and solving for $\Delta_{\dot{r}}$ gives

$$\Delta_{\dot{r}} = -\frac{90}{\pi^2 \alpha^4 P_0} \frac{\partial P(r_0, \dot{r})}{\partial \dot{r}} \tag{C6}$$

From propagation of errors,

$$\sigma_{\dot{r}} = -\frac{90}{\pi^2 \alpha^4 P_0} \sigma_{P'(r_0, \dot{r})} , \tag{C7}$$

where $P'(r_0, \dot{r}) = \partial P(r_0, \dot{r})/\partial \dot{r}$. Similarly, after using a finite difference estimate of the power derivative following Appendix B, we find

$$\sigma_{P'(r_0, \dot{r})} = \sigma_{P(r_0, \Delta_{\dot{r}})}/\Delta_{\dot{r}} , \tag{C8}$$

where $P(r_0, \Delta_{\dot{r}})$ represents the power as measured at $\dot{r} = \dot{r}_0 + \Delta_{\dot{r}}$.

Closely following Appendix B, we represent $[P(r_0, \Delta_{\dot{r}})]^{1/2}$ as a sum of the parallel components of the properly normalized time-series points y_j, as

$$\sqrt{P(r_0, \Delta_{\dot{r}})} = \sum_{j=0}^{N-1} y_j \cos(\phi_j + \delta_{\phi_j}) . \tag{C9}$$

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

$$\delta_{\phi_j} = \pi \alpha^2 \Delta_{\dot{r}} (u^2 - u + \tfrac{1}{6}) , \tag{C10}$$

where the u-term comes from keeping the response symmetric about r_0 (i.e., $q_r = -\alpha^2 \Delta_{\dot{r}}/2$) and the $1/6$ removes the accumulated phase error over the course of the observation (i.e., $\int_0^1 \pi \alpha^2 \Delta_{\dot{r}} (u^2 - u) du = -\pi \alpha^2 \Delta_{\dot{r}}/6$) and makes the vector summation of equation (C9) finish on the real axis.

Expanding the cosine term of equation (C9) gives

$$\sqrt{P(r_0, \Delta_{\dot{r}})} = \sum_{j=0}^{N-1} y_j \cos(\phi_j) \cos(\delta_{\phi_j}) - \sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}) . \tag{C11}$$

The first term shortens both $P(r_0, \Delta_{\dot{r}})$ and $P(r_0, -\Delta_{\dot{r}})$ 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 B), but its fluctuations are important. To calculate the fluctuations, we square the terms and obtain

$$\left\langle \left[\sum_{j=0}^{N-1} y_j \sin(\phi_j) \sin(\delta_{\phi_j}) \right]^2 \right\rangle = \left\langle \left[\sum_{j=0}^{N-1} y_j \sin(\phi_j) \right]^2 \right\rangle \langle \sin^2(\delta_{\phi_j}) \rangle \ . \tag{C12}$$

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

$$\langle \sin^2(\delta_{\phi_j}) \rangle = \int_0^1 \sin^2(\delta_{\phi_j})\, du \tag{C13a}$$

$$\simeq \int_0^1 \delta_{\phi_j}^2\, du \tag{C13b}$$

$$= \pi^2 \alpha^4 \Delta_{\dot{r}}^2 \int_0^1 (u^2 - u + \tfrac{1}{6})^2 du \tag{C13c}$$

$$= \frac{\pi^2 \alpha^4 \Delta_{\dot{r}}^2}{180} \ . \tag{C13d}$$

The fluctuations from the sine term are therefore $(\pi\alpha^2\Delta_{\dot{r}})^2/360$, and, since squaring equation (C9) doubles the errors, the standard deviation of $P(r_0, \Delta_{\dot{r}})$ is

$$\sigma_{P(r_0,\Delta_{\dot{r}})} = \sqrt{P(r_0, \Delta_{\dot{r}})} \frac{2\pi\alpha^2\Delta_{\dot{r}}}{\sqrt{360}} \simeq \sqrt{P_{\rm meas}} \frac{\pi\alpha^2\Delta_{\dot{r}}}{\sqrt{90}} \ . \tag{C14}$$

Substituting into equations (C8) and then (C7) as in Appendix B gives us the uncertainty in the frequency derivative,

$$\sigma_{\dot{r}} = \frac{3\sqrt{10}}{\pi\alpha^2\sqrt{P_{\rm meas}}} \ . \tag{C15}$$

APPENDIX D

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 § 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

$$s(u) = a \cos(2\pi r_0 u + \phi_0) W(u) \ , \tag{D1}$$

where a is the amplitude, $r_0 = f_0 T$ is the Fourier frequency, and ϕ_0 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

$$\langle u \rangle = -\frac{A'(0)}{2\pi i A(0)} \ , \tag{D2}$$

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

Equation (D2) is not directly applicable for our sinusoidal signal, since the information about the window function in equation (D1) 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 equation (D2), which gives us

$$\langle u \rangle_{s(u)} = -\frac{A'(r_0)}{2\pi i A(r_0)} \ . \tag{D3}$$

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

$$A'(r) = \frac{\partial a(r)}{\partial r} e^{i\phi(r)} + a(r) i \frac{\partial \phi(r)}{\partial r} e^{i\phi(r)} \ . \tag{D4}$$

At the frequency of our signal, the amplitude is $A(r_0) = a(r_0)e^{i\phi(r_0)}$ and the Fourier response is at its peak, making

$\partial a(r_0)/\partial r = 0$. Therefore

$$A'(r_0) = A(r_0) i \frac{\partial \phi(r_0)}{\partial r} \ . \tag{D5}$$

Substituting equation (D5) into equation (D3), we arrive at

$$\hat{C} = \langle u \rangle_{s(u)} = -\frac{1}{2\pi} \frac{\partial \phi(r_0)}{\partial r} \ , \tag{D6}$$

which is equivalent to equation (49).

We can also estimate the uncertainty on the measured value of the centroid. Following Appendix B from equations (B8)–(B15), we note that the same noise fluctuations, introduced when offsetting from the true pulsation frequency by a small amount Δ_r, that effect the Fourier amplitude will also effect the Fourier phases. When Δ_r is small and the powers are properly normalized, the amplitude fluctuations of equation (B13) with variance $\pi^2 \Delta_r^2/6$ correspond to an uncertainty in the phase measurement of $\sigma_{\phi(r_0+\Delta_r)} \simeq \pi \Delta_r / [6P(r_0)]^{1/2}$ radians, and therefore

$$\sigma_{\hat{C}} = -\frac{1}{2\pi} \frac{\sigma_{\phi(r_0+\Delta_r)}}{\Delta_r} \simeq \frac{1}{\sqrt{24 P(r_0)}} . \tag{D7}$$

APPENDIX E

DERIVATION OF PURITY

The "purity" of a signal (see § 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 equation (D1). The time variance is defined as

$$\langle (u - \langle u \rangle)^2 \rangle = \frac{\int_{-\infty}^{\infty} (u - \langle u \rangle)^2 W(u) du}{\int_{-\infty}^{\infty} W(u) du} \tag{E1}$$

but can be written as

$$\langle (u - \langle u \rangle)^2 \rangle = \langle u^2 \rangle - \langle u \rangle^2 \tag{E2a}$$

$$= -\frac{A''(0)}{4\pi^2 A(0)} + \frac{[A'(0)]^2}{4\pi^2 [A(0)]^2} \ , \tag{E2b}$$

using the moment theorem for Fourier transforms (see, e.g., Bracewell 1999). Since our signal is sinusoidal (see Appendix D), application of the modulation theorem gives

$$\langle (u - \langle u \rangle)^2 \rangle_{s(u)} = -\frac{A''(r_0)}{4\pi^2 A(r_0)} + \frac{[A'(r_0)]^2}{4\pi^2 [A(r_0)]^2} \tag{E3a}$$

$$= -\frac{A''(r_0)}{4\pi^2 A(r_0)} - \hat{C}^2 \ . \tag{E3b}$$

Using $A(r) = a(r) e^{i\phi(r)}$ and remembering that $\partial a(r_0)/\partial r = 0$, the second derivative of the Fourier amplitude is

$$A''(r_0) = iA(r_0) \frac{\partial^2 \phi(r_0)}{\partial r^2} - A(r_0) \left[\frac{\partial \phi(r_0)}{\partial r} \right]^2 + \frac{A(r_0)}{a(r_0)} \frac{\partial^2 a(r_0)}{\partial r^2} \ . \tag{E4}$$

From Appendix D, we know that $\partial \phi(r_0)/\partial r = -2\pi \hat{C}$, which makes the first term of equation (E4) equal to zero, since $\partial^2 \phi(r_0)/\partial r^2 = 0$, and the second term equal to $-4\pi^2 A(r_0) \hat{C}^2$. The second derivative of power at the peak response can be written as

$$\frac{\partial^2 P(r_0)}{\partial r^2} = \frac{\partial^2}{\partial r^2} [A^*(r_0) A(r_0)] = 2a(r_0) \frac{\partial^2 a(r_0)}{\partial r^2} \ , \tag{E5}$$

making the third term equal to $A(r_0)/2[a(r_0)]^2 \partial^2 P(r_0)/\partial r^2$. Substituting into equation (E3a) and simplifying gives

$$\langle (u - \langle u \rangle)^2 \rangle_{s(u)} = -\frac{1}{8\pi^2 P(r_0)} \frac{\partial^2 P(r_0)}{\partial r^2} \ . \tag{E6}$$

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

$$\langle (u - \langle u \rangle)^2 \rangle_{\text{pure}} = \langle u^2 \rangle - \langle u \rangle^2 = \tfrac{1}{3} - (\tfrac{1}{2})^2 = \tfrac{1}{12} \ , \tag{E7}$$

and then take the square root, we are left with

$$\alpha = \sqrt{\frac{\langle (u - \langle u \rangle)^2 \rangle_{s(u)}}{\langle (u - \langle u \rangle)^2 \rangle_{\text{pure}}}} = \frac{1}{\pi} \sqrt{-\frac{3}{2P(r_0)} \frac{\partial^2 P(r_0)}{\partial r^2}} , \tag{E8}$$

which is equivalent to equation (50).

In order to estimate the uncertainty in the measurement of α, we note that, by squaring equation (E8) and substituting the finite difference approximation of

$$\frac{\partial^2 P(r_0)}{\partial r^2} \simeq \frac{P(r_0 + \Delta_r) + P(r_0 - \Delta_r) - 2P(r_0)}{\Delta_r^2} = \frac{P^+ + P^- - 2P(r_0)}{\Delta_r^2} , \tag{E9}$$

where Δ_r corresponds to a small frequency offset from the measured peak power $P(r_0)$, we obtain

$$\alpha^2 \simeq -\frac{3}{2\pi^2 P(r_0)} \frac{P^+ + P^- - 2P(r_0)}{\Delta_r^2} . \tag{E10}$$

This means that

$$\sigma_\alpha \simeq \frac{3}{4\pi^2 \alpha \, \Delta_r^2 P(r_0)} \sigma_{\Delta_P} , \tag{E11}$$

where $\Delta_P = P^+ + P^- - 2P(r_0)$ and the extra factor of 2α comes from converting σ_{α^2} to σ_α.

We can then expand the Fourier amplitude around the peak of the signal as in equation (B10), where we see that the amplitude fluctuations from the sine term (which is antisymmetric about r_0) will cancel after the addition of $P(r_0 + \Delta_r)$ and $P(r_0 - \Delta_r)$ in the finite difference approximation shown above. Conversely, the fluctuations due to the cosine term (which is symmetric about r_0) will add. These fluctuations can be computed by taking the variance of the first nonconstant term in the expansion of cos (δ_j), which is $\delta_j^2/2 = 2\pi^2 \Delta_r^2 \left(u - \frac{1}{2}\right)^2$. The computation gives

$$\langle (\delta_j^2/2 - \langle \delta_j^2/2 \rangle)^2 \rangle = \langle (\delta_j^2/2)^2 \rangle - \langle \delta_j^2/2 \rangle^2 \tag{E12a}$$

$$= 4\pi^4 \Delta_r^4 \left\{ \int_0^1 \left(u - \frac{1}{2}\right)^4 du - \left[\int_0^1 \left(u - \frac{1}{2}\right)^2 du \right]^2 \right\} \tag{E12b}$$

$$= \pi^4 \Delta_r^4 \left[\frac{1}{20} - \left(\frac{1}{6}\right)^2 \right] \tag{E12c}$$

$$= \frac{\pi^4 \Delta_r^4}{45} , \tag{E12d}$$

and since the variations in the cos (ϕ_j) term average to 1/2, just like the sin (ϕ_j) term did in equation (B13), the standard deviation of $[P(r_0 + \Delta_r)]^{1/2}$ is equal to

$$\sigma_{\sqrt{P(r_0+\Delta_r)}} = \sqrt{P(r_0 + \Delta_r)} \frac{\pi^2 \Delta_r^2}{\sqrt{90}} \simeq \pi^2 \Delta_r^2 \sqrt{\frac{P(r_0)}{90}} . \tag{E13}$$

Doubling the error when converting from amplitudes to powers and then doubling it again because of the addition of the errors from the power offsets in equation (E9) gives

$$\sigma_{\Delta_P} = 4\pi^2 \Delta_r^2 \sqrt{P(r_0)/90} . \tag{E14}$$

Finally, by substituting into equation (E11), we get

$$\sigma_\alpha = \frac{3}{4\pi^2 \alpha \, \Delta_r^2 P(r_0)} 4\pi^2 \Delta_r^2 \sqrt{\frac{P(r_0)}{90}} = \frac{1}{\alpha \sqrt{10 P(r_0)}} . \tag{E15}$$

APPENDIX F

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 equation (D1), we initially assume a Fourier response equal to equation (18c) of

$$A_r = A_0 e^{-i\pi(r - r_0)} \, \text{sinc} \, [\pi(r - r_0)] , \tag{F1}$$

where $A_0 = Na/2e^{i\phi_0}$, ϕ_0 is the intrinsic phase of the signal, and r_0 is the true signal frequency (in FFT bins). This response is correct only for signals with a square window function [i.e., $W(u) = 1$ in the range $0 \leq u \leq 1$ and zero elsewhere].

From equation (F1) 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

$$\phi(r) = -2\pi r\hat{C} + c \; . \tag{F2}$$

When $r = r_0$, $\phi = \phi_0$, allowing us to solve for the constant of integration, $c = \phi_0 + 2\pi r_0\hat{C}$. Substituting into equation (F2), we see that the phase of the Fourier response is equal to

$$\phi(r_0) = \phi_0 - 2\pi\hat{C}(r - r_0) \; . \tag{F3}$$

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 rad.

The purity, as described in § 5 and Appendix E, 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 varies as sinc (πfT), 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 sinc $(\pi\alpha fT)$. This fact is also approximately true for more complicated window functions as long as $|r - r_0| \ll 1$. The Fourier response to a windowed sinusoid is therefore

$$A_r \simeq A_0 e^{-2\pi i\hat{C}(r-r_0)} \operatorname{sinc}\left[\pi\alpha(r - r_0)\right] \; . \tag{F4}$$

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

These same "effective duration" arguments also apply to the shape of the response in the $\dot{f}$ direction of the f-$\dot{f}$ plane. A change in the effective duration of a signal causes a change in the $\dot{f}$ response, since $fT^2 \rightarrow f\alpha^2T^2$, or equivalently, $\dot{r} \rightarrow \dot{r}\alpha^2$. The results of this change can be calculated by directly substituting $\dot{r}\alpha^2$ for $\dot{r}$ in equation (38), as was done in Appendix C.

REFERENCES

Abramowitz, M. & Stegun, I. A., ed. 1972, Handbook of Mathematical Functions (New York: Dover)
Anderson, S. B. 1992, Ph.D. thesis, Caltech
Anderson, S. B., Gorham, P. W., Kulkarni, S. R., Prince, T. A., & Wolszczan, A. 1990, Nature, 346, 42
Bailey, D. H. 1990, J. Supercomputing, 4, 23
Blackman, R. B., & Tukey, J. W. 1959, The Measurement of Power Spectra (New York: Dover)
Bracewell, R. N. 1999, The Fourier Transform and Its Applications (3d ed.; 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
Eikenberry, S. S., Fazio, G. G., & Ransom, S. M. 1998, ApJ, 492, 754
Fraser, D. 1976, J. Assoc. Comput. Machinery, 23, 298
Groth, E. J. 1975, ApJS, 29, 285
Israel, G. L., & Stella, L. 1996, ApJ, 468, 369
Johnston, H. M., & Kulkarni, S. R. 1991, ApJ, 368, 504
Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Sutherland, P. G., Kahn, S., & Grindlay, J. E. 1983, ApJ, 266, 160
Mardia, K. V., & Zemroch, P. J. 1975, Appl. Stat., 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, Ph.D. thesis, Univ. California Berkeley
Middleditch, J., & Córdova, 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: Cambridge Univ. Press), 372
Middleditch, J., & Kristian, J. 1984, ApJ, 279, 157
Middleditch, J., et al. 2000, NewA, 5, 243
Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes in C (2d ed.; Cambridge: Cambridge Univ. Press)
Ransom, S. M., Greenhill, L. J., Herrnstein, J. R., Manchester, R. N., Camilo, F., Eikenberry, S. S., & Lyne, A. G. 2001, ApJ, 546, L25
Sweet, R., & Wilson, J. 1995, Development of Out-of-Core Fast Fourier Transform Software for the Connection Machine (Denver: Cent. Comput. Math., Univ. Colorado)
Taylor, J. H., Manchester, R. N., Lyne, A. G., & Camilo, F. 1995, Pulsar Catalog
van der Klis, M. 1989, in Timing Neutron Stars, ed. H. Ögelman & E. P. J. van den Heuvel (NATO ASI Ser. C, 262) (Dordrecht: Kluwer), 27
Vaughan, B. A., et al. 1994, ApJ, 435, 362
Wood, K. S., et al. 1991, ApJ, 379, 295
Zhang, W., Jahoda, K., Swank, J. H., Morgan, E. H., & Giles, A. B. 1995, ApJ, 449, 930