THE ASTROPHYSICAL JOURNAL, 435:362–371, 1994 November 1

SEARCHES FOR MILLISECOND PULSATIONS IN LOW-MASS X-RAY BINARIES. II.

B. A. VAUGHAN,[1] M. VAN DER KLIS,[1] K. S. WOOD,[2] J. P. NORRIS,[3] P. HERTZ,[2] P. F. MICHELSON,[4] J. VAN PARADIJS,[1,5] W. H. G. LEWIN,[6] K. MITSUDA,[7] AND W. PENNINX[8]

Received 1993 November 2; accepted 1994 May 5

ABSTRACT

Coherent millisecond X-ray pulsations are expected from low-mass X-ray binaries (LMXBs), but remain undetected. Using the single-parameter Quadratic Coherence Recovery Technique (QCRT) to correct for unknown binary orbit motion, we have performed Fourier transform searches for coherent oscillations in all long, continuous segments of data obtained at 1 ms time resolution during *Ginga* observations of LMXB. We have searched the six known Z sources (GX 5−1, Cyg X-2, Sco X-1, GX 17+2, GX 340+0, and GX 349+2), seven of the 14 known atoll sources (GX 3+1, GX 9+1, GX 9+9, 1728−33. 1820−30, 1636−53 and 1608−52), the "peculiar" source Cir X-1, and the high-mass binary Cyg X-3. We find no evidence for coherent pulsations in any of these sources, with 99% confidence limits on the pulsed fraction between 0.3% and 5.0% at frequencies below the Nyquist frequency of 512 Hz.

A key assumption made in determining upper limits in previous searches is shown to be incorrect. We provide a recipe for correctly setting upper limits and detection thresholds. Finally we discuss and apply two strategies to improve sensitivity by utilizing multiple, independent, continuous segments of data with comparable count rates.

Subject headings: binaries: close — methods: numerical — stars: neutron — stars: oscillations — X-rays: stars

1. INTRODUCTION

Neutron star low-mass X-ray binaries (LMXBs) are believed to be systems comprised of a neutron star with a magnetic field $\lesssim 10^9$ G and a companion of mass $\lesssim M_\odot$ (Joss & Rappaport 1979; Lewin & Joss 1983; Webbink, Rappaport, & Savonije 1983; Lewin, van Paradijs, & van der Klis 1988; Bhattacharya & van den Heuvel 1991). Because of their weak magnetic fields and high accretion rates, the neutron star equilibrium spin periods in LMXBs are expected to be of order milliseconds (Ghosh & Lamb 1979a, b). Detection of coherent millisecond X-ray pulsations from an LMXB would constitute strong evidence for the evolution of millisecond radio pulsars from LMXB (Helfand, Ruderman, & Shaham 1983; Joss & Rappaport 1983; Paczyński 1983; Savonije 1983). It would support the currently favored beat frequency model of horizontal-branch QPO formation (Alpar & Shaham 1985; Lamb et al. 1985) in which quasi-periodic oscillations occur at the difference frequency between the neutron star rotation frequency and the Kepler frequency of the accreting matter as it falls onto the neutron star. It would provide a precise means to measure orbital periods and masses and to study accretion torques and neutron star dynamics on very short timescales.

The same weak magnetic fields and high accretion rates producing short neutron star spin period in LMXBs mean less channeling of accreted material along magnetic field lines and an accretion disk that may engulf the magnetosphere, resulting in smaller modulations, smearing from gravitational lensing (Wood, Ftaclas, & Kearney 1988; Mészáros, Riffert, & Berthiaume 1988), and further loss in observed modulation depth from scattering that both dephases the pulsar signal and spreads the pulsar beam (Brainerd & Lamb 1987; Kylafis & Klimmis 1987; Wang & Schlickeiser 1987; Bussard et al. 1988). See Wood et al. (1991, hereafter Paper I) for a more thorough discussion of spin period and modulation depth estimates.

It is likely that the spectral variability observed in all LMXB results from changes in the accretion environment. Most LMXBs fall into two classes: Z sources and atoll sources (Hasinger & van der Klis 1989), so-called because of the shapes they trace out in X-ray color-color diagrams. Z sources are more luminous and exhibit three states, called the horizontal, normal, and flaring branches, each with distinct, correlated spectral and time variability properties. The mass accretion rate onto the neutron star $\dot{M}$ is thought to increase as the source moves from the horizontal branch along the normal branch and onto the flaring branch (Vrtilek et al. 1990, 1991; Hasinger et al. 1990). Z sources are thought to have higher magnetic fields and accretion rates and, in general, longer orbital periods than atoll sources. (Hasinger & van der Klis 1989; van der Klis 1991). Atoll sources exhibit two spectral states, the island and the banana states, and show X-ray bursts.

Changes in accretion rate and inflow geometry within a given source can produce very different optical depths in the neutron star environment in different spectral states. If the horizontal-branch/normal-branch and normal branch/flaring-branch transitions correspond to the point at which the magnetosphere is engulfed by the accretion disk, and the Eddington limit is reached, respectively (Hasinger 1987; Lamb

[1] Astronomical Institute "Anton Pannekoek," University of Amsterdam, Center for High-Energy Astrophysics, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands. E-mail: brian@astro.uva.nl; michiel@astro.uva.nl; jvp@astro.uva.nl.

[2] Code 7621, E. O. Hulburt Center for Space Research, Naval Research Laboratory, Washington, DC 20375-5352. E-mail: 'SSD1::WOOD'@XIP.NRL.WAVY.MIL; hertz@xip.nrl.navy.mil.

[3] Code 690, Laboratory for High Energy Astrophysics, NASA/Goddard Space Flight Center, Greenbelt, MD 20771. E-mail: norris@grovx0.gsfc.nasa.gov.

[4] Physics Department, Stanford University, Stanford, CA 94305. E-mail: e7.f69@STANFORD.bitnet.

[5] Also University of Alabama in Huntsville.

[6] Massachusetts Institute of Technology, Center for Space Research, Room 37-627, Cambridge, MA 02139. E-mail: lewin@space.mit.edu.

[7] Institute of Space and Astronautical Science, 3-1-1 Yoshinodai, Sagamihara, Kanagawa 229, Japan. E-mail: mitsuda@astro. isas.ac.jp.

[8] TWI/SSOR, Mekelweg 4, 2628 CD Delft, The Netherlands.

1989; Hasinger & van der Klis 1989), the horizontal branch will have the smallest optical depth and the optimal conditions for detection of such pulsations. However, since there is no generally accepted explanation for the spectral states, it is desirable to search for coherent pulsations in all low-mass X-ray binaries in all spectral states. The set of searches presented here represents a complete survey of the Z sources, and a complete survey of all Z and atoll spectral states.

2. OBSERVATIONS

During its operational life, *Ginga* (1987 February–1991 November) was used for pointed observations of over 20 LMXBs with its Large Area Counter (LAC), a 4000 cm^2 proportional counter array sensitive in the range 1–37 keV (Makino et al. 1987; Turner et al. 1989). In its fastest telemetry mode, the PC mode, a pair of discriminators provided two spectral channels, the lower one accumulated 1024 times per second and the higher one 512 times per second. *Ginga*'s orbit and data recorder capacity allowed continuous data segments of up to roughly 40 minutes. *Ginga*'s large area, high time resolution, and relatively long gap-free observations make *Ginga* data well suited for pulsar searches. We scanned the *Ginga* database for PC mode observations of LMXB and found such observations for the sources Cir X-1, 1608−52, Sco X-1, 1636−53, GX 340+0, GX 349+2, GX 9+9, 1728−33, GX 3+1, GX 9+1, GX 5−1, GX 17+2, 1820−30, and Cyg X-2. We used the data accumulated 1024 times per second throughout our analyses, with integration times ranging from 256 through 2048 s. The energy range depended upon the discriminator settings used and varied between observations, typically covering a range of order 1–6 keV.

3. ANALYSIS

3.1. *The Coherence Recovery Technique*

The Quadratic Coherence Recovery Technique (QCRT) is described in detail in Paper I. We present here a summary of the method. We assume that the pulsating neutron star is in a binary system with a circular orbit of freqency $\Omega_{orb} = 2\pi/P_{orb}$. The approximations underlying QCRT are valid for general elliptical orbits, but we assume a circular orbit here to simplify the algebra. The pulsar has constant frequency ω_{pul} in its rest frame. The advance or retardation of the pulse phase due to binary motion is given by

$$\delta\phi_{pul} = \omega_{pul} a_\perp \sin(\Omega_{orb} t + \phi_0), \quad (1)$$

where $a_\perp$ is the projected orbital radius in light-seconds and ϕ_0 is the orbital phase at time $t = 0$. A periodic modulation in $\delta\phi_{pul}$ is equivalent to a periodic change in the observed pulse frequency. As a result, the pulsar signature in a power spectrum of the X-ray light curve is a broad feature, rather than a single peak. Expanding $\delta\phi_{pul}$ to second order in t yields

$$\delta\phi_{pul} = a_\perp \omega_{pul}[\sin\phi_0 + (\Omega_{orb}\cos\phi_0)t - \tfrac{1}{2}(\Omega_{orb}^2 \sin\phi_0)t^2]. \quad (2)$$

The second term in equation (2) is an average Doppler shift and does not broaden the pulsar peak, but rather shifts its apparent Fourier frequency. To compensate for the broadening produced by the quadratic term, we define a new time coordinate $t_\alpha = \alpha t^2$, where $\alpha = -\frac{1}{2} a_\perp \omega_{pul} \Omega_{orb}^2 \sin(\phi_0)$. Residual phase error leads to a loss of recovered power and limits the integration time. At most orbital phases, the first neglected term in the expansion of equation (1) gives the residual phase error, $E(\phi)$, and represents the typical error when the orbit is unknown:

$$E(\phi) = \tfrac{1}{6} a_\perp \omega_{pul} \Omega_{orb}^3 \cos(\phi_0) t^3. \quad (3)$$

More than 90% of the pulse fraction is recoverable at any orbital phase provided $T_{int} < P_{orb}/4\pi$. Thus, the maximum integration time is about 10% of the orbital period. If the integration time is significantly longer, phase residuals smear the signal over many frequency bins and severely diminish the sensitivity.

In all the sources we investigated, the orbital parameters are either unknown or ill constrained. We therefore searched a grid of trial α values. The minimum and maximum values of α, α_{min}, and α_{max}, and the increment $\delta\alpha$ are determined by orbital considerations, and by insisting that power split into two frequency channels at the Nyquist frequency by orbital motion is recovered into one channel, respectively, and are given by (Paper I)

$$\alpha_{max} = 1.79 \times 10^{-6} \mu_{comp} (P_{orb}/1\ \mathrm{hr})^{-4/3} (M_{tot}/M_\odot)^{1/3}\ \mathrm{s}^{-1}$$

$$\delta\alpha = \frac{1}{2\nu_{Nyq} T_{int}^2}, \quad (4)$$

where μ_{comp} is the ratio of the companion star mass to the total mass; M_{tot} is the total mass; ν_{Nyq} is the Nyquist frequency; T_{int} is the length, in seconds, of the FFT; and $\alpha_{min} = -\alpha_{max}$. Mismatch between the optimal value of α and the nearest value in the grid of α values leads to some loss in recovered power. In Appendix B of Paper I we show that the resulting loss in sensitivity is negligible at low frequencies and increases with frequency to about 7% at the Nyquist frequency.

The maximum Doppler shift of the pulsar frequency is the projected orbital velocity, $v_\perp$, divided by the speed of light, times the frequency ν:

$$\delta\nu_{max} = \beta_\perp \nu, \quad (5)$$

where $\beta_\perp = v_\perp/c$ is given by

$$\beta_\perp = 0.0021 \mu_{comp} (P_{orb}/1\ \mathrm{hr})^{-1/3} (M_{tot}/M_\odot)^{1/3}. \quad (6)$$

In calculating the power in the jth power spectrum bin, P_j, we adopt the normalization

$$P_j = \frac{2}{N_\gamma} |a_j|^2. \quad (7)$$

This normalization has become standard in X-ray astronomy. It differs from that used in Paper I by a factor of 2. The total number of photons is N_γ and the complex Fourier amplitudes are given by

$$a_j = \sum_{k=1}^{N} x_k e^{2\pi i jk/N}, \quad (8)$$

where x_k is the number of counts in data bin k of the light curve. To set detection thresholds, we need to consider the distribution of noise powers. To determine modulation depth limits in the absence of detection, we need to consider the interaction between signal and noise, and the amount of power we expect from a signal of modulation depth A and intensity r_0. Detection threshold, sensitivity, and upper limits are discussed in the next subsections.

To calculate the amount of power produced for a given average count rate and modulation depth, it is necessary to assume a pulse profile. Known X-ray pulsars in high-mass binaries have duty cycles of ~50% over a factor of ~1000 in spin period (Rappaport & Joss 1983), as compared with the ~4% duty cycles typical of radio pulsars (Srinivasan 1989). Like the pulsars we are searching for, pulsars in high-mass X-ray binaries are powered by accretion. The accretion process is not a function of neutron star spin period, unlike the electromagnetic processes powering radio pulsars. There is no reason to expect an additional factor of 100 in frequency to have a strong effect on pulse profile. It seems clear that the pulse profiles of slow X-ray pulsars are better assumption than the pulse profiles of millisecond radio pulsars.

For a sinusoidal profile $r(t) = r_0[1 + A \sin(w_{pul} t + \phi_0)]$, the expected signal power in the nearest Fourier frequency bin is (Leahy et al. 1983)

$$\langle P_j \rangle = 0.773 N_\gamma \frac{A^2}{2} \operatorname{sinc}^2 \left(\frac{\pi}{2} \frac{\nu_j}{\nu_{Nyq}} \right). \quad (9)$$

The frequency-dependent sinc function comes from binning. The multiplicative factor of 0.773 arises because the amount of power concentrated in a single bin depends upon the difference between the pulsar frequency and the nearest discrete Fourier frequency. On average, 0.773 of the power is concentrated in one bin.

If the pulse shape is not sinusoidal but instead has a duty cycle D (typical slow X-ray pulsars have $D \sim 0.5$), then the power at the fundamental frequency will be about D^2 times the power that would be observed for a sinusoidal profile with the same frequency and modulation depth.

Given a power, P, either a detection or an upper limit, equation (9) can be inverted to obtain the modulation depth, A:

$$A = 1.61 \sqrt{\frac{P}{N_\gamma}} 1 \Big/ \operatorname{sinc} \left(\frac{\pi}{2} \frac{\nu_j}{\nu_{Nyq}} \right). \quad (10)$$

The sinc function is unity at low frequency and reaches a value of $2/\pi$ at the Nyquist frequency. As a result, the sensitivity is degraded by about 50% at high frequencies.

3.2. *Detection Thresholds*

The detection threshold is determined entirely by the expected distribution of noise powers in the absence of any signal. It is by definition the power level that will not be exceeded by chance at any of the frequencies examined, with high confidence C, say 99% ($C = 0.99$). The false alarm probability is thus $1 - C$. If each of the powers is estimated from one frequency bin in one power spectrum, i.e., there has been no averaging of powers, then the noise power, P_{noise}, is distributed as χ^2 with 2 degrees of freedom and the probability that any single noise power exceeds a level P_0 is

$$\text{prob}_{single}(P_{noise} > P_0) = e^{-P_0/2}. \quad (11)$$

When examining N_{tot} frequency bins, the probability that at least one noise power exceeds P_0 is then

$$\text{prob}\,(P_{noise} > P_0) = 1 - (1 - e^{-P_0/2})^{N_{tot}}. \quad (12)$$

We set the detection threshold, P_{detect}, such that Prob $(P_{noise} > P_{detect})$ is $1 - C$. If C is close to 1, we obtain

$$1 - C = N_{tot} e^{-P_{detect}/2}, \quad (13)$$

or simply

$$P_{detect} = 2 \ln(N_{tot}) + 2 \ln\left(\frac{1}{1 - C} \right) \quad (14)$$

For a pulsar search performed using a single power spectrum, $N_{tot} = N_f$, the number of frequency bins examined, and

$$P_{detect} = 2 \ln(N_f) + 2 \ln\left(\frac{1}{1 - C} \right). \quad (15)$$

We show in Paper I that for QCRT with $\delta\alpha = 1/2\nu_{Nyq} T_{int}^2$, $N_{tot} = N_\alpha N_f/2$, where N_f is the number of frequency bins examined in each of the N_α power spectra. The detection threshold is then

$$P_{detect} = 2 \ln\left(\frac{N_\alpha N_f}{2} \right) + 2 \ln\left(\frac{1}{1 - C} \right). \quad (16)$$

3.3. *Sensitivities and Upper Limits*

In the following discussion we use the definitions for sensitivity and upper limits proposed by van der Klis (1989). Note that this approach is different from that taken by Leahy et al. (1983) and others, who call "upper limit" what we call "detection sensitivity." Sensitivity to a periodic signal, also called detection sensitivity or simply *sensitivity* and *the upper limit*, in the absence of a detection, are related but distinct quantities. Both require knowledge of the expected probability distribution of total power in a bin containing a signal and noise.

Detection sensitivity is a measure of the weakest signal a search procedure can confidently detect. Recall that to determine the C confidence level (say $C = 0.99$) detection threshold, P_{detect}, we need to consider all the independent frequency bins together. We assume no signal is present in any of the bins, only noise, and calculate P_{detect} such that the probability that none of the noise powers exceeds P_{detect} is C. The detection threshold thus depends upon the distribution of noise powers in all frequency bins in the absence of any signal. To determine the sensitivity power P_{sen}, we consider only a single frequency bin; call it j. We assume a signal of power P_{sen} is present in j and calculate P_{sen} such that the total power in j, arising from both signal and noise, will exceed P_{detect} with large probability C'. The sensitivity thus depends upon the distribution of total power in a single frequency bin containing a signal. Expressed as an amplitude of modulation, sensitivity depends additionally upon the assumed pulse shape and the total number of photons collected.

After computing a power spectrum and finding no power above the detection threshold, we can, for any given frequency bin j containing power P_j, determine an upper limit to the signal power. We assume a signal of power P_{UL} is present in j and calculate P_{UL} such that the total power in j will exceed P_j with large probability C'. To set an upper limit to an entire search in the absence of detection, we assume a signal of power P_{UL} is present in exactly one of the frequency bins investigated and calculate the upper limit in that bin. The most conservative upper limit is achieved by assuming the signal is present in the bin with the largest observed power, P_{max}.

It has become common practice in X-ray astronomy to assume that power is additive. In the context of pulsar searching, additivity of power means that the total power, P_{tot}, in a frequency bin containing both signal and noise is the sum of the signal power, P_{sig}, and the noise power, P_{noise}; $P_{tot} = P_{sig} + P_{noise}$.

Let us define $p(P_{tot}; P_{sig})$ to be the probability distribution of total power in a frequency bin containing a signal of power P_{sig}, and noise. Let $p_{noise}(P)$ be the probability distribution of noise power. P_{sig} is a fixed number. P_{noise} is a random variable. If we assume additivity of powers, then P_{tot} is the sum of a fixed number, P_{sig}, and a random variable, P_{noise}. The probability distribution of P_{tot} is thus the same as the probability distribution of P_{noise}, shifted by P_{sig}. Formally, the additivity of power assumption can be rewritten $P_{noise} = P_{tot} - P_{sig}$, and it follows immediately that $p(P_{tot}; P_{sig}) = p_{noise}(P_{tot} - P_{sig})$. For an individual power spectrum obtained with no averaging over frequencies, this means that in a frequency bin containing a signal of power P_{sig},

$$\text{Prob}\,(P_{tot} > P_0) = e^{-(P_0 - P_{sig})/2}, \quad P_0 > P_{sig}$$
$$= 1, \quad P_0 \le P_{sig}. \tag{17}$$

Assuming additivity of powers, the distribution of total power looks just like χ^2 with 2 degrees of freedom, but shifted in power by an amount P_{sig}. The commonly followed procedure laid out by Leahy et al. (1983) contains the additivity of powers assumption in equation (11), as does Paper I in equation (21).

In truth, the probability distribution of total power in a frequency bin containing both a signal and noise is not χ^2, but is more complicated. In § 3.5 we give a geometric argument for the reason that the distribution of total power in the presence of a signal is so different from the distribution in the absence of a signal. Groth (1975) exactly calculates $p_n(P_{tot}; P_{sig})$, the true probability distribution of total power, P_{tot}, for the sum of n independent power spectral bins containing Gaussian noise and signal power P_{sig}, where P_{sig} is the sum of the signal powers in the summed bins. He shows that

$$p_n(P_{tot}; P_{sig}) = (P_{tot}/P_{sig})^{(n-1)/2} \times \exp\,[-(P_{tot} + P_{sig})/2] I_{(n-1)}[(P_{tot} P_{sig})^{1/2}], \tag{18}$$

where I is a modified Bessel function of the first kind. Note that Groth's power normalization is one-half of ours. Formulae given here have been modified to our normalization. The probability that the total power is between 0 and P_{tot} is

$$f_n(P_{tot}; P_{sig}) = \int_0^{P_{tot}} p_n(x; P_{sig})dx = 1 - \exp\,[-(P_{tot} + P_{sig})/2] \times \sum_{m=0}^{\infty} \sum_{k=0}^{m+n-1} P_{tot}^k P_{sig}^m/(k!\, m!\, 2^{m+k}). \tag{19}$$

Figure 1 of Groth gives probability contours in the P_{tot}, P_{sig} plane. Groth points out that if a large number, n, of independent frequency bins are summed, the distribution of P_{tot} approaches a normal distribution, and additivity of powers becomes approximately correct. However, if n is small, or unity as in our case, equation (17) is simply wrong. Equation (11) in Leahy et al. is thus only correct for large n. Equation (21) in Paper I is wrong. Upper limits presented in Table 1 of the present paper supersede those in Table 2 of Paper I. It is likely that many additional published upper limits are similarly incorrect. We find that the upper limits presented in Paper I increase by about 30% when calculated correctly. The pulsar searches themselves remain valid as a detection threshold does not depend upon the interaction between signal and noise.

A correct procedure for determining upper limits following a single, unsuccessful pulsar search, with or without coherence recovery, is as follows:

1. Search the ensemble of power spectra for the largest observed power, P_{max}.

2. Decide upon a confidence level, C'. The upper limit is the amplitude of the signal that, if it were present in the bin with the largest power P_{max}, would, with probability C', have produced a power greater than P_{max}, in that bin. For example, a 99% confidence ($C' = 0.99$) upper limit A means there is a 1% chance a signal of strength A would result in a power less than or equal to P_{max}.

3. Invert equation (19) numerically to find P_{sig}, using $f_n = 1 - C'$ and $P_{tot} = P_{max}$. A simple FORTRAN program for doing so is included as an Appendix. Alternatively, use the probability contours in Figure 1 of Groth (1975). In using contours from Groth (1975), recall again that his normalization differs, and all his powers are a factor of 2 smaller. The FORTRAN routine in the appendix follows our normalization.

4. Convert the upper limit power from step 3 into an amplitude of modulation. Equation (10) provides a relation between modulation amplitude, $A(\nu)$, and signal power, for a sinusoidal signal. Note that it is frequency dependent. For a different (nonsinusoidal) pulse shape, it is necessary to determine a new relation. The frequency-dependent part of equation (10) arises from binning.

5. As mentioned in § 2, QCRT introduces losses due to incomplete signal recovery from approximating an elliptical orbit as locally quadratic, and due to mismatch between the value of α, from the grid of trial values, that best recovers the signal, and the truly optimum value. These losses increase with frequency. As a rule of thumb, $A(\nu)$ is correct below about half the Nyquist frequency. For frequencies near the Nyquist frequency, we find from simulations that it is necessary to increase $A(\nu)$ by 20% for a more realistic upper limit. Johnston & Kulkarni (1991) present an extensive set of simulations quantifying such losses for a variety of pulsar search strategies.

To determine the search sensitivity rather than the upper limit, follow the above procedure, but skip step 1 and use the threshold power P_{detect} (see eq. [16]) instead of P_{max} in step 3.

3.4. *Utilizing Multiple Searches to Improve Sensitivity*

In the previous section, we discussed setting upper limits for searches performed on one segment of data. Often the searches involve hundreds of power spectra, each with a different value of the QCRT parameter α. In this section, assume we have searched N_s nonoverlapping segments of data from the same source, with or without QCRT or other coherence recovery procedures. If the count rate is nearly the same in each segment, then, combined, these N_s searches can provide significantly better sensitivity than any of the individual searches.

When we search for peaks in all power spectra from all N_s segments, the number of frequency bins examined increases by a factor of N_s, so the detection threshold must rise by an additive factor of $2 \ln N_s$ to

$$P_{detect} = 2 \ln\left(\frac{n_\alpha N_r}{2}\right) + 2 \ln\left(\frac{1}{1 - C}\right) + 2 \ln N_s. \tag{20}$$

We define the search sensitivity as the minimum signal strength that, if present in all of the segments of data searched, will, with large probability C'', produce a total power above threshold in the power spectrum of at least one of the segments. Because the segments are independent, C'' is related to the probability that the total power will exceed threshold in any one individual segment, C', by

$$(1 - C')^{N_s} = 1 - C''. \tag{21}$$

For example, if $N_s = 4$ and $C'' = 0.99$, then $C' = 0.68$. To be 99% confidence that a signal present in each of the segments will be detected in at least one of them, it is sufficient to be 68% confident that the signal will be detected in any individual one. Because the probability distribution of total power is very broad for signal strengths near threshold, the sensitivity power P_{sen} is much lower for 68% confidence that it is for 99% confidence, and the sensitivity is consequently much better. If $N_s = 4$, $C = 0.99$, $C'' - 0.99$ (hence $C' = 0.68$ as above), and in each segment of data 10^6 independent frequency bins are investigated, then for an individual segment the detection threshold is $P_{detect} = 36.8$ and the sensitivity power, found using the procedure in the preceding section, is $P_{sen} = 69.2$. Taken together, the detection threshold is $P_{detect} = 39.6$, and the sensitivity power is $P_{sen} = 44.7$. The detection sensitivity, expressed as an amplitude of modulation, improves by 35%.

It has become standard practice in X-ray astronomy to average together power spectra of independent segments of data to improve statistics. As discussed in § 1, for sources in a binary system, if long FFTs are performed, averaging of power spectra does not help in pulsar detection because the signal will generally appear in different frequency bins within some range for different data segments because of Doppler shifting. The apparent signal frequency will vary with orbital phase, hence with data segment, but never by more than $\delta\nu_{max}$ as given by equation (6), in general less than about 0.5 Hz for a Nyquist frequency of 500 Hz. In all segments of data, the pulsar signal will be recovered by QCRT, at the Doppler-shifted frequency, in a single power spectral bin (in practice if may be recovered in adjacent bins) in the Fourier frequency range $(\nu - \delta\nu_{max}, \nu + \delta\nu_{max})$, where recall that $\delta\nu_{max}$ is proportional to ν. That is, the recovered peaks will be clustered about the true pulsar frequency. Although it is not possible to average power spectra from independent segments, it is possible to search for clusters of peaks. We call such a procedure "cluster analysis."

To determine the detection threshold and sensitivity for cluster analysis, divide the frequency axis of the power spectra of all the data segments searched into intervals of (varying) width $2\delta\nu_{max} = 2\beta_\perp \nu$. Recall that for each segment of data, QCRT involves computing multiple (in most cases hundreds of) power spectra, each with a different trial correction parameter, α, applied to the data. Both the value of α that optimally recovers the pulsar signal and the (Doppler-shifted) frequency at which it is recovered depend upon orbital phase. In general, the signal is recovered by different values of α, and at different frequencies, in each segment. Define detection as an interval containing a power above the detection threshold in at least one of the power spectra of every data segment searched. It is possible to show that the detection threshold becomes approximately

$$P_{detect} = 2\ln(N_a N_f) + \frac{2}{N_s}\ln\left[\frac{(2\beta_\perp)^{N_s - 1}}{2N_s(1 - C)}\right] \qquad (22)$$

provided $2\beta_\perp N_a N_f e^{-P_0/2} \ll 1$. For searches performed without QCRT, set $N_a = 1$ and replace $2N_s$ in the second logarithm with N_s. In implementing cluster analysis, the best procedure is not to define fixed frequency intervals but to search for clusters of peaks occurring anywhere. This will affect the probability. However, we find from simulations the equation (22) is approximately correct in this case.

Cluster analysis yields a lower detection thershold. For the same values as above ($N_s = 4$, $C = 0.99$, $N_f = 10^6$), and assuming $N_a = 1$ and $\beta_\perp = 0.001$, the detection threshold becomes $P_{detect} = 19.9$. To determine detection sensitivity, recall that a cluster analysis detection requires signal plus noise above threshold in the same frequency segment in a power spectrum of every one of the data segments. To claim with confidence C'' that a given signal will produce total power exceeding the detection threshold in all N_s segments requires that we can claim a confidence C' for each segment, where now, $(C')^{N_s} = C''$. If $N_s = 4$ and $C'' = 0.99$, then $C' = 0.9975$. Although it yields a lower detection threshold, cluster analysis requires higher confidence of detection for each segment than was necessary in searching individual segments. With $P_{detect} = 19.9$ and $C' = 0.9975$, the sensitivity power becomes $P_{sen} = 51.54$, better than for an individual search, but not as good as searching all the power spectra for a single peak above threshold as above.

3.5. *Noise-Signal Interactions*

For a 10^6 point power spectrum, the 99% confidence detection threshold power (P_{detect} in eq. [14]) is about 40. If a power greater than 40 is detected in the power spectrum, we can be confident that a signal is present. What is the signal power to which such a power spectrum is sensitive? According to equation (19) (at 99% confidence), this signal power is about 75. As the average noise power is only 2, it may seem odd that the signal power must be nearly twice the threshold power to be 99% confident that the total power, signal plus noise, will exceed the threshold power. To see why, recall that the total power in a frequency bin is the square of the sum of the signal and noise, each of which is a complex number. It is useful to draw these numbers in the complex plane (Fig. 1). In this plane, signal s and noise n appear as two-dimensional vectors.

Consider a frequency bin j containing signal and noise. The signal power is $P_{sig} = |s|^2$. The noise power is $P_{noise} = |n|^2$.

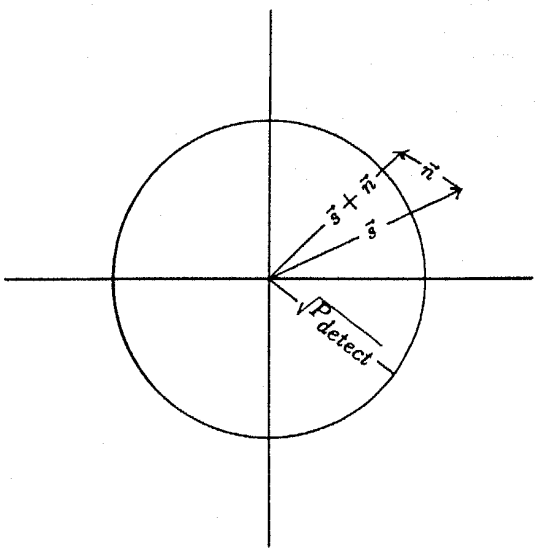

FIG. 1.—The total observed power is $P_{tot} = |s + n|^2$. For a typical search, $P_{detect} \sim 40$; thus $(P_{detect})^{1/2} \sim 6.3$. The real and imaginary parts of the noise, n, are normally distributed with null mean and unit variance. To be 99% confident that $|s + n| > (P_{detect})^{1/2}$, it is necessary to have $|s| > 8.6$, and hence $P_{sig} = |s|^2 > 75$.

The measure Fourier coefficient is $s + n$, and the total power in the bin is $P_j = |s + n|^2$. For the total power to exceed the detection threshold power $P_{\rm detect}$, $s + n$ must lie outside of a circle of radius r, where $r = (P_{\rm detect})^{1/2}$. For $P_{\rm detect} = 40$, we have $r = 6.3$. A signal power of $P_{\rm sig} = 75$ corresponds to a signal Fourier coefficient of length $|s| = (75)^{1/2} = 8.65$. In the presence of a signal power of 75, the total power can fall below the threshold of 40 only if the noise Fourier coefficient n has magnitude greater than $(75)^{1/2} - (40)^{1/2} = 2.35$, and if n points sufficiently far inward. A noise Fourier coefficient of 2.35 corresponds to a noise power of 5.5. The probability that the noise power exceeds 5.5 is 0.06. The probability that n points inward is 0.5; the probability that it points sufficiently far inward to bring $s + n$ inside the circle is even less. Thus, the probability that the total power falls below threshold is less than, but of order, 0.03. We find from equation (19) that it is 0.01. The signal power must be nearly twice the threshold power to be 99% confident that the total power will exceed threshold because (1) noise and signal combine as vectors, not as scalars, so that noise can either enhance or diminish the signal, and (2) the lengths of these vectors are equal to the square roots of the noise and signal powers, respectively.

4. RESULTS

Table 1 gives the modulation depth upper limits obtained for each individual data segment searched, a total of 53 segments from 15 sources. Included in the table is the observation date and time, source type (Z, atoll, or peculiar); the count rate and spectral state of the source during the search; the FFT length in seconds; the number of FFTs; 99% confidence upper limits on the modulation depth at 50 and 500 Hz; and the minimum orbital period, min $P_{\rm orb}$, for which the single parameter orbit correction is guaranteed to recover at least 90% of the pulsar signal. At quadrature, where the recovery problem is most severe, we recover 90% of the pulsar signal for $P_{\rm orb} = 4\pi T_{\rm int}$. The min $P_{\rm orb}$ values listed in the table are conservative estimates, valid even if the search is at or near quadrature. All of the upper limits reported in Paper I have been recalculated here, taking into account the interaction between signal and noise power described in § 4.3. By chance, upper limits for GX 9+1 reported in Hertz et al. (1990) remain approximately unchanged. A total of nine independent segments of data with comparable counting rates were searched using QCRT applied to equally long (2^{22} point) segments of data, and multiple search considerations discussed in § 4.4 happen to produce upper limits nearly identical to those originally obtained using the incorrect method of Paper I.

Six of the sources we searched have known orbital periods. The periods are given in the notes to the table. For two of these sources (Cir X-1 and Cyg X-3), X-ray properties are known to depend on orbital phase. The phases are thus also included in the caption.

The count rates given are total observed count rates, source plus background. The typical background rate in the energy channel used in the searches is $\sim 30\ {\rm s}^{-1}$. For dim sources, we took the background rate into account in determining upper limits. For bright sources, it is negligible.

The 2048 s observation of 1608−52 on 1989 August 23 and the 1024 s observation of 1728−33 on 1987 October 12 each contain an X-ray burst. Although the burst profiles dominate the power spectra at low Fourier frequencies, their effects are completely negligible in the range we are interested in, and necessitate no recalculation of the noise powers.

The orbital period of 1820−30 is 685 s. The formal limit on the validity of the single parameter coherence recovery technique is thus about 1 minute. However, it is believed that the companion star in 1820−30 is either a very low mass Roche lobe–filling white dwarf or a low-mass helium-burning star (van der Klis et al. 1993). In the case of a white dwarf, the companion mass is estimated to be 0.03 $M_\odot$, and in the case of a helium-burning companion, it is of order 0.24 $M_\odot$. It could be as low as 0.01 $M_\odot$. If so, a pulsar signal at 50 Hz would be spread over two frequency bins. Under the assumption that the companion mass in 1820−30 is 0.01 $M_\odot$, we may consider the search meaningful out to 50 Hz.

5. DISCUSSION

It is useful to think of a pulsar search as covering a space of four parameters whose dimensions are Fourier frequency, modulation depth, source and spectral state. We wish to be as complete as possible in each of the dimensions.

We have now performed an essentially complete survey of the millisecond low-mass X-ray binary data obtained using *Ginga*, which has the largest collecting area of all X-ray observatories flown so far. Our search covered 53 segments of data from 15 luminous X-ray binaries, and included data from all three Z source and both atoll source spectral states. However, it is far from complete. In all but two sources we searched only a single source state, and in only one, Cyg X-2, did we search all spectral states. Our completeness in frequency and modulation depth are comparable to that achieved in previous searches using *Ginga* data.

Although many of the upper limits listed in Table 1 are less than 0.75%, they are still an order of magnitude higher than the modulation depths we optimistically expect in low magnetic field neutron star binaries in the presence of unrecoverable signal smearing from scattering and gravitational lensing.

As discussed in § 1, it is likely that changes in accretion environment make some spectral states more favorable than others. The horizontal branch in Z sources is probably the most favorable, but in the absence of consensus on the causes of spectral variability, it is desirable to search all states. It is possible that at least some of the systems we have examined contain neutron stars with spin periods above the limited frequency range we searched, to 512 Hz. For such systems, no increase in viewing area or observation length will make it possible to observe pulsations if it is not accompanied by an improvement in time resolution. It seems unlikely, however, that the neutron star spin frequencies in all of the systems we have looked at are in excess of 512 Hz, considering the wide range of companion types, orbital periods, and likely evolutionary histories covered in our survey. Finally, it is possible that for some of the systems, the unknown orbital period is shorter than the limit given in Table 1. If that is the case, then the modulation depth limits for those sources are only valid over a part of our frequency range. In any event, considering the size of our survey, it is extremely improbable that a spin frequency in excess of 512 Hz, a very low orbital period, and/or an unfavorable spectral state has rendered all of our searches invalid.

New high count rate, millisecond X-ray data will have to wait for the next large area counter arrays. Two such arrays, XTE and USA, are scheduled for launch before the end of the century. It will be useful in the meantime to look for millisecond pulsations in as yet unsearched archival data, particu-

TABLE 1

LIMITS ON PULSE FRACTIONS IN LMXB[a]

SOURCE	OBSERVATION DATE	COUNT RATE	STATE[b]	T_{int}(s)	N_a	A 50 Hz	A 500 Hz	MINIMUM P^c_{orb}
Sco X-1 (1617-15)	1989 Mar 9	14000	F	2048	121	0.0028	0.0051	[d]
	1989 Mar 10	10000	F	2048	121	0.0031	0.0057	
GX 340+0 (1642-45)	1988 Mar 30	2200	N	1024	215	0.0096	0.018	3
	1988 Apr 6	2200	N	1024	215	0.0096	0.018	3
GX 5-1 (1758-25)	1987 Apr 20	7000	H	1024	215	0.0054	0.0099	3
	1987 Apr 27	5000	N[e]	2048	857	0.0045	0.0083	6
GX 9+1 (1758-20)	1988 Mar 29	3000	B	1024	215	0.0082	0.015	3
	1988 Mar 31	3250	B	1024	215	0.0079	0.015	3
GX 17+2 (1813-14)	1988 Mar 28	3500	F	1024	215	0.0076	0.014	3
	1988 Apr 1	2700	FV	1024	215	0.0087	0.016	3
Cyg X-2 (2142+38)	1987 Jun 7	3400	N	1024	3	0.0052	0.0095	[f]
	1987 Jun 8	3400	N	1024	3	0.0052	0.0095	
	1988 Jun 10	3000	H	2048	7	0.0050	0.0092	
	1988 Jun 10	2900	H	1024	3	0.0071	0.013	
	1988 Jun 11	3000	N	2048	7	0.0076	0.014	
	1988 Jun 11	3000	N	2048	7	0.0051	0.0094	
	1988 Jun 13	2900	N	1024	3	0.0073	0.013	
	1988 Oct 6	3600	F	1024	3	0.0063	0.012	
	1988 Oct 6	3600	FV	1024	3	0.0068	0.013	
	1988 Oct 7	4000	H	1024	3	0.0060	0.011	
	1989 Oct 22	4200	H	1024	3	0.0060	0.011	
	1989 Oct 22	6000	H	2048	7	0.0037	0.0068	
	1989 Oct 22	6100	HV	1024	3	0.0049	0.0090	
	1989 Oct 23	4600	H	1024	3	0.0055	0.010	
GX 9+9 (1728-16)	1988 Aug 8	2300	B	2048	961	0.0065	0.012	[g]
	1988 Aug 8	2300	B	1024	241	0.0089	0.016	
	1988 Aug 9	2200	B	2048	961	0.0071	0.013	
	1988 Aug 9	2100	B	1024	241	0.0093	0.017	
GX 349+2 (1702-36)	1989 Mar 5	5900	F	1024	301	0.0057	0.010	3
	1989 Mar 5	4300	FV	1024	301	0.0064	0.012	3
	1989 Mar 5	4000	FV	1024	301	0.0066	0.012	3
	1989 Mar 6	3900	F	1024	301	0.0071	0.013	3
Cir X-1 (1516-57)	1989 Mar 30	430	[h]	1024	301	0.013	0.024	[i]
	1987 Aug 25	2400	[j]	256	101	0.0099	0.018	
Cyg X-3 (2030+40)	1987 Aug 27	390	[k]	1024	201	0.023	0.043	[l]
	1987 Aug 27	280	[m]	1024	201	0.027	0.040	
	1987 Aug 27	330	[n]	1024	201	0.024	0.044	
GX 3+1 (1744-26)	1987 Apr 13	1250	B	1024	301	0.012	0.022	3
	1987 Apr 13	1400	B	1024	301	0.011	0.020	3
1608-52	1989 Aug 23	130	I	1024	301	0.042	0.078	3
	1989 Aug 23	130	I	1024	301	0.043	0.079	3
	1989 Aug 23	170	I	2048	961	0.024	0.044	6
	1990 Mar 7	310	I	512	161	0.034	0.063	1.5
	1990 Mar 7	270	I	512	161	0.037	0.068	1.5
	1991 Aug 13	300	I	512	161	0.034	0.063	1.5
	1991 Aug 13	400	I	512	161	0.029	0.054	1.5
1728-33	1987 Oct 12	650	I	512	161	0.023	0.043	1.5
	1987 Oct 12	770	I	1024	301	0.014	0.026	3
1636-53	1987 Aug 20	1150	B	1024	301	0.013	0.024	3
	1987 Aug 20	1300	B	1024	301	0.013	0.024	3
	1987 Aug 20	1200	B	1024	301	0.013	0.024	3
1820-30	1989 Apr 24	1150	B	1024	1	0.011	—[o]	[p]
	1989 Apr 24	1300	B	1024	1	0.011	—	

[a] 99% confidence.

[b] H = horizontal branch, N = normal branch, F = flaring branch, FV = flaring branch vertex, HV = horizontal branch vertex, B = banana state, I = island state.

[c] Minimum P_{orb} to recover >90% of power independent of orbital phase.

[d] Known period (19.2 hr).

[e] May exhibit newly discovered flaring branch behavior (Kuulkers et al. 1994).

[f] Known period ($9^d.8$).

[g] Known period (4.2 hr).

[h] Low-intensity soft state; orbital phase $\phi \sim 0.95$.

[i] Known period ($16^d.6$).

[j] Orbital phase $\phi \sim 0.75$.

[k] Low state, $\phi \sim 0.6$.

[l] Known period (4.8 hr).

[m] Low state, $\phi \sim 0.92$.

[n] Low state, $\phi \sim 0.2$.

[o] See § 5, last paragraph. Only valid if $M_c \leq 0.01\ M_\odot$ and for $\nu \leq 50$ Hz.

[p] Known period (685 s).

larly from the *EXOSAT* mission, in an attempt to search new areas of millisecond X-ray pulsar parameter space. Given the difficulty in obtaining new data, it is worthwhile to attempt new strategies, such as QCRT, to look even harder at data already analyzed. Even a low significance detection is interesting because it is then possible to go back and look again at other data from the same source, searching a much smaller frequency range with much better sensitivity to attempt to confirm the detection.

The authors wish to thank the *Ginga* team for their help in data extraction and analysis. P. H. and K. S. W. acknowledge partial support from the Office of Naval Research. W. H. G. L. was supported by the United States National Aeronautics and Space Administration under grant NAG 8-286. J. vP. acknowledges support from NATO through grant RG 331/88. This work was supported in part by the Netherlands Organization for Scientific Research (NWO) under grant PGS 78-277.

APPENDIX

Three FORTRAN functions are contained in this appendix. The first, FN, returns $f_n(P; P_s)$ given P, P_s, and n as input, using equation (19) (from Groth 1975, with a modified normalization). Double precision arithmetic is used. The variable DMIN specifies the precision. The outer sum (over m) is truncated when the incremental change in the partial sum is less than DMIN. All terms in the sum are calculated as logarithms and then exponentiated, to avoid overflow, underflow, and roundoff error. On return, the value EPS is subtracted from fn, for use in root search inversion of equation (19) to find P_s given P and f_n (performed by subroutine UL). To use FN to find $f_n(P; P_s)$, call FN with EPS equal to zero.

The second routine, FNC, simply calls FN with $P = -1$. When called with $P < 0$, FNC uses the previous values of P and n, and does not recalculate the array of logarithms LN.

The third routine, UL, returns P_s when called with P, n, and EPS. Given N summed frequency bins and total power (signal plus noise) P, UL returns the value of PS for which the probability that the total power is less than P is equal to EPS. To determine upper limits in the absence of incoherent summation, enter P_{max} for P, 1 for N, and one minus the confidence level (say $1 - 0.99 = 0.01$) for EPS. For sensitivity, replace P_{max} with P_{detect}.

In order to determine P_{sig}, UL first roughly brackets its position by stepping away from P, then calls a root finding routine, here generally called ROOT, to pin it down. ROOT is assumed to call FNC for function evaluation. Because, for given P_{tot}, $f_n(P_{tot}; P_{sig})$ is a monotonically decreasing function of P_{sig}, any simple root finding routine should converge rapidly. We found the bisection routine RTBIS from Numerical Recipes in C (Press et al. 1992) perfectly adequate.

The routines FN and UL are very general. The first, FN, returns $f_n(P_{tot}; P_{sig})$ given n, P_{tot}, and P_{sig}. And UL simply inverts $f_n(P_{tot}; P_{sig})$ to return P_{sig} given f_n, n, and P_{tot}. Thus, UL can also be used, for example, to determine error bars on the signal strength in the case of a detection. Or, FN and UL can be used to compare the sensitivity of searches carried out with and without incoherent summation, as in § 4.4, or different numbers of summed power spectra.

We cannot guarantee that the routines presented here are free of errors. We have used them to reproduce several of the figures included in Groth (1975), and have checked that in the limit of large n, $f_n(P_{tot}; P_{sig})$ approaches a normal distribution. Note that for large values of n, typically above 10, noise and signal power become essentially additive and the routines in this Appendix are unnecessary. Note also that P_{sig} does not exist (to be more correct, is not real valued) for all combinations of P_{tot}, n and EPS. These routines do not check that the calling values are possible and hence crash if they are not. Finally, let us point out yet again that the normalization used in these routines differs from that used in Groth (1975) by a factor of 2. To use the normalization in Groth (1975), simply remove the lines P = P/2.DO and UL = UL * 2. from UL.

```
FUNCTION FN(PO, PS, NO, EPSO)
PARAMETER(BIGM = 999, DMIN = 1.D-6)
PARAMETER(PMAX = 100, NMAX = 50)
IMPLICIT DOUBLE PRECISION(A-H,L,O-Z)
REAL*4 FN
DOUBLE PRECISION LN(2*PMAX + NMAX)
SAVE LNPS,N,LN,EPS
LNPS = DLOG(PS)
IF(PO.GE.O.DO) THEN
  P = PO
  LNP = DLOG(P)
  N = NO
  EPS = EPSO
  DO J = 1, 2 * PMAX + NMAX
    LN(J) = DLOG(DFLOAT(J))
  END DO
END IF
SUM = O.DO
LNMO = - P - PS
DO M = 0, BIGM
  OLDSUM = SUM
```

```
    LNMK = LNMO
    DSUM = DEXP(LNMK)
    DO K = 1, M + N - 1
      LNMK = LNMK + LNP - LN(K)
      DSUM = DSUM + DEXP(LNMK)
    END DO
    IF(SUM.GT.0.) THEN
      IF((DSUM/SUM).LE.DMIN) GO TO 88
    END IF
    SUM = SUM + DSUM
    LNMO = LNMO + LNPS - LN(M+1)
  END DO
  88 SUM = SUM + DSUM
  FN = 1.DO - SUM - EPS
  RETURN
  END

  FUNCTION FNC(PS)
  DOUBLE PRECISION PS, PDUMMY
  PDUMMY = -1.DO
  NDUMMY = -1
  FNC = FN(PDUMMY, PS, NDUMMY, EDUMMY)
  RETURN
  END

  FUNCTION UL(P,N,EPS)
  REAL*8 P
  PARAMETER(XACC=1,E-4)
  FACT = 1.1DO
  P = P / 2.DO
  PS = P
  X = FN(P, PS, N, EPS)
  XO = X
  IF(X.LT.0.) FACT = 1. / FACT
  DO J = 1, 9999
    PS = FACT * PS
    X = FNC(PS)
    IF(X*XO.LT.0.) THEN
      IF(XO.GE.0.) THEN
        UL = ROOT(PS/FACT,PS,XACC)
      ELSE
        UL = ROOT(PS,PS/FACT,XACC)
      END IF
      UL = UL * 2.
      RETURN
    END IF
  END DO
  END
```

REFERENCES

Alpar, M. A., & Shaham, J. 1985, Nature, 316, 239
Bhattacharya, D., & van den Heuvel, E. P. J. 1991, Phys. Rep., 203, 1
Brainerd, J., & Lamb, F. K. 1987, ApJ, 317, L33
Bussard, R. W., Weisskopf, M. C., Elsner, R. F., & Shibazaki, N. 1988, ApJ, 327, 284
Ghosh, P., & Lamb, F. K. 1979a, ApJ, 232, 259
———. 1979b, ApJ, 235, 296
Groth, E. J. 1975, ApJS, 29, 285
Hasinger, G. 1987, in IAU Symp. 125, The Origin and Evolution of Neutron Stars, ed. D. J. Helfand & J.-H. Huang (Dordrecht: Reidel), 333
Hasinger, G., & van der Klis, M. 1989, A&A, 225, 79
Hasinger, G., van der Klis, M., Ebisawa, K., Dotani, T., & Mitsuda, K. 1990, A&A, 235, 131
Helfand, D. J., Ruderman, M. A., & Shaham, J. 1983, Nature, 304, 423
Hertz, P., Norris, J. P., Wood, K. S., Vaughan, B. A., & Michelson, P. F. 1990, ApJ, 354, 267
Johnston, H., & Kulkarni, S. 1991, ApJ, 368, 504
Joss, P., & Rappaport, S. 1983, Nature, 304, 419
Kuulkers, E., van der Klis, M., Oosterbroek, T., Asai, K., Dotani, T., van Paradijs, J., & Lewin, W. H. G. 1994, A&A, in press
Kylafis, N. D., & Klimmis, G. S. 1987, ApJ, 323, 678
Lamb, F. K. 1989, in Proc. 23rd ESLAB Symp. on Two Topics in X-ray Astronomy, ed. J. Hunt & B. Battrick (Noordwijk: ESA Publications Division), 215
Lamb, F. K., Shibazaki, N., Shaham, J., & Alpar, M. A. 1985, Nature, 317, 681
Leahy, D., Darbro, W., Elsner, R., Weisskopf, M., Sutherland, P., Kahn, S., & Grindlay, J. 1983, ApJ, 266, 160
Lewin, W. H. G., & Joss, P. C. 1983, in Accretion Driven Stellar X-ray Sources, ed. W. H. G. Lewin & E. P. J. van den Heuvel (Cambridge: Cambridge Univ. Press), 41
Lewin, W., van Paradijs, J., & van der Klis, M. 1988, Space Sci. Rev., 46, 273
Makino, F., & the *Ginga* ASTRO-C Team, 1987, Ap. Lett. Comm., 25, 223

Mészáros, P., Riffert, H., & Berthiaume, G. 1988, ApJ, 325, 204
Paczyński, B. 1983, Nature, 304, 421
Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes in C (Cambridge: Cambridge Univ. Press), 334
Rappaport, S. A., & Joss, P. C. 1983, in Accretion Driven Stellar X-ray Sources, ed. W. H. G. Lewin & E. P. J. van den Heuvel (Cambridge: Cambridge Univ. Press), 1
Savonije, G. J. 1983, Nature, 304
Srinivasan, G. 1989, Astr. Ap. Rev., 1, 209
Turner, M. J. L., et al. 1989, PASJ, 41, 345
van der Klis, M. 1989, in Timing Neutron Stars, ed. H. Ogleman & E. P. J. van den Heuvel (Dordrecht: Kluwer), 27
van der Klis, M. 1991, in Frontiers of X-Ray Astronomy, ed. Y. Tanaka & K. Koyama (Tokyo: Universal Academic), 139
van der Klis, M., et al. 1993, MNRAS, 260, 686
Vrtilek, S. D., Penninx, W., Raymond, J. C., Verbunt, F., Hertz, P., Wood, K., Lewin, W. H. G., & Mitsuda, K. 1991, ApJ, 376, 278
Vrtilek, S. D., Raymond, J. C., Garcia, M. R., Verbunt, F., Hasinger, G., & Kürster, M. 1990, A&A, 235, 162
Wang, Y.-M., & Schlickheiser, R. 1987, ApJ, 313, 200
Webbink, R. F., Rappaport, S. A., & Savonije, G. J. 1983, ApJ, 270, 678
Wood, K. S., Ftaclas, C., & Kearney, M. 1988, ApJ, 324, L63
Wood, K. S., et al. 1991, ApJ, 379, 295 (Paper I)