X-Ray Timing of the Accreting Millisecond Pulsar SAX J1808.4−3658

by

Jacob M. Hartman

BS Mathematics, BS Physics
Pennsylvania State University (1999)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2007

Author .....
Department of Physics
August 31, 2007

Certified by.... .....
Deepto Chakrabarty
Associate Professor of Physics
Thesis Supervisor

Accepted by .....
Thomas J. Greytak
Professor of Physics
Interim Department Head

X-Ray Timing of the Accreting Millisecond Pulsar SAX J1808.4−3658

by

Jacob M. Hartman

Submitted to the Department of Physics
on August 31, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

We present a 7 yr timing study of the 2.5 ms X-ray pulsar SAX J1808.4−3658, an X-ray transient with a recurrence time of ≈2 yr, using data from the *Rossi X-ray Timing Explorer* covering 4 transient outbursts (1998–2005). Substantial pulse shape variability, both stochastic and systematic, was observed during each outburst. Analysis of the systematic pulse shape changes suggests that, as an outburst dims, the X-ray "hot spot" on the pulsar surface drifts longitudinally and a second hot spot may appear. The overall pulse shape variability limits the ability to measure spin frequency evolution within a given X-ray outburst (and calls previous $\dot{\nu}$ measurements of this source into question), with typical upper limits of $|\dot{\nu}| \lesssim 2.5 \times 10^{-14}$ Hz s^{-1} (2σ). However, combining data from all the outbursts shows with high (6σ) significance that the pulsar is undergoing long-term spin down at a rate $\dot{\nu} = (-5.6 \pm 2.0) \times 10^{-16}$ Hz s^{-1}, with most of the spin evolution occurring during X-ray quiescence. We discuss the possible contributions of magnetic propeller torques, magnetic dipole radiation, and gravitational radiation to the measured spin down, setting an upper limit of $B < 1.5 \times 10^8$ G for the pulsar's surface dipole magnetic field and $Q < 4.4 \times 10^{36}$ g cm^2 for the mass quadrupole moment. We also measured an orbital period derivative of $\dot{P}_{\rm orb} = (3.5 \pm 0.2) \times 10^{-12}$ s s^{-1}.

We identify a strong anti-correlation between the fractional amplitude of the harmonic (r_2) and the X-ray flux ($f_{\rm x}$) in the persistent pulsations of four sources: SAX J1808.4−3658, IGR J00291+5934, and XTE J1751−305, XTE J1807−294. These sources exhibit a power-law relationship $r_2 \propto f_{\rm x}^{\gamma}$ with slopes ranging from $\gamma = -0.47$ to -0.70. The three other accreting millisecond pulsars that we analyzed, XTE J0929−314, XTE J1814−338, and HETE J1900.1−2455, do not as fully explore a wide range of fluxes, but they too seem to obey a similar relation. We argue that these trends may be evidence of the recession of the accretion disk as the outbursts dim.

We examine the energy dependence of the persistent pulsations and thermonuclear burst oscillations from SAX J1808.4−3658. We confirm the soft phase lags previously discovered from this source, and we discover that these phase lags increase as the source flux decays slowly following its peak flux. When the source decay becomes rapid and the outburst enters its flaring tail stage, this relationship reverses, and the phase lags diminish as the flux dims further. This result, along with the pulse profile changes observed at the beginning of the flairing tail stage, suggests an abrupt change in the geometry of the accretion disk and column at this time in the outburst. In contrast, the thermonuclear burst oscillation timing does not show appreciable lags, and the burst oscillation phases and fractional amplitudes appear to be relatively independent of energy.

Thesis Supervisor: Deepto Chakrabarty
Title: Associate Professor of Physics

Acknowledgments

I can say without exageration that the completion of this PhD has been the most difficult challenge that I have ever undertaken. It would have not been remotely possible without the support of my collaborators at MIT and elsewhere and my wonderful friends and family.

My thanks to my advisor, Deepto Chakrabarty, who suggested the path of research that ultimately led to this thesis. I am a far better researcher thanks to his guidance, in particular his continuous emphasis on the precise presentation of data and scientific results. Thank you to Ed Morgan, whose always-open door and willingness to share his knowledge was invaluable, especially in the understanding and processing of data from the *RXTE*. Many thanks to Duncan Galloway, whose work on the MIT thermonuclear burst database enabled much of my research and who provided me with much help getting started on my research. And finally, thanks to Miriam Krauss, for being a great friend and helping me whenever I needed to understand anything that had to do with spectra.

I am also very grateful to my collaborators beyond MIT, in particular Anna Watts, whose insights into the meaning of my data were always useful; Alessandro Patruno, whose work parallel to mine at the University of Amsterdam provided many important verifications of my analysis, more than once helping me track down bugs; and Paul Ray, who provided critical help with the earliest stages of the research that became this thesis.

One of the most rewarding experiences at MIT was my involvement with the physics community and the many friends that I made there. Thank you to Brian Canavan, without whom the Physics Graduate Student Council would probably never have gotten off the ground. Thanks to Paul Nerenberg, for his friendship and his invaluable assistance, in particular with respect to organization—never one one my strongest points. And thanks to all my other friends who contributed their time to the physics community here at MIT: Andrew Grier, Murat Acar, Bonna Newman, and Dave Guarrera, to name just a few.

Beyond PGSC, thank you to all the new friends I met while at MIT. The weeks of hard work were made so much more tolerable by the "Friday beer" croud: Allyn Dullighan, Judd and Cassie Bowman, Matt Muterspaugh, and Tesla Jeltema. And thanks to my many friends on the on the sixth floor of Building 37, including John Fregeau, Adrienne Juett, and Robyn Sanderson. Many thanks to Jen Trebby, for her friendship and support.

Finally, for keeping me (generally) sane, I thank my friends and family back home. The ongoing online chat thread with my friends from Pennsylvania has been an endless provider of silly diversions, a forum for armchair political punditry, a source of much-needed perspective, and most importantly a way to stay in touch with my oldest friends. My brother, Steve, has been a true friend and confidant. And finally, thank you to my Mom and Dad for their endless love and support. You all have my deepest gratitude.

Contents

List of Figures

List of Tables

Chapter 1

Millisecond X-ray Pulsar Astrophysics

Accretion-powered millisecond pulsars (AMPs) are neutron stars that pulse in the X-ray at periods of less than 10 ms due to the anisotropic accretion of gas from a companion star (Alpar et al. 1982; Radhakrishnan and Srinivasan 1982). This growing class of objects provides a laboratory for an array of physics inaccessible anywhere else in the universe. The X-ray pulse shapes arising from the magnetically channeled accretion flow can constrain the compactness (and hence the equation of state) of the neutron star (e.g., Poutanen and Gierliński 2003; Bhattacharyya et al. 2005). Tracking the arrival times of these X-ray pulses allows us to measure the pulsar spin evolution, which directly probes magnetic disk accretion torque theory in a particularly interesting regime (e.g., Psaltis and Chakrabarty 1999) and also allows exploration of torques arising from other mechanisms such as gravitational wave emission (e.g., Wagoner 1984).

The neutron stars (NSs) in this class of objects are quite old: roughly a billion years have passed since their births in supernovae. A diagram of NS spin periods and period derivatives, shown in Figure 1-1, illustrates the evolution through which these objects are formed. The position on this plot also gives an estimate of the magnetic field: dipole spin down gives $B = 1.0 \times 10^{12}\ (P_0 \dot{P}_{-15})^{1/2}$ G, where P_0 gives the period in seconds and the $\dot{P}_{-15}$ the period derivative in 10^{-15} s s^{-1}. Newly formed neutron stars have spin periods of tens of milliseconds and magnetic fields ranging from 10^{12} G (typical) up to 10^{15} G (magnetars). These large magnetic fields produce large spin downs of $\sim 10^{-13}$, causing them to migrate quickly from the upper-middle region of the diagram to the upper-right, where the majority of known radio pulsars lie. Eventually, the continuing spin down causes them to reach sufficiently long periods ($\gtrsim 10$ s) that the radio emission mechanism shuts off (a threshold whimsically known as the "death line"), and they are no longer detectable.

That would be the end of the story — except that some of these stars are in binary systems. Eventually the companion star may turn off of the main sequence, expanding to fill its Roche lobe and accrete matter onto the NS. The infalling matter forms an accretion disk, and the interaction between the disk and the NS magnetic field spins the neutron star back up into the range of spin periods that emit pulsations (Alpar et al. 1982; Joss and Rappaport 1983). Such sources are aptly referred to as "recycled" pulsars. These systems are low-mass X-ray binaries (LMXBs), as massive companions would have long since evolved off the main sequence by this time. The accreting neutron stars quickly reach an equilibrium spin period, at which the magnetic accretion torques are minimized (Davidson and Ostriker 1973; Lamb

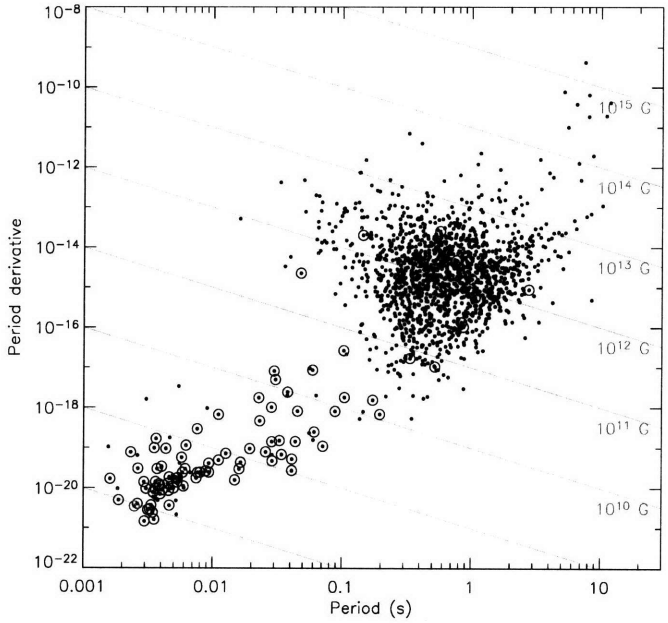

Figure 1-1 The P-$\dot{P}$ diagram, showing 1,640 radio pulsars and the accreting millisecond pulsar SAX J1808.4−3658. The grey diagonal lines show contours of dipolar magnetic field strength following $B = 1.0 \times 10^{12}\ (P_0 \dot{P}_{-15})^{1/2}$ G, ranging from 10^8 G in the lower left to 10^{15} G in the upper right. Binary systems are encircled. The 401 Hz AMP SAX J1808.4−3658 is marked in red. The work in this thesis allows us to place an AMP on the P-$\dot{P}$ diagram for the first time. The radio pulsar data is from the ATNF catalog[1] (Manchester et al. 2005).

et al. 1973). Over $\sim 10^9$ years, the magnetic field attenuates to $\sim 10^8$ G (possibly due to the accretion), shortening the equilibrium period to 10 ms or less and thereby producing an AMP. Eventually the accretion from the companion ceases, leaving the millisecond radio pulsars that occupy the lower left corner of the diagram. This thesis includes the first measurement of the spin period derivative of an AMP, SAX J1808.4−3658, which places it amid the millisecond radio pulsars, just as predicted.

The measurement of the spin period of AMPs relies on the timing analysis of the persistent X-ray emission or of thermonuclear X-ray bursts. In the case of persistent X-ray emission, the magnetic field of the neutron star funnels accreting matter onto hot spots near the magnetic poles, causing these regions to be brighter in the X-ray and producing spin-modulated pulsations. The pulsation mechanism is somewhat less clear for thermonuclear bursts, but again a hot spot on the NS surface produces oscillations (see Strohmayer and Bildsten 2006 for a recent review).[2] We consider sources with detected oscillations during bursts to be AMPs even if they lack persistent pulsations. The two sources that exhibit

[2]A note on jargon. When discussing persistent emission, the literature almost universally uses the term *pulsations*, while coherent variability during thermonuclear bursts is referred to as *oscillations*. Similarly, when an AMP begins actively accreting and brightens in the X-ray, it is in *outburst*, a state that typically lasts of order a month. This term is quite distinct from thermonuclear *bursts*, sudden and intense flares in the X-ray luminosity with timescales of tens of seconds. On behalf of X-ray astronomers everywhere, I extend my personal apologies to the reader.

both of these phenomena (SAX J1808.4−3658, Chakrabarty et al. 2003; and XTE J1814-338, Strohmayer et al. 2003) have the same frequency for both. By extension, the other burst oscillation sources presumably pulse at their spin periods as well.

This thesis studies these two modes of variability observed from millisecond pulsars. In this chapter, we review the physics and observational phenomenology of accretion and thermonuclear bursts. In Chapter 2, we detail our processing of timing data from the *Rossi X-ray Timing Analysis* satellite, which took all the observations discussed in this thesis. In Chapter 3, we review some general methods of timing analysis and their application to AMPs. We also motivate and develop a new minimum-variance estimator for pulse arrival measurement and a Monte Carlo technique for estimating phase model parameter uncertainties in the presence of the pulse shape noise that is common in AMPs. In Chapter 4, we apply these techniques to SAX J1808.4−3658. We make the first measurement of long-term spin down from an AMP and note characteristic pulse profile changes during the outbursts from this source. We discuss the application of these findings to our understanding of the evolution of AMPs. In Chapter 5, we discover a strong correlation between the pulse shapes and the source flux in the persistent pulsations of seven AMPs, and speculate on its implications for the accretion geometry of these systems. In Chapter 6, we report on the energy dependence of the timing of the persistent pulsations and thermonuclear burst oscillations from SAX J1808.4−3658. Finally, in Chapter 7, we summarize our findings and discuss interesting outstanding questions and future directions for timing studies of accretion-powered millisecond pulsars.

1.1 Accretion physics 101

Accretion onto a compact object is far and away the most efficient means of extracting energy from the infalling matter. Compared to the efficiency of fusing protons ($\approx$0.7%), the energy released by matter falling onto a neutron star is $GM/Rc^2 = 21\%$ of the rest mass. Here we have made the assumption of a "canonical neutron star" — $M = 1.4\ M_\odot$ and $R = 10$ km — dimensions that will be evoked repeatedly throughout this thesis. The velocity of the matter as it impacts the surface is thus $\approx\sqrt{2GM/R} \approx \frac{2}{3}c$. Needless to say, this accretion heats up the surface of the star. Accretion rates of $\dot{M} \sim 10^{-10}\ M_\odot\ \mathrm{yr}^{-1}$ produce luminosities of $L_\mathrm{x} \sim 10^{37}\ \mathrm{erg\ s}^{-1}$.

This accretion is not uniformly distributed over the surface of the star. The dipolar magnetic field is sufficiently strong to channel the infalling matter into accretion columns, deflecting its path toward the magnetic poles of the star and producing X-ray hot spots at the base of the columns. Our discussion of accretion theory depends entirely on the interaction between these gravitational and magnetic effects.

As plasma in the accretion disk spirals inward, eventually the growing magnetic pressure will surpass its kinetic energy density. At that point, the plasma will couple more strongly to the magnetic field lines, disrupting the disk and initiating its final fall onto the NS surface. The point at which this transition occurs is known (often interchangeably) as the Alfvén radius or magnetospheric radius. We denote it r_0. The dipolar magnetic field strength at r_0 is $B_0 \approx \mu/r_0^3$, where μ is the dipole moment of the NS. The free-fall radial velocity at this radius is

$$\frac{1}{2}\rho v_\mathrm{r}^2 = \frac{GM\rho}{r_0} \quad \Longrightarrow \quad v_\mathrm{r} = \sqrt{\frac{2GM}{r_0}}, \tag{1.1}$$

where ρ is the density of the plasma at r_0. It can be found by noting that the rate of matter

passing through radius r_0 is $\Omega r_0^2 v_r \rho = \dot{M}$ for an accretion disk of solid angle Ω. Equating the magnetic and kinetic energy densities,

$$\frac{B_A^2}{8\pi} = \frac{1}{2}\rho v^2 , \tag{1.2}$$

we find the Alfvén radius:

$$r_0 = (\Omega/4\pi)^{2/7}(2GM)^{-1/7}\dot{M}^{-2/7}\mu^{4/7} . \tag{1.3}$$

In practice, we drop the leading geometric factor as an order-unity correction.

A greater concern is the width of the transition region from Keplerian orbits to motion along the magnetic field lines. This theory was developed (Pringle and Rees 1972; Lamb et al. 1973; Davidson and Ostriker 1973) with high-field pulsars in mind, for which the magnetic field becomes dominant at distances of $\gtrsim 10^3$ km. In these cases, the Alfvén radius is far greater than both the radius of the star and the width of the transition region, producing a reasonably well-defined inner edge to the accretion disk. This is not at all the case for low-field pulsars. For a canonical NS with $B = 10^8$ G and $\dot{M} = 10^{-8}\ M_\odot\ \mathrm{yr}^{-1}$, we have $r_0 = 30$ km — merely two NS radii above the surface of the star, and most likely similar to the scale of the transition region. These sources require computationally expensive relativistic MHD simulations to more accurately understand the inner edge of the disk, such as those done by Romanova et al. (2004). Nevertheless, the scale of r_0 remains a useful tool for a first-order understanding.

The mass and spin period of the NS introduce three additional length scales. The light cylinder radius is defined by the point at which material can no longer rotate with the same angular frequency as the NS: $r_{lc} = c/2\pi\nu$, where ν is the spin frequency. Within the co-rotation radius, the magnetic field lines are closed; outside, they are necessarily open. The co-rotation radius is the point at which the Keplerian orbital frequency around the star is equal to the spin frequency: $r_{co} = (GM/4\pi^2\nu^2)^{1/3}$. Finally relativistic effects preclude stable orbits within six gravitational radii of a non-rotating object, a point known as the innermost stable circular orbit: $r_{ISCO} = 6GM/c^2$. For a canonical neutron star spinning at 400 Hz, we find $r_{lc} = 120$ km, $r_{co} = 31$ km, and $r_{ISCO} = 12.4$ km.

These simple concepts are all that one needs for a back-of-the-envelope approach to accretion physics. Things get quite messier when one attempts to better understand the precise where's and how's of angular momentum transfer between the magnetic field and the disk, but we sweep such concerns under the carpet by introducing order-unity factors. The exact details are not important for our discussion, and for low-field pulsars are best understood through relativistic MHD simulations — well beyond the scope of this introduction.

Consider the velocity of infalling matter as a function of its distance from the center of the NS. Figure 1-2 illustrates two scenarios, distinguished by having different magnetospheric radii. These configurations could easily describe the same source: Figure 1-2a arises during high $\dot{M}$, during which the high density of infalling matter has pushed r_0 inward; Figure 1-2b may occur during low $\dot{M}$, giving a higher r_0.

In Figure 1-2a, the magnetospheric radius is within the co-rotation radius, resulting in accretion onto the NS surface. Material in the accretion disk enters from the right, losing angular momentum through friction and drifting inward. At the magnetospheric radius, the matter attaches to the magnetic field. The resulting loss of some of its Keplerian orbital velocity imparts a torque to the NS. Finally, the material falls to the surface of the star

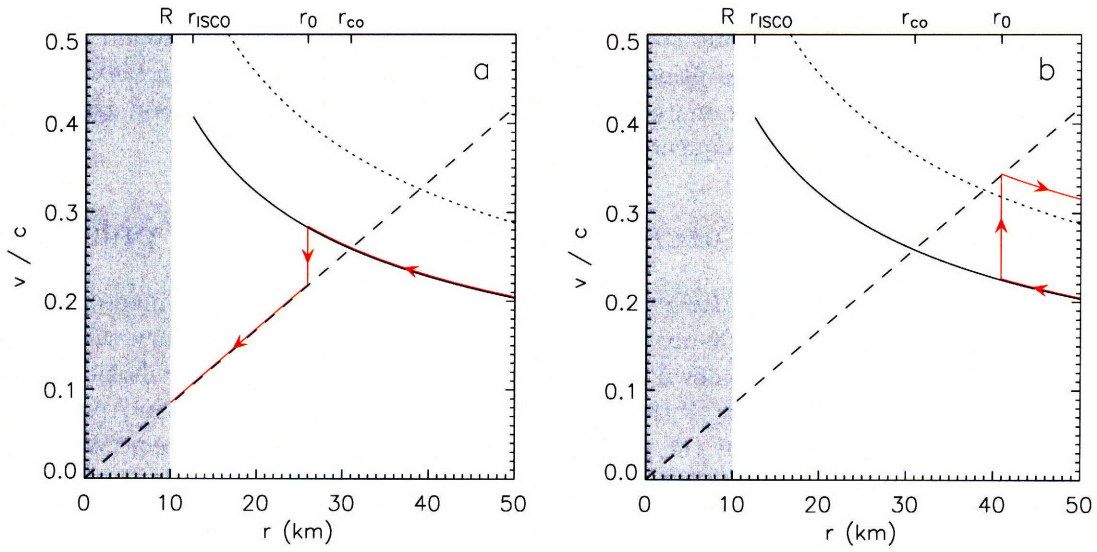

Figure 1-2 A diagram of the velocity vs. radial distance of infalling matter, illustrating the physics of accretion. The solid black lines show the Keplerian velocity of material in a circular orbit. The dashed line shows the co-rotational velocity, which is the velocity at which matter attached to the magnetic field lines travels. The dotted line shows the escape velocity. These curves all assume $M = 1.4\ M_\odot$ and $\nu = 400$ Hz. The NS itself, represented by the grey region, has a radius $R = 10$ km. The two figures are distinguished by their magnetospheric radii: $r_0 < r_{co}$ in (a), and $r_0 > r_{co}$ in (b). This difference decides the fate of matter in the accretion disk, shown by the red paths. These figures assume that the rotational sense of the accretion disk and the NS are the same.

along the co-rotating field lines.

In Figure 1-2b, the magnetospheric radius is outside the co-rotation radius, inhibiting accretion. The magnetic field lines have a greater velocity than the Keplerian orbits at magnetospheric radius, accelerating the matter and pushing it outward at the expense of the angular momentum of the NS. If the co-rotating velocity at r_0 exceeds the escape velocity (as shown in the figure), the magnetic field will fling matter out of the system, a result known as the "propeller effect" (Illarionov and Sunyaev 1975).

Note that the torque imparted to the star is sensitively dependent on the difference between the Keplerian velocity and the co-rotating velocity at r_0. In the conventional picture (Lamb et al. 1973; Davidson and Ostriker 1973), for $r_{co} > r_0$, the NS spins up, decreasing r_{co}; for $r_{co} < r_0$, the star spins down and r_{co} increases. These systems should favor spins near the $r_0 = r_{co}$ equilibrium. Equating these give an equilibrium spin frequency of

$$\begin{aligned} \nu_{eq} &= 0.18\ (GM)^{5/7} \dot{M}^{3/7} \mu^{-6/7} \\ &= 330 \left(\frac{\dot{M}}{10^{-10}\ M_\odot\ \mathrm{yr}^{-1}} \right)^{3/7} \left(\frac{B}{10^8\ \mathrm{G}} \right)^{-6/7} \mathrm{Hz}\,. \end{aligned} \tag{1.4}$$

We have assumed a canonical NS in deriving the second form. Rappaport et al. (2004) modifies the magnetic accretion torque theory for AMPs, allowing accretion to continue

even when $r_{co} < r_0$, but the qualitative picture of the torque reversing sign around $r_0 \sim r_{co}$ remains. Given our crude assumptions and handling of the transition region between the kinetically and magnetically dominated regions, this figure is in surprisingly good agreement with the observed AMP spin frequencies.

1.2 Observations of accretion-powered millisecond pulsars

The launch of the *Rossi X-ray Timing Explorer* (*RXTE*) in December 1995 made possible the observation of millisecond X-ray pulsations from LMXBs, an evolutionary stage long predicted as a precursor to the known millisecond radio pulsars. Millisecond pulsations were first detected during an outburst of the LMXB transient SAX J1808.4−3658 in April 1998: 401 Hz pulsations, modulated by the 2.01 hr orbital period of the binary system (Wijnands and van der Klis 1998; Chakrabarty and Morgan 1998).

Most AMPs are transient sources. Accretion onto the NS is typically episodic due to accretion-disk instabilities. After years of inactivity, during which matter from the companion star collects in the accretion disk, a source will then go into outburst, during which the disk becomes unstable and matter falls to the surface of the NS. During quiescence, the X-ray luminosity is low ($L_x \lesssim 10^{33}$ erg s^{-1}).

In the case of SAX J1808.4−3658, the outbursts last on order a month and recur roughly every two years. None of the other AMPs have undergone multiple outbursts during the lifespan of the *RXTE*, implying generally longer recurrence times of $\gtrsim 10$ yr. (Although XTE J1751−305 briefly brightened for a few days in April 2007; pulsations were not detected.) The lengths of the outbursts from all the sources are of order one month, with the notable exception of HETE J1900.1−2455. That source exhibited a particularly long outburst, beginning in June 2005 and lasting for two nearly years. Generally, pulsations are present throughout the outburst, although in HETE J1900.1−2455 they were intermittent (Galloway et al. 2007).

Figure 1-3 shows the behavior of the pulsations from SAX J1808.4−3658. The pulsations are highly coherent, as shown by the power spectrum on the left. During the 2002 outburst from this source, the measured coherence "quality factor" was $Q \equiv \nu/\Delta\nu \sim 10^5$–$10^6$, decisive evidence that the pulsations are coming from the surface of the NS. The arrival times of these pulsations are modulated by the orbit of the NS, allowing accurate measurements of the orbital parameters of the binary.

The pulse profiles — light curves folded at the pulsation period — have total fractional amplitudes of 1–15% rms. These profiles are quite sinusoidal, with the amplitude of the fundamental typically 2–10 times that of the second harmonic.[3] Higher harmonics are weaker still. These low-amplitude, sinusoidal pulse profiles markedly contrast with the narrower pulse peaks observed in non-accreting pulsars and are evidence of their substantially weaker magnetic fields.

1.3 Observations of thermonuclear X-ray bursts

For a wide range of accretion rates, the accreted gas creates an ocean of compressed hydrogen and helium, approximately 10 m thick, on the surface of the NS. After ∼10 hr of accretion, the density of this layer becomes sufficiently great that it ignites. The unstable burning

[3]Throughout this thesis, we refer to the harmonic at n times the fundamental frequency as the nth harmonic.

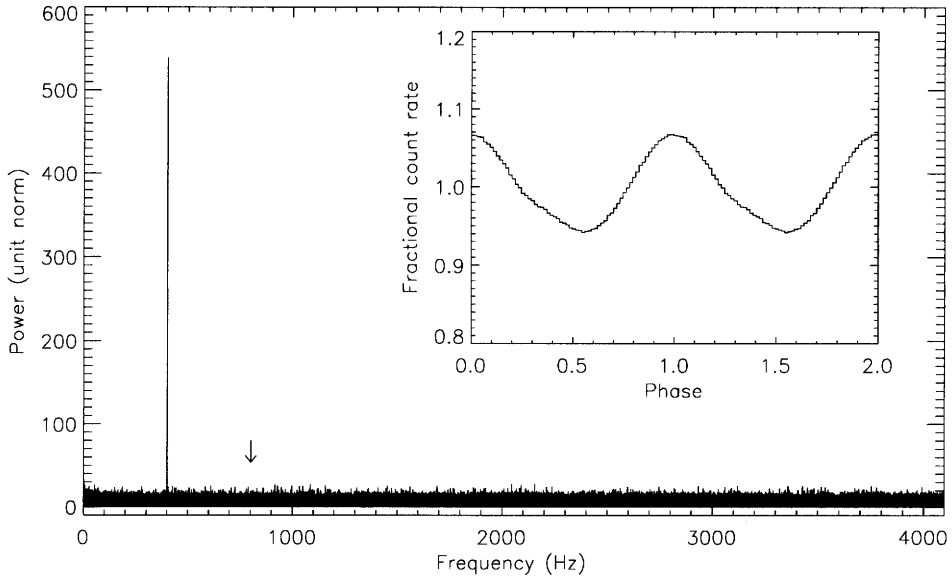

Figure 1-3 An example of a power spectrum and pulse profile from SAX J1808.4−3658, created with data from its 2002 outburst. Notice that the pulse profile is nearly sinusoidal, and the second harmonic (indicated by an arrow) does not even rise above the noise peaks in the power spectrum.

that results quickly spreads over the surface of the NS, burning off all the accreted on timescales of ∼10 s. Temperature anisotropies in the burning layer can produce X-ray oscillations during the bursts at approximately the spin frequency of the star. The detection of such burst oscillations from 4U 1728-34 shortly after the launch of the *RXTE*was the first measurement of the spin of a millisecond spin period from an LMXB. (See Strohmayer and Bildsten 2006 for a recent review.)

The light curves and energy spectra of thermonuclear bursts reveal much about the accreted material. The lightcurves have rapid rises (1–10 s), one or two (and, in rare cases, more) peaks, and roughly exponential decays with e-folding timescales of 3–30 s. These times scales are in good agreement with theoretical expectations (e.g., Woosley et al. 2004). Quick decays are indicative of primarily helium burning via the triple-α process, while longer decays are most likely due to the slower fusion rates of hydrogen burning via the CNO cycle. The spectra of the bursts are reasonably well-modeled by 2–3 keV blackbodies. The blackbody radii are of order 10 km, but they can increase substantially during the peaks of bright bursts that exceed their Eddington limits and cause the NS atmosphere to expand, a phenomenon referred to as photospheric radius expansion (PRE). See Galloway et al. (2006) for a comprehensive analysis of the 1000+ thermonuclear bursts observed by the *RXTE*.

Figure 1-4 shows two thermonuclear bursts that together illustrate many of the characteristic burst behaviors: a non-PRE burst from the 364 Hz oscillator 4U 1728-34, and a PRE burst from the 619 Hz source 4U 1608-52. In (a), the blackbody temperature peaks and decays along with the light curve, while the radius is relatively constant. (The rising $R_{\rm bb}$ toward the end of the burst is due to an increasingly poor blackbody fit due to the

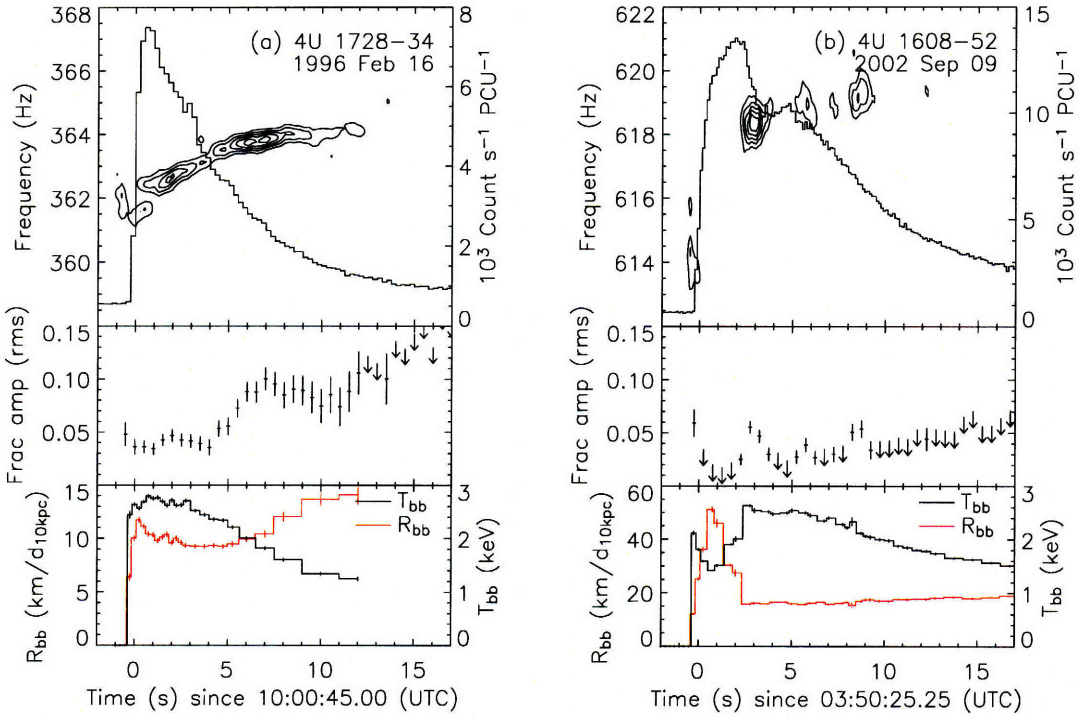

Figure 1-4 The dynamic power spectra, light curves, and blackbody fits of two thermonuclear bursts. In (a), the dynamic power spectra were computed using 2 s overlapping intervals separated by 0.25 s; in (b), 1 s intervals separated by 0.125 s were used. The contours show unit-norm powers of 8, 16, 24, etc. The fractional amplitudes are relative to the count rate above the persistent emission level.

non-blackbody spectral components of the persistent emission, not an actual blackbody radius increase.) In (b), the peak is marked by a drop in the temperature and a coincident rise in radius, the signature of PRE.

The burst oscillations of these bursts also follow the typical pattern. In both bursts, the oscillation frequency increases by $\approx$1% over the course of the burst. The oscillation frequencies from a single source plateau at roughly the same value (Muno et al. 2002), which is believed to be the spin frequency of the NS based on the similar behavior in SAX J1808.4−3658 (Chakrabarty et al. 2003; chapter 6 of this thesis). The fractional amplitudes vary throughout both bursts, sometimes subsiding and then re-appearing later. During PRE, oscillations are almost never detected, as is the case in (b). Fractional amplitudes may be particularly high during the burst rises (as high as 50% rms in some cases; Strohmayer et al. 1998), although that is not the case in the bursts shown here. It is important to note that oscillations are not always present. It is not uncommon for two bursts from the same source that appear quite similar in all other respects to be distinguished by one exhibiting burst oscillations while the other does not.

Our theoretical understanding of this diverse body of burst oscillation phenomenology is poorer than for the burst light curves or energy spectra. It is widely agreed that we are observing a hot spot in the burning layer that rotates (more or less) with the star,

but the details beyond that are murky. A hot spot of localized burning would explain the oscillations, and in particular the high fractional amplitude sometimes observed during the burst rise (Strohmayer et al. 1998), but the burning front moves rapidly to cover the entire star in seconds, raising the question of why a hot spot would persist and produce oscillations during the tail. From the perspective of this localized burning model, the gradual increase in frequency has been explained by hydrostatic contraction of the burning layer as it cools (Strohmayer et al. 1997), but this hypothesis underestimates some of the observed drifts (Cumming et al. 2002). MHD simulations by Spitkovsky et al. (2002) add the contribution of temperature and pressure gradients to explain the hot spot motion; they suggest that the persistence of oscillations during the tails may be due to inhomogeneous cooling. Or perhaps the the large frequency drifts are due to magnetic coupling between the NS and the differentially rotating burning layer (Lovelace et al. 2007). Finally, the "localized burning" picture may be jettisoned altogether during the burst tails in favor of oscillations produced by surface modes in the NS ocean (Piro and Bildsten 2005).

1.4 Table of sources

Table 1.1 lists all the reported AMPs, along with their frequencies, orbital periods (if known), and their observed modes of high-frequency variability. In the past few years, the number of accreting X-ray pulsars has increased dramatically, providing a sample of sufficient size to begin drawing interesting statistical inferences. (See, for instance, Muno et al. 2004 and §1.6.2.)

Four of the spin frequencies listed in Table 1.1 are marked as marginal detections. In my view, these sources should be viewed as "candidate AMPs," at most. They should not be the basis for theoretical work until their spins have been confirmed with second detections. Appendix A details the analysis on which these conclusions are based, as well as a closer look at the reported detection of 45 Hz oscillations in EXO 0748−676.

1.5 SAX J1808.4−3658

Of the above sources, SAX J1808.4−3658 deserves further comment, both because of its key role in the development of our understanding of AMPs and because of its centrality to this thesis.

The X-ray transient SAX J1808.4−3658 was discovered during an outburst in 1996 by the *BeppoSAX* Wide Field Cameras (in 't Zand et al. 1998). Timing analysis of *RXTE* data from a second outburst in 1998 revealed the presence of a 401 Hz (2.5 ms) accreting pulsar in a 2 hr binary (Wijnands and van der Klis 1998; Chakrabarty and Morgan 1998). The source is a recurrent X-ray transient, with subsequent ≈1 month X-ray outbursts detected in 2000, 2002, and 2005; it is the only known accreting millisecond pulsar for which pulsations have been detected during multiple outbursts. Faint quiescent X-ray emission has also been observed between outbursts, although no pulsations were detected (Stella et al. 2000; Campana et al. 2002; Heinke et al. 2007). A source distance of 3.4–3.6 kpc is estimated from X-ray burst properties (in 't Zand et al. 2001; Galloway and Cumming 2006). The pulsar is a weakly magnetized neutron star ($\sim 10^8$ G at the surface; Psaltis and Chakrabarty 1999) while the mass donor is likely an extremely low-mass ($\approx 0.05\ M_\odot$) brown dwarf (Bildsten and Chakrabarty 2001). SAX J1808.4−3658 is the only source known to exhibit all three of the rapid X-ray variability phenomena associated with neutron stars in LMXBs: accretion-

Table 1.1. Reported accretion-powered millisecond pulsars

Source	ν (Hz)[a]	$P_{\rm orb}$ (hr)	Type and references[b]
XTE J1739−285	1122?		N, K (Kaaret et al. 2007)
4U 1608−52	619		N (Hartman et al. 2003), K (Méndez et al. 1998)
GS 1826−238	611?		N (Thompson et al. 2005)
SAX J1750.8−2900	601		N, K (Kaaret et al. 2002)
IGR J00291+5934	598		A (Galloway et al. 2005)
MXB 1743−29	589		N (Strohmayer et al. 1997)
4U 1636−536	581	3.8	N (Strohmayer et al. 1998), K (Wijnands et al. 1997)
3A 1658−298	567		N (Wijnands et al. 2001)
Aql X−1	549	7.11	A (Casella et al. 2007), N, K (Zhang et al. 1998)
A 1744−361	530	19.0	N (Bhattacharyya et al. 2006)
KS 1731−260	524		N (Smith et al. 1997), K (Wijnands and van der Klis 1997)
SAX J1748.9−2021	442	8.7	A (Altamirano et al. 2007)
XTE J1751−305	435	0.71	A (Markwardt et al. 2002)
SAX J1748.9−2021	410?		N (Kaaret et al. 2003)
SAX J1808.4−3658	401	2.01	A, N, K (see §1.5 for references)
HETE J1900.1−2455	377		A (Kaaret et al. 2006)
4U 1728−34	364		N, K (Strohmayer et al. 1996)
4U 1702−429	330		N, K (Markwardt et al. 1999)
XTE J1814−338	314	4.27	A (Markwardt and Swank 2003), N (Strohmayer et al. 2003)
4U 1916−053	270	0.83	N (Galloway et al. 2001)
XTE J1807−294	191	0.68	A (Markwardt et al. 2003), K (Linares et al. 2005)
XTE J0929−314	185	0.73	A (Galloway et al. 2002)
SWIFT J1756.9−2508	182		A (Markwardt et al. 2007)
XB 1254−690	95?		N (Bhattacharyya 2007)
EXO 0748−676	45		N (Villarreal and Strohmayer 2004)

[a]Spin frequency, implied by persistent pulsations, burst oscillations, or both. Question marks denote marginal detections, as explained in the text.

[b]The types of high-frequency variability observed from each source: A = accretion-powered pulsations; N = thermonuclear burst oscillations; K = kilohertz QPOs.

Note. — This table was adopted from Lamb and Boutloukos (2007).

powered millisecond pulsations (Wijnands and van der Klis 1998), millisecond oscillations during thermonuclear X-ray bursts (Chakrabarty et al. 2003), and kilohertz quasi-periodic oscillations (Wijnands et al. 2003).

An optical counterpart has been detected both during outburst (Roche et al. 1998; Giles et al. 1999) and quiescence (Homer et al. 2001). The relatively high optical luminosity during X-ray quiescence has led to speculation that the neutron star may be an active radio pulsar during these intervals (Burderi et al. 2003; Campana et al. 2004), although radio pulsations have not been detected (Burgay et al. 2003). Transient unpulsed radio emission (Gaensler et al. 1999; Rupen et al. 2005) and an infrared excess (Wang et al. 2001; Greenhill et al. 2006), both attributed to synchrotron radiation in an outflow, have been reported during X-ray outbursts.

1.6 How is this interesting?

1.6.1 Constraining the NS equation of state

Apart from their intrinsic astrophysical interest, accreting X-ray pulsars provide a window looking into a particularly extreme corner of the universe that cannot be studied in any other way. First and foremost, the density and temperature regime found at the core of a NS exists nowhere else. Experiments such as the Relativistic Heavy Ion Collider[4], which smashes together gold nuclei to produce quark-gluon plasmas, can achieve super-nuclear densities, but they do so at near-GeV temperatures far more akin to the early universe than to stable matter in the NS core. Furthermore, theoretical attempts to probe such matter are stymied by the uncertainties and computational difficulty of QCD.

However, theory can provide equations of state if some assumptions are made about the species of particles that will exist at such densities, and each equation of state gives a relationship between the NS mass and radius (e.g., Lattimer and Prakash 2001). These can be constrained by observation. The most direct measurement of R/M is the observation of gravitationally red-shifted emission line from the NS surface. Unfortunately, the only reported measurement of this kind (Cottam et al. 2002) is somewhat uncertain and has not been seen in subsequent observations.

X-ray timing offers alternative (if somewhat less straightforward) methods of measuring the equation of state. The X-ray pulse profiles depend on the compactness of the NS as well as the geometry of the emitting regions. Attempts to model the pulse profiles of burst oscillations (Muno et al. 2002) and persistent pulsations (Poutanen and Gierliński 2003; Bhattacharyya et al. 2005) have yielded some limits on R/M but have not been particularly constraining. Additionally, Psaltis and Chakrabarty (1999) present a lower bound on M/R^3 for a source given an observed range of accretion rates over which pulsations are detected. However, this method is also very model-dependent, in this case relying on our limited understanding interaction of the interaction between the NS magnetic field and the accretion disk.

1.6.2 Searching for gravitational radiation

The distribution of AMP spins presents an intriguing puzzle. For a spherically symmetric, non-magnetized NS, the spin up due to accretion torques is limited only by the centrifugal break-up frequency of the star: $\nu_{\max} = \frac{1}{2\pi}(GM/R^3)^{1/2} = 2.2$ kHz for a canonical NS. The

[4] `http://www.bnl.gov/rhic/`

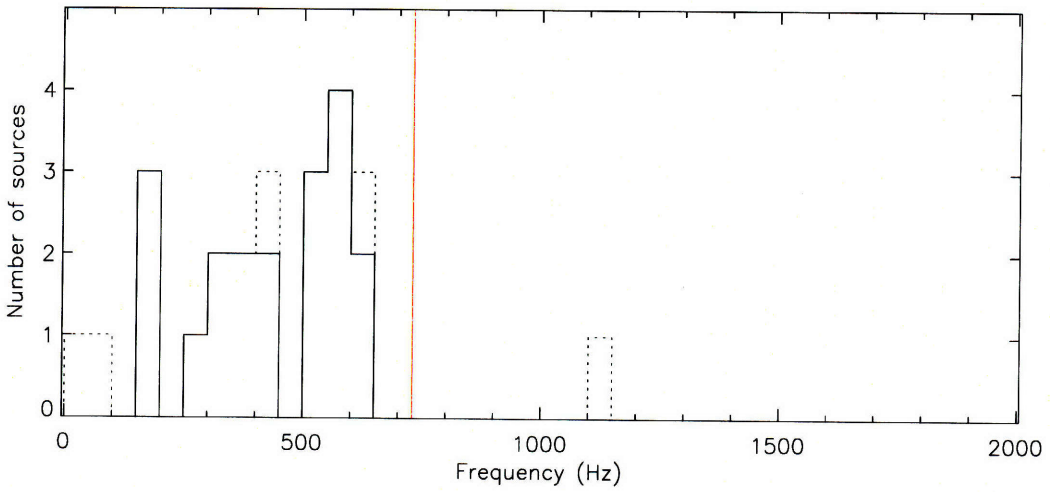

Figure 1-5 The distribution of observed spin frequencies of AMPs. The solid histogram includes all the sources for which a spin has been definitively measured; the dotted histogram also includes the measurements marked as questionable in Table 1.1. The vertical red line marks the observed spin frequency limit of 730 Hz, while the frequency axis extends all the way to the centrifugal limit of ≈2 kHz.

observed cut-off of the spins is far lower, as is clear from Figure 1-5. Bayesian analysis of the observed distribution yields an upper limit of 730 Hz (95% confidence; Chakrabarty 2005).

It is possible that this frequency cut off is due to a cut off in the distribution of NS magnetic fields (compare eq. [1.4]). The mechanism by which the magnetic fields decay is poorly understood, and it is possible that this mechanism abruptly becomes far less efficient once the field drops below $\sim 10^8$ G. Alternatively, gravitational radiation may provide a more natural explanation. For a star with a mass quadrupole Q, the spin-down torque due to the emission of gravitational radiation is $N_{\rm gr} = -\frac{32}{5} G Q^2 (2\pi\nu/c)^5$. This strong dependence on ν could cause the observed cut off.

A variety of mechanisms have been proposed in which rapidly rotating neutron stars can develop mass quadrupoles that give rise to gravitational radiation from the neutron star itself. These mechanisms include r-mode instabilities (Wagoner 1984; Andersson et al. 1999), accretion-induced variations in the density of the NS crust (Bildsten 1998; Ushomirsky et al. 2000), distortion of the NS due to toroidal magnetic fields (Cutler 2002), and magnetically confined mountains at the magnetic poles (Melatos and Payne 2005). Persistent x-ray pulsations can provide an accurate phase ephemeris, making feasible highly targeted searches using interferometric gravitational wave detectors.

Chapter 2

Timing analysis with the *Rossi X-ray Timing Explorer*

The *Rossi X-ray Timing Explorer* (*RXTE*; Bradt et al. 1993) satellite provides an unprecedented ability to measure the precise timing of photons from the X-ray sky. Launched into a low-Earth orbit in December 1995 and still operating as of this writing (far longer than its 5 year mission goal), the *RXTE* has proven a great success in measuring the rapid variability of X-ray sources. Its two pointed instruments, the Proportional Counter Array (PCA; Jahoda et al. 1996) and the High Energy X-ray Timing Experiment (HEXTE; Rothschild et al. 1998), record the arrival times of photons over a wide energy range with microsecond accuracy, while its All Sky Monitor (ASM) keeps watch for new X-ray transients and changes in variable sources. Its effective collecting area of $\approx$6000 cm^2 still greatly exceeds that of any other X-ray satellite in the 2.5–25 keV range. Figure 2-1 shows this satellite and its instruments.

2.1 The *RXTE* Proportional Counter Array

Of foremost importance to X-ray pulsar timing is the PCA. The PCA comprises five Xenon-filled proportional counter units (PCUs). The PCUs are sensitive to 2.5–60 keV photons and each had an effective area at the time of launch of 1300 cm^2 in the 2.5–15 keV range. The PCA has a 1° field of view, with each PCU aligned to within 5′ of the spacecraft science axis.

Figure 2-2 shows a schematic view of a PCU. The collimator and shielding restrict the angle at which an X-ray can enter the detector. An X-ray sufficiently aligned with the instrument axis will then pass through a thin mylar sheet to enter the gas-filled body of the detector. It will collide with and ionize a gas atom, and the liberated electron will collide other atoms, ionizing them as well. The cloud of electrons resulting from this avalanche will be attracted toward one of the anode wires running through the detector. For a properly chosen anode voltage (in this case, $\approx$+2000 V), the number of ionization events will be roughly proportional to the energy of the original photon. The resulting current is measured by a pulse-height analyzer, allowing the creation of spectra with an energy resolution of approximately 15% (at 7 keV). The anode wires are grouped into five layers: an anti-coincidence layer in front, surrounded by propane gas and separated from the other wires by a second mylar sheet, to eliminate non–X-ray events; three intermediate signal layers of wires to detect 2.5–60 keV X-rays; and farthest from the entry window and

on the sides, another anti-coincidence layer. "Good" events will be detected by only a single anode of a signal layer. The process of discharge and read-out takes approximately 10 μs, causing this amount of dead time during which later photons will not be recorded.

Exposure to the harsh space environment has somewhat degraded the capabilities of the PCA. (See Jahoda et al. 2006 for complete details.) Not long after launch, three of the PCUs began occasionally experiencing high-voltage breakdowns. To mitigate this problem, the instrument team stepped down the anode voltage between launch and May 2000, resulting in five different gain epochs, each with progressively harder photon responses. More critically for X-ray timing, the PCUs known to discharge are routinely cycled on and off. Figure 2-3 shows the result: through the lifetime of the mission, the average number of active PCUs — and hence the effective area of the PCA — has decreased markedly. Additionally, two of the PCUs have lost the propane in their outer anti-coincidence layers (PCU 0 in May 2000, PCU 1 in Dec 2006), causing increased background count rates in these PCUs.[1]

Each event (both X-rays and background events) from the PCA is sent to the Experiment Data System (EDS) for processing before its transmission down to Earth. The EDS comprises eight independently operating Event Analyzers.[2] Six of the EAs receive identical streams of data from the PCA; the other two receive the data from the ASM. Two of the PCA EAs compress the data stream into low-bandwidth "standard" data products. Standard 1 products providing 1/8 s resolution light curves from each PCU; Standard 2 products give far lower time resolution (typically 16 s) but provide spectra from each of the

[1]See the "PCA Digest" at `http://heasarc.gsfc.nasa.gov/docs/xte/pca_news.html` for the latest trials and tribulations of this instrument.

[2]At the heart of each EA is an 80C286 microprocessor (running at a whopping 8 MHz and equipped with 288 KB of RAM) and an equally antiquated Texas Instruments 320C25 digital signal processor. Please refer to the EDS Hardware Functional Description and Requirements, written by D. Gordon, at `http://xte.mit.edu/flight/eds.ps` for more details. It's a pretty fascinating read.

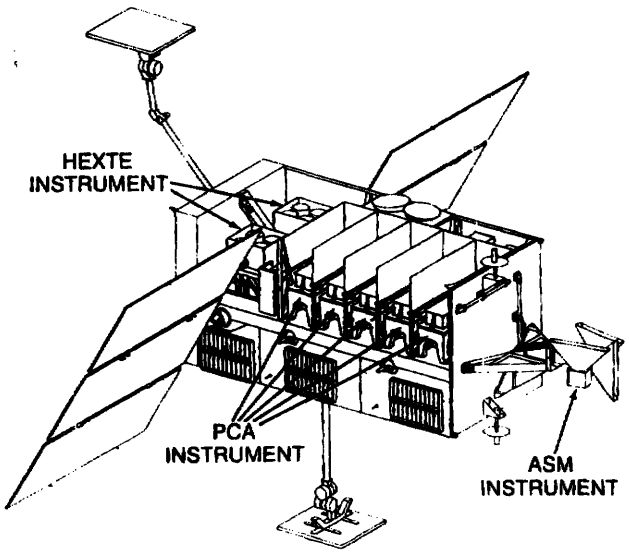

Figure 2-1 A cut-away view of the *RXTE* satellite, showing its three science instruments. Figure from Glasser et al. (1994).

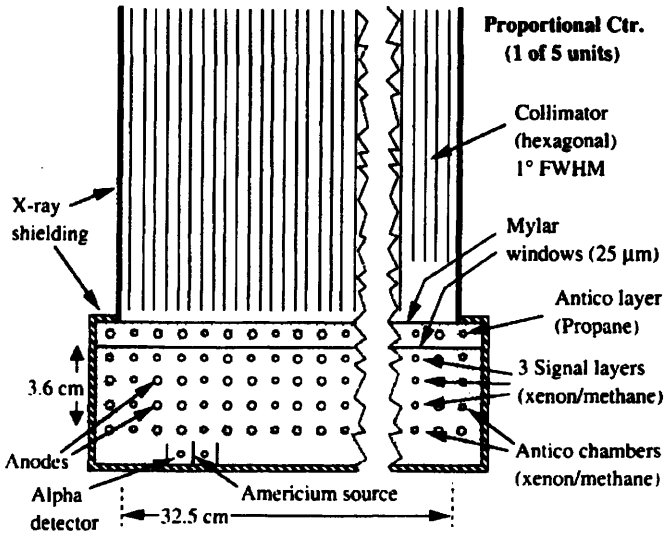

Figure 2-2 A cut-away view of a proportional counter unit. Figure from Bradt et al. (1993).

PCUs. The remaining four EAs devoted to the PCA are user programmable.[3]

The choice of data mode must weigh the attributes of the source (count rate, spectrum, potential for bursts) and the available telemetry, typically limited to 40 kbit s^{-1}. If the data products exceed the available telemetry, the buffers will overflow and data will be lost. Common options include the `GoodXenon` mode, which telemeters to Earth every event recorded by the PCA with 1 μs timing accuracy, but quickly runs into bandwidth limitations if the count rates get too high. Event modes also transmit every PCA event, but with less timing or spectral precision. The most common mode for observing accreting millisecond pulsars is `E_125us_64M_0_1s`, which provides 122 μs timing precision and 64 spectral channels distributed over the entire 2.5–60 keV range.) Single-bit configurations forgo all event details, simply recording whether or not a photon within a specified energy range arrives in each time bin. Since these configurations do not overflow, they are particularly useful for the timing of oscillations during thermonuclear bursts. Similarly, burst-catcher modes are quite useful for studying thermonuclear bursts, as they only take data when an EA detects that a burst is present. By taking only short intervals of data, they can capture more detail even when the count rate is high.

2.2 Processing *RXTE* timing data

Before timing data from the *RXTE* are ready for analysis, a few processing steps must take place. The importance of each step is quite dependent on the ultimate application: barycentering — let alone microsecond-level clock corrections — is of little concern when analyzing a 30 second burst, but is critical for multi-day analysis. Still, the following recipe remains relatively fixed throughout the work in this thesis. It is implemented using

[3]See the *RXTE* Technical Appendix for a thorough discussion of the EDS modes, as well as additional information on the science instruments. `http://heasarc.gsfc.nasa.gov/docs/xte/appendix_f.html`

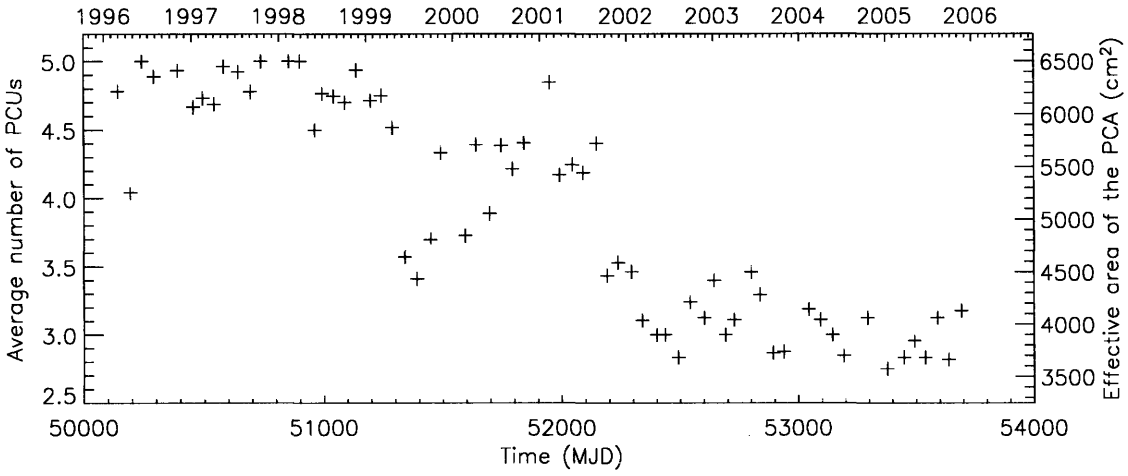

Figure 2-3 The average number of active PCUs over the *RXTE* mission. Each point shows the 50 day mean number of active PCUs during the observations of thermonuclear bursts in the MIT burst database (Galloway et al. 2006). The effective areas shown on the right axis assumes a constant 1300 cm^2 per PCU; it does not account for the decreasing anode voltages and other effects that further degrade the effective area of this instrument.

MIT's proprietary DS ("data stream") suite of processing tools for *RXTE* data.[4] Each DS program performs a single task, and complex processing pipelines can be built on the UNIX command line by piping the results of one program into the next; hence the name.

1. **Conversion of raw data products; energy cuts.** Because MIT built the EDS, we have tools to handle its native data formats and download them from Goddard Space Flight Center (where the *RXTE* control center is located) rather than starting with the FITS files used by most *RXTE* scientists. Thus the first step of processing is to convert these raw data files into the DS file format. When converting files from data modes with multiple energy channels, we typically combine the counts from select channels into a single number of counts per time bin. For persistent pulsation timing, we consistently used an energy cut of 2–15 keV. (During the later gain epochs, the 2 keV lower energy limit includes the lowest energy channel. Since this channel has a relatively high background rate, it was never included, pushing the lower bound up to ≈2.5 keV.) For thermonuclear burst studies, we sometimes used a softer range, such as 2–10 keV. Data product conversion and energy channel selection were performed with the `xenon2ds` or `E2ds` for event-mode data, and `CB2ds` or `sb2ds` for burst-catcher or single-bin modes.

2. **Filtering out bad time intervals.** Next, the data were filtered to remove Earth occultations (using the `ds_clean` tool) and times when the *RXTE* science axis was not pointed sufficiently close to the source (using `ds_filter`, typically removing data when the sensitivity to source photons fell below 60% of the on-axis sensitivity). In a

[4]See `/csr/xte/doc.ours/ds_README` on the `space.mit.edu` network for a description of the standard DS tools. Unfortunately, this document is not available on the Internet.

very few cases, this alignment requirement was relaxed to maximize the timing data taken during raster observations, in which the satellite is continuously panning across the sky.

3. **Barycentering / geocentering.** For the meaningful timing analysis of everything but short time intervals, we must correct for the time delay due to the Earth's orbit around the Sun and the *RXTE*'s 96 min orbit around the Earth. The motion of the Earth will shift the frequency of a 400 Hz source by as much as 0.04 Hz, while the motion of the *RXTE* will cause 0.01 Hz shifts and rapid decoherence times of $\approx$5 min. (See Appendix B for a derivation of the the timing errors caused by improper barycentering.) For timing analysis done solely with the DS tools, the photon arrival times were shifted to the solar system barycenter (using `ds_bary`). If the timing data were passed on to the TEMPO pulsar analysis program, the data were only shifted to the geocenter (via `ds_geo`), and TEMPO performed the barycentering.

4. **Fine-clock corrections.** Finally, the timing data were adjusted to correct for small variations in the rate of the on-board *RXTE* clock (`ds_fineclock`). This step provides absolute time measurements with errors of less than 3.4 μs (99% confidence; Jahoda et al. 2006).

Following the above steps, the data are ready for standard timing techniques, which will be discussed in the next chapter.

2.3 Sampling function of the *RXTE* observations

The observation of a source by the *RXTE* is far from continuous. Understanding the sampling function of the observations becomes important for multi-day timing studies, as the sampling can introduce periodicities in the data. In particular, these effects must be carefully considered during long-term frequency measurements, as we will see in §3.3.2.

Within each observation of longer than $\approx$4 ks, there will also be gaps due to the low-Earth orbit of *RXTE*. Observations are usually occulted by the Earth for $\approx$35 minutes once every 96 minute orbit. There will also be gaps if the orbit passes through the South Atlantic Anomaly (SAA), a region off the east coast of South America in which Earth's radiation belts dips down to the 580 km altitude at which the *RXTE* orbits. The enhanced proton fluxes wreak havoc on the high-voltage proportional counters, so the PCA must be shut off for 10–15 min.

Additionally, the schedule of the *RXTE* is almost never devoted to a single target for a long stretch of time, even during the outburst of transient sources. For instance, during the densely observed 2002 outburst of SAX J1808.4−3658, the 2–20 ks pointings at SAX J1808.4−3658 were interspersed with 1–15 ks pointings for the high-mass X-ray binary Cygnus X-1, the X-ray cluster RX J0658−5557, the dwarf nova WW Ceti, etc. To optimize scheduling, the observations of a source typically occur at roughly the same time each day such that that Earth occultations coincide with times of SAA passage.

Figure 2-4 shows a power spectrum of the resulting sampling function during the 2005 June observations of SAX J1808.4−3658. The 1 d and 96 min peaks due to the scheduling and orbital periodicities are clear, as are some of their harmonics. At periods shorter than 96 min, the power is low, because there are rarely gaps in the data at such short timescales. At periods longer than a few days, there may be appreciably more power due to changes

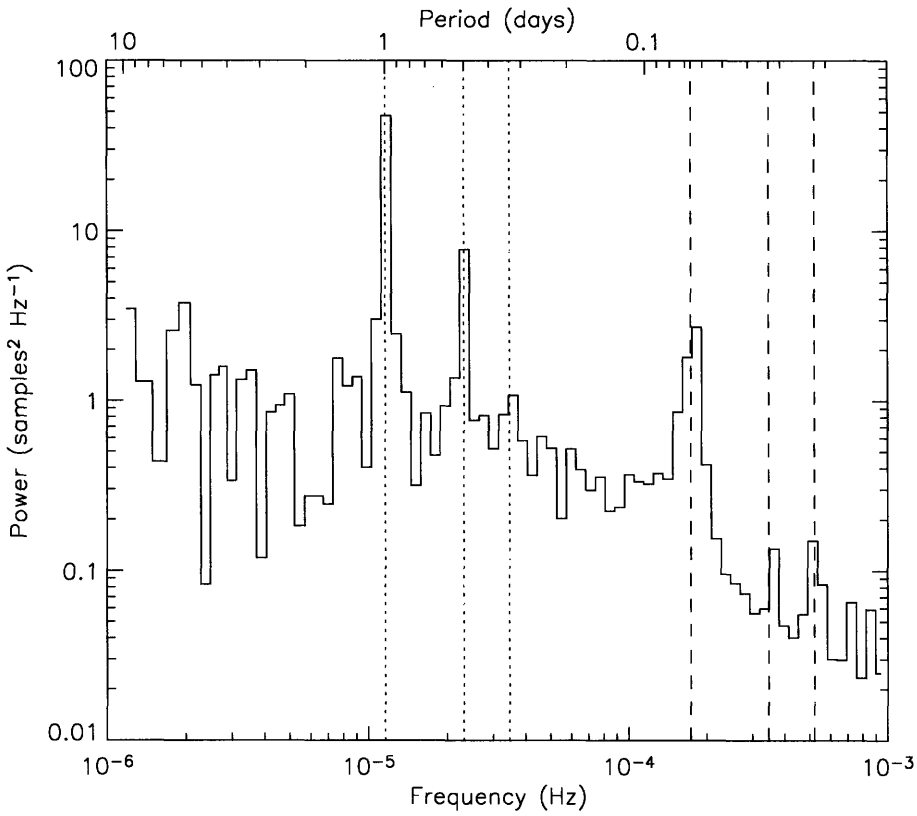

Figure 2-4 Power spectrum of the *RXTE* sampling function for the 2005 outburst from SAX J1808.4−3658. The vertical dotted lines mark the 1 d periodicity of the observation schedule and its 2nd and 3rd harmonics; the dashed lines mark the 96 min orbital period and its harmonics.

in the frequency of observation. (For instance, a source in outburst may be observed with greater frequency during the outburst peak than during the tail.)

2.4 Systems of time

Given the accuracy with which we are measuring time throughout this thesis, it is critical to understand the units with which it is measured. There are as many systems of time in use as there are ways to measure it, and the application of an improper time system can wreak havoc on a sensitive analysis.[5] The following table outlines the time systems in common use in X-ray astronomy. (Table 2.1 and the subsequent discussion draw from Seidelmann 1992.)

All but MET are expressed in days — in X-ray timing, most often Modified Julian Days (MJD), which defines its zero value at 1858 Nov 11 00:00:00 at the Greenwich meridian. (This date is significant only in its relation to the Julian date: MJD = JD - 2400000.5, which has its zero at noon on 4713 BCE Jan 1.) The *RXTE* MET is expressed in seconds

[5] I once spent three weeks tracking down a bug that ultimately was due to one program speaking UTC and another speaking TT. For the sake of future researchers, *please* thoroughly document your code.

Table 2.1. Systems of time used in X-ray astronomy

International Atomic Time	TAI
Coordinated Universal Time	UTC = TAI − leap seconds
Terrestrial Time	TT = TAI + 32.184 s
RXTE Mission Elapsed Time	MET = (TT − MJD 49353.000696574) × 86400 s d^{-1}
Barycentric Dynamical Time	TDB $\approx$ TT + $0^s.001658 \sin g + 0^s.000014 \sin 2g$ (where $g = 357.^\circ53 + 0.^\circ9856003$ (MJD − 51545) d^{-1})
Barycentric Coordinate Time	TCB = TDB + L_B × (MJD − 43144.5) × 86400 (where $L_B = 1.550505 \times 10^{-8}$)

since 1994 Jan 1 00:00:00 UTC.

TAI is an international standard giving the most precise available measurement of time, an average based on a large number of atomic clocks. UTC is the basis for civil timekeeping and is equivalent to the Greenwich Mean Time from which time zones are referenced. As such, it is tied to the rotation of the Earth. It differs from TAI by an integral number of leap seconds that are occasionally introduced to keep UTC synchronized to within 0.9 s of the observationally determined phase of the rotation of Earth.[6]

The barycentric times require special mention. The parametrization of ephemerides requires a time system not affected by the orbit of Earth. A clock on Earth will measure time at a periodically varying rate relative to a clock at the solar system barycenter due to the changing velocity of Earth as it moves between perihelion and apihelion. The periodic correction given in Table 2.1 accounts for this effect. (While this form of the correction is approximate, it is suitable for our accuracies; refer to Seidelmann 1992 for the full expression.) This correction is due entirely to special relativity. TCB goes one step further, also correcting for the slowing of clock rates due to our position within the gravitational well of the Sun. (It essentially improves upon the SI definition of a second, based on the frequency of the hyperfine transition of the cesium-133 atom at absolute zero, by placing the atom outside of a gravitational potential.) However, most pulsar timing to date uses TDB, a practice that we adopt throughout this thesis.

[6]Refer to `ftp://maia.usno.navy.mil/ser7/tai-utc.dat` for a current list of leap seconds.

Chapter 3

Techniques of timing X-ray pulsars

All the work in this thesis depends on the identification and characterization of oscillations in X-ray light curves. For the purposes of this brief review, we divide X-ray timing analysis into three general techniques: phase folding, in which a phase histogram is constructed from a light curve given a known oscillation frequency; Fourier analysis, which measures the periodic content at all frequencies; and phase connection, which combines aspects of phase folding and Fourier analysis to very precisely determine the frequency evolution of a signal. Finally, we discuss the application of these techniques to pulsations from AMPs and develop a method for reducing the effects of various types of timing noise present in these sources.

3.1 Time-domain analysis: phase folding

Time-domain analysis is the most straightforward method of measuring the X-ray pulsations. Given a phase model $\phi(t)$ that accounts for the spin of the NS and the orbit of the system, we can create a folded light curve by assigning each photon arriving at time t_i to one of $j = 1 \ldots n$ phase bins:

$$j = \left(\text{floor}\left[\, n\phi(t_i) \,\right] \bmod n\right) + 1\,. \tag{3.1}$$

The photon counts in each phase bin produce a pulse profile, such as the one shown in the inset of Figure 1-3.

Such analysis is very useful in studying the nature of the pulse profile and observing how it changes over time. If a pulse profile is reasonably stable, a "template profile" can be constructed by integrating a folded light curve over long periods of time. Any phase variation in the pulses can then be found by convolving shorter observed profiles with the template profile and measuring the offsets. (See Taylor 1993, Appendix A, for a detailed discussion of such analysis.) This technique is widely used in the analysis of rotation-powered pulsars, both in the radio and X-ray bands. However, it fails for AMPs, because their pulse profiles change markedly during the outbursts. Section 3.3.1 discusses how to measure phase offsets in this situation.

3.2 Frequency-domain analysis: Fourier transforms

Fourier analysis is the natural tool for analyzing periodic time series, such as the light curves of X-ray pulsars. In this section, we attempt only to provide an overview of this powerful technique, in particular of the deceptively complex topic of its application in estimating power spectra. For an excellent general review, see Press et al. (1992); for the application of these techniques to X-ray timing, we suggest van der Klis (1989).

Consider a series of sequential photon counts x_j, $j = 1 \ldots n$, in uniformly spaced bins of length Δt. The discrete Fourier transform of this series is

$$a_k = \sum_{j=1}^{n} x_j \exp\left(2\pi i j k/n\right) . \tag{3.2}$$

Each a_k measures the complex amplitude of our series in a basis of sinusoids with frequencies $\nu_k = k/n\Delta t = k/T$, where T is the length of our data. These a_k completely describe the original data; Fourier transforms are indeed transforms into a frequency basis, not simply projections. They can be considered as a real oscillation amplitude and a phase: $a_k = |a_k| \exp \phi_k$.

For real data, the a_k for $k = 0 \ldots n/2$ will be independent. a_0 simply measures the DC component of the data, while $a_{n/2}$ measures the most rapidly oscillating data that can be captured in n bins: alternating peaks and troughs. (It too is always real, as you can either alternate "up, down" or "down, up"; there is no phase.) This maximum frequency that can be represented by the data is the Nyquist frequency, $\nu_{\rm Nyq} = 1/2\Delta t$. The negative frequency components are not independent from their positive counterparts: $a_{-k} = a_k^*$. Counting the 2 independent real and $(n-1)/2$ independent complex components, we have n values, as needed to preserve the information in the original series.

The variables adopted here (Δt, ν) are suggestive of the application of the Fourier transform to time-domain data. For instance, *RXTE* photon events with time resolutions of 122 μs may be binned into a light curve of this same resolution, giving $\nu_{\rm Nyq} = 4096$ Hz. But we could just as meaningfully transform a *folded* light curve of x_j phase bins, as defined in equation (3.1), producing a_k that measure the harmonics of our pulse profile. Here, the number of bins n will of course be far smaller, and the Nyquist limit will impose a maximum harmonic to which we are sensitive. This application will be discussed in §3.3.

In our analysis of X-ray light curves, we most commonly use Fourier transforms to estimate the power spectrum of the time series. For a light curve with $N_{\rm ph}$ photons, the spectral power at frequency ν_k is

$$P_k = \frac{|a_k|^2}{N_{\rm ph}} . \tag{3.3}$$

Even in the absence of oscillations, P_k will not be zero; fluctuations due to Poisson noise will produce random oscillations. These noise powers follow an exponential distribution: the probability that random fluctuations will cause a power of at least P_k is $\exp(-P_k)$. As a result, the noise powers have a mean of 1. We refer to this normalization as the unit-norm powers, in contrast with Leahy norm powers, which are greater by a factor of 2 (Leahy et al. 1983).

We often calculate power spectra to search for oscillations. In doing so, we typically consider a large number of frequency bins, each of which is an independent test of the power. Say we look at the powers P_k of $N_{\rm tr}$ frequencies. The probability that none of these powers

will exceed some threshold $P_{\rm thresh}$ is

$$\begin{aligned} \mathrm{Prob}(P_k < P_{\rm thresh} \text{ for all } k) &= [1-\exp(-P_{\rm thresh})]^{N_{\rm tr}} \\ &\approx 1 - N_{\rm tr}\exp(-P_{\rm thresh})\,; \end{aligned} \tag{3.4}$$

conversely, the significance of a peak with power P is $N_{\rm tr}\exp(-P)$. A search of T seconds of data over a frequency bandwidth $\Delta\nu$ results in $T\Delta\nu$ independent trials.

A typical technique in the analysis of thermonuclear bursts is the construction of dynamic power spectra. Here we divide data of length T into $N_{\rm int}$ non-overlapping time intervals of length $T/N_{\rm int}$ and calculate the power spectrum of each. Again, the total number of independent trials will be $T\Delta\nu$ regardless of $N_{\rm int}$; choosing a large number of intervals will provide time resolution at the expense of frequency resolution, and vice versa. Because burst oscillations are virtually never continuously present at a single frequency throughout a burst,[1] trading some frequency resolution for time resolution is usually an advantageous search strategy. Dynamic power spectra also enable us to observe how the oscillation frequency changes over time (e.g., Fig. 1-4).

To improve the sensitivity of such searches, we often oversample in both the time and frequency. Time oversampling is straightforward: it amounts to using overlapping time intervals for the calculation of power spectra. Consider a short-lived burst oscillation that is present only for the length of a single time interval. If our time intervals are non-overlapping, the power of this oscillation will generally be split between two intervals, often greatly diminishing its chance of detection. Using overlapping intervals greatly reduces the chance that this will happen.

To oversample the frequency of a power spectrum, we sample at frequencies between the $\nu_k = k/T$. Because they analyze a finite length T of data, the power spectra estimated by Fourier analysis do not have a uniform frequency response. The power from an oscillation at a frequency of precisely ν_k will all be in bin k, but in the general case the power at some intermediate $\nu_{k+\epsilon}$ will divided among its surrounding bins and thus weakened. The remedy is simple enough: we sample these intermediate frequencies.

$$\begin{aligned} a_{k+\epsilon} &= \sum_{j=1}^{n} x_j \exp[2\pi i j(k+\epsilon)/n] \\ &= \sum_{j=1}^{n} [x_j \exp(2\pi i j\epsilon/n)] \exp(2\pi i j k/n)\,, \end{aligned} \tag{3.5}$$

the transform of the original x_j rotated by the complex phases $\exp(2\pi i j\epsilon/n)$. In practice, we evenly subdivide the frequency intervals by a factor of m (i.e., $\epsilon = l/m$ for $l = 0 \ldots m-1$), requiring m transforms. Alternatively, we can write

$$\begin{aligned} a_{k+l/m} &= \sum_{j=1}^{n} x_j \exp[2\pi i j(k+l/m)/n] \\ &= \sum_{j=1}^{mn} x'_j \exp[2\pi i j(mk+l)/mn)]\,, \end{aligned} \tag{3.6}$$

where $x'_j = x_j$ for $j < n$ and zero elsewhere. This form is a length-mn Fourier transform in

[1] The only case that comes to mind are the burst oscillations of XTE J1814−338 (Strohmayer et al. 2003).

which our original data has been padded by zeroes. It is somewhat slower ($\mathcal{O}[mn \log mn]$ vs. $\mathcal{O}[mn \log n]$) but generally easier to implement.

These two techniques augment the resolution of the dynamic power spectra, but at the expense of increasing the number of trials. Choosing a degree of oversampling involves balancing this with the improved sensitivity. In general, little additional power is gained by oversampling the frequency by a factor greater than $m = 6$. Similarly, overlapping intervals of length T hits the point of diminishing returns for step sizes much less than $T/8$.

These techniques also sacrifice independent sampling, which makes assessing the number of trials far more difficult. One must often resort to computationally expensive Monte Carlo simulations to assess the confidence in observed peaks. Appendix A discusses this issue in the context of evaluating the reliability of reported burst oscillation sources.

3.3 Phase connection

Phase connection is the most accurate way to measure the phase (or, equivalently, frequency) evolution of persistent pulsations. We first fold intervals of timing data according to an *a priori* phase timing model to obtain pulse profiles, as described in §3.1. We next compare the profile from each interval to a template profile in order to calculate the offsets between the observed and the predicted pulse times of arrival (TOAs). We then improved the phase model by fitting it to these TOA residuals. This procedure of phase model bootstrapping is the basis for all long-term pulsar timing.

For all the persistent pulsation analysis detailed in this thesis, we used the TEMPO pulsar timing program[2], version 11.005, to calculate pulse arrival times from a phase model and to improve a phase model by fitting it to arrival time residuals. TEMPO reads in a list of TOAs and a set of parameters describing the pulsar timing model. It then adjusts the model to minimize the timing residuals between the predicted and observed arrival times. The output files include a revised timing model, a covariance matrix for the fit parameters, and a list of the timing residuals. TEMPO also includes a predictive mode, which takes a timing model and generates a series of polynomial expansions that give the model's pulse arrival times during a specified time interval. TEMPO has been a standard tool of the radio pulsar community for decades and is well-tested at the microsecond-level accuracies with which we are measuring TOAs.

This section discusses the particulars of TOA measurement and phase model parameter fitting in the presence of pulse shape noise (both random and systematic) and timing noise with long-timescale correlations. These methods were developed for the timing of the four outbursts from SAX J1808.4−3658 (discussed in the next chapter), but they are of general application to the timing of AMPs and have been applied to all eight known sources with persistent pulsation.

3.3.1 TOA calculation in the presence of profile noise

Special care must be taken when measuring the pulse TOAs for rapidly rotating accretion-powered pulsars. In these systems, the pulse profiles exhibit variability well in excess of the Poisson noise expected from counting statistics. In this section, we develop a procedure to obtain a minimum-variance estimate of the timing residuals in the presence of such noise.

[2] `http://www.atnf.csiro.au/research/pulsar/tempo/`

To calculate the TOAs, we typically divided the timing data into 512 s intervals, with one pulse arrival time per interval. We chose this length because it provides sufficient counts to make accurate measurements during the dim tails of the outbursts, while it still is short enough to sample within the 2400 s orbital periods of the tightest AMP binaries. This is necessary to improve the binary model and resolve any additional short-timescale variability.

We then used TEMPO's predictive mode to generate a series of polynomial expansions predicting the times of pulse arrivals at the geocenter. Using the expansions, we calculated the expected phase for each photon arrival time. For each 512 s interval, we then divided the phases into n phase bins in order to create folded pulse profiles.

We next decomposed the profiles into their Fourier components. Let x_j designate the number of photons in the jth phase bin, and $N_{\rm ph} = \sum_{j=1}^{n} x_j$ is the total number of photons. The complex Fourier component of the kth harmonic is then

$$a_k = \sum_{j=1}^{n} x_j \exp\left(2\pi i j k/n\right) . \tag{3.7}$$

Throughout this paper, we number the harmonics such that the kth harmonic is k times the frequency of the 401 Hz fundamental. Since our analysis for the most part handles the phases and amplitudes of harmonics separately, we define these quantities explicitly as follows:

$$A_k \exp\left[2\pi i k\left(\phi_k + \Delta\phi_k\right)\right] = 2\sum_{j=1}^{n} x_j \exp\left(2\pi i j k/n\right) . \tag{3.8}$$

Here we are interested in the amplitude[3] A_k and the phase residual $\Delta\phi_k$, which we measure relative to a fixed phase offset ϕ_k. We include these offsets because we are principally interested in measuring the phase deviations from a fixed template profile, which we obtain by taking the Fourier transform of all $N'_{\rm ph} = \sum_{j=1}^{n} x'_j$ photons from an outburst:[4]

$$A'_k \exp\left(2\pi i k \phi_k\right) = 2\sum_{j=1}^{n} x'_j \exp\left(2\pi i j k/n\right) . \tag{3.9}$$

Note that we define phases as ranging from 0–1 throughout this paper, with positive phase residuals corresponding to positive time lags: $\Delta\phi_k > 0$ indicates that the kth harmonic arrived later than predicted by the model.

The uncertainties in the phase residuals $\Delta\phi_k$ due to Poisson noise follows from their definition in equation (3.8). Inverting to solve for $\Delta\phi_k$,

$$\Delta\phi_k = \frac{1}{2\pi i k}\left(\ln\sum_{j=1}^{n} x_j E_{jk} - \ln\frac{A_k}{2}\right) - \phi_k \,, \tag{3.10}$$

where we define the constants $E_{jk} \equiv \exp\left(2\pi i j k/n\right)$ for the sake of brevity. For relatively low

[3]We define A_k such that it is the actual amplitude, in photons, of the observed pulsations. We must therefore include both the positive and negative frequency components (which are equal for real signals), introducing the factor of 2 that precedes the summation.

[4]In the presence of sudden pulse profile changes, we may use multiple templates during a single outburst, thus using different values of ϕ_k on either side of the change. Further description is at the end of this section.

fractional amplitudes (certainly the case throughout this paper), each phase bin will contain approximately the same number of photons: $x_j \approx N_{\rm ph}/n$, with variances $(\sigma x_j)^2 \approx N_{\rm ph}/n$ due to Poisson counting statistics. These add in quadrature to give the variance in ϕ_k:

$$\sigma_k^2 = \left| \sum_{j=1}^{n} \left(\frac{\partial \phi_k}{\partial x_j} \right)^2 (\sigma x_j)^2 \right| = \left| \sum_{j=1}^{n} \left(\frac{1}{2\pi k} \frac{E_{jk}}{\sum_{j'=1}^{n} x_{j'} E_{j'k}} \right)^2 \left(\frac{N_{\rm ph}}{n} \right) \right| . \tag{3.11}$$

Summing the exponentials, we have $\left| \sum_{j=1}^{n} E_{jk}^2 \right| = \frac{1}{2} n$. From the definition of A_k in equation (3.8), $\left| \sum_{j=1}^{n} x_j E_{jk} \right| = \frac{1}{2} A_k$. Substituting these in, we reach our estimate of the phase uncertainty:

$$\sigma_k = \frac{\sqrt{2 N_{\rm ph}}}{2\pi k A_k} . \tag{3.12}$$

For our analysis, we rejected phase measurements with large uncertainties. We typically allowed a maximum uncertainty of 0.1 ms.

The measured fractional rms amplitudes are

$$r_k = \frac{A_k}{\sqrt{2}\,(N_{\rm ph} - B)} , \tag{3.13}$$

where B is the approximate number of background events within our energy range and time interval, estimated using the FTOOL `pcabackest`.[5] The r_k add in quadrature: the total rms fractional amplitude for a pulse profile described with m harmonics is $r = \left(\sum_{k=1}^{m} r_k^2 \right)^{1/2}$. Uncertainties on the fractional amplitudes are computed using the method described by Vaughan et al. (1994), which accounts for the inherently biased distribution of amplitudes — even in the absence of a signal, the measured amplitude will always be positive due to counting statistics. As discussed in the preceding section, the probability that the detection of a harmonic is due solely to Poisson noise is $\exp(-P_k)$, where $P_k = \frac{1}{4} A_k^2 / N_{\rm ph}$ is the unit-normalized detection power for the kth harmonic.

Since each harmonic provides an independent measurement of the phase residual, we can combine them to provide the overall phase residual for the sample pulse. We obtain the optimal estimator by weighting each measurement according to its variance:

$$\begin{aligned} \Delta\phi &= \sum_{k=1}^{m} \frac{\Delta\phi_k}{\sigma_k^2} \Big/ \sum_{k=1}^{m} \frac{1}{\sigma_k^2} \\ &= \sum_{k=1}^{m} \frac{k^2 A_k^2}{N_{\rm ph}} \Delta\phi_k \Big/ \sum_{k=1}^{m} \frac{k^2 A_k^2}{N_{\rm ph}} . \end{aligned} \tag{3.14}$$

Thus far, this analytical method closely parallels the work long done on spin-powered pulsars (e.g., Taylor 1993 Appendix A).

However, there are some essential differences that must be taken into account when dealing with accretion-powered pulsars. Unlike spin-powered pulsars, which usually show one or more sharp, asymmetric pulses per cycle, there is little harmonic content in the pulsations of SAX J1808.4−3658 beyond $m = 2$, so we truncate the series there. Additionally, while the individual pulses of radio pulsars show appreciable variability from one period to

[5] `http://heasarc.gsfc.nasa.gov/docs/xte/recipes/pcabackest.html`

the next, their integrated profiles are very stable (Manchester and Taylor 1977). In such cases, one expects that the pulse fractions of sample pulses are similar to the template ($A_k/N_{\rm ph} \approx A'_k/N'_{\rm ph}$) and that the harmonics reflect a common phase residual ($\Delta\phi_k \approx \Delta\phi$) that traces the rotational phase of the star. Indeed, the standard template-matching analysis is predicated on these assumptions. Furthermore, any variability is assumed to be due to Poisson noise, which is of equal magnitude at all timescales (i.e., it is white noise). In contrast, the accretion-powered pulsars show substantial pulse profile variability. Beyond the usual Poisson noise (σ_k from eq. [3.12]), three additional issues complicate the usual approach of template matching: long-timescale correlations (i.e., red noise) in the observed pulse fractions, often with each harmonic's A_k varying independently; red noise in the phase offsets $\Delta\phi_k$; and sudden pulse profile changes, in which the phase offset between the two measured harmonics changes drastically on the timescale of the observations.

In their timing analysis of 283 s accretion-powered pulsar Vela X-1, Boynton et al. (1984) partially address the issue of intrinsic pulse profile noise in the harmonics. In the most general case, the amplitude of the variability in the residuals $\Delta\phi_k$ is different for each harmonic k. This is the case for both Vela X-1 and SAX J1808.4−3658. They correct for the harmonic dependence of these fluctuations by applying a frequency-domain filter to "whiten" the noise.[6] This filter scales the amplitudes of each harmonic's residuals by empirically determined constants, such that all the residuals have approximately the same phase uncertainty (Boynton and Deeter 1985). Thus they were able to measure with much greater accuracy the underlying spin of Vela X-1.

Our approach was similar. For each outburst, we measured the total rms amplitude of the residuals for each harmonic with respect to a best-fit constant-frequency model. These residuals will represent the combined effect of the Poisson noise and any intrinsic profile noise:

$$\sigma^2_{k,\rm rms} = \langle\sigma^2_k\rangle + \sigma^2_{k,\rm int}\,. \tag{3.15}$$

We calculated $\langle\sigma^2_k\rangle$ as a weighted mean of the results from equation (3.12), giving us a value for $\sigma^2_{k,\rm int}$. We then incorporate this additional uncertainty into our weighting to determine $\Delta\phi$:

$$\Delta\phi = \sum_{k=1}^{m} \frac{\Delta\phi_k}{\sigma^2_k + \sigma^2_{k,\rm int}} \Big/ \sum_{k=1}^{m} \frac{1}{\sigma^2_k + \sigma^2_{k,\rm int}}\,. \tag{3.16}$$

σ^2_k changes from one TOA measurement to the next due to the variability of the pulse fraction and count rate; there is no assumption that these are constant, as there is in the case of standard template fitting. $\sigma^2_{k,\rm int}$ is a constant measured independently for each harmonic of each outburst. The result is a minimum-variance estimator for $\Delta\phi_k$. For instance, if the second harmonic has smaller intrinsic fluctuations than the fundamental, then our method presumes that it better reflects the spin of the NS and will weight it more strongly.

Sudden pulse profile changes are somewhat simpler to deal with. We use a different template pulse profile (and hence different measurements on either side of the change of ϕ_k, defined in eq. [3.9]). The decision to change the phase template to accommodate pulse profile shifts rather than modeling them as "shape noise" as described above must be made on a case-by-case basis. Rather than describe this process further here, we defer this

[6]Boynton and Deeter (1985) discuss noise coloration with respect to harmonics, generally finding lower harmonics to have intrinsically larger fluctuations. This language is distinct from the coloration of the timing spectra of those fluctuations themselves, which are also in general non-white and will be addressed later.

discussion until §4.2.2, in which we evaluate some particular examples during the outbursts of SAX J1808.4−3658.

The resulting phase residuals give the best estimator for the offset between the measured and predicted pulse arrival times. By adding these offsets to the phases predicted by the TEMPO ephemerides, we arrived at more accurate pulse arrival times for each interval.

3.3.2 Parameter fitting and uncertainty estimation

After measuring the pulse times of arrival, we input them into TEMPO to refit the timing solution. In order to interpret the resulting models, we must understand the nature of the noise in the TOAs and how it affects the model parameters. The harmonic weighting system described above makes the optimal choice to mitigate the phase variability due to a particularly noisy harmonic, but the TOA residuals are still of substantially greater magnitude than would be expected from Poisson noise alone. This leaves us with the task of estimating the fit uncertainties in the presence of such noise. These uncertainties are crucial to our construction of timing models, as they are needed to estimate the significance of fit components, such as frequency derivatives and instantaneous frequency changes.

We use the term "noise" with respect to spin timing analysis simply to mean harmonic phase variability that does not seem to be due to underlying spin frequency changes. While some of this profile variability may in fact be quite ordered (e.g., in SAX J1808.4−3658 we observed distinctive pulse shape changes that occurred in every outburst; see §4.3.2) we cannot model all of them and thus consider the unmodeled profile variability as "noise" from the phase-timing perspective. This is not a bad approximation: because we fit our models separately for each outburst, correlations in the pulse profile variability between outbursts are not relevant. Furthermore, in SAX J1808.4−3658 we found that the power spectra of the TOA residuals (see Fig. 4-4 in §4.3.3) resemble the power-law noise spectra typically observed in actual red noise processes, so treating it as such is reasonable.

When fitting the TOAs from an AMP outburst, we initially used the simplest possible timing model in TEMPO: a circular orbit and a constant pulse frequency. When this simple model proved insufficient to account for the phase residuals, we introduced a nonzero $\dot{\nu}$ and instantaneous frequency changes, as needed. However, there is a danger of overfitting the data. It is important to recognize that some of the features in the residuals are probably pulse profile variability rather than spin evolution. We took great care in our attempts to distinguish between the $\dot{\nu}$ measurements and the artifacts of intrinsic timing noise.

The colored nature of the timing noise in both harmonics is the primary difficulty in the interpretation of the parameter fits. TEMPO assumes that the TOA uncertainties one gives it are white and approximately Gaussian, as is the case of pure Poisson noise. As a result, it systematically underestimates the uncertainties in the fitted parameters in the presence of timing noise. Red timing noise is particularly problematic, because it dominates on the long timescales on which ν and $\dot{\nu}$ measurements depend.

Instead of adopting this white noise assumption of TEMPO, we estimated confidence intervals for ν and $\dot{\nu}$ using Monte Carlo simulations of the timing residuals of each outburst. After using TEMPO to obtain the best fit for a timing model, we calculated the power spectrum of the timing residuals that TEMPO output. This spectrum is a convolution of the true noise spectrum and the sampling function: most notably, there is excess power around 1 d, an artifact of *RXTE* observations often being scheduled approximately a day apart, and at the *RXTE* orbital period of 96 min due to Earth occultations, as discussed in §2.3. We applied a low-pass filter to allow through the red noise but remove these peaks due

to sampling. This filter was only partial, allowing through some higher frequency noise to reproduce the short-timescale scatter (principally but not entirely Poisson) that we observed within each observation.

We then created thousands of sets of artificial phase residuals using this filtered spectrum and random phases. To reproduce the sampling irregularities, we removed all points at times absent in the original data. The resulting Monte Carlo residuals were visually and statistically indistinguishable from their real counterparts after properly tuning the low-pass filter. For each set of residuals, we measured the frequency of the best linear fit (or, if our TEMPO model fit for $\dot{\nu}$, the frequency derivative of the best quadratic fit). The standard deviations of these measurements provided uncertainty estimates for the respective parameter, this time more accurately accounting for the noise spectrum.

In deriving improved phase model parameters, the new values sometimes differed considerably from the parameters with which we initially folded the data for TOA calculations. If our orbital model improved substantially, we iterated the above procedure, calculating new times of arrival for each 512 s interval and refitting. Because the rapid AMP orbits introduce periodic frequency modulation with amplitudes $\Delta\nu = \nu_0 \cdot 2\pi a_x \sin i / P_{\rm orb} > 1/512$ s, an inaccurate orbital ephemeris can significantly reduce detection strength. In contrast, AMP spin frequencies are remarkably stable through all the observations, so there was no need to recalculate TOAs upon the relatively minor revisions to the spin model.

Chapter 4

The long-term evolution of the spin, pulse shape, and orbit of SAX J1808.4−3658

This chapter is adapted from the paper "The Long-Term Evolution of the Spin, Pulse Shape, and Orbit of the Accretion-Powered Millisecond Pulsar SAX J1808.4−3658" by Jacob M. Hartman, Alessandro Patruno, Deepto Chakrabarty, David L. Kaplan, Craig B. Markwardt, Edward H. Morgan, Paul S. Ray, Michiel van der Klis, and Rudy Wijnands, which has been submitted to *The Astrophysical Journal*.

4.1 Introduction

The growing class of accretion-powered millisecond X-ray pulsars discovered by the *Rossi X-Ray Timing Explorer* has verified the hypothesis that old millisecond pulsars obtained their rapid spins through sustained accretion in X-ray binaries. These objects provide a versatile laboratory. The X-ray pulse shapes arising from the magnetically channeled accretion flow can constrain the compactness (and hence the equation of state) of the neutron star. Tracking the arrival times of these X-ray pulses allows us to measure the pulsar spin evolution, which directly probes magnetic disk accretion torque theory in a particularly interesting regime (Psaltis and Chakrabarty 1999) and also allows exploration of torques arising from other mechanisms such as gravitational wave emission (Bildsten 1998). There have been several reports of significant spin evolution in accreting millisecond pulsars, some with implied torques that are difficult to reconcile with standard accretion torque theory (Markwardt 2003; Morgan et al. 2003; Burderi et al. 2006, 2007). However, a variety of effects (including limited data spans, pulse shape variability, and non-Gaussian noise sources) can complicate the interpretation of these measurements. In this chapter, we address these difficulties using a comprehensive analysis of the most extensive data set.

We describe our application of phase-connected timing solutions for each outburst to study the spin history and pulse profile variability of SAX J1808.4−3658, providing the first look at the evolution of an accretion-powered X-ray pulsar. In section 4.2, we outline our analysis methods, noting the difficulties raised by pulse profile noise and describing a new technique to obtain a minimum-variance estimate of the spin phase in the presence of such noise. In section 4.3, we present the results of this analysis. In particular, we observe that

the source is spinning down between outbursts, the binary orbital period is decreasing, and the pulse profiles change in a characteristic manner as the outbursts progress. Finally, in section 4.4, we discuss the implications of these results to the properties of the neutron star and accretion geometry.

4.2 X-ray Observations and Data Analysis

4.2.1 Processing of *RXTE* observations

The *RXTE* Proportional Counter Array (PCA; Jahoda et al. 1996) has repeatedly observed SAX J1808.4−3658, primarily during outburst. These observations total 307 separate pointings and an exposure time of 1,371 ks from 1998 through 2005. All but three of the observations of SAX J1808.4−3658 were taken with the `E_125US_64M_0_1S` mode, which records the arrival of each photon with a time resolution of 122 μs and 64 energy channels covering the full range of the detectors. The other observations were rebinned to be compatible with the 122 μs resolution data; using higher time resolutions provides no benefit. We shifted the photon arrival times to the Earth's geocenter[1] using an optical position of R.A. = $18^{\rm h}08^{\rm m}27\overset{\rm s}{.}62$, Decl. = $-36°58'43\overset{''}{.}3$ (equinox J2000), with an uncertainty of $0\overset{''}{.}15$. We then applied the *RXTE* fine clock correction, which provides an absolute time accuracy of 5–8 μs (Rots et al. 2004). Finally, we filtered the data to remove Earth occultations, intervals of unstable pointing, and thermonuclear X-ray bursts. For three observations at the start and end of the 1998 outburst, we relaxed our requirement of stable pointing and included raster scanning data to extend our baseline for measuring the frequency evolution during this outburst. These observations provided additional valid phase measurements, but they were not used to calculate fractional amplitudes since the contribution of the source and background varied as the *RXTE* panned across the source. Table 4.1 lists all the observations that we included in our analysis.

We consistently used an energy cut of roughly 2–15 keV for our timing analysis. While the source is readily detectable in the PCA at higher energies, the background dominates above 15 keV, especially in the dimmer tails of the outbursts. Excluding these high-energy counts optimized the detection of pulsations when the source was dim, providing a longer baseline for our timing analysis.

4.2.2 Pulse timing analysis with tempo

To fit timing models to the SAX J1808.4−3658 outbursts, we followed the procedure outlined in §3.3 for performing phase connection in the presence of the random and systematic timing noise present in AMPs. These techniques were developed specifically to handle the four outbursts from this source.

Our timing models fit for the following parameters: the pulsar spin frequency and (if necessary) the first-order frequency derivative; the times and magnitudes of any instantaneous changes in the frequency; and the orbital parameters. Our models supplied, but did not fit, the position of the source. Because we fit the outbursts separately, the $\approx$1 month of data that each provides was not sufficient to improve the source position: the position

[1]We shifted the photons to the geocenter rather than the solar system barycenter since the TEMPO pulse timing program, which we used to fit phase models to the arrival times, was designed for radio timing and thus expects photon arrival times at some point on the Earth. TEMPO itself performed the barycentric corrections using the quoted position.

Table 4.1. Observations analyzed for each outburst

	Start (MJD)	Stop (MJD)	# obs.	Time (ks)	Avg. # PCUs	Observation IDs
1998 Apr	50914.8	50939.6	21	178.1	4.67	30411-01-*
2000 Feb	51564.1	51601.9	38	126.8	3.74	40035-01-01-00 – 40035-01-04-01
						40035-05-02-00 – 40035-05-18-00
2002 Oct	52562.1	52602.8	129	714.5	3.25	70080-01-*, 70080-02-*
						70080-03-05-00 – 70080-03-24-00
						70080-03-25-01, 70518-01-*
2005 Jun	53523.0	53581.4	55	284.3	2.84	91056-01-01-01 – 91056-01-04-01
						91418-01-01-00 – 91418-01-07-00

Note. — The ranges of observation IDs given here are for numerically sorted IDs, which do not always reflect temporal order.

of SAX J1808.4−3658 (in particular, its right ascension) is degenerate with the frequency and frequency derivative on such timescales.

To parametrize the orbit of SAX J1808.4−3658, we used TEMPO's ELL1 binary model, which employs the Laplace parameters $e \sin\omega$ and $e \cos\omega$, where e is the eccentricity and ω the longitude of periastron passage. This parametrization avoids the degeneracy of ω in low-eccentricity systems (Deeter et al. 1981). For most of the fits, we held $e = 0$ and solely fit the projected semimajor axis $a_x \sin i$, the orbital period $P_{\rm orb}$, and the time of ascending node[2] $T_{\rm asc}$. As a test of this assumption, we also repeated the fits allowing e to vary. It was always consistent with zero.

We relied solely on TEMPO to calculate the uncertainties in the binary orbit parameters. While the intrinsic pulse profile noise spectrum is colored on the timescale of days to weeks, it is approximately white on timescales equal to and shorter than the 2 hr orbital period. The amplitude of the short-timescale variability is roughly 1.5 times what one would expect from counting noise alone, so we scaled our Poisson-derived phase uncertainties accordingly when estimating the uncertainties of the orbital parameter fits. This rescaling makes TEMPO's uncertainty estimates for the orbital parameters reasonably accurate. One important consistency test of this simple approach worked nicely: the 1998, 2002, and 2005 measurements of $P_{\rm orb}$ and $a_x \sin i$, two parameters that should be the same for each outburst at our level of accuracy, had reduced χ^2 statistics close to unity.

In only one case did we model a sudden pulse profile change in SAX J1808.4−3658: at the end of the main body of the 2002 outburst, the fundamental phase ϕ_1 experienced a shift while the second harmonic, ϕ_2, remained constant. We changed the ϕ_1 value of our phase template accordingly at this point. The stability of ϕ_2 allowed us to phase connect across the feature, as Burderi et al. (2006) also noted.

Our distinction between sudden profile changes and pulse profile noise is admittedly somewhat arbitrary. We make it solely in the interest of best estimating the rotational phase of the star — we are not claiming to model some underlying difference in physical

[2]Past pulsar timing of SAX J1808.4−3658 uses the T_{90} fiducial, marking a time at which the mean longitude is 90°. Since its orbit is circular, $T_{90} = T_{\rm asc} + P_{\rm orb}/4$.

processes. In 2002, the phase residuals on either side of the modeled pulse profile change were quite stable, albeit with different values of ϕ_1. This stability makes it a good candidate for such treatment. In contrast, the phase residuals of both harmonics during the 2005 outburst show greater amplitude fluctuations at nearly all timescales. While its residuals and the 2002 residuals follow a similar pattern at the end of the main body of the outburst (the phase residuals of the second harmonic remain roughly constant, while the residuals of the fundamental drop appreciably), the fundamental continues to fluctuate wildly after this event rather than settling down on a "new" template profile. Therefore we elect to attribute these profile changes to intrinsic noise and weight the relatively stable second harmonic more strongly.

4.3 Results

The results of our pulse timing solutions are shown in Figure 4-1, which compares the lightcurves, phase residuals, and fractional amplitudes for each outburst. Inspecting the best-fit frequency lines in the phase residual plots, it is clear that a constant pulse profile attached to a constant-frequency rotator does not adequately describe the observed residuals. We consider five sources of phase residuals relative to a best-fit constant-frequency model: Poisson timing noise, intrinsic pulse profile noise, sudden and well-defined pulse profile changes, additional spin frequency derivatives, and instantaneous frequency changes in the underlying rotation of the star. In this section, we will consider all these possible contributions to the residuals and their relationships with each other and the other properties of each outburst.

4.3.1 Lightcurves of the outbursts

The lightcurves of each outburst are quite similar in shape. We divide them into four stages: the rise, which was only definitively captured in 2005 and took $\approx$5 days; the short-lived peak at a 2–25 keV flux of $(1.9-2.6)\times10^{-9}$ erg cm^{-2} s^{-1}, equal to a luminosity of $(4.7-6.4)\times10^{35}$ erg s^{-1} using the distance of 3.5 kpc and bolometric correction of $L_{\rm bol}/L_{2-25\ \rm keV}=2.12$ derived by (Galloway and Cumming 2006); a slow decay[3] in luminosity, lasting 10–15 days, until the source reaches approximately 8×10^{-10} erg cm^{-2} s^{-1}($=2.0\times10^{35}$ erg s^{-1}); and a sudden drop followed by low-luminosity flaring as the outburst flickers out, with the timescale between flares on the order of 5 days. Figure 4-2 shows a cartoon of a typical outburst from SAX J1808.4−3658, with each of these stages labeled.

The *RXTE* first collected high-resolution timing data from SAX J1808.4−3658 during the 1998 April outburst. Data from the *RXTE* All Sky Monitor (ASM) show that the peak luminosity occurred approximately three days before the first PCA observation. (See Figure 2 of Galloway and Cumming 2006 for a comparison of the ASM and PCA lightcurves; note that its 1998 plot does not include two raster scans analyzed here.) Unfortunately, PCA observations stopped shortly after the body of the outburst and do not sample the tail.

[3]Some authors (e.g., Cui et al. 1998; Burderi et al. 2006) refer to this part of the outburst as the "exponential decay" stage, based on the approximately exponential dimming of the 1998 and 2002 outbursts. However, the fall off of the 2005 outburst during this stage is closer to linear, so we simply refer to it as the "slow decay" stage to contrast it with the more rapid luminosity drop at its conclusion.

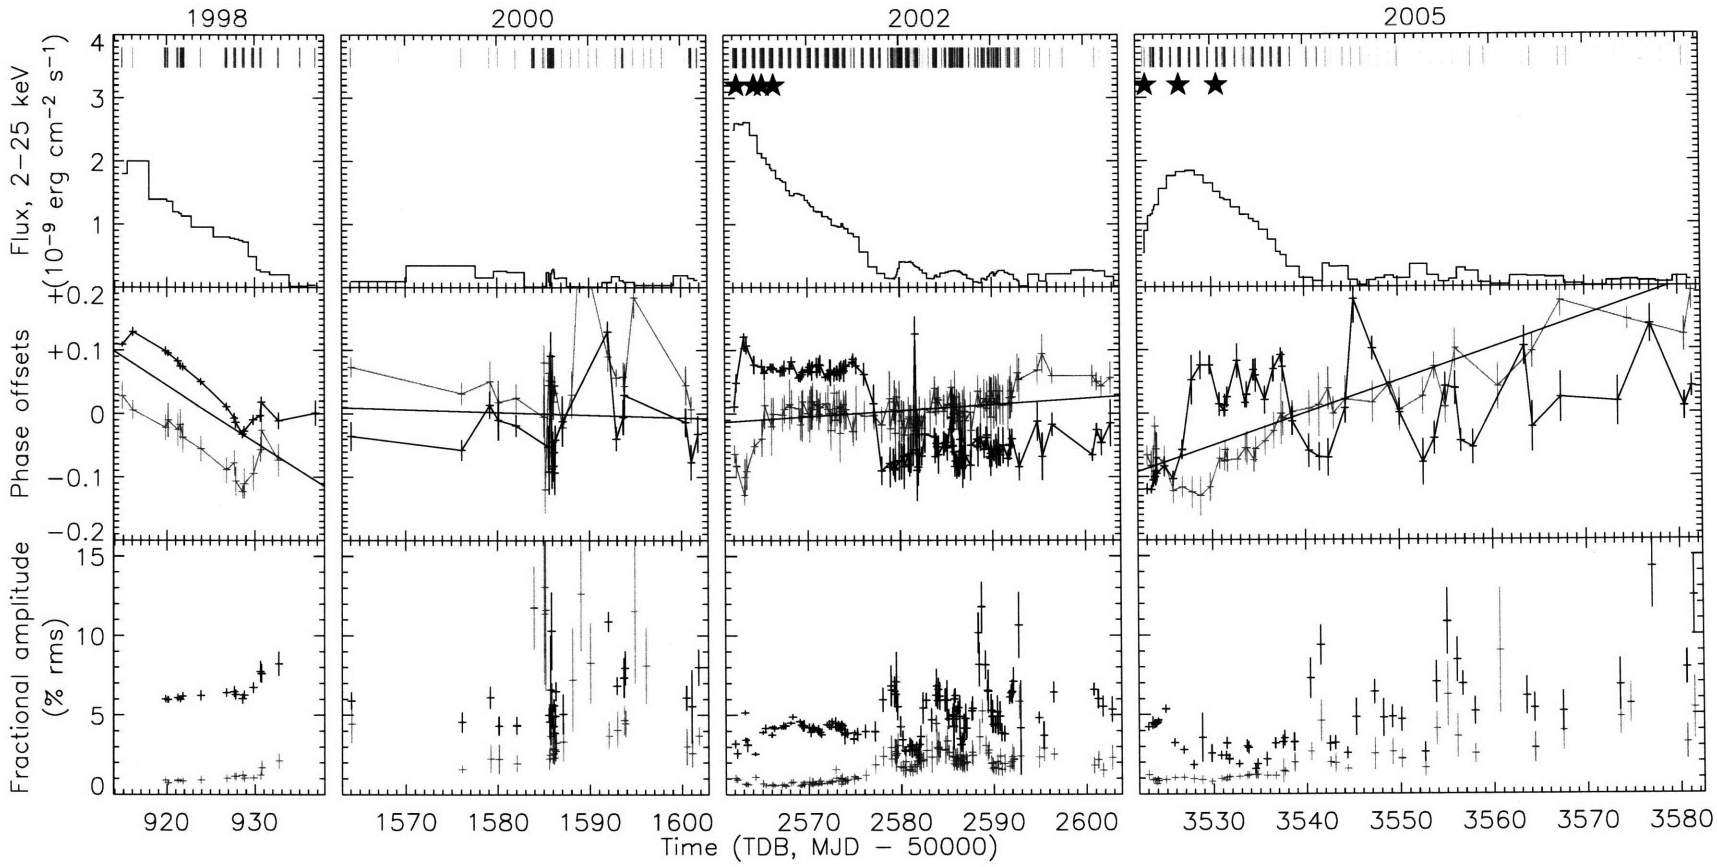

Figure 4-1 The lightcurves, phase residuals, and fractional amplitudes for all four outbursts. The top panels show the background-subtracted lightcurves for each outburst. The strips along the top of the graphs indicate the times of observations; stars indicate the times of thermonuclear X-ray bursts. The second panels show show the phase residuals relative to a constant frequency of 400.97521025 Hz, with black points giving the phases of the fundamental and grey points indicating the second harmonic. Positive phases indicate pulse arrivals later than predicted by the phase model. These data have been binned up so that there is only one point per observation. The black lines indicate the best-fit constant-frequency model for each outburst. The bottom plot shows the fractional amplitudes of the fundamental and harmonic.

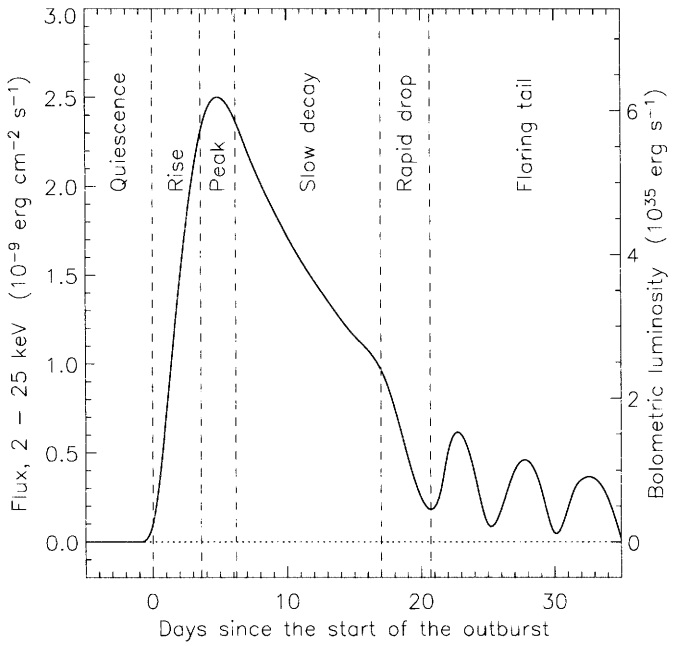

Figure 4-2 The anatomy of a typical outburst from SAX J1808.4−3658. The features of the lightcurve and their fluxes and timescales are similar to those observed during the 1998, 2002, and 2005 outbursts. Bolometric luminosities assume a distance of 3.5 kpc and a correction of $L_{\rm bol}/L_{2-25\ \rm keV} = 2.12$ (Galloway and Cumming 2006).

SAX J1808.4−3658 was discovered to be again in outburst when it emerged from behind the Sun in 2000 January. Coverage of the outburst was limited and included only the outburst tail. Wijnands et al. (2001) comment on the erratic nature of the flaring during the tail. ASM data indicate that the peak occurred 2 weeks prior to the first PCA observation and that we are observing the later, dimmer stage of the flaring tail. Comparison of the PCA data with the 2002 and 2005 outbursts suggests likewise.

The 2002 outburst, detected in mid-October and observed for the next two months, was the brightest, had the best PCA coverage, and included the detection of four extremely bright thermonuclear X-ray bursts during its peak. Its lightcurve was very similar in shape to the 1998 outburst.

In 2005 June, SAX J1808.4−3658 was again in outburst. This time, the detection preceded the peak by a few days, providing a full sampling of the lightcurve. This outburst was somewhat dimmer, with a peak luminosity of only 70% of the 2002 peak and a correspondingly shorter slow-decay stage. The subsequent rapid decay and flaring tail look quite similar to the other outbursts.

4.3.2 Characteristic pulse profile changes

Just as the lightcurves of each outburst were quite similar, the evolution of the pulse profile during each outburst was remarkably consistent. Figure 4-3 illustrates the full range of pulse profiles that we observed from SAX J1808.4−3658. In many instances, the similarity of the pulse profiles between outbursts is quite striking. In this section, we describe how these profile change throughout the outbursts.

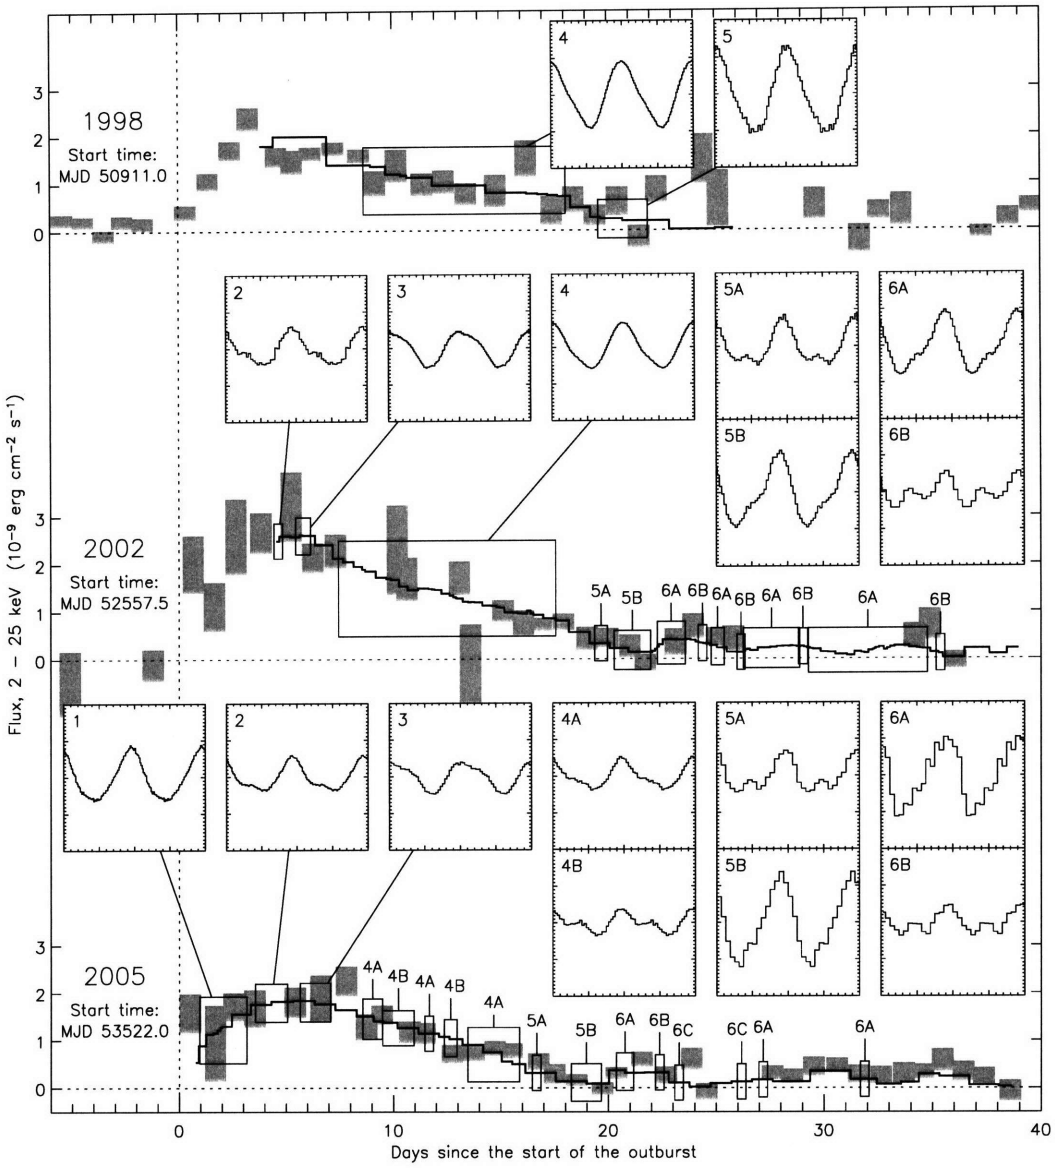

Figure 4-3 A comprehensive view of the 2–15 keV pulse profiles observed from SAX J1808.4−3658. Each pulse profile was calculated by folding the observations within the indicated time intervals using the best-fit constant-frequency model of each outburst; thus any movement of the peaks accurately reflects the phase offsets from the constant frequency. The profiles are background-subtracted and normalized such that the average source flux is constant, so the plotted profiles accurate reflect the change in fractional amplitude during the outburst. The profiles are numbered according their position within the outburst: 1 indicates the burst rise; 2, the beginning of the outburst maximum; 3, the end of the maximum; 4, the slow decay stage; 5, the steep luminosity drop marking the end of the main outburst; and 6, the flaring tail. During some parts of the burst, two pulse profiles are present, with the source switching between them. In these cases, we show both profiles and label the regions of the lightcurve in which they occurred accordingly. The solid black line shows the fluxes from the PCA observations; the grey boxes show the fluxes from the ASM daily averages.

We observed the burst rise, labeled as box 1 in Figure 4-3, exclusively during the 2005 outburst. The profiles are smooth and asymmetric, with a slow rise in flux followed by a more rapid drop-off after the peak. There is no sign of a second peak.

We observed the outburst maxima during 2002 and 2005. The similarity of the pulse profile evolution between the two outbursts was remarkable. During the first half of the maxima (labeled as profile set 2), the profiles show a secondary bump lagging the main pulse. Compared to the burst rise, the fractional amplitude has decreased somewhat. During the second half of the maximum, the pulse becomes broader, subsuming the lagging secondary bump. (See profile set 3.) This change appears to be gradual: in both outbursts, a mid-peak observation exhibited an intermediate pulse profile.

Profile set 4 shows the pulse profiles during the slow decay stages of the outbursts. The 1998 and 2002 profiles are quite similar: the pulses are somewhat asymmetric, rising more steeply than they fall. In both outbursts, this pulse profile is very stable during the approximately 10 days of the decay in luminosity. During the 2005 outburst, this asymmetry is more pronounced, and the profile varies between observations. Initially, the pulse exhibited a small lagging bump (box 4A), quite similar to the pulse profile during the first half of the outburst maximum. The relative size of that bump varied substantially, in some observations appearing as a small secondary peak (box 4B). Over the course of the decline, the source switched back and forth between a double-peaked and single-peaked profile as indicated in the figure. A given state would typically be seen for two or three observations (1–2 days) before switching to the other.

Profile set 5 covers the rapid drop in flux at the end of the outbursts. During 1998, the pulse profile was quite stable and did not appreciably change during this drop, although its fractional amplitude increased somewhat. In contrast, the 2002 and 2005 outbursts show a major pulse profile shift concurrent with the drop in luminosity. Prior to the drop, the pulses in set 4 show a quick rise and a slower fall. After the drop, the asymmetry of the 2002 and 2005 profiles reverses: boxes 5B show a slow rise and a quick drop. In terms of harmonic components, these changes represent a shift in the phase of the fundamental by approximately 0.15 cycles as it went from leading the second harmonic to lagging behind it. The phase of the harmonic did not change. In both outbursts, observations during the $\approx$2 days of rapid luminosity decline reveal an intermediate stage in which the main pulse is momentarily symmetric (boxes 5A). During this transition, small but significant secondary pulses are present.

During the flaring tail of the outburst (profile set 6), the pulse profile again showed substantial variability. In 2002, the profile repeatedly switched form between an asymmetric pulse (box 6A, identical to the pulse profile at the end of the rapid dimming stage) and a double-peaked profile (box 6B). The double peaked pulse profile occurs principally (but not exclusively) at the end of the flares, as their luminosity declines. These pulse profile changes are almost entirely the result of changing fractional amplitudes of the harmonic components; the phase offset between the fundamental and second harmonic remains for the most part constant. A notable exception occurs during the decay of the first flare at around MJD 52582. At this time the phase of the fundamental jumped by $\approx$0.2 cycles, indicating a sudden lag of this amount behind its previous arrival time. By the next observation, less than 2 hours later, the phase residual of the fundamental returned to its previous value.

The tail of the 2005 outburst is more chaotic. The fractional amplitudes and phases both exhibit strong red noise, producing a pulse profile that is sometimes asymmetric with a slow rise and quick fall (6A); at other times asymmetric with a quick rise and slow fall (6C; profile not shown, but basically just the reverse of box 6A); and in one instance clearly

Table 4.2. Noise properties of the outbursts

	Fundamental's noise			Second harmonic's noise		
	$\langle\sigma_1^2\rangle^{1/2}$	$\sigma_{1,\mathrm{int}}$	γ_{PLN}	$\langle\sigma_2^2\rangle^{1/2}$	$\sigma_{2,\mathrm{int}}$	γ_{PLN}
1998 Apr	0.007	0.014	0.96 ± 0.09	0.021	0.017	0.43 ± 0.12
2000 Feb	0.023	0.052	—	0.022	0.039	—
2002 Oct	0.012	0.016	0.67 ± 0.09	0.023	0.027	0.51 ± 0.09
2005 Jun	0.013	0.061	0.85 ± 0.07	0.019	0.024	0.77 ± 0.05

Note. — All phases are in cycles; i.e., fractions of the 2.5 ms spin period. We did not attempt to estimate power-law noise slopes for the 2000 outburst because of its low-quality data.

double-peaked (6B). The observations were sparse and generally short, so it was impossible to better characterize the evolution of these pulse profile fluctuations. The flaring tale of the 2000 outburst was quite similar, with a highly variable pulse profile that included double-peaked profiles and asymmetric single pulses of both orientations. We did not include it because the observations were few and sparse.

4.3.3 Noise properties of the timing residuals

To measure the spin phase of SAX J1808.4−3658 using the formalism developed in §3.3.1, we must characterize the variability harmonic components. This variability encompasses both the pulse profile changes discussed in the previous section as well as any noise in the spin phase of the star.

In our analysis of the phase residuals, we took into account the rms amplitude of the intrinsic pulse profile noise in each harmonic, $\sigma_{k,\mathrm{int}}^2$, defined in equation (3.15). A casual glance at the phase residuals of Figure 4-1 reveals that the magnitudes of these fluctuations vary substantially between outbursts. Table 4.2 summarizes these amplitudes for each outburst and compares them to the mean amplitudes of their Poisson noise, $\langle\sigma_1^2\rangle^{1/2}$. These values are then used in equation (3.16). For instance, in the 2002 outburst the fundamental is more heavily weighted in measuring the spin phase than the second harmonic, while in 2005 the opposite is true.

The scatter of the phase residuals between observations is generally greater than the scatter within an observation, suggesting that the pulse profile noise is red. Power spectra of the phase residuals, shown in Figure 4-4, confirm this. We estimated these power spectra using Fourier transforms of the residuals from equally spaced 512 s bins. Here we have not attempted to deconvolve the uneven sampling due to the schedule of observations, which produces a peak at 1 day (an artifact of *RXTE* observations often being spaced approximately a day apart) and a peak at the *RXTE* orbital period of 96 minutes due to Earth occultations. There are no peaks at the 2 hour binary orbital period, indicating that our modeling of the orbit is accurate and that the pulse profile is independent of orbital phase.

The resulting noise powers are around 2 decades higher at long periods ($\approx$3 days or longer) for 1998, and even more for 2005. The 2002 outburst spectra exhibit less profile

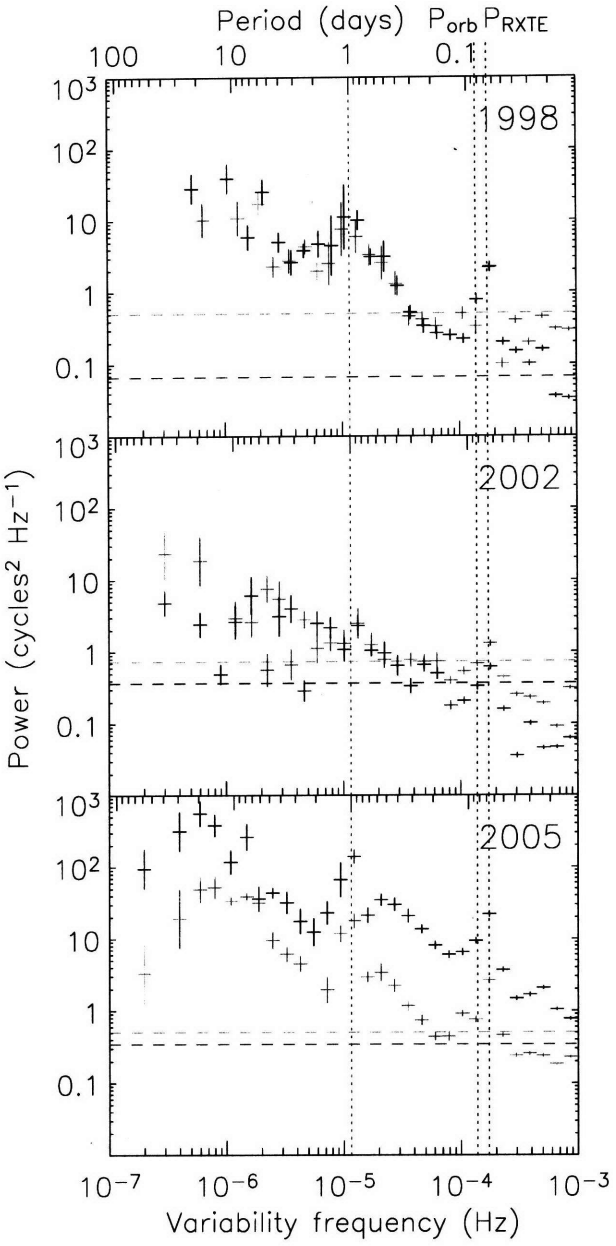

Figure 4-4 Power spectra of the phase residuals for the fundamental (black) and the second harmonic (grey) relative to the best-fit constant-frequency models. (The 2002 model also includes a phase shift in the fundamental to account for the profile change at MJD 52577.) The dashed lines show the power level due to counting statistics, a white noise contribution proportional to $\langle \sigma_k^2 \rangle$. The data points show the powers $P_k(f)$, from which we have subtracted the contribution of counting statistics. These powers are normalized such that $\int_{10^{-7}\ \mathrm{Hz}}^{10^{-3}\ \mathrm{Hz}} P_k(f)\, df = \left\langle \sigma_{k,\mathrm{int}}^2 \right\rangle$, as defined in equation (3.15). The solid lines give the best-fit power law noise curves.

noise at long timescales, but still are somewhat reddened. Poisson statistics produce an uncolored lower limit on noise. This white noise dominates at timescales shorter than the orbital period of 2 hours, except in the case of the particularly noisy fundamental of 2005. The spectra of the intrinsic profile noise (i.e., the spectra after subtracting off the Poisson contribution) roughly followed a power law noise spectrum, which we parametrized as $P_k(f) \propto f^{-\gamma_{\rm PLN}}$. The best-fit values of $\gamma_{\rm PLN}$, listed in Table 4.2, varied from roughly 0.4 to 1.

4.3.4 Fractional amplitudes of the harmonics

In our time-domain discussion of the pulse profiles, an apparent trend is the tendency of the pulses to become narrower, more asymmetric, or doubly peaked — generally speaking, to become less sinusoidal — as the outburst's flux decreases. In the frequency domain, the relation is striking: the fractional amplitude of the 802 Hz second harmonic, r_2 correlates strongly with the background-subtracted 2–25 keV flux, $f_{\rm x}$, as shown in Figure 4-5. This power-law dependency has a slope of -0.50 ± 0.01. The agreement with the data is excellent for such a simple model, giving a reduced χ^2 statistic of $\chi^2_\nu = 1.15$ with $\nu = 1816$ degrees of freedom. It spans two and a half decades in luminosity and includes every detected harmonic amplitudes from all four outbursts.

In terms of the pulse profile, the second harmonic contributes in two ways. If its peak is 45° out of phase with the peak of the fundamental, it will produce an asymmetric pulse profile (e.g., profile 6A in Figure 4-3). If it is in phase, a narrower primary pulse with a small second peak will result (as in profile 6B). If the components are 90° out of phase, the profile will be profile 6B flipped, but we never observed such a configuration.

To further understand the influence of flux on the pulse profile, we decomposed the second harmonic's fractional amplitude into its asymmetric and double-peaked components,

$$r_{2,\rm asym} = r_2 \left|\sin 4\pi(\phi_2 - \phi_1)\right| \quad \text{and} \tag{4.1a}$$

$$r_{2,\rm dp} = r_2 \left|\cos 4\pi(\phi_2 - \phi_1)\right| , \tag{4.1b}$$

where ϕ_1 and ϕ_2 are the observed phase offsets of the two harmonic components.[4] The resulting plots have substantially more scatter than Figure 4-3 due to the uncertainty of $\phi_2 - \phi_1$, which is considerable, particularly at low fluxes. However, they both roughly conform to the $r_2 \propto f_{\rm x}^{-1/2}$ power law. We conclude that the decrease in flux increases the asymmetry of the pulses and the presence of secondary pulses in approximately equal measure.

In contrast, the fractional amplitude of the fundamental behaves unpredictably. During the slow-decay stage of 1998, it is unvarying and strong, at a constant 5.5% rms. During this stage of 2002, it is weaker (4%) and somewhat variable; during 2005, it is weaker still and erratically changing by up to a full percent between observations. The story is more consistent in the tail. In all outbursts, the fractional amplitude of the fundamental varies widely, usually (but not always) having its maxima around the peaks of the flares and its minima during the fading portion of the flares.

For the most part, a pulse profile model only including the fundamental and second harmonic adequately describes the folded profiles. However, folding long stretches of data does

[4]These definitions of ϕ_k are slightly different from those used previously (e.g., in eq. [3.8]) since we are now interested in their phase offset to each other rather than their offset relative to a template profile.

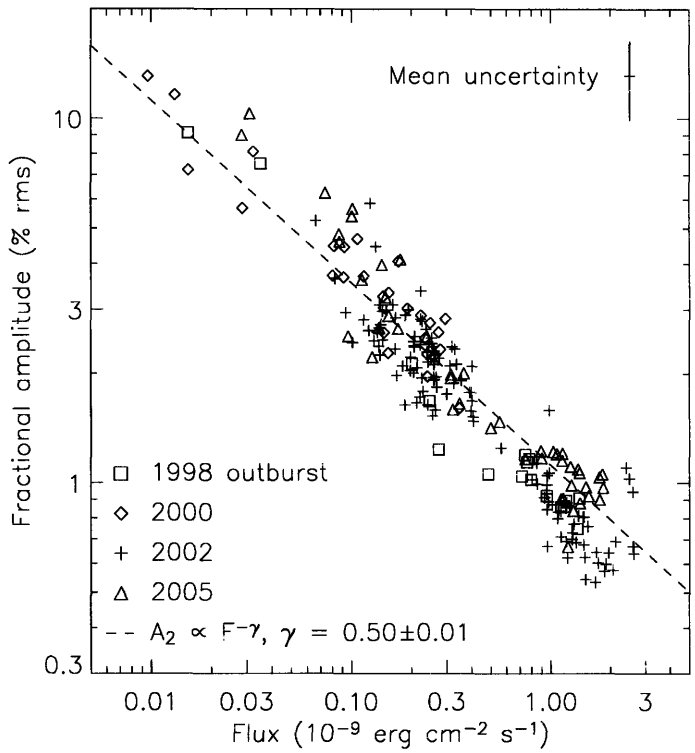

Figure 4-5 The fractional amplitude of the second harmonic scales with flux according to a power law of slope -0.50 ± 0.01, shown by the dashed line. Each point gives the mean amplitude and flux for a single observation. The scatter is commensurate with the mean uncertainty in fractional amplitude, which is shown by the error cross in the upper right.

Table 4.3. Upper limits on subharmonics and half-integral harmonics

Harmonic		Upper limit[a](% rms)		
Factor	Hz[b]	A	B	C
1/4	100.2	0.017	0.019	0.52
1/2	200.5	0.022	0.024	0.45
3/2	601.5	0.018	0.021	0.43
5/2	1002.4	0.026	0.024	0.42

[a]These background-corrected upper limits are quoted at the 95% confidence level.

[b]Frequencies listed here are approximate. The upper limits were obtained using exact multiples of the best-fit constant-ν models.

sometimes result in the detection of a third harmonic with fractional amplitudes ranging up to $\approx$0.25% rms. We do not reliably detect any higher harmonics.

4.3.5 Upper limits on the subharmonics

With some assumptions, we can strongly constrain the presence of subharmonics and half-integral harmonics. The most straightforward approach is to fold all the observations using multiples of the best-fit frequency models from each outburst. The amplitude of the resulting profile will give an upper limit. The resulting 95% confidence upper limits are listed in column A of Table 4.3. However, this approach is only statistically valid if the uncorrected fractional amplitude (i.e., the fractional amplitude relative to the source counts *and* the background) is constant. Clearly this assumption is false. Aside from the varying proportion of source photons, the background-subtracted fractional amplitudes of the fundamental and the second harmonic fluctuate throughout the outburst, spanning nearly an order of magnitude in the tails of the burst. There is no reason to believe that a subharmonic would not fluctuate similarly.

Column B of Table 4.3 takes the more moderate approach of only folding together observations during which the 2–25 keV flux exceeds 5×10^{-10} erg cm^{-2} s^{-1}, thereby only including the main body of the outbursts. Background photons are thus a much smaller contribution, and the fractional amplitudes of the observed two harmonics were relatively stable during these times. Nevertheless, we still are folding enough photons to obtain very stringent upper limits: in the case of the 200 Hz subharmonic, we get a 95% confidence upper limit of 0.024% rms. We feel that these numbers are our most reliable, not making unreasonable assumptions about the fractional amplitude fluctuations.

For completeness, we also include the most conservative upper limits, which make no assumptions whatsoever about the fractional amplitudes of the subharmonics. For instance, it would be possible in principle for the subharmonic to be present only during a single observation and zero-amplitude everywhere else. To constrain the resulting upper limits at least somewhat, we again only used observations during which the source was brighter than 5×10^{-10} erg cm^{-2} s^{-1} and that had at least 10^6 counts. These single-observation limits are tabulated in column C.

Table 4.4. Best-fit constant frequencies, and their $\dot{\nu}$ upper limits

	$\nu - \nu_0$ (μHz) [a]	$\dot{\nu}$ (10^{-14} Hz s^{-1}) [b]
1998 Apr	0.371 ± 0.018	$(-7.5, 7.3)$
2000 Feb	0.254 ± 0.012	$(-1.1, 4.2)$
2002 Oct	0.221 ± 0.006	$(-1.3, 2.5)$
2005 Jun	0.190 ± 0.016	$(-0.5, 2.4)$

[a]The frequency measurements are relative to $\nu_0 = 400.975210$ Hz.

[b]95% confidence intervals.

The stringent upper limits of column B provide the best evidence yet that the the spin frequency of the star is indeed 401 Hz. If the star was spinning at 200.5 Hz, with two antipodal hot spots each emitting pulses to produce the observed frequency, a 200.5 Hz subharmonic would almost certainly be present.

4.3.6 Spin frequency measurements and constraints

We initially performed the simplest possible fits to the phase residuals of each outburst: constant-frequency models. We did not include the data at the very beginning of the 2002 and 2005 outbursts, where pulse profile changes during the rise and peak obscure any variations in the phase. We also excluded the residuals during 2002's mid-outburst pulse profile change, but included the residuals of the fundamental on both sides of the shift by using different profile templates before and after it.

The resulting frequency measurements are shown in Figure 4-6 and summarized in Table 4.4. These data clearly indicate that the source is spinning down. The probability that the actual spin frequency is constant or increasing is less than 10^{-9} given the uncertainty estimates. These uncertainties do assume that our optical position is exact, but the position error is excluded because its effects are highly correlated; for instance, the 1998 and 2002 outbursts are six months apart on the calender, so a position offset would produce equal and opposite frequency displacements for the measurements from these outbursts. There is no position that would provide a statistically feasible constant or increasing frequency.

The linear fit through the measured frequencies is not particularly good: its χ^2 statistic is 9.7 with 2 degrees of freedom, yielding a probability of about 1% that the frequencies are drawn from a linear progression. Once again, changing the source position does not significantly change the result or improve the fit, and changes in the position by more than the $1\,\sigma$ uncertainty along the ecliptic substantially worsen the linear fit. To estimate the uncertainties of the linear slope in light of this poor fit, we rescaled the measurement uncertainties such that reduced χ^2 statistic would be unity. The resulting first-order spin derivative is $\dot{\nu} = (-5.6 \pm 2.0) \times 10^{-16}$ Hz s^{-1}.

Fitting second-order frequency models established that $\dot{\nu}$ is consistent with zero during all the outbursts. These measurements are particularly sensitive to pulse profile variations, so care must be taken to not overfit such features. We again exclude the initial observations of the 1998, 2002, and 2005 outbursts, because the pulse profile changes would induce large

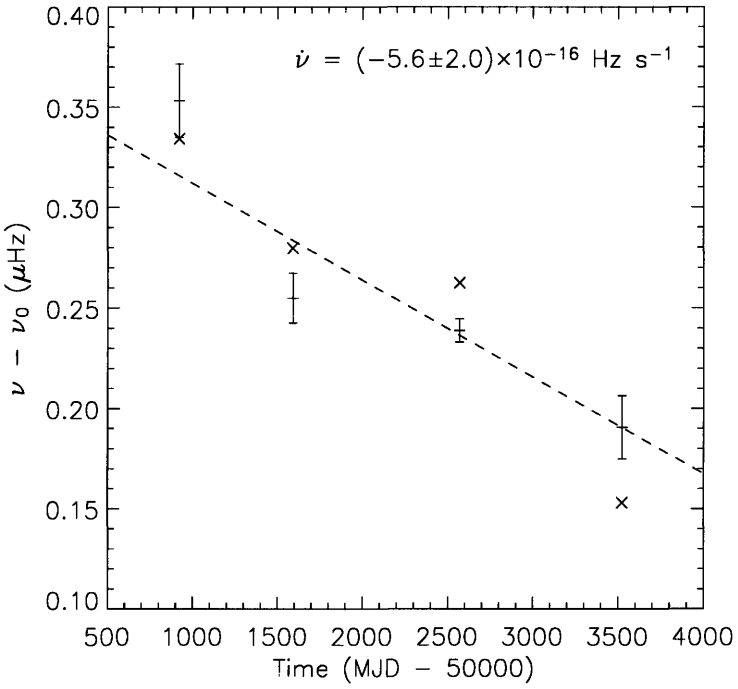

Figure 4-6 Constant-frequency measurements of the SAX J1808.4−3658 outbursts, showing the spin down of the star. The frequencies are relative to $\nu_0 = 400.975210$ Hz. The error bars are estimated using Monte Carlo simulations of phase residuals with the same noise properties as the actual outburst; they do not account for the uncertainty in the source position. The ×'s mark what the frequencies would be if the fit source position differed from the actual position by 2σ along the ecliptic plane in the direction of increasing RA. The same position error in the decreasing RA direction would move the frequency points by an equal amount in the opposite sense.

non-zero $\dot{\nu}$ measurements that most likely do not reflect the spin of the underlying neutron star. Using TEMPO to find the best-fit $\dot{\nu}$'s and applying Monte Carlos to estimate their uncertainties, we arrived at the 95% confidence intervals of Table 4.4. Excluding the 1998 outburst, which had the shortest span of timing data and thus the most poorly constrained $\dot{\nu}$, these 95% confidence upper limits were all of order $|\dot{\nu}| \lesssim 2.5 \times 10^{-14}$ Hz s^{-1}.

The uncertainties in the measurement of the frequency preclude phase connection between outbursts. During the 920 day gap between the 2002 and 2005 outbursts, the 6 nHz frequency uncertainty from 2002 would accumulate to a phase uncertainty of 0.5 cycles; the 2×10^{-16} Hz s^{-1} uncertainty in the long-term spin down would contribute 0.6 cycles. Worse, these estimates are best-case scenarios, since they assume that the spin down is constant.

During the 1998 and 2002 outbursts, we observed an abrupt change in the slope of the phase residuals at the end of the main outburst. We modeled these apparent instantaneous changes of frequency by including frequency glitches in our TEMPO fits. (While the TEMPO glitch models are useful in describing the data, we do not believe that we observed actual sudden changes in the spin frequency of the star, a point discussed in detail in §4.4.3.) Figure 4-7 shows the phase residuals of these outbursts and their best-fit glitch models. These models only employ an instantaneous change in frequency; including a phase jump or introducing a $\dot{\nu}$ after the events did not significantly improve the fits.[5]

In both outbursts, these apparent frequency changes coincide with the sudden drop in flux that marks the transition from the slow-decay stage to the flaring tail stage. At the same time, the fractional amplitudes of the fundamental and harmonic increase, and, in the case of 2002, the pulse profile shift occurs. (This pulse profile shift, discussed earlier in §4.3.2, is apparent in Figure 4-7 as the rapid advance of the fundamental phase.) If we view the phase residuals with respect to the pre-transition frequencies, as is the case in Figure 4-7, the residuals following the transition skew upward, indicating progressively increasing lags. This effect is more pronounced in 1998, but its coverage is far better in 2002. Interpreting these changes in slope as abrupt spin frequency changes, the frequencies would drop by 0.21 μHz and 0.03 μHz for 1998 and 2002, respectively. If we interpret them as the movement of a radiating spot, the rates of drift would be 6.5° day^{-1} and 1.0° day^{-1}, retrograde. The total observed shifts between the start of the flaring tail and the loss of the signal are substantial: 0.15 cycles (54°) in 1998 and 0.06 cycles (22°) in 2002. The data are not good enough to distinguish whether these drifts are continuous. For instance, it is possible that the hot spot made a retrograde jump ever time there was a flare.

We did not observe the main body of the 2000 outburst, so we cannot measure whether the apparent frequency decreased when it entered the flaring tail stage. But if it did, and if the decrease in the apparent frequency was of similar magnitude to that observed in 1998 and 2002, then including the main body of the 2000 outburst would raise the overall frequency of the outburst somewhat. This correction might put it in line with the other frequency measurements in Figure 4-6, reducing the large χ^2 statistic of the constant-$\dot{\nu}$ fit. Therefore we cannot conclude that the change in the observed frequency from one outburst to the next is incompatible with a linear progression.

During the 2005 outburst, the substantial pulse profile noise during the tail prevented us from measuring a change in apparent frequency. The uncertainty in the measurement of the frequency during the tail was 0.03 μHz, as estimated using Monte Carlo simulations of the profile noise, and the phase residuals jumped by as much as 0.1 cycles from one

[5]During the 2002 outburst, the absence of a phase jump refers only to the second harmonic, which we believe is a better tracer of the neutron star spin during this period of time (in agreement with the conclusions of Burderi et al. 2006).

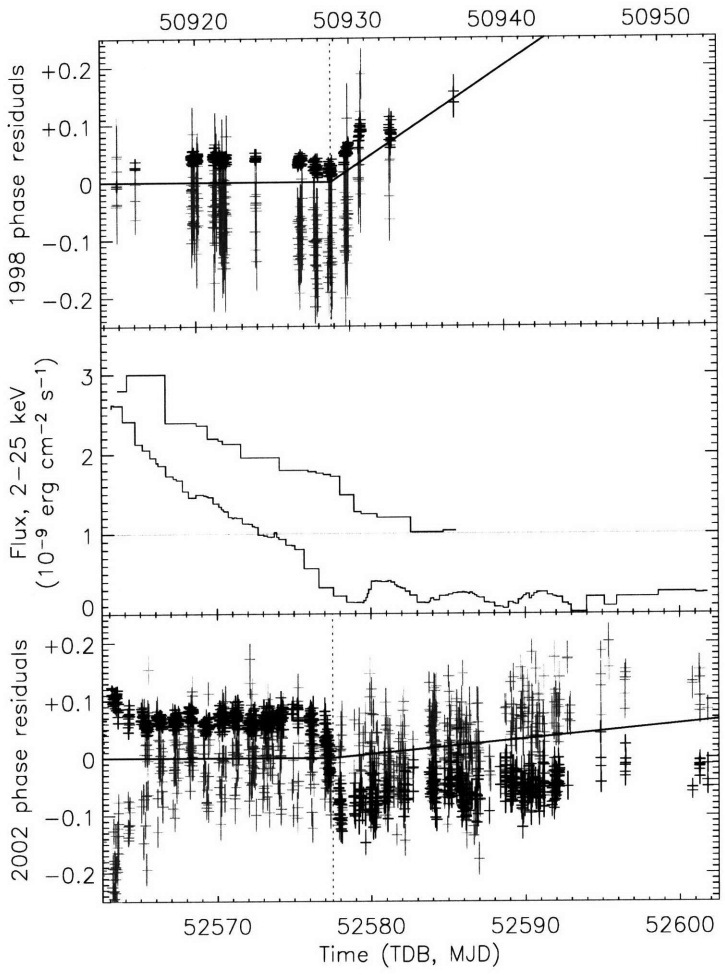

Figure 4-7 Comparison of the 1998 and 2002 glitch-like events. The phase plots show the phase residuals relative to a constant-frequency model for the fundamental (black points) and the harmonic (grey points), binned such that there is one point per observation. The black lines indicate the best timing models fit by TEMPO. The middle plot shows the 1998 and 2002 lightcurves for comparison. The 1998 lightcurve has been vertically offset by 1 erg cm^{-2} s^{-1} for clarity. The data are displayed such that the apparent changes in frequency are aligned. Notice that this alignment also has the effect of closely matching up the lightcurves.

Table 4.5. Binary parameter measurements

Outburst	$P_{\rm orb}$ (s)	$a_{\rm x} \sin i$ (light-ms)	$T_{\rm asc}$ (MJD, TDB)	$e \sin \omega$ (10^{-6})	$e \cos \omega$ (10^{-6})
1998 Apr	7249.1553(18)	62.8080(46)	50921.7584194(12)	-60 ± 64	-86 ± 64
2000 Feb	—	—	51591.8019861(40)	—	—
2002 Oct	7249.1565(6)	62.8147(31)	52570.0186514(9)	8 ± 57	41 ± 57
2005 Jun	7249.1547(24)	62.8282(109)	53524.9944192(32)	-173 ± 83	53 ± 83

Note. — We excluded the 2000 outburst when calculating everything but $T_{\rm asc}$ because its data were noisy and sparse.

Table 4.6. Combined timing parameters for SAX J1808.4−3658

Orbital period, $P_{\rm orb}$ (s) [a]	7249.156961(14)
Orbital period derivative, $\dot{P}_{\rm orb}$ (10^{-12} s s^{-1})	3.48(23)
Projected semimajor axis, $a_{\rm x} \sin i$ (light-ms)	62.8132(24)
Time of ascending node, $T_{\rm asc}$ (MJD, TDB)	52499.9602477(10)
Eccentricity, e (95% confidence upper limit)	$< 1.2 \times 10^{-4}$
Spin frequency, ν (Hz) [a]	400.975210240(11)
Spin frequency derivative, $\dot{\nu}$ (10^{-16} Hz s^{-1})	−5.6(2.0)

[a] $P_{\rm orb}$ and ν are specified for the time $T_{\rm asc}$.

observation to the next. If there was a smaller drift, as seen during the 2002 outburst, we would not necessarily detect it.

4.3.7 Evolution of the binary orbit

We fit the orbital parameters separately for each outburst. Table 4.5 lists the results. As expected, the values of $a_{\rm x} \sin i$ and $P_{\rm orb}$ were consistent among the outbursts. The fit parameters $e \sin \omega$ and $e \cos \omega$ were consistent with zero. We used them to improve significantly on previous upper limits on the eccentricity.

The measured time of ascending node advanced with each outburst, relative to the times expected if the period was constant. Figure 4-8 shows these $T_{\rm asc}$ residuals. This quadratic advance is consistent with a constant orbital period derivative of $\dot{P}_{\rm orb} = (3.5 \pm 0.2) \times 10^{-12}$ s s^{-1}, a significance of 15.6σ. Table 4.6 summarizes all the parameters for the pulse timing of SAX J1808.4−3658.

4.4 Discussion

Our analysis of multiple outbursts from SAX J1808.4−3658 allows us to greatly improve our understanding of the behavior of this low-mass X-ray binary. By comparing the observed frequency from each outburst, we can see the long-term spin down, which is too

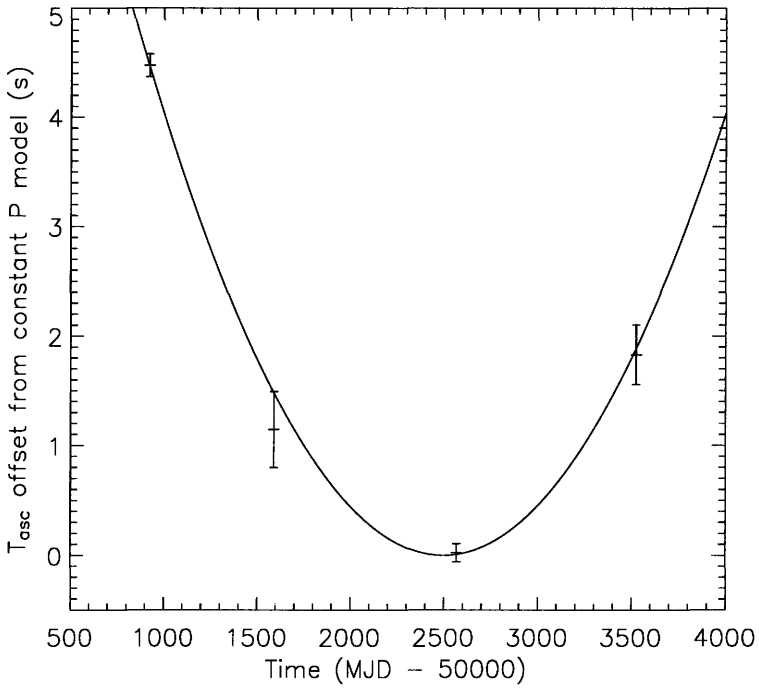

Figure 4-8 Measurement of an orbital period derivative. The points show the observed times of ascending node, relative to the expected times for a constant period. The $T_{\rm asc}$ of each outburst comes progressively later, indicating a period derivative of $(3.5 \pm 0.2) \times 10^{-12}$ s s^{-1}.

small to be detectable from a single outburst. Comparison of the pulse profiles from each outburst lead us to conclude that we are seeing characteristic, repeated profile changes as the outbursts progress, rather than a purely random noise process. Finally, fitting of the orbital parameters over the seven years of observation provides a greatly improved orbital ephemeris.

4.4.1 Long-term spin down

By observing the mean spin frequency of each outburst, we found that SAX J1808.4−3658 is spinning down at a rate of $\dot{\nu} = (-5.6 \pm 2.0) \times 10^{-16}$ Hz s^{-1}. This spin down results in a loss of rotational energy at a rate of $\dot{E} = 4\pi^2 I \nu \dot{\nu} = 9 \times 10^{33}$ erg s^{-1}, assuming a canonical value of $I = 10^{45}$ g cm^2 for the neutron star (NS) moment of inertia.

Most of this spin down occurs during X-ray quiescence; accretion torques during the outbursts play a minimal role. Over the seven years of our observations, the mean outburst frequency decreases by $\nu_{2005} - \nu_{1998} = -0.18 \pm 0.02$ μHz. Let us suppose that this frequency change happens only during the X-ray outbursts (which have a duty cycle of $\lesssim$5%). Since the outburst lightcurves are quite similar, it is reasonable to presume that each would contribute roughly the same frequency shift, thus splitting this frequency change into three equal steps. If the spin down is due to a constant $\dot{\nu}_{\rm outburst}$ that acts during the $\approx$20 days

of each outburst,[6] then

$$\dot{\nu}_{\rm outburst} \approx \frac{-0.18\ \mu{\rm Hz}}{3 \times 20\ {\rm d}} = -3.5 \times 10^{-14}\ {\rm Hz\ s^{-1}}\,. \tag{4.2}$$

By contrast, we were able to set stringent (95%-confidence) upper limits of $|\dot{\nu}| \lesssim 2.5 \times 10^{-14}$ Hz s^{-1} during the outbursts (Table 4.4). We conclude that the spin down is dominated by torques exerted during X-ray quiescence.

We will thus consider three possible sources of torque during quiescence: magnetic dipole radiation, the expulsion of matter by the magnetic field (i.e., the propeller effect), and gravitational radiation. In general, all three mechanisms can contribute to the observed spin down of SAX J1808.4−3658,

$$N_{\rm obs} = N_{\rm dipole} + N_{\rm prop} + N_{\rm gr}. \tag{4.3}$$

We discuss each below.

Magnetic dipole torque

A spinning dipolar magnetic field will produce a significant spin down during quiescence for the 10^8 G field strengths expected for a millisecond pulsar. Relativistic force-free MHD models of pulsar magnetospheres by Spitkovsky (2006) give a torque of $N_{\rm dipole} = -\mu^2(2\pi\nu/c)^3(1+\sin^2\alpha)$, where μ is the magnetic dipole moment and α is the angle between the magnetic and rotational poles. Pulse profile modeling of the 1998 outburst by Poutanen and Gierliński (2003) suggests that the magnetic hot spot is not far from the rotational pole, separated by an angle of $5-20°$. While other effects might also contribute to the spin down, the rotating magnetic field will always be present and provides an upper limit on the dipole moment:

$$\begin{aligned}\mu \ < \ & 0.77 \times 10^{26} \left(1+\sin^2\alpha\right)^{-1/2} \\ & \times \left(\frac{I}{10^{45}\ {\rm g\ cm^2}}\right)^{1/2} \left(\frac{\nu}{401\ {\rm Hz}}\right)^{-3/2} \\ & \times \left(\frac{-\dot{\nu}}{5.6 \times 10^{-16}\ {\rm Hz\ s^{-1}}}\right)^{1/2} {\rm G\ cm^3}\,.\end{aligned} \tag{4.4}$$

For $\alpha = 15°$, this upper limit on the dipole is 0.75×10^{26} G cm^3, yielding a field strength of roughly $B = 1.5 \times 10^8$ G at the magnetic poles.[7] We emphasize that this upper limit on the magnetic field is for a purely dipolar field. The presence of higher-order multipoles would require a stronger field at the NS surface to produce the observed $\dot{\nu}$.

If magnetic dipole torque is a significant contributor to the spin down of SAX J1808.4−3658, then the source may behave like a rotation-powered pulsar during quiescence, producing radio pulsations and a particle wind. The heating of the companion by a particle wind has been invoked as an explanation of why the companion is significantly brighter than expected

[6] In reality, $\dot{\nu}$ would almost certainly *not* be constant during as the accretion rate varies, but for argument's sake we make the most conservative assumptions possible. A varying $\dot{\nu}$ would require that it be sometimes greater than the value from equation (4.2), making it even less plausible that it would escape detection.

[7] The Spitkovsky (2006) formula for $N_{\rm dipole}$ differs substantially from the classically derived torque due to a rotating dipole in a vacuum, $N_{\rm vac} = \frac{2}{3}\mu^2(2\pi\nu/c)^3\sin^2\alpha$, especially for small α: for $\alpha = 15°$, the derived limit is approximately one fifth of the vacuum value.

in the optical. Burderi et al. (2003) predicted a dipole moment of $\mu = 5 \times 10^{26}$ G cm^3 based on the optical observations, somewhat higher than our approximate upper limit on μ, but most likely within the uncertainties of the model. A similar analysis by Campana et al. (2004) found the needed irradiation luminosity to be $L_{\rm x} = (4^{+3}_{-1}) \times 10^{33}$ erg s^{-1}, compatible with the observed $\dot{E} = 9 \times 10^{33}$ erg s^{-1} loss of rotational energy. No radio emission has been detected during quiescence. The upper limits of 0.5 mJy (Gaensler et al. 1999; Burgay et al. 2003) are not particularly constraining.

The X-ray luminosities of isolated millisecond pulsars, for which magnetic dipole radiation is the primary spin-down mechanism, shows a strong correlation with their rates of rotational energy loss. From the tables compiled in Zavlin (2006) and Cameron et al. (2007), the 5–10 keV X-ray luminosity goes as $L_{\rm x} \propto \dot{E}^{1.13}$ with less than a quarter decade of scatter. Based on this empirical relation, we would expect a quiescent luminosity for SAX J1808.4−3658 of 5×10^{30} erg s^{-1}. However, this prediction is a factor of ten lower than the observed quiescent fluxes of 8×10^{31} erg s^{-1} and 5×10^{31} erg s^{-1} (Campana et al. 2002; Heinke et al. 2007), suggesting other mechanisms for quiescent emission are at work.

Magnetic propeller torque

The propeller effect offers another possible explanation for the observed spin down during quiescence. If the Keplerian corotation radius ($r_{\rm co} = [GM/4\pi^2\nu^2]^{1/3} \approx 31$ km) is less than the magnetospheric radius r_0, at which point the infalling matter couples to the magnetic field, then the magnetic field will accelerate the matter, possibly ejecting it from the system (Illarionov and Sunyaev 1975). The torque exerted on the neutron star by propeller ejection of matter at a rate $\dot{M}_{\rm ej}$ depends on the details of the interaction between the pulsar magnetosphere and the accretion disk. However, we can parametrize this torque as

$$\begin{aligned} N_{\rm prop} &= -n\dot{M}_{\rm ej}(GMr_0)^{1/2} \\ &= -n(r_0/r_{\rm co})^{1/2}\dot{M}_{\rm ej}(GMr_{\rm co})^{1/2}\,, \end{aligned} \tag{4.5}$$

where the detailed physics determines the dimensionless torque n, which is zero for $r_0 = r_{\rm co}$ and of order unity for $r_0 \gtrsim 1.1\, r_{\rm co}$ (Ekşi et al. 2005).

We can then roughly estimate the rate at which matter would need to be ejected from the system during quiescence to account for the observed spin down:

$$\begin{aligned} \dot{M}_{\rm ej} < & -2.3 \times 10^{-12}\, n^{-1}(r_0/r_{\rm co})^{-1/2} \\ & \times \left(\frac{I}{10^{45}\ {\rm g\ cm^2}}\right)\left(\frac{M}{1.4\ M_\odot}\right)^{-2/3}\left(\frac{\nu}{401\ {\rm Hz}}\right)^{1/3} \\ & \times \left(\frac{-\dot{\nu}}{5.6 \times 10^{-16}\ {\rm Hz\ s^{-1}}}\right)\ M_\odot\ {\rm yr}^{-1}\,. \end{aligned} \tag{4.6}$$

As a consistency check, we note that this upper limit does not exceed the predicted long-term mass transfer rate for the binary, $1.0 \times 10^{-11}\ M_\odot$ yr^{-1}, which is driven by gravitational radiation emission due to the binary orbit (Bildsten and Chakrabarty 2001). Indeed, not all the mass lost by the donor star will necessarily reach the pulsar magnetosphere during quiescence and be propelled outward; most of it would queue up in the accretion disk and later reach the NS during an outburst. Galloway and Cumming (2006) found that the mass transfer is roughly conservative, albeit with enough uncertainty that propeller mass loss as large as the above $\dot{M}_{\rm ej}$ limit is not ruled out.

Even if propeller spin down provides the dominant quiescent torque, the resulting ejection of matter from the system would not greatly affect the binary orbit. The timescale for propeller spin down is proportional to the timescale for the ejection of mass: $\dot{P}_{\rm orb}/P_{\rm orb} \propto \dot{M}_{\rm ej}/M_{\rm c}$, where $M_{\rm c} \approx 0.05\ M_\odot$ is the mass of the companion. Applying the above $\dot{M}_{\rm ej}$ gives $\dot{M}_{\rm ej}/M_{\rm c} = (20\ {\rm Gyr})^{-1}$, far longer than the observed orbital evolution timescale of $\dot{P}_{\rm orb}/P_{\rm orb} = (66\ {\rm Myr})^{-1}$. More refined calculations using the arguments of Tauris and van den Heuvel (2006) yield a propeller timescale of 6 Gyr, still far too large. Clearly there are other contributions to the orbital evolution; we discuss some in §4.4.7.

Gravitational radiation torque

The observed spin down provides an upper limit on the spin-down due to gravitational radiation from this source. (See 1.6.2 for a discussion of the possible role of gravitational waves in NS evolution.) The mass quadrupole moment of the star, Q, determines the torque produced by gravitational radiation: $N_{\rm gr} = -\frac{32}{5}GQ^2(2\pi\nu/c)^5$. For our measured $\dot{\nu}$, this sets an upper limit of

$$Q < 4.4\times10^{36}\left(\frac{I}{10^{45}\ {\rm g\ cm^2}}\right)^{1/2}\left(\frac{\nu}{401\ {\rm Hz}}\right)^{-5/2} \times\left(\frac{-\dot{\nu}}{5.6\times10^{-16}\ {\rm Hz\ s^{-1}}}\right)^{1/2}\ {\rm g\ cm^2}, \quad (4.7)$$

or $Q \lesssim 10^{-8}\,I$. The strain amplitude of the resulting gravitational waves, averaged over all NS orientations, is $h_c = 115G\nu^2Q/dc^4$ (Brady et al. 1998), giving a characteristic strain at Earth of $h_c = 6\times10^{-28}$. This strain is undetectable by current or planned gravitational wave experiments. For Advanced LIGO, with a strain sensitivity of $\sim3\times10^{-24}\ {\rm Hz}^{-1/2}$ in the 100–400 Hz range (Fritschel 2003), even a search using an accurate phase model would require years of integration time. Note that the dependence of $N_{\rm gr}$ on the ν is very strong, so it is quite possible that gravitational wave emission produces larger spin downs in faster ($\approx$700 Hz) rotators.

4.4.2 Pulse profile variability

The evolution of the pulse profile is clearly not purely stochastic. With multiple outbursts, we are able to note for the first time that the pulse profile seems to take on similar shapes at similar times in the outbursts, as illustrated in Figure 4-3. These characteristic changes in the pulse profiles suggest that the emitting regions of the NS are changing shape and position as the outbursts progress. The consistency of these changes, along with the consistency of the outburst lightcurves, suggests that as the accretion disk empties onto the star, the geometry of the disk, the accretion funnels, and the resulting hot spots evolve in a similar manner for each outburst.

The most conspicuous trend in the changing pulse profiles during the outbursts is the strong anti-correlation between the fractional amplitude of the second harmonic and the source luminosity: $r_2 \propto L^{-1/2}$. In Chapter 5, we find that such anti-correlations are present in all the AMPs that have been well-observed over a wide range of fluxes, and we consider some theoretical implications of this surprising result.

4.4.3 Motion of the hot spot

During the 1998, 2002, and 2005 outbursts, we observed clear trends in the phase residuals that suggest that the emitting regions do not remain at a fixed longitude. During 2002 and 2005, an abrupt phase change in the fundamental at the end of the main body of the outburst produces an advance of the pulse peak that corresponds to a shift of the hot spot by $\approx 50°$ eastward.[8] These shifts are simultaneous with and occur on the same 3–4 day timescale as the sudden drops in luminosity at the end of the main outbursts. During 1998 and 2002, the phase residuals of both harmonics begin gradually increasing during the flaring tails of the outburst, corresponding to a westward drift of the hot spots. Motion of the hot spot has also been suggested to explain phase residuals in GX 1+4 and RX J0812.4−3114 (Galloway et al. 2001) and XTE J1814−338 (Papitto et al. 2007).

These trends in the phase residuals almost certainly represent motion of the observed hot spot rather than frequency glitches. Glitches are rapid increases in the spin frequency of the NS due to imperfect coupling between the crust and more rapidly rotating, superfluidic lower layers (e.g. Anderson and Itoh 1975). This interaction occurs well below the accretion layer, and it would not be expected to coincide with or have the same timescale as rapid changes in the accretion rate. Furthermore, the observed changes in the apparent frequency are drops in frequency, not increases.

When discussing the motion of the hot spots, the longitudes of the magnetic poles provide natural meridians from which to measure phase. Since their movement would require the realignment of currents in the core and crust, the magnetic poles remain at fixed positions for timescales far longer than the outbursts. The suppression of regions of the field due to accretion also occurs on long timescales (Cumming et al. 2001).

For high-field pulsars, the magnetospheric radius is far from the star, and the accretion column follows field lines that reach the NS surface near the magnetic pole. This is not necessarily the case for low-field pulsars. A closer accretion disk will intersect more curved field lines, which terminate farther from the poles. In the previous section, we described how the Alfvén radius can move outward from roughly $1.5R$ to $3R$ as the accretion rate drops. As the disk recesses, it will intersect decreasingly curved field lines that are rooted closer to the poles, causing the hot spots at the bases of the accretion columns to also approach the poles.

This simple picture can explain the observed phase shift as the luminosity rapidly drops during the end of the 2002 and 2005 outbursts. In both cases, the luminosity decreases by about a factor of 4. A change in $\dot{M}$ by this magnitude would cause the Alfvén radius to move outward by a factor of 1.5 and the inner edge of the accretion disk to move outward by a similar amount. This change will almost certainly cause material removed from the inner edge of the disk to attach to a different set of field lines, with the larger radius favoring lines that attach closer to the pole. If the hot spot tends to be to the west of the pole, as seen in Romanova et al. (2004) for a magnetic pole an angle of $\alpha = 30°$ from the rotational pole, then the attachment to different field lines would produce an eastward drift as observed. That said, these MHD simulations appear to have strong, chaotic dependencies on their parameters. (For $\alpha = 15°$, the hot spot is south of the magnetic pole; for 30°, west; and for 45°, north!) More work is needed to better model these observations.

The slow drifts seen during the tails of the 1998 and 2002 outbursts are difficult to explain. The flares during the tail cause the luminosity to change in a periodic manner,

[8] For a more natural description, we adopt the Earth-based convention of longitude: earlier pulse arrivals ≡ prograde hot spot motion ≡ eastward shift, and vice versa.

so we cannot expect a monotonic motion of the accretion disk's inner edge. The net drift during the tail of the 2002 outburst is of the same magnitude as the rapid phase shift that happens right before the tail begins, suggesting the drift may be a relaxation of the accretion column back to its original location.

4.4.4 Comparison with previous spin frequency measurements

There have been a number of previous reports of short-term $\dot{\nu}$ measurements made during outbursts of several accreting millisecond pulsars including XTE J0929−314 (Galloway et al. 2002), SAX J1808.4−3658 (Morgan et al. 2003; Burderi et al. 2006), XTE J1751−305 (Markwardt 2003), IGR J00291+5934 (Falanga et al. 2005; Burderi et al. 2007), and XTE J1814−334 (Papitto et al. 2007). Some of the reported $\dot{\nu}$ values have been surprisingly large given the estimates of $\dot{M}$ during the outbursts, possibly violating a basic prediction of magnetic disk accretion theory: that accretion torques cannot exceed the characteristic torque $N_{\rm char} = \dot{M}(GMr_{\rm co})^{1/2}$ exerted by accreting Keplerian material at the corotation radius (e.g., Ghosh and Lamb 1979).

In the particular case of SAX J1808.4−3658, spin derivatives as large as a few times 10^{-13} Hz s^{-1} near the outburst peak were reported (Morgan et al. 2003; Burderi et al. 2006), corresponding to accretion torques exceeding $N_{\rm char}$ for this source. However, these studies calculated pulse phase residuals using only a single harmonic; Morgan et al. (2003) reported $\dot{\nu}$ detections using only the fundamental, while Burderi et al. (2006) measuring the phase from the second-harmonic alone after noting the sudden phase shift of the fundamental in the middle of the 2002 outburst. Our results in §3.6 indicate that both of these approaches are likely to be contaminated by pulse shape changes, at least in the case of SAX J1808.4−3658.

Figure 4-9 illustrates this point. Taking the phases of the harmonic components as direct spin measurements can produce large values of $\dot{\nu}$ during the peak of the 2002 outburst. Fitting only the fundamental's residuals during the first 10 days of the outburst, we find $\dot{\nu} = (-1.2 \pm 0.4) \times 10^{-13}$ Hz s^{-1}. On the other hand, using only the second harmonic for the same interval, we find $\dot{\nu} = (5.3 \pm 0.1) \times 10^{-13}$ Hz s^{-1}, in good agreement with the Burderi et al. (2006) measurement. Because the pulse shape is changing rapidly during this part of the outburst, the pulse arrival times cannot be accurately determined. We therefore cannot reliably use this part of the outburst to measure the spin of the NS. Note that if we exclude this region of large pulse shape variability, the remaining phase residuals are consistent with a constant spin frequency over the outburst interval (§3.6). From an examination of all the outbursts of SAX J1808.4−3658 (excluding regions of large pulse shape variability), our work sets an upper limit of $|\dot{\nu}| \lesssim 2.5 \times 10^{-14}$ Hz s^{-1}.

We thus conclude that the past measurements of short-term $\dot{\nu}$ in SAX J1808.4−3658 are unreliable. The analysis technique we described in §2.3 can mitigate the effects of pulse shape variability to some extent, but attempts to measure $\dot{\nu}$ in accreting pulsars must properly account for these variability effects, and in some instances these effects may prevent such measurements. The $\dot{\nu}$ measurements reported in other accreting millisecond pulsars must all be reevaluated in this light; all the apparent violations of the $N \leq N_{\rm char}$ limit predicted by theory may be owing to spurious measurements caused by pulse shape variability. However, at least some accreting millisecond pulsars are observed to have relatively stable pulse shapes, indicating that accurate short-term $\dot{\nu}$ measurements are possible and that previous measurement of these sources should be reliable.

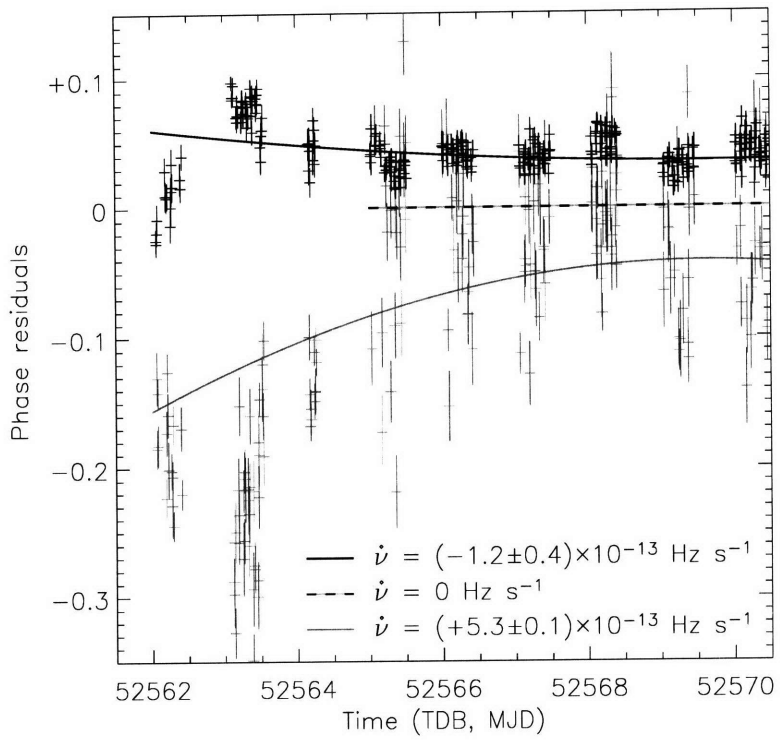

Figure 4-9 Fitting a frequency model using only the fundamental (black) or the harmonic (grey) produces false $\dot{\nu}$ measurements during the peak of the 2002 outburst. The data points are the 512 s phase residuals relative to the best constant-frequency model for the 2002 outburst. The solid lines give the best constant-$\dot{\nu}$ models, fit solely to the fundamental or the second harmonic. The dashed line shows the constant-frequency model derived using both, combined via equation (3.16); this fit did not use the points prior to MJD 52565.

4.4.5 Constraints on the magnetic field

We showed in §4.1.1 that the condition $N_{\rm dipole} \leq N_{\rm obs}$ implies that the magnetic dipole moment $\mu \lesssim 0.8 \times 10^{26}$ G cm^3. This limit is consistent with the range for μ implied by the observation of accretion-powered pulsations throughout the outbursts (Psaltis and Chakrabarty 1999). At low accretion rates, the field cannot be so strong that it centrifugally inhibits matter from reaching the NS; during times of high accretion, it must be strong enough to truncate the disk above the stellar surface in order for there to be pulsations. The dimmest observation in which we observed pulsations was in 1998, with a flux in the 2–25 keV band of 1.5×10^{-11} erg cm^{-2} s^{-1}; the brightest was at the peak of the 2002 outburst, 2.62×10^{-9} erg cm^{-2} s^{-1}. These fluxes, along with an improved estimate of the Eddington luminosity from observations of photospheric radius expansion bursts (Galloway and Cumming 2006), give us new limits on the range of accretion rates at which pulsations have been detected, relative to the Eddington rate $\dot{M}_{\rm E}$: $\dot{M}_{\rm min} = 1.8 \times 10^{-4}\ \dot{M}_{\rm E}$ and $\dot{M}_{\rm max} = 0.03\ \dot{M}_{\rm E}$. (We have made the usual assumption that $L \propto \dot{M}$.) These limits allow us to update the range for μ derived in Psaltis and Chakrabarty (1999), equations (11)

and (12):[9]

$$0.2 \times 10^{26} \text{ G cm}^3 \lesssim \mu \lesssim 6 \times 10^{26} \text{ G cm}^3 . \tag{4.8}$$

Taken together, we obtain a fairly narrow allowed range for the magnetic dipole moment,

$$0.2 \times 10^{26} \text{ G cm}^3 \lesssim \mu \lesssim 0.8 \times 10^{26} \text{ G cm}^3 , \tag{4.9}$$

which corresponds to a surface dipole magnetic field strength of (0.4–1.4) $\times 10^8$ G.

4.4.6 Constraints on accretion torques

Even though we did not detect an accretion-induced $\dot{\nu}$ during the outbursts, our new upper limits on $|\dot{\nu}|$ provide far stronger constraints on the accretion physics of low-B systems such as SAX J1808.4−3658 than previous measurements. Following the earlier analysis of the 1998 outburst by Psaltis and Chakrabarty (1999), the lower limit on the spin frequency derivative predicted by accretion torque theory during an outburst with an average accretion rate of $\dot{M}_{\rm avg} \approx \frac{1}{3}\dot{M}_{\rm max} \approx 0.01\ \dot{M}_{\rm E}$ is

$$\begin{aligned} \dot{\nu} \gtrsim\ & 2 \times 10^{-14}\, \eta \left(\frac{I}{10^{45} \text{ g cm}^2} \right)^{-1} \left(\frac{R}{10 \text{ km}} \right)^{3/2} \\ & \times \left(\frac{M}{1.4\ M_\odot} \right)^{1/2} \left(\frac{\dot{M}_{\rm avg}}{0.01\ \dot{M}_{\rm E}} \right) \text{ Hz s}^{-1} , \end{aligned} \tag{4.10}$$

where η is a dimensionless parameter encapsulating the disk-magnetosphere interaction. (Refer to Ghosh and Lamb 1979 for a discussion of the physics that goes into this parameter.) η is strongly dependent on the magnetospheric radius. For $r_0 \approx r_{\rm co}$, the NS will be in spin equilibrium with the accreted matter and η will be small. From the $\dot{\nu}$ confidence intervals in Table 4.4, the probability that we would have missed detecting the resulting 2×10^{-14} Hz s^{-1} spin up is 0.15%, suggesting that $\eta < 1$ and the source is near spin equilibrium during the outbursts.

4.4.7 Discussion of increasing $P_{\rm orb}$

Our seven year baseline for timing analysis provides the most precise measurements yet of the orbital period of SAX J1808.4−3658. We find that the orbital period is increasing at a rate $\dot{P}_{\rm orb} = 3.5(2) \times 10^{-12}$ s s^{-1}. This $\dot{P}_{\rm orb}$ lies somewhat outside the 90% confidence upper limit set by Papitto et al. (2005) using the 1998–2002 outbursts, most likely owing to the more limited baseline available in that analysis.

It is interesting to compare our measurement with theoretical expectations. For orbital periods $\lesssim$3 hr, mass transfer is LMXBs is driven by angular momentum losses due to gravitational radiation from the binary (Kraft et al. 1962), since magnetic braking torques are thought to be ineffective in this regime (Rappaport et al. 1983; Spruit and Ritter 1983). For SAX J1808.4−3658, the $\dot{M}$ predicted by this mechanism is consistent with observationally inferred long-term average value of $\dot{M} = 4 \times 10^{-11}\ M_\odot$ yr^{-1} (Bildsten and Chakrabarty 2001). For conservative mass transfer from a degenerate (brown dwarf) donor, this predicts orbital expansion on a time scale $P_{\rm orb}/\dot{P}_{\rm orb} = 3$ Gyr (see, e.g., Tauris and van

[9] In deriving this range for μ, we make the same conservative assumptions as Psaltis and Chakrabarty (1999): the Ghosh and Lamb (1991) boundary layer parameter ranges on $0.1 < \gamma_B(\dot{M}) < 1$; the NS mass is $1.4\ M_\odot < M < 2.3\ M_\odot$; and the NS radius is 10 km $< R <$ 15 km.

den Heuvel 2006). By contrast, our measured value of $P_{\rm orb}/\dot{P}_{\rm orb} = 66$ Myr is an order of magnitude more rapid.

The origin of the anomalously large $\dot{P}_{\rm orb}$ in SAX J1808.4−3658 is unclear, although we note that unexpectedly large $\dot{P}_{\rm orb}$ values have also been observed in several other LMXBs including 4U 1820−30 (van der Klis et al. 1993), EXO 0748−676 (Wolff et al. 2002), and 4U 1822−371 (Hellier et al. 1990). As pointed out by Chakrabarty and Morgan (1998), the binary parameters of SAX J1808.4−3658 are very similar to those of the so-called "black widow" millisecond radio pulsars, all of which are ablating their low-mass companions (see, e.g., Fruchter et al. 1990). If SAX J1808.4−3658 does indeed turn on as a radio pulsar during X-ray quiescence (Burderi et al. 2003; Campana et al. 2004; see also §4.1.1), it may be a black widow system as well, consistent with its very low donor mass. As such, it is interesting to note that a large and variable $\dot{P}_{\rm orb}$, both positive and negative, has been measured in two black widow pulsars (Arzoumanian et al. 1994; Doroshenko et al. 2001).

Although mass loss from the companion through an ablated wind would tend to increase $\dot{P}_{\rm orb}$, the mass loss rate required to explain the observed $\dot{P}_{\rm orb}$ in SAX J1808.4−3658 is $\sim 10^{-8} M_\odot$ yr^{-1} (Tauris and van den Heuvel 2006); this is unphysically large given our measured pulsar spindown rate (§4.1), which sets the pulsar luminosity available for irradiating the companion. This explanation for $\dot{P}_{\rm orb}$ is also inadequate in the black widow pulsars, where the orbital period variability is quasi-cyclic on a $\simeq$10 yr time scale (Arzoumanian et al. 1994; Doroshenko et al. 2001). In those systems, it has been suggested that tidal dissipation and magnetic activity in the companion is responsible for the orbital variability, requiring that the companion is at least partially non-degenerate, convective, and magnetically active (Arzoumanian et al. 1994; Applegate and Shaham 1994; Doroshenko et al. 2001). If this mechanism is active in SAX J1808.4−3658, we would expect quasi-cyclic variability of $P_{\rm orb}$ to to reveal itself over the next few years.

Chapter 5

A strong anti-correlation between harmonic amplitudes and flux

This chapter is adapted from the paper "A Strong Correlation Between Flux and Sinusoidal Pulse Profiles in Six Accreting Millisecond Pulsars" by Jacob M. Hartman, Alessandro Patruno, Anna Watts, and Deepto Chakrabarty, which is in preparation for submission to the *Astrophysical Journal*.

5.1 Introduction

Accreting millisecond pulsars (AMPs) provide a unique laboratory for studying the process of accretion in the presence of strong gravity and magnetic fields. The persistent pulsations arise as accreting material is channeled by the star's magnetic field, leading to X-ray–emitting hot spots at or near the magnetic poles that are modulated by the stellar rotation. Unlike pulsars with high magnetic fields, the scale heights for the accretion column and magnetosphere of the relatively low-field AMPs are of the same order as the neutron star radius (see §1.1 and references therein). This allows for the coupling between the magnetosphere and accretion disk to apply large torques to the neutron star, spinning it up to millisecond periods (Alpar et al. 1982; Radhakrishnan and Srinivasan 1982). Additionally, the pulsations have been observed throughout X-ray outbursts from these sources, which can span multiple orders of magnitude in flux, and hence inferred accretion rate. The result is a particularly complex challenge for the modeling of such systems, both analytical (e.g., Rappaport et al. 2004) and numerical (e.g., Romanova et al. 2004). The X-ray pulse timing of such sources can offer insight into this problem.

Given the wide range of accretion rates at which pulsations are detected from the AMPs, we expect the accretion geometry to vary, causing changes in the location and shape of the X-ray emitting hot spot at the base of the accretion column. These effects will cause systematic changes in the pulse profiles. In the previous chapter, we discussed the observed pulse profile variability in the 401 Hz accretion-powered pulsar SAX J1808.4−3658 (Wijnands and van der Klis 1998). In particular, we discovered a striking relationship between the fractional amplitude of the 802 Hz harmonic and the luminosity: $r_2 \propto L^{-1/2}$. This relation is quite robust, holding up over 2.5 orders of magnitude in flux and all four of the outbursts observed by *RXTE*, with the same proportionality constant for all the outbursts.

This result led us to examine the six other AMPs for which public data is available to search for a similar relationship: XTE J1751−305 (435 Hz; Markwardt et al. 2002),

Table 5.1. Observations analyzed for each AMP

	Data range (MJD)	# obs.	Time (ks)	Avg. # PCUs
XTE J1751−305	52368.6 – 52394.9	39	432.2	2.51
XTE J0929−314	52403.4 – 52456.1	42	137.5	2.97
XTE J1807−294	52697.4 – 52849.2	111	500.0	3.20
XTE J1814−338	52796.0 – 52855.5	90	447.0	3.26
HETE J1900.1−2455	53537.0 – 53927.9	112	405.5	3.07
IGR J00291+5934	53342.2 – 53359.6	48	392.9	2.85

Note. — All pointed observations within the given time intervals were included in our analysis.

XTE J0929−314 (185 Hz; Galloway et al. 2002), XTE J1807−294 (191 Hz; Markwardt et al. 2003), XTE J1814−338 (314 Hz; Markwardt and Swank 2003), HETE J1900.1−2455 (377 Hz; Kaaret et al. 2006), and IGR J00291+5934 (598 Hz; Galloway et al. 2005). In this chapter, we find that in every source for which a wide range of fluxes is well-sampled, there is indeed an anti-correlation between the fractional amplitude of the harmonic and the flux. We discuss some simple models that could produce such relationships.

5.2 Observations and analysis

The six AMPs were all observed repeatedly during their outbursts by the *RXTE* Proportional Counter Array. These sources were all observed in one or a combination of the `GoodXenon`, `E_62US_64M_0_1S`, and `E_125US_64M_0_1S` data modes. We used all the available pointed data for each source. We ignored raster scans, commonly conducted immediately after the discoveries of these sources to constrain their position, as such observations are of limited use for determining pulsation amplitudes due to their varying background rates. Table 5.1 lists the observations in our analysis. The average number of active PCUs during the observations was ≈3 for all the sources, providing an effective PCA area of roughly 3800 cm^2. To simplify the analysis, all the observations were binned to a time resolution 122 μs; using higher time resolutions provides no benefit. Only photons in the 2.5–15 keV range were included.

We shifted the photon arrival times to the Earth's geocenter using the positions listed in Table 5.2. (We shift the times to the geocenter rather than the solar system barycenter because TEMPO itself performs the barycentering.) We then applied the *RXTE* fine clock correction, which provides absolute time measurements with errors of less than 3.4 μs (99% confidence; Jahoda et al. 2006). Finally, we filtered the data to remove Earth occultations, intervals of unstable pointing, and (in the cases of XTE J1814−338 and HETE J1900.1−2455) thermonuclear X-ray bursts.

We then divided the data into 512 s intervals for the purposes of phase connection. We fit timing solutions to the fundamental component of the persistent pulsations from each source; i.e., time of arrival for each 512 s interval was determined by measuring the phase shifts relative to a pure sinusoid at the spin frequency. We adopted this simplified approach

Table 5.2. Timing solutions for the AMPs

	R.A.	Decl.	Data range (MJD)	ν (Hz)	$\dot{\nu}$ (10^{-13} Hz)	$a_x \sin i$ (light-ms)	$P_{\rm orb}$ (s)	$T_{\rm asc}$ (MJD, TDB)
XTE J1751−305	17:51:13.49	−30:37:23.4	52369.6 – 52379.8	435.31799356(9)	—	10.16(5)	2545.37(2)	52368.01290(1)
XTE J0929−314	09:29:20.194	−31:23:03.41	52403.4 – 52442.3	185.105254373(5)	—	5.99(1)	2614.751(3)	52405.48675(1)
XTE J1807−294	18:06:59.8	−29:24:30	52697.4 – 52793.4	190.6235070189(3)	+0.28(1)	4.82(2)	2404.417(2)	52720.67559(2)
XTE J1814−338	18:13:39.03	−33:46:22.3	52796.0 – 52844.2	314.3561086725(9)	−0.73(2)	390.62(4)	15388.723(3)	52794.960405(5)
HETE J1900.1−2455	19:00:08.65	−24:55:13.7	53538.7 – 53582.8	377.296172043(7)	—	18.44(1)	4995.260(6)	53549.13094(1)
IGR J00291+5934	00:29:03.06	+59:34:19.0	53342.2 – 53352.0	598.892130909(7)	+8.3(4)	64.99(2)	8844.078(2)	53345.161926(9)

Table 5.3. Best-fit power laws for the harmonic fractional amplitude vs. flux

	Slope, γ	χ^2	ν
SAX J1808.4−3658	-0.50 ± 0.01	206.2	231
XTE J1751−305	-0.66 ± 0.03	23.4	29
XTE J0929−314	-0.73 ± 0.17	409.9	30
XTE J1807−294	-0.47 ± 0.04	136.6	77
XTE J1814−338	-0.26 ± 0.07	306.1	75
HETE J1900.1−2455	-0.59 ± 0.16	202.4	61
IGR J00291+5934	-0.70 ± 0.04	184.0	33

rather than applying the harmonic weighting scheme discussed in §3.3.1 because our timing solutions are not of capital importance for this analysis. The harmonic fractional amplitudes are unaffected by phase errors, so long as the fit frequency evolution has errors of $\ll 0.5/512$ s. Using the fundamental alone provides more than sufficient accuracy for these purposes.

The timing solutions were constructed with the TEMPO pulsar timing program, as discussed in §3.3. We fit parameters for the spin frequency, the spin frequency derivative (if necessary), and circular orbits (using the `ELL1` orbital model). In no case did we have a sufficient baseline of data to improve upon the position, so we held fixed the positions from the literature. Parameter fitting with TEMPO and measurements of new TOAs with the timing solutions were bootstrapped until we arrived at the best timing solution. The parameters of these solutions are listed in Table 5.2. Note that the uncertainties in this table are estimated by TEMPO. They do *not* account for pulse profile variability, which is clearly evident in some of the sources; better uncertainty estimates require Monte Carlo simulations, as described in §3.3.2.

For each 512 s interval, the photons were folded according to our phase models. The fractional amplitudes were calculated using the methods outlined in §3.3.1. Backgrounds were estimated using the `pcabackest` FTOOL. Their uncertainties were estimated using the techniques of Vaughan et al. (1994).

2–25 keV fluxes were approximated following the technique described in §2.1 of Galloway et al. (2006). The observed fluxes were modeled with an absorbed blackbody and power-law model, and if necessary an iron line in the case of significant residuals above 6.4 keV.

Power law fits were calculated simply by a linear fits to the logarithms of the fluxes and the fractional amplitudes of the second harmonic and noted the resulting slope and its uncertainty. In all cases, the reduced χ^2 statistics ($\chi^2_\nu \equiv \chi^2/\nu$ for ν degrees of freedom) for these fits were somewhat greater than unity. To better estimate the uncertainty of the power law slopes, we rescaled the fractional amplitude uncertainties such that $\chi^2_\nu = 1$.

5.3 Results

Figure 5-1 shows the background-corrected harmonic fractional amplitude vs. flux for the other six accretion-powered millisecond pulsars. Power law slopes range from -0.26 to -0.73; these fits are listed in Table 5.3. Clearly some of the fits are much better than others: we do not have many points at low fluxes for XTE J0929−314, XTE J1814−338,

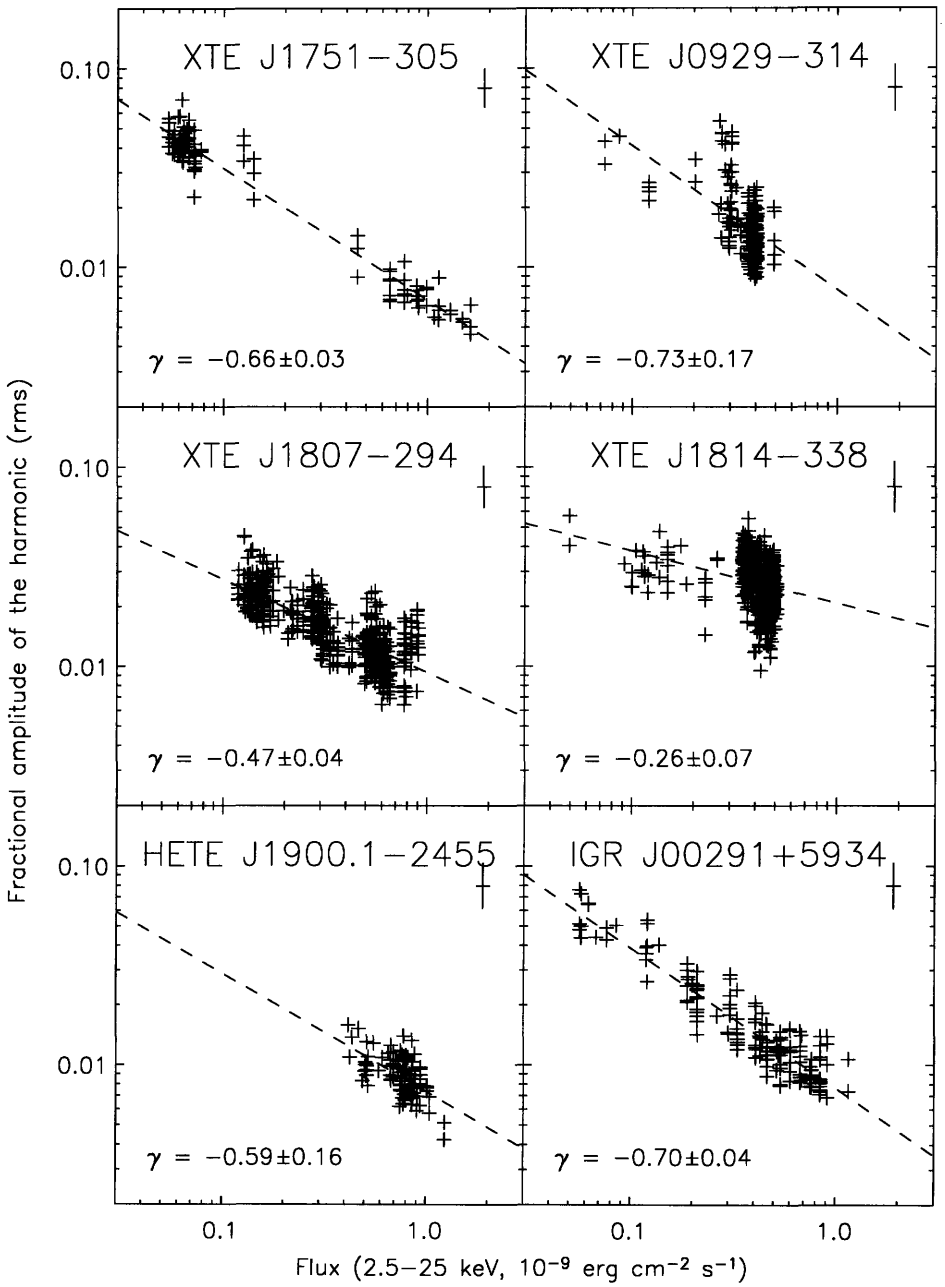

Figure 5-1 The fractional amplitude of the second harmonic vs. flux for six AMPs. The points show measurements of the fractional amplitude of the second harmonic from 512 s intervals of timing data. The dashed lines show the best-fit power laws, $r_2 \propto f_x^{\gamma}$. The error bars in the upper-right corner indicate the mean size of the uncertainties on each point.

Table 5.4. Derived distances, assuming the harmonic amplitude vs. flux relation is a standard candle

	$d_{\rm fit}$ (kpc)	$d_{\rm lit}$ (kpc)	Reference
SAX J1808.4−3658	—	3.5 ± 0.1	Galloway and Cumming (2006)
XTE J1751−305	4.6 ± 0.3	8 ± 5	Wijnands et al. (2005)
XTE J0929−314	4.2 ± 1.3	10 ± 5	Wijnands et al. (2005)
XTE J1807−294	4.6 ± 0.4	8 ± 5	Campana et al. (2005)
XTE J1814−338	2.5 ± 0.6	8 ± 2	Strohmayer et al. (2003)
HETE J1900.1−2455	5.5 ± 1.4	4.2 ± 1.2	Galloway et al. (2006)
IGR J00291+5934	4.4 ± 0.3	$4.2 \, ^{+\,1.2}_{-\,1.1}$	Galloway et al. (2005)

or HETE J1900.1−2455. However, the three sources for which we do have good data over at least a decade of flux are well-fit with power laws, albeit at significantly different slopes than the $\gamma = -0.50 \pm 0.01$ that we determined for SAX J1808.4−3658 in the preceding chapter (see Fig. 4-5).

The χ^2 statistics, while in some cases much greater than the number of degrees of freedom, are not prohibitive. For instance, IGR J00291+5934, one of the sources for which a power law appears to be a good fit, has $\chi^2_\nu = 5.6$. This statistic reflects a mean scatter of 0.26 decades in harmonic amplitude and a mean uncertainty of 0.12 decades. While the scatter is $\sqrt{\chi^2_\nu} = 2.2$ times larger than the mean fractional amplitude uncertainty, it is still small compared to the full decade of fractional amplitude that this source explores.

HETE J1900.1−2455 requires particular comment, since its pulsations turn on and off throughout its exceptionally long outburst (Galloway et al. 2007). The points plotted for this source are drawn solely from the initial stretch of time in which pulsations were detectable. Because these data do not include the observations of the source as it faded into quiescence, the flux range is particularly narrow, resulting in the large slope uncertainty.

These power law relations do not serve as a good standard candle. For SAX J1808.4−3658, the distance is well-determined: 3.5 ± 0.1 kpc (Galloway and Cumming 2006). If we assume that all the other sources have the same relationship (including the same pre-factor), then we can derive rough distances for them. This would be a very useful observational tool, given that the distances for most of these sources are poorly constrained. The derived distances are listed in Table 5.4. Some of these distances are in reasonable agreement for some sources, although this feat is not particularly notable given the large uncertainties. However, the distance derived for XTE J1814−338 using this relation is a factor of 3 too low. Additionally, the very significant deviations of the best-fit power law slope from the $\gamma = -\frac{1}{2}$ of SAX J1808.4−3658 would be a problem for such a distance estimation technique.

5.4 Discussion

Given the complexity of these systems, the simple power law agreement between the harmonic amplitude and the flux is quite surprising. Here we attempt to explain this result in terms of basic accretion theory.

One possible explanation is recession of the accretion disk as the accretion rate drops, revealing the star's previously occulted second hot spot. For rapidly rotating pulsars, partial occultation of the star by the accretion disk will be common. Assuming a mass of 1.4 $M_\odot$,

the co-rotation radius of SAX J1808.4−3658 is $r_{co} = 31$ km $\approx 3R$, where R is the NS radius. Following the standard pulsar accretion model (e.g., Ghosh and Lamb 1979), the inner edge of the accretion disk will be at roughly the Alfvén radius: $r_0 \approx r_A \equiv (2GM)^{-1/7}\dot{M}^{-2/7}\mu^{4/7}$. This truncation radius must be at $r_0 < r_{co}$ for infalling matter to reach the NS surface. There are clear problems with the application of this model, which was developed for higher-field pulsars with $r_0 \gg R$: the width of the transition region in which the magnetic field becomes dominant is on the same order as its distance to the star, muddling the definition of a truncation radius. Nevertheless, this simple model is still qualitatively instructive.

Since $r_{co} \approx 3R$, neutron stars in systems with inclinations $i \gtrsim 70°$ will always be partially occulted during outburst. During the outbursts of SAX J1808.4−3658, the peak fluxes at which the pulses are most sinusoidal are roughly a factor of 10 greater than the low fluxes at which the harmonics are more prevalent (cf. Figure 4-5). As a result, the Alfvén radius will increase by a factor of $r_{A,tail}/r_{A,peak} \approx 10^{2/7} = 2$ as the source dims. Because the minimal Alfvén radius is $3R$, the radius during the peak of the outbursts must be $r_{A,peak} \lesssim \frac{3}{2}R$. At this separation, the star will be partially occulted if $i \gtrsim 45°$. Thus the degree of occultation will depend on $\dot{M}$ for a wide range of inclinations. For $45° \lesssim i \lesssim 70°$, the disk will partially occult the NS above some critical $\dot{M}$. For $i \gtrsim 70°$, the NS will always be partially occulted, with the degree of occultation increasing as $\dot{M}$ increases. Pulse profile modeling by Poutanen and Gierliński (2003) suggests that the system is at an inclination of $i > 65°$.

The observations that show clearly double-peaked pulse profiles happen exclusively in the final, flaring tale stage of the outbursts, typically during the fading portion of a flare. In view of this model, one could imagine that the accretion disk is most recessed as the flares fade. One difficulty with this model is that the increased r_2 observed at low luminosities is not solely due to the appearance of doubly peaked pulse profiles; many profiles in this regime show single pulses, but with substantially greater asymmetry than typically seen at higher luminosities.

Another possible cause is the expansion of the hot spots during high accretion due to diffusive effects. Simulations of accretion flows by Romanova et al. (2004) demonstrate that as the fluence increases, the cross-sections of the accretion funnels grow. Modeling by Muno et al. (2002) establishes that the harmonic content of the pulsations decreases as the size of the hot spot increases.

Chapter 6

Comparing the accretion-powered pulsations and burst oscillations of SAX J1808.4−3658

This chapter is adapted from the papers "The flux dependence of soft lags in SAX J1808.4−3658" and "A lighthouse with two lights: comparing the accretion-powered pulsations and burst oscillations of the accreting millisecond pulsar SAX J1808.4−3658" by Jacob M. Hartman, Anna Watts, and Deepto Chakrabarty. Both papers are in preparation for submission to the *Astrophysical Journal*.

6.1 Introduction

The accreting millisecond pulsars SAX J1808.4−3658 (Wijnands and van der Klis 1998) and XTE J1814−338 (Markwardt and Swank 2003) are unique objects when it comes to studying the behavior of matter at the surface of neutron stars in LMXBs. They are the only two systems that show both persistent accretion-powered pulsations and burst oscillations (Chakrabarty et al. 2003; Strohmayer et al. 2003). Although there are six other AMPs that show persistent pulsations, only one has bursts, and they did not exhibit burst oscillations (Galloway et al. 2007). None of the other ∼10 neutron stars with burst oscillations have detectable persistent pulsations.

Both types of pulsations are generated by asymmetries at the neutron star surface that are modulated by the stellar rotation. The persistent pulsations arise as accreting material is channeled by the star's magnetic field, leading to hot spots at the magnetic poles. Burst oscillations are seen during X-ray bursts, which occur when accreted material burns unstably in a thermonuclear flash (see Strohmayer and Bildsten 2006 for a recent review). What causes the oscillations, however, is not yet fully understood. Possible mechanisms include asymmetric deposition of fuel that results in patchy burning, magnetic field effects, oscillations in the surface layers (e.g., Piro and Bildsten 2005) or hurricane-like vortices (Spitkovsky et al. 2002). However the burst and accretion pulsation processes are most likely inter-related, since fuel deposition and field geometry will affect the subsequent nuclear burning process. Comparing the two types of pulsation, and contrasting their properties with those of burst oscillations from the non-pulsing LMXBs, should therefore offer insight into the complex processes at work at the surface of an accreting neutron star.

Many phenomenological differences between these two modes of pulsations have already been noted in the literature. Additionally, the burst oscillations observed in the two millisecond pulsars in which this phenomenon has been observed, SAX J1808.4−3658 (Chakrabarty et al. 2003) and XTE J1814−338 (Strohmayer et al. 2003), differ in some notable ways from the burst oscillation in non-pulsing LMXBs.

The persistent pulsations of all the accretion-powered pulsars exhibit "soft lags," in which the pulse arrival time at soft energies lags behind the arrival of pulses at harder energies (Cui et al. 1998; Galloway et al. 2002; Kirsch et al. 2004; Galloway et al. 2005; Gierliński and Poutanen 2005; Watts and Strohmayer 2006; Galloway et al. 2007). The fractional amplitudes of these persistent pulsations typically peak at at 5–7 keV for the fundamental, and are relatively flat for the harmonic.

The burst oscillations of the non-pulsing LMXBs have amplitudes that fall with energy, and there is tentative evidence for hard lags (Muno et al. 2003). XTE J1814−338's burst oscillations are different. They have amplitudes that fall with energy, like the persistent pulsations, but without detectable lags (Watts and Strohmayer 2006). This chapter will investigate the energy-resolved timing of the burst oscillations in SAX J1808.4−3658.

Comparing the burst oscillation timing of SAX J1808.4−3658 and XTE J1814−338 to those of the non-pulsing LMXBs, several differences are apparent. Firstly, the amplitudes of the XTE J1814−338 burst oscillations, at ≈11% rms, are larger than those of most LMXBs. They also have greater harmonic content than the bursts from non-pulsing LMXB, with a second harmonic fractional amplitude of $\approx 2\%$ (Muno et al. 2002; Strohmayer et al. 2003). There is no evidence for amplitude or frequency variation in the XTE J1814−338 burst oscillations except during the one burst with possible photospheric radius expansin (PRE; Watts et al. 2005). LMXB burst oscillations, by contrast, show amplitude variability and frequency shifts of up to 5 Hz (Muno et al. 2002; see, for instance, burst [b] in Fig. 1-4).

As discussed in §1.5, SAX J1808.4−3658 was the first known AMP to exhibit burst oscillations, establishing the link between the pulsing AMPs and the non-pulsing burst oscillation sources. Three bursts were detected during the 1996 outburst, two during the peak and one later when the accretion rate had declined (in 't Zand et al. 1998, 2001). No bursts were seen during the 1998 or 2000 outbursts, observations of the latter being particularly sparse due to solar constraints (Wijnands et al. 2001). The 2002 and 2005 outbursts, fortunately, received extensive *RXTE* coverage. Four bursts were detected in 2002 (Chakrabarty et al. 2003) and three in 2005 (Galloway et al. 2006). The bursts are bright, with a fast rise, and have PRE and burst oscillations at the spin frequency. A study of the recurrence times and fluences suggests that they are triggered by unstable ignition of a pure helium layer underneath a stable burning layer of hydrogen (Galloway and Cumming 2006). Assuming that the resulting PRE reaches the Eddington limit for helium gives a distance of 3.4–3.6 kpc. This accords well with a lower limit of 3.4 kpc derived by equating the time averaged X-ray flux with the expected mass transfer rate as the binary orbit decays due to gravitational radiation. Using this estimate of source distance, measurements of the optical/IR counterpart suggest an inclination $i \lesssim 30^\circ$ (Giles et al. 1999; Wang et al. 2001).

Some initial analysis of the SAX J1808.4−3658 oscillations has already appeared in the literature. The accretion-powered pulsations were discussed for the 1998 outburst by Wijnands and van der Klis (1998) and Cui et al. (1998), and for the 2002 outburst by Chakrabarty et al. (2003) and Burderi et al. (2006). The 1998 results were used by Gierliński et al. (2002) and Poutanen and Gierliński (2003) to model the accretion geometry, although it should be noted that these models used the old distance estimate of 2.5 kpc (in 't Zand et al. 1998). A detailed analysis of the evolution of the accretion pulse profile across the

four outbursts observed by *RXTE* has been discussed in Chapter 4, as part of a study of the spin history. Burst oscillation properties for the 2002 outburst were first discussed in Chakrabarty et al. (2003). Additional analysis of the 2002 bursts, focusing on the burst rise and precursor activity, has since been carried out by Bhattacharyya et al. (2006) and Bhattacharyya (2007).

The persistent pulsations have amplitudes of 3–5% at the fundamental frequency, with the amplitude of the second harmonic being about one tenth of that value (see Ch. 4). The fundamental fractional amplitude of the 2002 burst oscillations is similar, apart from during the peak of the bursts, when PRE renders them undetectable. The second harmonic amplitude for the burst oscillations is lower than that of the persistent pulsations, with a 2 σ upper limit of 0.5%. These values are similar to those of the non-pulsing LMXB burst oscillations. The 2002 burst oscillations also show interesting variations in frequency and phase. During the rise the oscillations show a rapid and large ($\approx$5 Hz) increase in frequency, overshooting the spin frequency by $\approx$2 Hz. Bhattacharyya et al. (2006) have suggested that this frequency rise may be non-monotonic. The burst oscillations in the tail are offset in phase from the persistent pulsations, leading them by an average of 11% (Chakrabarty et al. 2003). The only energy dependent analysis reported to date is for the accretion-powered pulsations in the 1998 outburst. Like the other accreting millisecond pulsars the amplitude drops slightly with energy, and there are soft lags (Cui et al. 1998). Energy dependence during the other outbursts and for the burst oscillations has yet to be addressed.

In this chapter we present a comprehensive analysis of the properties of the pulsations for all the outbursts and bursts of SAX J1808.4−3658 detected by *RXTE*. For the persistent pulsations, the timing of which was discussed in Chapter 4, we study the energy dependence, including variation during the outbursts. For the burst oscillations, we analyze frequency and amplitude variation, energy dependence, and the relation to the accretion-powered pulsations. We then compare the pulsations of SAX J1808.4−3658 to those of the other accreting millisecond pulsars and burst oscillation sources. We conclude by discussing the implications for theoretical models of the accretion and burst process.

6.2 Data analysis

For our study of the energy-dependence of the persistent pulsations, we analyzed the same SAX J1808.4−3658 observations as in Chapter 4 (see Table 4.1). The `E_125US_64M_0_1S` PCA mode, which was used for nearly all the observations, provides 64 energy channels divided evenly over the detector response and 122 μs time resolution, a good combination for discerning the energy dependence of the persistent pulsations. The first thermonuclear burst observed from SAX J1808.4−3658 was one of the few exceptions, having been captured with the high-resolution `GoodXenon` mode. Its time resolution was binned down to 122 μs, and equivalent energy channel boundaries were applied.

Care must be taken due to the changing energy response of the channels as the voltages of the PCA anodes were stepped down as the spacecraft aged (Jahoda et al. 2006; see also §2.1). As a result, the energy response generally shifted to include harder photons. Throughout our analysis, we consistently used six channel ranges: 1–3, 4–6, 7–10, 11–14, 15–19, and 20–25. These channel numbers reflect the `64M_0` energy binning, with channel 0 being the lowest. Table 6.1 lists the associated energy ranges for each outburst. For short integrations, we sometimes combined them into three bands (labeled throughout this paper as the soft, medium, and hard bands) for better statistics. These energy channel

Table 6.1. Channel and energy ranges

	Channels		Energies (keV)		
	64M_0	GoodXenon	1998	2002	2005
Soft	1–3	5–8	1.94–3.35	1.95–3.67	1.95–3.67
	4–6	9–11	3.35–4.41	3.67–4.96	3.67–4.96
Medium	7–10	12–15	4.41–5.82	4.96–6.68	4.96–6.68
	11–14	16–23	5.82–8.68	6.68–10.16	6.68–10.16
Hard	15–19	24–33	8.68–12.27	10.16–14.54	10.16–14.54
	20–25	34–46	12.27–17.00	14.54–20.30	14.54–20.30

Note. — For the conversion between the 64M_0 channels and energies, consult Appendix 1 of the *RXTE* Technical Appendix,[1] and the online energy-channel conversion table.[2]

boundaries are identical to those used by Watts and Strohmayer (2006) in their analysis of XTE J1814−338, so our results for SAX J1808.4−3658 will be directly comparable.

After dividing the photons into these bands, we applied our standard *RXTE* processing pipeline, discussed in more detail in previous chapters: we shifted the photon arrival times to the Earth's geocenter, applied the *RXTE* fine clock correction, and filtered out data during Earth occultations and intervals of unstable pointing. For the pulsation analysis, we also removed any data within 5 min of thermonuclear bursts. To measures the phases and fractional amplitudes of the persistent pulsations, we folded 512 s intervals using the TEMPO timing solutions for the constant-frequency models derived in Chapter 4. Their values and uncertainties were calculated using the techniques of Chapter 3. In practice, these techniques give identical results to the Z_n^2 statistics used by Watts and Strohmayer (2006) for pulse profiles with minimal contributions from their higher harmonics (certainly the case for all known AMPs), making the results directly comparable.

Our measurement of the energy dependence of the burst timing properties was similar. We applied our TEMPO phase models to fold the burst data and measured the phases and fractional amplitudes. We did not include the data during the burst rises, as the frequency is changing so rapidly that accurate phase and amplitude measurements are not possible with the current data. We also excluded the timing data during PRE, because oscillations were not present. The exception is the single dim burst from the 2005 outburst, which did not exhibit PRE. For this burst, we used the data from the peak onward. We defined the end of the burst data to be when the number of photons dropped below twice the pre-burst persistent level. This cutoff resulted in integrations of ≈30 s per burst, during which we made a single phase and amplitude measurement for each of our three energy bands.

For shorter-interval timing within the bursts, in particular during the burst rise and peak, we had to supplement the event-mode data with single-bit data. The thermonuclear bursts from SAX J1808.4−3658 are the brightest ever observed by *RXTE*, and the resulting flood of photon events quickly overwhelms the spacecraft telemetry. As a result, the event mode data have numerous gaps during which events were not recorded. The single bit data accommodated these high count rates, but at the expense of any energy resolution. Single-bit data were not available for two of the bursts: during the first burst of 2002, the *RXTE*

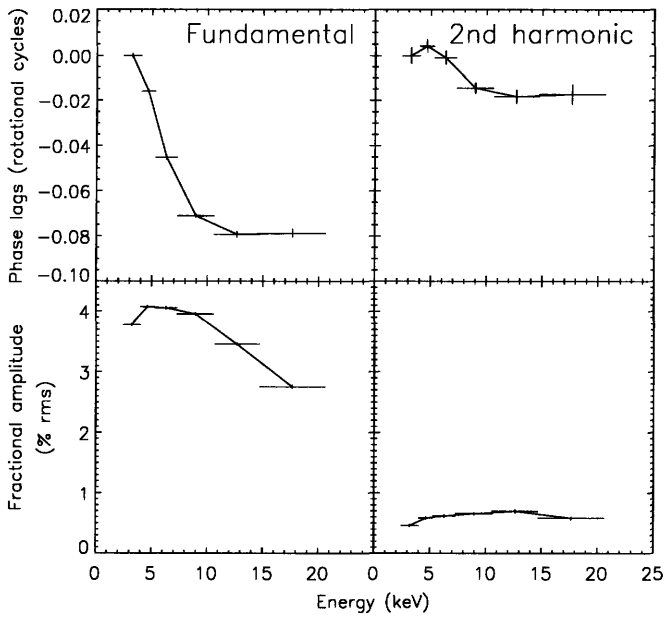

Figure 6-1 Average phase shifts (relative to the 2.0–3.7 keV band) and fractional amplitudes during the slow-decay stage of the 2002 outburst, MJDs 52564.0–52575.0.

Experimental Data System was improperly configured; and the third burst of 2005 occurred at the end of a slew preceding a SAX J1808.4−3658 observation, a mere 35 s before the single-bit data mode activated.[3]

6.3 Results

6.3.1 Persistent pulsations

Figure 6-1 shows the behavior of the phase shifts and fractional amplitudes against energy for both harmonics during 11 days of the 2002 outburst. Immediately apparent is the phase lag of the soft photons in the fundamental. Between 2 keV and 10 keV, harder photons arrive progressively sooner, with a slope of -0.012 ± 0.002 cycles keV^{-1}. Between 10 keV and 20 keV, this trend saturates at a lead of 0.08 cycles (200 μs). This result is in excellent agreement with soft lags previously reported for the 1998 outburst (Cui et al. 1998; Gierliński et al. 2002). Similar trends were seen in XTE J1814−338 (Watts and Strohmayer 2006), XTE J0929−314 (Galloway et al. 2002), XTE J1807−294 (Kirsch et al. 2004), IGR J00291+5934 (Galloway et al. 2005), and XTE J1751−305 (Gierliński and Poutanen 2005). The phase shifts of the second harmonic are not nearly as pronounced. There is not the same monotonic progression: the 3.7–5.0 keV band lags behind the softer 2.0–3.7 keV band. At higher energies, there is a pronounced lead, but it saturates at less than 0.02 rotational cycles ($\lesssim 50$ μs).

The fractional amplitude of the fundamental ranges from 2.8% to 4.1% rms, peaking at around 7 keV. The second harmonic amplitude ranges from 0.4% to 0.7% rms, peaking

[3]But had this burst been a mere 35 s *earlier*, SAX J1808.4−3658 would have been too far off-axis, and it would have been missed altogether. I argue that the glass is half-full.

at roughly 12 keV. The locations of these fractional amplitude peaks are quite similar to what is seen in XTE J1814−338 (Watts and Strohmayer 2006), although the fractional amplitudes in XTE J1814−338 were substantially greater.

Figure 6-2 illustrates how the energy-dependent phases and fractional amplitudes change over the course of the four observed outbursts. The soft lags in the fundamental are apparent throughout all four outbursts; for the second harmonic, the soft lags are apparent everywhere except during the burst peaks. However, the degree of these lags is clearly not constant. They are much more pronounced during the tails of the outbursts, when the flux is low. Watts and Strohmayer (2006) noted similar behavior in XTE J1814−338, although that source did not explore as wide a range of fluxes, so the relationship was less clear.

Interestingly, the observed fractional amplitudes tell a somewhat less consistent story. During 1998, the fractional amplitudes of the fundamental are greater than during the other outbursts, a discrepancy that we noted in Chapter 4. There is a considerable spread between the three energy bands, ranging from $\approx$7% rms for the soft band down to $\approx$4% for the hard band. However, for the 2002 and 2005 outbursts, the fractional amplitudes of the soft and medium bands are quite similar, with the hard band only slightly lower. But as the outbursts decay, the fractional amplitude vs. energy curves become steeper, resulting in a wider spread toward the end of the outbursts. Given the strong similarities between the light curves of the 1998 and 2002 outbursts, it is curious that their timing behaviors are markedly different.

To further understand the relationship between the soft lags and the flux, we measured the slope of the lags vs. 2–10 keV energies throughout the outbursts for both harmonics. Figure 6-3 plots the resulting slopes as a function of flux. For the fundamental, there is a clear increase in the magnitude of the soft lags as the 2–25 keV flux falls from its peak at $\sim 3 \times 10^{-9}$ erg cm^{-2} s^{-1} down to roughly 0.7×10^{-9} erg cm^{-2} s^{-1}. This range of fluxes covers the peak and slow decay stages of the outburst. The point at which this relation stops is significant: once the flux falls to 0.7×10^{-9} erg cm^{-2} s^{-1}, it drops off rapidly (see the light curves in Fig. 6-2) and the outbursts enter the "flaring tail" stage. The fundamental slopes suggest that in this part of the outburst the relation reverses, and the magnitude of the soft lags then *decrease* with flux below 0.7×10^{-9} erg cm^{-2} s^{-1}.

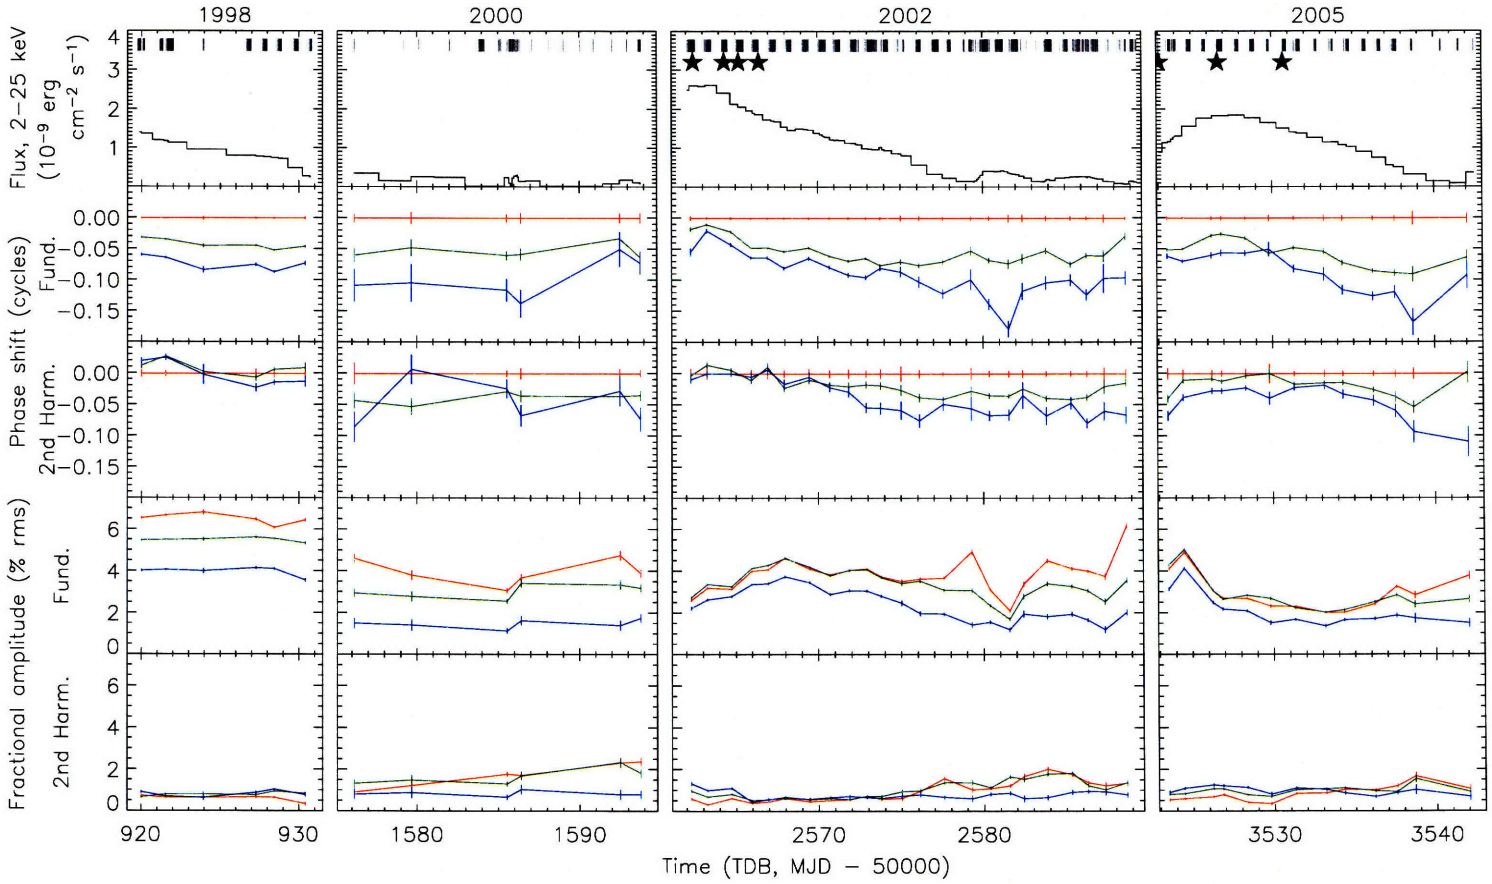

Figure 6-2 Phase shifts and fractional amplitudes as a function of energy for all the SAX J1808.4−3658 outbursts. The top panels show the background-subtracted light curves for each outburst. The strips along the tops of the graphs indicate the times of observations; stars indicate the times of the seven observed thermonuclear bursts. The second and third panels show the phase shifts of the medium energy band (green; roughly 5–10 keV) and hard energy band (blue; ≈10–20 keV) pulsations relative to the soft band (red; ≈2–5 keV). The negative phase lags indicate that the higher energies often lead the soft band. The bottom two panels show the background-subtracted fractional amplitudes for these three energy bands. Each point in the phase and fractional amplitude plots gives a day-long average. (They are not evenly spaced at 1 d intervals due to the uneven spacing of the *RXTE* observations.) Data within 5 min of thermonuclear bursts were excluded from these averages.

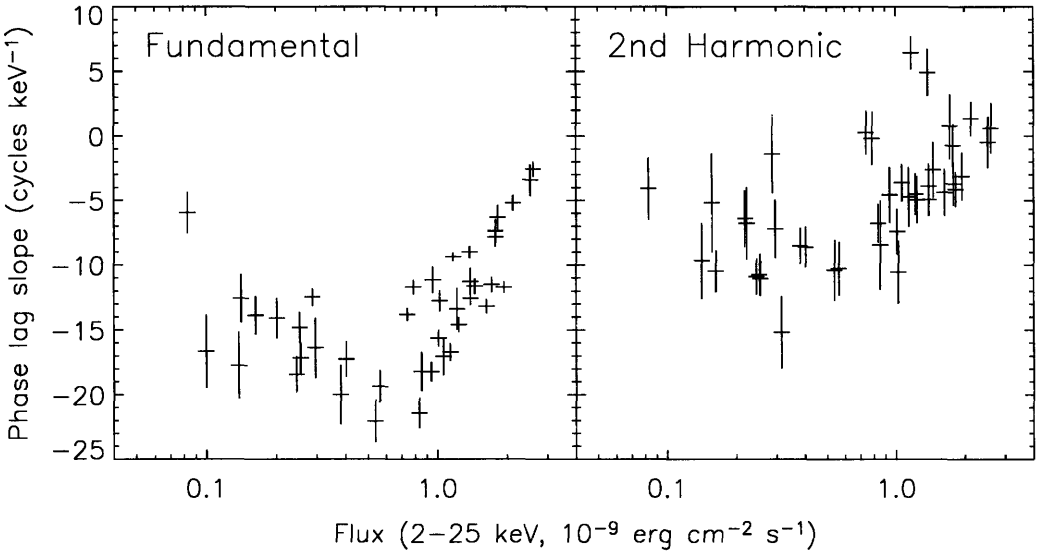

Figure 6-3 Correlation between the slope of the phase lags and flux. Each slope is determined by a linear fit to the phase lag vs. energy measurements over 2–10 keV, with each point representing a 1 d average. Data from all four of the observed outbursts are present.

The trend for the second harmonic is somewhat less clear. Measuring a slope through the 2.5–10 keV fractional amplitude points simply does not work as well since energy dependence is less linear, as can be seen in Fig. 6-1, so there is more scatter. Even still, the data suggest a similar relationship at high fluxes.

6.3.2 Energy dependence of the burst oscillations

There have been seven thermonuclear bursts detected from SAX J1808.4−3658; Table 6.2 lists some of their properties. Excepting burst 5, these are the six brightest thermonuclear bursts to be captured by *RXTE*.

The energy dependence of the burst oscillations observed during the tails of these bursts is shown in Figure 6-4. The energy dependence is far weaker than in the persistent pulsations. There are no evident soft lags. The hypothesis that the phase lags of the fundamental have no energy dependence can only be weakly rejected (2.6 σ confidence); the hypothesis that the fundamental's fractional amplitudes are independent of energy cannot be rejected at all (0.5 σ).

The second harmonic was too weak to detect in the individual energy bands. The 95% confidence upper limits are consistent with the fractional amplitudes in the bursts being the same as during the persistent pulsations ($\approx$1% rms).

6.3.3 Burst oscillation timing

Figure 6-5 shows the light curves dynamic power spectra for all the bursts observed from SAX J1808.4−3658. There are a number of features worthy of comment.

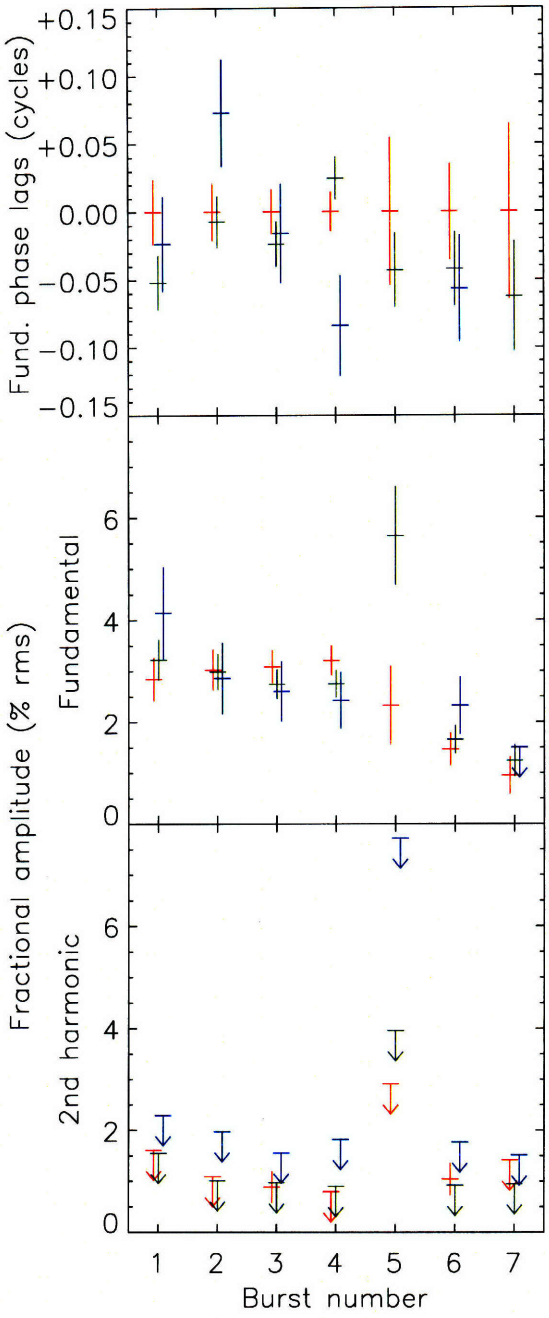

Figure 6-4 Timing properties of the thermonuclear bursts for the soft (red; ≈2–5 keV), medium (green; ≈5–10 keV), and hard (blue; ≈10–20 keV) energy bands. The phase shifts are relative to the soft band. Missing points indicate that the fundamental was not detected with 95% confidence. The fractional amplitudes are background-subtracted, but we have not attempted to correct for the presence of pulsed photons due to the underlying persistent pulsations. The upper limits indicate that pulsations were not detected with 95% confidence and give the 95% confidence upper limits on the fractional amplitudes. The upper limit for the fundamental's fractional amplitude in the hard band of burst 5 is above the top of the plot, at 10.9% rms.

Table 6.2. Thermonuclear bursts observed from SAX J1808.4−3658

Burst number	Time of start (UTC)	Peak flux[a] (10^{-9} erg cm^{-2} s^{-1})	PRE[a]
1	2002 Oct 15 09:55:33.500	162 ± 3	Y
2	2002 Oct 17 07:19:20.750	177 ± 3	Y
3	2002 Oct 18 04:25:17.000	179 ± 3	Y
4	2005 Oct 19 10:14:29.500	177 ± 3	Y
5	2005 Jun 02 00:41:23.000	14.8 ± 0.5	N
6	2005 Jun 05 15:19:00.125	173 ± 3	Y
7	2005 Jun 09 16:08:40.000	≈160[b]	Y[b]

[a]Peak fluxes and PRE determinations are from Galloway et al. (2006).

[b]These data are unavailable for burst 7 because its peak occurred during a slew. We estimate that its peak flux was ≈160 × 10^{-9} erg cm^{-2} s^{-1}, commensurate with the other similar bursts, and that it underwent photospheric radius expansion.

Burst rises. A unique feature of the bursts from SAX J1808.4−3658 is the rapid frequency increase during their rise (Chakrabarty et al. 2003). While it is typical for burst oscillations to rise in frequency by as much as 6 Hz over ∼10 s time, the oscillations during the rises of bursts 3 and 4 increase in frequency by roughly 6 Hz in approximately half a second, the time scale of the burst's notably rapid rise in flux. Equally puzzling, the oscillation frequency overshoots the spin frequency in these two bursts. Burst 1 appears to have similar behavior, although the short intervals of available data greatly limits the frequency resolution of the power spectra.

Burst 6, the sole PRE burst from 2005 with complete data, also shows a rapid increase in the oscillation frequency during the burst rise. However, the oscillations do not overshoot the spin frequency. After a rise of ≈3 Hz in 0.5 s, the oscillations reach 400.0 Hz just as the rise in the flux slows. Over the next three seconds, the frequency rises more slowly, reaching 400.5 Hz before the oscillations disappear during PRE. This slower evolution much more closely resembles the frequency changes sometimes seen in other thermonuclear burst sources.

Burst peaks. As is almost always the case in burst oscillation sources, there are no oscillations detected during photospheric radius expansion. This is not the case, however, for burst 5, the non-PRE burst. Weak oscillations are present throughout the burst at the spin frequency.

Burst tails. The character of the burst oscillations that reemerge following PRE is quite different from the oscillations during the rise. The frequency of the oscillations during the tails of bursts 2, 3, and 4 is quite stable (uncommonly stable for the tails of bursts; see, for instance, Muno et al. 2002), and it is equal to the spin frequency. Bursts 6 and 7 also exhibit oscillations during their tails, but they are of lower amplitude and disappear after less than 10 s.

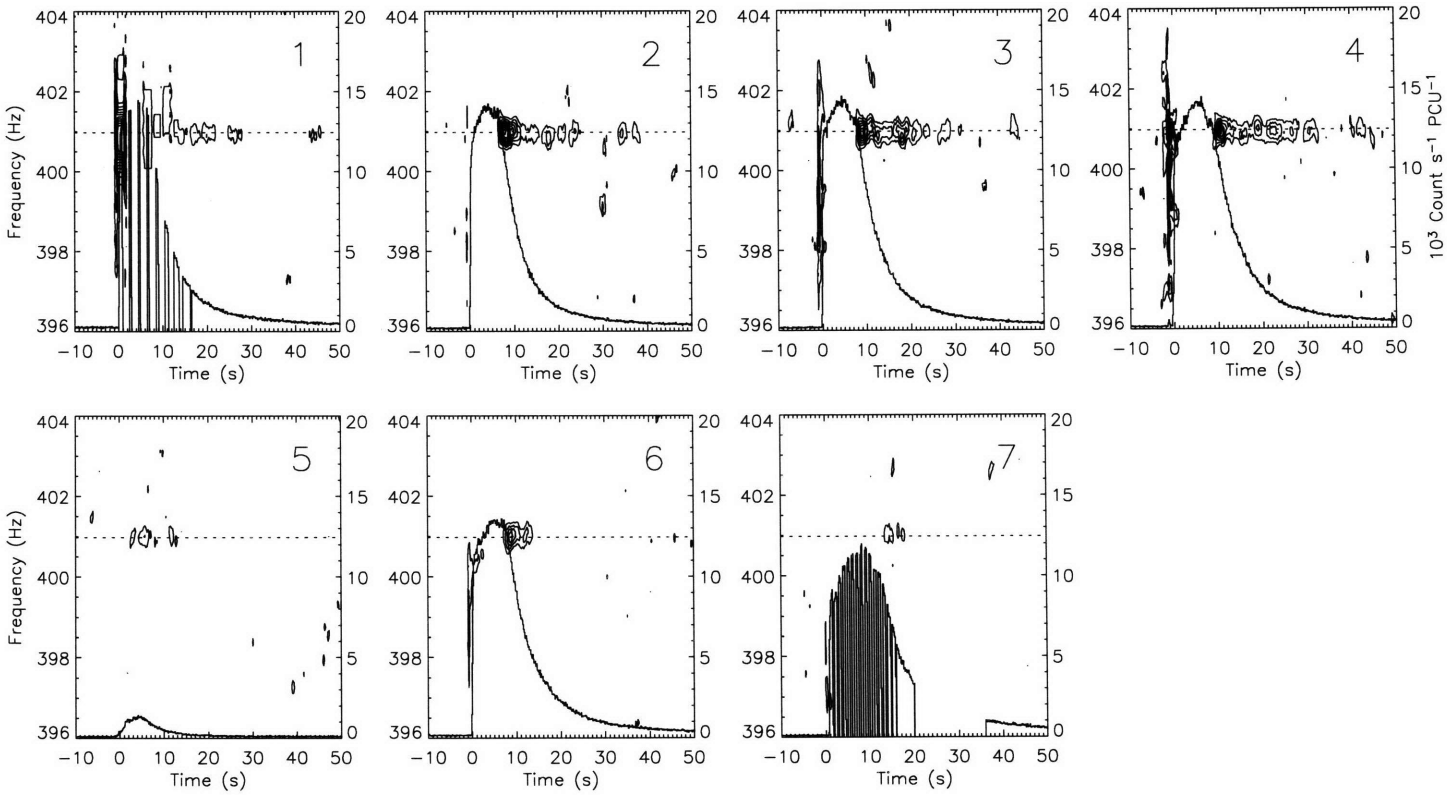

Figure 6-5 The light curves and dynamic power spectra of all seven thermonuclear bursts observed from SAX J1808.4−3658. The x-axes indicate the time since the start of the bursts, which are listed in Table 6.2. The spectra were created using 2 s long Fourier transforms spaced at 0.25 s intervals. The frequencies were oversampled by a factor of 4. Contour levels are at unit-norm powers of 6, 12, 18, etc. The light curves of bursts 1 and 7 show gaps in the data because the *RXTE* was not configured for observing high count rates at the times of these bursts.

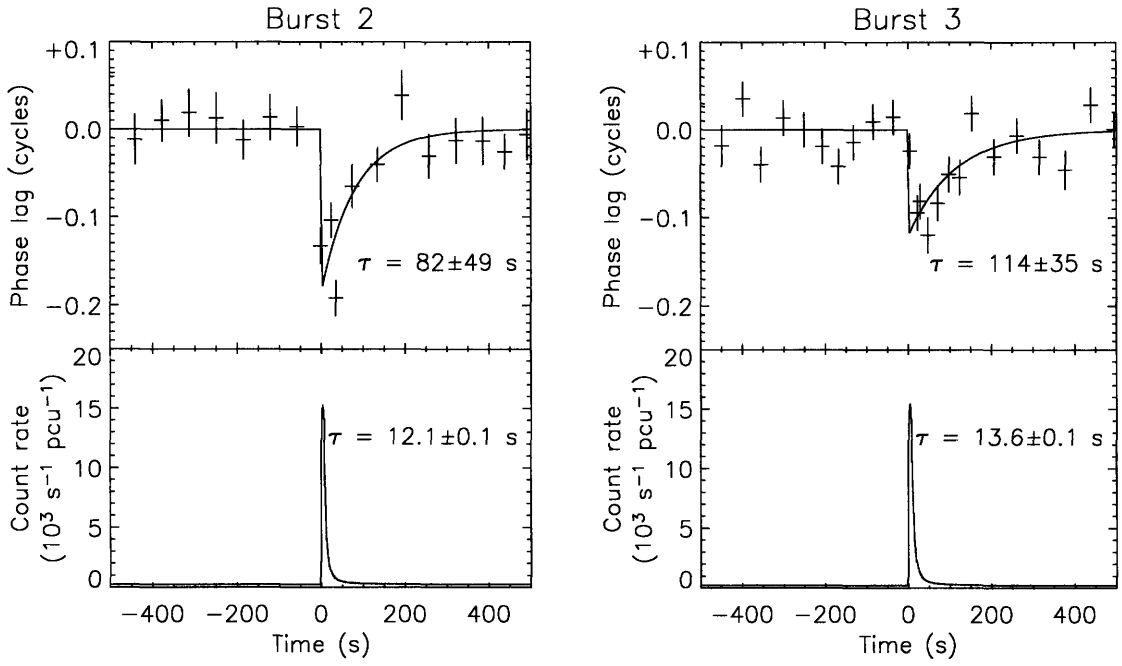

Figure 6-6 Phase connection through bursts 2 and 3. The points here show the phase residuals of the fundamental relative to the best-fit constant-frequency models for the persistent pulsations. The fits are exponentials starting at 7 s after the initial burst rise, which corresponds to the time at which the burst oscillations reappear following PRE.

The long trains of oscillations during the tails of bursts 2, 3, and 4 have the curious behavior of subside 25–35 s into the burst, then reappearing roughly 10 s later. In bursts 2 and 3, there is marginal evidence of briefly emergent 399.5 Hz oscillations in the middle of these intervals.

The phases of the oscillations during the bursts consistently lead the phases of the persistent pulsations by 0.10–0.15 cycles, as first noted by Chakrabarty et al. (2003). This phase difference is seen in all the bursts, including burst 5, the non-PRE burst.

The transition between this offset burst oscillation phase and the persistent pulastion phase was best captured during bursts 2 and 3, shown in Figure 6-6. (The other five bursts suffered from gaps in the data or lower-amplitude burst oscillations that resulted in poorer phase tracking.) The exponential fits to these residuals exhibit somewhat longer timescales than the decay timescales for the burst light curves. Following Galloway et al. (2006), we adopt a model-independent burst time scale of $\tau = E_{\rm b}/F_{\rm pk}$, where $E_{\rm b}$ is the integrated fluence during the burst and $F_{\rm pk}$ is the peak flux. For all the bursts except for burst 5, this time scale was 12–15 s. In contrast, the equivalently defined time scale for the exponential fits of the phase residuals was 82 ± 49 s for burst 2 and 114 ± 35 s for burst 3.

6.4 Discussion

6.4.1 The flux dependence of the soft lags

Our observation that the degree of phase lag depends strongly on the flux provides a powerful constraint on the theories that have been advanced to explain the soft lags in SAX J1808.4−3658 and other systems. In particular, the fact that the flux dependence changes sign at precisely the same flux at which the decay rate of the outburst light curve increases markedly suggests that these lags are directly tied to the hot spot and accretion column rather than static effects such as rotational Doppler shifts or a persistent halo of scattering material.

The flux dependence is problematic for models that explain the lags purely as Doppler boosts. Ford (2000) proposed that the rotational Doppler shifting of the hot spot would cause blue shifting while the hot spot is moving toward the viewer (at earlier phases) and red shifting while it is receding (at later phases). However, these shifts are totally independent of the flux. Alternatively, (Sazonov and Sunyaev 2001) suggested Doppler boosting by the Keplerian flow of the accretion disk, in which the light scattered from the disk would lead if the NS and disk are co-rotating and lag if they are counter-rotating. There are two issues with the application of this model to SAX J1808.4−3658. First, the disk will most likely downscatter the photons from the hot spot, in which case the observed soft lags imply a counter-rotating disk. But accretion from a counter-rotating disk imparts its angular momentum to the NS with near-unity efficiency (e.g., Ghosh and Lamb 1979), which would produce a substantially larger spin down than the upper limits derived in Chapter 4. Furthermore, as the accretion rate decreases, the Alfvén radius will move outward and the Keplerian velocity of the inner edge will decrease, decreasing the magnitude of this boost. However, we observed the opposite: the degree of the lags increases as the accretion rate drops.

In their paper reporting the discovery of the soft lags, Cui et al. (1998) suggested that they could result from the Compton downscattering of pulsed hard X-ray emission by a cooler electron cloud. In light of the flux dependence, the radius of this cloud would have to vary with the flux. Moreover, it would have to vary in a different way during the main body of the outburst and during the flaring tail to account for the patterns that we observed. This scenario seems unlikely.

Finally, Cui et al. (1998) also proposed that temperature variation across the hot spot, with the cooler edges producing the soft lags and resulting in less harmonic content at low energies. Gierliński et al. (2002) and Poutanen and Gierliński (2003) argue for a two-component model for the emission from the hot spot: a soft blackbody component emitting at a narrow range of angles close to the normal and a Comptonized component emitting harder photons at a broader range of angles. These approaches are the most promising, given our results. From Chapter 4, we observed that the hot spot is most likely changing its shape and its location on the star as the flux varies, so it would not be surprising, for instance, if the temperature of the hot spot is more uniform at high accretion rates (following the logic of Cui et al. 1998) or the opening angle of the hard, broadly emitting Compton component becomes smaller at high rates (following Poutanen and Gierliński 2003). These effects could account for the observations, although clearly more modeling is needed to put our speculations on firmer ground.

6.4.2 Reconnection of the burst oscillation phase

The phase offset of the burst oscillations raises a number of interesting questions. First off, what determines the time scale for the burst oscillation phase to reconnect with the persistent pulsation phase? It is clearly not the flux decay, which happens nearly an order of magnitude more quickly. This finding rules out the naïve model that the burst photons have one constant phase and fractional amplitude and the persistent pulsations have another, which reemerges as the burst flux subsides. Instead, it appears likely that the phase of the persistent pulsations is temporarily altered by the burst. Consider the observed phase ϕ, which is the combination of the burst phase $\phi_{\rm b}$ and the persistent phase, which we assume to be constant at zero:

$$\phi = \tan^{-1}\left(\frac{n_{\rm b}a_{\rm b}\sin\phi_{\rm b}}{n_{\rm b}a_{\rm b}+n_{\rm p}a_{\rm p}\cos\phi_{\rm b}}\right) \approx \frac{n_{\rm b}a_{\rm b}\phi_{\rm b}}{n_{\rm b}a_{\rm b}+n_{\rm p}a_{\rm p}}, \tag{6.1}$$

where n and a give the count rate and fractional amplitudes of the burst and persistent photons. If the persistent phase is indeed constant, $a_{\rm b}$ would have to become large to reproduce the observed ϕ; when $n_{\rm b}/n_{\rm p} = 10\%$, roughly 150 s into the burst, nearly all the burst photons would have to be pulsed. Alternatively, the phase difference between the burst oscillations and the persistent pulsation would have to *increase* in order to offset the rapidly fading burst flux. Both these scenarios seem unlikely, implying that our assumption of a constant persistent pulsation phase is wrong: the thermonuclear burst most likely temporarily changes the phase of the persistent pulsations.

It is feasible that the bursts from SAX J1808.4−3658 temporarily shift the accretion column and the hot spot at its base, resulting in the observed shifts of ≈45° prograde of the persistent hot spot. These bursts are the brightest yet observed by the *RXTE*. The photon pressure at the NS star surface for a peak flux of 1.8×10^{-7} erg cm^{-2} s^{-1}and a distance of 3.5 kpc is 6×10^{13} erg cm^{-3} and falls off as r^{-2}; for a 1×10^{-8} G dipolar field, the magnetic energy density at the surface is 4×10^{14} erg cm^{-3} and falls off as r^{-6}. These energy densities are equal at 16 km, half the Alfvén radius of this system of ≈30 km, roughly where the accretion columns begin. Given these energetics, it is easily imaginable that the bursts could move the accretion column by the observed 45° shift. In contrast, most of the bursts from XTE J1814−338 are far weaker, with peak fluxes of $\lesssim 1.2 \times 10^{-8}$ erg cm^{-2} s^{-1}, and do not show any difference of phase from the persistent pulsations (Watts et al. 2005). For $F_{\rm pk} \lesssim 1.6 \times 10^{-8}$ erg cm^{-2} s^{-1}, the photon pressure and magnetic pressure will be equal at a radius greater than the Alfvén radius, so we would expect less disruption of the accretion column. The one burst from XTE J1814−338 that does exhibit a departure in phase from the persistent pulsations is the brightest, with $F_{\rm pk} = 2.1 \times 10^{-8}$ erg cm^{-2} s^{-1}. If this picture is correct, then the observed ≈2 min time scale for the return to the original phase is particularly interesting, as it directly probes the relaxation time of the magnetic field.

Chapter 7

Conclusions

The discoveries of the *Rossi X-ray Timing Explorer* over the past decade have greatly enhanced our understanding of low-mass X-ray binaries. The population of known millisecond pulsars in LMXBs has grown from zero to ≈20, a sufficiently large sample that we can begin to consider them statistically (e.g., Muno et al. 2004; Chakrabarty 2005). In this thesis, we have explored this wealth of data in some novel ways — for instance, completing the first comparison of the spin frequencies and pulse shapes from multiple outbursts of a single source, or the pulse shape variability of multiple sources. These studies have uncovered many previously unknown aspects of the timing of these sources, raising new questions for the modeling of these systems.

The timing techniques long-applied to radio pulsars are only partially successful in their application to AMPs. Unlike spin-powered pulsars, the flux of AMPs will change by 2–3 orders of magnitude over the course of an outburst, reflecting the changing accretion rate. This changing $\dot{M}$ will change the accretion geometry and as a result the pulse profile, making the typical assumption of a fixed "template" profile of limited use. We developed a method of analysis that mitigates this problem by weighting the more stable harmonics of the pulse profile and by estimating the uncertainties of pulse phase models with Monte Carlo simulations.

We applied this method of analysis to four outbursts of SAX J1808.4−3658, which span seven years of *RXTE* observations. We observed substantial pulse shape variability during each outburst, necessitating such treatment. Much of the change in pulse shape is systematic: we observed remarkably similar pulse profiles from the same stages of different outbursts, suggesting systematic changes in the accretion column as the outbursts progress. In particular, we observed longitudinal drifts of the X-ray hot spots during the tails of the outbursts and an anti-correlation between the harmonic content of the pulses and the source flux. Attempts to explain these characteristic pulse profile changes will hopefully drive further modeling of the accretion hot spots, building on the work of Gierliński et al. (2002) and Poutanen and Gierliński (2003).

Additionally, we observed the spin frequency and orbital period of SAX J1808.4−3658 change from one outburst to the next. Combining data from all the outbursts, we showed with high (6σ) significance that the pulsar was undergoing long-term spin down at a rate $\dot{\nu} = (-5.6 \pm 2.0) \times 10^{-16}$ Hz s^{-1}, with most of the spin evolution occurring during X-ray quiescence. This spin-down establishes an upper limit of $B < 1.5 \times 10^8$ G for the pulsar's surface dipole magnetic field. This magnetic field strength is in good agreement with the expected field of an AMP; accordingly, the observed spin down places the source precisely

where one would expect it on a P-$\dot{P}$ diagram. However, this observed spin down may also reflect contributions of magnetic propeller torques and gravitational radiation, and it establishes an upper limit of $Q < 4.4 \times 10^{36}$ g cm^2 for the mass quadrupole moment. We measured an orbital period derivative of $\dot{P}_{\rm orb} = (3.5 \pm 0.2) \times 10^{-12}$ s s^{-1}. Given the short orbital period of 2.0 hr, this gives an unexpectedly short $P/\dot{P}$ time scale of 6×10^7 yr, adding SAX J1808.4−3658 to the list of LMXBs with anomalously high period derivatives. The relatively well-understood nature of this system (a NS accreting from a $\approx$0.05 $M_\odot$ brown dwarf at a known mass transfer rate) makes it a good target for the theoretical exploration of why we observe such rapid period evolution.

The anti-correlation between the fractional harmonic amplitude (r_2) and the flux ($f_{\rm x}$) that we observed in SAX J1808.4−3658 is also present in every other AMP we observed for which observations sampled a wide range of fluxes. These sources exhibited a power-law relationship $r_2 \propto f_{\rm x}^{\gamma}$ with slopes ranging from $\gamma = -0.47$ to -0.70. Such relations are quite clear in IGR J00291+5934, and XTE J1751−305, XTE J1807−294; the three other AMPs that we analyzed, XTE J0929−314, XTE J1814−338, and HETE J1900.1−2455, did not as fully explore a range of fluxes, but they too seemed to obey a similar relation. These trends may be evidence of the recession of the accretion disk as the outbursts dim, which could reveal the antipodal hot spot, thereby increasing the amplitude of the second harmonic. Alternatively, we may be observing the decreasing radius of the hot spot as the accretion rate falls, which should also increase the harmonic content. A better understanding of these striking trends may have to wait for better relativistic MHD modeling of accretion column (e.g. Romanova et al. 2004), given the complexities of their interaction in low-field pulsars.

We examined the energy dependence of the persistent pulsations and thermonuclear burst oscillations from SAX J1808.4−3658. We confirmed the soft phase lags previously discovered from this source, and we discovered that these phase lags increased as the source flux decayed slowly following its peak flux. When the source decay became rapid and the outburst entered its flaring tail stage, this relationship reversed, and the phase lags diminished as the flux dimmed further. This result, along with the pulse profile changes observed at the beginning of the flaring tail stage, suggested an abrupt change in the geometry of the accretion disk and column at this time in the outburst. In contrast, the thermonuclear burst oscillation timing did not show appreciable lags, and the burst oscillation phases and fractional amplitudes appeared to be relatively independent of energy.

One interesting thread that runs throughout this thesis is the discovery of abrupt changes in the patterns of X-ray emission when the accretion rate falls off rapidly between the main outburst and the "flaring tail" stage. The pulse shape can change dramatically during this transition, as seen in the 2002 and 2005 outbursts from SAX J1808.4−3658. The fractional amplitude of the second harmonic increases as the flux drops according to the anti-correlations discovered for most of the sources with persistent pulsations. And the soft lags observed in SAX J1808.4−3658 reverse their energy dependence, anti-correlating with flux prior to this transition and correlating with flux after it. All these phenomena point toward a critical accretion rate, across which the behavior of the accretion disk, column, and resulting hot spot change markedly. Developing a model that accommodates such changes represents a theoretical challenge that could provide significant progress in our understanding of these systems.

There is still much work to be done. When the *RXTE* ceases operation, it will have left behind an impressive legacy of data on the eight known millisecond pulsars. While archival searches for coherent X-ray pulsations in other LMXBs may yield additional sources, they have not yet (Chandler and Rutledge 2000; Dib et al. 2005). In my view, deeper exploration

of the outbursts from the known sources is a much more promising direction. For many of these sources, the phase and amplitude variations in the pulsations have not been explored in nearly the detail with which we studied SAX J1808.4−3658. It is my hope that the work in this thesis will spur on such investigation, and that such efforts will lead to improved models of the hot spot and accretion onto neutron stars.

Appendix A

Assessing the significance of reported burst oscillations

Determining the significance of candidate detections of oscillations during thermonuclear bursts requires careful analysis. In all cases, reported burst oscillations have been found using dynamic power spectra, in which a burst light curve is divided into short intervals (typically $T_{\rm int}$ =1–4 s) and a power spectrum is calculated for each. See §3.2 for a discussion of this technique.

Assuming a "unit normalization" of the power spectrum, in which the average power in the absence of a signal is 1, the probability that Poisson noise would cause the power in a single bin to exceed some power P is

$$\mathrm{Prob}(P_{\rm noise} > P) = \exp(-P)\,. \tag{A.1}$$

Note that many papers use the alternative normalization of Leahy et al. 1983, typically referred to as the "Leahy norm," in which the mean noise power is 2, and thus the probability of a given noise power is $\exp(-\frac{1}{2}P)$. Of course, there will be many power measurements in the dynamic power spectrum: the analysis of a light curve of length T over a frequency bandwidth of $\Delta\nu$ yields $N_{\rm tr} = T\Delta\nu$ independent trials if the time intervals do not overlap and the frequency range is not oversampled. These trials must be taken into account in reckoning the singificance of a maximum detected power $P_{\rm max}$:

$$\alpha \equiv \mathrm{Prob}(P_{\rm max} \text{ is due to noise}) = N_{\rm tr}\exp(-P_{\rm max})\,. \tag{A.2}$$

(In statistical jargon, this α is the probability of a type-I error, the detection of a false positive.) The resulting probability is typically presented as the σ associated with a two-sided test for a Gaussian distribution,

$$\alpha = \frac{2}{\sqrt{2\pi}}\int_{\sigma}^{\infty}\exp(-u^2/2)\,du\,. \tag{A.3}$$

However, the assumption in equation A.2 that the intervals are non-overlapping and the frequency is not oversampled is almost never correct. Time intervals are frequently made to overlap by factors of 4–16 (i.e., if the time step between the start of subsequent intervals is $T_{\rm int}/16$ – $T_{\rm int}/4$), and frequencies are sometimes oversampled. This enhances the sensitivity to signals that are split between time intervals or fall between frequency bins, as described in §3.2. This increased sensitivity comes at the expense of statistical

independence. Alternatively, power spectra from different bursts may be averaged together ("stacked") to enhance a peak that is not particularly significant in a single burst but is present throughout. The best way to estimate significance in these cases is to run Monte Carlo simulations.

One must also keep in mind other factors that drive up the number of trials. Proper studies will account for the number of bursts searched and the number of energy ranges used. Less commonly accounted for, however, are multiple searches done with different interval lengths: often multiple dynamic power spectra are created using different time interval lengths (e.g., 1 s to catch any short-lived oscillations or oscillations with rapidly changing frequencies, and 4 s to boost the power of longer-lived oscillations). Again, such methods require Monte Carlo simulations to reliably establish significances due to the the non-independence of the spectra. Finally, it is important to consider the sources that have been searched without success (and hence generally without a published report): in the MIT burst database, there are 13 bursters with definite burst oscillation detections and 30 without (Galloway et al. 2006). All 30 have been searched, to various extents. Given these additional factors, anything less than a 4 σ detection will usually be viewed with some suspicion.

I have marked four of the spin frequencies listed in Table 1.1 as uncertain detections. In three of the cases (XTE J1739−285, SAX J1748.9−2021, and XB 1254−690), the detection relied on a single time and frequency bin with an excess of power. One (GS 1826−238) relied on the cumulative power observed by adding together multiple spectra; this method was also applied in the more solid 45 Hz detection from EXO 0748−676. None of these oscillation detections were found using the more conservative search techniques applied to all the sources in the MIT burst database (Galloway et al. 2006).

With the exception of EXO 0748−676, I am unconvinced by these detections and believe that they should not be the sole drivers for new theoretical work. This caveat particularly applies to the surprising report of 1122 Hz oscillations from XTE J1739−285, which would otherwise upend the current thinking on the distribution of pulsar spin frequencies (see §1.6.2). This stance clearly requires some justification on a source-by-source basis.

A.1 XTE J1739−285

Kaaret et al. (2007) reported evidence of an 1122 Hz burst oscillation during the brightest of six observed bursts from XTE J1739−285. This discovery, if true, would be exceptional: it would make XTE J1739−285 the most rapidly rotating known neutron star by a wide margin, providing a meaningful constraint on the radius of this source in order for it to avoid centrifugal breakup. Given its potential importance, its significance must be addressed with great care.

Kaaret et al. (2007) report a peak power of 21.4 (unit norm),[1] from a dynamic power spectrum with 4 s intervals, each advancing by 0.125 s over 20 s of burst data. They derived a confidence of 99.96% (3.54 σ) by simulating 400,000 bursts with similar lightcurves to the burst with the detection and noting that only 14 of the bursts had four consecutive powers above 20.2, as did the real burst, within the search range of 20–2000 Hz. They also note that 58 simulated bursts had single-peak powers exceeding the observed power of 21.4. This confidence is adjusted for their use of two trial energy ranges. It also accounts for their analysis of multiple bursts.

[1] Kaaret et al. (2007) uses the Leahy norm, so powers listed therein are twice what we list here.

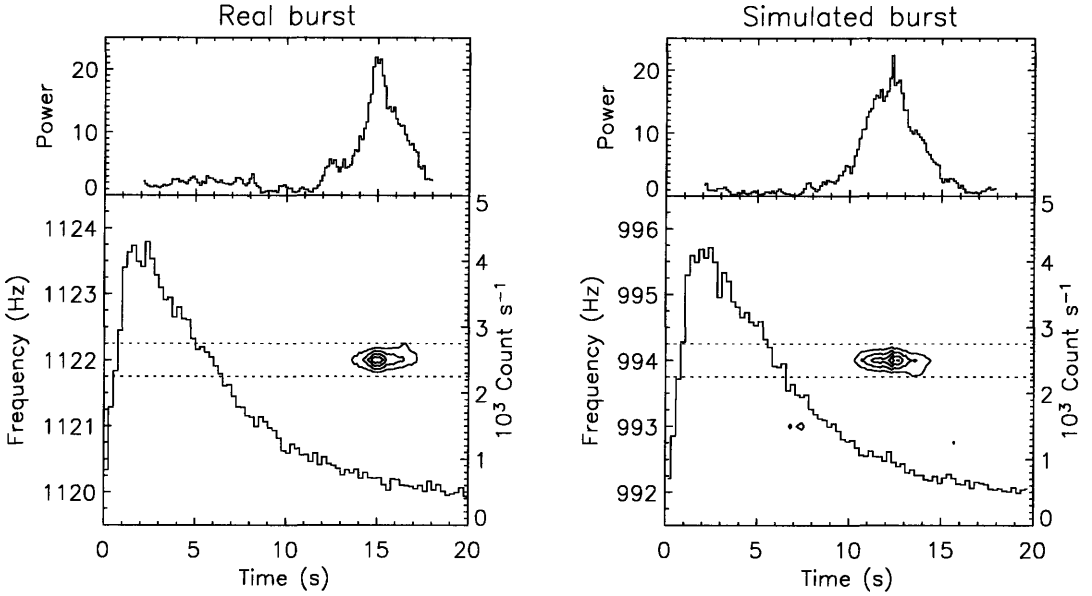

Figure A-1 Comparison of the reported burst oscillation from XTE J1739−285 and a noise peak from a simulated burst with no signal.

To verify this estimate of the significance, I created 50,000 artificial bursts with poisson noise conforming to the same light curve as the original burst when binned to 0.25 s. Each was 20 seconds long and binned at 125 μs. I searched each using 4 s intervals with 0.125 s between intervals, and I only noted noise peaks between 20 Hz and 2000 Hz. Of those 50,000 fake bursts, 15 had noise peaks with powers ≥ 21.4. This figure is in good agreement with the detection of 58/400,000 noise peaks at this level by Kaaret et al. (2007). Factoring in the search of 6 bursts and 2 different energy ranges, I arrive at a false-positive probability of $\alpha = 15/50,000 \times 6 \times 2 = 0.36\%$, or 2.9 σ.

This significance differs from the published value principally due to the different tests of a "detection" during the Monte Carlo simulations. Simply asking how many noise peaks exceed the actual power is the most straightforward test, and it is the one I adopted. The Kaaret et al. (2007) criterion of requiring four consecutive noise powers to exceed the observed maximum power for four consecutive powers is somewhat more contrived. Examining the detection power vs. time curves in Figure A-1, we see that the real burst has four consecutive peaks with powers ≥ 20.4, while the simulated burst does not. But if we instead defined a detection as eight consecutive bins with powers ≥ 15.0, the actual burst would note be detected, while the simulated burst would. Furthermore, if the Kaaret et al. (2007) search also employed other time interval lengths, that would drive the significance down further. In short, I consider the 2.9 σ figure to be a more reliable estimate.

A.2 SAX J1748.9−2021

The reported detection of 410 Hz burst oscillations from a single burst in SAX J1748.9−2021 by Kaaret et al. (2003) suffers from similar problems. They searched 16 bursts from this source using overlapping 3 s intervals offset by 0.25 s, looking for excess power in 100–

1000 Hz. They found a unit-norm power of 19.4 in a single time and frequency bin. Accounting only for their 3 s $\times$ 900 Hz = 2700 frequency bins, they estimated a significance of $\exp(-19.4) \times 2700 = 1.1 \times 10^{-5}$, or 4.4 σ. They did not account for multiple time intervals, multiple interval lengths, or multiple bursts.

Again, accurate accounting for the many non-independent trials requires Monte Carlo simulation. I created 1,500 synthetic bursts with the same light curves as the burst in question from SAX J1748.9−2021. Of these, 11 had noise peaks exceeding the detection power in the actual burst when searched in the manner described in Kaaret et al. (2003). Assuming that only a single time interval length and energy range were used, that gives a significance of $11/1,500 \times 16 = 12\%$, or 1.6 σ. This detection can be safely disregarded as noise.

A.3 GS 1826−238

Thompson et al. (2005) reported the detection of 612 Hz oscillations in the tails of 3 bursts from GS 1826−238, the "clocked burster." The oscillations were detected in the 10–30 keV data but not the full 2–60 keV PCA range. The number of energy bands tested was not given, but they do note that inclusion of higher or lower-energy photons weakened the result. The oscillation strength was achieved by stacking 180 0.25 s-long power spectra from 15–30 s after the peaks of the three bursts, which gave a peak 4.7 σ above the mean of the roughly Gaussian power distribution. Accounting for the 256 frequency bins in their 0–1024 Hz range, they gave a chance detection probability of 0.033%, or 3.6 σ.

While we cannot ascertain the significance without a better knowledge of the number of trials attempted, it seems quite safe to assume that many energy and time ranges were tested. Assuming (perhaps conservatively) that 3 independent energy ranges and 3 independent time ranges were analyzed, the significance becomes 0.3%, or just under 3 σ.

A.4 XB 1254−690

Bhattacharyya (2007) report a candidate oscillation at 95 Hz during the rise of a burst from XB 1254−690. The possible detection was only present in a single time and frequency bin, with a unit-norm power of 12.2. They used non-overlapping 1 s intervals to search 5 bursts, looking at powers up to the Nyquist frequency of 2048 Hz. They report a chance detection probability of $\exp(-12.2) * 1$ s $\times$ 2048 Hz $\times$ 5 = 5.4%, or 1.9 σ.

This calculation of the number of trials does not take into account the number of time intervals per burst. The bursts were weak and short, lasting 10–15 s each. For 1 s intervals, this makes the detection weaker still, driving the probability of detecting a noise peak of the reported power up to roughly 50%.

A.5 EXO 0748−676

Villarreal and Strohmayer (2004) reported the detection of a 45 Hz peak by stacking the power spectra of the decay intervals of 38 X-ray bursts. They correctly account for the number of trials and derive a significance level of 4×10^{-8}. This statistic would be sufficient to banish all doubt, except that it fails to pass a basic "sniff test" of timing studies: if every modification of the reported analysis *weakens* the detection, then it deserves further scrutiny. In this case, adding or removing bursts from the sample, not using the exact

combination of energy channels used by the authors, or not using the same time intervals within the burst tails all weaken the significance, in many cases weakening it to the degree that it is marginal or even not detectable.

However, even if one accounted for the additional number of trials needed to test multiple energy ranges and combinations of bursts, the detection would still be reasonably good. Furthermore, the possible detection of narrow iron lines during some bursts from EXO 0748−676 (Cottam et al. 2002; another result that is diminished by the inclusion of additional data) requires a slow spin ($\lesssim$100 Hz) to avoid excessive Doppler broadening. While this detection is not certain, the evidence that EXO 0748−676 is spinning at 45 Hz is compelling.

Appendix B

Frequency offsets produced by position errors

An accurate source position is essential for high-precision pulsar timing. An incorrect position results in errors during the barycentering of X-ray arrival times, producing a frequency offset due to improperly corrected Doppler shifts (see Manchester and Peters 1972 for its effects on pulsar timing). These errors will be particularly accute for sources near the ecliptic plane, as is the case for SAX J1808.4−3658. Defining β as the eccliptic lattitude, SAX J1808.4−3658 is at $\beta = -13.6°$.

Shifting the photon arrival times to the solar system barycenter requires the addition of a geometric delay to the photon arrival times, $t_{\rm delay} = \mathbf{r} \cdot \hat{\mathbf{s}}/c$, where $\mathbf{r}$ is a vector from the barycenter to the *RXTE* and $\hat{\mathbf{s}}$ is the unit direction vector of the source. For simplicity, we adopt a coordinate system in which the x-y plane defines the ecliptic, with the x-axis oriented toward the source. To convert coordinates from right ascension (α) and declination (δ), use

$$\sin\beta = \sin\delta\cos\epsilon - \cos\delta\sin\epsilon\sin\alpha \quad \text{(B.1)}$$

$$\cos\lambda\cos\beta = \cos\alpha\cos\delta\,, \quad \text{(B.2)}$$

where $\epsilon = 23.43°$ is the tilt of the Earth's rotational axis relative to the ecliptic, and λ gives the ecliptic longitude. In this system, $\hat{\mathbf{s}} = \hat{\mathbf{i}}\cos\beta + \hat{\mathbf{k}}\sin\beta$. We can write the Earth's orbit as $\mathbf{r} = a_\oplus(\hat{\mathbf{i}}\cos\tau + \hat{\mathbf{j}}\sin\tau)$, with $\tau = t/P_\oplus - \tau_0$ parameterizing time such that $\tau = 0$ when the Earth is closest to the source. Here we have ignored eccentricity and numerous other effects, all of which are utterly negligible for our calculations of frequency errors. Thus we can write the geometric delay as

$$t_{\rm delay} = \frac{a_\oplus}{c}\cos\beta\cos\tau\,. \quad \text{(B.3)}$$

Suppose that the source position that we assume during barycentering differs from the actual source position by some small $\Delta\tau$ and $\Delta\beta$. The calculated arrival times at the barycenter will be greater than the actual arrival times by

$$\Delta t = \frac{a_\oplus}{c}\cos(\beta+\Delta\beta)\cos(\tau+\Delta\tau) - t_{\rm delay}\,. \quad \text{(B.4)}$$

Applying small-angle approximates and removing second-order terms in $\Delta\tau$ and $\Delta\beta$, we

find

$$\Delta t = -\frac{a_\oplus}{c}(\Delta\tau \cos\beta \sin\tau + \Delta\beta \sin\beta \cos\tau)\,. \tag{B.5}$$

This error causes frequency and frequency derivative offsets of

$$\Delta\nu = -\nu\frac{d\Delta t}{dt} = \nu\frac{a_\oplus}{c}\left(\frac{2\pi}{P_\oplus}\right)(\Delta\tau \cos\beta\cos\tau - \Delta\beta \sin\beta \sin\tau) \tag{B.6}$$

$$\Delta\dot{\nu} = -\nu\frac{d^2\Delta t}{dt^2} = -\nu\frac{a_\oplus}{c}\left(\frac{2\pi}{P_\oplus}\right)^2(\Delta\tau \cos\beta \sin\tau + \Delta\beta \sin\beta \cos\tau)\,. \tag{B.7}$$

For a 400 Hz pulsar at $\beta = -13.6°$, a position error of $\Delta\tau = 0\farcs2$ parallel to the plane of the ecliptic produces frequency offsets of

$$\Delta\nu = 40\cos\tau \text{ nHz} \tag{B.8}$$

$$\Delta\dot{\nu} = -8\times10^{-15}\sin\tau \text{ Hz s}^{-1}\,. \tag{B.9}$$

These offsets are comparable with the expected timing uncertainties. Each outburst gives a baseline of about 2×10^6 s over which we can typically measure pulse arrival times with an accuracy of better than 25 μs, or 1×10^{-2} cycles, producing $\sim$5 nHz frequency uncertainties. By similar logic, we should be sensitive to $\dot{\nu}$'s as small as $\sim3\times10^{-15}$ Hz s^{-1}. In practice, the pulse shape noise observed in SAX J1808.4−3658 makes the actual uncertainties somewhat greater than these back-of-the-envelope values, increasing the ν uncertainty by a factor of $\sim$2 and the $\dot{\nu}$ uncertainty by a factor of $\sim$10, but the frequency uncertainty is still substantially less than the offsets due to a $0\farcs2$ position error.

Bibliography

Alpar, M. A., A. F. Cheng, M. A. Ruderman, and J. Shaham (1982). A new class of radio pulsars. Nature 300, 728–730.

Altamirano, D., P. Casella, A. Patruno, R. Wijnands, and M. van der Klis (2007). Intermittent millisecond X-ray pulsations from the neutron-star X-ray transient SAX J1748.9−2021 in the globular cluster NGC 6440. ApJ, submitted. astro-ph/0708.1316.

Anderson, P. W. and N. Itoh (1975). Pulsar glitches and restlessness as a hard superfluidity phenomenon. Nature 256, 25–27.

Andersson, N., K. D. Kokkotas, and N. Stergioulas (1999). On the Relevance of the R-Mode Instability for Accreting Neutron Stars and White Dwarfs. ApJ 516, 307–314.

Applegate, J. H. and J. Shaham (1994). Orbital period variability in the eclipsing pulsar binary PSR B1957+20: Evidence for a tidally powered star. ApJ 436, 312–318.

Arzoumanian, Z., A. S. Fruchter, and J. H. Taylor (1994). Orbital variability in the eclipsing pulsar binary PSR B1957+20. ApJ 426, L85–L88.

Bhattacharyya, S. (2007). Timing properties of XB 1254−690. MNRAS 377, 198–202.

Bhattacharyya, S., T. E. Strohmayer, C. B. Markwardt, and J. H. Swank (2006). The Discovery of a Neutron Star with a Spin Frequency of 530 Hz in A 1744−361. ApJ 639, L31–L34.

Bhattacharyya, S., T. E. Strohmayer, M. C. Miller, and C. B. Markwardt (2005). Constraints on Neutron Star Parameters from Burst Oscillation Light Curves of the Accreting Millisecond Pulsar XTE J1814−338. ApJ 619, 483–491.

Bildsten, L. (1998). Gravitational Radiation and Rotation of Accreting Neutron Stars. ApJ 501, L89–L93.

Bildsten, L. and D. Chakrabarty (2001). A Brown Dwarf Companion for the Accreting Millisecond Pulsar SAX J1808.4−3658. ApJ 557, 292–296.

Boynton, P. E. and J. E. Deeter (1985). X-Ray Pulse Timing Studies, II. Overcoming Limitations on Precision. In S. Hayakawa and F. Nagase (Eds.), Workshop on the Timing Studies of X-ray Sources, 13–28.

Boynton, P. E., J. E. Deeter, F. K. Lamb, G. Zylstra, S. H. Pravdo, N. E. White, K. S. Wood, and D. J. Yentis (1984). New evidence on the nature of the neutron star and accretion flow in VELA X-1 from pulse timing observations. ApJ 283, L53–L56.

Bradt, H. V., R. E. Rothschild, and J. H. Swank (1993). X-ray timing explorer mission. A&AS 97, 355–360.

Brady, P. R., T. Creighton, C. Cutler, and B. F. Schutz (1998). Searching for periodic sources with LIGO. Phys. Rev. D 57, 2101–2116.

Burderi, L., T. Di Salvo, F. D'Antona, N. R. Robba, and V. Testa (2003). The optical counterpart to SAX J1808.4−3658 in quiescence: Evidence of an active radio pulsar? A&A 404, L43–L46.

Burderi, L., T. Di Salvo, G. Lavagetto, M. T. Menna, A. Papitto, A. Riggio, R. Iaria, F. D'Antona, N. R. Robba, and L. Stella (2007). Timing an Accreting Millisecond Pulsar: Measuring the Accretion Torque in IGR J00291+5934. ApJ 657, 961–966.

Burderi, L., T. Di Salvo, M. T. Menna, A. Riggio, and A. Papitto (2006). Order in the Chaos: Spin-up and Spin-down during the 2002 Outburst of SAX J1808.4−3658. ApJ 653, L133–L136.

Burgay, M., L. Burderi, A. Possenti, N. D'Amico, R. N. Manchester, A. G. Lyne, F. Camilo, and S. Campana (2003). A Search for Pulsars in Quiescent Soft X-Ray Transients. I. ApJ 589, 902–910.

Cameron, P. B., R. E. Rutledge, F. Camilo, L. Bildsten, S. M. Ransom, and S. R. Kulkarni (2007). Variability of 19 Millisecond Pulsars in 47 Tucanae with Chandra HRC-S. ApJ 660, 587–594.

Campana, S., P. D'Avanzo, J. Casares, S. Covino, G. Israel, G. Marconi, R. Hynes, P. Charles, and L. Stella (2004). Indirect Evidence of an Active Radio Pulsar in the Quiescent State of the Transient Millisecond Pulsar SAX J1808.4−3658. ApJ 614, L49–L52.

Campana, S., N. Ferrari, L. Stella, and G. L. Israel (2005). XMM-Newton observations of two transient millisecond X-ray pulsars in quiescence. A&A 434, L9–L12.

Campana, S., L. Stella, F. Gastaldello, S. Mereghetti, M. Colpi, G. L. Israel, L. Burderi, T. Di Salvo, and R. N. Robba (2002). An XMM-Newton Study of the 401 Hz Accreting Pulsar SAX J1808.4−3658 in Quiescence. ApJ 575, L15–L19.

Casella, P., D. Altamirano, R. Wijnands, and M. van der Klis (2007). Discovery of coherent millisecond X-ray pulsations in Aql X-1. ApJ, submitted. astro-ph/0708.1110.

Chakrabarty, D. (2005). Millisecond Pulsars in X-Ray Binaries. In F. A. Rasio and I. H. Stairs (Eds.), Binary Radio Pulsars, Volume 328 of Astronomical Society of the Pacific Conference Series, 279–289.

Chakrabarty, D. and E. H. Morgan (1998). The two-hour orbit of a binary millisecond X-ray pulsar. Nature 394, 346–348.

Chakrabarty, D., E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt (2003). Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. Nature 424, 42–44.

Chandler, A. M. and R. E. Rutledge (2000). No Persistent Pulsations in Aquila X-1 As It Fades into Quiescence. ApJ 545, 1000–1006.

Cottam, J., F. Paerels, and M. Mendez (2002). Gravitationally redshifted absorption lines in the X-ray burst spectra of a neutron star. Nature 420, 51–54.

Cui, W., E. H. Morgan, and L. G. Titarchuk (1998). Soft Phase Lags in the Millisecond X-Ray Pulsar SAX J1808.4−3658. ApJ 504, L27–L30.

Cumming, A., S. M. Morsink, L. Bildsten, J. L. Friedman, and D. E. Holz (2002). Hydrostatic Expansion and Spin Changes during Type I X-Ray Bursts. ApJ 564, 343–352.

Cumming, A., E. Zweibel, and L. Bildsten (2001). Magnetic Screening in Accreting Neutron Stars. ApJ 557, 958–966.

Cutler, C. (2002). Gravitational waves from neutron stars with large toroidal B fields. Phys. Rev. D 66(8), 084025.1–6.

Davidson, K. and J. P. Ostriker (1973). Neutron-Star Accretion in a Stellar Wind: Model for a Pulsed X-Ray Source. ApJ 179, 585–598.

Deeter, J. E., S. H. Pravdo, and P. E. Boynton (1981). Pulse-timing observations of Hercules X-1. ApJ 247, 1003–1012.

Dib, R., S. M. Ransom, P. S. Ray, V. M. Kaspi, and A. M. Archibald (2005). An RXTE Archival Search for Coherent X-Ray Pulsations in the Low-Mass X-Ray Binary 4U 1820−30. ApJ 626, 333–342.

Doroshenko, O., O. Löhmer, M. Kramer, A. Jessner, R. Wielebinski, A. G. Lyne, and C. Lange (2001). Orbital variability of the PSR J2051−0827 binary system. A&A 379, 579–587.

Ekşi, K. Y., L. Hernquist, and R. Narayan (2005). Where Are All the Fallback Disks? Constraints on Propeller Systems. ApJ 623, L41–L44.

Falanga, M., L. Kuiper, J. Poutanen, E. W. Bonning, W. Hermsen, T. di Salvo, P. Goldoni, A. Goldwurm, S. E. Shaw, and L. Stella (2005). INTEGRAL and RXTE observations of accreting millisecond pulsar IGR J00291+5934 in outburst. A&A 444, 15–24.

Ford, E. C. (2000). Relativistic Effects in the Pulse Profile of the 2.5 Millisecond X-Ray Pulsar SAX J1808.4−3658. ApJ 535, L119–L122.

Fritschel, P. (2003). Second generation instruments for the Laser Interferometer Gravitational Wave Observatory (LIGO). Proc. SPIE 4856, 282–291.

Fruchter, A. S., G. Berman, G. Bower, M. Convery, W. M. Goss, T. H. Hankins, J. R. Klein, D. J. Nice, M. F. Ryba, D. R. Stinebring, J. H. Taylor, S. E. Thorsett, and J. M. Weisberg (1990). The eclipsing millisecond pulsar PSR 1957+20. ApJ 351, 642–650.

Gaensler, B. M., B. W. Stappers, and T. J. Getts (1999). Transient Radio Emission from SAX J1808.4−3658. ApJ 522, L117–L119.

Galloway, D. K., D. Chakrabarty, E. H. Morgan, and R. A. Remillard (2002). Discovery of a High-Latitude Accreting Millisecond Pulsar in an Ultracompact Binary. ApJ 576, L137–L140.

Galloway, D. K., D. Chakrabarty, M. P. Muno, and P. Savov (2001). Discovery of a 270 Hertz X-Ray Burst Oscillation in the X-Ray Dipper 4U 1916−053. ApJ 549, L85–L88.

Galloway, D. K. and A. Cumming (2006). Helium-rich Thermonuclear Bursts and the Distance to the Accretion-powered Millisecond Pulsar SAX J1808.4−3658. ApJ 652, 559–568.

Galloway, D. K., A. B. Giles, K. Wu, and J. G. Greenhill (2001). Accretion column eclipses in the X-ray pulsars GX 1+4 and RX J0812.4−3114. MNRAS 325, 419–425.

Galloway, D. K., C. B. Markwardt, E. H. Morgan, D. Chakrabarty, and T. E. Strohmayer (2005). Discovery of the Accretion-powered Millisecond X-Ray Pulsar IGR J00291+5934. ApJ 622, L45–L48.

Galloway, D. K., E. H. Morgan, M. I. Krauss, P. Kaaret, and D. Chakrabarty (2007). Intermittent Pulsations in an Accretion-powered Millisecond Pulsar. ApJ 654, L73–L76.

Galloway, D. K., M. P. Muno, J. M. Hartman, P. Savov, D. Psaltis, and D. Chakrabarty (2006). Thermonuclear (type-I) X-ray bursts observed by the Rossi X-ray Timing Explorer. ApJS, submitted. astro-ph/0608259.

Ghosh, P. and F. K. Lamb (1979). Accretion by rotating magnetic neutron stars. III — Accretion torques and period changes in pulsating X-ray sources. ApJ 234, 296–316.

Ghosh, P. and F. K. Lamb (1991). Plasma physics of accreting neutron stars. In Neutron Stars: Theory and Observation, 363–444. Dordrecht: Kluwer.

Gierliński, M., C. Done, and D. Barret (2002). Phase-resolved X-ray spectroscopy of the millisecond pulsar SAX J1808.4−3658. MNRAS 331, 141–153.

Gierliński, M. and J. Poutanen (2005). Physics of accretion in the millisecond pulsar XTE J1751−305. MNRAS 359, 1261–1276.

Giles, A. B., K. M. Hill, and J. G. Greenhill (1999). The optical counterpart of SAX J1808.4−3658, the transient bursting millisecond X-ray pulsar. MNRAS 304, 47–51.

Glasser, C. A., C. E. Odell, and S. E. Seufert (1994). The Proportional Counter Array (PCA) Instrument for the X-ray Timing Explorer Satellite (XTE). ITNS 41, 1343–1348.

Greenhill, J. G., A. B. Giles, and C. Coutures (2006). A transient I-band excess in the optical spectrum of the accreting millisecond pulsar SAX J1808.4−3658. MNRAS 370, 1303–1308.

Hartman, J. M., D. Chakrabarty, D. K. Galloway, M. P. Muno, P. Savov, M. Mendez, S. van Straaten, and T. Di Salvo (2003). Discovery of 619 Hz Thermonuclear Burst Oscillations in the Low-Mass X-Ray Binary 4U 1608−52. BAAS 35, 865.

Heinke, C. O., P. G. Jonker, R. Wijnands, and R. E. Taam (2007). Constraints on Thermal X-Ray Radiation from SAX J1808.4−3658 and Implications for Neutron Star Neutrino Emission. ApJ 660, 1424–1427.

Hellier, C., K. O. Mason, A. P. Smale, and D. Kilkenny (1990). The ephemeris of X 1822−371. MNRAS 244, 39P–43P.

Homer, L., P. A. Charles, D. Chakrabarty, and L. van Zyl (2001). The optical counterpart to SAX J1808.4−3658: observations in quiescence. MNRAS 325, 1471–1476.

Illarionov, A. F. and R. A. Sunyaev (1975). Why the Number of Galactic X-ray Stars Is so Small? A&A 39, 185–195.

in 't Zand, J. J. M., R. Cornelisse, E. Kuulkers, J. Heise, L. Kuiper, A. Bazzano, M. Cocchi, J. M. Muller, L. Natalucci, M. J. S. Smith, and P. Ubertini (2001). The first outburst of SAX J1808.4−3658 revisited. A&A 372, 916–921.

in 't Zand, J. J. M., J. Heise, J. M. Muller, A. Bazzano, M. Cocchi, L. Natalucci, and P. Ubertini (1998). Discovery of the X-ray transient SAX J1808.4−3658, a likely low-mass X-ray binary. A&A 331, L25–L28.

Jahoda, K., C. B. Markwardt, Y. Radeva, A. H. Rots, M. J. Stark, J. H. Swank, T. E. Strohmayer, and W. Zhang (2006). Calibration of the Rossi X-Ray Timing Explorer Proportional Counter Array. ApJS 163, 401–423.

Jahoda, K., J. H. Swank, A. B. Giles, M. J. Stark, T. Strohmayer, W. Zhang, and E. H. Morgan (1996). In-orbit performance and calibration of the Rossi X-ray Timing Explorer (RXTE) Proportional Counter Array (PCA). In O. H. Siegmund and M. A. Gummin (Eds.), Proc. SPIE Vol. 2808, p. 59-70, EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy VII, Oswald H. Siegmund; Mark A. Gummin; Eds., 59–70.

Joss, P. C. and S. A. Rappaport (1983). On the origin of the 6.1-ms pulsar. Nature 304, 419–421.

Kaaret, P., E. H. Morgan, R. Vanderspek, and J. A. Tomsick (2006). Discovery of the Millisecond X-Ray Pulsar HETE J1900.1−2455. ApJ 638, 963–967.

Kaaret, P., Z. Prieskorn, J. J. M. i. Zand, S. Brandt, N. Lund, S. Mereghetti, D. Götz, E. Kuulkers, and J. A. Tomsick (2007). Evidence of 1122 Hz X-Ray Burst Oscillations from the Neutron Star X-Ray Transient XTE J1739−285. ApJ 657, L97–L100.

Kaaret, P., J. J. M. i. Zand, J. Heise, and J. A. Tomsick (2002). Discovery of Millisecond Variability in the Neutron Star X-Ray Transient SAX J1750.8−2900. ApJ 575, 1018–1024.

Kaaret, P., J. J. M. i. Zand, J. Heise, and J. A. Tomsick (2003). Discovery of X-Ray Burst Oscillations from a Neutron Star X-Ray Transient in the Globular Cluster NGC 6440. ApJ 598, 481–485.

Kirsch, M. G. F., K. Mukerjee, M. G. Breitfellner, S. Djavidnia, M. J. Freyberg, E. Kendziorra, and M. J. S. Smith (2004). Studies of orbital parameters and pulse profile of the accreting millisecond pulsar XTE J1807−294. A&A 423, L9–L12.

Kraft, R. P., J. Mathews, and J. L. Greenstein (1962). Binary Stars among Cataclysmic Variables. II. Nova WZ Sagittae: a Possible Radiator of Gravitational Waves. ApJ 136, 312–314.

Lamb, F. K. and S. Boutloukos (2007). Accreting Neutron Stars in Low-Mass X-Ray Binary Systems. astro-ph/0608259.

Lamb, F. K., C. J. Pethick, and D. Pines (1973). A Model for Compact X-Ray Sources: Accretion by Rotating Magnetic Stars. ApJ 184, 271–290.

Lattimer, J. M. and M. Prakash (2001). Neutron Star Structure and the Equation of State. ApJ 550, 426–442.

Leahy, D. A., W. Darbro, R. F. Elsner, M. C. Weisskopf, S. Kahn, P. G. Sutherland, and J. E. Grindlay (1983). On searches for pulsed emission with application to four globular cluster X-ray sources — NGC 1851, 6441, 6624, and 6712. ApJ 266, 160–170.

Linares, M., M. van der Klis, D. Altamirano, and C. B. Markwardt (2005). Discovery of Kilohertz Quasi-periodic Oscillations and Shifted Frequency Correlations in the Accreting Millisecond Pulsar XTE J1807−294. ApJ 634, 1250–1260.

Lovelace, R. V. E., A. K. Kulkarni, and M. M. Romanova (2007). Torsional Magnetic Oscillations in Type I X-Ray Bursts. ApJ 656, 393–398.

Manchester, R. N., G. B. Hobbs, A. Teoh, and M. Hobbs (2005). The Australia Telescope National Facility Pulsar Catalogue. AJ 129, 1993–2006.

Manchester, R. N. and W. L. Peters (1972). Pulsar Parameters from Timing Observations. ApJ 173, 221–226.

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

Markwardt, C. B. (2003). Measurement of Accretion Torques on the Millisecond Pulsars SAX J1808.4−3658 and XTE J1751−305. BAAS 35, 657.

Markwardt, C. B., H. A. Krimm, and J. H. Swank (2007). SWIFT J1756.9−2508 is a 182 Hz Millisecond X-ray Pulsar. The Astronomer's Telegram 1108.

Markwardt, C. B., E. Smith, and J. H. Swank (2003). Discovery of a Fourth Accreting Millisecond Pulsar, XTE J1807−294. The Astronomer's Telegram 122.

Markwardt, C. B., T. E. Strohmayer, and J. H. Swank (1999). Observation of Kilohertz Quasi-periodic Oscillations from the Atoll Source 4U 1702−429 by the Rossi X-Ray Timing Explorer. ApJ 512, L125–L129.

Markwardt, C. B. and J. H. Swank (2003). XTE J1814-338. IAU Circ. 8144.

Markwardt, C. B., J. H. Swank, T. E. Strohmayer, J. J. M. i. Zand, and F. E. Marshall (2002). Discovery of a Second Millisecond Accreting Pulsar: XTE J1751−305. ApJ 575, L21–L24.

Melatos, A. and D. J. B. Payne (2005). Gravitational Radiation from an Accreting Millisecond Pulsar with a Magnetically Confined Mountain. ApJ 623, 1044–1050.

Méndez, M., M. van der Klis, J. van Paradijs, W. H. G. Lewin, B. A. Vaughan, E. Kuulkers, W. Zhang, F. K. Lamb, and D. Psaltis (1998). Discovery of a Second kHz QPO Peak in 4U 1608−52. ApJ 494, L65–L69.

Morgan, E. H., D. Chakrabarty, R. Wijnands, M. van der Klis, and C. Markwardt (2003). Spin Frequency Evolution of the Accreting Millisecond Pulsar SAX J1808.4−3658. BAAS 35, 629.

Muno, M. P., D. Chakrabarty, D. K. Galloway, and D. Psaltis (2002). The Frequency Stability of Millisecond Oscillations in Thermonuclear X-Ray Bursts. ApJ 580, 1048–1059.

Muno, M. P., D. K. Galloway, and D. Chakrabarty (2004). The Effect of Neutron Star Rotation on the Properties of Thermonuclear X-Ray Bursts. ApJ 608, 930–934.

Muno, M. P., F. Özel, and D. Chakrabarty (2002). The Amplitude Evolution and Harmonic Content of Millisecond Oscillations in Thermonuclear X-Ray Bursts. ApJ 581, 550–561.

Muno, M. P., F. Özel, and D. Chakrabarty (2003). The Energy Dependence of Millisecond Oscillations in Thermonuclear X-Ray Bursts. ApJ 595, 1066–1076.

Papitto, A., T. di Salvo, L. Burderi, M. T. Menna, G. Lavagetto, and A. Riggio (2007). Timing of the accreting millisecond pulsar XTE J1814−338. MNRAS 375, 971–976.

Papitto, A., M. T. Menna, L. Burderi, T. Di Salvo, F. D'Antona, and N. R. Robba (2005). Revised Orbital Parameters of the Accreting Millisecond Pulsar SAX J1808.4−3658. ApJ 621, L113–L116.

Piro, A. L. and L. Bildsten (2005). Surface Modes on Bursting Neutron Stars and X-Ray Burst Oscillations. ApJ 629, 438–450.

Poutanen, J. and M. Gierliński (2003). On the nature of the X-ray emission from the accreting millisecond pulsar SAX J1808.4−3658. MNRAS 343, 1301–1311.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (1992). Numerical Recipes in C. The Art of Scientific Computing, 2nd ed. New York: Cambridge University Press.

Pringle, J. E. and M. J. Rees (1972). Accretion Disc Models for Compact X-Ray Sources. A&A 21, 1–9.

Psaltis, D. and D. Chakrabarty (1999). The Disk-Magnetosphere Interaction in the Accretion-powered Millisecond Pulsar SAX J1808.4−3658. ApJ 521, 332–340.

Radhakrishnan, V. and G. Srinivasan (1982). On the origin of the recently discovered ultra-rapid pulsar. Current Science 51, 1096–1099.

Rappaport, S., F. Verbunt, and P. C. Joss (1983). A new technique for calculations of binary stellar evolution, with application to magnetic braking. ApJ 275, 713–731.

Rappaport, S. A., J. M. Fregeau, and H. Spruit (2004). Accretion onto Fast X-Ray Pulsars. ApJ 606, 436–443.

Roche, P., D. Chakrabarty, L. Morales-Rueda, R. Hynes, S. M. Slivan, C. Simpson, and P. Hewett (1998). SAX J1808.4−3658 = XTE J1808−369. IAU Circ. 6885.

Romanova, M. M., G. V. Ustyugova, A. V. Koldoba, and R. V. E. Lovelace (2004). Three-dimensional Simulations of Disk Accretion to an Inclined Dipole. II. Hot Spots and Variability. ApJ 610, 920–932.

Rothschild, R. E., P. R. Blanco, D. E. Gruber, W. A. Heindl, D. R. MacDonald, D. C. Marsden, M. R. Pelling, L. R. Wayne, and P. L. Hink (1998). In-Flight Performance of the High-Energy X-Ray Timing Experiment on the Rossi X-Ray Timing Explorer. ApJ 496, 538–549.

Rots, A. H., K. Jahoda, and A. G. Lyne (2004). Absolute Timing of the Crab Pulsar with the Rossi X-Ray Timing Explorer. ApJ 605, L129–L132.

Rupen, M. P., V. Dhawan, and A. J. Mioduszewski (2005). Radio detections of SAX J1808.4−3658. The Astronomer's Telegram 524.

Sazonov, S. Y. and R. A. Sunyaev (2001). Scattering in the inner accretion disk and the waveforms and polarization of millisecond flux oscillations in LMXBs. A&A 373, 241–250.

Seidelmann, P. K. (1992). Explanatory Supplement to the Astronomical Almanac. Mill Valley, California: University Science Books.

Smith, D. A., E. H. Morgan, and H. Bradt (1997). Rossi X-Ray Timing Explorer Discovery of Coherent Millisecond Pulsations during an X-Ray Burst from KS 1731−260. ApJ 479, L137–L140.

Spitkovsky, A. (2006). Time-dependent Force-free Pulsar Magnetospheres: Axisymmetric and Oblique Rotators. ApJ 648, L51–L54.

Spitkovsky, A., Y. Levin, and G. Ushomirsky (2002). Propagation of Thermonuclear Flames on Rapidly Rotating Neutron Stars: Extreme Weather during Type I X-Ray Bursts. ApJ 566, 1018–1038.

Spruit, H. C. and H. Ritter (1983). Stellar activity and the period gap in cataclysmic variables. A&A 124, 267–272.

Stella, L., S. Campana, S. Mereghetti, D. Ricci, and G. L. Israel (2000). The Discovery of Quiescent X-Ray Emission from SAX J1808.4−3658, the Transient 2.5 Millisecond Pulsar. ApJ 537, L115–L118.

Strohmayer, T. and L. Bildsten (2006). New views of thermonuclear bursts. In Compact Stellar X-ray Sources, 113–156. New York: Cambridge University Press.

Strohmayer, T. E., K. Jahoda, A. B. Giles, and U. Lee (1997). Millisecond Pulsations from a Low-Mass X-Ray Binary in the Galactic Center Region. ApJ 486, 355–362.

Strohmayer, T. E., C. B. Markwardt, J. H. Swank, and J. in't Zand (2003). X-Ray Bursts from the Accreting Millisecond Pulsar XTE J1814−338. ApJ 596, L67–L70.

Strohmayer, T. E., W. Zhang, J. H. Swank, A. Smale, L. Titarchuk, C. Day, and U. Lee (1996). Millisecond X-Ray Variability from an Accreting Neutron Star System. ApJ 469, L9–L12.

Strohmayer, T. E., W. Zhang, J. H. Swank, N. E. White, and I. Lapidus (1998). On the Amplitude of Burst Oscillations in 4U 1636−54: Evidence for Nuclear-powered Pulsars. ApJ 498, L135–L139.

Tauris, T. M. and E. P. J. van den Heuvel (2006). Formation and evolution of compact stellar X-ray sources. In Compact Stellar X-ray Sources, 623–665. New York: Cambridge University Press.

Taylor, J. H. (1993). Pulsar timing and relativistic gravity. In R. D. Blandford (Ed.), Pulsars as Physics Laboratories, 117–134. New York: Oxford University Press.

Thompson, T. W. J., R. E. Rothschild, J. A. Tomsick, and H. L. Marshall (2005). Chandra and RXTE Spectra of the Burster GS 1826−238. ApJ 634, 1261–1271.

Ushomirsky, G., C. Cutler, and L. Bildsten (2000). Deformations of accreting neutron star crusts and gravitational wave emission. MNRAS 319, 902–932.

van der Klis, M. (1989). Fourier Techniques in X-ray Timing. In H. Ögleman and E. P. J. van den Heuvel (Eds.), Timing Neutron Stars, Proc. NATO Advanced Study Institute, 27–69. Boston: Kluwer Academic Publishers.

van der Klis, M., G. Hasinger, F. Verbunt, J. van Paradijs, T. Belloni, and W. H. G. Lewin (1993). Further ROSAT measurements of the period of 4U 1820−30. A&A 279, L21–L24.

Vaughan, B. A., M. van der Klis, K. S. Wood, J. P. Norris, P. Hertz, P. F. Michelson, J. van Paradijs, W. H. G. Lewin, K. Mitsuda, and W. Penninx (1994). Searches for millisecond pulsations in low-mass X-ray binaries, 2. ApJ 435, 362–371.

Villarreal, A. R. and T. E. Strohmayer (2004). Discovery of the Neutron Star Spin Frequency in EXO 0748−676. ApJ 614, L121–L124.

Wagoner, R. V. (1984). Gravitational radiation from accreting neutron stars. ApJ 278, 345–348.

Wang, Z., D. Chakrabarty, P. Roche, P. A. Charles, E. Kuulkers, T. Shahbaz, C. Simpson, D. A. Forbes, and S. F. Helsdon (2001). The Optical Counterpart of the Accreting Millisecond Pulsar SAX J1808.4−3658 in Outburst: Constraints on the Binary Inclination. ApJ 563, L61–L64.

Watts, A. L. and T. E. Strohmayer (2006). The energy dependence of burst oscillations from the accreting millisecond pulsar XTE J1814−338. MNRAS 373, 769–780.

Watts, A. L., T. E. Strohmayer, and C. B. Markwardt (2005). Analysis of Variability in the Burst Oscillations of the Accreting Millisecond Pulsar XTE J1814−338. ApJ 634, 547–564.

Wijnands, R., J. Homan, C. O. Heinke, J. M. Miller, and W. H. G. Lewin (2005). Chandra Observations of the Accretion-driven Millisecond X-Ray Pulsars XTE J0929−314 and XTE J1751−305 in Quiescence. ApJ 619, 492–502.

Wijnands, R., M. Méndez, C. Markwardt, M. van der Klis, D. Chakrabarty, and E. Morgan (2001). The Erratic Luminosity Behavior of SAX J1808.4−3658 during Its 2000 Outburst. ApJ 560, 892–896.

Wijnands, R., T. Strohmayer, and L. M. Franco (2001). Discovery of Nearly Coherent Oscillations with a Frequency of ~567 Hz during Type I X-Ray Bursts of the X-Ray Transient and Eclipsing Binary X 1658−298. ApJ 549, L71–L75.

Wijnands, R. and M. van der Klis (1998). A millisecond pulsar in an X-ray binary system. Nature 394, 344–346.

Wijnands, R., M. van der Klis, J. Homan, D. Chakrabarty, C. B. Markwardt, and E. H. Morgan (2003). Quasi-periodic X-ray brightness fluctuations in an accreting millisecond pulsar. Nature 424, 44–47.

Wijnands, R. A. D. and M. van der Klis (1997). Discovery of Two Simultaneous Kilohertz Quasi-periodic Oscillations in KS 1731−260. ApJ 482, L65–L68.

Wijnands, R. A. D., M. van der Klis, J. van Paradijs, W. H. G. Lewin, F. K. Lamb, B. Vaughan, and E. Kuulkers (1997). Discovery in 4U 1636−53 of Two Simultaneous Quasi-periodic Oscillations near 900 Hz and 1176 Hz. ApJ 479, L141–L144.

Wolff, M. T., P. Hertz, K. S. Wood, P. S. Ray, and R. M. Bandyopadhyay (2002). Eclipse Timings of the Low-Mass X-Ray Binary EXO 0748−676. III. Orbital Period Jitter Observed with the Unconventional Stellar Aspect Experiment and the Rossi X-Ray Timing Explorer. ApJ 575, 384–396.

Woosley, S. E., A. Heger, A. Cumming, R. D. Hoffman, J. Pruet, T. Rauscher, J. L. Fisker, H. Schatz, B. A. Brown, and M. Wiescher (2004). Models for Type I X-Ray Bursts with Improved Nuclear Physics. ApJS 151, 75–102.

Zavlin, V. E. (2006). XMM-Newton Observations of Four Millisecond Pulsars. ApJ 638, 951–962.

Zhang, W., K. Jahoda, R. L. Kelley, T. E. Strohmayer, J. H. Swank, and S. N. Zhang (1998). Millisecond Oscillations in the Persistent and Bursting Flux of Aquila X-1 during an Outburst. ApJ 495, L9–L12.