

DETERMINING NEUTRON STAR MASSES AND RADII USING ENERGY-RESOLVED WAVEFORMS OF X-RAY BURST OSCILLATIONS

KA HO LO¹, M. COLEMAN MILLER^{2,3}, SUDIP BHATTACHARYYA⁴, AND FREDERICK K. LAMB^{1,5}

¹ Center for Theoretical Astrophysics and Department of Physics, University of Illinois at Urbana-Champaign,
1110 West Green Street, Urbana, IL 61801-3080, USA; fkf@illinois.edu

² Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, College Park, MD 20742-2421, USA

³ Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

⁴ Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research, Mumbai 400005, India

⁵ Department of Astronomy, University of Illinois at Urbana-Champaign, 1110 W. Green Street, Urbana, IL 61801-3074, USA

Received 2013 April 4; accepted 2013 August 13; published 2013 September 20

ABSTRACT

Simultaneous, precise measurements of the mass M and radius R of neutron stars can yield uniquely valuable information about the still uncertain properties of cold matter at several times the density of nuclear matter. One method that could be used to measure M and R is to analyze the energy-dependent waveforms of the X-ray flux oscillations seen during some thermonuclear bursts from some neutron stars. These oscillations are thought to be produced by X-ray emission from hotter regions on the surface of the star that are rotating at or near the spin frequency of the star. Here we explore how well M and R could be determined by generating and analyzing, using Bayesian techniques, synthetic energy-resolved X-ray data that we produce assuming a future space mission having 2–30 keV energy coverage and an effective area of 10 m², such as the proposed *Large Observatory for X-Ray Timing* or *Advanced X-Ray Timing Array* missions. We find that waveforms from hot spots within 10° of the rotation equator usually constrain both M and R with an uncertainty of about 10%, if there are 10⁶ total counts from the spot, whereas waveforms from spots within 20° of the rotation pole provide no useful constraints. The constraints we report can usually be achieved even if the burst oscillations vary with time and data from multiple bursts must be used to obtain 10⁶ counts from the hot spot. This is therefore a promising method to constrain M and R tightly enough to discriminate strongly between competing models of cold, high-density matter.

Key words: dense matter – equation of state – stars: neutron – X-rays: bursts

Online-only material: color figures

1. INTRODUCTION

The properties of matter at extremely high densities are among the most important currently unresolved questions in physics and astronomy. Neutron stars contain large quantities of cold matter at densities that are otherwise inaccessible. Studies of these stars can therefore help determine the properties of such matter. In particular, simultaneous measurements of the mass M and radius R of neutron stars could provide tight constraints on the equation of state of ultradense matter (see, e.g., Lattimer 2007; Lattimer & Prakash 2007; Read et al. 2009; Özel & Psaltis 2009; Hebeler et al. 2010).

Following the discovery of thermonuclear X-ray bursts from neutron stars in the mid-1970s (Grindlay et al. 1976; Lewin et al. 1976; for a review, see Lewin et al. 1993) and the discovery two decades later that some of these bursts produce X-ray flux oscillations at or near the star’s rotation frequency (Strohmayer et al. 1996; for a review, see Watts 2012), several approaches have been proposed for using observations of X-ray bursts to determine M and R .

The first attempts to constrain neutron star properties using X-ray bursts used their apparent peak luminosity and spectral temperature (van Paradijs 1979; Goldman 1979; Hoshi 1981; Marshall 1982). With improved understanding of the complexities of burst emission, this approach was developed into a method for measuring M and R (see Lewin et al. 1993). This method relies on six assumptions: (1) that the radius of the photosphere at “touchdown” (defined as the moment after the photosphere has reached its maximum inferred radius when

the “color” temperature T_c derived from fitting a Planck spectrum to the observed spectrum is highest) is the stellar radius; (2) that the emitting area at “touchdown” and during the cooling phase is the entire surface of the star; (3) that the flux seen at Earth at touchdown is the Eddington flux diluted by the distance to the burst source, i.e., $L_E/(4\pi d^2)$, where L_E is the Eddington luminosity and d is the distance to the star; (4) that the so-called “color correction factor” f_c (the ratio of T_c to the effective temperature T_{eff}) is accurately known—which requires knowing the chemical composition of the atmosphere and having an accurate stellar atmosphere spectral model—and that f_c remains constant in time during the burst; (5) that the distance d to the star is accurately known; and (6) that there is no other significant emission from the system.

This first method has been widely used to constrain M and R (see, e.g., van Paradijs 1982; Paczyński & Anderson 1986; van Paradijs et al. 1990). Tight constraints on M and R have recently been derived by using this method to analyze X-ray burst data obtained using the *Rossi X-Ray Timing Explorer (RXTE)*; Özel 2006; Özel et al. 2009, 2010; Güver et al. 2010a, 2010b). The derived constraints on M and R are substantially smaller than the uncertainties in the observed values of the input parameters because most of the observed values are discarded, since they would produce values of M and R that are complex numbers (Özel et al. 2010; Steiner et al. 2010; Suleimanov et al. 2011a). When this inconsistency is addressed and systematic errors are fully included, the uncertainties in M and R are likely to be substantially larger (see, e.g., Steiner et al. 2010; Miller et al. 2011; Suleimanov et al. 2011a; Miller et al. 2013).

A related but distinct method is to fit theoretical models derived from detailed burst atmosphere spectrum computations (Suleimanov et al. 2011b; Suleimanov et al. 2012) to an extensive portion of the “cooling track” in the color-correction factor versus X-ray flux plane (Suleimanov et al. 2011a). This method assumes (1) that burst spectra have been correctly computed using detailed stellar atmosphere models; (2) that the appropriate atmospheric model, including the chemical composition, can be determined by fitting to the cooling track; (3) that the emitting area is the entire surface of the star throughout the cooling period that is fitted; and assumptions (5) and (6) listed above. Testing whether the observed evolution of the blackbody normalization with flux is consistent with the theoretically predicted evolution for a passively cooling neutron star of constant radius provides a check on this approach. The model atmosphere spectra reported in Suleimanov et al. (2012) are in good agreement with other theoretical model atmosphere spectra, or the reasons for differences are understood (see Suleimanov et al. 2012), and they are also in good agreement with high-precision measurements of the spectra observed during a superburst from 4U 1820–30 (Miller et al. 2011; Miller et al. 2013). This method was applied to 4U 1724–307 and yielded a lower bound on the stellar radius of ~ 14 km, assuming the mass of this star is $< 2.3 M_{\odot}$ (Suleimanov et al. 2011a). In this application, the theoretical cooling track followed the observed cooling track only from $\sim 100\%$ of the Eddington flux at the stellar surface down to $\sim 50\%$. A concern with this method is that there is strong evidence that even during the decay phase of the superburst from 4U 1820–30, when—according to the detailed model atmosphere spectra—the flux at the stellar surface is substantially less than the Eddington flux, the emitting area decreases by $\sim 20\%$ and therefore is not equal to the area of the stellar surface (Miller et al. 2011; Miller et al. 2013).

A second approach to measuring M and R using X-ray burst observations is to fit detailed spectral models to high-precision measurements of X-ray burst spectra (Majczyna & Madej 2005; Miller et al. 2011; Miller et al. 2013). This approach assumes that the spectrum predicted by the atmospheric model is an accurate physical description of the observed spectrum. If it is, fitting a grid of spectral models to the observed X-ray spectra yields the radiation temperature T_{∞} in the locally comoving frame at the stellar surface as well as M , R , and the composition (Miller et al. 2011; Suleimanov et al. 2012). This approach can only be used if high-precision spectral data and accurate, high-precision spectral models are both available. At present, adequate data are available only for the 4U 1820–30 superburst (Miller et al. 2011; Miller et al. 2013).

A third approach for determining M and R using X-ray burst observations is to fit models of burst oscillation waveforms to observations of these waveforms (Strohmayer et al. 1997; Miller & Lamb 1998; see also Weinberg et al. 2001). As Miller & Lamb (1998) demonstrated, fitting energy-resolved waveform data can improve the constraints. Burst oscillations are thought to be produced by X-ray emission from a region on the surface of the star that is hotter than the rest of the surface and is rotating at or near the rotation frequency of the star (Strohmayer et al. 1996; for a review, see Watts 2012). Such a hotter region could be present either because a part of the stellar surface has retained more heat than other parts after thermonuclear burning has occurred or because a disturbance in the outer layers of the star, such as a surface normal mode, has made a localized region hotter. The observed amplitude of the waveform constrains the inclination of the hot spot relative to

the rotation axis, the observer’s inclination, and the compactness (M/R) of the star, the last primarily via general relativistic light-bending effects. The observed asymmetry and harmonic content of the waveform constrain the component of the velocity of the emitting region in the observer’s direction, primarily via special relativistic Doppler boosts and aberration. Because the rotation frequency of the emitting region is accurately known from the oscillation frequency, knowing the surface velocity constrains the stellar radius.

Nath et al. (2002) explored the constraints that could be derived on the compactness of 4U 1636–536 and 4U 1728–34 by analyzing *RXTE* observations of the bolometric flux oscillations that occur during the rise of X-ray bursts from these two neutron stars. They modeled the bolometric waveforms during the burst rise using small hot spots that expand linearly with time, neglected the relativistic Doppler shifts and aberration produced by the rotational motion of the hot spots, and assumed that the background was known. They found that they could not determine whether the oscillations are produced by a single spot or two antipodal spots, or constrain the hot spot and viewing geometry assuming either of these alternatives, using bolometric *RXTE* waveform data. The primary reason is that the changes in the waveform produced by changes in different parameters of the model are very similar. As a result of this degeneracy, they were unable to constrain the stellar compactness for models with a single hot spot, even if they assumed that the spot and the observer were known from other information to be in the rotation equator. They were able to obtain interesting upper bounds on the compactness for models with two antipodal hot spots and argued that it would be possible to simultaneously constrain the stellar compactness and the hot spot and viewing geometries with a count rate ~ 10 – 20 times higher than the *RXTE* rate. However, the 290 Hz subharmonic in the 4U 1636–536 waveform that appeared significant in one burst (Miller 1999; Strohmayer 2001), and suggested consideration of two antipodal spots, was subsequently found not to be significant in a more detailed analysis of additional data (Strohmayer 2001).

Bhattacharyya et al. (2005) investigated the constraints on the stellar compactness, hot spot properties, and system geometry that could be obtained by fitting model pulse profiles to an average energy-resolved oscillation profile produced by folding and stacking *RXTE* observations of 22 burst oscillations in three groups of bursts from XTE J1814–338. Stable oscillations occur throughout the bursts produced by this neutron star (Strohmayer et al. 2003). In generating their energy-resolved oscillation profile models, Bhattacharyya et al. assumed a single hot spot and an oscillation-phase-independent background chosen so that the model produces the observed total number of counts in each energy channel, when summed over oscillation phase. They included the relativistic Doppler shifts, aberration, and frame dragging produced by the rotation of the star, as well as gravitational redshift and light bending effects. In order to include frame dragging, Bhattacharyya et al. had to construct general relativistic stellar models numerically and therefore considered just two illustrative equations of state. They were able to achieve acceptable fits to the average oscillation profile using these equations of state, but obtained only weak constraints on the stellar compactness, hot spot properties, and system geometry.

Strohmayer (2004) explored the constraints on M and R that might be achieved by analyzing burst oscillation observations with much higher count rates than were achieved using *RXTE*. He did this by analyzing a synthetic bolometric waveform

similar to a waveform seen in 4U 1636–536. Strohmayer generated his synthetic and model waveforms using a code that is similar to the one used by Nath et al. (2002) but included the relativistic Doppler shifts and aberration produced by the rotational motion of the hot spots. The synthetic observed waveform included a background commensurate with the pre-burst emission in some bursts from 4U 1636–536 ($\sim 5\%$ of the maximum flux during strong bursts). The background was then fitted for as part of the waveform fitting process (T. E. Strohmayer 2013, private communication).⁶ Strohmayer assumed that the spot and the observer are both in the rotation equator, which is the most favorable possible geometry, and made the strong assumptions that the hot spot size and location, the radiation spectrum and beaming pattern, and the observer’s inclination are known independently of the waveform, so that only M and R have to be estimated using the waveform data. These assumptions eliminate the strong degeneracies between M , R , and the other model parameters that greatly increase the uncertainties in more realistic situations. Based on these assumptions, Strohmayer concluded that an X-ray timing mission with a collecting area ~ 10 times larger than *RXTE* would be able to determine M and R to within a few percent, thereby placing interesting constraints on the equation of state of neutron star matter.

Muno et al. (2003) examined averages of the energy-resolved flux oscillations observed from several X-ray burst sources using *RXTE*. They found that folded oscillations observed in their higher energy bands arrived later than those in their lower energy bands, although the energy dependence varied significantly with epoch and source. Their analysis of folded and averaged 4U 1636–536 oscillation profiles showed the clearest evidence for such a trend, which would be inconsistent with a simple rotating spot model of burst oscillations. Recently, Artigue et al. (2013) have analyzed the same 4U 1636–536 data in much more detail and find that the data are entirely consistent with a simple rotating spot model, although the parameters of this model are poorly constrained by these data.

New mission concepts are now being proposed that would provide much larger count rates than previous missions. Plans for the proposed *Large Observatory for X-Ray Timing (LOFT)* include a detector with an effective area ~ 10 m² at 8 keV and 2–30 keV energy coverage (Feroci et al. 2010; Mignani et al. 2012; Del Monte et al. 2012), more than an order of magnitude larger than the area of the *RXTE* Proportional Counter Array (PCA). The *Advanced X-Ray Timing Array (AXTAR)* concept includes a detector with an effective area greater than 3 m² and 2–50 keV energy coverage (Chakrabarty et al. 2008; Ray et al. 2011). One of the main motivations for these missions is to determine the properties of neutron stars with much higher precision than has been possible previously, by analyzing high-precision observations of the waveforms of burst oscillations. Although we focus here on exploring the constraints on M and

R that can be obtained by analyzing the waveforms of X-ray burst oscillations, our results are equally relevant for measuring M and R using the waveforms produced by X-ray emission from the heated polar caps of isolated rotation-powered millisecond pulsars (see, e.g., Bogdanov et al. 2007, 2008), which is the goal of the proposed NICER mission (Gendreau et al. 2012). This approach was originally developed to derive constraints on M and R by analyzing the accretion-powered X-ray oscillations produced by the rotation of some accreting neutron stars with millisecond spin periods (Poutanen & Gierliński 2003; see also Poutanen & Gierliński 2004; Poutanen & Beloborodov 2006; Cadeau et al. 2007; Leahy et al. 2008, 2009). Using methods similar to ours, Bogdanov (2013) has recently derived a lower limit of 11.1 km on the radius of the isolated pulsar PSR J0437–4715, assuming its mass is $1.76 M_{\odot}$ and that systematic errors can be neglected.

In this paper, we explore the constraints on M and R that could be derived by analyzing energy-resolved burst oscillation waveforms obtained using a future, satellite-borne detector with 2–30 keV energy coverage and an effective area 10 to 20 times larger than the *RXTE* PCA. We do this by first generating energy-dependent synthetic observed waveforms for a variety of neutron star and hot spot properties. We then use a Bayesian approach and Markov chain Monte Carlo (MCMC) sampling methods to determine the constraints on M and R that can be obtained by analyzing these synthetic observed waveforms. Specifically, we determine the joint posterior probability distribution of the parameters in our waveform model, given the synthetic waveform of interest, use this distribution to determine the joint posterior distribution of M and R by marginalizing the other parameters in the waveform model, and then use the joint distribution of M and R to determine the most probable values of M and R and Bayesian confidence regions in the M – R plane.

Our analysis applies to any waveforms produced by emission from a single region of hotter gas that is rotating about the star. Such a region could be present either because a part of the stellar surface has retained more heat than other parts after thermonuclear burning has occurred or because a disturbance in the outer layers of the star, such as a surface normal mode, has made a localized region hotter (see Watts 2012). Oscillations with the amplitudes ($\sim 5\%$ – 10%) required to obtain significant constraints on M and R are very probably produced by a single hotter region (see Lamb et al. 2009a).

We explore how the sizes and positions of these confidence regions depend on the inclinations of the hot spot and the observer and the background count rate. We also explore the effect on the confidence regions if the distance to the star or the inclination of the observer are known from other measurements, if a resonance scattering line is observed in the burst oscillation spectrum, or if the properties of the background are independently known. Finally, we explore the effects of deviations in the actual shape of the hot spot, radiation beaming pattern, and spectrum from those assumed in the fitted model. We assume that about 10^6 counts are collected from the hot spot and that all sources of background contribute about 0.3×10^6 , 10^6 , or 9×10^6 counts. Our treatment of the background is very conservative, in the sense that we usually make no assumptions about the magnitude or spectrum of the background. We do not even assume that the background is constant, only that it does not vary at frequencies commensurate with the burst oscillation frequency.

We find that the uncertainties in the measured values of M and R depend strongly on the inclination of the hot spot relative

⁶ We note that attempting to remove the background by subtracting the pre-burst count rate assumes that the pre-burst background, which is presumably produced at least in part by accretion onto the star, persists unchanged throughout the burst. This introduces possible systematic errors because luminous bursts are expected to alter the accretion flow significantly (see, e.g., Miller & Lamb 1996; Worpel et al. 2013). Even if the pre-burst background persists unchanged throughout the burst, subtracting it from the count rate during the burst, rather than including the background as a component of the model, incorrectly neglects the fluctuation in the number of counts produced by the background and the corresponding uncertainties in the model parameters. This conceptual error is built into the current version of the XSPEC analysis package, so all data analysis done using XSPEC makes this error.

to the stellar rotation axis. If the hot spot is within 10° of the rotation equator, both M and R can usually be determined with an uncertainty of about 10%. If instead the spot is within 20° of the rotation pole, the uncertainties are so large that waveform measurements alone provide no useful constraints on M and R . The uncertainties in M and R are affected little by background count rates less than or comparable to the count rate from the hot spot, but become significantly larger for higher background count rates. The precisions we report here can potentially be achieved even if the burst oscillations vary with time and data from multiple bursts must be combined to obtain 10^6 counts from the hot spot.

Observation of an identifiable atomic line in the hot-spot emission always tightly constrains M/R ; it can also tightly constrain M and R individually, if the spot is within about 30° of the rotation equator. Independent knowledge of the observer’s inclination can greatly reduce the uncertainties, as can independent information about the background. Knowledge of the star’s distance can also help, but not as much.

Modest deviations of the actual spectrum from that assumed in the fitted model have little effect on the accuracy or uncertainty of M and R estimates. Large deviations of the actual radiation beaming pattern from the pattern assumed in the waveform model can increase the uncertainties of M and R measurements substantially. In some cases, but not always, large deviations of the actual shape of the hot spot from the circular shape assumed in the waveform model increase the uncertainties of M and R estimates and bias them by moderate amounts. The physical conditions that produce tight constraints on M and R (relatively small spots far from the rotation pole) are the conditions in which the shape of the spot is unimportant.

Our results show that if a sufficient number of burst oscillations produced by hot spots at high inclinations are observed using the next generation of X-ray timing satellites, one can constrain M and R tightly enough to discriminate strongly between competing models of cold, high-density matter.

The remainder of this paper is organized as follows. In Section 2, we outline our approach to generating synthetic observed waveforms and producing the model waveforms that we fit to the synthetic waveform data. In Section 3, we describe the computational methods we use to produce synthetic waveform data and construct model waveforms, and the MCMC computational methods we use to determine the posterior probability distributions of the parameters in the model, given a synthetic waveform. In Section 4, we describe our results and in Section 5, we summarize our conclusions. In Appendix A, we summarize the suite of test problems and solutions that we use to validate our waveform and MCMC codes. In Appendix B, we discuss the constraints on system parameters that can be obtained by jointly fitting many segments of waveform data, from a single burst or from multiple bursts.

2. APPROACH

2.1. Generation of “Observed” and Model Waveforms

In the present work, we fit energy-resolved model waveforms to energy-resolved synthetic “observed” waveforms, using a Bayesian approach. For both waveforms, we assume that the oscillation is produced by emission from a single, uniformly emitting, hotter region on the stellar surface. Such a region could be present either because a part of the stellar surface has retained more heat than other parts after thermonuclear burning

occurred or because a disturbance in the outer layers of the star has made a localized region hotter (see Watts 2012).

We consider synthetic observed waveforms generated by circular and elongated hot spots (see Section 4.1.2), but to reduce the number of fitted parameters our model waveforms assume the hot spot is circular (see Section 4.1.1). For simplicity, we usually assume the hotter region emits radiation with a Planck spectrum and 100% efficiency but with the beaming pattern appropriate for an electron scattering atmosphere, assumptions that are mutually inconsistent. This inconsistency can be avoided when fitting real data by including an appropriate color factor in the emission model used to construct fitted waveforms. We also assume that the radiation propagating from the emitting area on the stellar surface reaches the observer without interacting with any ambient or intervening matter, but we do include background counts to illustrate the effects of possible background emission from the stellar surface, an accretion disk, and other sources in the field, as well as the instrumental background.

In generating synthetic and model waveforms, we assume that the hotter region has a constant size and shape, is located at a fixed stellar rotational latitude, and rotates at a constant frequency. These assumptions are not as restrictive as they might at first appear. The reason is that constraints on M and R similar to those obtained for waveforms that satisfy these assumptions can usually be obtained by appropriately analyzing data from a single burst or from multiple bursts, provided the data contains the same number of counts as assumed in our analysis. In particular, constraints on M and R similar to those found here can usually be achieved even if the oscillation frequency and other physical parameters vary during the burst or from burst to burst (see Section 4.2.7 and Appendix B).

In constructing the synthetic waveforms, we include a constant background component. This component is a catch-all for all counts not produced by radiation from the hot spot. These counts could be produced by emission from unassociated sources in the field, the accretion disk, the non-spot portion of the star, instrumental backgrounds, or any combination of these. For simplicity, we model this background by adding emission from the entire stellar surface with the beaming pattern expected for an electron scattering atmosphere and a spectrum having the shape of a Planck spectrum with a temperature lower than the temperature of the hot spot. We normalize the background spectrum to achieve the desired number of background counts. The number of counts contributed by this emission is important, but not their detailed properties.

In fitting model waveforms to the synthetic waveform, we treat the background component in the model waveform very conservatively, in the sense that we usually make no assumptions about its magnitude or spectrum. We do not even assume that the background is constant, only that it does not vary at frequencies commensurate with the hot spot rotation frequency. Any prior knowledge of the properties of the background can be used to restrict the background model and usually tightens the constraints on the values of the system parameters that can be derived from waveform observations.

We compute the time- and energy-resolved waveforms that would be seen by a distant observer using the Schwarzschild plus Doppler (S+D) approximation introduced by Miller & Lamb (1998). The S+D approximation treats exactly all special relativistic effects (such as relativistic Doppler boosts and aberration) produced by the rotational motion of the hot spot, but treats the star as spherical and uses the Schwarzschild spacetime to compute the general relativistic redshift, trace the

propagation of light from the stellar surface to the observer, and calculate light travel-time effects. The S+D approximation does not include the effects of stellar oblateness or frame dragging. However, for the stars considered here, and indeed for any stars that do not both rotate rapidly and have very low compactness, the effects of stellar oblateness and frame dragging are minimal and are negligible compared to the uncertainties in the X-ray emission (see Cadeau et al. 2007).

The S+D approximation allows us to exploit the spherical symmetry of the spacetime to reduce substantially the number of integrations needed to trace rays from the emitting area to the observer. As a result, this code is much faster than codes that take into account the oblateness of the star, approximate the star's exterior spacetime using the Kerr spacetime, or use numerically computed stars and exterior spacetimes. This speed is essential, because finding the high-probability region of the parameter space of the waveform model and then determining the posterior probability density in this region with some precision requires accurate computation of the likelihood at a large number of points in a high-dimensional parameter space, and therefore requires many ray tracings.

The waveform code we use is based on the code used and validated in the analysis of accretion-powered millisecond X-ray pulsar waveforms by Lamb et al. (2009a, 2009b). We report the results of further code validation tests in Appendix A.

To determine the constraints on M and R that can be derived from burst oscillation waveforms, we compute the posterior probability distribution of all the parameters in the waveform model, for each of a variety of synthetic observed waveform data sets, using a standard Bayesian approach. We use an MCMC algorithm to sample the parameters of the waveform model and compute the likelihood of each set of parameters, given the synthetic waveform being considered.

We construct each synthetic observed waveform data set by Poisson sampling the counts in each phase and energy bin of each synthetic observed waveform that we computed using the S+D approximation, to mimic the statistical fluctuations that would be present in actual data.

We generate the model waveforms that we compare with the synthetic waveform data using the same code that we use to generate the synthetic waveforms. The model waveforms include a possible background component with an arbitrary energy spectrum.

2.2. Burst Rise or Burst Tail?

Previous interest in using X-ray burst oscillations to determine the masses and radii of neutron stars has focused largely on using oscillations observed during the rise of bursts (see, e.g., Nath et al. 2002; Strohmayer 2004), primarily because of the large fractional modulation observed during the rise of some bursts and the evidence that emission early in the burst rise comes from a small, hotter region on the stellar surface. Although the oscillations observed during burst tails usually have smaller fractional amplitudes, tail oscillations usually last much longer than the oscillations during burst rises. We therefore consider the merits of using oscillations observed during burst tails as well as burst rises.

Key factors to consider in evaluating the relative merits of using burst rise oscillations and burst tail oscillations include (1) the total number of oscillating counts, (2) the total number of counts collected during the observation, (3) the information contained in these counts, including any information encoded in the constant X-ray flux from the star, (4) the variation of

the oscillation waveform with time during the measurement interval, and (5) the complexity of the emission pattern that produces the burst oscillations.

Here we show that analyses of oscillations observed during burst tails may provide uncertainties in M and R estimates comparable to or possibly even smaller than the uncertainties provided by analyses of oscillations observed during burst rises.

2.2.1. Statistical Uncertainties in Parameter Estimates

Consider first the uncertainties in M and R estimates produced by the fluctuations in the observed waveform caused by photon counting noise. We expect the fractional uncertainties in M and R estimates to decrease as the ratio \mathcal{R} of the number of modulated counts to the fluctuation in the number of counts increases. For the burst oscillation waveforms, we use as a figure of merit the quantity

$$\mathcal{R} \equiv N_{\text{osc}}/\sqrt{N_{\text{tot}}} = 1.4 f_{\text{rms}} \sqrt{N_{\text{tot}}}, \quad (1)$$

where N_{osc} is the number of counts in the oscillating component of the waveform, defined as the integral of the semi-amplitude of the oscillating count rate over the duration of the data segment; N_{tot} is the integral of the total count rate, including the background count rate, over the duration of the data segment; and f_{rms} is the fractional rms amplitude of the oscillation during the data segment. Our Bayesian analysis (see Section 4.2.1) shows that the fractional sizes of the confidence regions in the M - R plane decrease with increasing \mathcal{R} , approximately as \mathcal{R}^{-1} . The value of \mathcal{R} is therefore a useful figure of merit when comparing different data sets.

We illustrate the implications of this result for the relative merits of analyzing oscillations during burst rises and tails by using burst oscillation amplitudes observed with the *RXTE* PCA and scaling the *RXTE* count rates to estimate the count rates that would be observed using a next-generation large-area X-ray timing instrument like those mentioned in Section 1. The amplitudes, durations, and count rates of burst oscillations vary both from star to star and from burst to burst in a single star. We consider first the properties of the oscillations during the rise and tail of a well-observed X-ray burst from 4U 1636–536 (Strohmayer et al. 1998; see also Nath et al. 2002). We then consider the implications of the atlas of burst oscillations from 4U 1728–34 compiled by van Straaten et al. (2001). Next, we consider the implications of the 59 oscillation trains from 8 neutron stars that were studied by Munro et al. (2002b). Finally, we discuss the information on burst oscillation amplitudes and their occurrences during burst rises and tails in the comprehensive review by Galloway et al. (2008).

Oscillations during a burst from 4U 1636–536. During the first 1/16 s of the burst, the average total count rate (including a background assumed equal to the ~ 2000 counts s^{-1} pre-burst background) was ~ 4500 counts s^{-1} , while the average oscillation semi-amplitude was ~ 2000 counts s^{-1} . Thus, during the first 1/16 s of the burst, a total of ~ 280 counts was collected, including ~ 125 modulated counts. Our postulated next-generation X-ray timing instrument would therefore collect ~ 5600 total counts, including ~ 2500 modulated counts, yielding an \mathcal{R} -value ~ 33 when data from this part of the burst rise is used.

During the first 1/4 s of the burst, which includes most of the burst rise, the average total count rate (including the assumed background) was ~ 8000 counts s^{-1} , while the average oscillation semi-amplitude was ~ 2000 counts s^{-1} . Thus, during the first 1/4 s of the burst, a total of ~ 2000 counts was

collected, including ~ 500 modulated counts. Our postulated next-generation X-ray timing instrument would therefore collect $\sim 40,000$ total counts, including $\sim 10,000$ modulated counts, yielding an \mathcal{R} -value ~ 50 when all the data from the burst rise is used.

These results indicate that the 1/4 s data set would provide constraints similar to or even better than the 1/16 s data set, even though the average fractional amplitude of the oscillation was much smaller during the longer interval, because the total number of counts during the longer interval was much greater.

During 2 s in the tail of the burst, oscillations were observed with a semi-amplitude of 5.7%. The average *RXTE* PCA count rate during this 2 s interval was $\sim 25,000$ counts s^{-1} . Our postulated next-generation X-ray timing instrument would therefore collect $\sim 10^6$ total counts during this 2 s interval, yielding an \mathcal{R} -value ~ 55 , comparable to the \mathcal{R} -value obtained for the burst rise. \mathcal{R} scales as the square root of the number of counts, so combining data from the rise or tail of 25 bursts with properties like this burst from 4U 1636–536 would yield an \mathcal{R} -value ~ 250 .

Burst oscillations from 4U 1728–34. What do we expect from analyses of oscillations during the X-ray bursts from other neutron stars? van Straaten et al. (2001) compiled an atlas of the properties of 21 bursts with oscillations from 4U 1728–34. The fractional rms amplitudes of the oscillations that occur during burst rises range from $\sim 10\%$ up to $\sim 25\%$, but the count rates are low during some of the rises, the highest amplitudes typically last less than 1 s, and in some cases no oscillations were detected during the burst rise. The fractional rms amplitudes of the oscillations that occur at and after the peak range from 2% up to 15%. During the tail of one burst, the fractional rms amplitude was 8%–10% for 13 s and was $\sim 15\%$ for ~ 4 s from the peak well into the tail of two others. The fractional rms amplitudes in the tails of most bursts are $\sim 5\%$ and typically last ~ 5 s. The *RXTE* PCA count rates at the peaks of the 4U 1728–34 bursts are $\sim 30,000$ counts s^{-1} or sometimes even a little higher, and are typically $\sim 15,000$ – $20,000$ counts s^{-1} when tail oscillations are present. Assuming an rms amplitude of 5%, a duration of 5 s, and a mean count rate of 16,000 counts s^{-1} yields an \mathcal{R} -value ~ 20 for the *RXTE* PCA or ~ 88 for our postulated next-generation X-ray timing instrument. Combining data from 25 such oscillation trains would yield an \mathcal{R} -value ~ 440 .

Typical properties of burst oscillations from eight neutron stars. Munro et al. (2002b) searched the *RXTE* data in the public archive for 8 neutron stars that are known to produce burst oscillations and found 59 oscillation trains that persisted for more than 2 s, in a total of 159 X-ray bursts. The majority of these oscillation trains were observed during burst tails. Munro et al. reported that the typical rms amplitude of these oscillations was 5% and that the total number of counts during a typical oscillation train was 68,000, yielding a typical \mathcal{R} -value ~ 18 . Our postulated next-generation X-ray timing instrument would therefore collect $\sim 10^6$ counts during such oscillation trains, yielding an \mathcal{R} -value ~ 82 . Combining data from 25 such oscillation trains from a single star could yield an \mathcal{R} -value ~ 400 .

*Information on the occurrence and properties of burst oscillations from all *RXTE* observations.* Galloway et al. (2008) report that, except for a few bursts, the amplitudes of the oscillations seen in 4U 1702–429 and 4U 1728–34 during burst peaks and tails are comparable to or even greater than the amplitudes of the oscillations seen during burst rises. Only in 4U 1636–536 are the oscillation amplitudes observed during the burst rises typically much larger than the amplitudes observed during the

burst tails. According to Galloway et al. (2008), the stars that are most prolific in producing burst oscillations produce many more bursts with oscillations detectable only during the peak and the tail or only during the tail than bursts with oscillations detectable only during the rise, in addition to bursts in which oscillations are detectable during both the rise and the tail. They report, for example, that 2 bursts have been seen from 4U 1702–429 with oscillations only during the rise whereas 17 bursts have been seen with oscillations only during the tail, or the peak and the tail. Similarly, 5 bursts have been seen from 4U 1728–34 with oscillations only during the rise whereas 20 bursts have been seen with oscillations only during the tail, or the peak and the tail. In 4U 1636–536, only 5 bursts have been observed with oscillations only during the rise, whereas 20 bursts have been observed with oscillations only during the tail, or the peak and the tail.

In these estimates we have assumed that the unmodulated (constant in time) component of the X-ray flux time series contains no information about M and R , i.e., that only the modulated component of the waveform contains such information. If instead the unmodulated component contains some information about the system (such as the compactness of the star, which tends to increase the unmodulated count rate due to increased gravitational lensing) and one can extract this information, then using the data collected during the burst tail could be more useful than suggested by the preceding estimates.

To see this, suppose first that the X-ray waveform is constant in time (i.e., unmodulated) and that the magnitude of this constant count rate is important. Then the fractional uncertainties in estimates of M and R will depend only on the total number of counts. Now suppose that in addition to this same unmodulated component there is an infinitesimal modulated component. Clearly, this modulated component will not add much extra information about the system. This case may, in fact, apply to observations of the tails of some bursts. In this case, considering only the oscillations would underestimate our ability to constrain the model parameters using data from the burst tail compared to using data from the burst rise.

These results suggest that analyzing data collected from burst tails may provide constraints comparable to, or possibly even better than, those obtained by analyzing data collected from burst rises and that analyzing burst-tail data should therefore be considered. They also indicate that one does not need to rely on one or two particularly favorable stars to obtain constraints on M and R .

2.2.2. Other Considerations

In addition to the figure of merit \mathcal{R} of the data set, other important considerations for assessing the relative merits of burst rise and burst tail oscillations include (1) the presence of detectable higher harmonics in the waveform, (2) the time-dependence of the waveform, and (3) the size and possibly complex shape of the emitting region that produces the oscillations. Here we consider briefly these three factors.

Harmonic content of the waveform. The amplitudes of overtones of the fundamental oscillation frequency in the observed waveform play an important role in determining M and R by constraining the component of the velocity of the emitting region in the observer's direction, primarily via special relativistic Doppler boosts and aberration. Because the rotation frequency of the emitting region is accurately known from the oscillation frequency, knowing the surface velocity constrains the stellar

radius. Precise constraints on M and R can be derived only if at least one overtone is detected in the waveform.

Overtones have been detected in most oscillation trains observed during X-ray bursts from persistent accretion-powered X-ray pulsars, although not necessarily with the same strength as those seen in the accretion-powered oscillations from these stars (Chakrabarty et al. 2003; Strohmayer et al. 2003; Watts et al. 2005; Altamirano et al. 2010; for a review, see Watts 2012). Burst oscillations from stars in which X-ray oscillations have been detected only during X-ray bursts (Muno et al. 2002b) and from the intermittent accretion-powered X-ray pulsars (Muno et al. 2002b; Watts et al. 2009) tend not to have detectable overtones, except possibly for oscillations seen during the rising phases of some bursts from 4U 1636–536 (see Bhattacharyya & Strohmayer 2005).

Time-dependence of oscillation amplitudes and frequencies. Often, but not always, the amplitude of the burst oscillation varies during the burst. The variation is typically faster during the burst rise. For example, in the 4U 1636–536 burst studied by Strohmayer et al. 1998, the rms amplitude quickly reached $\sim 50\%$ during a 62 ms interval close to the onset of the burst, but averaged only $\sim 20\%$ during the first 0.25 s and $\sim 4\%$ during the burst tail. Bhattacharyya & Strohmayer (2005) reported a decrease in the rms amplitude of the fundamental from 40% to 15% within 0.75 s during the rise of a burst from 4U 1636–536. Chakrabarty et al. (2003) and Bhattacharyya & Strohmayer (2006) found that the amplitude of the oscillations during the rises of bursts from the accretion-powered millisecond X-ray pulsar SAX J1808.4–3658 can decrease by a factor ~ 5 within ~ 0.15 s. In the atlas of burst oscillations from 4U 1728–34 assembled by van Straaten et al. (2001), the rms amplitude varied by a factor ~ 2 – 5 within ~ 1 s during the rising portion of several bursts, but remained constant to within $\sim 30\%$ over intervals as long as 6–7 s during the rises and tails of others.

The oscillation frequency remains nearly constant during the rises and tails of some bursts, but increased by 3 Hz within ~ 1 s during the rise of a burst from 4U 1728–34 (van Straaten et al. 2001), by ~ 5 Hz within ~ 0.1 s during the rise of bursts from the accretion-powered millisecond X-ray pulsar SAX J1808.4–3658 (Chakrabarty et al. 2003; Bhattacharyya & Strohmayer 2006), and by 5 Hz within ~ 5 s during the tail of a burst from X1658–298 (Wijnands et al. 2001). Bhattacharyya & Strohmayer (2005) investigated 23 bursts from 4U 1636–536 that showed at least a moderately strong oscillation during the burst rise and found significant frequency evolution in only four. In all of these four bursts, the frequency increased by ~ 2 – 4 Hz over ~ 2 s. van Straaten et al. (2001) studied 21 bursts from 4U 1728–34 and found 5 burst rise oscillations that showed little frequency evolution. They also found five tail oscillations that showed little or no frequency evolution. The frequencies of five other burst rise oscillations increased by ~ 1 – 3 Hz over ~ 1 – 2 s, while the frequencies of seven other burst tail oscillations increased by ~ 1 – 2 Hz over 5–10 s. Chakrabarty et al. (2003) and Bhattacharyya & Strohmayer (2006) showed that the frequency of the oscillations during the rises of bursts from SAX J1808.4–3658 can increase by ~ 5 Hz within ~ 0.1 s and can also decrease by ~ 5 Hz within ~ 0.15 s.

As noted above, Muno et al. (2002a) found 59 oscillation trains that persisted for more than 2 s, the majority during burst tails. They found upward frequency drifts ranging from 0.4% in KS 1731–260 to 1.2% in 4U 1702–429. They also found three clear examples of downward drifts in frequency, namely 4U 1636–536 (see also Strohmayer 1999; Miller 2000),

KS 1731–260 (see also Muno et al. 2000), and 4U 1728–34. In all three cases, the frequency decrease was ~ 1 Hz. In KS 1731–260, the frequency decrease occurred during the peak of the burst, whereas in 4U 1728–34 and 4U 1636–536 the frequency decrease occurred in the tail. These frequency decreases occurred over ~ 2 – 4 s. Muno et al. found that 30% of the oscillation trains were inconsistent with a smooth model of frequency evolution. In particular, they found what appear to be discrete jumps of about 0.1 cycle in the phase of the oscillation from 7 s to 10 s into one burst from 4U 1636–536, and possible evidence for the simultaneous presence of two oscillations differing in frequency by about 1 Hz during the first 1–2 s of another burst from 4U 1636–536.

Clearly, the time dependence of burst oscillation amplitudes and frequencies varies greatly from star to star and from burst to burst in a given star, and large changes in the oscillation amplitude are common during the rising portion of bursts. Significant changes in the oscillation frequency are seen during the rises and tails of some bursts. Hence, in order to utilize all strong bursts, waveform models and procedures for fitting oscillations during the onset of bursts must be able to handle time-dependent waveforms.

Analyzing time-dependent waveforms. We have investigated the problem of analyzing burst oscillation waveforms that change with time (see Section 4.2.7 and Appendix B). If the oscillation frequency is changing but the change can be modeled accurately enough to maintain the correct oscillation phase when folding successive periods of the oscillation, then the constraints on M and R that can be obtained by analyzing the resulting folded waveform will be nearly the same as those that could be obtained by analyzing a similar waveform with a fixed oscillation frequency and the same number of counts. If the oscillation frequency varies too rapidly or irregularly during the burst rise or tail to be described accurately by a simple frequency model, the full burst oscillation data set can be divided into smaller time segments and analyzed using standard Bayesian techniques. This approach can also be used if other physical properties of the system, such as the size and inclination of the emitting region, vary significantly. The computational burden of this kind of analysis increases only linearly with the number of segments, so this is not a major problem. If, however, a waveform changes so rapidly that it must be divided into so many segments that the constraints derived from each segment are minimal, the constraints obtained by combining the results from the many individual segments can become unreliable. If the count rate is large enough, variations of the burst oscillation waveform on timescales shorter than the burst rise or burst tail (but substantially longer than the burst oscillation period) are not an insurmountable analysis problem.

Complexity of the emission pattern. Previous modeling of the waveforms of accretion-powered millisecond X-ray pulsars (see, e.g., Lamb et al. 2009a, 2009b) showed that the size and shape of the emitting region have only a weak effect on the amplitude and shape of the waveform, unless the emitting region is very large (angular radius $\gtrsim 50^\circ$ for inclinations $\gtrsim 60^\circ$). The hotter emitting regions that produce oscillations early in the rise of a burst are thought to cover a small fraction of the stellar surface (Watts 2012). In this case, the detailed shape of the emitting region does not significantly affect the properties of the waveform. Oscillations observed late in the rise of a burst or in burst tails are thought to be produced by larger emitting regions (Watts 2012), but these regions may still be small enough that their detailed shape does not significantly

affect the properties of the waveform. Our fits of waveform models that assume a circular emitting region to synthetic observed waveforms produced by large regions elongated in the east–west and north–south directions suggest that if the emitting region is large and distorted, this will increase the uncertainties in estimates of M and R , although not necessarily by a very large amount, but will not bias these estimates significantly (see Section 4.2.6). Our experience in fitting waveforms suggests that results similar to those we report here are likely to be achievable, even if there are moderate temperature variations across the hot spot.

3. COMPUTATIONAL METHODS

In this section, we describe the computational methods we use to determine the accuracy and precision with which M and R could be determined using the energy- and time-resolved burst oscillation waveform data that could be obtained by future large-area X-ray timing missions. We apply standard Bayesian inference methods to compute the best-fit values and confidence intervals for the parameters M and R in our waveform model, for a variety of synthetic observed waveforms of interest. We first explain how we compute the oscillation waveforms that we use. We then discuss our Bayesian analysis and sampling methods. Finally, we describe how we integrate over uninteresting parameters to determine the joint posterior distribution of M and R and determine their best-fit values and confidence regions.

3.1. Waveform Computation

The burst oscillation waveform model that we use here has eight parameters. The neutron star is described by its gravitational mass M and circumferential radius R . We represent the emitting area as a circular, uniformly radiating hot spot on the stellar surface, with an angular radius $\Delta\theta_{\text{spot}}$. We assume that the center of the spot is inclined at an angle θ_{spot} relative to the star’s rotation axis and that the spot rotates uniformly around the rotation axis of the star with a frequency ν_{rot} . We also assume that the spot emits radiation that has a blackbody spectrum with a temperature T_{co} when measured in an inertial frame at the stellar surface that is momentarily comoving with the surface. We assume that the values of these parameters are constant in time; the waveform is then perfectly periodic. Finally, we assume that the observer views the system from a distance d , at an inclination θ_{obs} relative to the star’s rotation axis.

We use as our global coordinate system Schwarzschild coordinates (r, θ, φ, t) centered on the star, with $\theta = 0$ aligned with the rotation axis and $\varphi = 0$ in the plane containing the rotation axis and the observer. We choose the zero of the Schwarzschild time coordinate t so that a light pulse that propagates radially from a point on the stellar surface immediately below the observer (i.e., at $\theta = \theta_{\text{obs}}$ and $\varphi = 0$) arrives at the observer at $t = 0$.

We use our waveform code to compute the phase- and energy-resolved photon number flux that would be seen by a distant observer during a single rotation of the hot spot. This code is based on the code we used previously to compute the waveforms of accreting millisecond X-ray pulsars (Lamb et al. 2009a, 2009b) and uses the Schwarzschild plus Doppler (S+D) approximation (Miller & Lamb 1998; see Section 2.1). In this approximation, the exterior spacetime is specified completely by the stellar compactness, GM/Rc^2 (hereafter, for conciseness, we write M/R for GM/Rc^2). The computational speed made

possible by this approximation makes it practical to compute the required likelihood distributions of the model parameters on a large, modern computer cluster.

It is convenient to specify the radiation from a point on the stellar surface in a local inertial frame located at the surface and momentarily comoving with it. We assume that the beaming pattern of the radiation emitted from the stellar surface is axisymmetric about the normal to the surface, when measured in the comoving frame. The specific intensity I'_0 in the comoving frame at the stellar surface can then be expressed as $I'_0(E'_0, \alpha')$, where E'_0 and α' are the photon energy and the angle between the photon direction and the normal to the surface measured in the comoving frame at the stellar surface. We assume that $I'_0(E'_0, \alpha')$ can be written as the product of a beaming function $g'(\alpha')$ and a spectral function $f'(E'_0)$, i.e.,

$$I'_0(E'_0, \alpha') = g'(\alpha')f'(E'_0). \quad (2)$$

In the present work, we usually consider the waveforms produced by the beaming pattern expected for emission from an electron scattering atmosphere, but we sometimes consider isotropic beaming. For the former beaming function, we use the quadratic expression

$$g'(\alpha') = a + b \cos \alpha' + c \cos^2 \alpha', \quad (3)$$

with

$$a = 0.42822, b = 0.92236, \text{ and } c = -0.085751. \quad (4)$$

This expression is a least-squares fit to the beaming function for emission from a uniform, semi-infinite, Thomson scattering atmosphere, taking polarization into account, and agrees with the actual beaming function (Chandrasekhar 1960, Table XXIV) to better than 1% for $0.02 \leq \cos \alpha' \leq 1$.

In the present work, we usually set $f'(E'_0)$ equal to the Planck function $B(E'_0, T')$, where T' is the radiation color temperature at the stellar surface, as measured in the comoving frame. We expect the actual spectral function to have a shape similar to, but not exactly the same as, a Planck spectrum and an efficiency $\sim 20\%$, rather than 100% (see, e.g., Suleimanov et al. 2012). Using a color factor to produce the appropriate lower efficiency would be important in analyzing real data.

In the few cases where we consider the presence in the spectrum of an atomic scattering line, we model the line by multiplying the continuum spectrum by the transmission factor $\exp[-\tau(E - E_c)]$, where E_c is the centroid energy of the line and $\tau(x)$ is a Gaussian profile having a maximum value $\tau(0)$ and a specified FWHM.

The specific intensity $I'_0(E'_0, \alpha')$ measured in the comoving frame at the stellar surface can be converted to the specific intensity $I_0(E_0, \alpha)$ measured in the static frame at the stellar surface using the invariance of $I(E)/E^3$ under a Lorentz boost. The result is

$$I_0(E_0, \alpha) = I'_0(E'_0, \alpha')\gamma^{-3}[1 - (\mathbf{v}/c) \cdot \hat{\mathbf{k}}]^{-3}, \quad (5)$$

where E_0 is the photon energy, α is the angle between the unit vector $\hat{\mathbf{k}}$ in the direction of the light ray and the normal to the surface, and \mathbf{v} is the linear velocity of the stellar surface at the point of emission, all measured in the static frame, and $\gamma = [1 - (v/c)^2]^{-1/2}$. The quantities E_0 and α are related to the corresponding quantities E'_0 and α' measured in the comoving frame by

$$E_0 = \delta E'_0 \quad (6)$$

and (Poutanen & Gierliński 2003)

$$\cos \alpha' = \delta \cos \alpha, \quad (7)$$

where

$$\delta = \frac{1}{\gamma[1 - (\mathbf{v}/c) \cdot \hat{\mathbf{k}}]} \quad (8)$$

is the Doppler factor. Here $\mathbf{v} \cdot \hat{\mathbf{k}} = v \cos \zeta$, where ζ is given by

$$\cos \zeta = \sin \alpha \sin \beta, \quad (9)$$

in terms of α and the angle β between $\hat{\mathbf{k}}$ and the direction to the star's rotation pole projected onto the plane tangent to the stellar surface at the point of the emission.

In order to compute the waveform seen by a distant observer located at an inclination θ_{obs} relative to the star's rotation axis, we divide the hot spot on the stellar surface into a fine grid in colatitude and longitude. We then consider the flux from an infinitesimal emitting area dA'_i in the comoving frame around each grid point. We transform this flux into the flux in the static frame using Equation (5) and the relation

$$dA_i = \delta dA'_i \quad (10)$$

between the infinitesimal area dA'_i measured in the comoving frame at the stellar surface and the infinitesimal area dA_i measured in the static frame at the stellar surface. Next we use spherical trigonometry and ray-tracing in the Schwarzschild spacetime to determine the direction of emission (α_i, β_i), measured in the static frame at the stellar surface, that is required for a light ray originating at a given grid point to reach the observer. We determine the angular separation ψ_i between the location (θ_i, φ_i) of the grid point and the direction to the observer using the spherical trigonometric relation

$$\cos \psi_i = \cos \theta_i \cos \theta_{\text{obs}} + \sin \theta_i \sin \theta_{\text{obs}} \cos \varphi_i \quad (11)$$

and then use the implicit relation

$$\psi_i \left(\alpha_i, \frac{M}{R} \right) = \int_0^1 \frac{\sin \alpha_i dx}{\sqrt{(1 - 2M/R) - (1 - 2Mx/R) x^2 \sin^2 \alpha_i}} \quad (12)$$

between ψ_i and α_i to determine α_i . We accurately evaluate the integral in Equation (12) using a combination of analytical and numerical methods (see Appendix A.1.1). Finally, we determine β_i using the spherical trigonometric relation

$$\cos \beta_i = (\cos \theta_{\text{obs}} - \cos \theta_i \cos \psi_i) / (\sin \theta_i \sin \psi_i). \quad (13)$$

In computing $\sin \beta_i$ from $\cos \beta_i$, we take into account the fact that $\sin \beta_i$ can be negative. Given the position of the observer and the emitting point on the star, this algorithm allows us to solve in a single step for the angles α_i and β_i at which the emitted ray leaves the stellar surface.

Knowing α_i and β_i , we can evaluate the energy-resolved photon number flux arriving at the distant observer from the infinitesimal element dA_i , using the expression (Poutanen & Gierliński 2003, Equation (13))

$$dF_i(E) = (1+z)^{-1} \frac{I_0(E_0, \alpha_i)}{E} \frac{dA_i \cos \alpha_i}{D^2} \left| \frac{\sin \psi}{\sin \alpha} \frac{\partial \psi}{\partial \alpha} \right|^{-1}, \quad (14)$$

where $E = (1+z)^{-1} E_0$ is the photon energy measured in the static frame at infinity and $1+z = (1 - 2M/R)^{-1/2}$ is the gravitational redshift from the stellar surface to infinity.

In order to compute efficiently the waveform produced by the emission from the entire hot spot, we proceed in two steps. First, we determine each of the waveforms produced by emission from a set of emitting areas $\{dA'_i\}$ with $\varphi = 0$ that span the colatitudes within the hot spot. In doing this, we take into account the fact that light emitted from different locations (R, θ_i, φ_i) takes different lengths of time to reach the observer. However, the waveform seen by the observer depends only on the *difference* in the light travel time from different emitting points. (A different choice for the zero of time or a different total propagation time would shift the arrival time of waveform but would not affect the shape of the waveform, which is what concerns us here.) Hence, we need to compute only the differences in the arrival times of the different rays that reach the observer. We choose to compute the arrival time of each ray relative to the arrival time of a radial ray emitted from a point on the stellar surface immediately below the observer, which in our time coordinate is $t = 0$ (see above). The arrival time of a photon that leaves the surface at an angle α to the normal as measured in the static frame is then

$$\Delta t = \int_0^1 dx \frac{R/c}{x^2 (1 - 2Mx/R)} \times \left[\frac{1}{\sqrt{1 - \sin^2 \alpha (1 - 2M/R)^{-1} (1 - 2Mx/R) x^2}} - 1 \right]. \quad (15)$$

We accurately evaluate this integral using a combination of analytical and numerical methods, as described in Appendix A. Once we have determined the waveform produced by emitting areas at $\varphi = 0$ that span the colatitudes within the hot spot, we compute the full waveform by adding the waveforms produced by emission from the grid points that have different values of φ . These waveforms can be generated quickly by appropriately shifting the phase of the waveform produced by the corresponding emitting element at $\varphi = 0$.

We have validated our waveform code using a suite of code tests. These tests, and the results, are discussed in Appendix A. We have carefully chosen values for the integration, hot spot, and angular resolution parameters that we use in our waveform code to provide resolutions fine enough to meet our accuracy requirements, but no finer, so that our code runs as fast as possible.

3.2. Bayesian Analysis and Sampling Methods

We wish to determine both the most probable (“best-fit”) values of the parameters in our waveform model, given an observed burst oscillation waveform, and the confidence regions for the values of these parameters. Both goals can be accomplished simultaneously and efficiently using Bayesian inference and an MCMC algorithm to sample the parameter space. Once we have computed the best-fit values of the model parameters, we can determine the accuracy of the fit by comparing them with the values that were used to generate the synthetic observed waveform.

The most probable values of the parameters \mathbf{y} in our waveform model and their confidence regions can be determined using the posterior probability distribution $p(\mathbf{y}|D, I)$, where D is the synthetic energy- and oscillation phase-resolved waveform

data of interest and I is any information obtained prior to the measurements under consideration. The desired posterior probability distribution can be obtained from the likelihood of the data, given the parameter values, using Bayes' theorem

$$p(\mathbf{y}|D, I) \propto p(D|\mathbf{y}, I)p(\mathbf{y}|I), \quad (16)$$

where $p(\mathbf{y}|I)$ is the prior probability distribution of the parameter values and the constant of proportionality is the inverse of the normalization factor. This constant of proportionality is irrelevant when estimating the values of the parameters in a given model. In the present analysis, we use the most uninformative prior, namely, we assume that $p(\mathbf{y}|I)$ is uniform for parameter values within the physical ranges we consider. Then, for parameter values within these ranges,

$$p(\mathbf{y}|D, I) \propto p(D|\mathbf{y}, I). \quad (17)$$

In performing the MCMC sampling of the parameter space, we use first-order Markov chains and start from a random point in the parameter space. At each step, we generate a proposed new set of parameter values \mathbf{y}' based on the current set of parameter values $\mathbf{y}^{(n)}$ by drawing from a proposal distribution $p_p(\mathbf{y}'|\mathbf{y}^{(n)})$. The proposed new set of parameter values is then accepted or rejected with specified probabilities. We use the basic Metropolis algorithm (see, e.g., von Toussaint 2011), drawing the new set of parameter values from a joint-normal distribution

$$p_p(\mathbf{y}'|\mathbf{y}^{(n)}) = p_p(\mathbf{y}^{(n)}|\mathbf{y}') \sim N(\mathbf{y}' - \mathbf{y}^{(n)}, \boldsymbol{\sigma}), \quad (18)$$

where the elements of $\boldsymbol{\sigma}$ are the standard deviations for each parameter, to be specified. We adopt the acceptance probability (which is also the transition probability)

$$T(\mathbf{y}', \mathbf{y}^{(n)}) = \min \left\{ 1, \frac{p(\mathbf{y}'|D, I)}{p(\mathbf{y}^{(n)}|D, I)} \right\}. \quad (19)$$

With this choice, we always accept the proposed new set of parameter values if its probability is higher than that of the current set; otherwise, we accept the proposed new set with probability $p(\mathbf{y}'|D, I)/p(\mathbf{y}^{(n)}|D, I)$. As required, this transition probability satisfies the detailed balance condition, as can be readily verified. Using relation (17), this ratio can be written

$$\frac{p(\mathbf{y}'|D, I)}{p(\mathbf{y}^{(n)}|D, I)} = \frac{p(D|\mathbf{y}', I)}{p(D|\mathbf{y}^{(n)}, I)}. \quad (20)$$

Hence, at each step the acceptance probability for the proposed set of parameter values is determined by calculating the ratio of the likelihoods of the data given the current parameters and given the proposed parameter values.

If Poisson noise is the only source of fluctuations in the data, the likelihood of the ‘‘observed’’ data, given a particular set \mathbf{y} of values for the model parameters, is

$$\mathcal{L} \equiv p(D|\mathbf{y}, I) = \prod_i \frac{m_i(\mathbf{y})^{d_i}}{d_i!} e^{-m_i(\mathbf{y})}, \quad (21)$$

where the product is over all the oscillation phase and energy bins, d_i is the measured number of counts in the i^{th} bin, and $m_i(\mathbf{y})$ is the number of counts in the i^{th} bin predicted by the model for the trial set \mathbf{y} of parameter values.

In our MCMC algorithm, we use the log likelihood only to determine the transition probability, which depends only on the *difference* of log likelihoods, so the $d_i!$ terms in the log likelihood cancel out, producing the expression

$$\log \mathcal{L} = \sum_i d_i \log m_i(\mathbf{y}) - N_{\text{model}}(\mathbf{y}), \quad (22)$$

where

$$N_{\text{model}}(\mathbf{y}) = \sum_i m_i(\mathbf{y}) \quad (23)$$

is the total number of photon counts in the model spectrum. Unlike many situations, where the normalization of the model is a free parameter and hence the total number of counts can be adjusted to be the same for every set of parameter values, here the total number of counts, and hence the normalization, depends explicitly on the distance to the star, the angular diameter of the spot, and other model parameters. The normalization of the model is therefore a key quantity for discriminating between different sets of parameter values.

3.3. Estimating M and R

We wish to determine the best-fit values and confidence regions of the parameters M and R in our waveform model, given a synthetic observed waveform of interest. With a uniform distribution over the allowed values of the model parameters as our prior, the posterior probability of a particular set of parameter values \mathbf{y} , given the data D , is proportional to $p(D|\mathbf{y})$, the likelihood of the data given those values of the parameters. There are two main computational tasks: (1) computing the relative likelihoods of the data over a set of trial model waveforms chosen to span and adequately sample the parameter space of the model, and (2) marginalizing the resulting posterior distribution by integrating over all the parameters in the model except M and R .

3.3.1. Construction of Synthetic Observed Waveforms

The data D representing a *synthetic observed waveform* is a list d_i of the (integer) number of counts observed in each oscillation phase-energy bin. We produce each synthetic observed waveform in three steps. First, we use the waveform generating code described in Section 3.1 to compute the oscillation-phase and energy-resolved waveform for a set of model parameter values of interest. Second, we add phase-independent (but energy-dependent) counts from our background model, which as we explained in Section 2.1 is a catch-all intended to mimic possible emission from the entire stellar surface, an accretion disk, and other sources in the field of view, as well as the instrumental background. Finally, we Poisson-sample the total number of photons in each of the phase-energy bins.

3.3.2. The Computational Problem

The straightforward way to determine the joint posterior probability distribution of M and R would be to integrate the joint distribution of all the parameters in our waveform model over every parameter except M and R . This integral,

$$p(M, R|D, I) = \int dV p(\mathbf{y}|D, I), \quad (24)$$

could in principle be computed using a Monte Carlo algorithm. Here V is the volume of the model parameter space when the

M – R subspace is excluded. If we were to sample the integrand at N points $\{x_i\}$ picked randomly but uniformly within V , the uncertainty in $p(M, R|D, I)$ would be (Press et al. 1992, Section 7.6)

$$\Delta p(M, R|D, I) \approx V \sqrt{\frac{\langle p^2 \rangle - \langle p \rangle^2}{N}}, \quad (25)$$

where

$$\langle p \rangle \equiv \frac{1}{N} \sum_{i=0}^N p(x_i|D, I) \quad (26)$$

and

$$\langle p^2 \rangle \equiv \frac{1}{N} \sum_{i=0}^N p(x_i|D, I)^2. \quad (27)$$

Expression (25) shows that the uncertainty in the marginalized posterior probability distribution obtained by Monte Carlo integration over the full posterior probability distribution decreases as the number N of sample points increases, but only as $1/\sqrt{N}$.

3.3.3. Computational Procedure

In this work, we seek to determine the most probable values of M and R and their confidence intervals by comparing our waveform model with a synthetic observed waveform, using a Bayesian approach. Each model waveform is a list of the expected number of photons m_i in each oscillation phase-energy bin. We construct a *complete* model waveform by computing a *hot spot* waveform (i.e., a model waveform without any background) and then adding a model of the background. For the waveform model we use here, specifying a complete waveform requires specifying 38 model parameters: M, R , the triplet of angles $\theta_{\text{spot}}, \Delta\theta_{\text{spot}}$, and θ_{obs} (which define the parameter subspace \mathbf{y}'), plus the color temperature T_{co} of the emission at the stellar surface measured in the comoving frame, the distance d to the star, the absolute phase of the oscillation ϕ_0 , and the background counts in 30 energy channels. Determining the joint posterior probability distribution of M and R using Equation (24) therefore requires accurate computation of the posterior probability distribution over a high-dimensional parameter space and subsequent computation of the marginalization integral over this space. We found that the computational effort needed to achieve sufficient accuracy using this approach was excessive. We therefore sought a more efficient approach, which we now describe. We presume that when the data from a large-area timing mission become available, the computational resources needed to do full Bayesian analyses of these data will also be available.

The purpose of our analysis here is not to reproduce all the steps that would be needed for a full Bayesian analysis of a real observed waveform, but rather to determine the precision and accuracy with which such a full analysis could determine M and R . Our initial exploration of the computational problem revealed several shortcuts that we could use for the current study that substantially reduce the computational burden but do not alter the results significantly. Using these shortcuts, the most probable values of M and R and their confidence intervals can usually be computed for a single synthetic observed waveform in 50–100 clock hours, running a parallel code on 150 nodes of a fast CPU cluster.

The procedure we use to determine the most probable values of M and R and their confidence intervals has seven steps: (1) construct an initial grid of points in the M – R plane; (2) for each M – R pair in this grid, choose values for the spot inclination

θ_{spot} , the spot angular radius $\Delta\theta_{\text{spot}}$, and the observer inclination θ_{obs} ; (3) determine the color temperature T_{co} of the emission from the hot spot measured in the comoving frame at the stellar surface, the absolute phase ϕ_0 of the waveform, the distance d to the star, and the background model that maximize the likelihood of the observed waveform; (4) sample the likelihood distribution $p(D|\mathbf{y}')$ over the three-dimensional parameter space \mathbf{y}' consisting of the three angles $\theta_{\text{spot}}, \Delta\theta_{\text{spot}}$, and θ_{obs} , using the most probable values of T_{co}, ϕ_0, d , and the background for each angle triplet; (5) integrate the posterior probability distribution $p(\mathbf{y}|D)$ over \mathbf{y}' , using the most probable values of the other model parameters except M and R , and thus associate an integrated probability density $p(M, R|D)$ with each point in the M – R grid; (6) refine and extend the M – R grid as needed to determine the most probable values of M and R and their confidence intervals with the desired accuracy; (7) use this final integrated probability density $p(M, R|D)$ to determine the most probable values of M and R and their 1σ and 3σ Bayesian confidence intervals. We now explain several aspects of this procedure in more detail.

Sampling M and R . To compute the posterior probability distribution $p(\mathbf{y}|D)$ efficiently, we first construct an initial grid of points in the M – R subspace surrounding the values of M and R that were used to generate the synthetic observed waveform being analyzed. This approach will of course not be possible when analyzing real data, because the actual values of M and R will not be known in advance. Instead, one will have to search the entire M – R parameter space, greatly increasing the computational effort required to adequately sample $p(\mathbf{y}|D)$. Once we have computed $p(\mathbf{y}|D)$ on the initial M – R grid, we use the results to extend and refine the grid, repeatedly if necessary, until the probability density has been determined over the relevant portion of the M – R plane with sufficient accuracy that we can determine the best-fit values of M and R and their uncertainties with the desired accuracy.

Determining the emission color temperature T_{co} . Our preliminary analysis showed that T_{co} is strongly constrained by the spectrum of the phase-dependent part of the observed waveform, and is much more strongly constrained than the angles $\theta_{\text{spot}}, \Delta\theta_{\text{spot}}$, and θ_{obs} . Hence the range of T_{co} values where the probability density is appreciable is very small and using the MCMC algorithm to sample this parameter is therefore very inefficient. Consequently we used a different approach to determine the value of T_{co} in our trial waveforms.

All the *trial* waveforms used in this paper were generated assuming that the spectrum of the emission from the hot spot has the same shape as the Planck spectrum. Most of the synthetic *observed* waveforms that we consider here were generated assuming that the spectrum of the emission from the hot spot has the same shape as a Planck spectrum. However, a few of the synthetic observed waveforms considered here were generated assuming a hot spot emission spectrum having the same shape as a Bose–Einstein spectrum, to explore the effects on the estimates of M and R of fitting a model that assumes a spectrum somewhat different from actual spectrum. We use different approaches to generate trial values of T_{co} in these two cases.

When the comoving emission spectrum used to generate the synthetic waveform has a Planck shape, we found that using the redshift relation for a non-rotating star, namely,

$$T_{\text{co}} = T_{\infty}(1 - 2M/R)^{-1/2}, \quad (28)$$

to estimate T_{co} using the observed radiation temperature T_{∞} gives T_{co} to within 5%, for spot rotation frequencies ≤ 600 Hz.

Consequently, to reduce the computational burden during our large-scale runs, for this synthetic spectrum we assume relation (28) and use as our trial value of T_{co} the value given by

$$T_{\text{co}} = T_{\text{co},0} \frac{(1 - 2M_0/R_0)^{1/2}}{(1 - 2M/R)^{1/2}}, \quad (29)$$

where the quantities with “0” subscripts have the values that were used in generating the synthetic waveform. Using this approach decreases the time required to compute $p(M, R|D)$ by a factor of five compared to the time required to solve for the most probable value of T_{co} .

When the emission spectrum used to generate the synthetic waveform has a Bose–Einstein shape, we included T_{co} as one of the parameters determined by maximizing the likelihood of the observed waveform, as we discuss below.

Changing the blackbody spectrum in the model waveform takes very little computational effort. The reason is that changing from one Planck spectrum with a given temperature and normalization to another with a different temperature and normalization can be done simply by mapping the spectrum to the different energy and multiplying it by the factor needed to renormalize it. Because photons of all energies are Doppler shifted, gravitationally redshifted, and deflected in the same way, once the energy-dependent waveform for one comoving emission temperature has been computed, the waveform for any other comoving emission temperature can be computed using the appropriate mapping and renormalization factor. Thus, the very time-consuming step of tracing rays needs to be done only once. To preserve accuracy when performing the mapping and renormalization, it is important to use tables that span the required energy range, so that interpolation can be used and extrapolation is not required.

If there is an atomic scattering line in the synthetic hot spot spectrum, we assume that we know the rest energy, width, and optical depth of the line and use this information to include a line with these properties in the spectrum of the waveform model.

Determining the most probable values of T_{co} , d , ϕ_0 , and the background. Given an M – R pair, there are 35 additional waveform parameters, if the synthetic emission has the shape of a Planck spectrum and T_{co} is determined as described above, or 36 additional parameters, if the synthetic emission has the shape of a Bose–Einstein spectrum and T_{co} has to be determined by a likelihood analysis.

As noted above, our preliminary analysis showed that T_{co} is strongly constrained by the spectrum of the phase-dependent part of the observed waveform. Hence, if T_{co} must be determined, we do so by maximizing the likelihood of the observed waveform, given the model waveform, using bisection.

In our preliminary investigation, we found that when the other waveform parameters are held fixed, the distance d to the star (which is equivalent to the normalization factor of the Planck spectral shape) is also very tightly constrained by the total number of observed counts. The total number of observed counts is in turn very tightly constrained by the data. Consequently, determining d by MCMC sampling and marginalization is very inefficient. We therefore determine d by maximizing the likelihood of the data using bisection. This method determines the most probable value of the distance to within 1%.

Our preliminary analysis showed that the phase ϕ_0 of the oscillation is very tightly constrained by the observed waveform. We therefore determine ϕ_0 by maximizing the likelihood using bisection. The first step is to apply the same, tentative phase shift to the model waveform in each photon-energy bin. We do

this by Fourier transforming the trial waveform at every photon energy, applying the tentative phase shift to all the harmonic components, and then inverting the Fourier transform to obtain the shifted model waveform. We perform the phase shift in the frequency domain rather than in the time domain to increase the accuracy: a phase shift applied in the time domain would have to be at least as large as the frequency times the width of a time bin, whereas even a very small phase shift can be applied to every harmonic in the frequency domain. Then, when the Fourier-transform is inverted the full waveform will be shifted by this small amount.

We compute the background model using bisection to determine the number of background counts at each energy that, when added to the phase-shifted hot spot waveform, produces a complete model waveform that maximizes the likelihood of the observed waveform.

Sampling the three angles. The values of the three angles θ_{spot} , $\Delta\theta_{\text{spot}}$, and θ_{obs} are much less tightly constrained than are T_{co} , d , ϕ_0 , and the background. Therefore, in our initial sampling of these angles we use the MCMC algorithm outlined above. In this step, we typically generate 30 chains, each with a length of 40, for a total of 1200 triplets. For each triplet of angles, we simultaneously determine the values of the remaining 32 (or 33) parameters that maximize the likelihood, using bisection. We then select the triplets that have a log likelihood within 20 of the maximum log likelihood. Finally, we use this collection of triplets to define the ranges of the three angles that we sample in the next step of our procedure.

For our second, final sample of the three angles, we typically choose 10,000 angle triplets randomly and uniformly within the ranges determined in the previous step. The number of triplets was chosen to provide enough resolution to obtain an acceptably accurate value of the integrated probability density. For each triplet of angles, we again determine the values of the remaining 32 (or 33) parameters that maximize the likelihood using bisection. We then sum the probabilities at all these triplets, to associate an integrated probability density $p(M, R|D)$ with this particular M – R pair.

Completing the sampling of M and R . Having computed the integrated probability density for a particular M – R pair, we then choose a different M – R pair in our grid and repeat the process. If necessary, we extend and refine the M – R grid, in order to determine the most probable values of M and R and their 1σ and 3σ confidence intervals with the desired accuracy. We use a common normalization for the integrated probability density of different M – R pairs, so it can be compared for different pairs. Finally, we use the resulting probability distribution in the M – R subspace to determine the most probable values of M and R and their 1σ and 3σ confidence intervals. We tested this algorithm in numerous cases and found that it does a good job of determining the most probable M – R pair and the corresponding confidence regions.

4. WAVEFORM ANALYSIS AND RESULTS

Our burst oscillation waveform analysis proceeds in the following three steps.

1. Generate a synthetic observed waveform, using the method described in Section 3.1.
2. Compute the joint posterior distribution of the mass and radius parameters in our waveform model, given this synthetic waveform, using the computational methods described in Sections 3.2 and 3.3.

3. Use this posterior distribution to determine the most probable values of M and R and the 1σ and 3σ confidence regions in the M – R plane, given this synthetic waveform.

We use this procedure to explore how the accuracy and precision of M and R estimates depend on the physical characteristics of the observed system by generating a number of different synthetic waveforms corresponding to systems with different physical characteristics, such as spot rotation rate, spot inclination, and viewer inclination. We then explore the effects on the uncertainties in M and R of having additional information about the observed system that is independent of that provided by the burst oscillation waveform. We consider knowledge of the distance, the observer’s inclination, and the size and spectrum of the background, measurement of the properties of an atomic scattering line in the emission from the hot spot, and high-precision measurements and modeling of the spectrum of the X-ray bursts from the star in question. We also explore the effects on M and R estimates if the actual spot shapes, X-ray spectra, and radiation beaming patterns differ from those assumed in our waveform model. Finally, we investigate the constraints that can be obtained by analyzing data on burst oscillations with changing waveforms, collected during a single burst or during multiple bursts.

4.1. Waveform Analysis

All the synthetic waveforms used in this work assume $M = 1.6 M_{\odot}$ and $R = 5.0 M$. We compute the joint posterior distribution of the parameters M and R in our waveform model, given the synthetic waveform, in the domain $1.0 M_{\odot} \leq M \leq 2.8 M_{\odot}$ and $5 \text{ km} \leq R \leq 30 \text{ km}$ for values of R/M in the range $3.5 \leq R/M \leq 6.8$. Given that all the synthetic waveforms assume $R/M = 5.0$, this range in R/M is adequate to evaluate all cases that provide relatively tight constraints on M and R .

4.1.1. Model Burst Oscillation Waveforms

The hot spot emission model that we usually use in fitting the synthetic waveform data assumes that the spot is a uniformly emitting circular area on the stellar surface that is rotating around the star and that each point on the surface of the hot spot radiates with a spectrum that has the shape of a Planck spectrum and the angular pattern expected for emission from a semi-infinite, nonrelativistic, electron scattering atmosphere (see Section 3.1 for further details).

The full waveform model has nine parameters that describe the waveform produced by the emission from the hot spot (the star’s gravitational mass M , circumferential radius R , and the spot rotation frequency ν_{rot} ; the spot inclination or colatitude θ_{spot} and angular radius $\Delta\theta_{\text{spot}}$; the color temperature T_{co} of the spot emission as measured in the comoving frame at the stellar surface; the observer’s inclination θ_{obs} relative to the rotation axis and the distance d to the star; and the phase of the waveform). The waveform model has an additional 30 parameters that describe the background model (the number of background counts in each of 30 energy channels). Thus there are in principle a total of 39 parameters in the model.

In practice, the spot rotation frequency can be determined separately and much more accurately than any of the other parameters in the waveform model. Therefore, to speed up the computations for the present study, we assume that ν_{rot} is already known. As discussed in Section 3.3.3, T_{co} can be determined very accurately using the color temperature observed at infinity when the spectrum of the emission from the hot spot that produces the observed waveform has the shape of a Planck

spectrum. As discussed there, in this case we use this fact to determine the value of T_{co} separately from the fitting procedure, to speed up the computations by a factor of five. When the spectrum from the hot spot does not have the shape of a Planck spectrum, T_{co} must be included as one of the fitted parameters. Thus, the hot spot component of the waveform model has seven adjustable parameters, if the spectrum has the shape of a Planck spectrum, or eight free parameters, if it does not. The full waveform model, which includes the background, therefore has either 37 or 38 adjustable parameters. Each synthetic oscillation waveform has 480 data points (see below), so the number of degrees of freedom in the fits described here is either $480 - 37 = 443$ or $480 - 38 = 442$, depending on the hot spot spectrum that was assumed in generating the synthetic waveform data.

If the properties of the observed neutron star are consistent with the properties assumed here, we expect the constraints on M and R to improve as the ratio \mathcal{R} of the number of modulated counts to the fluctuations in the total number of counts (see Equation (1)) increases. For a fixed value of \mathcal{R} , we expect the constraints to be tighter when the spot and observer are closer to the star’s rotational equator, because the relativistic Doppler shift and aberration are then greater, producing a waveform that depends more sensitively on the radius of the star.

4.1.2. Synthetic Burst Oscillation Waveforms

Each synthetic observed waveform is represented by the number d_i of counts in each of 30 energy bins of width 0.3 keV covering 3.5 keV to 12.5 keV at each of 16 rotational phases. There are thus a total of 480 data points in each synthetic waveform.

The physical models we use to generate the component of these waveforms that is produced by the hot spot are characterized by the gravitational mass M and circumferential radius R of the star; the spot rotation frequency ν_{rot} ; the location, size, and shape of the hot spot; the spectrum and beaming pattern of the emission from the hot spot; and the distance d to the observer and the observer’s inclination θ_{obs} relative to the star’s rotation axis. The waveforms used in this work assume that the emission from the hot spot has the shape of a Planck spectrum or a Bose–Einstein spectrum with a temperature $kT_{\text{co}} = 2.0 \text{ keV}$, as measured in a frame comoving with the surface of the rotating hot spot. To facilitate comparison of different cases, we adjust the distance to the star so that the expected number of counts from the hot spot is 10^6 in all cases.

To explore the effects of background counts, we add a background component to the waveform produced by the hot spot. As discussed in Section 2.1, this component is a catch-all for all counts not produced by radiation from the hot spot. These counts could be produced by emission from unassociated sources in the field, the accretion disk, the non-spot portion of the star, instrumental backgrounds, or any combination of these. We assume that the background does not vary at frequencies commensurate with the spot rotation frequency. We model the background by assuming uniform emission from the entire surface of the star with a spectrum having the shape of a Planck spectrum with a temperature of 1.5 keV as seen in a frame comoving with the surface of the star, which for this purpose is assumed to be rotating with the same frequency as the hot spot. To study the effects of different background levels, we adjust the normalization of the background component so that the number of counts expected from the background is either 0.3, 1.0, or 9.0 times the 10^6 counts expected from the hot spot. We refer to these as our low, medium, and high backgrounds.

Table 1
System Parameters for the Two Reference Cases^a

Case Label	ν_{rot} (Hz)	θ_{obs} (deg)	θ_{spot} (deg)
Low inclination	400	60	20
High inclination	600	90	90

Note. ^a Here ν_{rot} is the rotation frequency of the hot spot, θ_{obs} is the inclination of the observer relative to the rotation axis, and θ_{spot} is the inclination (colatitude) of the center of the hot spot.

Once the preliminary synthetic waveform is generated, it is Poisson sampled to produce the complete synthetic waveform. Because the waveforms are Poisson-sampled, the actual counts contributed to the complete waveform by the hot spot and the background are slightly different from their expected values.

Standard synthetic waveforms. We first analyze synthetic waveforms generated using burst oscillation models with the same qualitative properties as the model that we fit to the data. These “standard” synthetic waveforms assume that the hot spot is circular and uniformly emitting, that the spectrum of the emission from the hot spot has the shape of a Planck spectrum, as measured in the comoving frame at the stellar surface, and that the beaming pattern of the emission from the hot spot is that expected for a nonrelativistic electron scattering atmosphere, which is described by the Hopf beaming function (see Section 3.1). The standard waveforms we analyze here assume that the angular radius $\Delta\theta_{\text{spot}}$ of the hot spot is 25° .

Reference cases. The hot spot and observer inclinations have the biggest effects on the uncertainties in estimates of M and R . To facilitate comparisons, we constructed several standard waveforms for a “low-inclination” reference case and for a “high-inclination” reference case. The model parameters that define these two cases are ν_{rot} , θ_{obs} , and θ_{spot} . The values of these parameters for the two reference cases are listed in Table 1.

The low-inclination reference case is motivated by the evidence that the magnetic poles of many accreting neutron stars with millisecond rotation periods are close their rotation poles (see Lamb et al. 2009a, 2009b). This evidence includes the low fractional amplitudes of the X-ray oscillations of accretion-powered millisecond X-ray pulsars, their highly sinusoidal waveforms, and their intermittency in some stars. In the low-inclination reference case, the hot spot and the observer are at inclinations of 20° and 60° , respectively, and the rotation frequency is 400 Hz.

Even if the star’s magnetic poles are close to its rotation pole, the region heated by nuclear burning or surface modes may not be at or even near a magnetic pole, and hence may not be close to the rotation pole. As an example, the very high-amplitude X-ray oscillations found near the beginning of a burst from 4U 1636–536 by Strohmayer et al. (1998) are difficult to understand if the heated region is near the rotation pole. We therefore also consider a high-inclination reference case in which the hot spot and observer inclinations are both 90° (i.e., both are in the rotational equator) and the spot rotation frequency is 600 Hz. These choices produce the largest expected Doppler shifts and relativistic aberration, and should therefore illustrate the tightest possible constraints on M and R .

Nonstandard synthetic waveforms. In addition to our “standard” synthetic waveforms, which assume a circular hot spot with a 25° angular radius, a Planck emission spectrum, and the radiation beaming pattern appropriate for emission from a Thomson scattering atmosphere (described by the Hopf func-

tion), we also generated several other waveforms, to explore the consequences of fitting an incorrect waveform model to the data. These “nonstandard” waveforms were generated for spots elongated in the north–south and east–west directions, for a Bose–Einstein emission spectrum, or for isotropic beaming from the surface of the hot spot.

4.1.3. Independent Additional Information

We explore the improvements in constraints on M and R that become possible with independent additional information about important system properties. There are good prospects for determining the angle of our line of sight relative to the axis of the binary orbit using optical burst reflection mapping (Casares 2010) and Fe-line modeling (Cackett et al. 2010; Egron et al. 2011). The neutron stars in bursting systems are thought to have acquired nearly all their angular momentum via accretion (Lamb & Boutloukos 2008), so the rotation axes of these stars are likely to be closely aligned with the orbital axes of their systems, allowing us to infer the inclination of the observer relative to the stellar rotation axis. The distance of a bursting neutron star is often tricky to measure, but it can be constrained if the system is in a globular cluster. Parallax distances for some bursting neutron stars could become available in the future using *Gaia* (Eyer et al. 2013). The properties of the non-spot background are likely to be difficult to constrain independently of the waveform fitting process, but if the spectrum of the background is sufficiently distinct from the spectrum of the spot radiation (e.g., if both have Planck spectral shapes but have easily distinguishable temperatures), it may be possible to restrict the properties of the background using spectral measurements. It is also possible that an atomic resonance scattering line could be identified in the emission from the hot spot. We expect that having information like this would improve the constraints on M and R .

It may be possible to obtain other independent information by fitting high-precision spectral models to the observed spectra of long bursts or jointly fitting such models to the spectra of many bursts from the same source. At high burst temperatures, the burst atmosphere is expected to be scattering-dominated and to produce spectra with shapes that are very close to the shape of a Bose–Einstein spectrum (K. H. Lo et al., in preparation). This expectation is consistent with the observed spectra of such bursts (Boutloukos et al. 2010). If the observed spectra had exactly a Bose–Einstein spectral shape, they would provide no information about the properties of the star, because a thermodynamic equilibrium spectrum provides no such information. However, we have found that detailed, high-precision model atmosphere spectra provide a much better description of the best available spectral data (the *RXTE* PCA data on the superburst from 4U 1820–30) than do Bose–Einstein spectra, so some information about the properties of the star can be obtained. Information about the mass and radius of the star comes mainly from the small deviation of the spectrum from a Bose–Einstein shape at low photon energies that is caused by the free-free opacity. This deviation allows one to determine $(1+z)/g^{2/9}$, where z is the surface redshift and g is the surface gravity of the star. This relation between M and R complements the constraints on M and R obtained from fits to the burst oscillation waveform.

4.1.4. Errors in the Waveform Model

Most of the results we present here assume that we know the shape of the hot spot and the spectrum and beaming pattern of

Table 2
Synthetic Waveforms, Fitted Models, and Corresponding Figures^a

Figure	Inclinations ^b	Background ^c	Additional Description	δM_1^d (%)	δR_1^d (%)
1	High	Medium	All parameters known except M and R	1.9 ^e	1.3 ^e
2(a)	High	Low	All parameters unknown	2	5
2(b)	Low	Low	All parameters unknown	30	30
2(c)	High	Medium	All parameters unknown	3	4
2(d)	Low	Medium	All parameters unknown	50	40
2(e)	High	High	All parameters unknown	9	8
2(f)	Low	High	All parameters unknown	... ^f	... ^f
3(a)	High	Medium	$\theta_{\text{spot}} = 80^\circ$, all parameters unknown	6	6
3(b)	High	Medium	$\theta_{\text{spot}} = 60^\circ$, all parameters unknown	5	5
4(a)	High	High	Observer inclination known	5	5
4(b)	Low	High	Observer inclination known	... ^f	... ^f
4(c)	High	High	Distance known	7	6
4(d)	Low	High	Distance known	... ^f	... ^f
4(e)	High	High	Atomic line present in the data	4	4
4(f)	Low	High	Atomic line present in the data	... ^f	... ^f
5(a)	High	Medium	Background known	3	5
5(b)	Low	Medium	Background known	20	20
6(a)	High	Medium	Hot spot elongated north–south	40	40
6(b)	Medium	Medium	Hot spot elongated east–west	8	17
7(a)	High	Medium	Bose–Einstein spectrum in the data	6	6
7(b)	Low	Medium	Bose–Einstein spectrum in the data	... ^f	... ^f
7(c)	High	Medium	Isotropic beaming in the data	... ^g	... ^g
7(d)	Low	Medium	Isotropic beaming in the data	... ^f	... ^f

Notes.

^a All synthetic waveforms were normalized to produce $\sim 10^6$ counts from the hot spot.

^b See Table 1 for the parameters used in generating the high- and low-inclination synthetic waveforms. The elongated spots are defined in the text.

^c The backgrounds in all synthetic waveforms are phase-independent, with the spectral shape of a Planck spectrum having a temperature at infinity of $kT = 1.5$ keV. Our low, medium, and high backgrounds produce $\sim 3 \times 10^5$, $\sim 1 \times 10^6$, and $\sim 9 \times 10^6$ detected counts.

^d Approximate 1σ fractional uncertainties in M and R estimated by projecting the 1σ contours onto the M and R axes (see the text).

^e These purely statistical uncertainties assume that all relevant properties of the actual system are accurately described by the waveform model, that the values of these properties lie within the range of values covered by models considered in the waveform fitting process, and that the values of all the model parameters except M and R are known exactly by some means that is independent of the waveform fitting process, thereby eliminating the degeneracies of the model parameters.

^f Reliable estimates of the uncertainties are not possible because the confidence regions are too large.

^g The reduced χ^2 of this fit indicates that the model is incorrect.

the radiation from the spot. However, in practice these properties may not be accurately known. If any important properties of the actual star differ significantly from the properties assumed in the model being fit, we are attempting to fit the data using an incorrect model. In this situation, we expect the fit to be poor and the best-fit values of M and R to differ from their actual values. Of particular concern is the possibility of situations in which the fit is formally good but the best-fit values of M and R are significantly different from their actual values. A full analysis of possible systematic errors is beyond the scope of this work, but we briefly explore the effects of errors in the model of the spot emission.

4.2. Constraints on M and R

Consider now the constraints that can be obtained by fitting a waveform model to an observed burst oscillation waveform. Our results show that the uncertainties in M and R estimates produced by statistical fluctuations in the waveform are typically only a percent or two, if $\sim 10^6$ counts are collected from the hot spot of a given bursting neutron star and the values of all other parameters are known independently of the waveform analysis. However, for realistic situations in which some of the other parameters in the model must be determined from the waveform, the statistical uncertainties are usually much larger. The reason

for this is that the effects on the waveform of changing different parameters in the waveform model are often very similar. These degeneracies of the waveform with respect to changes in the values of the model parameters are an inherent property of any physical model based on a rotating hot spot and cannot be removed by “improving” the model. These degeneracies inflate the uncertainties in M and R estimates produced by statistical fluctuations.

The synthetic waveforms we consider here and the resulting 1σ uncertainties in M and R are summarized in Table 2. Plots of the 1σ and 3σ uncertainty contours for each of these waveforms are presented and discussed below. The χ^2/dof for the best-fit model is shown on each plot.

The uncertainties in M and R quoted in this section were estimated by projecting the extent of the two-dimensional uncertainty regions onto the M and R axes. This approach neglects the fact that the probability distributions of M and R are often correlated. It also somewhat overestimates the individual uncertainties in M and R , because the central peak in the distribution of the unwanted quantity would centrally weight the probability distribution of the quantity being considered, if the unwanted quantity were actually marginalized. This effect is neglected when the uncertainty region is simply projected. We note that the quoted uncertainties are themselves slightly




Figure 1. These contours are based on the highly unrealistic assumptions that all the relevant properties of the actual system are correctly included in the waveform model, that the actual values of these properties lie within the range of values possible in the waveform model, and that the values of all the parameters in the waveform model except M and R are known exactly, independent of the waveform fitting process. This procedure eliminates completely the effects of the degeneracy of the waveform with respect to variations in the values of the system parameters. The + symbol indicates the radius and mass that were used in generating the synthetic waveform and the star symbol shows where the marginalized posterior probability density is highest. The dotted and solid curves show, respectively, the 1σ and 3σ confidence contours. The reduced χ^2 indicates that the fit is acceptable. See the text for further details.

uncertain, due to incomplete sampling of the high-dimensional parameter space.

4.2.1. Effects of Statistical Fluctuations

We can estimate the uncertainties in estimates of M and R produced by the statistical fluctuations in the observed waveform by considering the (highly unrealistic) situation in which all the relevant properties of the actual system are correctly included in the waveform model, the actual values of these properties lie within the range of values possible in the model, and the values of all the parameters in the model except M and R are known exactly by some means that is completely independent of the waveform fitting process.

Figure 1 shows an example in which the assumptions in the model used to fit the synthetic waveform data are identical to the assumptions in the model that was used to produce the waveform data. The waveform data were generated using the parameter values of the high-inclination reference case and our medium background. In fitting the model to the waveform data, the values of all the parameters in the model except M and R were set to the values used to produce the synthetic waveform. This procedure completely eliminates the effects of the degeneracy of the waveform under variations of the model parameters. The most probable values of M and R were then determined using our standard fitting and marginalization procedures. The 1σ uncertainties in M and R are about 1.9% and 1.3%, respectively, whereas the 3σ uncertainties in M and R are about 5% and 3%. The value of χ^2/dof indicates that the fit is acceptable.

These results can be compared with those of Strohmayer (2004), who investigated the constraints on M and R that could be obtained by fitting a bolometric waveform model to synthetic bolometric waveform data. He assumed that the hot spot and the observer are both in the rotation equator (their most favorable locations) and that the size, shape, and location of the hot spot, the spectrum and beaming pattern of the radiation from the hot spot, and the inclination of the observer are all

Table 3
Scaling of Uncertainties in M and R with Total Counts^a

N_{tot}	δM_1 (%)	δM_3 (%)	δR_1 (%)	δR_3 (%)
10^4	15	42	12	27
10^5	6.8	15	4.7	11
10^6	1.8	5.0	1.3	3.1

Note. ^a N_{tot} is the total number of counts in the high-inclination, medium-background synthetic waveforms that were analyzed, δM_1 and δM_3 are the 1σ and 3σ uncertainties in M , and δR_1 and δR_3 are the 1σ and 3σ uncertainties in R . These uncertainties assume all system parameters are known except M and R (see text). The fractional rms amplitudes of all three waveforms are $\sim 54\%$ (see Table 4)

known independently of the waveform fitting analysis, so that only M and R have to be estimated from the waveform data. Strohmayer included in the synthetic observed waveform a background commensurate with the pre-burst emission in some bursts from 4U 1636–536 ($\sim 5\%$ of the maximum flux during strong bursts). The background was then fitted for as part of the waveform fitting process (T. E. Strohmayer 2013, private communication). With these assumptions and approximations, Strohmayer found that fitting a bolometric waveform model to bolometric waveform data from an X-ray instrument with an effective area $\sim 10 \text{ m}^2$ could provide estimates of M and R with uncertainties of a few percent.

Our simulations show that the statistical uncertainties in M and R scale roughly as $1/\mathcal{R}$, where $\mathcal{R} \equiv N_{\text{osc}}/\sqrt{N_{\text{tot}}}$ in terms of the number N_{osc} of counts in the oscillating component of the synthetic waveform and the total number N_{tot} of counts in the waveform (see Equation (1) and the discussion following it). This is illustrated in Table 3. In these particular examples, the expected number N_{back} of counts in the background is equal to the number N_{spot} of counts from the hot spot, so $N_{\text{osc}} \propto N_{\text{spot}} \propto N_{\text{spot}} + N_{\text{back}} = N_{\text{tot}}$ and hence \mathcal{R} is proportional to $\sqrt{N_{\text{tot}}}$.

4.2.2. Effects of Parameter Degeneracies

Comparison of the uncertainties listed in Table 2 and comparison of the confidence regions shown in Figure 1 with the confidence regions shown in the other figures in this section shows that when parameters in the waveform model other than M and R must be determined as part of the waveform fitting process, the uncertainties in the resulting estimates of M and R are generally much larger than when only M and R are unknown. The reason the uncertainties are so large is that the waveform is degenerate with respect to changes in many of the parameters in the waveform model, i.e., the effects on the waveform of changing the value of one of the system parameters can be partially or wholly compensated by changing the value of one or more other parameters.

1. The effect on the stellar compactness and surface redshift of an increase in R can be compensated almost exactly by an increase in M , because the special relativistic effects on the waveform of the rotational surface velocity are $\sim 5\%$ or less for the stellar models considered here (see Section 3). This is a very strong degeneracy.
2. Changes in the spot inclination, viewer inclination, and stellar radius affect the waveform in similar ways and can therefore compensate for one another. The increase in the

relativistic Doppler boost and aberration produced by an increase in the spot inclination can be compensated in part by a decrease in the observer’s inclination and/or a decrease in the stellar radius. The change in the projected emitting area seen by the observer produced by an increase in the spot inclination can be partially compensated by an increase in the observer’s inclination.

3. An increase in the stellar compactness can be partially compensated for by a decrease in the spot radius, but as shown by Lamb et al. (2009a), the dependence of the waveform on the spot radius is weak unless the spot is very large, while its dependence on the compactness is weak unless the star is very compact ($M/R > 0.25$). Hence this is a weak degeneracy.
4. Changing the number and spectrum of the background counts affects the inferred spectrum and modulation fraction of the waveform and therefore can partially compensate for changes in any of the other parameters that affect the spectrum and modulation fraction of the waveform, including the spot inclination and size, the stellar compactness, and the observer inclination.

As noted previously, these degeneracies are an inherent property of any physical model based on a rotating hot spot and cannot be removed by “improving” the model.

As a result of these degeneracies, the uncertainties in estimates of M and R depend sensitively on the physical properties of the system. If the hot spot and the observer are at high inclinations, the effects on the waveform caused by the rotation of the star will be maximal, reducing the effects of these degeneracies. If instead the hot spot and the observer are at low inclinations, the effects caused by the rotation of the star will be much less and these degeneracies will be much more important.

A large background exacerbates the degeneracy problem, because it increases the statistical fluctuations in the observed waveform. As a result, a wider range of models will adequately fit the waveform data. On the other hand, independent additional knowledge about the system can reduce or eliminate degeneracies. For example, accurate a priori knowledge of the observer’s inclination can significantly improve the constraints, if the spot and observer inclinations are high; a priori knowledge of the distance to the system can also help. Independent knowledge of the background can be used to tighten the constraints on M and R substantially. If the spot and observer inclinations are high, measurement of an identifiable atomic spectral line in the burst oscillation spectrum can be used to tighten the constraints.

4.2.3. Effects of Spot and Observer Inclinations

Figure 2 (see also Table 2) shows that the constraints on M and R are fairly tight when the hot spot and the observer are both at high inclinations but very loose when they are both at low inclinations, regardless of the background level. Comparison of Figures 2(c), 3(a), and 3(b) (again see Table 2) shows that for high observer inclinations, the constraints on M and R weaken slowly as the spot inclination decreases from 90° to 60° .

If the hot spot and observer inclinations are both high, the 1σ uncertainties in M and R range from a few percent for our low and medium backgrounds to about 10% for our high background (see Table 2). The 3σ uncertainties are much larger, ranging from about 15% for our low or medium backgrounds to about 50% for our high background. If instead the spot and observer inclinations are both low, the 1σ uncertainties in M and R are large, ranging from 30% to 50% or more

(again see Table 2), and the 3σ contours span much or all of the M – R domain that was searched. The importance of high spot and observer inclinations becomes especially clear if one compares the confidence regions for the high-inclination case shown in Figure 2(d) with the confidence regions for the low-inclination case shown in Figure 2(e). Although the fractional rms amplitudes of the oscillations are the same in these two cases, the inclinations are different, producing very different confidence regions (the rms amplitudes are the same because the backgrounds are different).

For high hot spot and observer inclinations, the oscillation amplitude and the effects on the energy dependence and harmonic content of the waveform of the relativistic Doppler boost and aberration are responsible for the lower bound on R . The upper bound on R is provided by the requirement that the amplitude be sufficiently large but the effects of rotational motion on the waveform must not be too large. This requirement also imposes a lower bound on the observer’s inclination.

For low spot and observer inclinations, the waveform is nearly sinusoidal and the modulation fraction is low. Therefore much wider ranges of the parameters other than M and R (such as the spot and observer inclinations) give a modulation fraction low enough to be consistent with the observed waveform, if the compactness is high. Consequently, when we marginalize over these other parameters, M – R pairs near the high compactness boundary are given more weight. As a result, the highest posterior probability density is typically near the highest compactness that was searched, which means it is typically close to the high-compactness boundary of the search domain, even though comparable fits would be possible for many values of M/R if we were to optimize the other parameters.

4.2.4. Effects of the Background

Figure 2 (see also Table 2) illustrates how the constraints on M and R are affected by background counts. The uncertainties in M and R are affected little by background count rates less than or comparable to the count rate from the hot spot, but become significantly larger for background count rates higher than the count rate from the hot spot. This is illustrated by Figures 2(a), (c), and (e). The background count rate for the case shown in Figure 2(a) is about 30% of the count rate from the hot spot, in Figure 2(c) it is about equal to the count rate from the hot spot, and in Figure 2(e) it is about nine times the count rate from the hot spot. The uncertainties in M and R for the cases shown in Figures 2(a) and (c) are therefore expected to be similar. The 3σ uncertainties in M and R are slightly smaller in the realization shown in Figure 2(c) than in the realization shown in Figure 2(a) because of a statistical fluctuation.

When the number of background counts exceeds the number of counts from the hot spot, the uncertainties increase as the background increases approximately as $1/\mathcal{R}$, where $\mathcal{R} \equiv N_{\text{osc}}/\sqrt{N_{\text{tot}}}$ in terms of the number of counts N_{osc} in the oscillation and the total number of counts N_{tot} (see Equation (1)). Thus, the 1σ and 3σ uncertainties for the high background cases, where the background contributes nine times as many counts as the hot spot, are several times larger than the uncertainties for the low and medium background cases, where the background contributes fewer and about the same number of counts, respectively, as the hot spot.

The \mathcal{R} -values, the fractional rms amplitudes of the first two harmonics, and the total fractional rms amplitudes for the synthetic waveforms generated in the low- and high-inclination reference cases for our low, medium, and high backgrounds

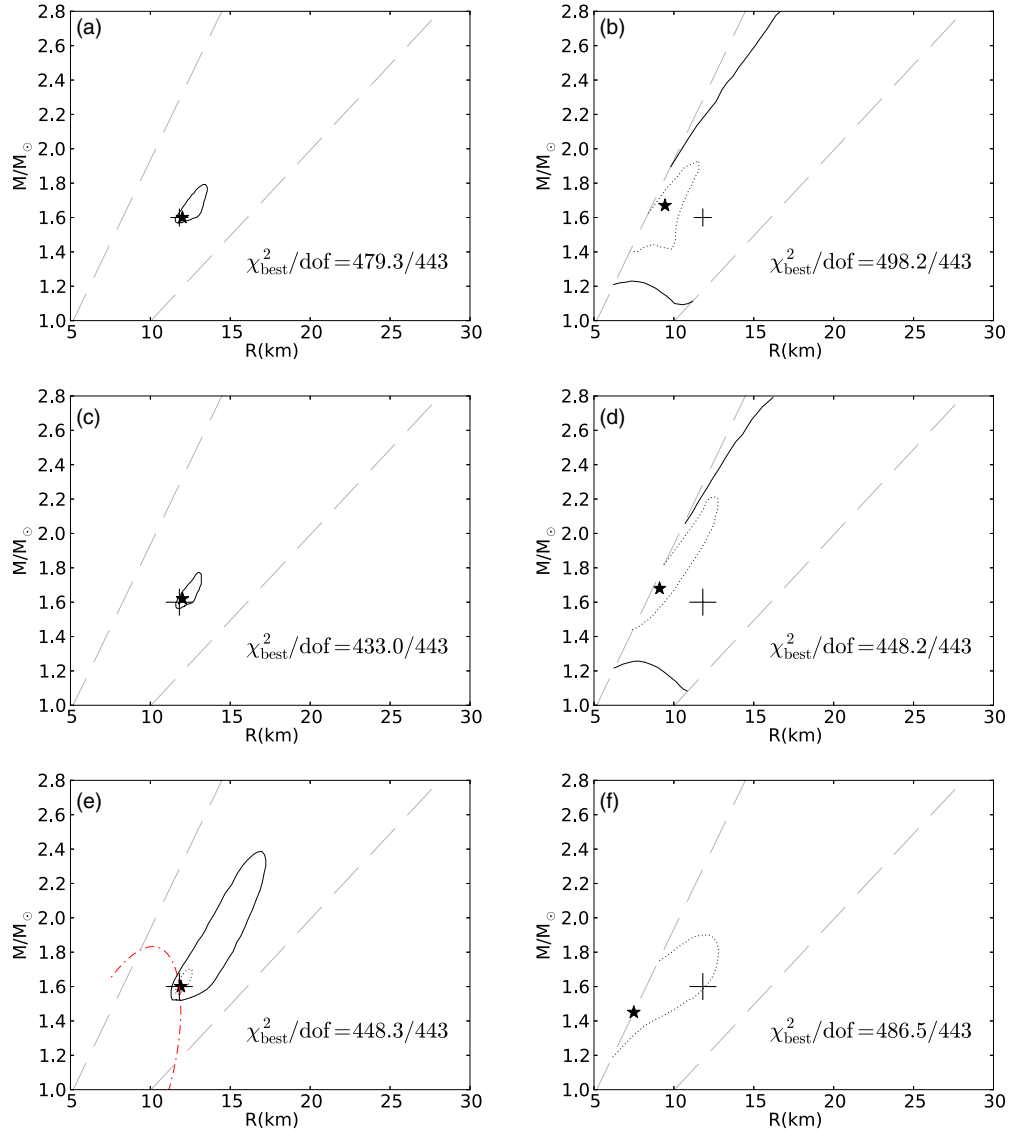


Figure 2. These results show that M and R are tightly constrained when the spot and observer inclinations are high and the total number of counts from the background is comparable to or less than the number from the hot spot. The contours in the left column are for the high spot and observer inclination model and (top to bottom) the low-, medium-, and high-background models (for the properties of these models, see Tables 1 and 2). The right column shows the corresponding contours for the low spot and observer inclination model. The dashed lines show the $R/M = 3.5$ and $R/M = 6.8$ boundaries within which the posterior probability distribution was computed. The dotted and solid curves show, respectively, the portions of the 1σ and 3σ confidence contours within this domain. The + symbol indicates the radius and mass that were used in generating the synthetic waveform and the star symbol indicates the values where the marginalized posterior probability density is highest. The red dash-dotted curve in Figure 2(e) shows points of constant $g^{2/9}/(1+z)$, illustrating the complementary constraint that can be obtained by fitting the spectra of X-ray bursts from the same star (see the text).

(A color version of this figure is available in the online journal.)

are listed in Table 4. Scaling from *RXTE* observations suggests $\mathcal{R} \sim 50$ for a typical burst rise or tail oscillation from 4U 1636–536 observed using a 10 m^2 detector, ~ 80 for a typical tail oscillation from 4U 1728–34, and ~ 90 for the oscillation trains from bursting neutron stars that persist for more than 2 s (see Section 2.2.1). \mathcal{R} scales as the square root of the number of counts, so combining the data from 25 such oscillations (see Section 4.2.7 and Appendix B) could give \mathcal{R} values ranging from ~ 250 for oscillations from 4U 1636–536 up to ~ 400 – 450 for the best oscillation trains from other bursting neutron stars. The latter \mathcal{R} values are comparable to those for the synthetic waveforms generated in the high-inclination, high-background cases. The typical total rms amplitudes of these observed oscillations are $\sim 5\%$, about half that of the $\sim 10\%$ fractional amplitudes of the oscillations in the synthetic

waveforms generated in the high-inclination, high-background cases. Second harmonics (first overtones) have been detected using *RXTE* only in burst oscillations from neutron stars that also produce accretion-powered millisecond oscillations at their rotation frequencies, except possibly for some oscillations observed during the rise of bursts from 4U 1636–536 (see Section 2.2.1).

4.2.5. Effects of Independent Knowledge of Relevant System Properties

The extent to which knowledge that is independent of the waveform fitting process improves the precision of M and R estimates depends strongly on the inclinations of the hot spot and the observer. If they are high, the uncertainties in M and R can be reduced by independent knowledge of the observer’s




Figure 3. Comparison of the contours shown here and in Figure 2(c) demonstrates that the constraints on M and R weaken slowly as the spot inclination decreases from 90° (in Figure 2(c)) to 80° (left panel) and to 60° (right panel). All three plots assume the observer inclination is 90° and our medium background ($\sim 10^6$ counts from the hot spot and $\sim 10^6$ background counts). The red dash-dotted curves show points of constant $g^{2/9}/(1+z)$, illustrating the complementary constraint that can be obtained by fitting the spectra of bursts from the same star (see the text). For the meanings of the line types and symbols, see Figure 2.

(A color version of this figure is available in the online journal.)

Table 4

\mathcal{R} Values and Fractional rms Amplitudes for Selected Synthetic Waveforms^a

Inclination	Background	\mathcal{R}	frms1 ^b	frms2 ^c	frms ^d
Low	Low	340	0.212	0.016	0.21
Low	Medium	280	0.141	0.011	0.14
Low	High	120	0.028	0.0019	0.028
High	Low	1360	0.748	0.293	0.81
High	Medium	1100	0.498	0.195	0.54
High	High	490	0.100	0.039	0.11

Notes.

^a Here $\mathcal{R} \equiv N_{\text{osc}}/\sqrt{N_{\text{tot}}}$, where N_{osc} is the number of counts in the oscillating component of the synthetic waveform and N_{tot} is the total number of counts in the waveform. See Equation (1).

^b Fractional rms amplitude of the first harmonic (fundamental).

^c Fractional rms amplitude of the second harmonic (first overtone).

^d Total fractional rms amplitude.

inclination or of the distance to the star. In this case, observing an atomic scattering line with a known rest energy tightly constrains the stellar compactness M/R and improves the constraints on M and R more than knowing the distance and approximately as much as knowing the observer's inclination. If instead the spot and observer inclinations are low, observing a scattering line with a known rest energy still tightly constrains the stellar compactness but improves the constraints on M and R very little. The uncertainties in M and R can be reduced substantially by independent knowledge of the size and spectrum of the background, especially if the spot and observer inclinations are high. Table 2 lists the 1σ fractional uncertainties for each of these cases. Analysis of high-quality observations of the spectra of the bursts using high-precision spectral models can help to further constrain M and R .

Figure 4(a) shows that if the spot and observer inclinations are both 90° , knowing the inclination of the observer reduces the individual 1σ uncertainties in M and R from $\sim 9\%$ to $\sim 5\%$ and the 3σ uncertainty from $\sim 50\%$ to $\sim 20\%$ for our high background (compare this figure with Figure 2(e)). Figure 4(c) shows that knowing the distance to the star reduces the individual 1σ uncertainties in M and R to $\sim 7\%$ and the 3σ uncertainties to $\sim 30\%$ (again compare this figure with Figure 2(e)). Additional computations show that if the spot is at 80° and the observer's inclination is known, the uncertainties in M and R are similar to the uncertainties obtained if the spot is at 90° . If the spot and

observer inclinations are high and the observer inclination and the distance are both known, the 1σ uncertainties in M and R are similar to the uncertainties obtained if only the observer's inclination is known. The similarity of the 1σ uncertainties in these two cases is not surprising, because independent knowledge of the observer's inclination is usually more helpful than independent knowledge of the distance. If instead the spot and observer inclinations are low, knowing the inclination of the observer or the distance to the star reduces the uncertainties only modestly (compare Figures 4(b) and (d) with Figure 2(f)).

Figures 4(e) and (f) illustrate the effect on the precision of M and R estimates if an atomic scattering line with a known rest energy is observed in the emission from the hot spot. The results shown here are for a line at 6.4 keV with a FWHM of 0.2 keV and an optical depth at the line center of 0.24. If the spot and observer inclinations are high, observing this line tightly constrains the stellar compactness and tightens the constraints on M and R , decreasing the 1σ uncertainty from $\sim 9\%$ to $\sim 4\%$ and the 3σ uncertainty from $\sim 50\%$ to $\sim 20\%$ for our high background, a greater improvement than would be achieved by independently determining the distance and approximately as much as would be produced by independently determining the observer's inclination (compare Figure 4(e) with Figure 2(e)). If instead the spot and observer inclinations are low, observing this scattering line still tightly constrains the stellar compactness but improves the constraints on M and R very little (compare Figure 4(f) with Figure 2(f); for the particular realization shown in Figure 4(f), the 1σ constraints on M and R are weaker with a scattering line present, probably due to a sampling fluctuation).

Figure 5 shows that knowing the size and spectrum of the background greatly improves the constraints on M and R . For our high-inclination case, knowing the background decreases the 1σ uncertainties in M and R from $\sim 9\%$ to $\sim 4\%$ and the 3σ uncertainties from $\sim 50\%$ to $\sim 9\%$ (compare Figure 5(a) with Figure 2(e)). For our low-inclination case, knowing the background leads to a 1σ uncertainty of $\sim 20\%$, whereas without this knowledge one obtains no useful constraints on M and R (compare Figure 5(b) with Figure 2(f)).

Fitting high-quality observations of the X-ray spectra of many bursts from the same star using detailed, high-precision model atmosphere calculations can produce a constraint on M and R that is complementary to that obtained by fitting burst oscillation waveforms. This is illustrated by the dash-dotted red curves in Figures 2(e), 3(a), 3(b), 4(c), 4(e), 4(f), 6(a),

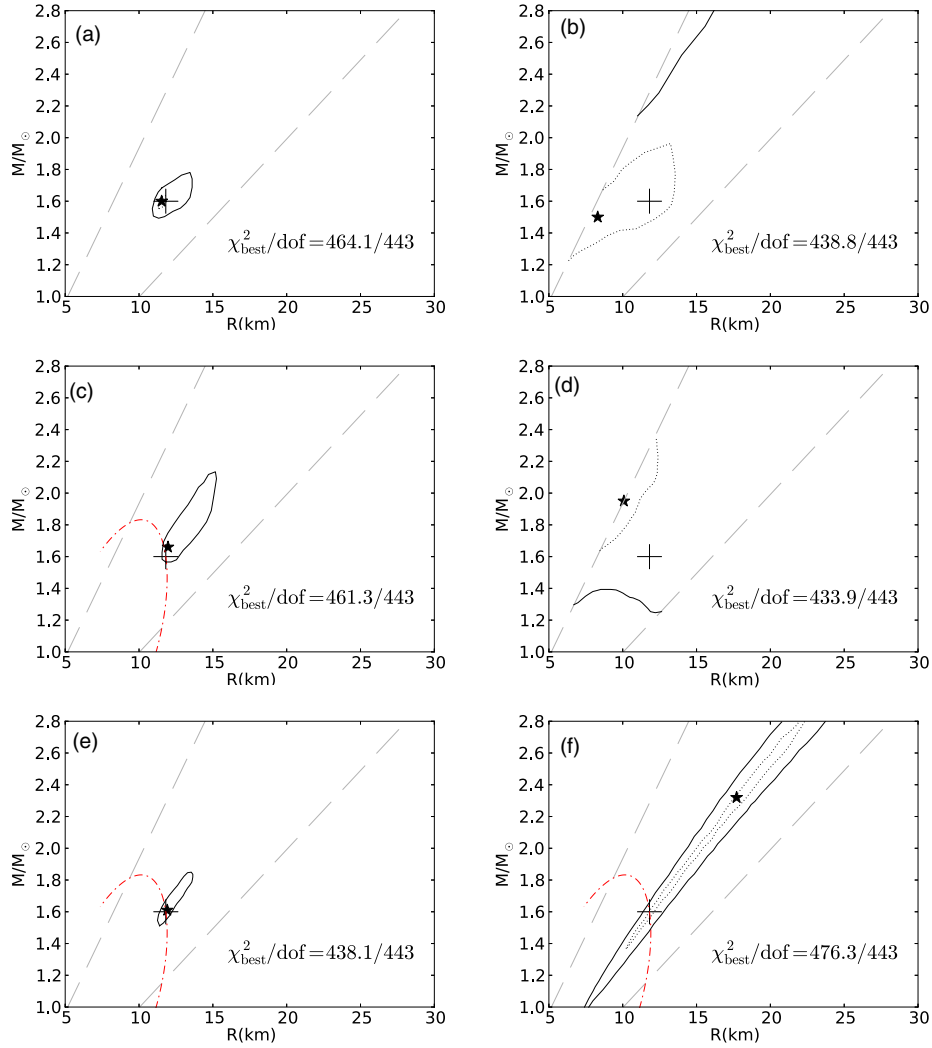


Figure 4. These results show that independent knowledge of some system parameters can significantly tighten the constraints on M and R . All results are for our high background. The contours in the left column are for the high spot and observer inclination case and (top to bottom) known observer inclination, known distance, and an identifiable atomic scattering line in the waveform spectrum (see the text for details); these contours should be compared with those in Figure 2(e). The right column shows the corresponding contours for the low spot and observer inclination case; these contours should be compared with those in Figure 2(f). For the meanings of the line types and symbols, see Figure 2.

(A color version of this figure is available in the online journal.)




Figure 5. These results assume that the properties of the background are known. Left: constraints for the high-inclination reference case (spot and observer inclinations of 90°). Right: constraints for the low-inclination reference case (spot and observer inclinations of 60° and 20° , respectively). Comparison of these contours with those shown in Figures 2(c) and (d), which assume the background is unknown and demonstrate that M and R are much more tightly constrained if the background is known. All these results assume our medium background. For the meanings of the line types and symbols, see Figure 2.




Figure 6. These results show that fitting the observed waveform using a model with an incorrect spot shape increases the uncertainty of the M and R estimates and can sometimes bias them. Left: constraints obtained when the actual emitting area is elongated in the north–south direction, represented here by two uniform, circular spots of angular radius 25° , centered at the same longitude, with inclinations of 40° and 90° , but the waveform is fit assuming a circular spot. Right: constraints obtained when the actual emitting area is in the rotational equator and elongated in the east–west direction, represented here by three uniform, circular spots of angular radius 25° , with centers at the same 90° inclination and spaced 22.5 apart in the azimuthal direction, but the waveform is fit assuming a circular spot. Both results are for our medium background. For the meanings of the line types and symbols, see Figure 2.

(A color version of this figure is available in the online journal.)

and 6(b), which show the relation $(1+z)/g^{2/9} = \text{const.}$ for the values of the redshift z and surface gravity g that correspond to the M and R values used in generating these synthetic waveforms. As discussed in Section 4.1.3, this combination of g and z can be determined by accurately measuring the spectra of the X-ray bursts. The red dash-dotted curve intersects the confidence region derived by fitting burst oscillation waveforms at a high angle. Combining these two complementary methods is therefore a promising way to further constrain M and R .

4.2.6. Effects of Errors in the Assumed Properties of the Hot Spot

In the results discussed previously, we assumed that we knew the shape of the hot spot and the spectrum and beaming pattern of the radiation from the spot. In practice, we may not know these properties. Any discrepancies between the actual properties of the hot spot and the properties assumed in the waveform model will produce systematic errors in the constraints on M and R , in addition to the statistical uncertainties, and these errors should be quantified. A full analysis of possible systematic errors is beyond the scope of this work, but our initial exploration shows that inaccuracies in modeling the waveform can sometimes increase the uncertainties in measurements of M and R .

Errors in the assumed spot shape. The effects of errors in the assumed shape of the hot spot are illustrated in Figure 6. For these examples, we generated synthetic waveforms using hot spots elongated in the east–west and north–south directions by superimposing two or three circular emitting areas, as described in the figure caption, and then fitting the resulting waveforms using our model, which assumes a circular hot spot.

Figure 6(a) shows the results for the moderately high-inclination spot elongated in the north–south direction by 50° and our medium background. The 1σ and 3σ uncertainties in M and R are, respectively, $\sim 40\%$ and $\sim 75\%$, much larger than the $\sim 4\%$ and $\sim 15\%$ uncertainties for the waveform produced by a circular spot in the rotational equator and the same background (see Figure 2(c)). The fit appears excellent ($\chi^2/\text{dof} = 432.5/443$), providing no warning that the model is incorrect. Although the best-fit values of M and R differ substantially from the values assumed in generating the waveform, the differences are only marginally significant.

Figure 6(b) shows the results for the spot in the rotational equator elongated in the east–west direction by 45° and our

medium background. The 1σ uncertainties in M and R are, respectively, $\sim 8\%$ and $\sim 17\%$, substantially larger than the $\sim 4\%$ uncertainties for the waveform produced by a circular spot in the rotational equator and the same background (see Figure 2(c)). The 3σ uncertainties in M and R are, respectively, $\sim 30\%$ and $\sim 40\%$, much larger than the $\sim 15\%$ uncertainties that were achieved for a circular spot (again see Figure 2(c)). The fit is good ($\chi^2/\text{dof} = 442/443$), providing no warning that the model is incorrect. However, in this case the best-fit values of M and R are very close to the values assumed in generating the waveform.

Elongation of the spot in the east–west direction smears only the phase of the waveform and consequently has a relatively small effect on the confidence regions, whereas elongation in the north–south direction smears the crucial (for determining M and R) effects of the Doppler boost and aberration and consequently has a much larger effect. Smearing in the east–west direction may be more likely than smearing in the north–south direction, because hot gas can spread relatively easily in the east–west direction, whereas gas can spread in the north–south direction only if its angular momentum can be changed. If so, elongation of hot spots may not be a major problem for determining M and R using X-ray burst oscillation waveforms.

Errors in the assumed spot spectrum. The effects of errors in the assumed spectrum of the emission from the hot spot are illustrated in Figures 7(a) and (b). For these examples, we generated synthetic waveforms assuming the spectrum of the radiation from the surface of the hot spot has the shape of a Bose–Einstein spectrum and then fitted the waveforms using our model, which assumes that the radiation has the shape of a Planck spectrum.

For the high-inclination reference case (Figure 7(a)), the best-fit values of M and R are very close to their actual values and the quality of the fit is acceptable (there would be a 10% probability of a reduced χ^2 this high if the model were correct). The 1σ uncertainties in M and R are both about 6%, whereas the 3σ uncertainties in M and R are 15% and 12%, respectively. These are comparable to the 3%–4% 1σ and about 15% 3σ uncertainties achieved by fitting a model with the correct spectral shape to the synthetic waveform (see Section 4.2.3).

For the low-inclination reference case (Figure 7(b)), the confidence regions are large, the 1σ region extends beyond the upper boundary of the search domain, and the best-fit M – R

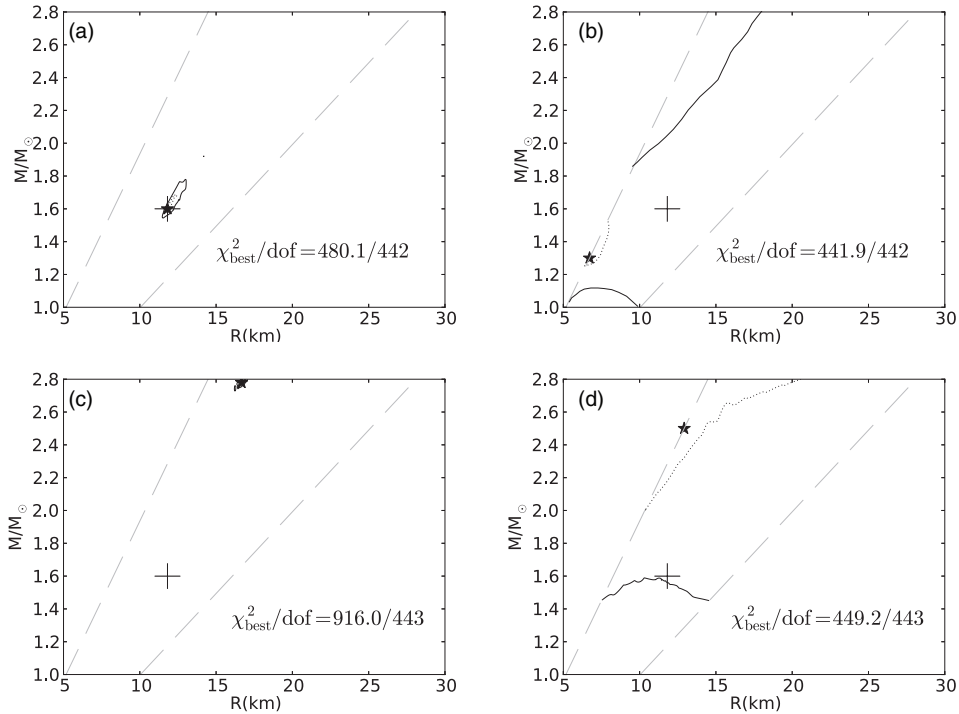


Figure 7. Top panels show that fitting the observed waveform using a model in which the emission spectrum differs by a modest amount from the actual emission spectrum has little effect on the estimated values of M and R or their uncertainties (compare Figures 2(a) and (b)). The bottom panels show that fitting the observed waveform using a model in which the beaming pattern differs substantially from the actual beaming pattern increases the uncertainty of M and R estimates and can sometimes bias them a moderate amount. The contours in the left column are for high spot and observer inclinations, whereas the contours in the right column are for low spot and observer inclinations. All results assume our medium background. The top row shows the constraints obtained when the actual emission spectrum has the shape of a Bose–Einstein spectrum with a chemical potential $\mu = -kT$, but the waveform model assumes the emission spectrum has the shape of a Planck spectrum. The bottom row shows the constraints obtained when the actual emission is isotropic but the waveform model assumes the emission has the beaming pattern described by the Hopf function. For the meanings of the line types and symbols, see Figure 2.

pair is on the boundary, all indicating that the most probable solution has not been determined. The reduced χ^2 value is close to unity and hence does not itself show that the waveform model is wrong, but this is of little consequence because the constraints on M and R are very weak.

Errors in the assumed spot beaming function. The effects of errors in the assumed beaming pattern of the radiation from the hot spot are illustrated in Figures 7(c) and (d). For these examples, we generated synthetic waveforms assuming that the surface of the hot spot radiates isotropically, but then fit these waveforms using our standard model, which assumes the beaming pattern from the surface of the hot spot is described by the Hopf function (the beaming pattern for radiation from a Thomson scattering atmosphere).

For the high-inclination reference case (see Figure 7(c)), the 1σ and 3σ contours are very small and far from the M – R pair used in generating the synthetic waveform, but the reduced χ^2 for this fit is extremely large, showing that the model being fit is wrong.

For the low-inclination reference case (see Figure 7(d)), the best-fit M – R pair is again far from the pair used in generating the synthetic waveform. Although the reduced χ^2 value is close to unity and hence does not itself show that the waveform model is wrong, this is of little consequence, because the constraints on M and R are very weak, the 1σ and 3σ confidence regions extend beyond one or both of the boundaries of the search domain, and the best-fit M – R pair is on the high-compactness boundary, all indicating that the most probable solution has not been determined.

4.2.7. Combining Different Segments of Data

We have seen that for favorable system properties, obtaining strong constraints on M and R usually requires $\sim 10^6$ counts from the hot spot. Acquiring this many counts will typically require collection of hundreds to thousands of seconds of data, even using an instrument with an effective area ~ 10 m². Thus, the required data must be accumulated over many bursts.⁷ We therefore analyzed how the constraints on M and R reported here would be affected if $\sim 10^6$ hot spot counts were collected by combining data from multiple bursts from the same neutron star. The same considerations apply if the data from a single burst are divided into several time segments that are then analyzed jointly.

We find that the constraints on interesting parameters obtained by jointly fitting many data sets are usually comparable to the constraints that could be obtained by fitting a single set of waveform data with the same average profile and the same total number of counts as the waveform obtained by combining the data sets (see Appendix B). Thus, if $\sim 10^6$ counts can be accumulated from the hot spots of many bursts produced by a given star, the prospects are good that constraints on M and R similar to those reported here can be obtained. This is true even though one might intuitively expect that the extra parameters involved in fitting many data sets with changing waveforms

⁷ An exception would be observation of a superburst with oscillations having an rms amplitude of 5%–10%, lasting $\sim 10^3$ s. An amplitude this large would be 5–10 times larger than the 1% amplitude of the oscillations seen during the only superburst in which oscillations have been detected (see Strohmayer & Markwardt 2002).

would compromise the constraints on the model parameters that are fixed (M , R , and θ_{obs}).

The effects on the waveform of the special relativistic Doppler boost and aberration play an important role in constraining M and R . These effects can be accurately captured even if the hot spot rotation frequency is changing, provided it can be tracked accurately enough to maintain the correct oscillation phase when folding successive periods of the oscillation. If this can be done, the constraints on M and R that can be obtained by analyzing the resulting folded waveform are often nearly the same as those that could be obtained by analyzing a similar waveform with a fixed oscillation frequency and the same number of counts; moderate changes in the values of nuisance parameters other than the spot rotation frequency often have only a small effect.

If the oscillation frequency varies so rapidly or erratically during a burst that it cannot be tracked or other important parameters (such as the spot inclination and spot radius) vary sufficiently during a single burst or from burst to burst that this has a substantial effect on the waveform, the full data set may have to be divided into a series of shorter segments and each segment analyzed separately. We find that the resulting constraints on M and R will usually be similar to the constraints obtained by analyzing a single data set in which these nuisance parameters do not vary and that has the same number of counts. Analyzing a sequence of data segments increases the computational burden only linearly with the number of segments, so this is not a major problem. If, however, a waveform changes so rapidly that it must be divided into so many segments that the constraints derived from each segment are minimal, the constraints obtained by combining the results from the many individual segments can become unreliable. If the count rate is large enough, variations of the burst oscillation waveform on timescales shorter than the burst rise or burst tail (but substantially longer than the burst oscillation period) are not an insurmountable analysis problem. If the segments are properly analyzed using a Bayesian approach, the order in which they are analyzed does not matter.

In summary, constraints similar to those reported here can usually be obtained if 10^6 hot spot counts are acquired by combining data from different segments of a single burst or from multiple bursts from the same neutron star, even if the oscillation frequency and other physical parameters vary during the burst or from burst to burst.

5. SUMMARY AND CONCLUSIONS

The goal of this paper was to explore the constraints on the masses and radii of neutron stars that could be inferred by analyzing observations of X-ray burst oscillations made using a future, satellite-borne detector with 2–30 keV energy coverage and an effective area 10 to 20 times larger than the *RXTE* PCA. This research was motivated by the new concepts for large-area X-ray timing space missions that are now being proposed. These include *LOFT* (Mignani et al. 2012; Del Monte et al. 2012) and *AXTAR* (Chakrabarty et al. 2008; Ray et al. 2011). Although we focused here on exploring the constraints on M and R that can be obtained by analyzing the waveforms of X-ray burst oscillations, our results are equally relevant to plans for analyzing the oscillations produced by X-ray emission from the heated polar caps of isolated rotation-powered millisecond pulsars, the goal of the proposed *NICER* mission (Gendreau et al. 2012).

Approach to the problem. We described and explained our approach in Section 2. There we considered the relative merits of using observations of oscillations during burst rises and

burst tails (see Section 2.2). We found that analyzing data collected from burst tails may provide constraints comparable to, or possibly even better than, those obtained by analyzing data collected from burst rises and that analyzing burst-tail data should therefore be considered. We also found that one probably does not need to rely on one or two particularly favorable stars to obtain interesting constraints on M and R .

Computational methods. We described our computational methods in Section 3. We explored the constraints on M and R that could be derived from X-ray burst oscillations by first generating energy-dependent synthetic observed waveforms for a variety of neutron star and hot spot properties. We then used a Bayesian approach and MCMC sampling methods to determine the constraints on M and R that can be obtained by analyzing these synthetic observed waveforms. We computed the joint posterior probability distribution of the parameters in our waveform model, marginalized the parameters other than M and R , and then determined the most probable values of M and R and their confidence intervals.

The purpose of the analysis presented here was not to reproduce all the steps that would be needed for a full Bayesian analysis of a real observed waveform, but rather to determine the precision and accuracy with which such a full analysis could determine M and R . We therefore made use of several shortcuts that reduced the computational burden substantially without significantly altering the results. Using these shortcuts, we could compute the most probable values of M and R and their confidence intervals for a single synthetic observed waveform in 50–100 clock hours, running the parallel version of our code on 150 nodes of a fast CPU cluster.

Assumed counts from the hot spot and background. We assumed that about 10^6 hot spot counts are available from the star. A 10 m^2 detector could collect this many hot spot counts by observing 20–25 bright X-ray bursts (see Section 2.2). We find that the constraints on M and R obtained by combining data from multiple segments of a single burst or from multiple bursts from the same star are usually similar to the constraints that would be obtained by analyzing an observation of a single, time-independent waveform that has the same number of counts and the same shape as the average of the waveform over the multiple data sets (see below, Section 4.2.7, and Appendix B).

We assumed that all sources of background, when combined, contribute about 0.3×10^6 , 10^6 , or 9×10^6 counts. Our treatment of the background was very conservative, in the sense that we usually made no assumptions about the magnitude or spectrum of the background. We did not even assume that the background is constant, only that it does not vary at frequencies commensurate with the burst oscillation frequency.

Our main results. We described our results in detail in Section 4. There we showed that if $\sim 10^6$ counts were collected from the hot spot of a single bursting neutron star and the values of all the parameters in the waveform model other than M and R were known independently of the waveform analysis, the uncertainties in M and R estimates produced by statistical fluctuations in the waveform would typically be only a percent or two (compare the analysis of a similar situation in Strohmayer 2004). For realistic situations in which some of the other parameters in the model must be determined from the waveform, the statistical uncertainties are usually much larger. The reason for this is that the effects on the waveform of changing different parameters in the waveform model are often very similar. These degeneracies of the waveform with respect to changes in the values of the model parameters are

an inherent property of any physical model based on a rotating hot spot and cannot be removed by “improving” the model. These degeneracies inflate the uncertainties in M and R estimates produced by the statistical fluctuations in the waveform.

We first explored how the uncertainties in the measured values of M and R depend on the inclinations of the hot spot and observer and the background count rate. We find that the uncertainties depend strongly on the inclination of the hot spot and observer relative to the stellar rotation axis. Specifically,

1. if the hot spot is within 10° of the rotational equator, both M and R can usually be determined with an uncertainty of about 10%;
2. if instead the spot is within 20° of the rotation pole, the uncertainties in M and R are so large that waveform measurements alone provide no useful constraints;
3. the uncertainties in M and R are affected little by background count rates less than or comparable to the count rate from the hot spot, but become significantly larger for higher background count rates.

Next, we explored the effects on the uncertainties in M and R if the distance to the star or the inclination of the observer are known from other measurements, if a resonance scattering line is observed in the burst oscillation spectrum, or if the properties of the background are independently known. We find that

1. independent information about the background can greatly reduce the uncertainties in estimates of M and R ;
2. independent knowledge of the observer’s inclination can also greatly reduce the uncertainties;
3. knowledge of the star’s distance can help, but not as much as knowledge of the background or the observer’s inclination;
4. observation of an identifiable atomic line in the hot-spot emission always tightly constrains M/R ; it can also tightly constrain M and R individually, if the hot spot is within about 30° of the rotation equator.

Finally, we explored the effects of deviations in the actual shape of the hot spot, radiation beaming pattern, and spectrum from those assumed in the our waveform model. We find that

1. modest deviations of the actual spectrum from that assumed in the waveform model have little effect on the accuracy or uncertainty of M and R estimates;
2. large deviations of the actual radiation beaming pattern from the pattern assumed in the waveform model can increase the uncertainties of M and R estimates substantially;
3. in some cases, but not always, large deviations of the actual shape of the hot spot from the circular shape assumed in the waveform model can increase the uncertainties of M and R estimates and bias them by moderate amounts. The physical conditions that produce tight constraints on M and R (relatively small spots far from the rotation pole) are the conditions in which the shape of the spot is unimportant (see Lamb et al. 2009a, 2009b).

Combining data from different bursts. Our results show that for favorable system properties, strong constraints on M and R can be obtained if one can analyze $\sim 10^6$ counts from the hot spot. This will typically require analyzing hundreds or thousands of seconds of data, even if the detector has an effective area ~ 10 m², and hence will require combining data from many bursts, unless a truly exceptional superburst is observed (see Section 4.2.7). We therefore analyzed how the constraints on M and R reported here would be affected if the $\sim 10^6$ hot spot

counts were obtained by combining data from multiple bursts from the same star. The same considerations apply if the data from a single burst are divided into several time segments and then analyzed jointly. We find that the constraints on M and R obtained by jointly fitting many data sets are usually comparable to the constraints that could be obtained by fitting a single segment of waveform data having the same profile and as the average profile of the combined data sets and the same total number of counts (see Section 4.2.7 and Appendix B).

Time varying waveforms. We also investigated the problem of analyzing burst oscillation waveforms that change with time. If the oscillation frequency changes during a burst but the change can be modeled accurately enough to maintain the correct oscillation phase when folding successive periods of the oscillation, the constraints on M and R obtained by analyzing the resulting folded waveform will be nearly the same as those that could be obtained by analyzing a similar waveform with a fixed oscillation frequency and the same number of counts. Even if the oscillation frequency varies too rapidly or irregularly during the burst rise or tail to be described accurately by a simple frequency model, the full burst oscillation data set can be divided into smaller time segments and analyzed using standard Bayesian techniques. This approach can also be used if other physical properties of the system, such as the size and inclination of the emitting region, vary significantly. The computational burden of this kind of analysis increases only linearly with the number of segments. Consequently, variations of the burst oscillation waveform on timescales shorter than the burst rise or burst tail (but substantially longer than the burst oscillation period) do not appear to pose an insurmountable analysis problem (again see Section 4.2.7 and Appendix B).

We thank Anthony Chan, Stephen Drake, José Garmilla, Aaron Hanks, John Hoffman, Miaotianzi Jin, David Kolschowsky, and Dave Kotan for assistance in developing and testing the ray-tracing code and Novarah Kazmi for help in assembling the references. We also thank Stratos Boutloukos, Juri Poutanen, and Tod Strohmayer for many useful discussions. This research was supported in part by the Simons Foundation (grant No. 230349) to M.C.M. and NSF grant AST0708424 at Maryland, and by the Fortner Endowed Chair, the Physics Department, and NSF grant AST0709015 at Illinois. M.C.M. gratefully acknowledges the hospitality of the Aspen Center for Physics, where part of this work was completed. M.C.M. also thanks the Department of Physics and Astronomy at Johns Hopkins University for their hospitality during his sabbatical.

APPENDIX A CODE TESTS

In this appendix, we summarize the suite of test problems we routinely solve to validate both our waveform code and our MCMC code (of which the waveform code is an essential part). We performed these tests not because the results are controversial, but rather to show that the codes used here correctly compute the results we present. We have been motivated to perform and document these tests by the importance of the results presented in this paper for the objectives of future large-area timing missions such as *LOFT* (Mignani et al. 2012; Del Monte et al. 2012) and *AXTAR* (Chakrabarty et al. 2008; Ray et al. 2011), and by the fact that the studies performed by the code testing subgroup of the LOFT Dense Matter Science Working Group, which were led by J. Poutanen, showed that all but two

of the codes used to produce recently published hot-spot waveforms had significant coding errors and that only one of the current hot-spot waveform codes that were initially submitted for testing by the subgroup was free of coding errors.

The waveform code used in the present work is based on the code developed and validated by Lamb et al. (2009a). We described some tests of this code in that paper. The waveform code used in the present work passed the code validation tests of the code testing subgroup of the LOFT Dense Matter Science Working Group. The code validation tests we describe here are designed to check many of the physical effects included in our waveform and MCMC codes, by using them to solve problems where the answer is known analytically. This suite of test problems includes the following problems.

1. Computing the deflection angle as a function of the star's compactness and the direction of the ray relative to the normal to the stellar surface (Appendix A.1.1).
2. Computing the light travel time as a function of the star's compactness and the direction of the ray relative to the normal to the stellar surface (Appendix A.1.2).
3. Computing the bolometric flux seen by a distant observer located directly above the center of an emitting spot on a non-rotating star (Appendix A.1.3).
4. Computing the radiation pattern produced by a small emitting spot on a non-rotating star in flat spacetime (Appendix A.1.4).
5. Computing the waveform produced by a small emitting spot on a rotating star in flat spacetime (Appendix A.1.5).
6. Computing the angular deflection of a pencil beam from a small emitting spot on a rotating star (Appendix A.1.6).
7. Computing the observed temperature of thermal emission from a small emitting spot on a rotating star in flat spacetime, as a function of rotational phase (Appendix A.1.7).
8. Determining the stellar compactness, spot temperature, and distance, given uniform thermal emission from the entire surface of a non-rotating star with no background (Appendix A.2.1).
9. Determining the spot inclination, observer inclination, and distance, given thermal emission from a small hot spot on a non-rotating star in flat spacetime with no background (Appendix A.2.2).
10. Determining the stellar compactness, spot temperature, and distance, given thermal emission from a small spot on a non-rotating star in flat spacetime with a background (Appendix A.2.3).

The values of the waveform code integration, hot spot, and angular resolution parameters that we use for the code tests reported here are the values listed in Table 5, unless otherwise noted. We use these same parameter values when computing M and R confidence regions (see Section 4.2), except that for that purpose we set N_{phase} to 16, which speeds up the computations while providing an accuracy that is sufficient for the waveform-fitting process. As discussed in Appendix A.3, we have carefully chosen the values of these code parameters to provide a resolution fine enough to meet our accuracy requirements, but no finer, so that our code runs as fast as possible.

We now discuss the suite of validation tests in detail.

Table 5
Values of the Waveform Code Resolution Parameters

Parameter	Value
N_{lat} (spot grid points in stellar latitude)	100
N_{long} (spot grid points in stellar longitude)	100
N_{alpha} (grid points in α)	1000
N_{phase} (grid points in phase w. time delays included)	1000
N_{loc} (grid points in phase w.o. time delays included)	1001
n_d (grid points used in computing the deflection angle)	100
ϵ_d (deflection angle integration parameter)	0.1
n_t (grid points used in computing the time delay)	100
ϵ_0 (first time delay integration parameter)	0.01
ϵ_1 (second time delay integration parameter)	0.01

A.1. Tests of the Waveform Code

A.1.1. Deflection Angle as a Function of Compactness and Direction

In the Schwarzschild spacetime, the total angle ψ by which a light ray from the stellar surface is eventually deflected is a function only of the initial angle α of the ray relative to the normal to the stellar surface (as measured in the static frame) and the compactness M/R of the star, and is given by the expression (see Pechenick et al. 1983, Equation (2.13) for a similar expression)

$$\psi\left(\alpha, \frac{M}{R}\right) = \int_0^1 \frac{\sin \alpha dy}{\sqrt{(1 - 2M/R) - (1 - 2My/R)y^2 \sin^2 \alpha}}, \quad (\text{A1})$$

where $y \equiv R/r$. The integrand in Equation (A1) diverges as $x \rightarrow 1$ and $\sin \alpha \rightarrow 1$. After our paper had been submitted, J. Poutanen (2013, private communication) kindly called our attention to the fact that these divergences can be removed by using the variable substitution $x = \sqrt{1 - y}$, which changes the integral (A1) to the form (J. Poutanen et al., in preparation)

$$\psi\left(\alpha, \frac{M}{R}\right) = 2 \frac{\sin \alpha}{\sqrt{1 - u}} \int_0^1 \frac{xdx}{\sqrt{\cos^2 \alpha + x^2 q}}, \quad (\text{A2})$$

where $u = 2M/R$ and $q = [2 - x^2 - u(1 - x^2)^2/(1 - u)] \sin^2 \alpha$.

In our waveform code, we integrate expression (A1) numerically, which requires special treatment as $y \rightarrow 1$ and $\sin \alpha \rightarrow 1$. We break the integral into two pieces, from $y = 0$ to $y = 1 - \epsilon_d$ and from $y = 1 - \epsilon_d$ to $y = 1$, with $\epsilon_d \ll 1$. Our code computes the first piece numerically, using Simpson's rule with n_d divisions, and computes the second piece using the approximate analytical expression obtained by expanding the square root in the integrand of expression (A1) in powers of ϵ_d and dropping terms $\mathcal{O}(\epsilon_d^3)$ and higher.

We test the deflection angles ψ given by our waveform code by comparing them with the values of ψ given by evaluating integral (A1) using Mathematica 8. Figure 8 shows the fractional difference $[(\psi_{\text{code}} - \psi_{\text{mathematica}})/\psi_{\text{mathematica}}]$ for three values of n_d and two values of the stellar compactness. Figure 9 shows the fractional difference for three values of ϵ_d and two values of the stellar compactness. The left and right panels of Figure 10 show, respectively, the relation $\psi(\alpha)$ and the fractional difference in ψ given by our code using $n_d = 100$ and $\epsilon_d = 0.1$, for four values of the stellar compactness that span the range considered in this work. Figures 8–10 show that the values $n_d = 100$ and $\epsilon_d = 0.1$ of these resolution parameters that we use in our




Figure 8. Fractional difference between the deflection angle ψ as a function of the initial ray direction α computed using our code and using Mathematica 8 for three different numbers n_d of integration steps and two values of the stellar compactness (left: $R/M = 5.0$; right: $R/M = 4.0$). This figure shows that the value $n_d = 100$ used in our waveform fitting is adequate. See Appendix A.1.1 for further details. (A color version of this figure is available in the online journal.)




Figure 9. Fractional difference between the deflection angle ψ as a function of the initial ray direction α computed using our code and using Mathematica 8 for three values of the integration parameter ϵ_d and two values of the stellar compactness (left: $R/M = 5.0$; right: $R/M = 4.0$). This figure shows that the value $\epsilon_d = 0.1$ used in our waveform fitting is adequate. See Appendix A.1.1 for further details. (A color version of this figure is available in the online journal.)




Figure 10. Left: deflection angle ψ as a function of the initial ray direction α computed using our code for four values of the stellar compactness. The deflection angle is larger for smaller values of R/M , i.e., for more compact stars. Right: fractional difference between the deflection angle ψ computed using our code and using Mathematica 8, as a function of α , for the same four values of the stellar compactness. The fractional difference is larger for smaller values of R/M , i.e., for more compact stars, but is smaller in magnitude than 2.5×10^{-4} for all these values of R/M . This figure shows that the resolution parameter values $n_d = 100$ and $\epsilon_d = 0.1$ used here provide sufficient accuracy for our waveform fitting. See Appendix A.1.1 for further details. (A color version of this figure is available in the online journal.)

waveform code when fitting synthetic waveform data provide sufficient accuracy for this purpose. We also test our waveform code by comparing the values of ψ given by our code using these values of the resolution parameters with the values of ψ given by Equation (A2). The fractional difference in ψ is always less than 4×10^{-4} for $R/M = 4.0$, the most extreme compactness we consider, while the fractional difference in the derivative

of ψ is always less than 10^{-3} and is much less than this for $\sin \alpha < 0.95$.

A.1.2. Light Travel Time as a Function of Compactness and Direction

The time it takes for a light ray emitted at an angle α relative to the surface normal (in the static frame) to travel from the




Figure 11. Fractional difference between the differential light travel time Δt as a function of the initial ray direction α computed using our code and using Mathematica 8, for three different numbers n_t of integration steps and two values of the stellar compactness (left: $R/M = 5.0$; right: $R/M = 4.0$). This figure shows that the resolution $n_t = 100$ used in our waveform fitting is adequate. See Appendix A.1.2 for further details.

(A color version of this figure is available in the online journal.)

surface to infinity, as measured at infinity, is given by

$$t = \int_0^1 dy \frac{R/c}{y^2(1-2My/R)} \times \left[\frac{1}{\sqrt{1-\sin^2\alpha(1-2M/y)^{-1}(1-2My/R)y^2}} \right]. \quad (\text{A3})$$

This expression accounts for the gravitational time delay in the Schwarzschild spacetime as well as the purely geometrical effect that the travel time is different for rays emitted from the stellar surface at different angles. The travel time given by Equation (A3) is formally infinite, because the ray is traced to infinity. However, what we are interested in is not the total travel time of each ray, but rather the *differential* travel times of different rays.

To be able to compute the differential travel times of any two rays, we calculate the differential travel time of every ray relative to the travel time of a reference ray. We choose as our reference ray the radial ray coming from the point on the stellar surface directly under the observer. The differential light travel time of an arbitrary ray compared to the reference ray is then

$$\Delta t \equiv t - [t]_{\alpha=0} = \int_0^1 dy \frac{R/c}{y^2(1-2My/R)} \times \left[\frac{1}{\sqrt{1-\sin^2\alpha(1-2M/y)^{-1}(1-2My/R)y^2}} - 1 \right]. \quad (\text{A4})$$

The integrand in this expression can diverge at the two limits of the integral. After our paper had been submitted, J. Poutanen (2013, private communication) kindly called our attention to the fact that these divergencies can be removed by the same variable substitution used in transforming Equation (A1) to Equation (A2), which changes the integral (A4) to the form (J. Poutanen et al., in preparation)

$$\Delta t = 2 \frac{R \sin^2 \alpha}{c(1-u)} \int_0^1 \frac{xdx}{\sqrt{x^2q + \cos^2\alpha(1 + \sqrt{x^2q + \cos^2\alpha})}}. \quad (\text{A5})$$

In our waveform code, we calculate the integral in Equation (A4) in three pieces. Near the upper endpoint, we

set $y = 1 - \epsilon_1$, with $\epsilon_1 \ll 1$, expand the expression inside the square root in powers of ϵ_1 , and drop the terms that are $\mathcal{O}(\epsilon_1^2)$ or higher. Near the lower endpoint of the integral, we set $y = \epsilon_0$, with $\epsilon_0 \ll 1$, expand the integrand in powers of ϵ_0 , keeping only terms up to $\mathcal{O}(\epsilon_0^4)$. Our waveform code computes the middle piece of the integral, from $y = \epsilon_0$ to $y = 1 - \epsilon_1$, numerically, using Simpson's rule with n_t divisions in this interval.

We test the differential light travel time Δt given by our waveform code by comparing it with the value of Δt given by evaluating integral (A4) using Mathematica 8. Figure 11 shows the fractional difference $[(\Delta t_{\text{code}} - \Delta t_{\text{mathematica}})/\Delta t_{\text{mathematica}}] \times 100\%$ for three values of n_t and two values of the stellar compactness. Figure 12 shows the fractional difference for three values of ϵ_0 and ϵ_1 and two values of the stellar compactness. The left and right panels of Figure 13 show, respectively, the relation $\Delta t(\alpha)$ and the fractional difference given by our code using $n_t = 100$ and $\epsilon_0 = \epsilon_1 = 0.01$, for four values of the stellar compactness that span the range considered in this work. Figures 11–13 show that the values $n_t = 100$ and $\epsilon_0 = \epsilon_1 = 0.01$ of these resolution parameters that we use in our waveform code when fitting synthetic data provide sufficient accuracy for this purpose. We also test our waveform code by comparing the values of Δt given by our code using these value of the resolution parameters with the values of Δt given by Equation (A5). The fractional difference in Δt is always less than 4×10^{-4} for $R/M = 4.0$, the most extreme compactness we consider.

A.1.3. Bolometric Flux Above the Center of an Emitting Spot on a Non-rotating Star

We test the accuracy of our waveform code's computation of the gravitational redshift and deflection angle and its integration over the area of the hot spot by comparing the *bolometric* flux that it predicts will be seen by an observer positioned on the radial vector from the center of a non-rotating star through the center of the spot, as a function of the compactness of the star and the angular radius of the spot, with the analytic result for this flux.

Consider a uniform circular spot of angular radius $\Delta\theta$ on the surface of a non-rotating star with compactness M/R , emitting radiation with an isotropic beaming pattern. In the Schwarzschild spacetime, the bolometric flux seen by a distant




Figure 12. Fractional difference between the differential light travel time Δt as a function of the initial ray direction α computed using our code and using Mathematica 8 for three different values of ϵ_0 and ϵ_1 and two values of the stellar compactness (left: $R/M = 5.0$; right: $R/M = 4.0$). This figure shows that the values $\epsilon_0 = \epsilon_1 = 0.01$ used in our waveform fitting are adequate. See Appendix A.1.2 for further details.

(A color version of this figure is available in the online journal.)




Figure 13. Left: differential light travel time Δt as a function of the initial ray direction α computed using our code for four values of the stellar compactness. The differential light travel time is smaller for smaller values of R/M , i.e., for more compact stars. Right: fractional difference between the light travel time Δt as a function of α computed using our code and using Mathematica 8 for the same four values of the stellar compactness. The fractional difference is larger for smaller values of R/M , i.e., for more compact stars, but is less than 1.5×10^{-2} for all these values of R/M . This figure shows that the resolution parameter values $n_t = 100$ and $\epsilon_0 = \epsilon_1 = 0.01$ used in our waveform fitting are adequate. We have assumed $M = 1.4 M_\odot$ to obtain R from R/M . See Appendix A.1.2 for further details.

(A color version of this figure is available in the online journal.)

observer positioned above the center of the spot is

$$F = \frac{I_0}{(1+z)^2} \left(\frac{R}{d}\right)^2 \int_0^{2\pi} d\phi \int_{\cos\Delta\theta}^1 \cos\alpha \left(\frac{d\cos\alpha}{d\cos\psi}\right) d\cos\theta, \quad (\text{A6})$$

where I_0 is the specific intensity measured at the stellar surface, $1+z = (1-2M/R)^{-1/2}$ is the redshift, R is the radius of the star, d is the distance to the star, α is the angle that a ray initially makes with the surface normal, as measured in the static frame, and $\psi = \psi(\alpha, M/R)$ is the deflection angle. In the limit of vanishing M/R , $z \rightarrow 0$, $\cos\alpha \rightarrow \cos\theta$, $d\cos\alpha/d\cos\psi \rightarrow 1$, and we recover the familiar expression

$$\begin{aligned} F_0 &= I_0 \left(\frac{R}{d}\right)^2 \int_0^{2\pi} d\phi \int_{\cos\Delta\theta}^1 \cos\theta d\cos\theta \\ &= \pi I_0 \left(\frac{R}{d}\right)^2 \sin^2\Delta\theta. \end{aligned} \quad (\text{A7})$$

A relationship between the angle θ of the emission point, measured from the symmetry axis of the spot, and the angle α a ray makes with the local surface normal is established by the requirement that such a ray must be deflected by $\psi = \theta$ in order

to be seen by the observer. In other words,

$$\theta = \int_0^1 \frac{\sin\alpha dx}{\sqrt{(1-2M/R) - (1-2Mx/R)x^2 \sin^2\alpha}}. \quad (\text{A8})$$

The condition $\psi = \theta$ allows us to simplify the expression for the flux considerably, namely,

$$\begin{aligned} F &= \frac{I_0}{(1+z)^2} \left(\frac{R}{d}\right)^2 \int_0^{2\pi} d\phi \int_{\cos\Delta\theta}^1 \cos\alpha \left(\frac{d\cos\alpha}{d\cos\theta}\right) d\cos\theta \\ &= \frac{I_0}{(1+z)^2} \left(\frac{R}{d}\right)^2 \int_0^{2\pi} d\phi \int_{\cos\alpha(\Delta\theta)}^1 \cos\alpha d\cos\alpha \\ &= \frac{\pi I_0}{(1+z)^2} \left(\frac{R}{d}\right)^2 \sin^2\alpha(\Delta\theta) \end{aligned} \quad (\text{A9})$$

where $\alpha(\Delta\theta)$ is the value of α for which $\psi(\alpha) = \Delta\theta$.

We perform this test of our code as follows. For a given value of M/R and a range of values of α , we compute the value of ψ that corresponds to each value of α . Next, for each $\psi(\alpha)$ (as long as $\psi < \pi$), we use the relation $\Delta\theta = \psi(\alpha)$ to construct the relation $\alpha(\Delta\theta)$ and then use this value of α in expression (A9) to determine the bolometric flux that is predicted by this analytical expression for each value of $\Delta\theta$. On the other hand, we use the angular radius $\Delta\theta(\alpha)$ to compute the bolometric flux for each




Figure 14. Scaled bolometric flux (left) and fractional error in the bolometric flux (right) seen by a distant observer as a function of $\Delta\theta$. The agreement between the results given by our waveform code (red dots) and the analytic results (blue curve) is excellent. These results are for $R/M = 4.0$, $M = 1.4 M_{\odot}$, and $\theta_{\text{spot}} = \theta_{\text{obs}} = 0$. See Appendix A.1.3 for further details.

(A color version of this figure is available in the online journal.)

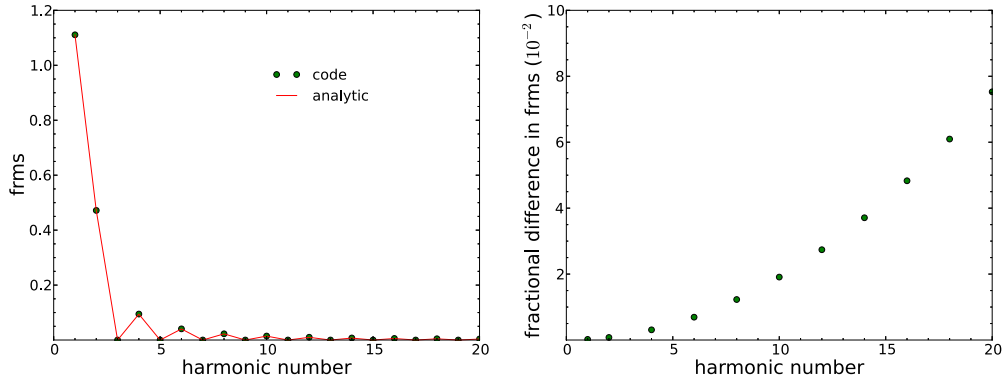


Figure 15. Fractional rms amplitudes of the first 20 harmonics of the test waveform, showing good agreement between the results given by our waveform code and the analytic results. Left: amplitudes given by our waveform code (green dots) compared with the analytic amplitudes (red curve). Right: fractional differences between the amplitudes given by our waveform code and the analytic amplitudes. Only the differences for harmonic components with nonzero analytic amplitudes are shown. See Appendix A.1.4 for further details.

(A color version of this figure is available in the online journal.)

value of α , using our waveform code. We then compare these two values of the flux at each value of $\Delta\theta$. Figure 14 shows an example of such a comparison. The results given by the code differ from the results predicted by the analytical expression by less than 5×10^{-4} for $\Delta\theta > 5^\circ$. Even for $\Delta\theta < 5^\circ$, the difference is less than 7×10^{-3} .

A.1.4. Radiation Pattern Produced by a Small Spot on a Non-rotating Star in Flat Spacetime

We test the accuracy with which the waveform code computes the radiation pattern from the star by using it to compute the flux seen by a distant observer moving around a non-rotating star in flat spacetime that has a very small spot that emits isotropically. We then compare this result with the analytic result for this situation. In addition to testing the computation of the radiation pattern, this comparison also tests the accuracy of the fast Fourier transform that we use to represent waveforms.

The radiation pattern for a star and hot spot with the assumed properties is

$$f(\phi) = \max\{\cos \phi, 0\}, \quad (\text{A10})$$

where ϕ is the phase of the waveform (equivalent to the azimuth of the distant observer, for a non-rotating star). The Fourier representation of this waveform is

$$f(\phi) = c_0 + \sum_{n=1}^{\infty} c_n \cos[n(\phi - \phi_n)], \quad (\text{A11})$$

where

$$\begin{aligned} c_0 &= \frac{1}{\pi} \\ c_1 &= \frac{1}{2} \\ c_n &= \begin{cases} 0, & n \text{ is odd } (n > 1) \\ \frac{2}{\pi} \frac{1}{n^2-1}, & n \text{ is even} \end{cases} \end{aligned} \quad (\text{A12})$$

and

$$\begin{aligned} \phi_1 &= 0 \\ \phi_n &= \begin{cases} \text{indeterminate (since } c_n = 0), & n \text{ is odd} \\ 0, & n = 2, 6, 10, \dots \\ \pi/n, & n = 4, 8, 12, \dots \end{cases} \end{aligned} \quad (\text{A13})$$

Figures 15 and 16 compare these analytic predictions for the fractional rms amplitudes and the phases of the first 20 harmonics with the results for these quantities given by our waveform code for $R/M = 10^6$, $M = 1.4 M_{\odot}$, $\theta_{\text{spot}} = 60^\circ$, $\Delta\theta = 1^\circ$, and $\theta_{\text{obs}} = 20^\circ$ and the values of the resolution parameters listed in Table 5. These figures show that the results given by our waveform code are in good agreement with the analytic predictions.

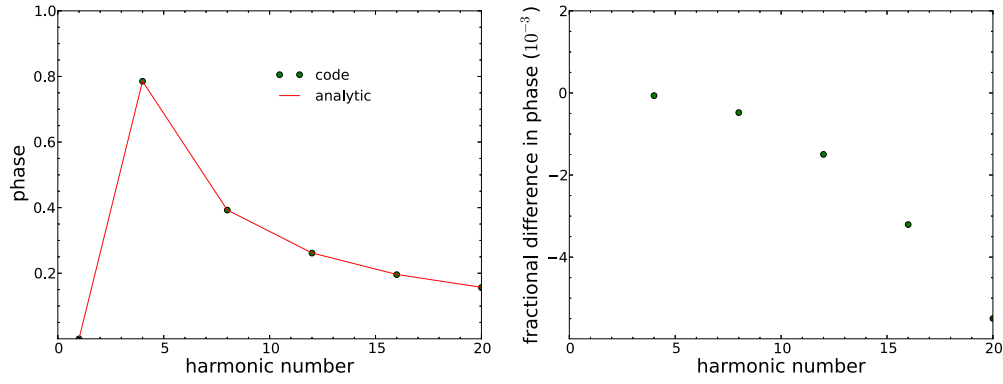


Figure 16. Phases of the first 20 harmonics of the test waveform, showing good agreement between the results given by our waveform code and the analytic results. Left: phase of the n^{th} harmonic component of the waveform as a function of harmonic number n (green dots) compared with the analytic phases (red curve). Right: fractional differences between the phases given by our waveform code and the analytic phases. Only the results for harmonic components with nonzero analytic amplitudes are shown. See Appendix A.1.4 for further details.

(A color version of this figure is available in the online journal.)

A.1.5. Waveform Produced by a Small Spot on a Rotating Star in Flat Spacetime

We test the accuracy with which our waveform code computes the Doppler boost and the photon travel time by using it to compute the bolometric waveform produced by a small spot on a rotating star in flat spacetime and comparing it with the analytic waveform for this case.

The analytic expression for this waveform is

$$f(\phi) \propto \max\{\delta^4 \cos \alpha, 0\}, \quad (\text{A14})$$

where ϕ is the observed phase (i.e., after correction for the time delay) and

$$\delta(\phi) = \frac{1}{\gamma[1 + (v/c) \sin \theta_{\text{spot}} \sin \theta_{\text{obs}} \sin \Phi]} \quad (\text{A15})$$

is the Doppler factor. Here v is the linear velocity at the rotational equator and $\Phi \equiv \phi - \Omega \Delta t$ is the phase prior to correction for time delay in terms of the angular rotation frequency Ω of the hot spot and the time delay

$$\Delta t = \frac{R}{c}(1 - \cos \alpha), \quad (\text{A16})$$

where α is the angle that the ray makes with the local normal to the stellar surface, as measured in the static frame, and is related to other angles via

$$\cos \alpha = \cos \theta_{\text{spot}} \cos \theta_{\text{obs}} + \sin \theta_{\text{spot}} \sin \theta_{\text{obs}} \cos \Phi. \quad (\text{A17})$$

To compute the flux at each phase ϕ analytically, one must simultaneously solve for Φ and $\cos \alpha$ using this set of equations.

Figure 17 shows that the bolometric waveform given by our waveform code for this case is in good agreement with the analytic result, both when the time delay set to zero and when it is included.

A.1.6. Angular Deflection of a Pencil Beam from a Small Spot on a Rotating Star

We test the accuracy with which our waveform code computes the relativistic aberration and gravitational light-bending effects by computing the rotational phase at which a beam of radiation emitted from the stellar surface in a given direction would be

seen by a distant observer and comparing this result with the phase given by Mathematica 8 (the “predicted” phase). If the star is not rotating, the rotational phase is undefined. In this case what is computed is the azimuth at which the beam of radiation would be seen by a distant observer.

In these tests, we assume that the radiation is emitted in a narrow “pencil” beam from a very small ($\Delta\theta = 1^\circ$) spot located on the star’s rotational equator ($\theta_{\text{spot}} = 90^\circ$) and that the observer is in the plane defined by the star’s rotational equator ($\theta_{\text{obs}} = 90^\circ$). Because of the spherical symmetry of the Schwarzschild spacetime, the radiation will remain in the equatorial plane as it propagates. The phase (azimuth) at which the radiation is seen by a distant observer depends on the angle α'_0 of the pencil beam relative to the local normal to the stellar surface, as measured in the corotating frame, the linear velocity v/c of the hot spot produced by its rotation, which determines the amount of relativistic aberration, and the compactness M/R of the star, which determines the amount of light bending.

We first use a Lorentz transformation to determine the angle α_0 of the ray in the static frame that corresponds to the angle α'_0 of the ray in the comoving frame. We then compute the angular deflection of this ray using Mathematica, which allows us to determine the predicted phase (azimuth) at which the beam would be seen by a distant observer. Next we use our waveform code to compute the phase (azimuth) at which such a beam would be seen by a distant observer. To do this we set the beaming function $g(\alpha')$ equal to 1 for $|\alpha' - \alpha'_0| < \sigma_{\alpha'_0}$ and equal to 0 for $|\alpha' - \alpha'_0| > \sigma_{\alpha'_0}$, with $\sigma_{\alpha'_0} = 0.01$. In order to test the computation of the aberration and the gravitational bending separately from the computation of the time delay, we set the light travel time to be zero for all rays. The phase (azimuth) at which the beam would be seen by a distant observer is then equal to the total angular deflection of the ray produced by the relativistic aberration and light bending. For simplicity, we trace only rays that propagate in the prograde direction relative to the rotation of the hot spot. Finally, we compare the phase (azimuth) given by the waveform code to the phase predicted using Mathematica.

Tables 6–8 show the results for the computed and predicted phases (or azimuths) at which the pencil beam would be seen by a distant observer. These results are for emission from the surface of a rapidly rotating star in flat spacetime, from the surface of a non-rotating relativistic star, and from the surface of a rapidly rotating relativistic star, for a range of stellar




Figure 17. Top left: bolometric flux waveform produced by a small spot on the surface of a rotating star in flat spacetime computed using our waveform code (green dots) and the corresponding analytic result (red curve) with the time delay set to zero. Top right: fractional differences (green curve) between the fluxes given by our code and the analytic fluxes. Bottom panels: same as in the top panels, but with the time delay included. These results are for $R/M = 10^6$, $M/M_\odot = 1.6$, $\theta_{\text{spot}} = 60^\circ$, $\theta_{\text{obs}} = 20^\circ$, $\Delta\theta = 1^\circ$, $\nu = 0.0141357027464$ Hz. See Appendix A.1.5 for further details.

(A color version of this figure is available in the online journal.)




Figure 18. Left: temperature of the Planck spectrum seen by a distant observer as a function of the star’s rotational phase computed using our waveform code (red dots) and the analytic expression (blue curve). Right: fractional difference (green curve) between the observed temperature computed using our waveform code and the analytic results are in excellent agreement. These results are for a star with a rotation rate such that the linear velocity at the rotational equator is $v/c = 0.7$, $R/M = 10^6$, for $M = 1.4 M_\odot$, $\theta_{\text{spot}} = 60^\circ$, $\theta_{\text{obs}} = 20^\circ$, and $\Delta\theta = 1^\circ$. See Appendix A.1.7 for further details.

(A color version of this figure is available in the online journal.)

compactnesses and hot spot rotation rates. All these results are for $M = 1.4 M_\odot$.

A.1.7. Observed Temperature of Thermal Emission from a Rotating Star in Flat Spacetime, as a Function of Rotational Phase

We test the accuracy with which our waveform code computes the relativistic Doppler effect by computing the observed temperature of thermal emission from the surface of a rotating star in flat spacetime as a function of the rotational phase and then comparing the result with the analytic result for the observed temperature as a function of phase.

For simplicity, we consider isotropic emission from a small spot at an inclination θ_{spot} on the stellar surface, with a Planck spectrum having a temperature T_0 as measured in the comoving

frame. Then the spectrum seen by a distant observer with an inclination θ_{obs} is a Planck spectrum with temperature

$$T(\phi) = \frac{T_0}{\gamma[1 - (v_0/c) \sin \theta_{\text{spot}} \sin \theta_{\text{obs}} \sin \phi]}. \quad (\text{A18})$$

Also for simplicity, we set the light travel time to zero. The observed phase is then the azimuthal position of the spot relative to the azimuth of the observer. To test a range of Doppler boosts, we choose values of the stellar rotational frequency that produce a linear velocity at the rotational equator of $0.1c$, $0.4c$, and $0.7c$. Figure 18 shows the comparison of the temperature between the code and the analytic result, for the case of $v/c = 0.7$. It is clear that the agreement is excellent.

Table 6
Comparison of Pencil-beam Phases for Different Surface Velocities^a

α_0	v/c	Phase (Code)	Phase (Predicted)
15°	0.1	0.958	0.958
30°	0.1	0.917	0.917
	0.4	0.917	0.917
45°	0.1	0.875	0.875
	0.4	0.875	0.875
	0.7	0.877	0.875
60°	0.1	0.833	0.833
	0.4	0.833	0.833
	0.7	0.833	0.833
75°	0.1	0.792	0.792
	0.4	0.791	0.792
	0.7	0.791	0.792
90°	0.1	0.751	0.750
	0.4	0.751	0.750
	0.7	0.751	0.750

Note. ^a Comparison of the computed and predicted phases at which a narrow beam of radiation emitted from a small spot on the equator of a rotating star at an angle α_0 relative to the local surface normal (as measured in the static frame) is seen by a distant observer who is also located in the equatorial plane. In this test, $R/M = 10^6$, so the spacetime is essentially flat. The stellar rotation frequency is chosen so the linear velocity of the stellar surface in the rotation equator is $0.1c$, $0.4c$, or $0.7c$. The values of v/c that have been omitted from the table would produce beams that point in the retrograde direction in the comoving frame, a geometry not considered in this test. See Appendix A.1.6 for further details.

Table 7
Comparison of Pencil-beam Azimuths for Different Stellar Compactnesses^a

α_0	R/M	Azimuth (Code)	Azimuth (Predicted)
0°	4.0	0.998	1.000
	5.0	0.998	1.000
	6.0	0.999	1.000
15°	4.0	0.941	0.941
	5.0	0.946	0.946
	6.0	0.948	0.949
30°	4.0	0.881	0.881
	5.0	0.891	0.892
	6.0	0.897	0.897
45°	4.0	0.818	0.818
	5.0	0.835	0.836
	6.0	0.844	0.845
60°	4.0	0.751	0.751
	5.0	0.777	0.777
	6.0	0.790	0.790
75°	4.0	0.674	0.674
	5.0	0.713	0.713
	6.0	0.733	0.733
90°	4.0	0.579	0.576
	5.0	0.642	0.640
	6.0	0.671	0.669

Note. ^a Comparison of the computed and predicted azimuths at which a narrow beam of radiation emitted from a small spot on the equator of a non-rotating star at an angle α_0 relative to the local surface normal (as measured in the static frame) would be seen by a distant observer who is also located in the equatorial plane, for three values of the stellar compactness. See Appendix A.1.6 for details.

A.2. Tests of the Waveform Fitting Code

In this subsection, we outline the tests we use to validate the set of codes that fit model waveforms to waveform data. An important part of this code set is the code that computes

Table 8
Comparison of Pencil-beam Azimuths for Different Stellar Compactnesses and Surface Velocities^a

α_0	R/M	v/c	Phase (Code)	Phase (Predicted)
30°	3.5	0.4	0.870	0.871
	6.0	0.4	0.897	0.897
45°	3.5	0.4	0.801	0.802
		0.7	0.802	0.802
	6.0	0.4	0.845	0.845
		0.7	0.790	0.790
60°	3.5	0.4	0.725	0.725
		0.7	0.725	0.725
	6.0	0.4	0.790	0.790
		0.7	0.790	0.790
75°	3.5	0.4	0.632	0.633
		0.7	0.632	0.633
	6.0	0.4	0.733	0.733
		0.7	0.732	0.733
90°	3.5	0.4	0.498	0.495
		0.7	0.498	0.495
	6.0	0.4	0.671	0.669
		0.7	0.670	0.669

Note. ^a Comparison of the computed and predicted phases at which a narrow beam of radiation emitted from a small spot on the equator of a rotating star at an angle α_0 relative to the local surface normal (as measured in the static frame) is seen by a distant observer who is also located in the equatorial plane, for two values of the stellar compactness. The stellar rotation frequency is chosen so the linear velocity of the stellar surface in the rotation equator is either $0.4c$ or $0.7c$. The values of v/c that have been omitted from the table would produce beams that point in the retrograde direction in the comoving frame, a geometry not considered in this test. See Appendix A.1.6 for further details.

the likelihood of an observed waveform, given a set of model parameters. The accuracy of this step depends crucially on the accuracy with which the model waveforms are computed. We have shown in the previous subsection that our code computes model waveforms with high accuracy. Here we discuss three relatively simple end-to-end tests that we use to validate the waveform fitting process.

A.2.1. Determining the Compactness, Temperature, and Distance Given Uniform Emission from a Non-rotating Star with no Background

In this exercise, we test our model-fitting code by fitting a model of emission from a non-rotating star to the synthetic spectrum produced by a non-rotating star that emits radiation uniformly over its entire surface, with a Planck spectrum. We assume there is no background. The observed spectrum is therefore a Planck spectrum with temperature

$$T_{\text{obs}} = T_{\text{co}} \left(1 - \frac{2M}{R} \right)^{1/2}, \quad (\text{A19})$$

where T_{co} is the temperature of the emission measured in the comoving frame at the stellar surface. In this case, the goodness of the fit is determined entirely by the shape and normalization of the model spectrum. If the distance is not independently known, the normalization of the spectrum can be adjusted by changing the distance. Hence the likelihood of the data given the model has the same value (its maximum value) for all $T_{\text{co}}-M/R$ pairs related by Equation (A19). Consequently, any change in the surface temperature of the emission in the model being fit can be perfectly compensated by appropriate changes in the compactness M/R of the model star and its distance. This is one of the strongest degeneracies in the waveform fitting problem.




Figure 19. Code test illustrating the strong degeneracy between the temperature T_{co} of the emission in the comoving frame and the stellar compactness M/R . Left: all parameters in the model except T_{co} and M/R were fixed in advance at the values used to generate the synthetic observed spectrum. The plotted points show the values of T_{co} and R/M that give log likelihoods within 1.15 (i.e., 1σ for 2 degrees of freedom) of the log likelihood given by the best-fit model. The models that give high likelihoods cluster tightly around $T_{\text{co}} = 2.0$ keV and $R/M = 5.0$, the values used in generating the synthetic observed spectrum. Right: all parameters except T_{co} , R/M , and d were fixed in advance at the values used to generate the synthetic spectrum. The plotted points show the values of the parameters T_{co} and R/M that give log likelihoods within 1.73 (i.e., 1σ for 3 degrees of freedom) of the log likelihood given by the best-fit model. In this case, the models that give high likelihoods cluster tightly along the curve given by Equation (A19) with T_{obs} equal to the observed value, since the normalization of each model can be freely adjusted by varying the distance to the star. The results shown here are for $M = 1.6 M_{\odot}$, $\Delta\theta = 180^{\circ}$, $d = 35$ kpc, $\approx 10^6$ counts from hot spot, and no background. See Appendix A.2.1 for further details.

(A color version of this figure is available in the online journal.)

This strong degeneracy can be broken for the synthetic data considered in this example if the mass M of the star and the distance d to the star are both known independently of the fitting process, because knowledge of these two quantities fixes the normalization of the observed spectrum. To see this, note that the normalization of the spectrum is proportional to A/d^2 , where A is the projected area of the emitting region. In the case considered here, $A = \pi R^2$ and is related to the stellar compactness and the mass through $R = (R/M) \times M$. Hence if M and d are both known independently, the temperature and flux of the observed radiation determines the flux at the stellar surface. This then fixes R and hence M/R , since M is assumed to be known. This result is illustrated in the left panel of Figure 19. If, however, d is not known, the normalization of the spectrum from the stellar surface is unknown and can be freely adjusted, even if M is known. In this case, we expect that all models having T_0 and M/R values that fall on the curve given by Equation (A19) will give good fits. The right-hand panel of Figure 19 illustrates this.

A.2.2. Determining the Spot Inclination, Observer Inclination, and Distance, Given Thermal Emission from a Small Spot on a Non-rotating Star in Flat Spacetime with no Background

In this exercise, we test our model fitting code by fitting a model of emission from a non-rotating star in flat spacetime to the synthetic flux produced by a non-rotating star that emits radiation from a small spot, with a Planck spectrum. We assume that all the parameters in the model are independently known except the inclination of the spot and the observer, or the inclination of the spot and the observer and the distance to the star.

In flat spacetime, if $\theta_{\text{spot}} + \theta_{\text{obs}} < \frac{\pi}{2}$ (so that the spot is not occulted by the star), the flux seen by an observer at the stellar azimuth ϕ is

$$f(\phi) = \cos \theta_{\text{spot}} \cos \theta_{\text{obs}} + \sin \theta_{\text{spot}} \sin \theta_{\text{obs}} \cos \phi. \quad (\text{A20})$$

This radiation pattern is a purely sinusoidal function of ϕ and can therefore be uniquely specified by the mean flux and the fractional rms amplitude of the flux variation with ϕ . The mean flux is

$$f_{\text{mean}} = \cos \theta_{\text{spot}} \cos \theta_{\text{obs}}, \quad (\text{A21})$$

while the fractional rms amplitude of the flux variation with ϕ is

$$f_{\text{rms}} = (1/\sqrt{2}) \tan \theta_{\text{spot}} \tan \theta_{\text{obs}}. \quad (\text{A22})$$

For a given fractional rms amplitude, Equation (A22) defines a curve in the $\theta_{\text{spot}}-\theta_{\text{obs}}$ plane along which the amplitude remains constant.

If the distance to the star is independently known, the model will provide a good fit to the data only close to two points on the curve given by Equation (A22). One of these is the point where θ_{spot} and θ_{obs} have the values that were used to generate the synthetic observed radiation pattern. The other is the point where the values of these two inclinations are interchanged, because the values of f_{mean} and f_{rms} are not changed by interchanging the two inclinations. Other points on the curve are strongly disfavored, because the value f_{mean} at these points will be very different from the observed value. If the distance is not independently known and must be included as a parameter in the fit, good fits can be found all along the curve, because the mean observed flux can be adjusted by varying the distance. These results are illustrated by the numerical results shown in Figure 20.

A.2.3. Determining the Compactness, Temperature, and Distance, Given Thermal Emission from a Small Spot on a Non-rotating Star in Flat Spacetime with a Background

This test is the same as the test discussed in Appendix A.2.1, except that here we add a background that is uniform in azimuth and has a thermal spectrum. In these fits, we treat the background in one of two ways.

1. We add to each model spectrum exactly the background that was used in generating the synthetic observed spectrum (i.e., we do not determine the background as part of the fitting procedure), or
2. we determine the background as part of the fitting procedure, using the code module that implements the procedure described in Section 3.3.3.

As subcases, we either set the distance to the value used in generating the synthetic observed spectrum during the fitting process (i.e., we assume that the distance is independently known), or we include the distance as one of the parameters




Figure 20. Code test illustrating the strong degeneracy between the spot and observer inclinations θ_{spot} and θ_{obs} . Left: all parameters in the model except θ_{spot} and θ_{obs} were fixed in advance at the values used to generate the synthetic radiation pattern. The plotted points show the values of θ_{spot} and θ_{obs} that give log likelihoods within 1.15 (i.e., 1σ for 2 degrees of freedom) of the log likelihood given by the best-fit model. The models that give high likelihoods cluster tightly around two points, one corresponding to the values of $\theta_{\text{spot}} = 20^\circ$ and $\theta_{\text{obs}} = 60^\circ$ used in generating the synthetic observed radiation pattern and the other given by interchanging the two values of θ_{spot} and θ_{obs} , which leaves the mean flux, f_{mean} , and the rms variation of the flux with ϕ , f_{rms} , unchanged (see Equations (A21) and (A22)). Right: all parameters in the model except θ_{spot} and θ_{obs} and d were fixed in advance at the values used to generate the synthetic radiation pattern. The plotted points show the values of θ_{spot} and θ_{obs} that give log likelihoods within 1.73 (i.e., 1σ for 3 degrees of freedom) of the log likelihood given by the best-fit model. In this case, the models that give high likelihoods cluster tightly along the curve given by Equation (A22), where f_{mean} and f_{rms} remain unchanged, but are spread along the curve, because the normalization of each model can be freely adjusted by varying the distance to the star. The results shown here are for $R/M = 10^6$ (to mimic flat space), $M = 1.6 M_\odot$, $\theta_{\text{spot}} = 60^\circ$, $\theta_{\text{obs}} = 20^\circ$, $\Delta\theta = 0.01$ radians, $T_{\text{co}} = 2.0$ keV, $d = 6000$ kpc (to produce the correct X-ray flux, given the large stellar radius), $\approx 10^6$ counts from the hot spot, and no background. See Appendix A.2.2 for further details.

(A color version of this figure is available in the online journal.)




Figure 21. Code test illustrating the increased uncertainty in the fitted parameters induced by the additional fluctuations in the number of counts produced by a background component in the observed flux. Left: all parameters in the model except T_{co} and M/R were fixed in advance at the values used to generate the synthetic observed spectrum. The plotted points show the values of T_{co} and R/M that give log likelihoods within 1.15 (i.e., 1σ for 2 degrees of freedom) of the log likelihood given by the best-fit model. The models that give high likelihoods cluster around $T_{\text{co}} = 2.0$ keV and $R/M = 5.0$, the values used in generating the synthetic observed spectrum, but not nearly as tightly as when there is no background. Right: all parameters except T_{co} , R/M , and d were fixed in advance at the values used to generate the synthetic spectrum. The plotted points show the values of the parameters T_{co} and R/M that give log likelihoods within 1.73 (i.e., 1σ for 3 degrees of freedom) of the log likelihood given by the best-fit model. In this case, the models that give high likelihoods cluster along the curve given by Equation (A19) with T_{obs} equal to the observed value, since the normalization of each model can be freely adjusted by varying the distance to the star. Again, they are not nearly as tightly clustered as when there is no background. The results shown here are for $M = 1.6 M_\odot$, $\Delta\theta = 180^\circ$, $d = 35$ kpc, $\approx 10^6$ counts from the hot spot, $T_{\text{bkg}} = 1.5$ keV, and $\approx 9 \times 10^6$ counts from the background. See Appendix A.2.3 for further details.

(A color version of this figure is available in the online journal.)

to be determined in the fitting process. As explained before, assuming that the distance is independently known amounts to fixing the normalization of the spectrum, which more strongly constrains the fit.

Figure 21 shows that the increase in the total number of counts due to the additional counts contributed by the background, increases the uncertainties in T_{co} and M/R . In producing the fits shown in this figure, we added to each model spectrum exactly the background that was used in generating the synthetic observed spectrum. The left panel in the figure shows the fit assuming that the distance to the star is independently known, whereas the right panel shows the fit obtained if the distance is determined as part of the fitting procedure. Compared to the case without any background (see Figure 19), the scatter around the expected relationship between T_{co} and M/R is larger when a background is present than when there is no background, even

though the exact background has been added to each model. The reason for the increased scatter is the increase in the *fluctuation* in the total number of counts, due to the counts contributed by the background (the fluctuation scales as the square root of the total number of counts).

A.3. Convergence Tests

We wish to use values for the waveform code resolution parameters listed in Table 5 that will provide a resolution fine enough to meet our accuracy requirements, but no finer, so that our code runs as fast as possible. In order to determine these values of the resolution parameters, we perform a series of convergence tests using different values of these parameters. In these tests, we compare bolometric waveforms computed for different values of the resolution parameters, assuming the




Figure 22. Percentage differences between three bolometric waveforms computed using three different values of the spot resolution parameters N_{lat} and N_{long} and a reference waveform computed using $N_{\text{lat}} = N_{\text{long}} = 100$, the value we use in our waveform analyses and the code tests described in this appendix. This figure shows that for the values of these resolution parameters that we use, the fractional error in the bolometric waveform is $\lesssim 3 \times 10^{-4}$. See Appendix A.3 for further details.

(A color version of this figure is available in the online journal.)




Figure 23. Percentage differences between three bolometric waveforms computed using three different values of the ray direction resolution parameter N_{α} and a reference waveform computed using $N_{\alpha} = 1000$, the value we use in our waveform analyses and the code tests described in this appendix. This figure shows that for the values of these resolution parameters that we use, the fractional error in the bolometric waveform is $\lesssim 10^{-4}$. See Appendix A.3 for further details.

(A color version of this figure is available in the online journal.)

following values of the physical parameters: $R/M = 5.0$, $M/M_{\odot} = 1.6$, $\nu_{\text{rot}} = 600$ Hz, $\theta_{\text{spot}} = 90^{\circ}$, $\Delta\theta = 25^{\circ}$, and $\theta_{\text{obs}} = 20^{\circ}$. We use the values of the resolution parameters listed in Table 5, except that for clarity in the figures we plot the flux at only 16 values of the waveform phase.

Figure 22 shows that our choice of 100 grid points in latitude and longitude ($N_{\text{lat}} = N_{\text{long}} = 100$) resolves the hot spot well enough that the bolometric flux as a function of phase has a fractional error $\lesssim 3 \times 10^{-4}$, which is adequate for our purposes. Figure 23 shows that 1000 grid points in the angle α' between the ray and the local normal ($N_{\alpha} = 1000$) resolves this angle well enough that the bolometric flux has a




Figure 24. Percentage differences between three bolometric waveforms computed using three different values of N_{loc} , the grid spacing in the waveform phase with time delays included, and a reference waveform computed using $N_{\text{loc}} = 1001$, the value we use in the code tests described in this appendix. This figure shows that the fractional error in the bolometric waveform for this value of N_{loc} is $\lesssim 10^{-5}$. See Appendix A.3 for further details.

(A color version of this figure is available in the online journal.)

fractional error $\lesssim 10^{-4}$, adequate for our purposes. Figure 24 shows that 1001 grid points in the final value of the waveform phase ($N_{\text{loc}} = 1001$) resolves the phase well enough that the bolometric flux has a fractional error $\lesssim 10^{-5}$, which is more than adequate for our purposes. Indeed, $N_{\text{loc}} = 101$ provides essentially the same accuracy.

The tests of the values of the resolution parameters used in computing the deflection angle and the light travel time have been presented in Figures 8, 9, 11, and 12.

APPENDIX B

JOINT FITS

In this Appendix, we discuss the constraints on the waveform model parameters that can be obtained by jointly fitting a collection of different, independent sets of waveform data. These sets could be different segments of a single burst or different bursts from the same star.

We show here that the constraints on the interesting parameters that can be obtained by jointly fitting many data sets are often comparable to the constraints that would have been obtained by fitting a single data set that has an oscillation profile which is the same as the average oscillation profile of the multiple data sets and has the same total number of counts as the multiple data sets have when combined. This is true even though one might intuitively expect that the extra parameters required to fit many data sets in which the uninteresting parameters change from set to set would compromise the constraints on the (unchanging) interesting parameters.

We assume that the values of some of the parameters in the model waveform, such as the mass M and radius R of the star, the inclination θ_{obs} of the observer, and the distance d to the star remain unchanged during a burst and from burst-to-burst, whereas the values of other parameters, such as the angular radius $\Delta\theta$ of the hot spot, the inclination θ_c of the spot center, and the color temperature T_{co} of the emission from the spot, may change during a burst and from burst-to-burst. The parameters

that remain unchanged must be treated differently from the parameters that may change.

Suppose that we have several independent data sets and wish to extract joint constraints on the subset of waveform parameters α_i that have the same values in all the data sets. We denote the other waveform parameters, which may change between data sets, by β_j . We can derive the correct approach to jointly fitting multiple data sets by starting from the Bayesian expression for the unnormalized posterior probability density

$$q(\alpha_i) = \int d\beta_j P(\alpha_i, \beta_j) \mathcal{L}(\alpha_i, \beta_j). \quad (\text{B1})$$

Here P is the joint prior probability density over all the parameters and \mathcal{L} is the likelihood of all the data, given the model and specific values for all the model parameters. If there are multiple data sets, then the joint likelihood is the product of the individual likelihoods, i.e.,

$$\mathcal{L} = \prod_k \mathcal{L}_k(\alpha_i, \beta_{jk}), \quad (\text{B2})$$

where here we write β_{jk} to indicate the set of potentially variable parameters that is being evaluated for data set k .

We assume that the prior probability density distribution for the fixed parameters α_i is the same for each data set. Then we can write the full prior probability distribution as

$$P(\alpha_i, \beta_{j1}, \beta_{j2}, \dots) = P(\alpha_i) P_1(\beta_{j1}|\alpha_i) P_2(\beta_{j2}|\alpha_i) \dots \quad (\text{B3})$$

where the product is over all data sets and the vertical bar indicates a conditional probability. For example, $P_1(\beta_{j1}|\alpha_i)$ is the prior probability of the variable parameters β_{j1} in the first data set given that the fixed parameters have the values α_i .

The posterior probability density of our fixed parameters is thus proportional to

$$q(\alpha_i) \propto P(\alpha_i) \prod_k \left[\int d\beta_{jk} P_k(\beta_{jk}|\alpha_i) \mathcal{L}_k(\alpha_i, \beta_{jk}) \right]. \quad (\text{B4})$$

Thus the posterior probability density is equal to the prior probability density times the product of the integrals over the variable parameters for each data set, of the prior probability density of the variable parameters times the likelihood of the data given the values of all the parameters. The process of integrating over the uninteresting parameters (called nuisance parameters in this context) is called marginalization. Marginalization over the nuisance parameters means that in some sense they do not count in the tally of parameters that are adjusted during the fitting process. It is this part of the procedure that ultimately yields constraints from many data sets that are not muddled by the introduction of many fit parameters.

It is important to note that the posterior probability density is usually *not* the product of the marginalized probability densities $Q_k(\alpha_i)$ for each of the data sets. That is,

$$q(\alpha_i) \neq \prod_k \left[\int d\beta_{jk} P(\alpha_i) P(\beta_{jk}|\alpha_i) \mathcal{L}_k(\alpha_i, \beta_{jk}) \right]. \quad (\text{B5})$$

For n data sets, this incorrect expression for the posterior probability density would yield

$$q(\alpha_i) \propto [P(\alpha_i)]^n \prod_k \left[\int d\beta_{jk} P_k(\beta_{jk}) \mathcal{L}_k(\alpha_i, \beta_{jk}) \right]. \quad (\text{B6})$$

In this incorrect expression for n data sets, there are n factors of the prior probability density, compared with only one in the correct expression. To see that n factors is incorrect, imagine a case in which the prior probability distribution is a multivariate Gaussian and that only one of a large number n of data sets is informative. The incorrect expression would nevertheless produce very tight constraints, because it raises the Gaussian to a high power.

The posterior probability density (B4) is not in general proportional to the posterior probability density one would obtain from an equivalent number of counts in a profile that has the average shape of the profile in the multiple data sets. To see this, let us perform a thought experiment in which we have obtained a single, continuous segment of data during which all the waveform parameters are known to be fixed, but we analyze it as some number of contiguous data sets where the values of the parameters α_i are fixed but the values of the β_{jk} need not be. Then the joint analysis gives Equation (B4), whereas the analysis that assumes all the parameters are fixed gives

$$q(\alpha_i) \propto P(\alpha_i) \int d\beta_j P(\beta_j|\alpha_i) \mathcal{L}(\alpha_i, \beta_j), \quad (\text{B7})$$

where as before $\mathcal{L} = \prod_k \mathcal{L}_k(\alpha_i, \beta_j)$. Note that instead of allowing the β_j to take on different values β_{jk} in each data set, here they are assumed to have the same values. In general, Equation (B7) does not give the same posterior as Equation (B4).

There are, however, circumstances in which the posteriors are identical. Consider a situation in which (1) the prior for the variable parameters is independent of the fixed parameters, and (2) the likelihood is the product of two independent functions, one for α_i and one for β_{jk} :

$$\mathcal{L}_k(\alpha_i, \beta_{jk}) = \mathcal{L}_{k\alpha} \mathcal{L}_{k\beta}. \quad (\text{B8})$$

Then

$$q(\alpha_i) \propto P(\alpha_i) \prod_k \left[\int d\beta_{jk} P_k(\beta_{jk}) \mathcal{L}_k(\alpha_i, \beta_{jk}) \right] = [P(\alpha_i) \prod_k \mathcal{L}_{k\alpha}] \prod_k \left[\int d\beta_{jk} P_k(\beta_{jk}) \mathcal{L}_{k\beta} \right]. \quad (\text{B9})$$

The second factor in Equation (B9) is independent of α and therefore enters the posterior as a constant, which will be normalized away. The product $\prod_k \mathcal{L}_{k\alpha}$ is just the total likelihood, and the posterior will therefore be unchanged by dividing the original single data set into multiple data sets.

In order to investigate whether this general result applies to the problem studied in this paper, namely, analysis of multiple segments of burst oscillation data, we have carried out a number of tests. In these tests we compared the constraints obtained by analyzing a set of waveform data as a single segment and as a sequence of segments. We find that in practice the constraints obtained by jointly analyzing multiple data sets are usually comparable to the constraints obtained by analyzing a single data set with an oscillation profile similar to the average profile of the oscillations in the multiple data sets and with the same total number of counts as there are in the multiple data sets.

We illustrate these results by the following two examples. Both compare the constraints on M and R obtained by analyzing a single data set that has about 10^6 counts from the hot spot with the constraints obtained by breaking up the single data set into five segments of equal length and then analyzing the segments jointly, using the method described above. All the confidence regions shown here are for the high-inclination reference case and our medium background. Figure 25 compares




Figure 25. Example showing that the constraints on M and R obtained by jointly analyzing multiple data sets as described in the text are usually comparable to the constraints obtained by analyzing a single data set with an oscillation profile similar to the average profile of the oscillations in the multiple data sets and with the same total number of counts as there are in the multiple data sets. The results shown here are for the high-inclination reference case with our medium background and no independent knowledge of the parameters in the model. Left: constraints obtained by analyzing a single data set with about 10^6 counts from the hot spot. Right: constraints obtained by breaking up the single data set into five segments of equal length and then analyzing the segments jointly.




Figure 26. Another example showing that the constraints on M and R obtained by jointly analyzing multiple data sets as described in the text are usually comparable to the constraints obtained by analyzing a single data set with an oscillation profile similar to the average profile of the oscillations in the multiple data sets and with the same total number of counts as there are in the multiple data sets. The results shown here are for the high-inclination reference case with our medium background and independent knowledge of the distance to the star. Left: constraints obtained by analyzing a single data set with about 10^6 counts from the hot spot. Right: constraints obtained by breaking up the single data set into five segments of equal length and then analyzing the segments jointly.

the constraints obtained when no independent information about the parameters in the model is available. Figure 26 compares the constraints obtained when the distance to the star is known independently of the waveform analysis. In both cases, the constraints on M and R obtained by jointly analyzing multiple data sets as described in the text are comparable to the constraints obtained by analyzing a single data set with an oscillation profile similar to the average profile of the oscillations in the multiple data sets and with the same total number of counts as there are in the multiple data sets.

REFERENCES

- Altamirano, D., Watts, A., Linares, M., et al. 2010, *MNRAS*, **409**, 1136
- Artigue, R., Barret, D., Lamb, F., Lo, K. H., & Miller, M. C. 2013, *MNRAS*, **433**, L64
- Bhattacharyya, S., & Strohmayer, T. E. 2005, *ApJL*, **634**, L157
- Bhattacharyya, S., & Strohmayer, T. E. 2006, *ApJL*, **642**, L161
- Bhattacharyya, S., Strohmayer, T. E., Miller, M. C., & Markwardt, C. B. 2005, *ApJ*, **619**, 483
- Bogdanov, S. 2013, *ApJ*, **762**, 96
- Bogdanov, S., Grindlay, J. E., & Rybicki, G. B. 2008, *ApJ*, **689**, 407
- Bogdanov, S., Rybicki, G. B., & Grindlay, J. E. 2007, *ApJ*, **670**, 668
- Boutloukos, S., Miller, M. C., & Lamb, F. K. 2010, *ApJL*, **720**, L15
- Cackett, E. M., Miller, J. M., Ballantyne, D. R., et al. 2010, *ApJ*, **720**, 205
- Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, *ApJ*, **654**, 458
- Casares, J. 2010, in *Highlights of Spanish Astrophysics V*, ed. J. M. Diego, L. J. Goicoechea, J. I. González-Serrano, & J. Gorgas (Berlin: Springer), 3
- Chakrabarty, D., Morgan, E. H., Munro, M. P., et al. 2003, *Natur*, **424**, 42
- Chakrabarty, D., Ray, P. S., & Strohmayer, T. E. 2008, in *AIP Conf. Proc. 1068, A Decade of Accreting Millisecond X-Ray Pulsars*, ed. R. Wijnands, D. Altamirano, P. Soleri, N. Degenaar, N. Rea, P. Casella, A. Patruno, & M. Linares (Melville, NY: AIP), 227
- Chandrasekhar, S. 1960, *Radiative Transfer* (New York: Dover)
- Del Monte, E., Donnarumma, I., & Consortium, L. 2012, *MmSAI*, **83**, 352
- Egron, E., di Salvo, T., Burderi, L., et al. 2011, *A&A*, **530**, A99
- Eyer, L., Holl, B., Pourbaix, D., et al. 2013, *Central European Astrophysical Bulletin*, **37**, 115
- Feroci, M., Stella, L., Vacchi, A., et al. 2010, *Proc. SPIE*, **7732**, 38
- Galloway, D. K., Munro, M. P., Hartman, J. M., Psaltis, D., & Chakrabarty, D. 2008, *ApJS*, **179**, 360
- Gendreau, K. C., Arzoumanian, Z., & Okajima, T. 2012, *Proc. SPIE*, **8443**, 13
- Goldman, I. 1979, *A&A*, **78**, L15
- Grindlay, J., Gursky, H., Schnopper, H., et al. 1976, *ApJL*, **205**, L127
- Güver, T., Özel, F., Cabrera-Lavers, A., & Wroblewski, P. 2010a, *ApJ*, **712**, 964
- Güver, T., Wroblewski, P., Camarota, L., & Özel, F. 2010b, *ApJ*, **719**, 1807
- Hebeler, K., Lattimer, J. M., Pethick, C., & Schwenk, A. 2010, *PhRvL*, **105**, 1102
- Hoshi, R. 1981, *ApJ*, **247**, 628
- Lamb, F. K., & Boutloukos, S. 2008, in *Short-Period Binary Stars: Observations, Analyses, and Results*, ed. E. F. Milone, D. A. Leahy, & D. W. Hobill (Astrophysics and Space Science Library, Vol. 352; Berlin: Springer), 87
- Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009a, *ApJ*, **706**, 417
- Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009b, *ApJL*, **705**, L36
- Lattimer, J. M. 2007, *AP&SS*, **308**, 371
- Lattimer, J. M., & Prakash, M. 2007, *PhR*, **442**, 109
- Leahy, D. A., Morsink, S. M., & Cadeau, C. 2008, *ApJ*, **672**, 1119
- Leahy, D. A., Morsink, S. M., Chung, Y., & Chou, Y. 2009, *ApJ*, **691**, 1235
- Lewin, W. H. G., Hoffman, J. A., Doty, J., et al. 1976, *MNRAS*, **177**, 83
- Lewin, W. H. G., van Paradijs, J., & Taam, R. E. 1993, *SSRv*, **62**, 223

- Majczyna, A., & Madej, J. 2005, *AcA*, **55**, 349
- Marshall, H. L. 1982, *ApJ*, **260**, 815
- Mignani, R. P., Zane, S., Walton, D., et al. 2012, in IAU Symp. 285, *New Horizons in Time-Domain Astronomy*, ed. E. Griffin, R. Hanisch, & R. Seaman (Cambridge: Cambridge Univ. Press), 372
- Miller, M. C. 1999, *ApJL*, **515**, L77
- Miller, M. C. 2000, *ApJ*, **531**, 458
- Miller, M. C., Boutloukos, S., Lo, K. H., & Lamb, F. K. 2011, in Proc. of Fast X-Ray Timing and Spectroscopy at Extreme Count Rates, ed. D. Barret, M. Méndez, & Stéphane Paltani, *PoS(HTRS 2011)024*
- Miller, M. C., Boutloukos, S., Lo, K. H., & Lamb, F. K. 2013, in IAU Symp. 290, *Feeding Compact Objects: Accretion on All Scales*, ed. C. M. Zhang, T. Belloni, M. Méndez, & S. N. Zhang (Cambridge: Cambridge Univ. Press), 101
- Miller, M. C., & Lamb, F. K. 1996, *ApJ*, **470**, 1033
- Miller, M. C., & Lamb, F. K. 1998, *ApJL*, **499**, L37
- Muno, M. P., Chakrabarty, D., Galloway, D. K., & Psaltis, D. 2002a, *ApJ*, **580**, 1048
- Muno, M. P., Fox, D. W., Morgan, E. H., & Bildsten, L. 2000, *ApJ*, **542**, 1016
- Muno, M. P., Özel, F., & Chakrabarty, D. 2002b, *ApJ*, **581**, 550
- Muno, M. P., Özel, F., & Chakrabarty, D. 2003, *ApJ*, **595**, 1066
- Nath, N. R., Strohmayer, T. E., & Swank, J. H. 2002, *ApJ*, **564**, 353
- Özel, F. 2006, *Natur*, **441**, 1115
- Özel, F., Baym, G., & Güver, T. 2010, *PhRvD*, **82**, 101301
- Özel, F., Güver, T., & Psaltis, D. 2009, *ApJ*, **693**, 1775
- Özel, F., & Psaltis, D. 2009, *PhRvD*, **80**, 103003
- Paczynski, B., & Anderson, N. 1986, *ApJ*, **302**, 1
- Pechenick, K. R., Ftaclas, C., & Cohen, J. M. 1983, *ApJ*, **274**, 846
- Poutanen, J., & Beloborodov, A. M. 2006, *MNRAS*, **373**, 836
- Poutanen, J., & Gierliński, M. 2003, *MNRAS*, **343**, 1301
- Poutanen, J., & Gierliński, M. 2004, *NuPhS*, **132**, 640
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C* (2nd ed.; Cambridge: Cambridge Univ. Press)
- Ray, P. S., Phlips, B. F., Wood, K. S., et al. 2011, in *Fast X-Ray Timing and Spectroscopy at Extreme Count Rates*, ed. D. Barret, M. Méndez, & Stéphane Paltani, *PoS(HTRS 2011)07*
- Read, J. S., Lackey, B. D., Owen, B. J., & Friedman, J. L. 2009, *PhRvD*, **79**, 124032
- Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, *ApJ*, **722**, 33
- Strohmayer, T. E. 1999, *ApJL*, **523**, L51
- Strohmayer, T. E. 2001, *AdSpR*, **28**, 511
- Strohmayer, T. E. 2004, in AIP Conf. Proc. 714, *X-Ray Timing 2003: RossiXTE and Beyond*, ed. P. Kaaret, F. K. Lamb, & J. H. Swank (Melville, NY: AIP), 245
- Strohmayer, T. E., & Markwardt, C. B. 2002, *ApJ*, **577**, 337
- Strohmayer, T. E., Markwardt, C. B., Swank, J. H., & in't Zand, J. 2003, *ApJL*, **596**, L67
- Strohmayer, T. E., Zhang, W., & Swank, J. H. 1997, *ApJL*, **487**, L77
- Strohmayer, T. E., Zhang, W., Swank, J. H., et al. 1996, *ApJL*, **469**, L9
- Strohmayer, T. E., Zhang, W., Swank, J. H., White, N. E., & Lapidus, I. 1998, *ApJL*, **498**, L135
- Suleimanov, V., Poutanen, J., Revnivtsev, M., & Werner, K. 2011a, *ApJ*, **742**, 122
- Suleimanov, V., Poutanen, J., & Werner, K. 2011b, *A&A*, **527**, A139
- Suleimanov, V., Poutanen, J., & Werner, K. 2012, *A&A*, **545**, 16
- van Paradijs, J. 1979, *ApJ*, **234**, 609
- van Paradijs, J. 1982, *A&A*, **107**, 51
- van Paradijs, J., Dotani, T., Tanaka, Y., & Tsuru, T. 1990, *PASJ*, **42**, 633
- van Straaten, S., van der Klis, M., Kuulkers, E., & Méndez, M. 2001, *ApJ*, **551**, 907
- von Toussaint, U. 2011, *RvMP*, **83**, 943
- Watts, A. L. 2012, *ARA&A*, **50**, 609
- Watts, A. L., Altamirano, D., Linares, M., et al. 2009, *ApJL*, **698**, L174
- Watts, A. L., Strohmayer, T. E., & Markwardt, C. B. 2005, *ApJ*, **634**, 547
- Weinberg, N., Miller, M. C., & Lamb, D. Q. 2001, *ApJ*, **546**, 1098
- Wijnands, R., Strohmayer, T., & Franco, L. M. 2001, *ApJL*, **549**, L71
- Worpel, H., Galloway, D. K., & Price, D. J. 2013, *ApJ*, **772**, 94