DETERMINING NEUTRON STAR PROPERTIES BY FITTING OBLATE-STAR WAVEFORM MODELS TO X-RAY BURST OSCILLATIONS

M. COLEMAN MILLER[1] AND FREDERICK K. LAMB[2,3]

[1] Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421 USA; miller@astro.umd.edu
[2] Center for Theoretical Astrophysics and Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL 61801-3080, USA
[3] Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, IL 61801-3074, USA

ABSTRACT

We describe sophisticated new Bayesian analysis methods that make it possible to estimate quickly the masses and radii of rapidly rotating, oblate neutron stars by fitting oblate-star waveform models to energy-resolved observations of the waveforms of X-ray burst oscillations produced by such stars. We find that a 25% variation of the temperature of the hot spot in the north-south direction does not significantly bias estimates of the mass M and equatorial radius $R_{\rm eq}$ derived by fitting a model that assumes a uniform-temperature spot. Our results show that fits of oblate-star waveform models to waveform data can simultaneously determine M and $R_{\rm eq}$ with 1σ uncertainties $\lesssim 7\%$ if (1) the star's rotation rate is $\gtrsim 600\,{\rm Hz}$; (2) the spot center and observer's sightline are both within $30°$ of the star's rotational equator; and (3) $\sim 10^7$ counts are collected during burst oscillations that have a fractional rms amplitude $\gtrsim 10\%$. A fractional amplitude of $\sim 10\%$ is realistic, and the accepted *NICER* and proposed *LOFT* and *AXTAR* space missions could collect 10^7 counts from a single star by combining data from many X-ray bursts from the star. Uncertainties $\lesssim 7\%$ are small enough to improve substantially our understanding of cold, ultradense matter.

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

1. INTRODUCTION

A currently unresolved fundamental question in physics and astronomy is the nature of cold matter at densities above the saturation density of nuclear matter. Such matter is inaccessible in the laboratory, but is present in large quantities in the interiors of neutron stars. Studies of neutron stars can therefore help determine the properties of such cold, ultradense matter. In particular, precise, simultaneous determinations of the gravitational mass M and equatorial circumferential radius $R_{\rm eq}$ of several neutron stars with different masses could provide tight constraints on the equation of state of this matter (see, e.g., Lattimer 2007; Lattimer & Prakash 2007; Özel & Psaltis 2009; Read et al. 2009; Hebeler et al. 2010).

The recent discovery of two neutron stars with masses of $\approx 2\,M_\odot$ (Demorest et al. 2010; Antoniadis et al. 2013) has placed an important lower bound on the stiffness of cold, ultradense matter, but still allows a wide range of proposed equations of state, because the radii of these two stars are unknown. Several methods have been proposed to estimate M and $R_{\rm eq}$ simultaneously by accurately measuring and interpreting the X-ray spectra of neutron stars. However, the estimates of M and $R_{\rm eq}$ made using these methods are currently dominated by systematic errors (for recent reviews, see Lo et al. 2013; Miller 2013).

An alternative approach is to determine M and $R_{\rm eq}$ by fitting waveform models to the X-ray oscillations of accretion-powered pulsars, the thermal X-ray oscillations produced by some rotation-powered (non-accreting) pulsars, or the X-ray oscillations observed during some thermonuclear X-ray bursts from accreting neutron stars. Estimates of M and $R_{\rm eq}$ made by fitting waveform models to observations of the oscillations produced by accretion-powered X-ray pulsars (see Poutanen & Gierliński 2003; Poutanen & Beloborodov 2006; Leahy et al. 2008, 2011; Morsink & Leahy 2011) may encounter significant systematic errors, because they depend on correctly modeling the complex, time-dependent thermal and nonthermal X-ray spectra and radiation beaming patterns of these pulsars, which are uncertain. For example, significant pulse profile variability has been seen in the accretion-powered millisecond pulsar SAX J1808–3658 (Hartman et al. 2008) and other accretion-powered X-ray pulsars, which could be due to disk-magnetospheric interactions (Kajava et al. 2011). Such interactions may produce complex time-varying emission patterns on the stellar surface (Romanova et al. 2003, 2004).

Estimates of M and $R_{\rm eq}$ can also be made by fitting waveform models to the X-ray oscillations produced by the heated polar caps of rotation-powered pulsars (see Braje et al. 2000; Bogdanov et al. 2007, 2008). Existing models give statistically acceptable fits to the X-ray oscillations of pulsars such as PSR J0437–4715 (Bogdanov 2013), although the constraints on M and $R_{\rm eq}$ obtained using current data are not tight. These estimates are likely to have fewer systematic errors than estimates obtained by fitting the oscillations of accretion-powered X-ray pulsars, but the temperature structure and radiation beaming patterns of the polar caps of these pulsars are uncertain. For example, current treatments assume that the energy that powers the X-ray emission comes from magnetospheric return currents, is deposited deep in the atmosphere, and propagates outward through a nonmagnetic, pure hydrogen atmosphere. However, other light element atmospheric compositions are equally consistent with the current data (see Miller 2013 for a discussion).

One of the most promising current methods for determining M and $R_{\rm eq}$ is to fit energy-dependent waveform models to the X-ray oscillations observed during some thermonuclear X-ray bursts from some bursting neutron stars (see Strohmayer et al. 1997; Miller & Lamb 1998; see also Weinberg et al. 2001). An advantage of this approach is that there is strong theoretical and observational evidence that the radiation from the hot spots that are created by thermonuclear burning is fully thermalized and that the spectra and radiation beaming

patterns from these spots are fairly well understood (Miller et al. 2011, 2013; Suleimanov et al. 2012). Simple hot-spot waveform models provide statistically acceptable descriptions of the observations of burst oscillations from neutron stars such as 4U 1636–536 (Artigue et al. 2013), although the constraints on M and R_{eq} that can be derived using currently available data are not very tight.

A recent study by Lo et al. (2013) analyzed the constraints on M and R_{eq} that could be obtained by fitting model waveforms to observations of X-ray burst oscillations carried out using a next-generation large-area X-ray timing instrument with an effective area $\sim 10\,\mathrm{m}^2$. Lo et al. found that fitting a standard waveform model to the oscillation produced by a hot spot located within $10°$ of the rotational equator can determine both M and R_{eq} with 1σ uncertainties of about 10% if the fractional rms amplitude is $\sim 10\%$ and $\sim 10^7$ counts are collected from the star. This is a realistic amplitude, and this many counts could be obtained from a given star by future space missions, such as the accepted *NICER* mission (Gendreau et al. 2012) and the proposed *LOFT* (Feroci 2012) and *AXTAR* (Ray et al. 2011) missions, by combining data from multiple bursts. If on the other hand the oscillation is produced by a spot that is located within $20°$ of the rotational pole, waveform fitting provides no useful constraints.

Importantly, Lo et al. (2013) demonstrated that fitting rotating hot spot models to energy-dependent X-ray waveforms gives results that are robust against several types of systematic error. In particular, they found that when their standard waveform model was fit to synthetic waveform data generated using spot shapes, energy spectra, or surface beaming patterns different from those assumed in the model, the fits did not simultaneously produce a statistically good fit, apparently tight constraints on the stellar mass and radius, and a significant bias in their inferred values. Thus, waveform models with these systematic deviations do not produce tight but misleading constraints.

In their analysis, Lo et al. (2013) used the Schwarzschild plus Doppler (S+D) approximation (Miller & Lamb 1998; Poutanen & Gierliński 2003; Viironen & Poutanen 2004) to speed the computation of synthetic waveform data and model waveforms. The S+D approximation treats exactly all special relativistic effects (such as relativistic Doppler boosts and aberration) produced by the rotational motion of the emitting gas, 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 therefore does not include the effects of stellar oblateness or frame dragging. Waveforms computed using the S+D approximation are accurate for stars that do not both rotate rapidly and have low compactness (Cadeau et al. 2007; Lo et al. 2013) and are expected to be fairly accurate even for rapidly rotating, oblate stars if the hotter region that produces the oscillation is near the rotational equator and the observer is at a high inclination, the geometry that is required for the waveform to have substantial harmonic structure (Poutanen & Beloborodov 2006), which is necessary to obtain tight constraints on M and R_{eq} (Lo et al. 2013).

Cadeau et al. (2007) studied the accuracies of various approximations for computing the waveform produced by a hot spot on the surface of a rotating neutron star. They did this by first constructing numerical models of rotating neutron stars and their exterior spacetimes using the rotating neutron star code `rns` (Stergioulas 2003) and then utilizing the results to compute the waveform produced by a hot spot on the surface of these model stars. Finally, they compared these accurate numerical waveforms with waveforms computed using various approximations, including the S+D approximation and a new approximation they called the oblate-star Schwarzschild-space-time (OS) approximation. In the OS approximation, the oblate surface of the spinning star is taken into account by embedding a surface with this oblateness in the Schwarzschild spacetime with a mass equal to the gravitational mass of the star. This approximation therefore does not include frame dragging or the effect of the stellar mass quadrupole on the spacetime. To simplify their assessment of the accuracies of waveforms computed using various approximations, they considered only waveforms produced by a hot spot of infinitesimal extent.

Cadeau et al. (2007) found that the most important effects on the waveform caused by rapid rotation are those produced by the oblateness of the star, and that as long as the correct shape of the star is used to formulate the initial conditions for ray-tracing, the waveforms produced by ray-tracing in the Schwarzschild spacetime are a very good approximation to the waveforms produced by ray-tracing using accurate numerical models of the rotating star and its exterior spacetime. Consequently, the OS approximation should be adequate for many purposes.

Cadeau et al. (2007) then carried out a preliminary investigation of the accuracies of M and R_{eq} estimates made using various approximate waveform models, by fitting these models to synthetic waveform data generated using their accurate numerical models. In order to make their fitting procedure tractable, Cadeau et al. made a number of simplifying assumptions. In addition to assuming that the emitting region is infinitesimal in extent, they used only bolometric waveforms and assumed that counts are produced only by photons from the hot spot, i.e., that there are no background counts from other parts of the star, the binary system, other sources in the field of view, or the detector (see Section 2.5.2 for a more detailed discussion of their assumptions and approach).

Based on a preliminary investigation in which they fit their S+D waveform model to synthetic waveform data generated using numerical models of rotating neutron stars and their exterior spacetimes, Cadeau et al. (2007) concluded that using the S+D waveform model can bias estimates of M and R_{eq} if the star has a large radius and is rotating rapidly. In this investigation they did not perform a full statistical analysis, but instead determined for each of their numerical synthetic waveforms the best-fit values of M, the circumferential radius $R(\theta_{spot})$ at the spot colatitude θ_{spot}, and $M/R(\theta_{spot})$ in their S+D waveform model, estimating the uncertainty in $M/R(\theta_{spot})$ using a $\Delta\chi^2$ approach. They found that the best-fit value of $M/R(\theta_{spot})$ was usually close to the true value, with a few exceptions. Despite this, the best-fit values of M and $R(\theta_{spot})$ often differed from their true values, especially if the rotational frequency is $\gtrsim 500\,\mathrm{Hz}$, the hot spot is at a medium to low rotational colatitude, and the observer's inclination to the rotational axis is small. However, for these hot-spot and observer geometries oscillation amplitudes are small and the effects on model waveforms of changes in the model parameters are highly degenerate, making estimates of M, R_{eq}, and M/R_{eq} highly uncertain (see Cadeau et al. 2007, Table

2; Lo et al. 2013). Cadeau et al. did not estimate the uncertainties in M and $R(\theta_{\mathrm{spot}})$ separately, nor did they determine the best-fit values or uncertainties of any of the other parameters in their S+D waveform model. Consequently, they could not determine whether the differences between the estimated and input values of M and $R(\theta_{\mathrm{spot}})$ are statistically significant. It is therefore not clear from their work whether, and if so under what circumstances, fitting an S+D waveform model to actual waveform data can produce fits that are good and constraining but yield parameter estimates that differ significantly from the true values of the parameters.

In this paper, as in Lo et al. (2013), we focus on X-ray burst oscillations, extending the work of Cadeau et al. (2007) and Lo et al. (2013) by performing a full Bayesian analysis of the constraints that can be obtained by fitting S+D and OS waveform models to the energy-resolved waveforms produced by rapidly rotating, oblate neutron stars. As we explain in Section 2.1, when analyzing burst oscillations the angular radius $\Delta\theta_{\mathrm{spot}}$ of the hot spot and a phase-independent background with an arbitrary energy spectrum must be included as part of the fit; to do otherwise is observationally incorrect and leads to misleadingly tight constraints on the mass and radius.

Our first step is to generate synthetic waveform data for stars with a variety of radii and rotational frequencies, using the OS approximation. We have chosen to use the OS approximation as the next step beyond the work of Lo et al. (2013), because the results of Cadeau et al. (2007) and Morsink et al. (2007) show that waveforms computed using the OS approximation are extremely close to those calculated using the much more computationally taxing approach of computing rotating neutron star models and their exterior spacetimes using a numerical code and solving for the required photon ray paths in these numerical spacetimes.

Next, we fit our standard S+D waveform model to the OS synthetic waveform data, to determine whether fitting this model to OS synthetic data produces statistically good fits, and if so, whether the best-fit values of the parameters in the S+D model are close to their true values or are biased by statistically significant amounts. If statistically acceptable but biased estimates are possible, we wish to determine the situations in which this occurs. Our final step is to fit our standard OS waveform model to the OS synthetic waveform data, to determine for the first time the constraints on M and R_{eq} that could be obtained by fitting the OS waveform model to the waveforms produced by rapidly rotating, oblate neutron stars.

The code that we utilize to generate synthetic waveform data using the OS approximation is based on the waveform code validated and used to compute accretion-powered millisecond X-ray pulsar waveforms by Lamb et al. (2009a, 2009b) and to compute X-ray burst oscillation waveforms by Lo et al. (2013, see their Appendix A), modified to implement the OS approximation using the approach developed by Morsink et al. (2007).

We find that if the neutron star has a large radius and is rotating rapidly and the hot spot that produces the oscillation is at a moderate to low rotational colatitude, fitting our standard S+D model to OS synthetic waveform data can produce fits that are statistically good but yield estimates of M and R_{eq} that have significant biases. However, this spot geometry generally does not produce tight constraints on M and R_{eq} (see, e.g., Cadeau et al. 2007, Table 2; Lo et al. 2013, Table 2), because it

produces waveforms in which the oscillation amplitude is low and overtones of the rotational frequency are very weak. If instead the hot spot is at a high rotational colatitude, fitting S+D models to OS waveform data can yield usefully tight constraints on M and R_{eq} with much smaller biases, even for rapidly rotating, oblate stars.

We have developed a sophisticated new Bayesian analysis procedure that makes it possible to fit OS waveform models to waveform data almost as quickly as S+D waveform models. Given the speed of our new procedure for fitting OS waveform models to waveform data and the risk that results obtained using the S+D approximation may be biased, OS waveform models should be used in preference to S+D waveform models in all future waveform analyses.

Using our new analysis procedure, we find that fitting our standard OS waveform model to OS synthetic waveform data produces tight constraints on M and R_{eq} if the star has a moderate to high rotation rate, the hot spot is at a moderate to high rotational colatitude, and the observer is at a moderate to high inclination to the rotational axis. As an example, our results show that if the star's rotation rate is $\gtrsim 600\,\mathrm{Hz}$, the spot center and the observer's sightline are both within $30°$ of the equatorial plane, the fractional rms amplitude of the oscillation is $\gtrsim 10\%$, and $\gtrsim 10^7$ total counts are collected from the star, then M and R_{eq} can both be determined with 1σ uncertainties $\lesssim 7\%$, comparable to the uncertainties we obtained when fitting S+D waveform models to S+D synthetic waveform data generated for these same situations (Lo et al. 2013). As noted earlier, this is a realistic fractional amplitude, and this many counts could be obtained from a single star by the accepted *NICER* and proposed *LOFT* and *AXTAR* space missions by combining data from many X-ray bursts. Simultaneous measurements of M and R_{eq} with these precisions would improve substantially our understanding of cold, ultradense matter.

We also find that no significant biases are introduced in the estimates of M and R_{eq} when we fit our standard OS model waveform, which assumes a uniform-temperature hot spot, to OS synthetic waveform data generated assuming a 25% north–south (latitudinal) variation in the surface temperature of the hot spot. This extends the important findings of Lo et al. (2013) that the waveform-fitting method provides unbiased estimates of M and R_{eq} for a broad range of systematic deviations of the actual properties of the hot spot from the properties assumed in computing the model waveform.

The remainder of this paper is organized as follows. In Section 2, we describe in more detail our assumptions and the ray-tracing and statistical methods we use and compare our approach to the approaches used by Lo et al. (2013) and Cadeau et al. (2007). In Section 3, we describe the synthetic waveforms we analyze, our standard S+D and OS waveform models, the fitting procedure we use, and the results obtained by fitting our standard S+D and OS waveform models to the OS synthetic waveform data. We also discuss the precisions of M and R_{eq} estimates obtained by fitting waveform models to waveform data. We summarize our conclusions in Section 4. Although we focus here on analyzing X-ray burst waveforms, our method can, with small changes, be used to analyze the X-ray waveforms produced by thermal emission from the polar caps of rotation-powered pulsars, a goal of the *NICER* mission.

2. METHODS

In this section, we first discuss burst oscillation phenomenology and the assumptions we make when generating synthetic waveform data and model waveforms. We then describe the analytical implementation of the OS approximation that we use, which was developed by Morsink et al. (2007), and explain the table-lookup method for light rays that we have introduced, which speeds up the waveform-fitting process by a factor of hundreds. Next, we discuss some of the fundamentals of Bayesian inference, which underlies our approach to estimating M and R_{eq}. We then describe the waveform data processing procedure we use to obtain the results we present in Section 3. We have found this procedure to be a reliable method for exploring the space of waveform model parameters and marginalizing the parameters that are not of interest to us here. We conclude this section by comparing our approach in this work to the approaches used by Lo et al. (2013) and Cadeau et al. (2007). The new, more sophisticated, and much faster Bayesian analysis procedures we introduce here allow us to determine the best-fit parameters of waveform models by a blind search of M–R_{eq} space, i.e., by sampling these parameters without using any knowledge of the values of M and R_{eq} that were used to generate the synthetic observed waveform data (compare Lo et al. 2013), despite the additional complexity of the OS approximation.

2.1. Burst Oscillation Phenomenology and Modeling

2.1.1. Hot Spot

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 stellar surface and is rotating at or near the rotation frequency of the star. Such a hotter region could be produced either by heating of part of the stellar surface by thermonuclear burning or because a disturbance in the outer layers of the star, such as a global surface mode, has made a localized region hotter (see Watts 2012).

Oscillations with the relatively high amplitudes ($\gtrsim 5\%$–10%) required to derive significant constraints on M and R_{eq} are probably produced predominantly by a single hotter region (see Lamb et al. 2009a). The dominant role of a single, localized hotter region is also indicated by the absence in the waveform of substantial second or higher harmonics of the fundamental oscillation frequency. Hence, regardless of whether the localized hotter region is produced by heating of the stellar surface by nuclear burning or by a global surface mode, the waveform it produces can be modeled using a single hot spot. Lo et al. (2013) have shown that the waveform produced by a hot spot is relatively insensitive to elongation of the spot in the east–west or north–south directions, so modeling the spot as a circular area is expected to be adequate.

An important question is whether oscillations observed during the rise of bursts or during burst peaks and tails are likely to be more useful in deriving constraints on M and R_{eq}. The nuclear-powered emission from the surface of the star is expected to be highly localized during the first fraction of a second of an X-ray burst. If adequate constraints could be obtained using data collected only when the emission comes from a hot spot with an angular radius $\Delta\theta_{spot} \lesssim 10°$, waveforms could be modeled assuming the hot spot is a point, and it would be unnecessary to include $\Delta\theta_{spot}$ as a parameter in fitting waveform models to waveform data. If, on the other

hand, data collected when the emission is larger must be used, $\Delta\theta_{spot}$ must be included as a parameter.

Lo et al. (2013) carried out a comprehensive analysis of the likely usefulness of data taken during burst rises versus burst peaks and tails (see their Section 2.2). They found that analyses of oscillations observed during burst tails may provide estimates of M and R_{eq} with uncertainties comparable to or possibly even smaller than the uncertainties provided by analyses of oscillations observed during burst rises. Our approach in this paper is based on their findings, which we therefore summarize here.

The precisions of M and R_{eq} estimates are most sensitive to the star's rotation rate and the inclinations of the hot spot and the observer, because these quantities strongly affect the special relativistic Doppler boost and aberration caused by the line-of-sight component of the surface rotational velocity, and therefore affect the harmonic content of the observed waveform. The uncertainties of M and R_{eq} estimates also depend on the amplitude of the oscillation and the total number of counts, because these determine the size of the statistical fluctuations in measurements of the waveform variation.

Using the results of their extensive parameter estimation study, Lo et al. (2013) showed that, other things being equal, the uncertainties in M and R_{eq} estimates obtained by fitting waveform models to waveform data scale approximately as $\mathcal{R}^{-1}$ (see their Equation (1), Section 4.2.1, and Table 3), where

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

Here $N_{osc} = 1.4\, f_{rms} N_{tot}$ is the number of counts collected from the oscillating component of the waveform during the observation, N_{tot} is the total number of counts collected, and f_{rms} is the average fractional rms amplitude of the oscillation during the observation. N_{osc}, N_{tot}, and f_{rms} are all directly observable. $\mathcal{R}^{-1}$ is the fractional variation in the shape of the waveform produced by the counting noise and is therefore a useful figure of merit for evaluating and comparing data sets.

N_{osc} is approximately equal to the integral of the semi-amplitude of the first harmonic (fundamental) component of the burst oscillation waveform over the duration of the data segment and f_{rms} is approximately equal to the rms amplitude f_{rms1} of this harmonic because the amplitudes of the higher harmonics in burst oscillation waveforms are much smaller than the amplitude of the first harmonic (see Watts 2012; Lo et al. 2013; and Table 3 below). $N_{tot} = N_{spot} + N_{back}$, where N_{spot} is the total number of counts detected from the hot spot and N_{back} is the total number of background counts, which we define as all counts not produced by photons from the hot spot. N_{spot} is not usually equal to N_{osc}, because for many geometries the hot spot contributes an unmodulated component to the waveform, as well as an oscillating component, i.e., $N_{spot} = N_{const} + N_{osc}$. The value of N_{back} during bursts cannot currently be measured or predicted (see Section 2.1.2). Consequently, although N_{tot} is directly observable, its components N_{spot} and N_{back} can only be determined by fitting models to the observed waveform.

Lo et al. (2013) found that for systems with properties that allow interesting constraints to be derived on M and R_{eq}, i.e., systems in which the stellar rotation rate is $\gtrsim 300\,\mathrm{Hz}$ and the spot center and the observer's sightline are both within $30°$ of the star's rotational equator, waveform data with $\mathcal{R} \gtrsim 400$ can provide estimates of M and R_{eq} with uncertainties $\lesssim 10\%$.

A hot spot small enough to be treated as a point source (i.e., with $\Delta\theta_{spot} \lesssim 10°$) spans $\lesssim 1\%$ of the stellar surface and therefore cannot produce enough counts to obtain interestingly tight constraints on M and R_{eq}, even if data are collected from many tens of bursts from a single star using a detector with a collecting area $\sim 10\,m^2$ and then combined. This is demonstrated by the existing observations of burst oscillations, which show that although the oscillations observed during the first fraction of a second of an X-ray burst typically have high fractional amplitudes, they are not bright enough and do not last long enough to provide the number of counts needed to derive tight constraints on M and R_{eq}, for a practical number of burst observations from a single star.

For example, Lo et al. (2013) showed that observation of a bright burst from the bright X-ray burst source 4U 1636–536 using a detector with an effective area $\sim 10\,m^2$ could achieve an $\mathcal{R}$-value ~ 33 during the first $1/16\,s$ of the burst, when the emitting area is small. $\mathcal{R}$ scales as the square root of the number of counts, so achieving an $\mathcal{R}$-value ~ 400, sufficient to provide estimates of M and R_{eq} with uncertainties $\lesssim 10\%$, would require combining data from the early rises of ~ 150 such bursts, which is impractical. Although the oscillation amplitude diminishes as the burst rises, using data from the entire $1/4\,s$ burst rise would yield a slightly larger $\mathcal{R}$-value ~ 50, because the longer duration of the observation yields a much larger number of counts. Achieving an $\mathcal{R}$-value ~ 400 would therefore require combining data from the entire rises of ~ 65 such bursts.

Although the oscillations observed during burst tails usually have smaller fractional amplitudes, they last much longer than the first fraction of a second of burst rises. They also have larger emitting areas. Using data collected during a $2\,s$ oscillation train observed in the tail of a burst from any one of the eight neutron stars that produce such oscillation trains would yield an $\mathcal{R}$-value ~ 80 (see Lo et al. 2013, Section 2.2.1). Hence, combining data from ~ 25 observations of such oscillation trains from a single star would yield $\mathcal{R} \sim 400$, sufficient to obtain $\sim 10\%$ constraints on M and R_{eq} for systems that have favorable properties (again see Lo et al. 2013).

These results show that obtaining constraints on M and R_{eq} with uncertainties $\lesssim 10\%$ will probably require using data from the peaks and/or tails of bursts, when the hot spot is not infinitesimal in extent.

Given that it will probably be necessary to include data taken when the hot spot is not infinitesimal, it is important to include the angular radius $\Delta\theta_{spot}$ of the spot as a parameter in waveform fits, because for a given observed waveform, the most probable value of $\Delta\theta_{spot}$ and the distance d to the star are related. Assuming point emission during the full rise, the peak of the burst or its tail would improperly remove this degeneracy, artificially reducing the uncertainties in M and R_{eq} estimates and possibly biasing them. As Lo et al. (2013) showed using a full Bayesian analysis, knowledge of the distance to the star can improve somewhat the precision of M and R_{eq} estimates by removing this degeneracy.

2.1.2. Backgrounds

An important factor that affects M and R_{eq} estimates made using burst oscillation waveforms is the difficulty of determining the background independently of the waveform-fitting process. Background counts could come from the non-spot portion of the star, the accretion disk, unassociated sources in the field, instrumental backgrounds, or any combination of these.

Independent knowledge of the background would improve the constraints on M and R_{eq} derived by waveform-fitting (see Lo et al. 2013), but both observational evidence (Yu et al. 1999; Kuulkers et al. 2003; Chen et al. 2011; in't Zand et al. 2011; Serino et al. 2012; Degenaar et al. 2013; Worpel et al. 2013; Peille et al. 2014) and theoretical arguments (Walker 1992; Miller & Lamb 1996; Ballantyne & Strohmayer 2004; Ballantyne & Everett 2005) indicate that the accretion-powered emission from neutron star systems is substantially different during a burst than before or after.

In principle, the accretion-powered emission could, at some times during a burst, be more luminous than before or after the burst, if the radiation from the burst significantly increases the radiation drag on the gas orbiting in the disk, causing the accretion rate to the stellar surface to increase, or it could be less luminous, when the increased radiation drag has depleted the inner disk. Unfortunately, the observed variations in the background are not understood theoretically and do not appear to be correlated with other properties of the bursts in any obvious way (see Peille et al. 2014 and references therein). Consequently, whether the background varies during a particular burst, and if so, by how much and in which direction, cannot currently be predicted. Hence, in order to obtain reliable estimates of M and R_{eq}, one must include in the fitted waveform model an oscillation-phase-independent background component with an arbitrary magnitude and energy spectrum.[4]

2.1.3. Other Aspects of the Waveform Models

In constructing synthetic waveform data and model waveforms, we assume that the hot spot has a constant size and shape, is located at a fixed stellar rotational latitude, and rotates at a constant frequency that is known a priori.

Our S+D and OS waveforms have seven primary parameters: the star's gravitational mass M; its equatorial circumferential radius R_{eq}; the colatitude θ_{spot} of the hot spot center; the angular radius $\Delta\theta_{spot}$ of the hot spot, which is assumed to be circular and uniform; the surface comoving temperature T_{spot} of the emission from the hot spot, which is assumed to have a blackbody spectrum and normalization; the inclination (colatitude) θ_{obs} of the observer relative to the hot spot rotational axis, which in this work we assume is also the stellar rotational axis; and the distance d to the star. In computing the shape of the star for our OS waveforms, we assume that the rotational frequency of the star is the same as the rotational frequency ν_{rot} of the hot spot.

In generating synthetic waveform data and computing model waveforms, we use the beaming function that describes radiation emerging from the surface of an electron-scattering atmosphere (see Equation (9)). As Suleimanov et al. (2012) have shown, for the 1–30 keV energy range and high surface fluxes, the beaming function for an electron-scattering atmosphere accurately describes not only the beaming pattern of radiation from such an atmosphere but also that from a pure hydrogen atmosphere, and deviates by at most 6% from the

---

[4] Even if the pre-burst background persisted 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 uncertainties in the model parameters these fluctuations induce.

beaming function of radiation from an atmosphere with solar composition. This beaming function therefore provides an excellent description of the beaming of radiation from burst atmospheres. When analyzing the X-ray waveforms produced by much cooler hot spots (such as the polar caps of rotation-powered pulsars), it will be necessary to use different beaming functions.

When generating synthetic waveform data, we include a constant background component. This component is a catch-all for all the counts not produced by the emission from the hot spot. As noted above, these counts could come from the non-spot portion of the star, the accretion disk, unassociated sources in the field, instrumental backgrounds, or any combination of these. For simplicity, we follow Lo et al. (2013), modeling 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 that is usually lower than the temperature of the hot spot. This is a reasonable approach, because the number of counts contributed by this emission is important, but not its detailed properties. We normalize the background spectrum to achieve the desired expected number of background counts.

In generating the synthetic waveform data we analyze here, we assume that $\sim 10^6$ counts have been collected from the hot spot and that the background is $\sim 9 \times 10^6$ counts. This corresponds to a realistic modulation fraction and is a sufficient number of counts to determine M and $R_{\rm eq}$ with uncertainties $\lesssim 7\%$ if the star's rotation rate is $\gtrsim 600$ Hz and the spot center and observer's sightline are both within $30°$ of the star's rotational equator. As noted earlier, a future space mission, such as the accepted *NICER* mission and the proposed *LOFT* and *AXTAR* missions, could obtain this number of counts by combining data from many bursts from a given star (see Lo et al. 2013).

2.2. The Oblate Schwarzschild Approximation and Ray-tracing

In the OS approximation (Cadeau et al. 2007; Morsink et al. 2007), the spacetime exterior to the star is Schwarzschild but the stellar surface is oblate. Thus, in effect, the star is treated as an oblate shell of infinitesimal mass surrounding a concentrated sphere of gravitational mass M. In Schwarzschild coordinates, the line element anywhere outside the star is therefore

$$ds^2 = -(1 - 2M/r)dt^2 + dr^2/(1 - 2M/r) + r^2\left(d\theta^2 + \sin^2\theta d\phi^2\right),$$ (2)

where t is the time as measured at infinity, r is the circumferential radius, and θ and ϕ are the standard polar and azimuthal angles. Here and henceforth we usually use geometrized units in which $G = c \equiv 1$. As in the S+D approximation discussed in Section 1, all special relativistic effects are treated exactly.

We assume that the rotating neutron star of interest is symmetric around its rotational axis. Morsink et al. (2007) found that for the families of stars they considered, if a star has an equatorial radius $R_{\rm eq}$ and an angular frequency Ω as seen at infinity, its radius as a function of colatitude θ is well described

by (see their Equations (8)–(10) and their Table 1)

$$\frac{R(\theta)}{R_{\rm eq}} \approx 1 + \sum_{n=0}^{2} a_{2n}(\zeta, \epsilon) P_{2n}(\cos\theta),$$ (3)

where $P_{2n}(\cos\theta)$ is the Legendre polynomial of order $2n$,

$$\zeta \equiv \frac{M}{R_{\rm eq}}, \qquad \text{and} \qquad \epsilon \equiv \frac{\Omega^2 R_{\rm eq}^3}{M}.$$ (4)

For neutron stars and hybrid quark stars,

$$a_0 = -0.18\epsilon + 0.23\zeta\epsilon - 0.05\epsilon^2,$$
$$a_2 = -0.39\epsilon + 0.29\zeta\epsilon + 0.13\epsilon^2, \text{ and}$$
$$a_4 = 0.04\epsilon - 0.15\zeta\epsilon + 0.07\epsilon^2.$$ (5)

The area of an infinitesimal surface element at colatitude θ on the surface of an oblate star is (see Morsink et al. 2007, Equations (2) and (3))

$$dS(\theta) = R^2(\theta)\sin\theta\left[1 + f^2(\theta)\right]^{1/2} d\theta d\phi,$$ (6)

where

$$f(\theta) \equiv \frac{(1 - 2M/R)^{-1/2}}{R}\frac{dR}{d\theta}.$$ (7)

The advantage of using the Schwarzschild spacetime rather than spacetimes that include the effect of the mass quadrupole of the star and frame dragging is that the spherical symmetry of the Schwarzschild spacetime guarantees that the path of any light ray in vacuum will lie in a plane. Thus, ray paths can be pre-computed and used in a lookup table, speeding up the computations enormously. We describe our table-lookup procedure later in this section.

When constructing OS synthetic waveforms and waveform models, we use the same ray-tracing codes described in Lo et al. (2013). Many of the equations and algorithms used in these codes were originally derived by Poutanen & Gierliński (2003).

We assume that the emission from the hot spot is azimuthally symmetric around the local surface normal, as seen in the surface comoving frame. The variation of the star's circumferential radius with colatitude creates an angle γ between this normal and the local outward radial vector given by (see Section 2.2 of Morsink et al. 2007)

$$\cos\gamma = \left[1 + f^2(\theta)\right]^{-1/2}.$$ (8)

We assume further that the emission as seen in the surface comoving frame extends from the surface normal to the tangent to the surface, with a beaming function that is usually given by that expected for radiation from a uniform, semi-infinite, Thomson scattering atmosphere, which we approximate by (Lo et al. 2013)

$$g'(\alpha') = 0.42822 + 0.92236\cos\alpha' - 0.085751\cos^2\alpha',$$ (9)

where α' is the angle between the emitted ray and the surface normal, as seen in the surface comoving frame.

In addition to the minor modification to the element of surface area given by Equation (6), there are two more

important differences between the OS approximation and the S+D approximation. The first is that the stellar radius varies with colatitude. As a result, if the equatorial circumferential radius is fixed, the maximum angular deflection of a photon leaving the star from a surface element not on the rotational equator is greater in the OS approximation than in the S+D approximation. The second difference is that in the S+D approximation, the maximum angle between the direction of an emitted photon and the local outward radial vector is $\pi/2$, assuming that the star is not more compact than the photon orbit $R = 3M$ (for more compact stars, the maximum angle for an escaping photon is less than $\pi/2$; however, we do not consider such compact stars). In contrast, for an oblate star the tilt of the normal vector from the radial vector means that in the direction toward the closer rotational pole, the maximum angle is greater than $\pi/2$, whereas in the direction away from that pole, the maximum angle is less than $\pi/2$. More generally, if ψ is an azimuthal angle in the stellar surface that is 0 for the direction toward the closer pole and π for the direction away from it, then the angle χ between the local outward radial vector and the photon ray in direction ψ in the plane tangent to the stellar surface is given by $\cos\chi = -\cos\psi \sin\gamma$.

The two effects just described both tend to increase the minimum flux in the OS waveform relative to the S+D waveform for the same spot location and size, equatorial radius, gravitational mass, and observer colatitude, for observers in the same rotational hemisphere as the hot spot. To see why, consider a point on the surface somewhere between the pole and equator. The minimum in the waveform at a given energy will occur when the observer is on the opposite side of the star (modulo some effects related to relativistic aberration). In the OS approximation, the surface emission comes from a smaller radius and a ray can start in a direction with a smaller impact parameter to the center of the star than in the S+D approximation. Consequently, the gravitational light deflection is greater and causes the minimum flux to be greater than in the S+D approximation (see, e.g., Figures 3 and 4 of Cadeau et al. 2007 and Figures 3 and 4 of Morsink et al. 2007). In contrast, for observers in the opposite rotational hemisphere from the hot spot these two effects tend to *decrease* the minimum flux in the OS waveform relative to the S+D waveform (see, e.g., Figures 6 and 8 of Cadeau et al. 2007; but see their Figure 5 for a counterexample). The OS and S+D waveforms are very nearly the same if the hot spot is in the rotational equator (see, e.g., Figure 7 of Cadeau et al. 2007), because $dR/d\theta$ vanishes there.

Our algorithm computes the required ray paths in advance and stores them in a table that is then read by our code prior to calculating the required waveforms. In the Schwarzschild spacetime, the angular deflection of a ray to infinity depends on R/M, and not on R and M independently. We find that a table of ray paths for 440 evenly spaced values of R/M, from 3.6 to 8.0, and 1000 evenly spaced values of the angle between the ray and the outward radial direction, from 0 radians to 1.82 radians (recall that for oblate stars, the maximum angle from the outward radial direction can exceed $\pi/2 \approx 1.57$) provides adequate accuracy: with this gridding, waveforms constructed using the ray table give a flux at any energy or phase that differs by no more than a few parts in 10^5 from the corresponding directly traced waveform.

For a given initial radius and angle from the outward radial vector, we compute (1) the deflection angle to infinity (see
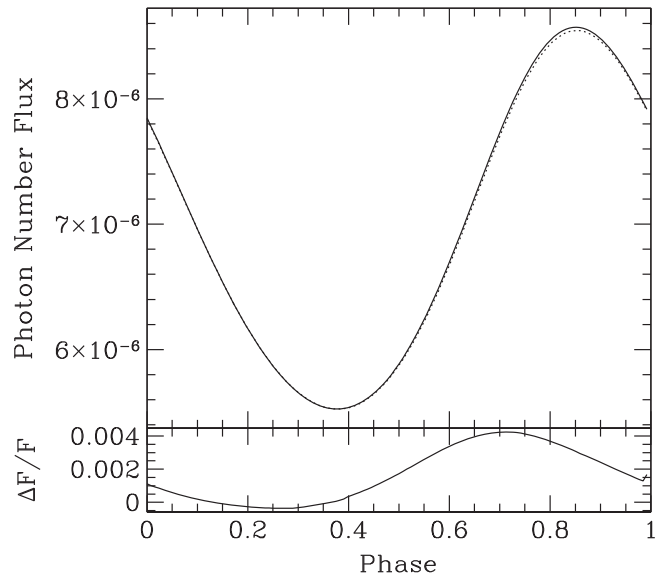
Figure 1. Absolute comparison of waveforms computed using two versions of the OS approximation, for a star with a gravitational mass of 1.4 $M_\odot$ and an equatorial circumferential radius of 16.4 km, a rotational frequency of 600 Hz as seen at infinity, and a hot spot with an angular radius of $0°.1$ centered at a colatitude of $41°$ and seen by an observer at an inclination of $20°$. Top: energy-integrated photon number flux in ph cm^{-2} s^{-1} as a function of rotational phase kindly computed by S. Morsink (2013, private communication; solid curve), using the stellar shape given by the rns code, and computed using our OS waveform code (dotted curve), which uses the analytical method for determining the stellar shape introduced by Morsink et al. (2007). In this figure only, the beaming pattern of the radiation from the hot spot was assumed to be isotropic in the surface comoving frame. Bottom: fractional difference between the solid and dotted curves in the top panel. The agreement between the two waveforms is excellent. Comparison of the two codes revealed that the difference between the two waveforms is due to the slight difference between the stellar shapes computed using the rns code and using the analytical method introduced by Morsink et al. (2007).

Poutanen & Gierliński 2003 and Section A.1.1 of Lo et al. 2013), (2) the time delay of the ray relative to a radial ray (see Braje et al. 2000; Viironen & Poutanen 2004; and Section A.1.2 of Lo et al. 2013), and (3) the gravitational lensing factor for a cluster of rays near the fiducial ray.

Figure 1 compares an energy-integrated photon number flux waveform computed using our OS waveform code with the energy-integrated waveform for the same model parameters kindly computed by S. Morsink (2013, private communication), using the stellar shape given by the Stergioulas rotating neutron star code rns (Stergioulas 2003). In this figure only, the beaming pattern of the radiation from the hot spot (i.e., the angular distribution of the intensity from any point on the hot spot) was assumed to be isotropic in the surface comoving frame. The agreement between the two waveforms is excellent. Comparison of the two codes revealed that the difference between the two waveforms is due to the slight difference between the stellar shape computed using the rns code and the shape computed using the analytical method introduced by Morsink et al. (2007).

2.3. Bayesian Inference and Marginalization

Our statistical approach to estimating M and $R_{\rm eq}$ is based on Bayesian inference techniques, and follows closely the

approach used in Lo et al. (2013). If we have a model with parameters y, and if the prior probability distribution over those parameters is $q(y|I)$ (where I represents prior information), then Bayes' theorem states that the posterior probability distribution after analyzing data D is

$$P(y|D, I) \propto p(D|y, I)q(y|I). \tag{10}$$

If Poisson noise is the only source of fluctuations in the data, the likelihood of the "observed" data, given a particular set y of values for the model parameters, is

$$\mathcal{L} \equiv p(D|y, I) = \prod_i \frac{m_i(y)^{d_i}}{d_i!} e^{-m_i(y)}, \tag{11}$$

where the product is over all the oscillation phase and energy bins, d_i is the measured number of counts in the ith bin, and $m_i(y)$ is the number of counts in the ith bin predicted by the model for the trial set y of parameter values. If the number of counts in a given bin exceeds a few tens, then the Poisson likelihood may be replaced by a Gaussian.

Unlike frequentist statistics such as χ^2, the value of $\mathcal{L}$ itself has no implication for whether the fit is good in an absolute sense. Instead, comparisons are made between different sets of parameter values and depend only on the ratio between different posterior probabilities, which means that they depend on the product of the ratio of the prior probabilities and the ratio of the likelihoods. The common factor $\prod_i (1/d_i!)$ therefore cancels out. In our analysis we adopt flat priors for all of our main parameters, within the range of values that we search for each parameter. In real situations, it is possible to have additional information about some of the parameters, and if so, that information should be included in the analysis via the prior. For example, Lo et al. (2013) found that knowledge of M/R_{eq} via an identified atomic line in the spectrum from the stellar surface, or of the observer's inclination angle, can tighten the constraints on M and R_{eq} considerably, whereas knowledge of the distance improves the constraints only modestly.

We adopt the common procedure of working with the log likelihood, which after removal of the $\sum \log(1/d_i!)$ term is

$$\log \mathcal{L} = \sum_i d_i \log m_i(y) - \sum_i m_i(y). \tag{12}$$

It is the ratio of the likelihoods and thus the difference between the log likelihoods that matters. When we quote uncertainties at a certain level of confidence, we use Wilks' theorem (Wilks 1938), which states that $\Delta \chi^2 \approx -2\Delta \log \mathcal{L}$ for a reasonably large total number of counts, and χ^2 tables.

In our problem, we are primarily interested in M and R_{eq}. If we designate the other parameters (called nuisance parameters in this context) by z, then the correct way to find the final posterior probability distribution for only M and R_{eq} is to marginalize over the other parameters, i.e.,

$$P(M, R_{eq}) \propto \int dz\, p(D|M, R_{eq}, z, I)q(M, R_{eq}, z|I). \tag{13}$$

However, our waveform model has a large number of nuisance parameters. As explained in Section 2.1.1 and Lo et al. (2013), one must include in the waveform model an oscillation-phase-independent background component with an arbitrary magnitude and energy spectrum, which is specified by the number of phase-independent counts in each energy channel. The number

of parameters in the background model is therefore equal to the number of energy channels. In addition, for a given set of candidate parameters one must consider an arbitrary shift in the phase of the model waveform relative to the waveform data, adding another parameter. It is not practical to marginalize fully over so many parameters. Consequently, we instead maximize the likelihood over the parameters in the background model and the start time of the model waveform, using a bisection method.

We have tested whether maximizing the likelihood of these nuisance parameters gives results comparable to marginalizing them, by fitting a Gaussian to the likelihood distributions found during the bisection procedure, analytically integrating over the Gaussian, and comparing the results with those obtained by simply using the maximum likelihood values of these parameters. We compared the log likelihood differences between these methods for combinations of $(M, R_{eq}, \theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$ ranging from the combination that maximized log $\mathcal{L}$ to combinations that gave log $\mathcal{L}$ values 20 less than the maximum. For five parameters, Wilks' theorem (Wilks 1938) suggests that $\Delta \log \mathcal{L} = 2.94$ is approximately equivalent to 1σ. We found that for the background model, the standard deviation in the difference of log $\mathcal{L}$ between the maximization and marginalization procedures was only 0.007. That is, even though the value of log $\mathcal{L}$ for a given parameter combination in the marginalization procedure has an offset from the value of log $\mathcal{L}$ for the same parameter combination in the maximization procedure, the offset is almost exactly constant from one parameter combination to the next. Thus, the differences between log likelihoods are preserved and hence maximization of the likelihood over the background parameters is functionally equivalent to marginalization. When we performed a similar comparison of maximization to marginalization for the shift in the start time of the waveform, we found that the standard deviation of the difference in log $\mathcal{L}$ between the two methods is 0.3. This is therefore a 0.1σ shift, which is too small to affect any of our results.

2.4. Data Processing Pipeline

A challenge faced in Lo et al. (2013) was that if the waveform data are informative, i.e., if they tightly constrain M and R_{eq}, a grid search over the values of the parameters in the waveform model requires a grid so fine that a truly blind search would require a prohibitive number of waveform computations, even if large computational resources are available.

In this section, we describe a new waveform data processing procedure that is based on the approach used by Lo et al. (2013) but performs blind searches far more efficiently. In this new procedure we first determine the volume of the waveform parameter space of interest by performing a Markov chain Monte Carlo (MCMC) search using ray-tracing. We then construct and interpolate in a table of waveforms to further localize the volume of interest and compute bounding ellipsoids that encompass the points in this reduced volume that have interestingly high likelihoods, finally sampling the volumes within these ellipsoids by performing a Monte Carlo integration using direct ray-tracing. The details are as follows:

1. We explore $(M, R_{eq}, \theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$ space using an MCMC search with ray-tracing (see Section 3.2 of Lo et al. 2013 for our MCMC approach and, e.g., von Toussaint 2011 for a discussion of the Metropolis

algorithm for MCMC sampling). In general, one would also need to explore a range of values of the surface comoving temperature T and the distance d to the star, but in the approach we use here we either assume that the redshifted temperature $T(1 - 2M/R_{eq})^{1/2}$ is known and maximize the likelihood over d, or assume that d is known and maximize the likelihood over T (see Section 3.3.3 of Lo et al. 2013 for an explanation and justification of this approach).

2. We determine the maximum log likelihood ($\log \mathcal{L}_{max}$) of the data given the waveform model, in the volume of the waveform parameter space explored in the previous step. We then select the $(M, R_{eq}, \theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$ combinations in this volume that have log likelihoods within 30 of the maximum and determine the range of each of these five parameters that spans the selected combinations. Finally, we generate a table of 10^5 template waveforms, one for each point on the five-dimensional grid of 10 evenly spaced values that span the previously determined range of each parameter.

3. We perform a second MCMC calculation using waveforms computed by interpolating in the table of template waveforms generated in the previous step. In this calculation, we pick an (M, R_{eq}) pair and marginalize by performing a Monte Carlo integration with 1000 points spanning $(\theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$. To do this, we first compute for the selected (M, R_{eq}) the log likelihood on an evenly spaced $10 \times 10 \times 10$ grid in the three angular variables, over the full range of each variable determined in step 1. Using the points on this grid with log likelihoods within 40 of the maximum found in the previous step, we determine, for each value of θ_{obs} on our grid, the minimum ellipse that contains all $(\theta_{spot}, \Delta\theta_{spot})$ pairs with $\log \mathcal{L} > \log \mathcal{L}_{max} - 40$ (see the Appendix for a description of the algorithm we used to find a minimum bounding ellipse). If, for our chosen (M, R_{eq}) pair, any of these ellipses are nonvanishing, we perform a 1000-point Monte Carlo integration over the $(\theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$ volume identified by the nonvanishing bounding ellipses. The use of bounding ellipses typically reduces the volume over which the integral must be performed by a factor ~ 30, which allows us to sample the relevant volume much more densely, reducing the fractional error in the integral by a factor $\sim 30^{1/2} \sim 5$ for a given number of evaluations.

4. We take all the (M, R_{eq}) pairs from the previous step that yielded nonvanishing marginalized posterior probabilities and use ray-tracing rather than waveform interpolation to compute the log likelihood in a full, uniformly spaced $10 \times 10 \times 10$ grid over $(\theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$. Then, as in the previous step, we construct bounding ellipses in $(\theta_{spot}, \Delta\theta_{spot})$ at each grid value of θ_{obs} for each (M, R_{eq}) pair and perform a Monte Carlo integration with 1000 points per (M, R_{eq}) combination, using direct ray-tracing rather than waveform interpolation. This yields the marginalized posterior probability at points in the (R_{eq}, M) plane that, because of the MCMC sampling procedure, are concentrated around the maximum posterior probability.

5. We normalize the marginalized posterior probability to its maximum and then determine and output the 1σ, 2σ, and 3σ contours in the (M, R_{eq}) plane.

2.5. Comparison with Previous Work

2.5.1. Lo et al. (2013)

In addition to using waveforms computed using the OS approximation, in this work we use a data processing procedure that is much more efficient than the one used in Lo et al. (2013). In that work, our marginalization over θ_{obs}, θ_{spot}, and $\Delta\theta_{spot}$ for each (M, R_{eq}) pair used Monte Carlo integration over a volume in the angular variables that was determined by MCMC sampling to contain all points in the angular space with log likelihoods within 20 of the maximum log likelihood, for the given values of M and R_{eq}. This volume was rectangular, which meant that many $(\theta_{obs}, \theta_{spot}, \Delta\theta_{spot})$ triplets within the volume gave very poor fits to the data. As a consequence, Lo et al. had to take 10^4 samples for a given (M, R_{eq}) pair in order to obtain a sufficiently precise result from the Monte Carlo integration. In contrast, the procedure we use here—finding minimum bounding ellipses in $(\theta_{spot}, \Delta\theta_{spot})$ for each $(M, R_{eq}, \theta_{obs})$ combination—reduces the integration volume by a factor ~ 30, which means that we are able to obtain better precision using 10^3 samples than we were previously able to obtain using 10^4 samples.

Lo et al. (2013) used a uniform grid of points in M and R_{eq}. The informative synthetic data sets that they considered produce small high-probability regions in the M–R_{eq} plane. In order to sample these regions adequately, the number of points that would have been required in this grid was so great that full, blind searches in the M–R_{eq} plane were computationally infeasible. They therefore created fine grids around the "true" values of M and R_{eq} (i.e., the values that were used to generate the synthetic waveform data), and argued that when real data become available from future larger-area X-ray detectors, the available computational power will have increased enough to permit blind searches. In contrast to the procedure used in Lo et al., the MCMC exploration of the M–R_{eq} plane used here automatically finds the regions of highest posterior probability and concentrates the sampling there. Thus, even for data that are highly informative, our blind search does a good job of exploring the high-probability regions.

The net result of these new procedures is that whereas the analysis procedure used in Lo et al. (2013) took 50–100 clock hours on a 150 core cluster and did not actually search the entire (M, R_{eq}) domain that we wished to consider, our current analysis procedure takes 50–100 clock hours on a 5 core desktop to do a blind search of the entire region of interest. This huge gain in efficiency allows us to produce significantly more precise and reliable uncertainty estimates than was possible previously.

2.5.2. Cadeau et al. (2007)

As we discussed in Section 1, Cadeau et al. (2007) performed a pioneering preliminary exploration of whether fitting S+D waveform models to the waveforms generated by a hot spot on a rotating neutron star produces estimates of M, R_{eq}, and GM/R_{eq} that are biased. They found that substantial differences between the actual and best-fit values of M and $R(\theta_{spot})$ are possible, especially if the neutron star has a large radius and is rotating rapidly. Here we extend and improve on this analysis in several ways.

Cadeau et al. (2007) considered only bolometric waveforms, whereas we consider energy-resolved waveforms, as did Lo

et al. (2013). In generating their synthetic waveforms, Cadeau et al. (1) assumed that the hotter region is infinitesimal in extent, whereas we assume a circular hot spot with an angular radius of 25°, which is more realistic (see Section 2.1); and (2) assumed that the beaming of radiation from the stellar surface is isotropic, whereas we use the Hopf beaming function (see Equation (9)), which is correct for burst atmospheres in which the opacity is dominated by electron scattering, is highly accurate for pure hydrogen atmospheres, and is fairly accurate even for atmospheres with solar composition (see Section 2.1). Cadeau et al. also (3) assumed that counts come only from the hot spot, i.e., that there are no backgrounds, whereas we include an appropriate background in our synthetic waveforms (see Section 2.1); and (4) did not Poisson-sample their waveforms, but instead assumed a fixed statistical error independent of the X-ray flux, whereas we Poisson-sample our waveforms, which gives appropriately greater statistical weight to the peaks of the waveforms.

In fitting the S+D waveform model to their synthetic waveform data, Cadeau et al. (2007) assumed (5) that the hot spot is known to be infinitesimal in extent, whereas we determine the best-fit angular radius of the spot; (6) that the beaming of radiation from the surface of the hot spot is known to be isotropic, whereas we use the Hopf beaming function; (7) that there are no background counts, whereas we include the magnitude and spectrum of the background in our model waveform and determine the background in the fitting process; (8) that the distance is known, whereas we determine the best-fit distance in the fitting process; and (9) that the phase of the model waveform relative to the phase of the synthetic observed waveform is known a priori, whereas we determine the relative phases of the model and synthetic waveforms as part of the fitting process, as would be necessary when fitting real data. Finally, Cadeau et al. (10) focused on the inferred value of GM/R_{eq} by minimizing χ^2 over all the other parameters in their waveform model, for each value of GM/R_{eq} they considered, whereas we perform an approximate Bayesian marginalization over the posterior probability space of all the parameters in our waveform model except the two parameters M and R_{eq} of interest to us here.

Our analysis therefore improves substantially on the already valuable results of Cadeau et al. (2007).

3. RESULTS

In this section, we first explain how we use the OS approximation to generate energy-resolved synthetic waveform data like the data that would be obtained by a next-generation, large-area X-ray detector when observing the X-ray oscillations produced by rotating, oblate neutron stars and hot spots with a variety of properties. Next we describe how we compute the joint posterior probability distribution of all the parameters in the waveform model being considered, given the waveform data, using standard Bayesian techniques. We then explain how we use these posterior distributions to compute confidence regions in the M–R_{eq} plane for each synthetic waveform we consider.

We present two categories of results. We first describe results obtained by fitting our standard waveform model computed using the S+D approximation to synthetic observed waveform data generated using the OS approximation, primarily to explore whether the estimated values of M and R_{eq} are significantly biased when our standard S+D waveform

model is fit to such data. We then present results obtained by fitting our standard OS waveform model to synthetic data generated using the OS approximation and compare the precision of the resulting constraints on M and R_{eq} with those obtained by Lo et al. (2013), who fit our standard S+D waveform model to waveform data generated using the S+D approximation.

Table 1 shows the 11 analyses discussed in this paper, listing in each case the waveform model and the values of the parameters in the model that were used to generate the synthetic waveform data, the model that was fit to the waveform data, the resulting minimum value of the total χ^2, the number of degrees of freedom (dof), and the figures that show the 1σ, 2σ, and 3σ contours in the M–R_{eq} plane for each case. Hereafter we use the figure label to identify each case.

3.1. Synthetic Observed Waveforms

All the synthetic observed waveforms we analyze here were generated using the OS approximation, assuming a stellar gravitational mass M of 1.6 $M_\odot$, a circular hot spot with an angular radius $\Delta\theta_{spot}$ of 25°, and a distance d to the neutron star of 10 kpc. Table 1 lists the rotational frequency ν_{rot} of the hot spot as seen at infinity, the stellar equatorial radius R_{eq}, the hot-spot inclination θ_{spot}, and the observer inclination θ_{obs} used to generate each synthetic observed waveform. We assumed ν_{rot} equal to either 300 or 600 Hz, R_{eq} equal to either 11.8 or 15 km, and θ_{spot} and θ_{obs} both equal to either 60° or 90°. We independently generate the synthetic waveform for each case that we analyze in Figures 2 and 3. In case 4 we used the same spot data as in case 3(a), but we independently generated new background data. In cases 5(a) and (b), we used the same waveform realization but two different phase binnings.

All the synthetic observed waveforms except one were generated assuming a uniform hot spot that emits radiation with a blackbody spectrum having a temperature of 2 keV as measured in the surface comoving frame. In case 3(d), we consider a hot spot with a temperature that varies with latitude (see below). Our assumption that the emission from the hot spot has a blackbody spectrum and normalization is formally inconsistent with our assumption that the beaming pattern of the radiation from the hot spot is that of an electron-scattering atmosphere (see Lo et al. 2013); when real data are analyzed, one should use emission spectra and normalizations computed using appropriate model atmospheres.

We assume that the background does not vary at frequencies commensurate with the spot rotational frequency. We include a placeholder background in the synthetic observed waveform by assuming uniform emission from the entire surface of a star with the same mass and radius as the star used in modeling the hot-spot waveform, with a spectrum having the shape of a Planck spectrum 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. In all cases except one, we assume that the background emission has a temperature of 1.5 keV. In case 4, we assume that the temperature of the background emission is the same as the temperature of the emission from the hot spot, i.e., 2 keV, in order to explore whether this weakens the derived constraints on M and R_{eq}. We normalize the background component to produce the desired number of expected background counts.

In all cases except one, we Poisson-sampled the synthetic waveforms to generate count data in 16 equally spaced phase

<div align="center">

Table 1
Synthetic Waveforms and Quality of Fit of Waveform Models

</div>

Figure (case)	Synthetic Waveform Data[a] and Fitted Waveform Model	$\nu_{\rm rot}$[b] (Hz)	$R_{\rm eq}$[c] (km)	$\theta_{\rm spot}$[d] (deg)	$\theta_{\rm obs}$[e] (deg)	$\chi^2_{\rm min}$/dof from Fit[f]
2(a)	OS data fit by S+D model	300	11.8	60	60	380.5/442
2(b)	OS data fit by S+D model	600	11.8	60	60	413.3/442
2(c)	OS data fit by S+D model	600	15	60	60	411.0/442
2(d)	OS data fit by S+D model	600	11.8	90	90	435.9/442
3(a)	OS data fit by OS model	600	11.8	90	90	440.1/442
3(b)	OS data fit by OS model	300	11.8	90	90	447.6/442
3(c)	OS data fit by OS model	600	15	60	60	475.7/442
3(d)	OS data with variation in $T_{\rm spot}$ fit by OS model	600	11.8	60	60	433.4/442
4	OS data with $T_{\rm back} = T_{\rm spot}$ fit by OS model	600	11.8	90	90	436.2/442
5(a)	OS data grouped in 32 phase bins fit by OS model	600	11.8	90	90	886.1/922
5(b)	OS data grouped in 16 phase bins fit by OS model	600	11.8	90	90	416.6/442

Notes.
[a] All synthetic waveform data were generated assuming $M = 1.6\,M_\odot$.
[b] Rotational frequency of the hot spot as seen at infinity.
[c] Equatorial circumferential radius of the star.
[d] Inclination (colatitude) of the hot spot center.
[e] Inclination of the observer.
[f] Minimized over all the parameters of the indicated model, given the data.

bins and 30 equally spaced energy channels. Each synthetic waveform therefore consists of the number of counts in each of $16 \times 30 = 480$ phase-energy bins. In case 5(a), we first Poisson-sampled the synthetic waveform in 32 bins. Thus the synthetic waveform in case 5(a) consists of the number of counts in each of $32 \times 30 = 960$ phase-energy bins. We then grouped the data in this sampled waveform into 16 bins (case 5(b)), to test whether using 16 phase bins provides adequate phase resolution. For all the synthetic waveforms except the one analyzed in case 3(d), the centroids of the energy channels are spaced 0.3 keV apart and run from 3.65 to 12.35 keV; for the waveform analyzed in case 3(d), the centroids are spaced 0.3 keV apart and run from 1.85 to 10.55 keV.

In all the synthetic observed waveforms, the expected number of counts from the spot is 10^6, whereas the expected number of counts from the background is 9×10^6 (the actual numbers vary because of the Poisson sampling). As noted in Section 2.1, these numbers produce a realistic modulation amplitude and a total number of counts comparable to the number that could be obtained by the accepted *NICER* mission and the proposed *LOFT* and *AXTAR* missions, by combining data from many bursts from a given star.

3.2. Model Waveforms and Fitting Procedure

The standard S+D and OS waveform models that we fit to the synthetic observed waveform data both assume that the temperature of the hot spot is uniform as seen in the surface comoving frame. In fitting these models to the synthetic waveform data, we assume that the rotational frequency of the star is the same as the rotational frequency $\nu_{\rm rot}$ of the hot spot and that $\nu_{\rm rot}$ is known. Both our standard waveform models have seven primary adjustable parameters (M, $R_{\rm eq}$, $\theta_{\rm spot}$, $\Delta\theta_{\rm spot}$, $kT_{\rm spot}$, $\theta_{\rm obs}$, and d) and 31 ancillary adjustable parameters (the phase-independent background in each of the 30 energy

channels and the overall time shift). Each waveform model therefore has 38 adjustable parameters, and hence there are $480 - 38 = 442$ (or, in case 5(a), $960 - 38 = 922$) dof. We assume uniform priors over the allowed ranges of all the parameters in our model waveforms. We perform a blind search over M, from 1.2 to 2.2 $M_\odot$; over $R_{\rm eq}/M$, from 4 to 8; and over $\theta_{\rm spot}$, $\Delta\theta_{\rm spot}$, and $\theta_{\rm obs}$, from 0.1 to $\pi/2$ radians. In principle, $\theta_{\rm obs}$ could range from 0 to π radians. However, we find that allowing this larger range does not change the confidence regions for M and $R_{\rm eq}$. Restricting $\theta_{\rm obs}$ to $\leqslant\pi/2$ radians allows us to sample the angular range of interest with a higher density of points.

In all cases except 3(d), we assume that kT_∞, the radiation temperature measured at infinity, is $kT_{\rm spot}(1 - 2M/R_{\rm eq})^{1/2}$. This assumption is justified because the energy of the spectral peak will in practice be very precisely measured. For all cases except 3(d), we maximize the likelihood over the distance d for a given combination of the other parameters as described in Section 2.3, rather than marginalizing over d. For a more detailed explanation and justification of this approach, see Section 3.3.3 of Lo et al. (2013). In case 3(d), the synthetic waveform data were generated with a surface comoving temperature that varies with latitude over the hot spot. Hence, in analyzing these waveform data we assume that we know d and maximize the likelihood over $kT_{\rm spot}$, rather than marginalizing over $kT_{\rm spot}$. When real data are analyzed, it will be necessary to marginalize the likelihood over both $kT_{\rm spot}$ and d.

In order to save computational time, we determine the magnitude and spectrum of the background and the time shift of the model waveform relative to the synthetic waveform by maximizing the likelihood, as described in Section 2.3, rather than by marginalizing these parameters.
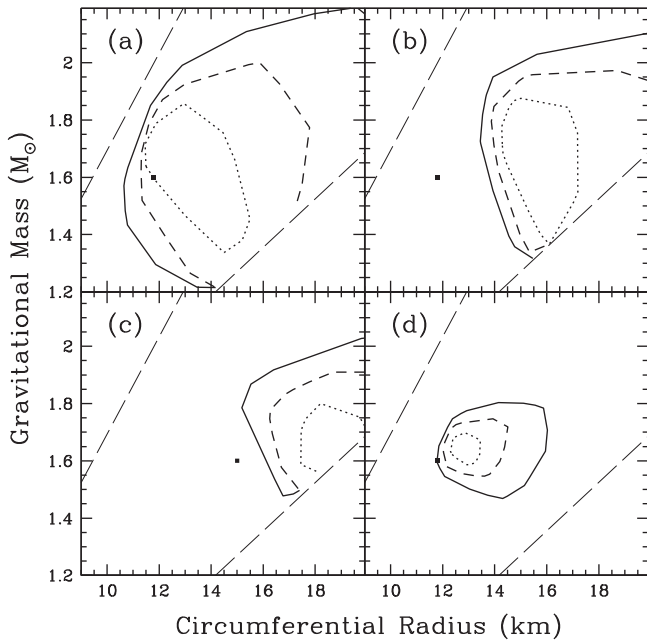
Figure 2. Constraints on M and R_{eq} obtained by fitting our standard S+D waveform model to synthetic observed waveform data generated using the OS approximation and assuming $M = 1.6\,M_\odot$. Table 1 lists the values of the other waveform parameters used to generate the data analyzed in each panel. The long-dashed lines show the $R_{eq}/M = 4$ and $R_{eq}/M = 8$ boundaries of the domain within which the posterior probability distribution was computed; the dotted, short-dashed, and solid curves show, respectively, the 1σ, 2σ, and 3σ confidence contours within this domain; and the black square indicates the values of M and R_{eq} used to generate the waveform data. In the cases shown, using the S+D model to analyze the data does not significantly bias the estimate of M. Panel (a) shows that when the center of the hot spot and the observer are at an intermediate latitude (here $60°$) and R_{eq} and ν_{rot} have intermediate values (here 11.8 km and 300 Hz), the inferred constraints on M and R_{eq} are weak and the estimate of R_{eq} is not statistically different from its true value. Panels (b) and (c) show that, in contrast, when the center of the hot spot and the observer are at an intermediate latitude and the star is rotating rapidly (here, at 600 Hz), the estimate of R_{eq} is significantly larger than its true value. Panel (d) shows that when the hot spot and the observer are on the rotational equator, the fractional bias in R_{eq} caused by using the S+D model is modest even if the star is rotating rapidly, but is nevertheless statistically significant because Req is tightly constrained.

3.3. Fits of the Schwarzschild+Doppler Model to Oblate Schwarzschild Data

Figure 2 shows the constraints on M and R_{eq} obtained by fitting our standard S+D waveform model to synthetic observed waveform data generated using the OS approximation. Figures 2(a)–(c) show the results obtained for two values of ν_{rot} and R_{eq} when the colatitude of the hot spot and the inclination of the observer are both $60°$, whereas Figure 2(d) shows a typical example of the results obtained when the colatitude of the hot spot and the inclination of the observer are both $90°$. In all these cases, fitting the data using our standard S+D model does not significantly bias the estimates of M but does significantly bias the estimates of R_{eq} in cases 2(b)–(d). The χ^2 values for these fits (see the last column of Table 1) indicate that they are all statistically good. We now discuss each of these cases in turn.

Figure 2(a) shows the constraints on M and R_{eq} obtained by fitting our standard S+D waveform model to the synthetic waveform produced by a hot spot at $\theta_{spot} = 60°$, rotating at $\nu_{rot} = 300$ Hz, on the surface of a star with a moderate radius

($R_{eq} = 11.8$ km). As noted above, in fitting our waveform models to synthetic waveform data we assume that the rotational frequency of the star is the same as the rotational frequency of the hot spot; consequently, in this case we assume that the stellar rotational frequency is 300 Hz. The oblateness of the stellar surface, and hence the deviation of the synthetic observed waveform from the waveform in the S+D model, scales as the square of the stellar rotational frequency. We therefore expect, and find, that for the moderate rotational frequency of this case, the bias in the estimated value of R_{eq} is not significant: the edge of the 1σ confidence region touches the values of M and R_{eq} that were used in generating the synthetic observed waveform, which are indicated by the black square. For the moderate spot colatitude, spot rotational frequency, and stellar radius of this case, the amplitudes of the overtones of ν_{rot} in the synthetic observed waveform are very low, and the constraints on M and R_{eq} are therefore expected to be weak (see, e.g., Lo et al. 2013). This expectation is confirmed by Figure 2(a): the 1σ contour is large, and the 2σ and 3σ contours are even larger, extending toward high M and R_{eq} and intersecting the lower boundary of our search domain at $M = R_{eq}/8$.

Figure 2(b) displays the constraints on M and R_{eq} obtained by fitting our standard S+D waveform model to the synthetic waveform produced by a hot spot that is again at $\theta_{spot} = 60°$ on a star with $R_{eq} = 11.8$ km, but now rotating at $\nu_{rot} = 600$ Hz. Even though the rotational frequency is twice that in the case featured in Figure 2(a), M and R_{eq} are still poorly constrained, because the hot spot is not near the rotational equator. The oblateness of the $R_{eq} = 11.8$ km star in Figure 2(b) should be ≈ 4 times larger than the oblateness of the 11.8 km star featured in Figure 2(a) (if the star featured in Figure 2(a) were spun up to 600 Hz, the increase in its rotational distention would cause its equatorial radius to be slightly larger than the 11.8 km radius assumed in Figure 2(b)). Figure 2(b) shows that the ≈ 4 times larger oblateness of this star is sufficient to introduce a significant bias in the values of M and R_{eq} estimated by fitting our standard S+D waveform model, which assumes the star is spherical. The values of M and R_{eq} that were used to generate the synthetic observed waveform are well outside the 3σ contour, partly because the contours in this case are modestly smaller than in the case featured in Figure 2(a), due to the star's higher rotational frequency.

Figure 2(c) shows the constraints on M and R_{eq} obtained by fitting our standard S+D waveform model to the synthetic waveform produced by a hot spot that is again at $\theta_{spot} = 60°$ and rotating at $\nu_{rot} = 600$ Hz, but now on the surface of a star with $R_{eq} = 15$ km. The contours are smaller than those in Figure 2(b) because R_{eq}, and hence the surface rotational velocity, is larger, but still extend to large M and R_{eq}, again intersecting the lower boundary of our search domain at $M = R_{eq}/8$. The estimated values of M and R_{eq} are again significantly biased.

Figure 2(d) illustrates the constraints on M and R_{eq} typically obtained by fitting our standard S+D waveform model to OS synthetic waveforms produced by a hot spot near the rotational equator, when it is viewed by an observer at a high inclination relative to the rotational axis. The confidence regions shown in this figure were obtained by fitting the synthetic OS waveform data produced by a hot spot with $\nu_{rot} = 600$ Hz, centered on the rotational equator of a star with $R_{eq} = 11.8$ km, and viewed by an observer who is in the plane defined by the star's rotational

equator. The constraints are much tighter than in the previous figures because of the high surface rotational velocity and the large projection of the velocity along the line-of-sight to the observer. The fractional biases in the estimates of M and $R_{\rm eq}$ are smaller than in the cases discussed previously but are still statistically significant, because the constraints on M and $R_{\rm eq}$ are much tighter. If the spot were infinitesimal in extent, the biases would be zero, because the stellar oblateness has no effect on the waveform produced by a point source if $dR/d\theta = 0$ at the source, and $dR/d\theta$ is zero on the rotational equator where this spot is located. However, the synthetic observed waveform used in this case was generated using a hot spot with an angular radius of $25°$, leading to a small but significant bias. Because this fit, like those discussed previously, is formally statistically acceptable, the quality of the fit does not by itself provide an indication that the M and $R_{\rm eq}$ estimates are biased.

These results show that fitting our standard S+D waveform model to OS synthetic waveform data tends to produce estimates of the star's equatorial radius that are significantly larger than its true radius, but produces estimates of the star's mass that do not differ from the true mass by a statistically significant amount. The oscillations produced by a hot spot near the star's rotational equator have the relatively strong overtones needed to provide tight constraints on M and $R_{\rm eq}$. As Figure 2(d) shows, the estimates of M and $R_{\rm eq}$ derived from such waveforms are also less susceptible to bias caused by the oblateness of the star.

We conclude that if a neutron star has a large radius or a rotational frequency $\gtrsim 300$ Hz, one should fit OS waveform models, rather than S+D waveform models, to the waveform data.

3.4. Fits of the Oblate Schwarzschild Model to Oblate Schwarzschild Data

Figure 3 shows examples of the constraints on M and $R_{\rm eq}$ obtained by fitting our standard OS waveform model to synthetic waveform data generated using the OS approximation. Figures 3(a) and (b) show the results obtained by fitting this model to the synthetic waveforms produced by a hot spot on the star's rotational equator, observed at an inclination of $90°$, for rotation rates of 600 and 300 Hz, whereas Figures 3(c) and (d) show the results obtained by fitting this model to the waveforms produced by a hot spot at a colatitude of $60°$, observed at an inclination of $60°$, for a rotation rate of 600 Hz (the waveform in case 3(d) was generated assuming a temperature variation across the hot spot, as discussed below). The χ^2 values for these fits (see the last column of Table 1) indicate that they are all statistically good. The approximate 1σ uncertainties in M and $R_{\rm eq}$ for each of these fits are listed in Table 2. Comparison of this table with Table 2 of Lo et al. (2013) shows that the constraints on M and $R_{\rm eq}$ obtained by fitting our standard OS waveform model to OS waveform data are similar to the constraints obtained there by fitting our standard S+D waveform model to S+D waveform data. We now discuss these results in more detail.

Figure 3(a) illustrates the constraints on M and $R_{\rm eq}$ typically obtained when our standard OS waveform model is fit to OS synthetic waveforms produced by a hot spot near the rotational equator, viewed by an observer at a high inclination relative to the rotational axis. The confidence regions shown in this figure were obtained by analyzing the OS waveform produced by a hot spot rotating at $\nu_{\rm rot} = 600$ Hz, centered on the rotational
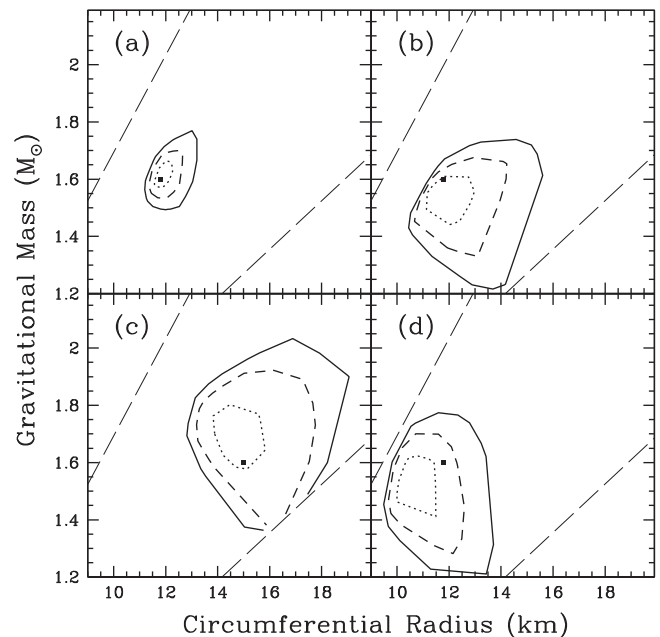
Figure 3. Constraints on M and $R_{\rm eq}$ obtained by fitting an OS waveform model to synthetic observed waveform data generated using the OS approximation. The line types and black square have the same meanings as in Figure 2. Table 2 lists the values of the parameters that were used to generate the waveform data for each panel. Panel (a) shows that when the center of the hot spot is on the rotational equator, the observer is in the plane of the star's rotational equator, and the star is rotating rapidly (here, at 600 Hz), M and $R_{\rm eq}$ are tightly constrained. The constraints on M and $R_{\rm eq}$ here are similar to those in Figure 2 (e) of Lo et al. (2013), which considers the same spot geometry and rotation rate. The parameter values used to generate the waveform data used in panel (b) are the same as in panel (a), except for the rotation rate, which is much lower (300 Hz), causing the constraints on M and $R_{\rm eq}$ to be substantially weaker. Panel (c) shows that when the spot and observer are at an intermediate colatitude (here $60°$), the constraints on M and $R_{\rm eq}$ are much weaker, even if the star has a large radius (here 15 km) and is rapidly rotating (here, at 600 Hz). Panel (d) shows the effect of fitting an OS waveform model that assumes a uniform hot spot to OS waveform data generated using a spot with a temperature that varies in the north–south (latitudinal) direction by 25% (see text for details). This result shows that using a hot spot model that differs from the actual spot in this way does not significantly bias the estimates of M and $R_{\rm eq}$.

equator of a star with $R_{\rm eq} = 11.8$ km, when seen by an observer who is in the plane defined by the star's rotational equator. The constraints in Figure 3(a) are tighter than in Figure 2(d) and the estimates of M and $R_{\rm eq}$ are not significantly biased. The constraints in Figure 3(a) (1σ uncertainties of 2.9% in $R_{\rm eq}$ and 2.8% in M; see Table 2) are much tighter than in Figure 3(b), which shows results for a much lower rotational frequency (300 Hz), and much tighter than in Figure 3(c), which shows results for a hot spot that is not near the rotational equator ($\theta_{\rm obs} = 60°$).

It is useful to compare Figure 3(a) with Figure 2(e) of Lo et al. (2013). In both cases, the synthetic waveform corresponds to a hot spot rotating at $\nu_{\rm rot} = 600$ Hz, centered on the rotational equator of a star with $R_{\rm eq} = 11.8$ km, and seen by an observer in the plane defined by the star's rotational equator. The 1σ confidence region in Figure 2(a) is comparable in size to the 1σ confidence region in Figure 2(e) of Lo et al. (2013), but the 3σ region in Figure 3(a) is much smaller than the 3σ region in Figure 2(e) of Lo et al. (2013).[5] There are at

---

[5] The definitions of the 1σ uncertainties in M and $R_{\rm eq}$ used here differ from those used in Lo et al. (2013) (see footnote j of Table 2).

Table 2
Constraints on M and R_{eq} Obtained Using Our Standard OS Waveform Model[a]

Figure (case)	ν_{rot}[b] (Hz)	θ_{spot}[c] (deg)	θ_{obs}[d] (deg)	R_{eq}[e] (km)	ΔR_{eq}[f] (km)	ΔM[g] ($M_\odot$)	$\delta R_{eq,1}$[h,j] (%)	δM_1[i,j] (%)
3(a)	600	90	90	11.8	11.57–12.27	1.57–1.66	2.9	2.8
3(b)	300	90	90	11.8	11.14–12.97	1.44–1.61	7.6	5.6
3(c)	600	60	60	15	13.83–15.83	1.58–1.80	6.7	6.5
3(d)	600	60	60	11.8	9.99–11.49	1.41–1.62	7.0	6.9
4	600	90	90	11.8	11.35–12.19	1.49–1.59	3.6	3.2
5(a)	600	90	90	11.8	11.63–12.03	1.56–1.64	1.7	2.5
5(b)	600	90	90	11.8	11.68–12.24	1.56–1.64	2.3	2.5

Notes.
[a] The 1σ M and R_{eq} uncertainties listed here somewhat understate the actual constraints on M (R_{eq}), because M/R_{eq} is usually better constrained than M or R_{eq} considered separately (see, e.g., Figure 3(a)). All the synthetic waveform data were generated assuming $M = 1.6 \, M_\odot$. The data analyzed in case 3(d) were generated assuming a temperature variation across the hot spot (see text).
[b] Rotational frequency of the hot spot as seen at infinity assumed in generating the synthetic observed waveform data.
[c] Inclination (colatitude) of the hot-spot center assumed in generating the synthetic observed waveform data.
[d] Inclination of the observer assumed in generating the synthetic observed waveform data.
[e] Equatorial radius assumed in generating the synthetic observed waveform data.
[f] Range of the 1σ contour projected onto the R_{eq} axis.
[g] Range of the 1σ contour projected onto the M axis.
[h] Approximate 1σ fractional uncertainty in R_{eq} computed by dividing one-half the 1σ range of R_{eq} by its central value.
[i] Approximate 1σ fractional uncertainty in M computed by dividing one-half the 1σ range of M by its central value.
[j] The definitions of the 1σ uncertainties in M and R_{eq} used here differ from those used in Lo et al. (2013), where they were estimated by projecting the *full* extent of the 1σ contour onto the M and R_{eq} axes, because in that work some of the 1σ confidence regions were highly asymmetric, with best-fit values of M and R_{eq} far from the center.

least four possible explanations for this difference: (1) a statistical fluctuation in the synthetic waveform data of Lo et al. (2013) made that waveform realization less constraining than it would typically be, (2) a statistical fluctuation in the synthetic waveform data used here made this waveform realization more constraining than it would typically be, (3) our new analysis pipeline does a better job of representing the true constraints, or (4) something about OS waveforms actually yields better 3σ (but not 1σ) constraints in this situation. It is not clear without further exploration which, if any, of these explanations is responsible for this difference.

Figure 3(b) shows the constraints on M and R_{eq} obtained by analyzing a synthetic waveform produced by the same hot spot and observer geometry as in Figure 3(a) (the hot spot is on the rotational equator of a star with $R_{eq} = 11.8$ km and is seen by an observer who is in the plane defined by the star's rotational equator), but for a rotation rate of 300 Hz, half the rotation rate assumed in Figure 3(a). The slower rotation rate decreases the harmonic content of the waveform, which increases the sizes of the confidence regions, but the 1σ uncertainties (5.6% in M and 7.6% in R_{eq}; see Table 2) are still interestingly small. This figure shows that fitting the waveforms of stars that rotate at a moderate rate can provide interesting constraints on M and R_{eq}, provided that the hot spot is near the rotational equator and the observer is at a high inclination.

Figure 3(c) displays the constraints on M and R_{eq} obtained by analyzing a synthetic waveform produced by a hot spot on a star with a larger radius ($R_{eq} = 15$ km) again rotating at 600 Hz, but with the spot at a colatitude of 60° and the observer at an inclination of 60°. The constraints are much less precise than for the case shown in Figure 3(a) and less precise overall than for the case shown in Figure 3(b), even though the star has a larger radius, because the hot spot is not near the rotational equator and the observer's inclination is not close to 90°. Even

so, the 1σ uncertainties (6.5% in M and 6.7% in R_{eq}; see Table 2) are still interestingly small.

Figure 3(d) shows a case in which the properties of the hot spot assumed in our standard OS waveform model are different from the properties used to generate the synthetic waveform data to which the model was fit. In this case, the fitted model assumes that the surface temperature is uniform across the hot spot, as seen in the frame comoving with the surface, whereas the synthetic waveform data were generated assuming that the temperature varies by 25% with colatitude, from 2 keV at the center of the spot to 1.5 keV at the top and bottom edges of the spot. This case tests whether such a variation, which is physically plausible, produces a significant bias in the estimated values of M and R_{eq}. As Figure 3(d) shows, there is a slight bias, but it is not statistically significant: the best-fit values of M and R_{eq} differ by only $\sim 1.5\sigma$ from the values used to generate the synthetic waveform data.

Figure 4 explores the effects on the mass and radius constraints if there is less contrast between the energy spectrum of the emission from the hot spot and the spectrum of the phase-independent background. For this analysis, we took the counts from the hot spot in each phase-energy bin generated for the analysis shown in Figure 3(a), but generated a new background having $\approx 9 \times 10^6$ counts and a spectrum identical to that of the hot spot (i.e., a Planck spectrum with a surface comoving temperature of 2.0 keV instead of the value of 1.5 keV used in generating our other synthetic waveforms). Hence, in this case the emission from the spot is simply extra emission on top of the background, rather than extra emission with a different spectrum. The constraints obtained are only marginally worse than in case 3(a). Given the randomness inherent in the generation of the synthetic data and in the MCMC analysis of the data, this difference may not be significant. This result shows that even if the phase-
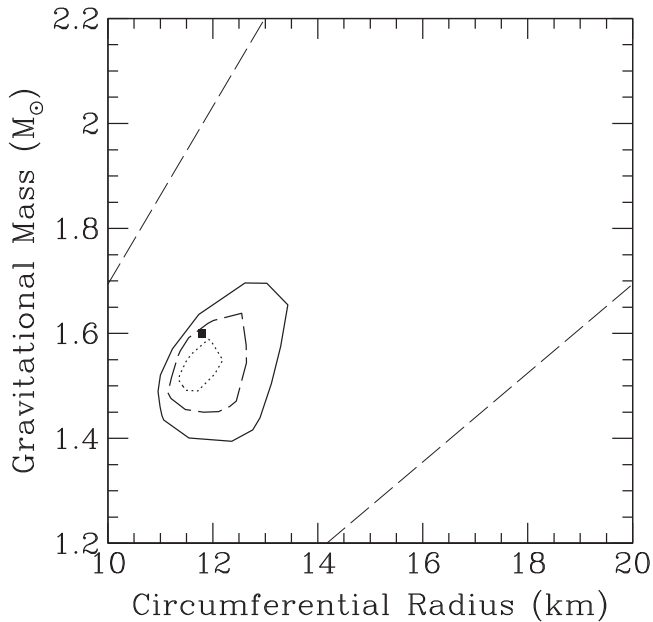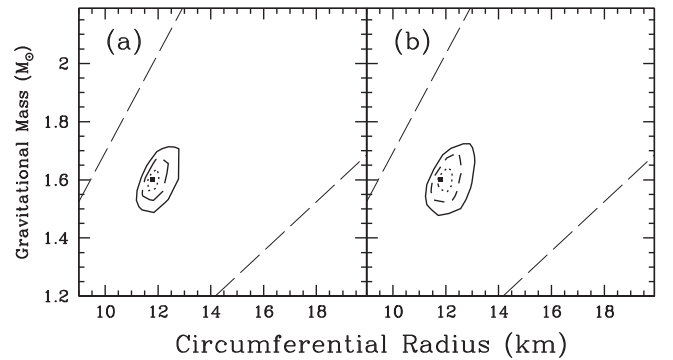
Figure 4. Constraints on M and R_{eq} obtained by fitting our standard OS waveform model to synthetic observed waveform data generated using the OS approximation with a background that has the same spectrum as the emission from the hot spot (a Planck spectrum with a surface comoving temperature of 2 keV). In this case the hot spot appears as extra emission from the surface with the same spectrum as the background. The synthetic waveform was generated assuming that the hot spot is centered on the rotational equator of a star rotating at 600 Hz and that the observer is in the plane of the rotational equator. The line types and black square have the same meanings as in Figure 2. This result shows that even if the phase-independent background has the same spectrum as the emission from the hot spot, one can obtain good constraints on M and R_{eq}.

Figure 5. Comparison of the constraints on M and R_{eq} obtained by fitting our standard OS waveform model to a single realization of synthetic observed waveform data generated using the OS approximation and then grouped into 32 (panel (a)) and 16 (panel (b)) phase bins. The synthetic waveform was computed assuming that the hot spot is centered on the rotational equator of a star rotating at 600 Hz and that the observer is in the plane of the star's rotational equator. The line types and black square have the same meanings as in Figure 2. This result shows that increasing the number of phase bins beyond 16 does not appear to improve significantly the precision of the constraints on M and R_{eq} (see also the discussion in the text).

independent background has the same spectrum as the emission from the hot spot, one can obtain good constraints on M and R_{eq}.

Figure 5 explores whether 16 phase bins are adequate for the cases considered here. We first generated a synthetic observed waveform using the OS approximation and 32 phase bins, for a hot spot centered on the rotational equator of a star rotating at 600 Hz and an observer in the plane of the star's rotational equator. We then regrouped these same data into 16 phase bins. Figure 5 shows that the constraints on M obtained by analyzing the data binned in these two different ways are essentially identical, whereas the constraint on R_{eq} appears slightly tighter when 32 phase bins are used. The randomness in the generation of each synthetic waveform and in the MCMC analysis of a given synthetic waveform produces variations in the derived constraints; for example, Table 2 indicates that the 1σ constraints derived in case 5(b) are slightly tighter than in case 3(a), even though both used waveform data with 16 phase bins. Thus, the apparent slight improvement in the constraint on R_{eq} when 32 phase bins are used may not be significant.

3.5. Origins and Sizes of the Uncertainties in M and R_{eq} Estimates

Understanding the uncertainties δM and δR_{eq} in M and R_{eq} estimates requires understanding how the properties of the observed waveforms constrain M and R_{eq}. This is explained in detail by Lo et al. (2013; see also Psaltis et al. 2014). The 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 rotational frequency of the emitting region is accurately known from the oscillation frequency, knowing the line-of-sight velocity of the emitting region constrains the stellar radius. The amplitude of the waveform constrains the colatitude of the hot spot, the observer's inclination, and the compactness (M/R) of the star, the last primarily via general relativistic light-bending effects.

Although the harmonic content of the waveform encodes information about the rotational velocity of the emitting element—and hence the cylindrical radius of the element (see Lo et al. 2013, Section 2.2.2)—other unrelated aspects of the system also affect the harmonic structure of the waveform. In particular, the semi-amplitude C_2 of the second harmonic in the waveform is not uniquely related to the rotational velocity of the emitting element. Consequently, there is in general no simple way to extract information about R_{eq} from the harmonic content of the waveform; model fitting is required.

Effects other than the rotational velocity that contribute to C_2 include (1) anisotropic beaming of the radiation from each emitting element and (2) occultation of part or all of the hot spot by the star, as it rotates (see Poutanen & Beloborodov 2006, which provides a useful guide to waveform properties, even though the results reported there assume the hot spot is infinitesimal in extent and were derived using an approximate analytic expression for the light deflection). To lowest order, the line-of-sight linear velocity of the emitting gas v_{los} contributes a second harmonic with semi-amplitude $C_2 \propto (v_{los}/c)C_1$, where c is the speed of light, C_1 is the semi-amplitude of the first harmonic (the fundamental), and the coefficient of proportionality depends on the spectrum of the emission (see Poutanen & Beloborodov 2006, Equation (66)). The linear velocity produced by the rotation of the emitting gas therefore contributes a second harmonic with semi-amplitude

$$C_2 \propto \left(2\pi\nu_{rot}R_{eq}/c\right)\sin\theta\,\sin\theta_{obs}, \qquad (14)$$

where θ is the colatitude of the surface element from which the radiation is emitted. Because this contribution to C_2 is $\mathcal{O}(v_{los}/c)C_1$ and $v_{los}/c \ll 1$, it is generally $\ll C_1$. In contrast,

Table 3
Waveform Harmonic Content, $\mathcal{R}$ Values, and Uncertainties in M and R_{eq}[a]

Figure (case)	ν_{rot} (Hz)	θ_{spot} (deg)	θ_{obs} (deg)	f_{rms1}[b] (%)	f_{rms2}[c] (%)	f_{rms3}[d] (%)	f_{rms}[e] (%)	$\mathcal{R}$[f]	$\delta R_{eq,1}$ (%)	δM_1 (%)
3(a)	600	90	90	10.0	3.9	1.4	11	479	2.9	2.8
3(b)	300	90	90	9.5	2.7	0.6	10	440	7.6	5.6
3(c)	600	60	60	8.1	2.6	0.7	9	377	6.7	6.5
4	600	90	90	10.0	3.8	1.3	11	479	3.6	3.2
5(b)	600	90	90	9.9	3.8	1.3	11	471	2.3	2.5

Notes.
[a] See the notes to Table 2 for the definitions of ν_{rot}, θ_{spot}, θ_{obs}, $\delta R_{eq,1}$, and δM_1. As noted there, the uncertainties listed here somewhat understate the actual constraints on M (R_{eq}), because M/R_{eq} is usually better constrained than M or R_{eq} considered separately. All the synthetic waveform data were generated assuming $M = 1.6\,M_\odot$.
[b] Fractional rms amplitude of the first harmonic (fundamental) component of the synthetic waveform.
[c] Fractional rms amplitude of the second harmonic component of the synthetic waveform.
[d] Fractional rms amplitude of the third harmonic component of the synthetic waveform.
[e] Fractional rms amplitude of the total variation of the synthetic waveform.
[f] Synthetic observed waveform figure of merit $\mathcal{R} \equiv 1.4 f_{rms}\sqrt{N_{tot}}$; see Equation (1).

the beaming of emission from the atmosphere of the hot spot (see Section 2.1.3) contributes a second harmonic with semi-amplitude $C_2 \approx C_1$. In particular, emission with the anisotropic beaming pattern $I(\alpha') = I_0(1 + h\cos\alpha')$ produces a second harmonic with semi-amplitude $C_2 \approx h\sin\theta\sin\theta_{obs}C_1$ (see Poutanen & Beloborodov 2006, Equation (50)); for a burst atmosphere, $h \approx 0.92$ (see Equation (9)). Because the component of the second harmonic contributed by the anisotropy of the radiation from the hot spot is independent of the star's rotational frequency, it is $\approx C_1$ even if the star is rotating slowly. Occultation of part or all of the hot spot by the star can also contribute a second harmonic component with a semi-amplitude $C_2 \approx C_1$. Thus, in burst oscillations the second harmonics generated by other effects can be comparable to or larger than the second harmonic produced by the rotational velocity of the star.

In order to obtain useful constraints on M and R_{eq}, the velocity of the hot spot surface must make a significant contribution to the harmonic content of the waveform. Equation (14) shows that this contribution is greater if the spot and the observer's sightline are closer to the star's rotational equator. Then the relativistic Doppler shift and aberration are greater, the waveform depends more sensitively on the radius of the star, and the constraints on M and R_{eq} are tighter.

The uncertainties in estimates of M and R_{eq} also depend on the fractional amplitude of the oscillation and the total number of counts, including the number of background counts, which we define as all counts not produced by photons from the hot spot. As explained in Section 2.1.1, if all the other properties of the system that affect the waveform are kept fixed, the uncertainties in estimates of M and R_{eq} obtained by fitting waveform models to waveform data scale approximately as $\mathcal{R}^{-1}$, where $\mathcal{R} = 1.4 f_{rms}\sqrt{N_{tot}}$. Here f_{rms} is the average fractional rms amplitude of the oscillation during the observation and N_{tot} is the total number of detected counts.

The dependence of the waveform harmonic content on the rotational frequency of the star, the inclinations of the spot and the observer, and $\mathcal{R}$ are illustrated by the results presented in Table 3 (compare Table 4 of Lo et al. 2013). The amplitudes of

the higher harmonics are substantially smaller in case 3(b) than in case 3(a) because the lower rotational frequency in the former case produces a line-of-sight velocity a factor of two smaller than in case 3(a). Comparing case 3(c) with case 3(a) demonstrates the sensitivity of the harmonic content of the waveform to the inclinations of the spot and the observer. Even though the line-of-sight velocity in case 3(c) is 95% of that in case 3(a), because the larger radius of the star in case 3(c) almost compensates for the higher inclinations of the spot and observer, the higher inclinations reduce the amplitudes of all the harmonics in case 3(c) compared to case 3(a). The $\sim$20% reduction in the total rms amplitude in case 3(c) reduces $\mathcal{R}$ by a comparable amount. The lower amplitudes of the higher harmonics in cases 3(b) and (c) increase the uncertainties δM and δR_{eq} in these cases (see the final columns of Table 3).

The dependence of δM and δR_{eq} on the rotational frequency of the star and the inclinations of the spot and the observer are also illustrated by the trends in the confidence regions for M and R_{eq} in Figures 3(a)–(c) (see also Section 4.2.3 of Lo et al. 2013). As noted above, the lower rotational frequency in case 3(b) produces a smaller line-of-sight velocity, and hence a larger confidence region for M and R_{eq} than in case 3(a). The smaller inclinations of the spot and the observer in case 3(c) produce a larger confidence region for M and R_{eq} than in case 3(a), even though the line-of-sight velocity in case 3(c) is almost the same as in case 3(a).

These results and those of Lo et al. (2013) show that δM and δR_{eq} are most sensitive to the colatitude of the hot spot and the inclination of the observer. They are sensitive, but less so, to the radius and the rotational frequency of the star. They depend more weakly on the background.

Three of the cases featured in Table 3 (cases 3(a), 4, and 5(b)) analyze different realizations of synthetic waveforms generated using identical values of the OS waveform parameters. Comparing these cases therefore provides insight into the sizes of the statistical and sampling errors in our results. The differences between the amplitudes of the harmonics in these three waveforms reflect the different shapes of the waveforms produced by fluctuations in the number of counts in each energy and phase bin; they indicate that the variations in the waveforms caused by these fluctuations are $\lesssim$2%. The fractional differences in the 1σ uncertainties in M

and $R_{\rm eq}$ are much larger and are probably caused by the limitations of our sampling of the parameters of the model waveform during the fitting process; these differences indicate that the sampling errors are ~20% in $\delta R_{\rm eq}$ and ~10% in δM.

As a specific example, suppose that (1) $\nu_{\rm rot} = 600$ Hz; (2) $\theta_{\rm spot} = 90°$ and $\theta_{\rm obs} = 90°$; (3) the number of counts from the hot spot is equal to the number of background counts, producing a fractional rms amplitude of ~54%; (4) the total number of detected counts is $\sim 2 \times 10^6$; and (5) the values of all the parameters in the waveform model except M and $R_{\rm eq}$ are known independently of the waveform-fitting procedure. The uncertainties in the estimates of M and $R_{\rm eq}$ would then be ~2% in M and ~1% in $R_{\rm eq}$ (see Section 4.2.1 and Table 2 of Lo et al. 2013).

Realistically, in addition to M and $R_{\rm eq}$, the values of some or all of the other parameters in the waveform model will have to be determined as part of the waveform-fitting procedure. Determining these additional parameters as part of the waveform-fitting process produces larger uncertainties in the estimates of M and $R_{\rm eq}$. The reason is that the effects on the waveform of changing different parameters in the waveform model are often very similar (see Section 4.2.2 of Lo et al. 2013 for a detailed discussion of these degeneracies). These degeneracies with respect to changes in the values of waveform model parameters are an inherent property of any physical model based on a rotating hot spot and cannot be removed by improving the model.

The number of background counts in observed burst oscillation waveforms is expected to be much greater than the number of counts collected from the hot spot (see Section 2.1.1 above and Section 2.2.1 of Lo et al. 2013). A large background increases the effects of the parameter degeneracies, because it increases the statistical fluctuations in the observed waveform. As a result, a wider range of model waveforms will adequately fit the waveform data. If the number of background counts is much greater than the number of oscillating counts, and the number of oscillating counts and the geometry remain unchanged, the uncertainties in M and $R_{\rm eq}$ increase with the number of background counts as $\mathcal{R}^{-1} \propto \sqrt{N_{\rm back}}$ (see Section 2.1.1).

The uncertainties in M and $R_{\rm eq}$ are much more sensitive to the inclinations of the hot spot and observer than to the background. This can be seen by comparing cases 2(a), (c), and (e) of Lo et al. (2013), which have the same inclinations but very different backgrounds. The relatively modest effect of an unknown background on the uncertainties in M and $R_{\rm eq}$ can be seen by comparing cases 5(a) and (b) of Lo et al. (2013), which assume the background is known exactly, with their corresponding cases 2(c) and (d), which assume the background is unknown and must be determined as part of the waveform-fitting procedure.

As an example of a realistic situation, suppose that (1) $\nu_{\rm rot} = 600$ Hz; (2) $\theta_{\rm spot} = 90°$ and $\theta_{\rm obs} = 90°$; (3) the number of background counts is 9 times the number of counts from the hot spot, producing a fractional rms amplitude of 11%; (4) the total number of detected counts is 10^7; and (5) the values of all the parameters in the waveform model, including M and $R_{\rm eq}$, must be determined as part of the waveform-fitting procedure. Then the uncertainties in the derived estimates of M and $R_{\rm eq}$ would be ~3% (see Table 2).

Independent knowledge of some of the system parameters can reduce or eliminate degeneracies, reducing the uncertainties in M and $R_{\rm eq}$. 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 (see Lo et al. 2013, Section 4.2.5).

Psaltis et al. (2014) proposed a simple formula for estimating the fractional uncertainty in estimates of $R_{\rm eq}$ obtained by fitting waveform models to waveform data. This formula is

$$\frac{\delta R_{\rm eq}}{R_{\rm eq}} = \left[\left(\frac{4\pi \nu_{\rm rot} R_{\rm eq}}{c} \right) \sin\theta_{\rm spot} \sin\theta_{\rm obs} \right]^{-1} \left(\frac{\sqrt{S+B}}{c_1 S} \right), \quad (15)$$

where S is the number of counts from the hot spot, B is the number of background counts, and c_1 is the fractional rms amplitude of the fundamental (first harmonic) in the waveform, defined in terms of S; the other symbols have their previous meanings. This formula is based on the following assumptions and approximations: (1) a single-energy or bolometric analysis is adequate; (2) the hot spot is infinitesimal in extent; (3) the distortion of the waveform produced by the surface rotational velocity is the dominant source of a nonzero second harmonic in the waveform; (4) the surface rotational velocity contributes a second harmonic with amplitude $C_2 = 2(v_{\rm los}/c)\, C_1$; (5) $\nu_{\rm rot}$, $R_{\rm eq}$, $\sin\theta_{\rm spot}$, and $\sin\theta_{\rm obs}$ are all known independently of the waveform-fitting process.

Several of the assumptions and approximations on which Equation (15) is based can be questioned for X-ray burst oscillations: (1) the use of energy-resolved waveforms provides better results than single-energy or bolometric waveforms (see, e.g., Lo et al. 2013; Psaltis et al. 2014); (2) it will probably be necessary to use data from the peaks and/or tails of bursts, when the hot spot is not infinitesimal in extent (see Section 2.1.1); (3) the distortion of the waveform produced by the surface rotational velocity is not the dominant source of a nonzero second harmonic for the radiation beaming pattern expected for X-ray burst oscillations (see above and Section 2.1.3); (4) the assumed coefficient of proportionality between C_2 and C_1 is not the one derived by Poutanen & Beloborodov (2006), which depends on the spectrum of the emission (see their Equation (66)); (5) in most cases, some or all of the parameters $R_{\rm eq}$, $\sin\theta_{\rm spot}$, and $\sin\theta_{\rm obs}$ will not be known a priori, and will have to be determined by the waveform-fitting process. Assuming that the values of these parameters are known a priori completely eliminates the degeneracies among them, and therefore greatly overestimates the precision with which $R_{\rm eq}$ can realistically be determined using a given data set. (Although the fractional rms amplitude c_1 of the first harmonic component of the waveform, the number of counts S from the hot spot, and the number of background counts B are not directly observable (see Section 2.1.2) and would have to be determined by waveform fitting, the total number of counts $N_{\rm tot} = S + B$ and the number of counts $c_1 S$ in the first harmonic component of the waveform are directly observable.)

How much do the simplifying assumptions made in deriving Equation (15) affect the computed uncertainty in $R_{\rm eq}$? To investigate this, we assume that the values of all the system parameters needed to evaluate Equation (15) are somehow known without doing any waveform fitting and evaluate

Equation (15) using the values of these parameters that we used to generate the synthetic waveforms in our cases 3(a)–(c). We then compare the values of δR_{eq} given by Equation (15) with the 1σ uncertainties in R_{eq} we obtained by fitting our standard OS waveform model to the corresponding synthetic waveform data, using our Bayesian approach. In case 3(a), $\nu_{rot} = 600$ Hz, $R_{eq} = 11.8$ km, $\theta_{obs} = \theta_{spot} = 90°$, $S = 10^6$, $B = 9 \times 10^6$, and $c_1 = 1.077$ (recall that the fractional rms amplitude can exceed unity). Inserting these values into Equation (15) yields $\delta R_{eq}/R_{eq} = 0.0099$, whereas our Bayesian statistical analysis yields $\delta R_{eq}/R_{eq} = 0.029$. In case 3(b), $\nu_{rot} = 300$ Hz, $R_{eq} = 11.8$ km, $\theta_{obs} = \theta_{spot} = 90°$, $S = 10^6$, $B = 9 \times 10^6$, and $c_1 = 0.990$. Inserting these values into Equation (15) yields $\delta R_{eq}/R_{eq} = 0.022$, whereas our Bayesian analysis yields $\delta R_{eq}/R_{eq} = 0.076$. Finally, in case 3(c), $\nu_{rot} = 600$ Hz, $R_{eq} = 15$ km, $\theta_{obs} = \theta_{spot} = 60°$, $S = 10^6$, $B = 9 \times 10^6$, and $c_1 = 0.856$. Inserting these values into Equation (15) yields $\delta R_{eq}/R_{eq} = 0.013$, whereas our Bayesian analysis yields $\delta R_{eq}/R_{eq} = 0.067$. Thus, in these cases the uncertainties in R_{eq} obtained by fitting our standard OS waveform model to synthetic waveform data using our Bayesian statistical approach are $\sim$3–5 times larger than the uncertainties given by Equation (15). In reality, the values of all the system parameters needed to evaluate Equation (15) usually will not be known.

Equations (14) and (15) suggest that the uncertainty in R_{eq} should be smaller for larger values of ν_{rot}, R_{eq}, $\sin \theta_{spot}$, and $\sin \theta_{obs}$, other things being equal. Our waveform-fitting results (see Table 2) are qualitatively consistent with this behavior but, as discussed above, reveal important quantitative differences from the scaling implied by these equations. This is not surprising, because Equations (14) and (15) do not include some important effects, such as occultation of the hot spot by the star. This can occur when the spot is near the rotational equator, and if it does, it will reduce further the uncertainties in M and R_{eq}. Thus, there is no assurance that the actual uncertainty in R_{eq} will scale with ν_{rot}, R_{eq}, $\sin \theta_{spot}$, and $\sin \theta_{obs}$ as simply as suggested by Equations (14) and (15). Consequently, choosing targets for observation and planning observational campaigns should be done using uncertainties in M and R_{eq} obtained by fitting waveform models to appropriate synthetic waveform data.

Tables 2 and 3 show the results of our Bayesian analysis of the uncertainties in M and R_{eq} that can be obtained by fitting our standard OS waveform model to burst oscillation waveform data, represented here by synthetic waveform data generated using the OS approximation. Figures 3, 4 and 5 shows the 1σ, 2σ, and 3σ confidence contours in the M–R_{eq} plane for these cases. We find that M and R_{eq} can both be estimated with 1σ uncertainties $\lesssim$7% if (1) the star's rotation rate is $\gtrsim$600 Hz, (2) the hot spot is located at a colatitude $\gtrsim$60°, (3) the star is observed at a rotational inclination $\gtrsim$60°, (4) the oscillations have a fractional rms modulation $\gtrsim$10%, and (5) $\gtrsim 10^7$ total counts are collected from the star.

4. CONCLUSIONS

We have extended the analysis by Lo et al. (2013) of the constraints on M and R_{eq} that can be achieved by fitting waveform models to burst oscillation waveform data, by fitting our standard S+D and OS waveform models to synthetic observed waveform data generated using the OS approximation.

We find that if the neutron star has a moderately large radius and is rapidly rotating and the hot spot that produces the oscillation is at a moderate to low rotational colatitude, fitting our standard S+D waveform model to synthetic waveform data generated using the OS approximation can produce fits that are statistically good but yield estimates of M and R_{eq} that have significant biases. However, this spot geometry generally does not lead to tight constraints on M and R_{eq} (see, e.g., Cadeau et al. 2007, Table 2; Lo et al. 2013, Table 2), because it produces waveforms in which the oscillation amplitude is low and overtones of the rotational frequency are very weak. If instead the hot spot is at a high rotational colatitude, fitting our standard S+D waveform model to OS synthetic waveform data can yield usefully tight constraints on M and R_{eq} with much smaller biases, even for rapidly rotating, oblate stars. However, our improved analysis procedure makes it possible to fit our standard OS waveform model to waveform data almost as quickly as our standard S+D waveform model. Consequently, even though our standard S+D waveform model is likely to be adequate when analyzing waveforms produced by the spots located near the star's rotational equator that will yield the tightest constraints on M and R_{eq}, there is now no reason not to use OS waveform models for all waveform analyses.

We find that fitting our standard OS waveform model to OS waveform data produces tight constraints on M and R_{eq} if the star has a moderate to high rotation rate, the hot spot is at a moderate to high rotational colatitude, and the observer is at a moderate to high inclination. Specifically, our results show that if the star's rotation rate is $\gtrsim$600 Hz, the spot center and the observer's sightline are both within 30° of the star's rotational equator, the fractional rms amplitude of the oscillations is $\gtrsim$10%, and $\gtrsim 10^7$ counts can be collected from the star, M and R_{eq} can both be determined with 1σ uncertainties $\lesssim$7%. This is a realistic fractional amplitude, and this many counts could be obtained from a single star by the accepted NICER and proposed LOFT and AXTAR space missions by combining data from many X-ray bursts. If the spot center and the observer's sightline are both close to the star's rotational equator, M and R_{eq} can be determined with 1σ uncertainities $\lesssim$3%. If the star's rotation rate is $\gtrsim$300 Hz and the spot center and the observer's sightline are both close to the star's rotational equator, M and R_{eq} can be determined with 1σ uncertainties $\lesssim$8%. Independent knowledge of the observer's inclination can reduce these uncertainties. Simultaneous measurements of M and R_{eq} with these precisions would improve substantially our understanding of cold, ultradense matter.

Comparison of the constraints on M and R_{eq} obtained by fitting our standard OS waveform model to a single realization of OS synthetic observed waveform data generated using the OS approximation and then grouped into 32 and 16 phase bins indicates that increasing the number of phase bins beyond 16 does not improve significantly the precision of the constraints on M and R_{eq}. We also investigated the constraints on M and R_{eq} obtained when the background has the same spectrum as the emission from the hot spot, by fitting our standard OS waveform model to synthetic observed waveform data generated using the OS approximation, assuming the emission from the hot spot and the background have the same spectrum. The result shows that one can obtain good constraints on M and

$R_{\rm eq}$ even if the phase-independent background has the same spectrum as the emission from the hot spot.

A key finding of Lo et al. (2013) was that M and $R_{\rm eq}$ estimates derived by fitting S+D waveform models to S+D synthetic waveform data were not significantly biased when a fit was both statistically good and highly constraining, even when the spectrum, beaming function, or spot shape assumed in the model differed substantially from those assumed in generating the waveform data. Here we extended this investigation by exploring the effect on estimates of M and $R_{\rm eq}$ of a 25% variation of the hot-spot temperature in the north–south direction. We find that such a temperature variation does not produce a significant bias in the estimated values of M and $R_{\rm eq}$. Thus, we still have not found a case in which a difference between the assumed properties of the hot spot in the fitted model and those of the hot spot that produced the observed waveform yields a fit that is both statistically good and highly constraining but gives M or $R_{\rm eq}$ estimates that are significantly biased. Consequently, we are cautiously optimistic that fitting model waveforms to burst oscillation data will provide measurements of neutron star masses and radii that are accurate as well as precise.

Finally, we comment that although the primary application of our work is to X-ray burst oscillation waveforms, our methods can, with small changes, be used to analyze the X-ray oscillations produced by the heated polar caps of rotation-powered pulsars, which is the focus of the *NICER* mission.

APPENDIX
CONSTRUCTION OF A MINIMUM BOUNDING ELLIPSOID

As we discussed in Section 2.4, the accuracy of our marginalization over $\theta_{\rm obs}$, θ_c, and $\Delta\theta_{\rm spot}$ is improved considerably when, for a given $\theta_{\rm obs}$, we construct the minimum ellipse that bounds the (θ_c, $\Delta\theta_{\rm spot}$) combinations that give good fits to the data then sample only this volume in our Monte Carlo integration. It is, however, not trivial to construct such a minimum ellipse. We therefore present in this appendix an algorithm for computing an almost minimum volume ellipsoid around a given set of points in any number of dimensions. This algorithm was originally derived by Kachiyan (1996), but we follow here the discussion of Todd & Yıldırım (2007). More details may be found in Todd & Yıldırım (2007); here we present only the essential formulae.

Suppose we have a set of points x in d dimensions. An ellipsoid that contains all of these points can be defined by its d-dimensional center c and a $d \times d$ symmetric, positive-definite matrix Q, where for any point x in the set

$$(x - c)^T \cdot Q \cdot (x - c) \leqslant 1. \qquad (16)$$

We can write $(x - c)$ in component notation as

$$(x - c) = \begin{pmatrix} x_1 - c_1 \\ x_2 - c_2 \\ \dots \\ x_d - c_d \end{pmatrix}, \qquad (17)$$

where the subscripts $1, 2, \dots, d$ represent each of the d dimensions, and $(x - c)^T$ is the transpose of $(x - c)$, i.e., $(x_1 - c_1, x_2 - c_2, \dots, x_d - c_d)$. The volume of this ellipsoid is $V_d (\det Q)^{-1/2}$, where V_d is the volume of the unit ball in d dimensions (i.e., π in two dimensions, $4\pi/3$ in three dimensions, etc.). The task set by Khachiyan is to find an ellipsoid whose volume is no more than a factor $(1 + \epsilon)$ times the minimum volume of an ellipsoid that encloses all the points.

We use the algorithm (see Algorithm 3.1 in Section 3 of Todd & Yıldırım 2007):

1. Input a set of m d-dimensional points (call these $a^1, \dots, a^m$) and a tolerance $\epsilon > 0$.
2. Let $k = 0$ and $n = d + 1$. Let p^0 be the m-dimensional vector $(1/m, 1/m, \dots, 1/m)$, i.e., the vector with elements $p_1^0 = 1/m$, $p_2^0 = 1/m$, $\dots$, $p_m^0 = 1/m$. Define q^i to be the n-dimensional set of vectors $q^i = ((a^i)^T, 1)^T$ for $i = 1$ to m.
3. Define $w_i(p) \equiv (q^i)^T \cdot \Lambda(p)^{-1} \cdot q^i$ for each $i = 1, \dots, m$, where $\Lambda(p)$ is the $n \times n$ matrix

$$\Lambda(p) \equiv \sum_{i=1}^m p_i q^i \left(q^i\right)^T.$$

4. If $w_i(p^0) \leqslant (1 + \epsilon)n$ for all $i = 1, \dots, m$, where

$$w_i(p) \equiv \left(q^i\right)^T \cdot \Lambda(p)^{-1} \cdot q^i,$$

then we are done. If not, iterate the following three steps until we have a p that satisfies $w_i(p) \leqslant (1 + \epsilon)n$ for all $i = 1, \dots, m$.

5. Let j be the index i that maximizes $(q^i)^T \cdot \Lambda(p^k)^{-1} \cdot q^i$, and let $\kappa \equiv (q^j)^T \cdot \Lambda(p^k)^{-1} \cdot q^j$.
6. Let $\beta \equiv \dfrac{\kappa - n}{n(\kappa - 1)}$.
7. Set $p^{k+1} = (1 - \beta)p^k$. Add β to p_j^{k+1}. Set $k \equiv k + 1$.

The output of these steps is a p^k that satisfies $w_i(p^k) \leqslant (1 + \epsilon)n$ for all i.

Once we have the desired p, we define a $d \times m$ matrix A whose ith column (recall that i runs from 1 to m) is a^i. Let P be a $m \times m$ diagonal matrix whose (i, i) component is p_i^k. Then the desired approximations to the center and defining matrix of the bounding ellipsoid are

$$c \equiv Ap^k \qquad (20)$$

and

$$Q \equiv \frac{1}{(1 + \epsilon)d} \left[A \cdot P \cdot A^T - A \cdot p^k \cdot \left(Ap^k\right)^T \right]^{-1}. \qquad (21)$$

In the $d = 2$ dimensional case that is of interest in our study here, Q and c can be used to define the axes and orientation of the bounding ellipse. In practice, we find that for some of the

a^i, $(a^i - c)^T \cdot Q \cdot (a^i - c)$ can be slightly larger than unity, so if it is critical for a given application that this product be strictly less than unity, one can replace the prefactor of Q by something like $1/[(1 + 1.1\epsilon)d]$.

REFERENCES

Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, Sci, 340, 448
Artigue, R., Barret, D., Lamb, F. K., Lo, K. H., & Miller, M. C. 2013, MNRAS, 433, L64
Ballantyne, D. R., & Everett, J. E. 2005, ApJ, 626, 364
Ballantyne, D. R., & Strohmayer, T. E. 2004, ApJL, 602, L105
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
Braje, T. M., Romani, R. W., & Rauch, K. P. 2000, ApJ, 531, 447
Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, ApJ, 654, 458
Chen, Y.-P., Zhang, S., Torres, D. F., et al. 2011, AAP, 534, A101
Degenaar, N., Miller, J. M., Wijnands, R., Altamirano, D., & Fabian, A. C. 2013, ApJL, 767, L37
Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, Natur, 467, 1081
Feroci, M., Stella, L., van der Klis, M., et al. 2012, ExA, 34, 415
Gendreau, K. C., Arzoumanian, Z., & Okajima, T. 2012, Proc. SPIE, 8443, 844313
Hartman, J. M., Patruno, A., Chakrabarty, D., et al. 2008, ApJ, 675, 1468
Hebeler, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2010, PhRvL, 105, 161102
in 't Zand, J. J. M., Galloway, D. K., & Ballantyne, D. R. 2011, AAP, 525, A111
Kachiyan, L. G. 1996, Math. Oper. Res., 21, 307
Kajava, J. J. E., Ibragimov, A., Annala, M., Patruno, A., & Poutanen, J. 2011, MNRAS, 417, 1454
Kuulkers, E., den Hartog, P. R., in 't Zand, J. J. M., Verbunt, F. W. M., Harris, W. E., & Cocchi, M. 2003, AAP, 399, 663
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., & Chou, Y. 2011, ApJ, 742, 17
Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, ApJ, 776, 19
Miller, M. C. 2013, arXiv:1312.0029
Miller, M. C., Boutloukos, S., Lo, K. H., & Lamb, F. K. 2011, in Fast X-ray Timing and Spectroscopy at Extreme Count Rates, ed. D. Barret, M. Méndez, & S. Paltani, PoS(HTRS2011)024
Miller, M. C., Boutloukos, S., Lo, K. H., & Lamb, F. K. 2013, in IAU Symp. 290, Feeding Compact Objects: Aceretion on All Scales, ed. C. M. Zhang et al. (Cambridge, UK: Cambridge Univ. Press), 101
Miller, M. C., & Lamb, F. K. 1996, ApJ, 470, 1033
Miller, M. C., & Lamb, F. K. 1998, ApJL, 499, L37
Morsink, S. M., & Leahy, D. A. 2011, ApJ, 726, 56
Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, ApJ, 663, 1244
Özel, F., & Psaltis, D. 2009, PhRvD, 80, 103003
Peille, P., Olive, J.-F., & Barret, D. 2014, A&A, 567, A80
Poutanen, J., & Beloborodov, A. M. 2006, MNRAS, 373, 836
Poutanen, J., & Gierliński, M. 2003, MNRAS, 343, 1301
Psaltis, D., Özel, F., & Chakrabarty, D. 2014, ApJ, 787, 136
Ray, P. S., Philips, B. F., Wood, K. S., et al. 2011, arXiv:1109.1309
Read, J. S., Markakis, C., Shibata, M., et al. 2009, PhRvD, 79, 124033
Romanova, M. M., Ustyugova, G. V., Koldoba, A. V., & Lovelace, R. V. E. 2004, ApJ, 610, 920
Romanova, M. M., Ustyugova, G. V., Koldoba, A. V., Wick, J. V., & Lovelace, R. V. E. 2003, ApJ, 595, 1009
Serino, M., Mihara, T., Matsuoka, M., et al. 2012, PASJ, 64, 91
Stergioulas, N. 2003, LRR, 6, 3
Strohmayer, T. E., Zhang, W., & Swank, J. H. 1997, ApJL, 487, L77
Suleimanov, V., Poutanen, J., & Werner, K. 2012, AAP, 545, A120
Todd, M. J., & Yıldırım, E. A. 2007, Discrete Applied Mathematics, 155, 1731
Viironen, K., & Poutanen, J. 2004, AAP, 426, 985
von Toussaint, U. 2011, RvMP, 83, 943
Walker, M. A. 1992, ApJ, 385, 642
Watts, A. L. 2012, ARA&A, 50, 609
Weinberg, N., Miller, M. C., & Lamb, D. Q. 2001, ApJ, 546, 1098
Wilks, S. S. 1938, Ann. Math. Stat., 9, 60
Worpel, H., Galloway, D. K., & Price, D. J. 2013, ApJ, 772, 94
Yu, W., Li, T. P., Zhang, W., & Zhang, S. N. 1999, ApJL, 512, L35