



UvA-DARE (Digital Academic Repository)

Quasi-periodic Oscillations in Short Recurring Bursts of the Soft Gamma Repeater J1550-5418

Huppenkothen, D.; D'Angelo, C.; Watts, A.L.; Heil, L.M.; van der Klis, M.; van der Horst, A.J.; Kouveliotou, C.; Baring, M.G.; Göğüş, E.; Granot, J.; Kaneko, Y.; Lin, L.; von Kienlin, A.; Younes, G.

DOI

[10.1088/0004-637X/787/2/128](https://doi.org/10.1088/0004-637X/787/2/128)

Publication date

2014

Document Version

Final published version

Published in

Astrophysical Journal

[Link to publication](#)

Citation for published version (APA):

Huppenkothen, D., D'Angelo, C., Watts, A. L., Heil, L. M., van der Klis, M., van der Horst, A. J., Kouveliotou, C., Baring, M. G., Göğüş, E., Granot, J., Kaneko, Y., Lin, L., von Kienlin, A., & Younes, G. (2014). Quasi-periodic Oscillations in Short Recurring Bursts of the Soft Gamma Repeater J1550-5418. *Astrophysical Journal*, 787(2), [128]. <https://doi.org/10.1088/0004-637X/787/2/128>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

QUASI-PERIODIC OSCILLATIONS IN SHORT RECURRING BURSTS OF THE SOFT GAMMA REPEATER J1550–5418

D. HUPPENKOTHEN¹, C. D'ANGELO¹, A. L. WATTS¹, L. HEIL¹, M. VAN DER KLIS¹, A. J. VAN DER HORST¹, C. KOUVELIOTOU^{2,3},
M. G. BARING⁴, E. GÖĞÜŞ⁵, J. GRANOT⁶, Y. KANEKO⁵, L. LIN⁷, A. VON KIENLIN⁸, AND G. YOUNES^{3,9}

¹ Astronomical Institute “Anton Pannekoek,” University of Amsterdam, Postbus 94249,
1090 GE Amsterdam, The Netherlands; D.Huppenkothen@uva.nl

² Astrophysics Office, ZP 12, NASA-Marshall Space Flight Center, Huntsville, AL 35812, USA

³ NSSTC, 320 Sparkman Drive, Huntsville, AL 35805, USA

⁴ Department of Physics and Astronomy, Rice University, MS-108, P.O. Box 1892, Houston, TX 77251, USA

⁵ Sabancı University, Orhanlı-Tuzla, İstanbul 34956, Turkey

⁶ Department of Natural Sciences, The Open University of Israel, 1 University Road, P.O. Box 808, Ra'anana 43537, Israel

⁷ François Arago Centre, APC, 10 rue Alice Domon et Léonie Duquet, F-75205 Paris, France

⁸ Max-Planck-Institut für extraterrestrische Physik, Giessenbachstrasse 1, D-85748 Garching, Germany

⁹ Universities Space Research Association, 6767 Old Madison Pike, Suite 450, Huntsville, AL 35806, USA

Received 2013 December 28; accepted 2014 April 8; published 2014 May 12

ABSTRACT

The discovery of quasi-periodic oscillations (QPOs) in magnetar giant flares has opened up prospects for neutron star asteroseismology. The scarcity of giant flares makes a search for QPOs in the shorter, far more numerous bursts from soft gamma repeaters (SGRs) desirable. In Huppenkothen et al., we developed a Bayesian method for searching for QPOs in short magnetar bursts, taking into account the effects of the complicated burst structure, and have shown its feasibility on a small sample of bursts. Here we apply the same method to a much larger sample from a burst storm of 286 bursts from SGR J1550–5418. We report a candidate signal at 260 Hz in a search of the individual bursts, which is fairly broad. We also find two QPOs at ~ 93 Hz, and one at 127 Hz, when averaging periodograms from a number of bursts in individual triggers, at frequencies close to QPOs previously observed in magnetar giant flares. Finally, for the first time, we explore the overall burst variability in the sample and report a weak anti-correlation between the power-law index of the broadband model characterizing aperiodic burst variability and the burst duration: shorter bursts have steeper power-law indices than longer bursts. This indicates that longer bursts vary over a broader range of timescales and are not simply longer versions of the short bursts.

Key words: methods: data analysis – methods: statistical – pulsars: individual (SGR J1550–5418) – stars: magnetic field – stars: neutron – X-rays: bursts

Online-only material: color figures

1. INTRODUCTION

Soft gamma repeaters (SGRs) represent a small class of neutron stars whose slow spin periods and high spin-down rates imply an unusually strong magnetic field in the excess of 10^{14} G. Duncan & Thompson (1992) and Thompson & Duncan (1995) predicted the existence of such objects, which they named magnetars. SGRs are believed to be one of two observational manifestations of neutron stars with an exceptionally strong magnetic field; anomalous X-ray pulsars (AXPs) form the other class of objects, although evidence suggests that there is no clear-cut line between them, and recently a low magnetic field source has been found (Rea et al. 2010).

The defining characteristic of SGRs are irregular bursts that vary in duration from tens to hundreds of milliseconds and span ~ 5 orders of magnitude in peak luminosity (10^{38} to 10^{43} erg s⁻¹) in hard X-rays < 100 keV. However, there is a very rare type of burst, the so-called giant flares, which have been only detected three times in the past 34 yr from three different sources. These reach peak luminosities of $\sim 10^{45}$ erg s⁻¹ and are believed to be powered by a catastrophic reordering of the magnetic field. Since this field is coupled to the solid crust, Duncan (1998) suggested that such large-scale reconfiguration might rupture the crust, triggering global seismic vibrations that would be visible as periodic modulations of the X-ray and γ -ray flux. This idea was confirmed by the detection of quasi-periodic oscillations (QPOs,

i.e., stochastic processes that vary on a characteristic timescale) in the expected range of frequencies (~ 10 –1000 Hz) in the tails of giant flares from two different magnetars (Israel et al. 2005; Strohmayer & Watts 2005, 2006; Watts & Strohmayer 2006). SGR giant flares thus present outstanding test cases for testing theories of neutron star structure and composition models. Several intermediate flares, in energy and duration between the short bursts and the giant flares, have also been observed, but no QPOs have been found in these bursts (Watts 2012).

To date, there have been few searches for QPOs in recurrent bursts of magnetars. El-Mezeini & Ibrahim (2010) reported QPO detections in a sample of bursts from SGR 1806–20 observed between 2 keV and 60 keV with the *RXTE*, however, a revised analysis incorporating variability in the burst envelope showed that the reported QPOs are not significant (Huppenkothen et al. 2013).

Finding QPOs in short SGR bursts is technically challenging: as shown in Huppenkothen et al. (2013), standard Fourier methods commonly used for this purpose fail when applied to the short, highly variable burst light curves. The major difficulty lies in the non-stationarity of magnetar bursts. The statistical distributions generally used in Fourier analysis in astronomy are strictly only valid for processes whose properties do not vary over the duration of an observation. This is clearly not true for an SGR burst: they are short events, exhibiting variability on

timescales roughly equivalent to the periods of QPOs observed in the giant flares. Below 100 Hz or so, many of the bursts exhibit a wealth of variability properties: to leading order, there is the rise and fall of the burst, i.e., a burst envelope. In most bursts, the envelope has a high degree of temporal substructure beyond this envelope. This substructure differs widely from burst to burst and is poorly understood. Perhaps what we call a burst is actually a superposition of many smaller events. Alternatively, the overall burst shape could be composed of an envelope combined with a stochastic process, leading to additional variability on shorter timescales. Finally, the complexity of the burst envelopes in general varies with energy, inserting another constraint in our interpretation of their structure. This lack of knowledge leads to two major problems when searching for QPOs. At low frequencies, very few cycles of a potential oscillation are sampled due to the short duration of the burst. A succession of peaks may look like a quasi-periodic signal to the naked eye, but could be a chance superposition of a stochastic process, without the characteristic timescale implied by a QPO. The other major difficulty is our lack of knowledge of the underlying statistical distribution that we must test against. The statistical distributions generally used in testing for QPOs are strictly defined for stationary stochastic processes. The presence of a burst envelope changes the observed distributions at low frequencies from those we know. This makes it difficult to derive inferences about the presence of a QPO at these frequencies.

In the absence of this knowledge, it is possible to make reasonable assumptions. In Huppenkothén et al. (2013), we introduced a Bayesian approach to deal with our uncertainty in the underlying burst processes by assuming a purely stochastic process with a power-law power spectrum, a so-called red noise process. While this assumption is strictly not true, either, we showed that it is a conservative choice: in practice, the presence of the burst envelope narrows the statistical distributions at low frequencies compared to the distribution we use to model the process. We are thus more likely to underestimate the significance of a signal at low frequency than overestimating our confidence in a detection. In Huppenkothén et al. (2013), we also analyzed a short bursting episode of the magnetar SGR J0501+4516, where we found one candidate detection out of 27 bursts. Our results were inconclusive with regard to the origin of this signal and showed where our method can potentially produce ambiguous results: the significant detections were all at integer multiple frequencies of a low-frequency signal at ~ 30 Hz, which was heavily affected by red noise and thus not significant itself. However, this signal corresponds to less than two full cycles at 30 Hz, given the short duration of the burst. We thus concluded that this signal was equally likely to be a chance occurrence of two red noise peaks close together as it was to be a QPO, and we deferred a more in-depth discussion to a later work with a larger sample of bursts.

In this paper, we perform a comprehensive search for QPOs in a much larger sample of bursts from a so-called burst storm observed from SGR J1550–5418 in 2009 January. SGR J1550–5418 (also 1E 1547.0–5408) was first observed with the *Einstein X-Ray Observatory* (Lamb & Markert 1981). Later observations with *XMM-Newton* revealed a soft X-ray spectrum and a possible association with a young supernova remnant, suggesting that it might be an AXP (Gelfand & Gaensler 2007). The AXP nature was confirmed by the subsequent detection of radio pulsations with a slow spin period of $P = 2.096$ s and a spin-down of $\dot{P} = 2.318 \times 10^{-14}$, implying a magnetic field of 3.2×10^{14} G (Camilo et al. 2007).

SGR J1550–5418 exhibited three major bursting episodes: in 2008 October, 2009 January, and 2009 March/April. The January episode was exceptional: the source showed hundreds of bursts within a single day, observed with several X-ray telescopes: the *Swift* Burst Alert Telescope (BAT; Israel et al. 2010; Scholz & Kaspi 2011), the *Fermi* Gamma-Ray Burst Monitor (GBM; Kaneko et al. 2010; von Kienlin et al. 2012; van der Horst et al. 2012), *RXTE* (Dib et al. 2012), and two main instruments on board the *INTEGRAL* spacecraft (Mereghetti et al. 2009; Savchenko et al. 2010).

Burst storms like the one observed from SGR J1550–5418 are rare and have been observed in only three other sources (SGR 1806–20, SGR 1900+14, and SGR 1627–41; Götz et al. 2006; Israel et al. 2008; Mazets et al. 1999), the first two of which have also exhibited a giant flare. During the first triggered observation recorded with *Fermi*/GBM on 2009 January 22, the source also showed an increase in persistent flux level up to ~ 100 keV (Kaneko et al. 2010) for around 150 s of intense bursting. A subsequent search for pulsations in this plateau of hard emission revealed a signal at the period of the neutron star, but no higher-frequency QPOs. The bursting episode ended in 2009 April, and there have been no subsequent bursts recorded since. A catalogue of magnetar bursts observed with *Fermi*/GBM is currently in preparation (A. C. Collazzi et al., in preparation).

Here we present the first large-scale robust QPO search from the 2009 January burst storm, observed with *Fermi*/GBM. In Section 2 we describe the sample in some detail, and in Section 3 we give a very brief overview of the Bayesian technique used in the QPO searches. Finally, in Section 4 we present our results and interpret both the QPO searches as well as a characterization of broadband variability in the bursts in Section 5.

2. DATA

X-ray bursts from SGR J1550–5418 triggered *Fermi*/GBM a total of 55 times between 2009 January 22 and 29, with 41 triggers on January 22 alone. Each trigger records data from 30 s before the trigger up until 300 s after the trigger. As a result, multiple (untriggered) bursts were observed in most triggers. van der Horst et al. (2012) identified a total of 286 bursts in this sample, which have time-tagged event (TTE) data available. The TTE data type has a time resolution of $2 \mu\text{s}$, needed for high-precision timing studies. We use data from the 12 NaI detectors, whose energy range of 8 keV to 4 MeV is sufficient, since SGR bursts rarely exhibit radiation above 200 keV. Additionally, we only used detectors with viewing angles to the source $< 60^\circ$ and checked whether the source was occulted by the spacecraft and the other instrument, the Large Area Detector (LAT).

The sample is the same as in van der Horst et al. (2012), and we use the burst durations, start times, and fluences from that paper in our analysis. We extracted TTE data between 8 keV and 200 keV around each burst, starting at $t_{\text{start}} - 0.1 \times T90$ (the burst duration, T90, is defined as the time in which the central 90% of the photons, starting at 5% and ending at 95%, reach the detector) and ending at $t_{\text{start}} + 1.1 \times T90$ in order to ensure that the entire burst is within our data set. The sample has a mean duration of 0.174 s and an overall asymmetric shape with a faster rise than decay. The estimated fluences range from 10^{-8} to 10^{-5} erg cm^{-2} . For a more detailed description of the data extraction process and sample definition, see van der Horst et al. (2012). An analysis of the first trigger, including a timing analysis of the inter-burst periods, was performed in Kaneko et al. (2010). Time-resolved spectroscopy and the

spectral evolution with burst flux are discussed in Younes et al. (2014). Of the bursts in the sample, 23 have saturated parts, where the detector cannot record all photons during periods of very high count rates. We excluded all 23 bursts from our analysis due to the rather complicated effect saturation has on the timing analysis. This gives us a total sample of 263 bursts.

3. ANALYSIS METHODS

QPOs are generally found by taking the Fourier transform of a light curve and looking for variability focused at a particular frequency. The square of the Fourier transform of the data is called the *periodogram*.¹⁰ Different types of variability have different frequency distributions. Our task becomes to disentangle the different components in the periodogram. While pure photon counting noise has a flat power spectrum with a well-behaved χ^2 distribution with two degrees of freedom about a constant mean, QPOs produce sharp coherent features. Stochastic processes with correlated frequencies, often termed “red noise” or “ $1/f$ noise,” are also often observed and follow power laws or broken power laws with stronger variability at low frequencies and decreasing power at higher frequencies.

The short duration of magnetar bursts means that this low-frequency variability has timescales similar to those of the QPOs observed in the giant flares. Thus, we must test for QPOs in a periodogram consisting of complicated variability. We adopt the method from Huppenkothen et al. (2013), first suggested for red-noise-dominated periodograms in Vaughan (2010). This method assumes an exponential distribution of powers about the underlying power spectrum, which we assume to be a power law or broken power law. To find QPOs, we fit a broadband noise power spectrum to each burst, which is then divided out. The highest outlier in the residuals is our candidate QPO detection. We then simulate fake periodograms using the broadband noise power spectrum and incorporating uncertainties in the model parameters, and we perform the same detection procedure on those simulations. We can thus compare the observed highest data/model outlier with a distribution of data/model outliers from the simulations, to infer the probability that the observed outlier is a significant QPO.

Below, we give a very brief overview of the QPO search strategy, and we refer the reader to Huppenkothen et al. (2013) for a detailed description, discussion of the method’s limitations, and tests on both simulated data and a smaller sample of magnetar bursts.

In more detail, for each burst:

1. We fit both a power law and a broken power law, i.e., the broadband noise model, to the periodogram of the burst observation. We fit the unnormalized posterior predictive distribution, consisting of a likelihood function following a χ^2 distribution around the broadband noise model and priors that are independent of each other and of the form $p(\theta) = 1/\theta$ (scale prior) for scale parameters (e.g., broadband noise amplitudes) and flat otherwise. As a result, we obtain the so-called maximum a posteriori (MAP) as a result of the numerical optimization step. The MAP estimate is the Bayesian equivalent of the maximum likelihood. For both models, we then construct the ratio of likelihoods at the parameter values corresponding to the MAP estimate.

2. We sample the posterior predictive distribution of the simpler broadband noise model—the power law—using a Markov chain Monte Carlo (MCMC) technique, in this case employing an affine-invariant MCMC ensemble sampler (Goodman & Weare 2010), as implemented in python by *emcee* (Foreman-Mackey et al. 2013). The resulting ensemble of parameter values will follow the posterior distribution of the assumed broadband model, thus allowing for statistical inferences over this distribution.
3. We simulate N_s artificial periodograms from the MCMC sample and fit each with the two broadband noise models considered such that we can construct the likelihood ratio for each of the fake periodograms. This will allow us to construct a distribution of likelihood ratios from a sample we know to be derived from the simpler model. If the likelihood ratio obtained for the observed periodogram is an outlier of the distribution of likelihood ratios, then the observed data are unlikely to be generated from this model. Note that this is strictly evidence *against* the power-law model; it is not direct evidence *in favor* of the more complex model. We use this approach to reject the simple power-law model for cases where the posterior predictive p -value (the ratio of samples in the posterior distribution of likelihood ratios lying above the observed values, divided by the total number of samples in this distribution) falls below 0.05. If this is true, we use the broken power law to model the broadband component of the periodogram; otherwise, the simple power-law model is adopted. A threshold of $p < 0.05$ is not very stringent, but desirable. We would rather reject the simpler broadband noise model in favor of a more complex one. It is preferable to overfit the periodogram, rather than underfit, because broadband noise features not adequately modeled by the broadband noise model may instead be mistaken for QPOs in the subsequent analysis.
4. We construct a second MCMC sample from the adopted broadband noise model in the same fashion as in step 2. We simulate N_s periodograms from this sample and fit with the adopted broadband model. For the observed periodogram and each fake periodogram, we divide the periodogram by the best-fit MAP parameter estimate and define the test statistic $T_R = \max_j (2I_j/S_j)$, where I_j is the observed power I at frequency j and S_j is the value of model S at frequency j . This test statistic is the maximum power in the residual periodogram after dividing out the broadband noise model. In the ideal case where the parameters θ defining the model S are perfectly known, $2I_j/S_j$ follows a χ^2_2 distribution. In reality, there is an uncertainty in S_j , since the parameters θ are not known exactly, leading to a deviation in the distribution of T_R from the theoretical expectation. Sampling the posterior probability distribution of the parameters given the data and the model via MCMC allows us to construct the actual distribution of T_R from simulated periodograms, taking into account all relevant uncertainties. We can thus construct a posterior distribution for the T_R statistic under the null hypothesis that the observed maximum power is due to a stochastic aperiodic process. Comparing T_R from the real data to this simulated distribution allows us to define a posterior predictive p -value for this null hypothesis. If the latter is small, then the observed maximum power is unlikely to be a product of a χ^2_2 distribution.

¹⁰ Note that throughout this paper, we use the term *periodogram* to denote the observed squared Fourier spectrum of a light curve, and we use the term *power spectrum* for the underlying (potentially stochastic) process that may have produced the observed data.

Although the giant flare QPOs were all very narrow, we cannot exclude broader signals in the shorter bursts. We search for these signals by performing exactly the same analysis as in step 4 on periodograms that are binned in frequency, where the bin widths are chosen from a logarithmic scale between 1 Hz and 200 Hz. Binning a broad QPO signal makes it easier to detect, since the QPO is grouped into a single bin, while random fluctuations from one frequency bin to the next are suppressed. We search for QPOs in the same way as in the unbinned periodogram: by defining T_R for the binned periodogram and comparing to the distribution of binned T_R values from the sample of simulated periodograms. In this case the p -values for different bins are not independent. To avoid excessive false positives, we accept significant detections only if they are detected in at least two different bin widths at the same frequency. In order to conserve computation time, we set the number of simulations $N_s = 10^4$. This implies that the significance can only be quoted to $p = 10^{-4}$ for a single trial. The detection limit we use depends on the number of trials: the more periodograms we search, the more likely it becomes to make a significant detection purely by chance, even if no signal is present. We thus require a more stringent detection limit for searching individual bursts, where we search hundreds of periodograms, than for searching averaged periodograms, where we only search 10. For searching individual bursts as in Section 4.1, we require $p < 10^{-4}$ for a single trial, corresponding to $p < 0.0263$ or roughly 2.3σ , given the number of bursts in our sample. For the 10 averaged periodograms we search in Sections 4.2 and 4.2.2, we choose 6×10^{-3} , or roughly 3.5σ . All p -values given below are trial-corrected: in the search of individual bursts, we correct by the total number of those bursts, i.e., 263, and for the averaged periodograms by the number of averaged periodograms searched, i.e., 10. The number of frequencies and bin widths we searched over is automatically taken into account by our methodology, by searching over the entire frequency and bin width range for the simulations as for the real data.

4. RESULTS

We searched light curves from 263 individual bursts for periodic signals and QPOs. In order to be sure to include the entire burst, we added 10% on either side of the burst duration (T90). Additionally, we constructed averaged periodograms from samples of bursts to explore whether a signal could be re-excited in several bursts. Finally, we characterized broadband variability for the sample as a whole, which may guide future work on emission and trigger mechanisms.

4.1. Individual Burst Searches

We searched all 263 bursts for QPOs over the complete range of available frequencies from ≤ 10 Hz to 4000 Hz. The maximum frequency was chosen to maximize computational efficiencies, while at the same time oscillatory modes are unlikely to occur at a much higher frequency.

Four bursts show detections significant with $p < 10^{-4}$ (single-trial) or $p < 0.0263$ corrected for $N_b = 263$ trials. Three candidates are significantly affected by dead time and pileup, that is, their count rate is close to the saturation count rate. This is the case when a significant number of photons arrive within less than $2.6 \mu\text{s}$ of each other (the dead time of the GBM recoding system) and are consequently recorded as a single photon. Here we used the highest intrinsic time resolution data from GBM: TTE data with $2 \mu\text{s}$ resolution. While $2.6 \mu\text{s}$

correspond to a higher frequency than we are interested in, the above effects can nevertheless influence the periodogram in nontrivial detector-dependent ways, which are not retrievable or quantifiable. A proper treatment of affected bursts is beyond the scope of this work; we thus consider the QPO search on these bursts as inconclusive and make no further statements about their properties.

The remaining burst, one of several in TTE data of trigger 090122218 with a burst duration of 0.49 s, has a significant detection of a broad feature at 260 Hz with $p < 0.0263$ (trial-corrected; also includes an uncertainty in the parameters of the broadband model). We plot the light curve and periodogram of this candidate in Figure 1. While there might be some red noise power left at these frequencies, the signal is largely dominated by white noise. We use the traditional (analytical) test against white noise for an upper limit on the detection probability (Groth 1975; van der Klis 1989). This would be a precise estimate if there was no red noise in the signal, but as we cannot exclude some contamination from red noise, this must be regarded as an upper limit instead. We find that the probability that the observed peak in the periodogram is due to Poisson counting noise alone is $p = 5.26 \times 10^{-6}$. The fractional rms amplitude is high, $\text{rms}_{\text{frac}} = 21\% \pm 3\%$, as estimated from integrating over the noise-level-subtracted periodogram. We estimated the error following Heil et al. (2012). This error calculation is somewhat too simplistic for the periodogram we consider here: there may be a residual contribution of aperiodic variability contaminating the powers we integrate over, which is not taken into account properly. However, our lack of knowledge about the burst processes involved precludes us from running simulations to establish the error to a higher degree of precision.

We measured the Q -value, defined as the centroid frequency divided by the width of the signal, by comparing the p -values for periodograms of this burst binned at several frequency resolutions and picking the frequency resolution that yielded the lowest p -value to reflect the most likely width of the signal. The QPO is extremely broad: the Q -value is $Q = \nu_0/\Delta\nu = 2.9$. This is at the lower boundary of what one would call a QPO ($Q > 2$) as opposed to a broadband noise feature. Due to the single occurrence of this signal in the sample, it is not possible to average periodograms, as is usually done to improve signal-to-noise ratio and estimate errors, such that a Lorentzian fit to the feature is possible for a precise estimate of the width and the rms amplitude (van der Klis 2006).

The periodogram of the same burst also shows a broad feature at 20 Hz. In order to understand the origin of this feature and its connection to the QPO at 260 Hz, we fit two fast-rise, slow-decay profiles to the light curve. We used the skew-normal distribution, a generalization of a simple Gaussian profile that allows for skewness and has the form

$$f(t) = \frac{2}{\sigma} \phi\left(\frac{t-\mu}{\sigma}\right) \Phi\left(\alpha\left(\frac{t-\mu}{\sigma}\right)\right), \quad (1)$$

where

$$\phi\left(\frac{t-\mu}{\sigma}\right) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right)$$

and

$$\Phi\left(\alpha\left(\frac{t-\mu}{\sigma}\right)\right) = 0.5 \left[1 + \text{erf}\left(\alpha\frac{t-\mu}{\sqrt{2}\sigma}\right) \right]. \quad (2)$$

Here μ is the location in time of the peak of the profile, σ is the width, and α is a skewness parameter (Azzalini 1985).

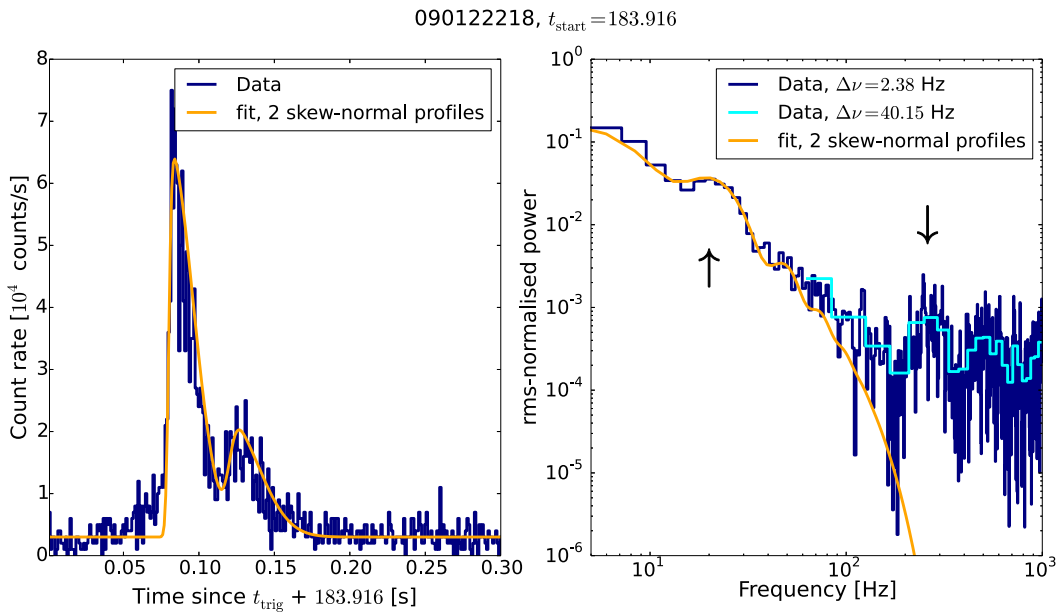


Figure 1. Light curve (left) and periodogram (right) at two different frequency resolutions for a burst in TTE data of trigger 090122218. There is a feature at ~ 20 Hz, which can be explained by the superposition of two individual peaks, modeled with the skew-normal function of Equation (1). A second feature at ~ 260 Hz is significant ($p < 0.0263$) in the binned periodogram (in cyan on the same plot), but very broad, with a Q -value $Q = \nu/\Delta\nu = 2.9$. We added arrows to guide the eye. (A color version of this figure is available in the online journal.)

We find that the signal at 20 Hz is easily reproduced by a superposition of two skewed peaks with a separation of 0.04 s and widths $\sigma_1 = 0.02$ s and $\sigma_2 = 0.016$ s. While the feature is easily reproduced by two non-periodic functions, there are too few cycles observed to make a strong statement about its nature (see Huppenkothen et al. 2013, for a similar feature). However, it cannot explain the highly significant signal at 260 Hz: the power spectrum of the two skew-normal functions fitted to the data turns over at lower frequencies and becomes negligible above 200 Hz. Beyond this frequency, there is very little power in this model, and the power spectrum at higher frequencies should be dominated by Poisson noise only. This implies that the QPO is not easily reproduced by a burst envelope and is likely a separate process producing variability at these frequencies. In order to confirm this observation, we have fit the observed light curve with both standard Gaussian profiles and Lorentzian profiles. Both alternatives give results very similar to the one presented above: a near-perfect fit to the low-frequency feature and a sharp drop in power around 200 Hz.

The sensitivity limits for signal detection vary strongly from burst to burst and with frequency, especially for the low-frequency part of the periodogram, where the contamination by broadband variability is strong. Below ~ 100 Hz, sensitivities range from $\sim 50\%$ fractional rms amplitude at 30 Hz to $\sim 10\%$ fractional rms amplitude at 100 Hz. Above ~ 150 Hz, the bursts are almost all dominated by photon detector noise, and a QPO should be the only source of non-white-noise variability in this regime. Our method converges toward standard Fourier methods in this frequency range. Instrumental effects such as dead time can still be an issue; neither method is equipped to deal with these effects without a large number of dedicated simulations. Above 150 Hz, sensitivities are generally in the range of 5%–10% fractional rms amplitude.

4.2. Averaged Periodograms

To increase sensitivity, we average the periodograms of a number of bursts. This assumes that the short bursts always

excite the same star quakes, which has also been seen in giant flares, where QPOs are detected to be present over many cycles.

4.2.1. Signal Grouped by Burst Duration

We sorted the bursts by duration (T90) into five groups: < 50 ms, 50–100 ms, 100–250 ms, 250–500 ms, and > 500 ms. To average periodograms, we picked the longest burst in each group and extracted light curves of the same duration for each burst in the sample, so that each periodogram would have the same number of frequencies. We then averaged the periodograms within a group to get the final periodogram. Since we use light curves of equal duration within each group, the shorter bursts in each group add noise into the final averaged periodogram, which reduces the QPO detection threshold somewhat. Limiting this effect is our main reason for dividing the bursts into groups, so that we can search for QPOs in the longest bursts without a strong noise component added by including the shortest bursts in the same sample.

There are no QPOs detected in the first four averaged periodograms. We report a candidate detection in the averaged periodogram of the longest bursts (T90 > 0.5 s, $N_{\text{bursts}} = 47$; see Figure 2 for the averaged periodogram). The strongest signal with $p < 2.5 \times 10^{-3}$ occurs at 10 Hz, with a width of ~ 5 Hz. Note that 10 Hz corresponds to a timescale of 0.1 s, close to the peak of the distribution of burst durations. However, we cannot exclude that this feature is actually an artifact caused by an inadequate characterization of the underlying power spectrum. Another process, such as a doubly broken power law or a combination of Lorentzians as often used in broadband noise modeling of X-ray binaries, may represent the shape of the power spectrum better, but requires more intricate model selection criteria than implemented here.

4.2.2. Averaged Periodograms per Trigger

We also search for QPOs in trigger data sets with high-resolution TTE data. Since data sets obtained with *Fermi*/GBM from an individual trigger are roughly 330 s, we searched those

Table 1
Signals from the Five Averaged Triggers, and Detection Sensitivities for Different Frequencies

Trigger ID	N_{bursts}	Min T90 (s)	Max T90 (s)	ν_0 (Hz)	$\Delta\nu$ (Hz)	Posterior p -value	Simulated p -value	Sensitivities in Fractional rms Amplitude			
								40 Hz	70 Hz	100 Hz	1000 Hz
090122037	32	0.0322	1.0724	99	27	$<4 \times 10^{-4}$	0.107	3.6	2.4	...	1.4
090122052	28	0.0364	1.4952	127	10	$<4 \times 10^{-4}$	0.016	4.8	2.5	1.9	1.5
090122194	20	0.0364	1.2124	93	12	$<4 \times 10^{-4}$	0.013	6.5	3.8	2.7	1.8
090122218	21	0.1176	1.3496	91	10	1.2×10^{-3}	0.009	5.2	3.0	2.3	1.6
090122283	30	0.0504	2.4724	61	20	8×10^{-3}	...	2.9	1.7	1.2	0.9

Notes. This table summarizes the results from the averaged periodograms of five triggers. The significant detections are shown in Figure 3. The last, 090122283, had no significant detections with $p < 6 \times 10^{-3}$, single trial probability. The second column gives the number of bursts averaged together, which equals the number of bursts in the trigger, excluding those that have saturated parts. The third and fourth columns give the minimum and maximum burst durations in the sample, respectively. Columns five and six present the centroid frequency ν_0 of the observed signal and the corresponding frequency bin width $\Delta\nu$ in which the signal is detected, respectively. We quote the detection threshold sensitivities where no detection has been made. The posterior p -value is estimated from simulations derived from the MCMC sample of the broadband model for each periodogram. The second p -value is derived from averaging random subsets of burst periodograms and extracting the highest outlier from the data-model residuals for each averaged periodogram, as described in the text.

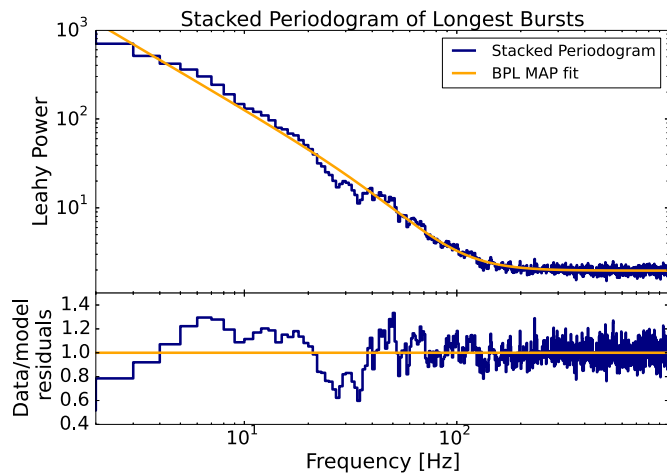


Figure 2. Top: averaged periodogram (blue) of the 47 longest bursts, and MAP fit of a broken power law to the periodogram (orange). The Leahy power is defined as $2|a_j|^2/N_{\text{ph}}$, where a_j is the Fourier amplitude at frequency ν_j and N_{ph} is the number of photons in a time series. Bottom: data/model residuals. The significant ($p < 2.5 \times 10^{-3}$) signal is at 10 Hz, with a width of ~ 5 Hz.

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

triggers with a large number of bursts (see Table 1 for an overview) in a short time span for long-lived signals. As for the duration-averaged periodograms, we extracted light curves around all bursts in those five triggers of the duration of the longest burst in that trigger data set. We then constructed the periodograms of these light curves and computed the average periodogram of the sample. The resulting periodograms do not all have the same frequency resolution; for those with a frequency resolution less than 1 Hz, we averaged neighboring frequency bins to achieve a resolution close to 1 Hz.

We searched data sets from five triggers with 20 to 32 bursts per trigger and excluded long-timescale variability below 60 Hz from the range of frequencies searched. Below 60 Hz, there will be a significant contribution from the overall shape of the short bursts (as in the 20 Hz feature discussed in Section 4.1), and thus our estimates are unreliable. We search both the unbinned periodogram and periodograms binned to different frequency resolutions between 1 Hz and 200 Hz, but considered only candidate signals with $Q > 7$. This is necessary, because at low frequencies, a candidate signal in a frequency bin that is wider than $0.2\nu_0$ likely incorporates power from the part of the periodogram below 60 Hz, where we believe estimates

to be unreliable. We find candidate detections in three of the triggers. The results are summarized in Table 1. The signals, at 93 Hz (trigger ID 090122194), 91 Hz (trigger ID 090122218), and 127 Hz (trigger ID 090122052), are fairly narrow, $Q \approx 7$ and $Q \approx 13$, respectively. Two of the signals are significant to $p < 4 \times 10^{-4}$ (roughly 3.7σ), and the third has a p -value of $p = 1.2 \times 10^{-3}$ (3.4σ). A fourth signal (trigger ID 090122037) is significant to $p < 4 \times 10^{-4}$, but fails to fulfill our criterion of $Q > 7$. At the same time, this signal is at a frequency of 99 Hz, close to the frequency where significant detections were made in two of the other triggers. We plot the periodograms for all three triggers in Figure 3.

The periodogram shape may change between different bursts, largely due to the wide spread in burst duration, fluence, and burst shape. The effects this may have on the averaged periodogram are hard to quantify without a large number of dedicated simulations of the overall burst variability, which is beyond the scope of this work. In order to test whether the observed QPOs could be due to the differences in duration, fluence, and shape in the averaged samples, we constructed a large number ($N_s = 10^3$) of averaged periodograms from randomly selected subsets of the burst sample, excluding the four triggers where candidates were observed. If the QPOs are due to effects of the varying burst properties, then these signals should appear in a large number of these simulations. We searched these periodograms in the same way we did for the averaged periodograms from individual trigger data and compared the resulting distribution of maximum powers > 60 Hz from the data-model residuals to the maximum powers from the averaged periodograms of individual triggers. Column 8 in Table 1 shows the p -values of observing the candidate signals presented above from a random subset of bursts.

In Table 1, we also show the detection sensitivities for all five averaged periodograms for various frequencies, for the ~ 1 Hz resolution of the periodogram. Note that we are even more sensitive to broader signals, as averaging neighboring frequency bins increases the signal-to-noise ratio. The numbers quoted here are upper limits for the most coherent signal we could have observed and were calculated for each quoted frequency from the distribution of maximum powers of the MCMC-derived sample of periodograms. The frequency dependence of the sensitivity is due to the fact that the low-frequency part of the periodogram is dominated by aperiodic red noise variability, and any quasi-periodic signal needs to introduce strong variations in order to be visible. Sensitivities for fractional rms amplitudes




Figure 3. Periodograms (blue, upper panels), MAP fits of a broken power law (orange), and data/model residuals (blue, lower panels) for the three triggers with candidate detections. Significant signals listed in Table 1 are indicated with black arrows.

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

are between 3% and 6% for the lowest frequencies and drop to 0.9%–1.7% at high frequencies. The differences in sensitivities between triggers are due to a combination of number of averaged bursts, number of averaged frequencies, and the average count rates of bursts included.

4.3. Broadband Variability

Magnetar bursts are a class of events with very diverse light curves: they differ vastly in duration and peak count rates, but also in overall burst shape (see Huppenkothén et al. 2013, for examples of burst light curves). How exactly this variability is produced is not well understood. Are all bursts a realization of fundamentally the same process? Are there characteristic rise or decay timescales? It is useful to characterize the variability properties of a large sample, which may answer some of these questions.

In the following, we give an overview of the broadband variability observed in the whole sample of bursts. Out of 263 burst periodograms, 193 were adequately fit with a simple power law plus a constant to account for the white-noise component; the remaining 70 rejected the null hypothesis to $p < 0.05$, and we thus adopted a broken power law for these periodograms. Burst duration and burst fluence could influence whether a simple power law or a broken power law fits the broadband noise. For example, for dim bursts, variability observed in the bright bursts may be hidden in the noise. In order to test this idea, we plot the fluence and burst duration distributions for bursts modeled with a power law and a broken power law in Figure 4.

Burst duration (T90) shows only a marginal difference in the T90 distribution ($p = 4.6 \times 10^{-4}$ for a two-sample Kolmogorov–Smirnov (K-S) test). There is an appreciable difference in fluence between the samples (two-sample K-S test: $p = 9.32 \times 10^{-11}$): bursts with broken power-law power spectra have higher fluence than bursts modeled with a simple power law. Note that the threshold for rejecting the power-law broadband model is not very high, $p < 0.05$. This is desirable for the main objective of our analysis, the search for QPOs: if the broadband noise is not adequately represented by the model, then broadband features may be attributed to QPOs instead, leading to false-positive detections. Setting a threshold of 5% is a compromise between reliability of QPO detection at the expense of potentially contaminating our sample of bursts fit with a broken power law with bursts that are consistent with a simple power law. In practice, this decreases the probability for rejecting the null hypothesis when performing the K-S test: the difference between the two distributions may be stronger than we report here.

As well as studying the overall properties of bursts with different broadband noise models, we can also study the broadband noise properties of the sample and see whether the noise properties change with fluence or burst duration. In Figure 5, we show the distribution of power-law indices for the various components. For the bursts modeled by a simple power law, the distribution of power-law indices varies between 1.5 and 4, with a median at $\mu_\gamma = 2.42$. The average low-frequency component of the two-component broken power law is flatter than for a single power law ($\mu_{\gamma_0} = 1.49$), while the higher frequency component is much steeper ($\mu_{\gamma_1} = 6.16$). Note that for several bursts, the second component is extremely steep. This may be caused by the contamination of this sample with bursts that were incorrectly classified as too complex for the simple power law. In this case, the second power-law index is often not well constrained and tends to high values. In Figure 5, we also plot the break frequency between the two components of the broken power law for those burst periodograms for which the simple power law was rejected. The distribution peaks around 100 Hz, below which the power-law index is fairly flat. At higher frequencies, it steepens considerably, as shown in panel 3 of Figure 5. The distribution is fairly broad, with the bulk of burst periodograms breaking between 30 Hz and 400 Hz.

We correlated the power-law indices with various burst properties to see whether there is a systematic effect due to burst duration or brightness, similarly to the reasoning for why some bursts require a more complex model than a simple power law. There is an indication of a correlation of the burst duration with the power-law index γ for bursts modeled with a simple power law (see Figure 6): shorter bursts seem to have slightly steeper power-law indices. A Spearman rank coefficient yields




Figure 4. Distributions of burst fluence and burst duration for the sample of bursts modeled by a simple power law (blue) and a broken power law (red). The burst duration is defined as the photon count T90. While there is only a marginal difference in T90 distributions between the two samples, there is a significant difference between the fluence distributions, $p = 9.32 \times 10^{-11}$: more complex bursts seem to have a higher fluence.

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

$R = -0.51$, indicating an anti-correlation with a probability for the null hypothesis (no correlation) of $p = 8.65 \times 10^{-16}$. In order to compute the slope of the correlation and incorporate the errors on the measurements of the power-law index, we binned the data set logarithmically into seven bins and computed the mean power-law index within each bin (following work in, e.g., X-ray binaries and active galactic nuclei, Gleissner et al. 2004). We estimated the error on the mean as a standard error, σ/\sqrt{n} , where σ is the standard deviation of the sample and n is the number of data points in each bin. The correlation can be fit with a power law, with index $\alpha = -0.1027 \pm 0.00523$.

In order to test whether the results presented in this section could be artifacts of systematic effects that we failed to take into account, we sampled randomly from the posterior distributions of the parameters of the broken power-law model from all bursts where the null hypothesis was rejected. We then sampled the distributions of T90 and fluence from the observed ensemble of bursts and created fake light curves with the power spectral properties of the broken power-law model with the randomly sampled parameter sets and combinations of burst T90 and burst fluence taken from the real sample. We simulated light curves following Timmer & Koenig (1995) and included Poisson statistics in the simulated light curves. For 1000 such fake bursts, we fit the periodograms and performed model selection between the power-law and broken power-law models. A total of 707 periodograms simulated from the broken power-law model were actually adequately fit with a simple power law instead, whereas only 293 bursts rejected the null hypothesis. This indicates that the fit is strongly dependent on burst duration and fluence. Especially for short bursts, where there are few frequencies below 100 Hz, there may not be enough data points to require a more complex model. The power-law and broken power-law samples are well separated in terms of both burst T90 and burst fluence. A two-sided K-S test reveals a separation in burst duration, $p = 1.0 \times 10^{-40}$, much stronger than observed in the data. The separation between the two samples in terms of the fluence distribution is comparable to that observed from the sample, $p = 2.2 \times 10^{-10}$. This indicates the possibility

that separating bursts by their preferred broadband model may not be meaningful: perhaps all bursts follow a broadband noise process that is closer to a broken power law than a simple power law, but for short bursts, there are not enough data points at low frequencies to confidently reject a simple power-law model.

We also tested for the presence of a correlation between burst T90 and power-law index for the sample of fake periodograms that accepted a power-law model. While there still seems to be a correlation (Spearman rank coefficient p -value for no correlation: $p = 1.35 \times 10^{-10}$), this correlation is completely flat, with a power-law slope of $\alpha = 1.77 \times 10^{-10}$, centered around a mean power-law index of $\langle \gamma \rangle = 2.56$.

5. DISCUSSION

We have searched both individual bursts and averaged periodograms from samples of bursts for QPOs. Our analysis is the most precise to date for fast transients while taking into account the effects of aperiodic variability, but it is also conservative: at low frequencies, real quasi-periodic features may be missed because we assume that the burst is a purely stochastic process, when it is not. Additionally, we do not model several effects that can significantly affect the outcome of a QPO search. Dead time can significantly affect especially the bright bursts and thus render inferences invalid even at high frequencies. At low frequencies, we have demonstrated that some of the power concentrated in broad bumps, which can be classified as QPOs with broad widths by the detection algorithm, can easily be modeled with a simple estimate of a burst envelope consisting of two functions with a faster rise than decay. We must hence be very careful when interpreting signals at frequencies comparable to the lowest QPO frequencies seen in the giant flares (18 Hz and 36 Hz).

5.1. Individual Bursts

We find no indication in individual bursts for QPOs at the frequencies and coherences seen in the giant flares. We detected one significant signal at 260 Hz, in the same burst where we

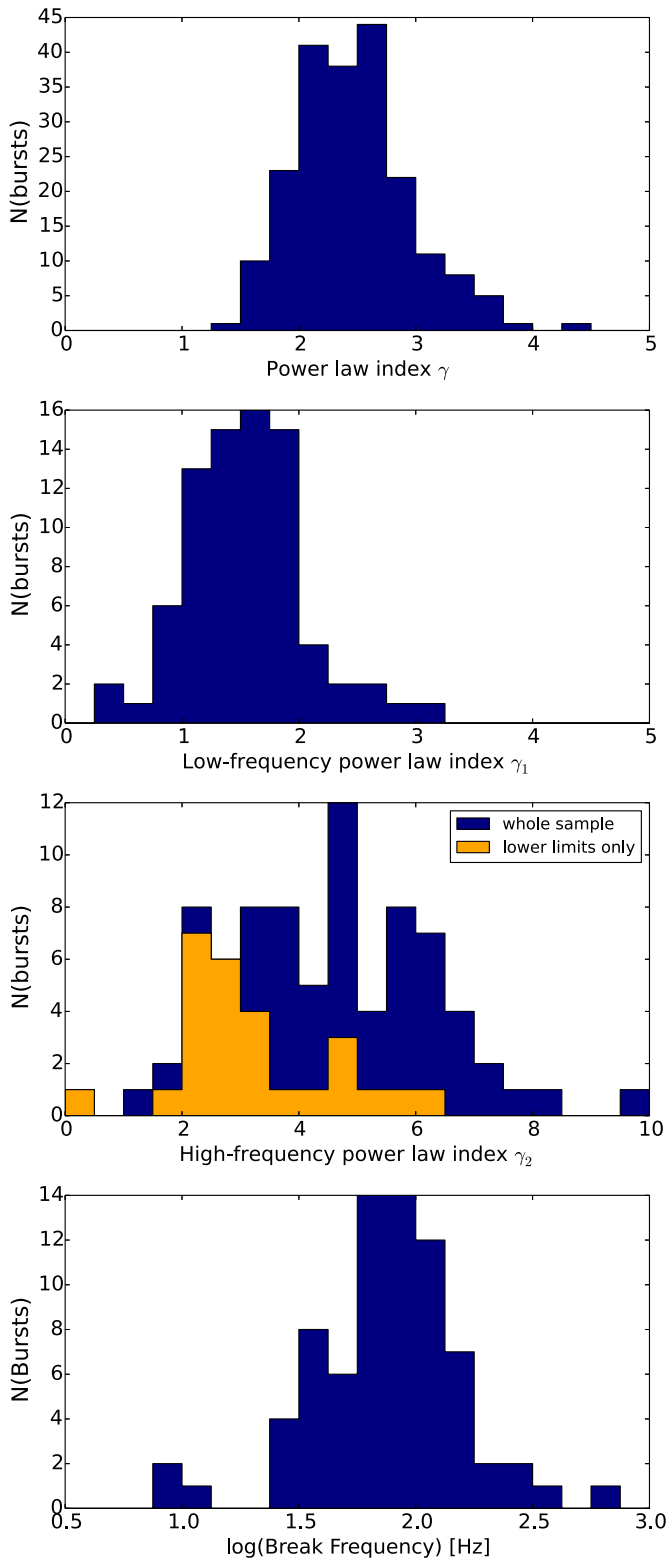


Figure 5. Distributions for the power-law index for bursts modeled with a simple power law (top), the low-frequency power-law index for bursts modeled with the broken power law (second from top), the high-frequency power-law index for the latter model (third from top), and the break frequency between the low-frequency and high-frequency power law indices (bottom). Note that the high-frequency power-law indices plotted here are the means of their posterior distributions, which, for some of the bursts, have a high variance. For the bursts where the high-frequency power-law index is largely unconstrained, we include the 5% quantile in the sample instead and include the fraction of these bursts in panel 3 in orange.

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




Figure 6. Burst duration (T_{90}) vs. power-law index for all bursts modeled with a simple power law. Errors on the power-law index are estimated from the marginalized posterior distribution approximated with the MCMC sample. We show the data binned logarithmically and averaged within each bin in orange. We tested for correlation using a Spearman rank coefficient, $R = -0.51$. The dashed, light blue line indicates a power-law fit ($\alpha = -0.1$) to the binned data points, as explained in the text.

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

found the broad feature at 20 Hz that the algorithm flagged as significant. The latter is just as likely to be caused by a superposition of two burst envelope profiles as it is to be a QPO. With less than two full cycles, it is impossible to tell both models apart. This is reminiscent of the candidate detection reported in Huppenkothén et al. (2013), where we noted that this candidate could be due to a chance occurrence of two such peaks close together. With a Q -value of 2.9, the signal at 260 Hz is far broader than anything seen in the giant flare QPOs ($Q > 10$), but very strong, with a fractional rms amplitude of $21\% \pm 3\%$. The burst is longer than average, $T_{90} = 0.48$ s, with a fluence at the lower end of the sample, $F = 3.94 \times 10^{-7}$ erg cm^{-2} . The detected QPO at 260 Hz is not present in any other burst in the entire sample, nor is it seen in the averaged periodogram of all bursts in this trigger, as described in Section 4.2.2.

5.2. Averaged Bursts

We find a candidate detection in the averaged periodogram of the longest bursts with durations $T_{90} > 0.5$ s. These bursts are highly structured and generally have multiple peaks. The detected signal at 10 Hz is quite broad and matches the position of the maximum in the distribution of burst durations in van der Horst et al. (2012). This suggests a characteristic timescale for individual peaks within highly structured bursts, rather than a crustal mode. This in turn raises the question of whether these many-peaked bursts are causally connected single events, or instead individual bursts that happen to appear close to each other. While our results favor the latter explanation, it has been argued that these should be causally connected events. One argument is based on the distribution of waiting times between bursts: while the waiting times between bursts generally follow a lognormal distribution, an excess of short waiting times has been observed when regarding each peak as an individual burst, rather than grouping these peaks into causally connected events (Gögüş et al. 1999, 2000). However, at this point, we cannot exclude that this candidate detection arose from an inadequate broadband noise model fit to the data.

A more complex periodogram shape, with additional power-law components, could explain the observed excess power. To check this would require a somewhat more complex setup of the model selection procedure, or should ideally be done with methods better suited to model broadband variability. We thus defer this task to a future work.

The most interesting results stem from averaging periodograms in individual triggers. This allows us to probe timescales on which QPOs were observed in the giant flares (tens to hundreds of seconds). We find signals in two of the averaged periodograms of five triggers, at a frequency of 93 Hz, very close to the strongest QPO reported in the 2004 giant flare at 92 Hz (Israel et al. 2005). In the periodogram of bursts from a third trigger, we find a significant detection at a slightly higher frequency, $\nu_0 = 127$ Hz, where no giant flare QPO has been observed.

When deriving inferences from an averaged periodogram, one makes the implicit assumption that all periodograms used to construct the average are realizations of the same underlying process. This need not be true for SGR bursts: we are averaging periodograms from bursts with vastly different durations and morphologies. While most bursts are described well by simple power laws or broken power laws in the Fourier domain, the parameters of these broadband noise models vary from burst to burst, as does the range of frequencies over which variability is observed (as seen, e.g., by the correlation between burst duration and power-law index reported in Section 4.3). Additionally, the results reported in Section 4.1 show that at low frequencies the burst envelope may dominate the power spectrum, which significantly alters both the shape and statistical distribution of the periodogram. In order to test this effect, we created averaged periodograms from randomly selected bursts of the sample, excluding the triggers where we detected a significant signal in the averaged periodograms. If we see many significant signals of the observed strength in these averaged periodograms from randomly selected bursts, then either the QPO we are interested in is re-excited in many of the bursts, or there are effects due to averaging vastly different bursts that our broadband noise model cannot take into account properly. The former case is unlikely: if a QPO at 93 Hz were present in many bursts, we would have likely observed it when averaging by duration, as in Section 4.2. All three signals observed at ~ 93 Hz and 127 Hz are fairly narrow and at comparatively high frequency. At ~ 90 Hz, the contribution by the burst envelope should be small (see Figure 1) and not cluster around a single frequency, but rather follow a power law. While the chance to observe one such signal is still too high to claim a strong detection ($p \approx 0.01$ for all three narrow candidates), the fact that it is observed twice out of five trials, at a frequency close to that observed in the 2004 giant flare ($\nu = 92$ Hz), strengthens the claim for a detection.

We note that even for those frequencies for which we do not detect any signal, we can quote stringent sensitivities that set quite tight upper limits on a signal that could have been there and go undetected by our algorithm. In the white-noise regime above ~ 150 Hz or so, where our algorithm approaches classical Fourier methods, the variability is lowest, and thus we have the highest sensitivity to weak signals. We can confidently exclude high-frequency QPOs at 625 Hz and 1840 Hz that have been observed in the giant flares (Strohmayer & Watts 2006; Watts & Strohmayer 2006). The QPO at 625 Hz was observed over a large fraction of the tail of the giant flare (>150 s) in two different energy ranges, with a high fractional rms amplitude of 8.5%. Conversely, the QPO at 1840 Hz is seen only in two

cycles, but with high significance and a large fractional rms amplitude of 18%. Since our sensitivities for all five averaged periodograms are much lower at these frequencies ($<1.7\%$), we can exclude a QPO of this type in the smaller flares from the burst storm to a high degree of confidence.

5.3. Aperiodic Variability

While the aperiodic variability in short magnetar bursts is a hindrance when searching for QPOs, it is interesting in its own right. Each burst has a unique temporal structure, which can sometimes be quite complex. Nevertheless, most burst periodograms can be modeled fairly well with simple empirical models, which allows us to draw a number of general conclusions and give indications where further work is required. The simplest question one can ask about is the separation between bursts modeled with simple power laws and those that require a more complex model, in our case a broken power law. Our results indicate that the differences between the two samples are largely due to systematic effects: for short bursts, fewer data points in the low-frequency part of the periodogram make it more difficult to constrain the shape of the power spectrum, and these bursts are thus more likely to accept a simple power law as a model for the underlying power spectrum. Similarly, a burst with a lower fluence will be more strongly affected by photon detector noise, rendering inferences about the shape of the power spectrum more difficult. This idea is strengthened by the fact that the fractions of bursts fit by a broken power law in both the observed sample and the sample of fake periodograms simulated from the broken power law are very close: 27% of observed bursts are fit by a broken power law, versus 28% bursts in the simulated sample. There is one striking discrepancy between the data and the simulations: in the simulated ensemble, the samples fit by the different models are very strongly separated in burst duration, whereas there is only mild evidence for this separation in the observed ensemble. One reason may lie in the lower number of bursts in the observed sample. Another reason may lie in the nature of our simulations: we simulated light curves from the posterior distributions of broken power-law parameters inferred for the real bursts following the method of Timmer & Koenig (1995). This method provides pure red-noise light curves, which are only an approximation to the real data, as discussed in Huppenkothén et al. (2013). It is possible that the separation in the simulated samples in burst duration is due to effects that are not adequately captured by this model. Without better knowledge of the underlying processes involved, however, we cannot construct a model more representative of the observed data, in order to increase confidence in our inferences.

We also extracted the distributions of broadband model parameters from the means of the MCMC samples of each individual burst. Note that the mean is a very simple estimate of the posterior probability distribution of the parameter. It can encode neither a skewness in the distribution nor correlations between parameters. The only way to encode the full information of both marginalized distributions of parameters and potential correlations between parameters is to plot the full posterior probability distribution, which, for three or more parameters, is impossible. For the purpose of this study, we accept the simple estimate and its limitations and refer a more nuanced analysis to future work. In general, the power-law index for bursts modeled by a simple power law is confined between 1.5 and 4. While the distribution is fairly broad, it peaks around $\mu_\gamma = 2.42$. This

is much higher than seen, for example, in gamma-ray bursts, where a similar analysis yields indices of 1.7 to 2.0 (Guidorzi et al. 2012; Beloborodov et al. 2000).

The search for correlations reveals two interesting observations. There is no correlation between power-law index and fluence for the bursts that can be modeled with a simple power law. This is not surprising: the highest-fluence bursts preferentially reject the simpler model and are consequently in the sample of bursts modeled with a broken power law. Secondly, we find an anti-correlation of power-law indices with burst duration: shorter bursts have steeper power laws. The anti-correlation can be modeled with a simple power law with $\alpha = -0.1$. The observed correlation is unlikely due to systematic effects: simulated periodograms indicate a distribution of power-law indices centered around $\langle \gamma \rangle = 2.56$, which does not change with burst duration or fluence. This observed anti-correlation would imply that magnetar bursts are not self-similar in burst duration: shorter bursts are not simply shorter copies of longer bursts. In the latter case, the power-law index would be the same, but shifted to higher frequencies. Shorter bursts are also not simply shorter snapshots of the same process. Instead, it implies that the relevant variability timescales in each burst depend on the duration of the burst. Longer bursts are variable over a larger range of timescales and variable at higher frequencies. Exactly how this difference between short and long bursts manifests, and what implications it might have for magnetar burst emission models, is unclear. Again, more nuanced methods are required to better understand the broadband variability in magnetar bursts. Understanding this variability, in turn, is valuable for performing QPO searches with more precision than is currently possible.

6. THEORETICAL DISCUSSION

The QPOs in the tail of the giant flare from SGR 1900+14 lie in the range 28–155 Hz (Strohmayer & Watts 2005). For the tail of the giant flare from SGR 1806–20, there are several QPOs in the range 18–150 Hz, and two isolated higher frequency signals at 625 Hz and 1840 Hz (Israel et al. 2005; Watts & Strohmayer 2006; Strohmayer & Watts 2006). Widths (FWHM) are in the range 1–20 Hz.

The most plausible explanation advanced for the giant flare QPOs is that they represent global seismic oscillations of the star, and it was swiftly realized that this would be a novel means of constraining not only the interior field strength (which is hard to measure directly) but also the dense matter equation of state (Samuelsson & Andersson 2007; Watts & Reddy 2007). The question of mode identification is therefore crucial.

In the original discovery papers, the QPOs were tentatively identified with torsional shear modes of the neutron star crust and torsional Alfvén modes of the highly magnetized fluid core. These identifications were based on the expected mode frequencies, which are set by both the size of the resonant volume and the relevant wave speed. For crustal shear modes, the appropriate speed is the shear speed $v_s = (\mu_s/\rho)^{1/2}$, where μ_s is the shear modulus and ρ the density. The shear modulus is of the order of the Coulomb potential energy $\sim Z^2 e^2/r$ per unit volume r^3 , where $r \sim (\rho/Am_p)^{-1/3}$ is the inter-ion spacing, while Z and A are the effective atomic number and mass number, respectively, of the ions in the crust. Using the shear modulus computed by Strohmayer et al. (1991) and scaling by typical values for the inner crust (Douchin & Haensel 2001), the shear

velocity as shown by Piro (2005) is

$$v_s = 1.1 \times 10^8 \text{ cm s}^{-1} \left(\frac{\rho}{10^{14} \text{ g cm}^{-3}} \right)^{1/6} \left(\frac{Z}{38} \right) \times \left(\frac{302}{A} \right)^{2/3} \left(\frac{1 - X_n}{0.25} \right)^{2/3}, \quad (3)$$

where X_n is the fraction of neutrons. This yields a rough estimate for the frequency for the fundamental crustal shear mode of $\nu \sim v_s/2\pi R = 18$ (10 km/R) Hz. Full-mode calculations find similar values, but with additional dependencies on the mass and radius of the star due to relativistic effects (see, e.g., Samuelsson & Andersson (2007)); it is this dependence that makes the modes potentially powerful diagnostics of the dense matter equation of state (Lattimer & Prakash 2007). Many of the lower QPO frequencies could be explained as angular harmonics with no radial nodes, while the two highest frequencies in the SGR 1806–20 giant flare were identified as radial overtones of these crustal modes.

For torsional Alfvén modes of the core, the appropriate wave speed is the Alfvén speed $v_A = B/\sqrt{4\pi\rho}$, where B is the magnetic field strength, giving

$$v_A = 10^8 \text{ cm s}^{-1} \left(\frac{B}{10^{16} \text{ G}} \right) \left(\frac{10^{15} \text{ g cm}^{-3}}{\rho} \right)^{1/2}. \quad (4)$$

This yields a very rough estimate for the frequency of the fundamental torsional Alfvén mode of $\nu \sim v_A/4R = 25$ (10 km/R) Hz (Thompson & Duncan 2001). Note, however, that the value of the field strength B in magnetar cores is highly uncertain, as is the appropriate value of the density ρ . In principle, only the charged component ($\sim 5\%$ – 10% of the core mass) should participate in Alfvén oscillations, reducing ρ ; however, there are mechanisms associated with superfluidity and superconductivity that can couple the charged and neutral components, leading to additional mass loading. As above, full-mode calculations that take into account relativistic effects lead to additional dependencies on neutron star mass and radius (see, e.g., Sotani et al. 2008). It should also be noted that the Alfvén modes constitute continua rather than a set of discrete frequencies, since the field lines within the core have a continuum of lengths. It has been suggested that the observed QPOs might be associated with a turning point in the continuum of Alfvén modes (Levin 2007; Sotani et al. 2008).

In fact, for a star with a magnetar strength field, crustal vibrations and core vibrations should couple together on very short timescales (Levin 2006, 2007). Considering them in isolation, as described above, is therefore not appropriate. The current viewpoint, based on more detailed modeling that takes into account the magnetic coupling between crust and core, is that the QPOs are in fact associated with global magneto-elastic axial (torsional) oscillations of the star (Glampedakis et al. 2006; Andersson et al. 2009; Steiner & Watts 2009; van Hoven & Levin 2011, 2012; Colaiuda & Kokkotas 2011, 2012; Gabler et al. 2012, 2013; Passamonti & Lander 2014, 2013; Glampedakis & Jones 2014).

Since magneto-elastic oscillations depend on the same physics described above, albeit now in a coupled system, they have frequencies in the same broad range as the simple estimates given above. The frequencies are set by many factors, including the dense matter equation of state (which sets mass and radius), field strength and geometry, superfluidity, superconductivity,

and crust composition. Current magneto-elastic torsional oscillation models have had some success in explaining the presence of oscillations at frequencies of 155 Hz and below. However, they struggle to explain the presence of the highest frequency oscillations, which should damp very rapidly (van Hoven & Levin 2012; Gabler et al. 2012). Various solutions to this problem are under investigation, including coupling to polar modes (Lander et al. 2010; Lander & Jones 2011; Colaiuda & Kokkotas 2012), and resonances between crust and core that might develop as a result of superfluid effects (Gabler et al. 2013; Passamonti & Lander 2014). However, until these issues are resolved, precise identification of the giant flare frequencies with specific global magneto-elastic modes remains a challenge.

The detection of frequencies in the smaller flares provides an entirely new viewpoint on this very difficult theoretical problem. It is therefore important to compare the properties of the giant flare QPOs to those detected in this study. We begin with the frequencies detected by averaging together multiple bursts from highly active episodes, at 93 Hz and 127 Hz. These frequencies are in the range found in the giant flares (indeed the strongest frequency found in the SGR 1806–20 giant flare was at 92 Hz). The widths are also comparable to the range observed in the tails of the giant flares. It therefore seems plausible that they are instances of the same phenomenon. If these frequencies do indeed represent global magneto-elastic oscillations, the implication is that such vibrations are excited not only by giant flares but also by trains of shorter bursts. This is important information for future theoretical studies of mode excitation.

The 260 Hz signal identified in one of the individual bursts is rather different. It is found in a frequency range where no signals were found in the giant flares. It is much broader than any of the oscillations seen in the tails of the giant flares and has very high fractional amplitude. Whether it is the same phenomenon as was observed in the giant flares is therefore far from clear. If it is the same phenomenon, and we are seeing a magneto-elastic oscillation mode, then a detection in this frequency range would be valuable. Magneto-elastic oscillation models, as outlined above, have difficulty in explaining the lifetime of higher frequency signals in the giant flares. This detection, with a much lower coherence, in a burst whose duration is comparable to the predicted lifetimes, provides a fresh perspective on this problem.

We may, however, be seeing something quite different. The giant flares consist of a short impulsive spike, followed by a long decaying tail as a trapped pair-plasma fireball slowly evaporates. In the smaller bursts, it is not clear whether a fireball forms: what we see may be more analogous to the impulsive spike of the giant flares. The variability that we have found in the short bursts in SGR 1550 (particularly the 260 Hz signal that appears to differ in properties) may instead be associated directly with the burst trigger process, be that a magnetospheric instability or the yielding of the crust. It is worth noting that there are tentative claims of variability at 43 Hz in the first 200 ms of the 1979 giant flare from SGR 0526–66 (Barat et al. 1983) and at 50 Hz in the first 500 ms of the SGR 1806–20 giant flare (Terasawa et al. 2006; Geotail paper). However, timing analysis of the peaks of the giant flares is heavily affected by dead time and saturation. In this respect the smaller flares, which are typically not saturated, may be more useful despite the lower count rates. However, the variability that would be expected in the initial trigger and yielding phase of a magnetar burst has not been studied in detail. There are nonetheless plausible mechanisms that might lead to quasi-periodic behavior.

If the burst trigger is magnetospheric, there may arise via variability may rise via plasma instabilities associated with magnetic reconnection (see, e.g., Kliem et al. 2000). If instead the trigger is crustal yielding, local effects and resonances may be significant. It is not clear, for example, whether locally excited shear waves would immediately couple to the entire crust (and from there to the core) rather than resonating, with low Q -value, in a smaller cavity that is temporarily coupled very poorly to the rest of the crust during the yielding process. Such local resonances would have frequencies quite different from those of global magneto-elastic oscillations. More detailed theoretical studies of the trigger process will be required to resolve both this question and the length of time required to establish global modes of any kind. However, the possibility that the 260 Hz signal is a new and direct signature of the trigger process is an exciting one.

7. CONCLUSIONS

We have searched 263 bursts from SGR J1550–5418 for QPOs. We find one candidate QPO in the individual burst searches. The signal is broad, but highly significant, and not close to any frequency observed in the giant flares. It is unclear whether this signal could come from the same phenomenon as the QPOs observed in the giant flares, or whether it may be associated directly with the burst trigger process.

Searching averaged periodograms reveals a significant signal at ~ 10 Hz in an averaged periodogram of all bursts with durations $T_{90} > 0.5$. This signal is comparable to the characteristic duration of a magnetar burst $T_{90,\text{max}} = 0.1$ s, but may be due to an inadequate broadband model for the periodogram. We find evidence for QPOs in periodograms averaging bursts from individual triggers, which are unlikely due to effects of averaging together bursts with vastly different timing properties. Two of these signals are located at 93 Hz, where QPOs in the giant flares have been observed. The third is at a higher frequency, 127 Hz. We consider these signals to be the best candidates of neutron star oscillations from short magnetar bursts to date. All three signals can be interpreted in the framework of magneto-elastic oscillations invoked to explain QPOs in magnetar giant flares. The possibility that not only giant flares but also smaller flares may excite these oscillations also provides an important new piece of information for future theoretical studies of mode excitation. For averaged periodograms, we can put constraints on weak signals that could have been there and would likely have been missed by our methods. For all but the lowest frequencies, our sensitivity to QPOs is lower than the observed fractional rms amplitudes in the giant flares. This is especially prominent for the high-frequency QPOs observed at 625 Hz and higher. We thus conclude that except for the signals at 93 Hz and 127 Hz, there are no giant-flare-like QPOs in this sample of small bursts.

Here we also characterized overall burst variability for the first time. We find a correlation between power-law index and burst duration. This implies that longer SGR bursts are variable over a broader range of timescales than short bursts and are not simply longer versions of the short bursts. Further work is required to disentangle overall variability in magnetar bursts. This is unlikely to be possible with Fourier methods, but would be very rewarding in terms of both understanding emission mechanisms and untangling possible QPO signals from the overall burst morphology.

The authors thank Phil Uttley and Yuri Levin for useful discussions and the referee for helpful suggestions. D.H.,

C.D., and A.L.W. acknowledge support from a Netherlands Organization for Scientific Research (NWO) Vidi Fellowship (PI: A. Watts). C.K. was partially supported by NASA grant NNH07ZDA001-GLAST. This publication is part of the GBM/Magnetar Key Project (NASA grant NNH07ZDA001-GLAST; PI: C. Kouveliotou). A.J.v.d.H. acknowledges support from the European Research Council via Advanced Investigator Grant No. 247295 (PI: R.A.M.J. Wijers).

REFERENCES

- Andersson, N., Glampedakis, K., & Samuelsson, L. 2009, *MNRAS*, **396**, 894
- Azzalini, A. 1985, *Scand. J. Stat.*, **12**, 171
- Barat, C., Hayles, R. I., Hurley, K., et al. 1983, *A&A*, **126**, 400
- Beloborodov, A. M., Stern, B. E., & Svensson, R. 2000, *ApJ*, **535**, 158
- Camilo, F., Ransom, S. M., Halpern, J. P., & Reynolds, J. 2007, *ApJL*, **666**, L93
- Colaiuda, A., & Kokkotas, K. D. 2011, *MNRAS*, **414**, 3014
- Colaiuda, A., & Kokkotas, K. D. 2012, *MNRAS*, **423**, 811
- Dib, R., Kaspi, V. M., Scholz, P., & Gavriil, F. P. 2012, *ApJ*, **748**, 3
- Douchin, F., & Haensel, P. 2001, *A&A*, **380**, 151
- Duncan, R. C. 1998, *ApJL*, **498**, L45
- Duncan, R. C., & Thompson, C. 1992, *ApJL*, **392**, L9
- El-Mezeini, A. M., & Ibrahim, A. I. 2010, *ApJL*, **721**, L121
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Gabler, M., Cerdá-Durán, P., Font, J. A., Müller, E., & Stergioulas, N. 2013, *MNRAS*, **430**, 1811
- Gabler, M., Cerdá-Durán, P., Stergioulas, N., Font, J. A., & Müller, E. 2012, *MNRAS*, **421**, 2054
- Gelfand, J. D., & Gaensler, B. M. 2007, *ApJ*, **667**, 1111
- Glampedakis, K., & Jones, D. I. 2014, *MNRAS*, **439**, 1522
- Glampedakis, K., Samuelsson, L., & Andersson, N. 2006, *MNRAS*, **371**, L74
- Gleissner, T., Wilms, J., Pottschmidt, K., et al. 2004, *A&A*, **414**, 1091
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 1999, *ApJL*, **526**, L93
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 2000, *ApJL*, **532**, L121
- Goodman, J., & Weare, J. 2010, *Commun. Appl. Math. Comput. Sci.*, **5**, 65
- Götz, D., Mereghetti, S., Molkov, S., et al. 2006, *A&A*, **445**, 313
- Groth, E. J. 1975, *ApJS*, **29**, 285
- Guidorzi, C., Margutti, R., Amati, L., et al. 2012, *MNRAS*, **422**, 1785
- Heil, L. M., Vaughan, S., & Uttley, P. 2012, *MNRAS*, **422**, 2620
- Huppenkothen, D., Watts, A. L., Uttley, P., et al. 2013, *ApJ*, **768**, 87
- Israel, G. L., Belloni, T., Stella, L., et al. 2005, *ApJL*, **628**, L53
- Israel, G. L., Esposito, P., Rea, N., et al. 2010, *MNRAS*, **408**, 1387
- Israel, G. L., Romano, P., Mangano, V., et al. 2008, *ApJ*, **685**, 1114
- Kaneko, Y., Gögüş, E., Kouveliotou, C., et al. 2010, *ApJ*, **710**, 1335
- Kliem, B., Karlicky, M., & Benz, A. O. 2000, *A&A*, **360**, 715
- Lamb, R. C., & Markert, T. H. 1981, *ApJ*, **244**, 94
- Lander, S. K., & Jones, D. I. 2011, *MNRAS*, **412**, 1730
- Lander, S. K., Jones, D. I., & Passamonti, A. 2010, *MNRAS*, **405**, 318
- Lattimer, J. M., & Prakash, M. 2007, *PhR*, **442**, 109
- Levin, Y. 2006, *MNRAS*, **368**, L35
- Levin, Y. 2007, *MNRAS*, **377**, 159
- Mazets, E. P., Aptekar, R. L., Butterworth, P. S., et al. 1999, *ApJL*, **519**, L151
- Mereghetti, S., Götz, D., Weidenspointner, G., et al. 2009, *ApJL*, **696**, L74
- Passamonti, A., & Lander, S. K. 2013, *MNRAS*, **429**, 767
- Passamonti, A., & Lander, S. K. 2014, *MNRAS*, **438**, 156
- Piro, A. L. 2005, *ApJL*, **634**, L153
- Rea, N., Esposito, P., Turolla, R., et al. 2010, *Sci*, **330**, 944
- Samuelsson, L., & Andersson, N. 2007, *MNRAS*, **374**, 256
- Savchenko, V., Neronov, A., Beckmann, V., Produit, N., & Walter, R. 2010, *A&A*, **510**, A77
- Scholz, P., & Kaspi, V. M. 2011, *ApJ*, **739**, 94
- Sotani, H., Kokkotas, K. D., & Stergioulas, N. 2008, *MNRAS*, **385**, L5
- Steiner, A. W., & Watts, A. L. 2009, *PhRvL*, **103**, 181101
- Strohmayer, T., van Horn, H. M., Ogata, S., Iyetomi, H., & Ichimaru, S. 1991, *ApJ*, **375**, 679
- Strohmayer, T. E., & Watts, A. L. 2005, *ApJL*, **632**, L111
- Strohmayer, T. E., & Watts, A. L. 2006, *ApJ*, **653**, 593
- Terasawa, T., Tanaka, Y. T., Yoshikawa, I., & Kawai, N. 2006, *JPhCS*, **31**, 76
- Thompson, C., & Duncan, R. C. 1995, *MNRAS*, **275**, 255
- Thompson, C., & Duncan, R. C. 2001, *ApJ*, **561**, 980
- Timmer, J., & Koenig, M. 1995, *A&A*, **300**, 707
- van der Horst, A. J., Kouveliotou, C., Gorgone, N. M., et al. 2012, *ApJ*, **749**, 122
- van der Klis, M. 1989, in *Timing Neutron Stars*, ed. H. Ögelman & E. P. J. van den Heuvel (New York, NY: Kluwer), 27
- van der Klis, M. 2006, in *Compact Stellar X-ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 39
- van Hoven, M., & Levin, Y. 2011, *MNRAS*, **410**, 1036
- van Hoven, M., & Levin, Y. 2012, *MNRAS*, **420**, 3035
- Vaughan, S. 2010, *MNRAS*, **402**, 307
- von Kienlin, A., Gruber, D., Kouveliotou, C., et al. 2012, *ApJ*, **755**, 150
- Watts, A. L. 2012, in *Neutron Star Crust*, ed. C. Bertulani & J. Piekarewicz (Hauppauge, NY: Nova Science Publishers), 265
- Watts, A. L., & Reddy, S. 2007, *MNRAS*, **379**, L63
- Watts, A. L., & Strohmayer, T. E. 2006, *ApJL*, **637**, L117
- Younes, G., Kouveliotou, C., van der Horst, A., et al. 2014, *ApJ*, **785**, 52