






On the Persistence of QPOs during the SGR 1806–20 Giant Flare

M. Coleman Miller¹ , Cecilia Chirenti² , and Tod E. Strohmayer³ 

¹Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421, USA

²Centro de Matemática, Computação e Cognição, UFABC, 09210-170 Santo André-SP, Brazil

³Astrophysics Science Division and Joint Space-Science Institute, NASA's Goddard Space Flight Center, Greenbelt, MD 20771, USA

Received 2018 August 28; revised 2018 November 21; accepted 2018 December 2; published 2019 January 24

Abstract

The discovery of quasi-periodic brightness oscillations (QPOs) in the X-ray emission accompanying the giant flares of the soft gamma-ray repeaters SGR 1806–20 and SGR 1900+14 has led to intense speculation about their nature and what they might reveal about the interiors of neutron stars. Here we take a fresh look at the giant flare data for SGR 1806–20, and in particular we analyze short segments of the post-peak emission using a Bayesian procedure, which has not previously been applied to these data. We find at best weak evidence that any QPO persists for more than ~ 1 s; instead, almost all the data are consistent with a picture in which there are numerous independently excited modes that decay within a few tenths of a second. This has interesting implications for the rapidity of decay of the QPO modes, which could occur by the previously suggested mechanism of coupling to the MHD continuum. The strongest QPOs favor certain rotational phases, which might suggest special regions of the crust or of the magnetosphere. We also find several previously unreported QPOs in these data, which may help in tracking down their origin.

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

1. Introduction

Quasi-periodic X-ray brightness oscillations in the 2004 December giant flare from SGR 1806–20 were first reported by Israel et al. (2005). They detected a strong QPO at 92.5 Hz in *Rossi X-Ray Timing Explorer* (*RXTE*) data and found evidence for lower-frequency QPOs at about 18 and 30 Hz. They suggested the QPOs could be related to torsional modes excited during the flare. Motivated by this, Strohmayer & Watts (2005) investigated the *RXTE* data obtained during the 1998 August event from SGR 1900+14. They found a sequence of rotation-phase-dependent QPOs at 28, 54, 84, and 155 Hz, and suggested an identification with a sequence of low-order torsional modes with different l values. These authors also re-examined the *RXTE* data from the SGR 1806–20 giant flare, and found evidence for additional oscillations at 150 and 625 Hz (Strohmayer & Watts 2006). They suggested that the 625 Hz QPO could be identified with torsional oscillation modes with at least one radial node in the crust, and that this could lead to a probe of the crust thickness.

These observational findings touched off a plethora of theoretical investigations. Levin (2006) noted that the coupling of the crust to the core by a strong vertical magnetic field should lead to damping of the crustal oscillations within a few tenths of a second. Subsequent investigations explored the magnetic coupling of the crust to the core in more realistic scenarios (Glampedakis et al. 2007; Levin 2007; Vietri et al. 2007; Colaiuda et al. 2009; Gabler et al. 2011, 2013; Levin & van Hoven 2011), and also explored how torsional mode identifications could be used to constrain the properties of the neutron star, including its mass and radius (Samuelsson & Andersson 2007; Sotani et al. 2007, 2016, 2017; Colaiuda & Kokkotas 2011; Miller & Lamb 2016), as well as its internal composition (Watts & Reddy 2007). Other authors have explored emission and X-ray modulation mechanisms for the QPOs (Timokhin et al. 2008; D'Angelo & Watts 2012), and the role of superfluidity in modifying the torsional oscillation

spectrum of neutron stars (Passamonti & Lander 2013, 2014). A focus of current theoretical modeling is the relative importance of magnetic and shear stresses, and to what extent the crustal torsional oscillation spectrum is present in the excitation spectrum, if at all (Gabler et al. 2014, 2016, 2017; Link & van Eysden 2016a, 2016b; Passamonti & Pons 2016; Rezzolla & Ahmedov 2016; Tews 2017).

Most of the data analysis and theoretical attention have focused on the frequencies of the QPOs (for a recent example, see Pumpe et al. 2018). The attraction of this focus is that, potentially, particular oscillation modes can be associated with particular QPO frequencies, and those associations can reveal aspects of the interiors of the oscillating neutron stars. But the *duration* of the signals also contains useful information, particularly about the expected damping mentioned above. Indeed, Huppenkothen et al. (2014) found evidence of such damping of the ~ 625 Hz QPO from SGR 1806–20. However, there has not yet been a systematic study of these QPOs over short intervals.

Here we perform a comprehensive analysis of QPOs from the SGR 1806–20 giant flare. We divide the data into intervals of one second, which is roughly one-eighth of the rotation period, to explore whether there are short-lived QPOs, or whether instead most QPOs are long-lived. This also allows us to determine whether there are QPOs that are present for many rotational periods, but only in some range of rotational phases. Somewhat surprisingly, we find that there is little evidence for long-lived QPOs; instead, there are many QPOs that come and go over times that are likely to be a few tenths of a second. This could provide further support for the idea that coupling of crustal modes to the core MHD continuum damps the modes rapidly. In Section 2 we describe the data set that we analyze, as well as our analysis method (which is a new Bayesian approach not previously used in this context). In Section 3 we present our results, and we conclude with a discussion of the results in Section 4.

2. Methods

2.1. Observations

The data we analyze are the same that were analyzed by Israel et al. (2005) and Strohmayer & Watts (2006). The data were recorded in the Good xenon_2s mode, which has a time resolution of $\approx 1 \mu\text{s}$ and a buffering time of 2 s, thus making it not saturate as easily as the Good xenon_16s mode that happened to be in use during the earlier giant flare from SGR 1900+14. We include all counts without making cuts on the energy channel; the motivation for this is that the burst actually came in through the side of the detector, which made particular energy channel assignments less reliable than normal.

2.2. Data Analysis

The rotational period of SGR 1806–20 is 7.56 s (e.g., Hurley et al. 2005). In order to keep rotational phase information it would be useful to analyze data in segments with durations that are evenly divisible into the rotational period; for example, one-eighth of a period has a duration of 0.945 s. However, because the *RXTE* data in this mode have time bins of duration 2^{-20} s, a fast Fourier transform of data from 0.945 s segments risks aliasing because 0.945 s is not a power of 2 times 2^{-20} s. We therefore divide data into segments that start every 0.945 s, but that are exactly 1 s long each; thus segments overlap each other by 0.055 segments. We have verified via spot checks that using 0.945 s segments or 1.0 s segments produces results that are qualitatively and quantitatively similar to each other. There are 368 such 1 s segments in the data, lasting a total of ≈ 348 s. We also explored whether a sliding window would reveal additional features, but the results were all qualitatively the same as those we obtained in the analyses that we report. We use data starting 20 s after the beginning of the burst (that is, our start time is 2004 December 27 at 21:30:51.378 UTC) to avoid saturation of the *RXTE* Proportional Counter Array detectors. A stacked light curve for the eight segments we use, based on the first 10 periods after our starting point, is in Figure 1; compare this figure with the bottom panel of Figure 1 from Strohmayer & Watts (2006), and note that the top panel of their same figure shows the full light curve of the giant flare, including the obvious rotational modulation.

For each segment, we divide our data into 4096 equal-length bins in time, which therefore last exactly 2^{-12} s each, and use the binned data to construct a power density spectrum from a fast Fourier transform. The resulting spectrum for each segment therefore extends to a Nyquist frequency of 2048 Hz in intervals of 1 Hz. Instead of the commonly used Leahy normalization of the power density spectrum (Leahy et al. 1983), in which the mean power for pure Poisson noise is 2, we use a normalization in which the mean for Poisson noise is 1, to match the assumption of Groth (1975), which we follow in our analysis of the power density spectrum. Our analysis does not consider frequencies below 10 Hz, because lower frequencies have greater red noise contributions from the pulse profile and the overall decay of the emission. We need to consider frequencies at least this low given the reports of some QPOs in the 10–20 Hz range. The contribution of the red noise in the vicinity of 10 Hz is small, which means that our conclusions are insensitive to the precise choice of the frequency floor.

Our Bayesian search for QPOs compares two models:




Figure 1. Light curve for one rotational period, produced by stacking the first 10 periods after the starting point (2004 December 27 at 21:30:51.378 UTC). The segments labeled 0 through 7 correspond to the segments that we analyze, i.e., we divide each of the 46 rotational periods in our data into eight parts, and analyze 1.0 s from the beginning of each segment. Because one-eighth of a rotational period is 0.945 s, this means that each segment overlaps the next by 0.055 s; we make this choice because unlike for a 0.945 s segment, a fast Fourier transform of a 1.0 s segment is commensurate with the intrinsic 2^{-20} s time binning of the *RXTE* data.

1. Red noise model: there is a red noise component, so that the non-Poisson power density at frequency f is

$$P_{\text{red}}(f) = A(f/15 \text{ Hz})^{-2}, \quad (1)$$

where we normalize the power to a fiducial frequency of 15 Hz. Thus, the amplitude A is the only parameter in this model. We chose an f^{-2} slope for our red noise model because this would correspond to an exponential decay. However, we found that the choice of the red noise slope makes no significant difference. In particular, we reanalyzed the strongest signals (listed in Table 1 below) using red noise power laws of f^{-1} , $f^{-1.5}$, $f^{-2.5}$, and $f^{-3.0}$ for both the red noise and Lorentzian models, and all of the Bayes factors were consistent with the values that we report here. We have also carried out an exploratory analysis in which we include the red noise slope as an extra parameter, and the results are generally consistent with the f^{-2} red noise fit.

2. Lorentzian model: there is a red noise component plus a Lorentzian, so the non-Poisson power density at frequency f is

$$P_{\text{Lorentz}}(f) = B(f/15 \text{ Hz})^{-2} + \frac{C}{(\Delta f)^2 + (f - f_0)^2}. \quad (2)$$

Thus, this model has four parameters: B , C , Δf , and f_0 . The red noise amplitude B for the Lorentzian model is determined independently of the parameter A in the red noise model.

We fit each of these two models to all of the 368 data segments, independently. That is, we allow the parameters of the models to vary freely from one data segment to the next.

Table 1
QPOs with Bayes Factors Larger than 1000 and $\Delta f < 10$ Hz

Start Time (s) ^a	Frequency (Hz) ^b	Freq Width (Hz)	Bayes Factor
301.455	31.499–32.724	1.324–2.640	1.80×10^7
102.060	91.239–92.036	0.748–1.593	2.47×10^6
210.735	92.967–93.465	0.566–1.066	9.77×10^5
138.915	21.148–21.605	0.610–1.097	5.09×10^5
177.660	87.870–91.119	2.787–7.296	2.30×10^5
165.375	92.109–92.792	0.631–1.324	2.09×10^5
223.020	23.302–25.214	1.957–4.915	1.97×10^5
192.780	27.574–28.662	0.913–1.956	1.68×10^5
164.430^c	94.731–100.485	2.380–7.281	1.60×10^5
132.300	16.621–17.052	0.556–0.893	6.28×10^4
113.400	156.103–157.109	0.781–1.916	2.21×10^4
89.775	25.857–27.221	1.273–2.820	1.10×10^4
102.060	22.908–23.506	0.557–0.961	9.76×10^3
238.140	28.831–32.546	3.148–9.308	5.63×10^3
83.160	148.988–149.803	0.621–1.172	4.57×10^3
286.335	20.202–21.262	1.702–4.525	4.30×10^3
218.295	27.267–28.475	0.775–1.643	3.09×10^3
172.935	92.024–93.510	0.702–2.698	2.49×10^3
218.295	94.308–97.290	2.000–6.738	1.83×10^3
156.870	91.605–92.506	0.604–1.159	1.48×10^3
132.300	88.422–89.257	0.615–1.380	1.36×10^3
203.175	30.329–31.999	1.077–3.396	1.32×10^3
123.795	50.962–51.984	0.641–3.249	1.02×10^3

Notes.

^a $t = 0$ corresponds to 2004 December 27 at 21:30:51.378 UTC.

^b The numbers for the frequency and frequency width show the 16%–84% credibility range.

^c Boldface indicates a frequency not previously reported.

We follow standard Bayesian procedure by computing an unnormalized posterior probability density p at a given parameter combination by multiplying the prior probability density q at that combination by the likelihood \mathcal{L} of the data given the model at that combination. Thus, for the Lorentzian model,

$$\begin{aligned} p_{\text{Lorentz}}(B, C, \Delta f, f_0) &\propto q_{\text{Lorentz}}(B, C, \Delta f, f_0) \\ \mathcal{L}_{\text{Lorentz}}(B, C, \Delta f, f_0). \end{aligned} \quad (3)$$

The priors are normalized so that their integral over the parameter space is 1, whereas the posteriors would have to be divided by their integral over all parameter space in order to be normalized.

We compute the likelihood using the formulae of Groth (1975), who showed that, for a single frequency bin ($n = 1$ in his Equation (15)), the likelihood of observing a power P when a non-Poisson power P_s is expected (and when the power is normalized such that the mean Poisson-only power is 1) is

$$\mathcal{L}(P; P_s) = e^{-(P+P_s)} \sum_{m=0}^{\infty} \frac{P^m P_s^m}{(m!)^2}. \quad (4)$$

In our analysis we use this formalism, along with the standard assumption that the amplitudes at different frequencies in the fast Fourier transform are statistically independent from each other. Note that in the limit of no non-Poisson power, i.e., $P_s \rightarrow 0$, only $m = 0$ contributes in the sum and thus $\mathcal{L}(P; 0) \rightarrow e^{-P}$, which is the familiar result for Poisson-only power with a mean of 1.

The likelihood of the entire power spectrum that we analyze for a given segment, which has 2038 independent frequencies (because 10 of the 2048 total frequencies are below 10 Hz), is the product of the likelihoods of the observed power at each frequency. For example, for the Lorentzian model,

$$\begin{aligned} \mathcal{L}_{\text{Lorentz}}(B, C, \Delta f, f_0) &= \prod_i \left[e^{-[P_i + P_{\text{Lorentz}}(f_i|B, C, \Delta f, f_0)]} \right. \\ &\quad \left. \times \sum_{m=0}^{\infty} \frac{P_i^m P_{\text{Lorentz}}^m(f_i|B, C, \Delta f, f_0)}{(m!)^2} \right]. \end{aligned} \quad (5)$$

Here, P_i is the observed power at the frequency f_i of bin i , and $P_{\text{Lorentz}}(f_i|B, C, \Delta f, f_0)$ is the non-Poisson power expected at frequency f_i given model parameter values B , C , Δf , and f_0 . We approximate the infinite sum at each frequency by truncating the series when the term being considered has a magnitude less than 10^{-20} of the running sum, which we find gives a fast and accurate measure of the likelihood.

Our prior for the red noise amplitude, for both models, is uniform between 0 and 10 at 15 Hz. Thus, the red noise component is $[0-10](f/15 \text{ Hz})^{-2}$; recall that we are fitting power densities, so, for instance, 0 would mean no non-Poisson noise. For no segment does the best-fit red noise amplitude approach 10, and by definition the amplitude cannot be negative. For similar reasons, our prior for the Lorentzian amplitude C is uniform between 0 and 30.

Our prior for the Lorentzian width Δf is uniform in log width, from $\log_{10}(\Delta f/\text{Hz}) = -0.3$ ($\Delta f \approx 0.5$ Hz) to $\log_{10}(\Delta f/\text{Hz}) = 2$ ($\Delta f = 100$ Hz). We choose a logarithmic prior so that we are not biased in the scale of the width, and we choose this minimum width because at our frequency resolution of ~ 1 Hz, narrower peaks could simply represent single-frequency fluctuations. Our maximum width encompasses the possibility that there could be an overall excess of noise, rather than a sharp peak.

Our prior for the Lorentzian centroid frequency f_0 depends on the search. We select this approach because it is possible that there are several QPOs present in a given data segment; if we had instead searched for a single Lorentzian over all frequencies, it is probable that only the strongest QPO would have been detected. In a given search, we assign a prior probability to f_0 that is uniform in log frequency within the bounds of that search. We base our frequency searches on previously reported QPOs, and thus our frequency bounds are 10–20 Hz; 20–40 Hz; 40–80 Hz; 80–120 Hz; 120–200 Hz; 200–550 Hz; 550–700 Hz; 700–1000 Hz; 1000–1400 Hz; and 1400–2000 Hz. We find that there are a few data segments that contain multiple strong QPOs, but that this is rare. We therefore also performed an exploratory analysis of the data segments in which we used the full frequency range 10–2000 Hz, which led to the same conclusions that we present in this paper.

In order to provide a framework for our later calculation of the significance of possible QPO signals, we also produce a synthetic data set with no red noise and no QPOs. We do this by generating each power density in the spectrum by random draws from the e^{-P} distribution expected for pure Poisson noise; we found that other methods (for example, generating and Fourier-analyzing a count rate curve that is similar to the real data but that has no periodic signals) produce comparable results. We then analyze our synthetic set in the same way that we analyze the real data.




Figure 2. Bayes factors, as a function of segment, for the real data (solid red squares) and for the synthetic data (solid black triangles). Recall that the segments start $1/8$ of a period apart, or 0.945 s, but last 1.0 s each. The different panels are for centroid frequencies in various ranges, which are given in Hertz in the top left of each panel. Note that the power floors for the real and the synthetic data are comparable, which suggests that there is no overall bias in the Bayes factors. Note also that there are far more segments in the real data than in the synthetic data that have Bayes factors greater than unity, or 10, or 100, or other thresholds. Most of these likely contain real signals, but at small Bayes factors individual claims to reality are highly uncertain.

We explore the parameter spaces of both models using an affine-invariant Markov chain Monte Carlo (MCMC) code that we wrote based on the approach of Goodman & Weare (2010). For both the red noise and the Lorentzian search, for every data segment, we perform 20 independent runs with 32 walkers each, and 100 candidate updates of each walker after convergence was established. We find this to be sufficient in all analyses.

After the parameter spaces of both models are explored for a given data segment, we estimate the significance of any possible QPOs by computing Bayes factors. In general, the Bayes factor in favor of some model A (with parameters α) over some model B (with parameters β) is the ratio

$$\mathcal{B}_{AB} = \frac{\int q_A(\alpha) \mathcal{L}_A(\alpha) d\alpha}{\int q_B(\beta) \mathcal{L}_B(\beta) d\beta}. \quad (6)$$

The integral in the numerator, which is over all possible values of the parameters α , is sometimes called the *evidence* for model A, and likewise the integral in the denominator is the evidence for model B. The odds ratio \mathcal{O}_{AB} of model A in favor of model B is simply \mathcal{B}_{AB} multiplied by the prior probability ratio between model A and model B. We set that prior probability

ratio to unity between the Lorentzian and the red noise model, so $\mathcal{O}_{\text{Lorentz,rednoise}} = \mathcal{B}_{\text{Lorentz,rednoise}}$ in our case.

Note that the Bayes factor depends on the priors as well as on the likelihood. For example, if we were to choose a very narrow prior on $\log_{10} f_0$ that happened to be near a strong signal, the Bayes factor would be much larger than if we chose a broad prior. This is why we perform the same analysis, using the same 10 frequency ranges, on the synthetic data (which has no signal) as we do on the real data.

3. Results

3.1. General Characteristics of the Signals

Figure 2 shows the Bayes factors in favor of the Lorentzian model as computed for both the real data (solid red squares and lines) and the synthetic data. There are several points to note:

1. Consistent with previous analyses, there are far more large Bayes factors in the real data than in the synthetic (no signal) data. The highest Bayes factor in the synthetic data is $\mathcal{B} \approx 450$. In contrast, there are a large number of segments and frequency ranges in the real data with $\mathcal{B} > 10^3$, with the highest Bayes factor exceeding 10^7 .




Figure 3. Plot of the best-fit centroid frequencies plus their 1σ uncertainties, for signals with Bayes factors $\mathcal{B} > 1000$ and frequency widths less than 10 Hz. The colors correspond to the magnitude of the Bayes factor: the thin blue bars are for $10^3 < \mathcal{B} < 10^5$, and the thick red bars are for $\mathcal{B} > 10^5$. The horizontal dotted lines show the boundaries between the frequency ranges that we search.

2. Virtually all of the signals last for only one of the 1 s segments. This suggests that almost all the signals damp rapidly. See Section 3.2 for an exploration of whether some signals last for multiple periods.
3. In addition to many signals that have been reported previously, there are other frequencies with strong to nearly conclusive significance. In Table 1 we list the signals found in our analysis with Bayes factors greater than 1000, i.e., larger than the highest Bayes factor found in the synthetic data, which have best-fit frequency widths less than 10 Hz at 84% credibility (so that these are likely to be QPOs rather than excess noise over a broad band). Many of these signals do not appear to have been reported earlier, e.g., the QPOs at 51, 97, and 157 Hz.

In Figure 3 we show a summary of the $\mathcal{B} > 1000$ segments with frequency width less than 10 Hz. The error bars are centered on the best-fit centroid frequencies and their widths represent the 16%–84% credibility range. The color indicates the magnitude of the Bayes factor: the thin blue error bars are for $10^3 < \mathcal{B} < 10^5$, and the thick red error bars are for $\mathcal{B} > 10^5$.

In Figure 4 we plot the rotational phases (as defined in Figure 1) at which QPOs with $\mathcal{B} > 1000$ were detected. The concentration at phases 4 and 7 is clear. This suggests that something about those phases makes it particularly easy to see QPOs. For example, it could be that the crust at those phases has properties that make modes especially easy to generate, or perhaps the magnetosphere at those phases facilitates the production of photons from modes. The favored phases are not quite offset by half of the period, which makes it tempting to think that a slightly off-center dipole magnetic field could produce the favoritism in phase.

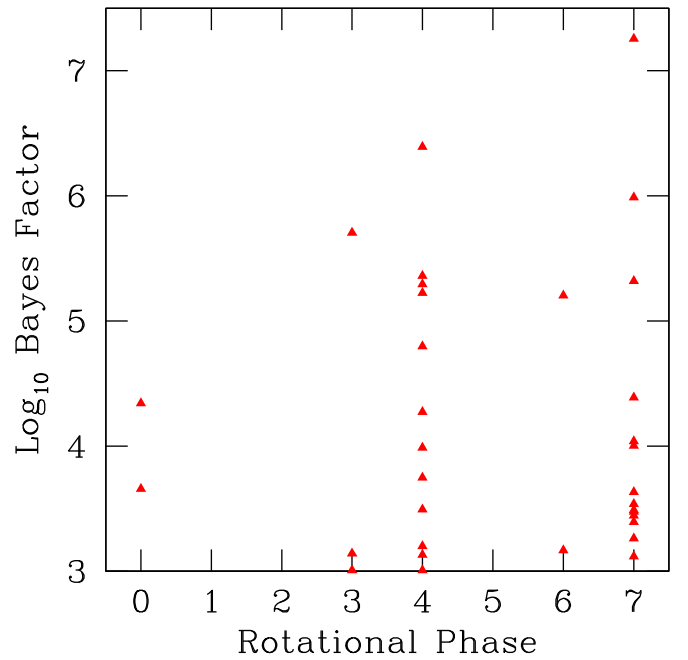


Figure 4. Rotational phases and log Bayes factors for all 1.0 s segments with Bayes factors greater than 1000 and frequency widths less than 10 Hz. Each solid red triangle represents a separate signal. Here, the rotational phase refers to the eight segments of the period; see Figure 1. The concentration in phases 4 and 7 is obvious; 26 of the 33 segments with a QPO that has a Bayes factor greater than 1000 are in those phases, and the remaining 7 are all in adjacent phases. This suggests that either the crust or the magnetosphere at those phases is particularly prone to the generation of the QPOs that we see.

3.2. Are There Persistent Signals?

So far we have focused on short segments of data: 1.0 s, which is close to 1/8 of the rotational period. Based just on the strengths of these signals as a function of segment number, there is no indication that the signals persist for more than one segment. For example, our strongest signal has a Bayes factor of $\mathcal{B} = 1.8 \times 10^7$ at a frequency $f = 32$ Hz. In the 20–40 Hz range, the previous segment has $\mathcal{B} = 0.018$ and the following segment has $\mathcal{B} = 0.432$. Thus, the duration of the signal is likely to be at most a few tenths of a second. More generally, in the entire data set there is only one pair of consecutive 1 s segments with strong signals, and those signals have frequencies that do not overlap at the 2σ level. The QPOs in this pair may therefore result from two completely independent excitations of modes.

There are, however, two additional possibilities to explore. First, as suggested by Strohmayer & Watts (2006), it could be that some modes persist for many rotational periods but are only visible during certain phases of the period. Second, perhaps in addition to the quickly damped oscillations that we detect with our 1.0 s analyses, there are weaker but longer-lasting oscillations at other frequencies.

To explore the first possibility we look at strong signals separated by eight segments, i.e., one full rotational period. Because we want to look for signals that are also related to each other, we apply the following cuts: (1) the signals must appear in consecutive periods rather than simply being separated by an integer number of periods, (2) the $\pm 1\sigma$ centroid frequencies of the two periods must overlap, (3) the frequency widths must be less than 10 Hz, and (4) at least one of the segments must have a Bayes factor greater than 1000. With these cuts, there are

only two possible pairs: the segment starting at 203.175 s and the subsequent segment, where the frequency is ≈ 31 Hz; and the segment starting at 165.375 s and the following segment, where the frequency is ≈ 92 Hz. In the second case, the following segment also has $\mathcal{B} > 1000$.

Thus, in those two cases it is possible that there is an oscillation that persists for a full period. However, given that the associated frequency ranges (20–40 Hz and 80–120 Hz) are the sources of the most common, and the strongest, QPOs (see Figures 2 and 3), it also seems plausible that in both of these cases there were independent excitations, by happenstance, which were separated by one rotational period. Some supporting evidence for this comes from the best-fit frequency widths, which are ≈ 2 Hz for the 31 Hz signal and ≈ 0.7 Hz for the 92 Hz signal. There are many possible origins of the frequency widths, but if they are related to decay rates, then over the 7.56 s rotational period we would expect the signals to drop to undetectability.

To explore the second possibility, of weaker persistent signals, we repeated our analysis for 8.0 s (roughly one period) segments that start every 7.56 s, and for 32.0 s (roughly four periods) segments that start every $4 \times 7.56 = 30.24$ s (again, we choose to analyze data segments with durations that are an integer power of two larger than the 2^{-20} s time resolution in the *RXTE* data). We are again looking for strong signals that have widths consistent with a damping time that is not much less than the duration of the segment.

From our 8.0 s analysis, our three best candidates are:

1. The segment starting 309.96 s from our zero-point has a best-fit frequency of 20.307 Hz, a best-fit Lorentzian width of 0.1884 Hz, and a Bayes factor of 13577.1 relative to the pure red noise model.
2. The segment starting 83.16 s from our zero-point has a best-fit frequency of 149.85 Hz, a best-fit Lorentzian width of 0.208 Hz, and a Bayes factor of 9097.86 relative to the pure red noise model.
3. The segment starting 60.48 s from our zero-point has a best-fit frequency of 240.547 Hz, a best-fit Lorentzian width of 0.4529 Hz, and a Bayes factor of 21350.8 relative to the pure red noise model. We note that no QPO at this frequency has been reported previously for these data.

From our 32.0 s analysis, we have no good candidates; all four of the segments that give Bayes factors that are formally larger than 1000 have best-fit frequency widths that are several times larger than would be consistent with a persistent signal.

We therefore conclude that almost all of the QPOs are produced by events that last a few tenths of a second, and that at most a few QPOs can last as much as a few seconds. Why, then, has it been possible for searches over much longer timescales to find signals?

We can provide some insight using a toy model. Suppose that in a given data set there are m independently excited flares that take the form of exponentially decaying periodic functions of angular frequency ω_0 ; for flare j , which starts at time t_j , we represent the signal by a count rate

$$c_j(t) = a_j e^{-\gamma_j(t-t_j)} e^{i\omega_0 t} e^{i\phi_j}, \quad (7)$$

where ϕ_j is the phase. The power density of the total signal, which we assume has a duration $T \gg 1/\gamma_j$, is

$$P(\omega) \propto \left| \int_0^T \sum_{j=1}^m c_j(t) e^{-i\omega t} dt \right|^2. \quad (8)$$

If we assume for simplicity that $\gamma_j = \gamma$ and $a_j = a$ for all j , then the power density becomes

$$P(\omega) \propto \frac{a^2}{\gamma^2 + (\omega - \omega_0)^2} \left| \sum_{j=1}^m e^{i\phi_j} \right|^2. \quad (9)$$

If all of the flares are in phase ($\phi_j = \phi$ for all j), then the squared factor is m^2 ; if the flares occur at random phases then the expectation value of the squared factor is m .

This indicates that even if $T \gg 1/\gamma$, the characteristic width of the peak in the power density spectrum will be $\sim \gamma$. This provides a possible explanation for a puzzle noted by Strohmayer & Watts (2006): that if the frequency widths are caused by exponential decay then the implied decay time would be far shorter than the lengths of the data segments being analyzed.

We also note that there are circumstances in which the signal could appear to be extremely significant over a relatively long interval of data. If, for example, there are several flares whose phases happen to line up, then the power density will be increased significantly compared to the power density when the flares have phases that are more evenly distributed between 0 and 2π . Moreover, because power density spectra have greater frequency resolution when the interval is longer, even a weak signal with a significant frequency breadth could be detected with high significance because for long-duration data sets there will be many independent frequency bins with powers in excess of the Poisson average. Thus, it seems possible that most of the apparently persistent QPOs in the giant flare of SGR 1806–20 are actually composed of multiple independent flares.

3.3. Can the Frequencies Be Identified?

We now compare our results for the strongest signals given in Table 1 with physical models that could explain the process by which they were generated. If a QPO results from damped oscillations of the generic form

$$\psi(t) = A e^{-\omega_I t} \sin(\omega_R t + \delta), \quad (10)$$

then its power density is given by

$$P_{\text{QPO}}(\omega) = |\Psi(\omega)|^2 = A^2 \left(\frac{(\omega_R \cos \delta + \omega_I \sin \delta)^2 + \omega^2 \sin^2 \delta}{(\omega_R^2 + \omega_I^2 - \omega^2)^2 + 4\omega^2 \omega_I^2} \right). \quad (11)$$

We stress that this is *not* the same as the Lorentzian power density given in the second term of our Lorentzian model (Equation (2)). In particular, the power density also depends on the phase δ . From Figure 3 we see that there are several signals with close but distinct frequencies. Could it be that they actually represent the same modes, only differing in phase?

It is possible to identify the parameters f_0 and Δf of the phenomenological Lorentzian model with the oscillation frequency ω_R and damping time $\tau \equiv 1/\omega_I$ of the damped oscillation. Choosing as limiting cases $\delta = 0$ and $\delta = \pi/2$ we




Figure 5. Frequency of the torsional modes of the crust of a neutron star as a function of the harmonic number ℓ . The exact value of the frequency for each mode depends on the details of the structure and composition of the star. The two limiting cases presented here correspond to a $2.4 M_{\odot}$ with the MS1 equation of state (red line) and to a $1.4 M_{\odot}$ star with the SLy equation of state (blue line), where we used data obtained by de Souza (2017). The points show the strongest QPO frequencies found in our analysis and the corresponding nearest integer value of ℓ according to each model.

have, to leading order,

$$f_0 \left[1 - \frac{1}{2} \left(\frac{\Delta f}{f_0} \right)^2 \right] < \frac{\omega_R}{2\pi} < f_0 \left[1 + \frac{1}{2} \left(\frac{\Delta f}{f_0} \right)^2 \right], \quad (12)$$

$$\Delta f \left[1 - \frac{1}{8} \left(\frac{\Delta f}{f_0} \right)^2 \right] < \tau^{-1} < \Delta f \left[1 + \frac{1}{8} \left(\frac{\Delta f}{f_0} \right)^2 \right]. \quad (13)$$

Consequently, for the strongest narrow QPOs reported in Table 1, where we find a largest value of $\Delta f/f_0 = 0.17$, the maximum fractional correction is ± 0.015 for the frequency and ± 0.004 for the damping time. Therefore, we can safely assume $\omega_R/2\pi \approx f_0$ and $\tau^{-1} \approx \Delta f$. This indicates that a possible difference in phase is not enough to explain the variation in the values of the frequencies.

Another possibility for the variation in the frequencies is a time evolution of the magnetic field strength or geometry after the giant flare. However, this should cause a clearer trend in the frequency behavior than we see. Moreover, a variation of a few Hz in the frequency of the mode would require a change of several times 10^{15} G in the magnetic field (Duncan 1998), which seems unreasonable.

Using a simple model to describe the QPOs as torsional oscillations of the crust, it can be found that the frequencies of modes with different harmonic number ℓ should scale approximately as $f_{\ell} = f_0 \sqrt{\ell(\ell+1)}$ (Hansen & Cioffi 1980). In Figure 5 we choose two limiting values of f_0 : one corresponding to a $2.4 M_{\odot}$ star with the MS1 equation of state, and the other corresponding to a $1.4 M_{\odot}$ star with the SLy equation of state. The expected frequencies in these two models, as a function of the harmonic number ℓ , are shown by the two solid lines. We use the resulting expressions to find the

corresponding nearest integer value of ℓ for our frequencies, which are shown with points. As the plot shows, the frequencies we found with our analysis of the data are compatible with torsional oscillations of the crust, but the mode identification would depend on the largely unknown features of the star. Although the mode identification is uncertain, our results are compatible with previous findings that show that not every ℓ seems to be strongly excited. This could be a consequence of the initial perturbation, i.e., the exact manner in which the crust was originally broken.

4. Discussion

Now we explore the implications our results could have for a model of the source star, the nature of the oscillations, and the mechanism responsible for the emission, that is, in what way the stellar oscillations can be coupled to the radiation emitted by the star and observed on Earth.

The results presented in Figure 4 show that the emission happened predominantly in two nearly opposite rotational phases labeled 4 and 7 in our notation; one full period goes from phases 0 to 7 in our analysis).⁴ A possible way of obtaining this symmetry would be by means of a slightly off-center dipolar magnetic field. The physical mechanism for the emission is not yet clearly understood (but see the models proposed by Timokhin et al. 2008; D’Angelo & Watts 2012; Gabler et al. 2014). Whether it comes from the crust or from the magnetosphere, the QPOs appear to be strong mostly in those two phases. If the giant flare indeed comes from a rearrangement of the magnetic field, the crust at the magnetic poles of the star could be the regions more likely to break; at the same time, the magnetosphere above the poles would have the largest magnetic energy density. Both conditions could amplify the signal in a way such that it would be mostly visible at these two phases, making it seem more beamed, even if it is a broad (thermal) emission.

Perhaps our most striking result is the indication that there are no obviously persistent oscillations in the tail of the giant flare. This is expected from theoretical studies of the properties of the crustal oscillations, which indicate that the crustal modes will couple to a continuum of MHD modes excited below the crust, which quickly damp the crustal oscillations. If, as we discussed in Section 3.3, we estimate the damping time τ of the oscillations by the inverse of the frequency width Δf , our results in Table 1 imply $\tau \approx 0.2$ – 2 s. This is consistent with the findings of Huppenkothen et al. (2014) for a higher frequency QPO ($\tau \approx 0.5$ s), and also roughly consistent with the expectations of Levin (2006) (τ at most 1 s), which take into account the damping resulting from the coupling with the continuum spectrum of MHD modes.

More sophisticated theoretical analyses have introduced the possibility that there are gaps in the continuous spectrum of MHD modes, which could allow for longer-lived oscillations, at least for the lower frequencies detected in the QPOs. Gaps in the continuum could be a consequence of more complicated magnetic field geometries (for example, including a toroidal component or a tangled magnetic field: see Link & van Eysden 2016a, 2016b). However, our analysis shows no compelling evidence for persistent oscillations in any part of the spectrum. This therefore suggests that the QPOs are independent oscillations with distinct

⁴ It is also worth noting that this phase dependence in the appearance of the QPOs is distinct from the phase dependence observed in the number of counts (see Figure 1), indicating that it should be a real effect, and not a bias resulting from the detection of more or less counts.




but close frequencies (see Figure 3). This could be understood as evidence of the existence of the continuous spectrum of MHD modes that was theoretically predicted. If the models for the continuum gap are correct, then the lack of persistent oscillations provides further support for a simple magnetic field configuration close to a pure dipole.

Another consequence of this picture is the need for a continued re-excitation of the modes after the giant flare. If the initial shock causes a starquake, then perhaps aftershocks in the crust provide the energy for the subsequent excitations. This could provide constraints on the nature of the crust, but other unknown mechanisms, perhaps including interactions with the perturbed magnetosphere, could be responsible for the continued input of energy.

Finally, we have not attempted to perform an identification of the modes we have obtained in our analysis. Even in the simplest scenario in which the QPOs are explained as torsional crustal oscillations, the exact frequencies will depend on the mass, compactness, equation of state, crust composition and shear modulus, magnetic field strength and geometry, and so on. This multitude of parameters makes it extremely challenging to identify the modes and to solve the inverse problem, particularly given that there is a lack of relevant analytical expressions or universal relations for these frequencies (work on these issues is in preparation by G. de Souza and C. Chirenti). Nonetheless, the detailed data available for these QPOs means that they can still serve as a rich source of information that can be used to constrain many aspects of the interiors of neutron stars.

We thank Fred Lamb, Yuri Levin, and Bennett Link for valuable discussions and comments on an earlier version of this manuscript. This work was supported in part by joint research workshop award 2015/50421-8 from FAPESP and the University of Maryland, College Park.

ORCID iDs

M. Coleman Miller  <https://orcid.org/0000-0002-2666-728X>
 Cecilia Chirenti  <https://orcid.org/0000-0003-2759-1368>
 Tod E. Strohmayer  <https://orcid.org/0000-0001-7681-5845>

References

- Colaiuda, A., Beyer, H., & Kokkotas, K. D. 2009, *MNRAS*, **396**, 1441
 Colaiuda, A., & Kokkotas, K. D. 2011, *MNRAS*, **414**, 3014
 D'Angelo, C. R., & Watts, A. L. 2012, *ApJL*, **751**, L41
 de Souza, G. H. 2017, PhD thesis, Univ. Campinas
 Duncan, R. C. 1998, *ApJL*, **498**, L45
 Gabler, M., Cerdá-Durán, P., Font, J. A., Müller, E., & Stergioulas, N. 2011, *MNRAS*, **410**, L37
 Gabler, M., Cerdá-Durán, P., Font, J. A., Müller, E., & Stergioulas, N. 2013, *MNRAS*, **430**, 1811
 Gabler, M., Cerdá-Durán, P., Font, J. A., Stergioulas, N., & Müller, E. 2014, *AN*, **335**, 240
 Gabler, M., Cerdá-Durán, P., Mate, A., et al. 2017, *AN*, **338**, 1105
 Gabler, M., Cerdá-Durán, P., Stergioulas, N., Font, J. A., & Müller, E. 2016, *MNRAS*, **460**, 4242
 Glampedakis, K., Samuelsson, L., & Andersson, N. 2007, *Ap&SS*, **308**, 607
 Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, **5**, 65
 Groth, E. J. 1975, *ApJS*, **29**, 285
 Hansen, C. J., & Cioffi, D. F. 1980, *ApJ*, **238**, 740
 Huppenkothen, D., Watts, A. L., & Levin, Y. 2014, *ApJ*, **793**, 129
 Hurley, K., Boggs, S. E., Smith, D. M., et al. 2005, *Natur*, **434**, 1098
 Israel, G. L., Belloni, T., Stella, L., et al. 2005, *ApJL*, **628**, L53
 Leahy, D. A., Darbro, W., Elsner, R. F., et al. 1983, *ApJ*, **266**, 160
 Levin, Y. 2006, *MNRAS*, **368**, L35
 Levin, Y. 2007, *MNRAS*, **377**, 159
 Levin, Y., & van Hoven, M. 2011, *MNRAS*, **418**, 659
 Link, B., & van Eysden, C. A. 2016a, *ApJS*, **224**, 6
 Link, B., & van Eysden, C. A. 2016b, *ApJL*, **823**, L1
 Miller, M. C., & Lamb, F. K. 2016, *EPJA*, **52**, 63
 Passamonti, A., & Lander, S. K. 2013, *MNRAS*, **429**, 767
 Passamonti, A., & Lander, S. K. 2014, *MNRAS*, **438**, 156
 Passamonti, A., & Pons, J. A. 2016, *MNRAS*, **463**, 1173
 Pompe, D., Gabler, M., Steinger, T., & Enßlin, T. A. 2018, *A&A*, **610**, A61
 Rezzolla, L., & Ahmedov, B. J. 2016, *MNRAS*, **459**, 4144
 Samuelsson, L., & Andersson, N. 2007, *MNRAS*, **374**, 256
 Sotani, H., Iida, K., & Oyamatsu, K. 2016, *NewA*, **43**, 80
 Sotani, H., Iida, K., & Oyamatsu, K. 2017, *MNRAS*, **464**, 3101
 Sotani, H., Kokkotas, K. D., & Stergioulas, N. 2007, *MNRAS*, **375**, 261
 Strohmayer, T. E., & Watts, A. L. 2005, *ApJL*, **632**, L111
 Strohmayer, T. E., & Watts, A. L. 2006, *ApJ*, **653**, 593
 Tews, I. 2017, *PhRvC*, **95**, 015803
 Timokhin, A. N., Eichler, D., & Lyubarsky, Y. 2008, *ApJ*, **680**, 1398
 Vietri, M., Stella, L., & Israel, G. L. 2007, *ApJ*, **661**, 1089
 Watts, A. L., & Reddy, S. 2007, *MNRAS*, **379**, L63