

This is the accepted manuscript made available via CHORUS. The article has been published as:

All-sky, narrowband, gravitational-wave radiometry with folded data

Eric Thrane, Sanjit Mitra, Nelson Christensen, Vuk Mandic, and Anirban Ain

Phys. Rev. D **91**, 124012 — Published 3 June 2015

DOI: [10.1103/PhysRevD.91.124012](https://doi.org/10.1103/PhysRevD.91.124012)

All-sky, narrowband, gravitational-wave radiometry with folded data

Eric Thrane,^{1, a} Sanjit Mitra,² Nelson Christensen,³ Vuk Mandic,⁴ and Anirban Ain²

¹*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

²*Pune University Campus, Pune 411007, India*

³*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057, USA*

⁴*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

Gravitational-wave radiometry is a powerful tool by which weak signals with unknown signal morphologies are recovered through a process of cross correlation. Radiometry has been used, e.g., to search for persistent signals from known neutron stars such as Scorpius X-1. In this paper, we demonstrate how a more ambitious search—for persistent signals from *unknown* neutron stars—can be efficiently carried out using folded data, in which an entire \sim year-long observing run is represented as a single sidereal day. The all-sky, narrowband radiometer search described here will provide a computationally tractable means to uncover gravitational-wave signals from unknown, nearby neutron stars in binary systems, which can have modulation depths of ≈ 0.1 –2 Hz. It will simultaneously provide a sensitive search algorithm for other persistent, narrowband signals from unexpected sources.

I. INTRODUCTION

Gravitational-wave radiometry is a technique by which two or more detectors are cross-correlated in order to identify signals in the form of excess coherence [1, 2]. It is especially well-suited to situations where it is either impossible or impractical to carry out a matched filter search due to theoretical uncertainty and/or the vastness of the signal parameter space. LIGO [3] and Virgo [4] radiometer searches for persistent gravitational waves have yielded limits on the gravitational-wave strain from targets including Scorpius X-1, the Galactic Center, and Supernova 1987A [5]. Radiometer searches are also sensitive to hot spots created from the superposition of many persistent sources [6–8]. By restricting the timescale of integration, the same method has been applied to search for relatively long-lived ~ 10 –1000 s gravitational-wave transients [9], e.g., associated with long gamma-ray bursts [10]. Radiometry has also been proposed in the context of gravitational-wave astronomy with pulsar timing arrays [11].

Due to its robustness and efficiency, there is strong motivation to extend the parameter space of radiometer searches. In this paper, we propose a method to extend the search for persistent narrowband signals to target all frequencies *and* all directions on the sky—not just ones associated with targets such as Scorpius X-1. An efficient all-sky search for persistent narrowband gravitational waves could facilitate the detection, e.g., of electromagnetically quiet spinning neutron stars in binary systems. The challenge is to overcome computational challenges that make the search difficult.

There is a significant burden from both data storage and computation time. In previous *broadband* radiometer searches for persistent gravitational waves [5], a Fisher matrix was used to characterize the covariance between

different patches of sky [12]. However, saving Fisher matrices for many frequency bins (and for many “jobs” associated with stretches of science quality data) creates a burdensome data storage problem. The computational cost of analyzing a full year of data to look at thousands of frequency bins is both inefficient and potentially prohibitive in terms of computation time.

Recent work [13] shows that by “folding” data into a single sidereal day, it is possible to represent an entire year of radiometer data with just one twenty-four long spectrogram. We show that this several-hundred-fold reduction in data volume facilitates efficient searches for persistent narrowband signals.

The remainder of this paper is organized as follows. In Section II, we discuss the motivation for narrowband all-sky searches. Section III describes how previous work on folded data [13] and searches for \approx day-long signals [14] can be combined to design an efficient all-sky narrowband radiometer. In Section IV, we demonstrate an all-sky narrowband radiometer search using folded data. Finally, in Section V, we discuss the astrophysical implications of these results.

II. MOTIVATION

One of the primary targets for an all-sky narrowband search for persistent gravitational waves is spinning neutron stars in binary systems. The motivation for such a search was recently laid out in [15]. In a variety of scenarios, accretion from a companion star can induce a time-varying quadrupole moment, which, in some cases, may persist after accretion abates. For example, persistent localized mass accumulation may occur due to magnetic fields [16] depending on unknown details of neutron star physics [17–21]. Magnetic fields may also induce deformations in the stellar interior [22]. Rotational instabilities such as *r*-modes may be sustained through accretion [23, 24].

When the neutron star in a binary is a pulsar, it is

^aElectronic address: eric.thrane@monash.edu

possible to carry out an optimal search using matched filtering. However, if it is electromagnetically quiet, then the space of possible signals becomes prohibitively large for a fully coherent search. When at least the sky location is known, one can apply a targeted radiometer search [1, 2, 5] or other semicoherent methods [15, 25–28]. If the sky location is unknown, the problem is significantly more challenging and the signal might be missed or vetoed as a noise artifact by non-specialized pipelines [15].

The search technique proposed here is highly complementary to the semicoherent method described in [15, 25]. The “TwoSpect” algorithm [15, 25] employs a canonical signal model in which a neutron star emits periodic waves, modulated by binary motion. The radiometer, on the other hand, employs only minimal assumptions; namely, that the signal is persistent and narrowband. By incorporating additional assumptions about the signal model, a tuned search gains sensitivity. The advantage of the radiometer, in contrast, is that it is extremely robust since it relies only on excess coherence in two or more detectors. Indeed, while we focus on spinning neutron stars in binary systems, the narrowband radiometer is sensitive to *any* persistent narrowband source, whether or not it is modulated by binary motion.

III. METHOD

A. Folding

We utilize the idea of folded data as described in [13], which we reformulate in the language of recent radiometer developments [9, 10, 14, 29–32]. During an observing run, typically lasting a few months to a few years, data is collected from two spatially separated detectors. (This discussion straightforwardly extends to three or more detectors, but our presentation focuses on the two-detector case for simplicity.) We parse the data for which both detectors are simultaneously operational into sidereal days (23 hr, 56 min, 4 s). Every sidereal day of radiometer data can be represented using a complex-valued estimator [14]:

$$\hat{\mathfrak{Y}}(t; f) \equiv \frac{2}{\mathcal{N}} \tilde{s}_1^*(t; f) \tilde{s}_2(t; f) \quad (1)$$

$$\sigma_{\mathfrak{Y}}(t; f) \equiv \frac{1}{2} \sqrt{P'_1(t; f) P'_2(t; f)}. \quad (2)$$

Here, $(t; f)$ are spectrogram indices: t is the start time of each data segment and f is the frequency of the Fourier transformed data associated with that segment. Each $\tilde{s}_I(t; f)$ refers to the discrete Fourier transform of the strain data from detector I . The variable \mathcal{N} is a Fourier normalization constant, $\sigma_{\mathfrak{Y}}(t; f)$ is an estimator for the uncertainty associated with $\hat{\mathfrak{Y}}$, and each $P'_I(t; f)$ represents the strain auto-power measured in detector I . (The prime denotes that the auto-power is calculated as an average of neighboring segments.) The direction of the

source is encoded in the phase of $\mathfrak{Y}(t; f)$, which depends on the time delay between the two detectors [30].

Since the signal-induced phase delay and detectors’ antenna factors are periodic on the timescale of a sidereal day, it is possible to define folded estimators by summing data over many sidereal days:

$$\hat{\mathfrak{Y}}^{\text{fold}}(t; f) = \sum_k \hat{\mathfrak{Y}}(t; f|k) \sigma_{\mathfrak{Y}}^{-2}(t; f|k) / \sum_k \sigma_{\mathfrak{Y}}^{-2}(t; f|k) \quad (3)$$

$$\sigma_{\mathfrak{Y}}^{\text{fold}}(t; f) = \left(\sum_k \sigma_{\mathfrak{Y}}^{-2}(t; f|k) \right)^{-1/2}. \quad (4)$$

Here, k runs over all the sidereal days of the observing run.

Following [30], we define complex signal-to-noise ratio for folded data:

$$\mathfrak{p}^{\text{fold}}(t; f) \equiv \hat{\mathfrak{Y}}^{\text{fold}}(t; f) / \sigma_{\mathfrak{Y}}^{\text{fold}}(t; f). \quad (5)$$

The variable $\mathfrak{p}^{\text{fold}}(t; f)$ is a complex-valued spectrogram; see Fig. 1. Gravitational-waves induce excess $\mathfrak{p}(t; f)$, which appears visually as brighter than usual spectrogram pixels.

By folding a year of data into a sidereal day, the volume of data is reduced by a factor of ≈ 365 , which dramatically reduces both storage needs and computation time. And, if we restrict our attention to persistent narrowband signals, then the folding operation is *lossless* in the sense that there is no more information (useful in a search for persistent signals) in a year of cross-correlated data than there is in a folded day.

B. Radiometry

By specializing the methods from [14] to consider approximately monochromatic signals lasting exactly one sidereal day, and then applying these methods to folded data [13], it is possible to construct a very efficient algorithm for the detection of electromagnetically quiet binary neutron stars (and other narrowband sources not matching a standard continuous wave model). Note that when we refer in this paper to signals as “approximately monochromatic,” we mean that they fit within single ≈ 1 Hz frequency bin. Such signals include neutron stars in binary systems, which are relatively broadband compared to isolated neutron stars. However, the signal is narrowband in the sense that it is contained within just one radiometer frequency bin. In reality, the modulation depth of neutron stars in binaries can range from ≈ 0.1 – 2 Hz [15], so a more nuanced definition of “narrowband,” employing variable frequency bin sizes, is appropriate. However, this is beyond our present scope.

Radiometer searches can be cast as a pattern recognition problems in which an algorithm looks for statistically significant clusters of pixels [9]. As we are focused

here on approximately monochromatic signals, the optimal pattern recognition algorithm is a simple sum over the frequency bins of \mathbf{p}^{fold} . In the language of seedless clustering [29, 30], each signal template corresponds to a different row in the \mathbf{p}^{fold} spectrogram. We assume that the signal starts at the beginning of the spectrogram and finishes at the end. To assume otherwise would imply an unphysical signal that starts and stops with a period matching Earth’s sidereal day. Likewise, it would not make sense to allow for any frequency evolution since this would imply that the signal repeats this evolution with the period of a sidereal day.

The following expression for the detection statistic $\text{SNR}_{\text{tot}}(f|\hat{\Omega})$ is specialized from [14] to focus on monochromatic signals:

$$\text{SNR}_{\text{tot}}(f|\hat{\Omega}) = \frac{\text{Re} \left[\sum_t e^{(2\pi i f \hat{\Omega} \cdot \Delta \vec{x}(t)/c)} \mathbf{p}^{\text{fold}}(t; f) \epsilon_{12}(t|\hat{\Omega}) \right]}{\left(\sum_t \epsilon_{12}^2(t|\hat{\Omega}) \right)^{1/2}}. \quad (6)$$

Here, $\hat{\Omega}$ is the direction of the source, $\Delta \vec{x}$ is the difference in detector position, and c is the speed of light. The phase factor $e^{(2\pi i f \hat{\Omega} \cdot \Delta \vec{x}(t)/c)}$ “points” the spectrogram toward the source by rotating the phase angle of \mathbf{p}^{fold} to zero so that the observed signal-to-noise ratio is real and positive [14, 30]. The variable $\epsilon_{12}(t|\hat{\Omega})$ is a time-dependent efficiency factor characterizing the fraction of gravitational-wave power measured due to the non-unity antenna factors of gravitational-wave interferometers:

$$\epsilon_{12}(t|\hat{\Omega}) \equiv \frac{1}{2} \sum_A F_1^A(t|\hat{\Omega}) F_2^A(t|\hat{\Omega}). \quad (7)$$

Here, $F_I^A(t|\hat{\Omega})$ is the antenna factor [33] for detector I and $A = +, \times$ are the different polarization states. As the Earth rotates, the detectors become more/less favorably aligned relative to the source. The sum over t is carried out over all the (typically ≈ 60 s) segments in a sidereal day. For additional information pertaining to Eq. 6, the interested reader is referred to [14] and reference therein.

For the sake of compact notation, Eq. 6 assumes that detector noise is approximately constant over the span of a sidereal day. However, we note that it is straightforward to generalize to the case of non-stationary noise by weighting each data segment appropriately using a standard inverse variance weighting method [34].

Previous work [14, 30–32] has shown that Eq. 5 and Eq. 6 can serve as the basis for “all-sky” searches for which the signal time and direction are unknown a priori. Many directions in the sky can be considered simultaneously using parallelized computer code run on suitable processors, e.g., GPUs and multi-core CPUs [30]. For each observation, we record $\text{SNR}_{\text{tot}}^{\text{max}}$ —the maximal value of SNR_{tot} for all the directions:

$$\text{SNR}_{\text{tot}}^{\text{max}}(f) = \max_{\hat{\Omega}} \left[\text{SNR}_{\text{tot}}(f|\hat{\Omega}) \right]. \quad (8)$$

Eq. 6 assumes a fairly constrained signal model: approximately monochromatic signals, which start at the

beginning of and finish at the of a sidereal day. Thus, it is possible to search every sky location (given some resolution) using a HEALPix [35] grid. In this work, we consider a number of sky locations chosen so as to be spaced at angular separations less than the diffraction limited resolution of the detector network. For the LIGO network,

$$\theta \approx \frac{c}{f} \frac{1}{|\Delta \vec{x}|} \approx \left(\frac{1000 \text{ Hz}}{f} \right) 5^\circ. \quad (9)$$

In the demonstration described below, we use HEALPix to efficiently sample the sky with 3072 tiles.

C. Comparison with targeted searches

In order to highlight the usefulness of this framework, it is useful to compare it to the procedure used for previous narrowband *targeted* radiometer searches [1, 5], applied, e.g., to a single direction such as Scorpius X-1. In such searches, one analyzes data in ≈ 500 – 10000 s-long “jobs,” defined as data stretches during which coincident data is available for both detectors. During initial LIGO’s fifth science run, there were ≈ 20000 jobs. For each job, one calculates cross-power (Eq. 1) and auto-powers (Eq. 2) by optimally combining the results from each of many (typically) 60 s segments in each job. Then, during post-processing, the results from each job are optimally combined to produce estimators analogous to $\hat{\mathfrak{V}}$ and $\sigma_{\mathbf{p}}$, but obtained by integrating over the entire observing run.

This procedure is repeated separately for each sky location. Using this scheme, each sky location requires the generation of tens of thousands of output files, and so there is significant overhead associated with the analysis of more than a few sky locations. Moreover, covariances between different patches on the sky—characterized by a Fisher matrix—lead to subtleties when interpreting the significance of an outlier [5]. In order to understand these covariances in this framework, one must therefore produce also a Fisher matrix for every frequency bin.

In the new paradigm we propose here, the data are first combined into a single sidereal day, which can be used to look in every sky direction. This data is small enough to be easily manipulated on a typical computer. Once the data are loaded, we can search over every sky location all at once without having to read or write additional files, thereby eliminating a significant bottleneck. Also, this method eliminates the need to store Fisher matrices. This is because the covariances between different sky locations are automatically encoded into the factors of $e^{(2\pi i f \hat{\Omega} \cdot \Delta \vec{x}/c)}$ and $\epsilon(t|\hat{\Omega})$.

Thus, one can use numerical simulations, in which \mathbf{p} is drawn from a Gaussian distribution, in order to determine the false-alarm probability of obtaining a given value of $\text{SNR}_{\text{tot}}^{\text{max}}$. (The assumption of Gaussianity is justified by invoking the central limit theorem, and it

has been born out in previous measurements [5, 36].) An example search, analyzing 750 Hz of bandwidth with $\Delta f = 1$ Hz resolution (see Fig. 1), and scanning over 3072 Healpixels, can be carried out in $\lesssim 100$ s using a single eight-core CPU. This computation time does not include the modest time required for folding the data.

IV. DEMONSTRATION

Our goal now is to demonstrate the recovery of a persistent narrowband signal in folded data without a priori knowledge of the sky location or frequency of the signal. To begin, we generate twenty days of simulated data for the LIGO Hanford and LIGO Livingston detectors operating at design sensitivity. The data consists of Gaussian noise plus a simulated monochromatic signal with frequency $f = 600$ Hz and strain amplitude $h_0 = 1.5 \times 10^{-24}$, located at $(\text{ra}, \text{dec}) = (18.5 \text{ hr}, 39^\circ)$. The signal is circularly polarized. For our purposes, a monochromatic signal is a reasonable proxy for a signal modulated by binary motion (discussed in Section II) so long as the binary modulation is $\lesssim 1$ Hz.

We generate coarse-grained spectrograms with a resolution of (79 s, 1 Hz) (with non-overlapping segments). This frequency bin width is relatively well-matched to the observed modulation depth of known neutron stars in binary systems, which can range from ≈ 0.1 –2 Hz [15]. (A more systematic study should employ variable bin widths to fully cover parameter space.) We study a band ranging from 500–1250 Hz, where many neutron stars in binary systems are expected to rotate [37]. Next, the data are folded into one sidereal day following the procedure outlined in Section III. As the data are folded, we analyze the integrated spectrogram with every accumulated day in order to study how the signal grows with time. Following the procedure described in Eqs. 6 and 8, we scan 3072 Healpixels and record the maximum signal-to-noise ratio at each frequency $\text{SNR}_{\text{tot}}^{\text{max}}(f)$.

The results are summarized in Fig. 1. In Fig. 1a, we show the real part of $e^{(2\pi i f \hat{\Omega} \cdot \Delta \vec{x}/c)} \mathbf{p}^{\text{fold}}(t; f)$ representing twenty days of data folded into a single sidereal day. Though the signal is virtually impossible to see with the naked eye in just one day of data, it is visible in this image at 600 Hz since the signal-to-noise ratio grows with continued integration. Fig. 1b shows the recovered signal ($\text{SNR}_{\text{tot}}^{\text{max}} \approx 12$) using just one day of data. The excellent match in the time-dependent modulation suggest that the sky location is well matched. Indeed, the best fit direction $(\text{ra}, \text{dec}) = (18.3 \text{ hr}, 38.9^\circ)$ is the closest possible pixel to the true source location.

In Fig. 1c, we show how the signal-to-noise ratio of the 600 Hz signal (blue) grows as expected like the square root of the total observation time. The dashed red curve shows the next-loudest frequency bin, which exhibits no tendency to grow or decline with continued integration. In Fig. 1d, we show the signal-to-noise ratio spectrum for the spectrogram in Fig. 1a, maximized over search

directions. That is, for each frequency bin, we scanned the entire sky and plot (in red) $\text{SNR}_{\text{tot}}^{\text{max}}$ for the brightest patch of sky as a function of frequency. The blue curves indicate typical one-sigma fluctuations for noise. The 600 Hz signal is visible as a dramatic spike. There is a slight tendency toward higher values of SNR with increasing frequency as angular resolution improves and there are more independent patches of sky.

From Fig. 1, we show that it is possible carry out a computationally efficient search for persistent narrowband signals from all directions in the sky and at all frequencies. Once the data have been folded, the entire search takes less than two minutes to carry out using a single eight-core CPU.

By studying the background distribution of $\text{SNR}_{\text{tot}}^{\text{max}}(f)$, we can estimate how the sensitivity of the radiometer search differs for a directed and an all-sky search. Naively, we expect the all-sky sensitivity to be worse, but only slightly, since the additional trial factors incurred by looking in many directions will only slightly increase the detection threshold given the expected rapidly falling background distribution. In order to identify a detection candidate with false alarm probability 1%, the targeted search requires $\text{SNR}_{\text{tot}}^{\text{max}}(f) \gtrsim 4.5$ while the all-sky search requires $\text{SNR}_{\text{tot}}^{\text{max}}(f) \gtrsim 5.8$ for the bandwidth and spectral resolution considered here. This corresponds to a change in strain sensitivity of just $(5.8/4.5)^{1/2} - 1 \approx 14\%$.

Upper limits can be set for each frequency bin independently, and so there are fewer trial factors to apply compared to the detection calculation, which must account for the existence of hundreds of frequency bins. On average, the targeted search will set limits corresponding to $\text{SNR}_{\text{tot}}^{\text{max}}(f) \approx 1.7$ whereas the all-sky search will set limits corresponding to $\text{SNR}_{\text{tot}}^{\text{max}}(f) \approx 4$ (around $f \approx 1000$ Hz where we scan ≈ 3072 sky positions). Thus, the all-sky limits are expected to be $\approx 50\%$ higher in strain than limits obtained from a targeted search. These estimates support the expectation that the all-sky search is only slightly less sensitive than the targeted search.

V. CONCLUSIONS

Rotating neutron stars in binaries are a very promising candidate for gravitational-wave detection by second-generation detectors like Advanced LIGO/Virgo. Previous work [25] has enabled the exploration of these promising sources, though, no detection candidates have been found with initial LIGO data [15]. Our presentation of an all-sky, narrowband radiometer using folded data is complementary since it makes only minimal assumptions about the source. In this way, it should be sensitive not only to neutron stars in binary systems, but also to other sources, which emit approximately narrowband gravitational waves and which may or may not conform to the canonical model of an isolated neutron star.

A back-of-the-envelope calculation provides promising

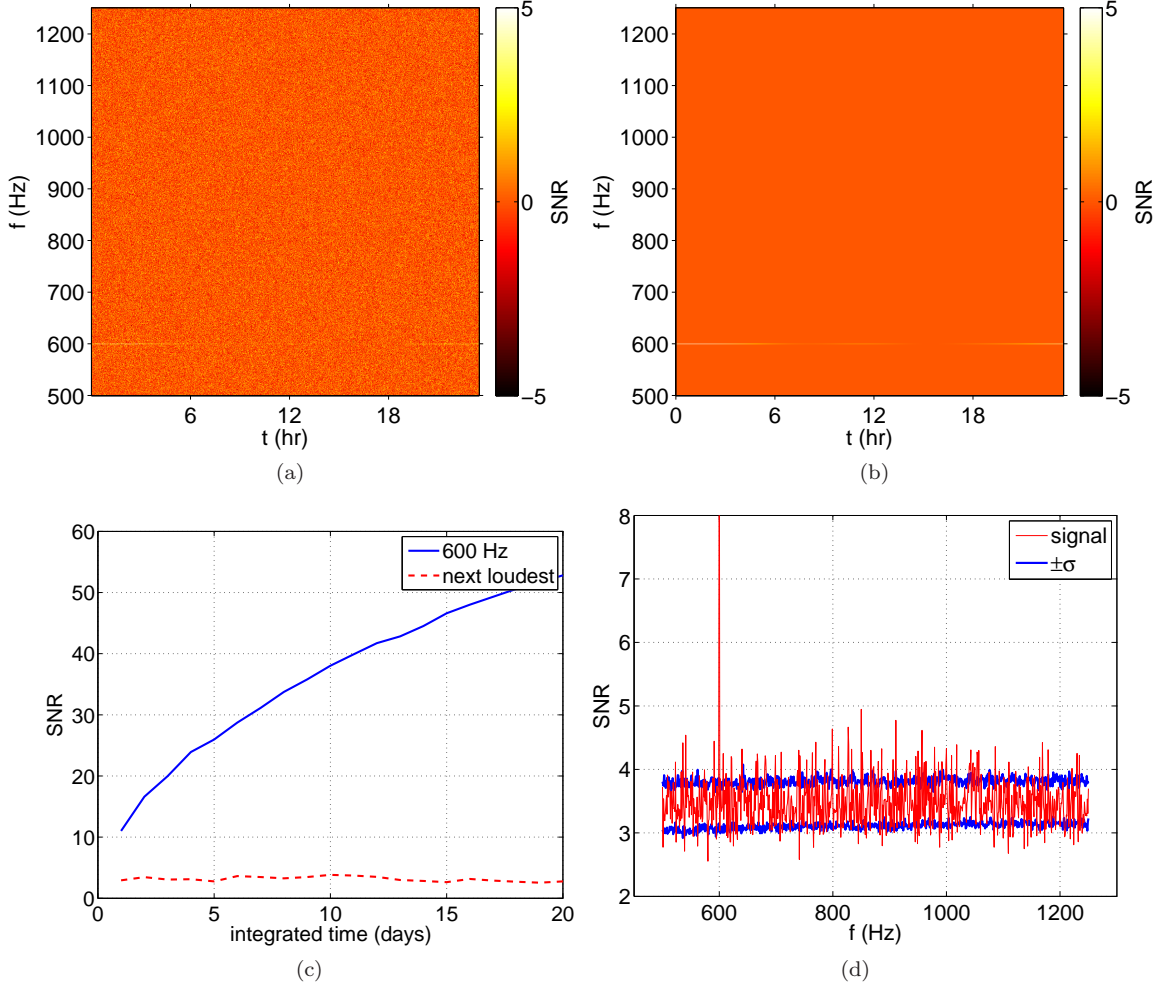


FIG. 1: Narrowband signals in folded data. Top-left: a $(\delta t, \Delta f) = (79 \text{ s}, 1 \text{ Hz})$ spectrogram of the real part of $e^{(2\pi i f \Omega \cdot \Delta \vec{x}/c)} \mathbf{p}^{\text{fold}}(t; f)$ consisting of twenty days of data, folded into one sidereal day. The data consist of Advanced LIGO [3] Monte Carlo noise plus a simulated signal with amplitude $h_0 = 1.5 \times 10^{-24}$ and frequency 600 Hz. The source is assumed to be face-on. The spectrogram is “pointed” in the correct direction using the appropriate phase factor to make the signal positive; see Eq. 1. The simulated signal is detectable in just one day of data, but we show the result of 20 days of folding to make the signal visible by eye. The top-right panel shows the recovered signal in just one day of data. The signal is easily identified without any prior assumption about the source location or frequency. The bottom-left panel shows how the folded SNR grows with the addition of folded data. Blue shows the result for the bin with the signal while red is the next loudest bin. The bottom-right panel shows the maximum SNR (scanning over the entire sky) as a function of frequency for one day of data (red). The injected signal is evident as a SNR ≈ 12 spike at 600 Hz. Typical fluctuations due to pure noise are indicated with blue. It is interesting to note that there is a slight trend toward higher SNR with increasing frequency due to increasing angular resolution.

results. In Subsection IV, we showed that a $h_0 = 1.5 \times 10^{-24}$ signal at $f = 600 \text{ Hz}$ can be easily recovered in one day of data. Scaling to one year of data, this becomes $h_0 \approx 2 \times 10^{-25}$. In terms of neutron-star ellipticity, this sensitivity can be restated [38] as

$$\epsilon \approx 1 \times 10^{-5} \left(\frac{0.4}{\beta} \right) \left(\frac{10^{45} \text{ g cm}^2}{I} \right) \left(\frac{r}{10 \text{ kpc}} \right) \left(\frac{600 \text{ Hz}}{f} \right)^2, \quad (10)$$

where β is an orientation factor, G is the gravitational constant, r is the distance to the source, and I is the

moment of inertia. (One arrives at a similar rough estimate by scaling the upper limits from targeted radiometer analyses with initial LIGO data [5] to Advanced LIGO sensitivity.) This is an order of magnitude above the expected strain at $f = 600 \text{ Hz}$ from Scorpius X-1 expected from accretion-torque balance [26], suggesting that a signal might be detectable from another somewhat closer neutron star—especially in light of the fact that additional optimizations are possible, and only a hint of a signal is needed to trigger a coherent matched filtering search, which can boost the significance.

VI. ACKNOWLEDGMENTS

This is LIGO document P1500015.

VM's work was supported by NSF grant PHY1204944.
NC's work was supported by NSF grant PHY-1204371.

-
- [1] B. Abbott et al., Phys. Rev. D **76**, 082003 (2007).
 - [2] S. W. Ballmer, Classical Quantum Gravity **23**, S179 (2006).
 - [3] Harry, G. M. for the LIGO Scientific Collaboration, Classical Quantum Gravity **27**, 084006 (2010).
 - [4] F. Acernese et al., Classical Quantum Gravity **32**, 024001 (2015).
 - [5] B. P. Abbott (The LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **107**, 271102 (2011).
 - [6] S. Dhurandhar, H. Tagoshi, Y. Okada, N. Kanda, and H. Takahashi, Phys. Rev. D **84**, 083007 (2011).
 - [7] N. Mazumder, S. Mitra, and S. Dhurandhar, Phys. Rev. D **89**, 084076 (2014).
 - [8] D. Talukder, E. Thrane, S. Bose, and T. Regimbau, Phys. Rev. D **89**, 123008 (2014).
 - [9] E. Thrane, S. Kandhasamy, C. D. Ott, et al., Phys. Rev. D **83**, 083004 (2011).
 - [10] J. Aasi et al., Phys. Rev. D **88**, 122004 (2013).
 - [11] M. Anholm, S. Ballmer, J. D. E. Creighton, L. R. Price, and X. Siemens, Phys. Rev. D **79**, 084030 (2009).
 - [12] E. Thrane, S. Ballmer, J. Romano, et al., Phys. Rev. D **80**, 122002 (2009).
 - [13] A. Ain, P. Dalvi, and S. Mitra (2015), <http://xxx.tau.ac.il/abs/1504.01714>.
 - [14] E. Thrane, V. Mandic, and N. Christensen (2015), <http://arxiv.org/abs/1501.06648>.
 - [15] Aasi et al., Phys. Rev. D **90**, 062010 (2014).
 - [16] L. Bildsten, Astrophys. J. Lett. **501**, L89 (1998).
 - [17] M. Vigelius and A. Melatos, MNRAS **395**, 1972 (2009).
 - [18] M. Priymak, A. Melatos, and D. J. B. Payne, MNRAS **417**, 2696 (2011).
 - [19] K. Wette, M. Vigelius, and A. Melatos, MNRAS **402**, 1099 (2010).
 - [20] M. Vigelius and A. Melatos, MNRAS **395**, 1985 (2009).
 - [21] M. Vigelius and A. Melatos, MNRAS **386**, 1294 (2008).
 - [22] A. Melatos and D. J. B. Payne, Astrophys. J. **623**, 1044 (2005).
 - [23] A. Reisenegger and A. Bonacic, Phys. Rev. Lett. p. 201103 (2003).
 - [24] *R-modes in accreting and young neutron stars*, vol. 575 of *AIP Conference Proceedings* (American Institute of Physics, Melville, NY, 2001).
 - [25] E. Goetz and K. Riles, Classical Quantum Gravity **28**, 215006 (2011).
 - [26] L. Sammut, C. Messenger, A. Melatos, and B. J. Owen, Phys. Rev. D **89**, 043001 (2014).
 - [27] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Phys. Rev. D **77**, 082001 (2008).
 - [28] S. van der Putten, H. J. Bulten, J. F. J. van den Brand, and M. Holtrop, Journal of Physics: Conferences Series **228**, 012005 (2010).
 - [29] E. Thrane and M. Coughlin, Phys. Rev. D **88**, 083010 (2013).
 - [30] E. Thrane and M. Coughlin, Phys. Rev. D **89**, 063012 (2014).
 - [31] M. Coughlin, E. Thrane, and N. Christensen, Phys. Rev. D **90**, 083005 (2014).
 - [32] M. Coughlin, P. Meyers, E. Thrane, J. Luo, and N. Christensen, Phys. Rev. D **91**, 063004 (2015).
 - [33] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1989).
 - [34] B. Allen & J. D. Romano, Phys. Rev. D **59**, 102001 (1999).
 - [35] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622**, 759 (2005).
 - [36] B. Abbott et al., Nature **460**, 990 (2009).
 - [37] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Winnands, M. van der Klis, and C. B. Markwardt, Nature **424**, 42 (2003).
 - [38] B. Abbott et al., Phys. Rev. D **76**, 042001 (2007).