

PAPER

Tuning into Scorpius X-1: adapting a continuous gravitational-wave search for a known binary system

To cite this article: Grant David Meadors *et al* 2016 *Class. Quantum Grav.* **33** 105017

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © © 2016 IOP Publishing Ltd.

During the embargo period (the 12 month period from the publication of the Version of Record of this article), the Accepted Manuscript is fully protected by copyright and cannot be reused or reposted elsewhere.

As the Version of Record of this article is going to be / has been published on a subscription basis, this Accepted Manuscript is available for reuse under a CC BY-NC-ND 3.0 licence after the 12 month embargo period.

After the embargo period, everyone is permitted to use copy and redistribute this article for non-commercial purposes only, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by-nc-nd/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions will likely be required. All third party content is fully copyright protected, unless specifically stated otherwise in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Tuning into Scorpius X-1: adapting a continuous gravitational-wave search for a known binary system

Grant David Meadors¹, Evan Goetz², Keith Riles²

¹Max-Planck-Institut für Gravitationsphysik, Am Mühlenberg 1, 14476 Potsdam and Callinstraße 38, 30167 Hannover, Germany

²University of Michigan, 450 Church Street, Ann Arbor, Michigan 48109, USA

E-mail: grant.meadors@aei.mpg.de

Abstract. We describe how the TwoSpect data analysis method for continuous gravitational waves (GWs) has been tuned for directed sources such as the Low Mass X-ray Binary (LMXB), Scorpius X-1 (Sco X-1). A comparison of five search algorithms generated simulations of the orbital and GW parameters of Sco X-1. Where that comparison focused on relative performance, here the simulations help quantify the sensitivity enhancement and parameter estimation abilities of this directed method, derived from an all-sky search for unknown sources, using doubly Fourier-transformed data. Sensitivity is shown to be enhanced when the source sky location and period are known, because we can run a fully-templated search, bypassing the all-sky hierarchical stage using an incoherent harmonic sum. The GW strain and frequency as well as the projected semi-major axis of the binary system are recovered and uncertainty estimated, for simulated signals that are detected. Upper limits on GW strain are set for undetected signals. Applications to future GW observatory data are discussed. Robust against spin-wandering and computationally tractable despite unknown frequency, this directed search is an important new tool for finding gravitational signals from LMXBs.

PACS numbers: 04.30.-w, 04.30.Tv, 04.40.Dg, 95.30.Sf., 95.75.Pq, 95.85.Sz, 97.60.Jd

1. Introduction

Continuous gravitational waves (GWs) from neutron stars in binary systems seem likely to be one of the most interesting types detectable by ground-based interferometric observatories. Binary systems constitute 237 of 578 (44%) of known pulsars in the ATNF catalog (v1.53, 2015) [1] with rotational frequency faster than 5 Hz, the approximate lower bound of the frequency range of Advanced LIGO and Advanced Virgo [2, 3]. TwoSpect [4, 5] is one data analysis method for targeting systems of Low Mass X-ray Binaries (LMXBs), including the prototypical Scorpius X-1 (Sco X-1). This method can apply to continuous-wave (CW) GWs of unknown frequencies from neutron stars in binary systems, with unknown sky locations, orbital periods, or projected semi-major axes. It was used in a prior all-sky search for CWs from unknown sources with data

from LIGO Science Run 6 and Virgo Science Runs 2 & 3 [6]. This paper describes analysis modifications, when sky location and orbital period are known, to run a *directed* search. The relative performance of this method was shown in a comparison of five algorithms [7]. More sensitive than the all-sky search, the new directed methodology is presented here, including discussion of detection criteria, parameter estimation, and upper limit strategy. Now, in the wake of the first GW observation [8], we soon hope to observe LMXBs such as Sco X-1.

CW analyses for isolated neutron stars are computationally-demanding [9] but conceptually simple. Ellipticity in a rotating star, or stellar r -modes [10, 11], would generate a time-varying mass quadrupole moment that could induce CW emission. When these emissions, with a dimensionless strain amplitude h_0 and a phase evolution described in Section 2, arrive at a GW observatory, they may be very faint. CW analyses include the \mathcal{F} -statistic, Hough, StackSlide and PowerFlux methods [12–17]. Many stars, possibly the best GW sources, are unknown or have ephemerides insufficiently precise to make a fully-coherent search tractable. Different strategies are therefore used depending on available information.

CW searches are categorized as *all-sky* (for unknown objects), *directed* (sky location known) and *targeted* (spin frequency also known). When sky location and other ephemerides are known, computational cost can often be reduced or reinvested in increased sensitivity. As recently reviewed [18], directed [19, 20] and targeted [21, 22] searches now exist.

Signals from neutron stars with rotational periods of milliseconds should populate the GW spectrum. Millisecond pulsars appear to have a speed limit somewhat higher than 700 Hz but below their expected relativistic break-up speed [23]: GW emission is a possible cause. Dedicated analyses are motivated by the large fraction of these pulsars that are in binary systems.

Binary systems intrinsically have more parameters to analyze than isolated stars. TwoSpect [5] uses doubly Fourier-transformed data in the first practical all-sky search for unknown neutron stars in binary systems [6]. Other binary searches include developments of the Sideband [24, 25], Radiometer [26, 27], Polynomial [28], and CrossCorr [29, 30]. A systematic comparison [7] challenged these five methods to detect simulated signals from Sco X-1. A stacked \mathcal{F} -statistic, derived from the coherent \mathcal{F} -statistic [31], is under investigation [32].

LMXBs, including Sco X-1, are believed to *spin-up* (increase rotational frequency) by accretion-driven recycling [33]. Accretion can also lead to non-axisymmetry that induces emission. These mechanisms suggest a prime GW source. In the *torque balance* hypothesis, spin-up would continue until it equaled and canceled *spin-down* from GW emission [34]. This hypothesis can be quantified in terms of GW strain.

Torque balance predicts a characteristic strain h_c given by Equation 1 (Equation 4 of Bildsten [35]; note $h_c/h_0 = 2.9/4.0$). For an LMXB with flux $\mathcal{F}_{X\text{-ray}}$ and NS that

Tuning into Scorpius X-1

3

rotates at a frequency ν_s , radiating at the quadrupolar GW frequency $f = 2\nu_s$,

$$h_c \approx 4 \times 10^{-27} \left(\frac{300 \text{ Hz}}{\nu_s} \right)^{1/2} \left(\frac{\mathcal{F}_{\text{X-ray}}}{10^{-8} \text{ erg cm}^{-2} \text{ s}^{-1}} \right)^{1/2}. \quad (1)$$

With X-ray flux ($3.9 \times 10^{-7} \text{ erg cm}^{-2} \text{ s}^{-1}$ [36]) but unknown frequency for Sco X-1, this limit can be evaluated for f between 50 and 1500 Hz [7], although large systematic uncertainties mean h_0 could be larger or smaller:

$$\begin{aligned} h_0 &= \frac{4.0}{2.9} h_c, \\ &\approx 3.5 \times 10^{-26} \left(\frac{600 \text{ Hz}}{f} \right)^{1/2}, \\ &\rightarrow h_0 < 2.2 \times 10^{-26} \text{ (} f = 1500 \text{ Hz)}, \quad h_0 < 1.2 \times 10^{-25} \text{ (} f = 50 \text{ Hz)}. \end{aligned} \quad (2)$$

Angular momentum from accretion would counterbalance that lost to GW emission. If true, f would remain stable for long durations, aside from spin-wandering due to variations in accretion. Computational costs are thus reduced by the restricted range of the frequency-derivative.

The torque-balance levels predicted by Equation 2 are near the estimated thresholds of detection, meriting effort in enhancing the sensitivity of our methods. Directed search improvements skip the computation-saving hierarchical steps of the all-sky search, testing all points in parameter space with templates for enhanced sensitivity to a signal from a known sky location. This paper details the directed, fully-templated method and shows its application to simulated data containing Sco X-1 signals.

2. Signal model

GW signals are defined by $h(t)$, strain as a function of time. Models of $h(t)$ in a binary search [5] depend on phase evolution, $\Phi(t)$,

$$\Phi(t) = \Phi_0 + 2\pi f_0 \cdot \tau(t) + \Delta f_{\text{obs}} \cdot P \cdot \sin(\Omega[t - T_{\text{asc}}]), \quad (3)$$

$$h(t) = h_0 F_+ \frac{1 + \cos^2 \iota}{2} \cos \Phi(t) + h_0 F_\times \cos \iota \sin \Phi(t). \quad (4)$$

Here solar-system barycentered time is $\tau(t)$ and $\Phi(t = 0) \equiv \Phi_0$. We also neglect spin-wandering, which could manifest as stochastic variation in Φ_0 .

The TwoSpect model does not currently search over *amplitude parameters*: h_0 is the GW strain amplitude, ι is the inclination angle of the neutron star with respect to the source, ψ is the GW polarization angle. While h_0 can be recovered, its estimation is confounded by ι . Both ι and ψ affect amplitude through detector response, which depends on angle via the plus- and cross antenna functions, F_+ and F_\times . Initial GW phase, Φ_0 , further specifies the signal, but is neither explored nor recoverable.

Our search is over unknown *Doppler parameters*, which drive signal evolution. Sky location (α, δ) is known for Sco X-1, as is orbital period $P = 2\pi/\Omega$. Initial orbital

phase is fixed by time of ascension T_{asc} , to which our method is currently insensitive. Frequency f and projected semi-major axis $a \sin i$ (expressed in light-seconds, ls) must be searched over. Because the signal is frequency-modulated by the orbital motion of the source in the circular binary system, the latter parameter manifests through modulation depth $\Delta f_{\text{obs}} = 2\pi f_0 \cdot (a \sin i)/(cP)$.

3. Analysis statistic

TwoSpect has already been described for the all-sky analysis [5, 6]. Only a brief summary is given here on the elements common to the directed analysis, with details in Appendix A.

Data containing calibrated GW strain, originally recorded by each observatory as time series, are read into the program from a sequence of short Fourier transforms (SFTs). These SFTs, on the order of minutes long, are shorter than the total observation time T_{obs} , on the order of months. Noise and antenna pattern weights are applied to enhance sensitivity by weighting those SFTs that are more sensitive to a putative source. Each SFT contains K frequency bins. The Earth's motion Doppler shifts the apparent frequency of a source, so these bins must be barycentered: their indices are shifted such that an unmodulated frequency from a given sky location remains in the same bin for all SFTs. These frequency bins represent the instantaneous frequency f of the signal. A second Fourier transform is performed, transforming the power in each frequency bin: the transform is from SFT time to f' . Effectively, f' represents the orbital frequency of the binary system.

Templates are the M pixel weights w_i in the 2-D image-plane (f, f') . Each template corresponds to a signal model from the $(f_0, \Delta f_{\text{obs}})$ astrophysical parameter space. After the second Fourier transform, the powers Z_i of data pixels in the (f, f') are measured and noise-background λ_i estimated. The test statistic, R , is the projection of the data (noise-subtracted powers) onto the templates, normalized by the templates:

$$R = \frac{\sum_{i=0}^{M-1} w_i [Z_i - \lambda_i]}{\sum_{i=0}^{M-1} [w_i]^2}. \quad (5)$$

Estimated p -values for this statistic allow determination of detection probability. Expositions of the detection statistic and signal model can be found in previous methods papers [5, 37].

4. Application to directed searches

4.1. Sensitivity and computational cost

The original design of the all-sky analysis employs hierarchical analysis to control computational costs while still maintaining good sensitivity to a broad parameter space [5]. Each narrow frequency band requires corrections for the Doppler effect caused by Earth's motion. Because this correction depends on sky position, an all-sky search

Table 1. Sco X-1 prior measured parameters from electromagnetic observations. Note that the projected semi-major axis value is derived from a velocity amplitude of $K_1 = 40 \pm 5 \text{ km s}^{-1}$ [7, 31].

Sco X-1 parameter	Value	Units
Distance [40]	2.8 ± 0.3	kpc
Eccentricity (ϵ) [7]	$< 0.068 (3\sigma)$	—
Right ascension (α) [41]	$16:19:55.067 \pm 0.06''$	—
Declination (δ) [41]	$-15^\circ 38' 25.02'' \pm 0.06''$	—
X-ray flux at Earth ($\mathcal{F}_{\text{X-ray}}$) [36]	3.9×10^{-7}	$\text{erg cm}^{-2} \text{ s}^{-1}$
Orbital period (P) [38]	68023.70 ± 0.04	s
Projected semi-major axis ($a \sin i$) [42]	1.44 ± 0.18	ls

for typical SFT lengths of several minutes would need $\mathcal{O}(10^{18})$ templates; in practice, an incoherent harmonic sum is used to reduce this number, only calculating templates for candidates with significant sums. This practice limits sensitivity [5]. The new, directed search in this paper can use $\mathcal{O}(10^8)$ templates at a single sky location with full sensitivity, because R -statistics are returned for every template.

For Sco X-1, distance, eccentricity, X-ray luminosity, sky location, and orbital period are known with good precision and projected semi-major axis is measured to be 1.44 ± 0.18 ls, (Table 1). NS spin frequency, however, is unknown [38]. Other targets, such as XTE J1751-305 [39], have known frequency, reducing computational costs substantially. Sco X-1, XTE J1751-305, and other LMXBs are the principal sources for the new, directed method, described next.

Searching over projected semi-major axis in the range $[-3\sigma_{a \sin i}, +3\sigma_{a \sin i}]$, around the measured value, and over GW signal frequency f from f_{\min} to f_{\max} , incurs a predictable computational cost. With a fixed rectangular parameter spacing of $1/(2T_{\text{coh}})$ in f and $1/(4T_{\text{coh}})$ in Δf_{obs} , using SFTs of coherence time T_{coh} and analysis bands of width f_{bw} , the search uses N_{template} templates per observatory (Appendix A):

$$N_{\text{template}} = [1 + 2f_{\text{bw}}T_{\text{coh}}] \times \sum_{j=1}^{j=\frac{f_{\max}-f_{\min}}{f_{\text{bw}}}} \left[1 + 2\pi (f_{\min} + jf_{\text{bw}}) \frac{4T_{\text{coh}}}{cP} 6\sigma_{a \sin i} \right], \quad (6a)$$

$$= 2 \left(T_{\text{coh}} + \frac{1}{f_{\text{bw}}} \right) \left[1 + \frac{4\pi T_{\text{coh}}}{cP} (6\sigma_{a \sin i})(f_{\max} + f_{\min} + f_{\text{bw}}) \right] (f_{\max} - f_{\min}). \quad (6b)$$

4.2. Participation in a Sco X-1 mock data challenge

Several analyses can search for CWs from known neutron stars in binary systems. TwoSpect can also seek unknown systems‡. The Sco X-1 Mock Data Challenge (MDC) compares five methods [7], including the new, directed version of TwoSpect. While the MDC published the results of the participants, this paper explains how our results were obtained and how similar analyses apply to forthcoming observations.

‡ Other all-sky binary search programs are based on the Polynomial [28] and Radiometer [27] methods.

The MDC simulates 50 *open*, unblinded and 50 *closed*, blinded signals, with signal parameters drawn from Sco X-1 astrophysical priors. Each signal is also known as an *injection*. Three observatories are simulated: H1 (Hanford), L1 (Livingston), and V1 (Virgo). Within 100 frequency bands of 5 Hz each, f is uniformly distributed. Bands cover a range from 50 to 1500 Hz and contain Gaussian noise simulating a detector with noise floor $4 \times 10^{-24} \text{ Hz}^{-1/2}$, the expected Advanced LIGO minimum. Projected semi-major axis $a \sin i$ is Gaussian-distributed with mean 1.44 ls and standard deviation 0.18 ls. Period P is Gaussian-distributed with mean 68023.70 s and standard deviation 0.04 s. Sky location is $(\alpha, \delta) = (16:19:55.067, -15^\circ 38' 25.02'')$. The uncertainty in period is small enough that it is not a free parameter.

Our directed search over f and $a \sin i$, with a fixed period, requires approximately $N_{\text{template}} = 5.0 \times 10^7$ templates per observatory, covering 100 bands of 5 Hz each, 500 Hz total. With three observatories, 1.5×10^8 templates are needed.

Figure 1 shows a wide analysis that is magnified in Figure 2. These analyses require minimal modification from that of the all-sky search: R -statistic and p -value code was identical, aside from looping over all template positions in the parameter grid.

5. Directed search demonstration and outlier follow-up

The new, directed method is demonstrated through the MDC described in Section 4. This MDC paper [7] presents the performance of each participating pipeline but not the details of our method. Here we highlight methodology not elaborated in the MDC paper, referencing the performance of the directed method for illustration and for contrast with the all-sky method [5].

After writing a loop over templates, post-processing is the main modification to the pipeline: new detection criteria and techniques for parameter and upper limit estimation are required, because the fully-templated output differs significantly from that of hierarchical method. In this section, we detail how these results have been obtained and establish a basis for future applications.

5.1. Overview of detection and parameter estimation

A set of extremal p -value outliers in 5 Hz bands is produced for each observatory, subject to a p -value threshold inferred from Gaussian noise. These sets are compared in pairwise coincidence (H1-L1, H1-V1, or L1-V1). Coincidence of outliers allows f or Δf_{obs} to differ by $\leq 1/T_{\text{coh}}$ between observatories, based on prior experience from the all-sky search [6]. This allowed difference is 2 steps in the f grid or 4 in the Δf_{obs} grid. Surviving outliers are classified as detections.

For a given detection in one band, the signal parameters are inferred from the values of the template with the highest (single-observatory) p -value. These parameters include f , $a \sin i$, and h_0 . Again, for open signals, the true parameters were known in advance. Total uncertainty in f and $a \sin i$, and non-systematic uncertainty (random)

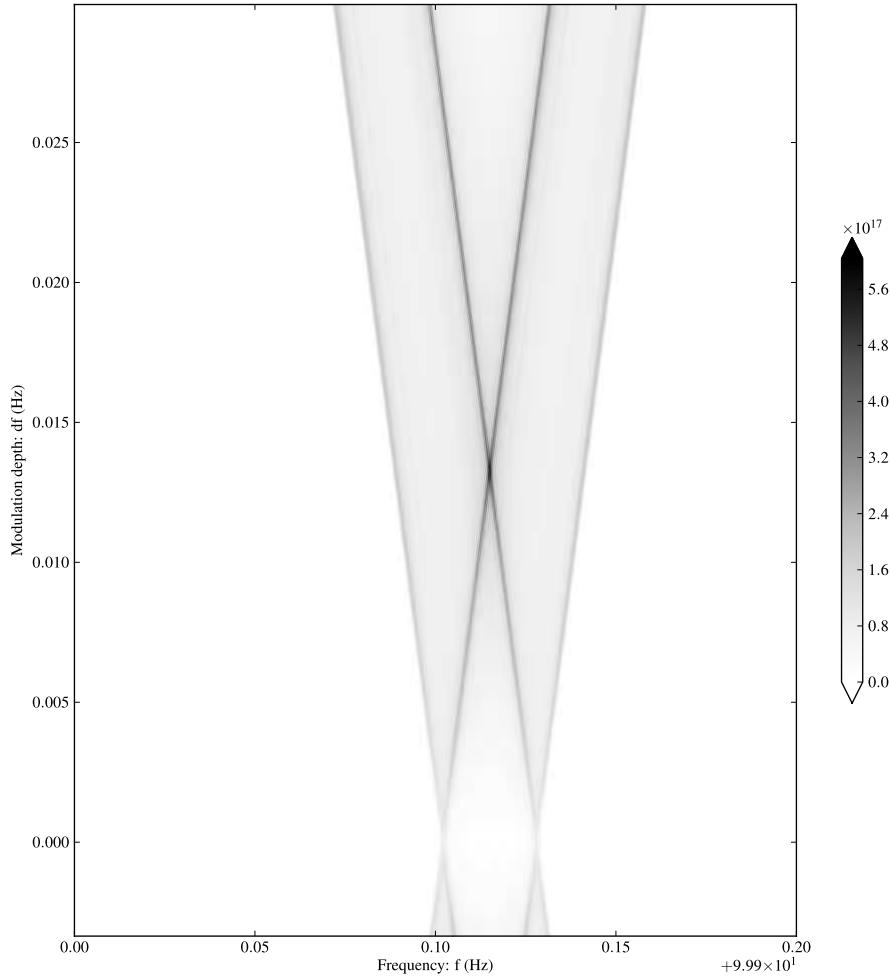


Figure 1. Simulated signal, 100.015 Hz frequency (f), 1.44 ls projected semi-major axis (modulation depth 0.0133 Hz, Δf_{obs}), showing single-template $-\log_{10} p$ -value. The graph spans 241 templates in Δf_{obs} , 721 in f , twice as dense in Δf_{obs} as in f . A signal with $h_0 = 4 \times 10^{-21}$ is injected into a Gaussian noise amplitude spectral density of $4 \times 10^{-24} \text{ Hz}^{-1/2}$, observed for 10^6 seconds. The p -value is extrapolated from the R statistic. For illustration, Δf_{obs} extends below zero to show that the code is well-behaved and that the algorithm gracefully mirrors results for negative input parameters. Most importantly, $-\log_{10} p$ is maximal at the true parameters. Signal is partly recovered when template f differs from the true signal by plus or minus Δf_{obs} , and also when template f and Δf_{obs} differ equally from the true parameters.

in h_0 , is determined from the standard deviation of the set of parameters of recovered open signals compared to their true parameters, as detailed in section 5.3.

The estimated strain is $h_{\text{rec}} \propto R^{1/4}$. Systematic uncertainty arises from the unknown inclination angle, $\cos \iota$, which dominates the total uncertainty in strain. This ambiguity cannot be resolved with the present algorithm and depends partially on the assumed prior distribution of signal amplitudes; the uncertainty is estimated by simulation in Section 5.4. With new algorithm enhancements since the MDC, $\cos \iota$ can become a searched parameter.




Figure 2. Heatmaps for H1, L1, and V1 observatory (top to bottom), of $-\log_{10} p$ -value for 11x11 templates centered on Sco X-1 MDC signal 8. This injection was detectable in a year of simulated data at $h_0 = 5.6 \times 10^{-25}$ in noise of $4 \times 10^{-24} \text{ Hz}^{-1/2}$ and $\cos \iota = 0.09$. Maxima for all observatories are within a template of the true parameters.

Due to the uniform noise floor and low number of injections, a single upper limit value is declared for all 100 bands based on the best estimate of the 95% confidence level of non-detected signals in the open set of injections.

In future applications to real data, the directed analysis can be post-processed using the detection criteria and parameter estimation methods described here. Improved upper limits methods are underway, and a strong candidate signal would likely be followed-up with other analyses, but the core pipeline is the same.

5.2. Detection claims

Studying the Gaussian noise in the Sco X-1 MDC open data set, we can set thresholds for detection. In real searches, detection candidates are followed-up, so the same thresholds developed here are used to mark interesting candidates.

To obtain a Gaussian noise sample, we excise injection signals, which are visible in the f vs Δf_{obs} plane. The excised region depends on injection frequency f_{inj} and modulation depth Δf_{inj} , on the Earth-orbital Doppler shift ($\approx 10^{-4}$), and additional bins (at least 10) to avoid spectral leakage. The half-width excised for each injection is δf_{inj} :

$$\delta f_{\text{inj}} = 2 \times (10^{-4} \times f_{\text{inj}} + \Delta f_{\text{inj}}) + \frac{10}{360} \text{Hz}. \quad (7)$$

The remaining noise sample is the data set with all intervals $[f_{\text{inj}} - \delta f_{\text{inj}}, f_{\text{inj}} + \delta f_{\text{inj}}]$ removed. In the remaining noise sample, the estimated p -value distribution is not perfectly uniform, due to gaps in the data. Nonetheless, the noise sample provides a distribution of template R -statistics and p -values in the absence of signals. This procedure provides an empirical measure of the estimated p -value that corresponds to an actual false alarm probability of 1% per 5 Hz frequency band. Taken together, these let us establish detection criteria.

If there is any candidate surviving the following criteria in a 5 Hz band, we mark it detected, else not detected:

- single-IFO candidates are the top 200 most extreme p -value outliers in a 5-Hz band, of those that pass a $\log_{10} p \leq \mathcal{T}$ threshold, where $\mathcal{T} = -7.75$ if $f < 360.0$ Hz (those that used 840-s SFTs) or -12.0 if $f \geq 360.0$ Hz (those that used 360-s SFTs). Note: The large discrepancy between the p -value thresholds in the MDC is a historical artifact from a configuration error. The discrepancy is much reduced when this is fixed, as done for future analyses. Our expectation remains that the threshold should be independent of coherence time.
- each candidate must survive at least one double-IFO coincidence test, involving a pairwise comparison of single-IFO candidates to see whether they are within $1/T_{\text{SFT}}$ in both frequency (f) and modulation depth (Δf_{obs}).

5.3. Parameter estimation and uncertainty for detected signals

Each template is associated with a particular $(f, \Delta f_{\text{obs}})$, so parameters are currently read off from the template with the extremal p -value corresponding to a detection. In the future, accuracy might be improved using interpolation, but the MDC validates that the existing method is highly-accurate. If signals were suspected in real data, this procedure, possibly extended with additional simulations, could generate a parameter space volume for follow-ups to examine.

The open set of signals, of which 31 of 50 were detected, are the foundation for understanding parameter estimation uncertainty. The reconstructed h_0 output is

Tuning into Scorpius X-1

10

$h_{\text{rec}} = CR^{1/4}$. The value of C is determined from the mean value of a large number of simulations for circularly-polarized waves over the whole sky and full range of f , P , and Δf_{obs} .

Then h_{rec} is rescaled twice, first by ρ_R for more accurate measurement at the Sco X-1 period and modulation depth, and second by $\rho_{\cos \iota}$ for unknown $\cos \iota$. Thus the final claimed value of h_0 for a signal j is $(h_0)_j$:

$$(h_0)_j = \rho_R \rho_{\cos \iota} h_{\text{rec},j}. \quad (8)$$

The first scale factor, $\rho_R = 1.11$, corrects the average values of h_{rec} (h_0 -reconstructed) in the open set to match the corresponding h_{eff} (h_0 -effective, given circular polarization weightings). That is, $\bar{h}_{\text{eff}} = 1.11 \times \bar{h}_{\text{rec}}$, where h_{eff} is defined *a priori* by h_{inj} (h_0 -injected) and $\cos \iota$:

$$h_{\text{eff}} = \frac{1}{\sqrt{2}} \sqrt{\left(\frac{1 + \cos^2 \iota}{2}\right)^2 + (\cos \iota)^2} \times h_{\text{inj}}. \quad (9)$$

In the MDC study [7], 4 of the 31 detected, open signals account for the largest frequency estimation error. This error arose from a misconfiguration that does not affect the other analyses; it was addressed by taking those 4 as one class and the remaining 27 as another, a step that should be unnecessary in future analyses. Then, the uncertainty due to random error for h_0 , f , and $a \sin i$ is estimated by the standard deviation between the recovered and true parameters:

$$\sigma_f^2 = \frac{1}{N_{\text{open}} - 1} \sum_{j \in \{\text{open}\}}^{N_{\text{open}}} (f_{\text{rec},j} - f_{\text{inj},j})^2, \quad (10)$$

$$\sigma_{a \sin i}^2 = \frac{1}{N_{\text{open}} - 1} \sum_{j \in \{\text{open}\}}^{N_{\text{open}}} ((a \sin i)_{\text{rec},j} - (a \sin i)_{\text{inj},j})^2, \quad (11)$$

$$\sigma_{h_0, \text{rand}}^2 = \frac{1}{N_{\text{open}} - 1} \sum_{j \in \{\text{open}\}}^{N_{\text{open}}} (\rho_R \times h_{\text{rec},j} - h_{\text{eff},j})^2, \quad (12)$$

where N_{open} is the number of open injections, σ_f , $\sigma_{a \sin i}$, and σ_{h_0} are the uncertainties we state for recovered f_{rec} , $(a \sin i)_{\text{rec}}$, and h_{rec} given injected f_{inj} , $(a \sin i)_{\text{inj}}$, and h_{inj} .

The error between injected and recovered parameters does not show any other clear correlation with p -value or signal frequency, at least in the 31 detected signals. Except for the most marginally detected signals, where noise fluctuations matter, uncertainty in f and $a \sin i$ is dominated by the template grid spacing. The σ_f and $\sigma_{a \sin i}$ error bars have been used uniformly for claiming uncertainties on the signals in the MDC.

The largest source of uncertainty for h_0 comes from correction for systematic underestimation, multiplying a factor of 1.74 into C . This uncertainty is the ambiguity in $\cos \iota$ discussed in Section 5.4. Parameter estimation uncertainty for f and $a \sin i$ is then just the random error; for h_0 , it is the quadrature sum of random error and $\cos \iota$ ambiguity.

5.4. Ambiguity from $\cos \iota$

The largest systematic uncertainty in h_0 comes from the unknown $\cos \iota$. The method is optimized for $|\cos \iota| = 1$ and computes the R statistic by weighting the SFTs assuming circularly-polarized GWs, which still provides good sensitivity for other polarizations. Recall that h_{rec} must be scaled by $\rho_R = 1.11$ to match h_{eff} . When a source is circularly polarized, the analysis estimates $\rho_R h_{\text{rec}} \approx h_{\text{inj}}$. In the case of linear polarization, Equation 9 indicates that h_{inj} will be about $2^{3/2}$ times larger than h_{eff} (and so h_{rec}). The aim is to find an average conversion factor from h_{rec} to h_{inj} and a robust estimate of the uncertainty.

Since h_{eff} of circularly-polarized signals is greater, those signals are more easily detectable than linearly-polarized signals of equivalent h_{inj} . Therefore the signals that are detectable are biased, near threshold, to being more likely circularly polarized. This “circularizes” the correction factor, depending on the detection efficiency of the pipeline and on the assumed prior distribution of strain amplitudes. Although the effect is minor, estimating its size requires simulation.

The simulation generates 2 million signal amplitudes between $h_{\text{inj}} = 3 \times 10^{-26}$ and $h_{\text{inj}} = 3 \times 10^{-24}$ with a distribution of $1/h_0$, the assumed prior distribution of h_0 values. This simulation code is independent of TwoSpect and should apply to similar directed searches. In this simulation, ρ_R is 1 for simplicity, so we can treat $h_{\text{rec}} = h_{\text{eff}}$ here. We model detection efficiency by assuming no signals are detected below $h_{\text{eff}} = 1 \times 10^{-25}$, all are detected above $h_{\text{eff}} = 3 \times 10^{-25}$, and the fraction detected is linear in h_0 between those values. Together with a uniform $\cos \iota$ distribution of $[-1, 1]$, this leads to a trapezoidal distribution of recovered, detected h_0 values with a curved lower (left) edge (Figure 3).

Part of the domain of the simulation must be excluded. To find the average $h_{\text{inj}}/h_{\text{rec}}$ for a given h_{rec} , every h_{rec} must correspond to a full sampling of the range of polarizations. A large h_{inj} with linear polarization or small h_{inj} with circular polarization could have the same h_{rec} . No linearly-polarized signals could produce h_{rec} above 1×10^{-24} , because Equation 9 shows that the largest signal, $h_{\text{inj}} = 3 \times 10^{-24}$, would be reconstructed a factor of $2\sqrt{2}$ smaller, at $h_{\text{rec}} = 1.06 \times 10^{-24}$ (again, $\rho_R = 1$ for the simulation). Above 1×10^{-24} , polarizations tend to be more circular, thus the average ratio must exclude this region or it will be biased by the limited range of the simulation h_{inj} . Expanding the domain would raise the cutoff, although the resulting ratio would no longer perfectly correspond to the MDC.

With the domain of the simulation determined, we compute that mean ratio of h_{inj} to h_{rec} to be 1.74. Fine-binning h_{rec} , an interval of $[1.74 - \sigma_{\cos \iota}, 1.74 + \sigma_{\cos \iota}]$ encloses 68% of corresponding h_{inj} when $\sigma_{\cos \iota}$ is 0.37 (found by manual optimization). Therefore our best estimate for the correction factor $\rho_{\cos \iota}$, with $\sigma_{\cos \iota}$ inferred as the standard deviation, is 1.74 ± 0.37 . This factor multiplies ρ_R , which is found to be 1.11.

The systematic uncertainty, being the uncertainty $\sigma_{\cos \iota}$ in the correction factor, scales with signal strength; the non-systematic (random) is fixed and is also multiplied by the correction factor. The final estimate of the uncertainty in h_0 for signal j is the




Figure 3. Histogram, 150 bins, of the distribution of 2 million simulated signals, strain between 3×10^{-26} and 3×10^{-24} under a log-uniform distribution, following application of $\cos \iota$ and detection efficiency cuts.

quadrature sum of the systematic and non-systematic uncertainties:

$$\sigma_{(h_0),j} = \sqrt{(\rho_{\cos \iota} \times \sigma_{h_0,\text{rand}})^2 + (\sigma_{\cos \iota} \times \rho_R \times h_{\text{rec},j})^2}. \quad (13)$$

For future data, a similar simulation could be run, with an updated detection efficiency model and prior distribution of strains, to find the uncertainty in h_0 due to $\cos \iota$ for promising signals.

5.5. Accuracy of parameter estimation uncertainty claims

The scheme described above reliably recovers parameters and states uncertainties consistent with the true distribution of errors, as shown in the MDC [7]. Verifying the calibration factors and confidence intervals once more, one can confirm that a conservative fraction of h_0 , f and $a \sin i$ are within their $1\text{-}\sigma$ error bars: 77.4% for σ_{h_0} , 74.2% for σ_f , and 67.7% for $\sigma_{a \sin i}$.

5.6. Upper limits for undetected signals and detection efficiency

Only upper limits have come from CW searches to date. Until GWs are detected from neutron stars, the scientific value of a CW analysis is to constrain the plausible h_0 and inferred ellipticity from those stars. For the MDC-simulated signals, we simplified upper-limit estimation. In a real detector, noise varies with frequency [43]; in this simulation, the noise floor is flat at $4 \times 10^{-24} \text{ Hz}^{-1/2}$. Given observational data, we would inject a large number of simulated signals into a number of smaller bands, in order to understand the upper limit as a function of frequency.

To measure detection efficiency, we calculate the h_{eff} for all signals and find the average detection rate for a given h_{eff} . Binomial uncertainty is also calculated and each $1\text{-}\sigma$ deviation (per 5-signal bin) is graphed in Figure 4, which shows a least-squares sigmoid fit. This detection efficiency curve maps from strain to probability of detection. Next we would like an upper limit function that takes a probability as an input and returns a strain that, with the given probability, is no less than the actual strain.

To characterize the upper limit, we plot the distribution of h_{rec} versus injected h_{eff} in Figure 5. We verify that 95% of non-detected open signals are covered by a naive upper limit of $h_{\text{rec}} = 2.19 \times 10^{-25}$. This claim does not rely on binning but rather on the sampling density of the injections. This number, when corrected by the ρ_R rescaling factor of 1.11 and $\cos \iota$ correction $\rho_{\cos \iota}$ of 1.74, yields the upper limit of $(1.74) \times (1.11) \times 2.19 \times 10^{-25} = 4.23 \times 10^{-25}$. Because of the flat noise floor, this is reported as a single upper limit for any non-detections in the MDC.

Frequency-dependent upper limits are well under development for actual observations. Sigmoid fits to the detection efficiency of a set of injections into real data are generated, one fit per frequency band. Upper limits at a given confidence can then be taken as the h_0 that yields a detection efficiency equal to that confidence. This advancement post-dates the MDC and is planned for future applications.

6. Conclusions

6.1. MDC results

TwoSpect analyses applied to the MDC data set detect more stars than the Radiometer, Sideband, or Polynomial pipelines; only the CrossCorr algorithm found more signals [7]. Each detection includes an estimate of $a \sin i$, which is not produced by Radiometer, Sideband, or Polynomial. The MDC did not model the spin-wandering of the neutron star that is expected in real data, although participants were told to assume its presence, and spin-wandering is planned for future MDCs. TwoSpect is also theoretically highly robust against spin-wandering. This method has already been applied to real data [6], though not using the directed search in a fully-templated mode. This experience validates the program's robustness with respect to non-Gaussian data artifacts. In all, 34 of 50 closed (and 31 of 50 open) signals are detected, and f , $a \sin i$, and h_0 are estimated. Strain upper limits of 4.23×10^{-25} noise of in 4×10^{-24} strain $\text{Hz}^{1/2}$ are




Figure 4. Detection efficiency for open and closed signals. Because only 100 signals are in the data set, this curve has fluctuations and large $1\text{-}\sigma$ error bars. Moreover, the binomial error for these error bars is zero in bins where no or all signals are detected, which is not necessarily realistic. The 95% level is just below 2×10^{-25} (before rescaling factors of 1.74 and 1.11), corroborating better techniques of estimating the upper limit.

determined for the 16 non-detected, closed signals. Although the distribution of h_0 values in the MDC was astrophysically optimistic, the MDC validated our ability to claim detections and recover orbital and GW parameters accurately.

6.2. Future directed CW binary searches

Algorithms such as TwoSpect are designed to find astrophysically-plausible strain from LMXBs. Torque-balance arguments suggest that strain could exceed 10^{-25} for Sco X-1 if it rotates at low frequencies.

The previously-published all-sky search in a year of S6 data set an overall upper




Figure 5. Detections and upper limit determination, with all 100 simulated signals. Injections are seen in three (black circle), one (gray triangle), or zero (open box) observatory pairs and plotted in recovered strain versus effective circular strain injected. The most significant template is shown, in each band, in each observatory, whether detected or not. (There are no injections seen with two detection pairs, because this plot shows only the loudest outlier from each 5 Hz band; if some injection were seen in two and not three pairs, it would mean two distinct coincidences were seen, only one of which would be the loudest). We identify a shelf of non-detected signals that are 95% covered by an upper limit about 2.19×10^{-25} . This number, when corrected, yielded the upper limit of $1.74 \cdot 1.11 \cdot 2.19 \times 10^{-25} = 4.23 \times 10^{-25}$. The unity-slope black line is shown to ascertain whether a further empirical rescaling factor is needed to match the dashed-and-dotted least-squares linear fit (it is: constant 1.11). The zero-slope horizontal dashed line is shown to indicate the ninety-five percent confidence upper limit in the absence of detection.

limit for circular polarization of 2.3×10^{-24} at 217 Hz [6], set for an H1 amplitude spectral density of $2.0 \times 10^{-23} \text{ Hz}^{-1/2}$. For random polarization, a multiplicative factor of 2.2 was applied, for an upper limit of 5.1×10^{-24} . This corresponds to a sensitivity depth (factor below the noise floor: Appendix A) of $8.7 \text{ Hz}^{-1/2}$ for circular polarization but $4.0 \text{ Hz}^{-1/2}$ for random. Extrapolating to Advanced LIGO design sensitivities of $4 \times 10^{-24} \text{ Hz}^{-1/2}$, this implies an upper limit around 5×10^{-25} for circular polarization and 1×10^{-24} for random. However, the all-sky paper included an opportunistic search for Sco X-1 on a narrow frequency range (20 to 57.25 Hz), setting a random polarization limit of 2×10^{-23} at 57 Hz in an L1 amplitude spectral density $1.8 \times 10^{-22} \text{ Hz}^{-1/2}$, a depth of $9 \text{ Hz}^{-1/2}$. This opportunistic search used 1800-s SFTs; longer SFT durations increase theoretical sensitivity. Significant improvement comes from focusing on one sky location; this can be viewed as a reduced trials factor. The directed search demonstrated in this paper achieved an 4.23×10^{-25} upper limit in simulated data at design sensitivity, a depth of $9.5 \text{ Hz}^{-1/2}$. It achieved this depth despite being tested with shorter (360-s and 840-s) SFTs. The directed search is also scalable over a much wider parameter range, like the all-sky method over which it gains twofold in sensitivity.

While real data complications may worsen this limit, several simplified and conservative steps were taken. The limit may improve with the enhancements now under development, when fully tested with injections as in the all-sky search. Even now, the method in this paper is more sensitive for random polarization than the all-sky method is for optimal, circular polarization. Additional improvements to the algorithm, such as coherent SFT summing, have been developed [37] and could further improve this limit in the future, pushing toward the torque-balance strain.

Directed TwoSpect analyses have been demonstrated in this paper. Comprehensively covering the parameter space of Sco X-1 at full sensitivity with the directed search, instead of hierarchically as before, does increase the probability of detection and improve upper limits. When detections do occur, the ability to determine the frequency and projected semi-major axis of the neutron star in the binary system will prove highly informative. Analyses of real data for signals from Sco X-1 and additional neutron stars in binary systems, such as XTE J1751-305, are underway. In the long term, we hope that the discovery of gravitational waves from neutron stars in LMXBs will provide a firm link between our observations and electromagnetic astronomy.

Acknowledgments

This work was partly funded by National Science Foundation grants NSF PHY 1205173 and NSF PHY 1505932. These investigations have taken place as part of the LIGO Scientific Collaboration. The Mock Data Challenge for Sco X-1 was organized by Chris Messenger. Thanks to Maria Alessandra Papa for providing extensive guidance and support, as well as to John Whelan for thorough proofreading, along with our referees for their helpful comments. Code for this paper is available online [44]. This document bears LIGO DCC number P1500037.

Appendix A. Mathematical details

Appendix A.1. Sensitivity depth

The relative sensitivity of a search can be quantified in terms of sensitivity depth $D(f)$ [32, 45],

$$D(f) \equiv \frac{S_H^{1/2}(f)}{h_0(f)}, \quad (\text{A.1})$$

for noise power spectral density S_H at frequency f and the GW strain $h_0(f)$ recoverable there. The depth generally depends on the observation time, because integrated signal-to-noise grows. This concept allows us to compare methods across data sets and extrapolate future performance.

Appendix A.2. Number of templates

Equation 6b is the integrated template density over parameter space dimensions. While we are interested astrophysically in $a \sin i$, the observable Δf_{obs} governs template placement. The equation decomposes into two number densities, $Q_{\text{frequency}}$ and Q_{mod} , and corresponding dimension length intervals, $L_{\text{frequency}}$, L_{mod} :

$$N_{\text{template}} = \Sigma_{(\text{frequency})} \Sigma_{(\text{mod})} (L_{\text{frequency}} \times Q_{\text{frequency}}) (L_{\text{mod}} \times Q_{\text{mod}}), \quad (\text{A.2})$$

where $Q_{\text{frequency}}$ is the inverse of the template spacing in frequency, which is $1/(2T_{\text{coh}})$, so $Q_{\text{frequency}} = 2T_{\text{coh}}$. Index the frequency dimension by j bands. One step j is made per band, width f_{bw} , so the length interval $L_{\text{frequency}} = f_{\text{bw}}$. Thus

$$L_{\text{frequency}} \times Q_{\text{frequency}} = (f_{\text{bw}}) \times (2T_{\text{coh}}). \quad (\text{A.3})$$

With fixed frequency bands, the number of templates per band does not change. The width in modulation depth, however, depends on $a \sin i$ and the frequency $f_j = (f_{\text{min}} + j f_{\text{bw}})$. Each template of Δf_{obs} is indexed by k . Analogous to before, $Q_{\text{mod}} = 1/(4T_{\text{coh}})$, so

$$L_{\text{mod}} \times Q_{\text{mod}} = (\Delta f_{\text{mod}, j, k}) \times (4T_{\text{coh}}), \quad (\text{A.4})$$

$$= \left(\frac{2\pi f_j}{cP} (a \sin i)_k \right) \times (4T_{\text{coh}}), \quad (\text{A.5})$$

when $a \sin i$ is given in light-seconds. Substituting Equation A.3 into Equation A.2, that term depends on neither the frequency nor modulation depth index, and so pulls out in front of the sums. Equation A.5 depends on k ; the sum $\Sigma_{(\text{mod})}$ is evaluated from $(a \sin i)_k = (a \sin i - 3\sigma_{a \sin i})$ to $(a \sin i + 3\sigma_{a \sin i})$ in practice, for an integrated length of $6\sigma_{a \sin i}$.

Combining these elements, including indexing frequency band steps by j ,

$$N_{\text{template}} = (f_{\text{bw}} \times 2T_{\text{coh}}) \Sigma_j \frac{2\pi(f_{\text{min}} + j f_{\text{bw}})}{cP} 6\sigma_{a \sin i} \times 4T_{\text{coh}}. \quad (\text{A.6})$$

Tuning into Scorpius X-1

18

Writing the limits of the sum of frequency bands, in addition to inserting ones to account for the edges of each sum, yields Equation 6b.

Appendix A.3. Test statistic calculation

The construction of this R -statistic can be described in several steps. The most important points from the original methods paper [5] are reiterated here, with some clarification. Science/observing runs are first parcelled into overlapping short Fourier transforms (SFTs), performed in the detector frame. The SFTs have typical coherence time T_{coh} (referred to as T_{SFT} in newer publications [37]) ranging from 60 s to 1800 s, depending on the hypothesized time-derivative of neutron star frequency [5]. The total number of SFTs with 50%-overlap for an observing time T_{obs} is N ,

$$N = \text{floor} \left(\frac{2T_{\text{obs}}}{T_{\text{coh}}} \right) - 1. \quad (\text{A.7})$$

SFT number in the observing run are indexed by $n \in [0, \dots, N - 1]$; SFT frequency bin is indexed by $k \in [0, \dots, K - 1]$, where $K = T_{\text{coh}} f_N$ for a Nyquist frequency f_N and only positive frequencies $k = T_{\text{coh}} f$ are used. Thus the transformation from time series to SFTs is a map from $x(t)$ ($= h(t) + n(t)$, signal strain plus noise) to \tilde{x}_k^n .

This array of n still depends on detector time t , and the analysis is to be done in Solar System Barycenter (SSB) time t_{SSB} . Travel from the source to SSB introduces an overall phase shift; uncertainty in the distance and proper motion is systematic and the same for gravitational and electromagnetic observations. Detector time is recorded in GPS time, running parallel with Terrestrial Time (TT), and SSB time runs parallel with Barycentric Dynamical Time (TDB). SSB time corrects t by δt_R for relativistic effects. Another overall phase shift is caused by Roemer delay Δt_R , the dot product of \hat{n}/c from the SSB to the sky location of interest with \vec{r} from the SSB to the detector [5, 12, 37]. Barycentering detector-frame data is equivalent to resampling in $\tau(t) = t_{\text{SSB}}(t) + \Delta t_R(t)$,

$$\tau(t) = t + \frac{\vec{r}(t) \cdot \hat{n}}{c} + \delta t_R. \quad (\text{A.8})$$

Each SFT frequency bin \tilde{x}_k^n is Doppler shifted to a frequency bin in the SSB frame corresponding to the sky location, frequency, and time t of the midpoint of the SFT n under investigation: $k(f(\alpha, \delta, t)) \rightarrow k(f_{\text{SSB}}(\tau))$. This barycentering procedure corresponds to the time-domain Equation A.8. Henceforth, barycentering is implicit in the k index.

Define the power $P_k^n = 2|\tilde{x}_k^n|^2/T_{\text{SFT}}$. Let $\langle P_k \rangle^n$ be the expected (estimated from a running mean over nearby n) noise-only power in a frequency bin k for SFT n . Also let $F_n^2 \equiv F_{+,n}^2 + F_{\times,n}^2$ for the antenna pattern at the chosen sky location and SFT n – taking this equal-weighted sum of F_+ and F_\times polarization components implies an assumption of circular polarization. Then the estimated power in a given bin \tilde{P}_k^n is normalized such that random, white, Gaussian noise will have an expectation value of 1 [5]:

$$\tilde{P}_k^n = \frac{F_n^2 (P_k^n - \langle P_k \rangle^n)}{(\langle P_k \rangle^n)^2} \left[\sum_{n'} \frac{F_{n'}^4}{(\langle P_k \rangle^{n'})^2} \right]^{-1}. \quad (\text{A.9})$$

REFERENCES

19

Then each row of barycentered frequency bins k is treated as a time series in n . Power for bin k in that time series is Fourier transformed by $\mathcal{F}_{f'}$ into $Z_k(f')$, where f' is the second Fourier transform frequency. During the transform, the background noise power $\lambda(f')$ is estimated from the noise in the SFTs, assuming the noise is Gaussian. This second Fourier power $Z_k(f')$ follows a χ^2 distribution with 2 degrees of freedom and mean 1.0, is proportional to h^4 , and is constructed by

$$Z_k(f') = \frac{|\mathcal{F}_{f'}[\tilde{P}_k^n]|^2}{\langle \lambda(f') \rangle}. \quad (\text{A.10})$$

When re-indexed by sorted template weight, i , $Z_k(f')$ becomes Z_i in Equation 5. Extensive discourse on the details of these calculations, as well as the estimation of background and calculation of template weights, is found in the original TwoSpect methods paper [5], and a paper on coherent addition of SFTs rigorously derives the SFT power by including Dirichlet kernel terms. Given the second Fourier power, we calculate the R -statistic its p -value, from which we seek to make a detection.

References

- [1] Manchester R, Hobbs G, Teoh A and Hobbs M 2005 *Astronom. J.* **129** 4
- [2] Harry G *et al.* 2010 *Class. Quant. Grav.* **27** 084006
- [3] Acernese F *et al.* 2009 Advanced Virgo baseline design Tech. Rep. VIR-0027A-09 Virgo
- [4] Goetz E A 2010 *Gravitational wave studies: detector calibration and an all-sky search for spinning neutron stars in binary systems* Ph.D. thesis University of Michigan
- [5] Goetz E and Riles K 2011 *Class. Quant. Grav.* **28** 215006
- [6] Aasi J *et al.* 2014 *Phys. Rev. D* **90** 062010
- [7] Messenger C, Bulten H J, Crowder S G, Dergachev V, Galloway D K, Goetz E, Jonker R J G, Lasky P D, Meadors G D, Melatos A, Premachandra S, Riles K, Sammut L, Thrane E H, Whelan J T and Zhang Y 2015 *Phys. Rev. D* **92**(2) 023006
- [8] Abbott B *et al.* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116**(6) 061102
- [9] Brady P, Creighton T, Cutler C and Schutz B 1998 *Phys. Rev. D* **57**(4) 2101
- [10] Shawhan P 2010 *Class. Quant. Grav.* **27** 084017
- [11] Owen B 2010 *Phys. Rev. D* **82** 104002
- [12] Jaranowski P, Królak A and Schutz B 1998 *Phys. Rev. D* **58** 063001
- [13] Krishnan B, Sintes A M, Papa M A, Schutz B F, Frasca S and Palomba C 2004 *Phys. Rev. D* **70** 082001
- [14] Abbott B 2008 *Phys. Rev. D* **77** 022001

REFERENCES

20

- [15] Abbott B *et al.* 2009 *Phys. Rev. Lett* **102** 111102
- [16] Dergachev V 2010 *Class. Quant. Grav.* **27** 205017
- [17] Abadie J *et al.* 2012 *Phys. Rev. D* **85** 022001
- [18] Riles K 2013 *Prog. in Particle & Nucl. Phys.* **68** 1
- [19] Wette K *et al.* 2008 *Class. Quant. Grav.* **25** 235011
- [20] Abadie J *et al.* 2010 *ApJ* **722** 1504
- [21] Dupuis R and Woan G 2005 *Phys. Rev. D* **72** 102002
- [22] Aasi J *et al.* 2014 *Astrophys. J* **785** 119
- [23] Chakrabarty D *et al.* 2003 *Nature* **424** 42
- [24] Messenger C and Woan G 2007 *Classical and Quantum Gravity* **24** S469
- [25] Sammut L, Messenger C, Melatos A and Owen B 2014 *Phys. Rev. D* **89** 043001
- [26] Ballmer S W 2006 *Classical and Quantum Gravity* **23** S179
- [27] Abadie J *et al.* 2011 *Phys. Rev. Lett.* **107** 271102
- [28] van der Putten S, Bulten H J, van den Brand J F J and Holtrop M 2010 *Journal of Physics Conference Series* **228** 012005
- [29] Dhurandhar S, Krishnan B, Mukhopadhyay H and Whelan J T 2008 *Phys. Rev. D* **77**(8) 082001
- [30] Whelan J T, Sundaesan S, Zhang Y and Peiris P 2015 *Phys. Rev. D* **91**(10) 102005
- [31] Abbott B *et al.* 2007 *Phys. Rev. D* **76** 082001
- [32] Leaci P and Prix R 2015 *Phys. Rev. D* **91**(10) 102003
- [33] Papaloizou J and Pringle J 1978 *MNRAS* **184** 501
- [34] Wagoner R 1984 *Ap. J.* **278** 345
- [35] Bildsten L 1998 *Astrophys. J. Lett.* **501** L89
- [36] Watts A, Krishnan B, Bildsten L and Schutz B 2008 *MNRAS* **389** 839
- [37] Goetz E and Riles K 2015 (*Preprint gr-qc/1510.06820*)
- [38] Galloway D K, Premachandra S, Steeghs D, Marsh T, Casares J and Cornelisse R 2014 *Ap J* **781** 14 (*Preprint 1311.6246*)
- [39] Markwardt C *et al.* 2002 *Astrophys. J* **575** L21–L24
- [40] Bradshaw C, Fomalont E and Geldzahler B 1999 *ApJ* **512** L121
- [41] Skrutskie M F *et al.* 2006 *The Astronomical Journal* **131** 1163–1183
- [42] Steeghs D and Casares J 2002 **568** 273–278 (*Preprint astro-ph/0107343*)
- [43] Abadie J *et al.* 2010 *NIM-A* **623** 223–240
- [44] The LIGO Scientific Collaboration LALApps repository Web: <http://www.lsc-group.phys.uwm.edu/daswg/>
- [45] Behnke B, Papa M and Prix R 2015 *PRD* **91** 064007