

Implementation of a new WEAVE-based search pipeline for continuous gravitational waves from known binary systems

Arunava Mukherjee^{1,2,3,*}, Reinhard Prix^{1,2,†} and Karl Wette^{4,5}

¹Max-Planck-Institut für Gravitationsphysik, Callinstrasse 38, D-30167 Hannover, Germany

²Leibniz Universität Hannover, Callinstrasse 38, D-30167 Hannover, Germany

³Saha Institute of Nuclear Physics, 1/AF Bidhannagar, Kolkata-700064, India

⁴Centre for Gravitational Astrophysics, Australian National University, Canberra ACT 2601, Australia

⁵ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), Hawthorn VIC 3122, Australia



(Received 30 July 2022; accepted 23 November 2022; published 17 March 2023)

Scorpius X-1 (Sco X-1) has long been considered one of the most promising targets for detecting continuous gravitational waves with ground-based detectors. Observational searches for Sco X-1 have achieved substantial sensitivity improvements in recent years, to the point of starting to rule out emission at the torque-balance limit in the low-frequency range ~ 40 – 180 Hz. In order to further enhance the detection probability, however, there is still much ground to cover for the full range of plausible signal frequencies ~ 20 – 1500 Hz, as well as a wider range of uncertainties in binary orbital parameters. Motivated by this challenge, we have developed BINARYWEAVE, a new search pipeline for continuous waves from a neutron star in a known binary system such as Sco X-1. This pipeline employs a semicoherent StackSlide \mathcal{F} -statistic using efficient lattice-based metric template banks, which can cover wide ranges in frequency and unknown orbital parameters. We present a detailed timing model and extensive injection-and-recovery simulations that illustrate that the pipeline can achieve high detection sensitivities over a significant portion of the parameter space when assuming sufficiently large (but realistic) computing budgets. Our studies further underline the need for stricter constraints on the Sco X-1 orbital parameters from electromagnetic observations, in order to be able to push sensitivity below the torque-balance limit over the entire range of possible source parameters.

DOI: [10.1103/PhysRevD.107.062005](https://doi.org/10.1103/PhysRevD.107.062005)

I. INTRODUCTION

Since the first direct detection of gravitational waves from the coalescence of two stellar-mass black holes [1], we have observed more than 90 more gravitational-wave events [2,3]. So far, all the observed signals originated from the coalescence of binary black-hole systems, binary neutron-star systems, and neutron-star–black-hole systems, each resulting in a short transient signal in the gravitational-wave detectors.

A different class of gravitational-wave signals, continuous gravitational waves (CWs) that are nearly monochromatic and long-lasting, is yet to be observed. Rapidly spinning neutron stars with some deviation from perfect

axisymmetry are promising sources of such CWs in the current generation of ground-based detectors, namely, Advanced-LIGO (aLIGO), Advanced-VIRGO, and KAGRA [4].

Different physical processes within a neutron star can produce CWs, resulting in different characteristics of the emitted signal. For example, a nonaxisymmetric deformation (or “mountain”) on a spinning neutron star emits CWs at twice the spin frequency, $f = 2f_{\text{rot}}$, while a freely precessing neutron star will additionally emit at $f \sim f_{\text{rot}}$ [4]. Oscillation modes of the internal fluids in a neutron star can also produce CWs; for example, inertial r-mode oscillations emit at approximately $f \sim \frac{4}{3}f_{\text{rot}}$ through the Chandrasekhar-Friedman-Schutz instability [4–7].

Accreting neutron stars in galactic low-mass x-ray binary (LMXB) systems are potentially strong emitters of CWs [8–11], as the accreting matter from the companion, channeled by the magnetic field of the neutron star, can result in a substantial degree of quadrupolar nonaxisymmetry of the spinning neutron star [12–14].

The accreting matter also exerts a spin-up torque on the neutron star, increasing its spin frequency f_{rot} . Interestingly, however, the observed distribution of neutron-star spin

*arunava.mukherjee@ligo.org; arunava.mukherjee@saha.ac.in

†reinhard.prix@aei.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

frequencies shows a pronounced cutoff above $f_{\text{rot}} \sim 700$ Hz, well below the theoretical breakup limit of realistic neutron-star equations of state [15,16]. Gravitational-wave emission is one of the conjectured braking mechanisms that could explain this surprising high-frequency cutoff in the spin distribution. According to the torque-balance scenario, the spin-down torque due to the emission of CWs would eventually counterbalance the accretion-induced spin-up torque. Thus, the larger the mass accretion rate, the stronger the expected gravitational-wave emission.

Sco X-1 is the brightest LMXB with one of the highest mass-accretion rates among the systems harboring a neutron star [17]. Moreover, it is relatively close to Earth, with a distance of only ~ 2.8 kpc [18], making it one of the most promising sources of detectable CWs [19,20].

Searching for CW signals in data from ground-based detectors is an active area of ongoing effort (e.g., see [4,21] for recent overviews). We typically classify these searches (in order of decreasing computational cost) into three main categories: all-sky searches for unknown sources over a wide range of source parameters; directed searches for sources with known sky locations and some unknown intrinsic parameters; and targeted searches for known pulsars, where the phase evolution of the system is assumed to be known. Searches for Sco X-1 fall into the directed category, with a known sky position and unknown frequency, and substantial uncertainties on some of the binary orbital parameters.

Sco X-1 has long been considered one of the high-priority targets for CW searches, starting with [22], with further searches on initial LIGO data [23–25], and more recently on data from the Advanced LIGO detectors [26–29]. Each successive search has improved constraints on the maximal strength of a putative CW signal from Sco X-1, with the latest constraints for the first time beating the above-mentioned torque-balance limit in a range of low spin frequencies $f_{\text{rot}} \sim 20\text{--}90$ Hz [29].

The past decade has seen the development and deployment of several pipelines for Sco X-1 searches (e.g., see [20] for an overview). Different pipelines tend to achieve different sensitivity per computing cost and degrees of robustness against the model assumptions, such as the effect of stochastic accretion torque on the spin-frequency evolution, i.e., the so-called spin wandering [30]. Recent advances in search techniques include the adaptation of the Viterbi “hidden Markov Model” methods to searches for Sco X-1 [31,32], and the cross-correlation *CrossCorr* pipeline [24,33], which was recently improved by using resampling techniques [34] as well as efficient lattice-based template banks [35].

One of the open challenges for finding CWs from Sco X-1 stems from the fact that most of the observed neutron-star rotation rates in accreting LMXBs fall above $f_{\text{rot}} \gtrsim 300$ Hz [15,16]. Given that the mass accretion rate of Sco X-1 is one of the highest observed among all LMXBs, the neutron star would have experienced a large amount of accretion-induced

spin-up torque and therefore have a high spin frequency f_{rot} . Unfortunately, the computational cost of such a CW search grows as a substantial power of frequency $\propto f^{3-6}$, depending on the assumed parameter space [36]. Therefore, reaching or surpassing the torque-balance limit at higher spin frequencies becomes increasingly challenging.

Here we present BINARYWEAVE, a new directed Sco X-1 search pipeline that employs a semicoherent \mathcal{F} -statistic StackSlide approach, as outlined and analyzed in Leaci and Prix [36]. This is achieved by extending the WEAVE framework [37], originally developed as an all-sky search for isolated neutron stars [38]. Using this framework enables us to use the fastest-available (resampling) \mathcal{F} -statistic algorithms and efficient lattice-based metric template banks for covering the parameter space and summing \mathcal{F} -statistics across segments. The tunable segment lengths and template-bank mismatch parameters allow this pipeline to translate increases in computing budget (e.g., by using Einstein@Home [39] or a large computing cluster) into improved sensitivity [40].

The BINARYWEAVE pipeline constructs a bank of a large number of templates originating from different values of the intrinsic source parameters, e.g., spin frequency of the neutron star and orbital parameters of the binary system. Construction of a reliable and efficient template bank maximizes the detection of a weak signal above a pre-determined threshold value. Often an increase in the number of templates increases the detection probability, but it also requires a higher amount of computational resources. Thus, our primary goal is to maximize the detection probability within the (varying) limitation in computing budgets. As discussed in detail in this paper (see Secs. IV and V C), the construction of a template bank is key to this idea.

There are two widely adopted general methodologies to construct template banks for GW searches: stochastic template banks and geometric template banks. A geometric template bank uses algorithms to place individual templates geometrically to tile the parameter space targeted for the search. The distance of any two adjacent templates in multidimensional parameter space is dictated by the maximum amount of affordable loss in signal-to-noise ratio, which in turn is determined by the metric of the parameter space locally. Thus, the knowledge of the parameter-space metric is crucial in order to construct any geometric template bank for any search pipeline.

The main difficulty stemmed from the fact that the Sco X-1 metric changes over the parameter space [37], while the lattice-tiling WEAVE template-bank construction requires a strictly constant metric. We have solved this problem by developing a local approximation to the binary-orbital coordinates resulting in an “effective” constant parameter-space metric allowing for efficient lattice tiling while satisfying good coverage and mismatch properties.

We present and characterize the sensitivity and computational performance of BINARYWEAVE, which essentially realizes the predicted sensitivities in Leaci and Prix [36]. We discuss its applications for different astrophysical Sco X-1 scenarios, observation setups, and computing budgets. We discuss some aspects of electromagnetic observations that would help to substantially alleviate the computational challenges and improve the chances for a Sco X-1 CW detection.

This paper is organized as follows: in Sec. II we introduce the signal waveform parameters, detection statistic, and template-bank construction. Section III describes the specifics of the implementation in BINARYWEAVE. In Sec. IV we present a detailed characterization of this new pipeline in terms of template-bank safety as well as computing-resource requirements. Section V presents the achievable sensitivities of this pipeline for different computational budgets, followed by a summary and an outlook in Sec. VI.

II. BACKGROUND

In this section we briefly introduce the concepts and notation required to understand the context of this paper (closely following [36]), namely, the CW signal waveform and its parameters, the detection statistics used, and the basics of metric template-bank construction.

A. Signal waveform and parameters

The time-dependent strain of a CW signal exerted on a gravitational-wave detector is $h(t; \mathcal{A}, \lambda)$, where t is the arrival time of a wavefront at the detector. The set of four amplitude parameters \mathcal{A} consists of the overall amplitude h_0 , the inclination angle $\cos \iota$, polarization angle ψ , and the initial phase ϕ_0 . The phase-evolution parameters λ determine the waveform phase $\phi(t; \lambda)$ as a function of time at a given detector. The phase evolution at the detector is determined by the source-frame frequency evolution $f(\tau)$ (dependent on the intrinsic spin evolution of the neutron star), and by the Rømer delay affecting the arrival time $\tau(t)$ at the detector, due to the relative motion of the detector and neutron star, and (if it is in a binary system) the star's intrinsic motion around its companion star.

The neutron-star spin typically changes slowly and can therefore be represented by a Taylor expansion around a reference time t_{ref} , resulting in a source-frame phase model of the form

$$\phi^{\text{src}}(\tau) = 2\pi \left[f(\tau - t_{\text{ref}}) + \frac{1}{2} \dot{f}(\tau - t_{\text{ref}})^2 + \dots \right], \quad (1)$$

with f denoting the source-frame gravitational-wave frequency at t_{ref} and its higher-order derivatives, or “spindown parameters,” $f^{(k)} \equiv df^k/d\tau^k|_{t_{\text{ref}}}$.

The waveform arrival time $\tau(t)$ from the source frame τ to the detector frame t is determined by the detector location and source sky position (e.g., right ascension and declination) and by the orbital parameters describing the intrinsic neutron-star motion if it is in a binary system [41]. These consist, at a minimum, of the orbital projected semimajor axis a_p , the period P_{orb} (or equivalently the mean orbital angular velocity $\Omega \equiv 2\pi/P_{\text{orb}}$), and a reference time of the orbit, such as the time of ascending node t_{asc} . These parameters would fully describe the time delay in a circular orbit, while for eccentric orbits we additionally require the eccentricity e and a rotation angle, such as the argument of periapsis ω . For systems with small eccentricity a common reparametrization uses Laplace-Lagrange parameters κ and η instead, defined as

$$\kappa \equiv e \cos \omega, \quad \eta \equiv e \sin \omega. \quad (2)$$

In the small-eccentricity limit one can relate the time of periapsis t_p to the time of ascending node t_{asc} via

$$t_p - t_{\text{asc}} = \frac{\omega}{\Omega}. \quad (3)$$

Explicit expressions for the resulting phase model can be found in [36,42].

B. Detection statistics

The gravitational-wave strain $x(t)$ observed in a detector in the presence of a signal and additive noise $n(t)$ can be written as $x(t) = n(t) + h(t; \mathcal{A}, \lambda)$. The detection problem therefore corresponds to distinguishing the pure-noise hypothesis, i.e., $h(t) = 0$, from the signal hypothesis with nonvanishing $h(t)$. The standard likelihood-ratio approach can be used to test different templates $h(t; \mathcal{A}, \lambda)$ against the data, with a common simplification consisting in the analytic maximization over amplitude parameters, first shown in Jaranowski *et al.* [43], resulting in the \mathcal{F} -statistic. While this approach is not strictly optimal compared to Bayesian marginalization [44], it requires far less computing cost per template λ and is therefore the current best choice for computationally constrained, wide parameter-space searches.

We denote the (coherent) statistic as $2\mathcal{F}(x; \lambda)$, which depends on the data x and the phase-evolution parameters λ of the waveform template for which the statistic is computed. In Gaussian noise this statistic follows a non-central χ^2 distribution with 4 degrees of freedom and a noncentrality parameter $\rho^2(\mathcal{A}_s, \lambda_s; \lambda)$, where \mathcal{A}_s and λ_s are the (unknown) signal amplitude- and phase-evolution parameters, while λ are the template phase-evolution parameters. The expectation value of the coherent \mathcal{F} -statistic is

$$E[2\mathcal{F}(x; \lambda)] = 4 + \rho^2(\mathcal{A}_s, \lambda_s; \lambda). \quad (4)$$

The noncentrality parameter ρ^2 characterizes the signal power in a given template, and in the coherent case its square root ρ is also known as the signal-to-noise ratio (SNR) for the coherent \mathcal{F} -statistic.

For wide parameter-space searches (such as for Sco X-1), some or all of the signal phase-evolution parameters λ_s are unknown, constrained only to fall in some astrophysically informed parameter space $\lambda_s \in \mathcal{P}$. The required number of templates to cover a parameter space \mathcal{P} using a coherent statistic \mathcal{F} grows rapidly as a function of the coherent integration time, which makes such searches effectively computationally impossible. Consequently, the best achievable sensitivity at a finite computational cost is typically achieved using semicoherent statistics, as first shown in Brady *et al.* [45] and analyzed in more detail in [40].

The BINARYWEAVE pipeline as an extension of WEAVE [37] is based on the standard StackSlide [46] semicoherent approach using summed \mathcal{F} -statistic over shorter coherent segments. The total observation live time T_{obs} is divided into N shorter segments of duration ΔT ; i.e., in an ideal uninterrupted observation one would have $T_{\text{obs}} = N\Delta T$. The semicoherent $\hat{\mathcal{F}}$ -statistic is defined as the sum of the coherent per-segment \mathcal{F} -statistics over all N segments, i.e.,

$$2\hat{\mathcal{F}}(x; \lambda) \equiv \sum_{\ell=1}^N 2\mathcal{F}_{\ell}(x; \lambda). \quad (5)$$

This statistic follows a (noncentral) χ^2 distribution with $4N$ degrees of freedom and a noncentrality parameter (or signal power) given by

$$\hat{\rho}^2(x; \lambda) = \sum_{\ell=1}^N \rho_{\ell}^2(x; \lambda), \quad (6)$$

so the expectation value of $2\hat{\mathcal{F}}$ is

$$E[2\hat{\mathcal{F}}(x; \lambda)] = 4N + \hat{\rho}^2(\mathcal{A}_s, \lambda_s; \lambda). \quad (7)$$

C. Template banks and parameter-space metrics

In order to systematically search a given parameter space \mathcal{P} , we need to populate it with a finite number of templates $\lambda_t \in \mathcal{P}$. The set of all templates $\{\lambda_t\}$ is referred to as the template bank, which is a discrete sampling of \mathcal{P} , i.e., $\{\lambda_t\} \subset \mathcal{P}$. Due to this discretization of \mathcal{P} , a signal with parameters $\lambda_s \in \mathcal{P}$ will not fall on an exact template, resulting in a loss of recovered signal power ρ^2 at a template λ_t , which is quantified in terms of the *mismatch* μ_0 , defined as the relative loss of signal power

$$\mu_0(\mathcal{A}_s, \lambda_s; \lambda_t) \equiv \frac{\rho^2(\mathcal{A}_s, \lambda_s; \lambda_s) - \rho^2(\mathcal{A}_s, \lambda_s; \lambda_t)}{\rho^2(\mathcal{A}_s, \lambda_s; \lambda_s)}, \quad (8)$$

which is a bounded function within $\mu_0 \in [0, 1]$.

Assuming a small offset $d\lambda \equiv \lambda_t - \lambda_s$ between the signal and template and neglecting the dependence on the (unknown) signal amplitude parameters \mathcal{A}_s [47], one can define the parameter-space (phase-) metric g_{ij} in terms of the truncated quadratic Taylor expansion:

$$\mu(\lambda_s; \lambda_t) = g_{ij}(\lambda_s) d\lambda^i d\lambda^j, \quad (9)$$

with implicit summation over the repeated indices $i, j = 1, \dots, n$, where n is the number of template-bank dimensions.

The mismatch μ represents the squared *distance* corresponding to the parameter offsets $d\lambda$, and the metric g_{ij} defines a distance measure on the parameter space. As a result, one can express the *bulk* number of templates $\mathcal{N}_{\mathcal{P}}$ in an n -dimensional lattice template bank covering the parameter space \mathcal{P} with maximum mismatch μ_{max} (corresponding to the squared covering radius of the lattice) [45,48] as

$$\mathcal{N}_{\mathcal{P}} = \theta_n \mu_{\text{max}}^{-n/2} \int_{\mathcal{P}} \sqrt{\det g(\lambda)} d^n \lambda, \quad (10)$$

in terms of the lattice-specific normalized thickness θ_n ; a thinner lattice will cover the same volume with fewer templates. This bulk template number ignores any extra padding typically required to fully cover the boundary $\partial\mathcal{P}$ of the parameter space \mathcal{P} , which tends to increase the total number of templates in practice (e.g., see [37,49]).

The metric allows for a simple estimate of the approximate *scale* of the template-bank resolution along single coordinates via

$$\delta\lambda^i = 2\sqrt{\frac{\mu_{\text{max}}}{g_{ii}}}, \quad (11)$$

which is obtained from Eq. (9) by assuming a single nonzero offset along one coordinate axis $\delta\lambda^i$. In a one-dimensional template-bank grid, the factor of 2 accounts for the fact that the maximum mismatch μ_{max} would be attained at the midpoint between two lattice templates. The true higher-dimensional grid spacings will typically be larger than this estimate, however, due to potential nonzero cross-terms g_{ij} that come into play when considering generic offsets, as well as using lattice structures other than a simple rectangular grid along coordinate axes.

A somewhat complementary grid-scale estimate can be obtained from considering the extents of the bounding box [36] around a metric ellipse of constant mismatch, Eq. (9), namely,

$$D\lambda^i = 2\sqrt{\mu_{\text{max}}(g^{-1})^{ii}}, \quad (12)$$

where g^{-1} is the inverse matrix of the metric g . Contrary to Eq. (11), this fully takes into account parameter

correlations but will generally result in an overestimate of the actual lattice grid spacing [47,50].

The coherent phase metric $g_{ij}(\lambda)$ at a parameter-space point λ [ignoring the (unknown) signal amplitude parameters \mathcal{A}_s] can be shown [45,47] to be expressible directly in terms of derivatives of the signal phase $\phi(t; \lambda)$, namely,

$$g_{ij}(\lambda) = \langle \partial_i \phi(\lambda) \partial_j \phi(\lambda) \rangle - \langle \partial_i \phi(\lambda) \rangle \langle \partial_j \phi(\lambda) \rangle, \quad (13)$$

where $\partial_i \phi(\lambda) \equiv \partial \phi / \partial \lambda^i$ and $\langle Q \rangle$ denotes time averaging of a quantity Q over the coherent duration ΔT , i.e., $\langle Q \rangle \equiv (1/\Delta T) \int_{t_0}^{t_0 + \Delta T} Q(t) dt$.

The corresponding semicoherent metric $\hat{g}_{ij}(\lambda)$ at a point λ can then be obtained [51] as the average over segments, namely,

$$\hat{g}_{ij}(\lambda) = \frac{1}{N} \sum_{\ell=1}^N g_{\ell,ij}(\lambda), \quad (14)$$

where $g_{\ell,ij}$ is the coherent metric of segment ℓ .

D. Sco X-1 parameter-space metric

A number of rapidly spinning neutron stars in LMXB systems are found to be in (approximate) spin equilibrium [52–54]. According to the gravitational-wave torque-balance hypothesis, the total amount of accretion-induced spin-up torque would be counterbalanced by the braking torques due to the emission of CWs and electromagnetic radiation [9]. This keeps the system in approximate torque balance, with random fluctuations in spin frequency due to the stochastic nature of the accretion flows, which is known as spin wandering [30,55].

Similar to previous studies, and following [36], we therefore assume a constant intrinsic signal frequency f with no long-term drifts, i.e., $f^{(k \geq 1)} = 0$, and we tackle the spin-wandering effect by limiting the maximal segment length ΔT such that the frequency resolution is still too coarse for any spin wandering effect to move the signal by more than one frequency bin.

We can therefore use the following physical phase-evolution parameters describing the CW waveforms

$$\lambda = \{f, a_p, t_{\text{asc}}, \Omega, \kappa, \eta\}, \quad (15)$$

and assuming the small-eccentricity limit for Sco X-1, i.e., $e \ll 1$, the approximate CW phase model [36,42] can be written as

$$\frac{\phi(t; \lambda)}{2\pi} \approx f \Delta t - f a_p \left[\sin \Psi + \frac{\kappa}{2} \sin 2\Psi - \frac{\eta}{2} \cos 2\Psi \right], \quad (16)$$

where $\Delta t \equiv t - t_{\text{ref}}$ and the orbital phase $\Psi(t)$ is given by

$$\Psi(t) = \Omega(t - t_{\text{asc}}). \quad (17)$$

From the explicit expression Eq. (16) of the phase, one can obtain the phase derivatives $\partial_i \phi$ with respect to the parameter-space coordinates λ^i , and time averaging yields the coherent metric components $g_{\ell,ij}$ for each segment ℓ according to Eq. (13). The semicoherent metric \hat{g}_{ij} is then obtained by averaging over segments following Eq. (14).

In the template-bank construction the metric will typically be computed numerically starting from the analytic expressions for the phase derivatives. However, it is important to also consider approximate analytic expressions for these metrics, in order to better understand their properties. As discussed in [36,42], analytic approximations can be found in the two limiting cases: short segments where $\Delta T \ll P_{\text{orb}}$, or long segments where $\Delta T \gg P_{\text{orb}}$. Longer segments will result in better sensitivity but also higher computational cost. The results in [36] indicate that using a realistically large computational budget, semicoherent $\hat{\mathcal{F}}$ -statistic searches for Sco X-1 can afford segments substantially longer than $P_{\text{orb}} \sim 19$ h. Therefore we will only discuss the long-segment limit here, for which the nonzero elements in the analytic approximation to the semicoherent metric are found [36] as

$$\begin{aligned} \hat{g}_{ff} &= \pi^2 \frac{\Delta T^2}{3}, \\ \hat{g}_{a_p a_p} &= 2\pi^2 f^2, \\ \hat{g}_{\Omega \Omega} &= 2\pi^2 (f a_p)^2 \left(\frac{\Delta T^2}{12} + \overline{\Delta_{\text{ma}}^2} \right), \\ \hat{g}_{t_{\text{asc}} t_{\text{asc}}} &= 2\pi^2 (f a_p \Omega)^2, \\ \hat{g}_{\Omega t_{\text{asc}}} &= \hat{g}_{t_{\text{asc}} \Omega} = -2\pi^2 (f a_p)^2 \Omega \overline{\Delta_{\text{ma}}}, \\ \hat{g}_{\kappa \kappa} &= \hat{g}_{\eta \eta} = \frac{\pi^2}{2} (f a_p)^2, \end{aligned} \quad (18)$$

where $\Delta_{\text{ma},\ell} \equiv t_{\text{mid},\ell} - t_{\text{asc}}$ is the time offset between the midpoint $t_{\text{mid},\ell}$ of segment ℓ and the ascending node t_{asc} , and where \overline{Q} denotes averaging over segments, i.e., $\overline{Q} \equiv (1/N) \sum_{\ell=1}^N Q_{\ell}$. Note that the coherent per-segment metric $g_{\ell,ij}$ can simply be read off these expressions as the special case $\ell = N = 1$.

There are two important aspects to consider about this metric:

- (1) The metric components still depend on the search parameters f, a_p, Ω and t_{asc} and are therefore not constant over the parameter space. This is an obstacle to constructing a lattice template bank, which will be dealt with in Sec. III.
- (2) There is little *refinement* of the semicoherent metric compared to the per-segment coherent resolution; in fact most components do not depend on the number of segments (i.e., the total duration of data used for the searches) for a fixed duration of coherent

segment ΔT , except for $\hat{g}_{\Omega\Omega}$ (via $\overline{\Delta_{\text{ma}}^2}$) and $\hat{g}_{\Omega t_{\text{asc}}}$ (via $\overline{\Delta_{\text{ma}}}$).

In order to simplify the expression, we can make use of the gauge freedom in t_{asc} , which is only defined up to an integer multiple of the period P_{orb} ; i.e.,

$$t'_{\text{asc}} = t_{\text{asc}} + nP_{\text{orb}}, \quad \text{for } n \in \mathbb{Z}, \quad (19)$$

describes the same physical orbit, as seen in Eqs. (16) and (17). Given the long-segment assumption $\Delta T \gg P_{\text{orb}}$, the total observation time will satisfy this even more strongly, i.e., $T_{\text{obs}} \geq N\Delta T \gg P_{\text{orb}}$. One can therefore choose a gauge $t_{\text{asc}} \approx \overline{t_{\text{mid}}}$ such that $\overline{\Delta_{\text{ma}}} \approx 0$, removing the only nonzero off-diagonal component $g_{\Omega t_{\text{asc}}}$. Further, assuming gapless segments, one can show [36] that in this case

$$\hat{g}_{\Omega\Omega} = \frac{\pi^2}{6} (fa_p)^2 (N\Delta T)^2. \quad (20)$$

In other words, only the semicoherent resolution in Ω increases with the number of segments, while all other parameters have the same metric resolution per segment and in the semicoherent combination. This point will be further discussed in Sec. III on the details of the BINARYWEAVE implementation.

E. Lattice-tiling template banks

The template-bank construction in BINARYWEAVE is directly inherited from WEAVE, described in full detail in [37,49]; therefore, we only provide a short overview here. The basic inputs to the lattice-tiling algorithm are the parameter-space coordinates $\{\lambda^i\}$, boundaries defining \mathcal{P} , and the corresponding template-bank metric, which must be constant over the search space. The code can use a coordinate transformation to *internal* coordinates if the metric is expressed in different coordinates than the standard CW waveform parameters described in Sec. II A. Based on these inputs, together with a maximum-mismatch parameter μ_{max} and a choice of lattice type, the algorithm constructs a template-bank lattice with covering radius $\sqrt{\mu_{\text{max}}}$ tiling the parameter space (and ensuring appropriate covering of the boundaries).

There are two main modes that semicoherent statistics can be computed over the set of segments: interpolating and noninterpolating. As mentioned in Sec. II D, the semicoherent template bank requires a finer resolution in Ω to compute $\hat{\mathcal{F}}$ than the per-segment template banks to compute \mathcal{F}_ℓ at given maximum mismatch μ_{max} . This can be used to save computing power by using coarser per-segment template banks together with a nearest-neighbor interpolation when picking per-segment \mathcal{F}_ℓ to sum in Eq. (5). The details and effects of such an interpolating StackSlide approach are discussed in [40]. The simpler, yet generally more computationally expensive method consists in using the same semicoherent fine grid over all segments,

such that Eq. (5) can directly be computed without any interpolation.

The amount of computing-cost savings due to interpolation depends on the refinement factor between coherent and semicoherent metrics, which in the case of the Sco X-1 metric is only linear in N if including Ω in the template bank, and unity otherwise, as discussed in Sec. II D. The expected sensitivity gains by using interpolation in this case would therefore be modest and partially reduced by the extra mismatch incurred due to interpolation itself (see [40]). Furthermore, it is more difficult to find optimal setup parameters for an interpolating setup, given there are two mismatch parameters $\{\mu_{\text{max}}, \hat{\mu}_{\text{max}}\}$ to tune rather than a single μ_{max} , in addition to the number N and length ΔT of the semicoherent segments.

The A_n^* lattice is a common choice [37,49,56,57] for a close-to-optimal template-bank lattice, based on earlier arguments about optimal covering lattices [48,58]. Recent work [59,60] has clarified, however, that finding the template-bank lattice that maximizes the expected detection probability at fixed number of templates is an instance of the quantizer problem [58] not the covering problem. This somewhat changes the choice of current “record holder” lattice in each dimension and reduces the relative advantage of A_n^* over the hypercubic lattice, but even in this paradigm A_n^* remains a close-to-optimal lattice and therefore continues to be a practically reasonable and sound choice.

F. Sco X-1 parameter space

Optical and x-ray observations tell us that Sco X-1 is an LMXB system [61]. Furthermore, x-ray spectral and timing characteristics indicate that the compact object in the Sco X-1 binary system is a neutron star [17]. Observations in optical and radio bands have constrained the three orbital parameters a_p , P_{orb} , and t_{asc} of Sco X-1 to different extents [18,62,63], namely, $a_p \in [1.45, 3.25]$ ls, $P_{\text{orb}} \sim 68\,023.86048 + -0.0432$ s, and $t_{\text{asc}} \sim 974\,416\,624 \pm 50$ GPS s (as used in a recent CW search [28]) while the spin frequency of the neutron star still remains practically unconstrained [64] to date. We provide a table of various Sco X-1 parameter-space ranges considered in this and past studies (and searches) in Table I, which will be discussed in more detail in Sec. V B.

The typical life cycle of an LMXB, along with one of the highest mass-accretion rate systems, indicates that the neutron star in Sco X-1 is likely to receive a large amount of accretion-induced spin-up torque [55] and is plausibly spinning rapidly. Most of the accreting neutron stars in LMXB systems are observed to be spinning in the range of ~ 300 – 600 Hz, although a few of them have also been observed at lower spin frequencies [15,65]. We explore the implications of different assumptions about the spin and orbital parameters of Sco X-1 for a wide-parameter search in Sec. V B.

TABLE I. Different parameter-space search regions considered for Sco X-1. Note that \mathcal{P}_0 has been used in this study as a test range for various Monte Carlo tests of BINARYWEAVE, and \mathcal{P}_{1-3} represent observational constraints considered in recent CW searches and studies. In addition, various combinations of parameter ranges are considered, \mathcal{P}_{4-21} , in order to explore the impact of improved observation constraints and reduced search ranges.

Search space \mathcal{P}	f [Hz]	a_p [ls]	P_{orb} [s]	t_{asc} [GPS s]	Reference(s)/comment(s)
\mathcal{P}_0	10–700	0.3–3.5	68023.7 ± 0.2	1124044455.0 ± 1000	BINARYWEAVE test range
\mathcal{P}_1	20–500	1.26–1.62	68023.70496 ± 0.0432	897753994 ± 100	Leaci and Prix [36]
\mathcal{P}_2	60–650	1.45–3.25	68023.86048 ± 0.0432	974416624 ± 50	Abbott <i>et al.</i> [28]
\mathcal{P}_3	40–180	1.45–3.25	68023.86 ± 0.12	1178556229 ± 417	Zhang <i>et al.</i> [29]
\mathcal{P}_4	600–700				
\mathcal{P}_5	1000–1100				
\mathcal{P}_6	1400–1500	1.45–3.25	68023.70496 ± 0.0432	974416624 ± 100	Different ranges in frequency with broad range in a_p
\mathcal{P}_7	20–250				
\mathcal{P}_8	20–1000				
\mathcal{P}_9	20–1500				
\mathcal{P}_{10}	600–700				
\mathcal{P}_{11}	1000–1100				
\mathcal{P}_{12}	1400–1500	1.40–1.50	68023.70496 ± 0.0432	974416624 ± 100	Different ranges in frequency with narrow range in a_p
\mathcal{P}_{13}	20–500				
\mathcal{P}_{14}	20–1000				
\mathcal{P}_{15}	20–1500				
\mathcal{P}_{16}	600–700				
\mathcal{P}_{17}	1000–1100				
\mathcal{P}_{18}	1400–1500	1.44–1.45	68023.70496 ± 0.0432	974416624 ± 100	Different ranges in frequency with well-constrained a_p
\mathcal{P}_{19}	20–500				
\mathcal{P}_{20}	20–1000				
\mathcal{P}_{21}	20–1500				

III. FLAT METRIC APPROXIMATION

As discussed in Sec. II D, the long-segment binary parameter-space metric Eq. (18) in physical coordinates is not constant over $\{f, a_p, \Omega, t_{\text{asc}}\}$, which prohibits its direct use for lattice tiling. This represents the main obstacle to applying the WEAVE framework to a directed binary search.

Regarding the frequency dependence, all metric components (except g_{ff}) scale as f^2 , as the signal phase at the detector is $\phi(t) \sim 2\pi f \tau(t)$ and the metric (13) is quadratic in phase. This scaling is similar to the metric over the sky position parameters and may be mitigated in the same way (e.g., [66]): a full search is typically broken into smaller work units distributed over nodes of a cluster (or Einstein@Home), where each work unit would analyze a relatively narrow frequency band $\lesssim \mathcal{O}(1 \text{ Hz})$. We can therefore deal with the frequency dependence by simply evaluating the metric at a fixed frequency within each narrow range, typically at the highest frequency to guarantee the given maximum-mismatch constraint over the search band, accepting small relative changes of the mismatch distribution over the frequency band.

A similar argument applies to t_{asc} in the long-segment regime: due to the gauge freedom Eq. (19), the maximal physical uncertainty for any system would be $\Delta t_{\text{asc}} < P_{\text{orb}} \ll \Delta T$ and can therefore be neglected in $g_{\Omega\Omega}$, as seen in Eq. (20), which is the only metric term that would be affected by this.

Given the narrow astrophysical uncertainties on P_{orb} for Sco X-1 (cf. Sec. II F), this approach could also be used for Ω , but it would be very specific to Sco X-1 and might not apply to other directed binary searches. Furthermore, ignoring the metric changes over the astrophysical range on a_p would not work well for Sco X-1, given that the currently uncertainty spans more than a factor of 2.

We observe that the metric Eq. (18) depends on a_p and Ω only via quadratic *scaling* of some components; i.e., the metric stretches or contracts along certain directions in parameter space. In order to absorb this scaling, we only need to assume that the metric change is negligible on the scale $\delta\lambda$ of a lattice cell, so we can resort to *local* rescaling via the following “pseudo” coordinate transformation of $\{\Omega, t_{\text{asc}}, \kappa, \eta\}$ into

$$\begin{aligned}
 v_p &\equiv \underline{a}_p \Omega, \\
 d_{\text{asc}} &\equiv \underline{a}_p \underline{\Omega} t_{\text{asc}}, \\
 \kappa_p &\equiv \underline{a}_p \kappa, \\
 \eta_p &\equiv \underline{a}_p \eta,
 \end{aligned} \tag{21}$$

where \underline{a}_p and $\underline{\Omega}$ will be treated as constant scaling parameters in derivatives. Substituting the new coordinates in Eq. (16) results in the (orbital) phase model

$$\frac{\phi_{\text{orb}}(t; \lambda)}{2\pi} = -f a_p \left(\sin \Psi + \frac{\kappa_p}{2a_p} \sin 2\Psi - \frac{\eta_p}{2a_p} \cos 2\Psi \right),$$

$$\Psi(t) = \frac{v_p}{a_p} \left(t - \frac{d_{\text{asc}}}{a_p \Omega} \right), \quad (22)$$

and the following *approximate* phase derivatives:

$$\begin{aligned} \partial_{v_p} \phi &= -2\pi f (t - t_{\text{asc}}) [\cos \Psi + \kappa \cos 2\Psi + \eta \sin 2\Psi], \\ \partial_{d_{\text{asc}}} \phi &= 2\pi f [\cos \Psi + \kappa \cos 2\Psi + \eta \sin 2\Psi], \\ \partial_{\kappa_p} \phi &= -\pi f \sin 2\Psi, \\ \partial_{\eta_p} \phi &= \pi f \cos 2\Psi. \end{aligned} \quad (23)$$

Applying the steps of Sec. II C this yields the following metric components [with \hat{g}_{ff} and $\hat{g}_{a_p a_p}$ unchanged from Eq. (18)]:

$$\begin{aligned} \hat{g}_{v_p v_p} &= 2\pi^2 f^2 \left(\frac{\Delta T^2}{12} + \overline{\Delta_{\text{ma}}^2} \right), \\ \hat{g}_{d_{\text{asc}} d_{\text{asc}}} &= 2\pi^2 f^2, \\ \hat{g}_{v_p d_{\text{asc}}} &= \hat{g}_{d_{\text{asc}} v_p} = -2\pi^2 f^2 \overline{\Delta_{\text{ma}}}, \\ \hat{g}_{\kappa_p \kappa_p} &= \hat{g}_{\eta_p \eta_p} = \frac{\pi^2}{2} f^2, \end{aligned} \quad (24)$$

which are constant over a_p and Ω and are therefore suitable for lattice tiling within the WEAVE framework.

We are applying the coordinate transformation Eq. (21) globally over the search parameter space, but we ignore the local changes in a_p , Ω -scaling within each lattice cell. This should be a good approximation as long as cells are small compared to the effects of changing a_p , Ω over their respective length scales.

We have thoroughly tested the safety and effectiveness of this metric approximation for a Sco X-1 search, which is discussed in the next section.

IV. TESTING AND CHARACTERIZATION

The semicoherent $\hat{\mathcal{F}}$ -statistic in BINARYWEAVE is computed by the well-tested WEAVE framework [37] using the standard LALSuite [67] \mathcal{F} -statistic implementation. The behavior of this statistic implementation in recovering signals in noise is therefore already well understood and tested. Therefore, the only new elements requiring careful testing and characterization are the template-bank mismatch and the computing cost. For this reason, we have limited the mismatch characterization studies for the signal-only cases without introducing any kind of GW detector noises.

A. Test setup and assumptions

The metric and template bank implemented in BINARYWEAVE can, in principle, handle eccentricity within the small-eccentricity approximation $e \ll 1$ of Eq. (16), which in [36] was seen to hold up to about $e \lesssim 0.1$. The orbital eccentricity of Sco X-1 is currently poorly constrained but expected to be close to zero due to Roche-lobe overflow accretion [63]. In order to simplify this first proof-of-concept study of BINARYWEAVE, we assume negligible eccentricity here and focus on purely circular orbits. Therefore we consider a Sco X-1 search parameter space \mathcal{P} that is (at most) four dimensional (4D), with search parameters $\{f, a_p, P_{\text{orb}}, t_{\text{asc}}\}$.

The orbital period P_{orb} for Sco X-1 is constrained to about $\Delta P_{\text{orb}} \sim 0.04$ s, compared to a period of $P_{\text{orb}} \sim 19$ h (cf. Sec. II F). The search resolution $\delta\Omega$ (and therefore also δP_{orb}) in Eq. (11) is determined by the metric [in particular, $\hat{g}_{\Omega\Omega}$ of Eq. (20)] and therefore depends on the search setup $\{N, \Delta T, \mu_{\text{max}}\}$, the search frequency f , and semimajor axis a_p . In particular, the resolution increases linearly with total search duration $T_{\text{obs}} = N\Delta T$, and for longer-duration searches (e.g., $T_{\text{obs}} \sim 6$ months) it will often fully resolve the parameter-space uncertainty in the period, i.e., $\delta P_{\text{orb}} < \Delta P_{\text{orb}}$. However, for coarser search setups, or assuming future improved observational constraints, it can also be sufficient to place a single template at the midpoint of the uncertainty range, resulting in a three-dimensional (3D) search space \mathcal{P} spanning only $\{f, a_p, t_{\text{asc}}\}$. In the following we will therefore consider both possibilities of 3D and 4D template banks.

In this study we are exclusively using the noninterpolating StackSlide WEAVE mode, which is simpler and easier to optimize for, while expected to yield similar sensitivity for directed binary searches, as discussed in Sec. II E. This means that the coherent segments and final semicoherent statistic use the same template grid and there is only a single mismatch parameter $\hat{\mu}_{\text{max}} = \mu_{\text{max}}$.

All subsequent simulations use \mathcal{F} -statistic input data split into short Fourier transforms (SFT) [68] of baseline $T_{\text{sft}} \leq 250$ s, which is a safe SFT length over the Sco X-1 parameter space, e.g., see Eq. (C2) in [36]. Furthermore, all simulations assume data from two detectors, namely, LIGO Hanford (H1) and LIGO Livingston (L1).

B. Template-bank mismatch

In order to ensure the validity of the constructed lattice template banks using the approximately flat metric constructed in Sec. III, we perform injection-recovery Monte Carlo tests. These tests are typically performed without noise, i.e., searching a data stream only containing the injected signal waveform. This allows one to directly measure signal power without noise bias and to accurately calculate the mismatch, which is the main purpose of template-bank tests. The signal parameters for the injections are drawn uniformly from the (wider) testing Sco X-1

parameter space \mathcal{P}_0 specified in Table I, with randomly drawn amplitude parameters \mathcal{A} , and a search grid is constructed around the injection point (randomly shifted to avoid systematic alignment effects).

A good template bank should satisfy the maximal mismatch criterion (e.g., [48]): the measured mismatch $\mu_0(\lambda_s; \lambda_t)$ of Eq. (8) for any injected signal $\lambda_s \in \mathcal{P}$ at its “closest” (i.e., highest signal power ρ^2) template λ_t should be less than the maximum mismatch μ_{\max} the template bank was constructed for, which can formally be written as

$$\max_{\lambda_s \in \mathcal{P}} \min_{\lambda_t} \mu_0(\lambda_s; \lambda_t) \leq \mu_{\max}, \quad (25)$$

where the “minimax” formulation (constraining the nearest-template mismatch at the worst-case signal location) implies that the mismatch is constrained for all possible signal locations.

Furthermore, an *efficient* template bank should ideally place only a single template within μ_{\max} of any signal, to avoid (computationally wasteful) over-resolution and producing excessive candidates per signal that would require some form of clustering or follow-up (see also [38]).

We have performed a number of signal injection-recovery tests of the BINARYWEAVE template banks for various different search setups $\{N, \Delta T, \mu_{\max}\}$. Here we only present a few representative examples in order to illustrate the main features of these template banks: in Sec. IV B 1 we illustrate the template grids for single-parameter (1D) and two-parameter (2D) searches, and in Sec. IV B 2 we provide examples of the mismatch distribution for 3D searches (for a nonresolved period uncertainty ΔP_{orb}) and full 4D searches.

1. Testing 1D and 2D lattice tilings

In order to illustrate and visualize the lattice tiling, we first consider simple one- and two-dimensional lattice cases, which also serve as a basic sanity check for the template-bank construction. The 1D searches are performed along all four coordinate axes in a neighborhood around the signal injection, with the three remaining parameters fixed to the injection values, with one example shown in Fig. 1. The 2D searches are performed along all six two-parameter combinations out of the four, with the remaining two parameters fixed to the injected signal location, with one example shown in Fig. 2.

These results illustrate the maximum-mismatch criterion of Eq. (25) being satisfied, as well as placing only one template in the “vicinity” $< \mu_{\max}$ of the signal as desired for an efficient template bank.

2. Testing 3D and 4D lattice tilings

Next we test the template-bank performance for the four possible combinations of *three* search parameters (3D searches) with the fourth one fixed to the signal injection

parameter, as well as 4D searches over all four parameters $\{f, a_p, P_{\text{orb}}, t_{\text{asc}}\}$. We perform several sets of simulations, using $\sim \mathcal{O}(100 - 1000)$ injections each, using varying search setups and maximum-mismatch values μ_{\max} , in order to obtain the resulting mismatch distribution of the template bank.

The injected signal parameters are randomly drawn from the test range \mathcal{P}_0 (cf. Sec. I), namely, $f \in [10, 700]$ Hz, and binary parameter ranges wider than the Sco X-1 constraints, namely, $a_p \in [0.3 - 3.5]$ ls, $P_{\text{orb}} = 68\,023.7 \pm 0.2$ s and $t_{\text{asc}} = 1\,124\,044\,455.0 \pm 1000$ GPS s.

Figure 3 shows an example for the mismatch distributions of coherent and semicoherent mismatches obtained for a set of 1000 injections and subsequent 3D searches in a small box around the injection in f , a_p , and t_{asc} , with P_{orb} fixed to the injected value. Figure 4 presents a corresponding example for the mismatch distributions obtained from 1000 4D-box searches around the injected signals.

We see that the means of the coherent and semicoherent mismatch distributions are $\langle \mu \rangle \approx \langle \hat{\mu} \rangle \approx 0.17 - 0.18$, and the highest observed semicoherent mismatch in the 3D case is $\max \hat{\mu}_0 \approx 0.4$, while in the 4D case it is $\max \hat{\mu}_0 \approx 0.35$. This is smaller than the imposed maximum mismatch of $\mu_{\max} = 0.5$, which is a common feature of the quadratic approximation Eq. (9) underlying the metric; namely, the measured mismatch values μ_0 tend to increasingly fall behind the predicted metric mismatch values with increasing mismatch (e.g., see [47, 56, 69]). Thus, in addition, we also test the metric mismatch implementations for a small mismatch value $\mu_{\max} = 0.05$ which is comparable to the realistic search setups relevant for Sco X-1 (discussed in detail in Sec. V). We see a good agreement for both 3D and 4D template banks with such small μ_{\max} values as shown in Fig. 5.

C. Required computing resources

1. Number of templates

As discussed in Sec. II C, the *bulk* template count for a parameter space \mathcal{P} (not counting any extra templates required for boundary padding of $\partial\mathcal{P}$) is given by Eq. (10).

Using the metric expressions in Eq. (14), this can be evaluated explicitly [36], and the bulk template count for 3D searches over $\{f, a_p, t_{\text{asc}}\}$ is found as

$$\begin{aligned} \hat{N}_{\text{3D}} &= \frac{\theta_3}{\mu_{\max}^{3/2}} \frac{\pi^3 \Delta T}{\sqrt{27}} \Omega(f_{\max}^3 - f_{\min}^3) (a_{p,\max}^2 - a_{p,\min}^2) \\ &\quad \times (t_{\text{asc},\max} - t_{\text{asc},\min}), \end{aligned} \quad (26)$$

while for a 4D template bank over $\{f, a_p, P_{\text{orb}}, t_{\text{asc}}\}$, one finds

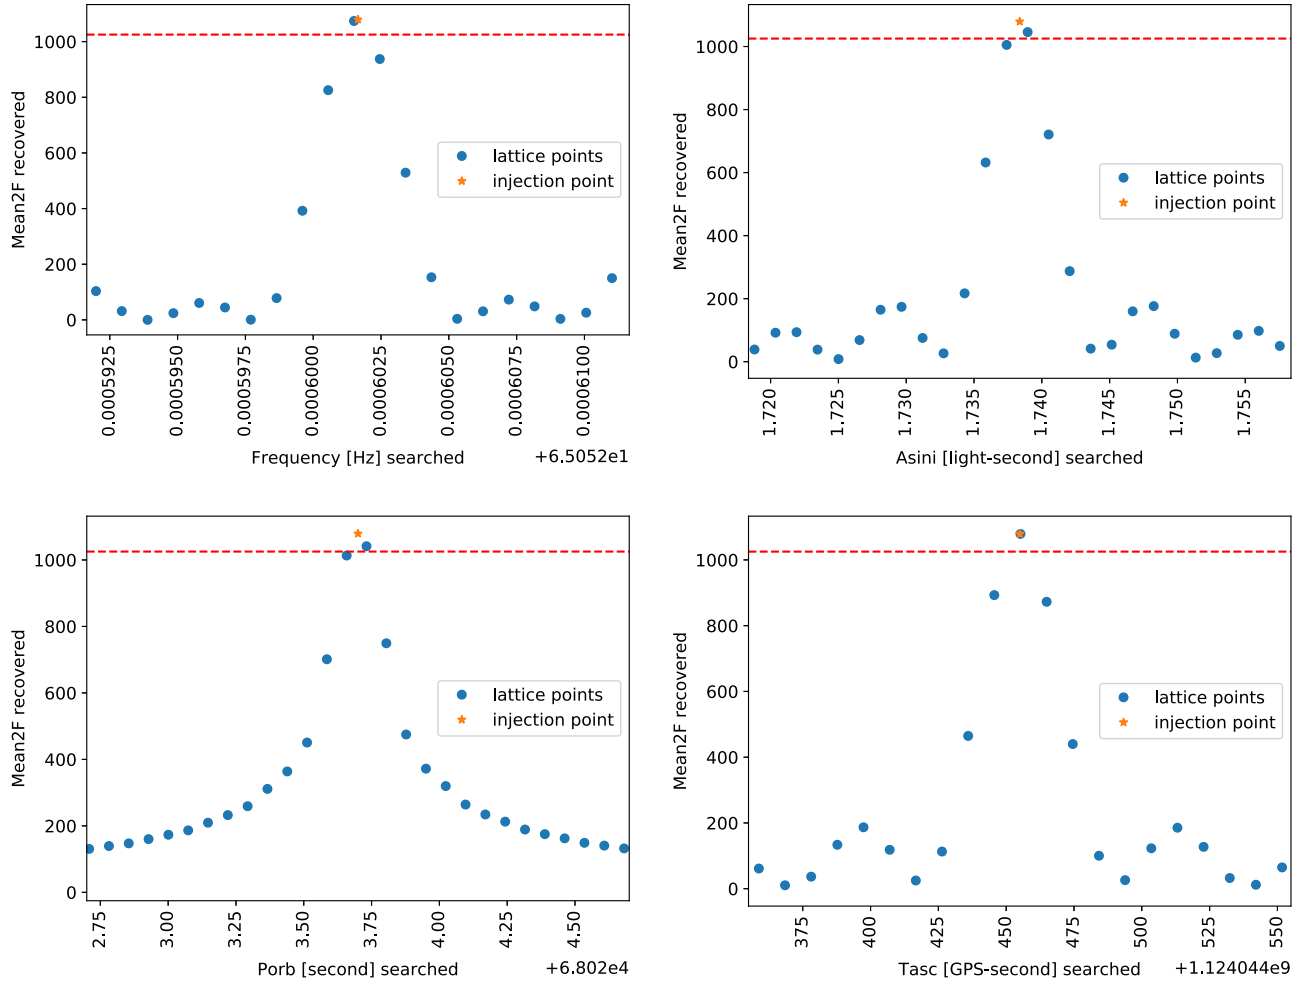


FIG. 1. Illustration of 1D template-bank searches around a noiseless signal injection, with the respective three remaining search parameters fixed to the injected signal. The filled circles mark the placement of templates and their corresponding measured $\hat{\mathcal{F}}$ -statistic values, while the star marks the signal injection point with its corresponding perfect-match $\hat{\mathcal{F}}$ -statistic. The template bank was constructed for a maximum mismatch of $\mu_{\max} = 0.05$, with $N = 120$ segments of $\Delta T = 3$ d. The horizontal dashed line denotes the $\hat{\mathcal{F}}$ -value corresponding to the maximum-mismatch criterion Eq. (25) relative to the injected signal power.

$$\hat{\mathcal{N}}_{4D} = \frac{\theta_4}{\mu_{\max}^2} \frac{\pi^4 \gamma \Delta T^2}{36\sqrt{2}} (f_{\max}^4 - f_{\min}^4) (a_{p,\max}^3 - a_{p,\min}^3) \times (\Omega_{\max}^2 - \Omega_{\min}^2) (t_{\text{asc},\max} - t_{\text{asc},\min}), \quad (27)$$

where the coordinate ranges are $\lambda^i \in [\lambda_{\min}^i, \lambda_{\max}^i]$, and γ is the semicoherent refinement factor associated with the P_{orb} (i.e., Ω), given by

$$\gamma = \sqrt{1 + 12 \frac{(\bar{\Delta}_{\text{ma}}^2 - \overline{\Delta_{\text{ma}}^2})}{\Delta T^2}}. \quad (28)$$

The refinement factor evaluates to $\gamma = N$ in the case of segments without gaps. We can use these theoretical expressions to test against the actual number of templates generated by the BINARYWEAVE code, which includes

boundary padding not accounted for in the above theoretical expressions.

In the following we consider two example search setups (cf. Table II), namely, search setup-I with $N = 180$ segments of duration $\Delta T = 1$ d and a maximum mismatch of $\mu_{\max} = 0.031$, and search setup-II with $N = 120$ segments of $\Delta T = 3$ d and maximum mismatch $\mu_{\max} = 0.056$.

We generate a BINARYWEAVE template bank for a small box around a randomly chosen point in f and a_p , drawn from the test range \mathcal{P}_0 of Table I. The box consists of $\mathcal{O}(10^5)$ frequency bins and a metric bounding-box extent $D\lambda^i$ [cf. Eq. (12)] along each binary-orbital parameter dimension. This is repeated 40 times, in order to obtain a representative sampling over a wide range of search parameters, and the resulting BINARYWEAVE template counts are compared to the theoretical predictions of (27), shown in Fig. 6.

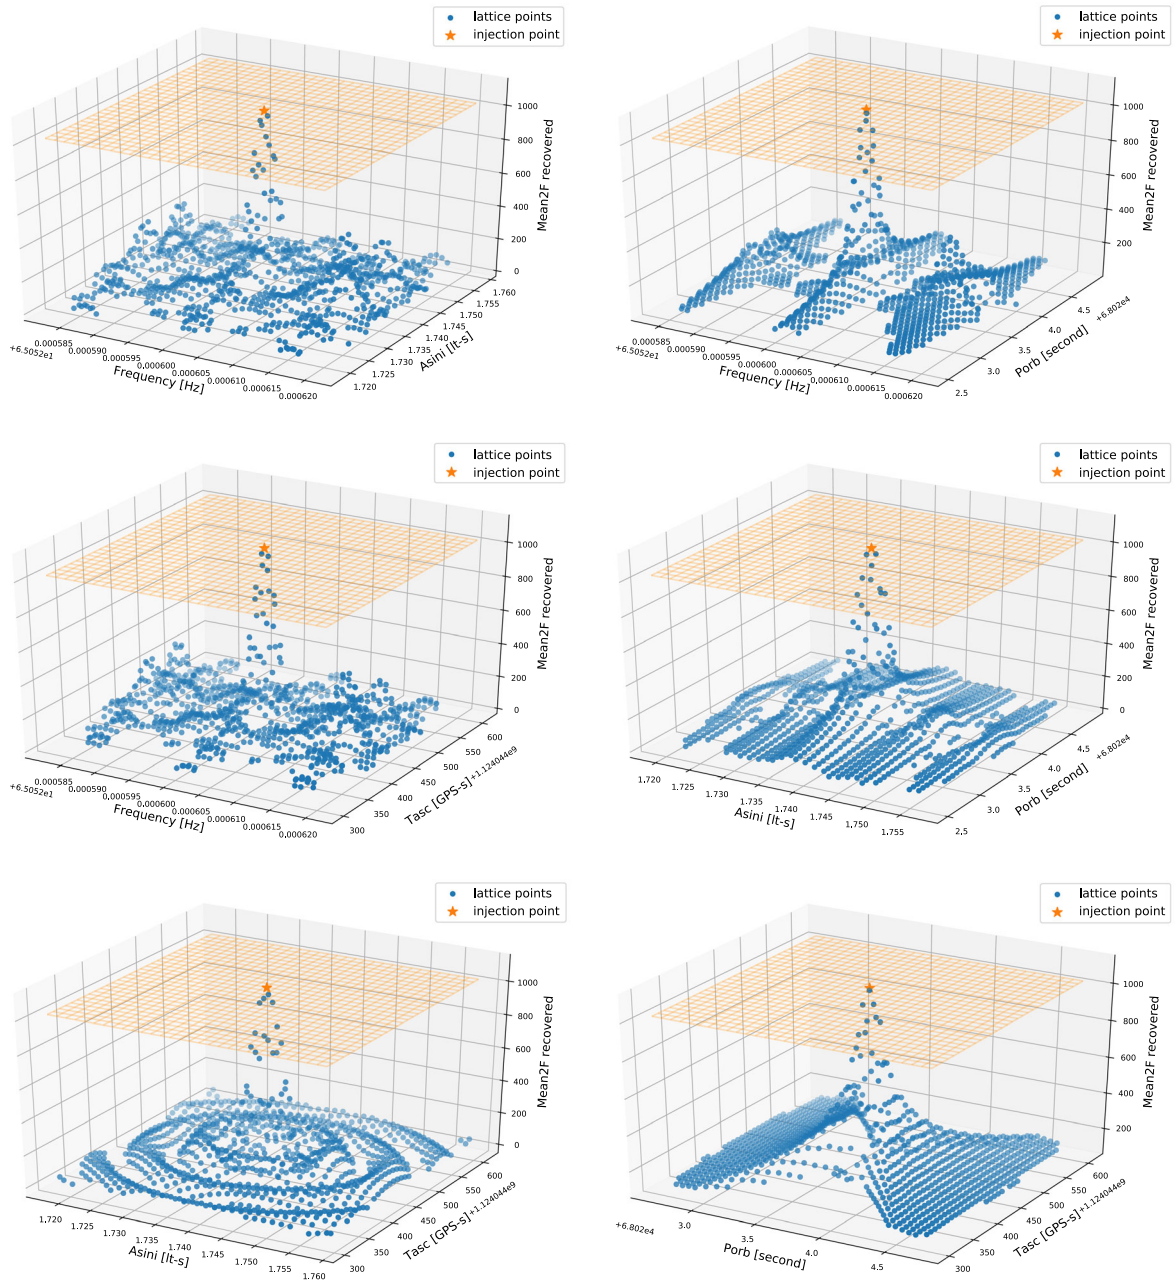


FIG. 2. Illustration of 2D template-bank searches around a noiseless signal injection, with the respective two remaining search parameters fixed to the injected signal. The filled circles mark the placement of templates and their corresponding measured $\hat{\mathcal{F}}$ -statistic values, while the star marks the signal injection point with its corresponding perfect-match $\hat{\mathcal{F}}$ -statistic. The template bank was constructed for a maximum mismatch of $\mu_{\max} = 0.05$, with $N = 120$ segments of $\Delta T = 3$ d. The horizontal mesh grid denotes the $\hat{\mathcal{F}}$ -value corresponding to the maximum-mismatch criterion Eq. (25) relative to the injected signal power.

We see that there is generally good agreement in the template counts, with the real template counts exceeding the theoretical *bulk* predictions by factors up to 2–3 at low template counts, with increasingly good agreement at higher template counts. The template counts exceed only at the lowest frequency regime (≤ 50 Hz) by a factor of ~ 2 –3, whereas they agree within 10% at intermediate

frequency (~ 200 Hz) and $< 5\%$ at higher frequency (~ 500 Hz). This effect is expected from the extra padding required to fully cover the parameter-space boundaries $\partial\mathcal{P}$, which decreases in relative importance for increasing total template counts (i.e., boundary effects are less important for template spacings that are small compared to the parameter-space extents).

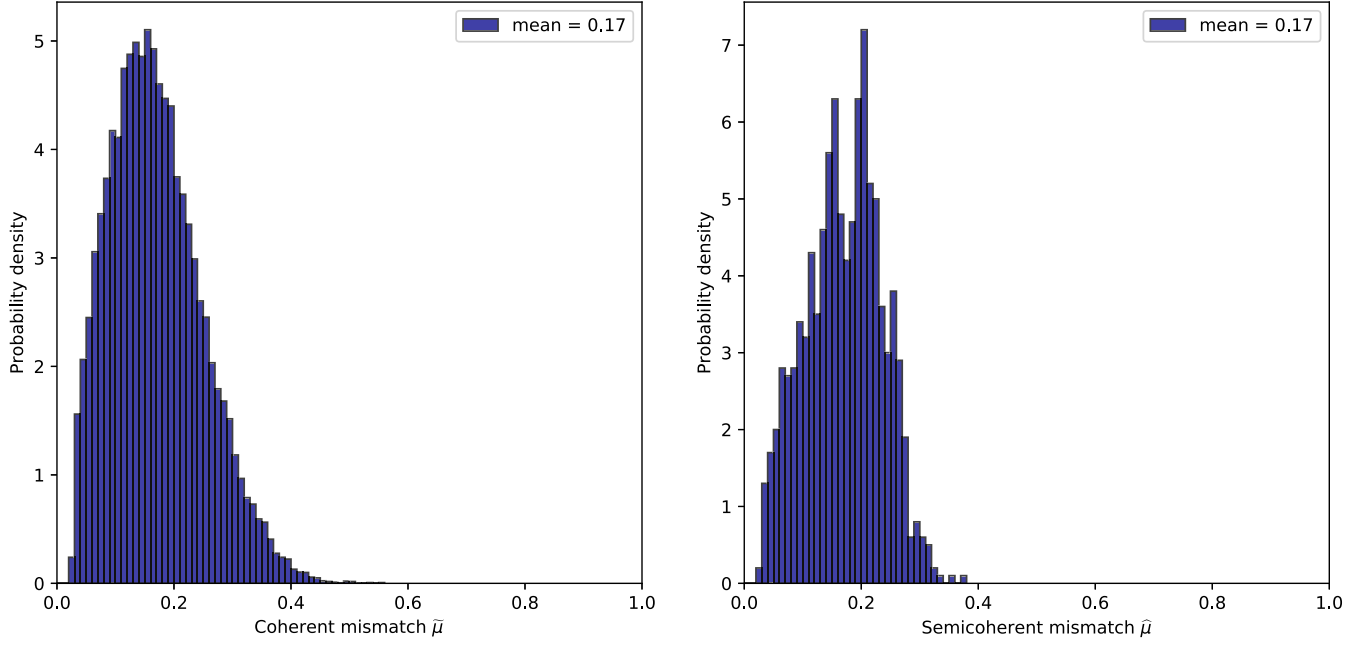


FIG. 3. Distribution of coherent per-segment mismatches μ_0 (left plot) and semicoherent mismatches $\hat{\mu}_0$ (right plot), obtained from 1000 simulated 3D searches over a small box in f , a_p , and t_{asc} around the injected signals (with P_{orb} fixed at its injection value), with parameters drawn randomly from the test range \mathcal{P}_0 defined in Table I. The template bank was constructed for a maximum mismatch of $\mu_{max} = 0.5$, with $N = 30$ segments of $\Delta T = 1$ d.

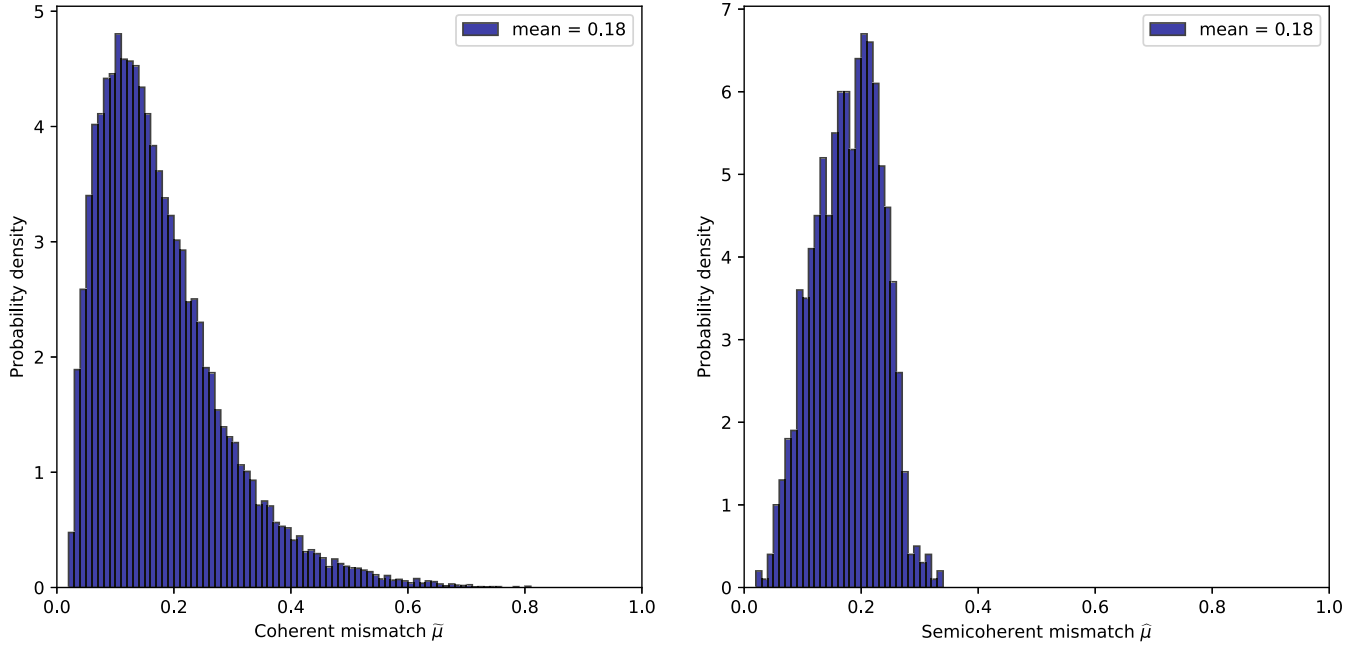


FIG. 4. Distribution of coherent per-segment mismatches μ_0 (left plot) and semicoherent mismatches $\hat{\mu}_0$ (right plot), obtained from 1000 simulated 4D searches over a small box in f , a_p , t_{asc} , and P_{orb} around the injected signals, with parameters drawn randomly from the test range \mathcal{P}_0 defined in Table I. The template bank was constructed for a maximum mismatch of $\mu_{max} = 0.5$, with $N = 30$ segments of $\Delta T = 1$ d.

2. Computing cost and memory usage

A detailed computing-cost (and memory) model exists for the semicoherent WEAVE implementation [37] as well as

for the underlying coherent \mathcal{F} -statistic implementation [70]. There are two different \mathcal{F} -statistic algorithms available, the resampling FFT algorithm (originally described

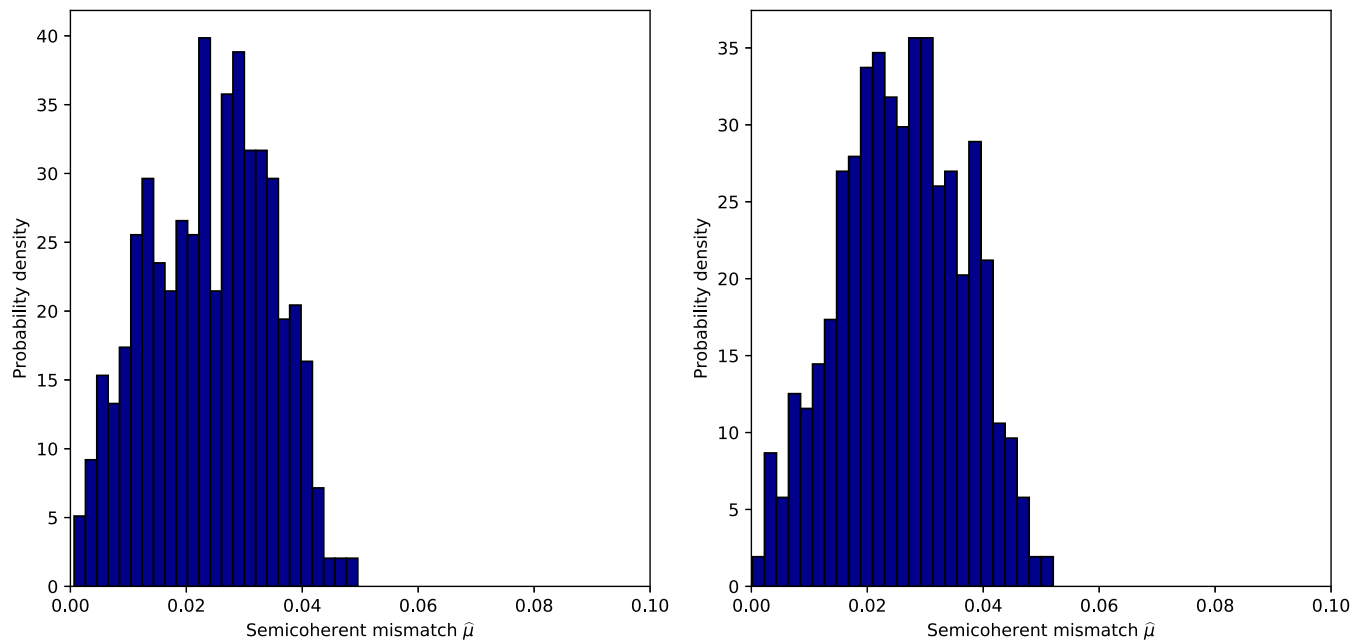


FIG. 5. Distribution of semicoherent mismatches $\hat{\mu}_0$ obtained from 500 simulated 4D searches (left plot) over a small box in f , a_p , and t_{asc} around the injected signals (with P_{orb} fixed at its injection value); and distribution of 500 simulated 4D searches (right plot) over a small box in f , a_p , t_{asc} , and P_{orb} around the injected signals, with parameters drawn randomly from the test range \mathcal{P}_0 defined in Table I. The template bank was constructed for a maximum mismatch of $\mu_{\text{max}} = 0.05$, with $N = 120$ segments of $\Delta T = 3$ d.

in [43]) and the so-called demodulation algorithm introduced in [68,71]. Because the resampling \mathcal{F} -statistic is substantially faster [i.e., $\mathcal{O}(100 - 1000)$] for large numbers of frequency bins [i.e., $\mathcal{O}(10^5)$] and SFTs, which is the relevant regime for the wide parameter-space search considered here, we will exclusively consider this algorithm for the following discussion of the Sco X-1 computing cost.¹

We performed the BINARYWEAVE tests and simulations on the LIGO Data Analysis System (LDAS) computing cluster at the LIGO Hanford Observatory, containing a combination of 2.4 GHz Xeon E5-2630v3, 2.2 GHz Xeon E5-2650v4, 3.5 GHz Xeon E3-1240v5, and 3.0 GHz Xeon Gold 6136 CPUs. We find that the resulting semicoherent timing coefficients measured on this hardware are essentially the same as given in Table III of [37], while the effective (resampling-FFT) \mathcal{F} -statistic time per template and detector is observed to fall in the range $\tau_{\mathcal{F}}^{\text{eff}} \approx (3.8 - 4.3) \times 10^{-7}$ s, consistent with the numbers obtained in [37].

We measure the CPU runtime per template \mathcal{C}_t and the maximum memory usage of BINARYWEAVE for the 80 box searches (two sets of 40 box searches each for setup-I and setup-II) described in the previous section (see Fig. 6). The maximum memory usage over all search boxes is found as ~ 2.2 GB, well below all-sky WEAVE numbers observed in [38], due to the fact that Sco X-1 has little refinement and

we can use a noninterpolating search setup, substantially alleviating memory requirements.

The runtime per template \mathcal{C}_t is found to be relatively constant over the search parameter space and for the two search setups considered, namely, $\mathcal{C}_t(\text{Searchsetup}) \approx 0.12 + -0.03$ ms and $\mathcal{C}_t(\text{Searchsetup}) \approx 0.14 + -0.03$ ms. Here we only consider the noninterpolating StackSlide method, in which the coherent segments and the semicoherent $\hat{\mathcal{F}}$ -statistic share the same template grid and number of templates \mathcal{N} , i.e., $\mathcal{N} = \hat{\mathcal{N}}$. This implies that both the coherent and semi-coherent contributions to the total computing cost are proportional to \mathcal{N} . Therefore we can use a simplified *effective* model for the *total* computing cost $\mathcal{C}_{\mathcal{P}}$ over a search space \mathcal{P} in the form

$$\mathcal{C}_{\mathcal{P}} = \mathcal{N}_{\mathcal{P}} \mathcal{C}_t, \quad (29)$$

where $\mathcal{N}_{\mathcal{P}}$ is the total number of templates covering the parameter space \mathcal{P} . Given the above timing measurements for the two setups, in the following we assume a (slightly conservative) effective CPU time per template of $\mathcal{C}_t = 0.145$ ms. This simplified effective cost model is plotted against the measured BINARYWEAVE runtimes in Fig. 7.

V. CHARACTERIZING POTENTIAL SCO X-1 SEARCHES

A. Sensitivity for different search setups

The sensitivity of a search is typically characterized by the weakest signal amplitude $h_{\text{fa}}^{\text{det}}$ detectable at a

¹A GPU port of the resampling \mathcal{F} -statistic [72], which yields speedup factors of $\mathcal{O}(10 - 100)$, was developed after this study was performed. A practical application of the GPU resampling \mathcal{F} -statistic with WEAVE can be found in [57].

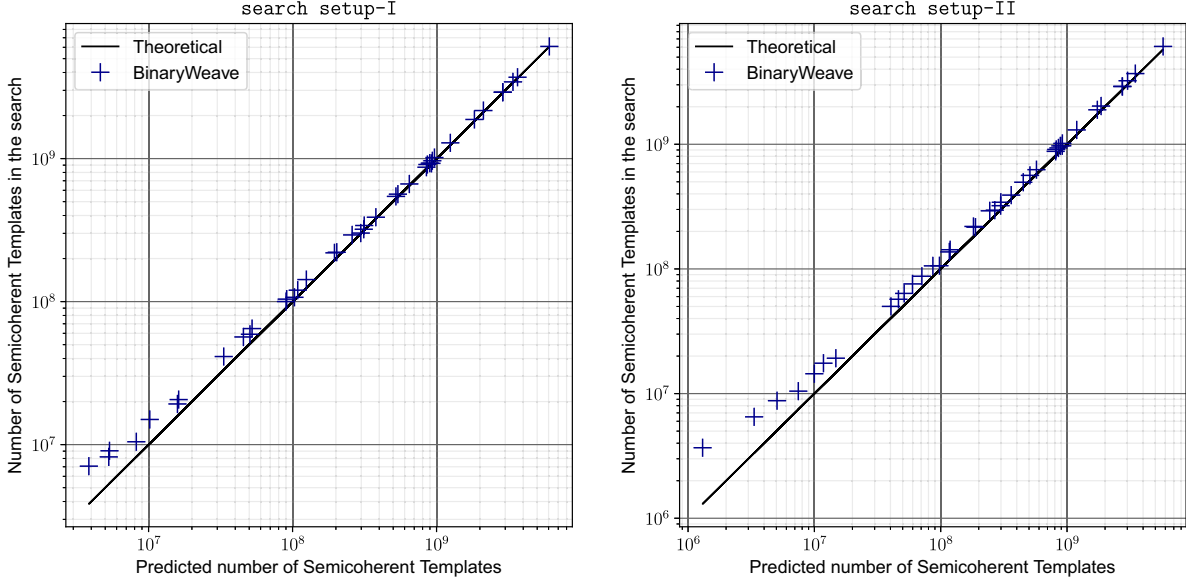


FIG. 6. Number of semicoherent templates \hat{N} constructed by BINARYWEAVE versus with the theoretical *bulk* predictions of Eq. (27). Each point ‘+’ corresponds to a simulated 4D-box search around a randomly chosen parameter-space location in $\{f, a_p\} \in \mathcal{P}_0$ [cf. Table I using either search setup-I (left plot) or search setup-II (right plot) defined in Table II].

false-alarm probability p_{fa} with detection probability (or “confidence level”) p_{det} . While this is astrophysically informative, for a given search method it is often more instructive [73] to use the sensitivity depth $\mathcal{D}_{p_{\text{fa}}}^{p_{\text{det}}}$ instead, defined as

$$\mathcal{D}_{p_{\text{fa}}}^{p_{\text{det}}} \equiv \frac{\sqrt{S_n}}{h_{p_{\text{fa}}}^{p_{\text{det}}}}, \quad (30)$$

which characterizes the sensitivity of a method independently of the noise floor (i.e., power spectral density) S_n .

As discussed in [73,74], the sensitivity of a semicoherent StackSlide $\hat{\mathcal{F}}$ -statistic search can be estimated quite accurately (to better than $\sim 10\%$) given the total amount of data used, the number N of semicoherent segments, and the mismatch distribution of the template bank. This algorithm is implemented in the OCTAPPS [75] function `SensitivityDepthStackSlide()`.

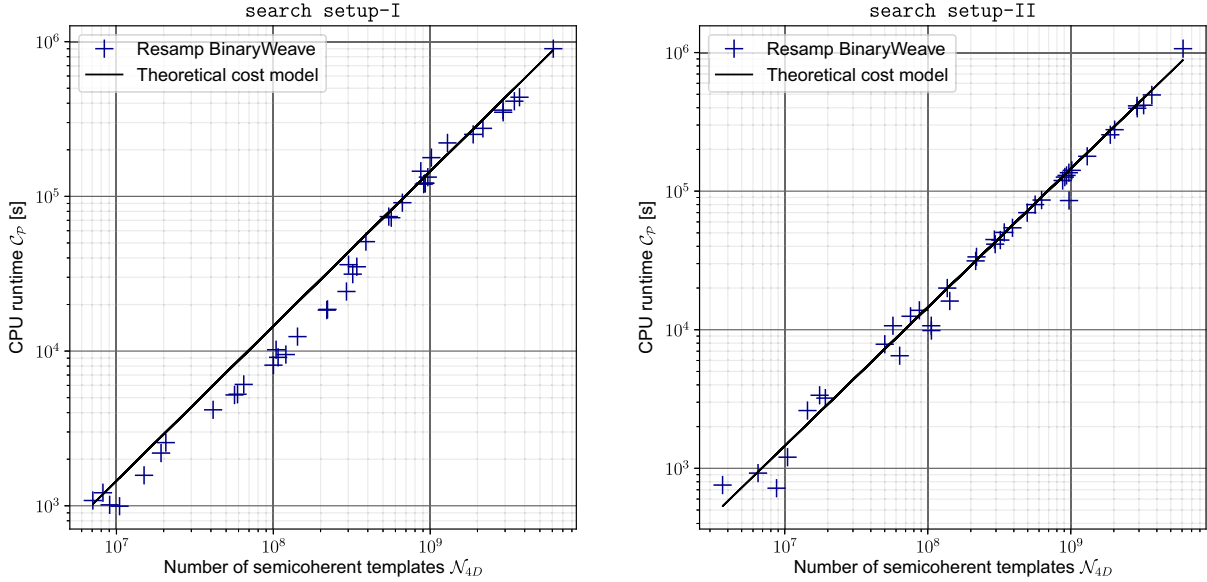


FIG. 7. CPU runtime C_P per search box as a function of the number of (semicoherent) templates \hat{N}_{4D} for that box, for search setup-I (left plot) and search setup-II (right plot), defined in Table II. The points ‘+’ mark the measured BINARYWEAVE runtimes, while the solid line indicates the effective cost model prediction of Eq. (29), using an effective cost per template of $C_t = 0.145$ ms.

TABLE II. Definition of example search setups with corresponding estimated sensitivity depth, discussed in Sec. VA. The sensitivity estimates assume a (per-template) false-alarm probability of $p_{\text{fa}} = 10^{-10}$ and detection confidences $p_{\text{det}} = 90\%$, 95% , and 99% , respectively, using the measured (4D) mismatch distributions obtained for each setup (cf. Sec. IV B 2).

Search setup	T_{obs} [months]	ΔT [days]	N	μ_{max}	$\mathcal{D}_{p_{\text{fa}}}^{90\%}$ [$1/\sqrt{\text{Hz}}$]	$\mathcal{D}_{p_{\text{fa}}}^{95\%}$ [$1/\sqrt{\text{Hz}}$]	$\mathcal{D}_{p_{\text{fa}}}^{99\%}$ [$1/\sqrt{\text{Hz}}$]
search setup-I	6	1	180	0.031	77	72	60
search setup-II	12	3	120	0.056	116	107	91
search setup-III	6	3	60	0.025	96	89	75
search setup-IV	12	1	360	0.025	93	86	73
search setup-V	6	10	18	0.025	120	111	94
search setup-VI	12	10	36	0.025	150	138	117

For each search setup listed in Table II, we obtain the mismatch distribution empirically by injection-recovery Monte Carlo simulation (cf. Sec. IV B 2) and use this to estimate the expected sensitivity depth for each setup. We use a canonical value of $p_{\text{fa}} = 10^{-10}$ (as was done in [36]) for the single-template false-alarm probability, which represents a somewhat typical false-alarm scale for wide parameter-space searches. We quote the sensitivity depth for $p_{\text{det}} = 90\%$, 95% , and 99% . The former two are typical confidence levels used for upper limits obtained in CW searches, while the last one might be interesting, for example, if one wants to reject the torque-balance hypothesis or a specific emission mechanism in some parameter range at high confidence.

In Table II, we summarize the sensitivity depths for a set of six different search setups. The sensitivity depths obtained from the empirical mismatch distributions corresponding to the well-studied setup-I and setup-II are presented in this table. In addition, we report the maximum achievable sensitivity depths for this BINARYWEAVE pipeline estimated from our simulated searches for four different setups that may be relevant for different cases of unknown spin wandering effects in Sco X-1.

B. Computing cost for different search scenarios

Here we present the CPU computing cost in terms of core hours, and million core hours (Mh), referring to the mix of CPU hardware used in the present study, cf. Sec. IV C 2. Another interesting unit used in [36] is Einstein@Home months (EM), which was defined as 12 000 (average) CPU cores running on Einstein@Home [39] for 30 days. If one assumes the (current) average Einstein@Home CPU to be roughly comparable to the one used here, one can convert $1 \text{ EM} \approx 8.6 \text{ Mh}$.

Let us first consider the example of the Sco X-1 parameter space \mathcal{P}_1 considered in Leaci and Prix [36] (cf. Table I) with two different search setups (I and II) of Table II. For search setup-I with $180 \times 1 \text{ d}$ segments and mismatch $\mu_{\text{max}} = 0.031$, the total number of (4D) templates given by Eq. (27) is $\mathcal{N}_{4\text{D}} = 5.84 \times 10^{14}$. Using the effective computing-cost model of Eq. (29), this results in a total CPU runtime of $\mathcal{C}_{\mathcal{P}_1}[\text{Search setup-I}] \approx 8.46 \times 10^{10} \text{ s} = 23.5 \text{ Mh}$. Using the above conversion factors, this would correspond

to 2.7 EM. Similarly, for search setup-II with $120 \times 3 \text{ d}$ segments and a mismatch of $\mu_{\text{max}} = 0.056$, we obtain a template count of $\mathcal{N}_{4\text{D}} = 1.07 \times 10^{15}$ and a corresponding total CPU runtime of $\mathcal{C}_{\mathcal{P}_1}[\text{Search setup-II}] \approx 1.56 \times 10^{11} \text{ s} = 43.2 \text{ Mh}$, which we can also express as 5.0 EM.

Next we consider a number of additional parameter-space scenarios, listed in Table I. The constraints from optical and radio emission observations come from different sources in the literature [18,62], with the most recent values given in [63,76]. Future observations are likely to further alter and improve these constraints. For a fully resolved period uncertainty, the total number of templates (and therefore computing cost) scales as a_{pmax}^3 for a wide parameter uncertainty in a_{p} [cf. Eq. (27)] but only as $a_{\text{pmax}}^2 \Delta a_{\text{p}}$ for narrow parameter uncertainty Δa_{p} .

In order to quantify the effects of future improved constraints on a_{p} , we consider three different scenarios: (i) $a_{\text{p}} \in [1.45, 3.25] \text{ ls}$ (search spaces $\mathcal{P}_4 - \mathcal{P}_9$), (ii) $a_{\text{p}} \in [1.40, 1.50] \text{ ls}$ (search spaces $\mathcal{P}_{10} - \mathcal{P}_{15}$), and (iii) $a_{\text{p}} \in [1.44, 1.45] \text{ ls}$ (search spaces $\mathcal{P}_{16} - \mathcal{P}_{21}$). Similarly, we consider six different frequency search ranges, three ‘‘deep-search’’ ranges covering only 100 Hz at different frequencies (600–700 Hz, 1000–1100 Hz, and 1400–1500 Hz), and three ‘‘broad-search’’ ranges within the LIGO/Virgo frequency band (20–500 Hz, 20–1000 Hz, and 20–1500 Hz). Finally, we consider both a 3D search (for an unresolved period uncertainty ΔP_{orb}) and a 4D search for all cases considered.

The resulting computing-cost estimates for all combinations of the two setups (I and II), 3D or 4D template bank, and different parameter spaces \mathcal{P}_{1-21} are given in Table III. We note that while some required computing budgets may seem unrealistically large, a recent GPU port of the \mathcal{F} -statistic and WEAVE [57,72] may yield speedup factors of tens to hundreds, making many more setups fall within reach of currently available computing resources.

C. Sensitivity versus computing cost

In addition to considering various fixed search scenarios as in the previous two subsections, it is also instructive to study how the achievable sensitivity varies as a function of

TABLE III. Computing-cost estimates \mathcal{C}_p (in million core hours [Mh]) for different parameter spaces \mathcal{P}_n defined in Table I. We consider two setups, search setup-I and search setup-II of Table II, assuming either a 3D or 4D template bank.

	(I,3D)	(I,4D)	(II,3D)	(II,4D)
\mathcal{P}_1	3.18	23.51	3.93	43.23
\mathcal{P}_2	28.50	466.48	35.22	857.69
\mathcal{P}_3	5.00	63.40	6.17	116.57
\mathcal{P}_4	26.38	577.57	32.60	1061.95
\mathcal{P}_5	68.76	2425.79	84.96	4460.17
\mathcal{P}_6	131.09	6381.48	161.97	11733.30
\mathcal{P}_7	3.24	20.42	4.01	37.54
\mathcal{P}_8	207.74	5226.87	256.69	9610.37
\mathcal{P}_9	701.14	26461.02	866.33	48652.49
\mathcal{P}_{10}	0.90	11.65	1.12	21.42
\mathcal{P}_{11}	2.36	48.94	2.91	89.97
\mathcal{P}_{12}	4.49	128.73	5.55	236.70
\mathcal{P}_{13}	0.11	0.41	0.14	0.76
\mathcal{P}_{14}	7.12	105.44	8.80	193.87
\mathcal{P}_{15}	24.03	533.80	29.70	981.46
\mathcal{P}_{16}	0.09	1.16	0.11	2.13
\mathcal{P}_{17}	0.23	4.86	0.29	8.93
\mathcal{P}_{18}	0.45	12.78	0.55	23.50
\mathcal{P}_{19}	0.01	0.04	0.01	0.08
\mathcal{P}_{20}	0.71	10.47	0.88	19.25
\mathcal{P}_{21}	2.40	52.99	2.96	97.43

the invested computing cost. This would generally involve a (3- or 4-dimensional) optimization problem over all search-setup parameters (see [36,40]) which is beyond the scope of this study, so we consider a simpler problem of varying the maximal template-bank mismatch μ_{\max} . In a sense, this provides a lower limit on the achievable sensitivity at any given cost, as one could always improve sensitivity further by varying all three setup parameters $\{\mu_{\max}, N, \Delta T\}$ at a fixed cost.

The search space is chosen as \mathcal{P}_2 , and we use again search setup-I (i.e., 180×1 d segments) and search setup-II (i.e., 180×3 d) as baselines, but now we vary the maximal template-bank mismatch in the range $0.025 \leq \mu_{\max} \leq 2.5$. For each mismatch, we can estimate the number of templates $\mathcal{N}_{4D} \propto \mu_{\max}^{-2}$ via Eq. (27) and obtain the corresponding computing cost \mathcal{C} from the simplified cost model Eq. (29). We use the corresponding theoretical mismatch distribution² for the A_n^* lattice, as well as the measured distribution from a set of 100 injection-recovery simulations using BINARYWEAVE, to estimate the expected sensitivity depth via `SensitivityDepthStackSlide()` from OCTAPPS.

This allows us to plot sensitivity depth versus computing cost, parametrized along μ_{\max} at fixed segment setup

²This will be a conservative overestimate of the mismatch (see Sec. IV B 2) and therefore an underestimate of the sensitivity.

$N \times \Delta T$, which is shown in Fig. 8. As expected, sensitivity improves as the invested computational cost increases and (equivalently) the maximum mismatch decreases; for $\mu_{\max} \lesssim 0.1$, however, further gains in sensitivity are minimal. We observe good agreement at small mismatches (i.e., large computing costs) between the theoretical estimates (using expected lattice mismatch distributions) and estimates using the measured mismatch distributions. The small loss of the measured \mathcal{F} -statistic value versus expected sensitivity in this regime from (well-known) additional intrinsic losses [$\hat{a}^{1/4}\mathcal{O}(1-3\%)$] of the high-performance \mathcal{F} -statistic implementation is compared to the exact calculation. At higher mismatches μ_{\max} , the measured mismatches tend to be smaller than the metric predictions, due to neglected higher-order terms in the metric approximation, as discussed previously in Secs. II C and IV B 2. This explains the measured sensitivity decreasing more slowly compared to the theoretical estimates at higher mismatches (i.e., smaller computing cost).

VI. SUMMARY AND OUTLOOK

In this paper, we presented the implementation and characterization of BINARYWEAVE, a new semicoherent search pipeline for CWs from neutron stars in binary systems with known sky position. This pipeline is based on the WEAVE framework [37], initially developed for all-sky searches of isolated sources, using the well-established semicoherent StackSlide $\hat{\mathcal{F}}$ -statistic.

The WEAVE framework requires a constant metric over the search parameter space for lattice tiling, and in order to apply the nonconstant binary metric of Leaci and Prix [36], we needed to develop a new internal coordinate system in which a constant approximation to the binary metric can be obtained. This is the basis for the BINARYWEAVE implementation. We performed extensive Monte Carlo tests for the safety (in terms of mismatches) of the resulting template banks and their template counts versus theoretical model expectations. Furthermore, we obtained a simplified timing model for the noninterpolating StackSlide mode used here, which allows easy estimates for the required computing cost of a given search, based on the known analytic template-count models.

Putting these pieces together, we illustrated expected sensitivity depths for BINARYWEAVE assuming different search setups, and we estimated the corresponding required computing costs for a number of different Sco X-1 parameter-space regions of interest.

Two other primary pipelines, CrossCorr and Viterbi, are presently used for searching CW signals from Sco X-1. The Viterbi pipeline aims to track the stochastic phase evolution model due to the spin-wandering effect of the neutron star in Sco X-1. It is thus more robust against this effect. The computational cost is also much less compared to the most sensitive searches of BINARYWEAVE. However, the maximum achievable sensitivity depth for

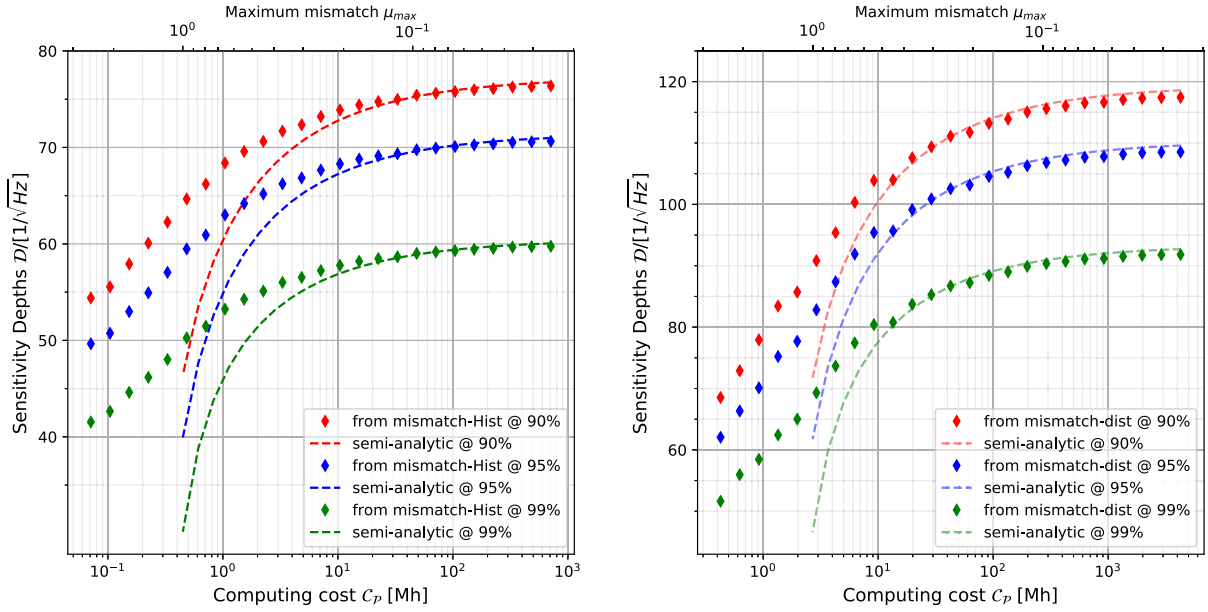


FIG. 8. Sensitivity depth $\mathcal{D}_{p_{\text{fa}}}^{p_{\text{det}}}$ as a function of (4D) computing cost $\mathcal{C}_{\mathcal{P}}$ for varying maximum mismatch μ_{max} at fixed segments ($N, \Delta T$), assuming Sco X-1 parameter space \mathcal{P}_2 of Table I. Sensitivity depth is estimated for a fixed (per-template) false alarm of $p_{\text{fa}} = 10^{-10}$ and different confidence levels of $p_{\text{det}} = 90\%$ (top), $p_{\text{det}} = 95\%$ (middle), and $p_{\text{det}} = 99\%$ (bottom). The segment setup is $N \times \Delta T = 180 \times 1$ d (left plot) and 120×3 d (right plot), corresponding to search setup-I and search setup-II, respectively (cf. Table II). The dashed lines correspond to the sensitivity estimate assuming a theoretical A_n^* lattice mismatch distribution, while the diamond markers correspond to using the measured BINARYWEAVE mismatch distributions. Computing cost is measured in million core hours (Mh).

Viterbi is also less as compared to the most sensitive search of BINARYWEAVE, provided the spin-wandering effect is not significantly large.

The sensitivity of the CrossCorr pipeline is expected to be comparable to BINARYWEAVE. The computing cost for *resampling* CrossCorr is also expected to be comparable to BINARYWEAVE. However, BINARYWEAVE can be adopted to utilize different grid spacing for coherent and semi-coherent template banks that can reduce the computing cost for a search. This extra amount of computing resource can be reutilized to further increase the sensitivity depth of BINARYWEAVE by either decreasing the mismatch or increasing the segment lengths.

One of the primary goals of developing BINARYWEAVE is to perform searches for Sco X-1 that can beat the torque-balance limit over as wide a frequency range as possible, and are able to take advantage of any large available computing budget. Still, at the current level of electromagnetic constraints on the Sco X-1 parameters, reaching the torque-balance limit over the full frequency range remains computationally prohibitive. Future improvements in these constraints will be immensely impactful to increase the chances of detecting a CW signal from Sco X-1 (or other LMXBs), as illustrated in Sec. V B.

ACKNOWLEDGMENTS

A. M. acknowledges Stuart Anderson, James Clark, Duncan Macleod, Dan Moraru, Keith Riles, Peter Shawhan, and several other members of the computing and software team of the LIGO Scientific Collaboration (LSC). A. M. is thankful to Heinz-Bernd Eggenstein for learning some of the advanced computational skills. A. M. also acknowledges computational assistance by Henning Fehrmann and Carsten Aulbert. A. M. is thankful to Grant David Meadors and several other past and present members of the continuous-waves working group of the LSC regarding general discussion on detectability of CW signals from Sco X-1. We thank Pep Covas and Paola Leaci for helpful feedback on the manuscript. This work has utilized the LDAS computing clusters at the LIGO Hanford Observatory (LHO) CalTech LIGO centre (CIT) and the ATLAS computing cluster at the MPI for Gravitational Physics Hannover. A. M. acknowledges support from the DST-SERB Start-up Research Grant No. SRG/2020/001290 for completion of this project. K. W. was supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav) through Project No. CE170100004.

- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [3] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
- [4] K. Riles, Recent searches for continuous gravitational waves, *Mod. Phys. Lett. A* **32**, 1730035–685 (2017).
- [5] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Gravitational waves from hot young rapidly rotating neutron stars, *Phys. Rev. D* **58**, 084020 (1998).
- [6] N. Andersson, A new class of unstable modes of rotating relativistic stars, *Astrophys. J.* **502**, 708 (1998).
- [7] B. Rajbhandari, B. J. Owen, S. Caride, and R. Inta, First searches for gravitational waves from r-modes of the Crab pulsar, *Phys. Rev. D* **104**, 122008 (2021).
- [8] R. V. Wagoner, Gravitational radiation from accreting neutron stars, *Astrophys. J.* **278**, 345 (1984).
- [9] L. Bildsten, Gravitational radiation and rotation of accreting neutron stars, *Astrophys. J. Lett.* **501**, L89 (1998).
- [10] G. Ushomirsky, C. Cutler, and L. Bildsten, Deformations of accreting neutron star crusts and gravitational wave emission, *Mon. Not. R. Astron. Soc.* **319**, 902 (2000).
- [11] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data, *Phys. Rev. D* **105**, 022002 (2022).
- [12] A. Melatos and D. J. B. Payne, Gravitational radiation from an accreting millisecond pulsar with a magnetically confined mountain, *Astrophys. J.* **623**, 1044 (2005).
- [13] N. K. Johnson-McDaniel and B. J. Owen, Maximum elastic deformations of relativistic stars, *Phys. Rev. D* **88**, 044004 (2013).
- [14] F. Gittins and N. Andersson, Modelling neutron star mountains in relativity, *Mon. Not. R. Astron. Soc.* **507**, 116 (2021).
- [15] D. Chakrabarty, The spin distribution of millisecond x-ray pulsars, *AIP Conf. Proc.* **1068**, 67 (2008).
- [16] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars, *Nature (London)* **424**, 42 (2003).
- [17] G. Hasinger and M. van der Klis, Two patterns of correlated x-ray timing and spectral behaviour in low-mass X-ray binaries, *Astron. Astrophys.* **225**, 79 (1989).
- [18] C. F. Bradshaw, E. B. Fomalont, and B. J. Geldzahler, High-resolution parallax measurements of Scorpius X-1, *Astrophys. J. Lett.* **512**, L121 (1999).
- [19] S. V. Dhurandhar and A. Vecchio, Searching for continuous gravitational wave sources in binary systems, *Phys. Rev. D* **63**, 122001 (2001).
- [20] C. Messenger, H. J. Bulten, S. G. Crowder, V. Dergachev, D. K. Galloway, E. Goetz, R. J. G. Jonker, P. D. Lasky, G. D. Meadors, A. Melatos, S. Premachandra, K. Riles, L. Sammut, E. H. Thrane, J. T. Whelan, and Y. Zhang, Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors, *Phys. Rev. D* **92**, 023006 (2015).
- [21] K. Riles, Searches for continuous-wave gravitational radiation, [arXiv:2206.06447](https://arxiv.org/abs/2206.06447).
- [22] C. Messenger and A. Vecchio, Searching for gravitational waves from low mass x-ray binaries, *Classical Quantum Gravity* **21**, S729 (2004).
- [23] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Directed search for gravitational waves from Scorpius X-1 with initial LIGO data, *Phys. Rev. D* **91**, 062008 (2015).
- [24] J. T. Whelan, S. Sundaresan, Y. Zhang, and P. Peiris, Model-based cross-correlation search for gravitational waves from Scorpius X-1, *Phys. Rev. D* **91**, 102005 (2015).
- [25] G. D. Meadors, E. Goetz, K. Riles, T. Creighton, and F. Robinet, Searches for continuous gravitational waves from Scorpius X-1 and XTE J1751-305 in LIGO's sixth science run, *Phys. Rev. D* **95**, 042005 (2017).
- [26] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, and R. X. Adhikari (LIGO Scientific Collaboration and Virgo Collaboration), Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model, *Phys. Rev. D* **95**, 122003 (2017).
- [27] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, D. Steeghs, and L. Wang (LIGO Scientific Collaboration and Virgo Collaboration), Upper limits on gravitational waves from Scorpius X-1 from a model-based cross-correlation search in Advanced LIGO data, *Astrophys. J.* **847**, 47 (2017).
- [28] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, and R. X. Adhikari (LIGO Scientific Collaboration and Virgo Collaboration), Search for gravitational waves from Scorpius X-1 in the second Advanced LIGO observing run with an improved hidden Markov model, *Phys. Rev. D* **100**, 122002 (2019).
- [29] Y. Zhang, M. A. Papa, B. Krishnan, and A. L. Watts, Search for continuous gravitational waves from Scorpius X-1 in LIGO O2 data, *Astrophys. J. Lett.* **906**, L14 (2021).
- [30] A. Mukherjee, C. Messenger, and K. Riles, Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches, *Phys. Rev. D* **97**, 043016 (2018).

- [31] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin, *Phys. Rev. D* **93**, 123009 (2016).
- [32] A. Melatos, P. Clearwater, S. Suvorova, L. Sun, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. III. Rotational phase tracking, *Phys. Rev. D* **104**, 042003 (2021).
- [33] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Cross-correlation search for periodic gravitational waves, *Phys. Rev. D* **77**, 082001 (2008).
- [34] G. D. Meadors, B. Krishnan, M. A. Papa, J. T. Whelan, and Y. Zhang, Resampling to accelerate cross-correlation searches for continuous gravitational waves from binary systems, *Phys. Rev. D* **97**, 044017 (2018).
- [35] K. J. Wagner, J. T. Whelan, J. K. Wofford, and K. Wette, Template lattices for a cross-correlation search for gravitational waves from Scorpius X-1, *Classical Quantum Gravity* **39**, 075013 (2022).
- [36] P. Leaci and R. Prix, Directed searches for continuous gravitational waves from binary systems: Parameter-space metrics and optimal Scorpius X-1 sensitivity, *Phys. Rev. D* **91**, 102003 (2015).
- [37] K. Wette, S. Walsh, R. Prix, and M. A. Papa, Implementing a semicoherent search for continuous gravitational waves using optimally constructed template banks, *Phys. Rev. D* **97**, 123016 (2018).
- [38] S. Walsh, K. Wette, M. Alessandra Papa, and R. Prix, Optimising the choice of analysis method for all-sky searches for continuous gravitational waves with Einstein@Home, *Phys. Rev. D* **99**, 082004 (2019).
- [39] Einstein@home project page, <https://einsteinathome.org/>.
- [40] R. Prix and M. Shaltev, Search for continuous gravitational waves: Optimal StackSlide method at fixed computing cost, *Phys. Rev. D* **85**, 084010 (2012).
- [41] A. E. Roy, *Orbital Motion*, 4th ed. (IOP Publishing, Bristol, 2005).
- [42] C. Messenger, Semicoherent search strategy for known continuous wave sources in binary systems, *Phys. Rev. D* **84**, 083003 (2011).
- [43] P. Jaranowski, A. Królak, and B. F. Schutz, Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection, *Phys. Rev. D* **58**, 063001 (1998).
- [44] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, *Classical Quantum Gravity* **26**, 204013 (2009).
- [45] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Searching for periodic sources with LIGO, *Phys. Rev. D* **57**, 2101 (1998).
- [46] G. Mendell and M. Landry, StackSlide and Hough search SNR and statistics, Technical Report No. T050003-x0, LIGO, 2005.
- [47] R. Prix, Search for continuous gravitational waves: Metric of the multidetector \mathcal{F} -statistic, *Phys. Rev. D* **75**, 023004 (2007).
- [48] R. Prix, Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces, *Classical Quantum Gravity* **24**, S481 (2007).
- [49] K. Wette, Lattice template placement for coherent all-sky searches for gravitational-wave pulsars, *Phys. Rev. D* **90**, 122010 (2014).
- [50] M. Shaltev and R. Prix, Fully coherent follow-up of continuous gravitational-wave candidates, *Phys. Rev. D* **87**, 084057 (2013).
- [51] P. R. Brady and T. Creighton, Searching for periodic sources with LIGO. II. Hierarchical searches, *Phys. Rev. D* **61**, 082001 (2000).
- [52] M. J. Coe, I. M. Monageng, J. A. Kennea, D. A. H. Buckley, P. A. Evans, A. Udalski, P. Groot, S. Bloemen, P. Vreeswijk, V. McBride, M. Klein-Wolt, P. Woudt, E. Körding, R. Le Poole, and D. Pieterse, SXP 15.6—an accreting pulsar close to spin equilibrium?, *Mon. Not. R. Astron. Soc.* **513**, 5567 (2022).
- [53] B. Haskell and A. Patruno, Spin equilibrium with or without gravitational wave emission: The case of XTE J1814-338 and SAX J1808.4-3658, *Astrophys. J. Lett.* **738**, L14 (2011).
- [54] I. Yi, J. C. Wheeler, and E. T. Vishniac, Torque reversal in accretion-powered x-ray pulsars, *Astrophys. J. Lett.* **481**, L51 (1997).
- [55] L. Bildsten, D. Chakrabarty, J. Chiu, M. H. Finger, D. T. Koh, R. W. Nelson, T. A. Prince, B. C. Rubin, D. M. Scott, M. Stollberg, B. A. Vaughan, C. A. Wilson, and R. B. Wilson, Observations of accreting pulsars, *Astrophys. J. Suppl. Ser.* **113**, 367 (1997).
- [56] K. Wette and R. Prix, Flat parameter-space metric for all-sky searches for gravitational-wave pulsars, *Phys. Rev. D* **88**, 123005 (2013).
- [57] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, Deep exploration for continuous gravitational waves at 171–172 Hz in LIGO second observing run data, *Phys. Rev. D* **103**, 083020 (2021).
- [58] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, A Series of Comprehensive Studies in Mathematics (Springer, New York, 1999).
- [59] B. Allen, Optimal template banks, *Phys. Rev. D* **104**, 042005 (2021).
- [60] B. Allen and A. A. Shoom, Template banks based on Z^n and A_n^* lattices, *Phys. Rev. D* **104**, 122007 (2021).
- [61] K. H. Prendergast and G. R. Burbidge, On the nature of some galactic x-ray sources, *Astrophys. J. Lett.* **151**, L83 (1968).
- [62] E. B. Fomalont, B. J. Geldzahler, and C. F. Bradshaw, Scorpius X-1: The evolution and nature of the twin compact radio lobes, *Astrophys. J.* **558**, 283 (2001).
- [63] L. Wang, D. Steeghs, D. K. Galloway, T. Marsh, and J. Casares, Precision ephemerides for gravitational-wave searches—III. Revised system parameters of Sco X-1, *Mon. Not. R. Astron. Soc.* **478**, 5174 (2018).
- [64] S. Galadage, K. Wette, D. K. Galloway, and C. Messenger, Deep searches for X-ray pulsations from Scorpius X-1 and Cygnus X-2 in support of continuous gravitational wave searches, *Mon. Not. R. Astron. Soc.* **509**, 1745 (2022).
- [65] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars, *Nature (London)* **424**, 42 (2003).

- [66] A. Singh, M. Alessandra Papa, H.-B. Eggenstein, S. Zhu, H. Pletsch, B. Allen, O. Bock, B. Maschenchalk, R. Prix, and X. Siemens, Results of an all-sky high-frequency Einstein@Home search for continuous gravitational waves in LIGO's fifth science run, *Phys. Rev. D* **94**, 064061 (2016).
- [67] LALSuite, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [68] R. Prix, The \mathcal{F} -statistic and its implementation in ComputeFstatistic_v2, Technical Report No. LIGO-T0900149-v6+, 2010.
- [69] B. Allen, Spherical ansatz for parameter-space metrics, *Phys. Rev. D* **100**, 124004 (2019).
- [70] R. Prix, Characterizing timing and memory-requirements of the \mathcal{F} -statistic implementations in LALSuite, Technical Report No. T1600531-v4+, 2017.
- [71] P.R. Williams and B.F. Schutz, An efficient matched filtering algorithm for the detection of continuous gravitational wave signals, *AIP Conf. Proc.* **523**, 473 (2000).
- [72] L. Dunn, P. Clearwater, A. Melatos, and K. Wette, Graphics processing unit implementation of the \mathcal{F} -statistic for continuous gravitational wave searches, *Classical Quantum Gravity* **39**, 045003 (2022).
- [73] C. Dreissigacker, R. Prix, and K. Wette, Fast and accurate sensitivity estimation for continuous-gravitational-wave searches, *Phys. Rev. D* **98**, 084058 (2018).
- [74] K. Wette, Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars, *Phys. Rev. D* **85**, 042003 (2012).
- [75] K. Wette, R. Prix, D. Keitel, M. Pitkin, C. Dreissigacker, J.T. Whelan, and P. Leaci, OCTAPPS: A library of Octave functions for continuous gravitational-wave data analysis, *J. Open Source Software* **3**, 707 (2018).
- [76] D. K. Galloway, S. Premachandra, D. Steeghs, T. Marsh, J. Casares, and R. Cornélisse, Precision ephemerides for gravitational-wave searches. I. Sco X-1, *Astrophys. J.* **781**, 14 (2014).