

Improved all-sky search method for continuous gravitational waves from unknown neutron stars in binary systems

P. B. Covas^{1,2} and R. Prix^{1,2}

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-30167 Hannover, Germany

²Leibniz Universität Hannover, D-30167 Hannover, Germany



(Received 5 August 2022; accepted 30 September 2022; published 20 October 2022)

Continuous gravitational waves from spinning deformed neutron stars have not been detected yet and are one of the most promising signals for future detection. All-sky searches for continuous gravitational waves from unknown neutron stars in binary systems are the most computationally challenging search type. Consequently, very few search algorithms and implementations exist for these sources, and only a handful of such searches have been performed so far. In this paper, we present a new all-sky binary search method, BinarySkyHough \mathcal{F} , which extends and improves upon the earlier BinarySkyHough method and which was the basis for a recent search [Covas *et al.*, *Astrophys. J. Lett.* **929**, L19 (2022)]. We compare the sensitivity and computational cost to the previous method, showing that it is both more sensitive and computationally efficient, which allows for broader and more sensitive searches.

DOI: [10.1103/PhysRevD.106.084035](https://doi.org/10.1103/PhysRevD.106.084035)

I. INTRODUCTION

Continuous gravitational waves (CWs) are long-lasting and nearly monochromatic gravitational waves, expected to be emitted by deformed spinning neutron stars (NSs) due to their time-varying quadrupole moment (e.g., [1]). Although many CW searches have been performed to date, using data from the LIGO (H1 and L1) and Virgo (V1) detectors, no detection has been achieved yet (see Ref. [2] for a recent review). The expected CW amplitudes are several orders of magnitude smaller than the compact-binary-coalescence signals currently being routinely detected. Therefore, the combined analysis of months to years worth of data is required to accumulate enough signal-to-noise ratio.

When searching for CWs from known pulsars, all the phase-evolution parameters are known from electromagnetic observations, which allows one to perform statistically optimal searches by coherent matched filtering with very little required computing power. All-sky CW searches for unknown neutron stars represent the opposite extreme, where no prior information about the signals is available, requiring an expensive explicit search over the phase-evolution parameters.

Furthermore, because the required parameter-space resolution increases rapidly with longer coherent integration time, the resulting computing cost explodes and makes it impossible to analyze longer stretches of data by coherent matched filtering. This computing-cost problem is pushed to the extreme when searching for unknown neutron stars in binary systems, as now we also need to search over the unknown binary orbital parameters [3]. Therefore, all-sky searches for unknown neutron stars in binary systems are the most computationally challenging type of searches.

The primary strategy followed by computationally limited CW searches is to break up the data into shorter segments that can be coherently analyzed individually and then combine these coherent results across segments. These are the so-called *semicoherent* methods (see Ref. [4] for a recent review of search methods). The resulting coarser parameter-space resolution entails a reduced computational cost, which allows for analyzing larger datasets and thereby regaining sensitivity. As a result, semicoherent methods are typically more sensitive than fully coherent matched filtering at a fixed computational budget [5,6].

The most commonly used coherent detection statistics are the \mathcal{F} -statistic and the Fourier power (for sufficiently short segments, where a simple sinusoid can approximate the signal). The \mathcal{F} -statistic is obtained by analytically maximizing the likelihood ratio over the four unknown amplitude parameters of a CW signal [7,8]. Although this statistic was initially thought to be optimal, its implicit amplitude priors have been shown to be unphysical [9]. Using better physical priors results in a more sensitive Bayes factor, albeit (currently) at increased computational

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). 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.

cost, which is why this is not yet a viable alternative to the \mathcal{F} -statistic for wide parameter-space searches. However, for short segments compared to a day, a new detection statistic has recently been found [10] that is more sensitive than \mathcal{F} at no extra computing cost, termed the *dominant-response* statistic \mathcal{F}_{AB} .

Using Fourier power over short segments directly as a coherent detection statistic is computationally cheaper, given that no phase demodulation or other additional calculations are needed. However, one limitation of this statistic is the constrained maximum coherent length (about $T_{\text{seg}} \lesssim 30$ min at around 500 Hz), resulting from the approximation of the signal as a simple sinusoid. Typically, this is expressed as the criterion that the signal power remains in a single frequency bin (of size $1/T_{\text{seg}}$). Furthermore, while demodulated statistics (such as \mathcal{F} and \mathcal{F}_{AB}) can naturally combine data from several detectors coherently [8,10], this is not straightforward to achieve for short Fourier transforms [11,12] and is not commonly used in the first stage of hierarchical searches. Therefore, constructing semicoherent statistics on demodulated coherent statistics is generally more sensitive and flexible than using Fourier power.

Only two previous all-sky binary pipelines have been used in searches before BinarySkyHou \mathcal{F} [13], namely, TwoSpect [14] and BinarySkyHough [15]. TwoSpect was the first pipeline for all-sky CW searches of unknown NSs in binary systems [16]. BinarySkyHough is an extension of SkyHough [17] (an all-sky pipeline for isolated systems) to searches for NSs in binary systems, which yields higher sensitivity compared to TwoSpect thanks to its more sensitive detection statistic and usage of GPU parallelization. Two recent all-sky binary searches deployed BinarySkyHough on data from Advanced LIGO's O2 and O3 observing runs [18,19], although over a reduced parameter space in frequency and binary parameters compared to the TwoSpect search.

BinarySkyHough uses short-segment Fourier power as its coherent detection statistic, limiting its attainable sensitivity (as discussed above). Here we present BinarySkyHou \mathcal{F} , an extension of BinarySkyHough, which features several improvements compared to the previous pipeline:

- (i) use of demodulated coherent statistics (such as \mathcal{F} - or \mathcal{F}_{AB} -statistic) instead of short-segment Fourier power;
- (ii) directly summing coherent detection statistics (the typical StackSlide approach [5,20]) instead of (thresholded) 1s and 0s as in the classical Hough algorithm [17];
- (iii) various code-implementation improvements (such as GPU coalesced memory access) and optimizations, increasing computational efficiency.

As will be shown in this paper, the new search pipeline is both more sensitive and more computationally efficient than BinarySkyHough; i.e., for the same coherent segment length

T_{seg} and mismatch distribution, it achieves higher sensitivity at lower computational cost.

A key ingredient of the new pipeline is the use of (low-order) Taylor-expanded phase parameters to describe the binary motion over the (short) coherent segments instead of the physical binary orbital parameters. These *Taylor coordinates* allow for a substantial dimensional reduction and solve the problem of covering the highly degenerate per-segment coherent parameter space with an efficient template bank. However, this approach limits the sensitivity to signals from binary systems with orbital periods substantially longer than the segment length T_{seg} .

The development of more sensitive all-sky binary search methods is of utmost importance, since more than half of all known millisecond pulsars are part of a binary system [21,22]. Furthermore, accretion from a companion gives a plausible mechanism to generate an asymmetry or excite an r-mode with a detectable amplitude in the current generation of gravitational-wave detectors, as recently discussed in [13].

This paper is organized as follows: In Sec. II, we introduce the approximate signal model used to compute the \mathcal{F} -statistic; in Sec. III, we present the new BinarySkyHou \mathcal{F} pipeline and we compare it to its predecessor; in Sec. IV, we show sensitivity comparisons of different detection statistics; in Sec. V, we summarize the main results of this paper and lay out some ideas for future work.

II. SIGNAL PHASE MODEL

A. Physical phase model

Assuming a slowly varying NS spin frequency, the phase of a CW signal in the source frame can be expressed in terms of the Taylor expansion around a reference time τ_{ref} , namely,

$$\phi(\tau) = \phi_0 + 2\pi \sum_{k=0}^s \frac{f_k}{(k+1)!} (\tau - \tau_{\text{ref}})^{k+1}, \quad (1)$$

where τ denotes the time in the source frame and s is the order of spin-down parameters f_k needed to accurately describe the intrinsic frequency evolution. The evolution of frequency $f(\tau)$ and higher-order spin-downs $f^{(k)'}(\tau)$ is given by

$$f(\tau) = \frac{1}{2\pi} \frac{d\phi(\tau)}{d\tau} \quad \text{and} \quad f^{(k)'}(\tau) = \frac{d^k f(\tau)}{d\tau^k}, \quad (2)$$

and the frequency and spin-down *parameters* f_k in the phase model of Eq. (1) are defined at the reference time τ_{ref} , i.e.,

$$f_k \equiv f^{(k)'}(\tau = \tau_{\text{ref}}). \quad (3)$$

In order to obtain the phase of the signal in the frame of a detector, we need to transform it from the source frame by

taking into account the movement of the NS and the movement of the detector with respect to the solar system barycenter (SSB). We absorb the unknown relative distance of the source with respect to the SSB into the reference time τ_{ref} , and here we neglect relativistic effects such as Shapiro and Einstein delays¹ and the transverse proper motion of the source. We can obtain the timing relation in two steps, first linking the wave-front-emission time τ in the source frame to its arrival time t_{SSB} in the SSB frame, namely,

$$\tau(t_{\text{SSB}}) = t_{\text{SSB}} - R(\tau), \quad (4)$$

where $R(\tau)$ is the radial distance (in light-travel time) of the source to the binary barycenter (BB) [3], with $R > 0$ when the source is further away from us than the BB. In the second step, we can relate the SSB time to the arrival time t at detector X by the Rømer-delay expression:

$$t_{\text{SSB}}^X(t) = t + \vec{r}^X(t) \cdot \hat{n}, \quad (5)$$

where $\vec{r}^X(t)$ is the position vector (in light-travel time) of detector X with respect to the SSB and \hat{n} is the sky-position unit vector pointing from the SSB to the BB. The radial distance R of the source to the BB can be expressed [3] as

$$R(\tau) = a_p [\sin \omega (\cos E - e) + \cos \omega \sin E \sqrt{1 - e^2}], \quad (6)$$

in terms of the eccentric anomaly E , given by Kepler's equation

$$E = \Omega(\tau - t_p) + e \sin E, \quad (7)$$

where a_p is the projected semimajor axis of the orbital ellipse (in light-travel time), $\Omega = 2\pi/P_{\text{orb}}$ is the (average) orbital angular velocity (corresponding to the period P_{orb}), e is the orbital eccentricity, t_p is the time of periape passage, and ω is the (angular) argument of periape.

For small-eccentricity orbits, this can be approximated by the (linear in e) ELL1 model [3,23,24], namely,

$$R(\tau) = a_p \left[\sin \Psi + \frac{\kappa}{2} \sin 2\Psi - \frac{\eta}{2} \cos 2\Psi \right] + \mathcal{O}(e^2) \quad (8)$$

(dropping a constant term $-3a_p\eta/2$) with the Laplace-Lagrange parameters defined as $\kappa \equiv e \cos \omega$ and $\eta \equiv e \sin \omega$ and the orbital phase

$$\Psi(\tau) = \Omega(\tau - t_{\text{asc}}), \quad (9)$$

using the time of ascending node t_{asc} instead of the time of periape passage t_p , which (to lowest order in e) are related

¹The numerically implemented phase model includes these effects for the solar system but not for the binary system.

by $t_{\text{asc}} = t_p - \omega/\Omega$. Expressions for R up to any order in e are found in Appendix C in Ref. [25].

B. Short-segment SSB Taylor coordinates $\{u_k\}$

Given all-sky binary CW searches need to cover a huge signal parameter space with finite computing resources, the longest coherent segment lengths T_{seg} that can be used are typically very short (i.e., much shorter than a day). If we further assume the orbital periods to be much longer than the short segments, i.e., $T_{\text{seg}} \ll P_{\text{orb}}$, then in this short-segment regime we can resort to Taylor expanding the phase (in the SSB) around each segment midtime t_m (translated to the SSB), similar to what was done in Refs. [3,24], namely,

$$\phi(t_{\text{SSB}}) = \phi_0 + 2\pi \sum_{k=1}^{k_{\text{max}}} \frac{u_k}{k!} (t_{\text{SSB}} - t_m)^k, \quad (10)$$

which defines the (SSB) *Taylor coordinates* $\{u_k\}_{k=1}^{k_{\text{max}}}$ as

$$u_k \equiv \left. \frac{1}{2\pi} \frac{d^k \phi}{dt_{\text{SSB}}^k} \right|_{t_m}. \quad (11)$$

Note that, for segments short compared to a day, one could also define Taylor coordinates in the detector frame instead of in the SSB, but this would result in detector-dependent coordinates that are not suitable for our present search method. The resulting expressions are given in the Appendix for reference.

Inserting the physical phase model of Eq. (1) in the form $\phi(t_{\text{SSB}}) = \phi(\tau(t_{\text{SSB}}))$, we obtain the phase derivatives²

$$\begin{aligned} u_1 &= [f(\tau)\dot{\tau}]|_{t_m}, \\ u_2 &= [f'(\tau)\dot{\tau}^2 + f(\tau)\ddot{\tau}]|_{t_m}, \\ u_3 &= [f''(\tau)\dot{\tau}^3 + 3f'(\tau)\ddot{\tau}\dot{\tau} + f(\tau)\dddot{\tau}]|_{t_m}, \\ &\vdots \end{aligned} \quad (12)$$

where the derivatives $\tau^{(k)} \equiv d^k \tau / dt_{\text{SSB}}^k$ of the source-to-SSB timing relation $\tau(t_{\text{SSB}})$ can be further expanded using Eq. (4), involving derivatives $R^{(k)} \equiv d^k R(\tau) / dt_{\text{SSB}}^k$ of the binary radial distance $R(\tau)$ of Eq. (6), which can be expanded in the same form as

²The general form of these successive chain- and product-rule expansions is governed by the so-called Faà di Bruno formula [26].

$$\begin{aligned}
\dot{R} &= R'\dot{\tau}, & \Psi_m &\equiv \Omega(t_m - t_{\text{asc}}), \\
\ddot{R} &= R''\dot{\tau}^2 + R'\ddot{\tau}, & & (19) \\
\dddot{R} &= R'''\dot{\tau}^3 + 3R''\ddot{\tau}\dot{\tau} + R'\dddot{\tau}, \\
&\vdots & &
\end{aligned}
\tag{13}$$

in terms of the source-frame time derivatives $R^{(k)'} \equiv d^k R(\tau)/d\tau^k$. This analysis is complicated by the fact that the binary radial distance $R(\tau)$ of Eq. (6) is a function of source-frame (emission) time τ , not the SSB (arrival) time t_{SSB} of a wave front. In Refs. [3,24], this difficulty could be neglected for the purposes of computing the parameter-space metric, where a slow-orbit approximation, i.e., $R(\tau) \approx R(t_{\text{SSB}})$, is sufficient. However, in the present application we want to preserve higher accuracy for the purpose of using these coordinates for coherent matched filtering.

Substituting into the timing derivatives of Eq. (4), we can now obtain the expressions

$$\begin{aligned}
\dot{\tau} &= [1 + R']^{-1}, \\
\ddot{\tau} &= [1 + R']^{-1}(-R''\dot{\tau}^2), \\
\dddot{\tau} &= [1 + R']^{-1}(-R'''\dot{\tau}^3 - 3R''\ddot{\tau}\dot{\tau}), \\
&\vdots
\end{aligned}
\tag{14}$$

which are explicit because of the iterative backward dependency of the $\tau^{(k)}$ on only lower-order derivatives, i.e., $\tau^{(k)} = \tau^{(k)}(\dot{\tau}, \ddot{\tau}, \dots, \tau^{(k-1)})$.

From these expressions we can obtain the explicit Taylor coordinates u_k via Eq. (12) as

$$u_1 = \frac{f_m}{1 + R'_m}, \tag{15}$$

$$u_2 = -\frac{f_m R''_m}{(1 + R'_m)^3} + \frac{f'_m}{(1 + R'_m)^2}, \tag{16}$$

where for the present application (such as Ref. [13]) at most first- or second-order u_k will be needed in practice, as discussed in Sec. III C 1. Here, we defined

$$\begin{aligned}
f_m &\equiv f(t_m) = f_0 + f_1(t_m - \tau_{\text{ref}}) + \dots, \\
f'_m &\equiv f'(t_m) = f'_1 + f'_2(t_m - \tau_{\text{ref}}) + \dots,
\end{aligned}
\tag{17}$$

and

$$\begin{aligned}
R'_m &= a_p \Omega [\cos \Psi_m + \kappa \cos 2\Psi_m + \eta \sin 2\Psi_m], \\
R''_m &= -a_p \Omega^2 [\sin \Psi_m + 2\kappa \sin 2\Psi_m - 2\eta \cos 2\Psi_m],
\end{aligned}
\tag{18}$$

where $R_m^{(k)'} \equiv R^{(k)'}(t_m)$ and

assuming the small-eccentricity approximation of Eq. (8) and (here) neglecting the NS-BB time delay of Eq. (4) as a higher-order correction, i.e., $\tau(t_m) \approx t_m$.

These u_1 and u_2 coordinates have units of Hz and Hz², respectively, and depend on the physical parameters $\{\{f_k\}, a_p, \Omega, t_{\text{asc}}, e, \omega\}$. Using these coordinates assumes that we have performed the standard SSB demodulation of the signal for any given sky position \hat{n} , which is typically expressed in terms of the right ascension α and declination δ in equatorial coordinates.

The resulting (constant) parameter-space metric for the Taylor coordinates $\{u_k\}$ [valid for any signal phase of the form Eq. (10)] is found in Eq. (57) in Ref. [3].

III. BINARYSKYHOUGH

In this section, we present a summary of the new BinarySkyHough pipeline and its main advantages over the previous BinarySkyHough.

A. Summary of the previous and new pipeline

The predecessor SkyHough and BinarySkyHough algorithms are described in more detail in Refs. [15,17]; here, we provide only a short overview summary. Both of these analyze the frequency-time matrix of short-Fourier-transform (SFT) power, by searching for “tracks” (corresponding to different source parameters) that are above the statistical expectation for noise.

SkyHough is limited to searches for signals from isolated systems, while BinarySkyHough is an extension designed for all-sky searches for unknown neutron stars in binary systems. Both are extremely fast model-based pipelines due to the highly efficient algorithms used to analyze the sky maps and their effective use of look-up tables (see Refs. [15,17] for details). Furthermore, BinarySkyHough leverages the computational advantages provided by GPUs by parallelizing the most expensive steps in the algorithm and, thus, further massively reducing the run-time of a search.

A BinarySkyHough search is divided in two consecutive stages, using different detection statistics. In the first stage, a “Hough” weighted sum of 1s and 0s (depending on whether the SFT power crossed a threshold or not) is calculated, and all of the templates are sorted by the resulting significance [a normalized Hough number count with normal distribution; see Eq. (25) in Ref. [15]]. The frequency-time pattern used for the tracks in the first stage is an approximation to the exact one, due to the usage of look-up tables (explained in Sec. IV B in Ref. [15]). In the second stage, the *refinement stage*, a fraction of the best-ranked templates is analyzed again, this time using a StackSlide weighted sum of SFT power, which has a higher sensitivity than summing weighted 1s and 0s

(e.g., see Ref. [20]), and using a more accurate expression for the frequency-time pattern.

The central new feature of the BinarySkyHough pipeline is to use a *demodulated* coherent detection statistic for the segments,³ such as the \mathcal{F} - [7,8] or \mathcal{F}_{AB} -statistic [10], instead of the number count or SFT power, but otherwise still benefit from the highly efficient GPU-based BinarySkyHough-type algorithm to combine the coherent results to a semicoherent statistic. Three main benefits arise from using a demodulated coherent statistic:

- (1) Demodulation removes one important constraint, as the signal is no longer approximated as a pure sinusoid. This allows the algorithm to turn increases in computing power into better sensitivity (shown in Sec. IV C).
- (2) The per-detector data are combined coherently, which reduces the number of coherent segments needed to combine in the semicoherent stage, improving sensitivity (shown in Sec. IV A) and reducing computational cost (shown in Sec. III D).
- (3) The parameter-space resolution and resulting mismatch can be controlled as a free parameter (rather than the fixed $1/T_{\text{seg}}$ Fourier resolution of SFTs), which will be discussed in Sec. III C 1.

When combining coherent results to calculate a semicoherent detection statistic, it has been shown that applying per-segment weights can improve the sensitivity [28]. While there is currently no analytic argument for using a weighted sum of \mathcal{F} - or \mathcal{F}_{AB} -statistics,⁴ empirically we find (shown in Sec. IV A) that using weights also improves the sensitivity for these detection statistics. The weight w_ℓ at segment ℓ is given by

$$w_\ell = K \frac{A_\ell + B_\ell}{S_{n;\ell}}, \quad \text{with} \quad (20)$$

$$A_\ell = \sum_{\alpha=1}^{N_{\text{SFTs}}} \frac{S_{n;\ell}}{S_{n;\alpha}} a_\alpha^2, \quad B_\ell = \sum_{\alpha=1}^{N_{\text{SFTs}}} \frac{S_{n;\ell}}{S_{n;\alpha}} b_\alpha^2, \quad (21)$$

where K is a normalization factor defined such that $\sum_{\ell=1}^{N_{\text{seg}}} w_\ell = N_{\text{seg}}$, $S_{n;\alpha}$ is the noise power spectral density of SFT α , $S_{n;\ell}$ is the noise power spectral density of segment ℓ (defined as the harmonic mean over the $\{S_{n;\alpha}\}$), and a_α and b_α are the antenna patterns of a detector (evaluated at the midtime of every SFT α) given by

³The \mathcal{F} -statistic has been used before in combination with the Hough algorithm, in Ref. [27], an all-sky search for isolated systems using day-long coherent segments, where the (single-stage) pipeline summed weighted 1s and 0s computed from thresholded \mathcal{F} -statistic values.

⁴This is left for future work but intuitively can be understood from the per-segment change in signal power, which can vary by around one order of magnitude between segments for such short coherent times.

Eqs. (12) and (13) in Ref. [7], where the sum goes over all the SFTs in segment ℓ . When the segment just has one SFT, we recover the weights given by Eq. (22) in Ref. [15]. For the dominant-response statistic, we use the following weights:

$$w_\ell = K \begin{cases} \frac{A_\ell + \frac{C_\ell^2}{A_\ell}}{S_{n;\ell}} & \text{if } A_\ell \geq B_\ell, \\ \frac{B_\ell + \frac{C_\ell^2}{B_\ell}}{S_{n;\ell}} & \text{otherwise,} \end{cases} \quad (22)$$

where $C_\ell = \sum_{\alpha=1}^{N_{\text{SFTs}}} \frac{S_{n;\ell}}{S_{n;\alpha}} a_\alpha b_\alpha$.

As discussed in Sec. II B, for computing-cost reasons, the coherent segments for all-sky binary searches need to be very short, which allows us to use a small number (currently one or two) of Taylor coordinates u_k to represent the spin-down and binary orbital motion in the coherent segments. Using physical parameters, we would need to build a (at least) six-dimensional parameter space grid⁵ over $\{\alpha, \delta, f_0, a_p, \Omega, t_{\text{asc}}\}$, while using the Taylor coordinates we effectively need to use only three (or four) parameter-space dimensions for the short segments currently considered, namely, $\{\alpha, \delta, u_1, u_2\}$. This reduces complexity (the parameter-space metric in physical parameters would be hugely degenerate for short segments; cf. [3,24]) and lowers the resulting computational cost.⁶

The u_k template bank is constructed as a hypercubic lattice in coordinate space. The code processes the sky in patches defined by an isotropic grid with cells of fixed solid angle, using partial Hough map derivatives [15] to process the semicoherent sky mapping. Coherent per-segment statistics are computed only for the center of each sky patch using the corresponding antenna pattern modulations and weights.

B. Semicoherent interpolation

In the previous section, we obtained the Taylor coordinates u_k , which together with the sky position coordinates will be used in the coherent stage in order to calculate the \mathcal{F} -statistic values over coordinates $\{\alpha, \delta, \{u_k\}\}$. In the semicoherent stage, on the other hand, we are using physical coordinates to combine the per-segment statistics, namely, $\{\alpha, \delta, \{f_k\}, a_p, \Omega, t_{\text{asc}}, e, \omega\}$. For every semicoherent template, we therefore need to calculate the appropriate *mapping* to the corresponding Taylor coordinates $\{u_k\}$ and coherent sky position.

⁵In comparison to BinarySkyHough, the new code is able to also search over spin-downs and eccentricity in the semicoherent stage.

⁶Using a Taylor phase approximation to lower the computational cost of an all-sky binary search has been first explored by the *Polynomial* pipeline [29], which did not use physical parameters for the semicoherent summation, however, resulting in reduced sensitivity.

In addition to using different signal parameters, the semicoherent template grids also generally need to be finer than the coherent per-segment grids, which results in the need to *interpolate* the coherent results when combining them semicoherently (typically using nearest-neighbor interpolation). This is a generic feature of the semicoherent approach (cf. [6,30]) and in SkyHough-derived pipelines takes the form of the so-called *master equation* [15,17], linking sky offsets to resulting effective frequency shifts of the signal.

The SkyHough-type sky interpolation works by breaking the sky into several sky patches, as mentioned above, where the center of each patch is used as the coherent sky template for every semicoherent sky template in the same sky patch. The resulting offset in sky position between the semicoherent and coherent template results in compensating offsets in the $\{u_k\}$ coordinates, generalizing the Hough master equation.

A simple way to derive the shift in u_k coordinates due to an offset $\delta\hat{n}$ in sky position is to use the full detector-frame Taylor coordinates u_k^X for each detector X , given in the Appendix in Eq. (A5). Using this, we can express the induced shifts δu_k^X as

$$\begin{aligned}\delta u_1^X &= u_1 \vec{v}_m^X \cdot \delta\hat{n}, \\ \delta u_2^X &\approx u_1 \vec{a}_m^X \cdot \delta\hat{n} + 2\vec{v}_m^X \cdot \delta\hat{n}(-u_1 R_m'' + f_m'),\end{aligned}\quad (23)$$

in terms of detector velocity \vec{v}_m^X and acceleration \vec{a}_m^X at the segment midtime t_m translated to the SSB.⁷ To remove the detector dependency, we average over detectors, which will be a good approximation for δu_1 (when using a grid with an $\sim 1/T_c$ resolution), given the detector velocity is dominated by the (detector-independent) orbital motion of Earth. On the other hand, the detector acceleration \vec{a}_m^X in δu_2^X is dominated by Earth's rotation, so averaging over detectors might be a less reliable approximation but should still work well as long as the detectors are not too far separated, such as for LIGO H1 and L1. Therefore, we arrive at the following generalized master equations:

$$u_{11} \approx u_1(1 + \vec{v}_m \cdot \delta\hat{n}),\quad (24)$$

$$u_{12} \approx u_2 + u_1 \vec{a}_m \cdot \delta\hat{n} + 2\vec{v}_m \cdot \delta\hat{n}(-u_1 R_m'' + f_m'),\quad (25)$$

with detector-averaged velocity $N_{\text{det}}\vec{v} \equiv \sum_X \vec{v}^X$ and acceleration $N_{\text{det}}\vec{a} = \sum_X \vec{a}^X$. Equation (24) agrees with the previous Hough-on- \mathcal{F} -statistic master equation found in Refs. [17,27] (with implicit detector averaging).

The u_{11} master equation is illustrated in Fig. 1, where the u_1 values with the highest signal power are plotted as a

⁷Strictly speaking, here it should be the segment midtime in the detector frame, but the maximal shift of ~ 500 s can be neglected for detector velocity and acceleration.




FIG. 1. Comparison between the u_{11} coordinate per segment given by Eq. (24) (green line) and the u_1 values per segment that maximize the signal power (orange points). The sky position is offset by 0.3 rad in both α and δ from the signal. This example assumes a single detector (H1) and one year of data, with segments of length $T_{\text{seg}} = 900$ s, and a constant-frequency signal of $f_0 = 100$ Hz without binary modulation.

function of segment midtime t_m for an offset $\delta\hat{n} = \hat{n} - \hat{n}_d$ between the signal sky position \hat{n} and the coherent demodulation point \hat{n}_d . In addition, we plot the predicted track of shifted u_{11} given by Eq. (24). These u_{11} values closely follow the path that minimizes the mismatch. The mismatch in Fig. 1 between the path followed by u_{11} and the maximum path is around 0.1%, whereas the mismatch between the nonshifted u_1 and the maximum path would be around 80%.

The parameter-space bounds for each of the Taylor coordinates are found using Eqs. (24) and (25), by calculating the maximum possible values over the given physical parameter space.

For computational efficiency reasons (namely, the lookup table approach used here; see Ref. [17]), the first-stage ‘‘Hough’’ semicoherent summation actually uses the following approximate expressions instead of Eqs. (24) and (25), namely,

$$u_{1H1} = f_m(1 - R_m') + f_H \vec{v}_m \cdot \delta\hat{n},\quad (26)$$

$$u_{1H2} = u_2,\quad (27)$$

with the ‘‘middle’’ frequency f_H given by

$$f_H \equiv f_{0H} + f_{1H}(t_{\text{mid}} - \tau_{\text{ref}}) + \dots,\quad (28)$$

where the f_{kH} denote the midpoints in frequency and spin-down ranges currently being searched over and t_{mid} is the midtime of the full dataset.

C. Mismatch

In this subsection, we describe and characterize different sources of mismatch for the BinarySkyHou \mathcal{F} pipeline. The mismatch is defined as the relative loss of signal power, namely,

$$\mu = 1 - \frac{\rho_r^2}{\rho^2}, \quad (29)$$

where ρ^2 is the full signal power (given by Eq. (20) in Ref. [10]) and ρ_r^2 is the signal power recovered by the search.

The total mismatch of the BinarySkyHou \mathcal{F} pipeline has several contributions, which can be separated in coherent and semicoherent mismatches. The main contributions to the coherent mismatch are offsets between signal and the closest template in the coherent template grid and the usage of the (truncated) Taylor coordinates u_k , while the semicoherent mismatch is produced by signal-template offsets in the semicoherent grid and approximations in the interpolation mapping (discussed in the previous section).

1. Mismatch due to Taylor-coordinate truncation

The usage of a limited number of Taylor coordinates u_k incurs an intrinsic mismatch due to the corresponding approximation of the signal waveform. In practice, we currently envisage using only u_1 or at most up to the order of u_2 , which turns out to be sufficient for currently considered practical all-sky binary searches (similar to the recent search [13] using only u_1) due to computational constraints. Therefore, we quantify the mismatch and corresponding constraints on the maximum coherent segment length T_{seg} due to truncation to the order of u_1 or u_2 .

The mismatch μ_{u_k} due to omission of the order of u_k (and higher) can be estimated as $\mu_{u_k} \sim g_{kk}u_k^2$ using the metric g_{kl} in u_k space, which is given in Eqs. (56) and (57) in Ref. [3]. Using this, we can express the mismatch due to truncation of $u_{k \geq 2}$ or $u_{k \geq 3}$ as

$$\begin{aligned} \mu_{u_2} &\approx g_{22}u_2^2 = \frac{\pi^2 T_{\text{seg}}^4}{180} u_2^2, \\ \mu_{u_3} &\approx g_{33}u_3^2 = \frac{\pi^2 T_{\text{seg}}^6}{4032} u_3^2. \end{aligned} \quad (30)$$

Using a time average $\langle \cdot \rangle$ over segments together with Eq. (16) for u_2 we obtain $\langle u_2^2 \rangle \approx \frac{1}{2} f_0^2 a_p^2 \Omega^4 + f_1^2$ (neglecting smaller corrections), and for u_3 we can use Eq. (58) in Ref. [3] as an estimate, which yields $\langle u_3^2 \rangle \approx \frac{1}{2} f_0^2 a_p^2 \Omega^6$, and so we obtain the (segment-averaged) mismatch estimates as

$$\langle \mu_{u_2} \rangle \approx \frac{\pi^2 T_{\text{seg}}^4}{360} (f_0^2 a_p^2 \Omega^4 + 2f_1^2), \quad (31)$$

$$\langle \mu_{u_3} \rangle \approx \frac{\pi^2 T_{\text{seg}}^6}{8064} f_0^2 a_p^2 \Omega^6, \quad (32)$$

which illustrate the fact that the segments must be short compared to the orbital period, i.e., $T_{\text{seg}} \Omega \ll 1$, in order for the Taylor coordinates u_k to be a good approximation, as discussed in Ref. [3].

We can rearrange these equations to obtain a constraint on the maximum coherent time T_{seg} allowed for a given acceptable mismatch $\langle \mu_u \rangle$ from Taylor truncation; namely, when using only u_1 , we find

$$T_{\text{seg}, u_1} \leq \left(\frac{360 \langle \mu_u \rangle}{\pi^2 (f_0^2 a_p^2 \Omega^4 + 2f_1^2)} \right)^{1/4}, \quad (33)$$

and similarly for truncation after u_2 we obtain the constraint

$$T_{\text{seg}, u_2} \leq \left(\frac{8064 \langle \mu_u \rangle}{\pi^2 f_0^2 a_p^2 \Omega^6} \right)^{1/6}. \quad (34)$$

These expressions for the maximum coherent time are illustrated in Fig. 2 as a function of frequency for different choices of binary orbital parameters. It can be seen that when u_2 is also used the maximum coherent time increases by a certain factor.

Figure 3 shows a test of the mismatch predicted by Eq. (31), by generating 1000 different signals with a constant frequency of 300 Hz and random orbital parameters log-uniformly distributed, with $P_{\text{orb}} \in [0.1, 1]$ day and $a_p \in [0.01, 1]$ ls. For each signal we measure the perfectly matched signal power when using physical coordinates and compare it to the signal power obtained with Taylor coordinates up to the order of u_1 . The corresponding measured mismatch is then compared to the model prediction of Eq. (31). The figure shows that these equations correctly predict the measured mismatch, in the range where the metric approximation is expected to be accurate (i.e., below mismatches of ~ 0.3).

2. Total mismatch

In the previous subsection, we discussed the mismatch contribution due to using a limited set of Taylor coordinates. Additionally, there will be template-bank mismatches incurred from the coherent and semicoherent template grids. If we count the Taylor-truncation mismatch μ_u as part of the coherent mismatch μ_c , then the total average (over the template bank) mismatch $\langle \mu \rangle$ will be given approximately [6] by the sum of the mean coherent $\langle \mu_c \rangle$ and mean semicoherent mismatch $\langle \mu_s \rangle$, namely,

$$\langle \mu \rangle \simeq \langle \mu_c \rangle + \langle \mu_s \rangle. \quad (35)$$

From this expression, it can be seen that if the mean coherent mismatch is reduced while the semicoherent mismatch is equal, the total mismatch would decrease.




FIG. 2. The left plot shows the maximum coherent time T_{seg} that can be used in a given region of the binary parameter space with a certain mismatch due to the Taylor truncation (assuming zero spin-down $f_1 = 0$). The left plot shows the results for $\langle \mu_{u_i} \rangle = 0.1$ (solid lines) and $\langle \mu_{u_i} \rangle = 0.4$ (dashed lines) in Eq. (33), using only the u_1 Taylor coordinate. The right plot shows the same when using both u_1 and u_2 Taylor coordinates.




FIG. 3. Measured mismatch $\langle \mu_{u_2} \rangle$ (blue points) from Taylor truncation (averaged over segments) against the predicted mismatch (solid line) given by Eq. (31), using 1000 injections and assuming one year of data from the H1 detector with coherent segments of $T_{\text{seg}} = 900$ s. The injections have a constant frequency of 300 Hz and random orbital parameters log-uniformly distributed, with $P_{\text{orb}} \in [0.1, 1]$ day and $a_p \in [0.01, 1]$ ls.

This is shown in Fig. 4, where the total measured mismatch can be seen for two different cases, which have different u_1 coherent grids but the same semicoherent grid (the coherent sky position is the same for both cases, and it is shifted from the signal value). A decrease in the total mismatch can be seen for the case where the coherent u_1 grid is finer, as predicted by the previous equation. This represents an improvement over the previous pipeline, where the coherent frequency grid was fixed to be equal to the Fourier transform spacing.




FIG. 4. Mismatch histograms for 1000 injections with a constant frequency of 100 Hz and random orbital parameters log-uniformly distributed, with $P_{\text{orb}} \in [15, 45]$ day and $a_p \in [10, 40]$ ls, assuming one year of data from both H1 and L1 detectors and coherent segments of $T_{\text{seg}} = 900$ s. The olive (right) histogram corresponds to a grid resolution of $\delta u_1 = 1/T_{\text{seg}}$, while the red (left) histogram uses a finer resolution of $\delta u_1 = 0.25/T_{\text{seg}}$, with the u_{c_1} factor in the legend indicating this. Only u_1 and a single sky position are searched in the coherent stage. The semicoherent grids are equal for both cases.

3. Maximum spin-down and eccentricity

Although BinarySkyHou \mathcal{F} is able to search the f_1 spin-down and eccentricity e parameters in the semicoherent summation, all-sky searches for neutron stars in binary systems are so computationally expensive that one would usually not explicitly search over these parameters at first.

For this reason, we want to estimate the value ranges in these omitted parameters to which the pipeline is still sensitive, which will depend on the particular setup (such as the amount of data and the coherent time).

The maximum covered spin-down value $|f_1|$ is important for interpreting astrophysical upper limits, since this limits the maximum “mountain height” (NS deformation) or r-mode amplitude that the search can find. The reason for this is that larger deformation (or mode amplitude) would produce a larger spin-down $|f_1|$ due to the emission of gravitational waves, which would potentially become undetectable by such a search if it was too large.

The semicoherent grid is constructed by requiring a maximum mismatch μ_M . We can estimate the maximum allowed values by calculating the required resolution for these parameters:

$$\begin{aligned} |f_{1,\max}| \equiv \delta f_1 &= \sqrt{\frac{\mu_M}{g_{f_1 f_1}}} = \sqrt{\frac{45\mu_M}{4\pi^2 T_{\text{seg}}^2 T_{\text{obs}}^2}}, \\ e_{\max} \equiv \delta e &= \sqrt{\frac{\mu_M}{g_{ee}}} = \sqrt{\frac{6\mu_M}{\pi^2 T_{\text{seg}}^2 f_0^2 \Omega^2 a_p^2}}, \end{aligned} \quad (36)$$

where $g_{f_1 f_1}$ is obtained from Ref. [31] assuming that the refinement factor is equal to N_{seg} (i.e., there are no gaps) and g_{ee} is found in Ref. [3].

It can be seen that for spin-down f_1 the maximum value depends on only the segment length T_{seg} and total amount of data, while the maximum eccentricity e depends on the frequency f_0 and orbital parameters a_p and Ω . To obtain a limit, one can take the parameters that produce the most conservative eccentricity or calculate a mean value over the parameter-space boundaries. The eccentricity equation has the exact same functional form as Eq. (43) in Ref. [15], while here we make explicit the dependence on the desired maximum mismatch.

The previous equations quantify the additional mismatch produced by a signal with nonzero spin-down f_1 and eccentricity e . However, another potential side effect would be a shift in the remaining estimated parameters due to correlations between the parameters. However, values exceeding the limits above would not automatically mean such signals are undetectable, only that the resulting mismatch would be larger, thus decreasing the sensitivity of the search.

We can compare the mismatch distribution obtained with the same grid, for four different cases: signals with $f_1 = 0$ and $e = 0$; signals with the maximum values; signals with values in between (with log-uniform distributions up to the maximum value); and signals with double the maximum value. This is shown in Fig. 5, where it can be seen that signals with parameters at the maximum value (the eccentricity maximum has been calculated using the binary parameters that give the largest eccentricity) increase the




FIG. 5. Measured mismatch for 1000 injections with random parameters, assuming one year of data from the H1 detector and segments of $T_{\text{seg}} = 900$ s. From left to right, the blue (A) histogram corresponds to signals with zero spin-down, $f_1 = 0$, and eccentricity $e = 0$; the red (B) histogram uses signals log-uniformly distributed in $f_1 \in [-4.7 \times 10^{-14}, -4.7 \times 10^{-11}]$ Hz/s and $e \in [0.2 \times 10^{-4}, 0.2 \times 10^{-1}]$; the black (C) histogram uses signals with f_1 and e at their maximum values given by Eq. (36); the yellow (D) histogram is for signals with f_1 and e given by twice the maximum values.

mean mismatch by ~ 0.08 , which would reduce the sensitivity of a search by $\sim 5\%$.

D. Computational model

In this section, we explain how the computational cost and random access memory (RAM) of our pipeline scale with different setup variables, updating and expanding Sec. V F in Ref. [15].

1. Coherent computational cost

Because of the calculation of the \mathcal{F} -statistic values, the coherent stage will have an additional computational cost, besides loading the input data and calculating the partial Hough map derivatives. In order to estimate this cost, we summarize the content of Ref. [32]. The cost of calculating the \mathcal{F} -statistic or its related quantities in segment ℓ at a single sky-patch scales with⁸

$$\tau_{\mathcal{F};\ell} = N_{T;\ell} (N_u N_{\text{dterms}} \tau_{\text{core}} + \tau_{\text{buffer}}), \quad (37)$$

⁸This assumes the so-called *demod* \mathcal{F} -statistic implementation, which is computationally favored for this type of search, but the pipeline can just as easily use the *resampling* implementation [32].

$$N_{T;\ell} = \sum_{X=1}^{N_{\text{det}}} N_{\text{SFT};\ell,X}, \quad (38)$$

where τ_{core} and τ_{buffer} are fundamental timing constants that depend on only the hardware and optimization settings (usually, τ_{buffer} is approximately one order of magnitude bigger), $2N_{\text{dterms}} + 1$ are the number of frequency bins that are used for the calculation of the Dirichlet kernel, N_u is the number of coherent u_k templates, and $N_{T;\ell}$ is the total number of SFTs in segment ℓ .

The total coherent computational cost of a single sky patch scales with the number of segments:

$$\tau_{\mathcal{F}} = \sum_{\ell=1}^{N_{\text{seg}}} \tau_{\mathcal{F};\ell} = N_{\text{SFT}}(N_u N_{\text{dterms}} \tau_{\text{core}} + \tau_{\text{buffer}}). \quad (39)$$

Since the calculations for each segment are independent from the others, these steps can be easily parallelized. We use an OpenMP loop to take advantage of multicore CPUs, which can speed up the calculation by approximately the number of used cores.

The total coherent computational cost will scale linearly with the number of sky patches.

2. Semicoherent computational cost

In the semicoherent stage, the coherent detection statistics are combined for every template that is searched.

The cost of the first stage $\tau_{H;j}$ over a sky patch j scales as

$$\tau_{H;j} = N_{\text{fs}} N_{\text{seg}} b(N_{\text{binary}}) g_j(N_{\delta}, N_{\alpha}, \delta_{u_1}, \delta_s) h(r) \tau_1, \quad (40)$$

where N_{fs} is the number of semicoherent frequency and spin-down templates, $b()$ is a function containing the nonlinear dependency on the number of binary templates N_{binary} , $g_j()$ is a function describing the effective number of semicoherent sky points needed (due to the SkyHough algorithm), N_{α} and N_{δ} are the number of right ascension and declination points, respectively, in each sky patch, δ_s is the semicoherent sky grid resolution, $h()$ is a function that depends on the threshold r set at the coherent stage, and τ_1 is a fundamental timing constant. The total semicoherent cost will scale with the number of sky patches. In the previous paper [15], it was assumed that $b = N_{\text{binary}}$ and $g_j = N_{\delta} N_{\alpha}$, which left some details out.

The function $b()$ depends on the GPU architecture. If we use a CPU, it would simply be equal to N_{binary} , but if we use a GPU, it depends nonlinearly on parameters such as the occupation of the GPU cores and the usage of shared memory. This can be seen in Fig. 7, where the nonlinear scaling with N_{binary} is clear.

The function $g_j()$ is equal to or less than $N_{\delta} N_{\alpha}$, and it encodes the SkyHough-type sky interpolation mechanism, which depends on the relation between the size of the

annulus produced by the Doppler modulations and the size of the semicoherent sky grid, as explained in Sec. IV B in Ref. [17]. At a given time stamp, the sky patches with \hat{n} more parallel to \vec{v} have wider annulus, which may contain several semicoherent sky pixels, thus lowering the number of sky points that need to be taken into account in the semicoherent loop. This effect will be different at each time stamp, and over a long observing run this will produce an average value between one and $N_{\delta} N_{\alpha}$ for the function $g_j()$. This effect gives the SkyHough algorithm a computational advantage.

The function h is different from one for a nonzero threshold r , which substitutes coherent values to 0 when below the threshold, thus reducing the computational cost. This function is given by $h = e^{-\frac{r}{\langle r \rangle}}$, where $\langle r \rangle$ is the expected value of the coherent detection statistic.

We define the average cost $\langle \tau_H \rangle$ of the first stage over different sky patches j :

$$\langle \tau_H \rangle = N_{\text{fs}} N_{\text{seg}} b(N_{\text{binary}}) \langle g_j(N_{\delta}, N_{\alpha}, \delta_{u_1}, \delta_s) \rangle h(r) \tau_1. \quad (41)$$

The cost τ_R of the second (i.e., ‘‘refinement’’) stage scales as

$$\tau_R = N_{\text{seg}} b(N_{\text{cand}}) N_a \tau_2, \quad (42)$$

where N_{cand} is the number of candidates that are passed to the second stage, N_a is the number of additional points around each candidate that are searched (when using a finer grid), $b()$ is the same function as before, and τ_2 is a fundamental timing constant.

The total semicoherent computational cost is the sum of the first and second stages.

3. Total computational cost

In order to estimate the total computational cost of a search, we add the coherent and semicoherent costs (neglecting other costs such as loading the data and writing output to files, since in a realistic scenario they are negligible):

$$\tau = \sum_{l=1}^{N_{\text{F}}} N_{\text{SP};l} (\tau_{\mathcal{F}} + \langle \tau_H \rangle_l + \tau_{R;l}), \quad (43)$$

where N_{F} is the number of frequency bands needed to cover a certain frequency range and $N_{\text{SP};l}$ is the number of sky patches at frequency band l . The values for $\langle \tau_H \rangle_l$ and $\tau_{R;l}$ depend on the frequency band, since N_{binary} , N_{δ} , N_{α} , and N_{cand} scale with frequency.

Figure 6 shows a comparison of the coherent and semicoherent costs as a function of the number of binary templates. It can be seen that the coherent cost stays constant, but the semicoherent cost increases, as expected. Past searches using BinarySkyHough and BinarySkyHou \mathcal{F} have used N_{binary} larger than 10^5 , so in such typical

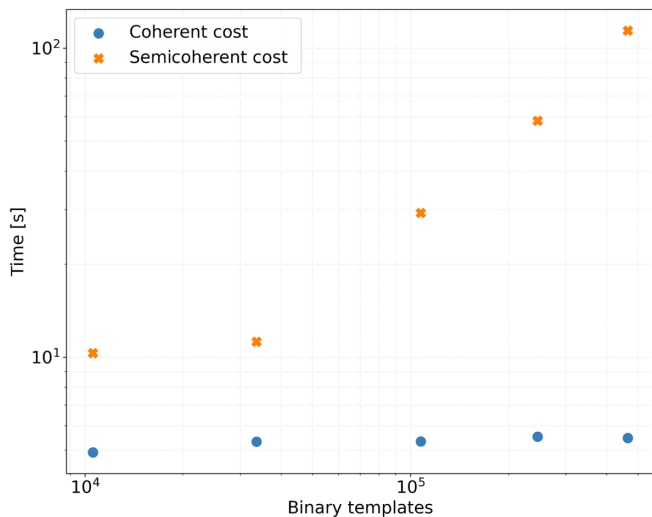


FIG. 6. Timing for a single sky patch of a 0.1 Hz frequency band as a function of the number of binary templates. The orange points show the computing time of the semicoherent stage, while the blue points are for the coherent stage (including the \mathcal{F} -statistic computation and generation of partial Hough map derivatives). These timings are obtained on an NVIDIA Tesla V100 and a single core of an Intel(R) Xeon(R) Silver 4215 CPU 2.50 GHz.

scenarios the coherent cost will be a small fraction of the total cost. If the number of binary templates is small, the coherent computational cost will have a non-negligible impact on the total cost. This also happens for isolated searches, where the search over binary parameters is substituted with a search over f_1 values, which usually is much less than 10^5 . In these two cases, calculating the \mathcal{F} -statistic might lower the sensitivity or the span of a search.

When comparing the computational cost to the previous pipeline, it is important to notice that, when using the \mathcal{F} -statistic as the coherent detection statistic, the total computational cost is reduced compared to using SFT power for multidetector searches. This is because the semicoherent summing of the power is done over $\sim N_{\text{seg}} N_{\text{det}}$ values, while the \mathcal{F} -statistic generates just N_{seg} coherent detection statistics. For this reason, in the semicoherent stage the combination of powers will take roughly N_{det} times longer than the combination of \mathcal{F} -statistic values. If the template grids are the same in both cases, the computational cost of a semicoherently dominated search using the \mathcal{F} -statistic will therefore be reduced. This improvement can be seen in Fig. 7, where the green points show the lowering of the computational cost compared to the orange points.

Furthermore, the computational efficiency of the code has been improved, mainly by better usage of CUDA's (Compute Unified Device Architecture) coalesced global memory access. This can be seen in Fig. 7, comparing the performance of the previous to the new code for the same setup as a function of the number of binary templates (increasing the RAM memory needs). The increased




FIG. 7. Timing for a single sky patch of a 0.1 Hz frequency band as a function of the number of binary templates. The blue points show the cost of the search for the BinarySkyHough pipeline using SFT power as the coherent detection statistic, the orange points show the cost for the new BinarySkyHough pipeline using the power statistic, while the green points refer to using the \mathcal{F} -statistic. All three searches use the same number of templates and the same amount of data (Gaussian noise with equal amounts of data from detectors H1 and L1). These timings are obtained on an NVIDIA Tesla V100 and a single core of an Intel(R) Xeon(R) Silver 4215 CPU 2.50 GHz.

efficiency of the new code corresponds to a lowering of the timing coefficient τ_1 in Eq. (40), manifesting as a weaker scaling with the number of binary templates.

4. RAM

In order to estimate the RAM required by BinarySkyHough \mathcal{F} , we find the scaling of the data structures in the code as a function of input parameters such as the maximum mismatch, the coherent time, and the amount of data used.

The biggest data structures in the code are

- (i) the partial Hough map derivatives, which hold the results from the coherent stage:

$$S_P = 6N_{\text{seg}}N_uN_{\text{sky}}, \quad (44)$$

where N_u is the total number of Taylor u -coordinate templates and $N_{\text{sky}} = N_\delta N_\alpha$ is the number of sky templates;

- (ii) the per-frequency bin semicoherent results:

$$S_R = 8N_{\text{binary}}N_{\text{sky}}. \quad (45)$$

These structures are orders of magnitude larger than the rest and are, therefore, enough to give a good estimate of the required RAM.

The main differences with the previous pipeline are as follows.

- (i) The number of frequency bins needed in the partial Hough map derivatives slightly decreases due to the coherent sky demodulation.
- (ii) The effective number of “segments” is reduced by the coherent multidetector combination.
- (iii) If more than one Taylor coordinate needs to be used to maintain the mismatch at a certain level, the RAM increases in order to hold the coherent results. This will limit the number of Taylor coordinates that can be used at a certain coherent time.

Another RAM limitation of the code is due to the usage of CUDA’s shared memory in the GPU kernel functions. This limits the size of the sky patches, which is dependent on the GPU architecture. The shared memory size is given by

$$S_S = 4TN_{\text{sky}}, \quad (46)$$

where T is the number of threads per block in the GPU kernel launch.

IV. SENSITIVITY AND PARAMETER ESTIMATION

In this section, we will estimate the sensitivity of the new pipeline and compare it to the previous one. In order to do this, we will compare the different detection statistics being used, showing the improvements in sensitivity that are possible due to the usage of demodulated statistics. We are not attempting to estimate a realistic sensitivity that these pipelines would achieve in an actual search, since this also depends on other postprocessing procedures, such as clustering or follow-up, which are beyond the scope of the present study.

To estimate the sensitivity, we will follow the same common procedure used in Ref. [15]. Namely, for different setups (encompassing the amount of data, detectors and their relative noise levels, maximum mismatch parameters, and coherent time) and detection statistics, we generate Gaussian noise and perform a search to obtain the threshold at a certain false-alarm probability (in the results shown, we use the top template as the threshold). Next, we add six groups (each with a different gravitational-wave amplitude) of 1000 randomly distributed signals to the previously generated Gaussian noise and perform a separate search for each signal. The resulting statistic values are compared to the threshold, and the detection probability is estimated by the fraction of signals detected (i.e., crossing that threshold) out of the total number of signals. This procedure is performed for every different detection statistic.

The injected signals have a random isotropic (NS spin) orientation, a random isotropic sky position, a random frequency f_0 between [100, 100.1] Hz, and a random period $P_{\text{orb}} \in [15, 60]$ day and semimajor axis $a_p \in [10, 40]$ ls.

The detection statistics that we compare are the original Hough number count (given by Eq. (25) in Ref. [15]), the SFT power (given by Eq. (26) in Ref. [15]), the \mathcal{F} -statistic (given by Eq. (23) in Ref. [10]), and the dominant-response statistic (given by Eq. (34) in Ref. [10]). For each of these statistics, we also compare their weighted versions discussed in Sec. III A. Here, we show the results for a single setup as an illustration, but we have tested various setups with different amounts of data and mismatch distributions and have obtained similar results.

A. Comparison of detection statistics

The left plot in Fig. 8 shows the sensitivity of two unweighted (\mathcal{F} and \mathcal{F}_{AB}) and three weighted detection

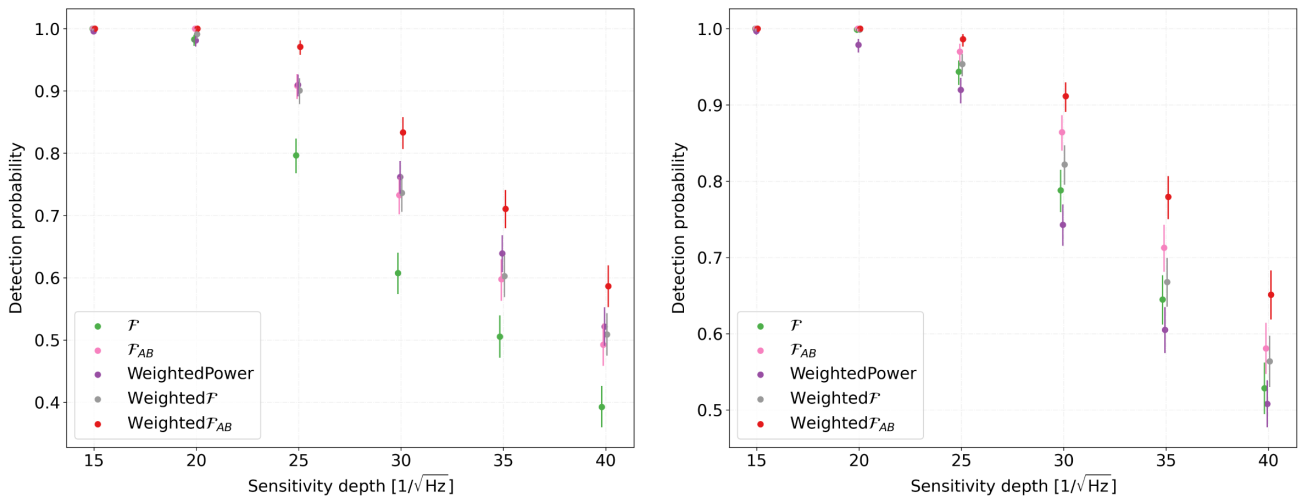


FIG. 8. Detection probability as a function of sensitivity depth $\sqrt{S_n}/h_0$ for different detection statistics in Gaussian noise. The left plot shows the results for the H1 and L1 detectors, while the right plot additionally includes the V1 detector. The error bars denote the 95% binomial confidence interval. We assume one year of data at the same noise floor from each of the detectors and a coherent segment length of $T_{\text{seg}} = 900$ s. (The points have been slightly displaced along the x axis to ease visibility.)




FIG. 9. Parameter estimation accuracy for different weighted detection statistics. We plot the cumulative distribution of parameter-space offsets between estimate and injected signal, over the six parameter-space dimensions searched. The plot uses 1000 injected signals assuming one year of data (with equal Gaussian noise floor) from both the H1 and L1 detectors and a coherent segment length of $T_{\text{seg}} = 900$ s.

statistics (SFT power, \mathcal{F} and \mathcal{F}_{AB}) for two detectors (H1 and L1), while the right plot shows the same comparison for three detectors (using V1 in addition). In both plots, we see that the most sensitive statistic is the weighted dominant-response statistic \mathcal{F}_{AB} . For the two-detector case, the sensitivities of the SFT power and \mathcal{F} -statistic are within the statistical errors of each other, but the right plot shows that for more than two detectors the \mathcal{F} -statistic is more sensitive. This is expected from the reduced number of effective segments (as discussed earlier), resulting in reduced χ^2 degrees of freedom for the background distribution (e.g., see Ref. [10]). We further see that the sensitivity of all statistics improves (on these short segments) when using weights. Overall, these results illustrate the sensitivity advantage of the demodulated (\mathcal{F} - or \mathcal{F}_{AB} -) statistics over using SFT power.

Next, we compare the parameter estimation accuracy of the different weighted detection statistics. To do this, we select the template with the highest detection statistic and compare its parameters with the parameters of the injected signal, if it was detected. The distance between injected signal and recovered “loudest” template is measured in terms of the number of grid bins along all six parameter-space dimensions. The results are shown in Fig. 9. It can be seen that the different detection statistics show a very similar behavior and that for all of them more than 90% of the signals are recovered within one bin.

B. Refinement stage

As discussed in Sec. III A, the new BinarySkyHough pipeline (like the previous BinarySkyHough) consists of




FIG. 10. Detection probability as a function of sensitivity depth $\sqrt{S_n}/h_0$ for different detection statistics in Gaussian noise. The plot shows the results for the H1 and L1 detectors, assuming one year of data at the same noise floor from each of the detectors and a coherent segment length of $T_{\text{seg}} = 900$ s. The error bars denote the 95% binomial confidence interval. (The points have been slightly displaced along the x axis to ease visibility.)

two main stages: After the initial search stage, a percentage of templates with the highest detection statistic are re-analyzed with more accurate u_{lk} interpolation expressions (see Sec. III B) and possibly with a finer mismatch. The previous BinarySkyHough pipeline also used a more sensitive detection statistic in the second stage, using the weighted power instead of the weighted Hough number count, while now we always use the most sensitive statistic (i.e., weighted \mathcal{F}_{AB}) in all stages.

The sensitivity of this refinement procedure depends on the percentage of templates that is passed to the second stage. If it is large enough, at a realistically low false-alarm probability the sensitivity of the search would be determined by the statistic used in the second stage, thus improving the overall sensitivity. In order to simulate realistic search conditions, for a candidate to count as detected in the second stage, we also require its first-stage detection statistic to be higher than the weakest candidate that was passed to the second stage.

Figure 10 shows the comparison of the weighted number count without a second stage with the result when the highest 1% of the candidates are passed to the second stage. We see that the sensitivity of the weighted number count (with the optimal threshold of⁹ 3.2) is within the uncertainty errors of the weighted power. This shows that the previous pipeline sensitivity was effectively determined by the second-stage weighted power statistic.

⁹Here the expectation value of the power statistic is 2, while in Ref. [17] it was 1, which explains the factor of 2 difference.

We further see that the sensitivity of the weighted dominant-response \mathcal{F}_{AB} -statistic is slightly improved in the refinement stage, due to the usage of the more accurate u_{lk} master equations (24) and (25). For more than two detectors, the sensitivities of the demodulated statistics are expected to be even better, similar to what we saw in the right plot in Fig. 8.

C. Increasing the coherent time

One of the advantages of the new pipeline is the ability to extend the coherent segment time T_{seg} while maintaining the same mismatch. Using a longer coherent time increases the computational cost while also increasing the sensitivity of the search. For a large number of segments, the sensitivity of a StackSlide search scales as $N_{\text{seg}}^{-1/4}$ (e.g., [6]), so at fixed false-alarm probability and amount of data, the coherent time would need to be increased by a factor of 16 in order to double the sensitivity of the search.

Here, we compare the difference in sensitivity between using a coherent time of $T_{\text{seg}} = 900$ s and $T_{\text{seg}} = 3600$ s, with fixed maximum mismatch parameters.

For this comparison, we add random signals from the binary parameter space region of $P_{\text{orb}} \in [0.785, 0.8]$ day and $a_p \in [0.5, 0.6]$ ls. In this region of the parameter space, several u_2 templates are needed in order to maintain the same coherent mismatch for the coherent time of $T_{\text{seg}} = 3600$ s, while the searches with $T_{\text{seg}} = 900$ s need only templates in u_1 .

Figure 11 shows the results by comparing the dominant-response \mathcal{F}_{AB} -statistic using these two coherent times. The improvement in sensitivity by using a longer coherent time is clear, for both the weighted and unweighted cases. For the unweighted detection statistics, the maximal expected improvement due to the 4-times increase in T_{seg} is around ~ 1.41 , which agrees approximately with what is seen in the figure.

An interesting point to observe here is that the relative improvement due to weighting of the statistics decreases for longer segment length T_{seg} (at constant Gaussian noise floors). This can be understood from the reduced differences in antenna-pattern responses between segments which therefore contribute more similar signal power.

We also show in Fig. 11 a search using the longer $T_{\text{seg}} = 3600$ s coherent segments, but without including u_2 templates. We see that this decreases its sensitivity, illustrating the need to include u_2 templates for this setup, as was expected from the mismatch estimates in Sec. III C 1.

V. CONCLUSIONS

In this paper, we have presented a new pipeline, BinarySkyHou \mathcal{F} , to search for continuous gravitational waves from unknown neutron stars in binary systems. This new pipeline is a descendant of the previous BinarySkyHough pipeline, improving several different




FIG. 11. Detection probability as a function of sensitivity depth $\sqrt{S_n}/h_0$ for different detection statistics in Gaussian noise, using two different coherent segment lengths, namely, $T_{\text{seg}} = 900$ s and $T_{\text{seg}} = 3600$ s. The plot shows the results for the H1 and L1 detectors, assuming one year of data at the same noise floor from each of the detectors. The error bars denote the 95% binomial confidence interval. (The points have been slightly displaced along the x axis to ease visibility.)

aspects. It can cover the same parameter space at a reduced computational cost, and the usage of demodulated (\mathcal{F} - and \mathcal{F}_{AB} -) statistics allows one to use longer coherent segments, which increases sensitivity and gives more flexibility to the pipeline. Furthermore, the per-detector data are combined coherently, thereby reducing the computational cost (by a factor of $\sim N_{\text{det}}$) and further improving sensitivity for searches with more than two detectors.

The new pipeline gives explicit control over the mismatch in the coherent stage, allowing one to perform lower-mismatch searches than before, for example, when following up an interesting candidate or targeting a particularly interesting smaller region of parameter space. One can now explicitly search over binary orbital parameters such as the eccentricity and argument of periaapse or higher-order frequency derivatives. Therefore, this pipeline can be used to follow up candidates from different other searches, such as from an all-sky isolated search that gives only constraints on the ranges of possible binary parameters to follow up.

BinarySkyHou \mathcal{F} also has some limitations. The Taylor coordinates used in the coherent stage allow searching only for orbital periods substantially longer than the coherent segment length, as discussed in Sec. III C 1. This limits the possibilities to search over the shortest orbital periods with longer coherent times. The corresponding RAM requirements limit the number of Taylor coordinates that can be used, which also limits the maximum coherent time that can be used in certain regions of the parameter space.

However, a number of future improvements to this pipeline can be envisaged, for example, the usage of a

nonzero threshold on the coherent statistics before summing or including only a subset of segments in the first stage depending on their sensitivity. Further enhancements could be achieved by reducing the second-stage mismatch by refining the search grid. It would also be interesting to develop an explicit optimization algorithm for the optimal search setup parameters given a certain amount of data and computational budget, similar to what was done in Refs. [6,33]. We also plan to explore the usage of further demodulated detection statistics, in particular, the line-robust statistics in Ref. [34], which should help mitigate against non-Gaussian artifacts in the data.

ACKNOWLEDGMENTS

This work has utilized the ATLAS computing cluster at the MPI for Gravitational Physics Hannover. This project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement No. 101029058.

APPENDIX: DETECTOR-FRAME TAYLOR COORDINATES

In Sec. II, we introduced short-segment SSB Taylor coordinates $\{u_k\}$, defined by the Taylor expansion of the signal phase in the SSB frame around each segment midpoint. Here, we show how similar *detector-frame* Taylor coordinates $\{u_k^X\}$ can be expressed for the phase evolution in the frame of detector X , which additionally includes the phase modulation due to the detector motion.

In complete analogy to Eq. (10), we can write the Taylor expansion in detector arrival time t around a midtime t_m :

$$\phi^X(t) = \phi_0 + 2\pi \sum_{k=1}^{k_{\max}} \frac{u_k^X}{k!} (t - t_m)^k, \quad (\text{A1})$$

with the detector-frame Taylor coordinates $\{u_k^X\}_{k=1}^{k_{\max}}$ defined as

$$u_k^X \equiv \left. \frac{1}{2\pi} \frac{d^k \phi^X}{dt^k} \right|_{t_m}, \quad (\text{A2})$$

in terms of the timing relation between source-frame (emission) time τ and detector-frame (arrival) time t , obtained by combining Eqs. (4) and (5):

$$\tau^X(t) = \tau(t_{\text{SSB}}^X(t)) = t + \vec{r}^X(t) \cdot \hat{n} - R(\tau^X). \quad (\text{A3})$$

The expressions Eq. (12) are formally identical, with the SSB derivatives $\tau^{(k)}$ replaced by detector-time derivatives $\tau^{X(k)} \equiv d^k \tau^X / dt^k$, which are found as

$$\begin{aligned} \dot{\tau}^X &= [1 + R']^{-1} (1 + \dot{\vec{r}}^X \cdot \hat{n}), \\ \ddot{\tau}^X &= [1 + R']^{-1} (\ddot{\vec{r}}^X \cdot \hat{n} - R'' \dot{\tau}^{X2}), \\ \dddot{\tau}^X &= [1 + R']^{-1} (\dddot{\vec{r}}^X \cdot \hat{n} - R''' \dot{\tau}^{X3} - 3R'' \ddot{\tau}^X \dot{\tau}^X), \\ &\vdots \end{aligned} \quad (\text{A4})$$

which generalizes Eq. (14).

We can write the first two orders explicitly as

$$\begin{aligned} \frac{u_1^X}{f_m} &= \frac{1 + \vec{v}_m^X \cdot \hat{n}}{1 + R'_m}, \\ u_2^X &= \frac{f_m}{1 + R'_m} \left(\ddot{\vec{a}}_m^X \cdot \hat{n} - R''_m \left(\frac{u_1^X}{f_m} \right)^2 \right) + f'_m \left(\frac{u_1^X}{f_m} \right)^2, \end{aligned} \quad (\text{A5})$$

with the definitions of Eqs. (17) and (18) (for small eccentricity and a single spin-down) and $\vec{v}_m^X \equiv \dot{\vec{r}}^X(t_m)$ and $\ddot{\vec{a}}_m^X \equiv \ddot{\vec{r}}^X(t_m)$ the detector velocity and acceleration at the segment midtime t_m , respectively. Again, these coordinates have units of Hz and Hz², respectively, but now also depend on the detector X and on the sky position \hat{n} of the signal.

Assuming no spin-down ($f_1 = 0$), a circular orbit ($e = 0$), and a nonrelativistic orbital velocity ($a_p \Omega \ll 1$), Eq. (A5) yields an approximate equation for the frequency-time pattern, namely,

$$\begin{aligned} u_1^X &\approx f_0 (1 + \vec{v}_m^X \cdot \hat{n}) (1 - a_p \Omega \cos \Psi_m) \\ &\approx f_0 + f_0 \vec{v}_m^X \cdot \hat{n} - f_0 a_p \Omega \cos \Psi_m, \end{aligned} \quad (\text{A6})$$

which is the same as Eq. (15) in Ref. [15], evaluated at segment midtime t_m .

[1] M. Sieniawska and M. Bejger, Continuous gravitational waves from neutron stars: Current status and prospects, *Universe* **5**, 217 (2019).
 [2] K. Riles, Searches for continuous-wave gravitational radiation, [arXiv:2206.06447](https://arxiv.org/abs/2206.06447).
 [3] 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).

[4] R. Tenorio, D. Keitel, and A.M. Sintes, Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era, *Universe* **7**, 474 (2021).
 [5] P.R. Brady and T. Creighton, Searching for periodic sources with LIGO. II. Hierarchical searches, *Phys. Rev. D* **61**, 082001 (2000).
 [6] R. Prix and M. Shaltev, Search for continuous gravitational waves: Optimal stackslide method at fixed computing cost, *Phys. Rev. D* **85**, 084010 (2012).

- [7] 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).
- [8] C. Cutler and B. F. Schutz, Generalized \mathcal{F} -statistic: Multiple detectors and multiple gravitational wave pulsars, *Phys. Rev. D* **72**, 063006 (2005).
- [9] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, *Classical Quantum Gravity* **26**, 204013 (2009).
- [10] P. Covas and R. Prix, Improved short-segment detection statistic for continuous gravitational waves, *Phys. Rev. D* **105**, 124007 (2022).
- [11] E. Goetz and K. Riles, Coherently combining data between detectors for all-sky semi-coherent continuous gravitational wave searches, *Classical Quantum Gravity* **33**, 085007 (2016).
- [12] V. Dergachev, Loosely coherent searches for sets of well-modeled signals, *Phys. Rev. D* **85**, 062003 (2012).
- [13] P. B. Covas, M. A. Papa, R. Prix, and B. J. Owen, Constraints on r-modes and mountains on millisecond neutron stars in binary systems, *Astrophys. J. Lett.* **929**, L19 (2022).
- [14] E. Goetz and K. Riles, An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems, *Classical Quantum Gravity* **28**, 215006 (2011).
- [15] P. B. Covas and A. M. Sintes, New method to search for continuous gravitational waves from unknown neutron stars in binary systems, *Phys. Rev. D* **99**, 124019 (2019).
- [16] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), First all-sky search for continuous gravitational waves from unknown sources in binary systems, *Phys. Rev. D* **90**, 062010 (2014).
- [17] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, Hough transform search for continuous gravitational waves, *Phys. Rev. D* **70**, 082001 (2004).
- [18] P. Covas and A. M. Sintes, First All-Sky Search for Continuous Gravitational-Wave Signals from Unknown Neutron Stars in Binary Systems Using Advanced LIGO Data, *Phys. Rev. Lett.* **124**, 191102 (2020).
- [19] R. Abbott *et al.*, All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems, *Phys. Rev. D* **103**, 064017 (2021).
- [20] G. Mendell and M. Landry, StackSlide and Hough Search SNR and Statistics, Technical Report No. LIGO-T05003, 2005, <https://dcc.ligo.org/LIGO-T050003/public>.
- [21] <http://www.atnf.csiro.au/research/pulsar/psrcat>.
- [22] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The Australia telescope national facility pulsar catalogue, *Astron. J.* **129**, 1993 (2005).
- [23] C. Lange, F. Camilo, N. Wex, M. Kramer, D. Backer, A. Lyne, and O. Doroshenko, Precision timing measurements of PSR J1012 + 5307, *Mon. Not. R. Astron. Soc.* **326**, 274 (2001).
- [24] C. Messenger, A semi-coherent search strategy for known continuous wave sources in binary systems, *Phys. Rev. D* **84**, 083003 (2011).
- [25] L. Nieder, B. Allen, C. J. Clark, and H. J. Pletsch, Exploiting orbital constraints from optical data to detect binary gamma-ray pulsars, *Astrophys. J.* **901**, 156 (2020).
- [26] Wikipedia Contributors, Faà di Bruno's formula—Wikipedia, The Free Encyclopedia (2022), [Online; accessed 24-July-2022].
- [27] J. Aasi *et al.*, Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data, *Phys. Rev. D* **87**, 042001 (2013).
- [28] B. Krishnan and A. M. Sintes, Hough search with improved sensitivity, Technical Report No. LIGO-T070124, 2007, <https://dcc.ligo.org/LIGO-T070124/public>.
- [29] S. v. d. Putten, H. J. Bulten, J. F. J. v. d. Brand, and M. Holthrop, Searching for gravitational waves from pulsars in binary systems: An all-sky search, *J. Phys. Conf. Ser.* **228**, 012005 (2010).
- [30] 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).
- [31] H. J. Pletsch, Parameter-space metric of semicoherent searches for continuous gravitational waves, *Phys. Rev. D* **82**, 042002 (2010).
- [32] R. Prix, Characterizing timing and memory-requirements of the \mathcal{F} -statistic implementations in LALSuite, Technical Report No. LIGO-T1600531, 2017, <https://dcc.ligo.org/LIGO-T1600531/public>.
- [33] J. Ming, M. A. Papa, B. Krishnan, R. Prix, C. Beer, S. J. Zhu, H.-B. Eggenstein, O. Bock, and B. Machenschalk, Optimally setting up directed searches for continuous gravitational waves in Advanced LIGO O1 data, *Phys. Rev. D* **97**, 024051 (2018).
- [34] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, Search for continuous gravitational waves: Improving robustness versus instrumental artifacts, *Phys. Rev. D* **89**, 064023 (2014).