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]

[1] *Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-30167 Hannover, Germany*
[2] *Leibniz Universität Hannover, D-30167 Hannover, Germany*

(Dated: 2022-08-02 17:20:50 +0200; commitID: 97b7687-CLEAN)

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, BinarySkyHou$\mathcal{F}$, which extends and improves upon the earlier BinarySkyHough method, and which was the basis for a recent search [Covas *et al.* [1]]. We compare the sensitivity and computational cost to previous methods, showing that it is both more sensitive and computationally efficient, which allows for broader and more sensitive searches.

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 resulting time-varying quadrupole moment [e.g. 2]. 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 [3] 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 [4]. 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 *semi-coherent* methods (see [5] 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, semi-coherent methods are typically more sensitive than fully-coherent matched filtering at a fixed computational budget [6, 7].

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 [8, 9]. Although this statistic was initially thought to be statistically optimal, its implicit amplitude priors were later shown to be unphysical [10]. Using more physical priors results in a more sensitive Bayes factor, albeit (currently) at an increased computational 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 was recently found [11] that is more sensitive than $\mathcal{F}$ at no extra computing cost, termed the *dominant-response* statistic $\mathcal{F}_{\text{AB}}$.

Using Fourier power over short segments directly as a coherent detection statistic is computationally cheaper still, given that no phase demodulation or other additional calculations are needed. However, one limitation of this statistic is the constrained maximum coherent length (about $T_{\text{seg}} \lesssim 30\,\text{minutes}$), 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 $\mathcal{F}_{\text{AB}}$) can naturally combine data from several detectors coherently [9, 11], this is not straightforward to achieve for short Fourier transforms [12] and is not a common practice. Therefore, constructing semi-coherent statistics on demodulated coherent statistics is generally more sensitive and flexible than using Fourier power.

Only two previous all-sky binary pipelines had been used in searches before BinarySkyHou$\mathcal{F}$ [1],

namely TwoSpect [13], and BinarySkyHough [14]. TwoSpect was the first pipeline to perform an all-sky CW search for unknown NSs in binary systems [15]. BinarySkyHough is an extension of SkyHough [16] (an all-sky pipeline for isolated systems) to NSs in binary systems, which yields higher sensitivity compared to TwoSpect thanks to its usage of GPU parallelization and a more sensitive detection statistic. Two recent all-sky binary searches [17, 18] deployed BinarySkyHough on data from Advanced LIGO's O2 and O3 observing runs, 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:

- Use of demodulated coherent statistics (such as $\mathcal{F}$- or $\mathcal{F}_{\text{AB}}$-statistic) instead of short-segment Fourier power.
- Directly summing coherent detection statistics (the typical StackSlide approach [6, 19]) instead of (thresholded) 1s and 0s as in the classical Hough algorithm [16].
- 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[1] [20]. 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 [1].

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

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

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

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

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

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

In order to obtain the phase of the signal in the frame of detector X, 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 delays[2] and the transverse proper motion of the source. We can break the timing relation into two steps, first linking the wavefront-emission time τ in the source frame to its arrival time t_{SSB} in the SSB frame, namely

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

where $R(\tau)$ is the radial distance (in light-travel time) of the source to the binary barycenter (BB) [4], with $R > 0$ when the source is further away from us than the BB. In

[1] http://www.atnf.csiro.au/research/pulsar/psrcat

[2] The numerically-implemented phase model includes these effects for the solar system but not for the binary system.

the second step we can relate the SSB time to the arrival time t at detector X by the Rømer-delay expression:

$$t^X_{\rm SSB}(t) = t + \vec{r}^X(t)\cdot\hat{n}, \tag{5}$$

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[4] as

$$R(\tau) = a_{\rm p}\left[\sin\omega(\cos E - e) + \cos\omega\sin E\sqrt{1-e^2}\right], \tag{6}$$

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

$$E = \Omega\,(\tau - t_{\rm p}) + e\sin E, \tag{7}$$

where $a_{\rm p}$ is the projected semi-major axis of the orbital ellipse (in light-travel time), $\Omega = 2\pi/P_{\rm orb}$ is the (average) orbital angular velocity (corresponding to the period $P_{\rm orb}$), e is the orbital eccentricity, $t_{\rm p}$ is the time of periapse passage and ω is the (angular) argument of periapse.

For small-eccentricity orbits, this can be approximated by the (linear in e) ELL1 model[21], namely

$$R(\tau) = a_{\rm p}\left[\sin\Psi + \frac{\eta}{2}\sin 2\Psi - \frac{\kappa}{2}\cos 2\Psi\right] + \mathcal{O}\left(e^2\right), \tag{8}$$

in terms of the Laplace-Lagrange parameters $\eta \equiv e\cos\omega$ and $\kappa \equiv e\sin\omega$ and the orbital phase

$$\Psi(\tau) = \Omega\,(\tau - t_{\rm asc}), \tag{9}$$

using the time of ascension $t_{\rm asc}$ instead of the time of periapse $t_{\rm p}$, which (to lowest order in e) are related by $t_{\rm asc} = t_{\rm p} - \omega/\Omega$. Expressions for R up to any order in e are given in Appendix C of [22].

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 $T_{\rm seg}$ 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., $T_{\rm seg} \ll P_{\rm orb}$, then in this short-segment regime we can resort to Taylor-expanding the phase (in the SSB) around each segment mid-time $t_{\rm m}$, similar to what was done in [4, 23], namely

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

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

$$u_k \equiv \frac{1}{2\pi}\,\frac{d^k\phi}{d\,t^k_{\rm SSB}}\bigg|_{t_{\rm m}}. \tag{11}$$

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

Inserting the physical phase model of Eq. (1) in the form $\phi(t_{\rm SSB}) = \phi\,(\tau(t_{\rm SSB}))$, we obtain the phase derivatives[3].

$$\begin{aligned}
u_1 &= \left[f(\tau)\,\dot\tau\right]\big|_{t_{\rm m}}, \\
u_2 &= \left[f'(\tau)\,\dot\tau^2 + f(\tau)\,\ddot\tau\right]\big|_{t_{\rm m}}, \\
u_3 &= \left[f''(\tau)\,\dot\tau^3 + 3f'(\tau)\,\ddot\tau\dot\tau + f(\tau)\,\dddot\tau\right]\big|_{t_{\rm m}}, \\
&\vdots
\end{aligned} \tag{12}$$

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

$$\begin{aligned}
R^{(1)} &= R'\,\dot\tau, \\
R^{(2)} &= R''\,\dot\tau^2 + R'\,\ddot\tau, \\
R^{(3)} &= R'''\,\dot\tau^3 + 3R''\,\ddot\tau\dot\tau + R'\,\dddot\tau, \\
&\vdots
\end{aligned} \tag{13}$$

in terms of the source-frame time derivatives $R^{(k)'} \equiv d^kR(\tau)/d\tau^k$. This analysis is complicated by the fact that the binary radial distance R of Eq. (6) is a function of source-frame (emission) time τ, not the SSB (arrival) time $t_{\rm SSB}$. In [4, 23] this difficulty could be neglected for the purposes of computing the parameter-space metric, where a slow-orbit approximation, i.e., $R(\tau) \approx R(t_{\rm SSB})$, 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

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

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

---

[3] The general form of these successive chain- and product-rule expansions is governed by the so-called Faà di Bruno's formula [24].

From these expressions we can obtain explicit Taylor coordinates u_k from Eq. (11) as

$$u_1 = \frac{f_{\rm m}}{1 + R'_{\rm m}}, \tag{15}$$

$$u_2 = -\frac{f_{\rm m}\, R''_{\rm m}}{\left(1 + R'_{\rm m}\right)^3} + \frac{f'_{\rm m}}{\left(1 + R'_{\rm m}\right)^2}, \tag{16}$$

where for our current applications (such as [1]) only up to the first- or second-order u_k will be needed in practice, as discussed in Sec. III C 1. Here we defined

$$\begin{aligned} f_{\rm m} &\equiv f(t_{\rm m}) = f_0 + f_1\,(t_{\rm m} - \tau_{\rm ref}) + \ldots, \\ f'_{\rm m} &\equiv f'(t_{\rm m}) = f_1 + f_2\,(t_{\rm m} - \tau_{\rm ref}) + \ldots, \end{aligned} \tag{17}$$

and

$$\begin{aligned} R'_{\rm m} &= a_{\rm p}\Omega\left[\cos\Psi_{\rm m} + \eta\cos 2\Psi_{\rm m} + \kappa\sin 2\Psi_{\rm m}\right], \\ R''_{\rm m} &= -a_{\rm p}\Omega^2\left[\sin\Psi_{\rm m} + 2\eta\sin 2\Psi_{\rm m} - 2\kappa\cos 2\Psi_{\rm m}\right], \end{aligned} \tag{18}$$

where $R^{(k)'}_{\rm m} \equiv R^{(k)'}(t_{\rm m})$ and

$$\Psi_{\rm m} \equiv \Omega\,(t_{\rm m} - t_{\rm asc})\,, \tag{19}$$

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

These u_1, u_2 coordinates have units of Hz and Hz^2, respectively, and depend on the physical parameters $\{\{f_k\}, a_{\rm p}, \Omega, t_{\rm asc}, e, \omega\}$. 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) of [4].

III. BINARYSKYHOU$\mathcal{F}$

In this section we present a summary of the new BinarySkyHou$\mathcal{F}$ pipeline and its main advantages over the former BinarySkyHough.

A. Summary of the previous and new pipeline

The predecessor SkyHough and BinarySkyHough algorithms are described in more detail in [16] and [14], here we only provide 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 [14, 16] 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 runtime 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) of [14]). The frequency-time pattern used for the tracks in the first stage is an approximation to the exact one, due to the usage of the look-up tables (explained in Sec. IV B of [14]). 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 [19]), and using a more accurate expression for the frequency-time pattern.

The main idea for the new BinarySkyHou$\mathcal{F}$ pipeline is to use a *demodulated* coherent detection statistic for the segments[4], such as the $\mathcal{F}$- [8, 9] or $\mathcal{F}_{\rm AB}$-statistic [11], 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 semi-coherent statistic.

Three main benefits arise from using a demodulated coherent statistic like the $\mathcal{F}$-statistic:

1. Demodulation removes the constraint on the maximum segment length, 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 is combined coherently, which also reduces the number of coherent segments needed to combine in the semi-coherent stage, again 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 $1/T_{\rm seg}$ Fourier resolution in SFTs), which will be discussed in Sec. III C 1.

[4] The $\mathcal{F}$-statistic has been used before in combination with the Hough algorithm, in [25], 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.

When combining coherent results to calculate a semi-coherent detection statistic, it has been shown that applying per-segment weights can improve the sensitivity [26]. We will not make a derivation here for the combination of $\mathcal{F}$[5], but in Sec. IV A we show that using weights also improves the sensitivity for these detection statistics. The weight w_ℓ at segment ℓ is given by:

$$w_\ell = K \frac{A_\ell + B_\ell}{S_{\mathrm{n};\ell}}, \tag{20}$$

$$K = \frac{N_{\mathrm{seg}}}{\sum_{\ell=1}^{N_{\mathrm{seg}}} w_\ell}, \tag{21}$$

$$A_\ell = \sum_{N_{\mathrm{SFTs}}} a^2, \quad B_\ell = \sum_{N_{\mathrm{SFTs}}} b^2, \tag{22}$$

where K is a normalization factor, $S_{\mathrm{n};\ell}$ is the noise power spectral density of segment ℓ, and a and b are the antenna patterns of a detector (evaluated at the mid-time of every SFT), 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 of [14]. For the dominant-response statistic, we use the following weights:

$$w_\ell = K \begin{cases} \frac{A_\ell + \frac{C_\ell^2}{A_\ell}}{S_{\mathrm{n};\ell}} & \text{if} A_\ell \geq B_\ell, \\ \frac{B_\ell + \frac{C_\ell^2}{B_\ell}}{S_{\mathrm{n};\ell}} & \text{otherwise}, \end{cases} \tag{23}$$

where $C_\ell = \sum_{N_{\mathrm{SFTs}}} a\, b$.

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 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) 6-dimensional parameter space grid[6] over $\{\alpha, \delta, f_0, a_\mathrm{p}, \Omega, t_{\mathrm{asc}}\}$, while using the Taylor coordinates we effectively only need to use three (or four) for the short segments currently considered, namely $\{\alpha, \delta, u_1\}$. This reduces complexity (the parameter-space metric in physical parameters would be hugely degenerate for short segments (cf. [4, 23]) and lowers the resulting computational cost[7].

The u_k template bank is constructed as a hyper-cubic 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 [14] to process the semi-coherent 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. Semi-coherent 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 semi-coherent stage, on the other hand, we are using physical coordinates to combine the per-segment statistics, namely $\{\alpha, \delta, \{f_k\}, a_\mathrm{p}, \Omega, t_{\mathrm{asc}}, e, \omega\}$. For every semi-coherent template, we therefore need to calculate the appropriate *mapping* to the corresponding Taylor coordinates $\{u_k\}$ and coherent sky position.

In addition to using different signal parameters, the semi-coherent 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 semi-coherently (typically using nearest-neighbor interpolation). This is a generic feature of the semi-coherent approach (cf. [7, 28]), and in SkyHough-derived pipelines takes the form of the so-called *master equation* [14, 16], 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 semi-coherent sky-template in the same sky patch. The resulting offset in sky-position between the semi-coherent 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, derived in the appendix and shown in Eq. (A5). These are valid as long as the coherent segments are short compared to a day. Using this we can express the induced shifts δu_k^X as

$$\begin{aligned} \delta u_1^X &= u_1\, \vec{v}^X_{\mathrm{m}'} \cdot \delta\hat{n}, \\ \delta u_2^X &\approx u_1\, \vec{a}^X_{\mathrm{m}'} \cdot \delta\hat{n} + 2\vec{v}^X_{\mathrm{m}'} \cdot \delta\hat{n}\, (-u_1\, R''_\mathrm{m} + f'_\mathrm{m}) \ , \end{aligned} \tag{24}$$

in terms of detector velocity $\vec{v}^X_{\mathrm{m}'}$ and acceleration $\vec{a}^X_{\mathrm{m}'}$ at the segment mid-time $t_{\mathrm{m}'}$[8]. To remove the detector

[5] Although we leave this for future work, the improvement of sensitivity with weights can be understood from the per-segment change in signal power, which can vary around one order of magnitude between segments for such short coherent times.

[6] In comparison to BinarySkyHough the new code is able to also search over spindowns and eccentricity in the semi-coherent stage.

[7] Using a Taylor phase approximation to lower the computational cost of an all-sky binary search has been first explored by the *Polynomial* pipeline [27], which did not use physical parameters for the semi-coherent summation, however, resulting in reduced sensitivity.

[8] The time $t_{\mathrm{m}'}$ where the detector-related quantities are evaluated is different than the time where the Taylor coordinates are evaluated: $t_\mathrm{m} = t_{\mathrm{m}'} + \vec{r}_{\mathrm{m}'} \cdot \hat{n}_\mathrm{d}$, where $\hat{n}_\mathrm{d}$ is the coherent demodulation point and we take a mean between the different detectors: $N_{\mathrm{det}}\, \vec{r}_{\mathrm{m}'} \equiv \sum_X \vec{r}^X_{\mathrm{m}'}$.

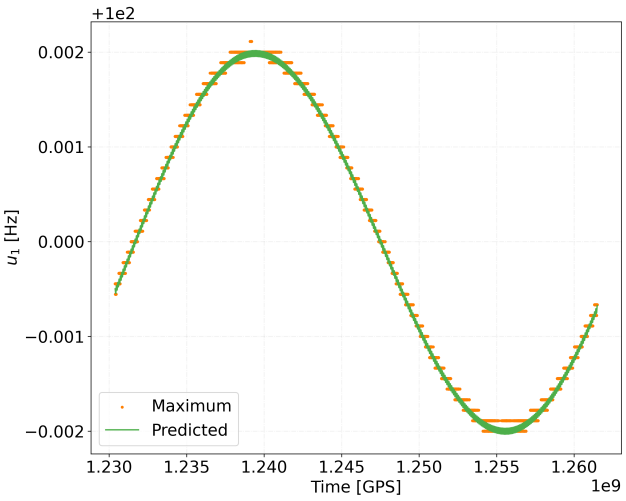

FIG. 1. Comparison between the u_{I1} coordinate per segment given by Eq. (25) (green line) and the u_1 values per segment that minimize the mismatch (orange points). The sky position has been shifted by 0.3 rad in both α and δ. This example uses one year of data from the H1 detector with segments of $T_{\mathrm{seg}} = 900$ s, and a signal with a frequency of $f_0 = 100\,\mathrm{Hz}$ without spindown and without binary modulation.

dependency we average over detectors, which will be a good approximation for δu_1, given the detector velocity is dominated by the (detector-independent) orbital motion of the Earth. On the other hand, the detector acceleration $\vec{a}^X_{\mathrm{m}}$ in δu^X_2 is dominated by the Earth's rotation, so averaging over detectors might be more dangerous, but should work well as long as the detectors are not too far separated, such as for H1 and L1.

Therefore we can arrive at the following generalized master equations

$$u_{\mathrm{I1}} = u_1\ (1 + \vec{v}_{\mathrm{m}'} \cdot \delta\hat{n})\,, \tag{25}$$

$$u_{\mathrm{I2}} = u_2 + u_1\,\vec{a}_{\mathrm{m}'} \cdot \delta\hat{n} + 2\vec{v}_{\mathrm{m}'} \cdot \delta\hat{n}\,(-u_1\,R''_{\mathrm{m}} + f'_{\mathrm{m}})\,, \tag{26}$$

with detector-averaged velocity $N_{\mathrm{det}}\,\vec{v} \equiv \sum_X \vec{v}^X$ and acceleration $N_{\mathrm{det}}\,\vec{a} = \sum_X \vec{a}^X$. Eq. (25) agrees with the previous Hough-on-$\mathcal{F}$-statistic master equation found in [16] and [25] (with implicit detector averaging).

This is shown in Fig. 1, where the u_1 values with the highest signal power are plotted as a function of segment time t_{m} for an offset $\delta\hat{n} = \hat{n} - \hat{n}_{\mathrm{d}}$ between the signal sky-position $\hat{n}$ and the coherent demodulation point $\hat{n}_{\mathrm{d}}$. In addition we plot the predicted track of shifted u_{I1} given by Eq. (25). These u_{I1} values closely follow the path that minimizes the mismatch. The mismatch in Fig. 1 between the path followed by u_{I1} and the maximum path is around 0.1%, whereas the mismatch between the non-shifted u_1 and the maximum path would be around 0.8.

The bounds in each of the Taylor coordinates are found using Eqs. (25) and (26), by calculating the maximum possible values over the given physical parameter space.

The Hough semi-coherent summation used in the first stage of the pipeline uses the following expressions for computational reasons (due to the look-up table approach [16]), which approximate Eqs. (25) and (26):

$$u_{\mathrm{IH1}} = f_{\mathrm{m}}(1 - R'_{\mathrm{m}}) + f_{\mathrm{H}}\vec{v}_{\mathrm{m}'} \cdot \delta\hat{n}, \tag{27}$$

$$u_{\mathrm{IH2}} = u_2, \tag{28}$$

where

$$f_{\mathrm{H}} \equiv f_{0\mathrm{H}} + f_{1\mathrm{H}}(t_{\mathrm{mid}} - \tau_{\mathrm{ref}}) + \ldots, \tag{29}$$

being $f_{k\mathrm{H}}$ the values at the middle of the range being searched and t_{mid} the mid-time of the analyzed 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

$$\mu = 1 - \frac{\rho^2_{\mathrm{r}}}{\rho^2}, \tag{30}$$

where ρ^2 is the full signal power (given by Eq. 20 in [11]) and ρ^2_{r} 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 semi-coherent 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 semi-coherent mismatch is produced by signal-template offsets in the semi-coherent 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 order u_2, which turns out to be sufficient for currently considered practical all-sky binary searches (similar to the recent search [1] 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 order u_1 or u_2.

The mismatch μ_{u_k} due to omission of order u_k (and higher) can be easily 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, 57) of [4]. Using this we can express the mismatch due to

truncation of $u_{k\geq 2}$ or $u_{k\geq 3}$ as

$$\begin{aligned}\mu_{u_2} &\approx g_{22}\, u_2^2 = \frac{\pi^2\, T_{\rm seg}^4}{180}\, u_2^2,\\ \mu_{u_3} &\approx g_{33}\, u_3^2 = \frac{\pi^2\, T_{\rm seg}^6}{4032}\, u_3^2.\end{aligned} \tag{31}$$

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

$$\langle \mu_{u_2} \rangle \approx \frac{\pi^2\, T_{\rm seg}^4}{360}\, \left(f_0^2 a_{\rm p}^2 \Omega^4 + 2 f_1^2 \right), \tag{32}$$

$$\langle \mu_{u_3} \rangle \approx \frac{\pi^2\, T_{\rm seg}^6}{8064}\, f_0^2 a_{\rm p}^2\, \Omega^6\,, \tag{33}$$

which illustrate the fact that the segments must be short compared to the orbital period, i.e., $T_{\rm seg}\, \Omega \ll 1$, in order for the Taylor-coordinates u_k to be a good approximation, as discussed in [4].

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

$$T_{{\rm seg},u_1} \leq \left(\frac{360\, \langle \mu_u \rangle}{\pi^2 \left(f_0^2 a_{\rm p}^2 \Omega^4 + 2 f_1^2 \right)} \right)^{1/4}, \tag{34}$$

and similarly for truncation after u_2 we obtain the constraint

$$T_{{\rm seg},u_2} \leq \left(\frac{8064\, \langle \mu_u \rangle}{\pi^2 f_0^2 a_{\rm p}^2 \Omega^6} \right)^{1/6}. \tag{35}$$

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.

Fig. 3 further shows a test of the mismatch given by Eq. (32), by generating 1000 different signals with a frequency of 300 Hz and random orbital parameters log-uniformly distributed between $[0.1, 1]$ days and $[0.01, 1]$ l-s for Ω and $a_{\rm p}$ respectively. 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 order u_1. The corresponding measured mismatch is then compared to the model prediction of Eq. (32). The figure shows that these equations correctly predict the measured mismatch, in the range where the metric approximation is expected to be accurate (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 semi-coherent template grids. If we count the Taylor-truncation mismatch μ_u as part of the coherent mismatch $\mu_{\rm c}$, then the total average (over the template bank) mismatch $\langle \mu \rangle$ will be given approximately by the sum of the mean coherent $\langle \mu_{\rm c} \rangle$ and mean semi-coherent mismatch $\langle \mu_{\rm s} \rangle$, as shown in [7], namely

$$\langle \mu \rangle \simeq \langle \mu_{\rm c} \rangle + \langle \mu_{\rm s} \rangle\,. \tag{36}$$

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

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 semi-coherent 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.

3. Maximum spindown and eccentricity

Although BinarySkyHou$\mathcal{F}$ is also able to search the f_1 and eccentricity parameters in the semi-coherent summation, all-sky searches for neutron stars in binary systems are so computationally expensive that one will usually not be able to explicitly search over these parameters. For this reason, we want to know the values to which the pipeline is still sensitive in this case, which will depend on the particular set-up (such as the amount of data and the coherent time). Here we provide expressions and some tests that help to quantify the limits of a search to these parameters.

The maximum value of f_1 allowed by the search is important when the astrophysical upper limits are interpreted, since this value sets a limit to the maximum ellipticity (deformation) or r-mode amplitude that can be found by a search. This is because a higher value of these quantities would produce a higher f_1 (due to the emission of gravitational waves), which would not be detectable by such search.

The semi-coherent grid is constructed by requiring a maximum mismatch $\mu_{\rm M}$. We can estimate the maximum allowed values by calculating the resolution for these parameters:

$$f_{1;\rm max} \equiv \delta f_1 = \sqrt{\frac{\mu_{\rm M}}{g_{f_1 f_1}}} = \sqrt{\frac{45 \mu_{\rm M}}{4\pi^2 T_{\rm seg}^2 T_{\rm obs}^2}}, \tag{37}$$

$$e_{\rm max} \equiv \delta e = \sqrt{\frac{\mu_{\rm M}}{g_{ee}}} = \sqrt{\frac{6 \mu_{\rm M}}{\pi^2 T_{\rm seg}^2 f_0^2 \Omega^2 a_p^2}}, \tag{38}$$

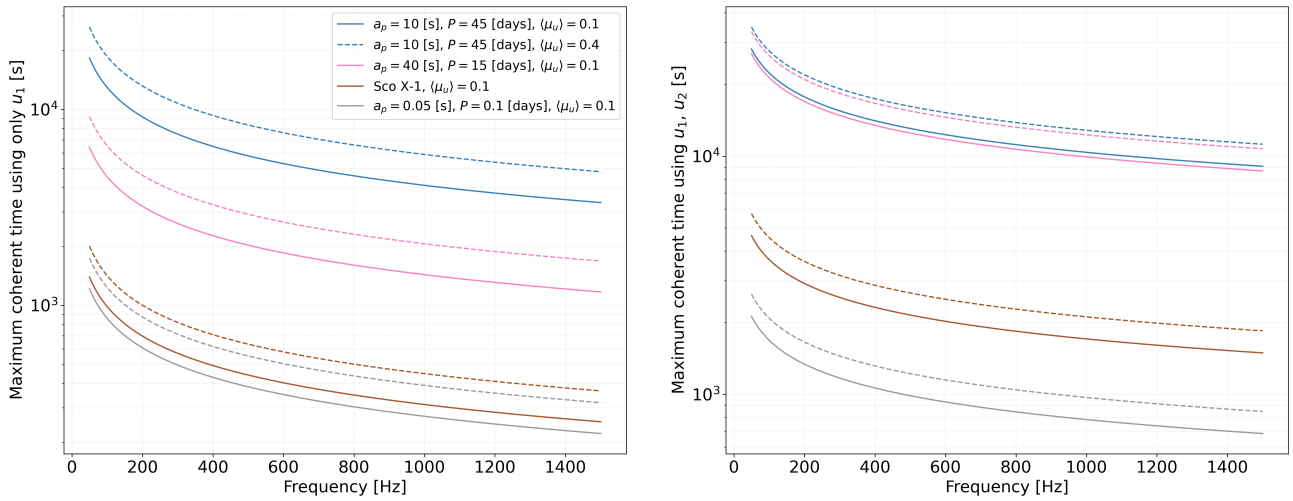

FIG. 2. The left plot shows the maximum coherent time that can be used at a given region of the binary parameters and for a certain mismatch due to the Taylor truncation (we show the results for $\langle\mu_u\rangle = 0.1$ with the lines and $\langle\mu_u\rangle = 0.4$ with the dashed lines) using only the u_1 Taylor coordinate for the coherent demodulation. The right plot shows the same but if the coherent demodulation is done using the u_1 and u_2 Taylor coordinates. All the traces have been calculated with $f_1 = 0$, which makes the resulting $T_{\rm seg}$ more conservative.

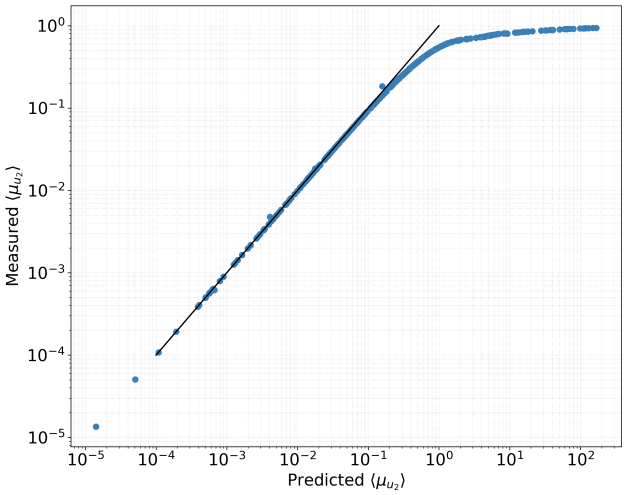

FIG. 3. Measured (averaged over the segments of the search) against predicted coherent mismatch given by equation (32) for 1000 injections, using the H1 detector with 1 year of data and segments of $T_{\rm seg} = 900$ s. The black line shows the values where the predicted and measured mismatch are equal.

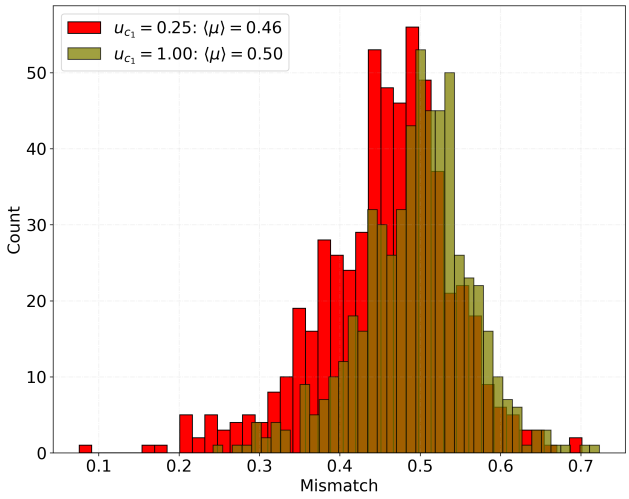

FIG. 4. Mismatch histograms for 1000 injections with random parameters, using the H1 and L1 detectors with 1 year of data and segments of $T_{\rm seg} = 900$ s. The green (rightmost) histogram shows the result with a grid resolution of $\delta u_1 = 1/T_{\rm seg}$, while the red (leftmost) histogram shows the result with a finer resolution of $\delta u_1 = 0.25/T_{\rm seg}$. Only u_1 and a single sky position are searched in the coherent stage. The semi-coherent grids are equal for both cases.

where $g_{f_1 f_1}$ is obtained from [29] assuming that the refinement factor is equal to $N_{\rm seg}$ and that there are no gaps, and g_{ee} is obtained from [4].

It can be seen that for f_1 the calculation is straightforward, while for e there is a dependence on the frequency and orbital parameters. To calculate a limit, one can take the parameters that produce the most conservative eccentricity, or calculate a mean value of the parameter space boundaries. The eccentricity equation has the exact same dependencies as Eq. 43 of [14], while now we make explicit the dependence on the desired maximum mismatch.

The previous equations only quantify the mismatch that would be produced by the f_1 and eccentricity values themselves. When this mismatch is added to the

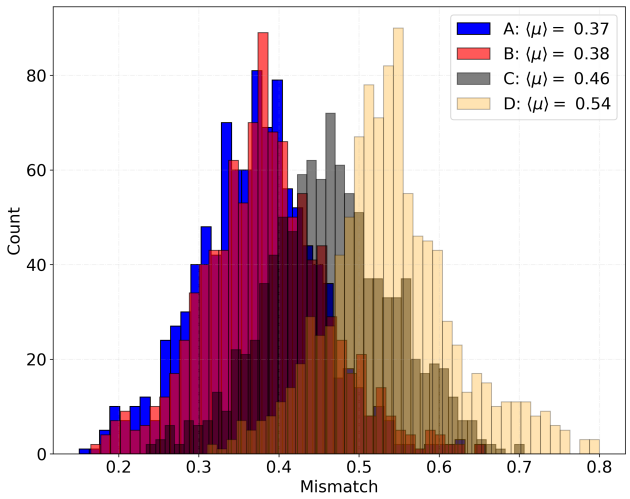

FIG. 5. This plot shows the measured mismatch for 1000 injections with random parameters, using the H1 detector with 1 year of data and segments of $T_{\rm seg} = 900$ s. From left to right, the blue (A) histogram shows signals with $f_1 = 0$ and $e = 0$; the red (B) histogram shows signals with f_1 and e log-uniformly respectively distributed between $[-4.7\times10^{-14}, -4.7\times10^{-11}]$ Hz/s and $[0.2\times10^{-4}, 0.2\times10^{-1}]$; the black (C) histogram shows signals with f_1 and e given by the maximum value; the yellow (D) histogram shows signals with f_1 and e given by twice the maximum value.

mismatch produced by the other parameters, the results could be slightly different, due to correlations between some of the parameters. Furthermore, a value higher than the one given by these expressions does not mean undetectability, only that the mismatch distribution would be shifted to higher values, 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); 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 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 set-up variables, updating and expanding Sec. VF of [14].

1. *Coherent computational cost*

Due to 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 [30]. The cost of calculating the $\mathcal{F}$-statistic or its related quantities in segment ℓ at a single sky-patch scales with[9]:

$$\tau_{\mathcal{F};\ell} = N_{\mathrm{T};\ell}(N_{u_1} N_{\rm dterms}\tau_{\rm core} + \tau_{\rm buffer}), \tag{39}$$

$$N_{\mathrm{T};\ell} = \sum_{X=1}^{N_{\rm det}} N_{\mathrm{SFT};\ell,X}\,, \tag{40}$$

where $\tau_{\rm core}$ and $\tau_{\rm buffer}$ are fundamental timing constants that only depend on the hardware and optimization settings (usually $\tau_{\rm buffer}$ is approximately one order of magnitude bigger), $2N_{\rm dterms}+1$ are the number of frequency bins that are used for the calculation of the Dirichlet kernel, N_{u_1} is the number of coherent u_1 templates, and $N_{\mathrm{T};\ell}$ is the total number of SFTs in segment ℓ. The previous equation assumed that all Taylor coordinates above u_1 are effectively 0. If a grid of templates is needed for those coordinates, the cost will scale linearly with the number of templates higher than u_1.

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

$$\tau_{\mathcal{F}} = \sum_{\ell=1}^{N_{\rm seg}} \tau_{\mathcal{F};\mathrm{l}} = N_{\rm SFT}(N_{u_1} N_{\rm dterms}\tau_{\rm core} + \tau_{\rm buffer}). \tag{41}$$

Since the calculations done for each segment are independent from the rest, this can be easily parallelized. We use an OpenMP loop to take advantage of multi-core 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. *Semi-coherent computational cost*

In the semi-coherent stage the coherent detection statistics are combined for every different template that is searched.

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

$$\tau_{\mathrm{H};j} = N_{\rm fs} N_{\rm seg}\, b(N_{\rm binary}) g_j(N_\delta, N_\alpha, \delta_{u_1}, \delta_s) h(r)\tau_1, \tag{42}$$

where $N_{\rm fs}$ is the number of semi-coherent frequency and spindown templates, b is a function containing the non-linear dependency on the number of binary templates

[9] Although here we explain the so-called *demod* implementation, our pipeline can also use the *resamp* implementation [30].

$N_{\rm binary}$, g is a function describing the effective number of semi-coherent sky points that are needed (due to the SkyHough algorithm), N_α and N_δ are the number of right ascension and declination points in each sky-patch, δ_s is the semi-coherent 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 semi-coherent cost will scale with the number of sky-patches.

In the previous paper [14] it was assumed that $b = N_{\rm binary}$ and $g = N_\delta N_\alpha$, which left some details out.

The function b depends on the GPU architecture. If we would use a CPU, it would simply be equal to $N_{\rm binary}$, but if we use a GPU it depends non-linearly 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 non-linear scaling with $N_{\rm binary}$ is clear.

The function g is equal 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 semi-coherent sky grid, as explained in Sec. IVB of [16]. At a given timestamp, the sky-patches with $\hat{n}$ more parallel to $\vec{v}$ have wider annulus, which may contain several semi-coherent sky pixels, thus lowering the number of sky points that need to be taken into account in the semi-coherent loop. This effect will be different at each timestamp, and over a long observing run this will produce an average value between 1 and $N_\delta N_\alpha$ for the function g. This effect gives the SkyHough algorithm a computational advantage.

The function h is different than 1 for a non-zero 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 of the first stage $\langle \tau_{\rm H} \rangle$ over different sky-patches j:

$$\langle \tau_{\rm H} \rangle = N_{fs} N_{\rm seg} b(N_{\rm binary}) \left\langle g_j(N_\delta, N_\alpha, \delta_{u_1}, \delta_s) \right\rangle h(r) \tau_1. \tag{43}$$

The cost of the second stage τ_R scales as:

$$\tau_{\rm R} = N_{\rm seg} b(N_{\rm cand}) N_{\rm a} \tau_2, \tag{44}$$

where $N_{\rm cand}$ is the number of templates that are passed to the second stage, $N_{\rm a}$ is the number of additional points around each template that are searched (in case a refinement is done), b is the same function as before, and τ_2 is a fundamental timing constant.

The total semi-coherent computational cost is the sum of the first and second stage.

3. *Total computational cost*

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

$$\tau = \sum_{l=1}^{N_{\rm F}} N_{{\rm SP};l} (\tau_{\mathcal{F}} + \langle \tau_{\rm H} \rangle_l + \tau_{{\rm R};l}), \tag{45}$$

where $N_{\rm F}$ is the number of frequency bands that are needed to cover a certain frequency range, and $N_{{\rm SP};l}$ is the number of sky-patches at frequency band l.

$\langle \tau_{\rm H} \rangle_l$ and $\tau_{{\rm R};l}$ depend on the frequency band since $N_{\rm binary}$, N_δ, N_α, and $N_{\rm cand}$ scale with the frequency.

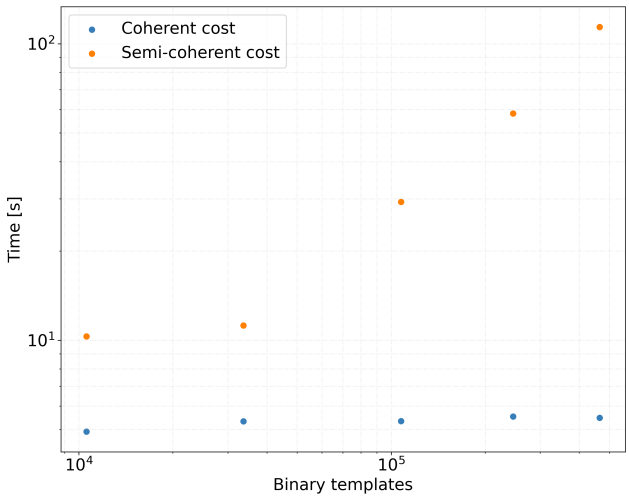

FIG. 6. Timing for a single sky-patch of a 0.1 Hz frequency band. The orange points show the cost of the semi-coherent stage, while the blue points show the cost of the coherent stage (including the $\mathcal{F}$-statistic computation and the generation of the partial Hough map derivatives). We have used a Tesla V100 and a Intel(R) Xeon(R) Silver 4215 CPU 2.50 GHz (using a single core) to obtain the timings.

Fig. 6 shows a comparison of the coherent and semi-coherent costs as a function of the number of binary templates. It can be seen that the coherent cost stays constant, but the semi-coherent cost increases, as expected. Previous searches using the BinarySkyHough or BinarySkyHou$\mathcal{F}$ pipelines have used $N_{\rm binary}$ larger than 10^5, so it can be seen that 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 $\mathcal{F}$-statistic as the coherent detection statistic, the total computational cost is reduced when compared to using SFT power for multi-detector searches. This is because the semi-coherent summing of the power is done over

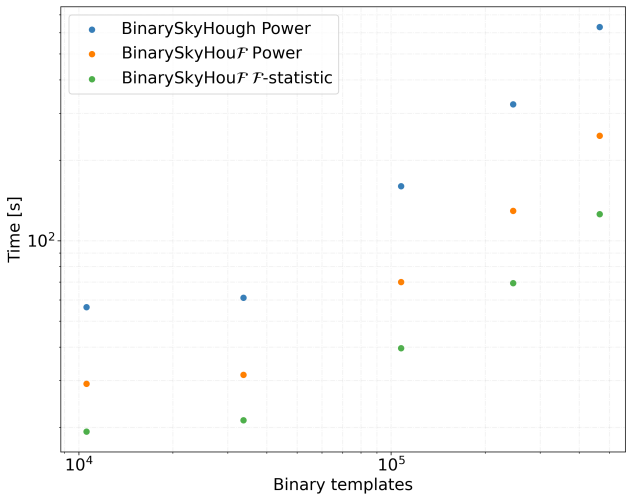

FIG. 7. Timing for a single sky-patch of a 0.1 Hz frequency band. The blue points show the cost of the search for the previous pipeline using the SFT power as the coherent detection statistic, the orange points show the cost for the new pipeline, and the green points show the cost for the new pipeline using the $\mathcal{F}$-statistic as the coherent detection statistic. The three searches have used the same number of templates, the same amount of data (Gaussian noise with equal amount of data from detectors H1 and L1), and a threshold of 0. We have used a Tesla V100 and a Intel(R) Xeon(R) Silver 4215 CPU 2.50 GHz (using a single core) to obtain the timings.

$\sim N_{\rm seg}N_{\rm det}$ values, while the $\mathcal{F}$-statistic generates $N_{\rm seg}$ coherent detection statistics (assuming equal timestamps for all detectors). For this reason, in the semi-coherent stage the combination of powers will take roughly $N_{\rm det}$ more times than the combination of $\mathcal{F}$-statistic values. If the template grids are the same in both cases, it can be seen that the computational cost of a semi-coherently dominated search using the $\mathcal{F}$-statistic will 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.

Besides this improvement, we have also modified the code in order to improve its efficiency[10]. Fig. 7 shows a comparison between the old and new code, using the exact same set-up. It can be seen that as the number of binary templates is increased (the size of the occupied RAM is increased), the new code becomes more efficient. This efficiency improvement is translated to a lowering of the fundamental timing constant τ_1.

[10] The speed-up of the code is mainly given by a better usage of CUDA's coalesced global memory access.

4. RAM

We want to estimate the RAM required by BinarySkyHou$\mathcal{F}$. To do so, 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 size of the data.

The biggest data structures in the code are:

- The partial Hough map derivatives, which hold the results from the coherent stage:

$$S_{\rm P} = 6N_{\rm seg}N_{u_1}N_{u_{k\geq 2}}N_{\rm sky}, \tag{46}$$

  where $N_{u_{k\geq 2}}$ is the number of templates for Taylor coordinates of order higher than 1, and $N_{\rm sky} = N_\delta N_\alpha$.

- The per-frequency bin semi-coherent results:

$$S_{\rm R} = 8N_{\rm binary}N_{\rm sky}. \tag{47}$$

These structures are orders of magnitude larger than the rest and are enough to give an estimate of the required RAM.

The main differences with the previous pipeline are:

- The number of frequency bins needed in the partial Hough map derivatives slightly decreases due to the coherent sky demodulation.
- Due to the coherent multi-detector combination, the number of segments is reduced.
- 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:

$$S_{\rm S} = 4TN_{\rm sky}, \tag{48}$$

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 different detection statistics, showing the improvements in sensitivity that are possible due to the usage of demodulated statistics.

We are not attempting to estimate the real sensitivity that these pipelines could have in an actual search, since that number also depends on other post-processing procedures, such as clustering or follow-up, which we do not study in this paper.

To estimate the sensitivity, we will follow the same common procedure that was used in [14]. For every different setup (amount of data, detectors and their relative sensitivity, maximum mismatch, coherent time) and detection statistic we generate Gaussian noise and do a search to obtain the threshold at a certain false alarm probability (in the results shown, we use the top template as the threshold). Afterwards, we add 6 groups (each with a different gravitational-wave amplitude) of 1000 randomly distributed signals to the previously generated Gaussian noise, and perform a separated search for each signal to calculate the detection statistic values, which are then compared to the threshold. The detection probability is then obtained by counting the number of detected signals and dividing by the total number of signals. This is done for every different detection statistic. The signals that we add have a random isotropic orientation, a random isotropic sky position, a random f_0 between $[100, 100.1]$ Hz, and a random $P_{\rm orb}$ and $a_{\rm p}$ between $[15, 60]$ days and $[10, 40]$ l-s, respectively.

The detection statistics that we want to compare are the Hough number count (given by Eq. 25 of [14]), the power (given by Eq. 26 of [14]), the $\mathcal{F}$-statistic (given by Eq. 23 of [11]), and the dominant-response statistic (given by Eq. 34 of [11]). For each of these statistics we also compare their weighted versions. We show the results for a single set-up, but we have tried set-ups with different amount of data and mismatch distributions and have obtained similar results.

A. Comparison of detection statistics

In this subsection we show a comparison of different detection statistics. The left plot of Fig. 8 shows the sensitivity of two unweighted and three weighted detection statistics when using two detectors, while the the right plot shows the comparison when using three. Both plots show that the most sensitive statistic is the weighted dominant-response statistic. For the two detectors case, the sensitivity of the power and $\mathcal{F}$-statistic is within the statistical errors, but the right plot shows that for more than two detectors the $\mathcal{F}$-statistic is more sensitive, as expected from the smaller number of degrees of freedom of the background distribution, since the power statistic combines semi-coherently the coherent per-detector results, while the $\mathcal{F}$-statistic combines coherently the coherent per-detector results. For this reason, for two detectors the number of degrees of freedom is equal for both statistics, while for more than two detectors it is less for the $\mathcal{F}$-statistic. The plot also shows that the detection statistics improve their sensitivity (for this short coherent time) when weights are applied. From these results the advantage of being able to use the $\mathcal{F}$-statistic or the dominant-response statistic is clear.

We also compare the parameter estimation of the weighted detection statistics. To do this, we select the template with the highest detection statistic, and compare its parameters with the parameters of the artificial signal. We do this for the six different searched dimensions, and only for the signals that have been detected. The results are shown in Fig. 9. It can be seen that the different detection statistics show a very similar behaviour. It can also be seen that for all the detection statistics, more than 90% of the signals are recovered within one bin.

B. Refinement stage

Our pipeline consists of two stages, where a percentage of the templates with the highest detection statistic are reanalyzed in the second stage with accurate $u_{\mathrm{I}k}$ expressions, and possibly with a mismatch refinement. The BinarySkyHough pipeline used the second stage to reanalyze the templates with a more sensitive detection statistic, using the weighted power instead of the weighted number count.

In order to do a fair comparison with the previous pipeline, we need to compare the new detection statistics with the weighted number count with a second stage where the weighted power is used. The sensitivity of this procedure depends on the percentage of templates that is passed to the second stage. If it is high enough, at a realistic low false alarm probability the sensitivity of the search would be given by the statistic used at the second stage, thus increasing the sensitivity. For a candidate to count as detected when using a second stage, we also require that the detection statistic of the first stage is higher than the threshold set by the last candidate that was sent to the second stage.

Fig. 10 shows the comparison of the weighted number count without a second stage with the result when 1% of the templates have been passed to the second stage. It can be seen that the sensitivity of the weighted number count (we have used a threshold of 3.2[11], the optimal value) is within the uncertainty errors of the weighted power. This shows that due to the second stage, the previous pipeline sensitivity was effectively given by the weighted power statistic. It can also be seen that the sensitivity of the weighted dominant-response statistic is slightly improved with the refinement stage, due to the usage of the exact $u_{\mathrm{I}k}$ equations. As the right plot of Fig. 8 showed, for more than two detectors the sensitivities of these detection statistics would be more different.

[11] In our case the expectation value of the power statistic is 2, while in [16] the expected value was 1. This is the reason for the factor of 2 difference.

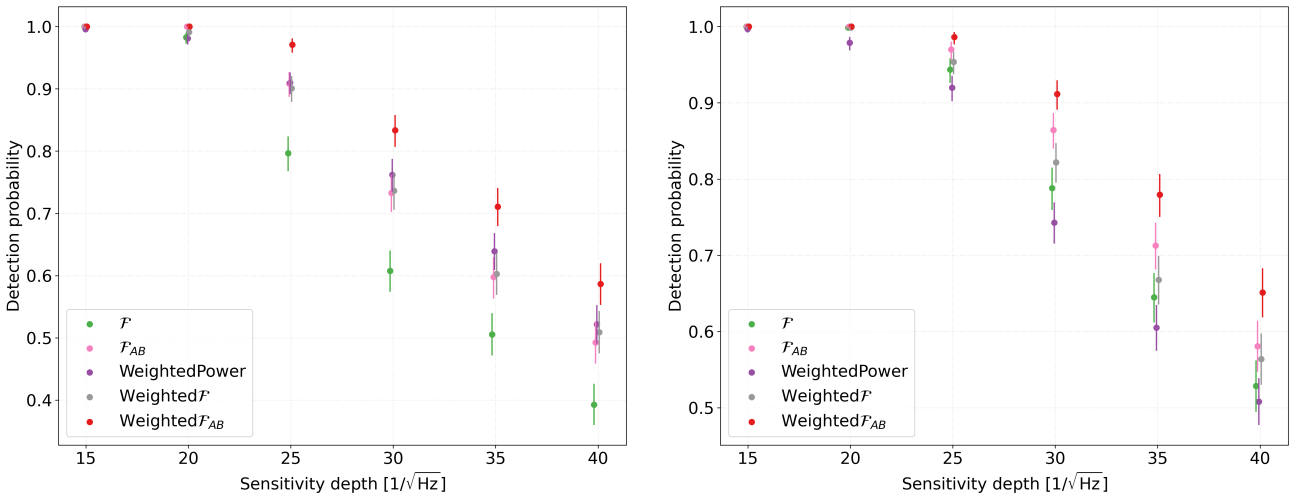

FIG. 8. Detection probability as a function of the sensitivity depth for 5 different detection statistics. The points have been slightly displaced in the x-axis to ease visibility. The error bars show the 95% binomial confidence interval. The left plot shows the results using Gaussian noise from the H1 and L1 detectors, while the right plot shows the results using Gaussian noise from the H1, L1, and V1 detectors. All detectors have the same amplitude spectral density and one year of data. We have used a coherent time of 900 s.

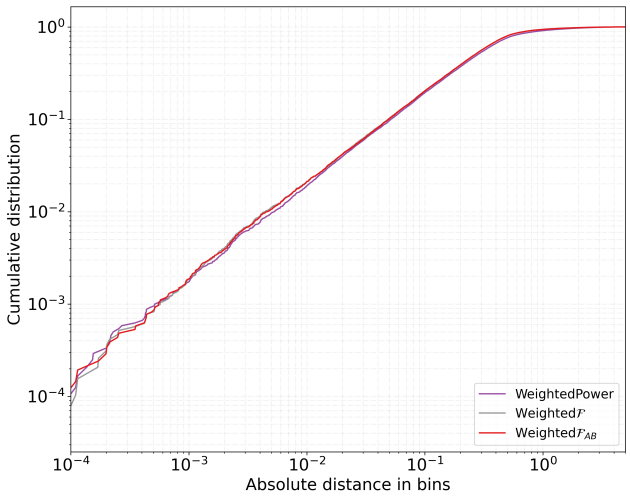

FIG. 9. Parameter estimation for 3 different weighted detection statistics. We show the cumulative distribution for 1000 artificial signals, joining together the 6 different searched parameters. The plot shows the results using Gaussian noise from the H1 and L1 detectors. All detectors have the same amplitude spectral density and one year of data. We have used a coherent time of 900 s.

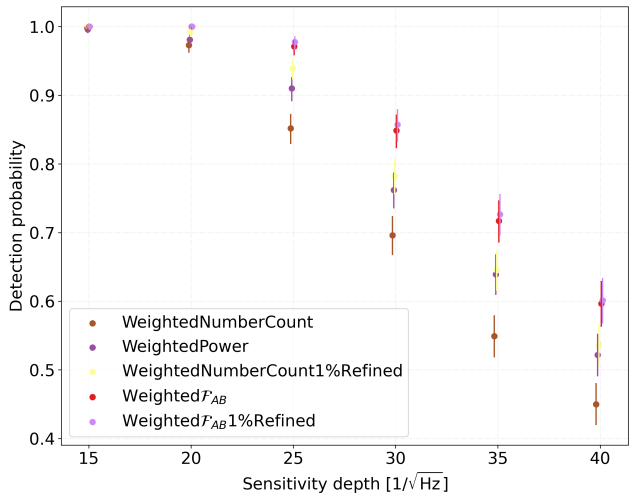

FIG. 10. Detection probability as a function of the sensitivity depth for five different detection statistics. The points have been slightly displaced in the x-axis to ease visibility. The error bars show the 95% binomial confidence interval. The plot shows the results using Gaussian noise from the H1 and L1 detectors. All detectors have the same amplitude spectral density and one year of data. We have used a coherent time of 900 s.

C. Increasing the coherent time

One of the main advantages of the new pipeline is the possibility of extending the coherent time while maintaining the same mismatch. While extending the coherent time increases the computational cost, it also increases the sensitivity of the search. For a large number of segments, the sensitivity of a StackSlide search scales as $N_{\rm seg}^{-1/4}$ [7]. If we assume that the false alarm probability stays constant, it can be seen that in order to double the sensitivity of a search the coherent time needs to be multiplied by 16.

In this section we compare the sensitivity between a

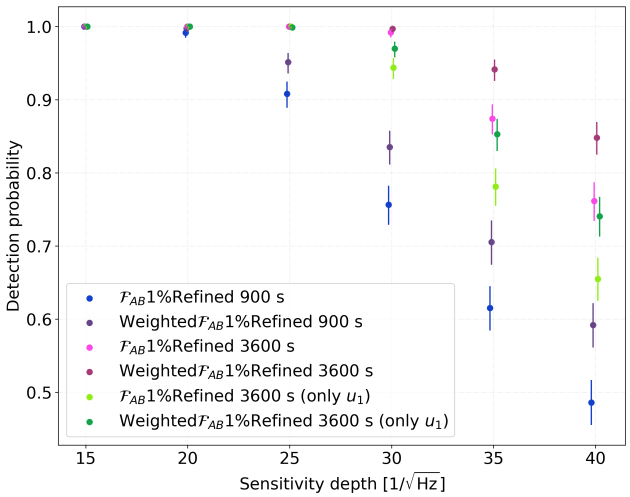

FIG. 11. Detection probability as a function of the sensitivity depth for 6 different detection statistics. The points have been slightly displaced in the x-axis to ease visibility. The error bars show the 95% binomial confidence interval. The plot shows the results using Gaussian noise from the H1 and L1 detectors. All detectors have the same amplitude spectral density and one year of data. We have used a coherent time of 900 s and 3600 s.

search with a coherent time of 900 s and a search with 3600 s, with the same maximum mismatch parameters. For this comparison we have added signals from a different region of the binary parameter space: random $P_{\rm orb}$ and $a_{\rm p}$ between $[0.785, 0.8]$ days and $[0.5, 0.6]$ l-s respectively. 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 3600 s, while all searches with 900 s do not need any u_2 template.

Fig. 11 shows the results by comparing the dominant-response detection statistic for these two coherent times, and also for a search with the long coherent time but without searching for additional u_2 templates. The improvement in sensitivity (taking into account the increased number of templates) due to using a longer coherent time is clear, both for the weighted and unweighted cases. For the unweighted detection statistics, the maximal improvement due to a 4 times increase in coherent time would be around 1.41, which is approximately what is observed in the figure. It can also be seen that as the coherent time increases, the improvement from using a weighted detection statistic decreases (for constant Gaussian noise), since the different segments contribute more equally to the total signal power as the segment duration approaches one sidereal day. The figure also shows that when the longer coherent time is used, but no u_2 templates are searched in the coherent stage, the sensitivity decreases due to the larger mismatch.

V. CONCLUSIONS

In this paper we have presented an improved pipeline to search for continuous gravitational waves from neutron stars in binary systems. The new BinarySkyHou$\mathcal{F}$ pipeline is an update of the BinarySkyHough pipeline and represents an improvement over many different aspects. The new pipeline can cover the same parameter space at a reduced computational cost. The usage of the $\mathcal{F}$-statistic or other related statistics allows the increase of the coherent time, which improves the sensitivity and gives more flexibility to the pipeline. Furthermore, the per-detector data can be coherently combined, which further reduces the computational cost (by a factor $N_{\rm det}$, assuming equal amount of data per detector), and improves the sensitivity for searches with more than two detectors.

The new pipeline is able to limit the mismatch in the coherent stage, thus allowing us to perform a low mismatch search if desired. The pipeline could be used to do a semi-coherent search on a reduced frequency or sky region with more sensitivity due to less mismatch and higher coherent time, while with the previous BinarySkyHough pipeline this was not possible. The eccentricity, argument of periapse, and first frequency derivative can now also be explicitly searched. For this reason, this pipeline can also be used to followup candidates from different types of searches, or from an all-sky isolated pipeline which gives a range of binary parameters to follow-up.

BinarySkyHou$\mathcal{F}$ still has some limitations. The Taylor coordinates used in the coherent stage only allow the search over orbital periods longer than a certain value related to the coherent time, as shown in Sec. III C 1, thus reducing the possibilities of a search over the shortest orbital periods with long coherent times. Furthermore, the RAM usage also limits the number of Taylor coordinates that can be used, thus limiting the maximum coherent time in certain regions of the parameter space.

Further research could lead to some improvements that might be applied to this pipeline. For example, the usage of the full Taylor expansion developed in App. A, or the usage of a non-zero threshold or a smaller amount of segments at the first stage. Further enhancements could be produced by reducing the mismatch in the second stage, by refining the grid and searching for additional templates. In the future it would be interesting to obtain an optimization algorithm that calculates the optimal maximum mismatch and coherent time given a certain amount of data and computational budget. Another venue for further research is to find how to optimally combine the semi-coherent results, since this is not known. We also want to explore the performance and advantages of other coherent detection statistics, such as ones that are robust to non-Gaussian noise features [31].

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 programme under the Marie Sklodowska-Curie grant agreement number 101029058.

Appendix A: 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:

$$\phi^X(t) = \phi_0 + 2\pi \sum_{k=1}^{k_\mathrm{max}} \frac{u_k^X}{k!} (t - t_\mathrm{m})^k \,, \tag{A1}$$

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

$$u_k^X \equiv \frac{1}{2\pi} \left.\frac{d^k \phi^X}{dt^k}\right|_{t_\mathrm{m}} \,, \tag{A2}$$

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):

$$\tau^X(t) = \tau(t^X_\mathrm{SSB}(t)) = t + \vec{r}^X(t)\cdot\hat{n} - R(\tau^X). \tag{A3}$$

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

$$\begin{aligned}
\dot{\tau}^X &= [1 + R']^{-1} \, (1 + \dot{\vec{r}}^X \cdot \hat{n}), \\
\ddot{\tau}^X &= [1 + R']^{-1} \left(\ddot{\vec{r}}^X \cdot \hat{n} - R'' \, \dot{\tau}^{X\,2} \right), \\
\dddot{\tau}^X &= [1 + R']^{-1} \left(\dddot{\vec{r}}^X \cdot \hat{n} - R'' \, \dot{\tau}^{X\,3} - 3R'' \, \ddot{\tau}^X \dot{\tau}^X \right), \\
&\vdots
\end{aligned} \tag{A4}$$

which generalizes Eq. (14).

We can write the first two orders explicitly as

$$\begin{aligned}
\frac{u_1^X}{f_\mathrm{m}} &= \frac{1 + \vec{v}_\mathrm{m}^X \cdot \hat{n}}{1 + R'_\mathrm{m}}, \\
u_2^X &= \frac{f_\mathrm{m}}{1 + R'_\mathrm{m}} \left(\vec{a}_\mathrm{m}^X \cdot \hat{n} - R''_\mathrm{m} \left(\frac{u_1^X}{f_\mathrm{m}} \right)^2 \right) + f'_\mathrm{m} \left(\frac{u_1^X}{f_\mathrm{m}} \right)^2,
\end{aligned} \tag{A5}$$

with the definitions of Eqs. (17) and (18) (for small eccentricity and a single spindown), and $\vec{v}_\mathrm{m}^X \equiv \dot{\vec{r}}^X(t_\mathrm{m})$, $\vec{a}_\mathrm{m}^X \equiv \ddot{\vec{r}}^X(t_\mathrm{m})$ and detector velocity and acceleration at the segment mid-time t_m, respectively. Again, these coordinates have units of Hz and Hz^2 respectively, but now also depend on the detector X and on the sky position $\hat{n}$ of the signal.

Assuming no spindown $f_1 = 0$, a circular orbit ($e = 0$), and a non-relativistic orbital velocity $a_\mathrm{p}\Omega \ll 1$, Eq. (A5) yields an approximate equation for the frequency-time pattern, namely

$$\begin{aligned}
u_1^X &\approx f_0 \, (1 + \vec{v}_\mathrm{m}^X \cdot \hat{n}) \, (1 - a_\mathrm{p}\Omega \cos\Psi_\mathrm{m}) \\
&\approx f_0 + f_0 \, \vec{v}_\mathrm{m}^X \cdot \hat{n} - f_0 \, a_\mathrm{p}\Omega \cos\Psi_\mathrm{m},
\end{aligned} \tag{A6}$$

which is the same as Eq. (15) of [14], evaluated at segment mid-time t_m.

REFERENCES

[1] 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, ApJL **929**, L19 (2022).

[2] M. Sieniawska and M. Bejger, Continuous Gravitational Waves from Neutron Stars: Current Status and Prospects, Universe **5**, 217 (2019).

[3] K. Riles, Searches for Continuous-Wave Gravitational Radiation (2022), arXiv:2206.06447 [astro-ph].

[4] 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).

[5] 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).

[6] P. R. Brady and T. Creighton, Searching for periodic sources with LIGO. II. Hierarchical searches, Phys. Rev. D. **61**, 082001 (2000).

[7] R. Prix and M. Shaltev, Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost, Phys. Rev. D. **85**, 084010 (2012).

[8] 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).

[9] C. Cutler and B. F. Schutz, Generalized $\mathcal{F}$-statistic: Multiple detectors and multiple gravitational wave pulsars, Phys. Rev. D. **72**, 063006 (2005).

[10] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, Class. Quant. Grav. **26**, 204013 (2009).

[11] P. Covas and R. Prix, Improved short-segment detection statistic for continuous gravitational waves, Phys. Rev. D. **105**, 124007 (2022).

[12] E. Goetz and K. Riles, Coherently combining data between detectors for all-sky semi-coherent continuous gravitational wave searches, Class. Quant. Grav. **33**, 085007 (2016).

[13] E. Goetz and K. Riles, An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems, Class. Quant. Grav. **28**, 215006 (2011).

[14] 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).

[15] J. Aasi *et al.* (LIGO Scientific Collaboration; Virgo Collaboration), First all-sky search for continuous gravitational waves from unknown sources in binary systems, Phys. Rev. D. **90**, 062010 (2014).

[16] 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).

[17] 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).

[18] 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).

[19] G. Mendell and M. Landry, *StackSlide and Hough Search SNR and Statistics*, Tech. Rep. LIGO-T05003 (DCC, 2005).

[20] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The Australia Telescope National Facility Pulsar Catalogue, The Astronomical Journal **129**, 1993 (2005).

[21] C. Lange, F. Camilo, N. Wex, M. Kramer, D. Backer, A. Lyne, and O. Doroshenko, Precision timing measurements of PSR J1012+5307, MNRAS **326**, 274–282 (2001).

[22] L. Nieder, B. Allen, C. J. Clark, and H. J. Pletsch, Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-ray Pulsars, ApJ **901**, 156 (2020).

[23] C. Messenger, A semi-coherent search strategy for known continuous wave sources in binary systems, Phys. Rev. D. **84**, 083003 (2011).

[24] Wikipedia contributors, Faà di Bruno's formula — Wikipedia, The Free Encyclopedia (2022), [Online; accessed 24-July-2022].

[25] J. Aasi *et al.*, Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data, Phys. Rev. D. **87**, 042001 (2013).

[26] B. Krishnan and A. M. Sintes, *Hough search with improved sensitivity*, Tech. Rep. LIGO-T070124 (DCC, 2007).

[27] S. v. d. Putten, H. J. Bulten, J. F. J. v. d. Brand, and M. Holtrop, Searching for gravitational waves from pulsars in binary systems: An all-sky search, Journal of Physics: Conference Series **228**, 012005 (2010).

[28] 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).

[29] H. J. Pletsch, Parameter-space metric of semicoherent searches for continuous gravitational waves, Phys. Rev. D. **82**, 042002 (2010).

[30] R. Prix, *Characterizing timing and memory-requirements of the $\mathcal{F}$-statistic implementations in LALSuite*, Tech. Rep. LIGO-T1600531 (DCC, 2017).

[31] 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).