PHYSICAL REVIEW D **87,** 084057 (2013)

Fully coherent follow-up of continuous gravitational-wave candidates

M. Shaltev and R. Prix

Albert-Einstein-Institut, Callinstraße 38, 30167 Hannover, Germany

(Received 11 March 2013; published 22 April 2013)

The search for continuous gravitational waves from unknown isolated sources is computationally limited due to the enormous parameter space that needs to be covered and the weakness of the expected signals. Therefore, semicoherent search strategies have been developed and applied in distributed computing environments such as Einstein@Home, in order to narrow down the parameter space and identify interesting candidates. However, in order to optimally confirm or dismiss a candidate as a possible gravitational wave signal, a fully coherent follow-up using all the available data is required. We present a general method and implementation of a *direct* (two-stage) transition to a fully coherent follow-up on semicoherent candidates. This method is based on a gridless Mesh Adaptive Direct Search algorithm using the $\mathcal{F}$-statistic. We demonstrate the detection power and computing cost of this follow-up procedure using extensive Monte Carlo simulations on (simulated) semicoherent candidates from a directed, as well as from an all-sky search setup.

DOI: 10.1103/PhysRevD.87.084057 PACS numbers: 04.30.−w

I. INTRODUCTION

Continuous gravitational waves (CWs) are expected to be emitted from rapidly spinning nonaxisymmetric compact objects, e.g., neutron stars. The computational cost of a coherent matched-filtering detection statistic, such as the $\mathcal{F}$-statistic [1], is small provided the parameters of the source (i.e., sky position α, δ, frequency f, frequency derivatives $\dot{f}$, ...) are known. However, wide parameter-space searches for unknown sources quickly become computationally prohibitive, due to the large number of points in parameter space (templates) that need to be searched [2].

In order to first reduce the parameter space to smaller, more promising regions, semicoherent search techniques have been developed [3–6] and are currently being used [7,8], for example, in the Einstein@Home distributed computing environment [9]. In a semicoherent search, the total amount of data T is divided into N shorter segments of duration ΔT. The coherent statistics from the individual segments are combined to produce a new semicoherent statistic. At a fixed computing cost these semicoherent methods are (typically) more sensitive than fully coherent searches [10].

Structuring a wide parameter-space search into hierarchical stages, which increasingly concentrate computational power onto the more promising regions of parameter space, was first described in Ref. [2] and elaborated further in Ref. [3], where a two-stage semicoherent hierarchical search was considered. An extended hierarchical scheme with an arbitrary number of semicoherent stages and a final fully coherent stage was studied numerically in Ref. [4], which concluded that three semicoherent stages will typically be a good choice. In Refs. [11,12], the use of an optimization procedure has been considered in the process of estimation of the source parameters, once a candidate is considered as a detection. In both cases, however, no practical method or implementation was provided for the systematic coherent follow-up of semicoherent candidates.

The aim of the present work is to introduce such a coherent follow-up search strategy and implementation. This is achieved by exploring the parameter space around a semicoherent candidate using a Mesh Adaptive Direct Search (MADS) algorithm. Using this method, we find that a fully coherent follow-up (using all of the available data) of initial semicoherent candidates is computationally feasible.

This paper is organized as follows: in Sec. II we describe the relevant basic concepts in CW searches, in Sec. III we propose a search strategy for the systematic follow-up of CW candidates, in Sec. IV we present a Monte Carlo study, and in Sec. V we discuss the results.

A. Notation

We distinguish a quantity Q when referring to a fully coherent stage using a tilde, $\tilde{Q}$, and when referring to a semicoherent stage using an overhat, $\hat{Q}$. Averaging over segments is denoted by an overbar, $\bar{Q}$.

II. CONTINUOUS GRAVITATIONAL WAVES

Continuous gravitational wave signals are quasi-monochromatic and sinusoidal in the source frame and undergo phase and amplitude modulation due to the diurnal and orbital motion of the detectors. The phase evolution of the signal at a detector can be approximated as [1]

$$\Phi(t) \approx \Phi_0 + 2\pi \sum_{k=0}^{s} \frac{f^{(k)}(t_0)(t-t_0)^{k+1}}{(k+1)!} + 2\pi \frac{\mathbf{r}(t)}{c}\mathbf{n} \sum_{k=0}^{s} \frac{f^{(k)}(t_0)(t-t_0)^{k}}{k!}, \quad (1)$$

where Φ_0 is the initial phase, $f^{(k)} \equiv \frac{d^k f}{dt^k}$ are the derivatives of the signal frequency f at the solar system barycenter (SSB) at reference time t_0, c is the speed of light, $\mathbf{r}(t)$ is the vector pointing from the SSB to the detector and $\mathbf{n}$ is the unit vector pointing from the SSB to the gravitational wave source.

A. Detection statistic

Following Refs. [1,13], the gravitational wave response of a detector can be expressed as a sum over four (detector-independent) amplitude parameters multiplying four (detector-dependent) basis waveforms. The amplitude parameters can be analytically maximized over, and the resulting detection statistic, known as the $\mathcal{F}$-statistic, is therefore a function only of the template "phase parameters" $\lambda \equiv \{\alpha, \delta, f, \dot{f}, \ldots\}$, where α (right ascension) and δ (declination) denote the sky position of the source.

In the presence of a signal, the fully coherent detection statistic $2\mathcal{F}$ follows a noncentral χ^2 distribution with four degrees of freedom and a noncentrality parameter given by the squared signal-to-noise ratio (SNR), ρ^2. The expectation value is therefore

$$E[2\mathcal{F}] = 4 + \rho^2, \tag{2}$$

with variance

$$\sigma^2[2\mathcal{F}] = 2(4 + 2\rho^2). \tag{3}$$

On the other hand, in the semicoherent approach, we divide the available data into N segments of duration ΔT and combine the individual coherent statistics of the segments to compute a semicoherent statistic, namely

$$\overline{2\mathcal{F}}(\lambda) = \frac{1}{N}\sum_{k=1}^{N} 2\mathcal{F}_k(\lambda), \tag{4}$$

where $2\mathcal{F}_k$ is the coherent $\mathcal{F}$-statistic in segment k. The quantity $N\overline{2\mathcal{F}}$ follows a noncentral χ^2 distribution with $4N$ degrees of freedom; thus the expectation value of $\overline{2\mathcal{F}}$ is

$$E[\overline{2\mathcal{F}}] = 4 + \overline{\rho^2}, \tag{5}$$

with variance

$$\sigma^2[\overline{2\mathcal{F}}] = \frac{2}{N}(4 + 2\overline{\rho^2}), \tag{6}$$

where $\overline{\rho^2}$ is the average SNR^2 over all segments, i.e.,

$$\overline{\rho^2} = \frac{1}{N}\sum_{k=1}^{N} \rho_k^2, \tag{7}$$

and ρ_k^2 denotes the SNR^2 in segment k.

B. Mismatch and Fisher matrix

A search for sources with unknown signal parameters implies a loss of detection power compared to the perfectly matched case. To quantify this, we use the notion of mismatch μ, as first introduced in Refs. [14,15]. This is defined as the fractional loss of expected SNR^2 at some parameter-space point λ compared to the expectation $\rho^2(\lambda_s)$ at the signal location λ_s, namely

$$\mu \equiv \frac{\rho^2(\lambda_s) - \rho^2(\lambda)}{\rho^2(\lambda_s)}, \tag{8}$$

such that $\mu \in [0, 1]$. Taylor expansion in small offsets $\Delta\lambda = \lambda - \lambda_s$ around the signal location yields

$$\mu \equiv g_{ij}(\lambda_s)\Delta\lambda^i\Delta\lambda^j + \mathcal{O}(\Delta\lambda^3), \tag{9}$$

where implicit summation over repeated parameter-space indices i, j applies, and the symmetric positive-definite matrix g_{ij} is commonly referred to as the parameter-space *metric*.

Neglecting higher-order terms, one often uses the "metric mismatch approximation," namely

$$\mu^* \equiv g_{ij}(\lambda_s)\Delta\lambda^i\Delta\lambda^j, \tag{10}$$

as a distance measure, with a range $\mu^* \in [0, \infty)$. This metric mismatch μ^* plays an important role in grid-based searches, where one typically constructs template banks in such a way that the mismatch of any putative signal and the "closest" template is bounded by a *maximal* mismatch m, i.e.,

$$\mu^* \leq m, \tag{11}$$

everywhere in the template bank.

In the presence of noise, μ as defined in Eq. (8) is not directly accessible, and we therefore introduce a related quantity, namely the fractional loss of *measured* SNR^2, namely

$$\breve{\mu} \equiv \frac{2\mathcal{F}(\lambda_s) - 2\mathcal{F}(\lambda)}{2\mathcal{F}(\lambda_s) - 4}. \tag{12}$$

Note that $\breve{\mu} \leq 1$, but contrary to Eq. (8), it can also be (slightly) negative, as we can have $2\mathcal{F}(\lambda_s) < 2\mathcal{F}(\lambda)$ due to noise.

For semicoherent searches, the metric is found [3] as the average of the fully coherent metrics over all the segments, namely

$$\hat{g}_{ij}(\lambda) = \frac{1}{N}\sum_{k=1}^{N} g_{ij,k}(\lambda), \tag{13}$$

where $\tilde{g}_{ij,k}$ is the coherent metric [Eq. (9)] in segment k.

A standard tool for parameter estimation is provided by the Fisher information matrix, which characterizes the statistical uncertainty of the maximum likelihood estimators (MLE) λ^i_{MLE} for the signal parameters λ^i_s. This can be formulated [16–18] as the well-known Cramer-Ráo lower bound on the variance of an unbiased MLE (i.e., $E[\lambda^i_{\mathrm{MLE}}] = \lambda^i_s$), namely

$$\sigma^2[\lambda^i_{\mathrm{MLE}}] \geq \{\Gamma^{-1}\}^{ii}, \tag{14}$$

where the matrix $\{\Gamma^{-1}\}^{ij}$ denotes the inverse of the Fisher matrix Γ_{ij}, which is closely related (e.g., Ref. [17]) to the metric g_{ij}, namely

$$\Gamma_{ij} = \rho^2 g_{ij}. \tag{15}$$

A semicoherent search over N segments can be considered as N different measurements; thus the semicoherent Fisher matrix yields [19]

$$\hat{\Gamma} = \sum_{k=1}^{N} \Gamma_{ij,k}. \tag{16}$$

Assuming constant SNR^2 for the different segments, we can rewrite Eq. (16) in terms of the semicoherent metric [Eq. (13)], namely

$$\hat{\Gamma} = N\overline{\rho^2}\hat{g}_{ij}, \tag{17}$$

and thus

$$\{\hat{\Gamma}^{-1}\}^{ij} = \frac{\hat{g}^{ij}}{N\overline{\rho^2}}, \tag{18}$$

where $\hat{g}^{ij}$ is the inverse matrix of $\hat{g}_{ij}$.

C. Computing cost

The computing cost C of a fully coherent (or an ideal semicoherent [10]) search is primarily due to the computation of the $\mathcal{F}$-statistic over all the templates. For a search over $\mathcal{N}$ templates using N segments of data from N_{det} detectors [10], the computing cost C is

$$C = N\mathcal{N}N_{\mathrm{det}}\, c_1, \tag{19}$$

where c_1 is the implementation-dependent computing cost for a single template, segment and detector. A method of $\mathcal{F}$-statistic computation based on short Fourier transforms (SFTs) [20] of length T_{SFT} is currently widely used in CW searches and will be considered in the present work. The cost per template in this case is proportional to the segment duration, namely

$$c_1^{\mathrm{SFT}} = c_0^{\mathrm{SFT}} \frac{\Delta T}{T_{\mathrm{SFT}}}, \tag{20}$$

where c_0^{SFT} is the implementation- and hardware-dependent fundamental computing cost per SFT. Using the total number of SFTs

$$N_{\mathrm{SFT}} = NN_{\mathrm{det}} \frac{\Delta T}{T_{\mathrm{SFT}}}, \tag{21}$$

we can write the total computing cost [Eq. (19)] of the SFT method as

$$C = \mathcal{N}N_{\mathrm{SFT}}c_0^{\mathrm{SFT}}. \tag{22}$$

In grid-based searches, the number of templates required to cover the search parameter space $\mathbb{P}$ is given by the general expression [21,22]

$$\mathcal{N} \equiv \theta_n m^{-n/2} \int_{\mathbb{P}} d^n\lambda \sqrt{\det g}, \tag{23}$$

where θ is the normalized lattice thickness, n is the number of search dimensions, m is the maximal template bank mismatch [Eq. (11)] and $\det g$ is the determinant of the parameter-space metric [Eq. (9)]. The normalized thickness is a constant depending on the grid structure, e.g., for a hypercubic lattice, $\theta_{\mathbb{Z}_n} = n^{n/2}2^{-n}$. The metric g_{ij} depends strongly on the duration ΔT and the number of segments N, in such a way that longer observation times typically require a (vastly) increased number of templates [2].

III. COHERENT FOLLOW-UP OF SEMICOHERENT CANDIDATES

A. Basic two-stage search strategy

Here we introduce a simple two-stage strategy for following up candidates from semicoherent searches. In the first stage, called *refinement*, we employ a finer search using the semicoherent statistic $\overline{2\mathcal{F}}$ to improve the initial maximum likelihood estimator. In the second stage, called *zoom*, we apply the fully coherent statistic $\widetilde{2\mathcal{F}}$ using all the data T, in order to test whether the candidate is inconsistent with Gaussian noise and if it further agrees with the signal model.

The motivation for this two-stage approach can be seen from an example 2D search grid shown in Fig. 1. The search templates are generally placed such that a putative signal λ_s will be recovered with a loss of SNR bounded by a maximal mismatch m, as given in Eq. (11), namely

$$g_{ij}\Delta\lambda^i\Delta\lambda^j \leq m, \tag{24}$$

where equality defines an (n-dimensional) isomismatch ellipse. The initial semicoherent search will yield "candidates" $\hat{\lambda}_c$ for which the statistic $\overline{2\mathcal{F}}$ exceeds a certain threshold and is higher than neighboring templates.

The initial *refinement* stage of our follow-up strategy therefore consists in finding the (nearby) parameter-space point $\hat{\lambda}_{\mathrm{MLE}}$ of the actual (local) maximum in the statistic $\overline{2\mathcal{F}}(\hat{\lambda})$ (which is a smooth function of $\hat{\lambda}$), referred to as the maximum likelihood estimator (MLE). This can be achieved simply by a denser placement of templates using the original statistic, i.e., by keeping the search setup unchanged in terms of the number and length of segments.

In the *zoom* stage, we fully coherently search the Fisher ellipse centered on the semicoherent MLE $\hat{\lambda}_{\mathrm{MLE}}$. This defines the parameter-space region that should contain the signal location λ_s with confidence corresponding to n_{B} standard deviations, i.e.,

$$\hat{\Gamma}_{ij}\delta\lambda^i\delta\lambda^j \leq n_{\mathrm{B}}^2, \tag{25}$$

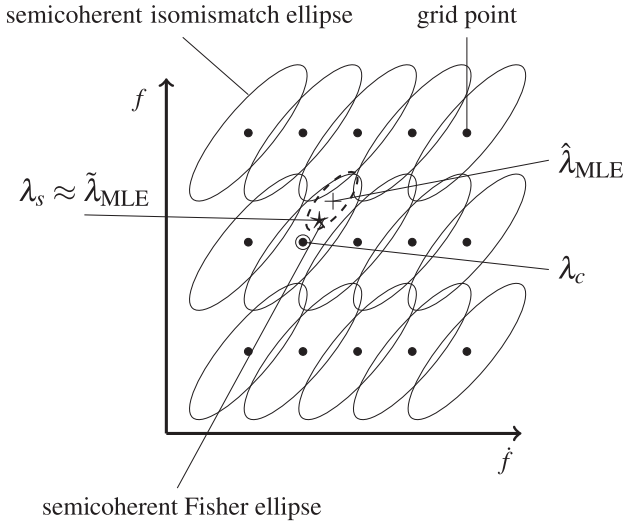

FIG. 1. Two-dimensional search grid in $\{f, \dot{f}\}$ space. The black dots are the search templates, placed such that the loss of SNR on any putative signal λ_s will be bounded by a maximal mismatch m, which defines the semicoherent isomismatch ellipses. The semicoherent Fisher ellipse centered on the MLE $\hat{\lambda}_{\mathrm{MLE}}$ is used to constrain the zoom parameter space. The aim of the zoom stage is to find $\tilde{\lambda}_{\mathrm{MLE}}$.

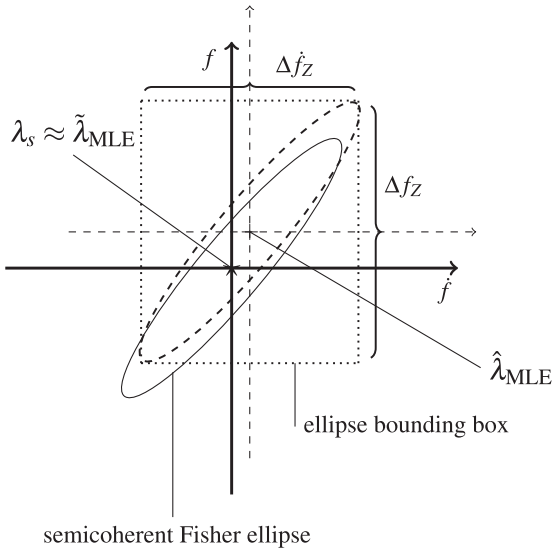

FIG. 2. Two-dimensional example: Fisher ellipse [Eq. (25)] defining the zoom search space, centered on the semicoherent MLE $\hat{\lambda}_{\mathrm{MLE}}$. The extents $\{\Delta f, \Delta \dot{f}\}$ of the bounding box are given by Eq. (27).

where $\delta\lambda^i = \hat{\lambda}^i_{\mathrm{MLE}} - \lambda^i_s$. Note that the Fisher ellipse actually describes the fluctuations of the maximum likelihood estimator $\hat{\lambda}_{\mathrm{MLE}}$ for a given signal location. However, provided the likelihood manifold is not strongly curved, this also describes our uncertainty of the signal location for a given MLE $\hat{\lambda}_{\mathrm{MLE}}$, as indicated in Fig. 2. The zoom stage will yield the fully coherent maximum likelihood estimator $\tilde{\lambda}_{\mathrm{MLE}}$, which represents our best estimate for the signal parameters λ_s. Thus, the two-stage search strategy corresponds to the transition

$$\hat{\lambda}_c \xrightarrow{\text{refinement}} \hat{\lambda}_{\mathrm{MLE}} \xrightarrow{\text{zoom}} \tilde{\lambda}_{\mathrm{MLE}} \approx \lambda_s.$$

In the following, we use a subscript R to denote quantities in the refinement stage and a subscript Z for quantities in the zoom stage.

The search volume for the refinement stage depends on the template bank construction of the original semicoherent search. Ideally, one isomismatch ellipse corresponding to the original template bank construction (see Fig. 1) should be sufficient. In practice, however, it might often be necessary to use several grid spacings in each direction, if the template bank was not originally constructed in a strictly metric way. In this case, the exact number of grid spacings will have to be empirically determined in a Monte Carlo study.

1. Bounding box and volume of n-dimensional ellipses

In the following discussion, it will be useful to express the bounding box and volume of an n-dimensional ellipse, namely for the isomismatch ellipse of Eq. (24) and the Fisher ellipse of Eq. (25). The general form of the n-dimensional ellipse equation is

$$G_{ij} d\lambda^i d\lambda^j = R^2, \tag{26}$$

where G_{ij} is a positive-definite symmetric matrix. The extents $\Delta\lambda^i$ of a bounding box along coordinate axes λ^i (as indicated in Fig. 2) can be obtained from the diagonal elements of the inverse matrix, $\{G^{-1}\}^{ij}$, namely

$$\Delta\lambda^i = 2R\sqrt{\{G^{-1}\}^{ii}}. \tag{27}$$

The ellipse coordinate volume is expressible via the matrix determinant, $\det G$, namely

$$V = \frac{R^n}{\sqrt{\det G}} \mathcal{V}_n, \tag{28}$$

where $\mathcal{V}_n = \frac{\pi^{n/2}}{\Gamma(1+n/2)}$ is the volume of the unit n-ball.

B. Classification of zoom outcomes

Assuming a real CW signal, we can estimate the range of expected values of the fully coherent zoom $\mathcal{F}$-statistic in $\tilde{\lambda}_s$. From Eq. (5), we can obtain a (rough) estimate of the average SNR^2 from the *measured* average SNR^2 of the semicoherent maximum likelihood estimator, namely

$$\overline{\rho^2}_{\mathrm{MLE}} \approx \overline{2\mathcal{F}}_{\mathrm{MLE}} - 4. \tag{29}$$

The SNR^2 of the fully coherent search is linear in the number of segments N, i.e.,

$$\tilde{\rho}^2 = N\overline{\rho^2}_{\mathrm{MLE}}. \tag{30}$$

Substitution of the above expression into Eq. (2) yields the expectation for the fully coherent matched filter in $\tilde{\lambda}_{\mathrm{MLE}}$, namely

$$\widetilde{2\mathcal{F}}_o \equiv E[\widetilde{2\mathcal{F}}] \approx 4 + N\overline{\rho^2}_{\mathrm{MLE}}. \tag{31}$$

Further substitution of Eq. (30) into Eq. (3) yields the corresponding variance as

$$\sigma_o^2 \equiv \sigma^2[\widetilde{2\mathcal{F}}] \approx 2(4 + 2N\overline{\rho^2}_{\mathrm{MLE}}). \tag{32}$$

These quantities are useful for defining what we mean by confirming a CW signal.

Note that the uncertainty in the original SNR estimation in Eq. (29) results in a distribution around the final estimate of Eq. (31) that is wider than estimated by Eq. (32). This effect can be computed analytically and empirically, and is found to amount to about a factor of 2.

Depending on the maximal $\widetilde{2\mathcal{F}}$ value found in the final zoom stage, we can distinguish three possible outcomes:

(1) *Consistency with Gaussian noise* (G).—The fully coherent $\widetilde{2\mathcal{F}}$ value does not exceed a threshold

$$\widetilde{2\mathcal{F}} < \widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}}, \tag{33}$$

where $\widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}}$ is chosen to correspond to some (small) false-alarm probability p_{fA} in Gaussian noise.

For example, a threshold $\widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}} = 60$ corresponds to a very small false-alarm probability of order 10^{-12} in a single template, as given by Eq. (43).

(2) *Non-Gaussian origin* ($\neg G$).—The candidate is loud enough to be inconsistent with Gaussian noise at the false-alarm probability p_{fA}, i.e.,

$$\widetilde{2\mathcal{F}} \geq \widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}}. \tag{34}$$

(3) *Signal recovery* (S).—A *subclass* of $\neg G$; the final zoomed candidate $\widetilde{2\mathcal{F}}$ exceeds the Gaussian noise threshold $\widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}}$ *and* falls into the predicted signal interval given by Eqs. (31) and (32) (at some confidence level). We can write this as

$$\widetilde{2\mathcal{F}}^{(S)}_{\mathrm{th}} < \widetilde{2\mathcal{F}} < \widetilde{2\mathcal{F}}^{(S)}_{\mathrm{max}}, \tag{35}$$

where $\widetilde{2\mathcal{F}}^{(S)}_{\mathrm{th}} \equiv \max\{\widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}}, \widetilde{2\mathcal{F}}_o - n_u\sigma_o\}$, and $\widetilde{2\mathcal{F}}^{(S)}_{\mathrm{max}} \equiv \widetilde{2\mathcal{F}}_o + n_u\sigma_o$, where n_u determines the desired confidence level. In this work we consider $n_u = 6$, which corresponds roughly to a confidence of ~99.6%.

Note that there can be cases where a zoomed candidate ends up in $\neg G$ but does not make it into the signal recovery (S) band, e.g., typically $\widetilde{2\mathcal{F}}^{(G)}_{\mathrm{th}} < \widetilde{2\mathcal{F}} < \widetilde{2\mathcal{F}}^{(S)}_{\mathrm{th}}$. There can be different reasons for this, e.g., the search algorithm converged to a secondary maximum in the refinement or zoom stage, the signal model deviates from reality and requires modification, or the "signal" found is of non-astrophysical origin (e.g., a detector noise artifact). Generally, further investigation will be required for all candidates falling into the non-Gaussian category ($\neg G$).

C. Grid-based computing cost of the zoom stage

We do not consider a grid-based follow-up method in this paper, but it is instructive to estimate the corresponding computing cost for later comparison. To estimate the number of templates required for the fully coherent search, we can use Eq. (28) to compute the volume of the follow-up Fisher ellipse, Eq. (25), and divide it by the volume covered by one coherent template, Eq. (24). Namely, the Fisher ellipse volume is given by

$$\hat{V} = \frac{n_{\mathrm{B}}^n}{(N\overline{\rho^2})^{n/2}\sqrt{\det \hat{g}}}\,\mathcal{V}_n, \tag{36}$$

while the coherent template volume at mismatch m is

$$\tilde{V} = \frac{m^{n/2}}{\sqrt{\det \tilde{g}}}\,\mathcal{V}_n. \tag{37}$$

Therefore, we can estimate the number of templates as

$$\mathcal{N} \approx \frac{\hat{V}}{\tilde{V}} = \frac{n_{\mathrm{B}}^n}{(N\overline{\rho^2})^{n/2}m^{n/2}}\frac{\sqrt{\det \tilde{g}}}{\sqrt{\det \hat{g}}}. \tag{38}$$

Consider a follow-up of a candidate from a directed $n = 2$ search in $\{f, \dot{f}\}$ (e.g., see Fig. 1). Assuming a semicoherent search using $N = 200$ segments of $\Delta T = 1$ d duration without gaps, and a fully coherent observation time of $T = 200$ d. Using the expressions found in Ref. [23], the determinants of the two-dimensional coherent and the semicoherent metrics are found as

$$\sqrt{\det \tilde{g}} = \pi^2 T^3 \frac{1}{540}, \tag{39}$$

$$\sqrt{\det \hat{g}} = \pi^2 \Delta T^3 \frac{\gamma(N)}{540}, \tag{40}$$

where $\gamma \approx \sqrt{5}N$ is the spin-down refinement factor. Putting everything together in Eq. (38), we obtain

$$\mathcal{N} \approx \frac{n_{\mathrm{B}}^2 N}{\sqrt{5}\overline{\rho^2} m}, \tag{41}$$

where we used $N = T/\Delta T$. For a signal with $\overline{\rho^2} = 1$, $n_{\mathrm{B}} = 24$,[1] and $m = 0.1$, the number of templates is therefore $\mathcal{N} \approx 5.1 \times 10^5$. Thus, using Eq. (22) for two detectors and the SFT method with $T_{\mathrm{SFT}} = 1800$ s, the zoom

[1]This large n_{B} value is found to contain the signal location in more than 98% of the cases even for weak signals, where the Fisher matrix may be a poor predictor; see Ref. [16].

computing cost is $C \approx 11$ min per candidate, where we used the fundamental computing cost $c_0^{\mathrm{SFT}} = 7 \times 10^{-8}$ s [10].

In the more general case where the sky position of the source is also unknown, the number of sky points typically scales at least quadratically with the observation time [23,24] (for coherent integration longer than a few days), thus generally resulting in completely prohibitive computational requirements for grid-based follow-up searches. In particular, extending the directed search example from the previous paragraph to an all-sky follow-up would require $\mathcal{N}_{\mathrm{sky}} \approx 1.3 \times 10^6$ sky points,[2] or a total of $\mathcal{N} \approx 6.8 \times 10^{11}$ templates.

For comparison, using the gridless search algorithm discussed in the next sections, it is possible to coherently follow up 2D directed candidates in less than 2 minutes (see Fig. 4(d)), and all-sky candidates in about 1 hour per candidate (see Fig. 5(d)).

D. Mesh Adaptive Direct Search (MADS)

A significant difference between the hierarchical search strategies discussed in Refs. [2–4] and in this work is the method of template bank construction at the different stages. Namely, we consider a gridless method for exploring the parameter space.

The MADS class of algorithms for derivative-free optimization was first introduced in Ref. [26] and further developed in Refs. [27,28], among others. In this subsection, we only introduce some of the control parameters of the algorithm required in the construction of MADS-based $\mathcal{F}$-statistic searches; for an in-depth treatment and proofs, we refer the reader to the cited publications.

MADS consists of the iteration of two steps, called *search* and *poll*, in which trial points are constructed and evaluated in order to find an extremum. In the search step, any strategy can be applied to construct trial points. In this work we use quadratic models (quadratic form) to approximate the objective function from a sample of objective values [28]. If the local exploration in the search step fails to generate a new solution, a set of poll points is generated using a stochastic or deterministic method. Stochastic means that the poll points are generated randomly [26], whereas deterministic refers to the usage of pseudorandom Halton sequences [27]. However, both methods generate points which form a dense set in the unit sphere after an infinite number of iterations. For a given starting point λ_c with parameter-space boundaries $\Delta\lambda_{\mathrm{B}}$, initial step sizes $d\lambda$ and a method for generation of poll points, the discretization of the parameter space Δ_k^m at iteration k is governed by a fixed rational number $u_b > 1$ and the coarsening $w^+ \geq 0$ and refining $w^- \leq -1$ exponents. If the current iteration generates a better solution, the discretization in the next iteration is coarser, namely $\Delta_{k+1}^m = u_b^{w^+} \Delta_k^m$; otherwise $\Delta_{k+1}^m = u_b^{w^-} \Delta_k^m$ [26]. The algorithm stops if an improved solution cannot be found or the total number of evaluated parameter-space points p reaches some given maximum $p_{\max}$.

[2]The number of sky templates has been estimated by numerical computation of the sky part of the metrics $\hat{g}$ and $\tilde{g}$ using FSTATMETRIC_V2 from LALSUITE [25]; see also Ref. [17].

E. MADS-based follow-up algorithm

From the point of view of the MADS algorithm, the function to optimize is a black box requiring some input to produce a single output value. The black box in our case is either the computation of the semicoherent $\mathcal{F}$-statistic $\overline{2\mathcal{F}}$ of Eq. (4) in the refinement, or the fully coherent $\mathcal{F}$-statistic $\widetilde{2\mathcal{F}}$ in the zoom stage. In order to minimize the possibility of convergence to secondary maxima, we run multiple instances of the MADS search in each stage, varying the mesh-coarsening exponent w^+. The minimal $w^+_{\min}$ and maximal $w^+_{\max}$ coarsening exponents determine the number of MADS steps in each pass, namely $n_{\mathrm{steps}} = w^+_{\max} - w^+_{\min} + 1$. Thus, we consider our search algorithm to be composed of several instances of MADS; see Fig. 3. The inputs of the search algorithm are the candidate λ_c to follow up, the search boundaries $\Delta\lambda_{R/Z}$ around the candidate, and a set of MADS input parameters, namely $\{d\lambda, u_b, w^+_{\min}, w^+_{\max}, w^-\}$. In the zoom stage, the search boundaries ($\Delta\lambda_Z$) are estimated from the bounding box of the Fisher ellipse, using Eq. (27). For the refinement stage, the search boundaries ($\Delta\lambda_R$) generally have to be determined depending on the template bank setup of the original semicoherent search. Note, however, that the bounding boxes $\Delta\lambda$ only serve as a necessary input parameter to the MADS search algorithm, while the effective search region can be further reduced by rejecting points that do not satisfy a given constraint. For example, the effective search region in the zoom stage always consists of the Fisher ellipse [Eq. (25)].

The initial step sizes $d\lambda^i$ are also empirically determined, typically as some fraction of the search boundary $\Delta\lambda^i_{R/Z}$.

We propose a four-pass algorithm with an equal (for simplicity) number of steps n_{steps} in each pass; however, with different starting point and method of trial point generation:

(1) *1st pass.*—Starting point λ_c, *deterministic* point generation.
(2) *2nd pass.*—Starting point λ_c, *stochastic* point generation.
(3) *3rd pass.*—Starting point from the loudest template from the first two passes, *deterministic* point generation.
(4) *4th pass.*—Starting point from the vicinity of the loudest point from the first two passes, *stochastic* point generation.

In the zoom stage, we terminate the search as soon as the loudest point of the current iteration satisfies the signal

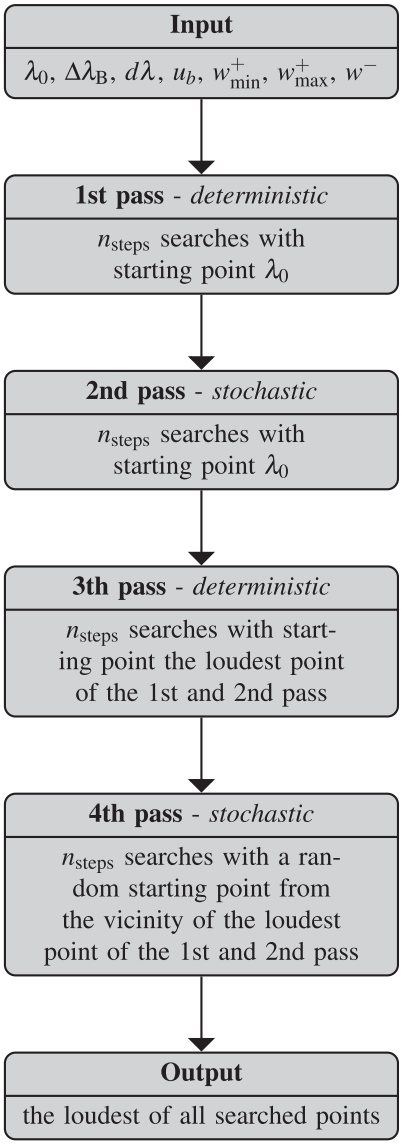

FIG. 3. MADS-based search algorithm with four passes, where $\lambda_0 = \lambda_c$ in the refinement stage and $\lambda_0 = \hat{\lambda}_{MLE}$ in the zoom stage.

confirmation condition (S) of Eq. (35). In lower-dimensional cases, such as the directed search considered later, a single pass is therefore often found to be sufficient. For later usage, we introduce the total number of MADS iterations n_I as the sum of the number of steps in each pass.

F. MADS follow-up computing cost

Contrary to grid-based searches, the computing cost of the MADS based algorithm is nondeterministic, due to the *a priori* unknown number of explored parameter-space points. To estimate the maximal computing cost of the refinement or the zoom stage using Eq. (19), we need the maximal number of possibly evaluated templates

$$\mathcal{N}_{max} = \sum_{i=0}^{n_I} p^i_{max}, \tag{42}$$

where p^i_{max} is the user-specified maximum of the number of computed templates at MADS iteration i. This maximal number is typically chosen to be large to avoid a too-early interruption of the MADS instance, e.g., when further improvement of the current solution is possible while the extremum is not yet found. However, if the extremum is found, a MADS iteration starting from this point terminates rapidly.

Note that the fundamental computing cost c_0^{SFT} in stochastic searches over the sky is typically larger than in a grid-based search, where a lot of templates with different spin-down components can be computed at fixed sky position. This results in a larger value of about $c_0^{SFT} \approx 3 \times 10^{-7}$ s instead of the number quoted in Sec. III C.

G. False-alarm and detection probability

After the final fully coherent zoom stage, we are left with a candidate falling into one of the three categories discussed earlier: namely, the candidate is consistent with the signal model (S), with Gaussian noise (G), or is of non-Gaussian origin but inconsistent with the signal model. An additional valuable piece of information is the false-alarm probability associated with the candidate. This is the probability of exceeding a threshold $2\mathcal{F}$ value in the absence of a signal, where the relevant distribution is the central χ^2 distribution with four degrees of freedom, denoted as $\chi^2_4(2\mathcal{F})$. The single-template false-alarm probability is

$$p^1_{fA} = \int_{2\mathcal{F}_{th}}^{\infty} d(2\mathcal{F})\chi^2_4(2\mathcal{F}) = (1 + \mathcal{F}_{th})e^{-\mathcal{F}_{th}}, \tag{43}$$

and for $\mathcal{N}$ independent templates, this results in

$$p_{fA} = 1 - (1 - p^1_{fA})^{\mathcal{N}}, \tag{44}$$

where for $\mathcal{N}p^1_{fA} \ll 1$, Taylor expansion yields $p_{fA} \approx \mathcal{N}p^1_{fA}$. For example, a threshold of $\widetilde{2\mathcal{F}}^{(G)}_{th} = 70$ for a search with $\mathcal{N} = 1 \times 10^5$ templates corresponds to a false-alarm probability of $p_{fA} \lesssim 2 \times 10^{-9}$, where the upper bound corresponds to $\mathcal{N}$ completely independent templates.

The overall detection probability of the follow-up method depends on the signal SNR. Higher SNR in the refinement stage yields better localization of the signal—i.e., a smaller Fisher ellipse—and thus also a higher probability of signal recovery [Eq. (35)]. In addition, the MADS algorithm parameters also affect the detection efficiency; e.g., an increased number of MADS iterations increases the detection probability, especially for signals with lower SNR. Because of this, the detection probability will have to be estimated empirically in a Monte Carlo study (see Figs. 4(c) and 5(c)).

IV. MONTE CARLO STUDIES

To demonstrate the capability of the systematic follow-up procedure proposed in Sec. III, we perform two different types of Monte Carlo (MC) studies.

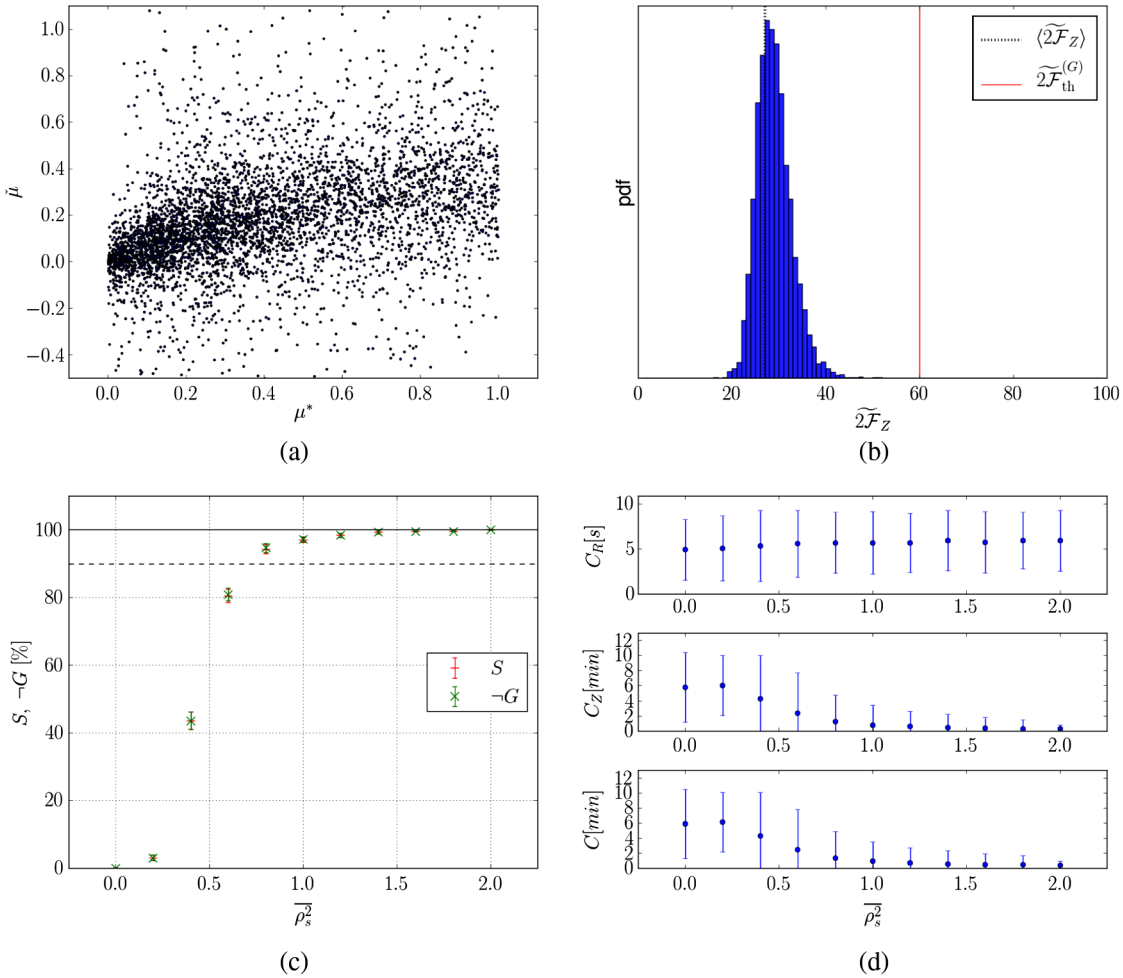

FIG. 4 (color online). Monte Carlo study of two-stage follow-up of candidates from a directed $\{f, \dot{f}\}$ semicoherent search pointed toward the Galactic center with $N = 200$ segments of duration $\Delta T = 1$ d. (a) SNR loss of the initial candidates $\check{\mu}$ versus semicoherent metric mismatch $\mu*$ to the closest template. (b) $2\widetilde{\mathcal{F}}_Z$ distribution after the fully coherent 2D $\{f, \dot{f}\}$ zoom stage of 5000 directed searches in pure Gaussian noise without injected signal. The maximal $2\mathcal{F}$ value found is $2\widetilde{\mathcal{F}}_Z^{\max} = 51.61$. The mean value $\langle 2\widetilde{\mathcal{F}}_Z\rangle = 29.00$ is plotted with a dotted line. (c) Percentage of the 5000 injected signals classified as recovered ($-S$) and of non-Gaussian origin ($\times\neg G$) as function of the noncentrality parameter $\overline{\rho^2}_s$, Eq. (7). The error bars are computed by using a Jackknife estimator. (d) *Upper plot*: computing cost of the semicoherent refinement stage. *Middle plot*: computing cost of the fully coherent zoom stage. *Lower plot*: total computing cost.

In the first case, we simulate a so-called *directed* search for a fixed sky position, where we follow up candidates in a two-dimensional spin-down space, i.e., $\{f, \dot{f}\}$. In the second case, we simulate an all-sky search over the four-dimensional parameter space $\{\alpha, \delta, f, \dot{f}\}$.

All MADS searches are implemented using the MADS reference library NOMAD [29], and the LAL library from the LALSUITE [25] is used for the $\mathcal{F}$-statistic computation [30]. The Gaussian data and signal injections are produced using the LALAPPS programs from LALSUITE. In particular, with LALAPPS_MAKEFAKEDATA_V4, we create data sets of total duration $T = 200$ d, with $N = 200$ segments of duration $\Delta T = 1$ d, using SFTs of length $T_{\rm SFT} = 1800$ s, for the two LIGO detectors H1 and L1. The noise level per detector is generated as Gaussian white noise with a power-spectral density S_n of $\sqrt{S_n} = 2 \times 10^{-23}$ Hz$^{-1/2}$.

Independently of the type of search, the initial candidates to follow up are prepared as follows: Rather than performing a semicoherent grid-based search using the Hough [5] or GCT method [6], we generate candidates by drawing a random point in the vicinity of the injection and consider it a candidate if the semicoherent metric mismatch μ^* is within the range

$$\mu^* \in [0, 1]; \tag{45}$$

see Figs. 4(a) and 5(a). This procedure for candidate preparation allows us to separate the study of the follow-up algorithm from the problem of how to set up a semicoherent search, which is a difficult question on its own.

Note that even if the original grid-based semicoherent search does not produce candidates that conform with Eq. (45), we can always increase the density of the grid until

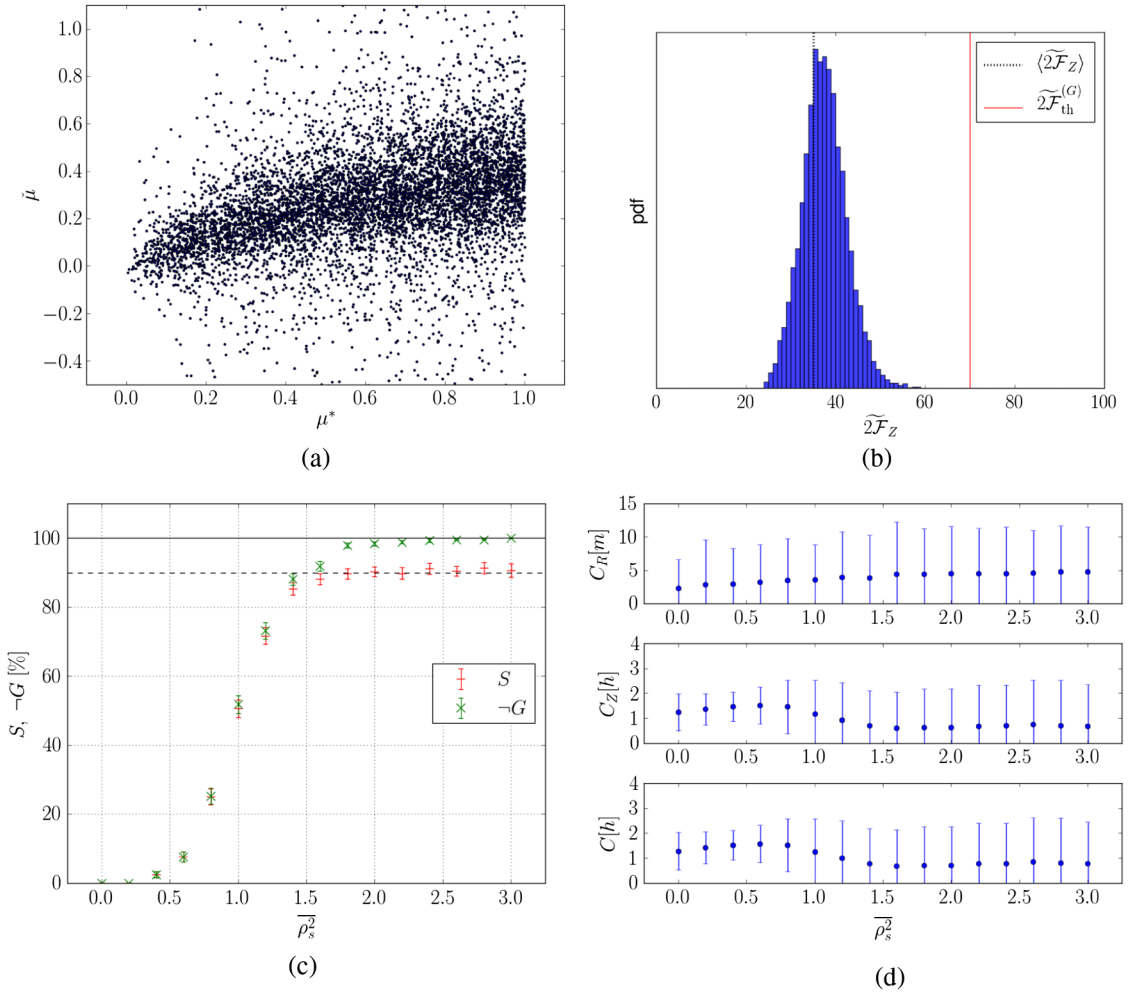

FIG. 5 (color online). Monte Carlo study of two-stage follow-up of candidates from an all-sky $\{\alpha, \delta, f, \dot{f}\}$ semicoherent search with $N = 200$ segments of duration $\Delta T = 1$ d. (a) SNR loss of the initial candidates $\check{\mu}$ versus semicoherent metric mismatch $\mu*$. (b) $\widetilde{2\mathcal{F}}_Z$ distribution after the fully coherent 4D $\{\alpha, \delta, f, \dot{f}\}$ zoom stage of 7500 searches in pure Gaussian noise, without injected signal. The maximal $2\mathcal{F}$ value found is $\widetilde{2\mathcal{F}}_Z^{\max} = 58.76$. The mean value $\langle\widetilde{2\mathcal{F}}_Z\rangle = 37.50$ indicated with dots. (c) Percentage of the 7500 injected signals classified as recovered ($-S$) and of non-Gaussian origin ($\times\neg G$) as function of the signal strength $\overline{\rho^2}_s$. The error bars are computed using a Jackknife estimator. (d) *Upper plot*: computing cost of the semicoherent refinement stage. *Middle plot*: computing cost of the fully coherent zoom stage. *Lower plot*: total computing cost.

Eq. (45) applies. This would amount to a (cheap) preprocessing stage inserted before the present follow-up procedure.

A. Follow-up of candidates from a directed search

For the directed type of searches, we fix the sky position to the coordinates of the Galactic center. This choice is arbitrary, and we could use any other point without qualitatively changing the results. We create 5000 data sets. Note that each data set has different Gaussian noise realization in which a CW signal from an isolated source is injected. In the process of injection, the original noise data set is also used to examine the behavior of the follow-up method in the absence of a signal.

The pulsar injection parameters λ_s are drawn uniformly in the range $f \in (50, 51)$ Hz, $\cos\iota \in (-1, 1)$, $\psi \in (-\pi/4, \pi/4)$ and $\phi_0 \in (0, 2\pi)$, where ι is the inclination angle of the source with respect to the line of sight, ψ is the polarization, and ϕ_0 is the initial phase of the signal [30]. The signal amplitude h_0 is chosen such that the expected average SNR^2 of Eq. (7) for a perfect match is distributed uniformly in the range $\overline{\rho^2}_s \in (0, 2)$. The spin-down $\dot{f}$ is chosen uniformly in the range $\dot{f} \in (-\frac{f_{\min}}{\tau_{\min}}, \frac{f_{\min}}{\tau_{\min}})$ with a minimal spin-down age $\tau_{\min} = 300$ yr at $f_{\min} = 50$ Hz. The MADS algorithm parameters used in the MC which have been found empirically to achieve good results are summarized in Table I. For this type of follow-up, we find that the first pass of the search algorithm in the refinement stage and only two repetitions of the second pass in the zoom stage are sufficient. We restrict the size of the search box for the refinement stage $\Delta\lambda_R$ by taking one frequency

TABLE I. Algorithm parameters for follow-up of candidates from directed searches.

Stage	w^-	w^+_{min}	w^+_{max}	u_b	p_{max}
R	-1	1	1	2	20 000
Z	-1	1	50	1.1	20 000

and two first spin-down metric extents. In the zoom stage, we constrain the parameter space to a Fisher ellipse [Eq. (25)] with $n_{\mathrm{B}} = 24$.

We first apply the follow-up chain to the pure Gaussian noise data without injected signals. The corresponding $\widetilde{2\mathcal{F}}_Z$ distribution of the resulting fully coherent zoom stage is plotted in Fig. 4(b). The maximal value found is $\widetilde{2\mathcal{F}}_Z^{\mathrm{max}} = 51.61$. We therefore use a threshold for the classification of non-Gaussian candidates ($\neg G$) of $\widetilde{2\mathcal{F}}_{\mathrm{th}}^{(G)} = 60$, which is safely above this level.

We next apply the follow-up chain to the Gaussian noise data with injected signal. In Fig. 4(c), we plot the percentage of injected signals that are classified as recovered signals (S) and non-Gaussian origin ($\neg G$) as a function of the injected signal strength $\overline{\rho^2}_s$. From this plot, we can read out the detection probability—namely, we reach 90% of signal recovery for candidates with $\overline{\rho_s^2} \approx 0.7$.

The computing cost as a function of $\overline{\rho_s^2}$ is plotted in Fig. 4(d). We notice that the cost of the refinement stage is negligible, and in the zoom stage, the averaged computing time decreases with higher signal strength.

B. Follow-up of candidates from an all-sky search

The data and signal preparation for the following all-sky Monte Carlo study is the same as in the directed search case; however, the sky position is drawn isotropically over the whole sky. We create 7500 data sets with uniformly distributed injected average SNR^2 in the range $\overline{\rho_s^2} \in (0, 3)$. The algorithm parameters used in the refinement and zoom stages which have been found empirically to yield good performance are given in Table II. We also find that here the zoom stage benefits from performing all four search passes shown in Fig. 3. The size of the search box for the refinement stage in the spin-down subspace has been defined exactly as in the directed search example. The sky subspace is constrained by using an $m = 1$ isomismatch ellipse. As in the previous example, we use $n_{\mathrm{B}} = 24$ in Eq. (25) to determine the size of the Fisher ellipse.

TABLE II. Follow-up algorithm parameters for full parameter-space searches.

Stage	w^-	w^+_{min}	w^+_{max}	u_b	p
R	-1	1	5	2	20 000
Z	-1	1	50	1.2	20 000

Similarly to the directed follow-up, we first test the pipeline using the Gaussian noise data without injections. The resulting distribution of final $\widetilde{2\mathcal{F}}_Z$ values is plotted in Fig. 5(b). The maximal value found is $\widetilde{2\mathcal{F}}_Z^{\mathrm{max}} = 58.76$, which is higher compared to the value found in the directed follow-up searches due to the increased number of evaluated templates. We therefore use a threshold for the classification of non-Gaussian candidates ($\neg G$) of $\widetilde{2\mathcal{F}}_{\mathrm{th}}^{(G)} = 70$, which is safely above this level.

Next, we search the data containing the injected signals. In Fig. 5(c), we plot the fraction of signals classified as recovered (S) and the percentage of MC trials found to be of non-Gaussian origin ($\neg G$) as a function of the injected signal strength $\overline{\rho^2}_s$. In order to achieve 90% signal recovery (S), we now need stronger signals, namely $\overline{\rho_s^2} \gtrsim 1.7$. However, for $\overline{\rho_s^2} \approx 1.5$, we can already achieve 90% "detection probability" in the sense of separating candidates from Gaussian noise ($\neg G$). This indicates that the zoom step sometimes converges on a secondary maximum. Given that any non-Gaussian ($\neg G$) candidates after zoom will receive further scrutiny, it would be straightforward to further explore the parameter space around such candidates to localize a potential primary maximum.

The computing cost as a function of $\overline{\rho_s^2}$ is plotted in Fig. 5(d). We notice that the total computing cost is dominated by the zoom stage, and the averaged computing time is rather independent of the signal strength.

V. DISCUSSION

We have studied a two-stage scheme for the fully coherent follow-up of semicoherent candidates. The first stage, called refinement, aims to find the maximum likelihood estimator of the initial semicoherent candidate. This allows us to better constrain the parameter space for the coherent zoom stage. The two-stage scheme is suitable for following-up candidates from all-sky or directed semicoherent searches. The proposed gridless optimization lowers the computing cost per candidate to acceptable levels. In Monte Carlo studies, we tested the efficiency of the algorithm for directed and all-sky follow-up searches.

In this paper, we restricted the all-sky follow-up optimization to four dimensions, namely sky, frequency, and first spin-down. Further work is required to extend the optimization to higher dimensions. A related attractive direction for further development is the extension and application of the search algorithm for follow-up of CW candidates in binary systems, which is a challenging higher-dimensional problem.

We also aim to extend the two-stage scheme presented here by including intermediate semicoherent zoom stages. This should allow us to further reduce the computing cost and increase detection efficiency.

ACKNOWLEDGMENTS

We are grateful to numerous colleagues for useful comments and discussions, in particular Karl Wette, Holger Pletsch, Paola Leaci, Vladimir Dergachev, David Keitel, Thomas Dent, Badri Krishnan, Maria Alessandra Papa and Bruce Allen. M. S. gratefully acknowledges the support of Bruce Allen and the IMPRS on Gravitational Wave Astronomy of the Max Planck Society. This paper has been assigned LIGO document No. LIGO-P1200185-v3.

[1] P. Jaranowski, A. Krolak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[2] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
[3] P. R. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
[4] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).
[5] B. Krishnan, A. Sintes, M. Papa, B. Schutz, S. Frasca, and C. Palomba, Phys. Rev. D **70**, 082001 (2004).
[6] H. J. Pletsch and B. Allen, Phys. Rev. Lett. **103**, 181102 (2009).
[7] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008).
[8] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **87**, 042001 (2013).
[9] Einstein@home, http://www.einsteinathome.org.
[10] R. Prix and M. Shaltev, Phys. Rev. D **85**, 084010 (2012).
[11] P. Jaranowski and A. Krolak, Phys. Rev. D **61**, 062001 (2000).
[12] A. Krolak, M. Tinto, and M. Vallisneri, Phys. Rev. D **70**, 022003 (2004).
[13] C. Cutler and B. F. Schutz, Phys. Rev. D **72**, 063006 (2005).
[14] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[15] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996).
[16] M. Vallisneri, Phys. Rev. D **77**, 042001 (2008).
[17] R. Prix, Phys. Rev. D **75**, 023004 (2007).
[18] S. M. Kay, *Estimation Theory*, Fundamentals of Statistical Signal Processing (Prentice-Hall, Englewood Cliffs, NJ, 1993), 1st ed., Vol. 1.
[19] D. Coe, arXiv:0906.4123.
[20] P. R. Williams and B. F. Schutz, arXiv:gr-qc/9912029.
[21] R. Prix, Classical Quantum Gravity **24**, S481 (2007).
[22] C. Messenger, R. Prix, and M. A. Papa, Phys. Rev. D **79**, 104017 (2009).
[23] H. J. Pletsch, Phys. Rev. D **82**, 042002 (2010).
[24] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, D. P. Anderson, S. B. Anderson, W. G. Anderson *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009).
[25] LALSuite, https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.
[26] C. Audet and J. E. Dennis, Jr., SIAM J. Optim. **17**, 188 (2006).
[27] M. A. Abramson, C. Audet, J. E. Dennis, and S. L. Digabel, SIAM J. Optim. **20**, 948 (2009).
[28] A. R. Conn and S. L. Digabel, Optim. Method. Softw. **28**, 139 (2013).
[29] S. Le Digabel, ACM Trans. Math. Softw. **37**, 44 (2011).
[30] R. Prix, LIGO Report No. LIGO-T0900149-v3, 2010.