Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. III. Rotational phase tracking

A. Melatos,^{1,2,*} P. Clearwater,^{1,2,3} S. Suvorova,^{1,2,4,5} L. Sun^(b),^{1,2,6,7} W. Moran,^{4,5} and R. J. Evans^{2,4}

¹School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia

²Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav),

University of Melbourne, Parkville, Victoria 3010, Australia

³Data61, Commonwealth Scientific and Industrial Research Organisation,

Corner Vimiera & Pembroke Roads, Marsfield, NSW 2122, Australia

⁴Department of Electrical and Electronic Engineering, University of Melbourne,

Parkville, Victoria 3010, Australia

⁵School of Electrical and Computer Engineering, RMIT University, Melbourne, Victoria 3000, Australia

⁶LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA

⁷OzGrav-ANU, Centre for Gravitational Astrophysics, College of Science, Australian National University, Australian Capital Territory 2601, Australia

(Received 15 June 2020; accepted 22 July 2021; published 18 August 2021)

A hidden Markov model (HMM) solved recursively by the Viterbi algorithm can be configured to search for persistent, quasimonochromatic gravitational radiation from an isolated or accreting neutron star, whose rotational frequency is unknown and wanders stochastically. Here an existing HMM analysis pipeline is generalized to track rotational phase and frequency simultaneously, by modeling the intrastep rotational evolution according to a phase-wrapped Ornstein-Uhlenbeck process, and by calculating the emission probability using a phase-sensitive version of the Bayesian matched filter known as the \mathcal{B} -statistic, which is more sensitive than its predecessors. The generalized algorithm tracks signals from isolated and binary sources with characteristic wave strain $h_0 \ge 1.3 \times 10^{-26}$ in Gaussian noise with amplitude spectral density 4×10^{-24} Hz^{-1/2}, for a simulated observation composed of $N_T = 37$ data segments, each $T_{drift} = 10$ days long, the typical duration of a search for the low-mass x-ray binary (LMXB) Sco X - 1 with the Laser Interferometer Gravitational Wave Observatory (LIGO). It is equally sensitive to isolated and binary sources and ≈ 1.5 times more sensitive than the previous pipeline, which achieves $h_0 \ge 2.0 \times 10^{-26}$ for a comparable search. Receiver operating characteristic curves (to demonstrate a recipe for setting detection thresholds) and errors in the recovered parameters are presented for a range of practical h_0 and N_T values. The generalized algorithm successfully detects every available synthetic signal in Stage I of the Sco X - 1Mock Data Challenge convened by the LIGO Scientific Collaboration, recovering the frequency and orbital semimajor axis with accuracies of better than 9.5×10^{-7} Hz (one part in $\sim 10^8$) and 1.6×10^{-3} lt s (one part in $\sim 10^3$) respectively. The Viterbi solver runs in $\approx 2 \times 10^3$ CPU-hr for an isolated source and $\sim 10^5$ CPU-hr for a LMXB source in a typical, broadband (0.5-kHz) search, i.e., $\lesssim 10$ times slower than the previous pipeline.

DOI: 10.1103/PhysRevD.104.042003

I. INTRODUCTION

Rapidly rotating neutron stars with time-varying mass and current quadrupole moments are promising targets of searches for continuous-wave gravitational radiation by long-baseline interferometers such as the Laser Interferometer Gravitational Wave Observatory (LIGO) and Virgo [1]. Several classes of isolated and accreting neutron stars are predicted to be approaching detection, if they emit at or near indirect amplitude limits derived from energy or angular momentum conservation arguments based on electromagnetic observations [2–4].

Among the challenges faced by such experiments is the fact that the signal frequency is often unknown or highly uncertain and wanders stochastically due to irregularities in the star's rotation, known as *spin wandering* or timing noise [5–7]. For some isolated targets, such as nonpulsating neutron stars in supernova remnants, the spin frequency f_* of the crust and corotating magnetosphere cannot be observed, e.g., central compact objects like Cassiopeia A

amelatos@unimelb.edu.au

or the putative neutron star in SNR 1987A [8-10]. In radio pulsars like the Crab, on the other hand, $f_*(t)$ is measured accurately as a function of time t by timing the radio pulsations, but there is no guarantee that the crust corotates exactly with the gravitational-wave-emitting quadrupole [11]. For accreting targets, such as low-mass x-ray binaries (LMXBs),¹ the accretion can drive electromagnetic signatures-thermal x-ray pulsations or type I x-ray burst oscillations—which allow $f_*(t)$ to be measured. However, $f_{*}(t)$ is unknown in some of the brightest sources, like Scorpius X - 1 (Sco X - 1), which exhibit neither signature. [12] Indirect upper limits on the characteristic gravitational wave strain h_0 , [13] based on energy conservation in isolated sources (i.e., the star spins down entirely due to gravitational radiation) and angular momentum conservation in binary sources (i.e., accretion torque balance), imply $h_0 \propto \tau^{-1/2}$ and $h_0 \propto F_X^{1/2}$ respectively, where $\tau =$ $f_*(2|\dot{f}_*|)^{-1}$ denotes the spin-down age, and F_X denotes the X-ray flux [1]. Hence the most promising targets—young, isolated objects and x-ray-luminous accretors-can be those for which the least is known about $f_*(t)$.

One powerful strategy for overcoming the challenge of spin wandering—especially in LMXB searches—is to track $f_*(t)$ with a hidden Markov model (HMM). [14] Given a time-ordered sequence of observations, a HMM relates each observation to the system's underlying, hidden state [e.g., $f_*(t)$] by an emission probability (e.g., a detection statistic of some type). The hidden state evolves through a concurrent sequence, whose stepwise transitions are modeled probabilistically as well (e.g., as a random walk).

In the gravitational wave context, a HMM solved by the fast, recursive, Viterbi algorithm [15] has been implemented as a general-purpose search pipeline and applied to look for the LMXB Sco X - 1 in Advanced LIGO data [16,17]. The pipeline exists in two versions.

(i) Version I calculates the emission probability by summing the maximum likelihood \mathcal{F} -statistic [13] at orbital sidebands incoherently without reference to the orbital phase [18,19]. Given Gaussian noise with one-sided amplitude spectral density $S_h(2f_*)^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$, representative of Advanced LIGO's design sensitivity, version I detects isolated sources with $h_0 \ge 2 \times 10^{-26}$ and binary sources with $h_0 \ge 8 \times 10^{-26}$ and finds 41 out of 50 injected signals in Stage I of the Sco X – 1 Mock Data Challenge (MDC) [18,20]. It was applied to data from Advanced LIGO's first observing run (O1) and returned the upper limit $h_0 \le h_0^{95\%} = 5 \times 10^{-25}$ (95% confidence) at 106 Hz for Sco

X - 1, noting that O1 did not reach full design sensitivity [16].

(ii) Version II tracks orbital phase as well as f_{*}(t) and sums the sideband power coherently using a Jacobi-Anger decomposition of the *F*-statistic [19]. Given S_h(2f_{*}) = 4 × 10⁻²⁴ Hz^{-1/2}, it detects isolated and binary sources with h₀ ≥ 2 × 10⁻²⁶ and finds all 50 injections in Stage I of the Sco X – 1 MDC. It is being applied to data from Advanced LIGO's second [17] and third observing runs.

In this paper, we extend version II of the HMM to track the rotational phase (i.e., the phase of the carrier wave) as well as the orbital phase. The result is an algorithm (version III) which performs nearly as well as a fully coherent matched filter like the \mathcal{F} -statistic, when the phase evolution is known electromagnetically. It maintains the same level of performance, when the phase evolution is unknown, as long as the HMM time-step is chosen to be shorter than the spin wandering timescale [7]. Ensuring that the latter condition is satisfied involves trial and error but is not taxing computationally for most realistic searches. Version III of the HMM is built on a phase-dependent version of the Bayesian matched filter called the \mathcal{B} -statistic used in loosely coherent and related continuous-wave searches [21–26]. It outperforms versions I and II because (i) the \mathcal{B} -statistic is more sensitive than the \mathcal{F} -statistic, and (ii) the in-built requirement of phase continuity reduces false alarms, as discussed in Sec. II. It leverages the existing, easy-to-use, thoroughly tested software infrastructure housed in the LIGO Scientific Collaboration Algorithm Library (LAL). Several of its subroutines and intermediate data products are shared by the \mathcal{F} -statistic and versions I and II of the HMM.²

The paper is structured as follows. In Secs. II-IV we describe how to modify the emission and transition probabilities of the HMM to track the rotational phase. The performance of the extended HMM is then tested by performing Monte-Carlo simulations with Gaussian noise for isolated and binary sources in Secs. V and VI respectively. Specifically, the sensitivity is calculated as a function of the user-selected false alarm and false dismissal probabilities and compared for versions I, II, and III of the HMM. The accuracy of frequency and phase recovery as part of a successful detection is also quantified. Finally we run the extended HMM on data from Stage I of the Sco X - 1 MDC in Sec. VII and confirm that it detects every injection easily. Implications for future gravitational wave searches and their astrophysical impact are discussed briefly in Sec. VIII. Among them is the tantalizing possibility that a gravitational wave detection of spin wandering (possibly in conjunction with radio/x-ray timing

¹In this paper, we follow the usual shorthand of using the term LMXB interchangeably to refer to either the binary system or the neutron star therein.

²A Viterbi-based algorithm has also been developed to perform nonparametric, all-sky searches [27]. Generalizing it to track phase as well as frequency lies outside the scope of this paper.

data) may clarify its physical origin, which remains a subject of debate in both isolated [5,28–32] and accreting [6,33–37] systems.

II. HMM TRACKING

HMM frequency tracking is exploited widely in engineering applications ranging from radar and sonar analysis [38] to mobile telephony [39,40] and has been extended to handle amplitude and phase information and multiple targets [41–43]. It delivers accurate estimation, when the signal-to-noise ratio (SNR) is low, but the sample size is large [14], as in continuous-wave gravitational wave data analysis. In this section we describe how to generalize a HMM that tracks $f_*(t)$ to one that tracks the rotational phase $\Phi_*(t)$ (and hence the carrier phase of the signal) as well as $f_*(t)$. Section II A sets out the tracking framework in its general form [14,44]. Section II B explains the central role played by stepwise phase continuity in reducing the HMM's false alarm rate. Section IIC discusses how to discretize the HMM's state space and the related challenges involved in enforcing phase continuity, when the emission probability is calculated from the output of a frequencydomain matched filter like the \mathcal{F} -statistic. Modified transition and emission probabilities are presented in Secs. III and IV.

A. General formulation and drift timescale

A HMM is a probabilistic finite state automaton defined by a hidden (unobservable) state variable, q(t), and an observable state variable, o(t). The automaton jumps through a time-ordered sequence of observations, $O = \{o(t_0), ..., o(t_{N_T})\}$, at discrete times $t_0 \le ... \le t_{N_T}$. In general there exist $N_Q^{N_T+1}$ possible hidden-state paths, $Q = \{q(t_0), ..., q(t_{N_T})\}$, which are consistent with O. Here N_Q counts the finite number of discrete values, that q(t) can take at time t.

Given *O*, some paths are more likely than others. If we assume that the automaton is Markovian, such that the transition probability from $q(t_n)$ to $q(t_{n+1})$ depends only on $q(t_n)$, then the probability that *Q* gives rise to *O* equals

$$\Pr(Q|O) = L_{o(t_{N_T})q(t_{N_T})}A_{q(t_{N_T})q(t_{N_{T-1}})} \times \cdots \times L_{o(t_1)q(t_1)}A_{q(t_1)q(t_0)}\Pi_{q(t_0)}.$$
 (1)

In (1),

$$A_{q_j q_i} = \Pr[q(t_{n+1}) = q_j | q(t_n) = q_i]$$
(2)

is the transition probability matrix;

$$L_{o_j q_i} = \Pr[o(t_n) = o_j | q(t_n) = q_i]$$
(3)

is the emission probability matrix, namely the probability that the system is observed in state $o(t_n)$ while occupying the hidden state $q(t_n)$; and

$$\Pi_{q_i} = \Pr[q(t_0) = q_i] \tag{4}$$

is the prior vector, namely the probability that the system occupies the hidden state $q(t_0)$ initially.

To solve the HMM, one seeks the most probable path $Q^*(O)$, which maximizes Pr(Q|O) given O, viz.

$$Q^*(O) = \arg \max \Pr(Q|O).$$
 (5)

The maximization can be done in many ways. In previous gravitational wave applications as well as in this paper, we employ the Viterbi algorithm, [14,15] whose logic and pseudocode are summarized briefly in Appendix A. The Viterbi algorithm is a dynamic programming algorithm. It is computationally efficient, executing of order $(N_T + 1)N_Q \ln N_Q$ floating point operations.

Table I summarizes how the general framework above maps onto versions I, II, and III of the HMM. For each version, it specifies the intended astrophysical target, the hidden astrophysical variables being tracked, the intermediate data inputs distilled from the raw observations (which go into calculating $L_{o_jq_i}$), as well as the forms of $A_{q_jq_i}$, $L_{o_jq_i}$, and Π_{q_i} , which define the probabilistic structure of the HMM. The entries in each column are discussed in detail when introduced in Secs. II–IV, together with full

TABLE I. Comparison of HMM versions I, II, and III: intended targets (column 2), hidden variables (column 3), intermediate data inputs (column 4), and probabilistic structure (columns 5–7). The entries in each column are discussed in detail in Secs. II–IV. In column 4, the terms Fourier and Bessel refer to ordinary and Bessel-weighted Fourier transforms of the raw interferometer data respectively, the latter to account for binary orbital phase, which go into calculating $L_{o_jq_i}$ as described in Sec. IV. In column 5, which defines $A_{q_jq_i}$, random walk refers to a discrete-time, simple random walk, and Ornstein-Uhlenbeck refers to continuous-time, damped Brownian motion, as described in Sec. III. The detection statistics \mathcal{F} , \mathcal{J} , and \mathcal{B} in column 6 are defined mathematically when first introduced in Secs. II–IV.

Version	Target	q(t)	o(t)	$A_{q_jq_i}$	$L_{o_jq_i}$	Π_{q_i}	Refs.
I	Isolated	$f_*(t)$	Fourier	Random walk	\mathcal{F} (max. likelihood)	Uniform	[16,18]
Π	Binary	$f_*(t)$	Bessel	Random walk	\mathcal{J} (max. likelihood)	Uniform	[17,19]
III	Isolated	$\Phi_{*}(t), f_{*}(t)$	Fourier	Ornstein-Uhlenbeck	\mathcal{B} (Bayesian)	Uniform	This paper
	Binary	$\Phi_*(t), f_*(t)$	Bessel	Ornstein-Uhlenbeck	\mathcal{B} (Bayesian)	Uniform	This paper

mathematical definitions of the various terms and symbols, e.g., \mathcal{F} , \mathcal{J} , and \mathcal{B} . In this paper, we take $q(t) = [f_*(t), \Phi_*(t)]$. We adopt a flat prior, as in previous work, [18,19] and track the phase difference $\Phi_*(t_{n+1}) - \Phi_*(t_n)$ across each HMM step; $\Phi_*(0)$ is the result of a historical accident, which obviates the need to track the absolute phase.

In gravitational wave applications, the underlying, stochastic evolution of q(t) is continuous. Nonetheless the discrete-time HMM defined by (1)-(5) provides an appropriate analysis framework, as long as the duration $T_{\text{drift}} = t_{n+1} - t_n$ of each HMM step is chosen wisely. A recipe for choosing T_{drift} in versions I and II of the HMM is given in previous papers [18,19]. The generalized recipe for version III is set out in Appendix B, where the key condition on T_{drift} is given by Eq. (B1). One always has $T_{\text{SFT}} \leq T_{\text{drift}} \leq T_{\text{obs}}$, where T_{SFT} denotes the duration of the short-time Fourier transforms (SFTs) [45] used to compute $L_{o_i q_i}$ (see Sec. II C and Appendix B), and $T_{obs} =$ $N_T T_{\text{drift}} \sim 1$ yr is the total observation time. The SFTs are a data management device to assist with storage and inputoutput. They divide the observing run into short stretches, typically $T_{\text{SFT}} = 1800$ s in length, during which one assumes that the antenna beam pattern is approximately constant (neglecting rotation of the Earth), and the detector noise is approximately stationary. They are knitted together to compute a detection statistic such as the \mathcal{F} -statistic coherently over an interval T_{drift} . By contrast, T_{drift} is a user-selected time interval which contains an integer number of SFTs, during which one assumes that the system stays within a single HMM state, if condition (B1) is satisfied. Detailed implementation instructions, explaining how the SFTs are converted into "data atoms" and hence values of the emission probability $L_{o,q}$, are provided in Ref. [46].

B. Phase continuity

In previous implementations of HMM-based gravitational wave searches, [16,18,19] $L_{o(t_n)q_i}$ is computed from the maximum-likelihood, frequency-domain matched filter called the \mathcal{F} -statistic [13] or a close variant, evaluated over the time interval $t_{n-1} \le t \le t_n$. For an isolated source, the \mathcal{F} -statistic concentrates all the signal power into a single frequency bin, of width $\Delta f_{\text{drift}} = (2T_{\text{drift}})^{-1}$, provided that the T_{drift} condition (B1) holds. For a binary source, the \mathcal{F} -statistic disperses the signal power into approximately 2M' + 1 = 2ceil $(2\pi f_*a_0) + 1$ orbital sidebands, separated by P^{-1} in frequency, where a_0 is the projected semimajor axis of the binary orbit, P is the orbital period, and ceil(...)returns the lowest integer greater than or equal to its argument. However, it is possible to redirect most of the signal power into a small subset ($\ll 2M' + 1$) of frequency bins by summing the \mathcal{F} -statistic values at the orbital sidebands with an appropriate weighting, namely Bessel coefficients arising from the Jacobi-Anger expansion of the waveform. If the coefficients are squared Bessel functions, the sum is incoherent, and $L_{o(t_n)q_i}$ exhibits a narrow, cuspy peak as a function of frequency, as in version I of the HMM (Bessel-weighted \mathcal{F} -statistic) [19]. If the coefficients include powers of $e^{i\phi_a}$, where ϕ_a is a reference phase (usually defined by the orbit's ascending node), and the \mathcal{F} -statistic is factorized into a product of complex numbers before summation, the sum is coherent with respect to orbital phase, and $L_{o(t_n)q_i}$ contains all the signal power in a single frequency bin, of width $\Delta f_{\text{drift}} = (2T_{\text{drift}})^{-1}$, as in version II of the HMM (\mathcal{J} -statistic) [19]. In summary, it is always possible to concentrate all the signal power into a single frequency bin, by calculating $L_{o(t_n)q_i}$ from the \mathcal{F} -statistic (isolated source) or \mathcal{J} -statistic (binary source). This result is confirmed by numerous Monte Carlo simulations in Ref. [19].

There is only one "correct" frequency bin at each HMM step, and $Q^*(O)$ either finds it or not. It is therefore natural to ask what extra advantage rotational phase tracking confers, when the optimal path $Q^*(O)$ in versions I and II of the HMM already captures the maximum signal power available to any HMM, for the reason set out in the previous paragraph. The answer is that phase tracking increases the detection probability by sharpening the HMM's ability to discriminate against spurious sequences. For example, if a strong noise event occurs in the *i*th frequency bin at the *n*th step, then $Q^*(O)$ is likely to contain $q(t_n) = q_i$, if frequency is the only hidden state variable. Yet if phase is tracked as well, the HMM is more likely to reject the spurious path containing $q(t_n) = q_i$ in favor of another path with lower $L_{o(t_n)q_i}$ $(j \neq i)$ but higher $A_{q(t_{n+1})q_i}$ and $A_{q_iq(t_{n-1})}$, i.e., a path whose transition probabilities into and out of the nth step are more consistent with phase continuity. This is equivalent to the distinction between a semicoherent and a coherent search. The latter is $\approx N_T^{1/4}$ times more sensitive than the former because it effectively reduces the denominator in the SNR by excluding false alarms that violate phase continuity.

We implement rotational phase tracking by enlarging the state vector to two dimensions for an isolated source, with $q(t) = [f_*(t), \Phi_*(t)]$, and four dimensions for a binary source, with $q(t) = [f_*(t), a_0(t), \phi_a(t), \Phi_*(t)]$. Under normal astrophysical circumstances, a_0 and ϕ_a are constant throughout a full search ($T_{obs} \leq 1$ yr), so there is no need to track them. Hence, for both target classes, the HMM reduces to two dimensions, with $q(t) = [f_*(t), \Phi_*(t)]$, except that it is computed on a grid of (a_0, ϕ_a) pairs for a binary source; see Sec. II A in Ref. [19]. This approach is readily parallelizable across (a_0, ϕ_a) pairs and sources.

C. Grid resolution

How do we select the number of hidden states, $N_Q = N_{f_*}N_{\Phi_*}$, with $N_{f_*} = B/\Delta f_{\text{drift}}$ and $N_{\Phi_*} = 2\pi/\Delta \Phi_{\text{drift}}$, where $B = \max f_* - \min f_*$ is the bandwidth, and $\Delta \Phi_{\text{drift}}$ is the

width of a phase bin? There are many valid ways to do this, as discussed in Appendix B, noting that Δf_{drift} and $\Delta \Phi_{\text{drift}}$ are related through $\Phi_*(t) = 2\pi \int_0^t dt' f_*(t')$. The choice comes down to how the HMM emission probability is calculated from the data, as foreshadowed in Sec. II A. In this paper, we seek to leverage the existing, easy-to-use, thoroughly tested software infrastructure for frequencydomain continuous-wave searches maintained in the LAL suite, including the \mathcal{F} -statistic [13,46], \mathcal{B} -statistic [21,23–25], and intermediate data products generated by the \mathcal{F} -statistic; see Sec. IV in this paper and Sec. III A in Ref. [19]. These software tools are built around Fourier transforms. We are therefore obliged to take Δf_{drift} to be the half-Nyquist bin width of the \mathcal{F} -statistic evaluated over a time interval of duration T_{drift} , viz. $\Delta f_{\rm drift} = (2T_{\rm drift})^{-1}.$

The half-Nyquist criterion creates a problem: small uncertainties in f_* of $\pm \Delta f_{\text{drift}}$ due to binning lead to large uncertainties in Φ_* of $\pm 2\pi T_{\text{drift}}\Delta f_{\text{drift}} = \pm \pi$ when propagated forward over one HMM time-step, degrading the HMM's ability to track $\Phi_*(t)$. One can circumvent this obstacle by abandoning the frequency domain, thereby surrendering its practical advantages. Alternatively, one can achieve sub-Nyquist frequency resolution ($\ll \Delta f_{\text{drift}}$ and hence $N_{\Phi_*} \gg 1$) by modeling the underlying evolution of $q(t') = [f_*(t'), \Phi_*(t')]$ within a HMM time-step $(t_n \le t' \le t_n + T_{\text{drift}})$. We adopt the latter approach. A simple, linear ramp does not improve the situation much, e.g., $f_*(t') = f_*(t_n) \pm (t' - t_n) \Delta f_{\text{drift}} / T_{\text{drift}}$ implies $\Phi_*(t_{n+1}) - \Phi_*(t_n) = 2\pi T_{\text{drift}} f_*(t_n) \pm \pi/2$, which is still a large fractional uncertainty. We find instead that evolving q(t') stochastically according to a phase-wrapped, Ornstein-Uhlenbeck process (i.e., Brownian motion that is 2π -periodic in phase) yields good practical results. The approach is described in Sec. III and Appendix C and tested against Monte Carlo simulations in Secs. V and VI. It is analogous to a vernier scale, in which the frequency bins yield a coarse first approximation to the frequency, and the phase bins yield a refined approximation. We find empirically that $N_{\Phi_*} = 32$ is adequate for the transition probabilities assumed in this paper (see Appendix B and footnote 10). Sub-Nyquist frequency resolution is routinely achieved in signal processing problems, where phase tracking is involved, using a variety of techniques [41].

III. TRANSITION PROBABILITIES

In this section we introduce an Ornstein-Uhlenbeck (Brownian) model of the stochastic, intrastep evolution of the star's rotation and hence the signal's frequency and phase. Transition probabilities $A_{q_jq_i}$ for frequency-phase tracking are presented in Sec. III A. The Ornstein-Uhlenbeck model is controlled by two auxiliary parameters. We explain how to set these parameters given T_{drift} in Sec. III B.

A. Stepping forward in frequency and phase

In versions I and II of the HMM, it is assumed that $f_*(t)$ jumps by -1, 0, or +1 frequency bins at every step with equal probability 1/3.³ In version III of the HMM, we again assume that $f_*(t')$ executes an unbiased random walk for $t_n \le t' \le t_n + T_{\text{drift}}$ and choose T_{drift} according to condition (B1), as discussed in Appendix B. However we model the intrastep random walk explicitly as an Ornstein-Uhlenbeck process that is 2π -periodic in phase. The aim is to derive $A_{q_jq_i}$ in a way that self-consistently relates the jumps in $f_*(t)$ and $\Phi_*(t)$ and allows adequate phase resolution ($N_{\Phi_*} \gg 1$), as discussed in Sec. II C.

The Ornstein-Uhlenbeck process is described by a pair of stochastic differential equations,

$$\frac{df_*}{dt} = -\gamma f_* + \sigma \xi(t), \tag{6}$$

$$\frac{d\Phi_*}{dt} = f_*. \tag{7}$$

It is controlled by two parameters: γ , a damping rate, and σ , a fluctuation amplitude. The fluctuating torque $\xi(t)$ has white noise statistics, viz.

$$\langle \xi(t) \rangle = 0, \tag{8}$$

$$\langle \xi(t)\xi(t')\rangle = \delta(t-t'), \qquad (9)$$

where $\langle ... \rangle$ denotes an ensemble average. We assume that there is no white noise forcing term in (7), i.e., the principal axes of the gravitational-wave-emitting quadrupole are fixed in the body frame rotating instantaneously at the frequency $f_*(t)$. In Brownian motion in thermal equilibrium, γ and σ are related by the fluctuation-dissipation theorem, with σ^2/γ proportional to the system temperature. Here, in contrast, γ and σ are independent. We explain how to choose them in practice in Sec. III B.

The stochastic differential equations (6) and (7) are equivalent to the forward Fokker-Planck equation [50]

$$\frac{\partial p}{\partial t} = \frac{\partial (\gamma f_* p)}{\partial f_*} - \frac{\partial (f_* p)}{\partial \Phi_*} + \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial f_*^2}, \qquad (10)$$

whose solution $p(t, f_*, \Phi_*)$ equals the probability density that the hidden state lies in the infinitesimal domain $(f_*, f_* + df_*) \cup (\Phi_*, \Phi_* + d\Phi_*)$ at time *t* if it started at $q(0) = [f_*(0), \Phi_*(0)]$ at t = 0, i.e., $p(0, f_*, \Phi_*) =$ $\delta[f_* - f_*(0)]\delta[\Phi_* - \Phi_*(0)]$. Hence evolving $p(t, f_*, \Phi_*)$ from $t = t_n$ to $t = t_{n+1}$ is exactly what one needs to

³As in previous papers, we exclude the possibility of impulsive rotational glitches with $f_*(t_{n+1}) - f_*(t_n) > \Delta f_{\text{drift}}$; [47,48] see footnote 3 in Ref. [18] and compare Ref. [49].

calculate the transition probabilities $A_{q_jq_i}$, as defined by (2). Specifically we write

$$A_{(f_{*j},\Phi_{*k})(f_{*l},\Phi_{*m})} = p(t_{n+1},f_{*j},\Phi_{*k})\Delta f_{\text{drift}}\Delta\Phi_{\text{drift}} \quad (11)$$

with $f_*(t_n) = f_{*l}$ and $\Phi_*(t_n) = \Phi_{*m}$, where the integers j, l and k, m index discrete frequency and phase bins respectively. Analytic formulas are derived for $p(t, f_*, \Phi_*)$ and its characteristic function in Appendix C [51].

In the Viterbi algorithm, it is sometimes more convenient to calculate the backward transition probabilities, $A_{q_jq_i}^{\text{back}} = \Pr[q(t_n) = q_j | q(t_{n+1}) = q_i]$. This can be done by solving the backward Fokker-Planck equation, which is adjoint to (10). Details and formulas are given in Appendix C. The resulting PDF is a 2π -wrapped Gaussian; see Eqs. (C10)–(C15).

B. Control parameters

How should the control parameters γ and σ be chosen? Two conditions must be satisfied during every HMM step: γ must be small enough, such that $\langle f_* \rangle$ does not drift by more than one frequency bin, Δf_{drift} ; and σ must be large enough, so that we have $\langle f_*^2 \rangle - \langle f_* \rangle^2 \approx (\Delta f_{\text{drift}})^2$, i.e., probability leaks significantly into the frequency bins on either side of the starting bin but not much further. From the moment formulas in Appendix C, typical of a diffusion process, the above conditions reduce to

$$f_*[1 - \exp(-\gamma T_{\text{drift}})] < \Delta f_{\text{drift}}$$
(12)

and



respectively for all f_* in the observation band. For a typical LMXB search with $T_{\rm drift} = 10$ d and $f_* \gtrsim 50$ Hz, we have $\Delta f_{\rm drift}/f_* \lesssim 1 \times 10^{-8}$, $\gamma T_{\rm drift} \ll 1$, and hence $\gamma < (2f_*T_{\rm drift}^2)^{-1}$ and $\sigma \approx (4T_{\rm drift}^3)^{-1/2}$.

Figure 1 presents an example of the transition probabilities for an illustrative choice of γ and σ satisfying the constraints in the previous paragraph and used subsequently in the validation experiments in Secs. V and VI. Contours of the PDF $A_{q_iq_i}$ in the $f_* - \Phi_*$ plane are plotted in Fig. 1(a). Three constant- f_* cross sections are plotted versus Φ_* in Fig. 1(b). We find that $p(t_{n+1}, f_*, \Phi_*)$ leaks significantly into the frequency bins on either side of the starting bin, with $A_{f_{*i\pm 1},f_{*i}} = 0.196$ and $A_{f_{*i},f_{*i}} = 0.608$ (normalized). In this implementation, the PDF is truncated to give $A_{f_{*i\pm 2,3,\ldots,f_{*i}}} = 0$ to achieve computational savings, but if one does not truncate one finds $A_{f_{*i\pm 2},f_{*i}} =$ 3.82×10^{-4} . The probabilities of jumping up or down in frequency are equal, as in version I of the HMM, while the probability of staying in the same bin is higher (0.608) than in version I (0.333).

In contrast, the PDF extends over many bins in phase, as is clear from Fig. 1(b), with full-width half-maximum ≈ 1.78 rad (nine bins). Phase wrapping ensures periodicity in Φ_* , but for the plotted parameters the PDF is tiny at the edges of the plot, and it is hard to verify the periodicity by eye. The initial state $q(t_n)$ determines whether the phase wraps or not. Figure 1(b) confirms that phase wrapping alternates between even and odd frequency bins (and depends on whether f_* jumps by zero or $\pm \Delta f_{drift}$),



FIG. 1. Forward transition probabilities $A_{q_jq_i} = \Pr[q(t_{n+1}) = q_j|q(t_n) = q_i]$ (not normalized) for $f_*(t_n) = 111.0$ Hz, $\Phi_*(t_n) = 0$ rad, $\gamma = 1.0 \times 10^{-16}$ s⁻¹, and $\sigma = 3.7 \times 10^{-10}$ s^{-3/2}. (a) Contour plot versus $f_*(t_{n+1}) - f_*(t_n)$ and $\Phi_*(t_{n+1}) - \Phi_*(t_n)$. The color scale is arbitrary; hot colors are high, cool colors are low. The white grid delineates frequency-phase bins. The horizontal and vertical axes are labeled by number of bins. A subset of the hidden state space is plotted for clarity. (b) Cross sections at fixed $f_*(t_{n+1}) - f_*(t_n) = 0, \pm \Delta f_{drift}$. The crosses, plus signs, and asterisks mark phase bins. The horizontal axes indicate $\Phi_*(t_{n+1}) - \Phi_*(t_n)$ in units of radians (top) and number of bins (bottom). The backward transition probabilities are identical but centered on $q(t_{n+1})$ instead of $q(t_n)$.

as discussed in Sec. II C; the phase jumps by π , when the frequency bin at t_n is odd, and by zero when the frequency bin at t_n is even. The contours slope diagonally, because f_* and Φ_* are correlated, with $\langle f_*\Phi_*\rangle - \langle f_*\rangle\langle\Phi_*\rangle \neq 0$; see Eq. (C14) in Appendix C. The shape of the contours is the same for the forward and backward transition probabilities, but the centroid shifts with $q(t_n)$ and $q(t_{n+1})$ respectively.

The above recipe for setting γ and σ is sensible but not unique. The optimal values of the control parameters (and $T_{\rm drift}$) depend on the waveform of the true signal, which is unknown in advance in an astronomical setting. Altering γ and σ does not introduce a systematic bias, because the Ornstein-Uhlenbeck process is symmetric with respect to positive and negative frequency jumps, but in general it increases or decreases the sensitivity modestly. It is found empirically that HMMs are robust to the exact form of $A_{q_iq_i}$, which is why the naive choice of $A_{q_iq_i}$ in version I of the HMM works well. [14] The extra sensitivity in version III comes from phase tracking, which depends weakly on γ and σ , because the PDF in Fig. 1 is broad in phase. When publishing searches with real data, it is important to emphasize that any upper limits are conditional on the signal model, which includes γ , σ , and T_{driff} .

IV. EMISSION PROBABILITIES

For the class of frequency-domain, continuous-wave searches considered in this paper, $L_{o_jq_i}$ in (3) can be expressed in terms of a suitable frequency-phase detection statistic $G(f_*, \Phi_*)$ as

$$L_{o(t_n)q_i} \propto \exp[G(f_{*i'}, \Phi_{*i''})].$$
 (14)

Here $G(f_{*i'}, \Phi_{*i''})$ is the log likelihood that $f_*(t_{n-1})$ lies in the *i*'th frequency bin $[f_{*i'}, f_{*i'} + \Delta f_{drift}]$, and $\Phi_*(t_{n-1})$ lies in the *i*"th frequency bin $[\Phi_{*i''}, \Phi_{*i''} + \Delta \Phi_{drift}]$, with $i = i'N_{\Phi_*} + i''$, given the data $o(t_n)$.⁴ Concretely $o(t_n)$ comprises a set of strain measurements, numbering T_{drift} multiplied by the interferometer sampling rate, or their Fourier-transformed counterparts, sampled during the interval $t_{n-1} \le t' \le t_n$. There exist many valid ways to construct $G(f_*, \Phi_*)$, depending on computational constraints, the data format, and the assumed model for the evolution of $q(t) = [f_*(t), \Phi_*(t)]$.

In this paper, we strive to exploit the easy-to-use, thoroughly tested software infrastructure in the LAL suite associated with the \mathcal{F} -statistic [13]. We are therefore led to build $G(f_*, \Phi_*)$ as a frequency-domain matched filter, using as many existing LAL components as possible. In versions I and II of the HMM, $G(f_*, \Phi_*)$ is constructed as a maximum likelihood estimator from the \mathcal{F} -statistic (isolated source) or a Bessel-weighted sum of \mathcal{F} -statistic values (binary source) [18,19]. In version III, we press into service the phase-dependent generalization of the Bayesian \mathcal{B} -statistic used in loosely coherent searches. [21–26] The latter choice is justified against maximum likelhood alternatives in Appendix D. We review briefly the signal model and its definitions in Sec. IV A, define the frequency domain intermediate data products that we need (e.g., complex Fourier amplitudes generated by the LAL) in Sec. IV B, and present a formula for $G(f_*, \Phi_*)$ in terms of the \mathcal{B} -statistic in Sec. IV C.

A. Signal model and likelihood

The gravitational wave signal measured at the Earth from a biaxial rotor can be written as a linear combination of eight independent components, [13]

$$h(t) = \sum_{i=1}^{4} A_{1i}h_{1i}(t) + A_{2i}h_{2i}(t).$$
(15)

In (15), A_{1i} and A_{2i} are arbitrary amplitudes set by the source, and $h_{1i}(t)$ and $h_{2i}(t)$ are defined in Ref. [13] as sinusoidal functions of $\Phi(t)$ and $2\Phi(t)$ respectively, where $\Phi(t)$ is the signal phase at the detector [note: $\Phi(t) \neq \Phi_*(t)$ in general]. The amplitudes of $h_{1i}(t)$ and $h_{2i}(t)$ are modulated diurnally by the antenna beam-pattern functions a(t) and b(t), defined by Eqs. (12) and (13) respectively in Ref. [13].

Following Eqs. (18) and (96) in Ref. [13], we split the signal phase into five terms,

$$\Phi(t) = 2\pi f_0[t + \Phi_m(t;\alpha,\delta)] + \Phi_s[t;f_0^{(k)},\alpha,\delta] - 2\pi f_0 a_0 \sin(2\pi t/P - \phi_a) + \Phi_w(t).$$
(16)

(1)

In (16), f_0 is the signal frequency at the detector,⁵ Φ_m is a time shift produced by the diurnal and annual motions of the detector and source relative to the Solar System barycenter, Φ_s is a phase shift combining the latter two effects with the intrinsic, deterministic, *secular* evolution of the source through the frequency derivatives $f_0^{(k)} = d^k f_0/dt^k$ ($k \ge 1$) (see Eq. (14) in Ref. [13]), the fourth term ($\propto a_0$) is the Doppler modulation produced by the source's orbital motion in a binary system, and $\Phi_w(t)$ is the phase accumulated from *stochastic* spin wandering. The sky position of the source (right ascension α , declination δ) enters Φ_m and Φ_s . Naturally it is possible to absorb the binary orbit and stochastic spin wandering into $f_0^{(k)}$, and

⁴Equally one can use some other reference time, e.g., t_n .

⁵One has $f_0 \neq f_*(t)$ in general. $f_*(t)$ is the true, underlying spin frequency of the star, which we cannot measure directly and which forms one component of the hidden state. f_0 is any arbitrary frequency, where the emission probability and associated phase model (16) are evaluated, which may or may not coincide with $f_*(t)$, depending on where in the parameter space we look.

hence absorb the fourth and fifth terms in (16) into Φ_s , but it is clearer to keep the contributions separate in what follows.

The output from a single interferometer is given by x(t) = h(t) + n(t), where n(t) denotes additive noise. The normalized log likelihood after measuring the time series x(t) over the interval $0 \le t \le T_{obs}$ is proportional to

$$\ln \Lambda' = (x||h) - \frac{1}{2}(h||h), \tag{17}$$

where we define the inner product

$$(x||y) = \frac{2}{T_{\text{obs}}} \int_0^{T_{\text{obs}}} dt x(t) y(t).$$
(18)

In versions I and II of the HMM, the emission probability is computed by maximizing $\ln \Lambda'$ in (17) with respect to the amplitudes A_{1i} and A_{2i} and evaluating the result on a grid of f_0 values to find the peak. (This procedure is not exactly the same as maximizing over A_{1i} , A_{2i} , and f_0 simultaneously.) The result is a sum of terms quadratic in $(x||h_{1i})$ or $(x||h_{2i})$, which can be computed from Fourier-transformed interferometer data as discussed in Sec. III D in Ref. [13] and Appendix D below; see also Ref. [46] and Sec. II A in Ref. [19]. In version III of the HMM, the emission probability is computed via the Bayesian \mathcal{B} -statistic, defined in Sec. IV C. The \mathcal{B} -statistic can be computed efficiently from the same, Fourier-transformed interferometer data used by versions I and II. We define the relevant Fourier integrals in Sec. IV B.

B. Fourier integrals

The waveforms $h_{1i}(t)$ and $h_{2i}(t)$ in (15) are amplitude modulated by the antenna beam pattern functions a(t) and b(t). The log likelihood ln Λ' in (17) is a function of (x||h), which reduces to calculating the Fourier transforms of x(t)a(t) and x(t)b(t), because one has $(x||h_{11}) =$ $(x||a(t)\cos \Phi(t))$ for example. For an isolated source $(a_0 = 0)$, let us define the Fourier integrals [13]

$$F_{1a}(f_0) = \int_0^{T_{\text{obs}}} dt_{\text{b}} x[t(t_{\text{b}})] a[t(t_{\text{b}})] e^{-i\Phi_{\text{s}}[t(t_{\text{b}})] - 2\pi i f_0 t_{\text{b}}}, \quad (19)$$

$$F_{1b}(f_0) = \int_0^{T_{\text{obs}}} dt_{\text{b}} x[t(t_{\text{b}})] b[t(t_{\text{b}})] e^{-i\Phi_{\text{s}}[t(t_{\text{b}})] - 2\pi i f_0 t_{\text{b}}}, \quad (20)$$

where $t_b = t + \Phi_m(t)$ defines a barycentered time coordinate t_b related implicitly to *t* through the time shift arising from the Earth's motion. In this paper, we neglect the secular frequency evolution of the source, e.g., due to electromagnetic braking, and set $f_0^{(k)} = 0$ for $k \ge 1$ and hence $\Phi_s[t(t_b)] = 0$. It is easy to keep $f_0^{(k)} \ne 0$ in (19) and (20) if desired. We also specialize without loss of generality to the case $A_{1i} = 0$, corresponding to a search for one signal frequency (as opposed to two simultaneously).

For a binary source $(a_0 \neq 0)$, the integrands in (19) and (20) feature an extra, Doppler-modulated phase factor $\exp[2\pi i f_0 a_0 \sin(2\pi t/P - \phi_a)]$, derived from (16). Upon expanding this factor using the Jacobi-Anger identity, we find that F_{1a} and F_{1b} should be replaced in $L_{o(t_n)q_i}$ by

$$J_{1a}(f_0) = \sum_{s=-M'}^{M'} J_s(2\pi f_0 a_0) e^{-is\phi_a} F_{1a}(f_0 + s/P), \qquad (21)$$

$$J_{1b}(f_0) = \sum_{s=-M'}^{M'} J_s(2\pi f_0 a_0) e^{-is\phi_a} F_{1b}(f_0 + s/P), \qquad (22)$$

where J_s denotes a Bessel function of order *s* of the first kind. Equations (21) and (22) add together the Fourier amplitudes in orbital sidebands coherently, by taking into account the relative orbital phases of the sidebands. [19] The infinite sums are truncated, because one has $J_s(x) \ll 1$ for $s \gg x$ to a good approximation [52,53].

It turns out that the emission probability $L_{o(t_n)q_i}$ in (14) can be calculated easily from F_{1a} and F_{1b} (isolated source) or J_{1a} and J_{1b} (binary source) in every HMM implementation we consider. The maximum likelihood formulas for $L_{o(t_n)q_i}$ in versions I and II of the HMM are quoted in Appendix D, where it is shown that they (and their phasedependent generalizations) are poorly suited to rotational phase tracking. The \mathcal{B} -statistic adopted in this paper for version III of the HMM is presented next in Sec. IV C.

The Fourier integrals (19) and (20) are taken formally over $0 \le t \le T_{obs}$. In practice, to facilitate data management, the integral is subdivided into "atoms" [46]. Each atom corresponds to one SFT, which is convolved with a sliding-window sinc function to increase the frequency resolution from $(2T_{SFT})^{-1}$ to $(2T_{obs})^{-1}$, as required by (19) and (20). The reader is referred to Sec. IV. 2 in Ref. [46] for full details; see also Sec. III A in Ref. [19]. In this paper, following Ref. [46], we approximate a(t) and b(t) as piecewise-constant during each SFT.

C. Phase-dependent *B*-statistic

The \mathcal{B} -statistic [21] is a Bayesian alternative to the maximum likelihood \mathcal{F} -statistic [13]. It is derived from the likelihood function Λ' in (17) combined with an isotropic prior on the source orientation (i.e., spin axis). Its detection efficiency is ≈ 5 per cent greater than that of the \mathcal{F} -statistic, and it is arguably motivated better astrophysically; the \mathcal{F} -statistic implicitly assumes a uniform prior on the amplitude, whereas the \mathcal{B} -statistic favors lower amplitudes, which is more realistic. [21] In practice, however, the \mathcal{F} -statistic has proved more popular than the \mathcal{B} -statistic, having been preferred in various published LIGO searches, e.g., [2,54] (targeted), [8,55] (directed), and [56,57] (all-sky), as well as forming the basis of versions I and II of the HMM [10,16,18,19]. This is because: (i) the advantage

held by the \mathcal{B} -statistic in terms of detection efficiency is small [21,23,24]; (ii) the \mathcal{F} -statistic software in the LAL was developed first and is now thoroughly tested; and (iii) the \mathcal{B} -statistic involves numerical integrals, which are relatively expensive computationally, although fast approximations do exist [23,24,26].

In this section, we present a phase-dependent version of the B-statistic formulated for loosely coherent searches [23]. The associated emission probability is calculated from (19) and (20) (isolated source) or (21) and (22) (binary source), i.e., the same intermediate data products as versions I and II of the HMM. We settle on this choice after testing several phase-dependent generalizations of the maximum likelihood \mathcal{F} -statistic, as discussed in Appendix D. Empirically we find that: (i) none perform as well as the \mathcal{B} -statistic nor offer any discernible improvement over versions I and II of the HMM; (ii) a HMM based on the \mathcal{B} -statistic approaches the theoretical sensitivity of a fully coherent search; and (iii) the sensitivity improvement exceeds the \approx 5 percent advantage of the \mathcal{B} -statistic over the \mathcal{F} -statistic without any phase dependence, [21] so phase tracking is clearly playing a role. Of course these empirical findings do not constitute a formal proof, that the phasedependent B-statistic always outperforms any phasedependent maximum likelihood estimator, cf. Ref. [21]. However such a formal proof lies outside the scope of this paper and is unnecessary at this stage given the excellent performance achieved in tests with synthetic data in Secs. V and VI. Other competing estimators will be tested in future work, e.g., the phase-relaxed \mathcal{F} -statistic [58].

Instead of maximizing Λ' in (17) with respect to A_{1i} and A_{2i} , we marginalize it (by Bayes's theorem) over uniform priors in three source-dependent variables: (i) the polarization angle, ψ ; (ii) the cosine of the inclination angle, $\cos \iota$; and (iii) the characteristic wave strain, h_0 [23–25]. Let us define

$$A_{+} = \frac{h_0}{2} (1 + \cos^2 \iota) \tag{23}$$

and

$$A_{\times} = h_0 \cos \iota \tag{24}$$

to be the real amplitudes of the plus and cross polarizations respectively, which can be related to A_{2i} as explained in Ref. [59]. (For simplicity we consider the popular case $A_{1i} = 0$ here.) Following Ref. [23], let us also define the auxiliary complex variables

$$w'_{1} = (2h_{0})^{-1}(A_{+}\cos 2\psi + iA_{\times}\sin 2\psi) \qquad (25)$$

and

$$w'_{2} = (2h_{0})^{-1} (A_{+} \sin 2\psi - iA_{\times} \cos 2\psi), \qquad (26)$$

which satisfy the identities

$$l = |w'_1 + iw'_2|^{1/2} + |w'_1 - iw'_2|^{1/2},$$
(27)

 $2h_0w'_1 = A_{21} - iA_{23}$, and $2h_0w'_2 = A_{22} - iA_{24}$. In terms of the above definitions, we obtain the following expression for the marginalized likelihood [23]:

$$\mathcal{B} = \int_0^{\pi} d\psi \int_{-1}^1 d(\cos \iota) \int_0^{h_0^{\max}} dh_0 \exp\left(h_0 U - \frac{h_0^2 V}{2}\right),$$
(28)

with

$$U = w_1'^* R_{1a}(f_0, \Phi_0) + w_2'^* R_{1b}(f_0, \Phi_0), \qquad (29)$$

$$V = A|w_1'|^2 + 2C \operatorname{Re}(w_1'w_2'^*) + B|w_2'|^2, \qquad (30)$$

$$R_{1a}(f_0, \Phi_0) = \operatorname{Re}[\exp(-i\Phi_0)F_{1a}(f_0)], \qquad (31)$$

and

$$R_{1b}(f_0, \Phi_0) = \operatorname{Re}[\exp(-i\Phi_0)F_{1b}(f_0)].$$
(32)

In (30), we have A = (a||a), B = (b||b), C = (a||b), and $C \ll \min(A, B)$ for most sky positions [26]. The *B*-statistic peaks, when the trial phase Φ_0 in (29), (31) and (32) matches the true signal phase at the detector, viz. $\Phi(t)$ in (16).

The h_0 integral in (28) is not normalized as it stands; the HMM disregards multiplicative constants. Hence we can take the limit $h_0^{\text{max}} \rightarrow \infty$ without loss of generality and express the h_0 integral in closed form as an error function. In loosely coherent searches, \mathcal{B} is maximized with respect to Φ_0 [23]. We cannot do the same here, because we track the rotational phase and therefore need \mathcal{B} to depend on Φ_0 . The final result, expressed again in the notation of Ref. [23], is given by

$$\mathcal{B}(f_0, \Phi_0) = \int_0^{\pi} d\psi \int_{-1}^1 d(\cos \iota) \left(\frac{\pi}{2V}\right)^{1/2} \\ \times \exp\left(\frac{U^2}{2V}\right) \left[1 + \operatorname{erf}\left(\frac{U}{\sqrt{2V}}\right)\right], \quad (33)$$

with

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x dy \exp(-y^2).$$
 (34)

The double integral in (33) is evaluated numerically by Simpson's rule in what follows.

Figure 2 presents examples of the emission probability for two signals from an isolated source injected into



FIG. 2. Logarithm of the emission probability, $G(f_0, \Phi_0) = \ln \mathcal{B}(f_0, \Phi_0)$ (not normalized), represented by its contours in the $f_0 - \Phi_0$ plane (color scale arbitrary). (a) Stronger injection; $h_0 = 4.0 \times 10^{-26}$. (b) Weaker injection; $h_0 = 1.7 \times 10^{-26}$. The injections are marked by crosses; their parameters are listed in Table II (isolated source). The observation $o(t_n)$ consists of $T_{obs} = 10$ days of data ($N_T = 1$). The white grid delineates frequency-phase bins with $\Delta f_{drift} = 5.8 \times 10^{-7}$ Hz and $\Delta \Phi_{drift} = \pi/16$. A subset of the hidden states is plotted for clarity. The noise is Gaussian, with $S_h(f_0)^{1/2} = 4 \times 10^{-24}$ Hz^{-1/2} as in Table II.

Gaussian noise with $S_h(f_0) = 4 \times 10^{-24}$. The contour plots in the figure are generated from $T_{\rm obs} = 10$ days of synthetic data. In Fig. 2(a) the stronger injection is clearly detectable, with $h_0 = 4.0 \times 10^{-26}$, and the emission probability peaks near the correct (f_0, Φ_0) bin. Figure 2(b) displays the weaker injection, with $h_0 = 1.7 \times 10^{-26}$. There is a hot spot near the correct bin, but it does not stand out visually from the other, noise-generated hot spots. The constant- f_0 cross section does not peak at the correct value of Φ_0 , although it still has roughly the same functional form as in Fig. 2(a). Note that the emission probability is not always unimodal in the vicinity of the injection as it is in Fig. 2(a). When the complex arguments of F_{1a} and F_{1b} differ sufficiently, \mathcal{B} develops two peaks as a function of Φ_0 (at fixed f_0), only one of which corresponds to the signal. The tests in Sec. V show that the HMM is effective at resolving this ambiguity and identifying the true peak for $N_T > 1$.

V. ISOLATED NEUTRON STAR

We begin by testing version III of the HMM on synthetic data generated by injecting the signal from an isolated neutron star into additive, Gaussian noise. Section VA describes the injection procedure. Tracking results are presented in Sec. V B for a representative sample of synthetic data. A systematic, threshold-based strategy for identifying signal candidates during an astrophysical search is described in Sec. V C and is applied to characterize the performance of the HMM in Sec. V D. The accuracy with which the HMM reconstructs the true hidden state sequence given a successful detection is quantified in Sec. V E. Versions I (isolated source) and III of the HMM are compared at each stage. versions II (binary source) and III are compared in Sec. VI.

A. Synthetic data

The signal phase corresponding to an isolated neutron star is given by (16) with $a_0 = 0$. In the tests below, the stochastic component of the injected phase evolves during the interval $t_n \le t' \le t_{n+1}$ according to $\Phi_w(t') = 2\pi [(t' - t_n) + t_n]$ $(t_n)^3 \dot{f}_*(t_n)/6 + (t'-t_n)^2 \dot{f}_*(t_n)/2 + (t'-t_n)f_*(t_n)] + \Phi_{\rm w}(t_n),$ where $\hat{f}_*(t') = \hat{f}_*(t_n)$ is drawn randomly from a uniform PDF while ensuring that $|f_*(t_{n+1}) - f_*(t_n)| \le \Delta f_{\text{drift}}$ is satisfied, and $f_{*}(t')$, $f_{*}(t')$, and $\Phi_{w}(t')$ are continuous from one HMM step to the next. This prescription is neither unique nor necessarily optimal; it is one of many, equally valid approaches. We assume for simplicity that there is no secular frequency drift, i.e., $\langle \dot{f}_*(t) \rangle = 0$. Incorporating $\langle f_*(t) \rangle \neq 0$ is straightforward; it is already part of LAL implementations of the \mathcal{F} -statistic, for example. However it is unnecessary in many astrophysical settings, because the HMM with the transition probabilities defined in Sec. III and Fig. 1 automatically handles secular spin evolution with $\langle |\dot{f}_*(t)| \rangle \lesssim \Delta f_{\text{drift}}/T_{\text{drift}}$ as a matter of course.

The stepwise evolution of \dot{f}_* differs deliberately from the stepwise evolution of \dot{f}_* modeled by the fluctuating torque $\xi(t)$ in (6), which underpins the transition probabilities in Sec. III and Appendix C. In general we do not know the functional form of the spin wandering in astrophysical sources. [7] Hence it is prudent to assume different forms of wandering in the test injections and transition probabilities, to double-check the robustness of the algorithm. The injection parameters are quoted in Table II and are the same as those in Ref. [18] to facilitate comparison, except that in this paper $f_*(t_0)$ and $\Phi_*(t_0)$ are chosen randomly (from uniform PDFs covering the ranges in Table II) as a self-blinding precaution. The synthetic data are generated using *Makefakedata_v4* in the LAL.

TABLE II.	Injection	parameter	's used	to cre	ate the	synthetic
data analyzed	l in Secs.	and VI.	Differen	nt tests	employ	/ different
subsets of the	e ranges ii	n the first	two lin	nes.		

Parameter	Value	Units
$\overline{\Phi_*(t_0)}$	$[0, 2\pi]$	rad
$f_{*}(t_{0})$	[111.0, 111.1]	Hz
$\dot{f}_*(t_0)$	0	$Hz s^{-1}$
Ψ	4.08407	rad
cosi	0.71934	_
α	4.27570	rad
δ	-0.27297	rad
$S_h(f_0)^{1/2}$	4×10^{-24}	$Hz^{-1/2}$

B. Representative example

Figure 3 illustrates the output of versions I and III of the HMM for three typical, injected signals with $h_0/10^{-26} = 1.7, 1.3$, and 1.1. The strongest signal is

detectable by both versions of the HMM, the intermediate signal is detectable by version III only, and the weakest signal is detectable by neither version. The figure displays the frequency path that best matches the injected $f_{*}(t)$. When the signal is detected, the best-matching frequency path is also the optimal HMM path, i.e., the frequency component of $Q^*(O)$. The frequency is recovered accurately, with root mean square errors of $\varepsilon_{f_*} = 6.5 \times 10^{-7}$ Hz = $1.1\Delta f_{\text{drift}}$ and $\varepsilon_{f_*} = 5.9 \times 10^{-7}$ Hz = $1.0\Delta f_{\text{drift}}$ for version III in Figs. 3(a) and 3(b) respectively. Note that the injected $f_*(t)$ traces a piecewise-parabolic path, because $\ddot{f}_*(t)$ is piecewise-constant (see Sec. VA). In contrast, the frequency path recovered by version III of the HMM, which obeys the Ornstein-Uhlenbeck transition probabilities in Appendix C, is piecewise-constant in the figure, because the HMM jumps between discrete frequency bins of width Δf_{drift} .

Figure 4 displays the absolute error between the injected and recovered phase as a function of time for the three



FIG. 3. Sample tracking output from versions I (purple curves) and III (green curves) of the HMM for three injected signals (black curves) from an isolated neutron star with source parameters drawn from Table II and $h_0/10^{-26} = 1.7$ [panel (a)], 1.3 [panel (b)], and 1.1 [panel (c)]. The purple and green curves are the best-matching frequency paths (with minimum path-integrated, root-mean-square error ε_{f_*} ; see Sec. V E), centred on $f_*(t_{N_T})$ and plotted in units of Δf_{drift} ; they are not necessarily the optimal path $Q^*(O)$. The optimal path matches well [i.e., within two frequency bins of $f_*(t)$ for all t] for versions I and III in (a) and version III in (b). The optimal path matches poorly for version I in (b) and versions I and III in (c); indeed it lies outside the border of the plot. We plot instead the paths with minimum ε_{f_*} , viz. the sixth, 20-th, and 411-th Viterbi paths respectively, which lie within a few frequency bins of $f_*(t)$ purely by chance but are of no practical use in an astrophysical search. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$.



FIG. 4. Accuracy of HMM phase reconstruction. Absolute pointwise phase error (in rad) between the injected phase and the phase recovered by version III of the HMM for the three isolated sources in Fig. 3, plotted in the range $[-\pi, \pi]$ versus time (in units of 10 days). Panels (a), (b), and (c) correspond one-to-one to the panels in Fig. 3.

isolated sources studied in Fig. 3. Superficially the phase reconstruction in Fig. 4 looks worse than the corresponding frequency reconstruction in Fig. 3. The B-statistic concentrates the signal power into at most two adjacent frequency bins yet spreads it out over multiple phase bins. As seen in Fig. 2(a), $\mathcal{B}(f_0, \Phi_0)$ is nearly a delta function in frequency (like the \mathcal{F} - and \mathcal{J} -statistics in versions I and II of the HMM) but has full-width half-maximum $\approx \pi$ in phase. Nevertheless, although the phase tracking is imperfect, it delivers improved sensitivity on balance, if one compares Fig. 3(a) with Fig. 3(b) for example. This improvement does not occur simply because version III of the HMM uses the \mathcal{B} -statistic, which in its phase-maximized form is ≈ 5 per cent more sensitive than the \mathcal{F} -statistic (see Sec. IV C) [21]. To verify this, we repeat the tests in Figs. 3 and 4 while artificially scrambling the phase, i.e., randomizing $\Phi_*(t_n)$ at every HMM step while keeping $f_*(t)$ continuous as in Sec. VA. Phase randomization converts the version III detection of the injection with $h_0 = 1.3 \times 10^{-26}$ into a nondetection while having no effect on the version I results.

When the signal is not detected, $Q^*(O)$ is clearly wrong, e.g., $\varepsilon_{f_*} = 1.2 \times 10^{-2}$ Hz = $2.1 \times 10^4 \Delta f_{\text{drift}}$ for version III in Fig. 3(c). The agreement looks better in the figure but artificially so. The minimum- ε_{f_*} paths plotted in the figure turn out to be the 6th, 20th, and 411th Viterbi paths [i.e., not $Q^*(O)$] for the nondetections using version I in Fig. 3(b) and versions I and III in Fig. 3(c) respectively. While these do lie within a few frequency bins of $f_*(t)$ by chance, they are of no practical use in an astrophysical search, where the true $f_*(t)$ is unknown. The optimal path is plotted whenever possible in Fig. 3 but it always lies far outside the border of the plot, when the signal is not detected.

The PDFs of ln \mathcal{B} in pure noise and for a relatively strong injection are compared in Appendix E 1 for completeness. They do not follow a chi-squared distribution, unlike the \mathcal{F} -statistic, because marginalizing over ψ , cos ι , and h_0 in (28) implicitly enforces constraints between the amplitudes in (15), so that ln \mathcal{B} is not a sum of independent squares.

C. Detection strategy

We assess the performance of version III of the HMM within the Neyman-Pearson framework applied to other continuous wave search pipelines developed by the LIGO Scientific Collaboration. [1] Specifically, we generate receiver operating characteristic (ROC) curves for a range of h_0 and N_T values, generalizing the tests in Ref. [51] to

include the time-dependent antenna beam-pattern functions a(t) and b(t). The aims of the exercise are (i) to characterize the sensitivity given user-selected false alarm and false dismissal probabilities, denoted by $P_{\rm a}$ and $P_{\rm d}$ respectively; and (ii) to develop a practical recipe for how to subdivide the full data set (duration $T_{\rm obs}$) into N_T segments of duration $T_{\rm drift}$.

To generate a ROC curve, i.e., a graph of $1 - P_d$ versus P_a , we must first define precisely what a detection means. This is not trivial for HMM-based algorithms. In versions I and II of the HMM, the probability that a Viterbi path terminates in a particular frequency bin is correlated with the termination probability for the $2N_T$ nearest bins, because HMM paths terminating in neighboring bins share common subpaths in general. The problem worsens in version III of the HMM, where the tails of $p(t_{n+1}, f_{*j}, \Phi_{*k})$ in (11) extend outside the range $|f_{*j} - f_*(t_n)| \le \Delta f_{drift}$ and wrap through 2π in phase, as calculated in Appendix C (see also Sec. II C). For example, the chance of encountering a false alarm within $\sim N_T$ bins of another false alarm is higher than encountering it elsewhere.

Several valid ways exist to handle the above correlations. In this paper, we adopt the following approach. First, we divide the full search space into disjoint parcels of width $2N_T\Delta f_{\rm drift}$ in frequency and 2π in phase, which we call "blocks." Each block contains $2N_TN_{\Phi}$ frequency-phase bins. (We check that the results do not change significantly, if the frequency width of the blocks is $kN_T\Delta f_{\rm drift}$ with $k \gtrsim 2$, in Appendix E.) Starting with multiple realizations of pure noise (i.e., $h_0 = 0$), we calculate

$$S_{i} = \max_{|i'-i| \le N_{T}} \max_{0 \le \Phi_{*''} \le 2\pi} \ln \Pr[Q^{*}(O)|O; q^{*}(t_{N_{T}}) = (f_{*i'}, \Phi_{*''})]$$
(35)

in the block centered on the *i*th frequency bin. In (35), S_i is the HMM log likelihood for the optimal path $q^*(t)$ terminating at a given frequency-phase bin, $q^*(t_{N_T}) = (f_{*i'}, \Phi_{*''})$, maximized over all the frequency-phase bins in the block centered at frequency f_{*i} , with $|f_{*i'} - f_{*i}| \leq N_T \Delta f_{\text{drift}}$ and $0 \leq \Phi_{*''} \leq 2\pi$. We call S_i the "block score" and write it as *S* henceforth as shorthand.⁶ We then define a threshold $S_{\text{th}}(f)$, where *f* is the central frequency of the block, such that an analyst-selected fraction P_a of the realizations are false alarms, i.e., they return $S > S_{\text{th}}(f)$. (The dependence on *f* is weak.) We then repeat the exercise after injecting a signal $h_0 > 0$ into multiple noise realizations. A block with $S > S_{\text{th}}(f)$ is flagged as a candidate. If any subset of the frequency component of the injected path, $\{f_*(t_1), \ldots, f_*(t_n)\}$, overlaps with the block, the candidate counts as a successful detection; otherwise the candidate is a false alarm.⁷ We check below that the results do not change significantly, if we require a minimum of (say) half the injected path to overlap with the block. Conversely, a false dismissal occurs, when zero candidates overlap even partially with the one or two blocks containing the injected signal.⁸

Sample histograms of the block score *S* in (35) are presented in Appendix E 1 as a validation test. Noise-only and noise-plus-injection histograms are visibly separate, when the detection threshold is exceeded, demonstrating the discriminating power of the HMM. The PDFs of *S* and $\ln \beta$ have different functional forms, brought about by the maximization steps in the Viterbi algorithm and (35) [19].

Continuous wave searches are typically subdivided into subbands of width $\Delta f_{sub} \sim 1$ Hz (0.6 Hz in this paper). Subbands are a housekeeping device to handle the practicalities of data management (e.g., storage and input-output overhead on a compute cluster). They are not the same as blocks, which are logical units in the detection strategy above. It is therefore necessary to convert the block-based false alarm probability, P_a, to a subband-based false alarm probability, $P'_{\rm a}$, using the binomial theorem, viz. $P'_{\rm a} = 1 - (1 - P_{\rm a})^{N'}$ with $N' = \Delta f_{sub} / (2N_T \Delta f_{drift})$ [52]. Note that $S_{th}(f)$ is a slow function of f over ~ 1 Hz, so it is usually good enough to use its midpoint value across the whole subband [52,53]. The above approach mimics the one adopted in previous searches for Sco X - 1 with the sideband algorithm, where frequency bins are correlated over windows of width $(2M'+1)\Delta f_{\text{drift}}$, i.e., the width of the Bessel comb of orbital sidebands [20,52,53,60]. The threshold $S_{\text{th}}(f)$ is also a function of N_T , as discussed in Appendix E 2.

D. ROC curves

A key question for any detection algorithm is how the trade-off between P_a and P_d adjusts, as the SNR changes. To this end, we present ROC curves in Fig. 5 for $h_0/10^{-26} = 1.7$, 1.3, and 1.1, $S_h(f_0)^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$, $T_{\text{drift}} = 10 \text{ d}$, $N_T = 37$, and the source parameters in Table II, adhering to the detection strategy in Sec. V C. Results from versions III and I of the HMM are plotted as solid and dashed curves respectively. The version III curve for $h_0 = 1.7 \times 10^{-26}$ overlaps with the top border of the figure and is invisible. The version III curve for $h_0 = 1.3 \times 10^{-26}$ gives $P_d \approx 0.1$ for

⁶The block score does not equal the Viterbi score used in previous work, [16,19] e.g., Eqs. (29)–(31) in Ref. [19]. The latter quantity is defined as the number of standard deviations that $\ln \Pr[Q^*(O)|O; q^*(t_{N_T}) = q_i]$ in the *i*th bin stands away from the mean, where the mean and standard deviation are computed over the full search band (width *B*) for one realization.

⁷There is no advantage in also testing for phase overlap with $\{\Phi_*(t_1), ..., \Phi_*(t_n)\}$, because $\mathcal{B}(f_0, \Phi_0)$ is a broad function of Φ_{02} see Fig. 2.

 $[\]Phi_{0s}$ see Fig. 2. It is always possible that the highest *S* value in a block is a false alarm, while the second-highest (say) is a real signal, because nearby HMM paths are correlated. In practice it is imprudent to claim a detection in a genuine, astrophysical search under such circumstances; the pragmatic response is to wait for more data.



FIG. 5. Comparative HMM performance for an isolated source. ROC curves for $h_0/10^{-26} = 1.7$ (blue curve; hidden under top border), 1.3 (green curve), and 1.1 (purple curve) and the source parameters in Table II. The false alarm probability P_a and detection probability $1 - P_d$ are plotted on the horizontal and vertical axes respectively. Solid and dashed curves correspond to HMM versions III and I respectively with $T_{drift} = 10$ d and $N_T = 37$. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$. Realizations: 10^4 per curve.

 $P_{\rm a} = 10^{-2}$, a popular combination in published LIGO searches, e.g., Ref. [16]. In comparison, version I of the HMM achieves the same $(P_{\rm a}, P_{\rm d})$ combination for $h_0 \approx 2 \times 10^{-26}$, i.e., its sensitivity is ≈ 1.5 times lower [18]. Version III of the HMM is a fairly reliable detection algorithm even at low false alarm probabilities, with $P_{\rm d} < 0.4$ for $P_{\rm a} \ge 10^{-4}$. The detection probability for $P_{\rm a} = 10^{-2}$ drops below $1 - P_{\rm d} = 0.5$ for $h_0 \le 1.1 \times 10^{-26}$.

A practical task when applying the HMM is to estimate in advance, how its performance scales with the volume of data available, and how the data and parameter space should be subdivided to maximize performance. Appendix E2 quantifies how the detection probability scales with N_T under two practical scenarions: (i) T_{drift} is fixed, so that the volume of data increases, as N_T increases; and (ii) $T_{\rm obs}$ is fixed, so that a fixed volume of data is subdivided into more coherent segments, as N_T increases. In scenario (i), $1 - P_d$ rises monotonically with N_T , as expected. In scenario (ii), $1 - P_d$ peaks, when $T_{\rm obs}/N_T$ matches the characteristic timescale over which $f_*(t)$ fluctuates intrinsically, also as expected. The block score threshold S_{th} is calculated versus N_T for both scenarios. Appendix E3 checks for completeness, that the ROC curves are insensitive to how the blocks are partitioned. It is found that P_d changes by ≤ 3 per cent at fixed P_a (with $10^{-4} \le P_a \le 1$) for block bandwidths $2kN_T\Delta f_{\text{drift}}$ in the range $0.243 \le k \le 2.00$, independent of the absolute position of the leftmost bin in the block.

E. Accuracy

Previous numerical experiments with versions I and II of the HMM demonstrate that the tracking accuracy is bounded by the Nyquist criterion [18,19,51]. When an

injected signal is detected successfully, the root mean square error integrated along the path satisfies $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$, whereas one typically finds $\varepsilon_{f_*} \gg \Delta f_{\text{drift}}$ for false alarms. The representative examples in Fig. 3 suggest that this remains true for version III of the HMM, with $\varepsilon_{f_*}/\Delta f_{\text{drift}} = 1.0$ (detection), 1.1 (detection), and 2×10^4 (nondetection) for $h_0/10^{-26} = 1.7$, 1.3, and 1.1 respectively. Versions I and III are equally accurate in Fig. 3(a), for example, with $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$. The tendency for version III to dwell somewhat longer in certain frequency bins follows from $A_{q;q_i}$ in Fig. 1.

We quantify the tracking accuracy systematically through Fig. 6, which displays ε_{f_*} for the optimal path in the highest-ranked block against the block score S. Versions III and I of the HMM are displayed in Figs. 6(a) and 6(b) respectively. The plotted symbols, each corresponding to one realization, separate into two clusters: detections at the bottom right, with $S \gtrsim S_{\rm th}(f)$ and $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$, and nondetections at the top left, with $S \lesssim S_{\rm th}(f)$ and $\varepsilon_{f_*} \gg \Delta f_{\rm drift}$. A handful of points form a bridge between the clusters, because a few realizations produce false alarms with $S > S_{\text{th}}(f)$ but $\varepsilon_{f_*} \gg \Delta f_{\text{drift}}$, e.g., the point with $S \approx -2.4$ and $\varepsilon_{f_*} \approx 2.5 \times 10^{-5}$ Hz in Fig. 6(a). These accidents are expected; phase consistency sometimes happens by chance in the noise along a path with fortuitously high \mathcal{B} values. Occasionally the tracker achieves a good match with $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$ even for $S \lesssim S_{\rm th}(f)$, corresponding to a false dismissal in a real search. About 5 percent of the latter events occur accidentally, when the signal block happens to rank highest (out of 20 blocks in Fig. 6) due to features in the noise (even with $h_0 = 0$). Note that no threshold is applied explicitly in constructing Fig. 6, although implicitly $S_{\rm th}(f)$ falls near the value of S below which $\varepsilon_{f_*} \gg \Delta f_{\text{drift}}$ typically occurs.

The significant uncertainty in phase tracking, exemplified by Fig. 4, does not impair the accuracy of frequency tracking reported in Fig. 6, as discussed in Secs. II C and V B. However, it does circumscribe the astrophysical questions that can be answered. Knowing the phase evolution more accurately can help distinguish between astrophysical emission mechanisms, in situations where the frequency evolution is not informative enough. A timedomain version of the HMM offers one possible way to achieve better phase tracking, at the cost of stepping outside the well-tested frequency-domain software infrastructure in the LAL suite. Designing a time-domain HMM is a goal of future work.

VI. NEUTRON STAR IN A BINARY

We now repeat the tests in Sec. V for a neutron star in a binary system. The HMM structure and search procedure remain unchanged, except that $F_{1a}(f_0)$ and $F_{1b}(f_0)$ are replaced by $J_{1a}(f_0)$ and $J_{1b}(f_0)$ respectively in the \mathcal{B} -statistic via (28)–(34). Appendix E4 verifies that this



FIG. 6. Tracking accuracy of the HMM. Root mean square frequency error ε_{f_*} (left axis in units of Hz, right axis in units of $\Delta f_{drift} = 5.8 \times 10^{-7}$ Hz) integrated along the optimal path in the highest-ranked block versus the block score *S*. (a) Version III (red, open circles) with $N_T = 37$ and $T_{drift} = 10$ d. (b) Version I (blue, filled circles) with $N_T = 37$ and $T_{drift} = 10$ d. Realizations: 3×10^2 per panel. Each realization comprises 20 contiguous, 37-bin blocks, one of which contains an injected signal with $h_0 = 1.3 \times 10^{-26}$ and the source parameters in Table II. The block scores in (a) and (b) should not be compared as they arise from different statistics (\mathcal{B} and \mathcal{F} respectively).

replacement leads to minimal loss of signal power; the Doppler sidebands collapse into a single frequency bin without discernible leakage into neighboring bins, just like for the \mathcal{J} -statistic. In Sec. VI A we present tracking results for a representative sample of synthetic data. ROC curves are discussed in Sec. VI B. In Sec. VI C we plot the \mathcal{B} -statistic as a function of the orbital parameters a_0 and ϕ_a , in order to inform the gridding strategy for future searches, e.g., for LMXBs. Versions II and III of the HMM are compared at each stage.

A. Representative example

The signal phase corresponding to a binary neutron star is given by (16) with $a_0 \neq 0$. Figure 7 illustrates the output of versions II and III of the HMM for three injected signals of the above form with the same h_0 values as in Sec. VA, viz. $h_0/10^{-26} = 1.7$, 1.3, and 1.1. The parameters of the binary orbit are quoted in Table III, with a_0 and ϕ_a set at the midpoints of their ranges. The stochastic component of the injected phase, $\Phi_w(t)$, evolves according to the algorithm in Sec. VA.

The results in Fig. 7 resemble those in Fig. 3. Both HMM versions detect the strongest signal, but only version III detects the intermediate signal. Neither detects the weakest signal. Version III is \approx 1.4 times more sensitive than version II, and its sensitivity is approximately the same for isolated and binary sources.⁹ Once the HMM fails to detect a signal, the optimal Viterbi path stands many bins away from the injected path and normally falls outside the plotted region. The agreement in Figs. 7(b) and 7(c) looks better than it

actually is, because we plot the minimum- ε_{f_*} paths, which turn out to be the 2nd, 2nd, and 408th Viterbi paths for the nondetections using version II in Fig. 7(b) and versions II and III in Fig. 7(c) respectively. Such coincidental successes are useless in an astrophysical search, where the true $f_*(t)$ is unknown. Similarly, it may seem that version II outperforms version III on the $h_0 = 1.1 \times 10^{-26}$ injection, because the minimum- ε_{f_*} paths are the 2nd (version II) versus the 408th (version III). Again this is misleading: paths other than the first are not ranked consistently by the Viterbi algorithm, and besides version III has 32 times more paths than version II (and a different bin numbering system) because it tracks both f_* and Φ_* .

The phase component of $Q^*(O)$ is discussed briefly for completeness in Appendix F.

B. ROC curves

In order to characterize the sensitivity of the HMM systematically, we compute ROC curves for the same three signal amplitudes in Fig. 7, viz. $h_0/10^{-26} = 1.7, 1.3,$ and 1.1. The results are plotted in Fig. 8, where solid and dashed curves correspond to versions III and II of the HMM respectively. In the regime of practical interest, viz. $5 \times 10^{-3} \le P_a \le 2 \times 10^{-1}$, version III of the HMM delivers a detection probability ≈ 0.05 higher than version II at the same $P_{\rm a}$, a significant advantage when operating near the detection limit. Replacing $F_{1a}(f_0)$ and $F_{1b}(f_0)$ with $J_{1a}(f_0)$ and $J_{1b}(f_0)$ in the *B*-statistic leads to similar tracking performance for isolated and binary sources, although there is some modest loss of sensitivity in the latter case. For example, a detection probability of ≈ 0.75 is achieved in Fig. 8 for a binary source with $h_0 = 1.3 \times 10^{-26}$, given $P_{\rm a} = 10^{-2}$, compared to ≈ 0.90 for an isolated source

⁹This is consistent with previous work: Version II of the HMM is sensitive down to the same h_0 value, $h_0 \approx 2 \times 10^{-26}$, for a binary source as version I is for an isolated source.



FIG. 7. Sample tracking output from versions II (purple curves) and III (green curves) of the HMM for three injected signals (black curves) from a source in a binary with parameters drawn from Tables II and III with $h_0/10^{-26} = 1.7$ [panel (a)], 1.3 [panel (b)], and 1.1 [panel (c)], plotted on the same axes as in Fig. 3. The purple and green curves are the best-matching frequency paths (with minimum ε_{f_*}); they are not necessarily the optimal path $Q^*(O)$. The optimal path matches well [i.e., within a few frequency bins of $f_*(t)$ for all t] for versions II and III in (a) and version III in (b). The optimal path matches poorly for version II in (b) and versions II and III in (c); indeed it lies outside the border of the plot. We plot instead the paths with minimum ε_{f_*} , viz. the 2nd, 2nd, and 408th Viterbi paths respectively, which lie within a few frequency bins of $f_*(t)$ purely by chance but are of no practical use in an astrophysical search. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$.

with the same h_0 in Fig. 5. [16] This is because the Jacobi-Anger decomposition (21) and (22) accounts for the binary motion imperfectly when combined with the \mathcal{B} -statistic, due to some covariance between the orbital and carrier phases in the orbital sidebands. For $h_0 \ge 1.7 \times 10^{-26}$, the performance is almost identical, as in Appendix E 4.

Monte Carlo simulations confirm that the performance of the HMM as a function of N_T for T_{drift} or T_{obs} fixed is the

TABLE III. Orbital parameters used to create the synthetic data for the binary sources analyzed in Sec. VI.

Parameter	Value	Units	Description
Р	68023.7	s	Orbital period
a_0	[1.26,1.62]	lt-s	Projected orbital semimajor axis
ϕ_{a}	$[0, 2\pi]$		Reference orbital phase
e	0.0	_	Orbital eccentricity

same as in the case of isolated sources (see Figs. 12 and 13 respectively in Appendix E). The results are not plotted to avoid repetition.

C. Sensitivity to orbital parameters

Electromagnetic observations normally supply prior constraints on LMXB orbital parameters. [12,61–63] For many objects, including Sco X – 1, the electromagnetic measurement of *P* through high-resolution optical spectroscopy is accurate enough, that a search over *P* is unnecessary. In contrast, searches over a_0 and ϕ_a are usually required [64].

Figure 9 displays $\ln \Pr[Q^*(O)|O]$ for version III of the HMM as a function of a_0 and $T_{\rm asc} = \phi_{\rm a} P/(2\pi) + {\rm constant}$, where $T_{\rm asc}$ is the time of ascending node. The log probability is evaluated at the true, injected value of f_* and maximized with respect to Φ_* , for a strong signal with $h_0 = 8 \times 10^{-26}$ tracked over $N_T = 37$ steps. Starting from



FIG. 8. Comparative HMM performance for a source in a binary. ROC curves for sources with $h_0/10^{-26} = 1.7$ (blue curve), 1.3 (green curve), and 1.1 (purple curve) and the source parameters in Tables II and III. The false alarm probability P_a and detection probability $1 - P_d$ are plotted on the horizontal and vertical axes respectively. Solid and dashed curves correspond to HMM versions III and II respectively. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$. Realizations: 10^4 per curve.



FIG. 9. HMM performance as a function of binary orbital elements. Log probability $\ln \Pr[Q^*(O)|O]$ versus a_0 and T_{asc} for a strong (SNR \gg 1) binary-star signal with constant f_{\star} observed during $N_T = 37$ 10-day segments. (Top left) Contours of $\ln \Pr[Q^*(O)|O]$ on the $T_{asc} - a_0$ plane, centered on the injected values T_{asc}^{true} and a_0^{true} . Hot colors stand for the highest log probabilities. (Top right) Cross section through the peak of $\ln \Pr[Q^*(O)|O]$ versus a_0 for $T_{asc} = T_{asc}^{true}$. (Bottom left) Cross section through the peak of $\ln \Pr[Q^*(O)|O]$ versus a_0 for $T_{asc} = T_{asc}^{true}$. (Bottom left) Cross section through the peak of $\ln \Pr[Q^*(O)|O]$ versus T_{asc} for $a_0 = a_0^{true}$. (Bottom right) Zoomed out version of the top left panel. Injection parameters: $h_0 = 8 \times 10^{-26}$, $f_{\star} = 111.1$ Hz, $a_0^{true} = 1.44$ lts, and $T_{asc}^{true} = 897753994$ s (arbitrary orbital phase), characteristic of Scorpius X – 1; see also Tables II and III.

the panel at the bottom right of the figure, we observe that $\ln \Pr[Q^*(O)|O]$ peaks strongly around the true, injected orbital elements a_0^{true} and $T_{\text{asc}}^{\text{true}}$. The top left panel zooms into the peak (note the magnified scale) and shows that it is encircled by "ripples" reminiscent of a diffraction pattern. The ripples are visible more clearly in the cross sections at

 $T_{\rm asc} = T_{\rm asc}^{\rm true}$ and $a_0 = a_0^{\rm true}$, graphed in the top right and bottom left panels respectively. Both cross sections are sinc-like, except that the nodes do not touch zero; $\Pr[Q^*(O)|O]$ is positive definite. Qualitatively the features in Fig. 9 match those observed in Fig. 4 in Ref. [19] for the \mathcal{J} -statistic HMM (version II), although the scales are not comparable of course.

In practice, in a search with real data, the grid spacings in a_0 and $T_{\rm asc}$ are set according to a parameter space metric and depend on the search frequency f_0 . [64] For example, the LIGO O2 search for Sco X - 1 with HMM version II employs 768 a_0 bins of width 2.3×10^{-3} lt s [with $1.45 \le a_0/(1 \text{ lt s}) \le 3.25$] at $f_0 = 60 \text{ Hz}$, compared to 8227 a_0 bins of width 2.2×10^{-4} lts at $f_0 = 650$ Hz. [17,63] The resolution is chosen to yield a mismatch of $\leq 10\%$ in the squared SNR, as defined by Eq. (5) in Ref. [64], the worst case being when the signal straddles the boundary between two bins. Without being comparable directly, the above approach is consistent with Fig. 9: the squared SNR is of the same order as $\ln \Pr[Q^*(O)|O]$, and $\ln \Pr[Q^*(O)|O]$ drops off by $\leq 10\%$ from its peak for $|a_0 - a_0^{\text{true}}| \lesssim 10^{-3}$ lt s in the top right panel of Fig. 9 and for $|T_{\rm asc} - T_{\rm asc}^{\rm true}| \lesssim 5$ s in the bottom left panel of Fig. 9. Convenient formulas for the number of a_0 and T_{asc} templates in terms of the desired mismatch are given in Sec. V of Ref. [64].

VII. SCO X – 1 MDC: A REALISTIC EXAMPLE

A. Synthetic data

The Sco X - 1 MDC is a project to compare systematically the performance of published continuous-wave search pipelines on a level playing field under simulated Advanced LIGO conditions [20]. The MDC predates HMM versions I and II. It evaluates the relative proficiency of five pipelines against criteria including sensitivity, computational cost, and accuracy in parameter estimation. The pipelines are based on the CrossCorr [65–67], TwoSpect [68,69], Radiometer [70–72], Sideband [52,60], and Polynomial [73] algorithms. Method papers describing each algorithm are cited in the previous sentence. Since the MDC was published, two of the pipelines have completed searches using Advanced LIGO data from O1 and O2 [3,74], as have HMM versions I and II [16,17]. Two other pipelines have completed searches using Initial LIGO data from Science Run 6 (S6) [53,75]. It should be noted that O1 and O2 do not achieve Advanced LIGO's design sensitivity, approximated in the MDC as $S_h(f_0)^{1/2} \approx 4 \times 10^{-24} \text{ Hz}^{-1/2}$ (Gaussian recolored).

The MDC enables an important check on the results in previous sections under realistic yet controlled conditions on a dataset generated by an independent party. Of course, the MDC is no longer closed, as it was in its original incarnation; the TwoSpect, Radiometer, Sideband, and Polynomial pipelines competed blindly in Ref. [20], before the injection parameters were revealed, and the CrossCorr pipeline analyzed the data in self-blinded mode, after the injection parameters were revealed. In this paper we preserve the etiquette of a self-blinded analysis but note in fairness that some of the authors participated in previous analyses of the same data with HMM versions I and II [18,19]. We also note that $f_{*}(t)$ does not wander for any of the injected signals, even though we allow for wandering in the HMM transition probabilities. Extensive testing in previous published work demonstrates, that the HMM delivers equal sensitivity, whether $f_*(t)$ wanders or not, as long as T_{drift} satisfies (B1), [18,19] in line with theoretical expectations [14]. Strictly speaking, however, the analysis in this section checks the sensitivity and accuracy of version III of the HMM; it does not test its robustness to spin wandering. (Indeed nor did the original MDC study with the five pipelines in Ref. [20].) A future incarnation of the MDC including spin wandering, drawing on the analysis in Ref. [7], is currently being prepared and should be encouraged.

The parameters of the 50 injected signals in Stage I (version 6) of the MDC are listed in Table III in Ref. [20]. They are designed to resemble Sco X – 1, with $0.050 \le 2f_*/(1 \text{ kHz}) \le 1.5$ and orbital elements similar to those measured electromagnetically [61–63]. Since the original MDC release, the data for three injections, with indexes 65, 66, and 75 in Ref. [20], are no longer accessible due to human error. They are omitted from the analysis below, which is restricted to 47 injections.

B. Search procedure

The analysis is conducted as follows in order to copy approximately some of the steps in a search with real LIGO data.

- (1) Starting from $f_0 = 50$ Hz and defining subbands in increments of 0.1 Hz, we identify the subband containing the injected signal. The partition is similar to the O2 Sco X – 1 search with version II of the HMM, which implemented 0.6-Hz subbands, without f_* being known of course. [17] In effect this step is self-blinded to a good approximation, because there are $(0.1 \text{ Hz})/\Delta f_{\text{drift}} = 1.7 \times 10^5$ frequency bins in the subband, any single one of which can contain the injected signal in principle.
- (2) An orbital grid is laid out in a_0 and $T_{\rm asc}$ as for the HMM O2 Sco X 1 search. The grid spacings in a_0 and $T_{\rm asc}$ are given by $1.2 \times 10^{-4} (f_0/0.3 \text{ kHz})^{-1}$ lt s and $0.89(f_0/0.3 \text{ kHz})^{-1}(a_0/1.44 \text{ lt s})^{-1}$ s within the electromagnetic priors $1.45 \le a_0/(1 \text{ lt s}) \le 3.25$ and $1164543014 \le T_{\rm asc}/(1 \text{ s}) \le 1164543614$ respectively [17,63]. The grid spacings are one quarter of what is predicted by the parameter space metric via Eqs. (70) and (71) in Ref. [64], assuming a squared-SNR mismatch of $\le 10\%$. The safety factor 1/4 is discussed further below. Strictly speaking the

grid spacing varies from one f_0 bin to the next, but in practice it is kept uniform within each 0.1-Hz subband, substituting the subband midpoint into the above formulas as a good approximation.

- (3) A grid is also laid out in orbital period P, with grid spacing $1.0(f_0/0.3 \text{ kHz})^{-1}(a_0/1.44 \text{ lt s})^{-1} \text{ s involv-}$ ing the same safety factor 1/4 from step 2 above, based on Eqs (70) and (71) in Ref. [64] in the regime $P \ll T_{\text{drift}}$. This is a new step. Some of the MDC injections are not exactly at P = 68023.7 s, the central value returned by electromagnetic observations, [61–63] although they are close to it. Previous MDC analyses ignore the slight mismatch, motivated by the parameter space metric which implies that one P template is sufficient, because the experimental uncertainty $(\pm 0.04 \text{ s})$ is less than the metric-based resolution ≈ 0.2 s. [19,64] They search P = 68023.7 s only and are still successful; for example, version II of the HMM finds all 50 injections thus. [19] However version III of the HMM, which is more sensitive to weaker signals, also depends more sensitively on P.
- (4) Version III of the HMM is executed on 4 × 4 × 4 adjacent triples (a₀, T_{asc}, P) centered on the injection. (The MDC analysis is executed on a subset of the grid for computational economy; in an astrophysical search, we scan the whole grid.) Each triple (a₀, T_{asc}, P) is accompanied by an f₀ scan divided into (0.1 Hz)/Δf_{drift}/N_T = 4671 blocks as described in Sec. V C. The highest log probability among these 64 × 4671 (f₀, a₀, T_{asc}, P) combinations becomes the block score according to (35).
- (5) The root mean square frequency error ε_{f_*} is calculated along the optimal, wandering Viterbi track as in Sec. V E. Absolute, signed errors ε_{a_0} and $\varepsilon_{T_{asc}}$ are also calculated for a_0 and T_{asc} respectively as the injected minus recovered values for the optimal Viterbi track. This approach is adopted deliberately to stay consistent with previous MDC analyses, which verify the accuracy of the top candidate in a block instead of quantifying the false alarm probability. [18,20] In a search with real data, one would instead compare the block score with a threshold set by P_a and follow up any candidates through a veto procedure [16,17].

C. Signal detectability

The results of analyzing the MDC data with version III of the HMM are presented in Table IV. Each line of the table corresponds to one injection, indexed as in Ref. [20] (first column). The injection parameters f_* , a_0 , and T_{asc} are quoted along with the respective errors ε_{f_*} , ε_{a_0} , and $\varepsilon_{T_{asc}}$ in the parameter values recovered by the HMM. Two simulated interferometers (H1 and L1) are employed, chiefly to preserve consistency with the previous MDC analysis involving version II of the HMM [19]. The data start at GPS time 1230338490 and are divided into $N_T = 37$ segments with $T_{\text{drift}} = 10$ d.

Version III detects 47 out of 47 available injections. The outcome is reassuring but not surprising. Version II also detects every signal, and version III is ≈ 1.5 times more

TABLE IV. Results of tracking the 47 available injections in the Sco X - 1 MDC, sorted by index from Ref. [20], using version III of the HMM to track phase and frequency.

Index	$h_0(10^{-25})$	$h_0^{\rm eff}(10^{-25})$	f_* (Hz)	$\varepsilon_{f_{\star}}$ (Hz)	<i>a</i> ⁰ (s)	ε_{a_0} (s)	$T_{\rm asc}$ (s)	$\varepsilon_{T_{\rm asc}}$ (s)
1	4.160	2.706	54.498391348174	4.342×10^{-7}	1.37952	-9.518×10^{-4}	1245967666.02	-11.12
2	4.044	2.511	64.411966012332	4.229×10^{-7}	1.76461	4.803×10^{-4}	1245967592.98	-5.27
3	3.565	3.463	73.795580913582	6.836×10^{-7}	1.53460	-1.585×10^{-3}	1245967461.35	-5.66
5	1.250	1.154	93.909518008164	5.104×10^{-7}	1.52018	-8.158×10^{-5}	1245966927.93	2.22
11	3.089	1.399	154.916883586097	3.464×10^{-7}	1.39229	4.297×10^{-5}	1245967559.97	2.67
14	2.044	1.286	183.974917468730	3.553×10^{-7}	1.50970	-7.066×10^{-4}	1245967551.05	-3.63
15	11.764	4.169	191.580343388804	3.612×10^{-7}	1.51814	-4.484×10^{-4}	1245967298.45	0.10
17	3.473	1.253	213.232194220000	2.244×10^{-7}	1.31021	-7.427×10^{-5}	1245967522.54	1.74
19	6.031	2.437	233.432565653291	3.189×10^{-7}	1.23123	-1.060×10^{-4}	1245967331.14	1.27
20	9.710	3.434	244.534697522529	3.941×10^{-7}	1.28442	-4.418×10^{-4}	1245967110.97	-1.10
21	1.815	0.792	254.415047846878	5.561×10^{-7}	1.07219	7.354×10^{-5}	1245967346.40	-1.24
23	2.968	1.677	271.739907539784	3.922×10^{-7}	1.44287	-2.731×10^{-4}	1245967302.29	-2.22
26	1.419	1.172	300.590450155009	3.342×10^{-7}	1.25869	-1.721×10^{-4}	1245967177.47	-1.87
29	4.275	3.131	330.590357652653	4.893×10^{-7}	1.33070	-6.673×10^{-5}	1245967520.83	-0.84
32	10.038	4.391	362.990820993568	1.870×10^{-7}	1.61109	-2.790×10^{-4}	1245967585.56	0.24
35	16.402	9.183	394.685589797695	3.466×10^{-7}	1.31376	-1.059×10^{-4}	1245967198.05	1.75
36	3.864	1.539	402.721233789014	5.075×10^{-7}	1.25484	-6.642×10^{-5}	1245967251.35	0.79
41	1.562	0.746	454.865249156175	2.651×10^{-7}	1.46578	-1.896×10^{-4}	1245967225.75	0.36
44	2.237	1.996	483.519617972096	8.346×10^{-8}	1.55221	-1.446×10^{-4}	1245967397.86	0.13
47	4.883	1.992	514.568399601819	2.824×10^{-7}	1.14020	-1.637×10^{-4}	1245967686.81	0.33
48	1.813	0.745	520.177348201609	6.614×10^{-7}	1.33669	-3.329×10^{-5}	1245967675.30	0.15
50	1.093	1.027	542.952477491471	5.178×10^{-7}	1.11915	-2.302×10^{-4}	1245967927.48	-1.47
51	9.146	3.372	552.120598886904	6.501×10^{-7}	1.32783	6.253×10^{-5}	1245967589.54	-0.94
52	2.786	1.550	560.755048768919	4.209×10^{-7}	1.79214	-6.193×10^{-5}	1245967377.20	0.61
54	1.518	1.256	593.663030872532	5.792×10^{-7}	1.61276	-3.115×10^{-5}	1245967624.53	0.30
57	1.577	0.788	622.605388362863	5.260×10^{-7}	1.51329	-5.596×10^{-5}	1245967203.21	-1.00
58	3.416	1.287	641.491604906276	6.158×10^{-7}	1.58443	-1.418×10^{-4}	1245967257.74	0.16
59	8.835	4.981	650.344230698489	7.830×10^{-7}	1.67711	-1.422×10^{-4}	1245967829.90	-0.69
60	2.961	2.467	664.611446618250	7.197×10^{-7}	1.58262	5.343×10^{-5}	1245967612.31	-0.41
61	6.064	2.158	674.711567789201	4.978×10^{-7}	1.49937	-1.037×10^{-4}	1245967003.32	-0.01
62	10.737	3.853	683.436210983289	8.223×10^{-7}	1.26951	-4.060×10^{-5}	1245967453.97	-0.00
63	1.119	0.745	690.534687981171	6.762×10^{-7}	1.51824	-3.958×10^{-5}	1245967419.39	-0.18
64	1.600	0.570	700.866836291234	5.143×10^{-7}	1.39993	-6.909×10^{-5}	1245967596.12	-0.96
67	4.580	1.623	744.255707971300	3.620×10^{-7}	1.67774	-1.551×10^{-4}	1245967084.30	0.27
68	3.696	1.844	754.435956775916	4.000×10^{-7}	1.41389	-8.960×10^{-5}	1245967538.70	0.38
69	2.889	1.053	761.538797037770	3.693×10^{-7}	1.62613	-1.239×10^{-4}	1245966821.55	0.03
71	2.923	1.232	804.231717847467	3.238×10^{-7}	1.65203	8.338×10^{-6}	1245967156.55	0.30
72	1.248	0.792	812.280741438401	4.597×10^{-7}	1.19649	-1.325×10^{-4}	1245967159.08	0.87
73	2.444	0.936	824.988633484129	9.533×10^{-7}	1.41715	-6.960×10^{-5}	1245967876.83	0.82
76	3.260	1.725	882.747979842807	4.813×10^{-7}	1.46249	-8.305×10^{-5}	1245966753.24	-0.17
79	4.681	1.656	931.006000308958	2.697×10^{-7}	1.49171	-7.243×10^{-5}	1245967290.06	0.14
83	5.925	2.186	1081.398956458276	7.176×10^{-7}	1.19854	-3.862×10^{-5}	1245967313.93	-1.02
84	11.609	7.184	1100.906018344283	7.529×10^{-7}	1.58972	-6.257×10^{-6}	1245967204.15	-0.35
85	4.553	1.633	1111.576831848269	8.018×10^{-7}	1.34479	-9.497×10^{-5}	1245967049.35	-0.90
90	0.684	0.618	1193.191890630547	4.053×10^{-7}	1.57513	-7.212×10^{-5}	1245966914.27	-0.21
95	4.293	3.059	1324.567365220908	5.198×10^{-7}	1.59169	-1.443×10^{-5}	1245967424.76	0.53
98	5.404	1.948	1372.042154535880	7.448×10^{-7}	1.31510	-7.340×10^{-5}	1245966869.92	-0.34

sensitive than version II according to the results in Secs. V and VI. The signal amplitudes are quoted in the second and third columns of Table IV in terms of h_0 and $h_0^{\text{eff}} =$ $h_0 2^{-1/2} \{ [(1 + \cos^2 i)/2]^2 + \cos^2 i \}^{1/2}$ respectively. The source inclination influences detectability through the relative weighting of the plus and cross polarizations, and h_0^{eff} serves as an amplitude proxy which normalizes for this effect, as verified in Ref. [19] (see the tests in Sec. VA and Fig. 5 of the latter reference). The quietest detected signals from the h_0 and h_0^{eff} perspectives are injections 90 ($h_0 = 6.8 \times 10^{-26}$) and 64 ($h_0^{\text{ref}} = 5.7 \times 10^{-26}$) respectively. Both lie well above the version III sensitivity limit $h_0 \ge 1.3 \times 10^{-26}$ established in Secs. V and VI. The conclusions are not affected by the absence of injections 65, 66, and 75, which are all relatively strong $[7.7 \le h_0/(10^{-25}) \le 9.3]$ and are detected easily by versions I and II with two interferometers.

D. Accuracy

Although versions II and III both detect all the injections, version III recovers the true signal parameters more accurately. The fifth column of Table IV indicates that version III recovers $f_*(t)$ with a root mean square error across $N_T = 37$ segments of $\varepsilon_{f_*} \leq 9.5 \times 10^{-7}$ Hz $\leq 2\Delta f_{\text{drift}}$. (Recall that the optimal Viterbi track is free to wander, whereas the injections are stationary.) Indeed 33 out of 47 injections are recovered with $\varepsilon_{f_*} \leq \Delta f_{\text{drift}}$. Essentially parameter estimation is limited by the spectral resolution. In contrast, version II of the HMM recovers 27 out of 50 injections with $\varepsilon_{f_*} \approx P^{-1} \gg \Delta f_{\text{drift}}$, much worse than the spectral resolution, viz. $1 \le \varepsilon_{f_*}/(10^{-5} \text{ Hz}) \le 2$; see Table IV in Ref. [19]. The step up from $\varepsilon_{f_*} \sim \Delta f_{\text{drift}}$ to $\varepsilon_{f_{-}} \sim P^{-1}$ occurs, because version II sometimes converges on the orbital sidebands $f_* \pm P^{-1}$, whereas version III always converges on the central peak f_* for the MDC injections. Interestingly, no strong correlation is found between ε_{f_*} and $h_0^{\rm eff}$ with version III. Once the HMM detects a signal, $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$ is grid-limited and essentially random. A similar lack of correlation is observed for version II [19].

Version III is also more accurate than version II when recovering the orbital elements. The seventh column of Table IV indicates that version III recovers a_0 with an absolute error of $|\varepsilon_{a_0}| \le 1.6 \times 10^{-3}$ lt s. This amounts to $\lesssim 5$ times the grid resolution, which decreases $\propto f_0^{-1}$ from 6.6×10^{-4} lt s at $f_0 = 54.5$ Hz (injection 1) to 2.6×10^{-5} lts at $f_0 = 1.37$ kHz (injection 98). Although the maximum value of $|\varepsilon_{a_0}|$ is comparable for versions II and III, version III recovers 26 out of 47 injections with $|\varepsilon_{a_0}| \le 1 \times 10^{-4}$ lt s, whereas version II only recovers eight out of 50 injections with $|\varepsilon_{a_0}| \le 1 \times 10^{-4}$ lt s. Interestingly version III underestimates a_0 41 out of 47 times. It is currently unclear why this happens, and more tests are needed to explore the behavior and check if it is a statistical fluctuation.

The ninth column of Table IV indicates that version III recovers $T_{\rm asc}$ with an absolute error of $|\varepsilon_{T_{\rm asc}}| \leq 11$ s, i.e., $\lesssim 5$ times the grid resolution, which decreases $\propto f_0^{-1}a_0^{-1}$ from ≈ 5 s at $f_0 = 54.5$ Hz (injection 1) to ≈ 0.2 s at $f_0 = 1.37$ kHz (injection 98). The $T_{\rm asc}$ estimates compare favorably with the orbital phase errors $|\varepsilon_{\phi_a}| = 2\pi |\varepsilon_{T_{\rm asc}}|/P \leq 1.0 \times 10^{-3}$ yielded by version II. The maximum ϕ_a error is comparable in versions II and III, but version III recovers 21 out of 47 injections with $|\varepsilon_{T_{\rm asc}}| \leq 0.5$ s, whereas version II recovers only five out of 50 injections with $|\varepsilon_{T_{\rm asc}}| \leq 0.5$ s. The $T_{\rm asc}$ results parallel the behavior observed in ε_{a_0} .

Accuracy of parameter estimation is a better diagnostic for illustrating the superiority of version III in the MDC context than (say) the minimum number of segments required to detect a signal. Version II detects 43 out of 50 injections with $N_T = 1$ and the remaining seven with $N_T \leq 13$ ($T_{\text{drift}} = 10$ d) [19]. There is not much room for version III to outperform against this measure but for the record it does: it detects every injection except the two weakest (indexes 64 and 90) with $N_T = 1$.

VIII. CONCLUSIONS

A HMM coupled with a stepwise matched filter provides an efficient, semicoherent way to detect and track the unknown signal frequency of a quasimonochromatic, continuous gravitational wave source with spin wandering driven by internal processes (isolated source) or accretion (binary source). In previous work HMMs have searched for the LMXB Sco X-1 in LIGO O1 and O2 data using frequency domain, maximum likelihood matched filters: the Bessel-weighted \mathcal{F} -statistic (version I), which does not track orbital phase, and the Jacobi-Anger \mathcal{J} -statistic (version II), which does. Here we generalize existing HMM pipelines to track rotational phase as well as orbital phase (version III). In the emission probability, the \mathcal{J} -statistic is replaced by a phase-sensitive version of the Bayesian \mathcal{B} -statistic introduced for loosely coherent searches. The data are input as SFTs, leveraging the well-tested software infrastructure in the LAL. In the transition probability, the intra-step spin wandering is modeled according to a phasewrapped Ornstein-Uhlenbeck process. A recipe for choosing the Ornstein-Uhlenbeck control parameters, γ and σ , is given in Sec. III B. A revised detection strategy based on block scores is described in Sec. VC.

The sensitivity of version III of the HMM is quantified. The ROC curves in Secs. V D and VIB give $P_d \ge 0.9$ (isolated source) and $P_d \ge 0.75$ (binary source), when the characteristic wave strain satisfies $h_0 \ge 1.3 \times 10^{-26}$, with $P_a = 10^{-2}$. Hence version III is ≈ 1.5 times more sensitive than version II. The requirement of phase continuity from one HMM step to the next lowers P_a at fixed h_0 and increases $1 - P_d$ at fixed P_a . Performance is optimized, when $T_{\rm drift}$ matches the source's spin wandering timescale. The results depend weakly on γ , σ , the block width when calculating the block score, and the location of the block boundary.

The tracking accuracy is quantified in Secs. VE and VIC. It is found that the root mean square frequency error is bounded spectrally and is therefore near-optimal, with $\varepsilon_{f_*} \lesssim \Delta f_{\text{drift}}$ when an injected signal is detected successfully and $\varepsilon_{f_*} \gg \Delta f_{\text{drift}}$ otherwise. The absolute errors in the orbital elements are limited to $\lesssim 5$ times the grid resolution in a_0 and ϕ_a (or equivalently T_{asc}) set by the parameter space metric. [64] The HMM log probability peaks unimodally at the correct value in the $a_0 - T_{asc}$ plane, with a sinc-like cross section (see Fig. 9). The accuracy is confirmed by the performance of version III of the HMM in the Sco X - 1 MDC (in self-blinded mode). It finds 47 out of 47 injections currently available (out of 50 originally) with $N_T = 37$, $T_{\text{drift}} = 10$ d, and two simulated interferometers, achieving accuracies of $|\varepsilon_{f_*}| \leq$ 9.5×10^{-7} Hz, $|\varepsilon_{a_0}| \le 1.6 \times 10^{-3}$ lt s, and $|\varepsilon_{T_{asc}}| \le 11$ s. Version III is less prone to converging on the sidebands $f_* \pm P^{-1}$ and is systematically more accurate, e.g., it recovers 26 out of 47 injections with $|\varepsilon_{a_0}| \le 1 \times 10^{-4}$ lts, whereas version II only achieves such accuracy eight times out of 50. The gridding strategy adopted here, which is to implement conservatively the parameter space metric in Ref. [64] as described in Sec. VII B, should be regarded as a first pass. Optimizing the gridding strategy is postponed to future work, in the context of a search with real data (which introduces other relevant constraints). Stage II of the MDC will test the robustness of the HMM and other algorithms like CrossCorr [67] and TwoSpect [69] to spin wandering. Previous studies demonstrate that the HMM handles signals with and without spin wandering with equal dexterity, as long as T_{drift} satisfies condition (B1) [18,19].

The HMM in this paper is solved by the Viterbi algorithm, which exploits dynamic programming. The additional phase tracking step inevitably slows down version III of the HMM compared to version II, with the number of operations scaling approximately $\propto N_O \ln N_O$ (see Sec. II A), and N_O increasing by a factor ~10. Overall, however, the implementation remains fast, processing ≈ 0.3 Hz per CPU-hr for one choice of (a_0, T_{asc}, P) , approximately 10 times slower than version II. Viterbi-based continuous wave searches have proved amenable to being implemented on graphical processing units, which can shorten the run time \approx 40-fold [76]. The computational savings from an optimized implementation on graphical processing units can be reinvested to extend the astrophysical ambition of an analysis, e.g., by targeting LMXBs other than Sco X - 1 [12]. Savings can also be reinvested to expand the scope of Viterbi-based, nonparametric, all-sky searches and searches for wandering instrumental lines [27].

What conclusions can we expect to draw about the astrophysical causes of spin wandering, when the HMM

ultimately detects a real signal? At present it is hard to say. Neutron star models involve a great deal of uncertain physics, which will blur the interpretation of any HMM detection, whether it involves versions I, II, or III, unless the detection itself reveals some unexpected and informative signature. Electromagnetic observations may improve the situation. Consider, for example, an LMXB where one observes simultaneously the X-ray flux $F_X(t)$ and the wandering spin $f_*(t)$. One might hope to crosscorrelate the fluctuations in F_X and f_* and thereby test the accretion physics [7]. However, the traditional assumption $F_X \propto \dot{M} \propto \dot{f}_*$, where \dot{M} denotes the mass accretion rate, does not always hold for various reasons, e.g., nonconservative mass transfer, hydromagnetic contributions to f_* , and unsteady dynamics due to magnetospheric instabilities [37,77]. Some of the relevant issues are canvassed in Ref. [78]. As a second illustrative example, suppose the HMM detects a steady tone with minimal spin wandering from a radio pulsar, that displays strong timing noise at radio wavelengths. Such an observation would arguably suggest, that the gravitational wave signal is emitted by the weakly coupled superfluid interior of the star as opposed to the crust (which is locked magnetically to the radio pulses). Furthermore, if f_* from the HMM approximately equals the time-averaged radio pulse frequency, it arguably represents partial evidence for pinning of the superfluid [32,79]. These and other possibilities will clarify themselves, once detections are made routinely.

ACKNOWLEDGMENTS

We would like to thank Paul Lasky, Chris Messenger, Keith Riles, Karl Wette, Letizia Sammut, John Whelan, Grant Meadors and the LIGO Scientific Collaboration Continuous Wave Working Group for detailed comments and informative discussions. We especially thank Karl Wette for alerting us to the existence of the phase extraction tool XLALEstimatePulsarAmplitudeParams in the LAL suite and Grant Meadors for pointing us to the phasesensitive formulation of the \mathcal{B} -statistic in Ref. [23]. The synthetic data for Stage I of the Sco X-1 MDC were prepared primarily by Chris Messenger with the assistance of members of the MDC team [20]. We thank Chris Messenger and Paul Lasky for their assistance in handling the MDC data. We thank the anonymous referees for their constructive feedback. P. Clearwater and L. Sun have been supported by Australian Postgraduate Awards. P. Clearwater was also a recipient of a scholarship from the Commonwealth Scientific and Industrial Research Organisation, Australia. L. Sun has been a member of the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-1764464. Advanced LIGO was built under Grant No. PHY-0823459. The research is supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav), Grant No. CE170100004.

APPENDIX A: VITERBI ALGORITHM

The Viterbi algorithm prunes the tree of possible hidden state sequences Q by appealing to Bellman's principle of optimality: if a subpath $\{q^*(t_i), ..., q^*(t_j)\}$ is optimal, then all of its subpaths are optimal as well [80]. Dynamic programming is exploited to implement the principle of optimality in an efficient, recursive fashion [14,15,18]. Pseudocode describing the implementation is presented below in abridged form for ease of reference.

At time t_k $(1 \le k \le N_T)$, let the vector $\boldsymbol{\delta}(t_k)$ store the N_Q maximum probabilities

$$\delta_{q_i}(t_k) = \max_{q_j} \Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}], \quad (A1)$$

with $1 \le i \le N_Q$, and let the vector $\mathbf{\Phi}(t_k)$ store the hidden states at t_{k-1} leading to the corresponding maximum probabilities in $\boldsymbol{\delta}(t_k)$, viz.

$$\Phi_{q_i}(t_k) = \operatorname*{arg\,max}_{q_j} \Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}],$$
(A2)

 $\Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}] = L_{o(t_k)q_i} A_{q_i q_j} \delta_{q_j}(t_{k-1}).$ (A3)

The components of $\delta(t_k)$ and $\Phi(t_k)$ are filled by running forward through the N_T observations, then the optimal path $Q^*(O)$ is reconstructed by backtracking.

(1) Initialization:

with $O^{(k)} = \{o(t_0), ..., o(t_k)\}$ and

$$\delta_{q_i}(t_0) = L_{o(t_0)q_i} \Pi_{q_i}, \tag{A4}$$

for $1 \le i \le N_Q$. (2) *Recursion*:

$$\delta_{q_i}(t_k) = L_{o(t_k)q_i} \max_{1 \le j \le N_Q} [A_{q_i q_j} \delta_{q_j}(t_{k-1})], \quad (A5)$$

$$\Phi_{q_i}(t_k) = \underset{1 \le j \le N_Q}{\arg\max} [A_{q_i q_j} \delta_{q_j}(t_{k-1})], \quad (A6)$$

for $1 \le i \le N_O$ and $1 \le k \le N_T$.

(3) *Termination*:

$$\max \Pr(Q|O) = \max_{q_i} \delta_{q_j}(t_{N_T})$$
(A7)

$$q^*(t_{N_T}) = \operatorname*{arg\,max}_{q_j} \delta_{q_j}(t_{N_T}) \tag{A8}$$

for $1 \le j \le N_Q$. (4) Optimal path backtracking:

$$q^*(t_k) = \Phi_{q^*(t_{k+1})}(t_{k+1}) \tag{A9}$$

for $0 \le k \le N_T - 1$.

APPENDIX B: DRIFT TIMESCALE

A practical recipe for choosing the drift timescale $T_{\text{drift}} = t_{n+1} - t_n$ (see Sec. II A) when tracking $f_*(t)$ is described in Refs. [18,19]. In this Appendix we generalize the recipe for the purpose of tracking $f_*(t)$ and $\Phi_*(t)$ in version III of the HMM.

The choice of T_{drift} is governed by the packaging of input data when computing the emission probability $L_{o_jq_i}$, which comes with implicit assumptions about the signal properties in the interval $t_{n-1} \leq t \leq t_n$. Importantly we require $L_{o_jq_i}$ to peak as sharply as possible in the neighborhood of the truly occupied hidden state $q(t_n)$, with $L_{o_jq_i} \approx \delta[q_i - q(t_n)]$ ideally, in order to maximize $\Pr[Q^*(O)|O]$. Typically $L_{o_jq_i}$ is computed from frequency-domain data covering the whole interval $t_{n-1} \leq t \leq t_n$, and q(t) does not contain frequency-drift variables like $f_*(t)$. Therefore the matched filter that computes $L_{o_jq_i}$ (e.g., the \mathcal{F} - or \mathcal{B} -statistic) assumes that $f_*(t)$ stays within a single, discrete bin during every HMM time-step. For this assumption to hold, one must choose T_{drift} to satisfy

$$\left| \int_{t}^{t+T_{\text{drift}}} dt' \dot{f}_{*}(t') \right| < \Delta f_{\text{drift}} \tag{B1}$$

for all t, where Δf_{drift} is the separation between adjacent frequency bins (which are assumed to be uniformly spaced in this paper, i.e., Δf_{drift} is independent of q_i). A different method of computing $L_{o_jq_i}$, e.g., from time-domain data, may impose a different constraint on T_{drift} .

It is tempting to extend the above argument to $\Phi_*(t)$ and insist that it should stay within a single bin too (of width $\Delta \Phi_{\text{drift}} = \pi/16$ in this paper),¹⁰ but this is unnecessary. Frequency-domain matched filters like the \mathcal{F} - and \mathcal{B} -statistic do not assume that $\Phi_*(t)$ is constant for $t_{n-1} \leq t \leq t_n$; they are well-behaved functions of $\Phi_*(t_{n-1})$ at the start of the HMM time-step. Confining $\Phi_*(t)$ to a single phase bin would shorten T_{drift} by a factor $\approx \pi/\Delta \Phi_{\text{drift}}$, widen every frequency bin by the same factor

¹⁰The analyst enjoys considerable freedom in setting $\Delta \Phi_{\rm drift}$, as long as the peaks in the transition probability in Fig. 1(b) are resolved. In contrast, $\Delta f_{\rm drift} = (2T_{\rm drift})^{-1}$ is determined by $T_{\rm drift}$. See Sec. II C for details.

(Nyquist theorem), and reduce proportionally the signal-tonoise ratio per frequency bin [13,22].

Naturally one does not know $f_*(t')$ in (B1) in advance, so there is some trial and error involved in choosing T_{drift} through (B1). In this paper we focus on gravitational wave searches for isolated and accreting neutron stars, whose rotatational irregularities have been studied extensively in radio [5,31] and x-ray [6,35] timing experiments, which yield autocorrelation timescales of days to months. These electromagnetic measurements therefore offer a starting point to estimate T_{drift} for other objects in the same class, where $f_{*}(t')$ is not measured [7]. For reasons of convenience described in Sec. II, we elect to work with Fouriertransformed data in this paper [18,19], which come packaged in calibrated, conditioned (antialias filtering, data drop-out), short-time Fourier transforms (SFTs) of duration $T_{\text{SFT}} = 30 \text{ min}$ [45]. Hence one has $T_{\text{SFT}} \leq T_{\text{drift}}$ as a practical matter, a constraint which would be absent in a time-domain analysis. Given the wide range of measured autocorrelation timescales, one can envisage a hierarchical search strategy, in which a search is repeated for several T_{drift} values in the range $T_{\text{SFT}} \leq T_{\text{drift}} \leq T_{\text{obs}}$, where $T_{\rm obs} \sim 1$ yr is the total observation time.

APPENDIX C: PHASE-WRAPPED ORNSTEIN-UHLENBECK PROCESS

In this Appendix we solve the Fokker-Planck equation corresponding to the stochastic differential equations (6) and (7) to obtain the probability density function (PDF) $p(t, f_*, \Phi_*)$ and hence the HMM transition probabilities over the interval $t_n \le t \le t_{n+1}$ given the initial state $q(t_n) =$ $[f_*(t_n), \Phi_*(t_n)]$ or the final state $q(t_{n+1}) = [f_*(t_{n+1}), \Phi_*(t_{n+1})]$. The discussion follows Appendix A in Ref. [51]. Equations (6) and (7) are equivalent to traditional, spatial Brownian motion, with f_* and Φ_* playing the roles of velocity and displacement respectively, except that Φ_* is 2π -periodic.

If the hidden state $q(t_n)$ occupied at the start of the HMM step $t_n \le t \le t_{n+1}$ is known with certainty, the PDF of the final state at $t = t_{n+1}$ is given by the solution $p^{F}(t, f_*, \Phi_*)$ of the forward Fokker-Planck equation [50]

$$\frac{\partial p^{\rm F}}{\partial t} = \gamma p^{\rm F} + \gamma f_* \frac{\partial p^{\rm F}}{\partial f_*} - f_* \frac{\partial p^{\rm F}}{\partial \Phi_*} + \frac{\sigma^2}{2} \frac{\partial^2 p^{\rm F}}{\partial f_*^2}, \qquad (C1)$$

evaluated at $t = t_{n+1}$ given $p^{\rm F}(t_n, f_*, \Phi_*) = \delta[f_* - f_*(t_n)] \times \delta[\Phi_* - \Phi_*(t_n)]$. If the final state $q(t_{n+1})$ is known with certainty, the PDF of the initial state is given by the solution $p^{\rm B}(t, f_*, \Phi_*)$ of the backward Fokker-Planck equation,

$$\frac{\partial p^{\rm B}}{\partial t} = \gamma f_* \frac{\partial p^{\rm B}}{\partial f_*} - f_* \frac{\partial p^{\rm B}}{\partial \Phi_*} - \frac{\sigma^2}{2} \frac{\partial^2 p^{\rm B}}{\partial f_*^2}, \qquad ({\rm C2})$$

evaluated at $t = t_n$ given $p^{B}(t_{n+1}, f_*, \Phi_*) = \delta[f_* - f_*(t_{n+1})]\delta[\Phi_* - \Phi_*(t_{n+1})]$. Equation (C2) is the adjoint of (C1). Upon multiplying (C1) by the integrating factor $\exp(-\gamma t)$, we find

$$p^{\mathrm{B}}(t, f_*, \Phi_*) \propto \exp(-\gamma t) p^{\mathrm{F}}(t, f_*, \Phi_*; \sigma^2 \mapsto -\sigma^2),$$
 (C3)

where $\sigma^2 \mapsto -\sigma^2$ denotes replacing σ^2 by $-\sigma^2$ in $p^{\rm F}$.

Upon Fourier analyzing p^{F} , as in Ref. [51], we find that the characteristic function

$$\tilde{p}^{\mathrm{F}}(t,\kappa,m) = \int_{0}^{2\pi} d\Phi_* \int_{-\infty}^{\infty} df_* \exp(-im\Phi_* - i\kappa f_*)$$
$$\times p^{\mathrm{F}}(t,f_*,\Phi_*) \tag{C4}$$

satisfies

$$\frac{\partial \tilde{p}^{\rm F}}{\partial t} = (-\gamma \kappa + m) \frac{\partial \tilde{p}^{\rm F}}{\partial \kappa} - \frac{\sigma^2 \kappa^2 \tilde{p}^{\rm F}}{2}, \qquad ({\rm C5})$$

subject to the initial condition

$$\tilde{p}^{\mathrm{F}}(t_n,\kappa,m) = \exp[-im\Phi_*(t_n) - i\kappa f_*(t_n)]. \quad (\mathrm{C6})$$

Equations (C5) and (C6) are solved by the method of characteristics to give

$$\tilde{p}^{\mathrm{r}}(t_{n+1},\kappa,m) = \exp\left[-im\Phi_{*}(t_{n}) - i\rho f_{*}(t_{n})\right] \\ \times \exp\left[\frac{\sigma^{2}}{4\gamma}\left(\rho - \frac{m}{\gamma}\right)\left(\rho + \frac{3m}{\gamma}\right)\right] \\ \times \exp\left\{-\frac{\sigma^{2}}{2}\left[\frac{m^{2}\tau}{\gamma^{2}} + \frac{2m}{\gamma^{2}}\left(\kappa - \frac{m}{\gamma}\right)\right]\right\} \\ \times \exp\left[-\frac{\sigma^{2}}{4\gamma}\left(\rho - \frac{m}{\gamma}\right)^{2}\exp(2\gamma\tau)\right],$$
(C7)

with $\tau = t_{n+1} - t_n$ and

$$\rho = \frac{m}{\gamma} + \left(\kappa - \frac{m}{\gamma}\right) \exp(-\gamma\tau) \tag{C8}$$

and hence

$$\tilde{p}^{\mathrm{F}}(t_{n+1}, f_*, \Phi_*) = (2\pi)^{-2} \sum_{m=-\infty}^{\infty} \exp(im\Phi_*) \\ \times \int_{-\infty}^{\infty} d\kappa \exp(i\kappa f_*) \tilde{p}^{\mathrm{F}}(t, \kappa, m).$$
(C9)

By completing the square in the argument of the exponential in (C7), one finds that (C9) can be written as a wrapped Gaussian [51].

The solution (C3) to the backward Fokker-Planck equation (C2) provides an efficient route to calculating the maximum probabilities at each HMM step, which are stored in the vector $\delta(t_k)$ in the Viterbi implementation described in Appendix A. Equation (C3), just like (C9), can be expressed as a wrapped Gaussian, viz.

$$p^{\mathrm{B}}(t_n, \mathbf{q}) = (2\pi)^{-1} (\det \mathbf{\Sigma})^{-1/2}$$
$$\times \sum_{m=-\infty}^{\infty} \exp[-(\mathbf{q} - \mathbf{Q}_m)\mathbf{\Sigma}^{-1}(\mathbf{q} - \mathbf{Q}_m)^{\mathrm{T}}],$$
(C10)

with $\mathbf{q} = (f_*, \Phi_*)$ and matrix elements

$$(\mathbf{Q}_m)_1 = f_*(t_{n+1}) \exp(-\gamma \tau),$$
 (C11)

$$(\mathbf{Q}_m)_2 = \Phi_*(t_{n+1}) + \frac{f_*(t_{n+1})}{\gamma} [1 - \exp(-\gamma\tau)] - 2\pi m,$$
(C12)

$$\boldsymbol{\Sigma}_{11} = \frac{\sigma^2}{2\gamma} [1 - \exp(-2\gamma\tau)], \qquad (C13)$$

$$\Sigma_{12} = \Sigma_{21} = \frac{\sigma^2}{2\gamma^2} [1 - \exp(-\gamma\tau)]^2,$$
 (C14)

$$\Sigma_{22} = \frac{\sigma^2}{2\gamma^3} \{ 1 + 2\gamma\tau - [2 - \exp(-\gamma\tau)]^2 \}.$$
 (C15)

We can then read off the moments $\langle f_* \rangle$, $\langle \Phi_* \rangle$, $\langle f_*^2 \rangle - \langle f_* \rangle^2$, $\langle f_* \Phi_* \rangle - \langle f_* \rangle \langle \Phi_* \rangle$, and $\langle \Phi_*^2 \rangle - \langle \Phi_* \rangle^2$ of p^{B} by inspection from (C11)–(C15) respectively [51].

APPENDIX D: MAXIMUM LIKELIHOOD ALTERNATIVES TO THE *B*-STATISTIC

In this Appendix, we review briefly the maximum likelihood formulas for $L_{o(t_n)q_i}$ used in versions I and II of the HMM, which do not depend on rotational phase. [18,19] We then present for completeness a natural, phase-dependent generalization of these maximum likelihood formulas. Empirical testing indicates, that the generalized formula yields no discernible improvement in performance over versions I and II of the HMM, unlike the \mathcal{B} -statistic presented in Sec. IV C.

In version I of the HMM, [18] for an isolated source $(a_0 = 0)$ with zero phase (cf. spin) wandering $(\Phi_w = 0)$, the log likelihood is just the \mathcal{F} -statistic, $G(f_0) = \mathcal{F}(f_0)$, viz.

$$\mathcal{F}(f_0) = \frac{4\mathbf{F}(f_0)\mathbf{H}^{-1}\mathbf{F}(f_0)^{\dagger}}{T_{\text{obs}}S_h(f_0)},$$
 (D1)

where a dagger denotes the Hermitian transpose, with

$$\mathbf{F}(f_0) = [F_{1a}(f_0), F_{1b}(f_0)], \qquad (D2)$$

$$\mathbf{H} = \begin{pmatrix} A & C \\ C & B \end{pmatrix}, \tag{D3}$$

A = (a||a), B = (b||b), and C = (a||b). In the general case $A_{2i} \neq 0$, Eqs. (D1)–(D3) contain additional, analogous terms involving F_{2a} and F_{2b} , obtained from F_{1a} and F_{1b} by replacing f_0 with $2f_0$.

For a binary source $(a_0 \neq 0)$ with zero phase wandering $(\Phi_w = 0)$, the log likelihood in version I of the HMM is approximated by the Bessel-weighted \mathcal{F} -statistic,

$$G(f_0) = \sum_{s=-M'}^{M'} [J_s(2\pi f_0 a_0)]^2 \mathcal{F}(f_0 - s/P), \quad (D4)$$

with $M' = \operatorname{ceil}(2\pi f_0 a_0)$. Equation (D4) adds together the power in orbital sidebands incoherently; it takes no account of the relative Fourier phases of the sidebands. This omission is corrected in version II of the HMM, [19] where F_{1a} and F_{1b} are replaced by J_{1a} and J_{1b} , defined by (21) and (22) respectively, in order to include orbital phase information. The log likelihood is calculated similarly to the binary-modulated \mathcal{F} -statistic and yields the \mathcal{J} -statistic, $G(f_0) = \mathcal{J}(f_0)$, with

$$\mathcal{J}(f_0) = \frac{4\mathbf{J}(f_0)\mathbf{H}^{-1}\mathbf{J}(f_0)^{\dagger}}{T_{\text{obs}}S_h(f_0)}$$
(D5)

and

$$\mathbf{J}(f_0) = [J_{1a}(f_0), J_{1b}(f_0)].$$
(D6)

Equation (D5) concentrates all the signal power in the orbital sidebands into one f_0 bin, unlike (D4), as verified in Fig. 1 in Ref. [19]. It is therefore as sensitive for binary sources, as (D1) is for isolated sources, i.e., (D1) and (D5) can detect the same h_0 value [19].

When the HMM tracks $\Phi_*(t)$ as well as $f_*(t)$, it is tempting to generalize $G(f_0)$ to $G(f_0, \Phi_0)$, where Φ_0 is the trial phase, by analogy with (D5). First, one may try to incorporate the phase into the amplitudes A_{1i} , as in Ref. [59], e.g., $A_{11} = A_+ \cos 2\psi \cos \Phi_w - A_\times \sin 2\psi \sin \Phi_w$. Unfortunately, maximizing the likelihood Λ' with respect to A_{1i} returns estimators \hat{A}_{1i} , which are rotated versions of the phase-independent estimators, e.g., \hat{A}_{11} becomes $\hat{A}_{11} \cos \Phi_w + \hat{A}_{13} \sin \Phi_w$. The resulting \mathcal{F} -statistic is independent of phase, as shown in Appendix A in Ref. [19] in the context of orbital phase. Instead, one may try to factorize the \mathcal{F} -statistic into a quadratic form constructed from complex amplitudes, multiply the complex amplitudes by the cosine of the phase, and reassemble the



FIG. 10. Maximum likelihood phase tracking. Estimated (vertical axis) versus injected (horizontal axis) phase for $h_0 = 8.0 \times 10^{-25}$ (purple points; 63 trials) and $h_0 = 8.0 \times 10^{-26}$ (green points; 63 trials), using the maximum likelihood estimate returned by the LAL function *XLALEstimatePulsarAmplitudeParams*.

quadratic form to obtain a real likelihood.¹¹ In this spirit, we define

$$G(f_0, \Phi_0) = \frac{4\mathbf{R}(f_0, \Phi_0)\mathbf{H}^{-1}\mathbf{R}(f_0, \Phi_0)^{\dagger}}{T_{\text{obs}}S_h(f_0)}, \quad (D7)$$

with

(

$$\mathbf{R}(f_0, \Phi_0) = [R_{1a}(f_0, \Phi_0), R_{1b}(f_0, \Phi_0)].$$
(D8)

Numerical experiments reveal that (D7) produces no improvement in sensitivity compared to version I of the HMM. Essentially this is because noise in the phase estimate defeats the HMM's ability to reject paths with inconsistent phase. This can be seen by plotting the output of the function XLALEstimatePulsarAmplitudeParams in the LAL suite, which returns maximum likelihood estimates of the source parameters (including phase) given F_{1a} and F_{1b} , against the injected phase.¹² Figure 10 demonstrates that the estimated and injected phases are strongly correlated for $h_0 = 8.0 \times 10^{-25}$. However, the correlation weakens appreciably for $h_0 = 8.0 \times 10^{-26}$ and even more so near the detection limit for version III of the HMM ($h_0 = 1.3 \times 10^{-26}$), where the points scatter randomly (not plotted). The Pearson correlation coefficient, computed versus h_0 in Table V, exhibits the same behavior.

Note that the HMM tracks the phase difference between HMM steps; the absolute phase enters through the prior and is not tracked explicitly. This differs subtly from a fully coherent \mathcal{F} -statistic search (without spin wandering), where \mathcal{F} is evaluated as a function of $\Phi_{w}(t_0)$ as well as $f_0^{(k)}$, α , and δ [13].

APPENDIX E: VALIDATION TESTS

In this Appendix, we present for completeness and reproducibility the results of several validation tests applied to version III of the HMM. The tests relate to the PDF of the \mathcal{B} -statistic after a single HMM step, the PDF of the block score after multiple HMM steps, the detection probability as a function of N_T for T_{drift} or T_{obs} fixed, the effect of the block definition on the detector's performance, and the conservation of signal power by the detection statistic. The tests will help to guide future refinements of the HMM.

1. PDF of the detection statistic

Figure 11(a) displays the PDF of $\ln \beta$ computed for a single HMM step in pure noise ($h_0 = 0$; purple histogram) and for a relatively strong injection ($h_0 = 5 \times 10^{-26}$; green histogram). The injection shifts the mode of the PDF to the right, as expected. Figure 11(b) investigates in more detail the functional form of the noise-only PDF. All the histograms and curves in Fig. 11(b) are normalized, and the results are independent of f_0 and Φ_0 .¹³ It is clear by inspection that the noise-only \mathcal{B} -statistic does not obey a central chi-squared distribution with four degrees of freedom (unlike the \mathcal{F} -statistic) nor with two to six degrees of freedom. The two statistics correspond to slightly different choices of amplitude priors within a Bayesian framework but are otherwise the same, with $|\ln \mathcal{B} - \mathcal{F}| \lesssim 0.05 \mathcal{F}$ for a wide range of signal and noise parameters [21,23–25]. However, by marginalizing over ψ , cos ι , and h_0 in (28), one implicitly enforces constraints between A_+ and A_{\times} and hence the four amplitudes A_{1i} in (15), so that the statistic is no longer the sum of four independent squares.

Detection with the HMM is performed using the block score S defined in (35) in Sec. V C. Figure 11(c) displays histograms of S after $N_T = 37$ steps of the HMM for pure noise ($h_0 = 0$; purple histogram) and an injection below the single-step detection threshold ($h_0 = 1.3 \times 10^{-26}$; green histogram). The peaks of the noise-only and noise-plusinjection histograms are clearly separated, demonstrating the discriminating power of the HMM. The PDFs of S are narrower than for ln \mathcal{B} and have thinner right-hand tails, because the nonlinear maximization step in the Viterbi algorithm produces an extreme value distribution similar to the Gumbel law. [19] The maximum is taken over all Viterbi paths terminating in a given frequency-phase bin, so paths terminating in neighboring bins are correlated because they share common subpaths. There is no analytic

¹¹In nongravitational-wave applications where the signal is an unmodulated sinusoid with a single polarization mode, and the antenna beam-pattern does not vary diurnally, this procedure yields the exact, maximum likelihood estimator [51].

¹²At the time of writing, *XLALEstimatePulsarAmplitudeParams* incorrectly adds π to the phase. The error is corrected here.

¹³There is a weak dependence on the width of the running median window applied to the power spectral density, as for the \mathcal{F} -statistic [16,53].

TABLE V. Estimated versus injected phase: Pearson correlation coefficient as a function of signal strength for the maximum likelihood estimator *XLALEstimatePulsarAmplitudeParams* with 10^3 realizations.

$h_0 (10^{-26})$	Coefficient		
80	0.978		
8.0	0.464		
1.7	0.156		
1.3	0.059		

expression for the PDF of $\ln \Pr[Q^*(O)|O]$ in the literature to the best of our knowledge. [19] We therefore rely on the empirical PDF in Fig. 11(c) to set $S_{\text{th}}(f)$ given P_{a} .

2. Detection probability versus N_T

Another important question is how the performance of the HMM scales with N_T . We formulate the question with respect to two practical scenarios: (i) T_{drift} is fixed, and

 $T_{\rm obs} \propto N_T$ varies; and (ii) $T_{\rm obs}$ is fixed, and $T_{\rm drift} \propto N_T^{-1}$ varies. Figure 12 presents data for scenario (i). As expected, the sensitivity of the HMM increases, as N_T and hence T_{obs} increase [51]. We observe in Fig. 12(b) that the detection probability rises with N_T at fixed $P_a = 10^{-2}$. The same trend occurs in Fig. 12(a) for $10^{-3} \le P_a \le 1$. Figure 12(b) corresponds to a vertical cut at constant $P_{\rm a} = 10^{-2}$ through the family of ROC curves in Fig. 12(a). One subtlety is that $S_{\rm th}$ depends on N_T through two countervailing factors. The number of frequency bins per block is proportional to N_T , so $S_{\rm th}$ should increase with N_T , ceteris paribus, to keep $P_{\rm a}$ per block fixed; but the product Pr(Q|O) in (1) decreases with N_T , as more factors $L_{o(t_n)q(t_n)}A_{q(t_n)q(t_{n-1})} \leq 1$ are appended, implying that $S_{\rm th}$ should decrease with N_T for fixed $P_{\rm a}$. The latter effect outweighs the former, as is evident in Fig. 12(c); the threshold decreases from $S_{\rm th} \approx 4.0$ for $N_T = 5$ to $S_{\text{th}} \approx -5.5$ for $N_T = 35$. In a genuine, astrophysical search one would typically set $P_{\rm a} = 10^{-2}$ for the whole search band ($B \sim 1$ kHz), or for sub-bands with $\Delta f_{sub} \sim 1$ Hz (to facilitate data handling), and hence



FIG. 11. Normalized PDF of the detection statistic for pure noise (purple histograms) and a detected injection with the source parameters in Table II (green histograms). (a) Logarithm of the \mathcal{B} -statistic, $\ln \mathcal{B}(f_*, \Phi_*)$, computed for a single HMM step in the bin (f_*, Φ_*) containing the injection (where present), with $h_0 = 0$ (purple histogram) and $h_0 = 5 \times 10^{-26}$ (green histogram). (b) Noise-only histogram from (a) rebinned over the domain [-3, 3] and overlaid with normalized, central, chi-squared distributions with 2, 3, 4, and 6 degrees of freedom (solid curves; color scheme in legend), in order to test for congruence with the functional form of the \mathcal{F} -statistic PDF. (c) Block score S_i defined by (35) for the block containing the injection (where present), with $h_0 = 0$ (purple histogram), $h_0 = 1.3 \times 10^{-26}$ (green histogram), and $N_T = 37$. Realizations: 2.5×10^3 per histogram.



FIG. 12. Detector performance as a function of N_T for $T_{\text{drift}} = 10$ d fixed, $T_{\text{obs}} = N_T T_{\text{drift}} \propto N_T$ variable, and the source parameters in Table II. (a) ROC curves for $h_0 = 1.7 \times 10^{-26}$ and $N_T = 5$ (purple curve), 15 (green curve), 30 (blue curve). (b) Detection probability $1 - P_d$ versus N_T for $h_0 = 1.3 \times 10^{-26}$ (green curve), 1.7×10^{-26} (purple curve), and $P_a = 10^{-2}$ per block. (c) Block score threshold S_{th} [see (35)] versus N_T for false alarm probability $P_a = 10^{-2}$ per block; the number of bins per block, $\Pr(Q|O)$, and hence S_{th} depend on N_T . All curves are calculated for version III of the HMM. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$. Realizations: 10^3 per curve.

have $P_a \ll 10^{-2}$ per block, with $N_T \Delta f_{\text{drift}} \ll \Delta f_{\text{sub}} \leq B$. The scalings with N_T are the same in this regime, but the ROC curves are time-consuming to generate by Monte Carlo simulations.¹⁴

Figure 13 presents data for scenario (ii) in the previous paragraph, i.e., fixed T_{obs} . The trend with N_T depends on whether $T_{drift} \propto N_T^{-1}$ is less or greater than the characteristic timescale over which the signal frequency wanders. [51] If T_{drift} is less than the wandering timescale, the detection probability decreases, as T_{drift} decreases; it is disadvantageous to shorten the coherent integration in a HMM segment, when the frequency wanders by less than one bin during a segment. We observe this behavior in Fig. 13(b) to the left of the peak. If T_{drift} is greater than the wandering timescale, the detection probability increases, as T_{drift} decreases; it is better to make the segments shorter, as required by condition (B1), up to the point where the frequency wanders by roughly one bin during a segment. We observe this behavior to the right of the peak in Fig. 13(b). The behavior in Fig. 13(b) for $P_a = 10^{-2}$ per block is consistent with the ROC curves in Fig. 13(a) over the range $10^{-3} \le P_a \le 1$. The threshold decreases with N_T in Fig. 13(c), just like in Fig. 12(c), because it is approximately independent of T_{drift} .

3. Block definition

What happens when a candidate straddles the boundary between two blocks? In this paper, we treat it as a special case, to be followed up through a veto procedure in a genuine astrophysical search. Straddlers represent a modest fraction $\sim N_T^{-1/2}$ of all signals or false alarms.¹⁵ Figure 14(a)

¹⁴Occasionally situations may arise, where it is desirable to hold the number of bins per block fixed while varying N_T , e.g., when comparing results from two data sets of different durations. We defer the analysis of such situations to future work.

¹⁵Alternatively one can record straddlers on a candidate list and consolidate candidates that share common subpaths, after all the data are analyzed. This complicates the statistical interpretation of the results, because HMM paths with common subpaths are correlated [19].



FIG. 13. Detector performance as a function of N_T for $T_{obs} = 360$ d fixed, $T_{drift} = N_T^{-1}T_{obs} \propto N_T^{-1}$ variable, and the source parameters in Table II. (a) ROC curves for $h_0 = 1.3 \times 10^{-26}$ and $N_T = 6$ (purple curve), 18 (green curve), 36 (blue curve), chosen to give an integer number of days per HMM step. (b) Detection probability $1 - P_d$ versus N_T for $h_0 = 1.3 \times 10^{-26}$ (purple curve), 1.5×10^{-26} (green curve), and $P_a = 10^{-2}$ per block. (c) Block score threshold S_{th} [see (35)] versus N_T for false alarm probability $P_a = 10^{-2}$ per block. The number of bins per block and hence S_{th} scale with N_T , with $T_{drift} = N_T^{-1}T_{obs} \propto N_T^{-1}$ variable (purple curve) and $T_{drift} = 10$ d = constant [green curve; copied from Fig. 12(c) for comparison]. All curves are calculated for version III of the HMM. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$. Realizations: 10^3 per curve.



FIG. 14. Effects of block definition on performance. (a) ROC curves for blocks of bandwidth $2N_T\Delta f_{drift}$, whose leftmost frequency bins are shifted 1, 2, ..., $N_T - 1$ bins to the right of an arbitrary reference bin, for $h_0 = 1.3 \times 10^{-26}$, $T_{drift} = 10$ d, and $N_T = 37$. The 36 curves overlap closely and cannot be distinguished by eye. (b) ROC curves for blocks of bandwidth $2kN_T\Delta f_{drift}$, with k = 0.243, 0.514, 0.757, 1.00, 1.24, 1.51, 1.76, 2.00 (chosen to give an integer number of bins per block). All curves are calculated for version III of the HMM. Source parameters: see Table II. Control parameters: $\gamma = 1.0 \times 10^{-16} \text{ s}^{-1}$, $\sigma = 3.7 \times 10^{-10} \text{ s}^{-3/2}$. Realizations: 10³ per curve.



FIG. 15. Conservation of signal power. \mathcal{B} -statistic at the correct phase bin versus frequency (in Hz) for a strong binary injection with $h_0 = 8 \times 10^{-25}$ and source parameters drawn from Tables II and III. (a) \mathcal{B} evaluated with F_{1a} and F_{1b} in (28)–(34). (b) \mathcal{B} evaluated with J_{1a} and J_{1b} (Jacobi-Anger version) in (28)–(34). Note the different vertical scales in (a) and (b).

verifies that the absolute position of the block boundary does not affect the ROC curves appreciably. It displays $N_T - 1$ individual ROC curves for $N_T - 1$ different block boundaries, in which the leftmost frequency bin is shifted right by $1, 2, ..., N_T - 1$ bins relative to an arbitrary, reference bin. The curves overlap closely and are barely distinguishable by eye.

Likewise we find that the performance of the HMM depends weakly on the bandwidth of each block. It is unlikely for a path to drift by $\approx N_T$ bins after $N_T \gg 1$ HMM steps, even when the tails in $A_{q_jq_i}$ with |j - i| > 1 are preserved, as in Appendix C (cf. truncated $A_{q_jq_i}$ with $|j - i| \le 1$ in Ref. [18]). Figure 14(b) verifies this property by plotting multiple ROC curves for block widths $2kN_T\Delta f_{drift}$ with $0.2 \le k \le 2$. Again the curves overlap closely. We use k = 1 henceforth in this paper.

4. Conservation of signal power

In version II of the HMM, based on the \mathcal{J} -statistic, $J_{1a}(f_0)$ and $J_{1b}(f_0)$ marshal the Doppler-shifted signal power into one frequency bin by coherently summing orbital sidebands weighted by $J_s(2\pi f_0 a_0)e^{-is\phi_a}$. It turns out that the same holds true empirically for the \mathcal{B} -statistic, although there exists no formal mathematical proof at the time of writing; it may not be possible to derive the \mathcal{B} statistic for a binary source exactly as a Jacobi-Anger expansion of the \mathcal{B} -statistic. This Appendix verifies numerically that minimal power is lost or dispersed into neighboring frequency bins, when the \mathcal{B} -statistic is evaluated using $J_{1a}(f_0)$ and $J_{1b}(f_0)$.

Figure 15(a) graphs $\mathcal{B}(f_0, \Phi_*)$ versus f_0 (evaluated for Φ_* in the injected bin) for a strong binary signal using F_{1a}

and F_{1b} to evaluate \mathcal{B} . As the orbital motion is not accounted for, \mathcal{B} displays a comb of orbital sidebands at $f_* + s/P$, which fill the band $111.09 \le f_0/(1 \text{ Hz}) \le$ 111.11. The comb exhibits the classic two-horned envelope familiar from the Sideband algorithm, [52,60] because the source spends more time moving perpendicular to the plane of the sky (when the orbital Doppler shift is a maximum) than moving perpendicular to the line of sight (when the Doppler shift is zero). Figure 15(b) shows the same thing as Fig. 15(a) but with F_{1a} and F_{1b} replaced by J_{1a} and J_{1b} when computing \mathcal{B} . The sidebands now merge into one peak, which is \approx 40 times higher than the tallest peak in the comb in Fig. 15(a) (note the different scales). Identical behavior is seen in Fig. 1 in Ref. [19] for the \mathcal{J} -statistic instead of the \mathcal{B} -statistic.

APPENDIX F: REPRESENTATIVE PHASE PATHS RECOVERED BY THE HMM FOR A SOURCE IN A BINARY

In this Appendix, we examine for completeness the optimal phase paths $\Phi_*(t)$ recovered by version III of the HMM for the representative examples of binary sources studied in Sec. VI A.

Figure 16 displays the absolute error between the injected and recovered phase for the three synthetic binary sources tracked in Fig. 7. The interpretation is the same as in Sec. V B. The phase error jumps around, even after unwinding the phase wrapping, because the \mathcal{B} -statistic spreads the signal power over multiple phase bins. On balance, though, the imperfect phase tracking delivers improved sensitivity, as evidenced by comparing Figs. 7(a) and 7(b) and the ROC curves in Sec. VI B.



FIG. 16. Phase tracking in a representative source in a binary. Layout as for Fig. 4 but for the three sources in Fig. 7.

- [1] K. Riles, Gravitational waves: Sources, detectors and searches, Prog. Part. Nucl. Phys. **68**, 1 (2013).
- [2] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, First search for gravitational waves from known pulsars with Advanced LIGO, Astrophys. J. 839, 12 (2017).
- [3] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.* (The LIGO Scientific and the Virgo Collaborations), Upper limits on gravitational waves from Scorpius X-1 from a model-based cross-correlation search in Advanced LIGO data, Astrophys. J. 847, 47 (2017).
- [4] G. Woan, M. D. Pitkin, B. Haskell, D. I. Jones, and P. D. Lasky, Evidence for a minimum ellipticity in millisecond pulsars, Astrophys. J. Lett. 863, L40 (2018).
- [5] J. M. Cordes and G. S. Downs, JPL pulsar timing observations. III—Pulsar rotation fluctuations, Astrophys. J. Suppl. Ser. 59, 343 (1985).
- [6] L. Bildsten, D. Chakrabarty, J. Chiu, M. H. Finger, D. T. Koh, R. W. Nelson, T. A. Prince, B. C. Rubin, D. M. Scott, M. Stollberg, B. A. Vaughan, C. A. Wilson, and R. B.

Wilson, Observations of accreting pulsars, Astrophys. J. Suppl. Ser. **113**, 367 (1997).

- [7] A. Mukherjee, C. Messenger, and K. Riles, Accretioninduced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches, Phys. Rev. D 97, 043016 (2018).
- [8] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, Searches for continuous gravitational waves from nine young supernova remnants, Astrophys. J. **813**, 39 (2015).
- [9] L. Sun, A. Melatos, P. D. Lasky, C. T. Y. Chung, and N. S. Darman, Cross-correlation search for continuous gravitational waves from a compact object in SNR 1987A in LIGO Science run 5, Phys. Rev. D 94, 082004 (2016).
- [10] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from young supernova remnants, Phys. Rev. D 97, 043013 (2018).
- [11] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain *et al.*, Beating the spin-down limit on gravitational

wave emission from the crab pulsar, Astrophys. J. Lett. **683**, L45 (2008).

- [12] A. L. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, Detecting gravitational wave emission from the known accreting neutron stars, Mon. Not. R. Astron. Soc. 389, 839 (2008).
- [13] 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).
- [14] B. G. Quinn and E. J. Hannan, in *The Estimation and Tracking of Frequency* (Cambridge University Press, Cambridge, England, 2001), p. 266.
- [15] A. Viterbi, Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, IEEE Trans. Inf. Theory 13, 260 (1967).
- [16] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, (The LIGO Scientific and the Virgo Collaborations), Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model, Phys. Rev. D **95**, 122003 (2017).
- [17] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.*, Search for gravitational waves from Scorpius X-1 in the second Advanced LIGO observing run with an improved hidden Markov model, Phys. Rev. D 100, 122002 (2019).
- [18] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin, Phys. Rev. D 93, 123009 (2016).
- [19] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking, Phys. Rev. D 96, 102006 (2017).
- [20] C. Messenger, H. Bulten, S. Crowder, V. Dergachev, D. Galloway, E. Goetz, R. Jonker, P. Lasky, G. Meadors, A. Melatos, S. Premachandra, K. Riles, L. Sammut, E. Thrane, J. Whelan, and Y. Zhang, Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors, Phys. Rev. D 92, 023006 (2015).
- [21] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, Classical Quant. Grav. 26, 204013 (2009).
- [22] V. Dergachev, On blind searches for noise dominated signals: A loosely coherent approach, Classical Quant. Grav. 27, 205017 (2010).
- [23] V. Dergachev, Loosely coherent searches for sets of wellmodeled signals, Phys. Rev. D 85, 062003 (2012).
- [24] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, New coordinates for the amplitude parameter space of continuous gravitational waves, Classical Quant. Grav. 31, 065002 (2014).
- [25] S. Dhurandhar, B. Krishnan, and J. L. Willis, Marginalizing the likelihood function for modeled gravitational wave searches, arXiv:1707.08163.

- [26] J. J. Bero and J. T. Whelan, An analytic approximation to the Bayesian detection statistic for continuous gravitational waves, Classical Quant. Grav. 36, 015013 (2019).
- [27] J. Bayley, C. Messenger, and G. Woan, Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals, Phys. Rev. D 100, 023006 (2019).
- [28] M. A. Alpar, R. Nandkumar, and D. Pines, Vortex creep and the internal temperature of neutron stars timing noise in pulsars, Astrophys. J. 311, 197 (1986).
- [29] K. S. Cheng, Outer magnetospheric fluctuations and pulsar timing noise, Astrophys. J. 321, 799 (1987).
- [30] P. Jones, The generation of timing noise by superfluid rotation in pulsars, Mon. Not. R. Astron. Soc. 246, 364 (1990), https://ui.adsabs.harvard.edu/abs/1990MNRAS.246 ..364J/abstract.
- [31] S. Price, B. Link, S. N. Shore, and D. J. Nice, Timecorrelated structure in spin fluctuations in pulsars, Mon. Not. R. Astron. Soc. 426, 2507 (2012).
- [32] A. Melatos, J. A. Douglass, and T. P. Simula, Persistent gravitational radiation from glitching pulsars, Astrophys. J. 807, 132 (2015).
- [33] R. E. Taam and B. A. Fryxell, On nonsteady accretion in stellar wind-fed X-ray sources, Astrophys. J. 327, L73 (1988).
- [34] A. Baykal, A. Alpar, and U. Kiziloglu, A shot noise model for a two-component neutron star, Astron. Astrophys. 252, 664 (1991), https://ui.adsabs.harvard.edu/abs/1991A%26A. ..252..664B/abstract.
- [35] A. Baykal and H. Oegelman, An empirical torque noise and spin-up model for accretion-powered x-ray pulsars, Astron. Astrophys. 267, 119 (1993), https://ui.adsabs.harvard.edu/ abs/1993A%26A...267..119B/abstract.
- [36] M. deKool and U. Anzer, A simple analysis of period noise in binary x-ray pulsars, Mon. Not. R. Astron. Soc. 262, 726 (1993).
- [37] M. M. Romanova, G. V. Ustyugova, A. V. Koldoba, and R. V. E. Lovelace, The propeller regime of disk accretion to a rapidly rotating magnetized star, Astrophys. J. 616, L151 (2004).
- [38] S. Paris and C. Jauffret, Frequency line tracking using HMM-based schemes, IEEE Trans. Aerospace Electron. Syst. 39, 439 (2003), https://ui.adsabs.harvard.edu/abs/ 2003ITAES..39..439P/abstract.
- [39] L. White and R. Elliott, A mixed MAP/MLSE receiver for convolutional coded signals transmitted over a fading channel, IEEE Trans. Signal Process. 50, 1205 (2002).
- [40] J. J. Williams and A. K. Katsaggelos, An HMM-based speech-to-video synthesizer, IEEE Trans. Neural Netw. 13, 900 (2002).
- [41] R. Barrett and D. Holdsworth, Frequency tracking using hidden Markov models with amplitude and phase information, IEEE Trans. Signal Process. 41, 2965 (1993).
- [42] X. Xie and R. Evans, Multiple target tracking and multiple frequency line tracking using hidden Markov models, IEEE Trans. Signal Process. 39, 2659 (1991).
- [43] X. Xie and R. Evans, Frequency-wavenumber tracking using hidden Markov models, IEEE Trans. Signal Process. 41, 1391 (1993).

- [44] R. Streit and R. Barrett, Frequency line tracking using hidden Markov models, IEEE Trans. Acoust. Speech Signal Process. 38, 586 (1990).
- [45] G. Mendell, Short-time Fourier transform (sft) specification, LIGO Report No. T020043, 2002.
- [46] R. Prix, The *F*-statistic and its implementation in Com puteFstatistic v2, LIGO Report No. T0900149, 2011.
- [47] A. Melatos, C. Peralta, and J. S. B. Wyithe, Avalanche dynamics of radio pulsar glitches, Astrophys. J. 672, 1103 (2008).
- [48] C. Espinoza, A. Lyne, B. Stappers, M. Kramer, M. Burgay, N. DAmico, P. Esposito, A. Pellizzoni, and A. Possenti, Glitches in the rotation of pulsars, in *Radio Pulsars:* An Astrophysical Key to Unlock The Secrets of the Universe, AIP Conference Proceedings Vol. 1357 (2011), pp. 117–120, https://ui.adsabs.harvard.edu/abs/2011AIPC .1357..117E/abstract.
- [49] A. Melatos, L. M. Dunn, S. Suvorova, W. Moran, and R. J. Evans, Pulsar glitch detection with a hidden Markov model, Astrophys. J. 896, 78 (2020).
- [50] C. W. Gardiner, Handbook of stochastic methods for physics, chemistry and the natural sciences, 2nd ed. (Springer, Berlin, 1994).
- [51] S. Suvorova, A. Melatos, R. J. Evans, W. Moran, P. Clearwater, and L. Sun, Phase-continuous frequency line track-before-detect of a tone with slow frequency variation, IEEE Trans. Signal Process. 66, 6434 (2018).
- [52] L. Sammut, C. Messenger, A. Melatos, and B. Owen, Implementation of the frequency-modulated sideband search method for gravitational waves from low mass x-ray binaries, Phys. Rev. D 89, 043001 (2014).
- [53] J. Aasi *et al.*, Directed search for gravitational waves from Scorpius X-1 with initial LIGO data, Phys. Rev. D 91, 062008 (2015).
- [54] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams *et al.*, Gravitational waves from known pulsars: Results from the initial detector era, Astrophys. J. **785**, 119 (2014).
- [55] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, Search for continuous gravitational waves from neutron stars in globular cluster NGC 6544, Phys. Rev. D **95**, 082005 (2017).
- [56] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.*, All-sky search for periodic gravitational waves in the O1 LIGO data, Phys. Rev. D **96**, 062002 (2017).
- [57] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.*, First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data, Phys. Rev. D **96**, 122004 (2017).
- [58] C. Cutler, An improved, "phase-relaxed" F-statistic for gravitational-wave data analysis, Phys. Rev. D 86, 063012 (2012).
- [59] R. Prix and J. T. Whelan, *F* -statistic search for white-dwarf binaries in the first Mock LISA Data Challenge, Classical Quant. Grav. 24, S565 (2007).

- [60] C. Messenger and G. Woan, A fast search strategy for gravitational waves from low-mass x-ray binaries, Classical Quant. Grav. 24, S469 (2007).
- [61] D. K. Galloway, S. Premachandra, D. Steeghs, T. Marsh, J. Casares, and R. Cornelisse, Precision ephemerides for gravitational-wave searches. I. Sco X-1, Astrophys. J. 781, 14 (2014).
- [62] S. S. Premachandra, D. K. Galloway, J. Casares, D. T. Steeghs, and T. R. Marsh, Precision ephemerides for gravitational wave searches. II. Cyg X-2, Astrophys. J. 823, 106 (2016).
- [63] L. Wang, D. Steeghs, D. K. Galloway, T. Marsh, and J. Casares, Precision ephemerides for gravitational-wave searches—III. Revised system parameters of Sco X-1, Mon. Not. R. Astron. Soc. 478, 5174 (2018).
- [64] 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).
- [65] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Cross-correlation search for periodic gravitational waves, Phys. Rev. D 77, 082001 (2008).
- [66] C. T. Y. Chung, A. Melatos, B. Krishnan, and J. T. Whelan, Designing a cross-correlation search for continuous-wave gravitational radiation from a neutron star in the supernova remnant SNR 1987A, Mon. Not. R. Astron. Soc. 414, 2650 (2011).
- [67] J. T. Whelan, S. Sundaresan, Y. Zhang, and P. Peiris, Modelbased cross-correlation search for gravitational waves from Scorpius X-1, Phys. Rev. D 91, 102005 (2015).
- [68] E. Goetz and K. Riles, An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems, Classical Quant. Grav. 28, 215006 (2011).
- [69] G. D. Meadors, E. Goetz, and K. Riles, Tuning into Scorpius X-1: adapting a continuous gravitational-wave search for a known binary system, Classical Quant. Grav. 33, 105017 (2016).
- [70] S. W. Ballmer, A radiometer for stochastic gravitational waves, Classical Quant. Grav. 23, S179 (2006).
- [71] B. Abbott *et al.*, Upper limit map of a background of gravitational waves, Phys. Rev. D 76, 082003 (2007).
- [72] J. Abadie *et al.*, Directional Limits on Persistent Gravitational Waves Using LIGO S5 Science Data, Phys. Rev. Lett. **107**, 271102 (2011).
- [73] S. van der Putten, H. J. Bulten, J. F. J. van den Brand, and M. Holtrop, Searching for gravitational waves from pulsars in binary systems: An all-sky search, J. Phys. Conf. Ser. 228, 012005 (2010).
- [74] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.*, Directional limits on persistent gravitational waves using data from Advanced LIGO's first two observing runs, Phys. Rev. D **100**, 062001 (2019).
- [75] G. D. Meadors, E. Goetz, K. Riles, T. Creighton, and F. Robinet, Searches for continuous gravitational waves from Scorpius X-1 and XTE J1751-305 in LIGO's sixth science run, Phys. Rev. D 95, 042005 (2017).
- [76] L. Dunn, P. Clearwater, A. Melatos, and K. Wette, Graphical processing unit implementation of the F-statistic for

continuous gravitational wave searches, Classical Quant. Grav. (to be published).

- [77] C. R. D'Angelo and H. C. Spruit, Episodic accretion on to strongly magnetic stars, Mon. Not. R. Astron. Soc. 406, 1208 (2010).
- [78] B. Haskell, M. Priymak, A. Patruno, M. Oppenoorth, A. Melatos, and P. D. Lasky, Detecting gravitational waves

from mountains on neutron stars in the advanced detector era, Mon. Not. R. Astron. Soc. **450**, 2393 (2015).

- [79] D. I. Jones, Gravitational wave emission from rotating superfluid neutron stars, Mon. Not. R. Astron. Soc. 402, 2503 (2010).
- [80] R. Bellman, *Dynamic programming* (Princeton University Press, Princeton, New Jersey, 1957), Vol. 70, p. 342.