Directed search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model

Dana Jones[1] and Ling Sun[2]

[1] *College of Arts and Sciences, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*
[2] *LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*

(Dated: November 1, 2019)

Results are presented from a semicoherent search for continuous gravitational waves from a nearby neutron star candidate, Fomalhaut b, using data collected in the second observing run of Advanced LIGO. The search is based on a hidden Markov model scheme, capable of tracking signal frequency evolution from the star's secular spin down and stochastic timing noise simultaneously. The scheme is combined with a frequency domain matched filter ($\mathcal{F}$-statistic), calculated coherently over 5-day time stretches. The frequency band 100–1000 Hz is searched. Any candidates surviving a hierarchy of vetoes will be considered for further scrutiny. We present the 95% confidence frequentist strain upper limits assuming that no detection is claimed in this search.

I. INTRODUCTION

Gravitational waves (GWs), perturbations in spacetime that propagate at the speed of light, were first directly observed in 2015 when the Laser Interferometer Gravitational-Wave Observatory (LIGO) detected a merging binary black hole system (GW150914) [1, 2]. The Virgo detector joined the observation at the end of the second observing run in 2017 [3]. In the years since the first detection, the sensitivity of these interferometers has been greatly improved, allowing for increasingly frequent detections of compact binary coalescences (CBCs) [4]. However, there are many other types of GW sources that also radiate with frequencies that fall within the observational band of ground-based interferometers that are worth studying. In particular, for this project, we will focus on persistent, well modelled continuous gravitational waves (CWs) produced by isolated spinning neutron stars. These CWs, if detected, will provide invaluable information regarding the structure of the neutron star sources as well as the nuclear equation of state in extreme-pressure situations [5]. Therefore, a great deal of work is being done to improve methods with which to detect CWs (e.g., Refs. [6–8]).

One of the most promising types of CW sources are young neutron stars that are born of core collapse supernovae. Young sources are of particular interest because they have had less time to settle out and thus are more likely to have certain nonaxisymmeteries. These nonaxisymmetries, coupled with the large spin-down also characteristic of young neutron stars, could produce detectable CWs [5, 9, 10]. Ref. [9] presents the latest search results for 15 supernova remnants and a nearby neutron star candidate, Fomalhaut b, in the first observing run of Advanced LIGO. We consult the list of sources in [9] and select the source we study here.

Fomalhaut b, one of the sources on the list, was originally hypothesized to be an exoplanet orbiting around the star Fomalhaut A [11]. However, certain peculiarities, namely its non-detection in the infrared and its potentially highly eccentric orbit, have led astronomers to speculate over whether or not it is in fact an exoplanet. Recently, a different hypothesis has been gaining momentum; rather than a companion object, Fomalhaut b may in fact be a background neutron star to Fomalhaut A [12]. This source has characteristics that could make it a less promising candidate for a CW search, namely its relatively old age and the remaining uncertainty regarding its identity. In spite of these things, however, Fomalhaut b remains a promising candidate due to its extremely close proximity. It is currently hypothesised to be just 11 pc away, which would make it, assuming the neutron star hypothesis to be true, the closest known neutron star to our solar system [12].

There are three types of searches for CW sources: targeted searches for pulsars whose sky positions and ephemerides are well measured electromagnetically, directed searches for neutron stars with known sky positions but unknown rotation frequencies, and all-sky searches, surveys done over the whole sky to search for emitting sources [5]. For this project, a directed search, more expensive than a targeted search but less expensive than an all-sky search, is conducted. Because the expected strain amplitudes of CWs are orders of magnitude smaller than those produced by CBCs, they are incredibly hard to detect without integrating data over very long periods of time. However, such integrations are very expensive and require vast computational resources [5]. Furthermore, there is an intrinsic, stochastic frequency wandering, or "timing noise," associated with these sources that may degrade the sensitivity of fully coherent searches [13]. Thus, although fully coherent searches are still of interest in certain systems with negligible timing noise, a semicoherent search can be of great use for sources in which the timing noise has significant effects [10].

This paper details the process of conducting a directed search of the neutron star candidate Fomalhaut b in the second Advanced LIGO observing run (O2) with the end goal of detecting a CW signal. The search is conducted using data from January 3, 2017 to August 25, 2017 in the frequency band 100–1000 Hz, divided into 1-Hz subbands to parallelize the computation. The data is split

into 5-day coherent segments in order to maintain search sensitivity in the presence of timing noise. The short Fourier transforms (SFTs) of the data are passed through a frequency domain matched filter (the $\mathcal{F}$-statistic). These coherent 5-day segments are then combined incoherently using a hidden Markov model (HMM), a tracking scheme equipped to deal with the evolving signal frequency. The tracking scheme has its origins in engineering, and it has recently been used in many CW searches to improve computational efficiency [8, 14, 15]. The output is then passed through a hierarchy of vetoes, and any candidates that remain will be further followed up in future studies.

The organization of the paper is as follows. Section II outlines the methods used in the search. Section III details the search setup, including source parameters, search configuration, threshold, and sensitivity estimates. In Section IV, we explain the four-step veto process performed on the over-threshold candidates, and discuss future work to follow up on any remaining candidates. Finally, the conclusion is given in Section V.

II. METHODS

This search is composed of two main procedures: (1) coherently summing up the signal power over 5-day time stretches using the $\mathcal{F}$-statistic, and (2) a HMM tracking scheme finding the most probable signal evolution path over the total observing run. The following subsections briefly describe the signal model and the methods.

A. Continuous Wave Signal

The phase of the signal as observed in the detector can be described as

$$\Phi(t) = 2\pi \sum_{k=0}^{s} \frac{f_0^{(k)} t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \hat{n} \cdot \vec{r}(t) \sum_{k=0}^{s} \frac{f_0^{(k)} t^k}{k!}, \quad (1)$$

where $f_0^{(k)}$ is the k-th time derivative of the signal frequency at $t = 0$, $\hat{n}$ is the unit vector directed outward from the solar system barycenter (SSB) to the neutron star, and $\vec{r}(t)$ is the position vector of the detector retative to the SSB. Then, the signal can be expressed as

$$h(t) = \mathcal{A}^\mu h_\mu(t), \quad (2)$$

where $\mathcal{A}^\mu$ represents the amplitude associated with each of four linearly independent components $h_\mu(t)$ that depend on the phase given above [10]. (See Eqs. 3-6 from Ref. [10] for the four components.)

B. $\mathcal{F}$-statistic

The $\mathcal{F}$-statistic is used to estimate the likelihood that the signal described above is present in the frequency domain data. The time-domain data $x(t)$ collected by the detector can be written as

$$x(t) = \mathcal{A}^\mu h_\mu(t) + n(t) \quad (3)$$

where $n(t)$ is stationary, additive noise [16]. First, a scalar product as a sum over single-detector inner products is defined as

$$(x|y) = \sum_X (x^X|y^X) \quad (4)$$

$$= \sum_X 4\Re \int_0^\infty df \frac{\tilde{x}^X(f)\tilde{y}^{X*}(f)}{S_h^X(f)}, \quad (5)$$

where X indexes the detector, $S_h^X(f)$ is the single-sided power spectral density of detector X, the tilde denotes a Fourier transform, and $\Re$ is the real part of a complex number [17]. Then, the $\mathcal{F}$-statistic can be expressed in the form

$$\mathcal{F} = \frac{1}{2} x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (6)$$

where $x_\mu = (x|h_\mu)$, and $\mathcal{M}^{\mu\nu}$ represents the matrix inverse of $\mathcal{M}_{\mu\nu} = (h_\mu|h_\nu)$ [10]. If the noise is Gaussian and the single-sided power spectral density, $S_h(f)$, is the same in all detectors, the probability of detection can be written such that it solely depends on one term: the signal-to-noise ratio, given by

$$\rho_0^2 = \frac{K h_0^2 T_{\text{coh}}}{S_h(f)}, \quad (7)$$

where K is a constant that depends on the sky location, orientation of the source, and number of detectors, h_0 is the characteristic gravitational-wave strain, and T_{coh} is the length of each coherent piece of data [16].

C. Hidden Markov Model

This semicoherent search is based on the HMM scheme solved by the Viterbi algorithm to estimate the most likely frequency evolution path of the signal. (See Ref. [18] for an explanation of the classic Viterbi algorithm.) It is computationally efficient and robust in the presence of timing noise [10].

A Markov chain is a stochastic process that transitions from one discrete state to another at discrete times. A hidden Markov chain is comprised of two variables: the state variable $q(t) \in \{q_1, \cdots, q_{N_Q}\}$, which is unobservable (in other words, it is hidden) and where N_Q is the total number of hidden states, and the measurement variable $o(t) \in \{o_1, \cdots, o_{N_O}\}$, which is observable and where N_O is the total number of observable states. For any time t_{n+1}, the hidden state is solely dependent on the state at time t_n and has a transition probability of

$$A_{q_j q_i} = \Pr[q(t_{n+1}) = q_j | q(t_n) = q_i]. \quad (8)$$

The hidden state q_i, present in the observable state o_j, has an emission probability defined as

$$L_{o_j q_i} = \Pr[o(t_n) = o_j | q(t_n) = q_i]. \tag{9}$$

The prior is written as

$$\Pi_{q_i} = \Pr[q(t_0) = q_i]. \tag{10}$$

The probability that an observed sequence $O = o(t_0), ..., o(t_{N_T})$, where N_T is the total number of time steps, is the result of a hidden state path $Q = q(t_0), ..., q(t_{N_T})$ via a Markov chain can be described by

$$\begin{aligned}\Pr(Q|O) =& L_{o(t_{N_T})q(t_{N_T})} A_{q(t_{N_T})q(t_{N_T-1})} \cdots L_{o(t_1)q(t_1)} \\ & \times A_{q(t_1)q(t_0)} \Pi_{q(t_0)}.\end{aligned} \tag{11}$$

The most probable path, calculated by maximizing $\Pr(Q|O)$ is

$$Q^*(O) = \arg\max \Pr(Q|O), \tag{12}$$

where arg max returns the argument that maximizes the function [10].

In this search, the one-dimensional state variable $q(t)$ is defined as $f_0(t)$. The discrete hidden states are mapped one-to-one to the frequency bins that make up the output of $\mathcal{F}(f)$ calculated over the span of length $T_{\rm coh}$ (see Sec. II B). Bin size is set such that $\Delta f = 1/(2T_{\rm coh})$. We choose $T_{\rm coh}$ to satisfy

$$\left| \int_t^{t+T_{\rm coh}} dt' \dot{f}_0(t') \right| < \Delta f \tag{13}$$

for $0 < t < T_{\rm obs} - T_{\rm coh}$, where $T_{\rm obs}$ is the total observation time. (See Ref. [10] for more details).

Assuming that the impact of timing noise is significantly smaller than that of secular spin down and that $\dot{f}_0$ follows a uniform distribution between zero and the maximum estimated frequency derivative $|\dot{f}_0|_{\rm max}$, and substituting $|\dot{f}_0|_{\rm max}$ into the above relation, Eqn. (8) becomes simply

$$A_{q_{i-1}q_i} = A_{q_i q_i} = \frac{1}{2}. \tag{14}$$

(All other entries here are zero.) Then, using the definition of the frequency domain estimator, the emission probability is defined as

$$\begin{aligned} L_{o(t)q_i} &= \Pr[o(t)|f_i \le f_0(t) \le f_i + \Delta f] && (15) \\ &\propto \exp[\mathcal{F}(f_i)], && (16) \end{aligned}$$

from t to $t+T_{\rm coh}$, where f_i represents the central frequency in the i-th bin. A uniform prior of $\Pi_{q_i} = N_Q^{-1}$ is selected because there is no independent knowledge of f_0 [10].

The algorithm outputs the most likely pathway of the frequency evolution $Q^*(O)$ over the course of $T_{\rm obs}$. This is called the Viterbi path and consists of a precise frequency measured at each time step. The Viterbi path has an associated Viterbi score S defined such that the log likelihood of the optimal Viterbi path is equal to the mean log likelihood of all possible paths plus S standard deviations at final step N_T. This is shown as follows:

$$S = \frac{\ln \delta_{q^*}(t_{N_T}) - \mu_{\ln\delta}(t_{N_T})}{\sigma_{\ln\delta}(t_{N_T})} \tag{17}$$

where

$$\mu_{\ln\delta}(t_{N_T}) = N_Q^{-1} \sum_{i=1}^{N_Q} \ln \delta_{q_i}(t_{N_T}) \tag{18}$$

and

$$\sigma_{\ln\delta}(t_{N_T})^2 = N_Q^{-1} \sum_{i=1}^{N_Q} [\ln \delta_{q_i}(t_{N_T}) - \mu_{\ln\delta}(t_{N_T})]^2. \tag{19}$$

Here, $\delta_{q_i}(t_{N_T})$ is the maximum probability of the path that ends in state q_i $(1 \le i \le N_Q)$ at step N_T and $\delta_{q^*}(t_{N_T})$ is the likelihood of the optimal Viterbi path (the probability $\Pr[Q^*(O)|O]$). If the Viterbi score S is significantly larger than the scores computed for all other possible frequency evolution pathways, it is more likely that a signal is present [10].

Figure 1 shows the results of tracking a synthetic signal with $h_0 = 8 \times 10^{-25}$ injected into Gaussian noise [amplitude spectral density (ASD) $S_h^{1/2} = 4 \times 10^{-24}\,{\rm Hz}^{-1/2}$] at 151.23 Hz, a frequency chosen randomly. The total observing time searched is 40 days, with $T_{\rm coh} = 1$ d. The frequency derivative $\dot{f}_0$ of the injected signal is set to $-1 \times 10^{-11}\,{\rm Hz\,s^{-1}}$. The signal evolution path reproduced by the HMM algorithm is plotted over the original signal path to show the accuracy with which the HMM algorithm is able to track signals, with a root-mean-square error 1.15×10^{-6} Hz between the two paths. The reproduced signal path is displayed in a stair-step pattern because the HMM method uses discrete frequency bins (shown as dashed lines). The injected signal, by contrast, evolves continuously. The HMM estimate matches the true signal well in that it never strays more than one bin from the true signal.

III. SETUP

A. Source Parameters

Fomalhaut b is located at 225739.1–293720 (J2000). In Ref. [9], Fomalhaut b was searched using both the most optimistic and pessimistic estimates of age and distance found in the literature: one search assumed a distance of 0.011 kpc and an age of 316 kyr, while the other assumed a distance of 0.02 kpc and an age of 3000 kyr. Both sets of estimates are taken into consideration when choosing an optimal value for $T_{\rm coh}$ in this search.

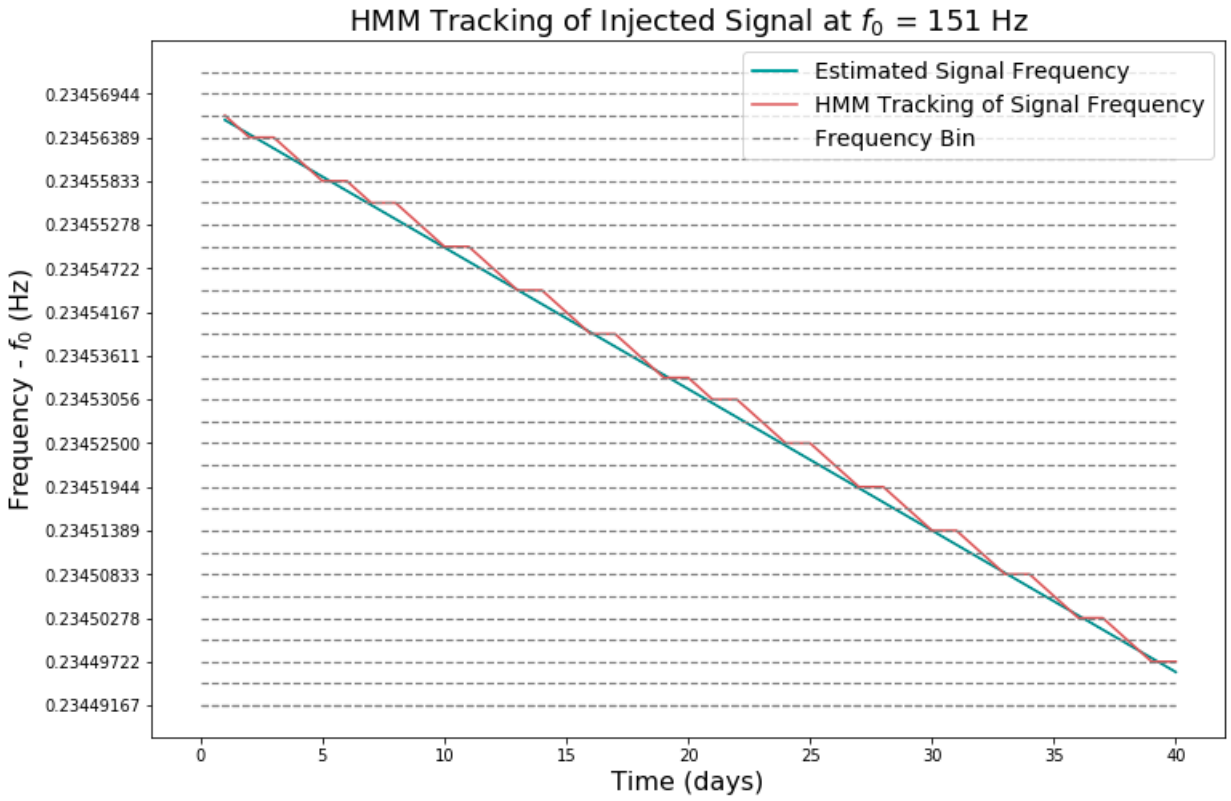

FIG. 1. HMM tracking of a synthetic signal with $h_0 = 8 \times 10^{-25}$ injected into Gaussian noise (ASD $S_h^{1/2} = 4 \times 10^{-24}\,\mathrm{Hz}^{-1/2}$) at a starting frequency of 151.23 Hz. The blue curve represents the injected signal (with $\dot{f}_0 = -1 \times 10^{-11}\,\mathrm{Hz\,s^{-1}}$), and the red curve represents the optimal Viterbi path reproduced by the HMM tracking. The dashed lines indicate the size of the frequency bins.

Fomalhaut b's first frequency derivative, though not directly known, is estimated using [9, 10]

$$-\frac{f_0}{(n_{\min}-1)t_{\mathrm{age}}} \leq \dot{f}_0 \leq -\frac{f_0}{(n_{\max}-1)t_{\mathrm{age}}}, \tag{20}$$

where t_{age} is the age of the source, and $n_{\min}$ and $n_{\max}$ represent the minimum and maximum breaking indices n, respectively. T_{coh} may then be chosen based on the following equation:

$$T_{\mathrm{coh}} \leq (2|\dot{f}_0|)^{-1/2}, \tag{21}$$

such that Eqn. (13) is satisfied.

Table I shows the range of estimates for $\dot{f}_0$, calculated using Eqn. (20), and T_{coh}, calculated using Eqn. (21). Here, the breaking index is assumed to be within the range $2 \leq n \leq 7$, so $\dot{f}_0$ and T_{coh} are calculated for three possible values within this range, 2, 5, and 7. t_{age} is estimated to be either 316 kyr or 3000 kyr. The frequency range considered is 100–1000 Hz. Although a larger band could be used, this band is where the detectors are most sensitive, and it is what is searched.

B. Search Configuration

The observation period used for this project is the duration of LIGO's O2 observing run, excluding the first few weeks when the data quality was less than optimal and was followed by a brief end-of-the-year break; that is, a GPS time from 1167545066 to 1187733592 (Jan. 03 to Aug. 25, 2017), a total duration of about 234 days. Based on the estimated ranges of $\dot{f}_0$ and T_{coh} (Table I), we choose $T_{\mathrm{coh}} = 5\,\mathrm{d}$. Plugging this into the relationship outlined in Eqn. (21), the $\dot{f}_0$ range covered in this search is $0 \leq |\dot{f}_0|/(\mathrm{Hz\,s^{-1}}) \leq 2.68 \times 10^{-12}$. We have the total number of steps $N_T = 47$.

C. Threshold

The Viterbi score threshold S_{th}, corresponding to a given false alarm probability (1% in this search), represents the Viterbi score needed such that the signal present in the data only has a 1% chance of having been falsely identified. To find this threshold, a series of 200 realizations of pure Gaussian noise are simulated and searched with the same configuration as the real search (i.e., $T_{\mathrm{coh}} = 5\,\mathrm{d}$ and $T_{\mathrm{obs}} = 234\,\mathrm{d}$) for each of five different

Breaking Index	Age (kyr)	$\lvert\dot{f}_0\rvert$ ($\mathrm{Hz\,s^{-1}}$)	T_{coh} (days)
2	316	$[1.00\times10^{-11}, 1.00\times10^{-10}]$	[0.82, 2.58]
	3000	$[1.06\times10^{-12}, 1.06\times10^{-11}]$	[2.52, 7.96]
5	316	$[2.51\times10^{-12}, 2.51\times10^{-11}]$	[1.63, 5.17]
	3000	$[2.64\times10^{-13}, 2.64\times10^{-12}]$	[5.04, 15.93]
7	316	$[1.67\times10^{-12}, 1.67\times10^{-11}]$	[2.00, 6.33]
	3000	$[1.76\times10^{-13}, 1.76\times10^{-12}]$	[6.17, 19.51]

TABLE I. Ranges of $\dot{f}_0$ and T_{coh} estimated for Fomalhaut b. These are estimated for three different breaking indices and two different age estimates. The frequency range considered is from 100 Hz to 1000 Hz.

1-Hz frequency sub-bands, starting from: 100 Hz, 300 Hz, 500 Hz, 700 Hz, and 900 Hz. The resulting scores are compiled and sorted and then the score at the 99th percentile is chosen, yielding $S_{\mathrm{th}} = 6.95$. Scores above this would be in the top one percent of the data, corresponding to the 1% false alarm probability.

This threshold is verified using interferometric data from O2. The above process is repeated in three 1-Hz sub-bands, starting from: 300 Hz, 600 Hz, and 900 Hz. This yields $S_{\mathrm{th}} = 6.97$, in good agreement with the threshold estimated using Gaussian noise. Hence we set $S_{\mathrm{th}} = 6.95$ in this search.

D. Sensitivity

The h_0 that yields a 95% detection efficiency, denoted by $h_0^{95\%}$, is obtained by injecting a series of 100 synthetic signals into Gaussian noise (ASD $S_h^{1/2} = 4\times10^{-24}\,\mathrm{Hz}^{-1/2}$) in four 1-Hz sub-bands, starting from 155 Hz, 355 Hz, 555 Hz, and 755 Hz. If more than 95% of the signals are recovered, this process is repeated with a smaller value of h_0, and vice versa. These simulations are marginalized over the orientation of the source. We obtain an average of $h_0^{95\%} = 1.56\times10^{-26}$ for these four sub-bands [10].

This process is repeated in real O2 data in a sample 1-Hz sub-band 715–716 Hz, chosen because it is a relatively clean band. This yields an upper limit of $h_0^{95\%} = 1.67\times10^{-25}$, an order of magnitude larger than $h_0^{95\%}$ obtained in Gaussian noise. This is because (1) the O2 noise level has not reached the design sensitivity, and (2) the duty cycle (i.e., the proportion of time that the data is in analyzable science mode) of O2 is only about 50%. In this sample sub-band, $h_0^{95\%}$ obtained from O2 data is slightly better than the existing upper limits presented in Ref. [9] (i.e. $h_0^{95\%} = 2\times10^{-25}$ at 715 Hz). The sensitivity in the full frequency band remains a work in progress.

IV. VETO PROCESS

We record the search results in the Appendices for the moment given that the analysis has not yet been fully completed. The results of the veto process that outlined below will be added in a future publication once the study is entirely finished. All candidates with $S > S_{\mathrm{th}}$ are followed by a four-stage veto process to eliminate the ones resulting from noise artifacts [14]. This section details the methods that have been implemented to eliminate candidates so far.

First, we eliminate any candidates whose Viterbi paths intersect any known instrumental lines present in either the Hanford or Livingston interferometers.

Second, additional candidates are vetoed due to the contamination of artifacts in a single detector. The criteria is as follows. For each candidate, if searching data from a single interferometer yields $S \geq S_{\cup}$, where $S_{\cup}$ is the original score obtained with both interferometers combined, while searching the other interferometer yields $S < S_{\mathrm{th}}$, and if the Viterbi path from the interferometer with $S \geq S_{\cup}$ intersects the original path, that candidate is vetoed.

Third, further candidates are vetoed by splitting the observation time into two halves and analyzing each half separately. A candidate is vetoed if the score obtained by searching one half of T_{obs} is higher than or equal to $S_{\cup}$ (corresponding to the full T_{obs}) but the score obtained in the other half is lower than S_{th}, and if the Viterbi path from the half with the higher score intersects the original path.

Fourth, the remaining candidates are further followed up by increasing T_{coh}. We increase T_{coh} from 5 days to 10 days for a candidate whose $\dot{f}_0$ is small enough such that Eqn. (21) holds true when $T_{\mathrm{coh}} = 10\,\mathrm{d}$. Then, for each candidate, the $\mathcal{F}$-statistic and the Viterbi path are recomputed. If the candidate is from astrophysical origin, the resulting Viterbi score with $T_{\mathrm{coh}} = 10\,\mathrm{d}$ should be higher than $S_{\cup}$, and the Viterbi paths of the original run and the run with the increased T_{coh} should match. Hence we veto a candidate if its score S corresponding to $T_{\mathrm{coh}} = 10\,\mathrm{d}$ falls below $S_{\cup}$ and if the Viterbi paths from these two separate runs are at least 0.01 Hz apart.

Any candidates that remain will be further followed up using a number of additional methods. Further scrutiny includes a deeper $\mathcal{F}$-statistic search estimating both f_0 and $\dot{f}_0$ in the remaining 1-Hz bands with longer T_{coh}, searching the data collected in the first and third observing

runs of Advanced LIGO, and using other CW search methods.

If no detection is claimed, we place an upper limit on the signal strain at the 95% confidence level, $h_0^{95\%}$, as a function of frequency. This remains a work in progress in the full band. The sample upper limit obtained in O2 in the 1 Hz sub-band starting from 715 Hz is $h_0^{95\%} = 1.67 \times 10^{-25}$, as detailed in Sec. III D.

V. CONCLUSION

We detail the methodology for a directed search for CW signals from the neutron star candidate Fomalhaut b in the frequency band 100–1000 Hz, using Advanced LIGO O2 data. An efficient, semicoherent search method, based on a HMM tracking scheme and a matched filter $\mathcal{F}$-statistic, is used to track the signal frequency. The above-threshold candidates are passed through a four-step veto procedure. Any candidate surviving these vetoes should also be cross-checked in more sensitive observing runs, O3 and beyond. The follow-ups of the candidates are underway. If all candidates are eventually eliminated, upper limits can be placed on h_0 with 95% confidence. This remains a work in progress. In a sample sub-band, 715–716 Hz, we obtain $h_0^{95\%} = 1.67 \times 10^{-25}$, assuming there is no detection.

In the future, the search may benefit from new information gathered from electromagnetic observations of the source, as well as the further upgrading of the instruments. Regardless of the eventual outcome of this search, this methodology represents an important development in the search for CWs and should be extended to other isolated neutron star sources.

VI. ACKNOWLEDGEMENTS

The authors would like to thank the LIGO Laboratory for providing the resources with which to conduct this search, as well as Alan Weinstein and the LIGO SURF program, the National Science Foundation, and the California Institute of Technology for sponsoring the project. L. Sun is 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–0757058. Advanced LIGO was built under award PHY–0823459.

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Observation of gravitational waves from a binary black hole merger," Phys. Rev. Lett. **116**, 061102 (2016).
[2] J. Aasi *et al.* (LSC), "Advanced LIGO," Classical and Quantum Gravity **32**, 074001 (2015).
[3] F. Acernese *et al.* (Virgo), "Advanced Virgo: a second-generation interferometric gravitational wave detector," Classical and Quantum Gravity **32**, 024001 (2015).
[4] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs," Phys. Rev. X **9**, 031040 (2019).
[5] K. Riles, "Recent searches for continuous gravitational waves," Mod. Phys. Lett. A **32**, 1730035 (2017).
[6] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "All-sky search for continuous gravitational waves from isolated neutron stars using advanced ligo o2 data," Phys. Rev. D **100**, 024004 (2019).
[7] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Narrow-band search for gravitational waves from known pulsars using the second ligo observing run," Phys. Rev. D **99**, 122002 (2019).
[8] B. P. Abbott *et al.*, "Search for gravitational waves from scorpius x-1 in the second advanced ligo observing run with an improved hidden markov model," (2019), arXiv:1906.12040 [gr-qc].
[9] B. P. Abbott *et al.*, "Searches for continuous gravitational waves from 15 supernova remnants and fomalhaut b with advanced LIGO," The Astrophysical Journal **875**, 122 (2019).
[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] T. Currie *et al.*, "Direct imaging confirmation and characterization of a dust-enshrouded candidate exoplanet orbiting Fomalhaut," The Astrophysical Journal **760**, L32 (2012).
[12] R. Neuhäuser, M. M. Hohle, C. Ginski, J. G. Schmidt, V. V. Hambaryan, and T. O. B. Schmidt, "The companion candidate near Fomalhaut – a background neutron star?" Monthly Notices of the Royal Astronomical Society **448**, 376–389 (2015).
[13] G. Hobbs, A. G. Lyne, and M. Kramer, "An analysis of the timing irregularities for 366 pulsars," Monthly Notices of the Royal Astronomical Society **402**, 1027–1048 (2010).
[14] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Search for gravitational waves from scorpius x-1 in the first advanced ligo observing run with a hidden markov model," Phys. Rev. D **95**, 122003 (2017).
[15] Ling Sun and Andrew Melatos, "Application of hidden markov model tracking to the search for long-duration transient gravitational waves from the remnant of the binary neutron star merger gw170817," Phys. Rev. D **99**, 123003 (2019).
[16] P. Jaranowski, A. Królak, and B. F. Schutz, "Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection," Physical Review D **58**, 063001 (1998).
[17] R. Prix, "Search for continuous gravitational waves: Metric of the multidetector F-statistic," Physical Review D **75**, 023004 (2007).

[18] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm,” IEEE Transactions on Information Theory **13**, 260–269 (1967).