Search technique to observe precessing compact binary mergers in the advanced detector era

Connor McIsaac, Charlie Hoy, and Ian Harry
University of Portsmouth, Portsmouth, PO1 3FX, United Kingdom

(Received 11 April 2023; accepted 22 November 2023; published 11 December 2023)

Gravitational-wave signals from compact binary coalescences are most effectively identified by matched filter searches. These searches match the data against a pregenerated bank of gravitational-wave templates. Currently, all modeled gravitational-wave searches use templates that restrict the component spins to be aligned (or antialigned) with the orbital angular momentum. This means that they are less sensitive to gravitational-wave signals from precessing binaries, implying that a significant fraction of signals may remain undetected. In this work, we introduce a matched filter search that is sensitive to signals generated from precessing binaries. We take advantage of the fact that a gravitational-wave signal from a precessing binary can be decomposed into a power series of five harmonics. This allows us to create a generic-spin template bank that is only ~3 times larger than existing aligned-spin banks. Our new search shows a ~100% increase in sensitive volume for neutron star black hole binaries with total mass larger than $17.5M_{\odot}$ and in-plane spins > 0.67, and improves sensitivity by ~60% on average across the full generic spin neutron-star black-hole parameter space. In addition, our generic spin search performs as well as existing aligned-spin searches for neutron star black hole signals with insignificant in-plane spins. We anticipate that this improved technique will identify significantly more gravitational-wave signals and help shed light on the unknown spin distribution of binaries in the Universe.

DOI: 10.1103/PhysRevD.108.123016

I. INTRODUCTION

In the last 10 years, gravitational-wave observatories such as Advanced LIGO, Advanced Virgo, and KAGRA have unlocked the gravitational-wave Universe [1–3]. At the time of writing, roughly 100 compact binary coalescences have been observed using data from these observatories by the LIGO-Virgo-KAGRA (LVK) collaborations [4–7]. In addition, the public release of LVK data via the Gravitational Wave Open Science Center [8] has enabled external groups to analyze the data and identify additional events [9–14].

These many observations have been made possible by the development of complex search algorithms to matched filter the gravitational-wave data against a set of filter waveforms representing our best knowledge of the gravitational-wave signal emitted by compact binary mergers [10,15–27]. However, while these searches have been undeniably successful, they are all limited in one important regard: They are all restricted to performing modeled searches with template banks that only contain aligned-spin templates, where the spin angular momentum of the two compact objects and the orbital angular momentum of the binary are aligned.[1] This is because the sky position and orientation of aligned-spin binaries have a simple relationship with the waveform observed by the gravitational wave detector, which allows us to analytically maximize over these extrinsic parameters [15,17].

In this work, we revisit the problem of searching for compact binary coalescences where the spins are misaligned with the orbital angular momentum. When the spins are misaligned, spin-orbit coupling will cause the orbital angular momentum and the spin angular momenta to precess around the direction of the total angular momentum [31]. This effect will modulate the phase and amplitude of the observed gravitational waves, with the exact form of the modulation depending on the orientation and sky position of the binary. This dependence means that we can no longer analytically maximize over the extrinsic

[1]We note that in addition to matched-filter searches, unmodeled, or semimodeled, pipelines have been developed to target the observation of sources that do not match well to our template waveforms [28–30]. However, while these can match the sensitivity of matched-filtering when searching for high-mass signals, the performance of such searches does not match matched-filtering when considering systems with relatively low chirp mass, as we will do in this work.

parameters, making the inclusion of precessional effects in a modeled search challenging.

There are two main formation channels for the production of compact binaries: through the isolated evolution of a pair of massive stars, or through the dynamical formation of binaries in dense stellar environments [32–35]. For compact binary coalescences that evolved from a pair of massive binary stars, we expect the spin angular momenta of the components to be roughly aligned with the orbital angular momentum, with some misalignment present due to kicks caused during core collapse of either component [36,37]. For binaries that are formed through dynamic capture in dense stellar environments, we expect to see isotropic spin distributions [38]. Observation of compact binary coalescences with large misaligned spins, or lack thereof, will therefore allow us to test the rate of binaries produced by different formation channels [39,40] and is therefore of great astrophysical interest. However, if current search methods are missing strongly precessing signals then this could introduce a bias to these measurements.

When using only aligned-spin templates it has been demonstrated that reasonable sensitivity is retained to templates with moderate effects due to precession [41,42], losing $\sim$17–23% of our sensitivity in the neutron-star–black-hole parameter space compared to an ideal search [41]. Aligned-spin templates are particularly effective when the component masses are close to equal, the magnitude of the spin angular momentum is small, or the orbit of the binary is close to face-on ($\iota = 0$) or face-off ($\iota = \pi$). In these cases the precession of the binary will only weakly effect the waveform, making it difficult to infer the presence or absence of precession in the observed signal [43]. For most individual events observed so far there are only weak constraints on the size of the misaligned spin components [5,7]. Recently, strong evidence for precession has been claimed in one observed compact binary coalescence—GW200129_065458 [44]. However, there may be some uncertainty in this measurement due to non-Gaussian noise at the time of the event and the uncertainty in the glitch model used to remove this noise [45]. In short, most of the compact binary mergers observed so far show no evidence for precession. While this is likely reflective of the underlying population, it is possible that precessing signals remain undetected in our data because we are only looking with aligned-spin waveforms.

Several methods have been proposed to search for precessing signals [46–53]. However, none of these methods has been applied to Advanced LIGO, Advanced Virgo and KAGRA data. This is because the methods either result in an increased rate of noise events that outweighs an increase in recovering precessing signals, or because they are computationally unfeasible.

In this work we introduce a new method that is similar in nature to [48] but uses the harmonic decomposition proposed in [54]. We show that we can minimize the unphysical freedom introduced by the maximization over extrinsic parameters by using a subset of the available harmonics, while still recovering the majority of signal power from precessing events. We demonstrate that a precessing bank containing only 355160 templates not only increases the observed signal-to-noise ratio for precessing injections compared to an aligned-spin bank, but also increases the sensitive volume by $\sim$100% for binaries with total mass larger than $17.5 M_\odot$ and in-plane spins > 0.67.

We will begin by reviewing the effect of precession on the evolution of compact binary coalescences and their signals in Sec. II. We review previous attempts to search for precessing compact binary mergers in Sec. III. We will then review the harmonic decomposition in Sec. IV and its use in modeling the precessing signal for different sky positions and orientations. In Sec. V we motivate why the harmonic decomposition offers a way to solve the precessing search problem. In Sec. VI we introduce our new modeled search using the harmonic decomposition to maximize over the set of intrinsic parameters. Finally, in Sec. VII we demonstrate that with an appropriate choice of detection statistic we can improve the sensitivity of modeled searches to neutron-star–black-hole signals by $\sim$60% on average across the full generic spin parameter space.

II. GRAVITATIONAL-WAVE SIGNALS OF PRECESSING BINARIES

In this section we provide an overview of the gravitational-wave signals produced by precessing binaries. For further details, we refer the reader to Refs. [31,46,48,55].

A binary consisting of two compact objects will slowly inspiral due to the emission of gravitational waves. The emitted gravitational waves carry away angular momentum from the binary along the direction of the orbital angular momentum $\boldsymbol{L}$. Assuming a quasicircular orbit, the binary's evolution can be fully described by 8 parameters: the masses, m_1 and m_2, and the spin angular momentum vector, $\boldsymbol{S}_1$ and $\boldsymbol{S}_2$, of each compact object.

If $\boldsymbol{S}_1$ and/or $\boldsymbol{S}_2$ are nonzero and aligned or antialigned with $\boldsymbol{L}$ we refer to the system as an "aligned-spin binary." In an aligned-spin binary, spin-orbit, and spin-spin couplings alter the rate of inspiral of the binary, adding a contribution to the overall phase of the observed signal, as well as the amplitude of the emitted gravitational waves [55,56]. If the total spin angular momentum $\boldsymbol{S} = \boldsymbol{S}_1 + \boldsymbol{S}_2$ is misaligned with $\boldsymbol{L}$ the binary will additionally undergo spin-induced orbital precession [31]. For the case when $|\boldsymbol{L}| \ll |\boldsymbol{J}|$, $\boldsymbol{L}$, $\boldsymbol{S}_1$, and $\boldsymbol{S}_2$ precess around the approximately constant total angular momentum $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$ [31], as illustrated in Fig. 1. Although the emitted gravitational-waves continue to carry away angular momentum along $\boldsymbol{L}$, the precession of $\boldsymbol{L}$ around $\boldsymbol{J}$ implies that on average the angular momentum is emitted parallel to $\boldsymbol{J}$, with any emission orthogonal to $\boldsymbol{J}$ averaging to zero.

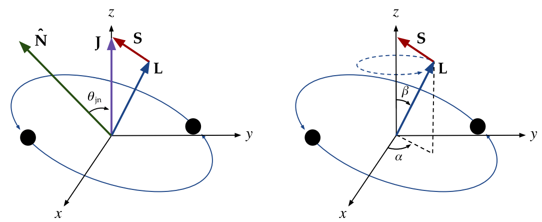

FIG. 1. Illustration of a binary with spins misaligned with the orbital angular momentum demonstrating the angles α, β, and θ_{JN}. In both panels the z axis is aligned with the total angular momentum. The solid blue arrow represents the orbital angular moment, the solid red arrow represents the combined spin angular momentum, and the solid purple arrow represents the total angular momentum. The dashed blue arrow illustrates the path of the orbital angular momentum vector as the binary processes. The solid green arrow shows the direction to the observer.

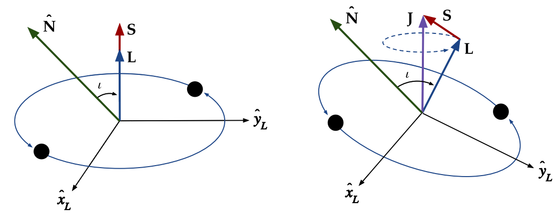

FIG. 2. Illustration of a binary with spins aligned (left) and misaligned (right) with the orbital angular momentum. The vectors $\hat{\boldsymbol{x}}_L$ and $\hat{\boldsymbol{y}}_L$ are defined to be orthogonal to $\boldsymbol{L}$ and form a basis for the instantaneous orbital plane of the binary. The solid blue arrow represents the orbital angular moment, the solid red arrow represents the combined spin angular momentum, and the solid purple arrow represents the total angular momentum. The dashed blue arrow illustrates the path of the orbital angular momentum vector as the binary processes. The solid green arrow shows the direction to the observer.

To simplify the modeling of a precessing binary, we utilize a source frame relative to the approximately constant vector $\hat{\boldsymbol{J}}$. We define this frame such that the z axis is parallel to the vector $\hat{\boldsymbol{J}}$ and the x axis is parallel to $\hat{\boldsymbol{J}} \times \hat{\boldsymbol{N}}$, where $\hat{\boldsymbol{N}}$ is the direction to the observer. The angle between the vectors $\hat{\boldsymbol{J}}$ and $\hat{\boldsymbol{N}}$ will be given by

$$\cos\theta_{\mathrm{JN}} = \hat{\boldsymbol{J}} \cdot \hat{\boldsymbol{N}}. \tag{1}$$

Following works such as [31] we then define two angles to track the precession of $\boldsymbol{L}$, the phase of the precession, α, and the opening angle β, both illustrated in Fig. 1. The opening angle is given by

$$\tan\beta = \frac{S_\perp}{|\boldsymbol{L}| + S_\parallel}, \tag{2}$$

where $S_\parallel$ and $S_\perp$ are the magnitudes of $\boldsymbol{S}$ parallel and perpendicular to $\boldsymbol{L}$ respectively.

As the precessing binary inspirals, $|\boldsymbol{L}|$ decreases while the magnitudes of the spin angular momenta $|\boldsymbol{S}_1|$ and $|\boldsymbol{S}_2|$ remain constant. Since $S_\parallel$ and $S_\perp$ also remain approximately constant throughout the inspiral, we see from Eq. (2) that the opening angle increases as the binary evolves. However, the rate of change of β will be small compared to the precession frequency $\Omega_p = \dot{\alpha}$ [57].

The orbital frequency Ω_{orb} of the binary is typically much larger than the precession frequency Ω_p, meaning that the binary can complete several orbits before $\boldsymbol{L}$ changes significantly [31]. For this case, the dynamics of the binary can be approximated as a set of quasicircular orbits within an orbital plane that is precessing. This condition is called the "adiabatic limit."

We can gain some insight into the effect of precession on the observed signal by defining an instantaneous orbital plane that is perpendicular to $\boldsymbol{L}$, and then modeling the dynamics within this plane using an aligned-spin waveform [48]. We can then examine how the observed signal changes as $\boldsymbol{L}$ changes.

We can define the vectors $\hat{\boldsymbol{x}}_L$ and $\hat{\boldsymbol{y}}_L$ to form a basis for the instantaneous orbital plane. Although $\hat{\boldsymbol{x}}_L$ and $\hat{\boldsymbol{y}}_L$ must be perpendicular to $\boldsymbol{L}$, we have the freedom to rotate them around $\boldsymbol{L}$. This rotation is degenerate with a change of the orbital phase, so we choose $\hat{\boldsymbol{x}}_L$ to be perpendicular to $\hat{\boldsymbol{J}}$, giving us

$$\hat{\boldsymbol{x}}_L = \frac{\hat{\boldsymbol{L}} \times \hat{\boldsymbol{J}}}{\sin\beta}, \qquad \hat{\boldsymbol{y}}_L = \frac{\cos\beta \hat{\boldsymbol{L}} - \hat{\boldsymbol{J}}}{\sin\beta}. \tag{3}$$

In the case that $\hat{\boldsymbol{L}}$ and $\hat{\boldsymbol{J}}$ are aligned, we will choose $\hat{\boldsymbol{x}}_L$ and $\hat{\boldsymbol{y}}_L$ to be aligned with the x and y axes of the source frame. These vectors are illustrated in Fig. 2.

For an aligned-spin binary inclined to the observer with angle ι, where ι is defined by

$$\cos\iota = \hat{\boldsymbol{L}} \cdot \hat{\boldsymbol{N}}, \tag{4}$$

the two polarizations of the emitted gravitational waves due to the dominant quadrupole can be written as

$$h_+(t) = \frac{1 + \cos^2\iota}{2r} A(t) \cos\left(2\phi(t) + 2\phi_0\right) \tag{5}$$

$$h_\times(t) = \frac{\cos\iota}{r} A(t) \sin\left(2\phi(t) + 2\phi_0\right). \tag{6}$$

Here the amplitude, $A(t)$, and orbital phase, $\phi(t)$, can be calculated using the post-Newtonian formalism [58]. The orbital phase is defined as the angle between the orbital separation vector, $\hat{\boldsymbol{r}}$, which points from m_1 to m_2, and the x axis

$$\hat{\boldsymbol{r}} = \hat{\boldsymbol{x}}_L \cos\phi(t) + \hat{\boldsymbol{y}}_L \sin\phi(t). \tag{7}$$

The detector observes a gravitational-wave, $h(t)$, which is a combination of the two polarizations,

$$h(t) = F_+(\Theta, \Phi, \Psi)h_+(t) + F_\times(\Theta, \Phi, \Psi)h_\times(t). \quad (8)$$

The detector's response to each polarization depends on the orientation of the detector, as defined by the angles Θ, Φ, and Ψ. The functions $F_+(\Theta, \Phi, \Psi)$ and $F_\times(\Theta, \Phi, \Psi)$ define the detector's response to each polarization. We can therefore write the observed gravitational-wave signal as

$$h(t) = \frac{1}{D_{\rm eff}} A(t) \cos(2\phi(t) + 2\phi_0), \quad (9)$$

where

$$D_{\rm eff} = r\left[F_+^2\left(\frac{1+\cos^2\iota}{2}\right)^2 + F_\times^2\cos^2\iota\right]^{-1/2} \quad (10)$$

and

$$2\phi_0 = 2\phi_c - \tan^{-1}\left(2\frac{F_\times}{F_+}\frac{\cos\iota}{1+\cos^2\iota}\right). \quad (11)$$

As $\boldsymbol{L}$ precesses around $\boldsymbol{J}$, the instantaneous orbital plane will rotate relative to the observer, and ι will become time-dependent. The overall amplitudes of h_+ and $h_\times$ are both dependent on ι, and are maximized when the binary is face-on ($\iota = 0$) and minimized when it is edge-on ($\iota = \pi/2$). A change in ι also changes the relative amplitudes of the two polarizations; the observed gravitational waves will be circularly polarized when face-on and linearly polarized when edge-on. This time-dependent change in ι will therefore produce a modulation effect on both the amplitude and phase of the observed signal.

The rate of change of the inclination angle will depend on the observer's viewing angle. For example, if the observer's line-of-sight is parallel to $\boldsymbol{J}$, the inclination angle will remain constant as $\boldsymbol{L}$ precesses around $\boldsymbol{J}$. However, if the observer's line-of-sight is initially parallel to $\boldsymbol{L}$, the inclination will oscillate between $\iota = 0$ and $\iota = 2\beta$ with a frequency of Ω_p. If β is large this will produce a strong modulation effect.

In the aligned-spin case, the gravitational wave phase is given by twice the orbital phase $\phi(t)$, as shown in Eqs. (5) and (6). This is simply the accumulated phase due to the orbital frequency $\Omega_{\rm orb}$. However, when the orbital plane is precessing, $\hat{\boldsymbol{x}}_L$ will also rotate around $\hat{\boldsymbol{L}}$. This means that the evolution of $\phi(t)$ becomes dependent on both the orbital and precession frequencies. This relationship is given by [59]

$$\dot{\phi} = \Omega_{\rm orb} - \Omega_p \cos\beta, \quad (12)$$

which introduces another modification to the phase of the observed signal.

III. PREVIOUS METHODS TO SEARCH FOR PRECESSING SIGNALS

In the case of an aligned-spin system, the binary can be parametrized by two component masses, m_1 and m_2, two spin magnitudes, s_{1z} and s_{2z}, two angles describing the orientation of the binary, (ι, ϕ_0), and three angles describing the orientation of the detector (Θ, Φ, Ψ). The two masses and spins determine the intrinsic evolution of the observed signal, while the five angles shift the observed waveform by constant amplitude, time and phase factors given by Eqs. (10) and (11). This allows us to analytically maximize over the five angles using a phase-maximized matched filter [15], leaving only the two masses and spins to maximize over. This is done using a large set of filter waveforms chosen to sufficiently cover the full range of masses and spins, which we refer to as a "template bank" [60,61].

In the case of a precessing binary we have six spin components, $\boldsymbol{S}_1$ and $\boldsymbol{S}_2$. We use three angles to define the orientation of the binary: two angles to define the initial orientation of the orbital plane, $(\theta_{\rm JN}, \alpha_0)$, where α_0 is the initial precession phase, and one angle to define the initial orbital phase, ϕ_0.[2] We also require the three angles describing the orientation of the detector (Θ, Φ, Ψ). In this case the two polarizations, h_+ and $h_\times$ have a more complex dependence on $\theta_{\rm JN}$ than just a phase and amplitude shift. The polarizations themselves are no longer related by a simple phase shift. Therefore, in addition to the four extra spin components, one also needs to consider the effect of $\theta_{\rm JN}$, Ψ, and ϕ_0 when developing a search to target precessing binaries.

Several methods have been proposed to tackle the problem of searching for precessing compact binary mergers. In [46,47] the authors introduce a small set of new parameters that modulate the phase of a nonprecessing waveform in order to mimic the effects caused by precession. However, adding these parameters to the template bank is computationally expensive and the resulting templates do not provide adequate match with precessing waveforms [62]. In [48] the authors extended this by adding a set of parameters that modulate the amplitude and phase of the waveform. The precessing signal is then expressed as a combination of modulated waveforms, each with a different amplitude and phase. The signal-to-noise ratio is then maximized analytically for the amplitude and phase of each modulated waveform, leaving only a few parameters to be added to the template bank. However, this approach allows for many unphysical combinations of amplitudes and phases, and the improvement in the match with precessing

[2] We note that α_0 is specified completely by $\boldsymbol{S}_1$ and $\boldsymbol{S}_2$ and the masses and is not an additional degree of freedom.

signals is outweighed by an increase in the observed noise due to the increased parameter space [63].

Another method was proposed in [48] and developed in [49,50]. This method models the binary and its precessional dynamics using a single spin $\boldsymbol{S}_1$ and only considers the $l = 2$, $|m| = 2$ harmonics in the instantaneous orbital plane. The gravitational wave polarizations in the radiation frame, h_+ and $h_\times$, are reexpressed as a combination of five waveforms, Q^I, using the $l = 2$ spherical harmonics as a basis. These five waveforms only depend on the intrinsic parameters of the binary and a constant time and phase offset. The observed signal is then expressed as a combination of the five waveforms, each multiplied by a coefficient, P_I, which is dependent on the sky position and orientation of the binary.

The sky location and orientation of the binary can be maximized over by maximizing over the coefficients P_I. The five components of P_I depend on six parameters in the case of a single detector, but the four parameters describing the location and orientation of the detector (Θ, Φ, Ψ, r) only enter as two independent combinations. This means that the five components of P_I depend on only four independent parameters and so they must be constrained somehow. In [49] a method was developed to maximize the signal-to-noise ratio over the constrained values of P_I, but this is computationally expensive. Instead, an unconstrained signal-to-noise ratio maximization of P_I is used, which allows the components to take unphysical values and increases the rate of noise triggers. The constraint problem becomes more difficult when we consider multiple detectors. For the 2-detector example we have the five P_I components measured independently in two detectors, but these 10 P_I values still depend on only 6 physical parameters. This means that if these values are not constrained considerable unphysical freedom is allowed. This method was extended to a targeted coherent search in [51]. The authors identified areas of the parameter space where precession effects were weak and restricted waveforms in these areas to only the dominant Q component, reducing the rate of noise triggers. This allowed an analysis to be targeted to specific areas of the parameter space. However, this method still faced issues due to the increase in the rate of noise triggers outweighing the increase in potential discovered signals.

In a third distinct approach, in [52], the authors proposed a search where templates are placed to cover all of the required parameters, including the two masses, six spin components and the inclination of the source.[3] Using this method for a single detector, the observed signal will always be consistent with a set of physical parameters. However, when we attempted to generate a template bank covering a physically meaningful parameter space for neutron-star–black-hole signals, 30 million templates were generated and the bank showed no sign of converging. Filtering gravitational wave data with this many templates, and constructing sets of filter waveforms of this size and larger, is computationally infeasible.

A common issue among most of these methods to search for precessing signals is that in order to properly model the observed signal, additional parameters must be added to our signal model in order to account for the effects of precession. However, any additional degrees of freedom in the signal model will not only increase the computational cost of the search, but will also increase the rate of noise triggers, even in the case of simple Gaussian noise. When constructing a precessing search we must therefore ensure that any increase in the sensitivity due to an improved signal model is not outweighed by a relatively larger increase in the noise rate.

[3]If only considering the $l = 2$, $|m| = 2$ harmonics in the instantaneous orbital plane, the orbital phase can still be analytically maximized over. This is not the case if higher-order modes are included.

IV. PRECESSING WAVEFORM HARMONIC DECOMPOSITION

In this section we will review the harmonic decomposition for precessing signals, as introduced in [54], before discussing how this formulation can be useful in solving many of the issues reducing the effectiveness of current precessing searches in the next section.

The gravitational waves emitted by a binary can be decomposed into a set of spin-weighted spherical harmonics. This decomposition is given by

$$h_+ - ih_\times = \sum_{l\geq 2} \sum_{-l\leq m\leq l} h_{l,m}(t) Y^{l,m}_{(-2)}(\theta_{\mathrm{JN}}, \varphi_0), \quad (13)$$

where $Y^{l,m}_{(-2)}(\theta_{\mathrm{JN}}, \varphi_0)$ are the spin-weighted spherical harmonics with weight -2 [64], $h_{l,m}(t)$ are the harmonic components for the binary and $\theta_{\mathrm{JN}}, \varphi_0$ are angles giving the direction to the observer in a source-centered coordinate system with the its z-axis along $\boldsymbol{J}$, as described earlier in Sec. II.

To calculate the emitted gravitational waves, $h'_+(t, \theta_{\mathrm{JN}}, \varphi_0)$ and $h'_\times(t, \theta_{\mathrm{JN}}, \varphi_0)$ for a binary rotated with respect to the original by the Euler angles (α, β, γ), we must calculate the new harmonic components. We do this by performing a rotation

$$h'_{l,m}(t) = e^{im\alpha} \sum_{-l\leq m'\leq l} e^{im'\gamma} d^l_{m',m}(-\beta) h_{l,m'}(t), \quad (14)$$

where $d^l_{m,m'}$ is the Wigner d-matrix [65].

An important property of the spin-weighted spherical harmonics is that, under rotation, the modes for a particular value of l will not couple with modes for other l values [66]. If we therefore start by restricting ourselves to the dominant $l = 2$, $|m| = 2$ mode, after performing a rotation we will have a maximum of five nonzero harmonic components with $l = 2$, $-2 \leq m \leq 2$.

In the case of a precessing binary under the adiabatic limit, we can therefore calculate the harmonic components, $h^{\mathrm{P}}_{2,m}(t)$, by performing a time-dependent rotation of the harmonic components of a nonprecessing binary, $h^{\mathrm{NP}}_{2,m}(t)$ [67]. The first two Euler angles of the rotation are given by $\alpha(t)$ and $\beta(t)$ as defined in Fig. 1. The third Euler angle is $\gamma(t)$ which is defined by [54,68]

$$\dot{\gamma} = \Omega_p \cos\beta. \tag{15}$$

By combining Eqs. (13) and (14) and using the definitions of the spin-weighted spherical harmonics and Wigner d-matrix, it is shown in [54] that the observed signal for a precessing binary can then be written as

$$h(t) = \Re\left[\frac{A_0(t)e^{2i(\phi_s(t)+\alpha(t))}}{(1+b^2(t))^2}\sum_{k=0}^{4}(be^{-i\alpha(t)})^k \times (F_+\mathcal{A}^+_k - iF_\times\mathcal{A}^\times_k)\right], \tag{16}$$

where the amplitude $A_0(t)$ is proportional to $|h^{\mathrm{NP}}_{2,2}(t)|$. The phase, $\phi_s(t)$, is a combination of the nonprecessing waveform's orbital phase and γ given by

$$\phi_s(t) = \phi(t) - \gamma(t), \tag{17}$$

the parameter $b(t)$ is defined as

$$b(t) = \tan(\beta(t)/2) \tag{18}$$

and the constants $\mathcal{A}^+_k$ and $\mathcal{A}^\times_k$ are defined as

$$\begin{aligned}
\mathcal{A}^+_0 &= \mathcal{A}^+_4 = \frac{1}{r}\left(\frac{1+\cos^2\theta_{\mathrm{JN}}}{2}\right)\\
\mathcal{A}^\times_0 &= -\mathcal{A}^\times_4 = \frac{1}{r}\cos\theta_{\mathrm{JN}}\\
\mathcal{A}^+_1 &= -\mathcal{A}^+_3 = \frac{2}{r}\sin\theta_{\mathrm{JN}}\cos\theta_{\mathrm{JN}}\\
\mathcal{A}^\times_1 &= \mathcal{A}^\times_3 = \frac{2}{r}\sin\theta_{\mathrm{JN}}\\
\mathcal{A}^+_2 &= \frac{3}{r}\sin^2\theta_{\mathrm{JN}}\\
\mathcal{A}^\times_2 &= 0.
\end{aligned} \tag{19}$$

We can see from Eq. (16) that the observed signal has five harmonic components forming a power series in $b(t)e^{-i\alpha(t)}$. For each value of k the amplitude is therefore scaled by an extra factor of b and the frequency increases by the precession frequency. Each individual harmonic's amplitude evolves proportionally to the aligned-spin waveform and will therefore not show the characteristic modulation of a precessing signal. The modulation effects are then generated by the interference between the different harmonics. This can be seen in Fig. 3, which shows a precessing signal and the 5 harmonics generated for the same binary.

Following [54] we will factorize out the dependence on the initial orbital phase and precession phase, defining

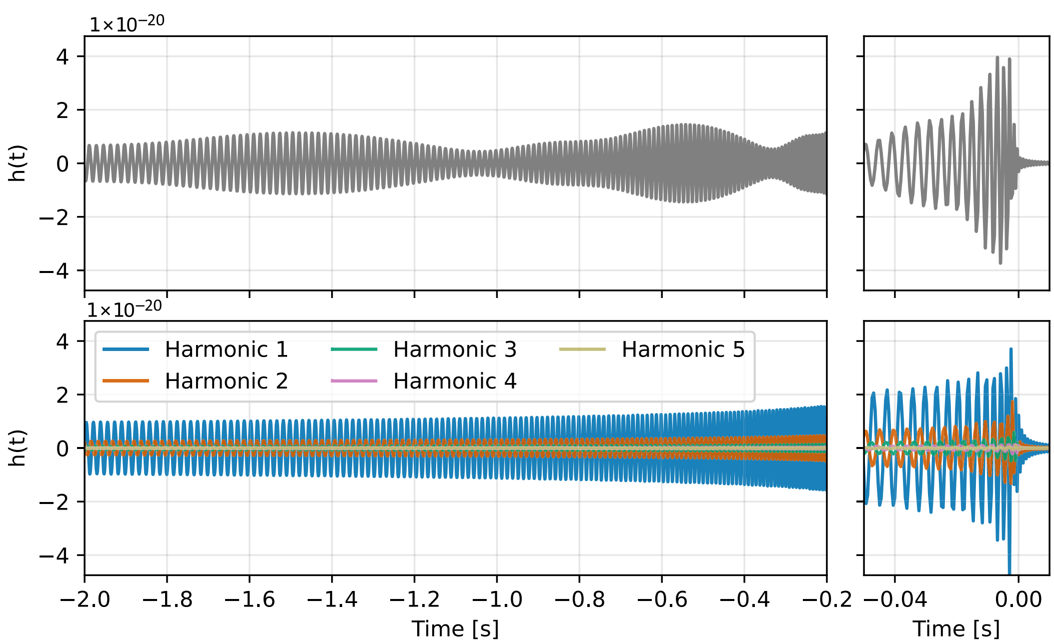

FIG. 3. An example of a precessing waveform for a binary with component masses $m_1 = 10M_\odot$, $m_2 = 1.5M_\odot$, and component spins $\boldsymbol{s}_1 = (0.5, 0.5, 0.5)$, $\boldsymbol{s}_2 = (0, 0, 0)$. The top panel shows the observed waveform for a signal with $\iota = \pi/4$ defined at 20 Hz viewed directly overhead. The bottom panel shows the 5 harmonics for this template. For both the top and bottom panels, the Left plot focuses on the inspiral, and the Right plot focuses on the merger and ringdown.

$$\hat{\phi}(t) = \phi_s(t) - \phi_0 + \alpha(t) - \alpha_0. \quad (20)$$

We can then rewrite the observed signal in the form

$$h(t) = \Re \sum_{k=0}^{4} A_k h_k(t) e^{i\phi_k}, \quad (21)$$

where

$$h_k(t) = \frac{A_0(t) b^k(t)}{(1 + b^2(t))^2} e^{i(2\hat{\phi}(t) - k(\alpha(t) - \alpha_0))}, \quad (22)$$

$$A_k = ((F_+ \mathcal{A}_k^+)^2 + (F_\times \mathcal{A}_k^\times)^2)^{1/2} \quad (23)$$

and

$$\phi_k = 2\phi_0 + (2 - k)\alpha_0 - \tan^{-1}\left(\frac{F_\times \mathcal{A}_k^\times}{F_+ \mathcal{A}_k^+}\right). \quad (24)$$

We can see that the signals' dependence on the extrinsic parameters is contained within the ten constant components of A_k and ϕ_k. Any change in the sky location or orientation will therefore simply correspond to a change in the overall amplitudes and phases of each harmonic, while the evolution of the amplitude and phase is unchanged.

Using the stationary phase approximation we can express the observed waveform in the frequency domain [54] as

$$\tilde{h}(f) = \Re \sum_{k=0}^{4} A_k \tilde{h}_k(f) e^{i\phi_k}, \quad (25)$$

where

$$\tilde{h}_k(f) = \frac{A_0(f) b^k(f)}{(1 + b^2(f))^2} e^{i(2\hat{\phi}(f) - k(\alpha(f) - \alpha_0))}. \quad (26)$$

V. APPLYING THE HARMONIC DECOMPOSITION

We now focus on formulating a procedure to search for precessing waveforms using the harmonic construction for the template filter waveforms. For a given set of intrinsic parameters (masses and spins), we can search using each harmonic individually and maximize over the ten parameters, D_k and ϕ_k, effectively maximizing over the extrinsic parameters of the binary. In order to generate the templates for each harmonic we linearly combine waveforms generated with different values of the extrinsic parameters $(\Theta, \Phi, \Psi, \theta_{\mathrm{JN}}, \alpha_0, \phi_0)$. A method for achieving this is laid out in [54].

However, using all 5 harmonics would be very similar to the method of [49] and would suffer the same issues of constraining the A_k and ϕ_k values to physically possible combinations. Nevertheless, the nature of the harmonic decomposition offers a way to solve this problem. As the harmonics form a power-series in the parameter b, if we have $b < 1$ then each subsequent harmonic in the series will be weaker than the previous. Likewise, if $b > 1$ then this will be reversed and the fifth harmonic will be the most significant. In [54] the authors show that for the majority of binaries in the sensitive range of current detectors the average value of b is below 0.4; it was shown that even for the most extreme population considered, the average value of b was 0.15 with over 90% of binaries having an average value of b below 0.3. In this case each subsequent harmonic after the first will be less significant than the previous when modeling the precessing signal. This provides a way to solve the problems faced in previous precessing searches. If precessing signals can be reliably modeled using fewer than five harmonics then we can use a smaller number of harmonics in the search. This will in turn reduce the freedom of the model to match with noise in the data. In [54], it is suggested to perform a search using only the $k = 0$, 1 harmonics. Here, we will investigate the best number of harmonics to be used in order to maximize the sensitivity of the search to precessing signals.

VI. SEARCH SETUP

In this section we will describe our implementation of the harmonic decomposition from [54] to perform a search for precessing binaries.

For the purpose of this work we will focus on the observation of neutron-star–black-holes as these systems are ones where the effects of precession are most observable [41,43,69,70]. Specifically, we will search for signals with black hole masses in the range $[5, 20] M_\odot$ and neutron star masses in the range $[1.2, 1.7] M_\odot$, with maximum spin magnitudes on each component of 0.99.

As a starting point we will use a two detector network consisting of the LIGO Hanford and LIGO Livingston detectors [1], but this method could be extended to a larger network in the future.

A. Waveform model

We use the `IMRPhenomXP` waveform model [71] to model both the filter waveforms and the simulated signals that we will add to the data. This two-spin model constructs precessing signals by taking an underlying aligned-spin waveform [72] and performing a time-dependent rotation to model the precession effects [67]. We note that the harmonic decomposition does not intrinsically rely on `IMRPhenomXP` and can be applied to other waveforms models, e.g., [73–76].

Previously, it has been shown that the four in-plane spin degrees of freedom can be mapped to a single parameter, which captures the dominant precession effects. This effective spin precession parameter, χ_p, is defined as [77],

$$\chi_p = \frac{1}{A_1 m_1^2} \max(A_1 S_{1\perp}, A_2 S_{2\perp}), \tag{27}$$

where

$$A_1 = 2 + \frac{3m_2}{2m_1}, \qquad A_2 = 2 + \frac{3m_1}{2m_2} \tag{28}$$

and $S_{1\perp}$, $S_{2\perp}$ are the in-plane component spins perpendicular to the orbital angular momentum. We note that other metrics have also been proposed [54,78–80].

Given that we use the same waveform model for both the filter waveforms and the simulated signals that we will add to the data, our results will not include the systematic effect of waveform inaccuracy [71,76]. This was chosen to demonstrate that our new method can be used to observe precessing signals, both now and in the future; as waveform models improve, the systematic errors will decrease, and hence the systematic uncertainty in current waveform models would not be a fair indicator of performance of this method in the future. Nevertheless, one should not ignore such systematics if using this method to perform an astrophysical search. The easiest way to approximately account for that would be to use different waveform models when generating simulated signals.

B. Matched filter

First let us consider a single detector and a signal with known intrinsic parameters. The log likelihood-ratio for a known signal, h in Gaussian noise, s, is given by

$$\lambda(h) = (h|s) - \frac{1}{2}(h|h), \tag{29}$$

where $(a|b)$ represents the commonly used noise-weighted inner product

$$(a|b) = 4\mathrm{Re} \int_0^{+\infty} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} df. \tag{30}$$

As shown in [15,17] for a signal with an unknown amplitude and phase we can maximize the log likelihood-ratio by using the phase-maximized matched-filter

$$\rho^2 = \frac{|(h|s)|^2}{(h|h)}. \tag{31}$$

If the 5 harmonics are independent we can matched-filter over each of them independently while maximizing over the phase (ϕ_k) and amplitude (A_k). However, the harmonics are not guaranteed to be orthogonal to one another. This will introduce covariance between the matched-filter outputs produced by each harmonic, making the maximization of the log likelihood ratio, or signal-to-noise ratio, more complicated to compute. In order to simplify this calculation we will first ensure that the harmonics are orthogonal and normalized such that

$$\mathcal{M}_{kl} = \delta_{kl} \quad \text{where } \mathcal{M}_{kl} = |\langle h_k | h_l \rangle|. \tag{32}$$

In order to diagonalize the matrix $\mathcal{M}_{kl}$, while maintaining the natural hierarchy of the harmonics we use the Gram-Schmidt process, where the first orthogonal harmonic is given by $\tilde{h}_{0\perp} = \tilde{h}_0$ and each orthogonal harmonic for $k > 0$ is given by

$$\tilde{h}_{k\perp} = \tilde{h}_k - \sum_{l=0}^{k-1} \langle h_l | h_k \rangle \tilde{h}_l \tag{33}$$

Finally, the orthonormalized harmonics are given by

$$\tilde{h}_{k\perp} = \frac{\tilde{h}_{k\perp}}{(h_{k\perp}|h_{k\perp})} \tag{34}$$

After completing this step we can maximize the log likelihood-ratio by summing the phase-maximized signal-to-noise ratios for each harmonic in quadrature

$$\rho_h^2 = \max_{A_k, \phi_k} [\rho^2] = \sum_{k=1}^{N} \rho_k^2, \tag{35}$$

where ρ_h is the total signal-to-noise ratio and N is the number of harmonics we choose to use, the choice of which we will discuss shortly. This gives a simple method to maximize the signal-to-noise ratio over the sky location and orientation, capturing the full (if $N = 5$) signal-to-noise ratio of the precessing signal. The phase-maximized signal-to-noise ratio for each harmonic, ρ_k^2, is the sum of two independent Gaussian random variables, each with a mean of 0 and variance of 1, in the absence of a signal [15]. The value of ρ_h^2 in the absence of a signal will therefore be the sum of $2N$ Gaussian random variables following a χ^2-distribution with 2N degrees of freedom.

C. Template bank generation

In order to generate our set of filter waveforms (or template bank) for this search we will use stochastic template bank generation [81,82].

In the case of a search with aligned-spin filter waveforms, stochastic template bank generation works as follows. Starting with an empty template bank, a random point is chosen from within the target parameter space and the corresponding signal, h_{prop}, is generated. For each template, h_i, within the current template bank, the match, $m(h_i, h_{\mathrm{prop}})$, is calculated and the maximum match across the template bank is defined as the fitting factor

$$\mathrm{FF}(h_{\mathrm{prop}}) \coloneqq \max_i [m(h_i, h_{\mathrm{prop}})]. \tag{36}$$

If the fitting factor is below a given threshold, usually 0.97, then the proposed template h_{prop} is not covered sufficiently by the current template bank and h_{prop} is added to the bank. This process is repeated iteratively until a set limit for the size or coverage of the template bank is reached.

In the case of aligned-spin filter waveforms the match is calculated using the phase-maximized matched-filter. In that case we do not need to maximize over the phase and time of the proposal h_{prop}, as the effect it would have on the match would be equivalent to a phase or time shift in the template h_i. The match therefore does not depend on the sky position or orientation.

In the case of a precessing signal, if we maximize over the values of A_k and ϕ_k for the bank template, h_i, the match will still depend on the values of A_k and ϕ_k of the proposal point, h_{prop}. Therefore, when proposing a new point, we choose a specific sky location and orientation, giving us specific values of A_k and ϕ_k for h_{prop}. We then calculate the match as

$$m(h_i, h_{\mathrm{prop}}) = \max_{t_i, t_{\mathrm{prop}}} \left[\sum_{k=1}^{5} \frac{|\langle \hat{h}_{k\perp,i} | h_{\mathrm{prop}} \rangle|^2}{(h_{\mathrm{prop}} | h_{\mathrm{prop}})} \right]^{1/2}, \quad (37)$$

noting that we explicitly include all five components during template bank generation. The match is then maximized across the template bank to get the fitting factor. If the fitting factor is less than the threshold, the point is then added to the bank, discarding the sky location and orientation parameters. We do note that this process may lead to a high density of templates in some regions of parameter space where for specific values of the binary orientation the waveform can vary significantly with small changes in other parameters.

We empirically evaluated the performance of the template bank as templates were added and chose to stop the stochastic generation at 358,866 templates. This was because choices we will discuss a little later in Sec. VI E regarding the number of harmonics we use for each template limit the coverage of the bank far more than the loss in coverage compared to creating a hypothetical fully-converged 5-component bank, which would by definition have a fitting factor ≥ 0.97 in all areas of parameter space. This also balances computational cost against the marginally increased sensitivity that is observed if we increase the size of the bank. We do note, however, that if someone were extending this method to use more harmonics than we allow, the fitting factor of the 5-harmonic template bank will become a limiting factor. In addition, we observed that a significant number of templates included unphysical features and did not follow the expected harmonic hierarchy. We discuss this more in Appendix B but these issues resulted in us removing 3706 templates from the template bank, resulting in a final bank of 355,160. We evaluate the performance of this template bank in detail in Sec. VI E.

D. How many components?

We now turn to the question of if we need to use all five harmonics when searching for precessing systems, or, if we can efficiently find precessing systems while using fewer harmonics, as suggested in [54]. We start this discussing by observing that the five amplitudes and phases recovered when filtering five harmonics are not independent. The five amplitudes, A_k, and five phases, ϕ_k, are dependent on a total of seven parameters (r, Θ, Φ, Ψ, ϕ_0, θ_{JN}, α_0). When considering a single detector, the three detector-response angles and the distance r enter as two independent quantities; an overall scaling factor and the ratio between the response factors F_+ and $F_\times$. This means that the ten components of A_k and ϕ_k are functions of five independent parameters and one would ideally place constraints on the allowed values of A_k and ϕ_k. Maximizing over all ten components of A_k and ϕ_k independently will allow for many unphysical combinations. This would in turn likely increase the rate of noise events, without accruing any additional signal power, reducing the sensitivity of such a search.

The five harmonics have a natural hierarchy and when $b < 1$ each harmonic after the first will be weaker by a factor of b and will, on average, contribute less signal-to-noise ratio to the observed signal. We do note though that for different sky positions and orientations the relative amplitudes of the harmonics will change as we can see in Eq. (23). In some cases, even when b is small, harmonics with larger k values will be able to contribute significantly to the signal-to-noise ratio. However, for most binary configurations this will be rare. We refer the reader to [54] for details.

To illustrate this we choose a selection of points with specific values of masses and spins. We then randomly generate a large number of systems with those mass and spins values, but isotropic sky positions and orientations. We calculate the match using a template with matching intrinsic parameters using different numbers of harmonics. In Fig. 4 we show the distribution of matches for four binaries with the following properties: $m_1 = 10M_\odot$; $m_2 = 1.5M_\odot$; component spins parallel to the orbital angular momentum, $s_{1\parallel} = 0.3$, $s_{2\parallel} = 0.3$; and perpendicular spin components of $\chi_p = 0.1$, $\chi_p = 0.3$, $\chi_p = 0.6$, $\chi_p = 0.9$ respectively. We find that when $\chi_p = 0.1$, a single harmonic is able to achieve a match of at least 0.97 for ~79% of the tested systems. When using two harmonics, this threshold is passed for 100% of the systems. In this case, using a third harmonic would not make a significant difference to the match. However, as χ_p increases, the match using the first two harmonics decreases. In the most extreme case considered, $\chi_p = 0.9$, using two harmonics only achieves a match of at least 0.97 for approximately 40% of the systems. However, using

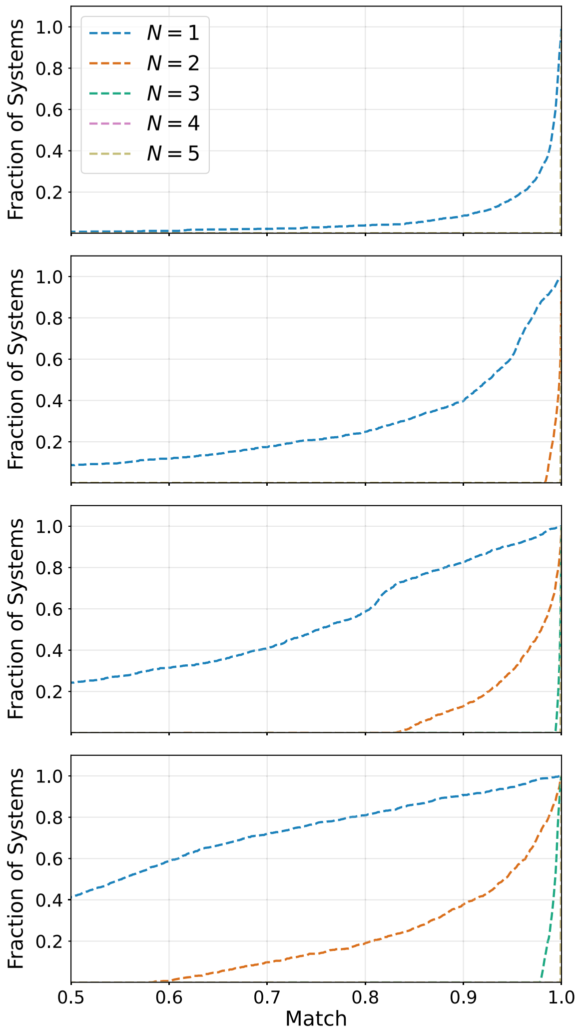

FIG. 4. Cumulative distribution of matches for four binaries with component masses $m_1 = 10M_\odot$, $m_2 = 1.5M_\odot$, component spins parallel to the orbital angular momentum $s_{1\|} = 0.3$, $s_{2\|} = 0.3$, and precessing spins (from top to bottom) of $\chi_p = 0.1$, $\chi_p = 0.3$, $\chi_p = 0.6$, $\chi_p = 0.9$ when using a different number of harmonics N. Each cumulative distribution is constructed by histogramming and summing the match based on 1 thousand randomly chosen systems. Each match is calculated by comparing a template with fixed intrinsic parameters and a randomly chosen sky position and orientation, with the waveform generated by summing the templates first N harmonics. We see that as the in-plane spin of the binary increases, a greater (lower) fraction of systems have match less (greater) than 0.97 for a fixed number of harmonics used.

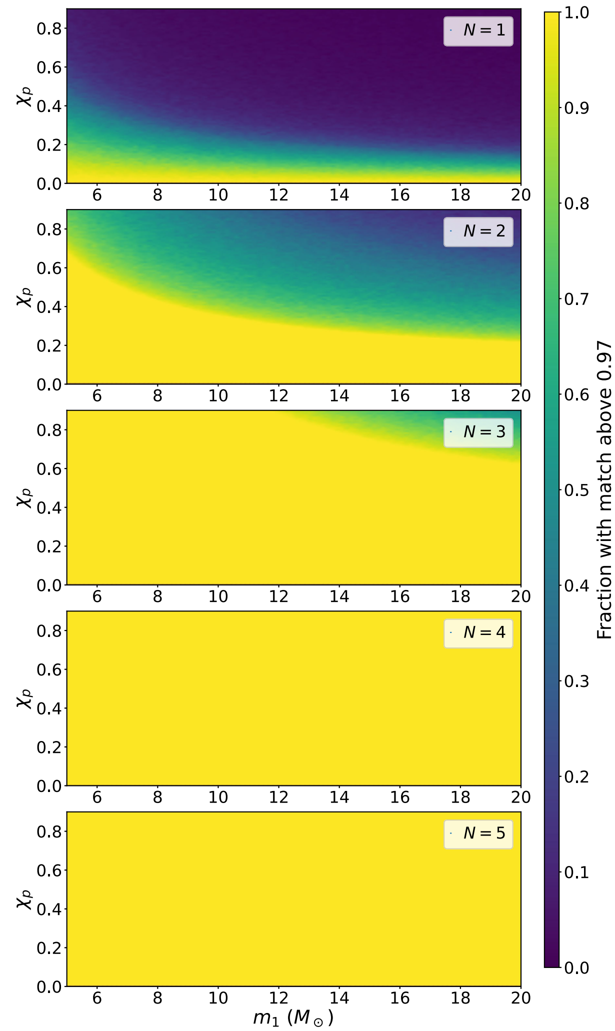

FIG. 5. The fraction of sky positions and orientations for which the match is over 0.97, using (from top to bottom) one harmonic, two harmonics, three harmonics, four harmonic and five harmonics. Matches are calculated for binaries with a fixed secondary mass, $m_2 = 1.5M_\odot$, and component spins parallel to the orbital angular momentum $s_{1\|} = 0.3$, $s_{2\|} = 0.3$, while the primary mass m_1 and in-plane spin, χ_p, are varied. Matches are calculated using templates with matching intrinsic parameters.

three harmonics results in matches larger than 0.97 for all points tested.

In Fig. 5 we consider systems where we fix $m_2 = 1.5M_\odot$ and component spins parallel to the orbital angular momentum: $s_{1\|} = 0.3$, $s_{2\|} = 0.3$. However, we allow the larger mass, $m1$ and the perpendicular spin χ_p to vary. As in Fig. 4 we randomly generate a large set of points with an isotropic distribution of sky location and orientation. We then compute the fraction of signals with matches greater than 0.97 when filtering with a varying number of harmonics. We see that there is only a small region of the parameter space—where $\chi_p \approx 0$—that a single harmonic is able to reliably recover the majority of the signal-to-noise ratio. A much larger region of the parameter space is covered by increasing to two harmonics, while using three harmonics has average matches above 0.97 for all but a small region with high mass ratios and large χ_p. Extending to four or five harmonics achieves an average match of > 0.97 for all configurations considered in this example. We release all of the code used to make these plots in [83] so that an interested reader can easily generate these figures of merit for different mass and/or spin configurations.

These results motivate a scenario where the number of harmonics N we filter with for any template is dependent on the importance of the subdominant harmonics for that template. In the search we demonstrate here, we will attempt to filter with the smallest number of harmonics for any template while still maintaining a match above 0.97 for the majority of binary configurations of that template. In this case, we will still have examples where there is unphysical freedom. Additionally, there may be some combinations of A_k and ϕ_k that are physically possible but statistically very unlikely. It would be best to impose suitable priors on A_k and ϕ_k such that we could include additional harmonics without penalizing the search. Doing this in practice is complicated, but we discuss this in more detail and present an approximate solution later in this work. However, our approximate solution is only valid for up to 3 harmonics, and therefore our method will lose some sensitivity to signals where the 4th and 5th harmonics are important.

E. Selecting the number of harmonics for each template

Now we describe our method to choose the optimal number of harmonics, N, to use for each template. We note that in practice our method is limited to no more than 3 harmonics. However, we describe here how we could select an optimal number of harmonics up to and including 5. For each template we generate a set of n samples with randomly drawn sky locations and orientations, $\mathbf{\Omega}$, and matching masses and spins. We then calculate an effective match [48]

$$m_{\mathrm{eff}}(h_i) = \left[\frac{\sum_{i=0}^{n} m(h_i, h_i(\mathbf{\Omega}))^3 \rho_{\mathrm{opt}}^3(h_i)}{\sum_{i=0}^{n} \rho_{\mathrm{opt}}^3(h_i)}\right]^{1/3}, \quad (38)$$

where

$$\rho_{\mathrm{opt}}^2(h) = \left[\sum_{k=0}^{N} |\langle \hat{h}_{k\perp} | h \rangle|^2\right]^{1/2} \quad (39)$$

The effective match is weighted by the observable volume of each signal. As a figure of merit, the effective match therefore favors the sky positions and orientations that we are most likely to observe.

We start by computing the effective match using only one harmonic in the match calculation and increase the number of harmonics until the effective match is greater than 0.97, recording the number of harmonics used as N for that template. This process is repeated for the full bank to identify the minimum number of harmonics that is required for every template.

The top row of Fig. 6 shows the number of templates that would ideally be filtered using 1, 2, 3, 4 or 5 harmonics. There is a reasonably even split between the 5 possibilities, with as many as 10% of the templates require 5 harmonics to achieve our figure of merit. As expected, the average number of harmonics selected increases as the in-plane spin and mass (and therefore mass-ratio) increase, where the effects of precession will be strongest (see, e.g., Ref. [43]).

The bottom row of Fig. 6 shows the effectiveness of our precessing template bank by computing fitting factors using a varying number of harmonics (either always using N harmonics, or limiting templates to no more than N harmonics). The simulated signals are drawn from a uniform distribution of component masses and spins within the parameter space considered. Spin, sky location and orientation angles are drawn isotropically. This is also compared to the performance with an aligned-spin template bank, which contains 118,837 templates covering aligned-spin signals with a minimum match of 0.97 within the target region. The aligned-spin template bank was generated with the IMRPhenomXAS [72] waveform model. When using all five harmonics, the precessing template bank performs very well, with the vast majority of fitting factors larger than 0.95. When restricting to two harmonics there is a noticeable tail of low fitting factors—as low as 0.7—but it is a significant improvement over the aligned-spin template bank, or only using one harmonic. Three and four harmonics respectively all provide additional significant improvement. Allowing a dynamic choice of the number of harmonics for each template results in a template bank where over 99% of the points have a fitting factor larger than 0.9. To see how the fitting factor varies across the parameter space, see Appendix C.

The exact threshold used for $m_{\mathrm{eff}}(h_i)$ when choosing N could be tuned in the future in order to make the optimal trade-off between an increase in the template banks' sensitivity to precessing signals and the increase in noise due to larger values of N. We note that the $m_{\mathrm{eff}}(h_i)$ threshold is distinct from the fitting factor of the template bank itself. Therefore, we incur a loss in sensitivity *both* due to the intrinsic fitting factor of the 5-harmonic template bank, and from the desire to use fewer harmonics for each template where possible.

We also considered choosing the value of N during the construction of the template bank so that the choice of N is taken into account when testing matches for templates with different intrinsic parameters. This could be achieved by calculating the effective match between each template and the proposal within the template bank generation loop. However, this would require some optimization to be computationally feasible and we leave this for future work. In the data release accompanying this paper we provide some additional thoughts, and a partial implementation, of how to achieve this.

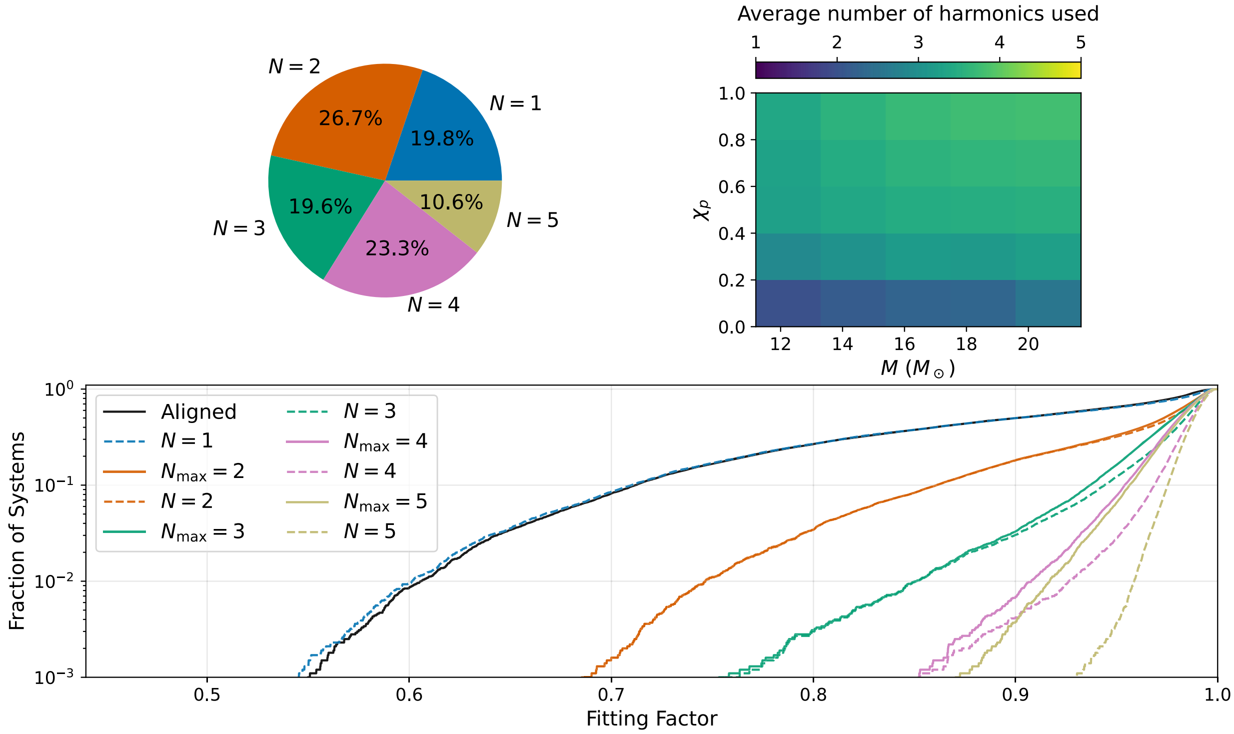

FIG. 6. Top row: the number of harmonics selected for the templates in the harmonic template bank. The left panel shows the percentage of the template bank that use one, two, three, four, or five harmonics. The right panel shows the average number of harmonics selected within a set of bins over the total mass and in-plane spin parameter. Bottom row: cumulative distribution of fitting factors when using a varying number of harmonics (either always using N harmonics, dashed lines, or allowing templates to vary the number of harmonics as required, but limiting them to use no more than N harmonics in total, solid lines) for a set of randomly chosen simulated signals within the target parameter space. The black line shows the fitting factor when using the aligned-spin template bank. Each other colour shows the fitting factor using a different number of harmonics. We see that a lower fraction of systems have fitting factor lower than a given value as the number of harmonics included in the bank increases.

VII. COINCIDENT SEARCH

We will assess the sensitivity of our method to precessing signals by applying it to a coincident modeled search. We use a stretch of ~8 days of data in the first half of the third LIGO-Virgo observing run from 14:42:36 21/05/2019 UTC to 10:38:20 29/05/2019 UTC. This data is available from GWOSC [8,84].

We search this data 3 times. First, using our aligned-spin template bank, to set a baseline. Second, using the precessing template bank, limited to a maximum of 2 harmonics. Third, using the precessing template bank, limited to a maximum of 3 harmonics. Using more than 3 harmonics would be desirable to maximize the recovered signal power, but there are two problems with this approach. First, the computational cost increases with the number of harmonics. Including both the increase in template bank size and the cost of filtering additional components, we expect the computational cost of our precessing search, compared to an aligned-spin search over the same parameter space, to increase by ~5 times, ~7 times, and ~8 times if using a maximum of 2, 3, and 5 harmonics respectively. Second, to effectively use more than 2 harmonics, we will show that one requires a ranking statistic that accounts for the expected signal and noise distribution of amplitudes and relative phases in the harmonics. This is a complex problem to solve, and we will demonstrate a novel method to include this information. However, it is only valid at present when filtering up to 3 harmonics.

In order to evaluate the sensitivity of the search, ~68, 000 simulated signals are added into the data and recovered using our search methods. The masses are drawn from a log-normal distribution over the target parameter space and the spin magnitudes are drawn uniformly with the larger body having spin magnitudes up to 0.99 and the second body up to 0.05. The spin orientations, sky location and orientation angles are drawn isotropically. The signals are then generated using `IMRPhenomXP` [71]. For all simulated signals the distance is drawn uniformly in chirp-mass weighted distance with a distance chosen uniformly in [1, 100] Mpc and then multiplied by $\mathcal{M}^{5/6}/1.2187$ (the chirp mass of a double neutron star with both components having a mass of 1.4). We consider the sensitivity to a highly precessing injection set in Appendix A.

A. Signal-consistency tests

There are many instances of non-Gaussian noise within the detector data. In order to mitigate their effects, the aligned-spin PyCBC search uses two χ^2 tests to distinguish between genuine astrophysical signals, and non-Gaussian noise artefacts [26].

First we apply the χ^2 test described in [85], which tests the distribution of power in different frequency bins. However, the power of the observed waveform as a function of frequency will change with the sky location and orientation due to the change in the relative amplitudes of the different harmonics. We will therefore reconstruct the combination of harmonics which maximized the signal-to-noise ratio for a particular trigger and use this signal to refilter the data and calculate the value of the χ^2 test. This method was previously used when developing a search for compact binary coalescences with higher harmonics in [86] and implemented in a search in [87].

In order to reconstruct the signal which maximized the signal-to-noise ratio we multiply the orthogonalized harmonics, $h_{k\perp}$, for a particular trigger by their complex signal-to-noise ratio and sum over the N harmonics

$$h = \sum_{k=0}^{N} \langle \hat{h}_{k\perp} | s \rangle \hat{h}_{k\perp}. \tag{40}$$

This signal can then be used to calculate the boundaries of the frequency bins such that the signal-to-noise ratio is evenly distributed between them and the reduced χ^2 value calculated as

$$\chi_r^2 = \frac{p}{2p-2} \sum_{i=1}^{p} \left(\frac{\rho}{p} - \rho_{bin,i} \right)^2. \tag{41}$$

Here p is the number of frequency bins to be used and $\rho_{\mathrm{bin},i}$ is the signal-to-noise ratio in bin i. The signal-to-noise ratio is then down-weighted for large values of χ_r^2 calculating a new reweighted signal-to-noise ratio

$$\tilde{\rho} = \begin{cases} \rho, & \text{if } \chi_r^2 \leq 1, \\ \rho[(1+(\chi_r^2)^3)/2]^{-\frac{1}{6}}, & \text{if } \chi_r^2 > 1. \end{cases} \tag{42}$$

If the waveform model used is not able to capture the full power of the observed signal then there will be residual power in the data due to this mismatch, increasing the value of the χ^2 test and reducing the significance of the signal. An improvement in the match from an improved template bank does not only increase the signal-to-noise ratio of the observed signal, but increases the robustness of the χ^2 test to signals in the data.

Figure 7 shows the observed signal-to-noise ratios and χ_r^2 values for a number of the simulated precessing signals in the data. We show results for the aligned-spin template bank and the harmonic template bank with a maximum of two and three harmonics. We see that the χ_r^2 value is reduced for simulated signals when using the harmonic template bank, particularly in the case of large in-plane spins. This is because the harmonic template bank provides a better match to the simulated signals in these cases. However, we do still observe increased values of χ_r^2 for large in-plane spins in the case of the harmonic template banks. This is likely due to the fact that the harmonic template banks do not have perfect coverage of the parameter space. This could be improved by using a larger template bank with better coverage of the parameter space or by using more harmonics.

We then apply the sine-Gaussian χ^2 test described in [88], which tests for excess power at frequencies above the final frequency of the search template. This step is performed using the same method as described in [88], reweighting

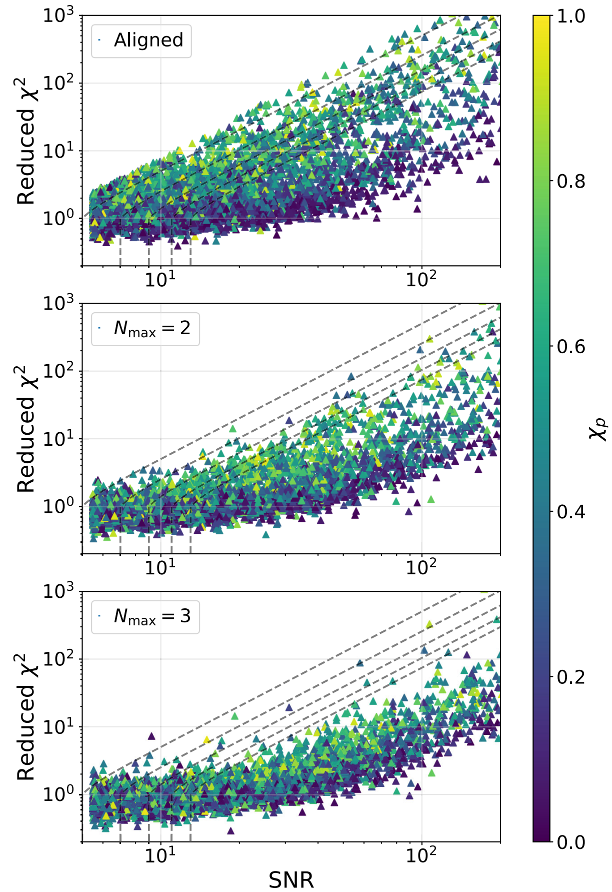

FIG. 7. Distribution of a set of simulated precessing signals in the signal-to-noise ratio-χ_r^2 plane. The dashed lines illustrate contours of constant reweighted signal-to-noise ratio. The top panel shows results for the aligned-spin template bank. The middle and bottom panels show results for the harmonic template bank with a maximum of two and three harmonics respectively.

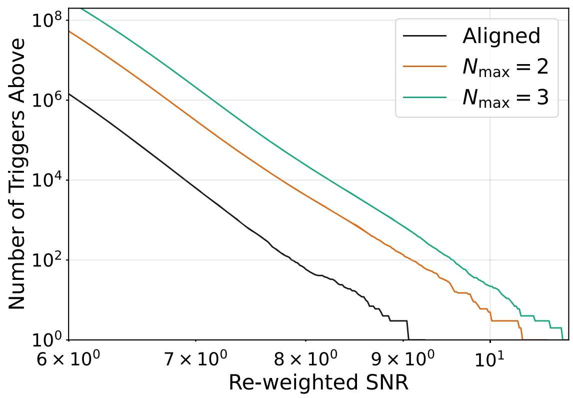

FIG. 8. Distribution of the single detector ranking statistic for the LIGO Livingstone detector, showing the number of trigger with reweighted signal-to-noise ratios above a given value. The black line shows the distribution for the aligned-spin search, while the orange and green lines show the distributions for the harmonic search using a maximum of two and three harmonics respectively.

the signal-to-noise ratio again to produce the final single-detector ranking statistic. This is given by

$$\tilde{\rho}_{sg} = \begin{cases} \tilde{\rho}, & \text{if } \chi^2_{r,sg} \leq 4, \\ \tilde{\rho}(\chi_{r,sg}/4)^{-1/2}, & \text{if } \chi^2_{r,sg} > 4, \end{cases} \tag{43}$$

where $\chi_{r,sg}$ is the output of the sine-Gaussian χ^2 test.

Figure 8 shows the distribution of reweighted signal-to-noise ratio for the three template banks. As expected, the noise rate is increased when using the harmonic template bank compare to the aligned-spin template bank. When increasing to a maximum of three harmonics we see a further increase in the noise rate. In order to achieve the same false-alarm rate with a single detector, if we were using this reweighted signal-to-noise ratio directly as the ranking statistic, we would need to observe larger signal-to-noise ratios. For example, in order to achieve the same false-alarm rate as a signal-to-noise ratio 7 trigger in the aligned-spin search, we would need to observe a signal-to-noise ratio of ∼7.9 in the case of the two harmonic template bank and a signal-to-noise ratio of ∼8.4 in the three harmonic template bank.

B. Coincident triggers

After computing the reweighted signal-to-noise ratio for each detector and identifying a set of single detector triggers, those triggers are compared across the detector network. Triggers are accepted as coincident triggers if they are generated by the same template in the bank and fall within a coincidence time window of each other. The coincidence time window is equal to the light travel time between the detectors, plus a small value to account for timing errors.

After generating a set of coincident triggers, we must evaluate the significance of each trigger. To assess the significance we perform time-slides of the data in order to generate a set of background coincidences [17]. Each coincident trigger is then assigned a false-alarm rate based on this background estimate.

We first compare the sensitivity of our three searches using the quadrature sum of the single detector trigger's signal-to-noise ratios as the ranking statistic. We apply a threshold on the false-alarm rate of 1 per 100 years to our set of injections and compute the detection efficiency using fifty distance bins. The volume contained in each distance bin is then multiplied by its detection efficiency and the results are summed across all bins to calculate the sensitive volume.

The relative sensitive volume as a function of the total mass and in-plane spin are shown in Fig. 9. When using a maximum of two harmonics, we find ∼31, 000 injections with a false-alarm rate of at least 1 per 100 years, and see a large improvement in the sensitivity of the search for signals with large in-plane spins, increasing the sensitive

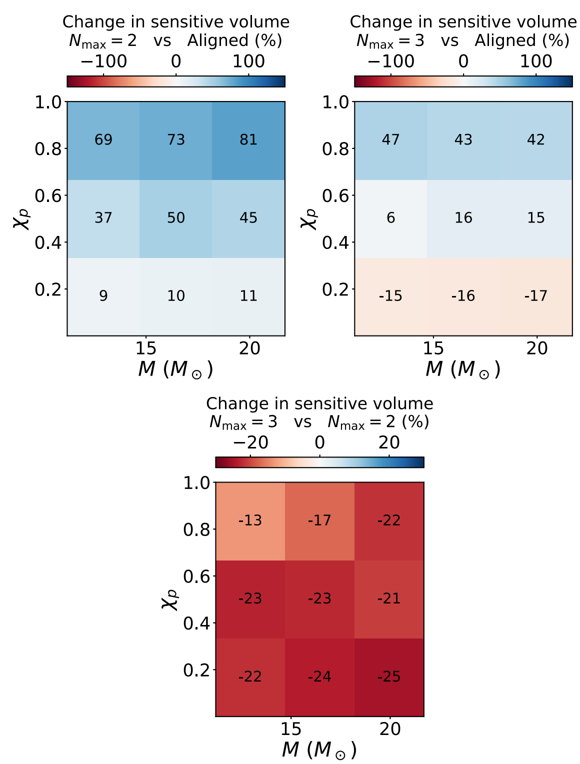

FIG. 9. The change in sensitive volume of the search due to the harmonic template bank when using the coincident signal-to-noise ratio. The top left and top right panels show the change in sensitive volume from the aligned-spin template bank to the harmonic template bank with a maximum of two and three harmonics respectively. The bottom panel shows the change in the sensitive volume between the harmonic template bank with a maximum of two and three harmonics respectively. The numbers in each square give the percentage change in sensitive volume.

volume by more than ∼50% for signals with $\chi_p > 0.67$. For small values of the in-plane spin, we see almost no change in the sensitive volume of the search with respect to the aligned-spin search. This is because the region with small in-plane spins is expected to show only weak precession effects so we do not expect to see a large improvement in the recovered signal-to-noise ratio. However, the increased rate of noise triggers when compared to the aligned-spin case will cause an increase in the false-alarm rate for a fixed signal-to-noise ratio.

When using a maximum of three harmonics, we find ∼28, 000 injections, and still see an improvement compared to the aligned-spin search for large in-plane spins. However, we now see a decrease in the sensitivity for small in-plane spins. This is due to the increase in the noise rate being more significant than the signal-to-noise ratio gained in this region of the parameter space. Across the full parameter space, the harmonic search using a maximum of three harmonics is less sensitive than the harmonic search using a maximum of two. This is again due to the increase in the rate of noise outweighing any increase in signal-to-noise ratio. In order to include more than two harmonics, we will therefore need to improve the ranking statistic being used to reduce the rate of coincident noise triggers.

C. A ranking statistic for our new precessing search

The quadrature sum of reweighted signal-to-noise ratios is not an optimal ranking statistic. It discards a lot of useful information about the coincident triggers, which could be used to differentiate between signals and noise, such as the relative signal-to-noise ratios and phases of each harmonic for each detector. For each coincident trigger we have a list of properties

$$\boldsymbol{\kappa} = ([\rho_{k,d}, \phi_{\mathrm{max},k,d}, \hat{\rho}_d, t_d], \boldsymbol{\xi}) \tag{44}$$

containing the reweighted signal-to-noise ratio and phase for each harmonic, the reweighted total signal-to-noise ratio, the time-of-arrival, and the parameters of the template. The parameters within the square brackets are recorded for each detector, d, in the set of detectors, $\{d\}$, that observed the trigger. In this work, the detector network is $\{H, L\}$, representing the LIGO Hanford and Livingston detectors respectively. The parameters with a subscript k are recorded for each harmonic used for the template that observed the trigger.

In [26], a ranking statistic is discussed based on the log-ratio of the signal event-rate density and the noise event-rate density. While this statistic is not formally an "optimal" statistic, it is more sensitive than just using the reweighted signal-to-noise ratio, and it was used to search for compact binary mergers with PyCBC in the O3 observing run [7,26,89]. We wish to develop a version of that statistic that can be applied to our precessing search. We do this using a slightly modified version of that detection statistic given by

$$R(\boldsymbol{\kappa}) = -\log A_{\{H,L\}} - \sum_{d} \log r_{n,d}(\hat{\rho}, \boldsymbol{\xi}) - \log p(\boldsymbol{\Omega}|N) + \log p(\boldsymbol{\Omega}|S), \tag{45}$$

where $A_{\{H,L\}}$ is the allowed time window for coincident triggers between the two detectors; $r_{n,d}(\hat{\rho}, \boldsymbol{\xi})$ is the expected noise-rate density for a trigger with reweighted signal-to-noise ratio $\hat{\rho}$, and template parameters $\boldsymbol{\xi}$, in detector d. Finally, $p(\boldsymbol{\Omega}|N)$ and $p(\boldsymbol{\Omega}|S)$ are the likelihoods of a trigger having a set of extrinsic parameters $\boldsymbol{\Omega}$ given that it is a noise trigger or signal, respectively. $\boldsymbol{\Omega}$ consists of the observed relative amplitudes, time delays and phase differences of the signal observed in each observatory.

For this work, we have not included the relative sensitivity for the detector network and template included in [26]. This term would need to be updated to take into account the sensitivity of the different harmonics and the relative amplitudes of each. This will be required if extending to a larger network of detectors so that the detection statistic is comparable across different combinations of detectors. However, as we are limiting ourselves to two detectors we do not consider this term.

The factor $A_{\{H,L\}}$ can be calculated in the same way as for the aligned-spin search [26], while the other terms will require modification. Let us consider first the noise-rate density for a single detector, $r_{n,d}(\hat{\rho}, \boldsymbol{\xi})$. This is estimated using a decaying exponential function. The model is parametrized by two parameters: the rate of triggers above the threshold μ_n, and the slope of the exponential α. These parameters are fit using a maximum likelihood method for each template and then averaged for templates with similar intrinsic parameters. Looking at Fig. 8, we see that the distribution of single detector triggers can still be well approximated by a decaying exponential in the case of the harmonic search. However, the rate of noise triggers will increase with the number of harmonics used. When averaging the values of μ_n and α, we must therefore ensure that we do not average over templates that use different numbers of harmonics. By doing this, the factor $r_{n,d}(\hat{\rho}, \boldsymbol{\xi})$ will naturally be larger for templates with more harmonics in order to account for the increased rate of noise triggers. When extending from two to three harmonics, we saw a decrease in the sensitivity across the full parameter space due to the increase in the noise-rate when including an additional harmonic. However, only a fraction of the templates in the bank use all three harmonics. By averaging the parameters of $r_{n,d}(\hat{\rho}, \boldsymbol{\xi})$ separately for different values of N, we should be able to increase the maximum number of harmonics in use without increasing the noise-rate for templates using fewer than the maximum number of harmonics.

The term $p(\mathbf{\Omega}|S)$ gives the likelihood of a signal having a set of extrinsic parameters $\mathbf{\Omega}$. In the case of the aligned-spin search, the extrinsic parameters include the estimated amplitude, phase and time-of-arrival for each detector in the network. The likelihood is estimated using a histogram of the expected amplitude ratios, phase differences and time differences with respect to a chosen reference detector. The histogram is generated by drawing sky positions and orientations isotropically, and then calculating the expected amplitudes ratios, phase differences and time differences across the network. These values are binned, and the histogram is then smoothed using a Gaussian kernel in order to account for the uncertainty in the measured parameters.

When using the harmonic template bank, the number of extrinsic parameters is increased as we now have an estimated amplitude and phase for each harmonic within each detector. The ideal solution would be to calculate $p(\mathbf{\Omega}|S)$ using the full set of these parameters. This would also have the effect of removing or down-weighting much of the unphysical parameter space introduced by maximizing independently over A_k and ϕ_k. However, this introduces a number of technical difficulties.

First, for a template with N harmonics in D detectors, a histogram including the full set of extrinsic parameters would include $(2D+1)(N-1)$ dimensions. For three harmonics in two detectors, this already gives a seven dimensional parameter space. Increasing to a network of three detectors would produce a fourteen dimensional parameter space, making the use of these histograms computationally impractical for use within a search.

Second, in the case of aligned-spin templates, the amplitude ratios and phase difference between detectors are only dependent on the extrinsic parameters of the binary, ignoring differences in power spectral density between detectors. This means that we can generate a single histogram that can be used for all templates within the bank. In the case of the harmonic template bank, this is still the case when comparing a single harmonic across multiple detectors. However, the relative amplitudes between different harmonics will depend on the parameter b. This means that we would therefore require a set of histograms covering different values of b in order to properly account for the uncertainty due to noise. This is further complicated by the fact that b will change as the binary evolves and the amplitude ratio will therefore depend on the value of b averaged over the sensitive window of the detector. As the frequency of each harmonic is higher than the last, the average factor of b measured between each subsequent harmonic may be different. In order to compare amplitudes across different harmonics, it will be important to study the size of this effect.

As a first step, we will compare the amplitudes and phases for a single harmonic across the two detectors in our network. In this case, the harmonics in the two detectors will have the same value of b^k, and the amplitude ratio will not depend on the value of b. This removes the dependence on the intrinsic parameters of the template. We choose the harmonic with the highest signal-to-noise ratio to be used as the reference, and compare the amplitude ratio, phase difference, and time difference with respect to the same harmonic in the second detector. This will require three histograms to be generated when using a maximum of three harmonics. The required histograms can be generated using the same method as the aligned-spin case, using Eqs. (23) and (24) to calculate the expected amplitude ratios and phase differences. In this case, the likelihood $p(\mathbf{\Omega}|N)$ for noise triggers will be assumed to be uniform, as in the aligned-spin search.

We can then calculate the coincident ranking statistic using Eq. (45) and produce updated false-alarm rates. We apply this coincident ranking statistic using the same template bank and simulated signals as the previous sections. For comparison, we use the full ranking statistic described in [26] when performing the search using the aligned-spin template bank.

D. Results with improved statistic

Figure 10 shows the change in sensitivity when applying this new ranking statistic. In the case of the two harmonic template bank, we find ∼31, 000 injections with a false-alarm rate of at least 1 per 100 years, and see a similar improvement over the aligned-spin template bank as we did before introducing the new ranking statistic.

However, when using the harmonic template bank with a maximum of three harmonics, we now find ∼31, 000 injections, see an improved sensitivity over using a maximum of two harmonics, and do not see a sensitivity loss for signals with small values of χ_p. For systems with $\chi_p > 0.67$ and total mass larger than $17.5 M_\odot$, we see an increase in sensitivity of ∼100%. This demonstrates that with the appropriate choice of ranking statistic, we can account for the increased noise rate caused by introducing additional harmonics while gaining sensitivity through the improved match with signals in the data.

It will certainly be possible to improve upon this ranking statistic by using a more complete treatment of $p(\mathbf{\Omega}|S)$ in the future. We have restricted ourselves to comparing a single harmonic across each detector. It would be beneficial to extend this to a larger set of harmonics in order to remove some of the unphysical freedom introduced by the maximization of A_k and ϕ_k. One possibility is to first compute $p(\mathbf{\Omega}|S)$ in each detector separately and include it as a signal-consistency test in order to reduce the rate of single detector triggers, before subsequently testing a reduced set of harmonics across the detector network. This would require a maximum of four phase differences and four amplitude ratios for the single detector case if using all five harmonics, making the dimensionality more manageable. However, this will still require careful treatment of the averaged value of b between different harmonics.

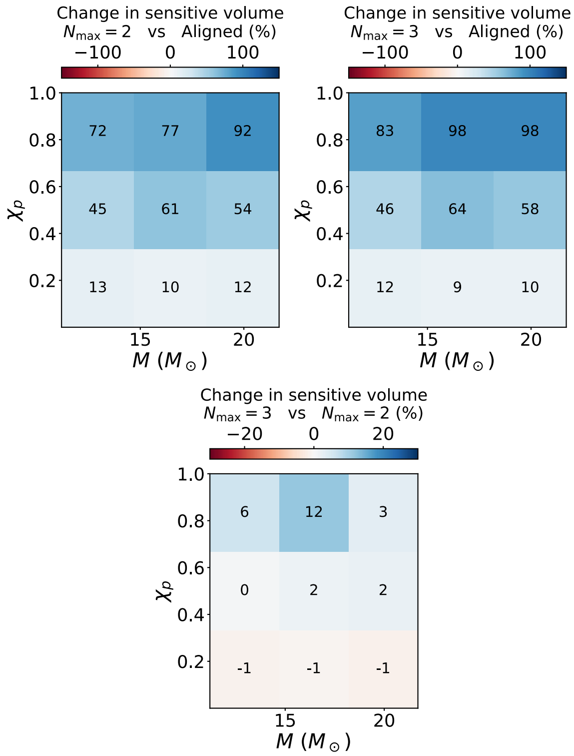

FIG. 10. The change in sensitive volume of the search due to the harmonic template bank when using the log signal-rate noise-rate ratio statistic. The top left and top right panels show the change in sensitive volume from the aligned-spin template bank to the harmonic template bank with a maximum of two and three harmonics, respectively. The bottom panel shows the change in the sensitive volume between the harmonic template bank with a maximum of two and three harmonics, respectively. The numbers in each square give the percentage change in sensitive volume.

No significant signals were observed in the stretch of data analysed during this test, with the most significant foreground trigger producing a false-alarm rate of $\sim 200\ \mathrm{yr}^{-1}$, which is consistent with the values we expect from noise triggers for this amount of data.

VIII. CONCLUSION AND OUTLOOK

In this work, we have demonstrated a method for searching for precessing binaries using a template bank utilizing a harmonic decomposition of precessing signals. We have shown that, given the natural hierarchy of the harmonics, we can often use fewer than the full set of five harmonics, and we have demonstrated a method for selecting the appropriate number of harmonics to be used for each template within a bank.

By introducing extra parameters to our template model, we not only increase the observed signal-to-noise ratio for signals, but we also increase the rate of noise triggers. We have shown that, by using an appropriate ranking statistic, this can be mitigated, and an effective search can be run using the first three harmonics. We have also shown that this results in a $\sim 100\%$ increase in the sensitive volume compared to the aligned-spin search when considering binaries with $\chi_p > 0.67$ and total mass larger than $17.5 M_\odot$.

There are two main improvements can be made to the current methods presented in this paper. First, we could construct a larger template bank in order to achieve a more complete coverage of the targeted parameter space, improving the signal-to-noise ratio of the observed signals. Second, and perhaps more importantly, we could extend the estimation of $p(\boldsymbol{\Omega}|S)$ to use more than one harmonic, allowing us to reduce the unphysical freedom caused by the maximization of A_k and ϕ_k and allowing us to employ more than 3 harmonics for templates that need it.

However, the results we present here demonstrate for the first time how a search on Advanced LIGO, Virgo, and KAGRA data can be performed using precessing compact binary mergers as waveform filter templates. This method will achieve a significant sensitivity improvement for such signals, ans with scope for further improvement, we believe that this method can be the key to uncovering a potential population of binaries with strong precessional effects.

ACKNOWLEDGMENTS

We thank Patricia Schmidt for insightful and helpful comments on this work while reviewing an earlier draft of the text, which appeared in C. M.'s thesis. We thank Stephen Fairhurst for motivating why the "two harmonic" approximation would be the right thing to use for precessing searches, even if it turns out that two harmonics may not be enough. We thank Mark Hannam for discussions concerning the inconsistencies with the `IMRPhenomXP` waveform model. C. M. acknowledges the DISCNet Center for Doctoral Training for support. C. H. thanks the UKRI Future Leaders Fellowship for support through the Grant No. MR/T01881X/1. I. H. thanks the STFC for support through the Grants No. ST/T000333/1 and No. ST/V005715/1. The authors are also grateful for computational resources provided by LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This research made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center ([90]), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

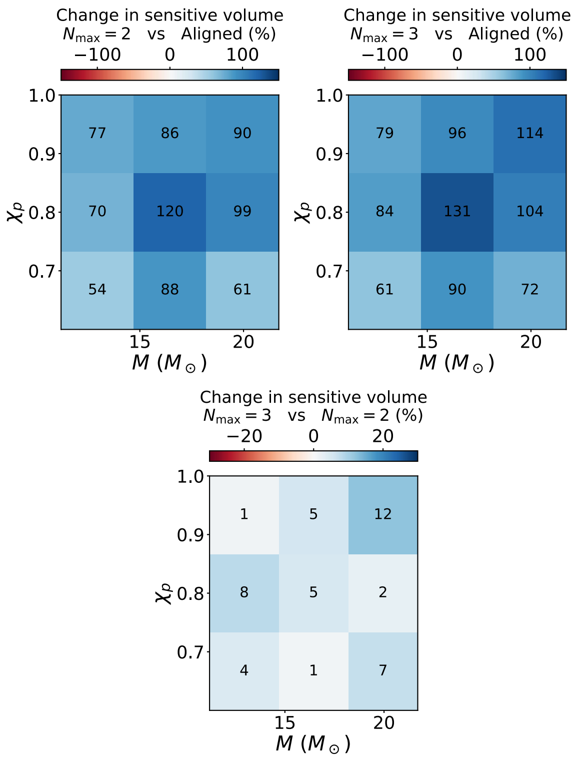

FIG. 11. Same as Fig. 10 but using a set of highly precessing simulated signals to estimate the sensitive volume.

APPENDIX A: HIGHLY PRECESSING INJECTION SET

In Sec. VII, we added a set of simulated signals into real gravitational-wave strain data to evaluate the sensitivity of our search. It was demonstrated that the sensitive volume improved by ~50%, and ~60% on average when using a precessing template bank limited to 2 and 3 harmonics respectively.

In Fig. 11, we show the improvement in sensitive volume for a set of simulated signals that are highly precessing: spins almost entirely within the plane of the binary. In comparison to the previous set, the new set of simulated signals modified only the primary spin distribution: the spin magnitudes were drawn uniformly between 0.7 and 0.99 and spin orientations drawn isotropically between $|\hat{\boldsymbol{S}}_1 \cdot \hat{\boldsymbol{L}}| < 0.3$.

We see that, when using our new ranking statistic introduced in Sec. VII C, the sensitive volume improves by ~80%, and ~90% on average when using a precessing template bank limited to 2 and 3 harmonics respectively, compared to an aligned-spin search. We also see that the 3 harmonic search out performs the 2 harmonic search for all χ_p and M.

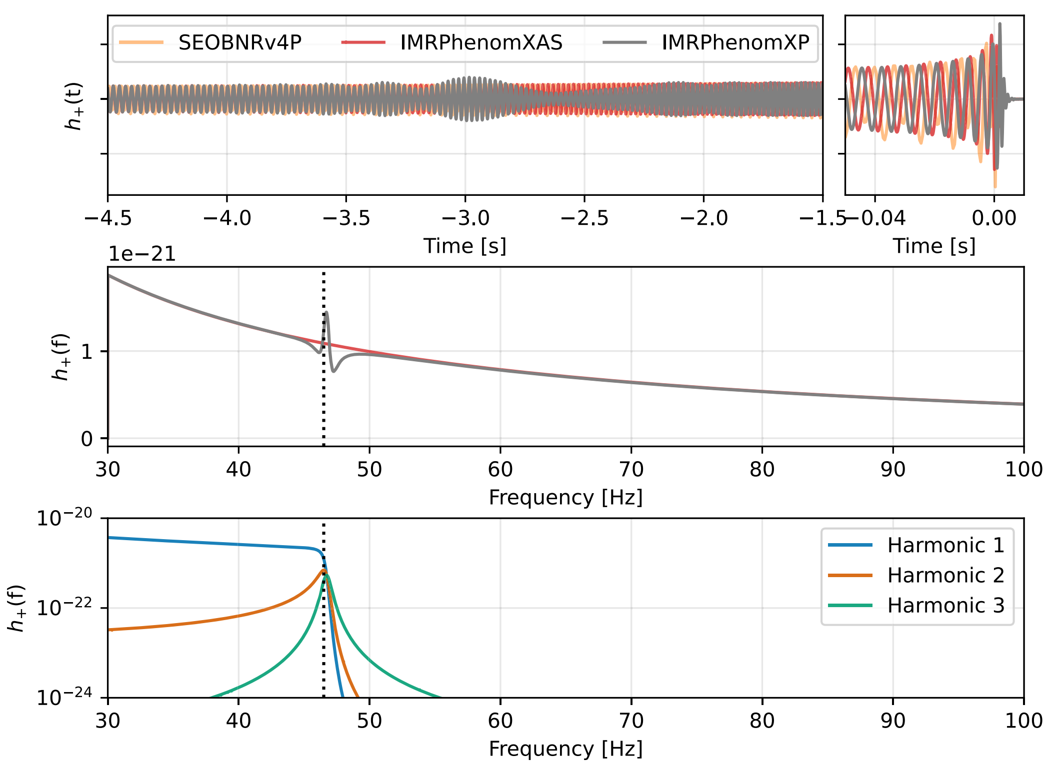

FIG. 12. Comparison of the gravitational-waves produced using different waveform models for a binary with masses $13.4M_\odot$ and $1.3M_\odot$, and component spins $s_1 = (0.0003, -0.0002, -0.5)$, $s_2 = (-0.02, -0.1, -0.2)$, viewed at an inclination angle of 1.5 radians. The top two panels show the gravitational-wave in the time-domain (Left: focusing on inspiral, Right: focusing on the merger and ringdown), the middle panel shows the gravitational-wave in the frequency-domain, and the bottom panel shows the harmonic decomposition of the gravitational-wave. The vertical black dotted line in the middle and bottom panels corresponds to 47 Hz.

APPENDIX B: WAVEFORM INCONSISTENCIES REDUCING THE BANK SIZE

After constructing the template bank through stochastic bank generation techniques (see Sec. VI C for details), we removed 1% of filter waveforms due to inconsistencies with the `IMRPhenomXP` waveform model. In our testing, we found that a several filter waveforms had nonphysical artifacts, particularly at low frequencies. This caused issues when deconstructing the waveform into the harmonic decomposition (see Sec. IV for details), and consequently, caused a background of large SNR events which biased our results.

We found that most of the inconsistent waveforms had low in-plane spins. In Fig. 12 we specifically show one of filter waveforms that caused issues; this waveform was generated for a binary with masses $13.4 M_{\odot}$ and $1.3 M_{\odot}$, with component spins $\boldsymbol{s}_1 = (0.0003, -0.0002, -0.5)$, $\boldsymbol{s}_2 = (-0.02, -0.1, -0.2)$, viewed at an inclination angle of 4.8 radians. For comparison we also show the gravitational-waves produced for the same binary configuration but with alternative models: `SEOBNRv4P` [74], which constructs precessing signals through the effective one-body approach [91], and the `IMRPhenomXAS` model, which constructs aligned-spin signals through the phenomenological approach. When generating the gravitational-wave with the `IMRPhenomXAS` model, we consider the aligned-spin projection of the binary (all in-plane spins reduced to exactly 0). From the frequency-domain, we can clearly see that the gravitational wave produced from the `IMRPhenomXP` waveform model has a nonphysical artifact at 47 Hz. Since this specific binary has $\chi_p = 0.008 \ll 1$, it should closely resemble an aligned-spin signal. As expected, wee see excellent agreement between the `IMRPhenomXAS` and `SEOBNRv4P` waveform models, suggesting that the issue is specific to the `IMRPhenomXP` model.

When deconstructing this specific binary into the harmonic decomposition, the nonphysical artifact at 47 Hz caused the leading harmonic to drop in amplitude from $O(10^{-21})$ to $O(10^{-29})$, and the 5th harmonic to increase in amplitude from (10^{-27}) to $O(10^{-21})$ to compensate. This means that when using only the leading two, or leading three harmonics, in the search analysis presented in Sec. VII, the reconstructed precessing waveform terminated at 47 Hz.

In order to construct a bank that neglected these inconsistent filter waveforms, we removed the waveforms that had either: (a) $\chi_p < 0.05$ (meaning that they closely resemble aligned-spin binaries) and a match between the `IMRPhenomXP` and `IMRPhenomXAS` < 0.85, or (b) a leading harmonic that was initially larger than the 5th harmonic, and then switching to having the 5th larger than the leading harmonic. These constraints removed ∼3700 filter waveforms from a total of 360, 000.

APPENDIX C: FITTING FACTOR ACROSS THE PARAMETER SPACE

After constructing our precessing template bank, we showed its effectiveness by computing fitting factors using a varying number of harmonics for a set of randomly chosen simulated signals, see Sec. VI E and Fig. 6 for details. We demonstrated that when using a dynamic choice for the number of harmonics, each template results in a template bank where over 99% of the points have fitting factor larger than 0.9.

In Fig. 13 we show how the fitting factor varies across the larger mass, m_1, and in-plane spin parameter space. We see that when using only one harmonic, the fitting factors are low for most of the parameter space. In fact, we find that ∼27% of points have fitting factors > 0.97. As the maximum number of harmonics increases, we see that a better overall performance with ∼54%, ∼64%, ∼74%, and ∼76% of points having fitting factors > 0.97 when using a

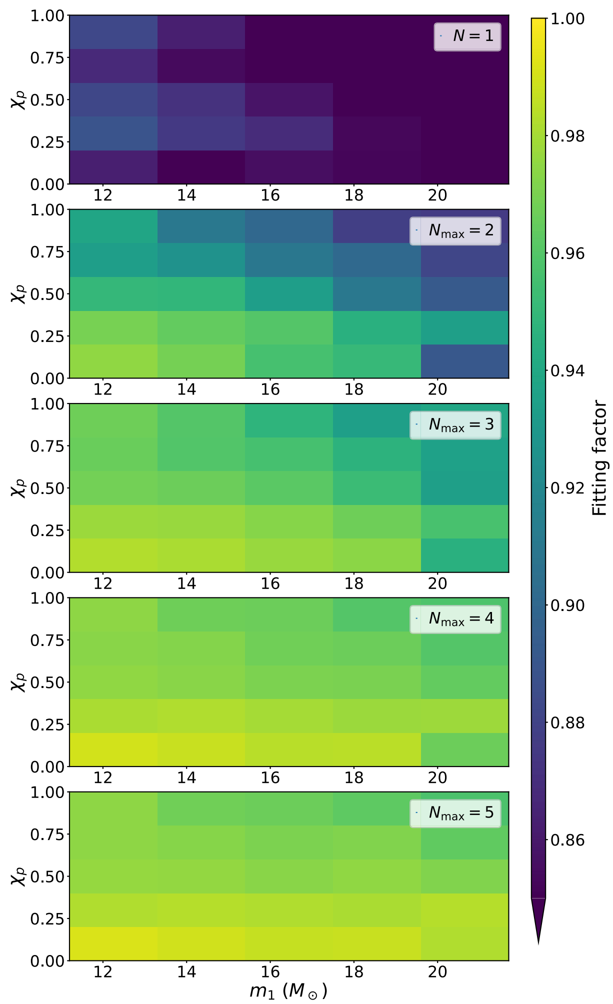

FIG. 13. The variation of fitting factor across the parameter space for the harmonic template bank with (top to bottom) one harmonic, a maximum of two, three and four harmonics respectively.

maximum of two, three, four, and five harmonics respectively. We see that in general, points with larger in-plane spins and larger primary mass have a lower fitting factor. Given that the precession modulations are more significant for high mass ratio and high in-plane spin configurations [43,54], this is understandable as more harmonics will be needed to fully describe the signal. Since we fail to obtain fitting factors > 0.97 in this extreme region of the parameter space when using a maximum of 5 harmonics, this shows that our bank is not fully converged. This arises from us stopping the template generation early at 358,866 templates. We note that for high masses, the lowest in-plane spin bin has a lower fitting factor than the adjacent bin for all searches with a maximum number of harmonics ≥ 2. This is unlike the rest of the parameter space, where the fitting factor decreases as the in-plane spin increases. We suspect this is because a lot of templates were removed in this region of parameter space due to the issues described in Appendix B, decreasing the overall sensitivity.

[1] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, Classical Quantum Gravity **32**, 074001 (2015).

[2] F. Acernese *et al.* (VIRGO Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical Quantum Gravity **32**, 024001 (2015).

[3] T. Akutsu *et al.* (KAGRA Collaboration), Overview of KAGRA: Calibration, detector characterization, physical environmental monitors, and the geophysics interferometer, Prog. Theor. Exp. Phys. **2021**, 05A102 (2021).

[4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), 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] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, Phys. Rev. X **11**, 021053 (2021).

[6] R. Abbott *et al.* (LIGO Scientific and VIRGO Collaborations), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, arXiv:2108.01045.

[7] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, arXiv:2111.03606.

[8] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, SoftwareX **13**, 100658 (2021).

[9] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, Astrophys. J. **872**, 195 (2019).

[10] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New search pipeline for compact binary mergers: Results for binary black holes in the first observing run of Advanced LIGO, Phys. Rev. D **100**, 023011 (2019).

[11] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo, Phys. Rev. D **101**, 083030 (2020).

[12] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, Astrophys. J. **891**, 123 (2020).

[13] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, Astrophys. J. **922**, 76 (2021).

[14] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3a data, Phys. Rev. D **106**, 043009 (2022).

[15] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries, Phys. Rev. D **85**, 122006 (2012).

[16] K. Cannon *et al.*, Toward early-warning detection of gravitational waves from compact binary coalescence, Astrophys. J. **748**, 136 (2012).

[17] S. Babak *et al.*, Searching for gravitational waves from binary coalescence, Phys. Rev. D **87**, 024033 (2013).

[18] K. Cannon, C. Hanna, and D. Keppel, Method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence, Phys. Rev. D **88**, 024025 (2013).

[19] T. Dal Canton *et al.*, Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors, Phys. Rev. D **90**, 082004 (2014).

[20] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era, Classical Quantum Gravity **33**, 175012 (2016).

[21] S. A. Usman *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, Classical Quantum Gravity **33**, 215004 (2016).

[22] K. Cannon, C. Hanna, and J. Peoples, Likelihood-ratio ranking statistic for compact binary coalescence candidates with rate estimation, arXiv:1504.04632.

[23] C. Messick *et al.*, Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data, Phys. Rev. D **95**, 042001 (2017).

[24] S. Sachdev *et al.*, The GstLAL search analysis methods for compact binary mergers in Advanced LIGO's second and Advanced Virgo's first observing runs, arXiv:1901.08580.

[25] Q. Chu *et al.*, SPIIR online coherent pipeline to search for gravitational waves from compact binary coalescences, Phys. Rev. D **105**, 024023 (2022).

[26] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, Extending the PyCBC search for gravitational waves from compact binary mergers to a global network, Phys. Rev. D **102**, 022004 (2020).

[27] F. Aubin *et al.*, The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run, Classical Quantum Gravity **38**, 095004 (2021).

[28] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Coherent method for detection of gravitational wave bursts, Classical Quantum Gravity **25**, 114029 (2008).

[29] R. Lynch, S. Vitale, R. Essick, E. Katsavounidis, and F. Robinet, Information-theoretic approach to the gravitational-wave burst detection problem, Phys. Rev. D **95**, 104046 (2017).

[30] V. Skliris, M. R. K. Norman, and P. J. Sutton, Real-time detection of unmodelled gravitational-wave transients using convolutional neural networks, arXiv:2009.14611.

[31] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries, Phys. Rev. D **49**, 6274 (1994).

[32] I. Mandel and A. Farmer, Merging stellar-mass binary black holes, Phys. Rep. **955**, 1 (2022).

[33] M. Mapelli, Astrophysics of stellar black holes, Proc. Int. Sch. Phys. "Enrico Fermi" **200**, 87 (2020).

[34] M. Mapelli, Formation channels of single and binary stellar-mass black holes, arXiv:2106.00699.

[35] D. Gerosa and M. Fishbach, Hierarchical mergers of stellar-mass black holes and their gravitational-wave signatures, Nat. Astron. **5**, 749 (2021).

[36] V. Kalogera, Spin orbit misalignment in close binaries with two compact objects, Astrophys. J. **541**, 319 (2000).

[37] D. Gerosa, E. Berti, R. O'Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, Spin orientations of merging black holes formed from the evolution of stellar binaries, Phys. Rev. D **98**, 084036 (2018).

[38] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, Illuminating black hole binary formation channels with spins in Advanced LIGO, Astrophys. J. Lett. **832**, L2 (2016).

[39] B. P. Gompertz, M. Nicholl, P. Schmidt, G. Pratten, and A. Vecchio, Constraints on compact binary merger evolution from spin-orbit misalignment in gravitational-wave observations, Mon. Not. R. Astron. Soc. **511**, 1454 (2022).

[40] S. Stevenson, C. P. L. Berry, and I. Mandel, Hierarchical analysis of gravitational-wave measurements of binary black hole spin–orbit misalignments, Mon. Not. R. Astron. Soc. **471**, 2801 (2017).

[41] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO, Phys. Rev. D **89**, 024010 (2014).

[42] J. Calderón Bustillo, P. Laguna, and D. Shoemaker, Detectability of gravitational waves from binary black holes: Impact of precession and higher modes, Phys. Rev. D **95**, 104038 (2017).

[43] R. Green, C. Hoy, S. Fairhurst, M. Hannam, F. Pannarale, and C. Thomas, Identifying when precession can be measured in gravitational waveforms, Phys. Rev. D **103**, 124023 (2021).

[44] M. Hannam *et al.*, General-relativistic precession in a black-hole binary, Nature (London) **610**, 652 (2022).

[45] E. Payne, S. Hourihane, J. Golomb, R. Udall, D. Davis, and K. Chatziioannou, Curious case of GW200129: Interplay between spin-precession inference and data-quality issues, Phys. Rev. D **106**, 104017 (2022).

[46] T. A. Apostolatos, Search templates for gravitational waves from precessing, inspiraling binaries, Phys. Rev. D **52**, 605 (1995).

[47] T. A. Apostolatos, Construction of a template family for the detection of gravitational waves from coalescing binaries, Phys. Rev. D **54**, 2421 (1996).

[48] A. Buonanno, Y.-b. Chen, and M. Vallisneri, Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit, Phys. Rev. D **67**, 104025 (2003).

[49] Y. Pan, A. Buonanno, Y.-b. Chen, and M. Vallisneri, A physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single spin binaries, Phys. Rev. D **69**, 104017 (2004).

[50] A. Buonanno, Y.-b. Chen, Y. Pan, and M. Vallisneri, A quasi-physical family of gravity-wave templates for precessing binaries of spinning compact objects. 2. Application to double-spin precessing binaries, Phys. Rev. D **70**, 104003 (2004); **74**, 029902(E) (2006).

[51] I. W. Harry and S. Fairhurst, A coherent triggered search for single spin compact binary coalescences in gravitational wave data, Classical Quantum Gravity **28**, 134008 (2011).

[52] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, Searching for gravitational waves from compact binaries with precessing spins, Phys. Rev. D **94**, 024012 (2016).

[53] D. Fazi, Development of a physical-template search for gravitational waves from spinning compact-object binaries with LIGO, Ph.D. thesis, Bologna University, 2009.

[54] S. Fairhurst, R. Green, C. Hoy, M. Hannam, and A. Muir, Two-harmonic approximation for gravitational waveforms from precessing binaries, Phys. Rev. D **102**, 024055 (2020).

[55] L. E. Kidder, Coalescing binary systems of compact objects to post-Newtonian 5/2 order. 5. Spin effects, Phys. Rev. D **52**, 821 (1995).

[56] M. Campanelli, C. O. Lousto, and Y. Zlochower, Spinning-black-hole binaries: The orbital hang up, Phys. Rev. D **74**, 041501 (2006).

[57] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation, Phys. Rev. D **86**, 064020 (2012).
[58] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, Living Rev. Relativity **9**, 4 (2006).
[59] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms, Phys. Rev. D **79**, 104023 (2009); **84**, 049901(E) (2011).
[60] B. S. Sathyaprakash and S. V. Dhurandhar, Choice of filters for the detection of gravitational waves from coalescing binaries, Phys. Rev. D **44**, 3819 (1991).
[61] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, Phys. Rev. D **53**, 6749 (1996).
[62] P. Grandclement, V. Kalogera, and A. Vecchio, Searching for gravitational waves from the inspiral of precessing binary systems: Problems with current waveforms, Phys. Rev. D **67**, 042003 (2003).
[63] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, B. S. Sathyaprakash, H. Tagoshi, and H. Takahashi, Template banks to search for compact binaries with spinning components in gravitational wave data, Phys. Rev. D **80**, 024009 (2009).
[64] L. E. Kidder, Using full information when computing modes of post-Newtonian waveforms from inspiralling compact binaries in circular orbit, Phys. Rev. D **77**, 044016 (2008).
[65] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Simple model of complete precessing black-hole-binary gravitational waveforms, Phys. Rev. Lett. **113**, 151101 (2014).
[66] L. Gualtieri, E. Berti, V. Cardoso, and U. Sperhake, Transformation of the multipolar components of gravitational radiation under rotations and boosts, Phys. Rev. D **78**, 044024 (2008).
[67] P. Schmidt, M. Hannam, and S. Husa, Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and nonprecessing inspiral signals, Phys. Rev. D **86**, 104063 (2012).
[68] M. Boyle, R. Owen, and H. P. Pfeiffer, A geometric approach to the precession of compact binaries, Phys. Rev. D **84**, 124011 (2011).
[69] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Detecting binary neutron star systems with spin in advanced gravitational-wave detectors, Phys. Rev. D **86**, 084017 (2012).
[70] C. Capano, I. Harry, S. Privitera, and A. Buonanno, Implementing a search for gravitational waves from binary black holes with nonprecessing spin, Phys. Rev. D **93**, 124007 (2016).
[71] G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, Phys. Rev. D **103**, 104056 (2021).
[72] G. Pratten, S. Husa, C. Garcia-Quiros, M. Colleoni, A. Ramos-Buades, H. Estelles, and R. Jaume, Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes, Phys. Rev. D **102**, 064001 (2020).
[73] S. Khan, K. Chatziioannou, M. Hannam, and F. Ohme, Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects, Phys. Rev. D **100**, 024059 (2019).
[74] S. Ossokine *et al.*, Multipolar effective-one-body waveforms for precessing binary black holes: Construction and validation, Phys. Rev. D **102**, 044055 (2020).
[75] E. Hamilton, L. London, J. E. Thompson, E. Fauchon-Jones, M. Hannam, C. Kalaghatgi, S. Khan, F. Pannarale, and A. Vano-Vinuales, Model of gravitational waves from precessing black-hole binaries through merger and ringdown, Phys. Rev. D **104**, 124027 (2021).
[76] A. Ramos-Buades, A. Buonanno, H. Estellés, M. Khalil, D. P. Mihaylov, S. Ossokine, L. Pompili, and M. Shiferaw, SEOBNRv5PHM: Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes, arXiv:2303.18046.
[77] P. Schmidt, F. Ohme, and M. Hannam, Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter, Phys. Rev. D **91**, 024043 (2015).
[78] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, When will we observe binary black holes precessing?, Phys. Rev. D **102**, 041302 (2020).
[79] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, A generalized precession parameter χ_{p} to interpret gravitational-wave data, Phys. Rev. D **103**, 064067 (2021).
[80] L. M. Thomas, P. Schmidt, and G. Pratten, New effective precession spin for modeling multimodal gravitational waveforms in the strong-field regime, Phys. Rev. D **103**, 083022 (2021).
[81] I. W. Harry, B. Allen, and B. S. Sathyaprakash, A stochastic template placement algorithm for gravitational wave data analysis, Phys. Rev. D **80**, 104014 (2009).
[82] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein, Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins, Phys. Rev. D **89**, 084041 (2014).
[83] Data release accompanying this paper, https://icg-gravwaves.github.io/precessing_search_paper. Please see here for more.
[84] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, The LIGO open science center, J. Phys. Conf. Ser. **610**, 012021 (2015).
[85] B. Allen, χ^2 time-frequency discriminator for gravitational wave detection, Phys. Rev. D **71**, 062001 (2005).
[86] I. Harry, J. Calderón Bustillo, and A. Nitz, Searching for the full symphony of black hole binary mergers, Phys. Rev. D **97**, 023004 (2018).
[87] K. Chandra, J. Calderón Bustillo, A. Pai, and I. Harry, First gravitational-wave search for intermediate-mass black hole mergers with higher order harmonics, Phys. Rev. D **106**, 123003 (2022).

[88] A. H. Nitz, Distinguishing short duration noise transients in LIGO data to improve the PyCBC search for gravitational waves from high mass binary black hole mergers, Classical Quantum Gravity **35**, 035016 (2018).
[89] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact binary mergers, Astrophys. J. **946**, 59 (2023).
[90] https://www.gw-openscience.org.
[91] A. Buonanno and T. Damour, Effective one-body approach to general relativistic two-body dynamics, Phys. Rev. D **59**, 084006 (1999).