

Searching for gravitational waves from compact binaries with precessing spins

Ian Harry,^{*} Stephen Privitera,[†] Alejandro Bohé,[‡] and Alessandra Buonanno[§]

*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1,
D-14476 Potsdam-Golm, Germany*

(Received 9 March 2016; published 5 July 2016)

Current searches for gravitational waves from compact-object binaries with the LIGO and Virgo observatories employ waveform models with spins aligned (or antialigned) with the orbital angular momentum. Here, we derive a new statistic to search for compact objects carrying generic (precessing) spins. Applying this statistic, we construct banks of both aligned- and generic-spin templates for binary black holes and neutron star–black hole binaries, and compare the effectualness of these banks towards simulated populations of generic-spin systems. We then use these banks in a pipeline analysis of Gaussian noise to measure the increase in background incurred by using generic- instead of aligned-spin banks. Although the generic-spin banks have roughly a factor of ten more templates than the aligned-spin banks, we find an overall improvement in signal recovery at a fixed false-alarm rate for systems with high-mass ratio and highly precessing spins. This gain in sensitivity comes at a small loss of sensitivity ($\lesssim 4\%$) for systems that are already well covered by aligned-spin templates. Since the observation of even a single binary merger with misaligned spins could provide unique astrophysical insights into the formation of these sources, we recommend that the method described here be developed further to mount a viable search for generic-spin binary mergers in LIGO/Virgo data.

DOI: [10.1103/PhysRevD.94.024012](https://doi.org/10.1103/PhysRevD.94.024012)

I. INTRODUCTION

On September 14, 2015, the Laser Interferometer Gravitational-Wave Observatory (LIGO) [1,2] made its first observation of gravitational waves, which were emitted by the binary black-hole merger dubbed GW150914 [3]. In the coming years, the field of gravitational-wave (GW) astronomy will begin to take off in earnest, with further upgrades to the LIGO detectors under way, and the expansion of the GW observatory network to include Advanced Virgo [4], KAGRA [5,6] and an additional LIGO observatory in India [7]. We expect not only to observe more binary black-hole mergers [8], but also signals from binary neutron-star and neutron star–black hole mergers [9].

Compact-object binary systems (henceforth, compact binaries) are thought to form through two channels: the coevolution of two massive stars in a binary [10–14], or by the dynamical capture of two preformed compact objects in dense stellar environments such as globular clusters [15–19]. The relative and absolute rates for these two potential formation channels are highly uncertain; they depend sensitively on a number of poorly constrained parameters, such as the typical stellar metallicity at formation, the distribution of supernova kicks, and the binding

energy of the common envelope. Indeed, clarification of the astrophysics of these formation channels is one of the great scientific promises of observing GW signals from compact-binary mergers [20–23]. The parameters of the merging binary are measurable through an observed GW signal [24,25], which can then provide information about the formation processes of the system.

In particular, it is thought that compact binaries formed by dynamical capture are more likely to have component angular momenta (spins) at large angles to the orbital angular momentum, while those formed by common evolution are more likely to have spins that are nearly aligned with the orbital angular momentum [19,26–28]. Present observations clearly indicate the potential for large spins on black holes in binaries, possibly close to the Kerr limit $|S/m^2| = |\chi| = 1$ [29–34]. Very few measurements of the angles between the spins and the orbital angular momentum exist from electromagnetic observations. In some cases, one can measure this spin misalignment via GW emission, as misalignment leads to precession of the orbital plane, which appears as phase and amplitude modulations in the observed signal [35–38]. However, for GW150914 it was not possible to constrain the spin misalignment [24].

We focus here on the effect of misalignment between the spin and the orbital angular momenta from the perspective of GW searches. The usual detection strategy for compact binaries is based on matched filtering of the data against a bank of templates spanning as densely as possible the full physical parameter space [39,40]. Covering the full

^{*}ian.harry@aei.mpg.de

[†]stephen.privitera@aei.mpg.de

[‡]alejandro.bohe@aei.mpg.de

[§]alessandra.buonanno@aei.mpg.de

parameter space is quite challenging. In fact, almost all searches of Initial LIGO and Initial Virgo data for compact-binary coalescences (CBCs) ignored the effect of spins in search templates [41–48]. The only search that did include spin effects in its templates [49] was a search using precessing-spin templates [50]. This search was later shown to perform, on average, no better than nonspinning searches [51], as the increased degrees of freedom in the signal space picked up extra noise, which offset the gains achieved in signal-power recovery. Several recent studies performed with aligned-spin templates have demonstrated that it is possible to pull in more signal than noise [52–55], and aligned-spin templates are currently used in searches with Advanced LIGO [56]. However, the effect of precession on the GW signal is not fully captured by aligned-spin templates [53,57,58], and it remains to be seen whether precessing-spin effects can be exploited to further improve the sensitivity of CBC searches. We note that while *detection* searches currently consider only aligned-spin templates, *parameter inference* algorithms, including the analysis of GW150914, do consider precessing waveforms [24,25]. Studies already exist in the literature demonstrating that for systems with strong precessional features it will be possible to break degeneracies that exist between physical parameters in the emitted gravitational-wave signal and therefore better measure the component spins, and potentially spin orientations [37,38]. However, to be able to make such measurements we must first be able to *detect* highly precessing systems, and the detection problem is the focus of this work.

In this work, we demonstrate a method for performing a detection search with precessing-spin templates. We restrict our attention to binary black hole (BBH) and neutron star-black hole (NSBH) systems, as neutron stars are generally expected to have smaller spins compared to black holes [59,60], and precessing-spin binary neutron star signals already match well with aligned-spin waveforms [57,61]. More specifically, we construct template banks of waveforms for precessing BBH systems with component masses $m_i \geq 3 M_\odot$, total masses $M_{\text{total}} \in [6, 100] M_\odot$ and mass ratios $m_1/m_2 \leq 5$, as well as for precessing NSBH systems with component masses $3 M_\odot < m_1 < 15 M_\odot$, $1 M_\odot < m_2 < 3 M_\odot$.

Our method makes use of a new statistic, which maximizes the matched-filter signal-to-noise ratio (SNR) over the detector-sky location of the source in addition to the phase and amplitude of the observed signal. We compare our new method to searches using aligned-spin templates and quantify the relative sensitivity between the two in a two-detector analysis of simulated Advanced LIGO noise. In doing so, we consider both the gain in signal power due to having templates that more accurately model precessing signals, as well as the increase in background from filtering the data with more templates. We find that, averaged across the parameter space, these two factors largely cancel out,

such that the expected rate of observations at a fixed rate of false alarms with the search presented here is roughly the same as that with an aligned-spin-only search. However, in certain regions of parameter space, namely at high mass ratios and large in-plane spins, we observe an increase in observation rate that is greater than 50%.

This paper is organized as follows. In Sec. II, we define the BBH and NSBH parameter spaces considered in this work, and justify the choice of waveform approximants used to model these systems. In Sec. III, we review the standard formulation of the phase-, amplitude- and time-maximized SNR currently used in searches with non-precessing templates, highlighting the assumptions which are invalid for precessing signals. Having thus laid out the mathematical formalism, we then derive in Sec. IV a phase-, amplitude-, time- and sky-location-maximized SNR statistic applicable to precessing templates, which we call the sky-maximized signal-to-noise ratio (sky-maxed SNR). We also present in Sec. IV a comparison between the sky-maxed SNR and previous methods, indicating in particular the reasons for which we find the sky-maxed SNR approach to be more promising. In Sec. V, we apply the sky-maxed SNR to construct banks of aligned-spin and precessing-spin template waveforms and demonstrate their coverage of the precessing signal space. In Sec. VI, we apply these template banks in a real pipeline analysis of simulated Gaussian noise in order to measure the increase in background incurred by using precessing template banks with the sky-maxed SNR. From the measured increase in background rate, we deduce the expected sensitivity improvement one could gain by employing such a statistic in Gaussian noise. Finally, in Sec. VII, we summarize the results, indicating the work remaining to turn the method presented here into a truly viable search method in real, non-Gaussian GW data.

II. PARAMETER SPACE AND WAVEFORM MODELS

In this paper, we consider two separate regions of the compact-binary parameter space, shown in Table I, corresponding to NSBH and BBH sources. We use different template waveform approximants in each region, based on the considerations below.

Let us first discuss the NSBH parameter space. For the purposes of detection with LIGO and Virgo, the signal from compact binaries with $M_{\text{total}} \lesssim 10\text{--}15 M_\odot$ are well modeled by the post-Newtonian (PN) approximation to the Einstein field equations [62], since only the inspiral portion of the signal is in band. Truncation of the various physical ingredients (energy, flux, precession equations, waveform amplitude) at different PN orders, together with distinct methods for reexpanding the balance equation when obtaining the frequency evolution equation, lead to a variety of so-called PN approximants, whose predictions for the signal can vary significantly [63,64]. In this paper,

TABLE I. Component mass and spin parameter ranges used to generate the waveforms used in this work for the BBH (left) and NSBH (right) parameter spaces. In both cases, the parameter space ranges are chosen identically for the template bank and the simulated signals used to test the coverage of the bank. For simulations, masses are drawn uniformly in component masses, spins are drawn uniformly in component spin magnitude, and all angular parameters are drawn isotropically. The parameter distribution for the stochastic placement of templates is detailed in the Appendix.

BBH parameter ranges	
m_1, m_2	$[3, 97]M_\odot$
M_{total}	$[6, 100]M_\odot$
m_1/m_2	$[1, 5]$
$ \chi_1 , \chi_2 $	$[0, 0.99]$
NSBH parameter ranges	
m_1	$[3, 15]M_\odot$
m_2	$[1, 3]M_\odot$
$ \chi_1 $	$[0, 1.0]$
$ \chi_2 $	$[0, 0.05]$

our goal is not to compare these different models but rather to understand the effect of adding precession to a search. While the details will certainly depend on the approximant that one considers, it seems reasonable to consider that the main effects will be captured by choosing one of them.

Specifically, when considering the NSBH parameter space, we use the TaylorT2 [63] approximant with all orbital phase-evolution terms up to 3.5PN order, all spin phase-evolution terms up to 2.5PN order, and using only the dominant amplitude term. The waveform is generated from a frequency lower than that used in the matched-filter integral and terminated at the frequency corresponding to the innermost stable circular orbit (ISCO) of a non-spinning black hole. In certain cases, the evolution reaches a minimum energy configuration before the ISCO is reached, and the waveform terminates at that point if this happens. For the purpose of bank placement, a faster way to evaluate waveform models is required. Thus, for template bank construction only, we use the closed-form, single-spin, frequency-domain precessing model introduced in Ref. [65]. This waveform is derived from the TaylorT2 reexpansion of the balance equation and has been shown to agree well with TaylorT2 in most cases [65].

For the larger masses considered in our BBH parameter space, the merger and ringdown portions of the signal enter the detectors' most sensitive band, and including these portions of the waveform becomes critical. Two approaches have been developed over the past years for the construction of approximate analytical models calibrated to numerical-relativity (NR) waveforms and covering the entire coalescence of the binary. One approach is the effective-one-body (EOB) formalism [66,67], which

combines a resummation of the available PN information and a description of the merger and ringdown phases with the calibration of a limited number of parameters (henceforth, EOBNR). The other approach is to construct phenomenological, frequency-domain models, which directly interpolate between NR waveforms hybridized to PN or EOB inspiral waveforms [68] (henceforth, IMRPhenom). Both approaches were successfully applied to the simpler problem of modeling aligned-spin systems [69–75] and have recently been extended to the case of precessing systems [36,76].

The template bank construction method that we use requires that the computational cost of generating waveforms be small. Precessing time-domain EOBNR models are currently orders of magnitude too slow for such a study; their use in these applications will require the development of fast frequency-domain reduced-order surrogate models [77], like those already constructed for aligned-spin EOBNR models [78,79]. Therefore, we compute template waveforms in our BBH parameter space using a phenomenological precessing approximant [36,80]. This approximant is constructed from an underlying aligned-spin model [74,75], which models the waveform in the coprecessing frame. Then, assuming a single in-plane spin, PN expressions are used to compute the precession angles as a function of orbital frequency. The $\ell = 2$ modes of the waveform in an inertial frame are then obtained by appropriately rotating the coprecessing aligned-spin waveform [36,80].

This IMRPhenom model has known pathologies in the region of parameter space where the projection of the total spin on the unit orbital angular momentum is large and negative as these configurations strongly violate one of the central assumptions in the model, namely that the direction of the total angular momentum remains approximately constant throughout the evolution [81]. A simple approximate and conservative boundary for the pathological region is given by $\chi_1 \cdot \hat{L} < -3/q$ where $q = m_1/m_2 > 1$ and \hat{L} denotes the normalized orbital angular momentum. As can be seen from this expression, this does not affect systems with $q < 3$ but covers an increasingly larger portion of the spin parameter space as q increases. For this reason, we decide to restrict our parameter space to $q \leq 5$, for which the problematic portion of parameter space (determined using the proper condition involving also the projection of the spin of the secondary body) is only a small fraction ($\approx 0.3\%$) of the full space. By this we mean that assuming the distributions described in the caption of Table I, only 0.3% of the simulations lie in this region. In the rest of this analysis, we excise this region from our parameter space both when assessing sensitivity improvement via sets of simulations and when constructing our template banks. This is also the reason why we do not use this model in the NSBH region of parameter space despite the obvious advantage of including the merger-ringdown portion and

therefore removing the uncertainties related to the termination of the waveform.

All the waveforms that we use in this paper are publicly available in the `lalsimulation` repository [82].¹

III. SKY-MAXED SNR: NONPRECESSING LIMIT

We now lay out the mathematical formalism we later use for deriving our sky-maxed SNR by first reviewing the phase-maximized matched filter SNR for nonprecessing templates. This method has been used in nearly all initial LIGO and Virgo CBC searches to date [39,41–47] and is well described in existing literature [39,40,83,84].

Consider the data output $s(t)$ of a GW detector, which consists of noise $n(t)$ and possibly a GW signal of known form $h(t)$. We wish to decide between the signal hypothesis and the noise hypothesis,

$$s(t) = \begin{cases} n(t) & \text{noise hypothesis} \\ n(t) + h(t) & \text{signal hypothesis} \end{cases},$$

given the observed data and predicted form of the signal. We assume that the noise is both stationary and Gaussian. Under these assumptions, the statistical properties of the noise are fully described by a single function, the one-sided ($f > 0$) noise power spectral density (PSD) $S_n(f)$, defined by

$$\frac{1}{2}\delta(f - f')S_n(f) = \mathbb{E}[\tilde{n}(f)\tilde{n}^*(f')], \quad (1)$$

where $\mathbb{E}[\cdot]$ denotes the expectation value over independent noise realizations. The PSD naturally induces a complex inner product $\langle | \rangle$ on the signal space,

$$\langle a|b \rangle = 4 \int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} df. \quad (2)$$

Using the Gaussian assumption, we can then express the probabilities $P(s|n)$ and $P(s|h)$ of the observed data given the signal and noise hypotheses, respectively, in terms of this inner product as

$$P(s|n) \propto e^{-\Re[\langle s|s \rangle]/2} \quad (3)$$

$$P(s|h) \propto e^{-\Re[\langle s-h|s-h \rangle]/2}, \quad (4)$$

¹The internal `lalsimulation` names for the waveforms described above are “SpinTaylorT2” and “SpinTaylorF2” for the two PN NSBH models, “SEOBNRv2” for the aligned-spin EOBNR approximant, “SEOBNRv3” for the precessing EOBNR approximant, “IMRPhenomPv2” for the precessing phenomenological approximant and “IMRPhenomD” for its aligned-spin counterpart.

where \Re denotes the real part. It follows that the likelihood ratio $\Lambda \equiv P(s|h)/P(s|n)$ between the signal and noise hypotheses is given by

$$\log \Lambda \equiv \lambda = \Re[\langle s|h \rangle] - \frac{1}{2}\Re[\langle h|h \rangle]. \quad (5)$$

By the Neymann-Pearson lemma, a search which thresholds on the matched-filter statistic given in Eq. (5) maximizes the probability of accepting the signal hypothesis whenever the signal hypothesis is true for any false-alarm probability. Equation (5) therefore gives the general prescription for searching for a GW signal of known form in stationary, Gaussian noise. One also defines the matched-filter SNR, ρ , by maximizing λ over an overall amplitude

$$\rho^2/2 = \max_{\text{amp}}(\lambda) = \frac{(\Re[\langle s|h \rangle])^2}{\langle h|h \rangle} = (\Re[\langle s|\hat{h} \rangle])^2, \quad (6)$$

where $\hat{x} \equiv x/\langle x|x \rangle^{1/2}$ is used to denote normalized waveforms.

In practice, the exact form of the signal is not known, and we must maximize Eq. (5) over the *a priori* unknown template parameters which determine the signal. A generic compact-binary coalescence GW signal is described by at least fifteen parameters²: the component masses, m_1 and m_2 ; the component dimensionless spin vectors χ_1, χ_2 ; the sky location of the signal with respect to the frame of the observer (θ, ϕ) ; the distance D to the source; the coalescence time t_c of the signal; the inclination of the binary with respect to the line-of-sight to the system i ; a polarization angle ψ ; and an orbital phase at coalescence ϕ_c . For some parameters, this maximization step can be performed analytically, or in a computationally efficient way using fast Fourier transform (FFT) algorithms, whereas for the remaining ones, one has to resort to discretizing the parameter space and repeating the matched-filter operation a large number of times. In this and the following sections, we are concerned with the possible analytic maximizations of Eq. (5). In Sec. V, we describe how we create banks of waveforms to optimize over the remaining parameters.

The observed signal $h(t)$ at the detector is the sum of the two GW polarizations, h_+ and h_\times , multiplied by the response function of the detector to each polarization, F_+ and F_\times [85], which encapsulate the full dependence of the signal on (θ, ϕ, ψ)

$$h(t) = F_+(\theta, \phi, \psi)h_+(t) + F_\times(\theta, \phi, \psi)h_\times(t). \quad (7)$$

It is this combination of h_+ and h_\times given in Eq. (7) that is used as the filter of Eq. (5).

²We restrict attention to compact binaries on circular orbits, removing parameters related to eccentricity. We also ignore any effects related to the internal structure of neutron stars.

In making the connection between the two waveform polarizations and the template waveform $h(t)$, aligned-spin searches rely on the simplifying assumption that only the dominant $(\ell, |m|) = (2, 2)$ modes of the waveform³ are taken into account. This allows us to write the dependence of the two GW polarizations on the physical parameters in the form [40]

$$\begin{aligned} h_+ &= \frac{1 + \cos^2 \iota}{2D} A(t - t_c; \xi) \cos[2(\Phi(t - t_c; \xi) + \phi_c)], \\ h_\times &= \frac{\cos \iota}{D} A(t - t_c; \xi) \sin[2(\Phi(t - t_c; \xi) + \phi_c)], \end{aligned} \quad (8)$$

where $A(t; \xi)$ and $\Phi(t; \xi)$ are functions of time and the parameters $\xi = (m_1, m_2, \chi_1, \chi_2)$, where χ_1 and χ_2 denote the constant projections of the spins in the direction of the orbital angular momentum. Inserting Eq. (8) into Eq. (7), we find that the full strain seen by a detector can now be written as

$$h(t) = \frac{A(t - t_c; \xi)}{D_{\text{eff}}} \cos[2(\Phi(t - t_c; \xi) + \phi_0)], \quad (9)$$

where

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

is the so-called effective distance and ϕ_0 , defined as

$$e^{2i\phi_0} = e^{2i\phi_c} \frac{F_+(1 + \cos^2 \iota)/2 - iF_\times \cos \iota}{[F_+^2 \left(\frac{1 + \cos^2 \iota}{2} \right)^2 + F_\times^2 \cos^2 \iota]^{1/2}}, \quad (11)$$

is the phase of the observed waveform at coalescence. Thus, in the aligned-spin case, the waveform $h(t)$ depends on the parameters ϕ_c , D , θ , ϕ , ψ and ι only through the combinations D_{eff} and ϕ_0 . Moreover, this dependence amounts only to an overall phase and an overall amplitude.

For search applications, we ultimately need the dependence of the Fourier transform of $h(t)$ on the physical parameters. Assuming that the time scale over which the amplitude A changes is much smaller than the orbital time scale we can apply the stationary phase approximation [39]. This allows us to conveniently factorize the dependence on ϕ_0 as

$$\tilde{h} = \frac{1}{D_{\text{eff}}} e^{2i\phi_0} \tilde{h}_0(f; t_c, \xi), \quad (12)$$

where we have defined

$$h_0(t - t_c, \xi) = A(t - t_c; \xi) \cos[2(\Phi(t - t_c; \xi))], \quad (13)$$

which depends only on $\xi = (m_1, m_2, \chi_1, \chi_2)$ and t_c .

We can now maximize Eq. (5) over the parameters ϕ_c , D , θ , ϕ , ψ , and ι by maximizing over the combinations of parameters D_{eff} and ϕ_0 . Inserting Eq. (12) into Eq. (5) and maximizing with respect to D_{eff} and ϕ_0 , we thus obtain

$$\max_{\iota, D, \theta, \phi, \psi, \phi_c} (\lambda) = \frac{1}{2} \max_{\phi_0} (\rho^2) = \frac{1}{2} |\langle s | \hat{h}_0 \rangle|^2. \quad (14)$$

The coalescence time t_c parametrizes time translations of h_0 . Therefore in the Fourier domain we can write $\tilde{h}_0(f; t_c, \xi) = \tilde{h}_0(f; \xi) e^{-2i\pi f t_c}$ and therefore

$$\langle s | h_0 \rangle = \int_0^\infty \frac{\tilde{s}^*(f) \tilde{h}_0(f; \xi)}{S_n(f)} e^{-2\pi i f t_c} df. \quad (15)$$

Evaluating $\langle s | h_0 \rangle$ over a range of t_c can be efficiently performed numerically by using widely available FFT routines.

To summarize, it is possible in the aligned-spin case to quickly maximize over all parameters describing the system except for $\xi = (m_1, m_2, \chi_1, \chi_2)$, provided subdominant modes can be neglected and that the stationary-phase approximation holds. The remaining parameters are searched over by repeating the matched-filter operation Eq. (14) over a discrete bank of templates, with time maximization handled by an efficient FFT implementation of Eq. (15). We discuss the construction of the discrete template banks in Sec. V.

IV. SKY-MAXED SNR: PRECESSING LIMIT

Consider now the case where we wish to conduct a search using waveforms with generically oriented spins. A first obvious difference with the case described above is that we now have to deal with the six dimensionless spin components: χ_1 and χ_2 .⁴ A more important difference is that in the precessing case the orientation of the source with respect to the detector varies as the orbit precesses. As a result, the two polarizations h_+ and h_\times cannot be written in the simple form of Eq. (8) where both are identical up to an amplitude rescaling that only depends on ι and a constant phase shift. Therefore, we return to Eq. (7) and derive a new statistic free of this assumption on the waveform. This statistic maximizes λ , not only over an amplitude and a phase, but also over the sky-location-dependent antenna factors. We then explore the statistical properties of this maximized form of λ and compare to previously proposed

³Defining these modes in a natural radiation frame where the z-axis, with respect to which the multipolar decomposition is performed, coincides with the direction of the angular orbital momentum.

⁴We do not need to specify here the frame used to define the spin components; note however that as the spins evolve with time in the precessing case, we define these values to be the spins at some reference time.

approaches, emphasizing the differences which make the current approach more promising.

A. Sky-maximized signal-to-noise ratio

We start by expressing the dependence in Eq. (7) on the detector related angles (θ, ϕ, ψ) and the distance D in terms of an overall amplitude and a phase between h_+ and h_\times ,

$$h = K(\theta, \phi, \psi, D)[h_+(t; t_c, \xi, \iota, \phi_c) \cos \kappa(\theta, \phi, \psi) + h_\times(t; t_c, \xi, \iota, \phi_c) \sin \kappa(\theta, \phi, \psi)], \quad (16)$$

where we have defined $\xi = (m_1, m_2, \chi_1, \chi_2)$ and

$$e^{i\kappa} = \frac{F_+ + iF_\times}{\sqrt{F_+^2 + F_\times^2}}, \quad (17)$$

$$K = \frac{1}{D} \sqrt{F_+^2 + F_\times^2}. \quad (18)$$

In addition, as in Eq. (12), we can factorize the dependence of ϕ_c in the Fourier domain as

$$\tilde{h} = K e^{2i\phi_c} [\tilde{h}_+(\phi_c = 0) \cos \kappa + \tilde{h}_\times(\phi_c = 0) \sin \kappa]. \quad (19)$$

As with the aligned-spin waveforms, this is not an exact symmetry. However, it is a particularly good approximation if one only considers a waveform containing the $(\ell, |m|) = (2, 2)$ modes in the corotating frame [36,50].

We could now perform the maximization over the amplitude K and the phase ϕ_c just as in Eq. (14) but with h_0 replaced by $\tilde{h}_+ \cos \kappa + \tilde{h}_\times \sin \kappa$. This would leave us with the two additional parameters ι and κ (plus the four new spin components) to be covered using a discrete bank. While it is possible to construct template banks in this manner, it would be desirable to further reduce the dimension of parameter space. Furthermore, different GW observatories, with different orientations and locations,

will not observe the *same* combination of sky angles κ . We therefore consider a scheme where we maximize not only over the overall amplitude K and phase ϕ_c , but also over the angle κ . Such a scheme removes all detector-dependent quantities from the parameters used when constructing the template bank, allowing the use of a simpler template coincidence method for a multidetector analysis.

Maximizing the log-likelihood defined in Eq. (5) over K and ϕ_c , straightforwardly leads to

$$\max_{K, \phi_c}(\lambda) = \frac{1}{2} \frac{u^2 |\hat{\rho}_+|^2 + 2u\hat{\gamma} + |\hat{\rho}_\times|^2}{u^2 + 2uI_{+\times} + 1}, \quad (20)$$

where we have defined

$$\hat{\rho}_{+, \times} = \langle s | \hat{h}_{+, \times} \rangle \quad (21)$$

$$\hat{\gamma} = \Re[\hat{\rho}_+ \hat{\rho}_\times^*] \quad (22)$$

$$\langle \hat{h}_+ | \hat{h}_\times \rangle = I_{+\times} + iJ_{+\times}, \quad (23)$$

with $I_{+\times}, J_{+\times} \in \mathbb{R}$ and we factorize the κ dependence in terms of

$$u \equiv \frac{1}{\tan \kappa} \sqrt{\frac{\langle h_+ | h_+ \rangle}{\langle h_\times | h_\times \rangle}}. \quad (24)$$

Taking the derivative of Eq. (20) with respect to u and solving for the roots leads to a quadratic equation in u ,

$$(I_{+\times} |\hat{\rho}_+|^2 - \hat{\gamma})u^2 + (|\hat{\rho}_+|^2 - |\hat{\rho}_\times|^2)u + (\hat{\gamma} - I_{+\times} |\hat{\rho}_\times|^2) = 0. \quad (25)$$

Substituting the roots of this equation back into Eq. (20), we obtain two extremal values for λ ,

$$\lambda = \frac{1}{4} \left(\frac{|\hat{\rho}_+|^2 - 2\hat{\gamma}I_{+\times} + |\hat{\rho}_\times|^2 \pm \sqrt{(|\hat{\rho}_+|^2 - |\hat{\rho}_\times|^2)^2 + 4(I_{+\times} |\hat{\rho}_+|^2 - \hat{\gamma})(I_{+\times} |\hat{\rho}_\times|^2 - \hat{\gamma})}}{1 - I_{+\times}^2} \right). \quad (26)$$

To take the maximal value of λ , we notice that the denominator of Eq. (26) is always positive, so the log-likelihood will always take a maximum value when the square-root term is positive. Therefore,

$$\max_{D, \phi_c, \theta, \phi, \psi}(\lambda) = \max_{K, \phi_c, u}(\lambda) = \frac{1}{4} \left(\frac{|\hat{\rho}_+|^2 - 2\hat{\gamma}I_{+\times} + |\hat{\rho}_\times|^2 + \sqrt{(|\hat{\rho}_+|^2 - |\hat{\rho}_\times|^2)^2 + 4(I_{+\times} |\hat{\rho}_+|^2 - \hat{\gamma})(I_{+\times} |\hat{\rho}_\times|^2 - \hat{\gamma})}}{1 - I_{+\times}^2} \right) \quad (27)$$

is the log-likelihood maximized over an overall phase, an overall amplitude and the sky location of the binary.

We notice that Eq. (27) is ill defined in the case that $I_{+\times} = \pm 1$. However, for compact-binary waveforms it is not possible for the \hat{h}_+ and \hat{h}_\times components to be identical

and so this case can never occur. Additionally, the terms within the square root of Eq. (27) will always take positive values and therefore the equation will always produce real, positive, values of λ . In analogy with the nonprecessing case discussed in the previous section, we define

$$\max_{K, \phi_c, u}(\lambda) = \frac{1}{2} \max_{\phi_c, u}(\rho^2) = \frac{1}{2} \rho_{\text{SM}}^2, \quad (28)$$

and we refer to the quantity ρ_{SM} as the sky-maxed SNR.

Additionally, Eq. (27) can be maximized over t_c in a similar way as in the aligned-spin case: both $\hat{\rho}_{+,x}$ can be efficiently computed for a discrete set of values of t_c using FFT algorithms and one can just pick the largest resulting value of λ . Note, however, that unlike in the aligned-spin case, maximizing the likelihood over the unknown coalescence time in the precessing case requires the computation of *two* inverse FFTs, which contributes to increasing the computational cost.

As a sanity check, we show how the sky-maxed SNR behaves in the aligned-spin limit. In this case, the simple relation between the polarizations Eq. (8) implies that in the frequency domain

$$\hat{h}_+ = \pm i \hat{h}_\times, \quad (29)$$

and therefore

$$\hat{\rho}_+ = \pm i \hat{\rho}_\times, \quad (30)$$

$$I_{+\times} \equiv \Re[\langle \hat{h}_+ | \hat{h}_\times \rangle] = 0. \quad (31)$$

Inserting these conditions into Eq. (27) results in the equation collapsing to the form of Eq. (14),

$$\max_{D, \phi_c, \theta, \phi, \psi}(\lambda) = \frac{1}{2} |\langle s | \hat{h}_+ \rangle|^2. \quad (32)$$

This is of course expected: in this case, the sky location enters only as a constant phase and amplitude shift and therefore by maximizing over these overall degrees of freedom, one has already maximized over κ .

B. Statistical distribution of the sky-maximized SNR in Gaussian noise

In Eqs. (27) and (28), we have defined a new statistic to be used in searches with precessing templates. This statistic has different statistical properties, in general, than the standard aligned-spin statistic defined in Eq. (6). Before applying the sky-maxed SNR, we wish to better understand these differences. Here, we investigate the distribution of the (squared) sky-maxed SNR ρ_{SM}^2 in Gaussian noise.

For a given template with a known value of $I_{+\times}$, the statistic defined in Eq. (27) is a combination of two

complex variables, $\hat{\rho}_+$ and $\hat{\rho}_\times$. We use the following notation for their real and imaginary parts,

$$\hat{\rho}_{+,x} = \epsilon_{+,x}^R + i \epsilon_{+,x}^I, \quad (33)$$

which are, for any point in time and in Gaussian noise, real Gaussian random variables with unit variance and zero mean. Thus,

$$\mathbb{E}[\epsilon_+^R] = \mathbb{E}[\epsilon_+^I] = \mathbb{E}[\epsilon_\times^R] = \mathbb{E}[\epsilon_\times^I] = 0, \quad (34)$$

$$\text{Var}[\epsilon_+^R] = \text{Var}[\epsilon_+^I] = \text{Var}[\epsilon_\times^R] = \text{Var}[\epsilon_\times^I] = 1. \quad (35)$$

However, while by definition the imaginary and real components of $\hat{\rho}_+$ (and separately of $\hat{\rho}_\times$) are statistically independent,

$$\mathbb{E}[\epsilon_+^R \epsilon_+^I] = \mathbb{E}[\epsilon_\times^R \epsilon_\times^I] = 0, \quad (36)$$

the correlation between $\hat{\rho}_+$ and $\hat{\rho}_\times$ will depend on the template and PSD being used.

In the aligned-spin case, we see directly from Eq. (32) that 2λ is the sum of two independent Gaussian variables squared and therefore has a χ^2 distribution with 2 degrees of freedom. In the generic case, the statistical distribution will depend on both the real and imaginary parts of $\langle \hat{h}_+ | \hat{h}_\times \rangle$: $I_{+\times}$ and $J_{+\times}$. In order to explore this, it is more convenient to reexpress Eq. (27) in terms of four *independent* standard normal variables. We therefore introduce a normalized linear combination,

$$\hat{h}_\perp = \frac{\hat{h}_\times - \langle \hat{h}_+ | \hat{h}_\times \rangle \hat{h}_+}{\sqrt{1 - |\langle \hat{h}_+ | \hat{h}_\times \rangle|^2}}, \quad (37)$$

such that $\langle \hat{h}_+ | \hat{h}_\perp \rangle = 0$. From this we obtain a new complex variable $\hat{\rho}_\perp$ whose real and imaginary parts $\epsilon_\perp^{R,I}$ are standard normal variables statistically independent of $\epsilon_+^{R,I}$. Reexpressing $\hat{\rho}_\times$ as a linear combination of $\hat{\rho}_+$ and $\hat{\rho}_\perp$,

$$\hat{\rho}_\times = \langle \hat{h}_+ | \hat{h}_\times \rangle \hat{\rho}_+ + \sqrt{1 - |\langle \hat{h}_+ | \hat{h}_\times \rangle|^2} \hat{\rho}_\perp, \quad (38)$$

and plugging this into Eq. (27), we rewrite the statistic in terms of four statistically independent, zero-mean, unit-variance vectors and the real and imaginary components of $\langle \hat{h}_+ | \hat{h}_\times \rangle$,

$$\begin{aligned} \lambda = & [H_-^2 |\hat{\rho}_+|^2 + H_+^2 |\hat{\rho}_\perp|^2 - 2J_{+\times} H_+ (\epsilon_+^I \epsilon_\perp^R - \epsilon_+^R \epsilon_\perp^I) + (H_+^4 (|\hat{\rho}_+|^4 + |\hat{\rho}_\perp|^4) \\ & + 2J_{+\times} H_+^3 (\epsilon_+^I \epsilon_\perp^R - \epsilon_+^R \epsilon_\perp^I) (|\hat{\rho}_+|^2 - |\hat{\rho}_\perp|^2) + H_+^2 \{8J_{+\times}^2 (\epsilon_+^I \epsilon_\perp^R - \epsilon_+^R \epsilon_\perp^I)^2 \\ & + 2H_-^2 ((\epsilon_+^R \epsilon_\perp^R)^2 - (\epsilon_+^I \epsilon_\perp^I)^2 + 4\epsilon_+^R \epsilon_\perp^I \epsilon_+^R \epsilon_\perp^I - (\epsilon_+^R \epsilon_\perp^I)^2 + (\epsilon_+^I \epsilon_\perp^R)^2)] / (4(1 - I_{+\times}^2)). \end{aligned} \quad (39)$$

Here we have defined for convenience

$$H_- = \sqrt{1 - (I_{+\times}^2 - J_{+\times}^2)}, \quad (40)$$

$$H_+ = \sqrt{1 - (I_{+\times}^2 + J_{+\times}^2)}, \quad (41)$$

and we remind the reader for completeness that

$$|\hat{\rho}_{+,\pm}|^2 = (\epsilon_{+,\pm}^R)^2 + (\epsilon_{+,\pm}^I)^2. \quad (42)$$

This fully explicit form, although not very elegant, allows us to easily identify some particular cases and symmetries. First, as already discussed, in the aligned-spin case for which $J_{+\times} = \pm 1$, $I_{+\times} = H_+ = 0$ and $H_- = \sqrt{2}$, the distribution trivially collapses to a χ^2 distribution with 2 degrees of freedom $2\lambda = (\epsilon_+^R)^2 + (\epsilon_+^I)^2$. Another interesting case arises when $J_{+\times} = 0$ since in this case $H_-^2 = H_+^2 = 1 - I_{+\times}^2$ and the dependence on $I_{+\times}$ completely cancels out. Furthermore, Eq. (39) allows us to show that the distribution of λ does not depend on the sign of $I_{+\times}$ and $J_{+\times}$. Indeed, Eq. (39) is left invariant by the transformation $I_{+\times} \rightarrow -I_{+\times}$ or by the transformation $(J_{+\times}, \epsilon_{\perp}^R, \epsilon_{\perp}^I) \rightarrow (-J_{+\times}, -\epsilon_{\perp}^R, -\epsilon_{\perp}^I)$.

As we have demonstrated, λ is formed from a combination of four orthogonal time series which, in Gaussian noise, will each be independent, and follow a normal distribution with zero mean and unit variance. We identify an upper bound on the probability density function (PDF) of the sky-maxed SNR at large SNRs by considering the case where one is free to capture the power in all four of these vectors, i.e. when the four components are added in quadrature. In this case 2λ would follow a χ^2 distribution with 4 degrees of freedom. However, λ as defined in Eq. (39) does not have the freedom to capture the power in all four of these vectors; it is constrained to the physical subspace. We find numerically that the PDF of the sky-maxed SNR takes the largest values at high SNRs when $\langle \hat{h}_+ | \hat{h}_\times \rangle = 0$. The PDF takes the smallest values at high SNRs in the case where $\langle \hat{h}_+ | \hat{h}_\times \rangle = \pm i$, as for nonspinning or aligned-spin restricted waveforms. This can be seen in Fig. 1 where we compare the distribution of $2\lambda = \rho_{\text{SM}}^2$ for both the lower and upper bound configurations with χ^2 distributions with 2, 4 and 6 degrees of freedom. We note that the lower bound at high SNRs for the ‘‘PTF’’ approach, which we discuss in the next subsection, follows a χ^2 distribution with 6 degrees of freedom.

Since the distribution of our sky-maxed SNR depends on the value of $\langle \hat{h}_+ | \hat{h}_\times \rangle$ for the template that one is considering, we find it informative to visualize the distribution of $\langle \hat{h}_+ | \hat{h}_\times \rangle$ corresponding to a set of precessing waveforms randomly drawn from our parameter spaces of interest, as shown in Fig. 2. The particular distribution that the simulations in the panels of this figure are drawn from

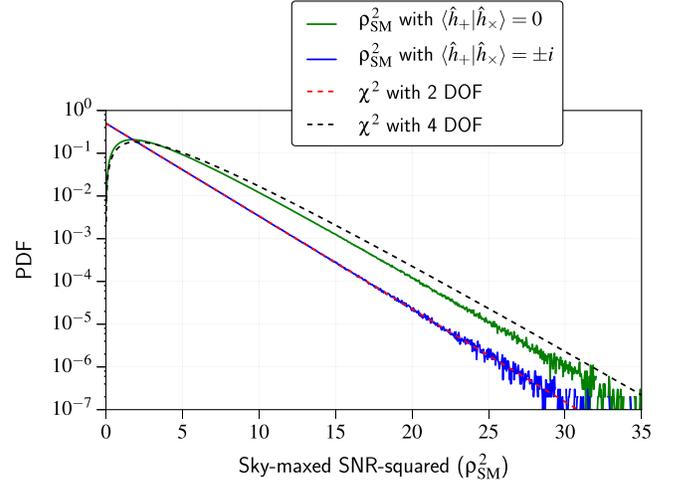


FIG. 1. The two limiting cases, $(\langle \hat{h}_+ | \hat{h}_\times \rangle = 0)$ and $(\langle \hat{h}_+ | \hat{h}_\times \rangle = \pm i)$, of the probability density function of the precessing sky-maxed SNR in Gaussian noise. Also plotted are χ^2 distributions with 2 and 4 degrees of freedom for direct comparison. As expected, when $\langle \hat{h}_+ | \hat{h}_\times \rangle = 0$ the statistic follows a χ^2 distribution with 2 degrees of freedom.

are listed in Table I. We see that even though the set of aligned-spin waveforms is of measure zero, the distribution is highly peaked around the aligned-spin value $\pm i$.

C. Comparison to previously proposed approaches

The first attempt to derive an efficient search for precessing waveforms in GW data analysis was presented in Ref. [86]. This approach involved adding an additional free parameter at the same order as the 1PN orbital term, to model the dominant effect of precession on the orbital phase. A later scheme for searching with precessing templates was introduced in Ref. [87], which used unphysical coefficients in the waveform expansion to model the effects of precession in a parameter space of reduced dimension. The approach described in Ref. [87] was used to search for precessing signals in LIGO data [44,49]. However, it was determined that due to the increased response to the noise background, this method did not offer any improvement in sensitivity compared to a nonspinning search pipeline [51]. The basic problem was that while the increased parameter space offered an improvement in the signal power recovered for precessing signals, the large, additional, unphysical parameter space being searched greatly increased the rate of background triggers of the search. Therefore, when estimating sensitivity at a constant rate of false alarm, the precessing search loses sensitivity compared to a nonspinning or aligned-spin search.

An alternative to these ‘‘unphysical template families’’ is to use a method that restricts to only physically possible template waveforms. A method for doing this was first proposed in Ref. [88] and then further explored in

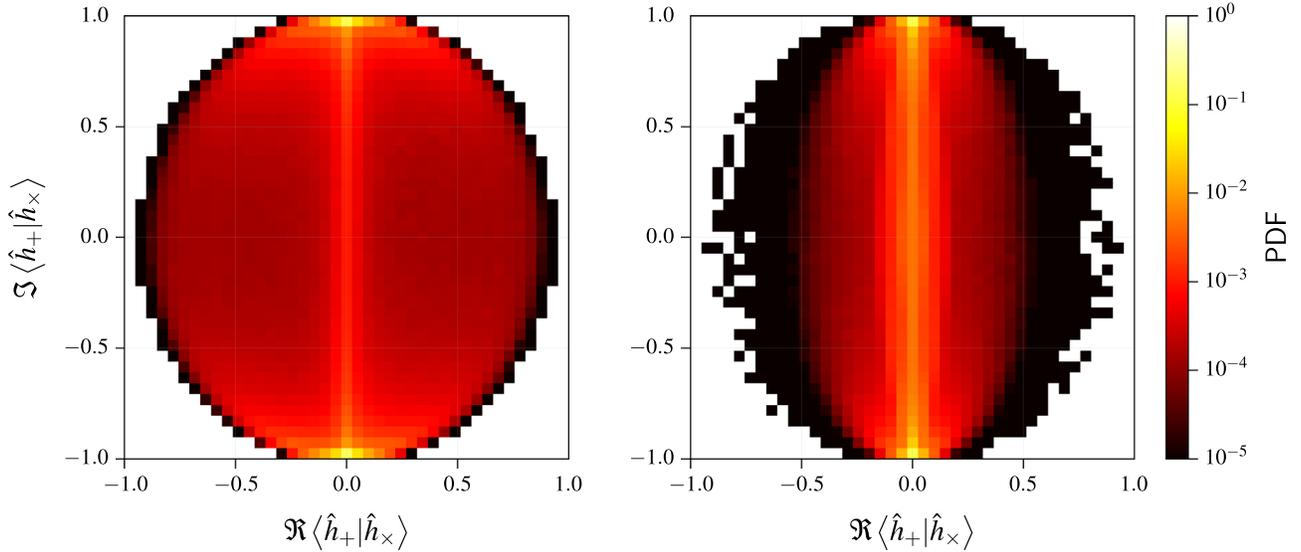


FIG. 2. The distribution of $\langle \hat{h}_+ | \hat{h}_\times \rangle$ for (left) a set of BBH signals modelled with the PhenomP approximant and (right) a set of NSBH signals modeled with the TaylorT2 approximant, each set drawn with the distribution described in Sec. IV B. Most waveforms cluster around the $\langle \hat{h}_+ | \hat{h}_\times \rangle = \pm i$ value achieved exactly by aligned-spin systems.

Refs. [50,89,90]. We will refer to this as the ‘‘PTF’’ (physical template family) approach. Our sky-maxed SNR is motivated from the PTF approach and we compare the two methods here.

The PTF approach considers single-spin systems and considers only the $(\ell, |m|) = (2, 2)$ modes in a corotating frame where the z direction tracks the orbital angular momentum. This is similar to our approach, except we do not restrict ourselves to considering single-spin systems. The single-spin approximation was explored, however, in terms of an effective spin for double-spin systems, and found to perform well, in Ref. [50]. The PTF approach parametrizes the single spin by the spin magnitude, χ_1 ; the cosine of the angle between the spin and the orbital angular momentum, $\kappa_1 = \hat{\mathbf{L}}_N \cdot \boldsymbol{\chi}_1$; and the azimuthal angle φ between the projections of the spin and the line of sight on the orbital plane, all these quantities being computed at some reference time. The PTF SNR is then constructed by reexpressing h_+ and h_\times as the sum of five basis waveforms constructed using a spherical harmonic ($\ell = 2$) basis to express the waveform as [88]

$$h(t) = P_I[D, \theta, \psi, \phi, \iota, \varphi] Q^I[m_1, m_2, \chi_1, \kappa_1; \phi_c, t_c; t], \quad (43)$$

where I takes values $\in [1, 5]$. Next, a maximization is performed over the five P_I constants, ϕ_c and t_c to obtain the PTF SNR as described in Refs. [50,90].

The statistical distribution of the PTF SNR, as described in Ref. [90], is template and noise curve dependent. However, in the best case scenario, when the PDF of the PTF SNR at high SNRs takes the smallest values, the distribution of SNRs follows a χ^2 distribution with 6 degrees of freedom. This can be compared to the distribution of values for the sky-maxed SNR, as shown in Fig. 1. However, as summarized in Table II, the PTF SNR maximizes over more parameters than the sky-maxed SNR. Specifically, considering a single-spin template, both the inclination angle and the azimuthal spin angle are maximized over in the PTF SNR, but not maximized over in the sky-maxed SNR.

The PTF SNR, as described above, is *not* restricted to physical values of the parameters. While the five P^I values

TABLE II. Comparison between the maximization schemes proposed in the PTF method [88] and in this paper. PTF lays a discrete bank in a 4d space and then ‘‘continuously’’ maximizes over the remaining $12 - 4 = 8$ parameters using a combination of analytic and numerical methods. Note that the total number of parameters is 12 instead of 15 due to their restriction to single-spin systems. In this paper, we essentially follow the same route except that (i) we do not restrict ourselves to single-spin systems and (ii) we include ι and φ in the discrete bank. Note that in PTF the parameter φ (together with χ_1 and κ_1) describes the orientation of the spin in the source frame. In our notation, these three parameters combine into $\boldsymbol{\chi}_1$.

Maximization scheme	Discrete parameters in the bank	Parameters continuously optimized over
Sky-maximized SNR	$m_1, m_2, \chi_1, \chi_2, \iota$	$\phi_c, t_c, \theta, \phi, \psi, D$
PTF SNR (constrained max)	$m_1, m_2, \chi_1, \kappa_1$	$\iota, \varphi, \phi_c, t_c, \theta, \phi, \psi, D$

do depend on six physical parameters, the values of D, θ, ψ, ϕ are degenerate in a single detector, and enter the waveform in only two combinations, an overall amplitude scaling, and F_+/F_\times . Therefore, only four independent physical values can be measured from the P^I and it then follows that allowing a free maximization over five values allows for unphysical combinations. The authors of Ref. [50] explored this and described a numerical method for constraining the P^I values to the physical manifold. However, that method is computationally expensive and was not included in attempts to use the PTF approach as a search method [89,90].

In Ref. [50] the authors considered searching for precessing signals in data from only a single detector. Here we wish to consider the case of a multidetector analysis. The standard approach is to measure the SNR from data in each detector *independently* and then look for times where both detectors obtain a SNR over some threshold, within some predefined time window and within some predefined window on the template parameters that are gridded over. No attempt is made to ensure consistency in the parameters that are maximized over. For a two-detector aligned-spin search one maximizes over time, an overall waveform phase, and an overall amplitude in each detector. As long as the time difference is within that allowed due to the light travel time between detectors it is always possible to find a physical solution for the maximized parameters. For more than three detectors this is not the case, and the aligned-spin coincident search will allow unphysical combinations of the maximized parameters. In that case the computationally more expensive coherent search [90–94] can offer some improvement in sensitivity by restricting the search to only physically possible values of the maximized parameters.

In the PTF approach, even considering the constrained statistic, one measures 6 maximized parameters ($P^{1..4}, \phi_c, t_c$) in each detector. With two detectors these are measured independently and so we obtain 12 independent quantities measuring only 8 physical parameters. For the nonconstrained statistic 14 independent quantities are maximized over, measuring the same 8 physical parameters. This results in a significantly large unphysical region of parameter space being searched over when considering multiple detector searches and makes it difficult to see sensitivity gains over the aligned-spin searches.

In the sky-maxed SNR, we reduce six physical parameters to four nondegenerate combinations of them that are maximized over. For a single detector, the resulting values are always consistent with some physical signal. When filtering in two detectors, however, we have obtained eight independent measurements for only six physical quantities. It therefore follows that, as in the PTF case, when filtering with two or more detectors, some degree of unphysical freedom is still allowed with this new statistic. However, this new statistic allows *less* unphysical freedom than in the

PTF case and therefore should offer a better chance to create a search that increases sensitivity to precessing systems.

V. TEMPLATE BANK CONSTRUCTION

We now apply the sky-maxed SNR to generate banks of generic-spin BBH and NSBH templates. We also quantify the expected improvement in SNR recovery from a simulated population of precessing-spin systems when using these generic-spin template banks, compared to aligned-spin-only banks. In each case, we first generate a bank of aligned-spin templates, and then add precessing templates to the aligned-spin bank in a second stage. By construction, since the aligned-spin bank is a subset of the generic-spin bank, the generic-spin template bank will outperform the aligned-spin bank towards any putative signal when performance is measured in terms of SNR recovery. We remind the reader that comparing the expected SNR recovery between the two banks does not include the effect of the increase in background event rate due to the increase in template bank size and increased degrees of freedom incurred by using the precessing templates. We consider this effect in the next section, in which we demonstrate the application of the precessing template banks in an analysis of simulated Advanced LIGO noise.

A. Precessing bank construction method

The construction of efficient template banks to search for nonspinning compact-binary mergers has been well explored in the literature [40,84,95–100]. These methods define a metric in the parameter space of the two masses, and use this metric to place a hexagonal lattice in appropriate coordinates. Recently, this method has been extended to place geometrical lattices of aligned-spin waveform templates [53,61]. However, for the parameter space of precessing compact-binary mergers, it is not clear how to choose a coordinate system in which it is appropriate to lay a lattice of waveform templates. Indeed, there is no reason to believe that such an intrinsically flat parameter space exists for the precessing parameter space. In addition, current geometric methods exclude the effects of merger and ringdown in the waveform model. This can cause these banks to be suboptimal in the BBH parameter space where merger and ringdown are important.

For these reasons, we make no attempt to employ a geometrical template bank for our precessing search. Instead, we use a “stochastic” template bank construction scheme [57,101,102]. The stochastic template scheme has the advantage that it is able to create a bank of waveform signals for any parameter space, but the disadvantage that it will require more templates to cover a given parameter space than a geometrical lattice, and can be considerably more computationally expensive [102,103]. In recent years,

a number of methods have been proposed to significantly speed up the generation of stochastic banks [55,57,104]. We use a number of these methods here, along with some new methods to optimize bank placement, which we describe in the Appendix.

The general stochastic approach works as follows. Begin with a seed bank B , which may be empty. Then randomly choose the parameters of a putative template signal and compute the corresponding template waveform g_{prop} . Then a match, $\mathcal{M}(g_{\text{prop}}, h)$, is computed between g_{prop} and all templates $h \in B$, where $\mathcal{M}(g_{\text{prop}}, h)$ defines the fraction of the signal power of g_{prop} that would be recovered if using the template h as a filter. The fitting factor FF is then defined as the match maximized over all templates in the bank,

$$\text{FF} \equiv \max_{h \in B} \mathcal{M}(g_{\text{prop}}, h). \quad (44)$$

If the fitting factor falls below a given threshold—the “minimal match”—then the proposed template is added to the bank and $B' = B \cup \{h_{\text{prop}}\}$ is set as the seed bank for the next iteration. Otherwise, the proposed template is discarded, and $B' = B$. The process repeats until a sufficiently high rejection rate of proposed templates is achieved.

For aligned-spin placement, the match $\mathcal{M}(g_{\text{prop}}, h)$ defined from the phase-maximized matched-filter SNR in Eq. (14) is

$$\begin{aligned} \mathcal{M}(g_{\text{prop}}, h) &\equiv \max_{\phi_c, t_c} \Re[\langle \hat{g}_{\text{prop}} | \hat{h}(\phi_c, t_c) \rangle] \\ &= \max_{t_c} |\langle \hat{g}_{\text{prop}} | \hat{h}(0, t_c) \rangle|. \end{aligned} \quad (45)$$

We note that, as written, the maximization over time and phase shift is performed only on the template h and not on the proposed waveform g_{prop} . However, given the assumptions that are used to construct Eq. (14), a phase or time shift in h can be modeled by an opposite phase or time shift in g_{prop} and therefore the form given in Eq. (45) serves to maximize over a phase and time shift in both h and g_{prop} . Therefore, when choosing seed points for aligned-spin stochastic template bank construction, one only needs to choose the masses and spins.

For precessing waveforms, we define the match as

$$\mathcal{M}(h, g) \equiv \max_{\phi_c, t_c, u} \Re[\langle \hat{g}_{\text{prop}} | \hat{h}(\phi_c, t_c, u) \rangle], \quad (46)$$

where the maximization is performed as described in Eq. (27). In this case, a variation in the value of u for \hat{g}_{prop} cannot be written as a corresponding shift of the value of u in h . Therefore, in the precessing case, we maximize over ϕ_c and u only in the template waveform. That is, when picking the putative template signal we choose a sky

location, construct $h(t)$ and compare that against the h_+ and h_\times components of all templates in the bank, using Eq. (27). If a putative point is accepted into the template bank, the sky location and $h(t)$ are discarded, and only the h_+ and h_\times components of that point are retained. Our choice not to maximize over ϕ_c and u in \hat{g}_{prop} is taken to allow us to use Eq. (27) directly when evaluating the fitting factor for potential precessing filter waveforms. This choice will result in an increase in the number of proposal points needed to construct a precessing template bank than if one were to consider also maximizing over ϕ_c and u in \hat{g}_{prop} . However, the final number of templates in the resulting template bank should not be affected by this choice.

B. Effective fitting factor

In order to quantify a template bank’s performance in terms of SNR recovery, we use the notions of *signal recovery fraction* and *effective fitting factor* as figures of merit. These notions were defined initially in Refs. [53,87], and we redefine them here for completeness. Consider a template bank and a model $p(\mathbf{v})$ for the distribution of source parameters. We assume that sources are uniformly distributed in volume, so that $p(\mathbf{v}) \propto r^2 p(\mathbf{v}')$ where \mathbf{v}' denotes all parameters other than distance. For a “perfect” template bank, where all fitting factors are unity, the expected total number of sources that would be observed above a SNR threshold ρ_0 is proportional to

$$N_{\text{opt}} \propto \int \sigma^3(\mathbf{v}') p(\mathbf{v}') d\mathbf{v}', \quad (47)$$

where $\sigma(\mathbf{v}')$ is the distance at which the expected SNR to the signal with parameters \mathbf{v}' is equal to ρ_0 . In reality, our template banks will not have a fitting factor of 1 for the entire parameter space, and therefore the number of observed signals above a SNR threshold ρ_0 will be smaller than N_{opt} according to

$$N_{\text{obs}} \propto \int \text{FF}^3(\mathbf{v}') \sigma^3(\mathbf{v}') p(\mathbf{v}') d\mathbf{v}', \quad (48)$$

where $\text{FF}(\mathbf{v}')$ denotes the fitting factor between the signal with parameters \mathbf{v}' and the template bank.

We then define the “signal recovery fraction” α as the ratio between N_{obs} and N_{opt} ,

$$\alpha \equiv \frac{N_{\text{obs}}}{N_{\text{opt}}} = \frac{\int \text{FF}^3(\mathbf{v}') \sigma^3(\mathbf{v}') p(\mathbf{v}') d\mathbf{v}'}{\int \sigma^3(\mathbf{v}') p(\mathbf{v}') d\mathbf{v}'}, \quad (49)$$

which takes values between 0 and 1. It is also convenient to express the bank performance in terms of the “effective fitting factor,” defined as

$$\text{FF}_{\text{eff}} \equiv \alpha^{1/3}, \quad (50)$$

which can be interpreted as the average SNR recovered for the observed population of sources. Including the σ factors means that signals that would not be seen at a large distance have only a weak effect on the signal recovery fraction. This includes signals that are poorly aligned with respect to the detector, and signals that have intrinsically low GW luminosity. In contrast, favorably oriented, intrinsically bright sources will have the largest effect on this measure.

We estimate the signal recovery fraction numerically by Monte Carlo integration, choosing a random set of source parameters $S = \{\mathbf{v}_i\}_{i=1}^N$ according to a sampling distribution $q(\mathbf{v}')$. The signal recovery fraction is then given by

$$\alpha \approx \frac{\sum_{i=1}^N \text{FF}_i^3 \sigma_i^3 (p_i/q_i)}{\sum_{i=1}^N \sigma_i^3 (p_i/q_i)}, \quad (51)$$

where $p_i/q_i = p(\mathbf{v}'_i)/q(\mathbf{v}'_i)$ corrects the sampling distribution q to match the desired astrophysical distribution p . Table I summarizes the distribution $q(\mathbf{v}')$ that we use when drawing signals to evaluate both our NSBH and BBH template banks. The template banks are constructed using the same limits on physical parameters.⁵

We wish to sample well all points in our parameter space, and so we choose sources uniformly distributed in component masses within the bounds for the BBH and NSBH parameter spaces given in Table I. However, such a choice leads to an effective fitting factor that is dominated by high mass systems since, to leading order, $\sigma \propto \mathcal{M}^{5/6}$. This would be the correct figure of merit if the distribution of masses for compact-binary mergers in the Universe were uniform in component masses. However, this distribution is not well understood [105,106]. To obtain a figure of merit for our banks that more evenly averages over the mass space, we therefore correct the sampling distribution by a factor $p/q = \mathcal{M}^{-5/6}$ to approximate a prior in mass such that the *observation rate* does not change as a function of component masses. Thus, we evaluate the signal recovery fraction from our simulations by the formula

$$\alpha \approx \left(\frac{\sum_{i=1}^N \text{FF}_i^3 \sigma_i^3 \mathcal{M}_i^{-5/6}}{\sum_{i=1}^N \sigma_i^3 \mathcal{M}_i^{-5/6}} \right). \quad (52)$$

To minimize issues of uncertainty in the real astrophysical mass distribution, we also report below results restricted to relatively small mass bins. However, when reporting results

⁵As discussed in Sec. II, the frequency-domain NSBH precessing waveform approximant that is used when placing the NSBH precessing template bank is a single-spin model. Therefore, NSBH precessing templates all have $\chi_2 = 0$, although the aligned-spin NSBH template bank contains templates with $\chi_2 \neq 0$. When calculating the effective fitting factor and associated quantities all NSBH templates and simulated signals are modeled using the double-spin time-domain TaylorT2 waveform model.

TABLE III. Sizes and effective fitting factors for the BBH and NSBH template banks. These cover the same range of parameters from which the corresponding signal set is drawn from, as described in Table I. When computing matches and constructing the template banks a lower frequency cutoff of 30 Hz is used and the Advanced LIGO “early” noise curve prediction [107]. The effective fitting factor reported in the last column is computed using Eq. (50).

Parameter space	Minimal match	Spin	Templates	Eff. FF
NSBH	0.97	Aligned	146,315	0.948
	0.90	Precessing	1,583,079	0.976
BBH	0.97	Aligned	23,948	0.984
	0.90	Precessing	237,909	0.988

in bins of spin, the chosen mass distribution will matter, and this distribution more evenly weights the simulations.

C. Effectualness of precessing template banks

We now generate template banks to cover the parameter spaces that are summarized in Table I. Our aligned-spin banks, which form the seed for the precessing banks, are generated using a minimal match of 0.97, which matches the value used for aligned-spin searches of Advanced LIGO and Advanced Virgo data [55]. The precessing banks become unmanageably large with our current methods if generated with a minimal match of 0.97. We therefore use a minimal match of 0.9 when completing our aligned-spin banks with precessing waveforms.

The sizes of the resulting template banks and the corresponding effective fitting factors of those banks are summarized in Table III, and the distribution of fitting factors for each bank is plotted in Fig. 3. We find that the precessing template banks are roughly an order of magnitude bigger than their aligned-spin counterparts, and that,

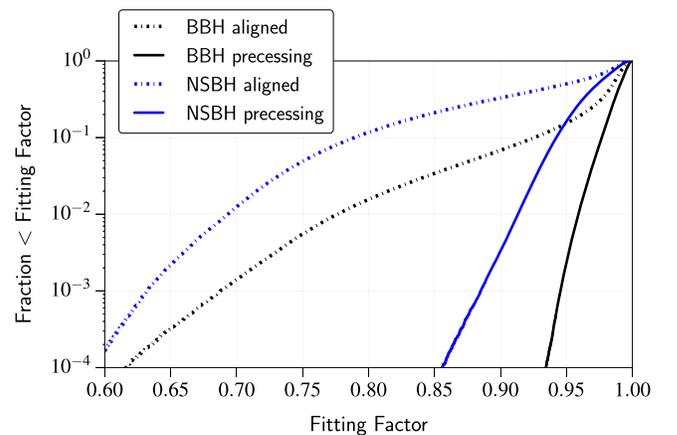


FIG. 3. The distribution of fitting factors for the aligned-spin and precessing-spin template banks summarized in Table III and covering the parameter spaces defined in Table I.

on average, aligned-spin banks are already performing reasonably well when searching for precessing systems. The BBH aligned-spin bank is more effective at recovering precessing BBH signals ($FF_{\text{eff}} = 0.984$) than the NSBH aligned-spin bank is at recovering precessing NSBH systems ($FF_{\text{eff}} = 0.948$). However, both aligned-spin bank distributions show long tails of precessing systems that are recovered with fitting factors less than 0.90. When using precessing template banks, these tails are significantly reduced. The effective fitting factor also increases to $FF_{\text{eff}} = 0.988$ for the BBH parameter space and $FF_{\text{eff}} = 0.976$ for the NSBH parameter space. While the increase in the overall effective fitting factor—averaged over the full parameter space—is small our precessing template banks seem to reduce an observational bias against highly precessing signals that will be present in current aligned-spin searches. We explore this further in the next section when we put these numbers into context by taking into account the increase in background incurred from filtering the data against a larger number of templates.

VI. ASSESSING THE SENSITIVITY OF THE PRECESSING SEARCH

The increase in templates from including the effects of precession, coupled with the fact that the precessing-spin templates on average produce larger background SNR values than the aligned-spin templates with the sky-maxed SNR, leads to an increase in the rate of false alarms at a given SNR. Correctly estimating the increase in observed signals with our precessing template banks requires assessing the sensitivity of the aligned-spin and precessing searches at a constant false-alarm rate, which we take on in this section. To do so, we incorporate the sky-maxed SNR into the PyCBC search pipeline described in [108,109] and, using each of the banks constructed in the previous section, perform a two-detector analysis of Gaussian noise. These analyses give us a direct measurement of the increase in background trigger rate, which we then combine with the fitting factor calculations above to estimate the change in detection rate of compact-binary systems at fixed false-alarm rate when using precessing templates instead of only aligned-spin templates.

A. Mapping between signal-to-noise ratio and false-alarm rate

In performing the precessing search, we split the original aligned-spin templates from the templates that were added to cover the precessing space and analyze these two sets separately. The “precessing-spin search” is then formed by combining the aligned-spin-only analysis with the precessing-spin-only analysis, giving equal weight to each “subsearch.” This choice amounts to assuming that an event is equally likely to appear in the aligned-spin set of templates as in the precessing set. A better choice for this could be

made, but other studies suggest it will not drastically affect our conclusions [110,111]; we leave this to future work. When assessing the sensitivity of the aligned-spin search alone, we use a false-alarm threshold of 1 in 100 years. When assessing the sensitivity of the combined search, we use a false-alarm threshold of 0.5 in 100 years for *both* the aligned-spin-only and precessing-spin-only subsearches, corresponding to a false-alarm threshold of 1 in 100 years for the full precessing-spin search.

The result of these analyses on the aligned-spin-only and precessing-spin-only template banks covering the BBH and NSBH parameter spaces can be seen in Fig. 4. The relevant thresholds are enumerated in Table IV. The table shows that a signal that appears in the BBH aligned-spin bin with a SNR of 9.64 would be deemed as significant in the BBH aligned-spin search as a signal that appears in the BBH precessing-spin bin with a SNR of 10.10 in the full BBH precessing-spin search. Thus, a precessing BBH signal would only be found with higher significance by the precessing-spin search if the precessing-spin templates increase its SNR by at least the factor $10.10/9.64 \approx 1.05$. Similarly, a precessing NSBH signal would only be found with higher significance by the NSBH precessing-spin search if the precessing-spin templates increase its SNR by at least the factor $10.44/9.79 \approx 1.07$. On the other hand, a given aligned-spin signal will always be found with a lower significance by the precessing-spin search, corresponding to a loss of volume to aligned-spin systems of about $(9.73/9.64)^3 \approx 1.03$ for BBH and $(9.92/9.79)^3 \approx 1.04$ for NSBH. Thus, the precessing search approach proposed here stands to lose at most 4% of signals that are already well covered by aligned-spin templates, but will improve the detection of any system whose fitting factor improves by more than 7% when going from the aligned-spin bank to the precessing-spin bank.

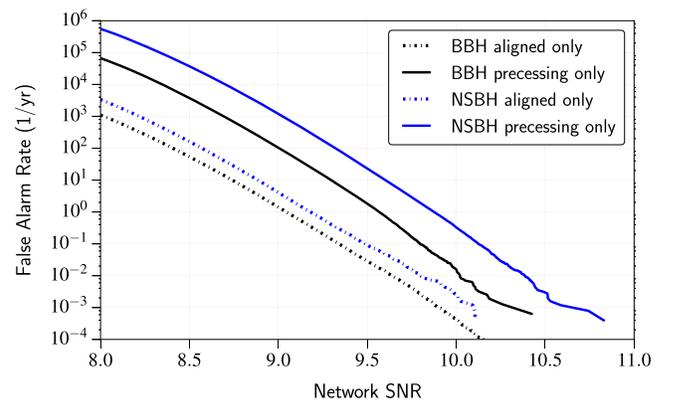


FIG. 4. The rate of coincident background events as a function of the network SNR $(\rho_H^2 + \rho_L^2)^{1/2}$ plotted for the BBH and NSBH aligned-spin-only and precessing-spin-only analyses. The full precessing search background is obtained by combining the aligned-spin-only and precessing-spin-only analyses, as described in the text.

TABLE IV. SNR thresholds at the false-alarm rate values used in this study for our BBH and NSBH aligned-spin-only and precessing-spin-only template banks.

Bank	SNR threshold at false-alarm rate of 10^{-2} yr $^{-1}$	SNR threshold at false-alarm rate of 0.5×10^{-2} yr $^{-1}$
BBH aligned	9.64	9.73
BBH precessing	...	10.10
NSBH aligned	9.79	9.92
NSBH precessing	...	10.44

B. Comparing sensitivity methodology

The signal recovery fraction and effective fitting factor defined in Sec. VB provide a useful measure of the fitting factor averaged across the parameter space. However, these measures do not take into account the background increase as we discussed in the section above. To do this, we replace the σ_i factor in the numerator of Eq. (52)—which we take here to denote the distance at which a simulation will be recovered with a SNR of 1, equivalent to setting $\rho_0 = 1$ in the notation of Sec. VB—with $\sigma_i/\rho_{\text{thresh}}$, where ρ_{thresh} is the threshold taken from Table IV. This gives the distance at which a simulation will be recovered with a SNR of ρ_{thresh} . For the aligned-spin search, the threshold-dependent signal recovery fraction is then written as

$$\alpha_{\text{aligned}} = (\rho_{\text{thresh}})^{-3} \alpha. \quad (53)$$

For the precessing-spin search, which consists of the aligned-spin-only and precessing-spin-only subsearches, the weighting factor depends on the two subsearches according to

$$\alpha_{\text{combined}} = \left(\frac{\sum_{i=1}^N (\text{FF}_{\text{weighted}})_i \sigma_i^3 \mathcal{M}_i^{-5/6}}{\sum_{i=1}^N \sigma_i^3 \mathcal{M}_i^{-5/6}} \right), \quad (54)$$

where we define

$$(\text{FF}_{\text{weighted}})_i = \max_j \{ (\rho_{\text{thresh}})_j^{-3} \text{FF}_{i,j}^3 \}, \quad (55)$$

and the index j runs over the two subsearches. This measure can be used to compute the *relative* sensitivities between search configurations. For the searches we describe above, the relative search sensitivity is computed as

$$\beta = \frac{\alpha_{\text{combined}}}{\alpha_{\text{aligned}}}. \quad (56)$$

Here α_{aligned} is computed using the aligned-spin bank results and the threshold from the aligned-spin bank at a false-alarm rate of 10^{-2} . α_{combined} is computed using both the precessing and aligned-spin results, taking the thresholds for both banks respectively at a false-alarm rate of 0.5×10^{-2} .

C. Results: BBH parameter space

When we include the varying SNR threshold, as discussed above, and evaluate the relative search sensitivity between the aligned-spin template bank and the generic-spin template bank as defined in Eq. (56) we find $\beta = 0.978$. This means that, given our assumed distribution of signals (reweighed by chirp mass), our combined search is on average slightly less sensitive than the aligned-spin search as far as the total number of detections is concerned in our BBH parameter space. However, this precessing search could still allow us to recover signals in specific “highly precessing” regions of parameter space that would not be observed with the aligned-spin search. As a first step in addressing this question, we need to understand which, if any, regions of parameter space are sufficiently precessing that we are gaining sensitivity when using our precessing search method.

Visualizing any quantity in the precessing parameter space is complicated by the large number of dimensions and one has to resort to choosing two dimensional slices and marginalizing over the remaining parameters. One slice that is traditionally shown considers the component masses, or equivalently the total mass and the mass ratio of the system (left panel of Fig. 5). For close to equal mass binaries, where precessional effects are expected to be small, the combined search leads to a loss in sensitivity (always smaller than 3%). Only for mass ratios close to the upper boundary of our parameter space do we obtain some mild sensitivity improvement of up to 7%.

The increase in sensitivity of the precessing search as the mass ratio increases is an expected consequence of the fact that the magnitude of the orbital angular momentum \mathbf{L} decreases as mass ratio increases, thereby allowing larger opening angles between \mathbf{L} and $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$. Larger mass ratios also imply more precessional cycles in the band, which contributes to our precessing search performing better. A similar effect is also obtained by lowering the total mass, which explains why the sensitivity improves to the left of the plot.

A more appropriate set of two variables to identify highly precessing regions has been used in the right panel of Fig. 5: the opening angle between \mathbf{J} and \mathbf{L} introduced above and the inclination with which the system is seen by the observer, defined as the angle between \mathbf{J} and the line of

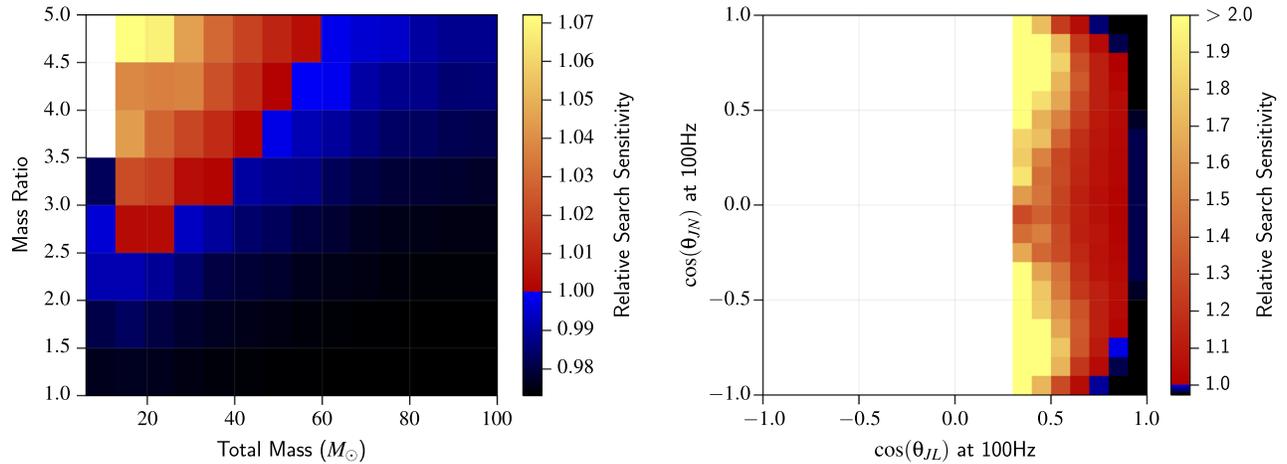


FIG. 5. The relative increase in sensitivity from performing an aligned-spin *and* precessing search compared to performing only the aligned-spin search is plotted in the (total mass, mass ratio) space in the left panel and as a function of opening angle and inclination (defined in the text) in the right panel. For BBH systems with the mass range considered here, the orbital angular momentum \mathbf{L} is always larger than the component angular momenta \mathbf{S}_i and therefore the angle between \mathbf{L} and \mathbf{J} can only take a restricted set of values, as can be seen in the right panel of this figure.

sight \mathbf{N} . The first quantity can be thought of as characterizing the intrinsic amount of precession in the system. The inclination modulates how much precession an observer would see [61]. Of course, both angles actually vary during the coalescence of the binary—the inclination varies on the radiation reaction time scale, whereas the opening angle can have modulations on the precessional time scale on top of the secular evolution on the radiation reaction time scale—and therefore the angles must be evaluated at some reference frequency. Here, we choose to use 100 Hz as it roughly corresponds to the peak sensitivity of the predicted Advanced LIGO noise curve [107]. We find regions where the sensitivity increases by a factor of larger than 2 with respect to the aligned-spin search. However, we note that it

is extremely rare for signals to lie in these regions given the simulation distribution we are using. We illustrate this in Fig. 6 where we show the number of simulations as a function of these spin angles.

The results in this section are dependent on the waveform model used. However, we expect that this dependence is only a weak one, and our results are still valid when using template banks constructed for other waveform approximants. To check this, we repeat our results using the EOBNR approximant, introduced earlier in Sec. II. Unfortunately, at the current time, producing a precessing template bank using the EOBNR approximant is not possible due to the time necessary to generate precessing EOBNR waveforms. However, as described in Ref. [55] we

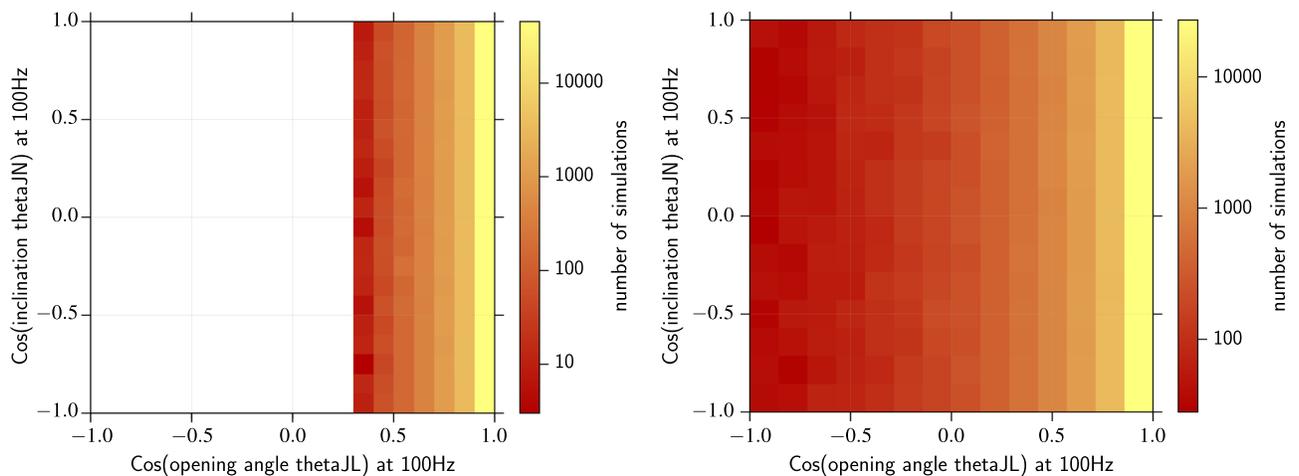


FIG. 6. The number of simulations performed as a function of opening angle and inclination (defined in the text) for our BBH (left) and NSBH (right) parameter spaces. For BBH systems with the mass range considered here, the orbital angular momentum \mathbf{L} is always larger than the component angular momenta \mathbf{S}_i and therefore the angle between \mathbf{L} and \mathbf{J} can only take a restricted set of values, as can be seen in the left panel of this figure.

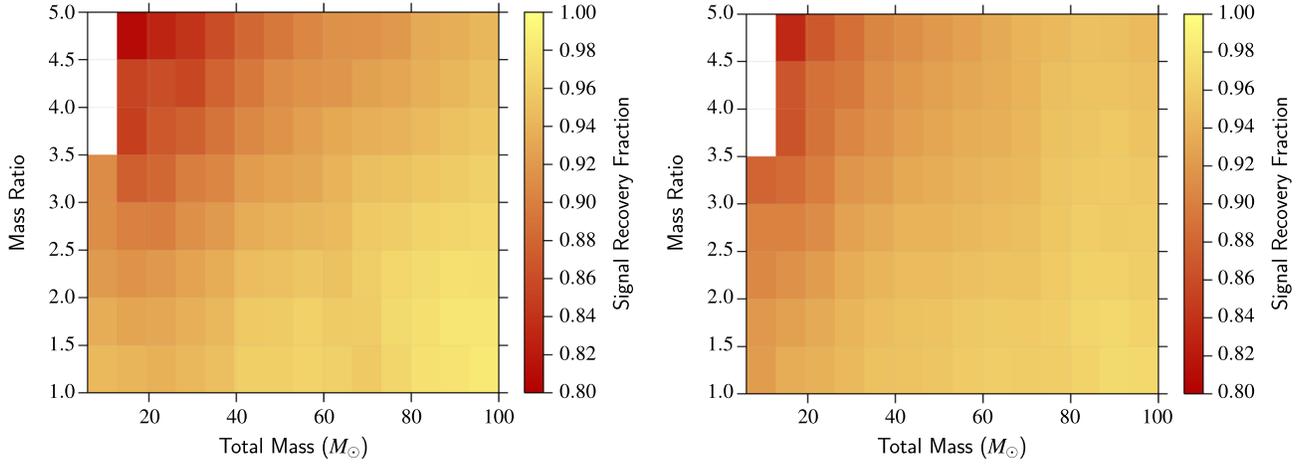


FIG. 7. Signal recovery fraction as a function of total mass and mass ratio comparing different waveform approximants. The left panel shows the signal recovery fraction of our set of precessing IMRPhenom simulations when using our aligned-spin IMRPhenom template bank. The right panel shows the signal recovery fraction of a set of precessing EOBNR simulations, with the same parameter distribution as the IMRPhenom simulations, recovered with an aligned-spin EOBNR template bank constructed as described in Ref. [55].

can generate an aligned-spin EOBNR template bank. In Fig. 7 we compare the ability to recover precessing EOBNR waveforms using an aligned-spin EOBNR template bank with our ability to recover IMRPhenom signals using our aligned-spin IMRPhenom template bank. We can see that the two panels in the figure are largely indistinguishable.

D. Results: NSBH parameter space

We now consider the NSBH parameter space, defined in Table I. When including the varying SNR threshold, using the values in Table IV and Eq. (56), we find the relative search sensitivity, evaluated at constant false-alarm rate, is $\beta = 1.014$. This means that we expect to recover 1.4% more signals averaged across the NSBH parameter space

when using the precessing search compared to the aligned-spin search. This number does, of course, depend on the distribution of parameters that we chose—in Table I—for our NSBH space.

In Fig. 8, we show the relative search sensitivity as a function of the mass ratio and total mass (left) and also the angles between \mathbf{J} and \mathbf{N} and between \mathbf{J} and \mathbf{L} (right). The relative search sensitivity as a function of total mass and mass ratio (left panel of Fig. 8) shows similar trends to the corresponding ones for BBH, but given mass ratio values up to 15, we see larger relative search sensitivities, up to 1.14 at a mass ratio of 15. On the right panels of Figs. 7 and 8 we notice that systems in our NSBH parameter space are able to cover all values of the angle between \mathbf{J} and \mathbf{L} , which was not the case for our BBH

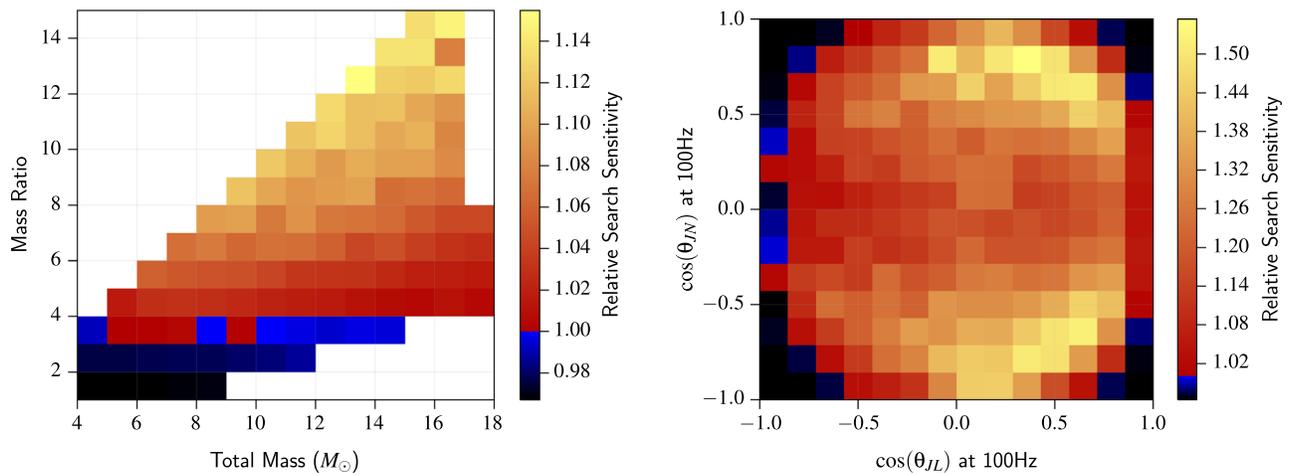


FIG. 8. The relative increase in sensitivity from performing an aligned-spin *and* precessing search compared to performing only the aligned-spin search is plotted in the (total mass, mass ratio) space in the left panel and as a function of opening angle and inclination (defined in the text) in the right panel.

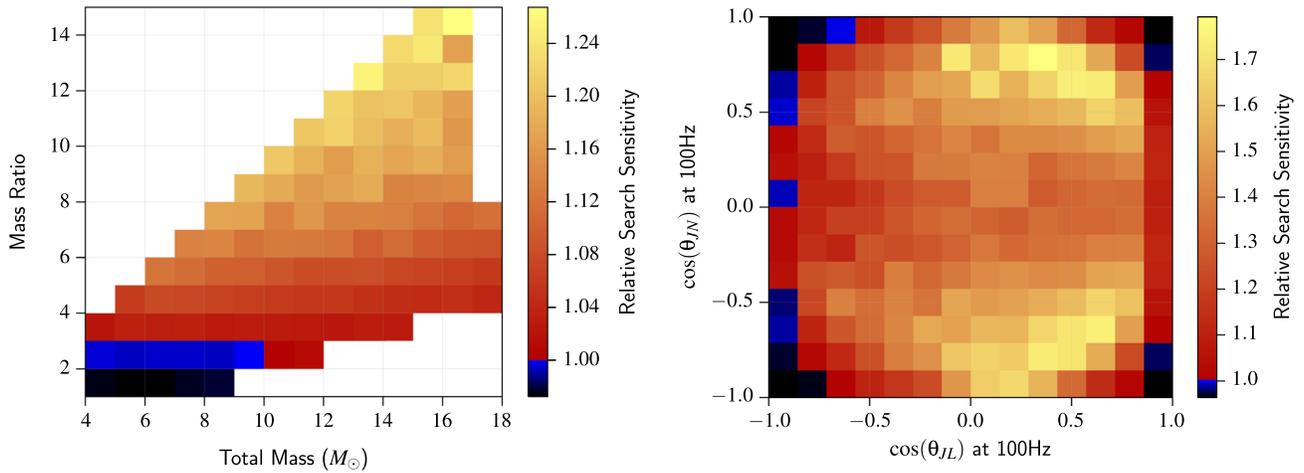


FIG. 9. This figure is equivalent to Fig. 8 except here we assume that for the precessing search *all* fitting factors are unity. This provides us with an upper limit on the sensitivity increase that could be gained if the template bank were placed much more densely. This is an upper limit because the increased background from the hypothetical denser precessing bank is not taken into account.

systems. This is due to the fact that the higher-mass ratios available to NSBH systems allow for cases where the black-hole spin angular momentum is larger than the orbital angular momentum. For the values of this angle that the BBH parameter space *can* produce, we see similar behavior between the BBH and NSBH parameter spaces. The main difference is that it is more likely for signals to have larger values of the angle between \mathbf{J} and \mathbf{L} , and therefore be more likely to show precessional effects, in the NSBH parameter space than in the BBH parameter space.

We can also ask if the fitting factor that we chose to place the precessing bank with is affecting our results, and if results would improve if the template bank were denser. To try to answer this we reproduce Fig. 8, except when constructing it we set all fitting factors for the precessing bank to unity. We use the same background increase as given in Table IV, which will be an underestimate, but this can provide an upper limit on the relative search sensitivity that would be obtained by using a denser template bank. The results of this can be seen in Fig. 9. Here the features are qualitatively similar to those in Fig. 8, but the relative sensitivity improves.

VII. DISCUSSION

In this work we have (i) derived a new method for detecting compact-binary coalescences when using waveform filters with generically oriented spins, (ii) constructed banks of generic-spin template waveforms using this method, (iii) demonstrated the method on a stretch of Gaussian noise and (iv) assessed the relative search sensitivity between our new generic-spin search and the current aligned-spin search. In general, averaged over the parameter spaces we consider, we have found that using our generic-spin search does not result in a net increase in the number of detections of compact-binary mergers—in our

NSBH parameter space we saw an increase of only 1%. However, we have demonstrated that in regions of parameter space where precessional effects are large, we can see improvements in detection rate that are larger than 50%. Systems where the precessional effects are strong may be rare, but these are also systems which offer us a better chance to disentangle the various parameters that describe the source, which in turn could allow us a better chance to understand the nature and origin of these systems [37,38]. Therefore, one might argue that an observation of a highly precessing system might be worth more than the one, or several, observations of systems that do not exhibit precessional effects.

We have also demonstrated that when considering systems with GW150914-like masses and generically oriented spins, we find signal recovery fractions that are larger than 0.95. This is consistent with what is expected due to the minimal match of 0.97 that is used to set the discreteness of the template bank in the aligned-spin parameter space. Current Advanced LIGO searches are therefore not missing systems with masses similar to GW150914 because the waveform filters do not consider misaligned spins.

It is foreseeable that in future work alternative precessing search methods might be proposed that could improve on the formulation we provide here. However, there is a fundamental difficulty we have observed in this work that leads us to believe that it will not be possible to significantly improve the relative sensitivity of such a hypothetical search over the one described here. Specifically, we notice that the number of templates needed to adequately cover a precessing parameter space is at least an order of magnitude larger than that required to cover the aligned-spin parameter space. This happens even though we are using a considerably smaller fitting factor for our precessing template bank than the

aligned-spin one. From this, one concludes that the *size* of the precessing parameter space is significantly larger than the region of that parameter space that is covered by a template bank placed in only the aligned-spin manifold. However, the majority of our signal waveforms, which assume an isotropic distribution of the spin directions, are recovered well by aligned-spin template waveforms. This implies that the density of astrophysical systems that lie close to the aligned-spin region of parameter space is also significantly larger than the density of astrophysical systems in the remaining, large, region of parameter space where precession is important. For these reasons, it is difficult to gain a significant increase in the number of signals observed when covering the significantly larger, sparsely populated, precessing parameter space unless the distribution of signals in the Universe strongly prefers highly precessing cases.

In this work, we have not discussed signal-based consistency tests. Real GW data are not Gaussian, and non-Gaussianities often produce a high SNR in matched-filter outputs. This will be equally true for aligned-spin as for precessing template banks. To be able to separate real events from non-Gaussianities with large SNRs we use a set of signal-based consistency tests [40,112]. Ranking statistics are then constructed that combine the SNR with the signal-based consistency tests such that non-Gaussianities are down-weighted if they do not match the features we expect of real signals [40]. To be able to apply our methods on real data, we need to extend these signal-based consistency tests to our precessing search. This will be our focus in future work.

It is worth pointing out that the signal-based consistency tests down-weight both non-Gaussianities *and* real signals that have a low overlap with the best matching template waveform [112]. It is therefore possible that our results will improve slightly when signal-based consistency tests are included. However, there is also the possibility that the noise background will increase for precessing templates relative to aligned-spin templates when applied in real noise even with signal-based consistency tests. This would decrease the improvement of using the precessing template bank.

Our results are also dependent on how well the models we are using match the waveforms that will be produced by real compact-binary mergers. For our BBH parameter space we are using a waveform model including inspiral, merger and ringdown phases and which has been tuned against NR simulations [36,75]. For our NSBH space, we are using an inspiral-only model, which is known to have some disagreement with other inspiral-only models, with inspiral-merger ringdown models and with NR [64,113]. In cases where there is a mismatch between our models and the real signal the signal recovery fractions will be lower than what we have calculated here, but this will be true for both aligned-spin and

precessing template waveforms. Nevertheless the methods presented here are equally applicable as waveform approximants improve.

For both our BBH and NSBH waveform models we have restricted them to only consider the dominant $l = 2$, $|m| = 2$ modes of the GW in the coprecessing frame. The effect of using higher-order modes for nonprecessing waveforms in searches has been previously explored [114], but no computationally feasible search method has been proposed. It would be possible to extend the method described here by including also the initial phase of the binaries in the orbital plane as a discrete template bank parameter. This would alter our sky-maxed SNR, but would enable us to filter with precessing, higher-order mode waveforms. Exploring how computationally costly this would be and whether this method would be feasible at all is a topic we leave for future work.

For both the BBH and NSBH parameter spaces we consider here we find an order of magnitude increase between the number of templates in the aligned-spin bank and the number of templates in the precessing bank. Coupled with the fact that filtering precessing templates using our scheme is a factor of 2 more expensive than filtering aligned-spin templates results in searches using our precessing banks being a factor of 20 more computationally expensive than searches using aligned-spin template banks. This additional computational cost could be used to, for example, increase the fitting factor of the aligned-spin bank and potentially also gain some increase in search sensitivity [115]. We would again make the point that the observation of a highly precessing signal could be very astrophysically rewarding, but exploring how to reduce this computational cost would be very beneficial. One possible approach we wish to consider in the future is, when building the template bank, to not include templates corresponding to systems with a low intrinsic luminosity. This might allow us to reduce the template count *and* the noise background. Other possibilities include some form of hierarchical approach to filtering the templates in the template bank. One could also consider schemes to make the bank construction process more efficient, such as by parallelizing not only in chirp mass bins, but also by using bins in mass ratio [116].

ACKNOWLEDGMENTS

The authors would like to thank Alex Nitz and Lijing Shao for useful and detailed comments, and Stas Babak, Sebastian Khan, Harald Pfeiffer, Michael Pürrer, Vivien Raymond, B. Sathyaprakash, Patricia Schmidt and Andrea Taracchini for helpful discussions. The authors would like to thank the anonymous referees for useful comments and feedback on this manuscript. The authors would also like to thank the Max Planck Gesellschaft for support.

APPENDIX: COMPUTATIONAL OPTIMIZATIONS OF PRECESSING BANK CONSTRUCTION

To generate our stochastic template banks of precessing signals, we use a number of recent optimizations described in Refs. [55,57,104], as well as some new methods. We briefly describe these optimizations in the next paragraphs.

When choosing h_{prop} the masses are chosen from a uniform distribution in the chirp time coordinates τ_0 - τ_3 [84]. These coordinates are optimal for nonspinning, inspiral-only signals. They are suboptimal in our case, but better than choosing mass parameters uniformly in m_1 - m_2 , since typically many more templates are needed at low mass compared to high mass. Spins are chosen isotropically with a uniform distribution of spin magnitude. Sky locations and orientation parameters are chosen isotropically. These distributions may not be the optimal choices for stochastic precessing template bank construction, but we leave a further exploration of this for future work.

When computing matches, those templates in the template bank that have values of chirp mass closest to the proposed template have matches computed first. This allows one to reject proposed templates more quickly. In addition a match is only computed if the difference in chirp mass between the template in the bank and the point being tested is less than 25%. This can lead to some over-coverage, especially at higher masses.

When computing matches, a match is first computed using a frequency spacing in the overlap integral of $df = 8$ Hz. If the resulting overlap is 0.05 less than the desired minimal match, the match at smaller values of the frequency spacing is assumed to also be less than the desired minimal match and is not computed explicitly. If the match using a frequency spacing of 8 Hz is not less than

the minimal match minus 0.05 then the match is computed using frequency spacings of $df = 0.5$ Hz and $df = 0.25$ Hz. If these matches agree to within 0.1% then the $df = 0.25$ Hz value is used. Otherwise an additional match at $df = 0.125$ Hz is computed and compared against $df = 0.25$ Hz. This process continues until either the match is in agreement to 0.1% or the value of df has exceeded the inverse of the waveform length. The first test using $df = 8$ Hz quickly rejects points that do not match well to each other. The subsequent tests at much smaller values of df , 0.5 Hz and 0.25 Hz, are because precessional features in the waveform can be missed when using larger values of df and we found some cases where matches would agree well at $df = 4$ Hz and $df = 2$ Hz, but then diverge when df became smaller.

A coarse bank is first produced with a lower convergence threshold to roughly map out the density of templates. Here, we choose a coarse threshold such that the algorithm terminates when fewer than one in X proposed templates are accepted to the bank over the last ten iterations. The space is then split into Y equally spaced bands in chirp mass, and a process launched for each chirp mass band, choosing points only in that band, and filling templates until fewer than one in Z proposed templates are accepted to the bank over the last ten iterations. It is possible that the parallel processes will cover the same region, but this effect is mitigated by beginning the process with a partially complete template bank over the full space before parallelizing. The parallelization also helps to ensure that all parts of the parameter space are being sampled by the stochastic process. The values X , Y and Z are chosen empirically and vary for the different banks generated here.

Further optimizations of this process will be greatly desirable if using the methods described in this paper in the future. We leave that to future work.

-
- [1] G. M. Harry (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084006 (2010).
 - [2] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
 - [3] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [4] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
 - [5] K. Somiya (KAGRA Collaboration), *Classical Quantum Gravity* **29**, 124007 (2012).
 - [6] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (KAGRA Collaboration), *Phys. Rev. D* **88**, 043007 (2013).
 - [7] B. Iyer *et al.*, <https://dcc.ligo.org/LIGO-M1100296/public>.
 - [8] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), [arXiv:1602.03842](https://arxiv.org/abs/1602.03842).
 - [9] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Classical Quantum Gravity* **27**, 173001 (2010).
 - [10] K. Belczynski, M. Dominik, T. Bulik, R. O’Shaughnessy, C. L. Fryer, and D. E. Holz, *Astrophys. J.* **715**, L138 (2010).
 - [11] Michal Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *Astrophys. J.* **759**, 52 (2012).
 - [12] T. Kinugawa, K. Inayoshi, K. Hotokezaka, D. Nakauchi, and T. Nakamura, *Mon. Not. R. Astron. Soc.* **442**, 2963 (2014).

- [13] R. D. Ferdman, I. H. Stairs, M. Kramer, G. H. Janssen, C. G. Bassa *et al.*, *Mon. Not. R. Astron. Soc.* **443**, 2183 (2014).
- [14] P. Marchant, N. Langer, P. Podsiadlowski, T. Tauris, and T. Moriya, *Astron. Astrophys.* **588**, A50 (2016).
- [15] C. L. Rodriguez, M. Morscher, B. Pattabiraman, S. Chatterjee, C.-J. Haster, and F. A. Rasio, *Phys. Rev. Lett.* **115**, 051101 (2015); **116**, 029901(E) (2016).
- [16] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. Coleman Miller, and M. Walczak, *Astrophys. J.* **789**, 120 (2014).
- [17] D. Clausen, S. Sigurdsson, and D. F. Chernoff, *Mon. Not. R. Astron. Soc.* **428**, 3618 (2013).
- [18] N. Ivanova, C. Heinke, F. Rasio, K. Belczynski, and J. Fregeau, *Mon. Not. R. Astron. Soc.* **386**, 553 (2008).
- [19] D. Pooley, W. Lewin, S. Anderson, H. Baumgardt, A. Filippenko *et al.*, *Astrophys. J.* **591**, L131 (2003).
- [20] K. Belczynski, S. Repetto, D. E. Holz, R. O’Shaughnessy, T. Bulik, E. Berti, C. Fryer, and M. Dominik, *Astrophys. J.* **819**, 108 (2016).
- [21] S. Vitale, R. Lynch, P. Graff, and R. Sturani, *arXiv*: 1503.04307.
- [22] S. Stevenson, F. Ohme, and S. Fairhurst, *Astrophys. J.* **810**, 58 (2015).
- [23] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), *Astrophys. J.* **818**, L22 (2016).
- [24] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), *arXiv*:1602.03840.
- [25] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [26] K. Belczynski, R. E. Taam, E. Rantsiou, and M. van der Sluys, *Astrophys. J.* **682**, 474 (2008).
- [27] J. F. Steiner and J. E. McClintock, *Astrophys. J.* **745**, 136 (2012).
- [28] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski, *Astrophys. J.* **719**, L79 (2010).
- [29] L. Gou, J. E. McClintock, M. J. Reid, J. A. Orosz, J. F. Steiner, R. Narayan, J. Xiang, R. A. Remillard, K. A. Arnaud, and S. W. Davis, *Astrophys. J.* **742**, 85 (2011).
- [30] L. Gou, J. E. McClintock, R. A. Remillard, J. F. Steiner, M. J. Reid, J. A. Orosz, R. Narayan, M. Hanke, and J. García, *Astrophys. J.* **790**, 29 (2014).
- [31] A. Fabian, D. Wilkins, J. Miller, R. Reis, C. Reynolds *et al.*, *Mon. Not. R. Astron. Soc.* **424**, 217 (2012).
- [32] M. A. Nowak, J. Wilms, K. Pottschmidt, N. Schulz, D. Maitra, and J. Miller, *Astrophys. J.* **744**, 107 (2012).
- [33] C. S. Reynolds, L. W. Brenneman, A. M. Lohfink, M. L. Trippe, J. M. Miller, R. C. Reis, M. A. Nowak, and A. C. Fabian, *AIP Conf. Proc.* **1427**, 157 (2012).
- [34] L. Brenneman, C. Reynolds, M. Nowak, R. Reis, M. Trippe *et al.*, *Astrophys. J.* **736**, 103 (2011).
- [35] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [36] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [37] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, *Phys. Rev. Lett.* **112**, 251101 (2014).
- [38] R. O’Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, V. Raymond, C. Kim, and C.-H. Lee, *Phys. Rev. D* **89**, 102005 (2014).
- [39] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [40] S. Babak, R. Biswas, P. Brady, D. Brown, K. Cannon *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [41] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 122001 (2004).
- [42] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **73**, 062001 (2006).
- [43] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **77**, 062002 (2008).
- [44] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **79**, 122001 (2009).
- [45] J. Abadie *et al.* (Virgo, LIGO Scientific Collaborations), *Phys. Rev. D* **83**, 122005 (2011); **86**, 069903(E) (2012).
- [46] J. Abadie *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. D* **85**, 082002 (2012).
- [47] J. Aasi *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. D* **87**, 022002 (2013).
- [48] J. Abadie *et al.* (LIGO Scientific Collaboration), *Astrophys. J.* **760**, 12 (2012).
- [49] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **78**, 042002 (2008).
- [50] A. Buonanno, Y.-b. Chen, Y. Pan, and M. Vallisneri, *Phys. Rev. D* **70**, 104003 (2004).
- [51] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, B. S. Sathyaprakash, H. Tagoshi, and H. Takahashi, *Phys. Rev. D* **80**, 024009 (2009).
- [52] S. Privitera, S. R. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, *Phys. Rev. D* **89**, 024003 (2014).
- [53] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
- [54] T. Dal Canton, A. P. Lundgren, and A. B. Nielsen, *Phys. Rev. D* **91**, 062010 (2015).
- [55] C. Capano, I. Harry, S. Privitera, and A. Buonanno, *Phys. Rev. D* **93**, 124007 (2016).
- [56] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), *Phys. Rev. D* **93**, 122003 (2016).
- [57] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
- [58] D. A. Brown, A. Lundgren, and R. O’Shaughnessy, *Phys. Rev. D* **86**, 064020 (2012).
- [59] D. R. Lorimer, *Living Rev. Relativ.* **11**, 8 (2008).
- [60] D. Chakrabarty, *AIP Conf. Proc.* **1068**, 67 (2008).
- [61] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev. D* **86**, 084017 (2012).
- [62] L. Blanchet, *Living Rev. Relativ.* **17**, 2 (2014).
- [63] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [64] A. H. Nitz, A. Lundgren, D. A. Brown, E. Ochsner, D. Keppel, and I. W. Harry, *Phys. Rev. D* **88**, 124039 (2013).
- [65] A. Lundgren and R. O’Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014).
- [66] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [67] A. Buonanno and T. Damour, *Phys. Rev. D* **62**, 064015 (2000).
- [68] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).

- [69] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **86**, 024011 (2012).
- [70] A. Taracchini *et al.*, *Phys. Rev. D* **89**, 061502 (2014).
- [71] A. Nagar, T. Damour, C. Reisswig, and D. Pollney, *Phys. Rev. D* **93**, 044046 (2016).
- [72] T. Damour and A. Nagar, *Phys. Rev. D* **90**, 044018 (2014).
- [73] L. Santamaria, F. Ohme, P. Ajith, B. Bruegmann, N. Dorband *et al.*, *Phys. Rev. D* **82**, 064016 (2010).
- [74] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Jiménez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [75] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016).
- [76] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **89**, 084006 (2014).
- [77] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, *Phys. Rev. X* **4**, 031006 (2014).
- [78] M. Pürrer, *Classical Quantum Gravity* **31**, 195010 (2014).
- [79] M. Pürrer, *Phys. Rev. D* **93**, 064041 (2016).
- [80] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015).
- [81] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, [arXiv:1604.08253](https://arxiv.org/abs/1604.08253).
- [82] LIGO Scientific Collaboration, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.
- [83] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987), pp. 330–458.
- [84] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [85] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
- [86] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
- [87] A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003); **74**, 029904(E) (2006).
- [88] Y. Pan, A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **69**, 104017 (2004); **74**, 029905(E) (2006).
- [89] D. Fazi, Ph.D. thesis, Università di Bologna, 2009.
- [90] I. W. Harry and S. Fairhurst, *Phys. Rev. D* **83**, 084002 (2011).
- [91] S. Bose, A. Pai, and S. V. Dhurandhar, *Int. J. Mod. Phys. D* **09**, 325 (2000).
- [92] L. S. Finn, *Phys. Rev. D* **63**, 102001 (2001).
- [93] A. Pai, S. Dhurandhar, and S. Bose, *Phys. Rev. D* **64**, 042004 (2001).
- [94] D. M. Macleod, I. W. Harry, and S. Fairhurst, *Phys. Rev. D* **93**, 064004 (2016).
- [95] E. Poisson and C. M. Will, *Phys. Rev. D* **52**, 848 (1995).
- [96] R. Balasubramanian, B. Sathyaprakash, and S. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996).
- [97] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [98] B. J. Owen and B. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [99] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. Sathyaprakash, *Classical Quantum Gravity* **23**, 5477 (2006).
- [100] T. Cokelaer, *Phys. Rev. D* **76**, 102004 (2007).
- [101] S. Babak, *Classical Quantum Gravity* **25**, 195011 (2008).
- [102] I. W. Harry, B. Allen, and B. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [103] G. M. Manca and M. Vallisneri, *Phys. Rev. D* **81**, 024004 (2010).
- [104] H. Fehrmann and H. J. Pletsch, *Phys. Rev. D* **90**, 124049 (2014).
- [105] F. Ozel, D. Psaltis, R. Narayan, and J. E. McClintock, *Astrophys. J.* **725**, 1918 (2010).
- [106] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, *Astrophys. J.* **741**, 103 (2011).
- [107] J. Aasi *et al.* (Virgo, LIGO Scientific Collaborations), *Living Rev. Relativ.* **19**, 1 (2016).
- [108] T. Dal Canton *et al.*, *Phys. Rev. D* **90**, 082004 (2014).
- [109] S. A. Usman *et al.*, [arXiv:1508.02357](https://arxiv.org/abs/1508.02357).
- [110] R. Prix and B. Krishnan, *Classical Quantum Gravity* **26**, 204013 (2009).
- [111] T. Dent and J. Veitch, *Phys. Rev. D* **89**, 062002 (2014).
- [112] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
- [113] P. Kumar, K. Barkett, S. Bhagwat, N. Afshari, D. A. Brown, G. Lovelace, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **92**, 102001 (2015).
- [114] C. Capano, Y. Pan, and A. Buonanno, *Phys. Rev. D* **89**, 102003 (2014).
- [115] D. Keppel, *Phys. Rev. D* **87**, 124003 (2013).
- [116] N. Indik *et al.* (to be published).