Stochastic template bank for gravitational wave searches for precessing neutron-star–black-hole coalescence events

Nathaniel Indik, K. Haris, Tito Dal Canton, Henning Fehrmann, Badri Krishnan, Andrew Lundgren, Alex B. Nielsen, and Archana Pai

Phys. Rev. D **95**, 064056 — Published 30 March 2017

DOI: 10.1103/PhysRevD.95.064056

A stochastic template bank for gravitational wave searches for precessing neutron star - black hole coalescence events

Nathaniel Indik,[1,2] K Haris,[3] Tito Dal Canton,[1,4,2] Henning Fehrmann,[1,2]
Badri Krishnan,[1,2] Andrew Lundgren,[1,2] Alex B Nielsen,[1,2] and Archana Pai[3]

[1] *Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Callinstr. 38, D-30167 Hannover, Germany*
[2] *Leibniz Universität Hannover, Welfengarten 1-A, D-30167 Hannover, Germany*
[3] *Indian Institute for Science Education and Research Thiruvananthapuram, CET Campus, Trivandrum 695016, India*
[4] *NASA Postdoctoral Program Fellow, Goddard Space Flight Center, Greenbelt, MD 20771, USA*

Gravitational wave searches to date have largely focused on non-precessing systems. Including precession effects greatly increases the number of templates to be searched over. This leads to a corresponding increase in the computational cost and can increase the false alarm rate of a realistic search. On the other hand, there might be astrophysical systems that are entirely missed by non-precessing searches. In this paper we consider the problem of constructing a template bank using stochastic methods for neutron star - black hole binaries allowing for precession, but with the restrictions that the total angular momentum of the binary is pointing towards the detector and that the neutron star spin is negligible relative to that of the black hole. We quantify the number of templates required for the search, and we explicitly construct the template bank. We show that despite the large number of templates, stochastic methods can be adapted to solve the problem. We quantify the parameter space region over which the non-precessing search might miss signals.

I. INTRODUCTION

Binary systems consisting of neutron stars and black holes are key targets for the present generation of gravitational wave detectors such as Advanced LIGO [1] and Advanced Virgo [2]. The LIGO detectors have to date observed two such events with high significance, labeled GW150914 [3] and GW151226 [4] and a third LVT151012 with lower significance [5]. All three of these are binary black hole coalescence events. The searches for these events *a priori* cover a wide range of masses and spins magnitudes, but use only waveforms for which the spins of the individual compact objects are assumed to be completely aligned or completely anti-aligned with the orbital angular momentum [6]. Misalignments between the spin and orbital angular momentum generally causes precession of the orbital plane and additional modulations of the gravitational waveforms [7]. While follow-up studies for accurate parameter estimation do include precession [8], including these waveforms directly in the initial search pipelines is challenging; the dimensionality of the parameter space to be searched is increased, implying a significant increase in the total number of templates. This is challenging not only for computational reasons, but also because a larger number of independent templates leads to a larger probability for false alarms. Nevertheless, if a significant number of NSBH systems in our universe display precessional modulations that cannot be accurately recovered by spin-aligned templates, the search pipeline could potentially detect more events if precession effects were to be included [9].

For searches based on matched filtering with modeled waveforms, the traditional method of constructing a template bank was to use the parameter space metric [10, 11] for determining the spacing between adjacent templates. This method has been successfully used to search for non- spinning systems [12] and has also been applied to aligned- spin systems [13, 14]. For precessing waveforms however, the parameter space metric is not yet sufficiently well understood for it to be directly used to place templates. The main issue is that to place a lattice of templates, one needs a coordinate system on the parameter space where the metric is explicitly flat. It is not clear whether such a coordinate system exists (even in any approximate sense) for the space of precessing waveforms. In situations where such geometric template placement methods are not available, stochastic methods are commonly employed [15, 16]. This also includes the most recent searches over the first Advanced LIGO observing run [5, 6].

The basic idea of stochastic methods is to place templates at random points in the parameter space and to remove templates which happen to lie very close to other templates. These stochastic methods are generally applicable but they are typically less efficient than the geometric methods, i.e. they require more templates than a geometric bank to achieve the same coverage over the same parameter space (however, stochastic methods become more efficient in higher dimensions and can be competitive with geometric methods [17]). Moreover, the stochastic template placement procedure can be computationally demanding for large parameter spaces, which is in fact the case for precessing waveforms.

In this paper we shall meet this computational challenge and show how stochastic methods can be applied to cover the space of precessing waveforms. The main computational problem we face is that for every proposed template, one typically compares it with *all* previously accepted templates to decide whether or not it should be accepted. We shall see that with an appropriate choice of coordinates, it is possible to break up the parameter space into smaller regions, and treat each region independently. This paper presents the largest template bank

constructed thus far with stochastic methods for binary coalescence searches.

Specifically, we shall focus on neutron-star – black-hole (NSBH) systems, but we expect that our method would apply to other source systems as well. We shall consider NSBH binaries with a black hole of mass M_1 and a neutron star mass of M_2 such that $2M_\odot < M_1 < 16M_\odot$, and $1M_\odot < M_2 < 3M_\odot$. Since neutron star spins are expected to be small we shall ignore them, but the black hole spin will be allowed to take any magnitude which is meaningful in the Kerr metric and any direction [18]. We shall use the frequency domain waveform model presented in [19]. This waveform model does not, so far, contain the merger and ringdown portions. For the parameter space above, and for the expected sensitive frequency range of the Advanced LIGO and Virgo detectors, the inspiral portion of the waveform will have the largest contribution to the signal-to-noise ratio. Thus, the merger and ringdown phases will not be important for our purposes. However, the methods used in our study should be useful also for higher mass systems where merger effects are more important.

The most effectual implementation of a stochastic NSBH template bank constructed to date [20] required approximately 1.6 million, assuming the detector to be in the "early Advanced LIGO" configuration [21]. This construction used a new detection statistic based on maximizing the signal-to-noise ratio (SNR) over sky-positions and required a minimal match criteria of 90% when comparing each proposed template with previously accepted ones, as opposed to the more conventional 97%. Using the conventional 97% value would lead to a much larger number of templates. Moreover, as the detector improves its low frequency sensitivity over the next few years, the number of templates increases further. The method used in this paper could be used to deal with both of the above issues. We shall use the conventional 97% minimal match value and, for simplicity, we use the conventional SNR rather than the detection statistic introduced in [20], but we expect that our method can be adapted to that detection statistic as well.

The plan for the rest of the paper is as follows. Sec. II briefly sets up notation and the parameters describing a precessing binary system and the gravitational waveform, and outlines the stochastic template placement algorithm. Sec. III describes the stochastic template bank. Sec. IV compares this precessing template bank with the aligned spin bank and studies how well it recovers injected signals.

II. BACKGROUND

A. Precessing binaries

Consider an NSBH system consisting of a black hole with mass M_1, spin $\mathbf{S}$, and a neutron star of mass M_2 and zero spin. Let $\hat{\mathbf{N}}(\theta,\phi)$ be the unit vector along the line-of-sight from the detector to the binary system, and let $\mathbf{L}$ be the orbital angular momentum of the binary. Define the dimensionless spin of the black hole as $\chi = |\mathbf{S}|/M_1^2$. The component of $\mathbf{S}$ along $\mathbf{L}$ will be determined by the quantity $\kappa = \hat{\mathbf{S}}\cdot\hat{\mathbf{L}}$, and the component of $\mathbf{S}$ orthogonal to $\mathbf{L}$ is

$$\mathbf{S}^\perp = \mathbf{S} - (\mathbf{S}\cdot\hat{\mathbf{L}})\hat{\mathbf{L}}\,. \tag{1}$$

It can be shown that the direction of the total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$ is approximately conserved [7], and that $\hat{\mathbf{L}}$ and $\hat{\mathbf{S}}$ precess around $\mathbf{J}$. The magnitude of $\mathbf{L}$ decreases steadily because of the emission of gravitational radiation but the magnitude of $\mathbf{S}$ remains constant as does the angle between $\mathbf{L}$ and $\mathbf{S}$. The opening angle β of the precession cone is given by

$$\cos\beta \;= \hat{\mathbf{J}}\cdot\hat{\mathbf{L}}\,. \tag{2}$$

As the magnitude of $\mathbf{L}$ decreases, β should increase in order to maintain the direction of $\mathbf{J}$ and the angle between $\mathbf{L}$ and $\mathbf{S}$ [7]. However, the precession time-scale is smaller than the radiation reaction time scale (which determines the rate at which L decreases). It can be shown [22] that for the advanced LIGO and Virgo detectors, to a reasonable approximation, $\mathbf{L}$ and $\mathbf{S}$ precess steadily around $\mathbf{J}$ with a constant opening angle β. The rare case of transitional precession occurs when $\mathbf{J}\sim 0$ at some point during the evolution of the binary system. Finally, α_0 is an azimuthal angle that expresses the orientation of $\hat{\mathbf{L}}$ relative to $\hat{\mathbf{J}}$ in the inertial detector frame and we shall define the angle θ_J as $\cos\theta_J = \hat{\mathbf{J}}\cdot\hat{\mathbf{N}}$.

For a plane gravitational wave traveling in a direction $\hat{\mathbf{z}}$, and a frame $(\hat{\mathbf{x}},\hat{\mathbf{y}})$ in the plane orthogonal to $\hat{\mathbf{z}}$, we define the tensors

$$\mathbf{e}^+_{ab} = \hat{\mathbf{x}}_a\hat{\mathbf{x}}_b - \hat{\mathbf{y}}_a\hat{\mathbf{y}}_b\,, \quad \mathbf{e}^\times_{ab} = \hat{\mathbf{x}}_a\hat{\mathbf{y}}_b + \hat{\mathbf{y}}_a\hat{\mathbf{x}}_b\,. \tag{3}$$

The gravitational wave can be written as a sum of two transverse polarizations

$$h_{ab}(t) = h_+(t)\mathbf{e}^+_{ab} + h_\times(t)\mathbf{e}^\times_{ab}\,. \tag{4}$$

It is always possible to find a frame $(\hat{\mathbf{x}},\hat{\mathbf{y}})$ such that

$$h_+(t) = A_+(t)\cos 2\Phi(t)\,, \qquad h_\times(t) = A_\times(t)\sin 2\Phi(t)\,, \tag{5}$$

where $A_{+,\times}$ are slowly varying amplitudes and $\Phi(t)$ is a rapidly varying phase. For the case of a binary system, the wave-frame $(\hat{\mathbf{x}},\hat{\mathbf{y}})$ is tied to the direction of the orbital angular momentum, and $\hat{\mathbf{x}}$ is taken to be $\pm\hat{\mathbf{N}}\times\hat{\mathbf{L}}$. The direction of $\mathbf{x}$ in the detector frame defines a polarization angle ψ and, following [7], we choose the convention:

$$\psi(t) = \tan^{-1}\left(\frac{\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{z}} - (\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}})(\hat{\mathbf{z}}\cdot\hat{\mathbf{N}})}{\hat{\mathbf{N}}\cdot(\hat{\mathbf{L}}(t)\times\hat{\mathbf{z}})}\right). \tag{6}$$

Note that because of precession, the direction of $\mathbf{L}$ changes in time and thus ψ also changes with time. With

these conventions, the expressions for $h_{+,\times}$ are:

$$h_+(t) = -\frac{2\pi M}{rD}\left[1 + (\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}})^2\right]\cos 2\Phi(t)\,, \quad (7)$$

$$h_\times(t) = -\frac{2\pi M}{rD}\left[-2\ \hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}}\right]\sin 2\Phi(t)\,, \quad (8)$$

where D is the distance to the binary system, r is the binary orbital diameter and $M = M_1 + M_2$ is the total mass. Throughout the paper we shall also use the *chirp mass*,

$$\mathcal{M}_{\rm C} = \eta^{\frac{3}{5}} M\,, \quad (9)$$

and the following quantities related to the mass ratio of the components,

$$\nu = \frac{M_1}{M} \quad \eta = \frac{M_1 M_2}{M^2} \quad q = \frac{M_1}{M_2}\,. \quad (10)$$

It is also convenient to express the black hole spin via the dimensionless vector $\chi := \mathbf{S}/M_1^2$ and decompose it into components parallel and perpendicular to $\widehat{\mathbf{L}}$, $\chi^{\parallel}$ and $\chi^{\perp}$ respectively. The total dimensionless spin magnitude is thus $\chi = \sqrt{(\chi^{\parallel})^2 + (\chi^{\perp})^2}$.

The detector response functions to these polarizations are denoted by $F_+(\mathbf{N},\psi)$ and $F_\times(\mathbf{N},\psi)$. If the signal is parallel to one of these configurations, it is said to be linearly polarized. In contrast, a signal that can be decomposed into an equal linear combination of these two principal directions is circularly polarized. In general, the signal seen by the detector $h(t)$ will be a linear combination of the two polarizations:

$$\begin{aligned} h(t) &= h_+(t)F_+(\hat{\mathbf{N}},\psi(t)) + h_\times(t)F_\times(\hat{\mathbf{N}},\psi(t)) \\ &= A(t)\cos[2\Phi(t) + \varphi(t)]\,, \end{aligned} \quad (11)$$

where

$$\begin{aligned} A(t) = \frac{2\pi M}{rD}\Bigg(&\left[1 + (\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}})^2\right]^2 F_+(\theta,\phi,\psi(t)) \\ &+4\ [\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}}]^2 F_\times(\theta,\phi,\psi(t))\Bigg)^{1/2}\,, \end{aligned} \quad (12)$$

and

$$\varphi(t) = \tan^{-1}\left(\frac{2\ (\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}})\ F_\times(\theta,\phi,\psi(t))}{[1 + (\hat{\mathbf{L}}(t)\cdot\hat{\mathbf{N}})]^2\ F_+(\theta,\phi,\psi(t))}\right)\,. \quad (13)$$

In summary gravitational wave signals from an NSBH precessing binary system can be expressed in terms of the following parameters: the component masses (M_1, M_2), the black hole spin vector $\mathbf{S}$, the overall constant amplitude A, the polar angles of total angular momentum vector (θ_J, ψ_J), the location of the source (θ, ϕ), the time of arrival of the signal t_0 and the initial phase ϕ_0. For the purposes of this paper we have chosen to focus on "face on" systems, i.e. we assume that $\mathbf{J}$ is either aligned or anti-aligned with $\mathbf{N}$ so that $\theta_J = 0^\circ$ or 180°. For such cases, ψ_J will disappear from the waveform expression. These systems will be, on the average, more luminous than edge-on systems and thus more likely to be detected [9].

For our purposes, the post-Newtonian (PN) formalism provides a reasonable approximation to the observed gravitational waveform. There are a variety of PN approximants available which differ in how one deals with the energy, flux and balance equations (see e.g. [23] for a recent review). Our goal in this paper is not to study the differences between various approximants, but is rather to understand how precession affects the size of the template bank. For this purpose, most approximants should give similar results and our main consideration is computational efficiency. For this purpose, since most of our computations are in the frequency domain, it turns out to be very useful to work directly with the Fourier transform of $h(t)$. We shall use the frequency domain model introduced in [19]. An implementation of this waveform model is publicly available in [24], where it is called the "SpinTaylorF2" model.

B. Matched Filtering

Matched filtering is a methodology used to determine if data from a gravitational wave detector $x(t)$, contains some signal of known form, $h(t)$, or only Gaussian noise $n(t)$. Thus, in the absence of a signal,

$$x(t) = n(t)\,, \quad (14)$$

and in the presence of a signal

$$x(t) = h(t) + n(t)\,. \quad (15)$$

If the noise is stationary, we can characterize it by the single-sided power-spectral-density (PSD) $S_n(f)$ according to

$$\langle \tilde{n}^\star(f)\tilde{n}(f')\rangle = \frac{1}{2}S_n(f)\delta(f - f')\,. \quad (16)$$

Here the brackets $\langle\cdot\rangle$ denote an average over many realizations of the noise, and $\tilde{n}(f)$ denotes the Fourier transform of $n(t)$.

The PSD is used to define the inner product between two time-series $x(t)$ and $y(t)$:

$$(x|y) := 4{\rm Re}\int_0^\infty \frac{\tilde{x}^\star(f)\tilde{y}(f)}{S_n(f)}df\,. \quad (17)$$

This inner product is used to define the norm of a time series $x(t)$ and a normalized time series $\hat{x}$ in the usual way:

$$||x|| := (x|x)^{1/2}\,, \qquad \hat{x} = x/||x||\,. \quad (18)$$

The likelihood function Λ can be shown to be [25, 26]

$$\log\Lambda = (x|h) - \frac{1}{2}(h|h)\,. \quad (19)$$

The idealized procedure to search for a signal with unknown parameters is to compute $\log\Lambda$ for all points (suitably discretized) in a given parameter space and to find the point where $\log\Lambda$ is maximum. The likelihood can be analytically maximized for certain parameters (such as the initial phase ϕ_0 and an overall constant amplitude) or by a Fast-Fourier transform (such as the time of arrival t_0) (see e.g. [27]), while other parameters (the so-called intrinsic parameters) must be explicitly maximized over. These intrinsic parameters we denote as λ_i. (We shall consider only binary systems with circular orbits and we shall also not consider any parameters associated with the internal structure of the neutron star.)

A template bank is a collection of waveforms $\{h_I\}$ labeled by the index I. Given a template bank, we would like to know how effective it is in recovering a given signal h. This is quantified in terms of a number, namely the *fitting-factor* (FF) defined as,

$$FF(h,\{h_I\}) = \max_I \mu(h,h_I)\,, \tag{20}$$

where

$$\mu(h,h_I) = \max_{t_0,\phi_0}(\hat{h}|\hat{h}_I(t_0,\phi_0)) \tag{21}$$

is the *match* between h and h_I. $\mu(h,h_I)$ represents the fraction of the optimal SNR of signal h captured by the template h_I. The fitting factor depends on a particular template bank and a particular *target* waveform h. Since we will compute this for a fixed template bank, we usually drop its dependence on $\{h_I\}$ and write $FF(h)$.

The loss in SNR can be quantified by the match between a signal and the nearest template and can be formulated geometrically [10, 28]. The match between nearby points in parameter space can be approximated as

$$\mu(\hat{h}(\lambda),\hat{h}(\lambda+d\lambda)) = 1 - g_{ij}d\lambda^i d\lambda^j + \dots \tag{22}$$

with the metric

$$g_{ij} = -\frac{1}{2}\left.\frac{\partial^2\mu\left(\hat{h}(\lambda),\hat{h}(\lambda')\right)}{\partial\lambda'_i\partial\lambda'_j}\right|_{\lambda'=\lambda}. \tag{23}$$

This metric[1] is useful in quantifying the density of templates. The higher the metric determinant, the higher the required template density for a fixed given allowed SNR loss (which corresponds to a given minimal match).

For aligned waveforms, there is an analytic expression for the metric [13, 14]. However for the NSBH precessing parameter space, we do not have an analytic way to calculate this metric [29]. We can approximate the precessing metric by calculating the numerical derivative for small perturbations in the waveform parameters ($\delta\lambda \to 0$). The numeric metric is useful because it provides an independent validation of the stochastic bank. For validation, we can qualitatively compare the distribution of the determinant of the metric, and the distribution of templates in the stochastic bank. As we shall see later, the density of templates placed in the face-on bank will correlate with the regions of the NSBH parameter space where the invariant volume element, i.e. the square root of the determinant of g_{ij} is high.

[1] For Gaussian stationary noise, one can show that the metric g_{ij} is equivalent to constructing the scalar product $\frac{1}{2}\left(\frac{\partial\hat{h}}{\partial\lambda_i}\middle|\frac{\partial\hat{h}}{\partial\lambda_j}\right)$ and projecting out the parameters t_0 and ϕ_0.

C. Stochastic Placement

The points in our stochastic template bank are populated using the following steps. The starting point for these could be either an empty template bank or an existing "seed" template bank [15, 16]:

1. Propose a physically viable point in parameter space p following some probability distribution (we call this distribution the *proposal distribution*). If we are starting with an empty bank, then the first proposed point will always be accepted.

2. Calculate the match of the waveform at p with all the waveforms previously accepted into the bank.

3. Append the candidate to the bank if all the matches are below some threshold, known as the *minimal match*. We shall take the minimal match to be 97%.

4. Repeat the previous steps until a convergence condition is achieved. We shall take the convergence criteria to be: continue the process until, in the previous 1000 trials, only 30 or fewer points have been accepted.

The resulting template bank will, of course, depend on the proposal distribution that we start with. If we had sufficiently reliable astrophysical information on spin orientation, mass distributions etc., we could tailor the template bank appropriately. In the absence of any such prior information, we need to apply some other criteria for the proposal distribution. A well motivated choice is to choose the distribution according to the value of the determinant of the metric g_{ij} (this is the choice made in e.g. [17]); indeed a geometric placement algorithm would satisfy this condition. However, this is not the only possibility and we shall discuss our choice below.

D. Proposal Distribution

We shall start with the assumption that the binary system is face-on, i.e. **J** is pointing either directly towards or directly away from the detector and we shall fix the sky-location to be directly overhead the detector. With these

assumptions, we are left with a five dimensional problem: the two masses M_1 and M_2, and the three components of the black hole spin $\mathbf{S}$.

Even with these assumptions, the full problem is a significant computational challenge. An important issue is that the stochastic placement algorithm is not easy to parallelize. Imagine trying to divide the full parameter space into smaller sub-regions and applying the procedure outlined above to each of these sub-regions. Note that in the second step of the procedure outlined in the previous section, we need to check the match of a new waveform with *all* of the previously accepted waveforms in the template bank. Thus, in principle, each sub-region needs to be aware of the points that have been accepted in the other sub-regions. Dealing with each sub-region independently could lead to a significant over-coverage, i.e. accepting many more points than necessary.

If we could find sub-regions which are uncorrelated from each other (by a suitable choice of coordinates) and if the sub-regions were sufficiently large, then the parallelization would be close to optimal. While we do not have the optimal coordinates for this purpose, it turns out that the so-called *chirp times* (τ_0, τ_3) [30, 31] are a good approximation:

$$\tau_0 = {\mathcal{M}_\mathrm{C}}^{-5/3} \tag{24}$$

$$\tau_3 = {\mathcal{M}_\mathrm{C}}^{-2/3}(\nu(1-\nu))^{-3/5}(4\pi - \beta_C) \tag{25}$$

where

$$\beta_C = \frac{1}{12}(38\nu^2 + 75\nu)\chi^{\parallel}\,. \tag{26}$$

The chirp-time was first introduced in [32] as the time taken for the GW signal to reach coalescence starting from some initial frequency. Chirp times are also the coordinates typically used in geometric methods for tempate placement [12] and is also the coordinate where the parameter space metric for binary inspiral systems is most easily understood (see e.g. [33]).

We wish to cover the (τ_0, τ_3) space uniformly. In particular, while constructing the template bank for a particular rectangular region, we would like to ensure that we generate templates only for that rectangular region. If we were to pick values of $M_1, M_2, \mathbf{S}$ directly, this would not be guaranteed. We therefore follow the following steps:

1. Generate values of τ_0 and τ_3 randomly within the rectangular region under consideration following a uniform distribution.

2. The value of τ_0 determines the chirp mass $\mathcal{M}_\mathcal{C}$, but τ_3 depends on both ν and $\chi^{\parallel}$. Our strategy is to then pick a value of q which, along with the chosen value of τ_3 determines $\chi^{\parallel}$. The value of ν is chosen randomly assuming that M_1 and M_2 are uniformly distributed in their allowed ranges. In practice, we draw random values of M_1 and M_2 from uniform distributions, which determines ν. Given q and τ_3, we solve Eq. (25) for $\chi^{\parallel}$.

3. To fix the component of $\mathbf{S}$ perpendicular to $\widehat{\mathbf{L}}$, we note that the total spin magnitude χ is bounded below by $\chi^{\parallel}$. We pick a value of $\chi^{\perp}$, such that $||\chi||$ is uniformly distributed between $\chi^{\parallel}$ and 1.

4. Finally, we choose α_0 uniformly between 0 and 2π.

This procedure ensures that the proposal distribution covers all possible precessing binary configurations. Lower values of τ_3 get mapped to the more aligned systems, i.e. larger values of $\chi^{\parallel}$, while lower values of τ_0 are mapped to systems with higher total mass, M. While the resulting distribution of points in the physical parameters $(M_1, M_2, \chi^{\parallel}, \chi^{\perp}, \alpha_0)$ will not be completely physical, our final results are not very sensitive to this choice of distribution.

III. THE PRECESSING FACE-ON TEMPLATE BANK [FOB]

The total range of chirp times corresponding to our parameter space is broken up into 938 smaller "chirp time boxes". A stochastic template bank is constructed for each box independently and the 938 template banks are then concatenated. Parallelizing the stochastic placement is computationally advantageous, because the stochastic placement algorithm's efficiency scales with the square of the number of templates placed [13, 14]. By splitting up these regions, we limit the number of comparisons needed for each stochastic template bank candidate to decide whether it should be accepted or not. However, speeding up the algorithm comes at the cost of over-coverage between neighboring boxes which we shall discuss towards the end of this section.

The parameter space metric discussed earlier also plays a role in reducing the computational cost, and in particular we use the metric for the space of *aligned-spin* waveforms [14]. The goal is to minimize the number of times that we need to calculate the match. For a proposed parameter space point, we consider only those waveforms which have a match of better than 70% with the proposed waveform as computed by the aligned spin metric. The full match is computed only for the waveforms in the template bank which cross this threshold. The 70% threshold was found by trial and error and is low enough that we do not miss any templates close to the proposed waveform. Fig. 1 shows the convergence of the match for three particular boxes in (τ_0, τ_3) space.

Before presenting the result of the above procedure and discussing some properties of the precessing face on template bank (FOB), we briefly describe an aligned-spin bank (ASB) which we shall use as a reference for comparison. Such bank covers the same space of masses and aligned spin components, but ignores precession. It is constructed via stochastic placement using nonprecessing, inspiral-only post-Newtonian templates (namely the "TaylorF2" model in LALSimulation [24]) and contains 130,646 templates. The template density is shown in

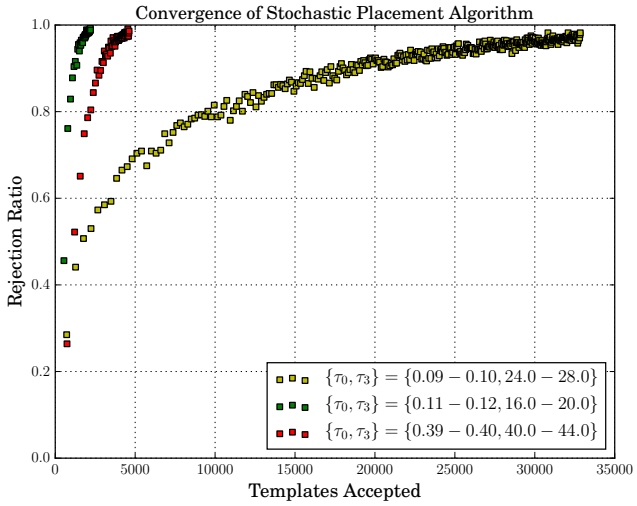

FIG. 1. Convergence curves of three different boxes used to construct the FOB.

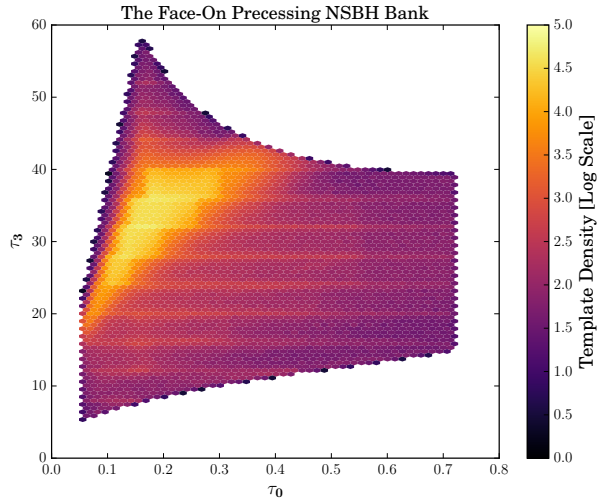

FIG. 3. The FOB in chirp time coordinates. Each hexbin has dimensions $\{\Delta\tau_0 = 0.014, \Delta\tau_3 = 1.0\}$.

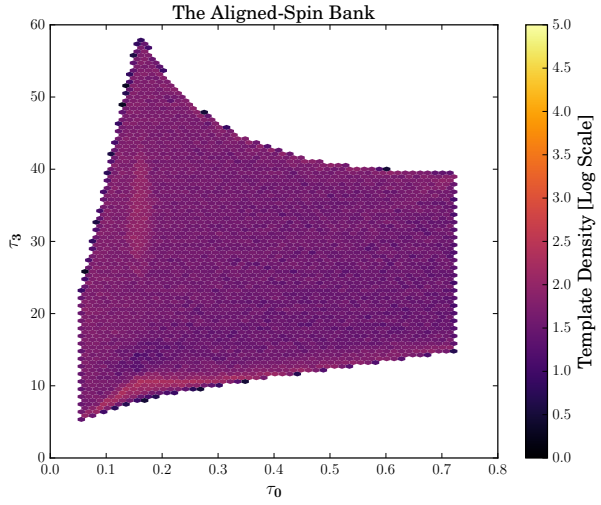

FIG. 2. The aligned spin bank (ASB) in chirp time coordinates. The color bar density scale is the same as in Figure 3 for ease of comparison. Each hexbin has dimensions $\{\Delta\tau_0 = 0.014, \Delta\tau_3 = 1.0\}$.

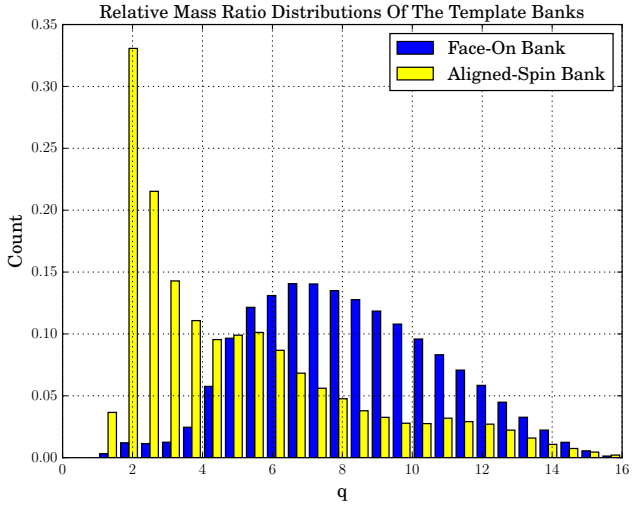

FIG. 4. Normalized distribution of the mass-ratio q for the face-on precessing template and ASB.

Figure 2 in chirp time coordinates (τ_0, τ_3) and in such coordinates it is approximately constant. The ability of a similar bank at detecting aligned-spin and precessing NSBH systems has been characterized in previous studies [9, 31]. In contrast, the template bank for precessing face-on systems is shown in Figure 3. It contains $6,908,681$ templates – a dramatic increase compared to the ASB. The densest parts are in the *high mass ratio, and highly anti-aligned spin ($\kappa < -0.5$) region* of the bank. More than half of the total number of templates are placed in this region. Figure 4 shows the distribution of the mass ratio in the FOB and ASB, thereby demonstrating that the vast majority of points in the FOB consist of asymmetric systems (with mass-ratio $q > 4$) in contrast to the ASB which is dominated by more symmetric systems.

Figures 6, 7 and 8 display the precessing template bank in different slices of the parameter space. Figure 6 shows the template bank density in the $(q, \chi^{\parallel})$ plane. Figure 7 shows the distribution of templates in the $(\chi^{\perp}, \chi^{\parallel})$ plane. Finally Figure 8 gives the template bank distribution in the (q, β) plane. In Figure 7, we note that, higher template densities occur in the higher values of spin-orbit misalignment, which in-turn indicates higher precession.

As a result of breaking up the parameter space into independent boxes, it is to be expected that the algorithm will place more templates than necessary at the

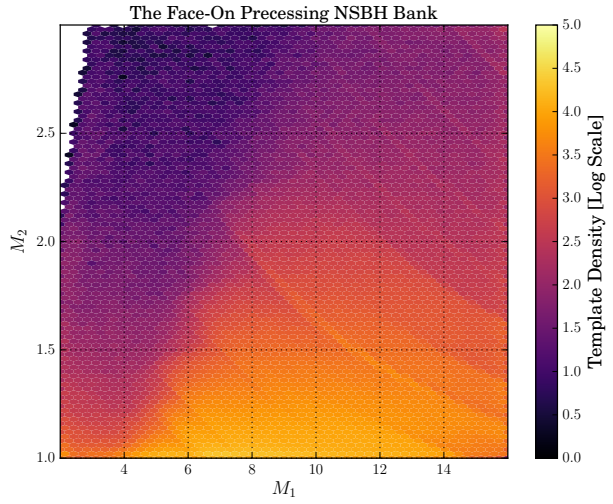

FIG. 5. The FOB in solar mass (M_1, M_2) coordinates. As before, the color bar is scaled with respect to the density of templates per bin $\{\Delta M_1 = 0.28, \Delta M_2 = 0.02\}$.

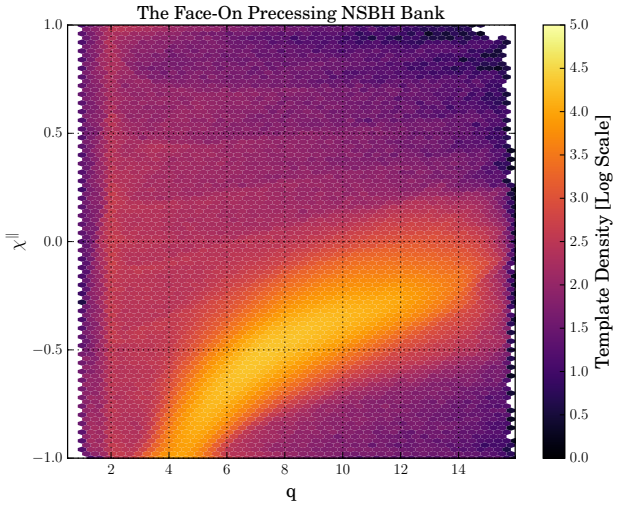

FIG. 6. The FOB in $(q, \chi^{\parallel})$ coordinates. As before, the color bar is scaled with respect to the density of templates per hexbin $\{\Delta q = 0.3, \Delta\chi^{\parallel} = 0.04\}$.

borders between adjacent boxes. This creates so-called "gridlines" in the bank which are clearly visible in Figure 3. These are an artifact of splitting up the parameter space into independent regions. It results in having a larger number of templates than necessary. However, we shall see that this is not a large effect for the chirp time boxes that we have chosen.

The gridlines were most pronounced at the edges of the boxes along the vertical direction, which implies that there is a degeneracy along the τ_3 direction. The gridlines along the τ_0 direction are much less pronounced. This is not surprising since τ_0 is determined entirely by the chirp mass $\mathcal{M}_C$, and it is well known that $\mathcal{M}_C$ is the parameter best determined from the inspiral phase [34].

This suggests also that it should be possible to replace τ_3 by a better coordinate leading to fewer correlations. Regardless, we shall now quantify the correlations between adjacent boxes in the τ_3 direction. Also, in order to optimize the size of the chirp time boxes, it was crucial to estimate how far these gridlines overlapped into adjacent boxes. To study this issue, we looked at two adjacent boxes in τ_3. By taking points in the lower box and calculating the overlap with every point in the above box, we determined the extent of the overcoverage. Fig. 9 displays the templates in adjacent boxes which have an overlap greater than 95% with templates in the adjacent box. The extent of these templates extends to about 25% of the box in the τ_3 direction. However, the number of such templates is only about 7% of the total number of templates in the upper box, and 1% for the lower box.

To conclude this section, we validate the distributions obtained above by a numerical calculation of the Fisher matrix. If one were able to carry out a geometric template placement procedure, the density of templates would be proportional to the invariant volume element,

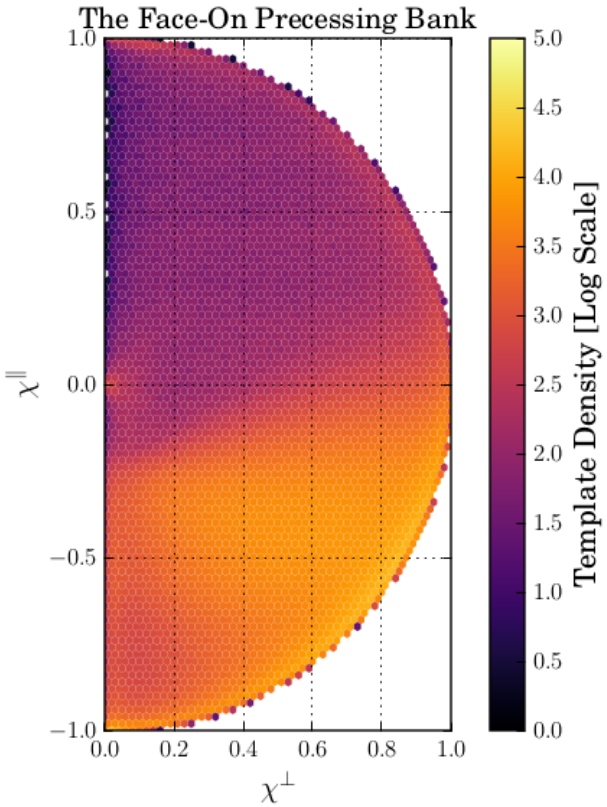

FIG. 7. The spin distribution of the FOB. The y-axis is the component of spin parallel to the orbital angular momentum **L** and the x-axis is the component of the spin perpendicular to **L**. Each hexbin has dimensions $\{\Delta\chi^{\parallel} = 0.04, \Delta\chi^{\perp} = 0.02\}$.

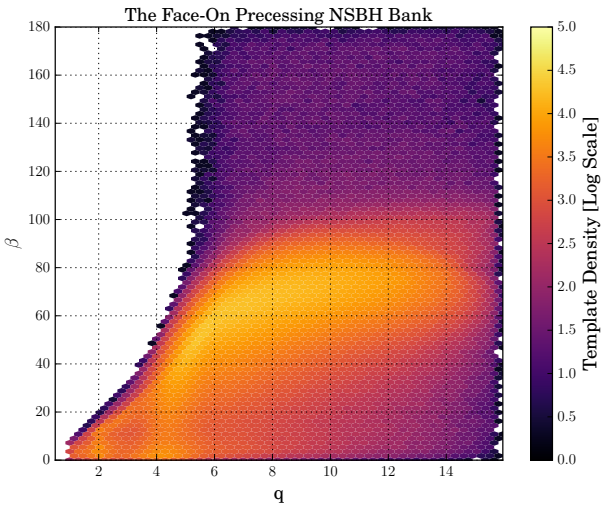

FIG. 8. The β and mass ratio, q, distribution of the FOB. Each hexbin has dimensions $\{\Delta q = 0.3, \Delta\beta = 4^\circ\}$.

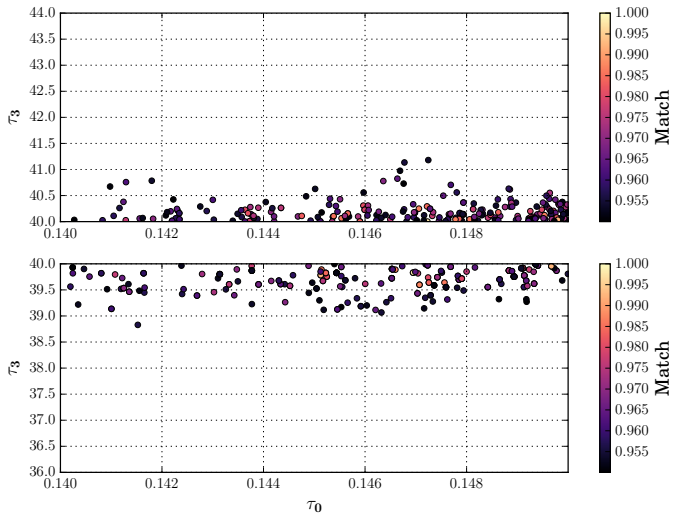

FIG. 9. Plot of the templates in the box $(\tau_0, \tau_3) = (0.14 - 0.15, 40 - 44)$ that had an overlap greater than 95% with templates placed in the box, $(\tau_0, \tau_3) = (0.14 - 0.15, 36 - 40)$.

i.e. to the square root of the determinant of g_{ij}. The same is generally true for probabilistic methods of template placement. We compute g_{ij} and its determinant directly by numerically computing the overlap between the derivatives of neighboring waveforms and compare this with the actual distribution of templates obtained in the template bank. Figure 10 shows the contour plot of $\log\sqrt{|g|}$ for the $\{15M_\odot, 1.4M_\odot\}$ case. Also shown are the points in the template bank whose masses are within 1% of these mass values demonstrating qualitative agreement between the two entirely different calculations. Similar results are obtained for other values of the masses and other slices of the parameter space. This agreement between the two independent calculations provides a sanity check and indicates that the great increase in the number of templates is a real feature of the space of precessing waveforms. Using a different detection statistic as in [20] helps ameliorate the problem somewhat, but does not eliminate it.

IV. EFFECTUALNESS OF THE TEMPLATE BANK

In this section we estimate the effectualness of the FOB for different populations of NSBH systems and compare it with the ASB. When calculating matches, all simulations use a lower frequency cutoff of 30 Hz and an upper cutoff of $4400/(M_1 + M_2)$ Hz, which is the frequency corresponding to the innermost stable circular orbit of a Schwarzschild black hole with mass equal to $M_1 + M_2$.

First, we want to consider the waveforms which were used to construct the FOB, namely the SpinTaylorF2 waveforms and we want to compare the ASB with the FOB for precessing waveforms. In order for a bank to recover signals effectively, it must be able to recover NSBH

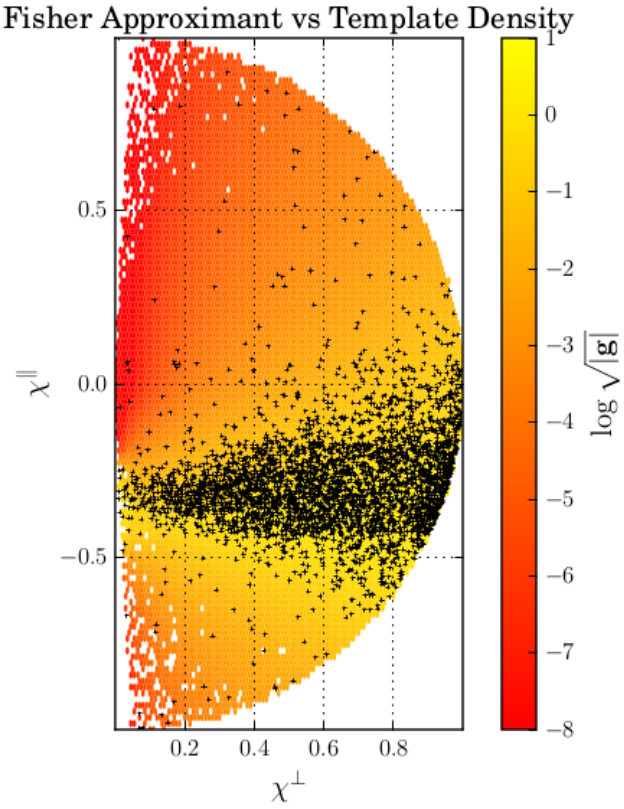

FIG. 10. Comparison of the stochastic bank and the metric approximant χ distributions for fixed masses $\{15M_\odot, 1.4M_\odot\}$. The color map represents $\log\sqrt{|g|}$ and the scattered points denote the templates in the stochastic bank with masses within 1% of $\{15M_\odot, 1.4M_\odot\}$. Each hexbin has dimensions $\{\Delta\chi^\parallel = 0.04, \Delta\chi^\perp = 0.02\}$.

systems over a range of mass and spin values and orientations of $\hat{\mathbf{J}}$. We consider two cases: i) by constraining the injections to be face-on NSBH systems, we look at how well the FOB and ASB could recover SpinTaylorF2 injections from the same proposal distribution used to construct the FOB, and ii) for arbitrary orientations of the total angular momentum (i.e. $0^\circ < \theta_J < 180^\circ$). In both cases, we considered injections over the same $\{M_1, M_2\}$ parameter space as before, i.e. $2M_\odot < M_1 < 16M_\odot$ and $1M_\odot < M_2 < 3M_\odot$. Figs. 11 and 12 show the recovered fitting factors for the ASB and FOB banks for these two cases. Figure 11 shows the case when the injections are face-on. This is what the FOB was built for and indeed, the plot shows that the FOB greatly outperforms the ASB. The fitting factors are worse than 97% for no more than 1% of the injections. Figure 12 shows the corresponding result when the injections are not constrained to be face-on. The recovered matches are reduced, but the FOB still outperforms the ASB over the full mass range.

To further investigate the differences between the FOB and ASB template banks, we now calculate the difference between the fitting factor obtained for the FOB and the ASB (we compute $FF_{FOB} - FF_{ASB}$) and plot the result over different slices of the parameter space. These plots break up the relative performance of the two banks over different portions of the parameter space. Figs. 13 and 14 plot the difference in the fitting factors over (τ_0, τ_3) space. Figs. 15 and 16 show the difference in the fitting factor in $q, \chi^{\parallel}$ coordinates for face-on and arbitrary injections respectively. Here, in Figure 15, as expected, we see that the FOB always performs better. Further, in the regions where the metric has highest density (see. Figure 6), the FOB shows the most improvement. Finally, in what is possibly more illuminating, Figure 17 shows the fitting factors for the FOB in the space of θ_J and the precession cone opening angle β. We quote the value of β at a reference frequency of 100 Hz. While in principle β evolves in time and thus has a frequency dependence, it was shown in [22] that it is roughly constant over the inspiral regime for the frequency range of interest for ground based detectors. Figure 17 shows a clear correlation between the spin orientation and the opening angle. To a good approximation, the figure shows a circle in the θ_J, β plane i.e. cone around the $\beta = 90^\circ$ direction. This relation was found analyticaly in [22] and we refer the reader to this paper for further discussion.

To quantify the improvement that a precessing face-on bank would bring to a CBC search, we calculated the relative improvement in detection volume [31] of the FOB and ASB banks. In the absence of any prior astrophysical likelihood distribution of NSBH systems, the detector volume, $\mathcal{V}$, is proportional to the sum of the cube of the product of the optimal SNR of the injections, ρ_i, with the fitting factor, m_i, obtained from attempting to recover a set of injected NSBH signals into the bank,

$$\mathcal{V} \propto \sum_{\mathrm{i}} (\mathrm{m_i}\rho_\mathrm{i})^3 \,. \tag{27}$$

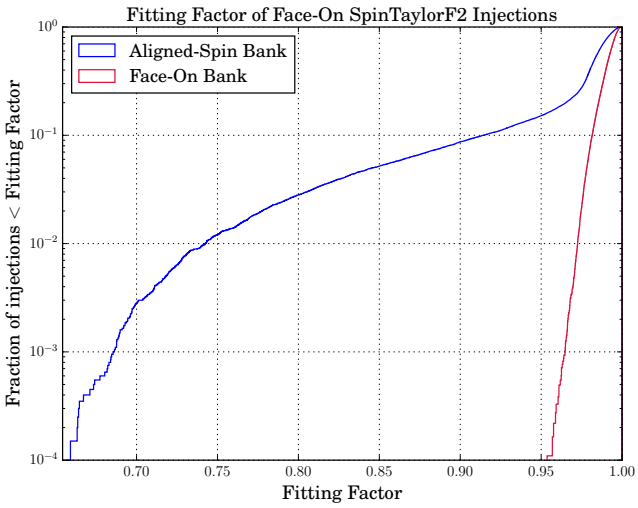

FIG. 11. Cumulative histogram showing the recovered fitting factor of the face-on-precessing and aligned spin template banks for face-on SpinTaylorF2 injections.

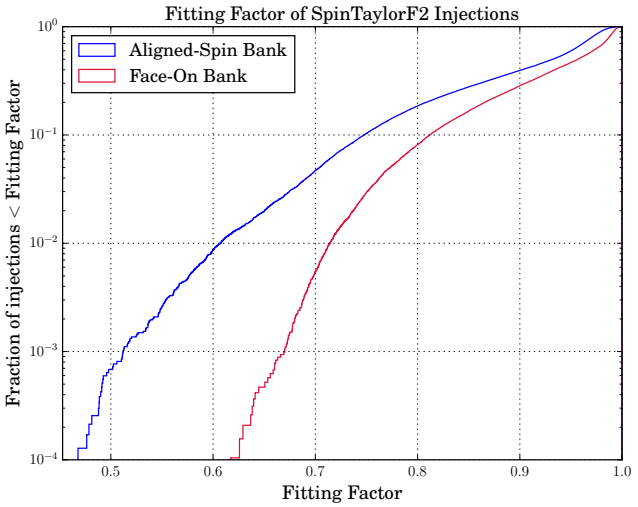

FIG. 12. Cumulative histogram showing the recovered fitting factor of the face-on-precessing and aligned spin template banks for SpinTaylorF2 injections with the component masses distributed uniformly within their respective ranges, spins distributed uniformly in κ, and $\widehat{\mathbf{J}}$ distributed uniformly over the sphere.

By taking the ratio of the detection volumes of the FOB and ASB, $\mathcal{V}_{\mathrm{FOB}}$ vs $\mathcal{V}_{\mathrm{ASB}}$, we get a measure of the relative improvement the FOB could bring to the search. Results are shown in Table I.

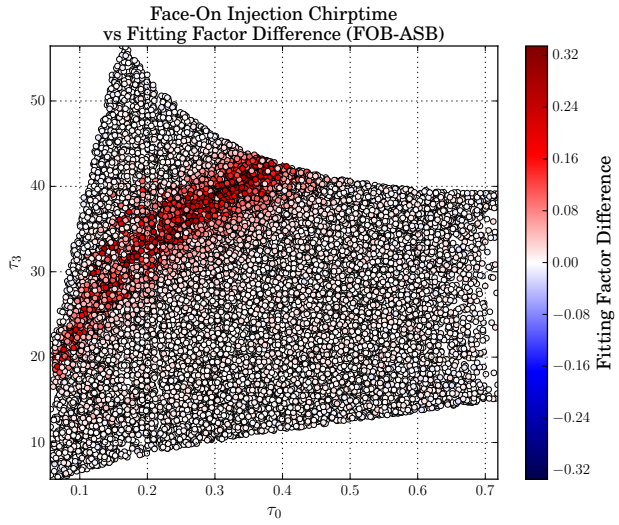

FIG. 13. A plot of the difference in the recovered fitting factor between the precessing and aligned template banks over the $\{\tau_0, \tau_3\}$ parameter space for face-on SpinTaylorF2 injections.

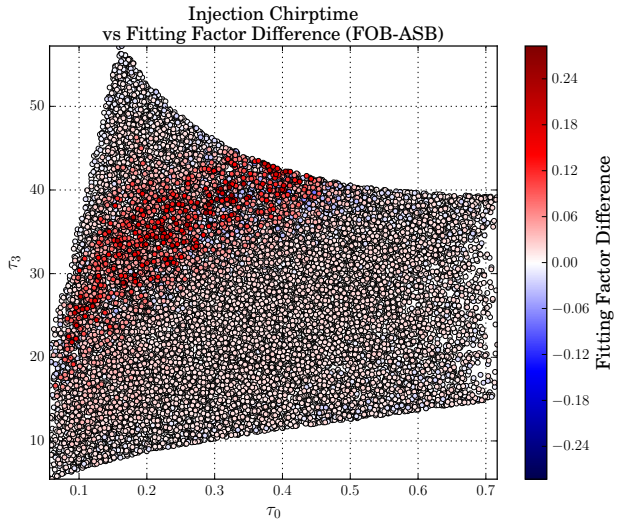

FIG. 14. A plot of the difference in the recovered fitting factor between the precessing and aligned template banks over the $\{\tau_0, \tau_3\}$ parameter space for SpinTaylorF2 injections that are distributed uniformly in chirp time, $\{\tau_0, \tau_3\}$, with $\widehat{\mathbf{J}}$ distributed uniformly over the sphere.

V. CONCLUSIONS

In this paper we have presented a template bank for gravitational wave searches for precessing NSBH systems. The template bank assumes that the total angular momentum vector is pointing directly towards or away from the detectors. It covers the mass ranges $2M_\odot < M_1 < 16M_\odot$, $1M_\odot < M_2 < 3M_\odot$ and the black hole spin vector can have arbitrary orientation. The template bank ends up having $6,908,681$ templates assuming the early Advanced LIGO noise curve. We have shown that the sensitive volume for systems with large spin misalignments (i.e. large precession cone angles) for this template bank is roughy twice as large as for the aligned spin bank (see third row of Table I).

We use the frequency domain, inspiral-only, SpinTaylorF2 waveform for our study. The aligned spin template bank over the same mass range has only $130,646$ templates and this great increase in the number of templates is validated by an independent numerical evaluation of the determinant of the parameter space metric. Despite this large increase in the number of templates, we show that stochastic methods can still be implemented. It requires us to break up the parameter space into smaller, approximately independent regions and we found that the chirp times provide a suitable coordinate choice with which to do this. The template bank could be pruned by removing templates near the boundaries of the chirp time boxes but this would only reduce the number of templates by about 5-10%. Using a different detection statistic as in [20] should further help in decreasing the number of templates somewhat, but it is still an open issue whether the 97% minimal match condition should be kept as gravitational wave detectors improve their low

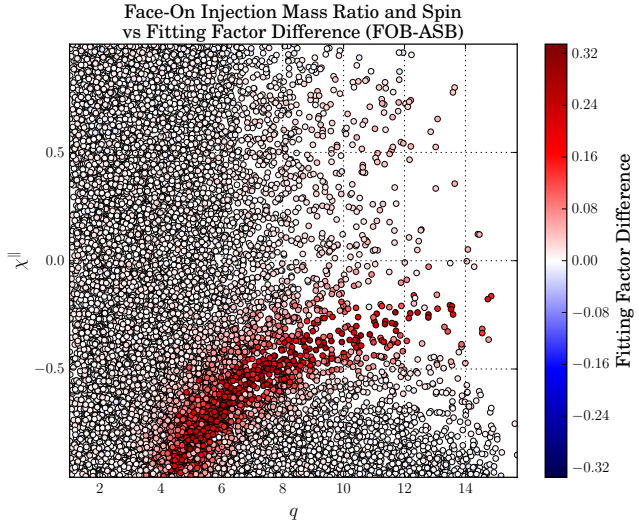

FIG. 15. A plot of the difference in the recovered fitting factor between the precessing and aligned template banks over the $\{q, \chi^{\parallel}\}$ parameter space for face-on SpinTaylorF2 injections.

frequency sensitivity. In either case, working in chirp time coordinates should allow us to deal with the computational problem.

A large fraction of the templates of our bank are in the anti-aligned part of parameter space (with $\kappa < -0.5$). If one believes that such systems are disfavored astrophysically, it is straightforward to construct a precessing template bank for restricted values of κ. Depending on how restricted we would like the black hole spin orientation to be, this might provide a useful compromise between

TABLE I. Table of the improvement in the relative detection volumes calculated from each injection set. The values in the third row represent injections with the component masses distributed uniformly within their respective ranges, spins distributed uniformly in κ, and $\widehat{\mathbf{J}}$ distributed uniformly over the sphere. Results are grouped into three different regions $\{All\,, HP\,, High\beta\,\}$. *All* is the entire NSBH parameter space spanned by the injection set. $High\beta$ is defined as the region of parameter space that contains recovered injections with $\beta \in \{60^\circ, 120^\circ\}$. HP is the "High Precession" region of parameter space examined by [9] that contains recovered injections with $||\chi|| > 0.7$ and $45^\circ < \theta_J < 135^\circ$.

Injected Waveform	θ_J	Mass Range $M_\odot$	$\frac{\mathcal{V}^{\rm All}_{\rm FOB}}{\mathcal{V}^{\rm All}_{\rm ASB}} - 1$	$\frac{\mathcal{V}^{\rm HP}_{\rm FOB}}{\mathcal{V}^{\rm HP}_{\rm ASB}} - 1$	$\frac{\mathcal{V}^{\rm High\beta}_{\rm FOB}}{\mathcal{V}^{\rm High\beta}_{\rm ASB}} - 1$
SpinTaylorF2	0°	$\{2-16, 1-3\}$	3.26%	3.41%	14.2%
SpinTaylorF2	0°	$\{15, 1.4\}$	6.42%	4.66%	23.9%
SpinTaylorF2	$0-180^\circ$	$\{2-16, 1-3\}$	23.7%	9.88%	134%
SpinTaylorF2	$0-180^\circ$	$\{15, 1.4\}$	11.3%	3.22%	14.4%

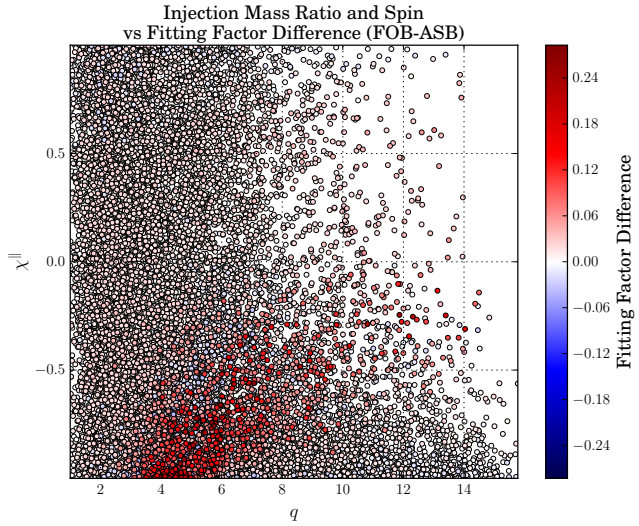

FIG. 16. A plot of the difference in the recovered fitting factor between the precessing and aligned template banks over the $\{q, \chi^{\parallel}\}$ parameter space for SpinTaylorF2 injections that are distributed uniformly in chirp time, $\{\tau_0, \tau_3\}$, with $\widehat{\mathbf{J}}$ distributed uniformly over the sphere.

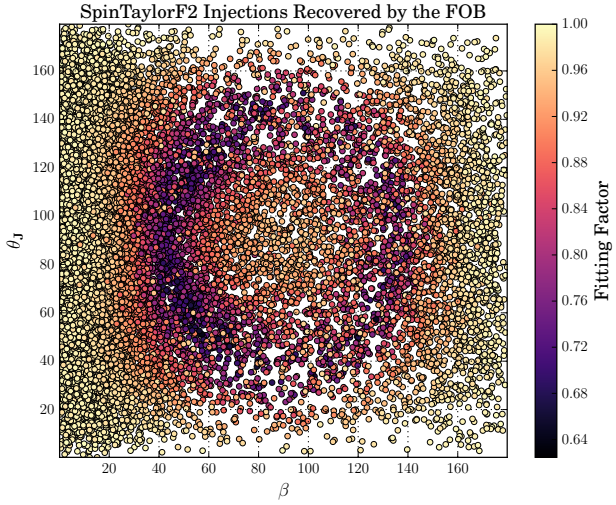

FIG. 17. Recovered fitting factor of the precessing template bank over the $\{\theta_J, \beta\,\}$ parameter space with SpinTaylorF2 injections that are distributed uniformly in chirp time, $\{\tau_0, \tau_3\}$, with $\widehat{\mathbf{J}}$ distributed uniformly over the sphere.

computational cost and astrophysical priors. It would also be desirable to be able to apply traditional geometric methods and to place a lattice of templates, but this requires us to find suitable coordinates for the space of precessing signals.

ACKNOWLEDGMENTS

TDC is supported by an appointment to the NASA Postdoctoral Program at the Goddard Space Flight Center. KH acknowledges DST-MPG Max Planck Partner Group at IISER TVM for the travel support for the visit to the Albert Einstein Institute during which part of the work was carried out. This paper has LIGO document number LIGO-P1600330.

[1] J. Aasi *et al.*, Classical and Quantum Gravity **32**, 074001 (2015).

[2] F. Acernese *et al.*, Classical and Quantum Gravity **32**, 024001 (2015).

[3] B. P. Abbott *et al.* (Virgo, LIGO Scientific), Phys. Rev. Lett. **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].

[4] B. P. Abbott *et al.* (Virgo, LIGO Scientific), Phys. Rev. Lett. **116**, 241103 (2016), arXiv:1606.04855 [gr-qc].

[5] B. P. Abbott *et al.* (Virgo, LIGO Scientific), Phys. Rev. **X6**, 041015 (2016), arXiv:1606.04856 [gr-qc].

[6] B. P. Abbott *et al.* (Virgo, LIGO Scientific), Phys. Rev. **D93**, 122003 (2016), arXiv:1602.03839 [gr-qc].

[7] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Phys. Rev. **D49**, 6274 (1994).

[8] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. X **6**, 041014 (2016).

[9] T. Dal Canton, A. P. Lundgren, and A. B. Nielsen, Phys. Rev. **D91**, 062010 (2015), arXiv:1411.6815 [gr-qc].

[10] B. J. Owen and B. S. Sathyaprakash, Phys. Rev. **D60**, 022002 (1999), arXiv:gr-qc/9808076 [gr-qc].

[11] S. V. Dhurandhar and B. S. Sathyaprakash, Phys. Rev. **D49**, 1707 (1994).

[12] T. Cokelaer, Phys. Rev. **D76**, 102004 (2007), arXiv:0706.4437 [gr-qc].

[13] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Phys. Rev. **D86**, 084017 (2012), arXiv:1207.6406 [gr-qc].

[14] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, Phys. Rev. **D89**, 024010 (2014), arXiv:1307.3562 [gr-qc].

[15] I. W. Harry, B. Allen, and B. S. Sathyaprakash, Phys. Rev. **D80**, 104014 (2009), arXiv:0908.2090 [gr-qc].

[16] S. Babak, Class. Quant. Grav. **25**, 195011 (2008), arXiv:0801.4070 [gr-qc].

[17] C. Messenger, R. Prix, and M. A. Papa, Phys. Rev. **D79**, 104017 (2009), arXiv:0809.5223 [gr-qc].

[18] D. R. Lorimer, Living Rev. Rel. **11**, 8 (2008), arXiv:0811.0762 [astro-ph].

[19] A. Lundgren and R. O'Shaughnessy, Phys. Rev. **D89**, 044021 (2014), arXiv:1304.3332 [gr-qc].

[20] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, Phys. Rev. **D94**, 024012 (2016), arXiv:1603.02444 [gr-qc].

[21] J. Aasi *et al.* (VIRGO, LIGO Scientific), (2013), 10.1007/lrr-2016-1, [Living Rev. Rel.19,1(2016)], arXiv:1304.0670 [gr-qc].

[22] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, Phys. Rev. **D86**, 064020 (2012), arXiv:1203.6060 [gr-qc].

[23] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Phys. Rev. **D80**, 084043 (2009), arXiv:0907.0700 [gr-qc].

[24] LIGO Scientific Collaboration, `https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html`.

[25] L. S. Finn, Phys. Rev. **D46**, 5236 (1992), arXiv:gr-qc/9209010 [gr-qc].

[26] P. Jaranowski, A. Krolak, and B. F. Schutz, Phys. Rev. **D58**, 063001 (1998), arXiv:gr-qc/9804014 [gr-qc].

[27] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Phys. Rev. **D85**, 122006 (2012), arXiv:gr-qc/0509116 [gr-qc].

[28] B. J. Owen, Phys. Rev. **D53**, 6749 (1996), arXiv:gr-qc/9511032 [gr-qc].

[29] R. O'Shaughnessy, P. Nepal, and A. Lundgren, (2015), arXiv:1509.06581 [gr-qc].

[30] A. S. Sengupta, S. Dhurandhar, and A. Lazzarini, Phys. Rev. **D67**, 082004 (2003), arXiv:gr-qc/0301025 [gr-qc].

[31] T. Dal Canton *et al.*, Phys. Rev. **D90**, 082004 (2014), arXiv:1405.6731 [gr-qc].

[32] B. S. Sathyaprakash and S. V. Dhurandhar, Phys. Rev. **D44**, 3819 (1991).

[33] D. Keppel, A. P. Lundgren, B. J. Owen, and H. Zhu, Phys. Rev. **D88**, 063002 (2013), arXiv:1305.5381 [gr-qc].

[34] F. Ohme, A. B. Nielsen, D. Keppel, and A. Lundgren, Phys. Rev. **D88**, 042002 (2013), arXiv:1304.7017 [gr-qc].