

Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors

Tito Dal Canton,¹ Alexander H. Nitz,² Andrew P. Lundgren,¹ Alex B. Nielsen,¹ Duncan A. Brown,² Thomas Dent,¹ Ian W. Harry,² Badri Krishnan,¹ Andrew J. Miller,^{3,1} Karl Wette,¹ Karsten Wiesner,¹ and Joshua L. Willis^{3,1}

¹*Max-Planck-Institut für Gravitationsphysik, Callinstrasse 38, D-30167 Hannover, Germany*

²*Department of Physics, Syracuse University, Syracuse, New York 13244, USA*

³*Abilene Christian University, Abilene, Texas 79699, USA*

(Received 9 June 2014; revised manuscript received 19 September 2014; published 21 October 2014)

We study the effect of spins on searches for gravitational waves from compact binary coalescences in realistic simulated early advanced LIGO data. We construct a detection pipeline including matched filtering, signal-based vetoes, a coincidence test between different detectors, and an estimate of the rate of background events. We restrict attention to neutron star-black hole (NS-BH) binary systems, and we compare a search using nonspinning templates to one using templates that include spins aligned with the orbital angular momentum. To run the searches we implement the binary inspiral matched-filter computation in PYCBC, a new software toolkit for gravitational-wave data analysis. We find that the inclusion of aligned-spin effects significantly increases the astrophysical reach of the search. Considering astrophysical NS-BH systems with nonprecessing black hole spins, for dimensionless spin components along the orbital angular momentum uniformly distributed in $(-1, 1)$, the sensitive volume of the search with aligned-spin templates is increased by $\sim 50\%$ compared to the nonspinning search; for signals with aligned spins uniformly distributed in the range $(0.7, 1)$, the increase in sensitive volume is a factor of ~ 10 .

DOI: [10.1103/PhysRevD.90.082004](https://doi.org/10.1103/PhysRevD.90.082004)

PACS numbers: 04.80.Nn

I. INTRODUCTION

We present here the first realistic gravitational wave (GW) search pipeline for coalescing compact binaries containing a neutron star (NS) and a black hole (BH) with spin aligned with the orbital angular momentum. Our pipeline includes a physical template bank, signal-based vetoes and coincidence between multiple detectors. We show that a simple extension of traditional search methods to include the effects of aligned spin can lead to an appreciable improvement in detection efficiency, even during the early observational runs of advanced GW detectors, before they reach full sensitivity. See [1,2] for descriptions of the advanced LIGO and Virgo detectors. In addition, the KAGRA detector is currently under construction in Japan [3] and an advanced detector has also been proposed in India [4].

In this paper we shall focus on neutron star-black hole (NS-BH) binary systems which are promising sources for the advanced detectors and pose a computational challenge. Based on our current understanding of the population and evolution of binary systems, it is expected that the coalescence rate for NS-BH systems within the sensitive volume of the advanced detectors is in the range 0.2–300/year [5]. To achieve this detection rate, we must be able to distinguish signals from noise at a matched filter signal-to-noise ratio (SNR) of 8 or above in the LIGO detectors. Thus we require accurate models of the signal waveforms for matched filtering, as well as effective methods to exclude false alarms due to non-Gaussian artifacts in the data.

With one exception [6], previous searches of initial LIGO data did not incorporate the effect of the compact objects' angular momentum (spin) in the waveforms used for filtering the data.¹ This was because the search methods and detector sensitivity at the time did not warrant the inclusion of spin effects [6,8]. In general, including extra parameters such as spin in the search increases the size of the template bank, making the search computationally more demanding and increasing the false-alarm rate. An important question therefore is whether more accurate waveform models can offset this increase in false-alarm rate. It has been recently demonstrated that, for the case of BH-BH binaries, including a single effective aligned-spin parameter in the search space does improve the detection rate [9], but the question remains open for NS-BH systems. The initial LIGO detectors had relatively low sensitivity at higher frequencies meaning that the modified phase in the expected signal due to spin was less visible. The situation will be different in the advanced detector era. The advanced LIGO detectors will be able to discern the extra features in the waveform due to the effects of spin for a significant number of events [10,11]. Furthermore, the increased computational requirements for spinning searches can be met by improving the analysis software used to process the

¹However, some other LIGO searches (see e.g. [7]) have quantified how well they could detect spinning systems, even though they conducted the primary search using waveform models without spin.

data and exploiting modern computational platforms such as graphics processing units (GPUs).

The aim of this paper is twofold. First, we show that it is indeed important to incorporate spin in searches for NS-BH compact binary coalescence (CBC) events, even in the early advanced detector era. Second, we describe a new software package for CBC searches known as PyCBC, which is designed to meet the computational challenges of the advanced detector era.

Since the NS-BH merger rate is uncertain by about three orders of magnitude, it is clear that much remains unknown about the population of compact binary systems. A measurement of this rate would constrain models of the formation and evolution of stellar binaries [12,13]. NS-BH systems are also of interest astrophysically because they (along with double NS systems) are expected to be progenitors of short-hard gamma-ray bursts [14,15]. A detection of NS-BH coalescences would allow us to explore the behavior of compact objects in the strong field regime and observation of the merger phase would provide important information about the tidal disruption of NS and their equation of state [16,17]. NS-BH systems have thus been of significant interest for numerical relativity simulations [18–22].

The spin angular momentum of binary objects affects the intrinsic evolution of their orbits due to spin-orbit and spin-spin couplings in the post-Newtonian orbital energy and GW flux. If the spin of the objects in the binary system is not aligned with the orbital angular momentum, then the orbits will also precess [23]. Searches for such precessing signals are computationally demanding as such signals are described by many independent parameters. Previous investigations of the effect of spin on GW searches largely focused on the precessing case, for which a number of phenomenological search templates have been proposed [24–27]. However, none of these attempts were successful when applied to real data. Here we focus instead on the simpler problem of “aligned-spin” systems, where the spin angular momenta are aligned with the orbital angular momentum.

Including the effect of aligned spins still increases the size of the template bank. The larger number of templates increases the number of false alarms in pure noise. False alarms from non-Gaussian transients (glitches) triggering spinning templates had an adverse affect on previous attempts to include spin effects in searches [8]. To counter this problem we also include here a signal-based veto, the χ^2 test [28,29] used in previous LIGO searches [7,30–39]. This veto reduces the significance of glitch-induced triggers in the search, and thus greatly reduces the threshold on signal SNR that must be applied to achieve a desired false-alarm rate. In order to simulate the behavior of real advanced LIGO data, we analyze two months of real data from the two 4 km initial LIGO detectors at Hanford (H1) and Livingston (L1), recolored to a spectrum typical of the

sensitivity that advanced detectors are expected to achieve in 2016–2017 [40].

A comparison of searches with and without including spin effects depends critically on the expected distribution of spin magnitudes and orientations in the target astrophysical population. The maximum theoretical spin for an isolated Kerr BH is given by $\chi = 1$ where χ is the dimensionless ratio cJ/Gm^2 between the spin angular momentum J and the mass m . The maximum value of χ to which a BH can spin up due to accretion of matter from a thin accretion disk is thought to be very close to this limit [41]. A number of stellar mass BHs have been discovered using x-ray techniques. Observations suggest that many of them have quite large spins, even close to this maximum limit [42]. This is especially true of the BHs in high-mass x-ray binaries whose measured spins are all above 0.85. These high-mass systems are the most likely to form NS-BH binaries and it is likely that the BHs were born with these high spins since they have had insufficient time to spin up due to accretion [42].

Binary systems that are potential sources for advanced LIGO are expected to have undergone a hypercritical common envelope (HCE) phase [43]. The available modeling of this phase suggests that hypercritical accretion onto the BH will further spin up the BHs [13] from their spin values before HCE. Taken in conjunction with the x-ray data, this suggests that many of the BHs observable to the LIGO detectors will have large spins. We test our analysis with the full range of spin values from -1 to 1 (where negative values indicate spins antialigned with the orbital angular momentum) but we also display results for restricted ranges, including a high spin range 0.7 to 1 .

The maximum possible spin for a NS is set by the break-up velocity, which for expected equations of state corresponds to $\chi \sim 0.7$ [44]; realistic NS values are thought to lie below this, braked by r-mode instabilities and perhaps by GW emission [45–47]. The maximal spin observed for accreting millisecond pulsars corresponds to $\chi \sim 0.4$, but the maximum value observed in a binary of two compact objects is only $\chi \sim 0.03$ [48]. In NS-BH binary systems the BH is likely to form first due to its larger mass and is therefore unable to contribute matter to spin up the NS. We therefore largely ignore the spin of the NS in NS-BH binaries and concentrate on a single spin, that of the BH.

While some studies suggest that an appreciable fraction of NS-BH systems may have significant spin misalignments [49], others suggest a small misalignment for most systems [50]. For small misalignments, the aligned-spin search would be close to optimal. A closed form for the waveform of a single spin precessing system has recently been provided [51], but a full search based on this method has not yet been implemented and may, as an initial step, require an efficient single-spin aligned search similar to that presented here. We may therefore view this investigation as the first step towards a fully precessing search.

A complete gravitational waveform includes the merger and post-merger ringdown signal as well as the inspiral signal. For simplicity, we ignore the merger and ringdown part of the waveform and ignore the possibility that the NS may be tidally disrupted and destroyed during the inspiral phase [52]. A particular recent model for a complete waveform is EOBNRv2, based on the effective one-body (EOB) framework calibrated by numerical relativity waveforms [53]. Standard inspiral-only waveforms were found to match EOBNRv2 waveforms for total masses below 11.4 solar masses for the advanced LIGO detectors [54]. Most of our simulated signals have total masses below this limit, although it is not yet known what the effect of spin is on this limit. The effect of merger and ringdown will be studied in detail elsewhere.

The aligned-spin search pipeline employed in this paper is based on the PyCBC software package [55]. PyCBC is a newly developed toolkit for CBC searches in the advanced detector era written in the PYTHON programming language. It is based on modular software libraries: modules in isolation are quite simple, but can be put together in useful and sophisticated ways. PyCBC allows scientists to create complicated entire end-to-end pipelines for performing CBC searches. PyCBC also enables scientists to use GPUs in a transparent manner. PyCBC builds on software from the LIGO Algorithms Library [56] used in previous LIGO searches.

The rest of this paper is organized as follows. Section II introduces our search pipeline. Section III introduces the PyCBC toolkit and the computational details of the pipeline; this section can be read independently. The template banks are described in Sec. IV and further details of the search are in Sec. V. Section VI presents the main results and Sec. VII provides a summary and directions for future work.

II. SEARCH METHOD

A fair comparison of the effects of spin in a search needs to take into account all the details of a search of real GW data. Therefore we implement a prototype pipeline which can search for both spinning and nonspinning systems. This prototype pipeline is applied to a synthetic data set obtained by recoloring initial LIGO data as described in [57].

We first summarize the basic matched-filtering method employed, which is described in more detail in [58]. Let $s(t)$ be the data stream from a GW detector. Let $n(t)$ be the noise and $h(t)$ a GW signal which may or may not be present in the data stream. Thus, $s(t) = n(t)$ in the absence of a signal, and $s(t) = n(t) + h(t)$ otherwise. We denote the Fourier transform of a time series $x(t)$ as $\tilde{x}(f)$ defined as

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt. \quad (1)$$

With the assumption that $n(t)$ is a stationary noise process, we define its single-sided power spectral density (PSD) $S_n(f)$ as

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

where $\langle \cdot \rangle$ denotes the expectation value over an ensemble of noise realizations. While the assumption of stationarity is not a good one for realistic data, this definition of the PSD is still applicable over short time scales. The non-stationarity is handled by continuously estimating $S_n(f)$ from the data using a modification of the Welch method [59] as described in [58].

The signal $h(t)$ as seen in the detector is a linear combination of the two polarizations $h_+(t)$ and $h_\times(t)$:

$$\begin{aligned} h(t) &= F_+(\mathbf{n}, \psi; t_0) h_+(t - t_0, \phi_0) \\ &\quad + F_\times(\mathbf{n}, \psi; t_0) h_\times(t - t_0, \phi_0) \\ &= A(t) \cos(\phi_0 + \phi(t - t_0)). \end{aligned} \quad (3)$$

Here the beam pattern functions $F_{+,\times}$ depend on the sky position given by a unit-vector \mathbf{n} pointing towards the source, and on the polarization angle ψ (see e.g. [23]). The beam pattern functions $F_{+,\times}$ can be taken to be constant for the duration that the signal is seen by the detector. t_0 is a suitably defined arrival time: in this case we will use an inspiral waveform described by the post-Newtonian approximation, then a convenient choice for t_0 is the termination time, such that the frequency of a signal with GW phase evolution $\phi(t - t_0)$ formally becomes infinite at t_0 . ϕ_0 is the corresponding termination phase. In the restricted post-Newtonian approximation, the slowly varying amplitude $A(t)$ is given by

$$A(t) = - \left(\frac{G\mathcal{M}}{c^2 D_{\text{eff}}} \right) \left(\frac{t_0 - t}{5G\mathcal{M}/c^3} \right)^{-1/4}, \quad (4)$$

with $\mathcal{M} = M\eta^{3/5}$ being the chirp mass of the binary, $M = m_1 + m_2$ the total mass, $\eta = m_1 m_2 / M^2$ the symmetric mass ratio, $D_{\text{eff}} = D / \sqrt{F_+^2 (1 + \cos^2 \iota)^2 / 4 + F_\times^2 \cos^2 \iota}$ the effective distance, ι the angle between the line of sight from the binary system to Earth and the orbital angular momentum, and D the distance to the binary system (see e.g. [58]).

The termination time t_0 can be searched over by an inverse fast Fourier transform (FFT) and the search over ϕ_0 can be handled by an analytic maximization. As shown in [58], this results in having to compute the complex statistic

$$z(t_0) = 4 \int_0^\infty \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(f)} e^{-i2\pi f t_0} df, \quad (5)$$

where \tilde{h} is a suitably normalized inspiral waveform template expressed in the frequency domain. The SNR is then defined as $\rho = |z|/\sigma$, where

$$\sigma^2 := 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)} df. \quad (6)$$

With this normalization, in Gaussian noise in the absence of a signal we would have $\langle \rho^2 \rangle = 2$. For practical purposes the integrations in (5) and (6) are limited to a lower frequency cutoff below which the detector is dominated by seismic noise and an upper frequency cutoff beyond which the post-Newtonian waveform becomes unreliable.

This work focuses on a sensitivity curve that could reasonably represent the early (2016) runs of advanced LIGO [40] and uses a lower frequency cutoff $f_L = 30$ Hz. However, in Sec. IV we also investigate three different sensitivities, namely, (i) the typical sensitivity of initial LIGO during its sixth scientific run [60] with $f_L = 40$ Hz, (ii) the projected sensitivity from [61], used for ease of comparison with the results of [25], with $f_L = 10$ Hz, and (iii) the zero-detuned, high-power design sensitivity of the mature advanced LIGO detectors [40] with $f_L = 10$ Hz. These curves are shown in Fig. 1.

The SNR ρ works well as a detection statistic in Gaussian noise. To deal with non-Gaussian noise of realistic detectors and veto non-Gaussian transients of nonastrophysical origin, other statistics have been developed. A widely used signal-based veto is the reduced χ^2 statistic [28], which computes the partial SNRs ρ_ℓ in p nonoverlapping frequency bands and combines them as

$$\chi_r^2 = \frac{p}{2p-2} \sum_{\ell=1}^p \left(\rho_\ell - \frac{\rho}{p} \right)^2. \quad (7)$$

The bands are chosen so that a true signal with total SNR ρ would have a partial SNR of ρ/p in each band; the union of

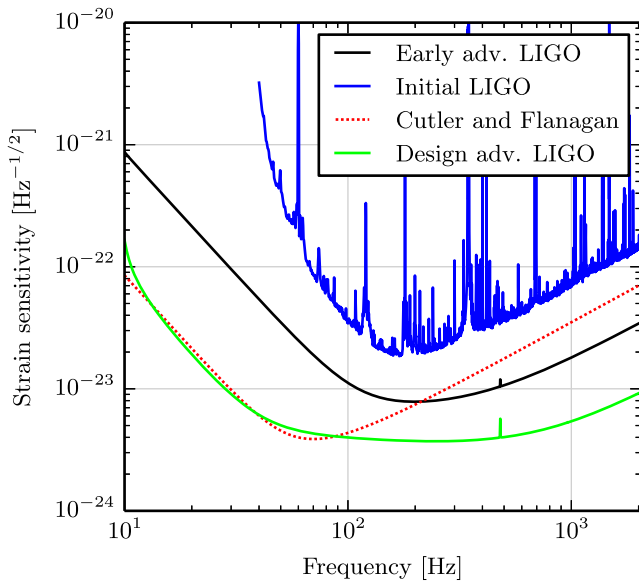


FIG. 1 (color online). Sensitivity curves used in this work. The black solid curve corresponds to the recolored data used for testing the search.

the bands must cover the full frequency range used to compute ρ . We note that computing χ_r^2 for each time sample requires p inverse FFTs and is thus computationally expensive. The exact computational method of calculating ρ and χ_r^2 given a discretely sampled time series $x(t)$, the FINDCHIRP algorithm, is described in [58]. We continue to use the same algorithm in this work.

In order to mitigate the effect of non-Gaussian transients that plagued previous spinning studies [8] we use a modified detection statistic that extends the usual SNR using the χ^2 -veto, known as the *reweighted SNR* statistic [29,31]:

$$\hat{\rho} = \begin{cases} \rho[(1 + (\chi_r^2)^3)/2]^{-1/6} & \text{if } \chi_r^2 > 1 \\ \rho & \text{otherwise} \end{cases} \quad (8)$$

We threshold on both the SNR and reweighted SNR when generating candidate events, and rank them via reweighted SNR; this choice was found to be sufficient for our purposes, although it is possible that other choices of ranking statistic would perform even better.

Our prototype search pipeline is sketched in Fig. 2. We choose “standard” values for most parameters in the

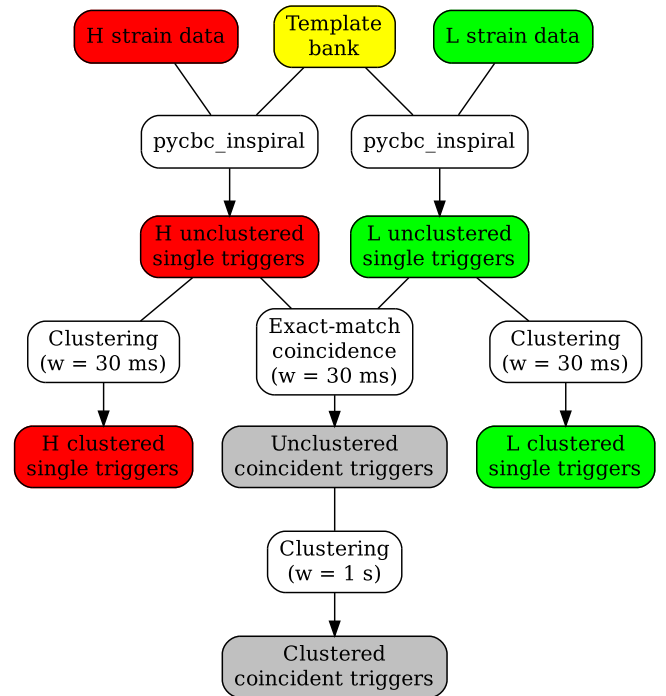


FIG. 2 (color online). Flowchart of the search pipeline. Data from the Hanford (H) and Livingston (L) detectors are processed by the main search engine PYCBC_INSPIRAL which computes the SNR and χ_r^2 time series for a common template bank which is also fixed in time. This results in a list of unclustered single-detector triggers which then pass through a coincidence step followed by clustering over a suitable time window. The single-detector triggers can additionally be clustered independently in each detector.

pipeline (such as the number of χ^2 bands p , the coincidence windows etc.) that have been commonly used in other searches [29]. Notable differences are the coincidence method, and the use of a template bank which is common for all detectors and fixed for the whole data set. A more detailed tuning of the pipeline could improve the sensitivity further. The next sections introduce the PyCBC toolkit and describe the main components of the pipeline in detail.

III. THE PYCBC TOOLKIT

Since 2004, when the first result of a CBC search on LIGO data was published, the bulk of the data analysis for CBC searches has been carried out using software from the LIGO Algorithms Library (LAL) [56]. This is a set of tools and applications written in the C programming language. The computational landscape has diversified significantly over the last 10 years. In particular, the use of GPU for general purpose computing is now more widespread and even central processing unit (CPU) design is moving towards parallel architectures. It is important that the software infrastructure for GW searches is flexible enough to keep up with this diversity.

As we saw in the previous section, the computational cost for CBC analyses is typically dominated by the cost of performing Fourier transforms, primarily in computing the SNR and the χ_r^2 statistics, described earlier, for each inspiral waveform in a large template bank. One strategy would then be to move the FFT calculations to GPUs and keep the remaining computations on the usual CPU of a computer. This would require relatively minor modifications to the existing software in LAL and we could continue using LAL without major modifications. While this does speed up the analysis time somewhat, a detailed profiling of the code reveals that the GPUs are very under-utilized and significant time is spent in data transfers between the CPU and the GPU. This suggests that further development should allow the majority of the computation to run on the GPU.

In this section we describe a new toolkit, PyCBC [55], which builds on the tools available within LAL and makes it easier to assemble complex end-to-end pipelines, and also enables the use of GPUs in a transparent and user-friendly manner. PyCBC is written in the PYTHON programming language [62], a convenient high-level scripting language with a large user community. There are extensive collections of external libraries in PYTHON for a wide variety of tasks, including interfaces to GPUs and general purpose scientific computing. The PYTHON modules of PyCBC need to be able to access the existing LAL software written in C. This is important firstly because C can often be computationally more efficient and secondly because LAL has an extensive collection of GW-specific functionality which has been well tested and widely used within the LIGO and Virgo Collaborations.

PyCBC uses the SWIG framework [63] to access LAL software for CPU computations. This enables one to

perform computations within PyCBC without sacrificing computational speed. PyCBC supports GPU computation via either the CUDA [64] or OPENCL [65] architectures, using, respectively, the PYCUDA [66] or PYOPENCL packages [67].

An example will help us illustrate how these design choices lead to a toolkit that is flexible and maintainable, easy for users to code in, and transparently provides the performance capability of GPUs, while also allowing the same code to run optimally on a CPU when that platform is chosen instead. A simplified script for the basic matched-filtering operation that performs the convolution of a template with a data segment in PyCBC is as follows:

with CUDAScheme:

```
for data in segments:
    for params in bank:
        make_wavform(params, template)
        template *= data
        ifft (template, snr_time_series).
```

While a real code is somewhat more complex, particularly due to the thresholding, clustering, and χ^2 vetoes mentioned earlier, the above sample code shows how the design of PyCBC achieves several important goals:

- (1) Transferring data between the CPU and GPU is transparent to the author of the scripts: he or she need only perform the relevant calculations inside the WITH CUDAScheme block (a context block in PYTHON) and memory will automatically be transferred as it is used in computations within the block. In actual scripts, the context (in the example above, CUDAScheme) is a variable determined at run-time, so that the same script may execute on any CPU, CUDA, or OpenCL platforms.
- (2) We leverage PYTHON's object oriented capabilities to "make simple things simple." In the example above, the multiplication of the template by the data requires only the single $*$ operator, though in reality it represents an element-by-element multiplication of two frequency series, which is also transparently sanity-checked first to ensure the two series have the same length and frequency resolution.
- (3) Simplicity for the user is mirrored by comparative simplicity for PyCBC developers, because the basic PyCBC objects (vectors, time-series, and frequency-series) leverage the uniform interface for arithmetic and basic mathematical operations presented by Numpy (used for CPU), PyCUDA, and PyOpenCL. Considering the wide variety of basic operations, many of which can have multiple instances depending on the precision and type (real or complex) of their inputs, this is a huge saving in development overhead, and immediately provides a functionality not present in LAL.
- (4) The inverse FFT is transparently dispatched to the appropriate library (CUFFT [68] for CUDA, FFTW

[69] or MKL [70] for CPU) which again is not written by the PyCBC developers. In the end, very little code must be separately written for the three supported platforms; in the example listing above, only the generation of the frequency domain waveform would be written and maintained by the PyCBC project directly.

As a result of this design, it is also simple to change which parts of the computation are performed on the CPU, and which on the GPU. The code listing above shows part of a script where the entire matched filter computation, and not only the inverse FFT, is performed on the GPU. It therefore makes more efficient use of the GPU while at the same time requiring very little additional coding.

Given the large number of templates in the spinning template bank, our search is computationally costly; the vast majority of the cost is represented by the matched-filtering stage, while coincidence and clustering are comparatively trivial. The Atlas cluster [71] at the Albert Einstein Institute in Hannover is equipped with Nvidia Tesla C2050 GPUs and PyCBC's flexibility allows us to accelerate the search by running the matched-filter stage on these GPUs. Our implementation of the CBC matched filtering engine uses roughly 35% of the GPU time as reported by the NVIDIA-SMI tool [72]. A detailed profiling of the code, a performance comparison between CPUs and GPUs, and further optimizations will be presented elsewhere.

IV. TEMPLATE BANKS

To describe the template bank used in our search, we establish some standard notation. The inner product between two signals $h_1(t)$ and $h_2(t)$, also known as the *overlap*, is defined as

$$(h_1|h_2) = 4\text{Re} \int_0^\infty \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)} df. \quad (9)$$

We define the normalized signal as $\hat{h}(t) := h(t)/\sqrt{(h|h)}$. The *match* between the two waveforms is defined by maximizing the inner product between the two normalized waveforms over the time of arrival and the phase of, say, h_2 :

$$m(h_1, h_2) = \max_{t_0, \phi_0} (\hat{h}_1|\hat{h}_2(t_0, \phi_0)). \quad (10)$$

Consider a template bank of N waveforms h_I with $I = 1 \dots N$ that is meant to cover a particular parameter space of masses and spins. The *fitting factor* for any waveform h in the parameter space with the template bank is defined as:

$$FF = \max_I m(h_I, h). \quad (11)$$

In constructing a template bank, a common requirement is that any waveform h in the target parameter space must

have a fitting factor larger than 0.97 [29]; thus, any waveform h in the parameter space must match some waveform in the template bank by at least 0.97. In the actual spinning template bank employed we find that matches can fall as low as 0.94. This small deterioration of the minimal match condition occurs only in a small region of parameter space for low values of $\eta \sim 0.05$ and will not greatly affect signals with BH masses below $15M_\odot$ and NS masses around $1.35M_\odot$. The cause of these lower match values is discussed towards the end of this section.

In this section we compare a spinning template bank with a nonspinning template bank. Both banks are constructed using a stochastic placement procedure that was previously presented in [10]; a general introduction to stochastic template banks can be found in [73,74]. We use the stochastic bank algorithm implemented within the PyCBC framework. The template waveforms use the restricted frequency-domain TaylorF2 approximant containing 3.5 pN nonspinning phase corrections [75,76] and 2.5 pN spinning phase corrections [61,77–79]. When calculating the matched-filter SNR, our template waveforms terminate at a frequency corresponding to the innermost stable circular orbit (ISCO) of a Schwarzschild BH of the same total mass as the template, i.e. $f_{\text{ISCO}} := c^3(6\sqrt{6}\pi GM)^{-1}$. This was the standard choice in past CBC searches. However, in the construction of our banks, templates are assumed to terminate at a fixed frequency of 1000 Hz, which is close to the maximum ISCO frequency in our parameter space. Past searches also made a fixed-frequency assumption. Although PyCBC has the ability to construct banks with a varying termination frequency, we do not explore the effect of this choice in this study.

The template bank for the nonspinning search has a BH mass m_{BH} ranging from 3 to $15M_\odot$ and a NS mass m_{NS} ranging from $1M_\odot$ to the equal-mass boundary $m_{\text{BH}} = m_{\text{NS}}$. We also impose the constraint $M \leq 18M_\odot$. Both spins are constrained to zero. This results in ~ 28000 templates. The bank for the spinning search is constructed instead with $m_{\text{BH}} \in [3, 15]M_\odot$, $m_{\text{NS}} \in [1, 3]M_\odot$, $\chi_{\text{BH}} \in [-1, 1]$ and $\chi_{\text{NS}} \in [-0.4, 0.4]$. Such settings produce ~ 150000 templates, which turn out to be mostly clumped around extremal values of χ_{BH} .

Figure 3 shows the mass boundary of the two banks. As can be seen, the nonspinning bank has a larger mass range for the NS than the spinning bank, in particular it includes part of the binary BH region. We make this choice partly because this is how a traditional low-mass nonspinning search would be carried out and partly to allow spinning signals to be recovered by nonspinning templates with similar chirp mass but closer to the equal-mass boundary, thanks to a degeneracy between spin and symmetric mass ratio [80]. In other words, we explicitly favor the nonspinning search by tolerating a bias in the recovered symmetric mass ratio. The fraction of templates in the nonspinning bank with $m_{\text{NS}} > 3M_\odot$ is $\sim 6\%$.




FIG. 3. Mass boundaries used in constructing our nonspinning and spinning template banks. The nonspinning bank includes templates with NS masses above the usual NS mass range. As explained in the text, these templates are able to detect spinning NS-BH signals with a NS mass in the usual range.

Figure 4 shows the template density of the spinning bank in the (τ_0, τ_3) plane, where

$$\tau_0 = \frac{5}{256\pi\eta f_0} (\pi M f_0)^{-5/3} \quad (12)$$




FIG. 4 (color online). Template density of the spinning stochastic bank in (τ_0, τ_3) coordinates. The black contour delimits the nonspinning region of NS-BH parameter space and the dashed lines show the additional NS mass range allowed by the nonspinning bank. Templates above and below the black contour correspond to $\chi_{\text{BH}} < 0$ and $\chi_{\text{BH}} > 0$, respectively.

$$\tau_3 = \frac{1}{128\pi\eta f_0} (\pi M f_0)^{-2/3} \times \left[16\pi - \frac{\chi_{\text{BH}}}{6} (19\delta^2 + 113\delta + 94) \right] \quad (13)$$

are the *chirp times* [81] extended to include spin-orbit effects, $\delta = (m_{\text{BH}} - m_{\text{NS}})/M$ and $f_0 = 20$ Hz is a fiducial frequency. The nonspinning part of the NS-BH parameter space is shown as a black contour in the figure and the region covered by the nonspinning bank corresponds to the black contour plus the small area delimited by the dashed contour. As can be seen, including the effect of spin broadens the covered region significantly. Moreover, although the density remains approximately constant inside the black contour, it increases noticeably outside; in particular, a large amount of templates is concentrated at small τ_3 values. Better coordinates for representing spinning templates in which the template density is nearly constant are given in [10,82].

A. Fitting-factor calculations

The behavior of a template bank with respect to various signals can be studied without the effect of detector noise by numerically evaluating the fitting factors defined in (11), which can be done by PYCBC. In order to get a first characterization of the effect of spin on a few nonspinning banks associated with the different sensitivity curves and lower-frequency cutoffs shown in Fig. 1, we calculate the fitting factors for such banks using simulated signals with fixed masses ($m_{\text{BH}} = 7.8M_\odot$ and $m_{\text{NS}} = 1.35M_\odot$) and a full range of physical spins for the heavier object ($-1 < \chi_{\text{BH}} < 1$) and zero spin for the lighter object ($\chi_{\text{NS}} = 0$). The signals are simulated using the standard time-domain SpinTaylorT2 approximant available in LALSimulation [56]. In principle one could choose other approximants such as SpinTaylorT1 or SpinTaylorT4, which treat the Taylor expansions of the energy and flux differently. However, we choose SpinTaylorT2 because it is essentially the time-domain version of our frequency-domain templates, reducing issues related to agreement between signal and template approximants which are outside the scope of this paper. The waveform generation starts at 20 Hz (well outside the integration range of the matched filter) and terminates at the minimum-energy circular orbit (MECO) after which the evolution of the orbit is no longer expected to be adiabatic (see e.g. [83]). This choice is different from the termination condition assumed in the construction of the template banks (1000 Hz) as well as the upper frequency limit used in matched filtering (the template ISCO frequency). In reality, a physical NS-BH waveform terminates with the merger and ringdown, typically at frequencies higher than ISCO, so any choice of abrupt termination of the signal is artificial. Given that we do not consider NS-BH merger

and ringdown in this study, MECO is a good choice both for implementation reasons and because it is also almost always greater than ISCO. Nevertheless, as discussed later in this section, this discrepancy can affect the fitting factor of binaries at high mass or high positive BH spin. The performance of our template banks for more realistic signal models including merger and ringdown will be assessed in a future study.

The results for the different nonspinning banks are given in Fig. 5, showing similar behavior over different sensitivity curves and lower frequency cutoffs. As can be seen, in all banks there is a range of low BH spins for which the nonspinning bank is able to match the spinning signals fairly well, but then a sharp fall-off in the fitting factor occurs above $|\chi_{\text{BH}}| \sim 0.4$. The shading (color online) of the points shows the recovered value of η in the nonspinning bank. Although the signals are simulated with $\eta = 0.126$, as the spin is increased the recovered value of η also increases, compensating for the larger spin. It can be seen that the sharp fall-off in the fitting factor for positively aligned systems is associated with the maximum physical value of $\eta = 1/4$, corresponding to equal mass templates. Thus, if we had injected signals with a different value of η , the fall-off in match could happen at different values of χ . For the equal mass case $\eta = 1/4$, for instance, we are already at the boundary and η cannot increase any further to compensate for the spin. The match then starts to decrease sharply even for small positive spins. In the rest of the paper




FIG. 5 (color online). Fitting factor of SpinTaylorT2 NS-BH signals with fixed masses ($m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$) and different BH spins and nonspinning template banks constructed for different choices of sensitivity and lower cutoff frequency f_L (see Fig. 1). The nonspinning bank considered in the rest of this paper corresponds to the early advanced-LIGO sensitivity (top right plot).

we will only consider the early advanced LIGO sensitivity curve (top-right panel of Fig. 5) and references to “the (non) spinning bank” will denote banks built for this case.

The loss of match at high BH spins can be further understood by comparing the true and recovered values of the masses in the nonspinning bank, as is done in Fig. 6. The shading now shows the match and again it is clear that the rapid fall-off in match for the positively aligned waveforms is due to the boundary at $\eta = 1/4$. The rapid fall-off in the match for antialigned waveforms is due to a different effect, namely the fact that the minimum m_{NS} in all template banks is $1M_{\odot}$. Unlike the $\eta = 1/4$ case, this is not a physical boundary and one could obtain better matches for highly spinning antialigned systems by lowering the minimum NS mass in the template bank.

The recovered mass and spin parameters in the spinning bank are given in Fig. 7. Here, as expected the templates are all well-matched, although there is a slight bias in the recovered masses and χ_{BH} values. The recovered χ_{NS} value is seen to be widely scattered and it is clear that this does not have any significant impact on the match, nor is its value well recovered by the bank. In other aligned-spin search investigations [9] a single effective spin parameter was used and the minimal impact of the χ_{NS} value seen here is consistent with that approach. The fact that the matches descend below 0.95 despite the design choice that the bank should have a minimal match of 0.97 was also noted in [11] and was explained there by an inconsistency between the termination condition of template and signal waveforms.

As an overall test of the performance of the two banks over the NS-BH parameter space, we calculate fitting




FIG. 6 (color online). Mass parameters recovered in the nonspinning bank for SpinTaylorT2 signals with fixed masses ($m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$, dashed lines) and variable dimensionless BH spin (x axes).

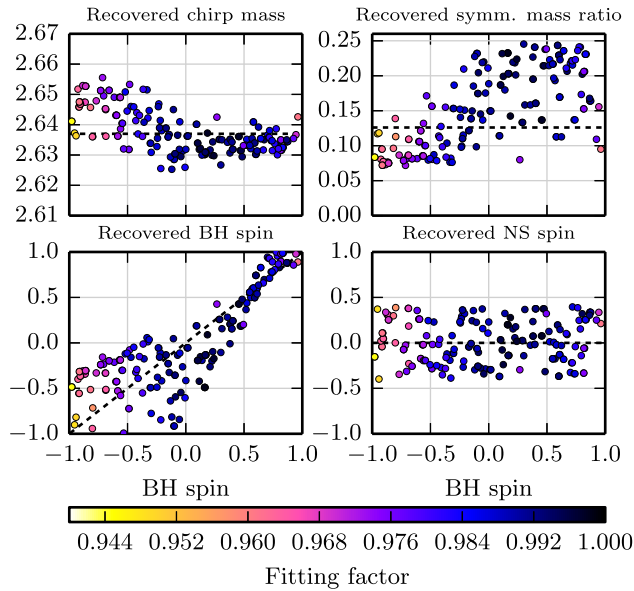


FIG. 7 (color online). Mass and spin parameters recovered in the spinning bank for SpintaylorT2 signals with fixed masses ($m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$) and variable dimensionless BH spin (x axes). Dashed lines show the true parameters.

factors with SpintaylorT2 signals uniformly distributed across the parameter space. The result is shown in Fig. 8, illustrating the deficiency of the nonspinning bank over the parameter space. It can be clearly seen that the values of χ_{BH} at which the match suddenly drops are a function of η ; for equal-mass systems the match starts to drop for $\chi_{\text{BH}} \gtrsim 0$. Interestingly, the spinning bank can have a mismatch as large as 15% in some parts of the parameter space. As noted

already, this is likely an effect of the different (mass- and spin-dependent) termination conditions of the template and test waveforms. In fact, a template terminating before the signal loses the signal power contained between ISCO and MECO, while in the case of a template terminating after MECO the SNR normalization defined in Eq. (6) is too large; both cases result in an effective SNR loss. This hypothesis is supported by the fact that most of the residual mismatch on the right plot of Fig. 8 covers the region of (η, χ_{BH}) plane where a large difference exists between the MECO frequency and either the fixed termination at 1000 Hz assumed in constructing the bank or the ISCO termination used for the template matched filtering.

Although our test signals and templates both use spinning phase corrections up to 2.5 pN order, 3.5 pN corrections have been implemented during the development of this paper and are now ready to be used in searches [84,85]. Unfortunately, including 3.5 pN spinning terms in the MECO definition can lead to very different termination frequencies for our SpintaylorT2 signals, introducing technical difficulties which complicate our fitting factor calculation. Nevertheless, as a rough characterization of the effect of 3.5 pN terms, we test our nonspinning, 2.5 pN TaylorF2 bank against ISCO-terminated TaylorF2 signals with 2.5 pN and 3.5 pN spinning terms. We find that the largest variation in fitting factor when going from 2.5 to 3.5 pN signals is ~ 0.05 , which is comparable to the maximum mismatch used for constructing the bank and well below the loss due to neglecting spinning terms altogether. A more detailed characterization of the inclusion of 3.5 pN spin terms represents a separate study, but we see no reason for not using the best available phasing in future searches.

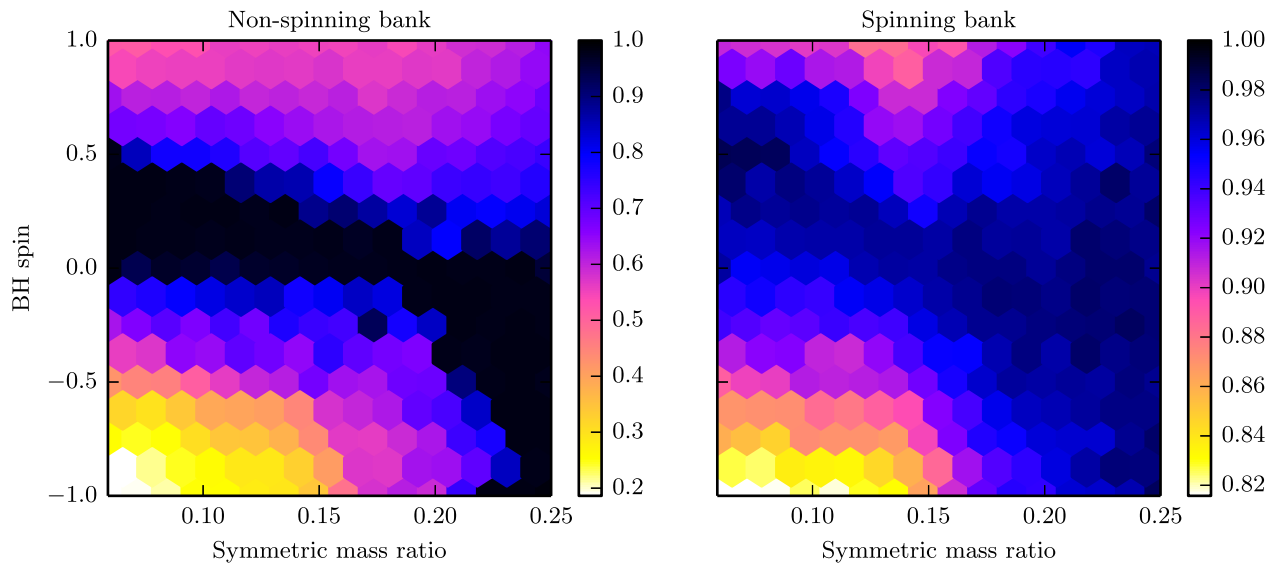


FIG. 8 (color online). Performance of the two template banks across the whole nonprecessing NS-BH parameter space. The color scale shows the lowest fitting factor found in each hexagonal bin (note the different color scales). The region of poor match in the spinning bank is likely due to the different termination conditions of templates and test waveforms.

V. DETAILS OF THE PIPELINE

After having described the template banks, the next step is evaluating the performance of a full search pipeline running on realistic data. This section describes in detail the various components of the pipeline that we implement (see Fig. 2) and the corresponding parameter choices.

A. Inspiral trigger generation

Strain data are first processed by a PYCBC implementation of the standard FINDCHIRP algorithm [58] used in previous CBC searches. After conditioning the data [29] and estimating the noise PSD, the SNR time series is computed for each template with a low-frequency cutoff of 30 Hz. This is lower than past searches because we are targeting the early advanced-LIGO sensitivity. Local maxima of the SNR time series that satisfy the condition $\rho > 5.9$ are identified. The χ_r^2 -statistic is then computed for each surviving maximum via (7) using 16 frequency bands, as is typical in CBC searches [7,30–39], and combined with the SNR via (8) to obtain $\hat{\rho}$. In order to reduce the very large number of maxima produced by glitches, only those with $\hat{\rho} > 5.9$ are kept as candidate triggers. Such triggers are then stored in a MongoDB database [86], where they can be conveniently accessed by the next processing stages and also queried to investigate the features of the data and the search.

The thresholds on SNR and reweighted SNR used here are higher than past CBC searches (e.g. [29,31–34]) as they are chosen to fit the triggers into the available database storage space, which is limited in our prototype setup. There is, in principle, no technical barrier to extending the storage space of the database to handle a larger number of events, which would allow one to lower the threshold back to the usual value.

Even if the noise PSD is continuously estimated from the data in order to evaluate the SNR, our template banks are constructed using the early advanced-LIGO model PSD and thus are constant for the whole data set. This is also a notable difference with past searches, where template banks were regenerated on a time scale of ~ 30 minutes to account for the variability of the noise PSD. Our choice is based on simplicity and the relatively high computational cost of template bank generation. It is also partly justified by the fact that the synthetic data we analyze are recolored to the same noise curve used for constructing the banks. Although the impact of a fixed or varying bank on the sensitivity of a search is not yet fully understood, we expect our choice to have a small effect on the result of our comparison.

B. Coincidence and clustering

The next stage of the pipeline is the identification of triggers in coincidence between the Hanford and Livingston detectors. Although the recolored strain data we analyze cover a two-month period, the different duty

cycles of the detectors reduce the amount of data analyzed in coincidence to about 25 days.

Because we store the triggers in a centralized database, different coincidence methods can be applied to them. We choose an exact-match method, where a trigger in detector A can only form a coincidence with a trigger in detector B if the two triggers share the same template, similar to [9]. This method has the advantage of simplicity and is straightforwardly applicable to parameter spaces of any dimensionality. It requires however a common template bank for all detectors, another difference with respect to past CBC searches, as discussed above. A systematic comparison of different coincidence methods is outside the scope of this paper. Considering that the maximum arrival time delay between Hanford and Livingston is ~ 10 ms, and that the uncertainty in coalescence time is of the order of a few milliseconds [87], we choose a conservative coincidence window of ± 15 ms. Each pair of coincident triggers is stored in the database and is tagged with the *combined* SNR and reweighted SNR, defined, respectively, by summing in quadrature the single-detector SNRs and reweighted SNRs.

In order to keep only the most representative trigger among all the triggers produced by a single inspiral signal or glitch, a final clustering step is performed on coincident triggers. A trigger is defined as representative if no other triggers with higher combined reweighted SNR exist within ± 0.5 s.

Finally, in order to study the distribution of single-detector false alarms, we also perform clustering of single-detector triggers. This works in the same way as coincidence clustering, but it uses a window of ± 15 ms only.

C. Background and sensitivity estimation

In order to estimate the sensitivity of the search, we need to determine (i) the background rate of candidates in the absence of astrophysical signals and (ii) how well the pipeline is able to detect simulated NS-BH signals.

As usual in CBC searches [31], we sample the background distribution via time slides, i.e. by repeating the coincidence step many times, each with a different time delay applied to triggers from one of the detectors. We use 800 time delays, all multiples of 5 seconds. To avoid the possibility of true signals contaminating the background, coincident triggers with zero time delay are excluded from the sample.

We estimate the sensitivity to a population of NS-BH binaries by simulating each binary's gravitational waveform, adding it to the strain data, analyzing the data and recovering coincident triggers (if any) corresponding to each coalescence. We perform three separate analysis runs with simulated signals spaced over the full duration of recolored data at intervals of ~ 10 minutes, resulting in $\sim 3 \times 10^4$ signals in total. The source population is chosen to cover the parameter space reasonably broadly, while being astrophysically plausible. The BH mass is assigned a

Gaussian distribution centered on $7.8M_{\odot}$ with a standard deviation of $3M_{\odot}$, truncated to the $[3, 12]M_{\odot}$ range. The mean value is motivated by [88] which suggests a mass distribution $(7.8 \pm 1.2)M_{\odot}$ for low-mass x-ray binaries; we choose a broader distribution with the same mean. The NS mass is also Gaussian distributed with mean $1.35M_{\odot}$ and standard deviation $0.13M_{\odot}$ (following [89]), truncated to $[1, 2]M_{\odot}$. Since this study ignores precession, both spins are aligned with the orbital angular momentum. Both χ_{BH} and χ_{NS} are distributed uniformly, over ranges $[-0.99, 0.99]$ and $[-0.05, 0.05]$, respectively; as described later, however, we also consider three subsets of the χ_{BH} range. The orbital angular momentum is distributed isotropically.

For the distance distribution, it is useful to introduce the notion of a *chirp distance*.² In the frequency domain, the amplitude of the signal in the restricted post-Newtonian approximation is proportional to $\mathcal{M}^{5/6}/D$ with \mathcal{M} and D being, respectively, the chirp mass and distance defined in Sec. II. The chirp distance is then defined as

$$\mathcal{D} = D \left(\frac{\mathcal{M}_{\text{BNS}}}{\mathcal{M}} \right)^{5/6} \quad (14)$$

with $\mathcal{M}_{\text{BNS}} \approx 1.22M_{\odot}$ being the chirp mass of a canonical binary NS system. This quantity conveniently absorbs all the mass dependent terms in the amplitude: to a first approximation, the detection efficiency should have no additional mass dependence. We then simulate a uniform distribution of sources over chirp distance, in the interval $[1, 160]$ Mpc. Though unphysical, the choice of uniform chirp distance ensures that (i) the efficiency-vs-distance curve is sampled accurately across its variation from 1 to 0 and (ii) the most massive sources do not dominate the recovered sample simply because of their high mass. As for our fitting-factor calculations in Sec. IV, the signal waveforms are simulated via the standard SpinTaylorT2 approximant from LALSimulation [56], starting at 20 Hz and terminating at the MECO.

The sensitivity of the searches is estimated by applying a window around the parameters of each source and recovering the most significant coincident trigger within that window. Based on the results of the fitting-factor simulations in Sec. IV, we choose a coalescence-time window of ± 0.5 s and a chirp-mass window of $\pm 0.6M_{\odot}$. The figure of merit we compute to compare the sensitivity of the two searches is

$$V(\rho^*) = \frac{\sum_i \mathcal{D}_i^2 P_i(\rho^*)}{\sum_i \mathcal{D}_i^2}, \quad (15)$$

where $P_i(\rho^*) = 1$ if source i is recovered with a ranking statistic larger than ρ^* and equals 0 otherwise, and \mathcal{D}_i is the

²See for instance [30], but note that their definition uses the *effective* rather than physical distance.

chirp distance of source i . Here we use the quadrature sum of reweighted SNRs $\hat{\rho}$ over coincident triggers as ranking statistic. The \mathcal{D}^2 weighting corrects the figure of merit for the unphysical distance distribution of the simulated binaries, such that $V(\rho^*)$ is proportional to the sensitive volume of the search, which in turn is proportional to the expected rate of detections [90].

VI. RESULTS

A. Background

Due to the increased dimensionality of the parameter space when going from nonspinning to spinning templates, we expect a higher false-alarm rate for the spinning search both in single-detector triggers as well as in triggers coincident between the two detectors.

Single-detector background triggers associated with SNR and reweighted SNR are shown in Fig. 9. The spinning search clearly has a higher false-alarm rate for both detection statistics. As is well known from past CBC searches [91], the SNR background exhibits a large tail associated with non-Gaussian transient glitches.³ The spinning search seems to be affected more by glitches, as can be seen from the much larger tail at high SNR. Thanks to the effectiveness of the χ^2 test, however, the reweighted SNR is almost tail-free, although we find that strong glitches can still lead to false alarms noticeably stronger than what is typical in stationary Gaussian noise. The increase of false-alarm rate associated with the reweighted SNR is proportional to the increase in number of templates ($\sim 5\times$) for almost all values of the threshold. Applying a fixed threshold in false-alarm rate implies an increase in single-detector reweighted SNR of 0.5 or less when going from the nonspinning to the spinning search. The background distribution of reweighted SNR falls approximately like $\exp(-k\hat{\rho})$ with $k \sim 4$, such that if the total rate of triggers increases by a factor α , the increase in statistic threshold required to compensate this increase (and thus preserve the same false-alarm rate) is only $\Delta\hat{\rho} \sim \log(\alpha)/k$.

The coincident background distribution over the combined (quadrature sum) $\hat{\rho}$ statistic is shown in Fig. 10. As for single-detector backgrounds, the larger false-alarm rate of the spinning search is consistent with the increase in template bank size except at very low rate, where our background sample is likely affected by a small number of loud glitches. Nevertheless, the increase in ranking statistic required to maintain a fixed false-alarm rate from nonspinning to spinning search is only about 0.3.

³Note that in our test we analyze all available science-mode data, including a few poor-quality data segments which a real search would exclude via data-quality flags [29,92]. Thus, the tail in our SNR background is likely exaggerated.




FIG. 9 (color online). Rate of single-detector false alarms for the spinning and nonspinning searches as a function of the threshold on SNR and reweighted SNR.

B. Signal recovery and sensitivity

As a check of the correct behavior of the search, we first calculate the *optimal* SNR of each simulated source, i.e. the SNR obtained for a vanishing noise realization and a perfectly matched template. For each source producing a coincidence in both searches, we compare its combined optimal SNR with the combined SNR and reweighted SNR actually recovered by the searches (the χ^2 value for a zero noise realization and ideal template is also zero, thus the optimal reweighted SNR is equal to the optimal SNR). We find that the nonspinning search fails to recover a noticeable fraction of both SNR and reweighted SNR for $|\chi_{\text{BH}}| \gtrsim 0.5$, which is roughly consistent with the fitting factor calculations (Fig. 11, top and middle rows). The

impact of the χ^2 test on spinning signals is particularly dramatic, as the loss in reweighted SNR is much larger than the loss in SNR. The spinning search, instead, recovers the expected SNR almost completely, for all values of the BH spin (Fig. 11, bottom-left plot). Note however that sources with optimal SNR larger than ~ 100 have a significant loss in reweighted SNR even in the spinning bank; in fact, the reweighted SNR appears to asymptote to a finite value when the optimal SNR becomes very large (Fig. 11, bottom-right plot). This can be explained by the small but nonzero residual mismatch which is also present in the spinning bank. In fact, with any nonzero mismatch, at some (large) value of ρ the χ^2 statistic eventually starts growing like ρ^2 [58]. Combining this fact with the definition of reweighted SNR [Eq. (8)] results in a finite reweighted SNR for arbitrarily large SNR.

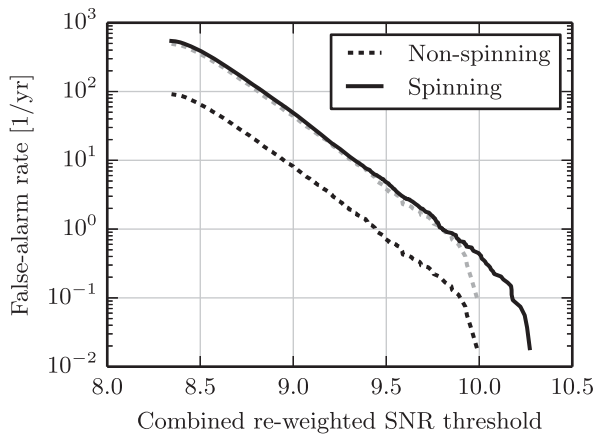


FIG. 10. Rate of coincident false alarms for the spinning and nonspinning searches as a function of the threshold on combined reweighted SNR. The shaded dotted curve shows the nonspinning curve multiplied by the relative number of templates of the two searches ($\sim 5\times$).

Considering the relative sensitivity of the two searches at fixed false-alarm rate, we find that it depends strongly on the distribution of BH spins. Figure 12 shows the receiver operating characteristic (ROC) curves for four populations of NS-BH binaries associated with different BH spin distributions. As done throughout this paper, all cases assume alignment between the BH spin and the orbital angular momentum.

Assuming BHs can have any spin magnitude within the limits of the Kerr bound, we obtain an increase in sensitivity of the spinning search between 40% and 60% depending on the false-alarm rate. A slightly larger improvement is obtained if χ_{BH} is restricted to be positive. If all BHs are highly spinning and positively aligned with the orbital angular momentum, however, the spinning search can be $\mathcal{O}(10)$ times more sensitive than the nonspinning one at interesting false-alarm rates. This large difference can be understood by considering the dramatic loss in reweighted SNR of the nonspinning search, which is

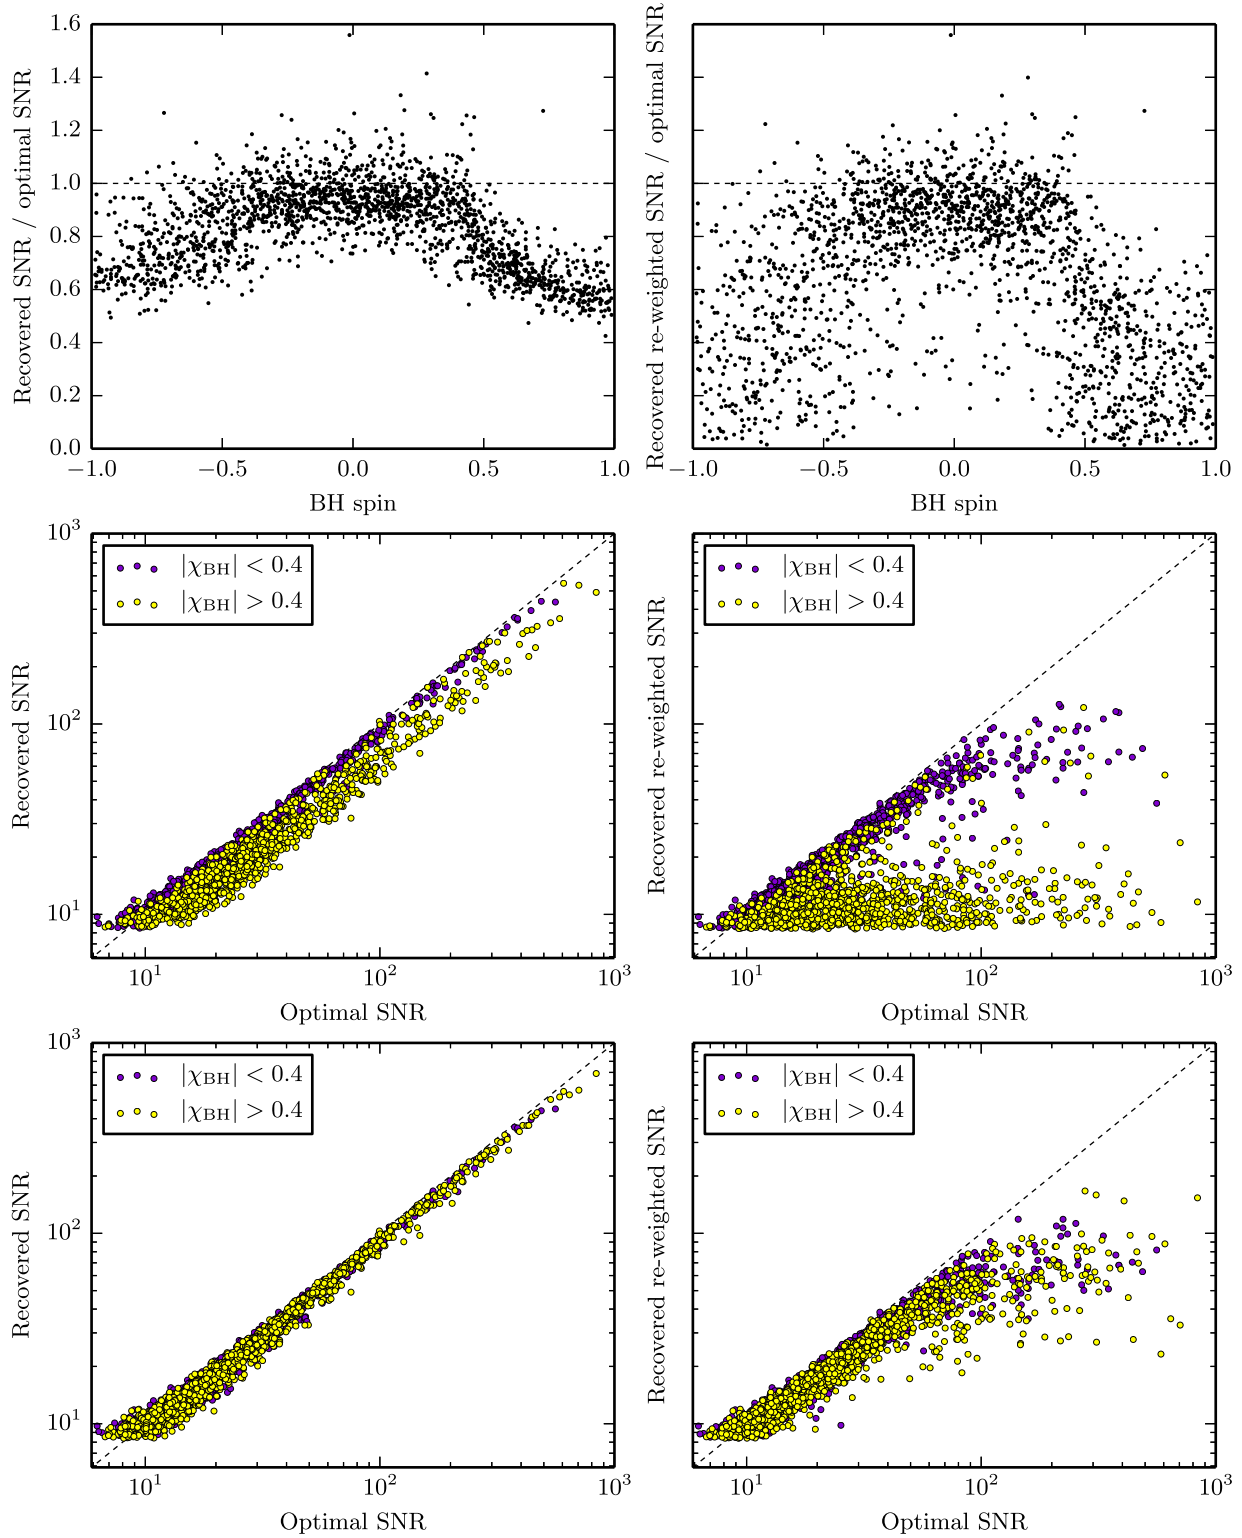


FIG. 11 (color online). Top row: fraction of optimal combined SNR and reweighted SNR recovered by the nonspinning search for each simulated (and found) source, as a function of the BH spin (compare with Fig. 5). Middle row: combined SNR and reweighted SNR recovered by the nonspinning search for each found source vs the optimal combined SNR. The color distinguishes between high (>0.4) and low (<0.4) BH spin magnitude. Bottom row: same as middle row, for the spinning search.




FIG. 12. ROC curves for the spinning and nonspinning searches, comparing the relative sensitive volume (or number of detections) at fixed false-alarm rate. The four panels assume NS-BH systems with different limits on a uniform distribution of BH spins.

due in part to the large SNR loss and in part to the poor χ^2 value of highly-spinning signals. If the search could be carried out using the standard SNR as the ranking statistic, or if we tuned the χ^2 veto differently, the improvement could be significantly less dramatic, but likely still interesting; Fig. 8 (left plot) and Fig. 11 (top-left plot) both suggest a factor of $0.6^{-3} \approx 4.6$ when using the SNR as the ranking statistic. We also note that the improvement could be less dramatic if precession is included in the simulated binaries, but this will be studied in a forthcoming paper. For weakly spinning BHs ($|\chi_{\text{BH}}| < 0.4$), the spinning search is a few percent less sensitive, as can be expected from the larger background, although the difference is comparable with the statistical fluctuations of our ROC curves.

Our signals provide insufficient statistics for studying the case of very small BH spins. Nevertheless, we can conclude with a back-of-the-envelope comparison of the searches assuming a worst-case population of exactly nonspinning BHs. Using the background curves from Fig. 10 and assuming a nonspinning search with detection threshold $\rho^* \gtrsim 9.5$, the relative sensitive volume of the spinning search would be

$$V(\rho^*) \approx \left(\frac{\rho^*}{\rho^* + 0.3} \right)^3 > 90\%. \quad (16)$$

In the worst case, therefore, the spinning search would lose 10% or less of the signals; the major burden would be the larger computational cost.

VII. CONCLUSIONS

We show for the first time how an aligned-spin search for NS-BH binaries can be successfully implemented in the advanced detector era. We demonstrate a prototype search in the PyCBC framework which contains all the essential elements of a realistic CBC search: matched filtering with a template bank, signal-based vetoes, a suitable ranking statistic, coincidence, clustering, background estimates using time slides and sensitivity estimates via a simulated population of signals. The use of reweighted SNR as event ranking statistic [31] is sufficient to reduce the background to a level that makes the search more sensitive than traditional nonspinning searches over the full range of BH spins.

An important element in making this work is running the analysis on GPUs using PyCBC. The design of PyCBC

and the available tools make it easy to put together such a pipeline, and the use of GPU speeds up the analysis relative to using CPUs. Such an analysis would have been much harder with earlier technology.

Our spinning search has improved sensitivity relative to a traditional nonspinning search. The improvement in sensitivity depends strongly on the exact distribution of BH spins. If the magnitude of the dimensionless BH spins (taken to be parallel to the orbital angular momentum) is mostly below ~ 0.4 , the spinning and nonspinning searches have approximately the same sensitivity, despite the fact that the spinning template bank contains many more templates. If the BH spin is distributed uniformly in the range $(-1, 1)$, then the spinning search has approximately 50% greater astrophysical reach as measured by sensitive volume. The increase is one order of magnitude if the spin is uniformly distributed in the interval $(0.7, 1)$. We stress that these results assume that (i) all systems are nonprecessing, and (ii) the BH and NS mass distributions are Gaussians with means $7.8M_{\odot}$ and $1.35M_{\odot}$ and standard deviations $3M_{\odot}$ and $0.13M_{\odot}$, respectively, with the additional constraints $m_{\text{BH}} \in [3, 12]M_{\odot}$ and $m_{\text{NS}} \in [1, 2]M_{\odot}$. Although a careful study of the effect of different mass distributions is outside the scope of this paper, different distributions are unlikely to change the fact that a spinning search is more sensitive than a nonspinning search; this would require an unrealistic distribution restricted to the parameter-space region where the nonspinning bank performs well (see Fig. 8).

The available x-ray data and population synthesis studies suggest that the spin parameters of BHs may be reasonably large, greater than ~ 0.7 in the mass range we used. If this is the case, then for aligned systems an aligned spin search offers a significant improvement in sensitive volume and hence event rate, relative to traditional nonspinning searches. The improvement in search sensitivity could then mean the difference between detection and nondetection, depending on the astrophysical rate of NS-BH coalescence events.

We base our conclusions on an idealized noise PSD which could represent the early runs of advanced LIGO and we employ template banks fixed in time and identical between different detectors. We also show that the SNR loss of a nonspinning bank relative to a spinning one depends on the noise PSD. If advanced LIGO's sensitivity has a significantly smaller bandwidth than our model, or if its noise PSD turns out to have a large variability over a time scale of a few months, the sensitivity of a search to spin effects could be smaller and thus our spinning search could be less beneficial. We argue however that these are unlikely scenarios, as the evolution of advanced LIGO will drive towards the large bandwidth of the final design sensitivity. In addition, if the bandwidth is so narrow that spin effects are less important than we find here, the spinning search would be at least as sensitive as the nonspinning one for a uniform distribution of BH spins and thus would only produce a larger computational cost.

The search methodology we present is straightforward and based on previous CBC analyses, but it has not been fully optimized. We expect that further improvements to the analysis will be possible. These include (i) constructing a better template bank using the full 3.5 pN phasing, a geometric placement algorithm and a mass- and spin-dependent upper frequency cutoff, (ii) correcting the event ranking statistic to reflect the nonuniform distribution of templates over component masses and spins, as described in [93], (iii) improving the ranking statistic by accounting for event distributions over extrinsic parameters such as coalescence time and amplitude, and (iv) tuning the coincidence and clustering steps. Including merger and ringdown effects should further improve the search sensitivity at the higher mass end of the parameter space. The impact of poor data quality on the computing time of inspiral jobs deserves further research. Improved data conditioning techniques such as gating, i.e. appropriately windowing out the data segment in the vicinity of high-amplitude glitches, are under investigation. In developing and testing such improvements to the search, PyCBC will be an essential tool. Finally, investigating the effects of precession will be important as well. This search should be seen as an intermediate step towards a full precessing search; a study is underway to quantify how well the current pipeline performs in detecting precessing signals.

Our conclusions remain robust towards further tuning: when using a template bank that includes the effect of spin, with a signal-based veto such as χ^2 and the infrastructure required to run a search to completion, the gain in signal sensitivity easily outweighs the increase in background. Thus we advise an aligned-spin search rather than a nonspinning search for NS-BH binaries, even for the early advanced detectors.

ACKNOWLEDGMENTS

We thank Gergely Debreczeni, Mate Ferenc Nagy, Frank Ohme, Gianluca Guidi and Richard O'Shaughnessy for useful discussions and comments, and Carsten Aulbert and Oliver Bock for valuable support in using the Atlas cluster and MONGODB at the Albert Einstein Institute. We also thank the LIGO Collaboration for providing the recolored synthetic strain data we analyzed. T. D. C. is supported by the International Max-Planck Research School on Gravitational-Wave Astronomy. A. J. M. and J. L. W. are supported in part by the Pursuit program and Office of Undergraduate Research of ACU. This work is supported by National Science Foundation Awards No. PHY-0847611 and No. PHY-0854812 and a Cottrell Scholar award from the Research Corporation for Science Advancement. Part of the computations used in this work were performed on the Syracuse University Gravitation and Relativity cluster, which is supported by NSF Awards No. PHY-1040231 and No. PHY-1104371. This paper has LIGO document No. LIGO-P1400053.

- [1] G. M. Harry, *Classical Quantum Gravity* **27**, 084006 (2010).
- [2] T. Accadia *et al.*, *JINST* **7**, P03012 (2012).
- [3] K. Somiya, *Classical Quantum Gravity* **29**, 124007 (2012).
- [4] B. Iyer, T. Souradeep, C. S. Unnikrishnan, S. Dhurandhar, S. Raja, and A. Sengupta, <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=75988>, 2011.
- [5] J. Abadie *et al.*, *Classical Quantum Gravity* **27**, 173001 (2010).
- [6] B. Abbott *et al.*, *Phys. Rev. D* **78**, 042002 (2008).
- [7] J. Aasi *et al.*, *Phys. Rev. D* **87**, 022002 (2013).
- [8] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, B. Sathyaprakash, H. Tagoshi, and H. Takahashi, *Phys. Rev. D* **80**, 024009 (2009).
- [9] S. Privitera, S. R. P. 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).
- [10] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev. D* **86**, 084017 (2012).
- [11] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
- [12] I. Mandel and R. O'Shaughnessy, *Classical Quantum Gravity* **27**, 114007 (2010).
- [13] R. W. O'Shaughnessy, J. Kaplan, V. Kalogera, and K. Belczynski, *Astrophys. J.* **632**, 1035 (2005).
- [14] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, *Nature (London)* **340**, 126 (1989).
- [15] R. Narayan, B. Paczynski, and T. Piran, *Astrophys. J.* **395**, L83 (1992).
- [16] V. Ferrari, L. Gualtieri, and F. Pannarale, *Phys. Rev. D* **81**, 064026 (2010).
- [17] F. Pannarale, L. Rezzolla, F. Ohme, and J. S. Read, *Phys. Rev. D* **84**, 104017 (2011).
- [18] Z. B. Etienne, J. A. Faber, Y. Tung Liu, S. L. Shapiro, K. Taniguchi, and T. W. Baumgarte, *Phys. Rev. D* **77**, 084002 (2008).
- [19] Z. B. Etienne, Y. Tung Liu, S. L. Shapiro, and T. W. Baumgarte, *Phys. Rev. D* **79**, 044024 (2009).
- [20] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. D* **78**, 104015 (2008).
- [21] M. Shibata and K. Uryū, *Phys. Rev. D* **74**, 121503 (2006).
- [22] F. Löffler, L. Rezzolla, and M. Ansorg, *Phys. Rev. D* **74**, 104018 (2006).
- [23] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [24] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
- [25] T. A. Apostolatos, *Phys. Rev. D* **54**, 2421 (1996).
- [26] A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003).
- [27] P. Grandclement, V. Kalogera, and A. Vecchio, *Phys. Rev. D* **67**, 042003 (2003).
- [28] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
- [29] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, K. Cannon *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [30] B. P. Abbott *et al.*, *Phys. Rev. D* **79**, 122001 (2009).
- [31] J. Abadie *et al.*, *Phys. Rev. D* **85**, 082002 (2012).
- [32] J. Abadie *et al.*, *Phys. Rev. D* **83**, 122005 (2011).
- [33] J. Abadie *et al.*, *Phys. Rev. D* **82**, 102001 (2010).
- [34] B. P. Abbott *et al.*, *Phys. Rev. D* **80**, 047101 (2009).
- [35] B. Abbott *et al.*, *Phys. Rev. D* **73**, 102002 (2006).
- [36] B. Abbott *et al.*, *Phys. Rev. D* **69**, 122001 (2004).
- [37] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082001 (2005).
- [38] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082002 (2005).
- [39] J. Abadie *et al.*, *Astrophys. J.* **715**, 1453 (2010).
- [40] J. Aasi *et al.* (to be published).
- [41] K. S. Thorne, *Astrophys. J.* **191**, 507 (1974).
- [42] J. E. McClintock, R. Narayan, and J. F. Steiner (to be published).
- [43] M. Dominik, K. Belczynski, C. Fryer, D. Holz, E. Berti, T. Bulik, I. Mandel, and R. O'Shaughnessy, *Astrophys. J.* **759**, 52 (2012).
- [44] K.-W. Lo and L.-M. Lin, *Astrophys. J.* **728**, 12 (2011).
- [45] R. V. Wagoner, *Astrophys. J.* **278**, 345 (1984).
- [46] L. Bildsten, *Astrophys. J.* **501**, L89 (1998).
- [47] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, *Nature (London)* **424**, 42 (2003).
- [48] T. Damour, A. Nagar, and L. Villain, *Phys. Rev. D* **85**, 123007 (2012).
- [49] K. Belczynski, R. E. Taam, E. Rantsiou, and M. van der Sluis *Astrophys. J.* **682**, 474 (2008).
- [50] V. Kalogera, *Astrophys. J.* **541**, 319 (2000).
- [51] A. Lundgren and R. O'Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014).
- [52] F. Foucart, M. Brett Deaton, M. D. Duez, L. E. Kidder, I. MacDonald, C. Ott, H. Pfeiffer, M. Scheel, B. Szilagyi, and S. Teukolsky, *Phys. Rev. D* **87**, 084006 (2013).
- [53] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **84**, 124052 (2011).
- [54] D. A. Brown, P. Kumar, and A. H. Nitz, *Phys. Rev. D* **87**, 082004 (2013).
- [55] PyCBC, <https://www.lsc-group.phys.uwm.edu/daswg/projects/pycbc.html>.
- [56] The LIGO Algorithms Library, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [57] J. Aasi *et al.*, *Classical Quantum Gravity* **31**, 115004 (2014).
- [58] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [59] D. B. Percival and A. T. Walden, *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques* (Cambridge University Press, Cambridge, England, 1993).
- [60] J. Abadie *et al.* (to be published).
- [61] C. Cutler and É. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [62] The PYTHON Programming Language, <http://www.python.org>.
- [63] SWIG, <http://www.swig.org>.
- [64] CUDA, http://www.nvidia.com/object/cuda_home_new.html.
- [65] OPENCL, <https://www.khronos.org/opencv>.
- [66] PyCUDA, <http://mathematician.de/software/pycuda>.
- [67] PyOPENCL, <http://mathematician.de/software/pyopencl>.
- [68] cuFFT, <https://developer.nvidia.com/cuFFT>.
- [69] M. Frigo and S. G. Johnson, in *Proc. IEEE* **93**, 216 (2005).
- [70] MKL, <http://software.intel.com/en-us/intel-mkl>.
- [71] The Atlas Computational Cluster, <https://wiki.atlas.aei.uni-hannover.de/foswiki/bin/view/ATLAS/WebHome>.
- [72] NVIDIA System Management Interface, <https://developer.nvidia.com/nvidia-system-management-interface>.

- [73] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [74] S. Babak, *Classical Quantum Gravity* **25**, 195011 (2008).
- [75] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman, *Phys. Rev. Lett.* **74**, 3515 (1995).
- [76] L. Blanchet, T. Damour, G. Esposito-Farèse, and B. R. Iyer, *Phys. Rev. Lett.* **93**, 091101 (2004).
- [77] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009).
- [78] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995).
- [79] L. Blanchet, A. Buonanno, and G. Faye, *Phys. Rev. D* **74**, 104034 (2006).
- [80] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, *Phys. Rev. D* **87**, 024035 (2013).
- [81] A. S. Sengupta, S. Dhurandhar, and A. Lazzarini, *Phys. Rev. D* **67**, 082004 (2003).
- [82] F. Ohme, A. B. Nielsen, D. Keppel, and A. Lundgren, *Phys. Rev. D* **88**, 042002 (2013).
- [83] Y. Pan, A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **69**, 104017 (2004).
- [84] S. Marsat, A. Bohe, G. Faye, and L. Blanchet, *Classical Quantum Gravity* **30**, 055007 (2013).
- [85] M. Wade, J. D. E. Creighton, E. Ochsner, and A. B. Nielsen, *Phys. Rev. D* **88**, 083002 (2013).
- [86] Mongo DB, <http://www.mongodb.org>.
- [87] A. B. Nielsen, *Classical Quantum Gravity* **30**, 075023 (2013).
- [88] F. Ozel, D. Psaltis, R. Narayan, and J. E. McClintock, *Astrophys. J.* **725**, 1918 (2010).
- [89] B. Kiziltan, A. Kottas, M. De Yoreo, and S. E. Thorsett, *Astrophys. J.* **778**, 66 (2013).
- [90] L. Samuel Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
- [91] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji *et al.*, *Classical Quantum Gravity* **25**, 184004 (2008).
- [92] J. Slutsky, L. Blackburn, D. A. Brown, L. Cadonati, J. Cain *et al.*, *Classical Quantum Gravity* **27**, 165023 (2010).
- [93] T. Dent and J. Veitch, *Phys. Rev. D* **89**, 062002 (2014).