

Hierarchical search strategy for the efficient detection of gravitational waves from nonprecessing coalescing compact binaries with aligned-spins

Bhooshan Gadre,^{*} Sanjit Mitra,[†] and Sanjeev Dhurandhar[‡]
*Inter-University Centre for Astronomy and Astrophysics (IUCAA),
 Post Bag 4, Ganeshkhind, Pune 411007, India*

 (Received 20 July 2018; revised manuscript received 14 January 2019; published 25 June 2019)

In the first three years of gravitational wave (GW) astronomy, more than ten compact binary coalescences (CBCs) have been detected. As the sensitivities and bandwidths of the detectors improve and new detectors join the network, many more sources are expected to be detected. The goal will not only be to find as many sources as possible in the data but to understand the dynamics of the sources much more precisely. Standard searches are currently restricted to a smaller parameter space which assumes aligned spins. Construction of a larger and denser parameter space, and optimizing the resultant increase in false alarms, pose a serious computational challenge. We present here a two-stage hierarchical strategy to search for CBCs in data from a network of detectors and demonstrate the computational advantage in real life scenario by introducing it in the standard PyCBC pipeline with the usual restricted parameter space. With this implementation, we gain an enormous computational speed up, by a factor of ~ 20 , over the flat search on LIGO's first observation run (O1) data. The saving in the computational cost will, in turn, may allow us to search for precessing binaries, will provide more options to search for sources of different kinds and help us to support the never ending urge for extracting more science out of the data with limited resources.

DOI: [10.1103/PhysRevD.99.124035](https://doi.org/10.1103/PhysRevD.99.124035)

I. INTRODUCTION

During the first (O1) and the second (O2) observation runs, the twin LIGO (Laser Interferometric Gravitational-wave Observatory) detectors observed gravitational wave (GW) signals from 11 events with confidence—10 mergers of binary black holes (BBHs) and one double neutron star coalescence [1]. The neutron star coalescence had electromagnetic counterparts in almost every band and is even now being followed by many electromagnetic (radio) telescopes. For the last two of the observations of compact binary coalescences (CBCs), the data from the VIRGO detector also was used supplementing the LIGO [1–9] data. We soon expect to have a larger network of such interferometric detectors with KAGRA coming online soon, and LIGO-India following in few years [10–12]. CBCs are perhaps going to be the most abundant sources for the current and next generation terrestrial interferometric GW detectors [13].

However, GW signals are usually buried deep into noisy interferometric strain data. To extract the signals from CBCs, where phase can be precisely modeled, the method of matched filtering is generally used [14–16] which is optimal in several ways. The signal waveform for a

particular set of signal parameters is obtained from the general theory of relativity by using various techniques involving analytical approximations, perturbation theory, numerical relativity, etc. [17–22]. The modeled signal is then cross-correlated with the inverse noise power weighted data from each of the detectors. This correlation is in fact the maximum likelihood estimator. If the signal with a loud enough correlation is simultaneously present in a pair of detectors with matching parameters, we consider it as a possible astrophysical signal, where the significance of detection needs to be estimated from the statistical properties of the data. This is a simplistic picture however of how coincident detection works. The difficulty lies in the fact that we do not know the signal parameters *a priori* and therefore a search must be carried out in the deemed parameter space.

For these searches, we assume quasicircular orbits for the CBCs. For circular orbits, the GW waveforms depend upon 15 parameters which can be split into two distinct classes: 8 intrinsic and 7 extrinsic. The intrinsic parameters are the component masses (m_1, m_2), individual spin angular momenta (s_1, s_2), and the extrinsic parameters are sky location (θ, ϕ), luminosity distance (d_L), orbital inclination (i), time and phase of coalescence (t_c, ϕ_c), and polarization angle (ψ). The dynamics of the source depends only upon the intrinsic parameters. We can model the generic GW signal in the source frame using essentially the intrinsic parameters and then transform it subsequently to the wave

^{*}bug@iucaa.in

[†]sanjit@iucaa.in

[‡]sanjeev@iucaa.in

frame. For data analysis also, the intrinsic and extrinsic parameters are dealt with differently. One makes use of the symmetries in the signal model to efficiently search over the parameters t_c and ϕ_c (this will be described later in the text). Similarly, the other extrinsic parameters can be dealt with in a quick way [23,24]. However, for the intrinsic parameters, we need to discretize the deemed parameter space. This set of GW signal waveforms at discrete points systematically sampled over the intrinsic parameters is known as a template bank [25–29]. We then search for the signal by correlating all the templates in the bank with the detector data. To search for CBCs in current data from LIGO-Virgo detectors, typically few hundred thousand templates are needed to sample the parameter space with sufficient density, which requires a formidable amount of computation.

The computational cost for a matched filtering search in the full parameter space is too large given the available resources and hence is not feasible. The current searches make a simplifying assumption—to reduce the dimensionality of the parameter space—the spins of the binary components are assumed to be aligned with the orbital angular momentum. These nonprecessing templates can detect a good part of the full parameter space when precessional effects are not dominant [30]. With this set up, matched filter based LIGO pipelines use the template bank with minimal match (MM), the minimum value of scalar product between any two normalized templates, of 0.97. For a search up to a total mass of the $100 M_\odot$, $\sim 250\,000$ templates are required [26,31]. Further, as the low frequency limit of the sensitive band of the detectors is reduced, the number of cycles of the CBC signals in the detector bandwidth increases rapidly, which demands an increase in the template density in the parameter space. Further, better sensitivity at lower frequencies means that we can also observe the heavier binaries, resulting in extension of the detectable parameter space. Both these effects together tend to increase the nonprecessing template bank by at least a few times.

Computational cost is orders of magnitude larger when searching for GW signals from precessing CBC systems. It has been shown that, even for the restricted parameter space of mass ratio less than 5, the precessing template bank with $MM = 0.9$ is more than 10 times larger than the corresponding nonprecessing template bank with $MM = 0.97$ [32]. Fortunately, precession of the binary becomes important only when masses are unequal and orbital inclination is not nearly face-on [30]. Since less power in GW is emitted if the line of sight lies in the orbital plane, the chances of detection of such binaries have been low, which is why one could justify restricting the current matched filtering searches for CBC to dominant mode(s) of nonprecessing signal models only [33,34]. However, with progressively increasing sensitivities of the detectors and the addition of more detectors to the network, one can no longer afford to

miss precessing binaries and the interesting science that they have to offer. While there are claims that, through secular evolution, the component spins of the compact binaries are more likely to align or antialign to the orbital angular momentum when they enter LIGO’s sensitive band, sensitive searches for precessing binaries are needed to test such claims through null detections. Such searches are clearly not feasible using the standard matched filtering scheme with available computing resources. This makes a strong case to develop cost reducing algorithms.

In general, due to the constant demand to extract more science out of a given amount of data, computational costs could get very high and perhaps out of reach of the current available computational resources. The present matched filtering searches employ the coincident detection strategy, instead of the more detection efficient coherent strategy because the coherent strategy is significantly computationally more expensive than the coincident one [35]. It is therefore very important to develop cost effective algorithms for matched-filter based searches, which will allow us to provide more computing resources to search for GWs from other astrophysical sources, e.g., from millisecond pulsars, and will enable us to perform more sophisticated searches, e.g., the precessing coherent search online which is the holy grail of the CBC searches.

In this paper, we propose a hierarchical strategy to search for CBCs in data from a network of GW detectors, the goal being to reduce the computational cost of the analysis. We demonstrate the benefit of this method using spin-aligned template banks. These banks have the advantage that there are fewer parameters over which the search needs to be carried out—there are only four intrinsic parameters to be reckoned with, the two masses and two component spins parallel to the orbital angular momentum. Also since the systems do not precess, the orbital inclination parameter i becomes redundant.

Current searches like `gstLAL`, use singular value decomposition (SVD) like algorithms to numerically reduce the size of the nonprecessing template banks. This makes the matched filtering part of the search computationally significantly cheaper, however the reconstruction from the SVD basis to the actual binary template filtered output requires extra computation [36]. Therefore, there is no overall significant saving in terms of computation as compared to standard search. There are, still, some extra benefits to be gained from `PyCBC` like standard searches and speeding them up so that we can try to accommodate higher modes and precessing effects in the templates. The standard search is also called the “flat” search in the literature [37–39], we use this terminology in this paper.

Here we introduce the hierarchical detection strategy to speed up the matched filtering search using the `PyCBC` pipeline [40–43]. We only consider a 2-stage hierarchical search and compare it with the matched filtering search similar to the one used for the analysis of advanced LIGO’s

first observation run (O1) data. We also present a scheme to estimate the false-detection background, that is necessary to assign confidence levels to detected events.

The layout of the paper is as follows. In Sec. II, we briefly describe the standard flat search or the single stage search for CBCs with matched filtering. In Sec. III, we review the previous use of hierarchical algorithm and we discuss our current implementation of the 2-stage hierarchical search. Then in Sec. IV, we compare the results of our implementation of the hierarchical search with the flat search using aLIGO like simulated data. In Sec. V, we show that the method performs nearly as well without any special optimization. Finally in Sec. VI, we summarize and discuss future directions and the procedure we would like to adopt in these strategies. The method presented in this paper not only shows a proof of concept, but its potential is also demonstrated by applying it to real data.

II. PRELIMINARIES

A. The matched filter

The matched filter (MF) is noise weighted correlation of the modeled GW signal (the template) with the data. It is an optimal detection statistic (in the Neyman-Pearson sense), surrogate of the maximum likelihood statistic when the noise is stationary and Gaussian [14,16,44]. The mathematical form of the MF statistic, which is same as the SNR (usually denoted by ρ) for normalized templates, is maximized over phase of coalescence analytically and also all other parameters of the signal for nonprecessing waveforms, and is given by,

$$\begin{aligned} \rho &\equiv \max_{\lambda} (x, h(\lambda)) \\ &= \max_{\lambda} \left(4\text{Abs} \int_0^{\infty} \frac{\tilde{x}^*(f)(\tilde{h} + i\tilde{h})(f; \lambda)}{S_n(f)} df \right), \end{aligned} \quad (2.1)$$

where x is the time series strain data, $h(\lambda)$ is the normalized expected GW signal for the source parameters of the binary, given by λ and $S_n(f)$ is the noise power spectral density (PSD). The round brackets denote a scalar product on the space of data trains, which has been defined in Eq. (2.1) in the Fourier domain. Tilde (\sim) above a quantity denotes the Fourier space representation of the time series representation of the function. Because of the maximization over phase in the MF, in stationary Gaussian noise, the detection statistic follows a Rayleigh probability distribution in absence of the signal and a Rician distribution when a signal is present in the data [45]. In general, we have no knowledge of the signal parameters λ and therefore we must search over the full parameter space to carry out the maximization. The search over the time of coalescence t_c is performed in a quick way by using fast Fourier transform (FFT) and for ϕ_c a basis of waveforms with $\phi_c = 0, \pi/2$ is used to search over ϕ_c efficiently. For the rest of the parameters, namely, the intrinsic parameters, as discussed

in the introduction, we require a template bank. The template bank is constructed with MM of 0.97. In the next subsection we describe how we construct the template bank.

B. The template bank

The discrete sampling of the intrinsic parameters has to be done with due care. Otherwise we may miss out signals due to the loss of SNR because of the mismatch in the template and signal parameters. There can be many reasons for loss in SNR, mainly it is the phase mismatch which matters the most, which may be due to inaccurate modeling of the signal, etc. But one of the reasons is the mismatch due to the discrete nature of the template bank. As the templates are normalized, $(h(\lambda), h(\lambda)) = 1$, a match between any of the two waveforms with slightly different parameters can be written as follows:

$$\begin{aligned} \mathcal{H}(\lambda, \lambda + \Delta\lambda) &\equiv (h(\lambda), h(\lambda + \Delta\lambda)) \\ &= 1 - ds^2 = 1 - g_{ab}(\lambda) \Delta\lambda^a \Delta\lambda^b, \end{aligned} \quad (2.2)$$

where we have kept lowest order terms in $\Delta\lambda$ and defined the metric g_{ab} as:

$$g_{ab}(\lambda) = -\frac{1}{2} \left(h(\lambda), \frac{\partial^2 h}{\partial \lambda^a \partial \lambda^b}(\lambda) \right). \quad (2.3)$$

The distance ds and template space metric g_{ab} can be used to systematically place templates in the bank with a given value of MM, provided $1 - \text{MM}$ is small. Usually the mismatch $1 - \text{MM}$ is chosen at the level of 3% [26], which means $\text{MM} = 0.97$. This corresponds to a maximum loss of about 10% of the astrophysical events within the detectable range. The metric can be analytically calculated for inspiral waveforms given by the post-Newtonian expansion. But here, we use the full IMR waveforms with nonprecessing component spins in the search. For such waveforms, there is no sufficiently accurate analytic or semianalytic form of the metric which can be used to construct a geometric template bank. Therefore, the current searches use a different approach which employs stochastic methods in order to obtain a template bank [27,29], for which the match is directly computed to obtain a stochastic placement of the templates. If the match is close to unity, then the metric is being used implicitly. If the match is not close to unity as in the case of the coarse bank as explained in Sec. III B 2, then the metric approximation fails.

A template bank depends on the PSD of the noise present in the detector. When we have more than one detector, in general, we have to deal with more than one PSD. However, it is convenient to have a common template bank, which facilitates the coincident detection approach [26,33]. For the two LIGO detectors, we combine the two PSDs into a single effective PSD by taking the harmonic mean of the two PSDs. This effective PSD is used to

construct a common template bank for the search. As the strain noise from the LIGO detectors is neither stationary nor Gaussian—there are glitches in the data—the coincident detection approach is preferred which naturally rules out glitches and facilitates signal consistency checks for astrophysical trigger selection and GW detection.

C. Coincidence and vetoes

Below we describe the criteria for coincident detection. A coincident trigger must satisfy the following: (i) there are corresponding triggers in each detector—the SNRs must cross the preset thresholds for the same data segment, (ii) the intrinsic parameters recovered independently for each detector are such that they match (the same template clicks) and (iii) the difference in the estimated times of coalescence is not be more than light travel time between the two detectors. This difference is allowed a small margin of error because the noise can throw the triggers a little away from their true coalescence times. This is the procedure followed for the current searches in the two aLIGO detectors.

To further reduce the false alarms, χ^2 dependent vetoes are applied in the form of newSNR [23,33] to triggers from each of the detectors. These collected individual detector triggers along with the coincident newSNR statistics are used to estimate the noise background and to assign the statistical significance to the detected GW triggers. We escalate a candidate trigger to a detection if the trigger passes a sufficiently high threshold for which the false alarm probability is very small. These ideas will be made precise later.

III. HIERARCHICAL SEARCH

The idea of a two stage hierarchical search is fairly straightforward. First we search over the parameter space by using a coarse grid with a lower threshold on SNR or the detection statistic. The candidate triggers from the first stage are then followed up by finely sampled the parameter space around the neighborhood (nhbd) of each trigger. The goal is to effectively reduce the number of matched filter computations needed to find a GW signal if it is present in the data. This may also help in reducing the background arising due to false alarms caused by noise artefacts. The speed-up one gets depends on the coarseness of the first stage bank and the false alarm rate which is related to the choice of first stage signal-to-noise-ratio (SNR) threshold. This procedure is optimized by adjusting the first stage threshold to yield minimum computational cost for a fixed search sensitivity usually defined in terms of sensitivity distance or volume for CBC searches [41].

In principle, one could also increase the number of stages of hierarchy, though so far we have restricted ourselves only to two stages.

A. Review of the nonspinning hierarchical search

It has been shown previously that a two stage hierarchical search algorithm can be used to speed-up the nonspinning CBC searches by more than an order of magnitude in simulated initial LIGO (iLIGO) like data [38] and by factor of 7–8 in real data from the second science run (S2) of iLIGO. The first such study was carried out by Mohanty and Dhurandhar [46]. They used Newtonian waveforms and the detector noise was assumed to be stationary and Gaussian. The hierarchy was performed over just one parameter, namely, the chirp mass. This work was extended to hierarchy over both the masses for 1.5 post-Newtonian (PN) inspiral waveforms by Mohanty and Dhurandhar [46], Mohanty [47]. This was then followed up by Sengupta *et al.* [37,38] which further extended the hierarchy to three parameters, namely, the masses and time of coalescence. To incorporate the hierarchy in time of coalescence the data was down sampled in the coarser first stage. 2PN inspiral-only waveforms were used in their analysis. This most recent work used a geometric template bank placement [37,48]. The full details of the previous hierarchical searches with nonspinning GW signal waveforms over simulated and initial LIGO second science run (S2) data are given in [49].

In the latest two stage hierarchical search proposed in [37], chirp times τ_0 and τ_3 were used instead of individual component masses to create fine and coarse template banks. The template space Fischer-Rao metric depends very weakly on the chirp times in the parameter space considered. The geometric fine bank with mismatch less than 3% was created using 2PN inspiral-only metric using hexagonal closed packing template placement scheme with iLIGO noise PSD for masses in the range of $(1, 30)M_\odot$. In the first stage of the search the data were sampled at a lower rate of 512 Hz and the coarse template bank was created with mismatch less than 20%, that is, MM of 0.8. For such large values of mismatch, the metric approximation breaks down. Therefore, the coarse bank is created numerically by a rectangular placement of the templates along the τ_0 axis. In the first stage, the lower MM reduces the number of templates in the bank significantly. Moreover, downsampling reduces the cost of each FFT in each MF operation. However, this reduction in computational cost comes at the cost of reduced SNR of the recoverable signal. Hence, in order to ensure that we do not lose an otherwise detectable GW signal, the applied SNR threshold must be lower than the one used in the single stage flat search which is the usual search with the bank of $MM > 0.97$. With the individual detector SNR thresholds of 6 and 8 for the first and second stage respectively, the search showed computational cost reduction by few orders of magnitude for simulated data with Gaussian noise [37,38] and almost by an order of magnitude during search with iLIGO S2 data.

All the earlier works mentioned above considered only a single detector and did not use any signal consistency tests

such as the χ^2 discriminator. Apart from introducing those essential components in the search to make the implementation applicable for real data, there are two primary routes to further extend the hierarchical search strategies, either by increasing the number of stages in the hierarchy or by including more parameters in the two-stage hierarchy or both. Since the current CBC waveforms include spins, we have opted for the latter. We may explore the feasibility of the former option in future.

B. Hierarchical search with aligned-spin waveforms

In this work we explore the possibility of a hierarchical algorithm for CBC searches with non-precessing template waveforms in the modern set up. We use the PyCBC pipeline [41] with LIGO’s O1 type of search setup [33]. We use the full inspiral-merger-ringdown (IMR) aligned spin waveforms with dominant (2,2) mode. Both coarse and fine template banks are generated using stochastic template placement algorithm [27].

1. Guidelines for tuning the search

Before moving on to the full-fledged pipeline, we present our initial study with “zero-noise” BBH injection case. Studying this simpler case helps in choosing of thresholds in the first and second stages of the hierarchy, the size of the fine bank nbhds for each of the coarse bank trigger template etc. This study of injections without noise is equivalent to the averaging over a very large number of ensemble of detectors with additive Gaussian noise. In simple terms, here we look at the peak of the distribution of the matched filter statistic by assuming zero-noise in the detector but with sensitivity determined by the PSD. We use the AdvLIGO PSD for computing the match using the inner product described in Eq. (2.1).

We consider 2000 binary black hole (BBH) injections in H1-L1 detectors with single detector optimal matched filter SNR in the range of 5 to 15 for each of the detectors. Both the BHs have masses uniformly sampled in the range of $(5, 10) M_\odot$ and spin components along the orbital angular momentum uniformly sampled in the range $(-0.98, 0.98)$. Further, the injections are uniformly spread over the full sky. For this study, we use the actual matched filter SNRs and coincident SNRs without χ^2 (the χ^2 weighted SNRs are not applicable here).

We construct coarse and fine templates bank with MM of 0.9 and 0.97 respectively for the parameter space in component masses and spins as used for the BBH injections. Both the banks are created using template space metric as described in Sec. II B. This is possible because we have used the `TaylorF2RedSpin` approximant in this study for which the analytic metric is available. The coarse bank has 1200+ templates while the fine bank contains 10000+ templates, that is, the fine bank is about eight times denser than the coarse bank. The banks and the proof that

there are no holes (i.e., the prescribed MM condition is satisfied in both the cases) are shown in Fig. 1. We choose a sampling rate of 512 Hz for the construction of the coarse bank (stage I search) and a sampling rate of 1024 Hz for the fine bank (stage II search). We choose these specific sampling rates, because the ISCO frequencies for BHs having masses in the range $10\text{--}20 M_\odot$ are in range 220–440 Hz. With the sampling rates of 512 Hz and 1024 Hz, the Nyquist frequencies are 256 Hz and 512 Hz respectively. The reduction in sampling frequency leads to a loss of SNR for some of the injections in the coarse stage when compared with the generic flat search. The template duration for all the signals under consideration is less than 8 sec. Therefore, data segments have been chosen to be 16 sec in duration for computing matched filters.

We then match filter the data, using stage I and stage II banks for each of the injections in both H1 and L1 detectors and compute coincident SNRs. We compare the SNR in stage I with the SNR in stage II in Fig. 2. The loss of SNR is due to the coarse sampling of the parameter space and




FIG. 1. The figure at the top shows the non-spinning template banks in $\tau_0 - \tau_3$ plane and the one at the bottom shows that the bank does not have any holes since the fitting factor (FF) values are greater than the MM values used to construct each of the banks.

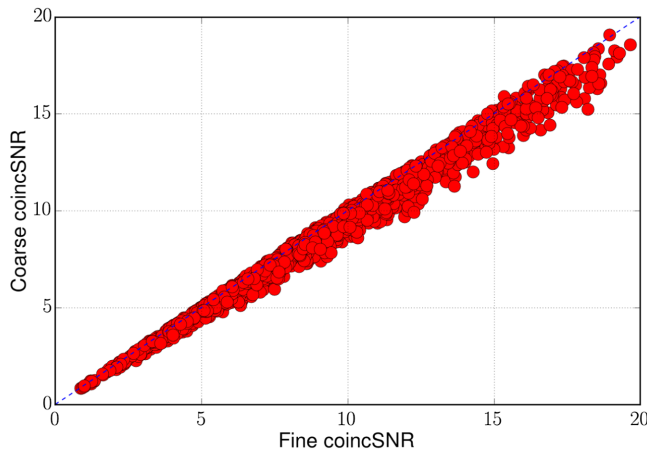


FIG. 2. The coincident SNR for the fine search is plotted against and coarse search coincident SNR with red circles. The figure shows how much SNR is lost in stage I in which a lower sampling rate and a coarse bank are used.

reduced sampling frequency as mentioned above. Then, in Fig. 3, we plot the maximum possible loss in SNR on the horizontal axis and on the vertical axis we plot the match between the templates that correspond to maximum SNR (trigger templates) in stage I and II. The match between stage I and stage II trigger templates corresponding to the same injection tells us how large the nhbd of the particular template of coarse bank should be. The figure tells us that even for zero-noise case, we have to consider nhbd large enough with a match as low as 85% in order to recover the full SNR in stage II of the hierarchy. With noisy data, it is prudent to choose a lower value of the match—we choose this value to be 75%.

Now we come to the choice of thresholds. The maximum SNR loss is 85% in stage I. This guides our choice of stage I SNR thresholds (with reference to the flat search): (i) for

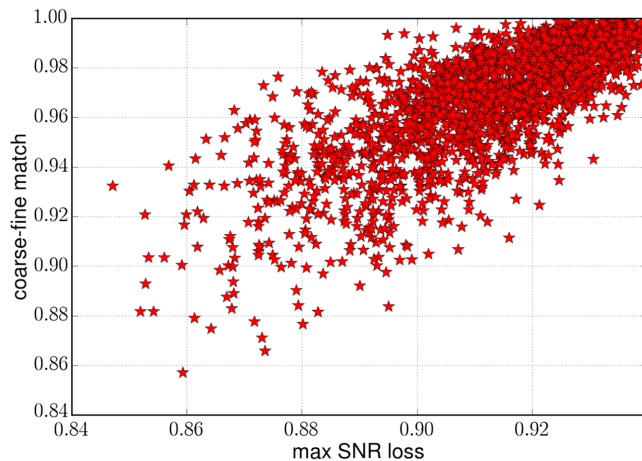


FIG. 3. The match between the stage I (coarse) trigger template and stage II (fine) trigger template is plotted versus the maximum loss of the SNR (normalized) in stage I. Each injection is represented by a red star.

individual detectors ($\rho_{\text{single},I}$ and (ii) for coincidence $\rho_{\text{coinc},I}$). We keep these thresholds at about 90% of the respective thresholds for the flat search as discussed below. For example, we may choose $\rho_{\text{single,flat}} = 5$ and $\rho_{\text{coinc,flat}} = 8$ for the flat search and $\rho_{\text{single},I} = 4.5$ for stage I and $\rho_{\text{coinc},I} = 7.2$ for stage II of the hierarchical search.

Figure 4 shows the recovered stage I SNRs for injections in each of the detectors H1 and L1. Variation in SNR is observed because the two detectors have different orientations (thus different antenna patterns) and therefore different responses to the same signal coming from a given location in the sky. The SNR variation in a single detector also shows that there could be signals that the hierarchical search may miss but which the flat search may detect. If the signal is barely above the single detector SNR threshold ($\rho_{\text{single,flat}}$) in one of the detectors, then, even with a reduced SNR threshold ($\rho_{\text{single},I}$) we may miss it as the noise may not trigger the correct template in stage I in one of the detectors. Also, we may not detect few borderline injections because of SNR loss of more than 90% as is seen in Fig. 3. Vetoes can further aggravate the problem by pushing down the newSNR because the trigger templates in the coarse bank can have a larger mismatch. Because of this, we would not have a coincident trigger to go to stage II even if the signal was otherwise loud enough. But to be fair this could also happen in a coincident flat search depending upon the choices of the threshold. For the current choices, the flat search may miss fewer signals compared the hierarchical search. So this means that there is a chance that the hierarchical search may miss out some signals which have a low SNR in one of the detectors as compared to the flat search. This is demonstrated in Fig. 5. If we however consider more than two detectors and a coincidence analysis, we believe this is unlikely to happen.




FIG. 4. The figure shows the recovered SNR for each injection as observed in detectors H1 and L1. The large variation in recovered SNRs is due to the different orientations of the detectors which have different antenna pattern functions. The dashed line is the line of equal SNR.

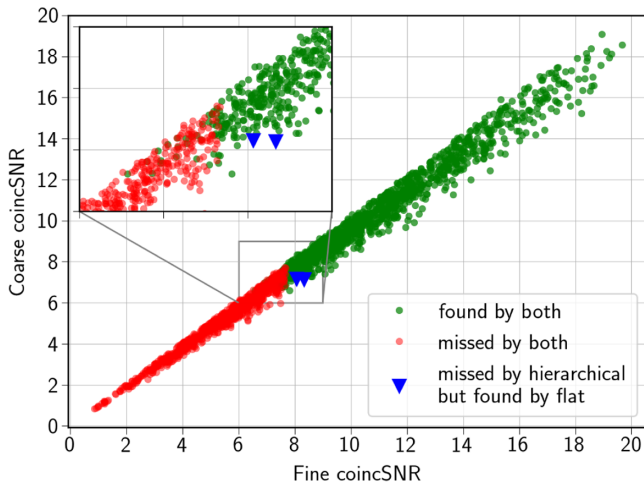


FIG. 5. Injections found and missed in the flat and hierarchical searches for the zero-noise case. The red dots show injections missed by both searches while the green ones show those that are found by the both. The blue inverted triangles correspond to the injections which are likely to be missed by hierarchical search but recovered by the flat search with the current configuration.

This simplified case guides our choice of thresholds and nbhds for the full fledged hierarchical search strategy.

2. Formalism

Here we describe the 2-stage hierarchical search pipeline for coincident detection with two detectors. The full search is illustrated by the flowchart in Fig. 6. We start by creating stochastic coarse and fine banks for the intrinsic parameters (these are masses and aligned spins) which have MM of 0.9 and 0.97 respectively. For template banks we use the harmonic PSD which is the harmonic mean of PSDs of H1 and L1 detectors during O1. We use the same PSD to generate simulated Gaussian noise for both the H1 and L1 detectors. These data are then divided into smaller chunks of 4096 sec each for estimating the local PSD which is required for the matched filtering computations. The matched filtering is done with data segments of duration 256 sec and with 128 sec overlap with the previous segment. This overlap is needed because we must discard data from both the ends of a data segment due to the circularity property of the FFT algorithm and also get rid of other numerical artefacts [41,50]. We therefore actually search only 128 sec of data in one matched filter computation.

The hierarchical search begins with the first stage, where data are sampled at a lower rate and with a coarse bank for each detector. Single detector events are recorded if the statistic crosses a predetermined threshold for each detector. The statistic employed here is the power χ^2 re-weighted new-SNR [33,41]. In the first stage the threshold is lower than the second stage threshold. We then compare parameters of the triggers from each detector and select only those

triggers whose parameters match—these are the coincident triggers for the first stage. We then follow up these candidate triggers with a fine search using a fine subbank constructed around each coincident trigger.

To create the fine subbank from the zero-lag coincident 1st stage triggers, we take the union of the fine bank nbhds (which have been precomputed) around 1st stage templates which have been triggered. Then we repeat the search with a higher sampling rate and with time segment specific fine subbank and stage II threshold. We then collect the individual detector triggers from each of the detectors H1 and L1. From these we select the coincident triggers for stage II. Coincident triggers crossing the stage II threshold give us the foreground triggers.

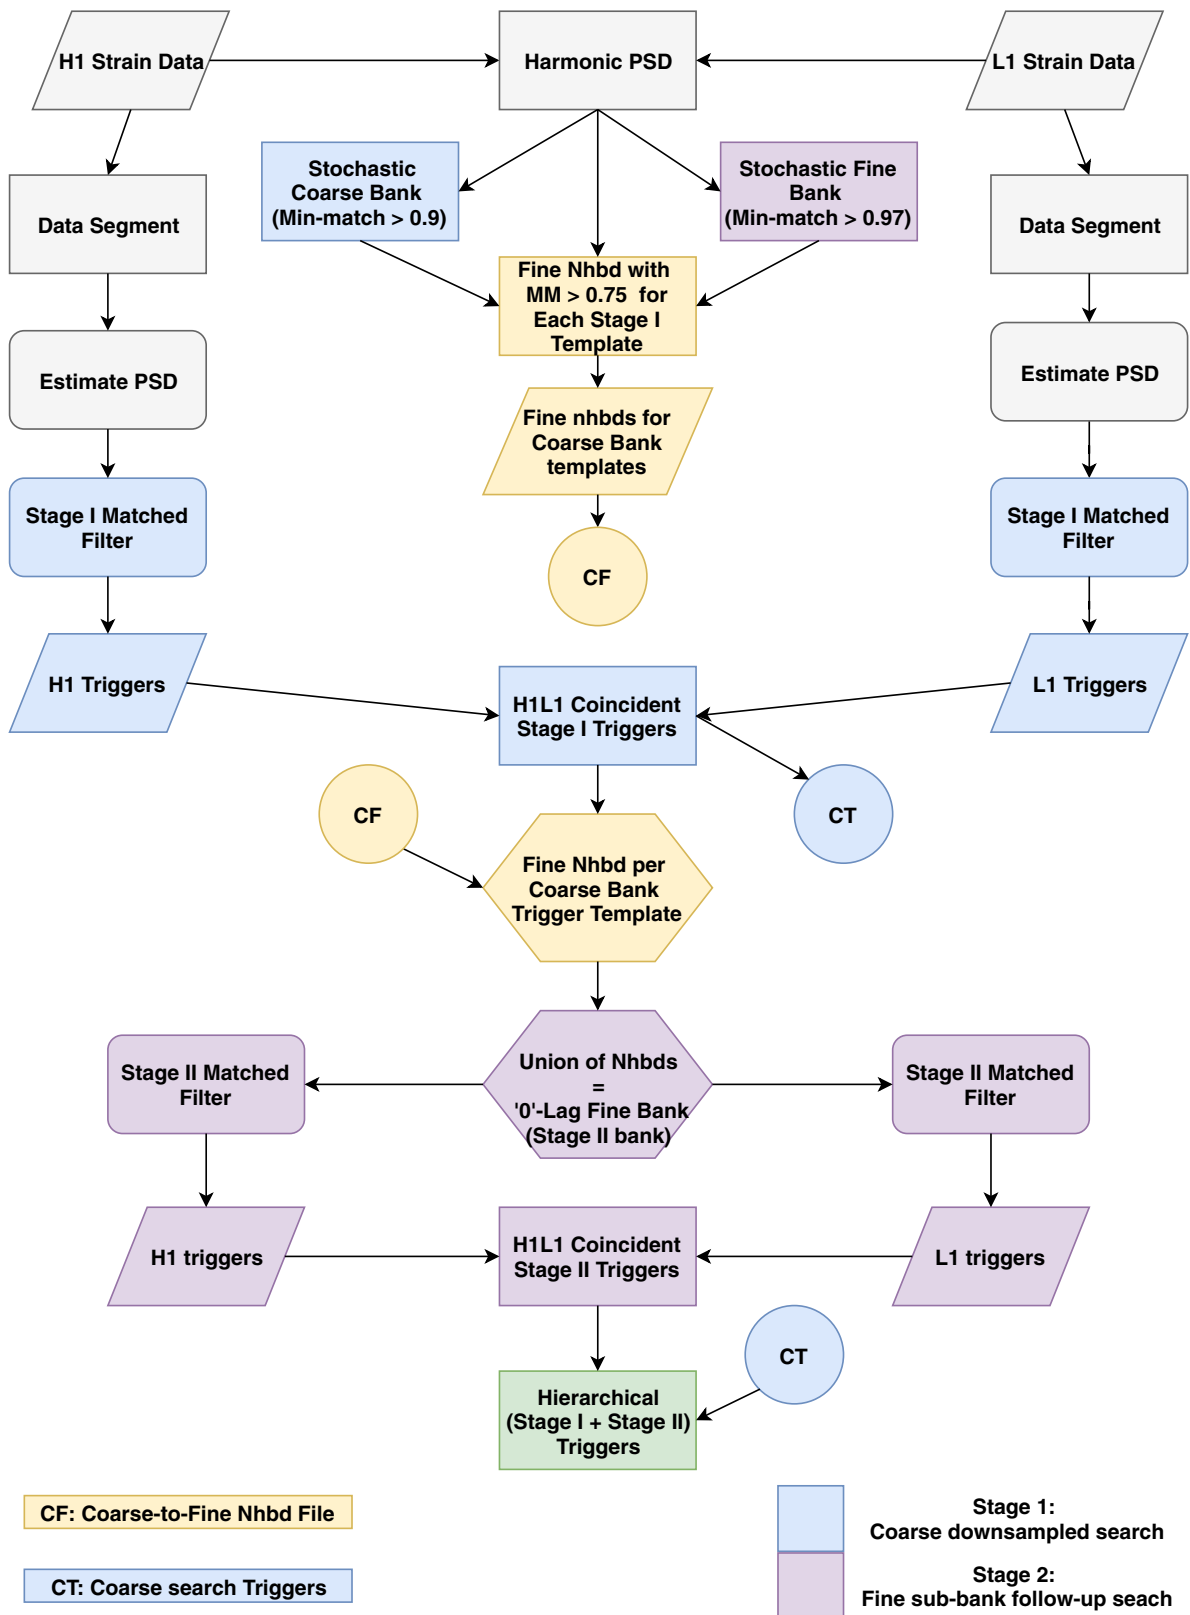
The noise background is obtained as follows: We use combined (coarse + fine subbank) individual detector triggers in the time-slides to get the hierarchical search noise background. This background is usually lower than that for the flat search [41] background as it comes from significantly reduced number of templates. The fine subbank contributes negligibly to the noise background because it has much fewer templates. This procedure for estimation of the noise background for the hierarchical search differs from that of the flat search.

In principle, the noise background for the hierarchical search can be used to assign significance to the foreground triggers using the scaling as will be described later. But since we use only zero-lag subbank in the second stage, it may show some bias. We have therefore used the usual flat search background to assign significance and estimate the sensitive distance in this work. With the hierarchical background, the sensitive distance for the hierarchical search will slightly improve than what has been shown in Sec. IV. Further analysis is required to improve the estimate of the noise background and investigate whether there is a generic way to obtain in a robust way an equivalent background. But our preliminary study with real data in Sec. V shows that it is possible.

We can of course reproduce a background for the hierarchical search analogous to the flat search. We must then consider 1st stage triggers arising from also nonzero time lags and then construct fine subbanks around these triggers. But then the union of these subbanks will be almost as large as the fine bank, and thus we stand to lose any computational advantage that we may have otherwise obtained.

3. Parameters used in the hierarchical search

In this subsection we provide the detailed description of the parameters used in our hierarchical search. We consider the same ranges of masses and spin parameters as were employed in the search during the first aLIGO science run (O1). We also use the same waveforms as those employed in O1. For aligned-spin GW signals, our set of intrinsic parameters are component masses and spins along orbital



Hierarchical Search: Flow Chart

FIG. 6. Flow chart describing the set up of the current 2-stage hierarchical search.

angular momentum. For the individual masses, the heavier mass is in the range $(1, 100)M_\odot$ and secondary mass is in the range $(1, 50)M_\odot$. In case of neutron stars we have taken the masses to be in the range $(1, 3)M_\odot$ and spin components (dimensionless) along the orbital angular momentum from -0.4 to 0.4 . Black holes have spin components ranging from -0.9895 to 0.9895 . The fine and coarse banks are generated using stochastic template placement algorithms (überbanks) with mismatch of 3% and 10% respectively. The choice of 10% mismatch for the coarse bank is somewhat arbitrary. With these numbers and with O1 harmonic power spectral density (PSD), we obtained ~ 60000 templates for the coarse bank and ~ 250000 templates for the fine bank. One observes that the fine bank is little over 4 times larger than the coarse bank. For search templates, we have used `TaylorF2` approximant for total mass less than $4M_\odot$ and `SEOBNRv2_ROM_DoubleSpin` for the rest of the parameter space.

Even though the choice of the $MM = 0.9$ may look somewhat arbitrary, we have tried other values of MM , for example, 0.8 (as employed in the previous hierarchical searches) and 0.85 . We found that the overall loss in the SNR is unacceptable because of loss in sensitivity. Therefore, we fix MM at 0.9 which gives a coarse bank with about quarter the number templates as compared to the usual fine bank with $MM = 0.97$.

We use the harmonic PSD of O1 to generate template banks and for simulating data with the lower cut-off frequency set at 30 Hz. For the first stage in the search, we sample the data at a reduced rate of 512 Hz while for the second stage, we sample the data at 4096 Hz. Because of the reduced sampling rate of 512 Hz in the first stage, we must cut off the signal at 256 Hz. However, even after applying this upper cutoff, we recover more than 90% of the signal SNR, for all the 10855 nonprecessing injections. We also ensure that the banks do not have ‘‘holes.’’

As mentioned earlier, 5 days of simulated coincident data for the two LIGO detectors H1 and L1 are used assuming both of them have the harmonic PSD of O1 run. We inject more than 10855 nonprecessing CBC signals in the data with the parameter ranges as mentioned earlier. The injections were uniform in volume, orbital inclination, and coalescence phases. The injections were distributed as follows: ~ 2171 double neutron star (DNS), 4342 neutron star-black hole (NSBH), and 4342 binary black hole (BBH). Neutron star masses were in the range $(1, 3)M_\odot$. Further, the injections were uniformly distributed in the total mass. All the injections were with aligned-spin. The optimal SNRs for the injections were in the range $(8, 30)$. For DNS injections, we have used `SpinTaylorT4` approximant for injection and `IMRPhenomD` and `SEOBNRv2` for NSBH and BBH injections, respectively.

Apart from the above, we injected 8684 precessing signals with total mass in the range of 5 to $150M_\odot$ with

the dominant mass ranging from 4 to $100M_\odot$. For the precessing injections, we have used the `IMRPhenomPv2` approximant.

We now go on to stage I search and describe the corresponding triggers with their associated fine subbanks which will be used in stage II.

4. Stage I triggers

The goal of the first stage is to obtain candidate triggers which will then be followed up in the second stage of the search. To obtain these, we need to decide a threshold on the detection statistics, which is the chi-square weighted newSNR [33] for the single detector statistic and coincident newSNR, i.e., newSNR of single detectors added in quadrature for the pair of LIGO detectors. This statistic is used for coincident triggers for both the stages and also for the flat search. We decide on the individual detector thresholds $\rho_{\text{single,flat}} = 5.0 = \rho_{\text{single,II}}$ where $\rho_{\text{single,flat}}$ is the threshold for the flat search and $\rho_{\text{single,II}}$ is the threshold for the second stage search. We decide to keep single detector newSNR ($\rho_{\text{single,I}}$) to be 4.5 which is 90% of $\rho_{\text{single,flat}}$. We have chosen these values because we expect SNR loss to be less than 10%. The amplitude of the GW signal in frequency domain scales as $f^{-7/6}$ and the SNR in the first stage is reduced both because we are employing a coarse template bank and an upper frequency cutoff (lowered sampling rate), which we denote by ρ_{reduced} . The results are as follows:

$$\rho_{\text{reduced}} = MM_I \frac{\rho(f_l, f_{u,I})}{\rho(f_l, f_{u,II})} \quad (3.1)$$

where we have defined ρ as:

$$\rho(f_l, f_u) = \int_{f_l}^{f_u} \frac{f^{-7/3}}{S_n(f)} df. \quad (3.2)$$

The recovered SNR relative to the full injected SNR due to upper frequency cutoff (stage I of the hierarchical search) is shown in Fig. 7 as a function of the mass parameters. It can be seen that the least recovered SNR is $\sim 94\%$ which corresponds to a loss of 6% in the worst case scenario. For the values we have chosen, we get $\rho_{\text{reduced}} > 88\%$ for MM_I of 0.9 for the 1st stage bank (coarse bank). But if we use the factor MM_I/MM_{II} instead of MM_I in equation (3.1), we get $\rho_{\text{reduced}} > 91\%$.

The procedure as described in Sec. III B 2 has been followed.

5. Stage II search

After obtaining the coincident triggers obtained in stage I, we proceed to stage II. Here we construct a fine bank in a small neighborhood around each stage I trigger. The neighborhood for the subbank is so chosen that the




FIG. 7. The figure shows the reduced SNR because of the lowered sampling rate, over the parameter space under consideration. The maximum SNR loss is $\sim 6\%$.

templates in the fine bank have a match more than 0.75 with the trigger template from the coarse bank. A smaller value than 0.85 is chosen for the match in order to compensate for the noise effects and other factors. These fine subbanks are precalculated for each coarse template. Figure 8 shows a typical stage I trigger template (red star) with its associated fine subbank of templates (blue dots) with MM > 0.75. The template masses are given in terms of two equivalent mass parametrizations. Also, plots in the Fig. 9 show the number of templates in the fine bank in the nhbd of each of the coarse bank templates. The left plot shows the templates with masses parametrised in terms of chirp times and the histogram on the right shows that the nhbd with MM > 0.75 few tens to 100 templates per nhbd considered independently.

We first obtain coincident 1st stage triggers for each data segment and then create a fine subbank corresponding to

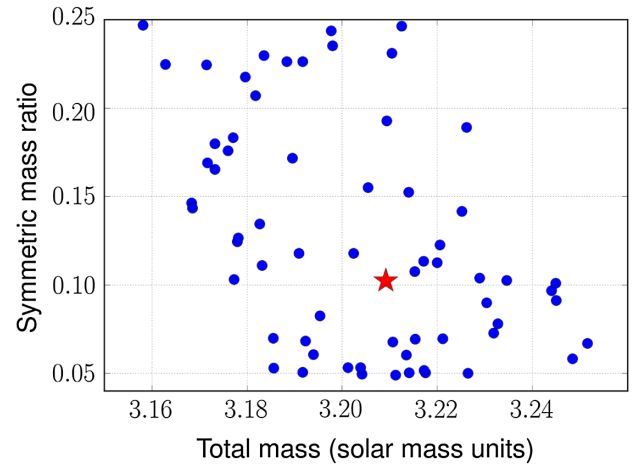



FIG. 8. A typical fine-bank neighborhood around a coarse template in terms of various mass parametrizations. The red star shows one such coarse bank template, while the blue dots depict the templates in the fine bank neighborhood with MM > 0.75. There are 65 templates in this fine-bank neighborhood.




FIG. 9. The figure shows the number of templates in different fine-bank neighborhoods as one scans the parameter space with the coarse bank. The left plot denotes number of templates in the fine neighborhood in terms of chirp times. On the right, we plot a histogram showing the number of coarse templates having different sizes of fine subbanks. The figure gives an idea of the differing sizes of the fine subbanks.

that segment. Then we take the union over all the data segments of all the fine subbanks. This unified subbank depends upon the single detector statistic used, threshold for that statistics and both the coarse and fine banks used. Then for each data segment and the corresponding fine subbank, we perform a search with full sampling rate of 4096 Hz as used in single stage flat search. Then we follow the same procedure as in the first stage of again collecting single detector triggers from the fine subbank and obtaining the second stage coincident triggers by matching parameters. We then cluster these coarse and fine subbank triggers together and obtain the final triggers. Now these final triggers need to be compared with the noise background for estimating their statistical significance. For this second stage, we use the $\rho_{\text{single,II}}$ which is the same as that for the flat search. To estimate the noise background, we have used the single detector trigger time slides but with some caveats which we will discuss in the next section. We now present our results.

IV. COMPARISON WITH THE FLAT SEARCH

In this section, we compare the results of the hierarchical search with the flat search. For this analysis we assume stationary Gaussian noise. We begin by comparing the noise background and noise foreground without injections. For each individual detector, on an average, we found ~ 53 triggers per second from the flat search with $\rho_{\text{single,flat}} = 5.5$ and $\sim 7\text{--}8$ triggers per second from the coarse search with $\rho_{\text{single,I}} = 5.5$. But for $\rho_{\text{single,I}} = 5.0$, we obtained ~ 111 triggers per second. For both the flat and stage I (of the hierarchical) searches we use their respective full banks and obtain the corresponding triggers. We observe less number of triggers for stage I of hierarchical search although we maintain the same threshold. This is because of the reduced number of the templates in the coarse bank and also because of the reduced sampling rate. We expect that with only a quarter of the templates and 1/8th of the sample points, we can at best get a factor of ~ 32 reduction in the computational cost as compared to that of the flat search. This is because the main cost of the search comes from matched filter computations which in turn is due to FFT used in the data analysis. The cost of a FFT scales as $N \log N$ where N is the number of data points in the data segment. First, we have about quarter the number of matched filter computations because of the reduced number of the templates in the coarse bank and second, the cost of each FFT goes down by the factor of about 8 because N is reduced by this factor due to lower sampling rate. However, the computational gain or speed-up is much less than this number ~ 32 , because the candidate triggers from the stage I need to be followed up with the stage II fine search. The computational cost incurred in the second stage depends on the number of stage I candidate triggers which in turn depends on stage I threshold $\rho_{\text{single,I}}$, and also on the size of

the stage II fine subbank. The size of the fine subbank depends upon the choice of the relevant neighborhood for each 1st stage template that is triggered. For example, choosing a very low $\rho_{\text{single,I}} \sim 3.5$ will fetch a huge number of coincident triggers in the first stage from the entire coarse bank, which will lead to searching over almost all of the fine bank in the stage II taking away most of the computational benefits of the hierarchical search. While increasing the $\rho_{\text{single,I}}$ will reduce the candidate triggers in the I stage, but will lead to the increase in the size of the coarse bank and hence the matched filtering cost at stage I. Therefore a compromise must be sought if the hierarchical strategy has to succeed.

With the help of time slides, we compute the noise background for more than 260 years of coincident data. This is done for the flat search in the usual way and also for both the stages of the hierarchical search which uses the coarse bank and the zero-lag fine subbanks. Foreground is computed for 5 days of the coincident data for the flat and the 2-stage hierarchical search. Note that the thresholds and other parameters (clustering etc.) used for the background, foreground, and injection recovery in stage II of the hierarchical search are the same as that of the flat search. Figure 10 shows backgrounds and foregrounds for the flat (red), stage I (green), and stage II (blue) in terms of cumulative number of coincident events per year with newSNR plotted on the horizontal axis. From the figure, we make the following observations about the noise background:

- (i) Stage I (green) backgrounds and foregrounds are lower by more than an order of magnitude than that of the flat (red) search.
- (ii) Stage I background generated from the fine subbanks is negligible as compared to the Stage I background generated by the coarse bank—something like two orders of difference in magnitude.




FIG. 10. Full search background and foreground event rates (per year). The foreground data is of ~ 5 days and background amounts to more than 140 years after time slides.

- (iii) Since the hierarchical search background is the union of stage I and stage II background, the hierarchical search background is essentially determined by the stage I background generated by the coarse bank.
- (iv) When the background for stage II (blue) is calculated using time slides, the single detector triggers from two different data segments are likely to have very few or no common templates in the stage II fine subbank. Hence stage II background does not match the foreground. This may lead to some bias in the total hierarchical background but it should be negligible compared to stage I

In practice, we may use the stage I background to assign statistical significance to the triggers and it may be simply scaled up to obtain the flat search background for the simulated data. This scaling argument follows later.

In principle for the hierarchical search we can estimate the background using time slides as in the flat search but then we lose out on the computational benefits—we find that even with 2000 time slides, the full fine bank is covered.

However, stage II foreground contributes comparatively much more to the overall foreground. This is because we use fine subbanks in stage II constructed using only “0”-lag (foreground) coincident triggers from stage I. All templates in the subbank can contribute to the foreground evaluation which give better chances for noise coincidences.

This utilization of the full fine bank in stage II implies that when we do time slides with the nonzero lag triggers from stage I, we recover the same background as the flat search. Hence, at least with the simulated colored detector noise, we came up with the idea of scaling the background obtained from the coarse bank to recover the flat search background. Interestingly, if we scale the stage I background by the speed-up factor, we recover the flat search background over the coincident new-SNR threshold of 8. This can be seen in Fig. 11. Moreover, the same scaling factor does match the noise-only foreground of stage I of the hierarchy with that of the flat search. A little excess in the low new-SNR region between 7.5 to 8 is because of the reduced single detector and coincident thresholds for the stage I. The scaling of the stage I to the flat background is the same as the speed-up factor because it is exactly the ratio of the number of independent random variables that the matched filtering operation produces in each case. This speed-up factor is explained in detail later.

Thus, for simulated data, we are able to use the scaling argument to obtain the equivalent flat search background from the stage I background. From this equivalent background we are then able to assign the correct statistical significance to the foreground triggers.

We now demonstrate the above argument by using 100 time slides with nonzero time lags corresponding to 2.6 years of background. This is shown in Fig. 12 by the




FIG. 11. Full search background and foreground event rates (per year). The first coarse stage rates are scaled by the factor of ~ 20 to match all the events. The scale factor to get back the same event rates is same as the speed-up factor for simulated data.

magenta line. If we scale this hierarchical background we get a background equivalent to ~ 260 years. We observe that the scaled hierarchical search background matches well with the flat search and the scaled stage I background (except for the low SNR region). However, even for just 100 nonzero lag time slides, we need to perform the matched filtering operation for more than 30 000 fine bank templates in stage II for each data segment. With real data, this exercise needs to be carried out with even more care and caution as will be described later in Sec. V.

We now investigate the recovery of injected CBC signals by both types of searches, flat and hierarchical.

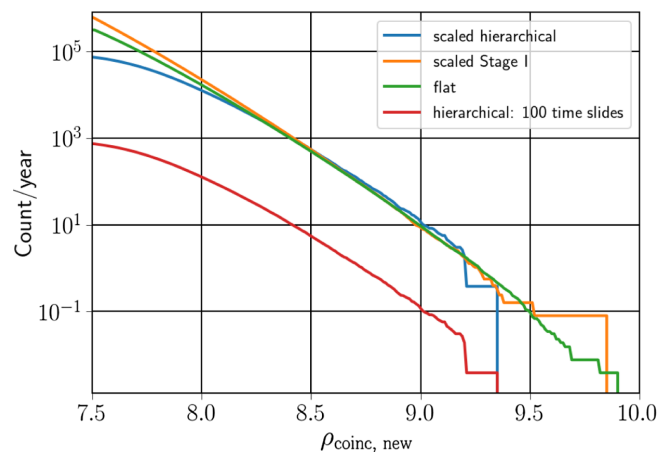


FIG. 12. Figure shows the flat search background (red), stage I background scaled by the speed-up factor (green) and hierarchical search (flat search equivalent) background calculated using 100 time slides (blue) and scaled with the constant factor to get a background equivalent to a duration ~ 141 years. It can be seen that both the scaled backgrounds match well with the flat search background. With 100 time slides one needs to match filter > 30000 fine bank templates per data segment.

As discussed in Sec. III B, we inject 10000 aligned spin CBC (DNS, NSBH, and BBH) signals. In addition, we inject more than 8000 precessing CBC (NSBH and BBH) ones. Figure 13 shows the newSNR as found in stage I of the hierarchical search and as found in the flat search. The flat search newSNRs are the best values that stage II of the hierarchical search can recover to perform as well as the flat search. Figure 14 shows all the above mentioned injections which are missed or found by both searches. For each of the subplots, we have plotted injection with flat coincident SNR against the coincident newSNR for the stage I of the hierarchy. The red dots show the injections missed by the both flat and hierarchical search while green dots show the injections found by both the searches. The blue triangles denote the injections missed only by the hierarchical search but found by the flat search. For all the 3 aligned spin cases, the injections missed by the hierarchical search are $\sim 2\%$ of those found by the flat search. For the precessing case, the hierarchical search loses $\sim 6\%$ of the injections compared to those recovered by the flat search. This somewhat larger loss in the hierarchical case may be because of the higher dimensionality of the parameter space required to describe precessing systems—the injections have more “room” to distribute themselves. More specifically, this penalizes the coarser search more because the ratio of volumes of the fine to the coarse neighborhoods is smaller for the precessing case because of the higher number of dimensions. Figures 15 and 16 show the sensitivity distance for both searches, hierarchical and flat, with a varying coincident newSNR




FIG. 13. The figure shows coincident recovered newSNRs for all the injections. It shows how much SNR is lost in stage I with the lower sampling rate and the coarse bank. Stage II may recover the SNR with the fine subbank and a full sampling rate.

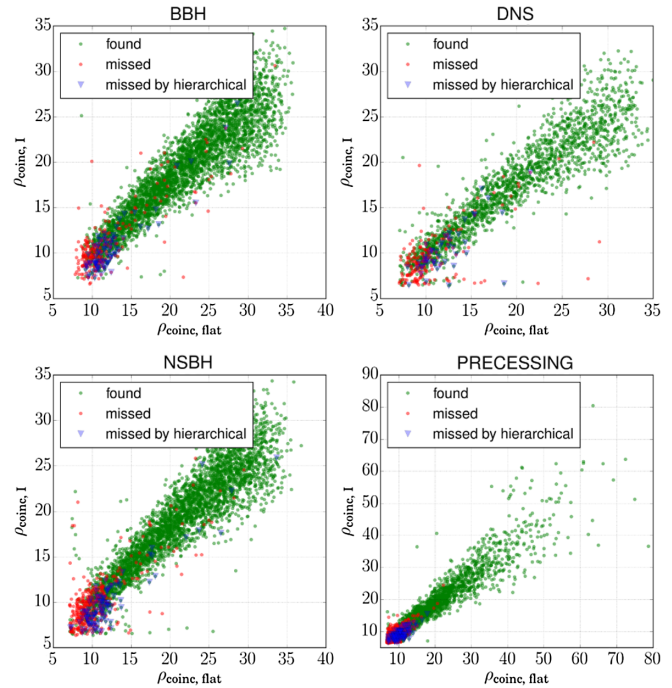


FIG. 14. The plot shows the recovered coincident newSNR of signals missed (red dots) and found (green dots) by the hierarchical search. The blue triangles show the injections missed by the hierarchical search but found by the flat search. These are $\sim 2\%$ of the total injections.

threshold. The newSNR threshold corresponds to a false alarm rate as can be seen from Fig. 10. We see that both hierarchical and flat searches have almost similar sensitivity distances as a function newSNR as shown in Fig. 17. This implies that both the searches perform almost equally well. For the calculation of the sensitivity distance, we have used all the CBC injections. It can be seen that only the BNS search has slightly lower sensitivity for the hierarchical search than for the flat search. This is expected as we are using truncated waveforms with much lower MM and BNS signals are of long duration and contribute significantly to the SNR at higher frequencies which means that the fractional loss in SNR is more. The lower recovery of signals is due to reduced stage II SNR compared to flat search for a few BNS sources. This is because the SNR of these sources in the one of the detectors is slightly lower so that some false trigger templates are contributed from the stage I. Thus we see that, the hierarchical search recovers almost all the injections as those recovered by the flat search. There is a slight advantage to the hierarchical search over the flat search, because we can choose a slightly lower detection threshold with hierarchical search for the same false alarm rate. We have not addressed this question here. Thus for the hierarchical procedure we have proposed, we conclude that both the searches have almost similar distance sensitivity for the injected set of signals. Next we consider the computational cost of each kind of search.

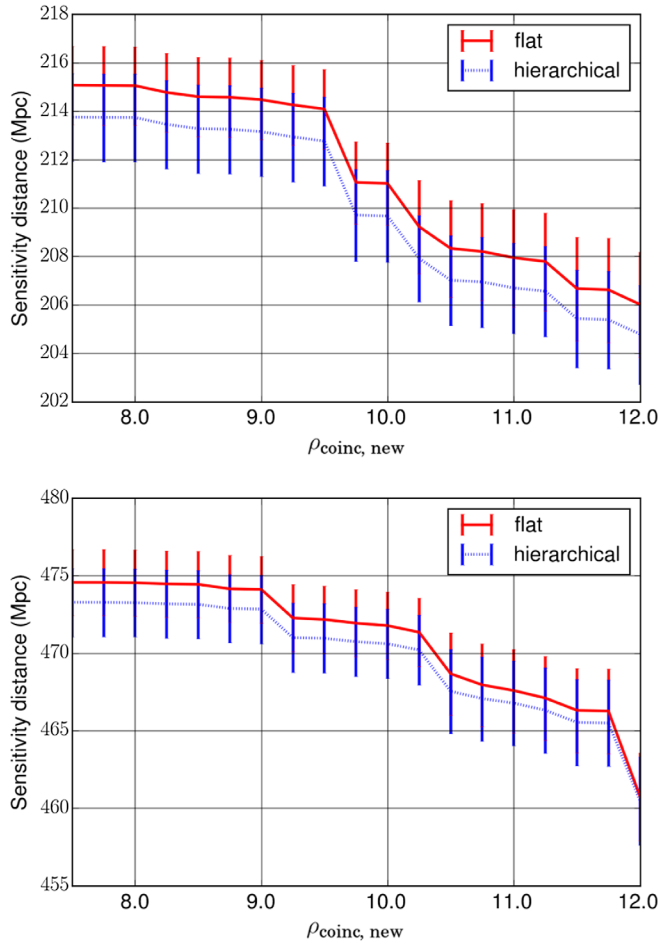


FIG. 15. Comparison of flat and hierarchical searches for distance sensitivities. Top: Aligned-spin BNS injections; bottom: aligned-spin NSBH injections.

We now look more closely at the computational costs and the computational gain from the hierarchical search. We also explain how the same is related to the background estimation. For a data segment of length 256 sec, we have few hundreds of coincident first stage triggers (no additional coincident threshold is applied). On an average, we have 40-90 templates of the fine bank in the neighborhood of the each coarse trigger template as can be seen from Fig. 8. We then obtain a second stage fine subbank which is the union of these neighborhoods. This subbank has 1000–4000 templates on an average per data segment. We now compute the average number of floating point operations (FLO) per data segment. We do 60 000 MF calculations at 512 Hz sampling rate in the first stage and at most 4000 MF calculations at 4096 Hz sampling rate in the second stage. Each MF computation involves a complex FFT corresponding to the two phases of the waveform. On the other hand, in the flat search, we do 250 000 MF calculations at the sampling rate of 4096 Hz. Each MF calculation uses discrete Fourier transform (DFT). Each DFT with N data points requires $\alpha N \log N$ FLO, where $\alpha \sim 3$ for a real DFT




FIG. 16. Comparison of flat and hierarchical searches for distance sensitivities: Top: aligned-spin BBH injections; bottom: precessing injections.




FIG. 17. The plot shows the relative sensitivities of volume and distance (in %) of the hierarchical search compared with the flat search. Solid lines show relative volume sensitivity and the dash-dotted lines show relative distance sensitivity.

and double this number for a complex FFT and depends on the algorithm used. Thus, roughly, discarding the α factor which is common to both the searches, the flat search requires 250×4.096 mega-FLO while the hierarchical search strategy adopted here, requires $60 \times 0.512 + 4 \times 4.096$ mega-FLO. Thus one obtains a computational gain of ~ 20 in the matched filter computations. Other data conditioning require the same computation.

We look back at the estimation of the noise background for the hierarchical search. We argue that the hierarchical background is just scaled down from the flat search background roughly by the speed up factor, which in this case is ~ 20 . The noise background arises from the number of triggers which essentially stem from the number of independent Gaussian random variables in the matched filter output. The Gaussian variables in the matched filter output are however correlated. For the flat search we get roughly $256 \times 250000 \times 4096$ data points (Gaussian variables not necessarily independent) per segment. But for the hierarchical search we must consider both stage I and stage II data points. For the hierarchical search we have $256 \times 60000 \times 512 + 256 \times 4000 \times 4096$ data points per segment. We may expect the effect of correlation between Gaussian variables to be about the same in both the flat case and the hierarchical case. Ignoring the effect correlations and except for the slowly varying factor of $\log N$, the ratio of independent Gaussian variables in the two situations is roughly the same as the ratio of matched filtering operations required for each of the searches. This is in fact the speed up factor. This is evident from Figs. 10 and 11. However real data contains non-Gaussian artefacts and we basically sample the tail of the noise distribution (rare events) to estimate the background. Therefore, this scaling exercise needs to be carried out carefully in order to obtain the correct scaling. The scaling may depend upon template duration as very short duration templates are more susceptible to the glitches and artefacts in the real data.

We now make a few remarks. First of all, the non-preprocessing injections we used are in the H1-L1 coincident SNR range 8 to 30 and our preprocessing injections are linearly distributed in distance. Second, we get our noise background for the hierarchical search by scaling up essentially the background from the first stage of the hierarchy. As explained before, obtaining the full background equivalent to the flat background will compromise computational advantage that is expected to be gained from the hierarchical search.

Our analysis of simulation data shows that we may be able to employ the same procedure on real data. We may use the stage I background to infer the significance of detection, after further investigations with real data. This may not be exactly equivalent to the flat search background, but it can be considered as a separate hierarchical background. A few hundred time-slides can be performed to get the noise background on a shorter duration of data

and then it could be scaled up to obtain an estimate of the full flat search background. The background so obtained could be used to decide on the statistical significance of triggers. We propose to address this issue of the background and second, also tune the pipeline for injection recovery per mass bin with the real data in a future work.

V. SEARCHES WITH REAL DATA

In this section we demonstrate how our two stage hierarchical search works on real data. For this purpose we have used four segments of data, each of 4096 sec duration, from the first observing run O1 of the twin LIGO detectors [51]. The four segments are chosen such that each one contains one detected event, including the trigger ‘‘GW151012’’ [1] whose status was escalated to a true GW event.

We employ identical template banks to those for simulated data described in Sec. IV. Also, we have used sampling rate of 512 Hz and 4096 Hz for stage I and stage II respectively. We decided to keep roughly the same number of stage I triggers as for the simulated data case in order to have ready comparison in speed up factors. Further, we take the clustering window to be of 1 sec over a template as is used in the real data search [41]. Also we increase the single detector threshold slightly; $\rho_{\text{single,I}} = 5.0$ and $\rho_{\text{single,flat}} = 5.5$. We use the newSNR statistic (SNR weighted with power and sine-Gaussian based χ^2 vetoes) [23,52,53] which makes the data behave as close to Gaussian as possible. We decided not to use the phase-time statistic [42] here, as it requires specific tuning which is outside the scope of this paper. We further test our method by injecting 2000 CBC signals. For the above mentioned thresholds, we obtain 100–105 triggers per second for stage I hierarchical search and 80–90 triggers

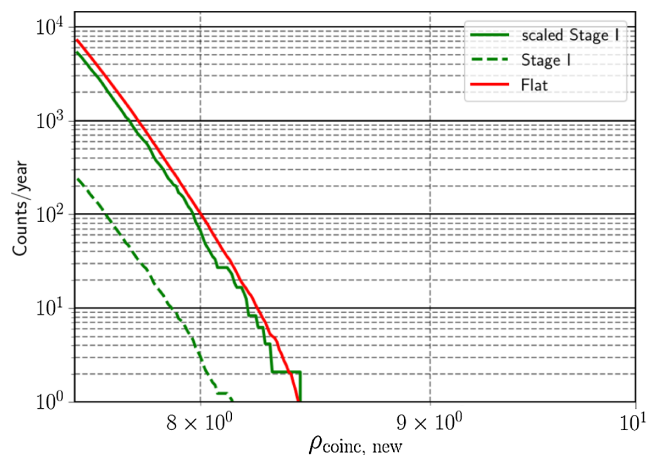


FIG. 18. The plot shows the noise only backgrounds: the solid red curve is corresponds to the flat search, while the green dashed curve corresponds to the stage I of the hierarchical search. The solid green curve is the scaled background by the speed-up factor of ≈ 20 .

TABLE I. The table shows the recovered values of the detection statistic in our searches using the data chunk containing each of the GW event. IFAR for all the events in more than 10 yrs which is maximum for the data we used.

Event	Stage I SNR	Flat/Stage II SNR
LVT151012	8.1	8.9
GW150914	16.23	19.47
GW151226	7.9	9.1
GW170104	8.1	9.2

per second for the flat search. With these choices, we use 4000 to 6000 templates in the stage II of the hierarchical search per segment for the follow-up of the candidate triggers. These numbers are about 150% of those corresponding to simulated data. This is due to the higher trigger rate in the real data. However, this does not significantly reduce the speed up factor. We believe further tuning can be performed for real data, with optimized statistics and thresholds, to obtain even better results.




FIG. 19. Distance sensitivities for hierarchical and flat searches. The top figure shows that the distance sensitivity is just above 40 Mpc for BNS while the bottom figure shows that it is more than 80 Mpc for NSBH for a coincident SNR of 8.

We present the results of this *test application* on real data. Figure 18 shows background from flat search and stage I of hierarchical search. For simulated data it was observed that we get the background pertaining to the flat search, if we scale up the stage I background by the speed up factor of ≈ 20 . The same scaling seems to work in the case of real data. Thus, we can, in principle, use the scaled stage I background (solid red line) to assign the correct significance to the foreground triggers which would be equivalent to the flat search. This can be achieved by fixing this scale factor by doing a flat search run with a small fraction of real data, like some pipelines do, to fix the background scale [1]. In Table I we mention the SNRs obtained for the four events in stage I and stage II of the hierarchical search and the flat search.

We now turn to the sensitivity of the searches with the 2000 injections which are uniformly distributed with distances in the range 30–750 Mpc and other parameters

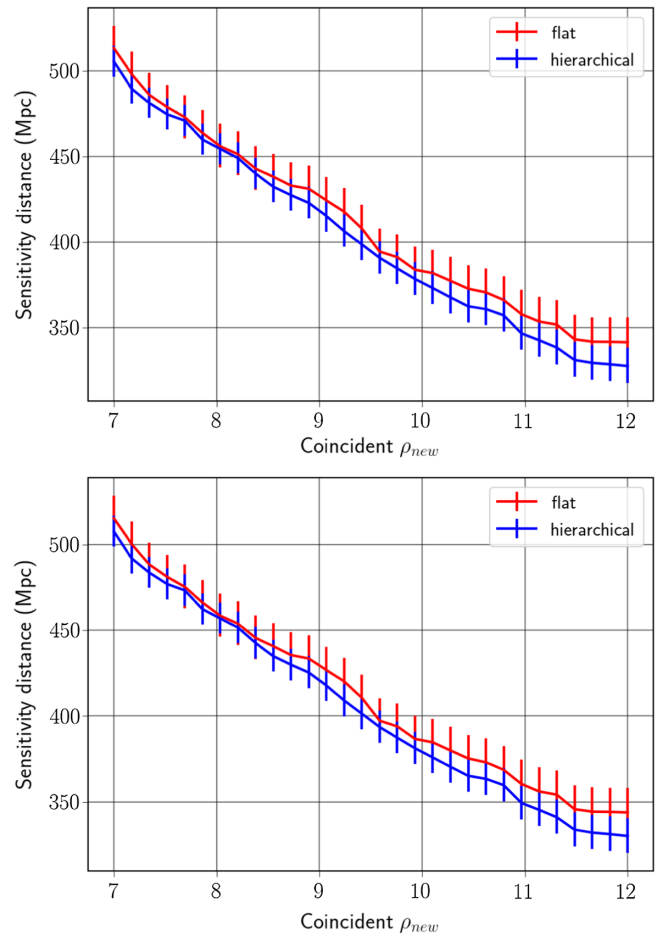


FIG. 20. Distance sensitivities shown for hierarchical and flat searches. The upper plot is for BBH and the bottom plot shows when all BNS, NSBH, and BBH are taken together. The search sensitivities are just above 400 Mpc for coincident SNR of 8 for both cases. The hierarchical search performs systematically a little worse than the flat search—we lose about 3% more injections than in the flat search.

uniformly distributed (identical to what was used for simulated data). The plots in Figs. 19 and 20 show the sensitivities with error bars for the flat and the 2-stage hierarchical search with small chunks of the real data. We can see that the hierarchical search sensitivity (blue line) is always slightly less as compared to that of the flat search (red line) which is optimized for all the 3 types of the CBC sources: BNS, NSBH, and BBH. The top plot in Fig. 19 shows the sensitivity distance for BNS which is just over 40 Mpc while the bottom plot in the figure shows the same for NSBH which is over 80 Mpc. These numbers are for a coincident SNR of 8 in both plots. In Fig. 20 the distance sensitivity is shown in the top plot for BBH, while the bottom plot shows the same when all the injections BNS, NSBH, and BBH are taken into account. The distance sensitivities are just over 400 Mpc for both cases. This analysis shows that hierarchical search performs almost as well as on real data as the flat search—we just about lose less than 3% of the total injections in this trial search. This is expected due to more triggers, inefficient clustering, and choices of the thresholds and statistics made for each stage in an *ad hoc* manner. We emphasize that this is only a demonstration and the hierarchical search needs to be tuned further to obtain almost full sensitivity with a sizable speed-up. One of the caveats here is that a small full run may be needed to determine the exact scaling factor to get the correct background with optimal choices of the parameters as clustering at various levels may be slightly different as compared to the optimized flat search.

VI. DISCUSSIONS AND FUTURE PROSPECTS

In this work, we have demonstrated that the two stage hierarchical search is ~ 20 times faster than the flat search which has been used in LIGO O1 analysis. This factor of reduction in computational cost has been obtained without any optimization. With a judicious choice of parameters we have shown that it can be almost as good as the single stage flat search in sensitivity—that is, given a set of injections, this search detects as many signals as the flat search. In the future, we propose to run and optimize our 2-stage hierarchical search on O1 and O2 data so as to improve the performance. We have already demonstrated the same in a limited way on real data, albeit with a slightly lower sensitivity.

As pointed out before, the computational effort saved by doing a hierarchical search can be used elsewhere. It can be used to do more detailed analysis of the detected CBCs such as test of general theory of relativity by comparing waveforms predicted by other theories of gravity etc. The saved CPU time could be used to search for other astrophysical sources. This issue will become all the more important when detectors become more sensitive in the future. The demand for computation will increase because the event rate will go

up with the corresponding requirement of a much denser template bank covering the parameter space.

The two stage hierarchical method can be readily employed for online searches where we do not care about assigning the exact significance to the triggers (the exact significance would only be obtained from the estimation of the full background). We may roughly scale the stage I background to arrive at a crude estimate of the significance of online triggers. Using a rough estimate of the background is a prevalent practice for obtaining promising triggers quickly [54]. In the first stage, we can adjust the false alarm rate to a desired level by varying the threshold and get online triggers much faster. This is the consequence of the speed up we get from the hierarchical algorithm.

Another important direction to follow is the implementation of a hierarchical search with precessing waveforms. We believe that the order of magnitude reduction in the computational cost will allow us to make at least partial inroads into searches for precessing binaries. But creating a template bank with precessing templates is also very difficult as it has to be done stochastically [32]. We plan to explore the possibility of performing multistage hierarchical searches using hybrid (nonprecessing + partial precessing) template banks.

ACKNOWLEDGMENTS

We would like to thank Albert Lazzarini (Caltech) for suggesting us the original idea. We like to acknowledge Alex Nitz (AEI-Hannover), Remya Nair (Kyoto University) for reading the manuscript and suggestions. We would also like to thank Anand Senpupta (IITG), Sukanta Bose (IUCAA), Badri Krishnan (AEI-Hannover), Ian Harry (AEI-Potsdam), Sumit Kumar (ICTS), and Patrick Brady (UWM) for the useful discussions. We acknowledge the use of IUCAA LDG cluster Sarathi for the computational/numerical work. B. G. acknowledges the support of University Grants Commission (UGC), India. This research benefited from a grant awarded to IUCAA by the Navajbai Ratan Tata Trust (NRTT). S. M. acknowledges support from the Department of Science & Technology (DST), India provided under the Swarna Jayanti Fellowships scheme. S. V. D. acknowledges the support of the Senior Scientist NASI Platinum Jubilee Fellowship. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), [arXiv:1811.12907](https://arxiv.org/abs/1811.12907).
- [2] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, *Astrophys. J.* **872**, 195 (2019).
- [3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [5] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **118**, 221101 (2017).
- [6] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J.* **851**, L35 (2017).
- [7] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **96**, 022001 (2017).
- [8] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 141101 (2017).
- [9] B. P. Abbott *et al.* (GROND, SALT Group, OzGrav, DFN, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAASTRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), Dark Energy Survey, MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, and ALMA Collaborations), *Astrophys. J.* **848**, L12 (2017).
- [10] B. Iyer, T. Souradeep, C. S. Unnikrishnan, S. Dhurandhar, S. Raja, A. Kumar, and A. Sengupta, LIGO-India, Technical Report No. LIGO-M1100296, 2011.
- [11] B. S. Sathyaprakash and S. Fairhurst, Technical Report, 2012, <https://dcc.ligo.org/public/0091/T1200219/001/LIGO-T1200219-v1.pdf>.
- [12] T. Akutsu *et al.* (KAGRA Collaboration), in *15th International Conference on Topics in Astroparticle and Underground Physics (TAUP 2017) Sudbury, Ontario, Canada, 2017* (2017) [[arXiv:1710.04823](https://arxiv.org/abs/1710.04823)].
- [13] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J.* **833**, L1 (2016).
- [14] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [15] S. V. Dhurandhar and B. S. Sathyaprakash, *Phys. Rev. D* **49**, 1707 (1994).
- [16] S. V. Dhurandhar and B. F. Schutz, *Phys. Rev. D* **50**, 2390 (1994).
- [17] G. Faye, S. Marsat, L. Blanchet, and B. R. Iyer, *Classical Quantum Gravity* **29**, 175004 (2012).
- [18] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [19] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **86**, 024011 (2012).
- [20] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer *et al.*, *Phys. Rev. D* **89**, 061502 (2014).
- [21] <http://www.black-holes.org/SpEC.html>.
- [22] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [23] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [24] A. Pai, S. Dhurandhar, and S. Bose, *Phys. Rev. D* **64**, 042004 (2001).
- [25] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [26] T. Dal Canton and I. W. Harry, [arXiv:1705.01845](https://arxiv.org/abs/1705.01845).
- [27] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [28] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, B. S. Sathyaprakash, H. Tagoshi, and H. Takahashi, *Phys. Rev. D* **80**, 024009 (2009).
- [29] S. Roy, A. S. Sengupta, and P. Ajith, *Phys. Rev. D* **99**, 024048 (2019).
- [30] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
- [31] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **93**, 122003 (2016).
- [32] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, *Phys. Rev. D* **94**, 024012 (2016).
- [33] C. Capano, I. Harry, S. Privitera, and A. Buonanno, *Phys. Rev. D* **93**, 124007 (2016).
- [34] S. Babak *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [35] D. Macleod, I. W. Harry, and S. Fairhurst, *Phys. Rev. D* **93**, 064004 (2016).
- [36] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, *Phys. Rev. D* **82**, 044025 (2010).
- [37] A. S. Sengupta, S. V. Dhurandhar, A. Lazzarini, and T. Prince, *Classical Quantum Gravity* **19**, 1507 (2002).
- [38] A. S. Sengupta, S. Dhurandhar, and A. Lazzarini, *Phys. Rev. D* **67**, 082004 (2003).
- [39] S. Mitra, S. V. Dhurandhar, and L. S. Finn, *Phys. Rev. D* **72**, 102001 (2005).
- [40] T. Dal Canton *et al.*, *Phys. Rev. D* **90**, 082004 (2014).
- [41] S. A. Usman *et al.*, *Classical Quantum Gravity* **33**, 215004 (2016).
- [42] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, *Astrophys. J.* **849**, 118 (2017).
- [43] PyCBC, <https://github.com/ligo-cbc/pycbc/compare/v1.9.2..v1.9.4>.
- [44] J. D. E. Creighton and W. G. Anderson, *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis* (Wiley-VCH, Berlin, 2011).
- [45] H. Mukhopadhyay, H. Tagoshi, S. Dhurandhar, and N. Kanda, *Phys. Rev. D* **80**, 123019 (2009).
- [46] S. D. Mohanty and S. V. Dhurandhar, *Phys. Rev. D* **54**, 7108 (1996).
- [47] S. D. Mohanty, *Phys. Rev. D* **57**, 630 (1998).
- [48] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [49] A. Sengupta, Ph.D. thesis, IUCAA, Pune, 2004.
- [50] D. A. Brown, Ph.D. thesis, Wisconsin University, Milwaukee, 2004.

-
- [51] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, *J. Phys. Conf. Ser.* **610**, 012021 (2015).
- [52] A. H. Nitz, *Classical Quantum Gravity* **35**, 035016 (2018).
- [53] S. Dhurandhar, A. Gupta, B. Gadre, and S. Bose, *Phys. Rev. D* **96**, 103018 (2017).
- [54] A. H. Nitz, T. Dal Canton, D. Davis, and S. Reyes, *Phys. Rev. D* **98**, 024050 (2018).