

Hierarchical search for compact binary coalescences in the Advanced LIGO's first two observing runs

Kanchan Soni^{1*}*Inter-University Centre for Astronomy and Astrophysics, Pune 411007, India*Bhooshan Uday Gadre^{2†}*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476 Potsdam, Germany and Gravitational and Subatomic Physics (GRASP), Utrecht University, Princetonplein 1, 3584 CC, Utrecht, The Netherlands*Sanjit Mitra^{1‡}*Inter-University Centre for Astronomy and Astrophysics, Pune 411007, India*Sanjeev Dhurandhar^{1§}*Inter-University Centre for Astronomy and Astrophysics, Pune 411007, India*

(Received 22 June 2021; accepted 31 January 2022; published 2 March 2022)

Detection of many compact binary coalescences (CBCs) is one of the primary goals of the present and future ground-based gravitational-wave (GW) detectors. While increasing the detectors' sensitivities will be crucial in achieving this, efficient data analysis strategies can play a vital role. With given computational power in hand, efficient data analysis techniques can expand the size and dimensionality of the parameter space to search for a variety of GW sources. Matched filtering-based analyses that depend on modeled signals to produce adequate signal-to-noise ratios for signal detection may miss them if the parameter space is too restrained. Specifically, the CBC search is currently limited to nonprecessing binaries only, where the spins of the components are either aligned or antialigned to the orbital angular momentum. A hierarchical search for CBCs is thus well motivated. The first stage of this search is performed by matched filtering coarsely sampled data with a coarse template bank to look for candidate events. These candidates are then followed up for a finer search around the vicinity of an event's parameter space found in the first stage. Performing such a search leads to enormous savings in the computational cost without much loss in sensitivity. Here we report the first successful implementation of the hierarchical search as a PyCBC-based production pipeline to perform a complete analysis of Laser Interferometer Gravitational Wave Observatory (LIGO) observing runs. With this, we analyze Advanced LIGO's first and second observing run data. We recover all the events detected by the PyCBC (flat) search in the first GW catalog, GWTC-1, published by the LIGO-Virgo collaboration, with nearly the same significance using a scaled background. In the analysis, we get an impressive factor of 20 reduction in computation compared to the flat search. With a standard injection study, we show that the sensitivity of the hierarchical search remains comparable to the flat search within the error bars.

DOI: [10.1103/PhysRevD.105.064005](https://doi.org/10.1103/PhysRevD.105.064005)

I. INTRODUCTION

Gravitational-wave (GW) astronomy began with the detection of GW signal from a binary black hole (BBH) merger, GW150914 [1], using the Advanced Laser Interferometer Gravitational Wave Observatory (LIGO) [2,3] observatories. With the latest advancements in the sensitivities of detectors

and search techniques like cWB [4], Gstreamer wrappings of the LIGO Algorithm Library (GstLAL) [5], PyCBC [6], LIGO-Virgo (LV) collaboration detected GW signals from ten BBHs and one binary neutron star (BNS) coalescence in the first two, O1 and O2, observing runs [7]. During this period, several independent searches [8–10] over publicly available data detected a few additional BBH events. A paradigm shift in the number of detections occurred with the improvement in the sensitivities of Advanced LIGO [11] and Advanced Virgo [12] detectors in the third observing run. This has led to the detections of many GW events [13–15] including GW190425 [16], the second BNS event, GW190412 [17],

*kanchasoni@iucaa.in

†b.u.gadre@uu.nl

‡sanjit@iucaa.in

§sanjeev@iucaa.in

and GW190814 [18], the first two highly asymmetric compact binary coalescences (CBCs) that emit a significant amount of gravitational radiation beyond the quadrupole moment, and GW190521 [19], the first binary merger to form an intermediate-mass black hole.

Matched filtering [20–24], a primary and most sensitive algorithm, is used to detect signals that can be well modeled. Since the GW signals from merging binaries in circular orbits can be modeled using their intrinsic¹ and extrinsic² parameters, the matched-filtering method is employed for their detection. The method involves correlating an interferometer’s output, time-series data, with the modeled waveforms (*templates*) for each detector in the network. If a GW signal is present in the detector’s output, the correlation results in a peak (*trigger*) in the signal-to-noise ratio (SNR) corresponding to the best-matching template. Since the prior knowledge of the source parameters, like its component masses, spins, and location in the sky, remains unknown to the observers, the search is required to be performed over a wide range of several source parameters using a “bank of templates.” The templates in the bank are closely placed to ensure that the search does not miss any signal. Since the data contain non-Gaussian noise, a coincidence search over the time of arrival, phase, and other source parameters is performed between different detectors to reduce the rate of false alarms. The coincident candidates obtained are then assigned significance based on the noise background.

The above procedure for detecting GW signals from CBCs is followed by the search pipelines like GstLAL [5,25,26], MBTA [27,28], PyCBC [6,8,9,29], and SPIIR [30]. These pipelines perform a *one-step search*³ for the nonprecessing coalescing binaries in quasicircular orbits.

One of the challenges that template-based search pipelines face is the high computational cost of matched filtering, typically a year’s worth of data over $\sim\mathcal{O}(10^5)$ templates. Since this process, especially in the PyCBC (or flat) search, involves fast Fourier transform (FFT) of the product of uniformly sampled time-series data and a template, the number of floating-point operations scales as $N \log_2 N$, where N is the number of data points. These operations repeat over $\sim\mathcal{O}(10^5)$ templates [31], even in the restricted parameter space of nonprecessing binaries with quasicircular orbits, which amounts to sizable computational cost. The cost further increases when a search is envisaged for precessing binaries where the orbital plane precesses due to the misalignment of component spins with the orbital angular momentum. In such cases, the number of templates and the matched-filtering operations increases at

least tenfold [32], thus making the search computationally expensive to pursue with the current capabilities. Furthermore, the search for primordial black holes in the subsolar region requires templates $\sim\mathcal{O}(10^5-10^6)$ [33,34], which makes the search more expensive. To reduce the search’s cost, matched filtering over the data is performed above a frequency of 45 Hz [33,34] while compromising with the overall reduction of $\sim 24\%$ in the sensitive volume. While another [35] search still uses a lower frequency of 20 Hz with waveforms having low eccentricities, it assumes very low nonprecessing spins to make searches computationally manageable. These limitations can be reduced if faster matched-filtering search algorithms are developed.

With the advancements in current detectors and upcoming new detectors, e.g., KAGRA [36] and LIGO-India [37], the CBC detection rate is bound to increase, and finer details of the detected sources would be sought to unravel their exact dynamics, formation, and evolution scenarios. However, this would significantly increase the volume of the search parameter space. The increment in volume would happen in two ways: first, the number of parameters (dimension of the parameter space) of different CBC sources would increase, and secondly, their ranges may increase. Nevertheless, a comprehensive matched-filter based search is important because one would like to capture the nontrivial dynamics of interesting astrophysical sources. Therefore, to facilitate this quest, we assert that a matched-filter based search pipeline needs speeding up by orders of magnitude.

One way to speed up the search is by performing matched filtering *hierarchically using multiple banks of varying densities*. The idea of performing matched filtering in hierarchical steps was formally introduced in Mohanty and Dhurandhar [38], where hierarchy was performed over the chirp mass of binaries using Newtonian waveforms. This work was later extended to the post-Newtonian waveforms [39], where hierarchy was performed over the component masses of a binary system. A further improvement was realized by reducing the sampling rate in the first stage of the hierarchy. In the recent work of Gadre *et al.* [40] hierarchy was performed over all the three intrinsic parameters, including the effective spin of the binary. This algorithm had used two-detector coincidence analysis and had provided an order of magnitude speed-up compared to the flat analysis.

In this paper, we revisit the hierarchical search formulated in Gadre *et al.* [40], and for the first time, implement it as a working PyCBC-based pipeline to analyze the data from an entire LIGO observing run. We describe an efficient two-stage hierarchical search pipeline to search for GW signals from CBCs in the two detectors. This pipeline improves the hierarchical search sensitivity by incorporating better detection statistics for the single-detector and coincident triggers, as used by the flat analysis in Abbott *et al.* [13]. For this work, we construct two template banks—*coarse* and *neighborhood(nbhd)* bank, to target GW signals from nonprecessing CBC sources that have

¹Component masses (m_1, m_2) and individual spins (\vec{s}_1, \vec{s}_2) vectors of the coalescing binary.

²Sky location (ζ, ϕ), luminosity distance (d_L), orbital inclination (i), polarization angle (ψ), and time and phase of coalescence (t_c, ϕ_c) of the coalescing binary with respect to the detector.

³Search involving match-filtering data, sampled at a fixed rate using a bank of templates.

quasicircular orbits. Using the former bank in the first stage and a dynamical subset of the latter bank in the second stage of the hierarchical search, we test the potency of the pipeline by applying it to the data from the first two observing runs of Advanced LIGO. Our pipeline recovers all the GW events observed by the flat search from the first gravitational wave catalog (GWTC-1) [7].

In our work, we assign the significance to the detected events using a *scaled background* [40], constructed by scaling the background obtained in the first stage by time sliding the filtered output across detectors using a coarse bank. To justify the accuracy of this background, we perform simulations that involve the injection of the GW signals into the data and compare their recoveries with the hierarchical and flat search separately. Furthermore, we compare the sensitivities of the two searches through “volume-time” (*VT*) ratio curves. Lastly, we conclude our findings from the injection study by discussing the two searches’ matched-filter computational costs.

The paper is organized as follows:

- (i) In Sec. II we state the prerequisites and describe the search methodology for the hierarchical search. The section segregates into subparts. In Sec. II A, we describe the generation of template banks. Section II B elaborates on the matched-filtering process and selection criteria for the generated triggers in two stages. The strategy to collect coincident triggers is described in Sec. II C. The final step in the pipeline is to assign significance to the coincident candidates. We describe this process in Sec. II D.
- (ii) In Sec. III, we implement the hierarchical search pipeline over the first two observing runs of Advanced LIGO and present our findings.
- (iii) In Sec. IV, we compare the sensitivities of the hierarchical search with flat search. We also discuss the computational advantages of the former search with the latter.
- (iv) In Sec. V, we summarize our findings and point out the directions of future research.

II. PREREQUISITES AND SEARCH METHODOLOGY

The idea of the hierarchical search is straightforward; the flat search algorithm is divided into two stages, *stage 1* and *stage 2*, such that the number of matched-filter operations reduces successively. Stage 1 search ensures matched filtering of the data sampled at the lower sampling rate (512 Hz) using a sparsely sampled bank called the coarse bank. Having fewer templates in a coarse bank significantly reduces the computational cost of matched filtering. Further reduction in the computation is achieved by sampling data at a lower rate than the value used in the flat search. The coarse bank can reduce the SNRs for a good fraction of events because of the sparsely placed templates. To compensate for the loss in SNRs, we identify triggers in each detector above *coarse*

thresholds, set at lower values than those used in the flat search. We then perform a coincidence test on these identified triggers, using optimal detection statistics and obtain the zero-lag (or foreground) candidates. These foreground candidates are then followed up in stage 2 to ascertain whether they are signals or false alarms.

In stage 2, we again perform matched filtering over the data segments containing followed-up foreground candidates from the stage-1 search. These data segments are sampled at a flat search sampling rate (2048 Hz) and filtered using a dynamic union of nbhds of mismatch extending up to 0.75 around the followed-up stage-1 trigger templates. We refer to this union of nbhds as the *stage-2 bank* from now on. The triggers generated for each detector in this stage are identified above *fine thresholds*, equal to the thresholds set for SNRs in the flat search. These triggers

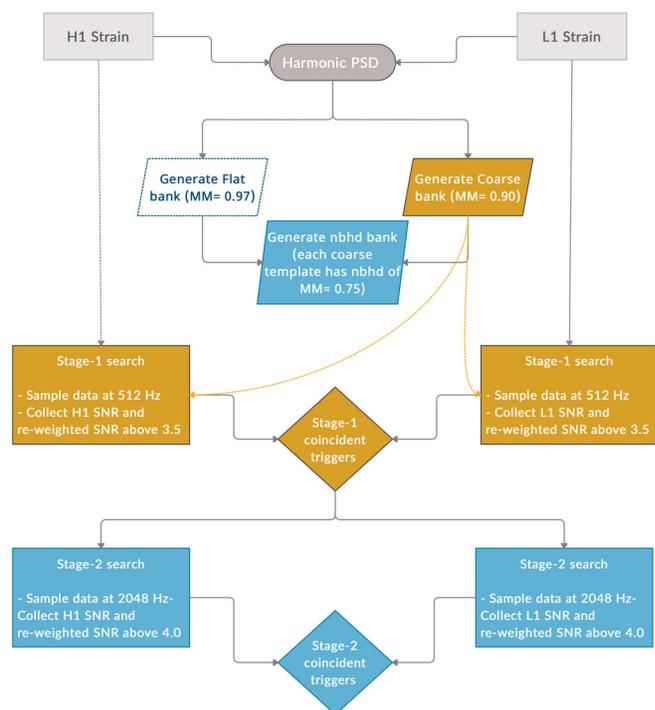


FIG. 1. A flowchart depicting the working of a two-stage hierarchical search pipeline. The choice of color describes the stage: yellow for stage 1 and blue for stage 2. The first step generates the harmonic-mean power spectral density (PSD) using the strain data from the two detectors. The generated PSD is used to create flat (in dotted box) and coarse banks. Using these two banks, a nbhd bank is thus constructed. Stage-1 search begins with matched filtering of the strain data from the two detectors using a coarse bank. The generated triggers are then identified above SNRs and reweighted SNRs of 3.5. Next, a coincidence test is made to collect the foreground candidates (in diamond box), which are then followed up in stage 2 for a finer search. In stage 2, a search over the segments containing these followed-up candidates is performed using a subset of nbhd bank, stage-2 bank. The triggers generated are then identified above SNR and reweighted SNR of 4. At last, the selected triggers are then subjected to a coincidence test to obtain a final list of foreground candidates.

are then subjected to a coincidence test before generating the final list of foreground candidates.

Our two-stage hierarchical search pipeline is described through the flowchart in Fig. 1.

A. Template banks

One of the most crucial steps in a template-based search is to adequately grid up the parameter space. A pragmatic approach suggests densely populating the search space to minimize the loss in SNR. However, such dense placement of templates makes the search computationally expensive and limits the volume and dimensionality of the parameter space that can be covered, given a fixed amount of computation power. Generally, the templates are placed such that the “match” (\mathcal{M}) does not fall below a certain minimum value called the minimal match (MM). For instance, if MM is chosen as 0.97, it means that the expected SNR for a signal does not fall more than 3% ($1 - \text{MM} = 0.03$), corresponding to a loss of $\sim 10\%$ ($\approx 1 - \text{MM}^3$) in the astrophysical events.

The match between two normalized templates is their scalar product, maximized over the extrinsic parameters, namely, the time t_c and phase ϕ_c at coalescence. If $h(t_c, \phi_c, \vec{\theta})$ and $h(t'_c, \phi'_c, \vec{\theta}')$ are normalized templates defined by the intrinsic parameters $\vec{\theta}$ and $\vec{\theta}'$, where $t'_c = (t_c + \Delta t_c)$, $\phi'_c = (\phi_c + \Delta \phi_c)$, $\vec{\theta}' = (\vec{\theta} + \Delta \vec{\theta})$ for $\vec{\theta} = \{m_1, m_2, s_{1z}, s_{2z}\}$, then the match is

$$\mathcal{M}(\vec{\theta}, \Delta \vec{\theta}) = \max_{\Delta t_c, \Delta \phi_c} (h(t_c, \phi_c, \vec{\theta}), h(t'_c, \phi'_c, \vec{\theta}')), \quad (1)$$

where the scalar product of arbitrary data trains $x(t)$ and $y(t)$ is defined as

$$(x, y) := 4\mathcal{R} \left\{ \int_{f_{\text{low}}}^{f_{\text{high}}} \frac{\tilde{x}(f)\tilde{y}^*(f)}{S_n(f)} df \right\}. \quad (2)$$

Note that the match does not depend (or weakly depends) on the absolute values of the extrinsic parameters t_c and ϕ_c , and hence they have been dropped as arguments of \mathcal{M} .

In Eq. (2), \mathcal{R} denotes the real part of a complex quantity evaluated under the sensitive frequency band, i.e., f_{low} to f_{high} of the detector and weighted by the detector’s one-sided noise PSD $S_n(f)$ defined by

$$\langle \tilde{n}(f)\tilde{n}(f') \rangle = \frac{1}{2} S_n(f) \delta(f - f'). \quad (3)$$

The angular brackets denote the ensemble average of the noise ($n(f)$) realizations. The tilde in Eqs. (2) and (3) represents Fourier transform of the quantity in question, e.g., $\tilde{x}(f)$ is the Fourier transform of $x(t)$ and is given by

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt. \quad (4)$$

Assuming a slowly varying metric $g_{mn}(\vec{\theta})$ around the targeted templates, we Taylor expand $\mathcal{M}(\vec{\theta}, \Delta \vec{\theta})$ to the lowest order of $\Delta \theta$ as

$$\mathcal{M}(\vec{\theta}, \Delta \vec{\theta}) \approx 1 - g_{mn}(\vec{\theta}) \Delta \theta^m \Delta \theta^n. \quad (5)$$

where the parameter-space metric is defined as

$$g_{mn} = -\frac{1}{2} \frac{\partial^2 \mathcal{M}}{\partial \Delta \theta^m \partial \Delta \theta^n}(\vec{\theta}). \quad (6)$$

Therefore, by varying the source parameters $\vec{\theta}$ and calculating the metric g_{mn} , templates can be effectually placed in the bank. However, generally, the metric does not have a closed-form expression for aligned-spin waveforms having inspiral, merger, and ringdown phases for a wide range of source parameters, e.g., SEOBNR [41]. Moreover, in some cases, metric changes rapidly across the parameter space, making the sphere-covering problem [42] highly nontrivial. Therefore, techniques like stochastic placement [43] are used to construct the bank, where a direct match is computed between templates for varying source parameters. This technique efficiently places the templates in a bank. However, if the volume of the parameter space (as defined via the metric) is large, then the template bank also becomes large and increases the computational cost for bank generation. In such a case, techniques like hybrid geometric-random placements [44,45] efficiently generate a full nonprecessing bank.

The density of templates in a bank relies on time-average noise PSD across all the detectors. Since the search pipeline uses a common template bank for all the detectors, a time-averaged noise PSD for each detector is estimated. These time-averaged PSDs are then combined as a harmonic mean [46–48] for the bank’s construction.

In this work, we construct a coarse and nbhd bank for targeting GW signals from nonprecessing sources with quasicircular orbits, using Advanced LIGO-Virgo noise PSD as used in GWTC-2 [13]. We describe the construction of banks for the parameter ranges provided in Table I in the following sections.

1. Coarse bank

We construct a coarse bank with a mismatch of 10% (or $\text{MM} = 0.90$) following Gadre *et al.* [40], using the hybrid geometric-random method [44,45]. The templates in the bank are generated at a minimum frequency of 15 Hz. We discard the templates with a duration of less than 150 ms to avoid artifacts in the matched-filtering steps. The bank is designed to search for quadrupolar, quasicircular, and nonprecessing CBC sources with the redshifted total mass (M_{tot}) of the binary in the range $[2 M_{\odot}, 500 M_{\odot}]$. We restrict the primary (m_1) and secondary (m_2) mass observed in the detector’s frame in the ranges $[1 M_{\odot}, 500 M_{\odot}]$ and $[1 M_{\odot}, 120 M_{\odot}]$, respectively. The ranges for individual dimensionless spins of the binaries comprising a black hole

TABLE I. Table summarizing the minimal match values and the ranges of the source parameters for the coarse, nbhd, and flat banks. The χ_{BH} and χ_{NS} are the dimensionless effective spins for a black hole and neutron star, respectively.

Bank	MM	$M_{\text{tot}}(M_{\odot})$	χ_{BH}	χ_{NS}	f_{min} (Hz)
Coarse	0.90	2–500	−0.998–0.998	−0.05–0.05	15
Flat and nbhd	0.97	2–500	−0.998–0.998	−0.05–0.05	15

(χ_{BH}) and a neutron star (χ_{NS}) are provided in Table I. Thus, we construct a nonprecessing coarse bank (see Fig. 2) consisting of 85,080 templates.

To check whether the bank does not possess holes in the parameter space, we test the bank’s performance in terms of fitting factor (FF) [49]. In this test, we estimate FF for $\sim 80,000$ quasicircular, quadrupolar, spin-aligned, and nonprecessing CBC signals that span the bank’s search parameter space. We use `TaylorF2RedSpin` [50] with M_{tot} in the range $[2 M_{\odot}, 5 M_{\odot}]$ and $|\chi_{\text{eff}}| \leq 0.05$, and `SEOBNRv4_ROM` [41] in $[5 M_{\odot}, 500 M_{\odot}]$ with $|\chi_{\text{eff}}| \leq 0.998$. We recover FF greater than 0.90, as can be seen in Fig. 3. This result signifies that our bank is effectual and suffices the design criteria as per Table I.

2. Neighborhood bank

For a template corresponding to stage-1 coincident trigger, the template nbhd is the region in parameter space where mismatch with neighboring templates can be up to 25%, as described in Sec. IIIB2 of *Gadre et al.* [40]. To sample these nbhds, we use a pregenerated flat bank with $MM = 0.97$ with the search parameter space provided in Table I. We include flat bank’s templates having $MM \geq MM_{\text{nbhd}} \equiv 0.75$ with the trigger template. We calculate nbhds for all the coarse templates. This precomputed bank with assigned nbhds is referred to as a nbhd bank, and a dynamic subset of it is termed as a stage-2 bank. The stage-2 bank is dynamic because the number of templates

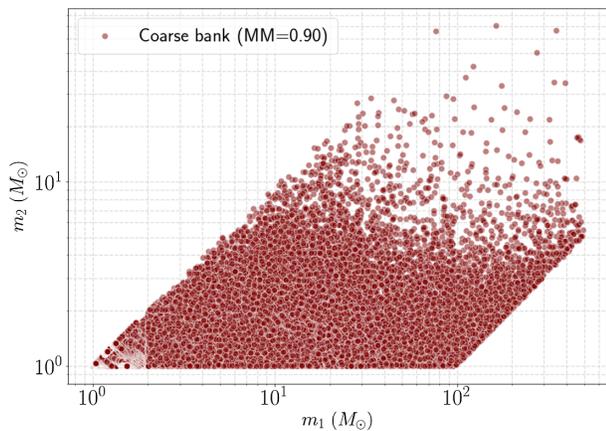


FIG. 2. The distribution of coarse bank templates in $m_1 - m_2$ mass plane. Each plot point represents a template with an $MM = 0.90$ with the neighboring templates.

residing in this bank changes depending on the noise realization of each segment.

To identify the nbhds of coarse templates, we adopt the following strategy. For coarse templates with $M_{\text{tot}} \geq 12 M_{\odot}$, we perform an exact match calculation with all the flat bank templates. For templates with $M_{\text{tot}} < 12 M_{\odot}$ [50], we first shortlist a set of templates that may be able to satisfy the nbhd criteria. For that, we define a minimal match ellipsoid with $MM = MM_{\text{nbhd}}$ in the following way: Consider a coarse template of $M_{\text{tot}} < 12 M_{\odot}$ for which nbhd has to be calculated. We first construct a minimal-match ellipsoid centered at this template in a coordinate system where the metric varies slowly over the parameter space, i.e., the metric is almost constant, and the signal manifold is almost flat. Therefore we choose chirp-time coordinates $\tau_0, \tau_3, \tau_{3s}$, collectively labeled as τ^α . These coordinates are given by scaling $\theta_0, \theta_3, \theta_{3s}$, described in Ref. [44], with $(2\pi f_o)^{-1}$ at $f_o = 20$ Hz. In these coordinates, we estimate the metric components using `TaylorF2RedSpin` waveform model. Once the metric is known, we, following [44], diagonalize it by an orthogonal transformation \mathcal{O} , and obtain the eigenvalues γ_α with new coordinates $\xi^\alpha = \mathcal{O}^\alpha_\beta \tau^\beta$. The metric in these coordinates is in a diagonal form and is given by

$$ds^2 = \sum_{\alpha=1}^3 \gamma_\alpha (d\xi^\alpha)^2. \quad (7)$$

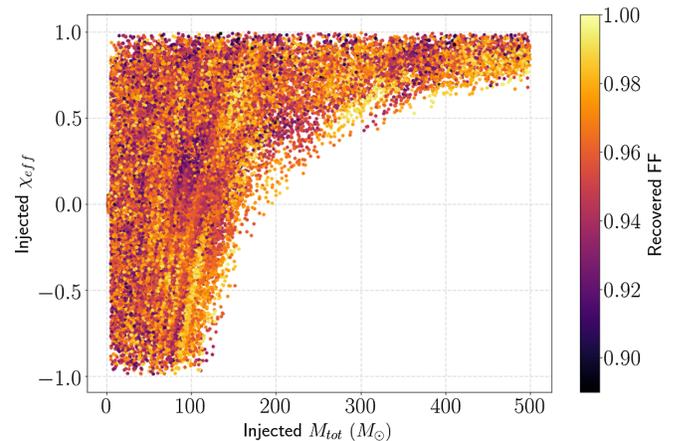


FIG. 3. Recovered fitting factor as a function of effective spin (χ_{eff}) and redshifted total mass (M_{tot}) plot for injected signals with M_{tot} in the range $[2 M_{\odot}, 500 M_{\odot}]$ with the signal duration cutoff of 150 ms.

This is just a principal axis transformation to an orthogonal basis. Along the eigendirections, the lengths of the semiaxes [$r_\alpha(\text{MM})$] of the ellipsoid for a given value of MM are given by

$$r_\alpha(\text{MM}) = \sqrt{\frac{1 - \text{MM}}{\gamma_\alpha}}. \quad (8)$$

As MM reduces from its maximum value of unity, the ellipsoid increases in size. In the ξ^α coordinates, let the coarse and flat templates be labeled by ξ_0^α and ξ^α , respectively. Let $\Delta\xi^\alpha = \xi^\alpha - \xi_0^\alpha$, and define the distance $d(\xi^\alpha, \xi_0^\alpha)$ by the equation

$$d^2(\xi^\alpha, \xi_0^\alpha) = \sum_{\alpha=1}^3 \gamma_\alpha (\Delta\xi^\alpha)^2. \quad (9)$$

Then, the relation $d(\xi^\alpha, \xi_0^\alpha) \leq \sqrt{1 - \text{MM}_{\text{nbhd}}}$ defines the ellipsoid in ξ^α coordinates. We use this ellipsoid to guide our selection of flat templates. Note that the metric approximation is extrapolated beyond its validity regime, so the ellipsoid is only a crude estimate of the nbhd. In any case, since we have made a conservative choice of $\text{MM} = \text{MM}_{\text{nbhd}}$, we do not expect to miss out on any signals. We choose templates accordingly in this region and compute the match between a flat template and a given coarse template. If the match is above the stipulated MM_{nbhd} , we retain the template in the nbhd. Thus, the final list of templates in the nbhd is obtained by the actual computation of the match between coarse and fine templates inside the ellipsoid.

In general, we find that a single nbhd around a coarse template (not very close to the boundary of the parameter space) contains ~ 40 – 150 templates. Since the match falls gradually with an increasing mismatch in the τ_3 mass parameter (as compared to τ_0), the nbhd tends to extend considerably along with this coordinate (see Fig. 3 in Sengupta *et al.* [51]). Therefore, a large portion of the nbhd can extend outside the physical parameter space considered for the search, especially for higher τ_0 . This causes a significant variation in the number of templates in the nbhd, as reflected in the top panel of Fig. 4. It is also interesting to note that the variation in the number of templates in the nbhd (bottom panel of Fig. 4) resembles the actual template density of the flat bank plotted in τ_0 and χ_{eff} coordinates (Fig 5). The figure indicates that there is a higher template density around high χ_{eff} and low τ_0 .

B. Matched filter

The model-dependent search for GWs from CBCs using templates in the LIGO-Virgo data exploits the matched filtering [20] technique rigorously. This technique correlates discretely sampled time-series data $s(t)$ with the normalized templates $h(t_c, \phi_c, \vec{\theta})$ for the source parameters

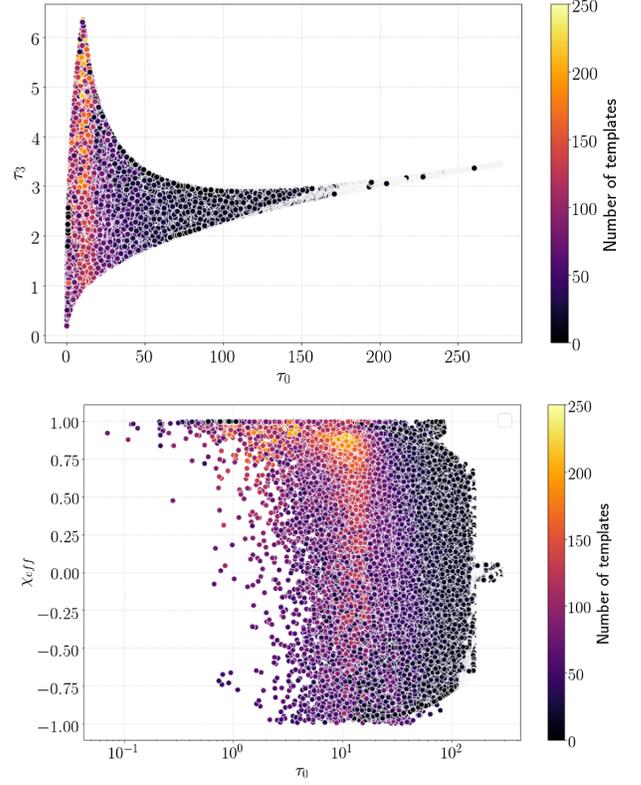


FIG. 4. The distribution of nbhd bank templates in $\tau_3 - \tau_0$ (top) and $\chi_{\text{eff}} - \tau_0$ plane (bottom). The color scale represents the number of templates in the nbhd of each coarse template. Typically, there are ~ 40 – 150 flat templates in the nbhd of a coarse template.

$(\vec{\theta})$ within the detectors' sensitive band. The correlation generates matched-filter SNR time series $\rho(t_c)$ maximized over the coalescence phase ϕ_c , and it is defined as

$$\rho(t_c; \vec{\theta}) \equiv |(s, (1 + i)h(t_c, \phi_c = 0, \vec{\theta}))|. \quad (10)$$

Generally, the data obtained from the detectors are non-stationary and non-Gaussian [52–54]. Preprocessing steps

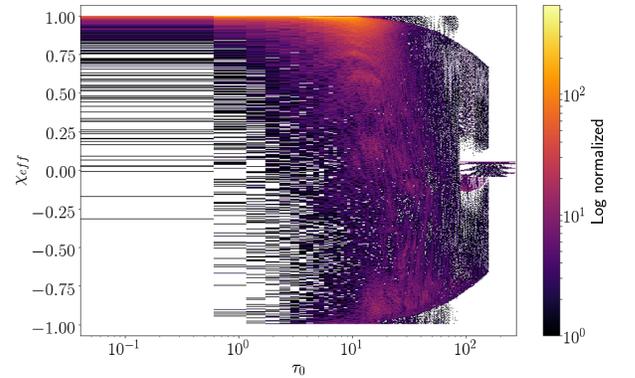


FIG. 5. Plot showing the density of flat bank's templates in $\chi_{\text{eff}} - \tau_0$ plane. The color scale represents the log-normalized number density.

involving data-quality checks and application of vetoes flag most of the artifacts present in the data [55]. Nevertheless, the short-duration glitches or long-duration correlations, as described in Venumadhav *et al.* [10], still remain in it. Matched filtering over these noise transients often leads to high SNRs. These short-duration noise transients are removed from the standard search pipelines by nullifying noise contributions in the time-series data via gating [6]. We, therefore, apply a similar gating priory to the matched filtering for each detector to remove the nonstationary transients from the strain data in our analysis.

Matched filtering the data produces several triggers with varying SNRs for each template in the bank. These triggers are first clustered within a time window of 0.5 s to retain only the ones with high SNRs [24]. In the second step, the SNRs of triggers due to loud noise artifacts are suppressed using signal consistency tests like the standard chi square (χ_r^2) [6,56], and sine-Gaussian chi square (χ_{sg}^2) [57].

Like in flat search, the trigger SNRs (ρ) generated in both the stages of the hierarchical search are down-weighted with their reduced chi-square values using χ_r^2 [6,56] veto defined as

$$\tilde{\rho} = \begin{cases} \frac{\rho}{[(1+(\chi_r^2)^3)/2]^{1/6}} & \text{if } \chi_r^2 > 1, \\ \rho & \text{otherwise.} \end{cases} \quad (11)$$

Usually, χ_r^2 veto is ineffective in the region where signals are too short. In such cases, the short-duration templates ring with ‘‘blip’’ glitches present in the data. Therefore, we further down-weight $\tilde{\rho}$ for the templates with $M_{\text{tot}} > 30 M_\odot$ using χ_{sg}^2 [57] veto defined as

$$\hat{\rho} = \begin{cases} \tilde{\rho}(\chi_{r,sg}^2)^{-1/2} & \text{if } \chi_{r,sg}^2 > 6, \\ \tilde{\rho} & \text{otherwise.} \end{cases} \quad (12)$$

In each stage of the hierarchical search, the triggers that surpass the two tests above specific thresholds on ρ and $\tilde{\rho}$ (see Sec. III) are subjected to a coincidence test to recover the real GW events. The coincident events are obtained based on the optimal detection statistics as defined in [29,58], which we elaborate on in the following section.

C. Ranking statistics

A pair of triggers from the two detectors is coincident if it simultaneously occurs within the light travel time between them and is recovered with identical template parameters. The coincidence is evaluated based on optimal detection or ranking statistics (Λ_{opt}), defined as the ratio of the likelihood for data containing signal to the likelihood for data having noise [59]. These likelihoods are the functions of the template parameters ($\vec{\theta}$) and $\hat{\rho}$, χ_r^2 .

In the recent works [7–9,29], the optimal detection statistics were approximated by taking the ratio of coincident event rate densities due to signal ($p(\vec{\kappa}|S)$) and noise

($p(\vec{\kappa}|N)$). Therefore, for an unknown coincident with template parameters $\vec{\kappa} = \{\hat{\rho}_H, \hat{\rho}_L, \chi_H^2, \chi_L^2, \delta t_c, \delta \phi_c, \vec{\theta}\}$, where $\delta t_c, \delta \phi_c$ is the time and phase difference in between two detectors, Λ_{opt} is given as

$$\Lambda_{\text{opt}} = \frac{p(\vec{\kappa}|S)}{p(\vec{\kappa}|N)} \equiv \frac{p(\vec{\kappa}|S)}{r_\theta^{\text{HL}} p(\vec{\theta}, \delta t_c, \delta \phi_c|N)}. \quad (13)$$

For the statistics, $p(\delta t_c, \delta \phi_c|N)$ is expected to be uniform over $(\vec{\theta}, \delta t_c, \delta \phi_c)$ [59]; thereby, it is marginalized and treated as a constant. If the noise is uncorrelated between detectors, $p(\vec{\kappa}|N)$ ($\approx r_\theta^{\text{HL}}$) can be safely written as a product of single-detector noise rate densities [29] ($r_{\vec{\theta},H}, r_{\vec{\theta},L}$) given by

$$r_\theta^{\text{HL}} = 2\tau_{\text{HL}}(r_{\vec{\theta},H}(\hat{\rho}_H)r_{\vec{\theta},L}(\hat{\rho}_L)), \quad (14)$$

where, τ_{HL} is the allowed time window for a coincidence of trigger in twin LIGO detectors at Hanford (H) and Livingston (L).

Thus, by estimating r_θ^{HL} and $p(\vec{\kappa}|S)$ through accurate modeling [29,57], one can obtain Λ_{opt} for the coincident triggers.

In each stage of the hierarchical search, we model r_θ^{HL} and $p(\vec{\kappa}|S)$ separately to obtain the ranking statistics of coincident and time-shifted events. In the first stage (and flat search), we adopt a similar methodology of modeling coincident signal and noise rate densities for a two-detector configuration as described in Davies *et al.* [29]. However, we model coincident noise rate density slightly differently in the second stage. In the following sections, we first review the existing modeling procedure for signal and noise rate densities used by the flat and stage-1 search and then elaborate on modeling noise rate densities for stage-2 search.

1. Signal model: For flat, stage 1, and stage 2

To model $p(\vec{\kappa}|S)$, one requires the probable astrophysical distribution of the binary sources that Advanced LIGO detectors can detect. In reality, the exact distribution is unknown to the observers. Nevertheless, the source population can be approximated as uniform in volume and isotropic in the sky location and orientation of the binary. Assuming these distributions for sources, we can estimate how their detection parameters like signal amplitudes, time, and phase differences vary with respect to the pair of the LIGO detectors.

As described in [29,58], $p(\vec{\kappa}|S)$ is precomputed by performing Monte Carlo simulations assuming fixed detector sensitivity. Then the corresponding $p(\vec{\kappa}|S)$ is used to rank each coincident trigger with parameter closed to $\vec{\kappa}$ [58].

We use the above recipe to generate $p(\vec{\kappa}|S)$ in the flat and both stages of the hierarchical search.

2. Noise model: For flat and stage 1

The coincident noise event rate density, $r_{\vec{\theta}}^{\text{HL}}$, for the flat and stage-1 search is obtained by first estimating the single-detector noise rate densities ($r_{\vec{\theta},d}|_{d=\{H,L\}}$). Like in Davies *et al.*[29], this quantity in the flat and stage-1 search is calculated by modeling the tail of the trigger distribution for each detector (d) and template with a falling exponential function as

$$r_{\vec{\theta},d}(\hat{\rho}_d, N) = \mu(\vec{\theta})p(\hat{\rho}_d|\vec{\theta}, N), \quad (15)$$

given

$$p(\hat{\rho}_d|\vec{\theta}, N) = \begin{cases} \alpha(\vec{\theta}) \exp[-\alpha(\vec{\theta})(\hat{\rho}_d - \hat{\rho}_{\text{th},d})] & \text{if } \hat{\rho}_d > \hat{\rho}_{\text{th},d}, \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

where $\mu(\vec{\theta})$ and $\alpha(\vec{\theta})$ denote trigger count above the threshold ($\hat{\rho}_{\text{th},d}$) and exponential decay rate, respectively.

The fit parameter $\alpha(\vec{\theta})$ is obtained by maximum logarithmic likelihood fitting method. For discrete samples of $\hat{\rho}_d$ of k th trigger, we maximize

$$\ln p(\hat{\rho}_d|\alpha, n) = n \ln \alpha - \alpha \sum_k^n (\hat{\rho}_{k,d} - \hat{\rho}_{\text{th},d}), \quad (17)$$

at a fixed $\hat{\rho}_{\text{th},d}$ ($\equiv 6$) to obtain $\alpha_{\text{max}} = (\bar{\hat{\rho}}_d - \hat{\rho}_{\text{th},d})^{-1}$. Here, $\bar{\hat{\rho}}_d$ is the mean of $\hat{\rho}_d$ and the variance (σ_d) in fit parameter is given by $1/\sqrt{n}$, where n denotes the number triggers generated for a particular template.

In the flat and stage-1 search, we calculate α_{max} and n for each flat and coarse template, respectively. Generally, not all the templates have sufficient triggers above 6 to fit the trigger distribution's exponential tail. In such cases, the low number of triggers gives a high variance to the fit parameter values. To avoid such problems, we take the moving average of the fit parameters and smooth $\mu(\vec{\theta})$ by taking mean over nearby templates with similar values of effective spin, template duration, and symmetric mass ratio as performed in Davies *et al.* [29].

3. Noise model: For stage 2

In principle, the procedure for obtaining single-detector noise rate densities described previously can be applied in stage 2. However, it cannot be implemented, as this stage possesses insufficient triggers above $\hat{\rho}_{\text{th},d}$ to obtain meaningful fit parameters. The reason is, we follow only foreground candidates from stage 1 that have $\Lambda_{\text{opt}} \geq 7$. Matched filtering over these followed-up triggers utilizes

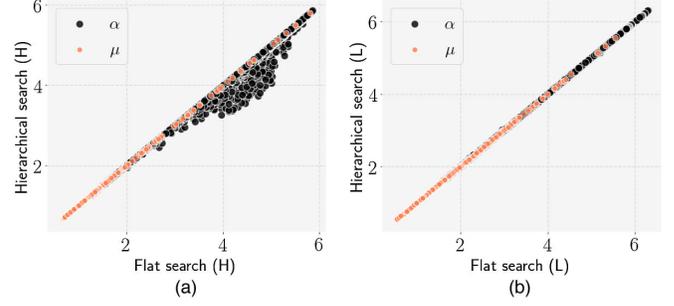


FIG. 6. Comparison plot for the fit coefficients, α and μ , obtained from the flat and hierarchical search for (a) Hanford (H) and (b) Livingston (L) detectors.

fewer nbhds and corresponding templates to generate fewer triggers. Having an inadequate and biased set of triggers for a template can give a significant variance in the values of fit parameters, leading to overestimating single-detector noise rates if we only use stage-2 triggers. We, therefore, *do not* explicitly calculate the fit parameters in stage 2. Instead, we reuse the fit values of the “closest” coarse template to the stage-2 trigger template. The “closeness” relies on the highest match value between the coarse and stage-2 bank templates.

To verify the applicability of the above procedure, we perform a flat and hierarchical search on 14 days and obtain fit parameters. Figure 6 compares the fit parameters obtained in both the searches. The scatter points in the diagonal signify that the values are comparable for the two searches in both the detectors. A few templates in the Hanford detector show low α indicating small fluctuations in their values due to noise. These small fluctuations can appear at different periods of observational time. However, these variations in α negligibly affect the modeling of single-detector noise rate density, as can be seen later in Sec. IV.

D. Assigning significance

The significance of any event is evaluated based on their false alarm rate (FAR) estimate above a fixed statistic Λ^* as

$$\text{FAR}(\Lambda^*) = \int d^n \vec{\kappa} r_{\vec{\kappa}}^{\text{HL}} \Theta(\Lambda_{\text{opt}}(\vec{\kappa}) - \Lambda^*), \quad (18)$$

where $r_{\vec{\kappa}}^{\text{HL}} \equiv r_{\vec{\theta}}^{\text{HL}}$ by construction. FAR signifies the rate of occurrence of a nonastrophysical coincident candidate with a similar or higher Λ_{opt} [see Eq. (13)] in the observing period. FAR is estimated in the flat and stage-1 search with respect to a noise background constructed by time-sliding data by a minimum of 100 ms across the detectors. Such a procedure omits all the possibilities to have a coincidence due to a real GW signal. At each time shift, Λ_{opt} is recomputed to rank the candidates above a certain threshold (Λ^*). Performing several time shifts generates many

plausible candidates that could be cumbersome to store. In order to mitigate the storage problem, the background computation is optimized in the standard PyCBC search. At first, a clustering over time is performed such that candidates with the highest statistic value, falling within 10 s, are kept. In the next step, candidates are selectively chosen with all or few time slides falling in the ranking statistic value's bin. For instance, candidates with all possible time slides with ranking statistics greater than 9 are chosen, but only some are selected with time slides of 30 s for which statistic value lies between 8 and 8.5.

In principle, a similar strategy can be implemented to assign FARs to the detected candidates in stage 2 of the hierarchical search. However, the background constructed by time-sliding stage-2 triggers using a union of stage-2 banks can bias the detected candidates' FAR estimates, as shown in Gadre *et al.* [40]. Therefore, we avoid such biases by constructing an approximate background that would mimic a background constructed in the flat search. As proposed in Gadre *et al.* [40], we construct a scaled stage-1 background for assigning significance to the final list of coincident triggers. First, we construct a stage-1 background by time-sliding stage-1 triggers by 100 ms across the detectors as done in the flat search. We then scale this background by a factor equal to the ratio of the number of templates times the sampling frequency used in a flat search to stage 1. This number turns out to be close to the computational gain and is approximately 20.

To justify our argument on mimicking a flat background, we compare the foreground and background obtained from the flat and hierarchical search performed over 14 days of data around the first BBH, GW150914 [1], event. We find that the foregrounds due to noise candidates match their respective backgrounds for both the searches, as shown in Fig. 7. We observe that the noise background is higher in

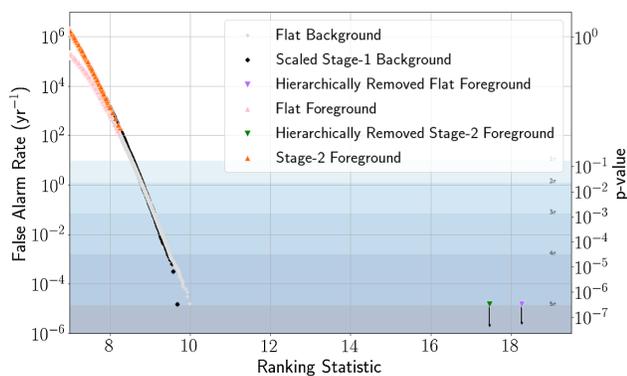


FIG. 7. Plot showing FAR vs ranking statistic curves for the foreground candidates (foreground) and the time-shifted candidates (background) from a flat and hierarchical search. The foreground (triangle) overlays the background (circle) in each search. The loudest event, GW150814, is hierarchically removed from the background in both searches. Note that the scaled stage-1 background (black) roughly matches the flat background (gray) above ranking statistic value 8.

the lower ranking statistics region than that of flat. This observation is expected as the scaling factor linearly increases the number density of noise triggers in a particular ranking statistics bin. We also notice that the scaled stage-1 background roughly matches the flat background above ranking statistic value 8. Therefore, the reliability of the FARs will be limited to the ranking statistic value $\gtrsim 8$. Another way to justify the reliability of the background is by looking at the effects of variations in sampling rates and the number of templates for each pair of the search. Figure 8 compares the backgrounds obtained from the flat search at 512 Hz and the stage-1 search at 2048-Hz sampling rates. We show that if these backgrounds are scaled with a factor of 4 and 5, respectively, both nearly match the standard flat search background at 2048-Hz sampling. Thus, the factor of 5 reductions in the number of templates and 4 reductions in the sampling rate, whose product gives us the final scaling factor of 20, are valid scaling factors on their own. While the scaling

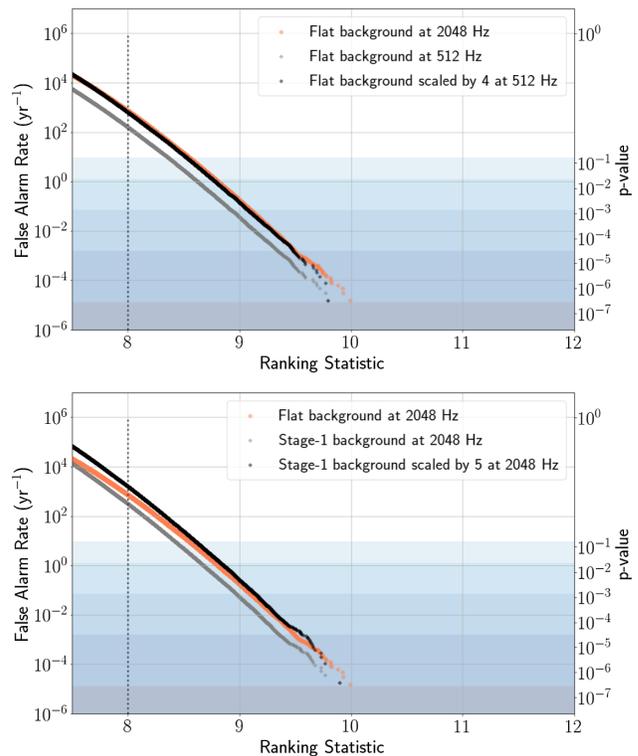


FIG. 8. Plot depicting FAR vs ranking statistic curves for the backgrounds obtained for the flat and stage-1 search at different sampling rates. The top panel shows the backgrounds obtained for the flat search performed at sampling rates 512 Hz (gray) and 2048 Hz (orange). The bottom panel shows the background obtained at a sampling rate of 2048 Hz for the flat (orange) and stage-1 search (gray). On scaling the flat background obtained at a sampling rate of 512 Hz by a factor of 4 (top), and stage-1 background obtained at 2048 Hz by 5 (bottom), the resultant backgrounds (black) approximately match the standard flat background (orange) obtained at 2048-Hz sampling rate.

argument still lacks concrete proof, we think it makes our argument much more robust, at least for the standard search with a bank of mismatch 0.97.

While the scaled background may not precisely match the background of the flat search, it is still a monotonic function of the detection statistic and reasonably close to the flat search background. Hence, the FAR estimate based on the scaled background can be used for detection, as long as a reasonable FAR threshold to claim a detection is decided by comparing it with the corresponding flat search.

III. SEARCH FOR CBC IN O1 AND O2 DATA

We search for CBCs using the two-stage hierarchical search [60] over the data from the first (O1) and second (O2) observing runs of twin LIGO detectors. We use 21.39 days of coincident data from O1 and 31.4 days from O2.

The periods of poor-quality data are marked and removed from the analysis using data-quality flags, Category 1 (CAT-1) and Category 2 (CAT-2) [55]. CAT-1 vetoes remove the times during which at least one of the key components of a detector was not operational in the nominal configuration due to critical issues. The duration over which excessive noise is observed due to instrumental artifacts is marked and removed by CAT-2 flags.

As described in Sec. II B, the data undergo preprocessing before entering the matched filtering step. In both stages, we use 512 s of overlapping data segments for matched-filter computation. We pad data segments with zeros 144 s at the beginning and 16 s at the end to avoid the artifacts generated from the discrete Fourier transform. Once the data segment is prepared, we perform a hierarchical search in two stages.

We begin the search by matched-filtering data segments sampled at 512 Hz with a coarse bank (see Sec. II A 1) and

obtain a list of stage-1 triggers above coarse thresholds on ρ and $\tilde{\rho}$. Triggers with $\rho > 3.5$ that pass the χ_r^2 test with $\tilde{\rho} > 3.5$ get further reweighted by χ_{sg}^2 veto. The choice of coarse thresholds for stage-1 search may seem arbitrary. However, we tested out different values for ρ and $\tilde{\rho}$ and found that setting both values at 3.5 gives the optimal computational cost of handling bulk triggers.

The surviving single-detector triggers then undergo a coincidence test (see Sec. II C) to obtain foreground candidates. These foreground candidates are then followed up in stage 2.

Stage 2 of the hierarchical search begins with matched-filtering data segments sampled at 2048 Hz that contain foreground candidates with $\Lambda_{\text{opt}} \geq 7$ [40] from stage 1. Each segment is filtered using a unique stage-2 bank (see Sec. II A 2) constructed from the dynamic union of the nbhds around each followed-up trigger template. The matched-filter SNRs generated in this stage are then reweighted with fine thresholds on ρ and $\tilde{\rho}$ of 4. As described in Sec. II C, the resultant triggers are then subjected to a coincidence test to obtain the second stage's foreground candidates.

The final step in the search involves assigning significance to the potential foreground candidates ($\Lambda_{\text{opt}} > 8$) obtained in stage 2. We assign FARs to these candidates using a scaled stage-1 background, as described in Sec. II D. Based on this background, we present the results from the analysis in the next section.

We report the recovery of all ten confirmed GW events with FAR below 1 per year in stage 2 of hierarchical search. These events were previously detected by the flat analysis in GWTC-1. Although the detection statistics used in both the stages of hierarchical search are more recent than those used in the flat analysis of GWTC-1, we still detect these events with nearly similar detection confidence levels in stage 2 but with a computational gain in the matched

TABLE II. List of GW events detected via hierarchical search. The events are arranged in the ascending order of their event time. We report and compare detected events' FARs, network SNRs ($\hat{\rho}_T$), and redshifted chirp masses ($\mathcal{M}_{\text{chirp}}$) in the stage-1 and stage-2 search with the events reported by the flat search in Abbott *et al.* [7]. We see an improvement in the FAR and network SNR values for the events, with network SNR varying between 9 and 12 from stage 1 to stage 2.

Sl. no.	Event	UTC	Flat			Hierarchical					
			FAR (yr ⁻¹)	$\hat{\rho}_T$	$\mathcal{M}_{\text{chirp}} (M_{\odot})$	Stage 1			Stage 2		
						FAR (yr ⁻¹)	$\hat{\rho}_T$	$\mathcal{M}_{\text{chirp}} (M_{\odot})$	FAR (yr ⁻¹)	$\hat{\rho}_T$	$\mathcal{M}_{\text{chirp}} (M_{\odot})$
1	GW150914	09:50:45.4	$\leq 1.53 \times 10^{-5}$	23.6	32.75	1.52×10^{-5}	23.3	29.71	1.52×10^{-5}	24.0	31.96
2	GW151012	09:54:43.4	0.17	9.5	18.47	0.42	8.9	18.68	0.05	9.8	18.31
3	GW151226	03:38:53.6	$\leq 1.70 \times 10^{-5}$	13.1	9.70	1.69×10^{-5}	11.9	9.89	1.69×10^{-5}	13.1	9.72
4	GW170104	10:11:58.6	$\leq 1.39 \times 10^{-5}$	13.0	20.19	1.37×10^{-5}	12.2	18.37	1.37×10^{-5}	12.9	29.17
5	GW170608	02:01:16.5	$\leq 3.09 \times 10^{-4}$	15.4	8.61	3.08×10^{-4}	8.9	8.65	3.08×10^{-4}	14.8	9.03
6	GW170729	18:56:29.3	1.36	9.8	40.27	1.68	9.3	54.41	0.05	10.6	47.51
7	GW170809	08:28:21.8	1.45×10^{-4}	12.2	23.53	0.56	11.3	29.71	1.70×10^{-3}	12.1	23.65
8	GW170814	10:30:43.5	$\leq 1.25 \times 10^{-5}$	16.3	25.20	1.25×10^{-5}	16.0	25.09	1.25×10^{-5}	17.2	26.58
9	GW170817	12:41:04.4	$\leq 1.25 \times 10^{-5}$	30.9	1.20	2.51×10^{-5}	28.7	1.20	1.25×10^{-5}	31.5	1.20
10	GW170823	13:13:58.5	$\leq 3.29 \times 10^{-5}$	11.1	23.61	3.30×10^{-5}	11.3	32.32	3.30×10^{-5}	11.1	46.85

filtering by a factor of ~ 20 . A comparison of the recovered events' FARs, network SNRs ($\hat{\rho}_T \equiv \sqrt{\hat{\rho}_H^2 + \hat{\rho}_L^2}$), and redshifted chirp mass from the flat search in GWTC-1 and both the stages of the hierarchical search, is given in Table II.

In our analysis, we recover the loudest events—GW150914, GW151226, GW170104, GW170608, GW170814, GW170817, and GW170823, with comparable FARs in both the stages of the hierarchical search. However, the network SNRs of these events improve in stage 2. The remaining events—GW151012, GW170729, and GW170809, see improvements in their FARs and network SNRs in the stage-2 search.

IV. COMPARISON WITH THE FLAT SEARCH

In this section, we compare the search sensitivities of hierarchical and flat search pipelines using similar detection statistics as defined in Sec. II C. We also highlight the computational advantages of using the former pipeline over the latter.

A. Comparison of sensitivities

The sensitivity of a search pipeline is measured in terms of the total number of astrophysical signals detected at a given detection statistics and a fixed FAR threshold. In order to measure this quantity, a population of simulated GW signals is injected into the real data and recovered using the search pipeline. For a population of binary mergers, uniformly distributed over comoving volume (V), one can compute the sensitive reach of the detectors in terms of the volume covered in the given observable time. Suppose that a binary's merger rate is defined by μ_m ; then, the number of detection that one can make above a certain FAR threshold in T_{obs} observation time is the product of volume, time, and merger rate $\mu_m \langle VT \rangle$ [61]. The sensitive volume-time $\langle VT \rangle$ over here is defined as

$$\langle VT \rangle_{\{\vec{\theta}\}} = T_{\text{obs}} \int_0^\infty p(z|\{\vec{\theta}\}) \frac{dV}{dz} \frac{1}{(1+z)} dz, \quad (19)$$

where $p(z|\{\vec{\theta}\})$ is the probability of recovering a signal with parameters $\vec{\theta}$ at a redshift z . For a constant value of μ_m , the ratio of VT can be exploited to compare the sensitivities of any two search pipelines [7,29].

In our study, we compare the search sensitivities of the hierarchical and flat search pipelines by computing the ratio of their VT for a common injection set.

To calculate VT for each pipeline, we inject quadrupolar GW signals from the nonprecessing BBH, BNS, and NSBH like sources into the data. These signals are generated using waveform models `SpinTaylorT4` [62] for BNS and `SEOBNRv4_opt` [41] for BBH and NSBH systems. To remain agnostic about the binary merger population, we distribute the signals obtained from these

TABLE III. Table depicting the ranges for redshifted component masses, total mass, and dimensionless effective spins for each compact object of injected BBH, BNS, and NSBH sources.

Parameter	BBH	BNS	NSBH
$m_1(M_\odot)$	2.5–150	1–2.5	2.5–97.5
$m_2(M_\odot)$	2.5–150	1–2.5	1–2.5
$M_{\text{tot}}(M_\odot)$	5–300	2–5	3.5–100
$\chi_{1,z}$	0–0.998	0–0.4	0–0.998
$\chi_{2,z}$	0–0.998	0–0.4	0–0.4

models uniformly over the chirp distance between 50 and 400 Mpc. We uniformly distribute the component masses for BNS and distribute the logarithms of component mass of BBH and NSBH injections in the ranges provided in Table III. Thus, we generate 12,203 BNS and 16,271 BBH and NSBH injections each.

We inject the generated signals in 5 days of coincident data in O1 observed from September 12, 2015, to September 26, 2015, and analyze it using the flat and hierarchical search pipelines separately. The matched-filtering and coincidence studies in the hierarchical search are carried out as per Sec. III. In the case of flat search, we perform matched filtering over data segments sampled at 2048 Hz and identify triggers with ρ and $\tilde{\rho}$ above 4 in each detector. We run a coincidence test over the collected single-detector triggers with the appropriate clustering in time as defined in Sec. II B. Here, triggers observed within 100 ms of a time window in two detectors are identified and ranked according to their statistic values (see Sec. II C).

The foreground candidates obtained in both the searches are assigned FARs based on their respective noise backgrounds using similar ranking statistics described in Sec. II C. In the flat search, we estimate the background by time-sliding triggers across the detectors. Each trigger is shifted by 100 ms in time, and then again, the statistic is estimated. A time slide of 100 ms can generate a large number of triggers. Therefore, we first cluster the candidates within a time window of 10 s and then apply decimation to the background as performed in the flat search. In the case of hierarchical search, we assign FARs to the detected candidates after scaling the stage-1 background, as described in Sec. II D. The recovered candidates via clustering over statistic values are then sorted with respect to their FARs. A highly ranked candidate with a FAR value below 1 per year [7] and falling within 1 s of merger time is marked as a detected injection in both the searches.

Figure 9 compares the sensitivities of the hierarchical and flat search. As can be seen in the top panel, most of the injections are recovered with comparable FARs by both searches. We infer this result from the high density of scattered points lying near the diagonal of the plot. We also see that some injections are only recovered by one search. However, these stand-alone recoveries in the majority have a low astrophysical significance. A few of these injections

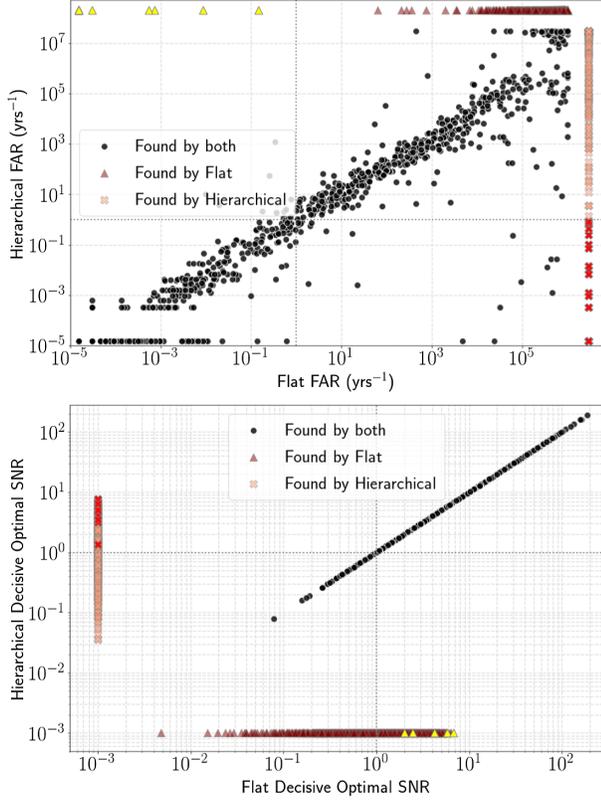


FIG. 9. Scatter plots of FARs (top) and decisive optimal SNRs (bottom) for the injections recovered in the hierarchical search vs flat search. The injections found by both searches are represented by circular points (black). The other markers, cross and triangle, show the injections found by only one search. A few of these points for the flat (yellow) and hierarchical (red) search show low FARs and injected optimal SNRs. The concentration of the points near the diagonal in the top panel implies that the estimated FARs by both the searches are reasonably close. A vertical and horizontal line in the plot shows FAR of 1 per year in the top panel and a decisive optimal SNR of 1 in the bottom panel. The bottom panel confirms that the injections which are not detected by either of the searches were for low ($\sim < 8$) decisive optimal SNR.

show low FARs, for instance, the injections recovered by only hierarchical search represented in color in Fig. 9. A follow-up study on these significant detections showed that these injections were made at very low optimal SNRs (see bottom panel Fig. 9) and were likely recovered due to coincidence with noise fluctuation around the injection time. In the other case where injections are recovered by only flat search, hierarchical search misses these injections because stage-1 search fails to recover them.

The FAR comparison in Fig. 9 shows that both flat and hierarchical search performs similarly for loud CBC injection. However, the sensitivity towards detecting fainter injections varies for both searches. This conclusion is further supported by the VT comparison in Fig. 10. In the top panel of Fig. 10, we see that the sensitivity of stage-1 search is lower

than flat search across all the chirp mass and IFAR⁴ bins. This result is expected as the loss in matched-filter SNRs is bound to happen in stage 1 due to low sampling rates and the use of a coarse bank. However, performing a stage-2 search on the potential foreground candidates from stage 1 retains the overall sensitivity of the search pipeline, which can be viewed in the bottom panel of Fig. 10. In this plot, we see that the sensitivity of hierarchical search remains consistent with the flat search with VT ratio varying between a factor of 1 ± 1.042 and 1 ± 0.08 for IFAR of 10 y depending on the chirp mass bins.

B. Comparison of computational efficiencies

Now we estimate the computational cost of matched filtering for the flat and hierarchical search.

The computational cost of matched filtering relies on the number of FFT operations performed on a segment using a bank of templates. As defined previously, FFT operations scale as $N \log_2 N$. In the case of flat search, we filter a data segment sampled at 2048 Hz with the entire flat bank. If the segment is of length 512 s, then N in the flat search is 512×2048 , and the number of matched-filter operations is $512 \times 2048 \times 428,725 \times \log_2(512 \times 2048)$, where 428,725 represents the number of templates in the flat bank.

In the case of hierarchical search, the total number of matched-filter operations is the sum of the number of FFTs performed in stage 1 and stage 2. Since in stage-1 search we matched filter a data segment sampled at 512 Hz using the coarse bank containing 85,080 templates, the number of matched-filter operations becomes $512 \times 512 \times 85080 \log_2(512 \times 512)$. If the same segment gets followed up to the stage-2 search, the number of matched-filter operations reduces due to fewer templates in a stage-2 bank. The number of templates in this bank can vary for each segment and detector, as can be seen from Fig. 11. Thus, we compute the total number of FFT operations for all the flat and combined stages of the hierarchical search segments. To estimate the overall gain in the computational speed, we take the ratio of the computed FFT operations for the flat to the hierarchical search.

We first define the following quantities:

Quantity	Description
N_{seg}	Total number of data segments in two detectors
t_{seg}	Duration of each segment
f_{flat}	Sampling rate for flat and stage-2 search
f_{coarse}	Sampling rate for stage 1
$N_{\text{temp}}^{\text{flat}}$	Number of templates in the flat bank
$N_{\text{temp}}^{\text{stage1}}$	Number of templates in the stage-1 bank
$N_{\text{temp}}^{\text{stage2}}$	Total number of templates for <i>all</i> the segments used in the stage-2 search

⁴Inverse false alarm rate (IFAR = $\frac{1}{\text{FAR}}$).

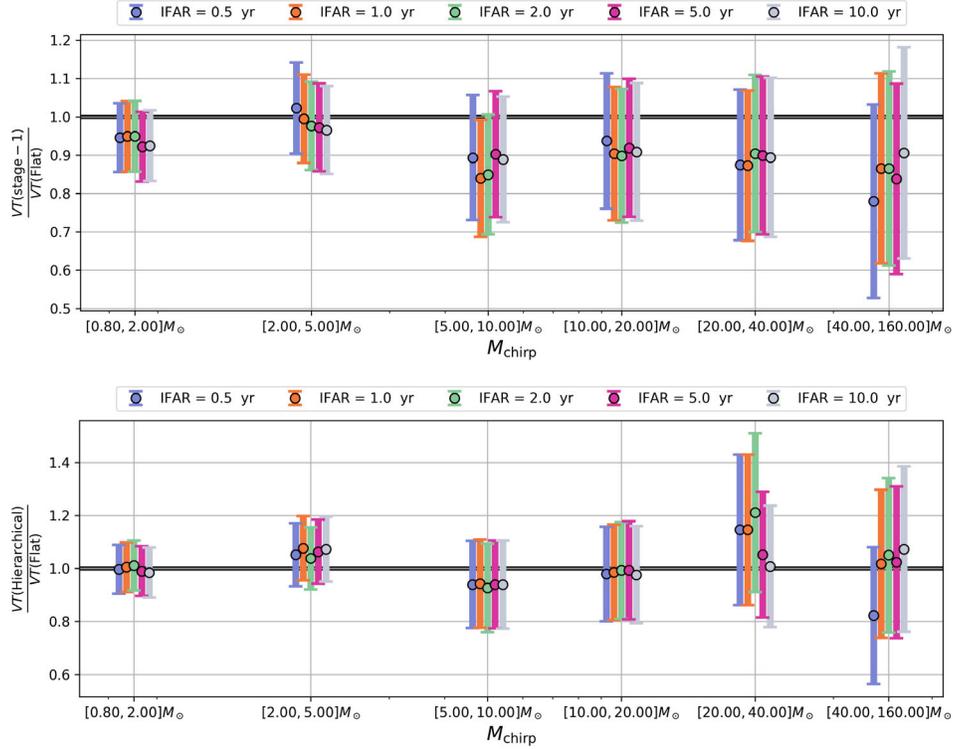


FIG. 10. Comparison of volume \times time (VT) ratio of (top) stage 1, and (bottom) stage 2 (or hierarchical) with flat search. The sensitivity of stage-1 search drops for higher chirp mass bins across all IFAR bins in the top panel. In the bottom panel, the VT ratio improves across entire chirp mass and IFAR bins, maintaining the overall sensitivity of hierarchical search comparable to flat.

Let

$$\begin{aligned} O_{\text{flat}} &= k f_{\text{flat}} t_{\text{seg}} \log_2(f_{\text{flat}} t_{\text{seg}}), \\ O_{\text{coarse}} &= k f_{\text{coarse}} t_{\text{seg}} \log_2(f_{\text{coarse}} t_{\text{seg}}), \end{aligned} \quad (20)$$

where O_{flat} and O_{coarse} are the number of floating-point operations required to perform a FFT for a segment at the flat and coarse sampling rates, respectively. k is a factor of

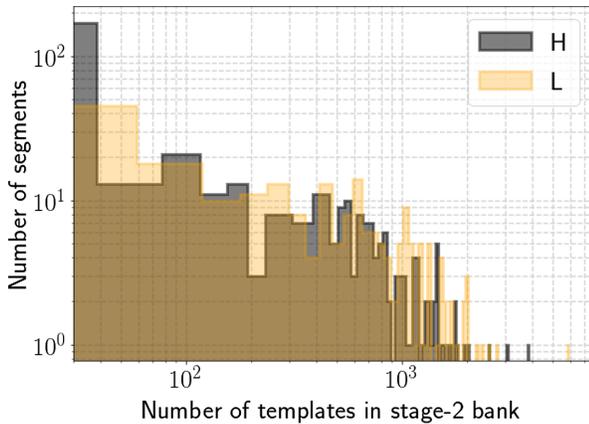


FIG. 11. Histogram depicting the number of templates in stage-2 bank generated for each data segment.

few which cancels out from the numerator and denominator. Thus, the gain is given by

$$\text{gain} \approx \frac{N_{\text{seg}} N_{\text{temp}}^{\text{flat}} O_{\text{flat}}}{N_{\text{seg}} N_{\text{temp}}^{\text{stage1}} O_{\text{coarse}} + N_{\text{temp}}^{\text{stage2}} O_{\text{flat}}}. \quad (21)$$

While the number of templates in flat search and stage 1 is fixed for all the segments, it varies for each segment in stage 2 as only specific triggers are followed up and filtered using a stage-2 bank. The total area of the histograms for the two detectors together shown in Fig. 11 provides us with $N_{\text{temp}}^{\text{stage2}}$. Since $N_{\text{temp}}^{\text{stage2}}$ is much smaller than $N_{\text{seg}} N_{\text{temp}}^{\text{stage1}}$, the computation in stage 1 dominates the cost, so the stage-2 cost does not affect the gain.

Substituting the numerical values, $N_{\text{seg}} = 390[\text{H}] + 225[\text{L}] = 615$, $t_{\text{seg}} = 512$ s, $f_{\text{flat}} = 2048$ Hz, $f_{\text{coarse}} = 512$ Hz, $N_{\text{temp}}^{\text{flat}} = 428,725$, $N_{\text{temp}}^{\text{stage1}} = 85,080$, and $N_{\text{temp}}^{\text{stage2}} = (132,036[\text{H}] + 132,134[\text{L}]) = 264,170$, the gain yielded is 22 for the analysis. We do not expect this number to change significantly for different observing runs. We also compare the actual CPU core hours used by the flat and hierarchical search for performing the matched-filtering operations. We found that the total CPU core hours used by the hierarchical search are around 824.16 and 547.37, respectively, for the Hanford and Livingston detectors. These numbers are nearly

1/19 times the number obtained for the flat search, i.e., 15,471.81 for Hanford and 10,478.64 for Livingston. Thus, we conclude that with the present setting, *the hierarchical search provides an overall computational speed-up by a factor of ~ 20* .

V. CONCLUSION AND DISCUSSION

Efficient searches for GWs originating from CBCs can expand the size and dimensionality of the search parameter space to detect interesting sources with present and future detectors. The hierarchical search is perhaps the most straightforward approach that brings more than one order of magnitude enhancement in the efficiency without compromising the robustness of the search. In this work, we successfully demonstrate the efficiency of hierarchical search by applying the analysis on the first two observing runs of Advanced LIGO. By introducing essential modifications to the previously developed codes, we transform them into a complete analysis pipeline [60]. We improve the selection criteria for single-detector triggers using chi-square and sine-Gaussian vetoes to reweigh matched-filter SNRs. We also implement coincident detection statistics formulated in [29,58] in the hierarchical search that utilizes phase and time differences between detectors and detection parameters, significantly reducing false alarms due to noise events. With our pipeline, we recover all the events in the LIGO-Virgo Collaboration’s official transient catalog, GWTC-1, detected by the standard PyCBC analysis with nearly the same statistical confidence and a whopping factor of 20 computational speed-up. This work also demonstrates that hierarchical search is at hand for production analysis of the present and upcoming datasets from ground-based detectors.

Following Gadre *et al.* [40], we estimate the detected candidates’ significance by scaling the noise background obtained in stage 1 with a factor close to the speed-up factor. Although the argument on assigning significance to detected candidates using this background may not be so rigorous, our work shows that this prescription works. The background estimation for the hierarchical search needs more scrutiny, and our future goal is to address this issue. It is outside the scope of the present investigation because an in-depth mathematical and statistical analysis of the empirical background estimation using time slides will be

required. While the outcome of this exercise builds enough confidence for application in production runs that are otherwise restrictive due to computational cost, we plan to carry out an extensive study focused on accurate background estimation for the hierarchical search.

In our opinion, the hierarchical search pipeline can be used for ambitious searches that are currently deferred due to computational limitations. For instance, a search for binaries with nonaligned spins and subsolar sources requires an enormous number of templates. With hierarchical search, we can attempt to carry out their search at feasible computation cost without compromising the accuracy of sensitivity of the search. The hierarchical strategy could also reduce the computation cost of low-latency searches, which we plan to demonstrate in the future. Developing a comprehensive offline or a low-latency search for such sources is an arduous task ahead, and more sophisticated techniques will have to be brought in, in the coming years. Nevertheless, the hierarchical search is a major step in this direction that should be exploited.

ACKNOWLEDGMENTS

The authors acknowledge the computational resources provided by the IUCAA LDG cluster Sarathi, LIGO Laboratory, and are supported by National Science Foundation Grants. The authors acknowledge support from Soumen Roy in providing code and helping in the generation of banks. The authors are grateful for the valuable discussions from Shreejit P. Jadhav at various stages of this work. K. S. acknowledges technical support for cluster-related issues from Deepak Bankar. The hierarchical search pipeline uses PyCBC version 1.16.13 and is built upon LALSuite [63], NumPy [64], SciPy [65], and Astropy [66]. K. S. acknowledges the Inter-University Centre of Astronomy and Astrophysics (IUCAA), India, for the fellowship support. B. G. acknowledges the support of the Max Planck Society. S. M. acknowledges support from the Department of Science and Technology (DST), Ministry of Science and Technology, India, provided under the Swarna Jayanti Fellowships scheme. S. V. D. acknowledges the support of the Senior Scientist Platinum Jubilee Fellowship from National Academy of Sciences India (NASI), India. This manuscript has been assigned a LIGO Document No. LIGO-P2100202.

[1] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **116**, 061102 (2016).

[2] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **116**, 131103 (2016).

[3] B. P. Abbott *et al.*, *Phys. Rev. D* **93**, 112004 (2016).

[4] S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. F. Da Silva, and G. Mitselmakher, *Phys. Rev. D* **93**, 042004 (2016).

- [5] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. E. Creighton, R. Everett *et al.*, *Phys. Rev. D* **95**, 042001 (2017).
- [6] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. D. Canton, T. Dent, S. Fairhurst *et al.*, *Classical Quantum Gravity* **33**, 215004 (2016).
- [7] B. P. Abbott *et al.*, *Phys. Rev. X* **9**, 031040 (2019).
- [8] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, *Astrophys. J.* **872**, 195 (2019).
- [9] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, *Astrophys. J.* **891**, 123 (2020).
- [10] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
- [11] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, *Classical Quantum Gravity* **32**, 074001 (2015).
- [12] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, *Classical Quantum Gravity* **32**, 024001 (2015).
- [13] R. Abbott *et al.*, *Phys. Rev. X* **11**, 021053 (2021).
- [14] R. Abbott *et al.*, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [15] R. Abbott *et al.*, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
- [16] B. P. Abbott *et al.*, *Astrophys. J. Lett.* **892**, L3 (2020).
- [17] B. P. Abbott *et al.*, *Phys. Rev. D* **102**, 043015 (2020).
- [18] R. Abbott *et al.*, *Astrophys. J. Lett.* **896**, L44 (2020).
- [19] R. Abbott *et al.*, *Astrophys. J. Lett.* **900**, L13 (2020).
- [20] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [21] S. V. Dhurandhar and B. S. Sathyaprakash, *Phys. Rev. D* **49**, 1707 (1994).
- [22] S. V. Dhurandhar and B. F. Schutz, *Phys. Rev. D* **50**, 2390 (1994).
- [23] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [24] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [25] S. Sachdev, S. Caudill, H. Fong, K. L. R. Lo, C. Messick, D. Mukherjee, R. Magee, L. Tsukad *et al.*, [arXiv:1901.08580](https://arxiv.org/abs/1901.08580).
- [26] C. Hanna, C. Caudill, C. Messick, A. Reza, S. Sachdev, L. Tsukada, K. Cannon, K. Blackburn, J. D. E. Creighton, H. Fong *et al.*, *Phys. Rev. D* **101**, 022003 (2020).
- [27] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, *Classical Quantum Gravity* **33**, 175012 (2016).
- [28] F. Aubin, F. Brighenti, R. Chierici, D. Estevez, G. Greco, G. M. Guidi, V. Juste, F. Marion, B. Mours, E. Nitoglia *et al.*, *Classical Quantum Gravity* **38**, 095004 (2021).
- [29] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, *Phys. Rev. D* **102**, 022004 (2020).
- [30] Q. Chu, Ph.D. thesis, The University of Western Australia (2017).
- [31] T. D. Canton and I. W. Harry, [arXiv:1705.01845](https://arxiv.org/abs/1705.01845).
- [32] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, *Phys. Rev. D* **94**, 024012 (2016).
- [33] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **123**, 161102 (2019).
- [34] R. Abbott *et al.*, [arXiv:2109.12197](https://arxiv.org/abs/2109.12197).
- [35] A. H. Nitz and Y.-F. Wang, *Astrophys. J.* **915**, 54 (2021).
- [36] T. Akutsu, M. Ando, K. Arai, Y. Arai, S. Araki, S. Araya, N. Aritomi, Y. Aso, S. Bae, Y. Bae *et al.*, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [37] B. Iyer *et al.*, LIGO-India Technical Report No. LIGO-M1100296, 2011, <https://dcc.ligo.org/LIGO-M1100296/public>.
- [38] S. D. Mohanty and S. V. Dhurandhar, *Phys. Rev. D* **54**, 7108 (1996).
- [39] S. D. Mohanty, *Phys. Rev. D* **57**, 630 (1998).
- [40] B. Gadre, S. Mitra, and S. Dhurandhar, *Phys. Rev. D* **99**, 124035 (2019).
- [41] C. Devine, Z. B. Etienne, and S. T. McWilliams, *Classical Quantum Gravity* **33**, 125025 (2016).
- [42] R. Prix, *Classical Quantum Gravity* **24**, S481 (2007).
- [43] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [44] S. Roy, A. S. Sengupta, and N. Thakor, *Phys. Rev. D* **95**, 104045 (2017).
- [45] S. Roy, A. S. Sengupta, and P. Ajith, *Phys. Rev. D* **99**, 024048 (2019).
- [46] I. W. Harry and S. Fairhurst, *Phys. Rev. D* **83**, 084002 (2011).
- [47] D. Keppel, *Phys. Rev. D* **87**, 124003 (2013).
- [48] D. Keppel, [arXiv:1307.4158](https://arxiv.org/abs/1307.4158).
- [49] A. Buonanno, Y. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003).
- [50] C. Kalaghatgi, P. Ajith, and K. G. Arun, *Phys. Rev. D* **91**, 124042 (2015).
- [51] A. S. Sengupta, S. V. Dhurandhar, A. Lazzarini, and T. Prince, *Classical Quantum Gravity* **19**, 1507 (2002).
- [52] B. P. Abbott *et al.*, *Classical Quantum Gravity* **33**, 134001 (2016).
- [53] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson *et al.*, *Classical Quantum Gravity* **34**, 064003 (2017).
- [54] Y. Huang, H. Middleton, K. K. Y. Ng, S. Vitale, and J. Veitch, *Phys. Rev. D* **98**, 123021 (2018).
- [55] B. P. Abbott *et al.*, *Classical Quantum Gravity* **35**, 065010 (2018).
- [56] B. Allen, J. K. Blackburn, P. R. Brady, J. D. E. Creighton, T. Creighton, S. Droz, A. D. Gillespie, S. A. Hughes, S. Kawamura, T. T. Lyons *et al.*, *Phys. Rev. Lett.* **83**, 1498 (1999).
- [57] A. H. Nitz, *Classical Quantum Gravity* **35**, 035016 (2018).
- [58] A. H. Nitz, T. Dent, T. D. Canton, S. Fairhurst, and D. A. Brown, *Astrophys. J.* **849**, 118 (2017).
- [59] R. Biswas, P. R. Brady, J. Burguet-Castell, K. Cannon, J. Clayton, A. Dietz, N. Fotopoulos, L. M. Goggin, D. Keppel, C. Pankow, L. R. Price, and R. Vaulin, *Phys. Rev. D* **85**, 122008 (2012).
- [60] Hierarchical search pipeline, https://github.com/Kanchan-05/pycbc/tree/Hierarchical-search/hierarchical_search (2021).
- [61] V. Tiwari, *Classical Quantum Gravity* **35**, 145009 (2018).

- [62] R. Sturani, S. Fischetti, L. Cadonati, G. M. Guidi, J. Healy, D. Shoemaker, and A. Viceré, *J. Phys. Conf. Ser.* **243**, 012007 (2010).
- [63] Ligo Scientific Collaboration, ligo algorithm library-LALSuite, free software (gpl) (2018).
- [64] S. v. d. Walt, S. C. Colbert, and G. Varoquaux, *Comput. Sci. Eng.* **13** (2011).
- [65] P. Virtanen *et al.*, *Nat. Methods* **17**, 261 (2020).
- [66] A. M. Price-Whelan *et al.*, *Astron. J.* **156**, 123 (2018).